Estimating the Long-Run Effects of Resource Allocation Mechanisms

Kleoniki Vlachou
AvMet Applications Inc.
Reston, Virginia, USA
Email: Vlachou@avmet.com

David J. Lovell
Department of Civil and Environmental Engineering and Institute for Systems Research
University of Maryland
College Park, Maryland, USA
Email: lovell@umd.edu

***Abstract*—A potentially powerful way to improve the planning of en route traffic management initiatives would be to grant carriers some ability to express their preferences for allocation details that the service provider would otherwise be ambivalent to. In this paper we propose an idea of how airlines can express their preferences without revealing too much information about their business models and cost functions and we test this idea for its validity. We also examine if airlines in the long-run will be treated equally by the allocation mechanism chosen, by receiving on average the number of slots close to what would have been considered their fair share.**

Keywords-flow constrained areas; airspace flow programs; resource allocation; simulation; fairness

I. Introduction

The air transportation system in the United States is one of the most complex systems in the world. Every day approximately 60,000 flights of commercial, military and general aviation aircraft occupy the National Airspace System (NAS). Despite a series of events that made the aviation industry suffer during recent decades - the terror attacks of September 11th, very high fuel prices and the worldwide recession - the Federal Aviation Administration (FAA) forecasts show that air traffic will continue to grow [1]. With more than 750 million passengers carried in 2014, the FAA forecasts that more than 1.1 billion passengers will be flown by 2035, an average increase of 1.9 percent a year.

Balanced against the increasing demand for air transportation we have the capacity of the system, which has significant variations, due to factors such as fluctuating weather conditions and equipment outages. Thus, the Federal Aviation Administration (FAA), who is responsible for alleviating the problems associated with imbalanced demand and capacity, is using various Air Traffic Flow Management (ATFM) initiatives to do so. Two of the most commonly used ATFM initiatives are Ground Delay Programs (GDPs) and Airspace Flow Programs (AFPs).

The Ground Delay Program (GDP) is a mechanism implemented when it is projected that the arrival demand at an airport would exceed capacity, usually because of adverse weather conditions around the airport area [2]. The goal of a GDP is to decrease the arrival rate and this is achieved by intentionally delaying the take-offs of inbound flights for specified amounts of time [3]. The motivation for doing so is that it is safer and cheaper for flights to absorb the required delay on the ground before take-off rather than when airborne.

GDPs are used to deal with capacity constraints around a destination airport, but in many cases bad weather impacts the capacity of an en-route sector and this problem has to be addressed differently. To this end, the FAA introduced the Airspace Flow Program (AFP) in 2006. To initiate the AFP, the first step is to identify the problematic area by creating a Flow Constrained Area (FCA), which is a geometric description of the region of constrained capacity. Then the Enhanced Traffic Management System (ETMS) takes the FCA description and produces a list of the flights and the times that they are expected to pass through the FCA [4], based on the carriers' declared schedules and flight plans. The available capacity is rationed among these flights, at the same time allowing such collaborative decision-making (CDM) tools as self-cancellations and intra-carrier slot swaps to take place. The resulting allocation is sent to the Flight Schedule Monitor (FSM) where a controlled departure time is assigned to each flight so that the flow through the FCA will not exceed the declared capacity.

A major shortcoming of this ATM initiative is the inability of carriers to express their preferences based on their business and economic models. The Collaborative Trajectory Options Program (CTOP) is a new concept that was introduced to fill this gap [4]. This initiative manages en-route congestion and allows NAS customers to submit cost-weighted sets of alternative trajectory options for their flights. There are two key enabling ideas at the core of CTOP. The first is that NAS customers can submit cost-weighted sets of trajectories, and they are able to update them as often as needed. The second is that traffic managers run allocation algorithms to adjust the demand to the reduced capacity while attempting to place each flight on the lowest cost option available. The newest concept, which essentially is an improvement of CTOP, is called Collaborative Airspace Constraint Resolution (CACR), which will extend the capabilities of CTOP by managing flights within 45 minutes prior to departure and with adequate automation assistance provided to traffic managers for defining airspace constraints [5]. The participation of carriers in these programs is optional since they can submit just a single flight plan. The level of their involvement depends on how much advanced planning they are willing to do, because of the extra time that is required to create all the alternative trajectory options sets [6]. Even the FAA anticipates that not all carriers will participate from the very beginning in this new method of providing flight path information and a transitional phase will need to exist [7].

II. PROBLEM DESCRIPTION

As discussed above, the initialization of AFP planning begins with the identification of the FCA and the produced list of flights scheduled to pass thought that FCA. A very important thing in the initiatives previously mentioned is that airlines can express their priorities. This is mainly the motivation for this research. Airlines tend to be reserved about how much information they want to share. So, one goal of this paper is to present an idea of how airlines can express their preferences without revealing too much information about their business models and cost functions. Additionally we want it to be simpler but just as powerful at injecting carrier preferences in the allocation process. We also examine if this way of expressing priorities is valid in the sense that airlines will have the incentive of being truthful to the priorities they state, without trying to game the system.

During an AFP the number of available slots is less than the number of flights scheduled to pass through the FCA, which means that these slots must be divided in a fair way among the carriers. So it is important to have an equitable resource allocation mechanism to do so. But even then, on a given day the slots allocated to an airline will not match exactly its fair share. Some days they will get more and some others less, so it is important to see if in the long run they will get on average what they deserve. If we consider the difference of the fair share from the actual allocation as an error, another goal of this paper is to measure this error.

Another aspect of this problem that we are mindful of is the variety in the sizes of carriers, or more precisely, the number of flights they have planned through the FCA. This does not stem only from the size of the carrier itself, but also takes into account the fact that FCAs are geographically specific, and carriers have definite geographic patterns with which they operate, regardless of their size. Nevertheless, in a given FCA, there will be "big" airlines, which will have many of their flights planned to pass through the affected area, and there will also be "smaller" ones with fewer flights. We would like to see what would be the impact of allocation procedures on those two different categories of carriers. For example, there might be cases where an airline has one flight scheduled to pass through the FCA. This means that its share in the available slots will be less than one, so the resulting slot allocation might omit this carrier altogether. It might be "fair" to do this on some fraction of days, but certainly not on every day, if the underlying pattern were to be repeated. So equity between carriers is a major concern, and perhaps more so for those who would expect to have a small presence in the schedule affected by a given FCA.

This work was supported by the Federal Aviation Administration under contract number WA1007261.

III. METHODOLOGY

For the purposes of this paper we will use an allocation mechanism previously developed by Vakili and Ball [8], which is proven to be fair, equitable, and immune to gaming. It is a two-step process:

Step 1: Determine the fair share of the constrained resource set for each carrier, using the original schedule as the basis of fairness.

Step 2: Allocate flights to slots in a manner consistent with the fair share determined in Step 1.

For the first step we make sure that each slot will be used and that the numbers of slots are assigned proportionally to each carrier. Since the number of available slots is less than the number of scheduled flights, it means that the fair share for each airline will be a smaller number than their initial number of flights, and most likely non-integer.

In the second step, airlines are essentially chosen randomly in proportion to their fair share. First the allocation begins with the fractional shares and this happens for a specific reason. Carriers with small presence in an area are very likely to have fair shares that will be less than 1. Allocating the fractional shares first increases the chances of small carriers to actually have a slot assigned to them. When an airline is chosen, it is assigned its highest priority flight that is still feasible. In this phase the airlines' preferences are taken into account. Some of the details of the algorithm are omitted here for the sake of brevity; for more information please see Vakili and Ball [6] and for the extension of this work, where variations of the two-step process are investigated, see Vlachou and Lovell [9].

One of the issues surrounding incorporating carriers' preferences into a collaborative decision-making setting is that, while they would like the final allocation to be sensitive to their wishes, they would prefer not to articulate those wishes in such a clear manner that their internal business models might be discerned. It is presumed, for example, that they could construct actual cost estimates associated with a number of different routing and delay options, but they would unlikely be willing to share such proprietary information. For this paper, we assume that their preference information is provided in a simpler and more obscure manner. Each airline will give to each of their flights a priority number ranging from 1 to 4. The greater the number assigned to a flight the more important this flight is. In TABLE I we can see an example of how the list of airlines' preferences might look. In this list, airline *A*, for its first flight (f_{A1}), will give a priority number 3, and for its second flight (f_{A2}), a priority number 4, which means that it considers the second flight more important than the first. So when airline *A* is selected to receive a slot, if both of these flights can be assigned to it, the allocation mechanism will pick the second flight.

TABLE I. EXAMPLE OF LIST WITH AIRLINES PREFERENCES

Flight	Priority Index
f_{A1}	3
f_{A2}	4
f_{B1}	2
f_{B2}	4
f_{A3}	1
f_{C1}	3
f_{C2}	4

As mentioned before, one of the goals of this paper is to examine if this way of expressing priorities is valid. We wanted to see if at the end most of the higher priority flights have been assigned to slots. Also we wanted to check if the delays incurred by the highest priority flights are less than the rest. We want airlines to give truthful preferences, which they will be more willing to do if they actually see that it makes sense delays-wise for some of their flights to have higher priority. If, for example, airlines give a 4 to all of their flights in order to game the system, since they will not be able to assign all of these flights to slots, they might miss the opportunity represented by a flight that is in reality more important than others to get a slot. Giving all flights the same priority number would cause them to be assigned following the Ration-By-Scheduled (RBS) method, as it is being currently used. This essentially would cancel the allocation mechanism proposed here and the potential benefits for the airlines to give priority to flights that are more important. Also it is important to understand that airlines do not compete with each other for a specific slot, so they gain nothing by submitting erroneous information. The priority numbers for flights are used to sort the flights of each airline separately. A flight of an airline with priority 4 does not compete with a flight with priority 4 of another airline.

When considering delay costs, we weighted each flight's delay by the priority assigned to it by its respective carrier. In our analysis we estimated the average delay per flight per priority number. Delays were estimated for the flights that were captured by the AFP program and had a slot assigned to them within the duration of the program. The flights that were not assigned to a slot, for which each airline must decide whether to reroute, cancel, or delay them further, were not included in the calculations. Another thing we examined is the average delay per flight per airline and in these calculations we did not weight the delays according to the priority number assigned to them.

In this problem there are two levels of randomness that we can identify and that we want to examine. The first is due to the random selection of airlines in the allocation algorithm. Even when the number of flights and slots stay the same, each time the allocation scheme is implemented the selection of airlines can differ. The second level of randomness comes from the fact that each time an AFP is implemented the number of flights that each airline has will vary. For the purpose of this research and in order to be able to have multiple repetitions of the allocation mechanism with varying input data, we used Monte Carlo simulation.

The first thing we measured from the simulation output was, for each airline, how much variance there was in the number of assigned slots. Since each airline does not get the exact same number of slots each time, it would be desirable for these numbers to be quite close. For example, if an airline gets on average 10 slots but one time gets 5 and the next time 15, this is something that might be quite challenging for the airline's dispatchers.

Also we wanted to see how much on average the slots assigned to the airlines deviated from their initially calculated fair share. In order to measure the deviation of the actual allocation to the fair share of each airline we used the following indicator:

$$\textit{Fair Share Deviation} = \frac{\textit{Computed Fair Share} - \textit{Actual Share}}{\textit{Number of Slots}} \quad (1)$$

The fair share deviation indicator, as it was similarly used by Carr et al. [10], shows how much the actual average share (number of slots allocated to each airline) deviates from the fair share estimated before the allocation. If the numbers are zero then the actual share matches the fair share computed. The bigger this number becomes (positive), the more the actual share deviates and essentially the particular airline receives fewer slots than its fair share. If the number is negative the particular airline receives more slots than its fair share. What would be desirable is for these numbers to be very close to zero.

IV. SIMULATION AND RESULTS

For the purposes of this paper we broke the simulation into two different parts, in order to isolate the variance that will appear in the outputs into two different sources. In the first part, we wanted to test the effect that the random allocation procedure has by itself. In order to do that, we used a deterministic set of input data, so the only randomness in the simulation is in the proportional allocation procedure itself. In the second part, recognizing that there are stochastic fluctuations in demand input data due to various causes (e.g., schedule changes, seasonality, unexpected cancellations and delays due to crew issues and maintenance, etc.) we wanted to add randomness to the input data.

A. Monte Carlo Simulation for Deterministic Set of Flights

In the first set of simulations we considered a deterministic set of input data to our Monte Carlo simulation. Rather than work with a particular geographic scenario and its associated FCA, we worked with a hypothetical AFP whose magnitude was commensurate with what tends to be observed in reality. We set the AFP time to be 2 hours, and tested 3 different scenarios for capacity reduction: 25, 20, and 15 aircraft per hour, respectively. This is out of a nominal flow of approximately 30 aircraft per hour. We employed synthetic carriers, but in order to properly represent the variety of share sizes that each might have in an AFP, we used actual data on flights into Boston Logan airport, from the Bureau of Transportation Statistics (BTS) [11]. From this we found the mean number of flights per airline, which is broken down as shown in TABLE II.

TABLE II. AVERAGE NUMBER OF FLIGHTS PER AIRLINE

Airline	Average Number of Flights Per Airline
1	3
2	2
3	6
4	5
5	10
6	16
7	6
8	11
9	4

In other words, the total number of flights scheduled initially to pass from the affected airspace area is 63 flights in 2 hours and the scenarios we considered are the capacity to drop to 50, 40, and 30 aircraft, respectively, for the two hours of AFP time. Each simulation ran for 1000 replications and the results are presented in the following tables and figures.

In Figure 1 we can see the fair share for each airline, computed as described in section III. After we ran the allocation procedure we estimated the average number of slots that each airline actually received, as can be seen in Figure 2. A first comparison of those two figures shows that, in the long run, airlines will be assigned numbers of slots that are close to their estimated fair share.

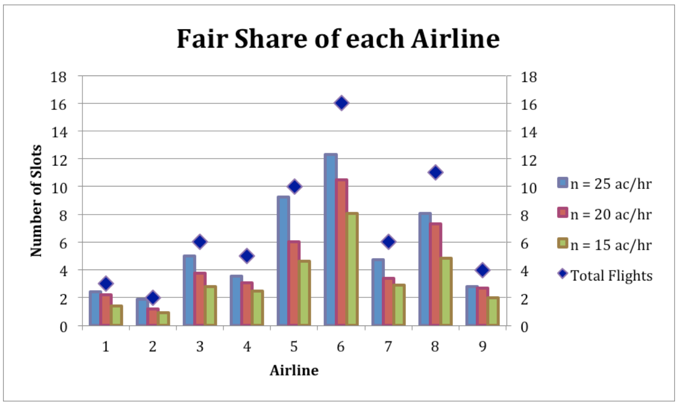

Figure 1. Computed fair share for each capacity reduction scenario

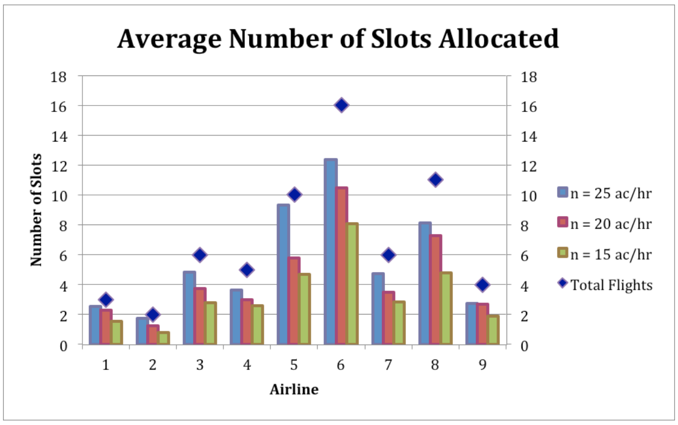

Figure 2. Average number of slots for each airline

We also present the median number of slots for each airline, in Figure 3, to have a better idea of how many slots airlines will usually get. Here it is clearer that smaller airlines have a good chance of getting slots. In TABLE III we present the variance in slot allocation for each airline. As we can see the variance for small and bigger airlines is approximately the same, although it tends to be slightly larger for bigger carriers than for smaller ones. Probably this is caused by the fact that the resource allocation mechanism is specifically designed to be protective of small airlines' slot claims.

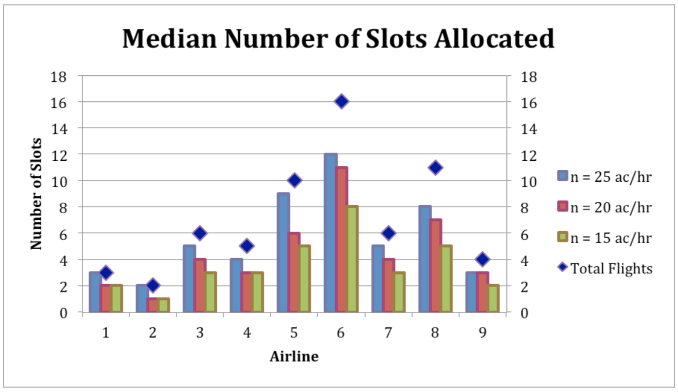

Figure 3. Median number of slots for each airline

TABLE III. VARIANCE OF SLOTS ALLOCATED TO EACH AIRLINE

		Capacity Reduction to n ac/hr			
		25	20	15	Total Flights
Airline	1	0.2486	0.2635	0.2483	3
	2	0.1857	0.3074	0.1559	2
	3	0.1484	0.2343	0.1684	6
	4	0.2331	0.1209	0.2425	5
	5	0.2098	0.1966	0.2126	10
	6	0.2354	0.3382	0.0803	16
	7	0.2031	0.3116	0.1484	6
	8	0.1125	0.3967	0.1571	11
	9	0.1964	0.2474	0.1140	4

In TABLE IV we can see the results for the fair share deviation indicator. As explained before, the fair share deviation indicator shows how much the actual average share deviates from the fair share estimated before the allocation. As can be seen in this table, the indicator for each airline is quite close to zero, which means that in the long run airlines will receive numbers of slots that are very close to their fair share.

In TABLE V we present the average delay per flight per airline incurred by flights that actually got assigned to a slot. For some airlines as the number of slots available gets smaller, their average delay increases and for some others it decreases. Smaller airlines tend to have the lowest delays, because they get fewer flights, so fewer flights are included in the calculations. The second airline has the least flights and when the capacity is reduced to 15ac/h, many times it would not get a slot at all, which means the delay is accounted as zero. There is no significant difference in the delays among the airlines.

TABLE IV. DEVIATION OF THE FAIR SHARE FROM THE ACTUAL AVERAGE ALLOCATION

		Reduced capacity, ac/hr		
		25	20	15
Airline	1	-0.0024	-0.0029	-0.0046
	2	0.0022	-0.0019	0.0036
	3	0.0035	0.0013	0.0012
	4	-0.0011	0.0007	-0.0034
	5	-0.0013	0.0055	-0.0028
	6	-0.0017	-0.0008	-0.0012
	7	0.0007	-0.0025	0.0023
	8	-0.0008	0.0006	0.0004
	9	0.0010	0.0000	0.0043

Finally in TABLE VI we present the results of the weighted average delay per priority number. Although the delays of flights with priority 4 were weighted more, the average delay per flight is consistently less than the average delay of flights with priority 3. From the simulation we observed that most of the flights assigned to slots were of priority 4 and 3 and consequently most of the flights left unassigned had a priority of 2 and 1. This explains the fact that the delays for flights with priority 2 and 1 are less than the ones with higher priorities. It is not that they were assigned to slots that were closer to the initial scheduled times, but that they weren't assigned to any slot at all.

TABLE V. AVERAGE DELAY (IN MIN) PER FLIGHT PER AIRLINE

		Reduced capacity, ac/hr			
		25	20	15	Total Flights
Airline	1	17.5	25.9	20.4	3
	2	9.9	9.9	3.5	2
	3	6.1	15.5	22.4	6
	4	10.3	22.6	17.1	5
	5	4.0	13.1	17.8	10
	6	7.9	10.4	14.3	16
	7	11.5	18.2	17.9	6
	8	7.5	12.9	26.0	11
	9	18.0	12.9	14.2	4

TABLE VI. WEIGHTED AVERAGE DELAY (IN MIN) PER FLIGHT PER PRIORITY NUMBER

		Reduced capacity, ac/hr		
		25	20	15
Priority Number	4	16.4	35.3	67.2
	3	22.4	47.6	76.6
	2	21.5	28.8	13.6
	1	15.5	20.4	12.1

B. Monte Carlo Simulation for Random Set of Flights

For the next part of our simulation we added variability to our input data (the number of flights for each airline). In reality, the schedule of airlines fluctuates, and since we wanted to comport with that, we looked at the traffic at Boston Logan for every Monday of two consecutive months – February and March of 2011. For the same airlines as mentioned before, we observed the number of flights and incorporated similar levels of variation into our simulation. For example, for Airline 2 it was observed that the number of flights in those consecutive months was consistently 2. So for this airline there was essentially no variation in the number of flights and we kept it this way in the simulation. For Airline 4, we observed that the number of flights within the two month period varied from 4 to 6. On average we would get 5 flights – same number as the previous experiment – but now we consider the number of flights to be uniformly distributed between 4 and 6. The mean number of flights matches the number of flights used in the previous set of experiments, but now the number of flights for each airline for each run varies following a uniform distribution whose extremes are identical to what was observed from the real data.

In Figure 4 we present the average fair share computed for each airline. Since the number of flights for each airline fluctuated from run to run, a new fair share was calculated in each run and in this table we have the average fair share from these runs.

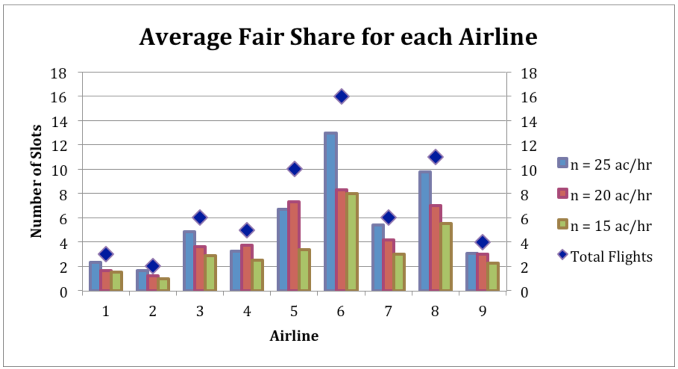

Figure 4. Computed average fair share for each airline and for each capacity reduction scenario

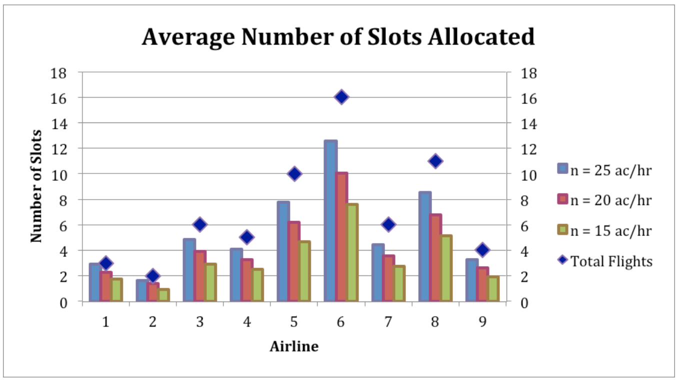

Figure 5. Average number of slots for each airline

In Figures 5 and 6, we present the average and median number of slots that each airline actually got after the resource allocation mechanism was implemented. Again here we can see that the number of slots allocated to each airline matches very well its estimated fair share, and also smaller carriers have good chances to get slots.

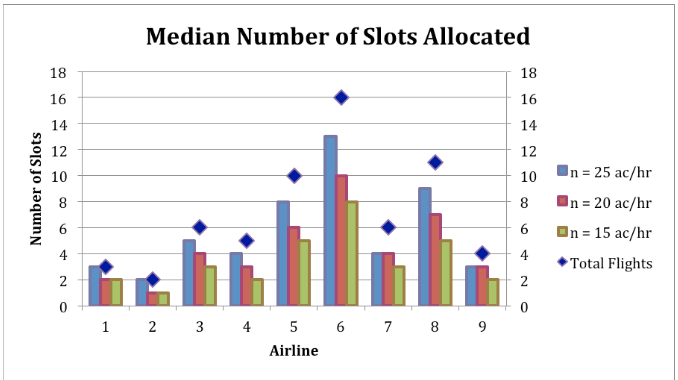

Figure 6. Median number of slots for each airline

In TABLE VII we see the variance of slots allocated to each airline. As we can see, the variance of slots for airlines with larger numbers of flights tends to be greater than for those with fewer flights. This is partly because the bigger airlines have greater fluctuation in their schedules on a day-to-day basis. Also the variances, compared to the results with deterministic flights, are greater. This is expected, and the difference represents the marginal contribution of the noise in the schedule to the observed variation. The added effect of variability on number of flights has contributed to that. This can be better observed in the following Figures 7-9.

TABLE VII. VARIANCE OF SLOTS ALLOCATED TO EACH AIRLINE

		Capacity Reduction to n ac/hr			
		25	**20**	**15**	**Total Flights**
Airline	**1**	0.4604	0.4335	0.2336	3
	2	0.2429	0.2614	0.1365	2
	3	0.4616	0.3600	0.1921	6
	4	0.6174	0.5093	0.3399	5
	5	1.0365	0.7458	0.4474	10
	6	1.1467	0.8742	0.4793	16
	7	0.5051	0.4534	0.2766	6
	8	0.6727	0.6434	0.3426	11
	9	0.3503	0.3307	0.1726	4

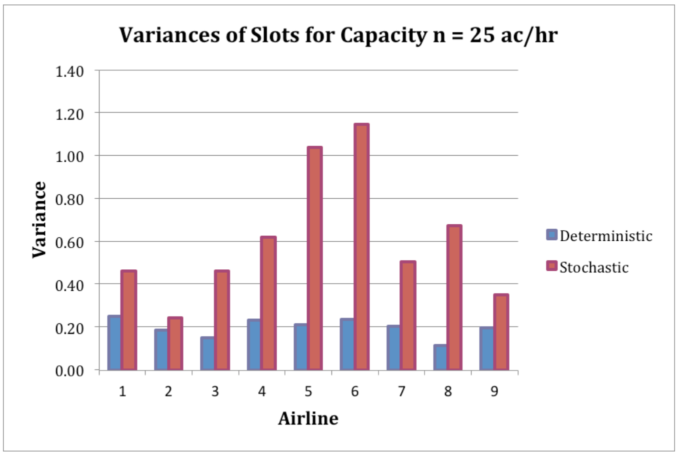

Figure 7. Variance of slots when capacity is reduced to 25 ac/hr

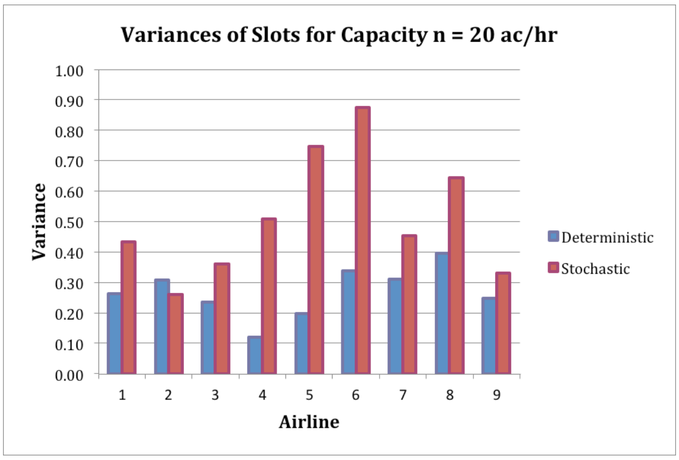

Figure 8. Variance of slots when capacity is reduced to 20 ac/hr

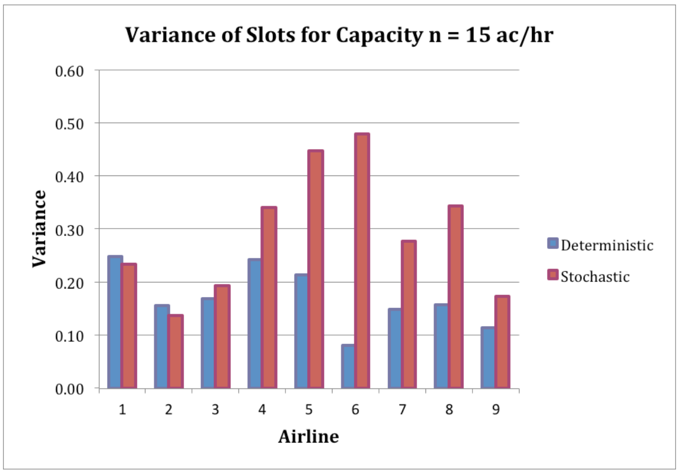

Figure 9. Variance of slots when capacity is reduced to 15 ac/hr

In Figures 7-9 we can see that airlines 5 and 6 have greater difference in their variances for each scenario, between the deterministic and the stochastic cases. This is due to the fact that they are bigger carriers and the range of flights they have is greater than the other airlines. For example airline 6 has on average 16 flights but it was observed that there were days that had 14 flights and other days up to 17. Airline 2, which represents the smaller carrier with only 2 flights, has small variance and the difference of it between the deterministic and stochastic cases is also small, because it was observed that the number of flights it has does not fluctuate with time.

In TABLE VIII we present the results for the fair share deviation indicator. Here also, the indicator for each airline is very close to zero, which means that the airlines on average will be getting numbers of slots that are close to their fair share.

TABLE VIII. DEVIATION OF THE AVERAGE FAIR SHARE FROM THE ACTUAL AVERAGE ALLOCATION

		Reduced capacity, ac/hr		
		25	**20**	**15**
Airline	**1**	-0.00093	-0.00008	-0.00181
	2	-0.00043	-0.00173	0.00245
	3	0.00009	0.00074	0.00152
	4	-0.00007	0.00009	-0.00116
	5	-0.00024	0.00018	-0.00111
	6	0.00043	0.00085	-0.00041
	7	0.00077	-0.00011	-0.00136
	8	-0.00011	0.00095	-0.00021
	9	0.00050	-0.00089	0.00209

TABLE IX. AVERAGE DELAY (IN MIN) PER FLIGHT PER AIRLINE

		Reduced capacity, ac/hr			
		25	**20**	**15**	**Total Flights**
Airline	**1**	12.2	17.1	21.2	3
	2	13.0	19.9	3.8	2
	3	10.0	15.6	19.9	6
	4	10.8	17.4	21.2	5
	5	8.2	13.3	18.9	10
	6	6.5	11.4	16.1	16
	7	10.6	16.1	20.9	6
	8	7.7	13.0	18.5	11
	9	11.6	18.2	19.6	4

TABLE X. WEIGHTED AVERAGE DELAY (IN MIN) PER FLIGHT PER PRIORITY NUMBER

		Reduced capacity, ac/hr		
		25	**20**	**15**
Priority Number	**4**	23.2	36.8	48.1
	3	24.8	44.0	64.4
	2	20.1	35.6	47.3
	1	11.6	17.1	18.0

Finally in tables IX and X we present the average delays per flight per airline and per priority number assigned, respectively. In the first of these two tables, we can see it is clear that when capacity is reduced the amount of delay accrued by each airline increases. With the exception of airline 2, which has the least slots (and when capacity is reduced much it often ends up being assigned zero slots) the differences in delays among the airlines are reasonable. In the other table we can see again the same trend as before. The weighted average delay for flights with priority 4 is consistently less than the delay of flights with priority 3. As the priority level decreases it is observed that the number of flights left without being assigned increases, which causes the reduced weighted delay compared to the higher priority flights.

V. Conlusions

From the above results and analysis we can see that during AFPs, if a preference structure and allocation mechanism such as we modeled is implemented, then airlines in the long run will be getting on average what they deserve. As we saw the smaller carriers have a good chance of actually getting slots in the constraint areas. For smaller airlines the variance of slots allocated tends to be smaller than the variances for the bigger carriers. It was also shown that for the flights with high priority numbers, most of them were assigned to slots and most of the ones with lower priorities were not. The weighted delays for the flights with priority 4 were less than the ones with priority 3.

References

[1] Federal Aviation Administration. FAA Aerospace Forecast, Fiscal Years 2015-2035. U.S. Department of Transportation, Federal Aviation Administration Policy and Plans.

[2] M. O. Ball and G. Lulli, “Ground delay programs: optimizing over the included flight set based on distance”, Air Traffic Control Quarterly, Vol. 12, 2004, pp 1-25.

[3] R. deNeaufville and A. Odoni, Airport Systems: Planning, Design, and Management. The McGraw Hill, New York, 2003.

[4] N. Vakili, “Preference based fair allocation of limited airspaces resources”. Ph.D. Thesis, University of Maryland, 2009.

[5] Metron Aviation official website. http://www.metronaviation.com/ . Last access January 2012.

[6] National Business Aviation Association official website. http://www.nbaa.org . Last access February 2016.

[7] CSC North American Public Sector, “Traffic Flow Management System (TFMS) Collaborative Trajectory Options Program (CTOP) Interface Control Document (ICD) for the Traffic Flow Management-Modernization (TFM-M) Program”, prepared for the Federal Aviation Administration, Contract number: DTFAWA-04-C-00045, 2013.

[8] N. Vakili and M. O. Ball, “Equitable allocation of enroute airspace resources”, 8th USA/Europe Air Traffic Management Research and Development Seminar, Napa, CA, 2009.

[9] K. Vlachou and D. J. Lovell, “Mechanisms for equitable resource allocation when airspace capacity is reduced”, Transportation Research Record 2325, 2013, pp 97-102.

[10] G. C. Carr, H. Erzberger and F. Neuman, “Airline arrival prioritazation in sequencing and scheduling”, 2nd USA/Europe Air Traffic Management Research and Development Seminar, Orlando, FL, 1998.

[11] Bureau of Transportation Statistics official website. www.bts.gov. Last access February 2012.