IMPROVING METROPLEX OPERATIONS EFFICIENCY USING SPEED SEGREGATION AND TRAJECTORY FLEXIBILITY

Husni Idris
Engility Corporation, Billerica, MA USA

Keywords: *Airport systems, metroplex, airport capacity, trajectory flexibility*

Abstract

The growth of air traffic demand is increasing the number of operations at major and secondary airports. As a result, the interdependencies between nearby airports are also increasing leading to the emergence of multi-airport systems (metroplexes). Often, these metroplexes constitute bottleneck capacities of the national air transportation network and hence a major cause of delay. The metroplex capacity limitations are caused by several inefficiencies. One such inefficiency is assigning the metroplex airspace among the competing airports based on procedures that segregate traffic by destination airport. These procedures limit the opportunity to share scare airspace resources dynamically between the airports and increase the mixing of slow and fast aircraft in single flows. This paper proposes assigning airspace and segregating traffic according to aircraft speed as opposed to the destination airport, resulting in sharing of airspace resources and potential throughput and flexibility gains. The proposed approach is analyzed, using simulation of hypothetical scenarios, in terms of airport throughput and aircraft trajectory flexibility.

1 Introduction

The Next Generation Air Transportation System (NextGen) is expected to receive up to three times the current traffic demand by the year 2025 [1]. In order to handle this traffic volume it is essential to increase the capacity at the bottlenecks of the air transportation network, which are typically the airports. One key capacity limitation at airport systems is the increasing interdependence between nearby airports. This interdependence leads to the emergence of multi-airport systems where one airport has to limit its operations is order to accommodate the operational needs of a neighboring airport. Therefore, the operation of these airports requires coordination between the traffic managers to maximize the overall throughput and ensure fairness. A typical example of a meroplex is the New York metropolitan airport system, which includes four major airports, John F. Kennedy (JFK), Laguardia (LGA), Newark (EWR), and Teterboro (TEB), within 20 miles of each other, in addition to a number of secondary surrounding airports.

The capacity limitations at a metroplex are caused by several inefficiencies often resulting from procedures that predate the high density environment and hence are not optimized for current and future high density operations. One such inefficiency is the assignment of airspace among competing airports based on destination airports, which is often not optimal from a throughput perspective. These procedures limit the opportunity to share scare airspace resources dynamically between the airports and increase the mixing of slow and fast aircraft in single flows. This paper proposes assigning airspace and segregating traffic according to aircraft speed as opposed to the destination airport, resulting in sharing of airspace resources among the airports and potential throughput and flexibility gains. The proposed approach is analyzed, using a simulation of hypothetical scenarios, in terms of airport throughput and trajectory flexibility.

First, in order to provide operational context, the specific types of metroplex interdependencies and inefficiencies that are addressed in this paper are presented in Section 2 along with the proposed solutions. Then, the analysis approach is presented in Section 3, including the simulation that was used for the analysis. Preliminary results and insights gained from the analysis are discussed in Section 4, concluding with final remarks and future extensions in Section 5.

2 Metroplex Inefficiencies and Proposed Solutions

The airspace in terminal radar control (TRACON) areas is assigned largely based on destination airport, resulting in segregating the arrival traffic flows by destination. This segregation extends often to the Air Route Traffic Control Centers (ARTCC)) such that the traffic in each en route sector consist mainly of flows destined to one or few destinations. This destination-based segregation helps the air navigation service providers in their largely manual process of merging and metering the flows based on the conditions at the destination. The destination airports are often the origin of the restrictions imposed on the traffic flows. However, the segregation by destination airport causes inefficiencies in high-density conditions, for example:

(1) Capacity is reduced due to the early mixing of aircraft with different speeds in a single destination airport/runway flow. Sequencing a slow aircraft behind a fast aircraft requires a large spacing between the two aircraft (due to wake vortex separation requirements). If the slow aircraft is merged early behind the fast aircraft at the minimum separation requirement, this separation opens up with time because of the speed difference. On the other hand, sequencing a fast aircraft behind a slow aircraft requires excessive spacing if the merging is established early to avoid closing the spacing to a value below the separation requirement. Once the separation requirement is established the faster trailing aircraft has to slow down to match the speed of the slower leading aircraft, and thus travel at a lower speed than it would otherwise.

(2) Aircraft often travel excessive travel distance due to the procedural separation of the routes used by the different airports, as explained in the example below.

(3) Capacity is often lost due to the procedural switching of delegation of shared airspace to flows destined to different airports. For example, an airport may lose the usage of certain runways when a shared airspace is delegated to another airport in the metroplex. This is also explained in the example below.

One example of cases (2) and (3) above is the interdependency between LGA and JFK, based on the standard operating procedures (SOP) as shown in Fig. 1. Specifically, when JFK is forced to perform instrument landing (ILS) on runway 13 left (13L), LGA is obligated procedurally to land only ILS runway 13.

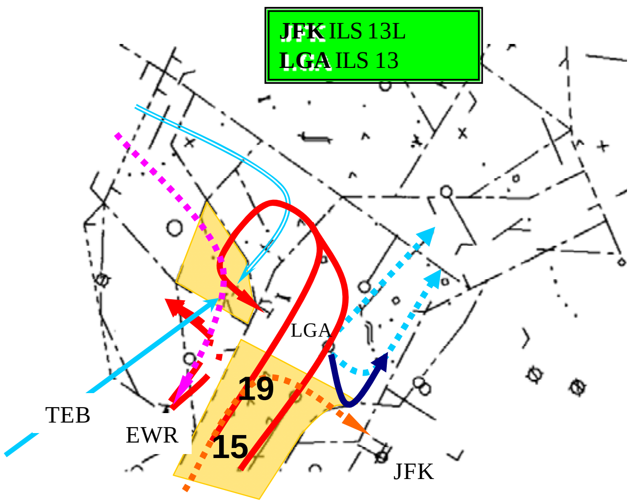

Fig. 1. Example of metroplex interdependency

In addition to limiting the available arrival runways at LGA to only one, this procedure often reduces LGA to single runway operations when they can only depart on runway 13. While this happens few times a year, it is known to be the most limiting situation at LGA. The reason for this procedure, as clarified in the LGA SOP, is sharing airspace areas 15 and 19. Typically, LGA owns area 15 from 10000 feet and below and area 19 from 12000 feet and below, except when JFK lands ILS runway 13L. Under this condition, LGA has to give up altitudes 4000 feet and below to JFK, to be used by the JFK arrivals. As shown in the diagram in Fig. 1, LGA flights in this condition have to fly higher above the JFK arrivals and perform an additional loop during which they also stay above the EWR and TEB traffic, then descend and approach runway 13. Therefore, LGA flights fly a longer distance at higher altitude, inefficiently.

Two potential methods to help mitigate these types of inefficiencies are:

(1) Reducing the effect of speed mixing by segregating the traffic by speed where possible and delaying the merge establishment as much as possible. Ideas of segregation by speed for a single airport were published in Idris and Simpson 1998 [2]. Generalizing this approach to a metroplex environment, traffic may be segregated by speed as opposed to, or in addition to, by the destination airport. This is shown in Fig. 2 notionally. For example, by using multiple parallel downwind and base leg segments the fast aircraft (larger triangles) are assigned on outer downwind/base legs relative to the slower aircraft (smaller triangles) which are assigned on inner downwind/base legs. Airports are shown as crosses in the figure.

(2) Allowing sharing of approach segments (such as downwind and base legs) by arrivals to neighboring airports. This sharing can be combined with speed segregation where, for example, aircraft of the same speed category but heading to different airports are assigned the same approach segments. An example is shown in Fig. 2 where the fast aircraft heading to two different airports (blue and red large triangles) share a downwind leg and diverge to different base legs heading to their respective airports.

The hypothesized benefits of these approaches include:

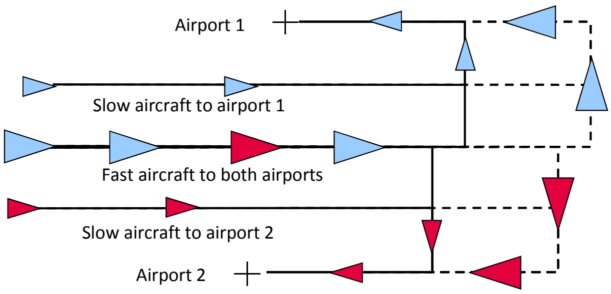

Fig. 2. Segregation by speed instead of by destination airport

(1) Increasing throughput by delaying the merging of slow and fast aircraft.

(2) Increasing throughput by enabling the use of runways otherwise unused due to interdependencies. For example, LGA could use runways other than 13 (such as runway 4) when JFK is landing ILS runway 13L, by sharing airspace areas 15 and 19 below 4000 feet with JFK. Flights for LGD and JFK would share approach segments in areas 15 and 19 and perform an exit to the respective airport/runway.

(3) Increasing throughput and reducing travel distance by sharing of airspace and routes currently delegated procedurally to different airports. In the example above, LGA flights landing on runway 13 could share airspace areas 19 and 15 below 4000 feet with the JFK arrivals to runway 13L. Then each will perform a late exit to the respective airport. If prioritization is needed, this sharing can be allowed for some of the LGA flights only when there is a lull in the JFK flow to runway 13L, thus reducing the longer and higher travel for at least some of the LGA flights.

It must be noted that such sharing methods require new controller procedures and possibly automation support. These requirements are not addressed in this paper, focusing mainly on making the benefit case.

3 Analysis Approach and Simulation

This paper presents preliminary proof of concept analysis for the traffic organization approach described in Section 2. The analysis compares the scenarios listed in Table 1:

Table 1. Analysis scenarios

Single airport	Airport has one approach path with mixed aircraft speeds (baseline)
	Airport has multiple approach paths, with different speed category per path
Two airports	Each airport has one approach path with mixed aircraft speeds (baseline)
	Airports share approach segments for aircraft with same speed

A single airport scenario is analyzed to demonstrate the effect of segregating the arrivals to the airport by speed. The airport in a baseline case has one approach path with aircraft of all speeds sharing the path. In another case the airport has multiple approach paths where aircraft of different speeds are segregated on different paths up to the merge on a final approach leg. A two-airport scenario is analyzed to demonstrate the effect of adding sharing of approach paths to the speed segregation. In the baseline case each airport has its independent approach path with all aircraft speeds assigned to it. In the comparison case the two airports share some approach segments for aircraft with the same speed. The analysis compares the scenarios in terms of the following metrics: (1) Throughput and (2) Trajectory flexibility.

A Matlab simulation was used for the analysis. It was developed for trajectory flexibility planning in the presence of controlled arrival time (CTA) and hazard/traffic avoidance constraints [3]. Trajectory flexibility is defined as the ability of a trajectory to accommodate disturbances while meeting constraints such as controlled time of arrival constraints and traffic/hazard avoidance constraints. Disturbances are events that pose risk of constraint violation, such as the uncertainty in the traffic/hazard dynamics. Relevant trajectory characteristics to measuring flexibility were identified; they included robustness and adaptability. Robustness is defined as the ability of a trajectory to remain feasible given disturbances, while adaptability is defined as the ability to regain feasibility if feasibility is lost due to disturbances. Formal mathematical definitions of these metrics and estimation techniques are given in Idris et al [3].

The tool implements a dynamic programming approach to estimate the trajectory flexibility metrics and to use them (along with other objectives) for trajectory planning. These methods are described in Idris et al [3]. Briefly, the method is based on discretizing space (2 dimensions in this paper) into square cells and time into steps. Then a solution space of all trajectories is built as a reachability tree connecting the resulting discrete nodes (each node represents the location of the center of a cell and time). The tree is based on reachability given discrete degrees of freedom, namely allowable speed and heading changes with discrete increments and within given ranges. At each node of the tree, adaptability is measured by the number of feasible trajectories that reach from that node to the destination (defined as a location with a time of arrival constraint). To estimate this number of feasible trajectories, a convolution process is used that starts at the destination and proceeds backwards, adding up feasible trajectories within the reachability bounds from each node. The convolution process is preceded at each time step by a filtering process that zeros out the number of trajectories at nodes that violate any constraints. The constraints include violation of separation from hazards and from other traffic or violation of time of arrival constraints at a node. After building the solution space tree with the adaptability metric at each node, a trajectory is computed using a dynamic program that optimizes an objective function. For more details see Idris et al [3]. The objective function used in this analysis maximized adaptability at each node along the trajectory, ignoring robustness.

To simulate a terminal/metroplex environment, the aircraft were forced to follow typical trombone approach patterns. A structured pattern was assumed in order to isolate the effect of the speed segregation only. Relaxing this structure introduces other interactions that are out of the scope of the concept of speed segregation and resource sharing among airports. For example, relaxing structure may introduce spatial interactions between aircraft because they would path stretch along the downwind as well as the base leg.

This may reduce throughput due to a factor other than speed mixing. Therefore, to isolate the speed factor and not deal with the effect of spatial interaction between trajectories, the aircraft were assumed to strictly follow the route structure (in this case the trombone approach).

Fig. 3 gives an example screen capture for the two airport scenario described in Fig. 2. Trombone approach patterns were forced by introducing polyhedral hazard constraints that blocked out the ability of the aircraft to path stretch along the downwind and runway centerline, allowing path stretching only in a base leg region. Polygons were also introduced to limit the extremities of the airspace to about 40 nautical miles (nmi) representing typical terminal airspace size, and the parallel downwind legs were 5 nmi apart. Each aircraft may be forced to a particular approach pattern by applying specific hazards to it. This enabled giving slow aircraft different patterns than fast aircraft in some experiments.

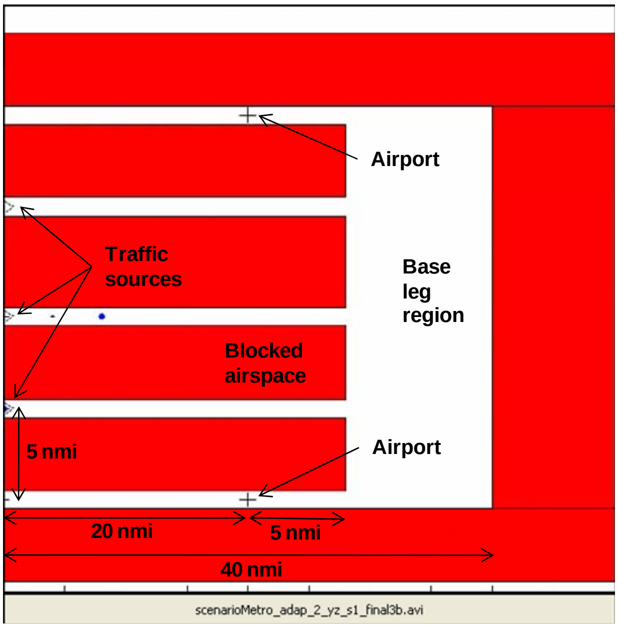

Fig. 3. Using blocked polygons to establish a trombone pattern

In the terminal area aircraft mostly only reduce their speed under strict control and in gradual steps. To model this behavior, the speed was forced to monotonically decrease along the approach and within limits representative of the speed step downs from the terminal entry speed towards the landing speed. Heading constraints were also added: If the airport is to the left of the downwind segment the aircraft are allowed to turn only counterclockwise. If the airport is to the right of the downwind segment, then the aircraft are allowed to turn only clockwise.

4 Simulation Results and Observations

The scenarios listed in Table 1 were analyzed in terms of throughput and trajectory flexibility. Preliminary insights are presented in this section, first for the single airport scenario followed by the two-airport scenario.

4.1 Single airport scenario

The impact of speed segregation on throughput and flexibility is first analyzed for a single airport case, using two aircraft speed categories: Fast aircraft with landing speed of 120 knots and slow aircraft with landing speed of 80 knots. Two cases are compared as shown in Fig. 4.

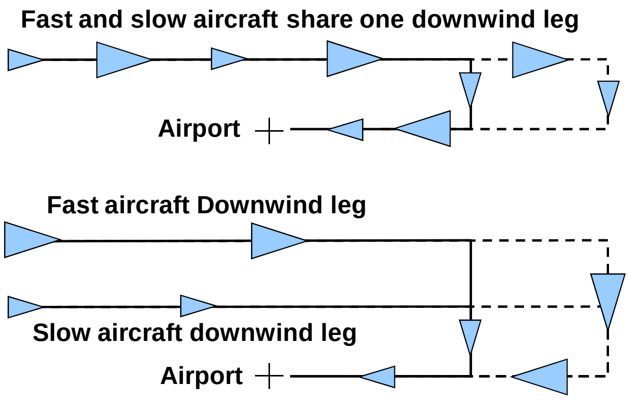

Fig. 4. Geometry of the single airport scenario

Case (a): All aircraft approach the runway using a single trombone approach path. A downwind leg is shared by both fast and slow aircraft. Each aircraft adjusts its base leg (by stretching the downwind leg), as needed, to meet its CTA at the runway while maintaining the separation requirements with all other aircraft. Case (b): Two downwind legs are used to segregate the fast from the slow aircraft. An outer downwind leg is used for the fast aircraft and an inner downwind for the slow aircraft. The two downwind legs are 5 nmi apart (larger than the minimum requirement of 3 nmi) such that no separation assurance is required between aircraft on the two legs. Similarly to case (a), each aircraft adjusts its base leg (by stretching the downwind leg), as needed, to meet its CTA at the runway while maintaining the separation requirements with all other aircraft. The following parameters were applied to both cases:

- The landing speed is met with a plus/minus 30 knot tolerance.
- Maximum demand is assumed available at all time in order to determine the maximum throughput. In other words all aircraft are available at the entry point at the simulation start and are introduced one by one as early as allowed by the separation constraints.
- Fast and slow aircraft are alternated in the demand stream to maximize the speed mix effect.
- To maximize throughput, the CTA at the runway is minimized using the following algorithm with one minute time resolution: For each aircraft i
    1. Assign CTA(i) = CTA(i-1) + 1 minute
    2. Increment the entry time of aircraft i by 1 minute until CTA(i) is met; if not met before the entry time is larger than CTA(i), then
    3. Increment CTA(i) by 1 minute and go to step 2.
- To model the procedure of stepping the aircraft speed down gradually by the air traffic controller, the speed limits are set to [max min] = [210 170] until CTA/2 then to [max min] = [170 landing-speed] from CTA/2 until landing. These limits ensure that the speed profile selected by the trajectory generation algorithm remains above 170 knots for half the duration and is reduced below 170 knots in the second half. The landing speed is 120 knots for fast or 80 knots for slow aircraft.
- The heading was limited to the downwind heading along the downwind leg and to the runway heading along the final approach. On the base leg the heading was unlimited as long as the turns are in the direction towards the runway.
- Given the CTA and the heading and speed limits, the speed and heading profiles are selected to maximize adaptability with speed increments of 10 knots and heading increments of 10 degrees.
- Separation requirement is set to 3 nmi, independent of the aircraft type due to simulation limitation.
- Path stretching is allowed only using trombone of the base leg within 40 nmi (see Fig. 3).
- Time is discretized using 2 minute time increments and space is discretized using 1x1 nautical mile square cells.

Throughput and trajectory flexibility are compared in Figures 5 and 6 respectively between cases (a) and (b). Fig. 5 shows the trajectories of the 10 aircraft in the scenario in terms of the X-location along the downwind direction versus time. X increases as the aircraft moves along the downwind, becomes flat or continues to increase slowly when the aircraft moves along a base leg (depending on the heading used) and then decreases as the aircraft moves along the runway centerline until landing. The slope of the curve represents the speed along the x-axis. The slope at the end of the curve shows the alternating of the fast and slow aircraft in the scenario.

As sown in Fig. 5, throughput was higher in case (b), segregating speed categories onto two trombones, than in case (a) which combined the two speed categories in one trombone. In this example the throughput was higher by about 20%. The 10 aircraft in the scenario landed in 1020 seconds in case (b) compared to 1320 in the baseline case (a). However, it should be noted that this is a theoretical benefit estimate relative to a baseline that is not validated against current operations. Therefore, the increase in throughput should not be interpreted as a potential increase relative to current operations, since current operations may include some degree of segregation by speed and delay of merging fast and slow aircraft, to the extent practiced by controllers.

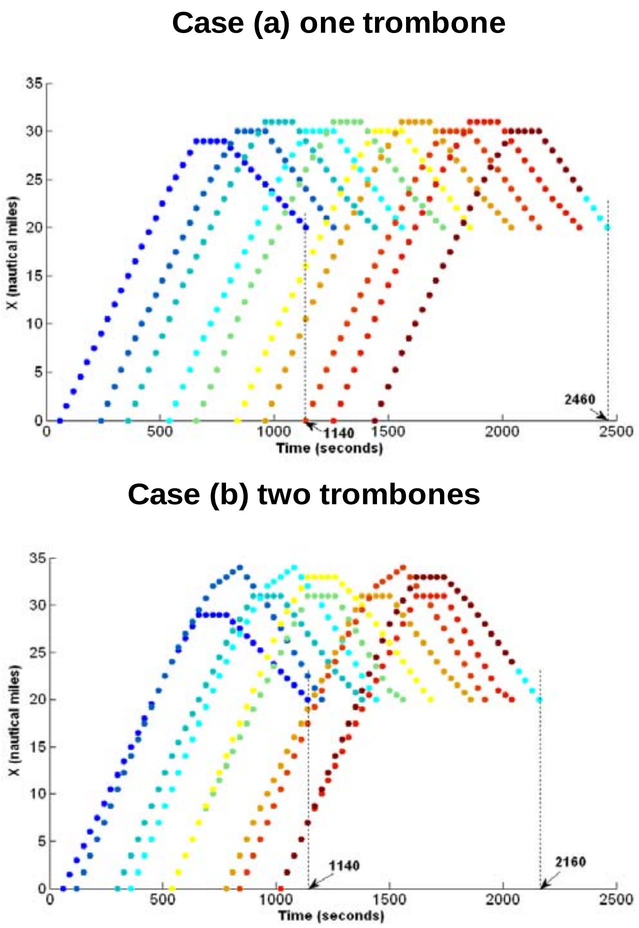

Fig. 5. Throughput analysis in single-airport scenario

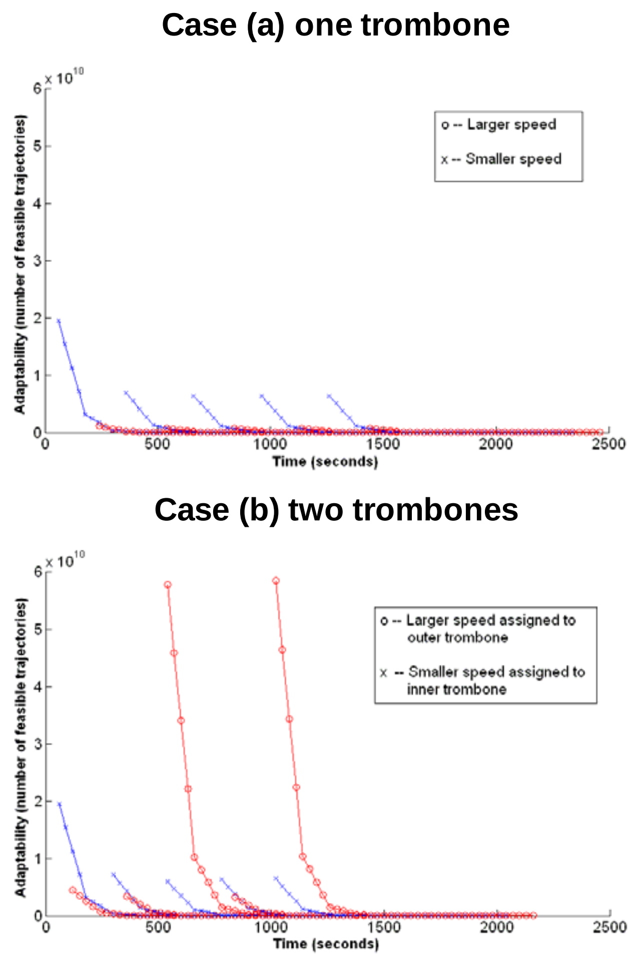

Fig. 6. Flexibility analysis in single-airport scenario

The trajectory flexibility analysis is shown in Fig. 6, which plots the adaptability along each trajectory over time. The following observations can be made from this analysis:

(1) It was observed that the faster aircraft have lower adaptability relative to slower aircraft for most aircraft in both cases (a) and (b), as shown in Fig. 6. This difference is mainly due to the smaller speed range that is available to the faster aircraft, between the maximum speed and the landing speed.

(2) It is clear in Fig. 6 that the first aircraft, which is a slow aircraft, has higher adaptability than the aircraft that followed it, because it is not impeded by any aircraft ahead of it.

(3) Signs were observed that adaptability of the fast aircraft increase in case (b) relative to case (a) of Fig. 6 when they are separated from the slower aircraft and placed on an outer trombone. This increase in adaptability is due to at least two factors: (a) providing more airspace for maneuvering to the aircraft that are placed on the outer trombone, and (b) providing more speed range for maneuverability to the faster aircraft which can stay at higher speed for a longer duration of time in case (b) relative to case (a). The increases speed maneuverability is evident from the speed profiles in Fig. 5, where in case (b) the faster aircraft exhibited faster slopes relative to case (a) and some fast aircraft passed the slower aircraft on the inner trombone. Two of the fast aircraft (third and fifth) managed to find a heading/speed profile that exhibited substantially large adaptability.

4.2 Two airport scenario

The scenario analyzed in this section includes two airports with speed segregation over multiple approach paths. Two cases are compared as shown in Fig. 7:

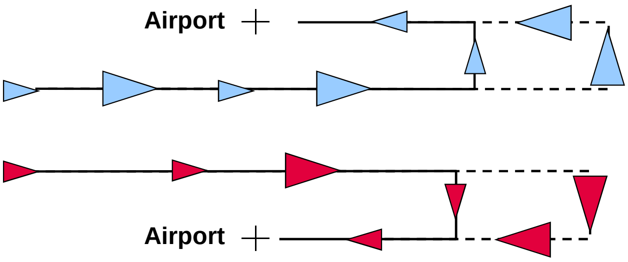

Case (a): Each airport has one independent trombone for fast and slow aircraft

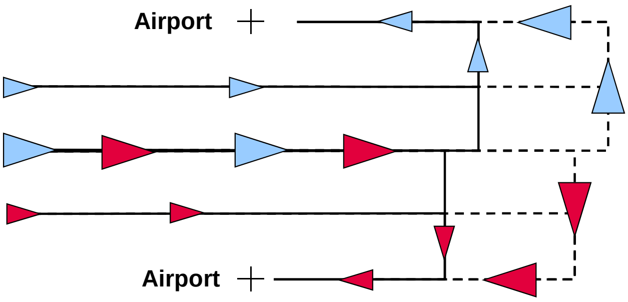

Case (b): Both airports share downwind leg for fast aircraft and each airport has separate downwind for slow aircraft

Fig. 7. Two airport scenario

In Case (a), each airport has a single downwind leg shared by all speed categories heading to the airport. In this case the two airports operate independently with separate approach paths. In Case (b) the two airports use three downwind legs: Two of which are dedicated to the slow traffic of each of the airports, with each airport using the leg closer to it. The third central downwind leg is shared by the fast traffic heading to either airport. The dimensions are given in Fig. 3.

The scenario includes the same parameters as in the single-airport scenario, where the aircraft speed categories are alternated in the demand stream for each airport, to maximize the speed mix effect. They are also alternated between airports to maintain equal loads on the two airports. However, because all aircraft are available to enter at the beginning of the simulation, unless separation requirements are violated aircraft heading to different airports may be introduced simultaneously. The CTA is minimized for each aircraft (to maximize throughput) using the same algorithm as in the single airport scenario. Time is discretized using 1 minute time increments and space is discretized using 1x1 nautical mile cells

Throughput and trajectory flexibility (adaptability) were compared between cases (a) and (b) in Figures 8 and 9 respectively.

As sown in Fig. 8, throughput was higher in case (b), segregating speed categories onto two trombones with a shared central downwind for the fast aircraft, than in case (a) which combined the two speed categories in one trombone for each airport. In this example the throughput was higher by about 13% for the first airport and by about 6% for the second airport. The 10 aircraft landed in 480 seconds less for airport 1 and in 240 seconds less for airport 2, in case (b) compared to case (a). The difference between the two airports is attributed mainly to the order in which the aircraft were introduced: The second airport lost some benefits because its first fast aircraft had to platoon along the shared centerline behind the first fast aircraft heading to the first airport. The gain for either airport is not as high as the 20 percent observed in the single airport scenario: which is the gain of both airports if each had two independent trombones for their slow and fast aircraft without sharing resources with the other airport. The reduction in the gain is because while speed segregation increased throughput, the central downwind leg was shared between the two airports reducing the benefits relative to a case where each airport has its own downwind leg for the fast aircraft as well. Despite sharing the fast downwind leg for the fast aircraft, the speed segregation resulted in a throughput benefit in the range between 6 and 13 percent.

Again, it should be noted that this benefit estimate is theoretical simulation that is not validated against current operations. Therefore, the increase in throughput should not be interpreted as a potential increase if such a procedure is adopted relative to current operations, since the current operations may include some degree of segregation by speed and airspace sharing, to the extent practiced by controllers.

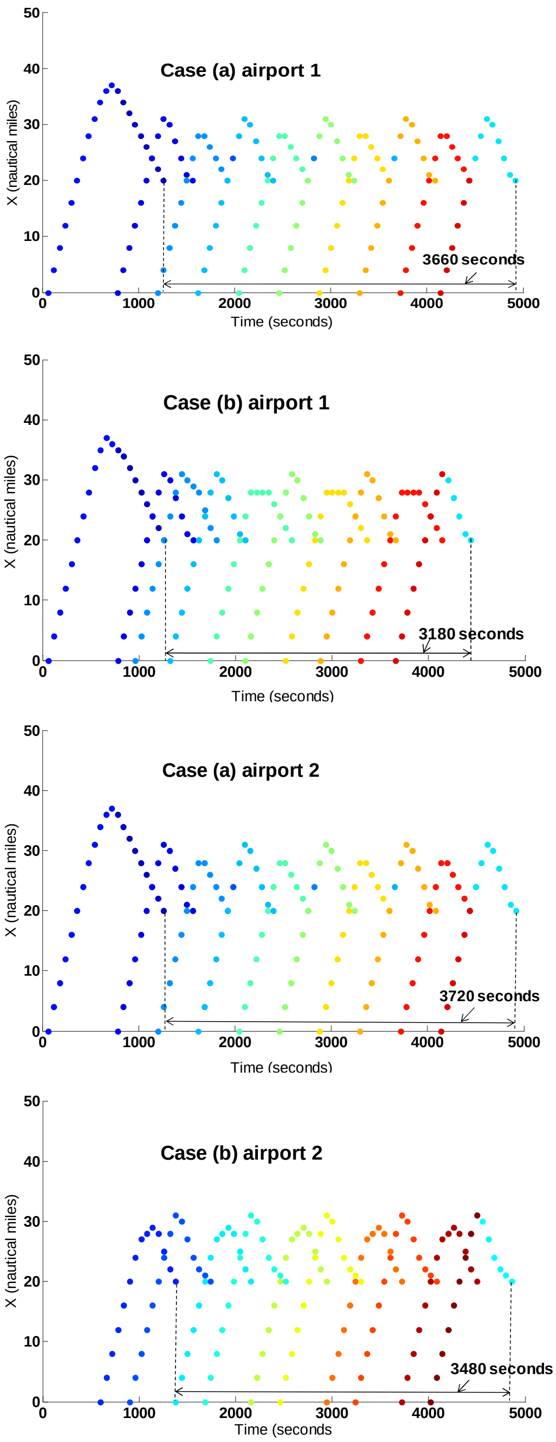

Fig. 8. Throughput analysis in two-airport scenario

It should also be noted that sharing the downwind segment enables other benefits that are not measured in this analysis. For example, enabling the airports to use additional runways or travel less distance, as described in Section 2.

In terms of trajectory flexibility, signs were also observed that adaptability of the fast aircraft increased (as shown in Fig. 9) in case (b) relative to case (a) when they are separated from the slower aircraft and placed on a separate downwind leg.

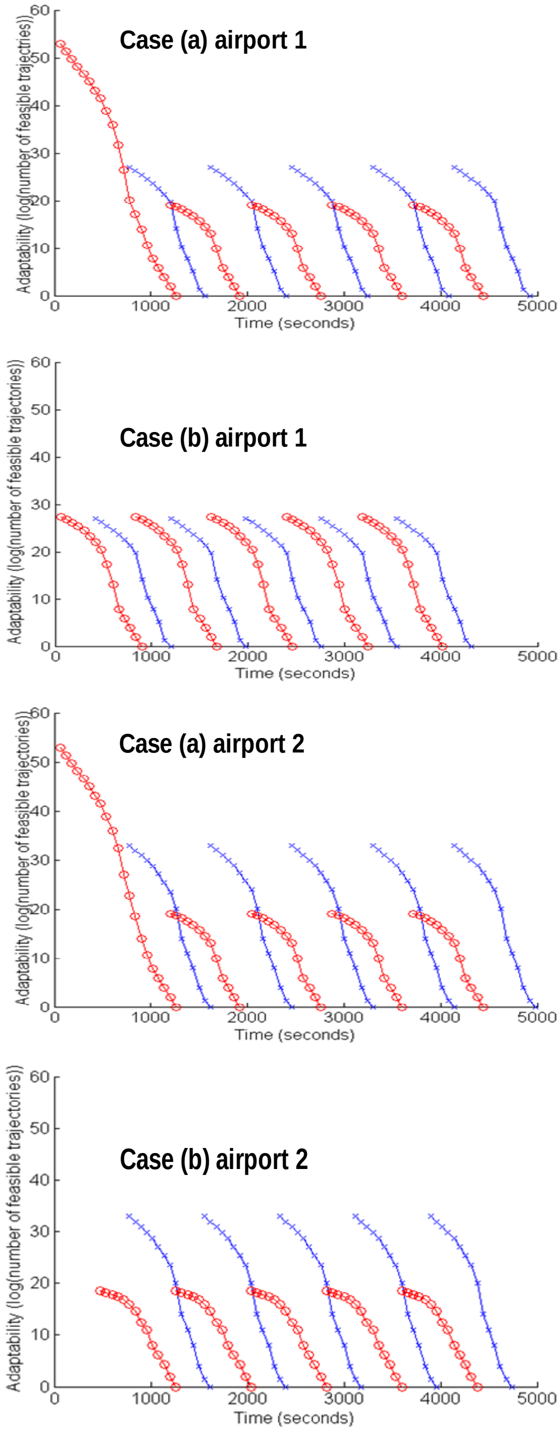

Fig. 9. Flexibility analysis in two-airport scenario

The log of adaptability (log of the number of feasible trajectories at each point) is plotted for better viewing because adaptability grows exponentially with negative time.

The increase in adaptability is due to at least two factors: (1) providing more airspace for maneuvering to the aircraft that are placed on the outer trombone for each airport. (2) providing more speed for maneuverability to the faster aircraft which can stay at higher speed longer in case (b) than in case (a) when they are mixed with the slower aircraft on single trombone. The gain in adaptability for the second airport is not as noticeable as for the first airport. One reason may be because of the order in which the aircraft are introduced. The fast aircraft of the second airport are introduced behind those of airport 1 on the shared downwind leg. This may have impeded the aircraft of airport 2 and resulted in less adaptability, in addition to less throughput, for airport 2.

5 Concluding Remarks

A preliminary analysis was described in this paper of a concept for organizing the arrival traffic to multiple airports in a metroplex based on speed segregation. Segregating the arrival traffic by speed may replace the current practice of segregating the traffic by destination airport or used in addition to it where feasible and beneficial. This concept involves two techniques: (1) Increasing the sharing of airspace and route resources among airports, particularly by aircraft of similar speed, thus opening up airspace and runway capacity and shortening travel distance and (2) segregating aircraft by speed, thus reducing the mixing between fast and slow aircraft in the same flow leading to further increase in capacity.

This preliminary analysis demonstrated encouraging signs of benefiting from the solutions proposed. The benefits were demonstrated in terms of increasing throughput and trajectory flexibility, which allows better mitigation of the risk of constraint violation. For a single-airport scenario the theoretical increase in throughput was up to 20 percent because of speed segregation. For a two-airport scenario the increase in throughput was less and ranged between 6 and 13 percent, because of sharing resources between the two airports in addition to the speed segregation.

The analysis presented in this paper is preliminary and further analysis of additional scenarios is needed. Future extensions of this research include identifying operational rather than theoretical benefits by using a baseline that represents current operations and the level of speed segregation and airspace sharing that is currently practiced. Future research may also investigate the added benefits of increasing runway usage and reducing travel distance because of sharing resources among airports more effectively than in current procedures.

6 References

[1] Joint Planning and Development Office. *Next Generation Air Transportation System Integrated Plan.* URL: http://www.jpdo.gov/library/NGATS_v1_1204r.pdf

[2] Idris H.R. and Simpson R.W. *New Approach to the Planning and Control of Air Traffic in the Terminal Area.* Presented at the American Institute of Aeronautics & Astronautics (AIAA) Guidance, Navigation and Control Conference, Boston, MA, August 10-12, 1998.

[3] Idris, H., D. Delahaye, and D. Wing. *Distributed Trajectory Flexibility Preservation for Traffic Complexity Mitigation.* Proceedings of the 8th USA/Europe Air Traffic Management R&D Seminar, Napa valley, California, 2009.

7 Contact Author Email Address

husni.idris@engilitycorp.com or husni@alum.mit.edu

Copyright Statement