JOURNAL OF AIRCRAFT
Vol. 36, No. 5, September–October 1999

Delay Exchanges in Arrival Sequencing and Scheduling

Gregory C. Carr* and Heinz Erzberger†
NASA Ames Research Center, Moffett Field, California 94035-1000
and
Frank Neuman‡
Raytheon STX Corporation, Moffett Field, California 94035-1000

Air traffic management (ATM) must often place restrictions on arriving flights transitioning from en route airspace to highly congested terminal airspace. The restriction of arrival traffic, or arrival flow management, is performed without regard for the relative priority that airlines may be placing on individual flights. The development of new arrival flow management techniques that consider priorities expressed by air carriers will reduce the economic impact of ATM restrictions and lead to increased airline economic efficiency by allowing airlines to have greater control over their individual arrival banks of aircraft. NASA is exploring the possibility of allowing airlines to express relative arrival priorities to ATM through the development of new sequencing and scheduling algorithms that take airline arrival preferences into account. This paper introduces the concept of "delay exchange," which is defined as a fair method of accommodating an airline request for an earlier arrival by advancing the landing time of one aircraft while simultaneously delaying another aircraft from the same airline. Fast-time simulation is used to evaluate the feasibility of scheduling these delay exchanges among individual arriving aircraft. Results show that the probability of successfully time advancing an aircraft is highest for an algorithm that allows delay exchange between aircraft arriving at any feeder fix. Results also show that the success of this algorithm varies with airport acceptance rate, indicating that the performance of this algorithm is a function of traffic density as well as the position of the aircraft within the traffic rush interval.

Introduction

THE continued growth of air traffic within the United States, combined with the use of "hub and spoke" operations by air carriers, has led to increased congestion and delays in the terminal airspace surrounding the nation's busier airports. The problem of congestion is exacerbated at hub airports, where air carriers schedule large numbers of flights to arrive/depart at the same time or within a very short time period. These arriving and departing groups of aircraft are commonly referred to as "banks," and the simultaneous arrival of several banks of aircraft can easily strain the capacity of an airport. The capacity of an airport at any given moment is defined by the airport acceptance rate (AAR), or the maximum number of aircraft that can land at the airport in 1 h. The AAR is set by air traffic control and is based on factors such as weather conditions, controller workload, available runways, departing traffic, and landing direction. To ensure that the safe capacity of the terminal area is not exceeded, air traffic management (ATM) often places restrictions on arriving flights transitioning from en route airspace to terminal airspace. The constraint of arrival traffic is commonly referred to as arrival flow management, and includes techniques such as metering, vectoring, and the imposition of miles-in-trail restrictions. Arrival flow management is performed without consideration for the relative priority that airlines may be placing on individual flights based on factors such as crew criticality, passenger connectivity, critical turnaround times, gate availability, on-time performance, fuel status, or runway preference. To air carriers, "hubbing" makes good economic and competitive sense.[1] At the same time, however, hubbing operations often strain the arrival capacity of an airport, precipitating delays that can directly impact the economic efficiency of an air carrier's flight operations.

Received 8 July 1998; revision received 22 January 1999; accepted for publication 22 January 1999.

*Research Engineer. Member AIAA.
†Senior Scientist. Fellow AIAA.
‡Research Engineer.

NASA and the Federal Aviation Administration (FAA) have designed and developed a suite of software decision support tools (DSTs) to improve the efficiency of high-density airspace.[2] Collectively known as the Center–TRACON Automation System (CTAS), operational evaluation of these DSTs has shown them to be effective in improving airport capacity and reducing delays while maintaining controller workload at a reasonable level.[3] One of these tools, the traffic management advisor (TMA), is currently being used at the Fort Worth Air Route Traffic Control Center (ZFW) to perform arrival flow management of traffic into the Dallas/Fort Worth airport (DFW). The TMA is a time-based planning tool that assists traffic management coordinators and en route air traffic controllers in efficiently balancing arrival demand with airport capacity.[3] The primary algorithm in the TMA is a real-time scheduler that generates efficient landing sequences and times for arrivals within about 200 n miles from touchdown.[4] This scheduler sequences aircraft so that they arrive in a first-come-first-served (FCFS) order.

While FCFS sequencing establishes a fair order based on estimated times of arrival (ETAs), it does not take into account individual airline priorities among incoming flights. The ability to specify the preferred arrival order within the user's own arrival bank is useful for maximizing bank integrity and minimizing bank time; i.e., exchange of passengers/cargo and aircraft servicing.[5] The development of new arrival flow management techniques that consider priorities expressed by air carriers will likely reduce the economic impact of ATM restrictions on the airlines. This leads to increased airline economic efficiency by allowing airlines to have greater control over their individual arrival banks of aircraft. As part of its collaborative arrival planning research and development program, NASA Ames Research Center is exploring the possibility of allowing airlines to express relative arrival priorities to ATM through the development of new sequencing and scheduling algorithms that take into account airline arrival preferences.[6]

Delay Exchanges

When an aircraft enters an Air Route Traffic Control Center, or "Center" airspace, on its way to a hub airport, it may be sufficiently late so that some of its passengers would miss a connecting flight. However, the magnitude of the delay might be such that if aircraft could make up some time while in Center airspace, its passengers would not miss their connecting flight. Under these circumstances, the airline may request an earlier arrival for this aircraft, and ATM may accommodate this request by dictating speed and/or heading changes to the aircraft, which will lead to its arriving earlier at the "feeder fix," or entry point to the terminal radar approach control (TRACON) airspace. In terms of sequencing and scheduling, this time-advanced (TA) aircraft would have changed positions within the sequence of aircraft arriving at the same feeder fix, and would have an earlier scheduled time of arrival at the feeder fix (STAff). If arrival traffic is light, there may be large enough "gaps" in traffic that the TA aircraft could be fit into without impacting other arrival traffic. However, this putative case is of lesser importance, because for light traffic conditions delays would be small, and it is unlikely that ATM would be using sequencing and scheduling to perform time-based metering of arrival traffic. Of greater interest is the case in which time advance is attempted during heavy traffic such as that experienced during a rush period, when delays are large and ATM is using time-based metering tools to perform sequencing and scheduling of arrival traffic. In this case, the reordering of the TA aircraft from the FCFS or "natural" order of arrival would almost certainly lead to increased delays for aircraft belonging to other airlines as well as other aircraft belonging to the same airline. The resulting additional delays would be considered unfair by the other airlines, making this kind of time advance scheme unacceptable.

In the interest of fairness as well as scheduling efficiency, a method of scheduling that time advances an aircraft without adversely impacting other air carriers must be defined. Given that airlines view their own flights in terms of relative priority, i.e., one flight may be more important to the airline than another flight, it is assumed that an airline may be willing to accept additional delay for one flight in return for earlier arrival of another flight. The concept of a "delay exchange" is defined as a fair method of accommodating an airline request for an earlier arrival by advancing the landing time of one aircraft while simultaneously delaying another aircraft from the *same* airline. For purposes of this study, it is assumed that any scheduled Center delay of the TA aircraft can be reduced to zero, and that an additional 5 min of time advance can be gained in the Center airspace if a speed change is initiated at least 45 min from the feeder fix. Any time advance gained by the TA aircraft must be offset by an equivalent amount of additional delay for another aircraft. In this manner, delay exchange will not cause additional Center delays for aircraft from other airlines.

When two aircraft that belong to different weight classes trade delays, the TRACON schedule will be altered in some manner because scheduling in the TRACON is contingent on required runway separations that are a function of aircraft weight class. Because the focus of this study is to determine the feasibility of performing delay exchanges in Center airspace in terms of how often a delay exchange can be performed successfully and how much time advance can be gained through a successful delay exchange, the impact of delay exchange on the TRACON schedule is not modeled. However, the potential impact on TRACON scheduling is minimized by restricting delay exchanges to jet aircraft only, and by requiring that the scheduled aircraft arrival times at the feeder fixes be preserved when performing a delay exchange. In other words, the scheduled times of arrival for all aircraft in the arrival rush except the delay exchange pair remain unchanged, and the scheduled times of arrival for the pair of aircraft involved in the delay exchange are simply interchanged. Thus, AAR is maintained during a delay exchange, and the TRACON schedule is affected only by differences in required separations at the runway that may result from the reordering of the delay exchange aircraft.

Three different delay exchange algorithms are examined in this study:

1) The first algorithm performs delay exchanges between aircraft belonging to the same stream class.

2) The second algorithm performs delay exchanges between aircraft belonging to different stream classes.

3) The third algorithm performs delay exchanges between aircraft belonging to any stream class.

Stream class is a CTAS-specific term that refers to a grouping of arrival aircraft based on engine type and feeder fix for the purposes of scheduling. Each feeder fix will typically have two stream classes: one for jet aircraft and one for combined turboprop and propeller (piston-engined) aircraft. The two stream classes at a feeder fix are vertically separated to allow the scheduling of simultaneous arrivals of jet and turboprop (or propeller) aircraft at the feeder fix.

Delay Exchange Algorithms

Delay Exchange Between Aircraft in the Same Stream Class

The first algorithm performs delay exchange between two aircraft that belong to the same stream class. In this scenario, the advancing aircraft and the aircraft with which it exchanges delay are scheduled to arrive at the same feeder fix. Figure 1 is a graphical representation of this scenario. Aircraft estimated time of arrival at the feeder fix (ETAff) and STAff are depicted along vertical timelines. ETAff represents the earliest possible time of arrival for an aircraft, provided that the aircraft could fly to the feeder fix with no delay; i.e., this represents the earliest time at which an aircraft could arrive at the feeder fix if there were no other aircraft in the airspace. STAff is the aircraft scheduled time of arrival at the feeder fix as determined by the sequencing and scheduling algorithm, and includes delays necessitated by constraints such as airport acceptance rate, runway loading, and separation criteria at both the runways and the feeder fixes (sequencing and scheduling is described in detail in Ref. 4). Time is increasing upward along each timeline, i.e., earlier arrivals are toward the bottom of the timelines, and aircraft are labeled in order of ETAff. Aircraft number 4 is the aircraft to be advanced by delay exchange; as shown on the timeline depicting STAff with time advance, its maximum possible time advance is calculated by first reducing its scheduled delay to zero (4a), and then further reducing its STAff by 5 min (4b) through speedup. This point (4b) represents the earliest possible arrival time or the minimum STAff for aircraft 4. However, because there is no aircraft scheduled precisely at this time with which to exchange delays, the algorithm searches for the first aircraft (from the same airline) having a STAff greater than a minimum STAff that aircraft 4 could meet (4b). Although aircraft 1 has the minimum STAff that is greater than the minimum STAff aircraft 4 could meet, aircraft 1 does not belong to the same airline as aircraft 4, and as such it cannot be considered for delay exchange.

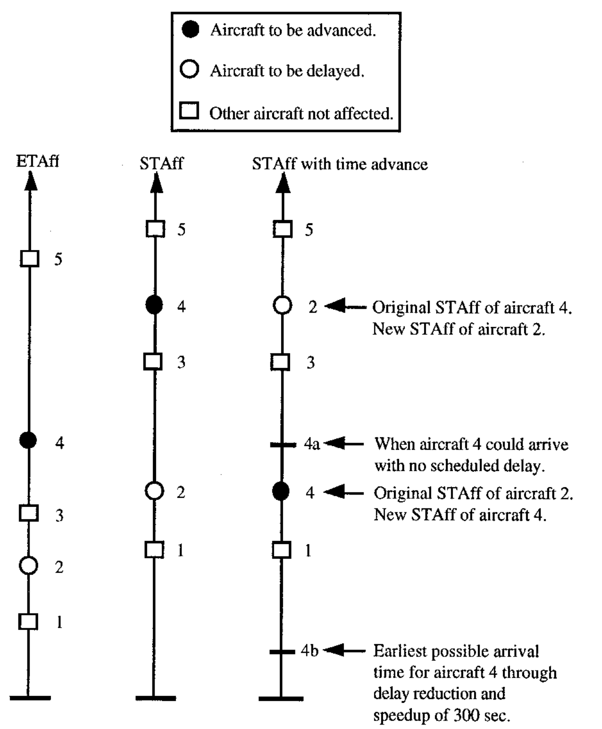

Fig. 1 Delay exchange between aircraft in the same stream class.

Because aircraft 2 belongs to the same airline as aircraft 4, its delay can be exchanged with that of aircraft 4. Now aircraft 2 will be delayed by an amount such that its STAff is equal to the original STAff of aircraft 4, and aircraft 4 will be assigned an STAff equal to the original STAff of aircraft 2. If there were no aircraft from the same airline having STAffs between the original STAff of aircraft 4 (4) and the minimum STAff that aircraft 4 could achieve through delay reduction and speedup (4b), then the algorithm would produce no time advance for aircraft 4.

From this example it is seen that the amount of time advance gained through delay exchange in the same stream class will range from a minimum value of zero to a maximum value equal to 300 s plus the advancing aircraft's original scheduled delay. It is also seen that the success of the delay exchange algorithm will vary with the traffic mixture and traffic density at each feeder fix. Because the amount of delay exchange is directly proportional to the originally scheduled delay for each aircraft, and because delays increase with increasing traffic density, the amount of time advance that can be gained through delay exchange should be greatest when the traffic is most dense.

Delay Exchange Between Aircraft in Different Stream Classes

The second algorithm performs delay exchanges between two aircraft that belong to different stream classes. In this scenario, the advancing aircraft is arriving at one feeder fix, while additional delay is assigned to an aircraft arriving at a different feeder fix. A detailed description of how this algorithm works is given later in this section. To provide a better understanding of this algorithm, the constraints for successful delay exchange between aircraft in different stream classes are summarized next:

1) To advance an aircraft, there must, in the same stream class, be a STAff gap between aircraft of at least 120 s for the TA aircraft to fit into. (For purposes of this simulation, the 5-mile separation criterion in Center airspace is translated into a time-based separation requirement of 60 s or 1 min. Therefore, a gap in traffic of at least 120 s is required to enable the advancing aircraft to meet this separation requirement.)

2) The amount of time advance must be such that the TA aircraft can be positioned in the STAff gap in its own stream class with at least 60 s to the aircraft ahead and 60 s to the aircraft behind.

3) A gap in STAffs of at least 120 s must exist in one of the three other stream classes at the original STAff of the TA aircraft.

4) The amount of delay must be such that the delayed aircraft can be positioned in the STAff gap in its own stream class with at least 60 s to the aircraft ahead and 60 s to the aircraft behind.

5) The aircraft to be delayed must have a STAff equal to, or earlier than, the new STAff of the TA aircraft. In this manner the aircraft will be delayed by an amount greater than or equal to the amount of time advance. Currently, a range of 1 min is allowed from the new STAff of the TA aircraft to an earlier STAff, so that on the average the additional delay is slightly larger than the amount of time advance.

Figure 2 is a graphical representation of delay exchange between two aircraft in different stream classes. Aircraft STAffs for each of the four stream class are depicted along vertical timelines, with time increasing upward along each timeline; i.e., earlier arrivals are toward the bottom of the timelines. In this example, an aircraft in stream class 2 meets all of the conditions required to perform a delay exchange with the TA aircraft in stream class 1.

As in the previous algorithm, it is assumed that the aircraft to be advanced can reduce its scheduled delay to zero, and can gain an additional 5 min by speeding up if the speed changes are initiated 45 min from the feeder fix. This represents the minimum STAff that the advancing aircraft could achieve. After calculating the minimum STAff for the TA aircraft, the algorithm checks the same stream class to determine whether there is a gap in traffic at the minimum STAff into which the TA aircraft could fit. To meet the Center separation requirement, the advancing aircraft must be placed at least 60 s ahead of one aircraft, and at least 60 s behind another aircraft. If a gap in the same stream class does exist that is at least 120 s in length, the amount of time advance must be such that it reaches at least 60 s into the gap; thus meeting the minimum spacing criterion for the aircraft that will be trailing the TA aircraft. Similarly, if advance places the TA aircraft within a gap of sufficient size, but the new STAff falls within 60 s of the lead aircraft in the gap, the amount of time advance will be reduced to meet the separation criterion. If there is no gap of sufficient size at the minimum STAff of the TA aircraft, the algorithm will search for the next later gap in traffic into which the TA aircraft could fit. If there are no gaps of sufficient length in the same stream class as the TA aircraft, the algorithm will produce no time advance.

If a gap meeting the constraints is found in the same stream class as that of the TA aircraft, the algorithm will then search the other three stream classes for a gap in traffic at the original STAff of the TA aircraft. This gap will be used to accommodate an aircraft that

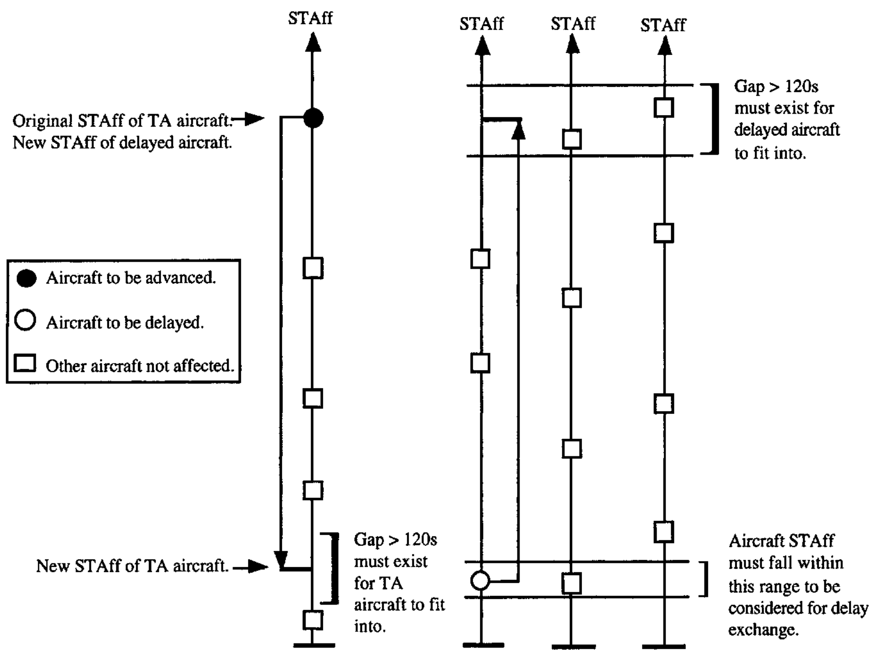

Fig. 2 Delay exchange between aircraft in different stream classes.

will be delayed by an amount greater than or equal to the reduction in delay of the TA aircraft. It is imperative that this gap fall at the original STAff of the TA aircraft so that the AAR will be maintained. As before, the gap must be at least 120 s in length to meet the minimum separation criterion. If no gap of sufficient size is found at the original STAff of the TA aircraft, the algorithm will produce no time advance. If a gap of sufficient size is found at the original STAff of the TA aircraft, the algorithm will search the stream class for an aircraft belonging to the same airline as the TA aircraft, which, if delayed to fit into the gap, would absorb at least as much additional delay as the amount of time advance. To meet this requirement, the aircraft to be delayed must have a STAff equal to or earlier than the new STAff of the TA aircraft. If no such aircraft is found that meets these requirements, the algorithm will produce no time advance. If multiple aircraft are found that meet these requirements, the algorithm will select the aircraft that yields the greatest amount of time advance.

Delay Exchange Between Aircraft in Any Stream Class

The third algorithm that is evaluated in this study is actually the combined application of the previous two algorithms. In this scenario, an advancing aircraft can trade delay with an aircraft arriving at the same feeder fix or at a different feeder fix; i.e., any stream class. For the aircraft that is to be TA, the amount of time advance that can be gained by delay exchange with another aircraft from the same stream class, and the amount of time advance that can be gained by delay exchange with an aircraft from a different stream class are determined. The time-advance values are then compared, and the greater of the two values is taken to be the resultant time advance. By applying the two delay exchange algorithms in concert, the likelihood of a successful delay exchange should be increased over that for either one of the two previous algorithms alone.

Fast-Time Simulation

A fast-time simulation originally developed for statistical evaluation of CTAS sequencing and scheduling algorithms has been modified for use in this investigation. In contrast to real-time simulation or field tests, which would require on the order of 90 min to examine a single traffic rush period, the fast-time simulation allows examination of multiple rush periods in a matter of minutes. Fast-time simulation permits examining the outcome of each delay exchange algorithm for many traffic rush periods having the same statistical parameters. For each of the three delay exchange algorithms, fast-time simulation determines the number of times an aircraft can be TA, as well as the amount of time advance gained through delay exchange. It should be noted here that because this simulation does not provide any information about the controller workload required to meet the calculated schedule or to perform a delay exchange, the output of the simulation is used only to determine the feasibility of scheduling delay exchanges. The fast-time simulation is composed of three major components: 1) an airport model, 2) a statistical model of the arrival traffic, and 3) a scheduler model.

Airport Model

The arrival airspace at DFW is divided into Center and TRACON regions, with the TRACON encompassing the airspace within ~40 n miles of the airport. Arrival traffic is merged at four waypoints on the Center–TRACON boundary, referred to as "feeder fixes." These feeder fixes correspond to the four primary arrival directions, and during heavy traffic periods traffic is funneled through these fixes as a means of controlling or metering the flow rate into the terminal area.[4] Traffic flowing to each feeder fix is separated into two independent streams ("stream classes" in CTAS jargon), which are vertically separated by 2000 ft at the fix. This allows jet and turboprop aircraft, which have significantly different airspeed ranges, to cross the feeder fixes independently and avoid conflicts due to overtakes.

The airport is modeled according to the landing practices at DFW, with four feeder fixes and three runways available for landing. The runways are considered to be independent so that no stagger requirements are necessary for scheduling. The airport model is composed of the minimum flight times from each feeder fix to all landing runways. These TRACON transition times were obtained from an analysis using the minimum flight times measured for several traffic samples.[7] The TRACON transition times vary with feeder fix, aircraft type, runway assignment, and airport configuration. The airport model contains transition times for both airport configurations at DFW: "north flow" with arrival traffic arriving/departing in a northerly direction, and "south flow" with traffic arriving/departing in a southerly direction. It should be noted that because the data used in this simulation were collected, a fourth arrival runway has been added at DFW. However, the three-runway model and traffic data are sufficient for purposes of this investigation.

Traffic Model

The traffic model is based on actual traffic data recorded during six rush periods at DFW. Although the traffic data were recorded over a span of several months, the mix of aircraft type remained nearly constant for each traffic sample. The data were recorded during the "noon balloon," a daily arrival rush lasting ~90 min. The noon balloon was chosen as the basis for the traffic model because during this arrival rush demand exceeds airport capacity and air traffic managers impose time-based metering restrictions through CTAS sequencing and scheduling algorithms. Data recorded during the six rush periods include the aircraft type, aircraft identification stream class (aircraft engine type and assigned feeder fix together define a stream class), and the ETAff as calculated by CTAS.

The traffic model itself is composed of a list of the average of the ETAs at the feeder fix for each aircraft during the six recorded rush periods. This average value is taken as the nominal ETAff for each aircraft in the arrival rush period. Statistical variations in ETAff between traffic samples are generated by adding an approximately Gaussian distribution to the nominal ETAff. The maximum range of the variation in the ETAff is specified as an input to the simulation. Whether an aircraft appears in a given traffic sample depends on the number of times the aircraft appeared in the six recorded rush periods. For example, an aircraft that appeared in two of the six recorded traffic periods would have a probability of $\frac{2}{6}$ of appearing in a simulated traffic sample. In this manner, the number of aircraft modeled in the arrival rush is also varied between traffic samples.

Scheduler Model

The scheduler model is intended to approximate the sequencing and scheduling algorithms presently used in CTAS at the Fort Worth Center. A detailed description of the actual scheduling algorithm can be found in Ref. 4. Aircraft are sequenced and scheduled to be FCFS at both the feeder fixes and runways while meeting feeder fix and runway threshold separation constraints. Because scheduling is done in time rather than distance, the prescribed minimum separation criteria are translated into minimum time separations at both the feeder fixes and the runway thresholds. For aircraft in the Center, the minimum in-trail separation requirement for aircraft is 5 n miles, which is translated to a 60-s time separation for purposes of this simulation. The separation criteria at the runway threshold are a function of both aircraft weight class and landing order as determined by the FAA's wake vortex safety rules. Airport acceptance rate is taken into consideration by limiting the number of aircraft that are allowed to land in sliding 10-min intervals, and the scheduler balances flights among runways to minimize overall delay.

The minimum ETA at the runway for each aircraft is calculated by adding the feeder fix-to-runway threshold crossing time for each of the three runways to the ETAff, and selecting the minimum of the three resulting values. This ETA at the runway represents the earliest possible time of arrival for an aircraft, provided that the aircraft could fly to the runway with no delay. The FCFS sequence is established by time-ordering aircraft according to increasing ETA at the runway. Beginning with the first aircraft in the sequence, each aircraft is tentatively scheduled to each of the three runways, while ensuring that the prescribed minimum time separation between aircraft at the runway thresholds is met for subsequent aircraft. The runway that results in the earliest time of arrival for the aircraft is chosen as the landing runway. Scheduling to the runway automatically provides the correct amount of traffic to load the runways equally when traffic is heavy (runway balancing), and directs aircraft to the closest available runway. The STAff is determined by subtracting the sum of the TRACON transition time and any TRACON delay from the previously calculated scheduled time of arrival at the runway. Finally, if the scheduled times of arrival at the feeder fix for two flights are less than the required 60 s apart, the times will be altered to meet the required separation at the feeder fix.

Simulation Inputs/Outputs

Inputs to the fast-time simulation include the number of traffic samples to be generated, the range in the random variation of the ETAffs generated by the traffic model, airport configuration, and AAR. To determine the statistical performance of the algorithms, 500 traffic samples are generated for each of the 77 jet aircraft contained in the noon-balloon model, resulting in a total of 38,500 traffic samples for each algorithm. With one aircraft advanced per traffic sample, 38,500 also represents the total number of times each algorithm attempts time advance. The AAR is set to be 96 aircraft/h, or 32 aircraft per runway per hour on the first run of the simulation; it is increased to 108 aircraft/h, or 36 aircraft per runway per hour for the second run. For both runs or AARs, the airport configuration modeled is south flow for DFW, and the range in the random variation of the ETAffs is set to ±250 s.

The output of the fast-time simulation includes the number of times each delay exchange algorithm results in a nonzero time advance, frequency distributions of the amount of time advance, average time advance, and the maximum individual and maximum average time advance. Because of the large number of trials for each algorithm, results can be interpreted as empirically determined probabilities. For the sake of simplicity, most data are normalized and measures of success are presented as a percentage of the total number of TA aircraft.

Results and Discussion

One measure of the success of the delay exchange algorithms is how often each of the algorithms results in a nonzero time advance. Figure 3 shows the number of times each algorithm yields a nonzero time advance for two different AARs. The success rate is given as a percentage of the total number of TA aircraft, and as expected, shows that for both AARs, the algorithm that allows delay exchange with any stream class has a higher success rate than either the algorithm that only allows delay exchange with the same stream class or the algorithm that only allows delay exchange with a different stream class. The any-stream class (ASC) algorithm can be thought of as a combination of the same-stream class (SSC) and different-stream class (DSC) algorithms. As such, if either the ASC or DSC algorithms yields a nonzero time advance, this nonzero value will be used for the delay exchange by the ASC algorithm. For an AAR of 96, the ASC algorithm yields a nonzero time advance 54% of the time; the SSC algorithm yields a nonzero time advance 43% of the time, and the DSC algorithm results in a nonzero time advance 18% of the time. The success rate of the ASC algorithm in producing a nonzero time advance is ~26% greater than that of the SSC algorithm, and it has three times the success rate of the DSC algorithm. Similar results are seen for AAR = 108, with the ASC algorithm yielding a nonzero time advance 52% of the time, and the SSC and DSC algorithms yielding a nonzero time advance 41 and 18% of the time, respectively.

For both AARs, the DSC algorithm has a lower rate of success than the SSC algorithm. This is because the conditions for performing a delay exchange with an aircraft from a different stream class are more restrictive than those for performing a delay exchange within the same stream class. It is interesting to note that the percentage of nonzero, TA aircraft for the delay exchange within the same stream class decreased from ~43 to ~41% with the increase in AAR. Because a higher AAR results in decreased traffic density with aircraft STAffs packed less tightly together, the probability of successful delay exchange within the same stream class is actually decreased when the AAR is increased.

Of equal importance to how often a nonzero time advance results from one of the delay exchange algorithms is how much time advance results from each of the algorithms. Figure 4 shows the cumulative frequency distribution of the amount of time advance gained through delay exchange for an AAR of 96. Each curve is the cumulative frequency for which the amount of time advance per aircraft is greater than or equal to the value given on the abscissa. (Note that the x axis starts at a time advance value of 60 s. The zero point is excluded for clarity because 100% of the TA aircraft have a time advance of 0 or more.) Again, it is seen that the ASC algorithm performs better than both the SSC and DSC algorithms. For the ASC algorithm, the amount of time advance gained through delay exchange is greater than or equal to 60 s ~54% of the time. For the SSC algorithm, the TA aircraft is advanced 60 s or more ~43% of the time; whereas for the DSC algorithm, the aircraft is advanced 60 s or more 18% of the time. Because it is assumed that each aircraft can gain a maximum time advance equal to its scheduled delay plus 300 s through speedup, it is of interest to determine how often the resultant time advance for the TA aircraft is 300 s or more. Figure 4 shows that for the ASC algorithm, ~25% of the TA aircraft receive 300 s or more time advance through delay exchange. For the SSC algorithm, 11% of the TA aircraft receive 300 s or more time advance; and for the DSC algorithm, 15% of the TA aircraft receive 300 s or more time advance.

For an AAR of 108, the trends in cumulative frequency distribution of the time advance for each of the algorithms are similar to those for the lower AAR. Table 1 lists several important points from the time-advance cumulative frequency distribution for both AARs. Table 1 shows again that the performance of the DSC algorithm remains unaffected by a change in AAR. The DSC algorithm is highly constrained, requiring not only candidate aircraft in other stream classes to exchange delays, but also requiring gaps in traffic in both the stream class of the TA aircraft and the stream class of the delayed aircraft. While a less-restrictive AAR may increase the number of gaps in traffic because the aircraft will be less closely packed, the decreased traffic density also reduces the number of

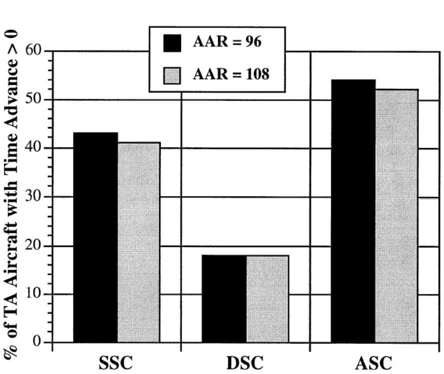

Fig. 3 Success rate for each algorithm.

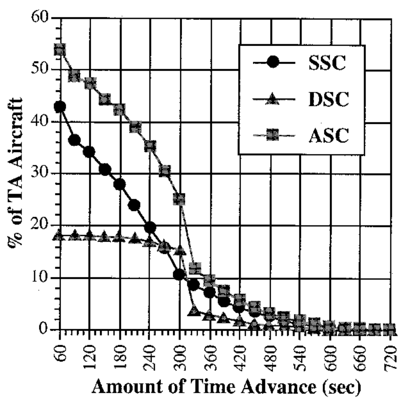

Fig. 4 Time advance cumulative frequency distribution (AAR = 96).

Table 1 Time advance cumulative frequency data

Algorithm	Percent of TA aircraft with time advance ≥60 s		Percent of TA aircraft with time advance ≥300 s	
	AAR = 96	AAR = 108	AAR = 96	AAR = 108
SSC	43	41	11	5
DSC	18	18	15	15
ASC	54	52	25	19

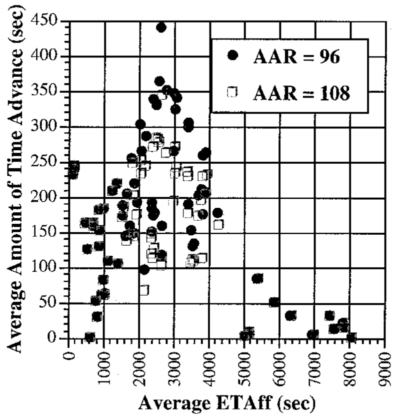

Fig. 5 Average time advance for the ASC algorithm.

candidate aircraft that meet the constraints to be considered for delay exchange. As a result of these two competing factors, a change in the AAR has a negligible impact on the performance of the DSC algorithm. On the other hand, the increased AAR causes a measurable decrease in the performance of the SSC algorithm and an attendant decrease in the performance of the ASC algorithm. For the SSC algorithm, the percentage of TA aircraft receiving 300 or more seconds of time advance decreases from 11 to 5% when the AAR is increased from 96 to 108. For the ASC algorithm, the percentage of aircraft that are TA more than 300 s decreases from 25 to 19% when the AAR is increased from 96 to 108. Here again, the decrease in performance with an increase in AAR is because the less-restrictive AAR leads to decreased traffic density and smaller scheduled delays, reducing the frequency and amount of time advance that can be gained.

It is instructive to examine the variation in average time advance with traffic density. Because the ASC algorithm works by choosing the greater of the two time advance values determined by the SSC and DSC algorithms, the average time advance determined by the ASC algorithm must be greater than or equal to the average time advance determined by either of the remaining two algorithms. In light of this we only present results on average time advance for the ASC algorithm. Figure 5 is a plot of the average time advance for the ASC algorithm vs the average ETAff of each jet aircraft in the noon balloon. The average amount of time advance for an aircraft is defined as the sum of the time advance per traffic sample divided by the total number of traffic samples. Figure 5 clearly shows that the average time advance gained through delay exchange is greatest in the region where the aircraft ETAffs are most closely bunched. Again, with time advance varying directly with scheduled delay, the amount of average time advance increases with increased traffic density because the scheduled delays and opportunities for delay exchange increase with increasing traffic density.

Figure 6 is a plot of the difference in the average time advance calculated by the ASC algorithm for the two AARs. The vertical axis is the difference between the average time advance for an AAR of 96 and an AAR of 108, and the abscissa is the average ETAff of each of the TA aircraft. Because the difference is positive, it shows that the average delay exchange produced by the ASC algorithm is larger for a more restrictive or lower AAR. Once again this behavior is attributable to the fact that the lower AAR leads to larger delays, thus increasing the likelihood for larger time advance. Figure 6 also shows that the difference in time advance for the two AARs is greatest where the traffic is most dense, confirming that the amount of time advance that can be gained through delay exchange increases with increasing traffic density.

Table 2 Maximum average and individual time advance for the ASC algorithm

Time advance, s	AAR = 96	AAR = 108
Maximum average	438.6	342.7
Maximum individual	823	693

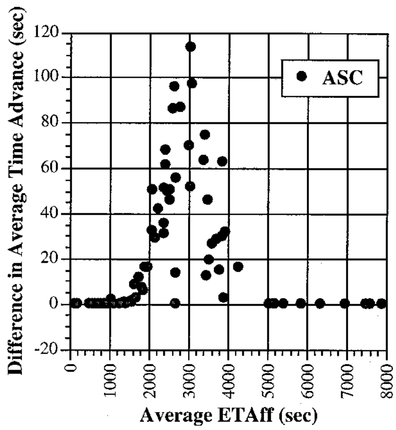

Fig. 6 Change in average time advance with AAR.

Table 2 lists the maximum individual and maximum average time advance gained through delay exchange using the ASC algorithm for both AARs tested. It is seen that the maximum average time advance is substantially less than the maximum individual time advance for both AARs. This is because the individual time advance for any aircraft may vary widely between traffic samples. As expected, both the maximum individual and maximum average time advance are greater for the lower or more restrictive AAR.

Conclusions

The concept of a delay exchange is defined as a fair method of accommodating an airline request for an earlier arrival by advancing the landing time of one aircraft while simultaneously delaying the landing time of another aircraft from the same airline. Any time advance gained by the TA aircraft must be offset by an equivalent amount of additional delay for another aircraft. In this manner, delay exchange will not cause additional Center delays for aircraft from other airlines. For the purposes of this study, it is assumed that any scheduled Center delay of the TA aircraft can be reduced to zero, and that an additional 5 min of time advance can be gained in the Center airspace if a speed change is initiated at least 45 min from the feeder fix. Delay exchanges are restricted to jet aircraft to minimize potential impact on TRACON scheduling. Three different delay exchange algorithms are evaluated in this study: the first algorithm performs delay exchanges between aircraft belonging to the same stream class, the second algorithm performs delay exchanges between aircraft belonging to different stream classes, and the third algorithm allows delay exchanges to be performed between aircraft belonging to any stream class.

A fast-time simulation originally developed for statistical evaluation of CTAS sequencing and scheduling algorithms has been modified for use in this investigation. For each of the three delay exchange algorithms tested, fast-time simulation determines the number of times an aircraft can be TA, as well as the amount of time advance gained through delay exchange. Results show that the algorithm that allows delay exchange between aircraft arriving at any feeder fix has a higher rate of successful time advance than either the algorithm that only allows delay exchange between aircraft arriving at the same feeder fix, or the algorithm that only allows delay exchange between aircraft arriving at different feeder fixes.

For an airport acceptance rate of 96 aircraft per hour, the any-stream class algorithm yields a time advance of 300 s or more for 25% of the traffic samples. For an airport acceptance rate of 108 aircraft per hour, the any-stream class algorithm yields a time advance of 300 s or more for 19% of all traffic samples. A decrease in the success rate with an increase in airport acceptance rate is due to the fact that with a smaller number of aircraft receiving delays for the higher airport acceptance rate, fewer opportunities for delay exchange are found by the algorithm.

Because this simulation does not provide any information about the controller workload required to meet the calculated schedule or perform the delay exchange, the output of the simulation is used only to investigate the feasibility of performing delay exchanges in terms of scheduling. By statistically determining how often and how much time advance can be gained through delay exchange, this simulation provides a starting point for evaluating the potential benefit of performing delay exchanges in scheduled arrival traffic. Given an airline's estimate of how often a specific aircraft may need to request an earlier arrival, and how much earlier the aircraft needs to arrive to meet the airline's objective, the results of the simulation could be used to estimate how often an airline's request for delay exchange would be successful. With the ASC algorithm yielding a 25% chance for an aircraft to gain more than 5 min of time advance through delay exchange, future work will focus on improving the success rate of the delay exchange algorithms. One possibility is to relax the constraint requiring that an equivalent amount of additional delay be traded for any time advance. Relaxation of this constraint should increase the success rate of the delay exchange algorithms. For example, assume that an airline may be willing to accept an additional 10 min of delay, distributed among several arrivals, if it can gain 5 min of time advance for a particular high-priority aircraft. Other areas that will be explored in the future include determining the impact of delay exchanges on controller workload, as well as examining how much of any time advance gained in the Center airspace may be lost in the TRACON.

References

[1]Bond, L., "Global Positioning Sense II: An Update," *The Journal of Air Traffic Control*, Vol. 39, No. 4, 1997, pp. 51–55.

[2]Erzberger, H., Davis, T. J., and Green, S. M., "Design of Center-TRACON Automation System," *Proceedings of the AGARD Guidance and Control Panel 56th Symposium on Machine Intelligence in Air Traffic Management*, AGARD, Neuilly Sur Seine, France, 1993, pp. 52-1–52-14.

[3]Swenson, H. N., Hoang, T., Engelland, S., Vincent, D., Sanders, T., Sanford, B., and Heere, K., "Design and Operational Evaluation of the Traffic Management Advisor at the Fort Worth Air Route Traffic Control Center," *Proceedings of the 1st USA/Europe Air Traffic Management Research and Development Seminar*, Federal Aviation Administration, Washington, DC, 1997.

[4]Erzberger, H., "Design Principles and Algorithms for Automated Air Traffic Management," *AGARD Lecture Series No. 200, Knowledge-Based Functions in Aerospace Systems*, AGARD, Neuilly Sur Seine, France, 1995.

[5]Green, S. M., Goka, T., and Williams, D. H., "Enabling User Preferences Through Data Exchange," AIAA Paper 97-3682, Aug. 1997.

[6]Quinn, C., and Zelenka, R. E., "ATC/Air Carrier Collaborative Arrival Planning," *Proceedings of the 2nd USA/Europe Air Traffic Management Research and Development Seminar*, Federal Aviation Administration, Washington, DC, 1998.

[7]Ballin, M., and Erzberger, H., "An Analysis of Landing Rates and Separations at the Dallas/Fort Worth International Airport," NASA TM-110397, July 1996.