CONCEPTS TO IMPROVE AIR TRAFFIC MANAGEMENT SYSTEM PERFORMANCE[1]

Michelle. J. Blucher, Wayne W. Cooper, Jr., Christopher T. DeSenti, Dr. Alvin L. McFarland,
Dr. Satish C. Mohleji, Arthur P. Smith, III

Center for Advanced Aviation System Development
The MITRE Corporation

ABSTRACT

Demand on the air traffic management system continues to grow. The Federal Aviation Administration is faced with the challenge of increasing capacity to meet the rising demand. Several concepts are being researched that seek to improve system performance in the areas of capacity and efficiency. This paper presents these research endeavors and how they effect air traffic system operations. The four concepts under investigation are focused on various system domains – airports, terminal/transition airspace, and en route airspace. The concepts presented include the Departure Enhanced Planning and Runway/Taxiway-Assignment System (DEPARTS), relaxation of the altitude-for-direction (AFD) rule at higher altitudes, the Transition Airspace Controller Tools (TACT), and capacity enhancements for converging runway configurations. The research objectives for each of the above concepts is discussed including results from studies and experiments and the expected operational impacts these enhancements have on air traffic management system performance.

INTRODUCTION

The nation's air traffic management (ATM) system has faced a significant challenge over the last several years. Increasing demand has led to increasing congestion and delays, not only at major airports, but also in a number of en route airspace sectors. While bad weather has always affected airports, the extent and severity of delays has grown markedly in the last few years. Similarly, thunderstorms have affected the en route air traffic control (ATC) system, in addition to their impacts on airports. Increases in demand have also compounded the problem in the en route system: congestion "backs up" more quickly, more "crossing traffic" compounds this, and there is less capacity in nearby sectors where traffic previously could be off-loaded. Recent approaches to dealing with the congestion have been to impose more restrictions that cause aircraft to fly longer routes, or to fly less efficient flight profiles. These factors all underscore the need and urgency for ATM system capacity, performance, and efficiency improvements.

---

[1] This work was produced for the U.S. Government under Contract Number DTFA01-01-C-00001, and is subject to Federal Aviation Administration Acquisition Management System Clause 3.5-13, Rights in Data-General, Alt. III and Alt. IV (Oct., 1996). The contents of this document reflect the views of the authors and The MITRE Corporation and do not necessarily reflect the views of the Federal Aviation Administration or the Department of Transportation. Neither the FAA nor the DOT makes any warranty or guarantee, expressed or implied, concerning the content or accuracy of these views.

This paper presents several concepts that The MITRE Corporation's Center for Advanced Aviation System Development (CAASD) is investigating. CAASD works in close partnership with the Federal Aviation Administration (FAA), as well as other aviation stakeholders. As the FAA's Federally Funded Research and Development Center (FFRDC), CAASD brings together the expertise and outlook of government, industry, and academia to solve complex technical problems that cannot be solved by any one group alone.

The four concepts discussed in this paper are designed to increase the capacity and quality of service provided by the ATM system. Each of these concepts is at a different stage of research with varying factors addressed in the research. Once evaluated, these concepts will require further testing in actual operational settings. These concepts are focused on different ATM system domains (i.e., airports, terminal areas, en route airspace, traffic flow management) and entail procedural improvements, ATM decision support tools, and avionics improvements. These concepts, as depicted in Figure 1, include the Departure Enhanced Planning and Runway/Taxiway-Assignment System (DEPARTS), relaxation of the altitude-for-direction (AFD) rule at higher altitudes, the Transition Airspace Controller Tools (TACT), and capacity enhancements for converging runway configurations.

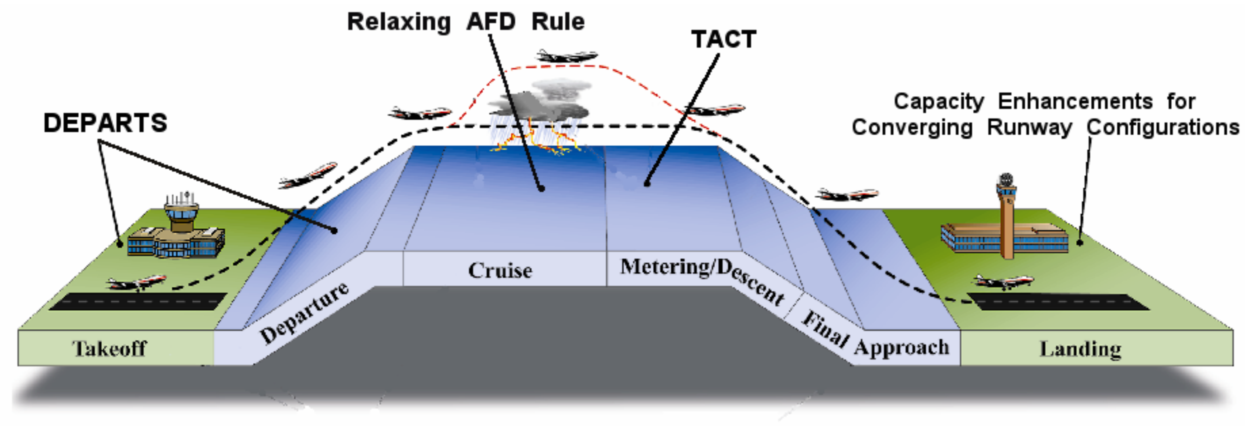

Figure 1. By Domain, Concepts to Improve ATM System Performance

The research behind each of these improvements will be discussed in this paper, together with a description of their expected operational impacts on the ATM system. Each of the following sections represent independent work program areas within CAASD. At this stage, no attempt has been made to present an integrated view of all four concepts. Given the constraints of this publication, some details of the research and findings are not discussed directly, rather references are made to other papers and documents on the given subject.

DEPARTURE ENHANCED PLANNING AND RUNWAY/TAXIWAY-ASSIGNMENT SYSTEM

Problem Being Addressed

Aircraft often encounter significant departure delays at major airports in the United States (US) during periods of high demand. These delays have an adverse impact on airlines and passengers in terms of disrupted flight connections and increased operating costs. Airlines need improved accuracy in predicting departure schedules to plan their operations efficiently, especially at hub airports. The ATC system also needs better information on aircraft push-back times to minimize departure queues and maximize departure runway throughput. Optimization of runway departure sequences to maximize runway utilization is difficult due to limited and inaccurate information on gate departure times, ramp area movement, and taxi times. As a result, aircraft are delayed waiting in departure queues longer than necessary. This lack of gate departure time predictability also adversely affects the predictability of the traffic levels in the airspace immediately around the airport managed by the Terminal Radar Approach Control (TRACON) and the adjoining Air Route Traffic Control Center (ARTCC). This lack of predictability, in turn adversely affects the ability to manage the flow of this traffic proactively.

Figure 2 depicts the information exchange that occurs from the time the original Instrument Flight Rule (IFR) flight plan is submitted until the departing flight takes off. The flight plan is typically submitted 30 minutes or more prior to the time of push-back. The flight plan specifies an estimated runway departure or “wheels-off” time, however schedule push-back times are not entered or updated in the ATC system computers. This flight plan may be amended with new runway departure times; this often occurs, especially if the flight is included in a national Ground Delay Program (GDP), which is managed by the Air Traffic Control System Command Center (ATCSCC) in Herndon, VA (not depicted).

The FAA Air Traffic Control Tower (ATCT) determines the runway configurations in use and the assignment of flights to runways, and passes this information back to the ramp control tower(s) at the airport. In some cases, the ATCT controls the gate and ramp area as well, but at most major airports the gate and ramp area control is under a separate ramp control tower. The controllers at the ramp control towers authorize push-backs and issue instructions to proceed to specific ramp exit “spots.” The flight first in line at a spot will then request clearance to taxi from a ground controller in the ATCT. The ground controller then sequences these flights for the departure runway. After the aircraft has been cleared to taxi into take-off position, a local controller in the ATCT will then give the departure (take-off) clearance to the flight that is first in line for that runway.

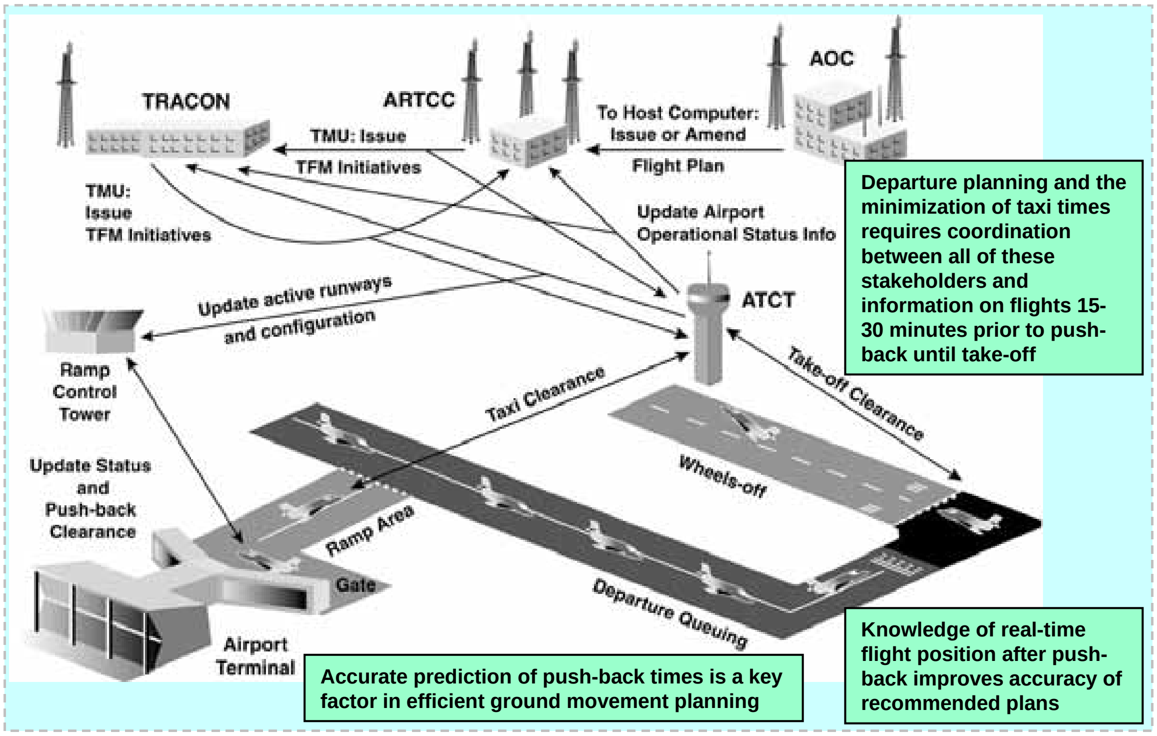

Departure planning and the minimization of taxi times requires coordination between all of these stakeholders and information on flights 15-30 minutes prior to push-back until take-off

Accurate prediction of push-back times is a key factor in efficient ground movement planning

Knowledge of real-time flight position after push-back improves accuracy of recommended plans

Figure 2. Information Flow For Departure Planning and Execution

Research: Difficulties in Predicting Departure Times

A key input to a tactical departure planning model is a schedule of when specific flights plan to push-back from the gate. This information should be known approximately when a flight's IFR flight plan is filed with the FAA. Unfortunately, last minute updates to push-back times are usually not made available to the FAA for input to departure planning.

The largest domestic airlines provide Airline Service Quality Performance (ASQP) data to the FAA for each domestic flight to or from the busiest US airports. This information includes actual "out" (gate push-back), "off" (wheels off), "on" (touchdown at the arrival airport), and "in" (gate arrival) times, which are often referred to as Out, Off, On and In (OOOI) data. A study of flights at major US airports was conducted to analyze the difference between the originally scheduled gate push-back time and the actual push-back time. The comparison, push-back prediction error, showed that the scheduled departure times and actual departure times differ greatly with mean differences in the range of 9 to 25 minutes and standard deviations in the range of 28 to 52 minutes. These statistics, based on August 2000 ASQP data for 15 major US airports, are depicted in Figure 3. Clearly, these differences can be reduced if users provide updates to predicted push-back times as conditions change prior to push-back. These updates will likely be done more frequently in the future through FAA-user information exchange capabilities such as the Collaborative Decision Making Network (CDM-Net).[2]

---

[2] CDM-Net is currently being used by the FAA and airlines to collaboratively manage departure times for flights that are included in GDPs.

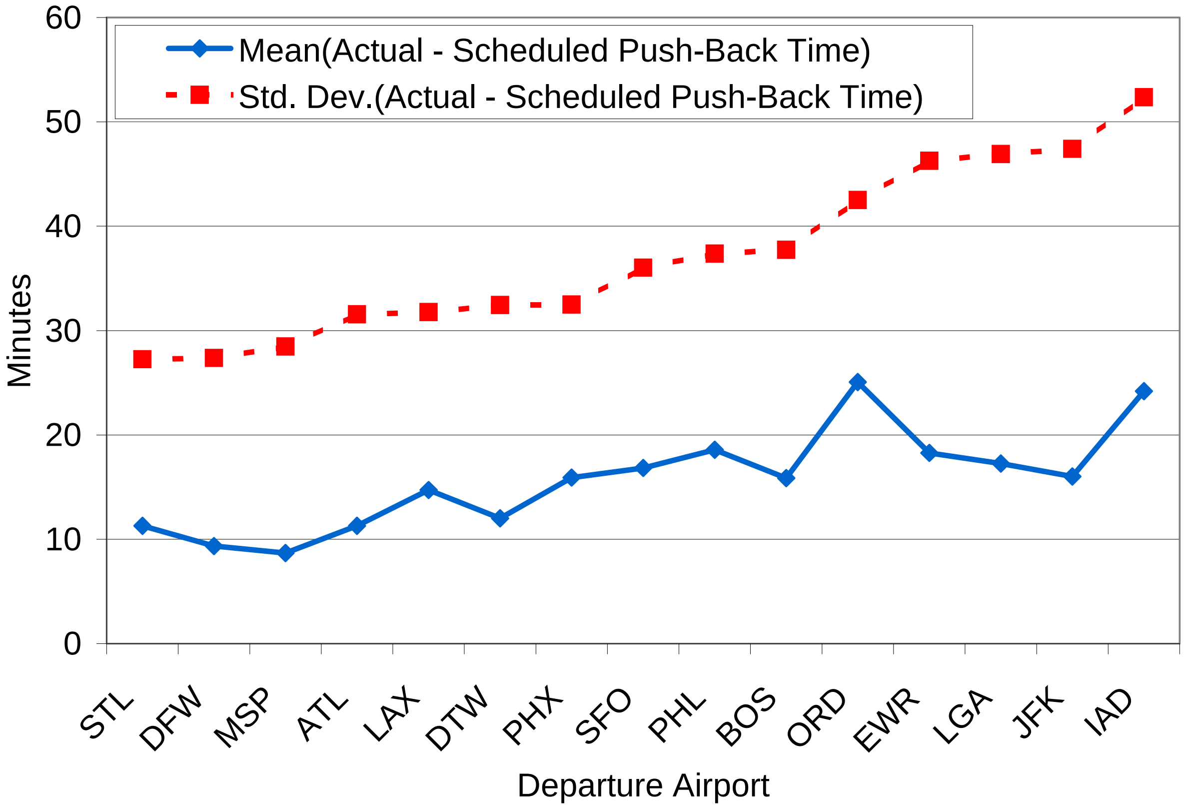

Figure 3. Scheduled vs. Actual Push-Back Times For 15 Major US Airports

In addition to the difficulties in predicting departure times, the taxi clearance times are not widely recorded nor made available in the current ATC system in an automated form. The first time that the FAA ground controller receives reliable information that a flight is ready to depart is when the pilot calls the ATCT for a taxi clearance. In order to optimize the runway departure sequence and to reduce the size of the taxiway queues, accurate and immediate reporting of both actual push-back time and taxi clearance time is needed.

A number of researchers are investigating the design and use of decision support tools to improve departure operations [Anagnostakis et al., 2000, Idris et al., 1998, Idris, 1999, Pujet, 1999, Schumsky, 1997].[3] CAASD has done research in this area, and developed a lab prototype for La Guardia departure operations [Barrer et al., 1989]. The Departure Enhanced Planning and Runway/Taxiway-Assignment System (DEPARTS), a departure planning tool being developed at CAASD as a laboratory prototype, is intended to improve planning accuracy for departures thereby helping users to reduce taxi-out times. Other research has also been done in this area, in

---

[3] Pujet developed and calibrated a stochastic model of taxi-out times that considers the departure runway configuration in use and departure demand. Idris developed a statistical model of the departure queueing process for an airport, including a model of the resequencing dynamics from first-come first-served ordering, based on a knowledge of how air traffic controller clearances are translated into the ground movements of specific aircraft. Both Pujet and Idris used Boston Logan for their research.

particular, Idris [Idris, 1999] has researched measures of reliability in controlling departure operations. Pujet [Pujet, 1999] has researched the area of modeling and minimizing taxi-out times.

Details of Study

DEPARTS is currently being developed as a lab prototype as it is not in the scope of our research to develop the required real-time data infrastructure required for eventual fielding. DEPARTS is being developed to (1) create the algorithms required for optimization of departure schedules as part of a future surface air traffic management system, and (2) to analyze the requirements for data accuracy and timeliness to produce operationally useful results. The operational context and potential benefits of DEPARTS algorithms under a variety of different input data quality scenarios has been discussed in detail elsewhere [Cooper, Cherniavsky et al., 2001, Cooper, Mohleji et al., 2001].

The objectives of the DEPARTS research effort are:

1. To establish the requirement for minimum predictability of future push-back times 30 minutes ahead of time. Thirty minutes is used as most flight plans are filed in this timeframe and it provides a reasonable look-ahead time for a supervisor in the ATCT;
2. To develop algorithms and models to assign flights optimally to the appropriate runway (in the case of multiple departure runways) and sequence them for each runway in order to minimize total taxi-out times; and
3. To provide a tool to plan the adjustment of scheduled gate push-back times or the rerouting of departures via a different departure fix to alleviate the impact of traffic flow management (TFM) initiatives (e.g., miles-in-trail (MIT) restrictions) on these departure fixes.

The information required from the airport being modeled by DEPARTS includes:

- Flight plan information for each flight: route of flight, departure fix and projected runway departure time and arrival time. The departure fix is needed to determine the path of flight in the terminal area, so as to avoid conflicts with other departures. The arrival airport is needed to apply certain TFM restrictions.
- Gate number, planned and actual push-back times, actual taxi clearance time, actual runway used, and actual wheels-off time for each departure. This information is needed to determine the path and unimpeded taxi time from the current position of the flight to the departure runway threshold.
- Planned and actual wheels-on times, runway used, gate number, and in-gate time for each arrival. This information is needed to determine the path and unimpeded taxi time from the current position of the flight to the arrival gate.
- TFM restrictions in effect: type, time issued, duration, and specific flights affected. This information is needed to add the required departure time constraints.

- Hourly airport arrival and departure rates, runway configurations used, and projected runway configuration change times (if any): this information is needed to conform the algorithms to the appropriate airport configuration and weather conditions.

The DEPARTS algorithms would receive the above information from an existing real-time information management system. Only a limited number of airports have such information available in real-time. Several systems currently in use at selected airports, or currently under development, have most or all of this information: ARMT,[4] SMA,[5] DDTC,[6] DSP,[7] and SMS.[8] CAASD is investigating the potential future integration of DEPARTS algorithms into one or more of these initiatives, focusing particularly on SMS. There is no system currently being used in the US air traffic control system to plan the runway assignment and departure sequence of future push-backs.

The DEPARTS prototype is a pre-departure planning tool that focuses on a short-time planning horizon of up to 30 minutes. It allows quick updates of the recommended runway and take-off time assignments based on changing conditions at the airport or anticipated push-back times of near-term departures.

The mathematical optimization problem that must be solved is inherently difficult since it must make the optimal runway departure time assignments for *all* departures in the planning horizon, considering *all* of the resource constraints. If these algorithms were used as part of a fielded decision support system, they would also have very strict limits on solutions times, typically of one minute or less for total run time. The mathematical model used to solve this problem is formulated as a Mixed-Integer Program (MIP). It is a model that "mixes" both binary (i.e., integer) decisions, such as assigning individual flights to a specific order in the departure sequence for a runway, with continuous decision variables, such as the optimal take-off time. This type of mathematical model has been extensively researched, as have efficient techniques for formulating them [Nemhauser et al., 1999, Williams, 1999]. The commercial optimization code CPLEX, a product of ILOG, Inc., is used to solve these MIP models.

---

[4] Airport Resource Management Tool (ARMT): a locally developed prototype system in use at Atlanta that uses a central data base to provide a real-time display of information on departure and arrival operations to the ATCT.

[5] Surface Movement Advisor (SMA): a prototype system developed jointly by the FAA and NASA for use at Atlanta which provides real-time information, predicted arrival times, and recommended departure runway sequencing. SMA was deployed as part of Free Flight Phase 1 and is operationally used to provide predicted arrival times.

[6] Data Link Delivery of Taxi Clearances (DDTC): a system developed by ARINC and Northwest Airlines to improve the communication of clearances and situational awareness of the airline's ramp at Detroit.

[7] Departure Spacing Program (DSP): a prototype system developed by the Computer Sciences Corporation, to provide departure metering capabilities to manage departures from multiple airports to a common departure fix. DSP is currently in use at centers and TRACONs in the New York area.

[8] Surface Management System (SMS): the NASA Ames Research Center is developing a prototype system that will include all of this information as well as decision support capabilities for departure planning similar to the ones being developed by DEPARTS in the laboratory.

In DEPARTS, the computational difficulty of this problem is addressed by breaking down the mathematical model into two smaller linked sub-problems, the "master problem" and the "detailed problem," which are solved in sequence. This is motivated by a basic assumption that the primary constraint for departures, under most conditions (ignoring TFM restrictions), is the capacity of the runways being used, not the taxiway structure.

The objective of the master problem solver is to assign the optimal runway and take-off time to each flight, considering the following operational parameters and restrictions:

- The predicted ready to push-back time for each flight
- The unimpeded taxi time from the departure gate to the ramp exit threshold (i.e., the edge of the ramp area where the taxiway network begins)
- The unimpeded taxi time from the ramp exit threshold to the departure runway threshold
- The minimum practical total queuing time from push-back to take-off, including waiting time at the departure gate, waiting for taxi clearance, and waiting in the final runway departure queue
- The minimum departure spacing on each runway due to wake-vortex spacing requirements
- Restrictions on runway assignments (e.g., based on departure route to avoid crossing traffic flows after take-off)
- Departure TFM restrictions, including MIT restrictions

The current objective is to minimize the sum of the expected time from scheduled push-back to estimated departure clearance time for all departures. However, since each flight is modeled individually, other considerations for each flight, perhaps based on airline preferences, could be used. Equity in re-sequencing flights from their "first come, first served" order across different airport users must be maintained for any decision support tool to be accepted.

The master problem is broken down further into sub-problems, with a unique set of candidate runway assignments specified for each sub-problem. This can readily be done for an airport such as Atlanta, as there is a set of operationally feasible assignments of flights by departure fix to specific departure runways. The assignment of the so-called departure fix "splits" is done by the ATCT supervisor to avoid having departure flight paths from different runways cross in the terminal airspace [Cooper, Mohleji et al., 2001]. The master problem solver defines a series of candidate departure splits, and then solves each set as a separate MIP model. The departure split solution with the smallest objective function value would then typically be chosen, although all solutions would be available. For Atlanta, a typical problem during a departure push for one candidate split with a 30-minute time horizon would have 48 push-backs during the time horizon, 31 taxiing flights at the beginning of the horizon, with about 600 total binary decision variables and 3800 linear constraints. The solution of the master problem provides the essential information that would be required by a tower supervisor (i.e., runway assignments).

The detailed problem solver takes the solution of the master problem solver as input; thus, the ready to push-back times and departure clearance times for each flight are given. The solution of the master problem (e.g., departure runway assignment and sequencing) constrains the solution of the detailed problem. The detailed problem solver recommends the best taxi path for each departure, including its planned taxi clearance time and the amount of time spent waiting at the departure gate, in the ramp area, on taxiways, and in the final departure queue. The objective of the detailed problem solver is to determine an operationally feasible taxi path for each planned departure that considers the relative priorities given for waiting at the gate, in the ramp area, and on the taxiways [Cooper, Mohleji et al., 2001]. It also must have the flexibility to "resolve" the solution from the master model if it finds that the master model solution cannot be implemented due to taxi path constraints. This re-solution feature is necessary for algorithmic robustness, although the master problem solver should avoid generating any infeasible solutions under almost all conditions.

The detailed problem solver has not yet been fully developed. The current DEPARTS master problem does not include the modeling of arrivals, as the flow of Atlanta arrival traffic is largely separated from departure traffic except in the ramp area.

The use of DEPARTS algorithms is intended to provide the following operational benefits to the various positions involved in near-term departure planning, through the use of a single, shared departure plan:

- ATC tower supervisor/traffic management coordinator (TMC): an advisory of optimized runway assignments for departures that are ready to push-back in the next 15-30 minutes. Following this advisory will reduce average taxi-out times, and increase departure throughput during peak demand periods.
- ATC ramp controller: an advisory of desired departure runway assignments and push-back times and the sequence of flights leaving each ramp exit. Following this advisory will set up the sequence of aircraft expected by the ground controller, thus reducing the ground controller's workload.
- ATC ground controller: an advisory of taxi clearance request times for flights routed via each ramp exit and the optimized sequence for each departure runway. By having the ramp controller provide flights in the desired sequence and time from each ramp exit point, the workload needed by the ground controller to set up the optimal sequence is reduced. Using the sequencing advisory provided by DEPARTS will also improve the ultimate departure throughput and reduce taxi-out times, particularly during peak periods, as the optimization model may consider more alternatives than the ground controller can during busy periods.
- Airline Operations Center (AOC) flight dispatcher: advance information to support rerouting of flights prior to push-back to balance the traffic load dynamically across multiple departure fixes, reducing taxi-out delays due to departure fix MIT restrictions.
- All positions named above: an approved departure plan shared among all positions by a common data infrastructure will provide: (1) increased awareness of future expected delays and of the effects of TFM restrictions, (2) better planning information for traffic management of these departures after take-off (e.g., future traffic loads at busy sectors), and (3) reduce the need for routine voice communication of departure planning information (e.g., runway assignments) among these positions.

It should be noted that the DEPARTS functionality is intended to be first implemented at the tower supervisor level, where less detailed decisions are made (e.g., runway assignments), and later at the ground and ramp controller levels. The more detailed decisions (e.g., final departure sequence) on the movement of individual departures require higher data integrity and model accuracy.

Results

DEPARTS has been developed as a laboratory simulation prototype for current operations at Atlanta using several representative days of historical flight operations. Benefits accrue during periods of peak departure demand, when taxi-out delays are greatest: the so-called departure "push" periods. The simulation results show that DEPARTS would reduce average daily taxi-out delays by about one to two minutes per flight, with greater benefits during departure push periods [Cooper, Cherniavsky et al., 2001]. This estimated benefit is based on a comparison to the current (quite efficient) operation at Atlanta (ATL) for two days in August 2000. No departure planning decision support tool was in use. Other airports experiencing significant daily peak departure delays could also have comparable benefits. Benefits accrue with increasing accuracy of predicted push-back times and with the number of flights pushing back in a time period. More flights pushing back in short time periods provides the algorithms a greater opportunity to improve a first come, first served solution.

Achieving the maximum potential benefit of the DEPARTS capability requires both that

- The airlines provide real-time updates to the anticipated ready-to-push-back times as flight plan changes occur (e.g., expected delays due to connecting bags or passengers)
- The ATC system changes future planned runway assignments due to anticipated departure mix changes

ATL has departures from independent parallel runways. The ongoing research effort includes modifying the DEPARTS algorithms to incorporate characteristics of a large number of airports, such as (1) dependent/converging/intersecting runways, (2) shared runways for arrivals and departures, (3) the explicit modeling of arrivals from touchdown to gate, and (4) additional flexibility in assigning departures to runways during off-peak periods, when operational constraints that prohibit the crossing of departure flight paths (even if separated by time) in the terminal airspace can be lifted. These enhancement are need to ensure that DEPARTS provides a robust solution and significant estimated benefit that can be applied to a wide variety of airports.

Summary

The use of real-time optimization-based departure planning has the potential to provide significant operational benefits at major airports by reducing taxi-out delays and providing better predictions of runway departure times for downstream TFM. Such has been the result of the use of the DEPARTS algorithms as a planning tool as part of an integrated information management system.

RELAXING THE ALTITUDE-FOR-DIRECTION RULE AT HIGHER ALTITUDES

There have been several types of flight restrictions applied by the FAA over the years to provide structure so that air traffic controllers could handle a significant growth in traffic. These restrictions include: altitude restrictions for departures and arrivals, required arrival and departure routes, en route high and low altitude preferred routes, one-way airways, miles-in-trail restrictions, and the Altitude-for-Direction (AFD) rule. The objective of this study was to investigate if and where it would be feasible to relax the latter of these restrictions, the AFD rule. There are several ways in which the AFD rule might be eliminated, but certain restrictions on cruise altitude still retained. Also, the rule might be eliminated only in certain geographical areas or at certain altitudes. The term, "relaxing the AFD rule," is intended to convey that this study did not consider eliminating the rule everywhere.

Problem Being Addressed

The AFD rule causes longer duration flights to burn more fuel than if the rule were not in place, because the rule causes flights above Flight Level (FL) 290 to approximate the optimal cruise altitude profile with 4000 foot step climbs. The rule is in place to structure traffic for air traffic controllers. It prevents a number of traffic conflicts that otherwise would have to be resolved by controllers. And it avoids many opposite-direction, same-altitude, high-closing-speed conflicts, where late or missed conflict recognition might occur if the controller’s attention is focused elsewhere. RTCA Special Committee 192 suggested that the feasibility of eliminating the AFD rule be investigated [RTCA, 2000]. In 1999, FAA and industry participants outlined a program called Free Flight Phase 1 (FFP1) to expedite limited deployment of automation capabilities that had been demonstrated in field trials. In 2000, the FAA committed to enhancing these systems and deploying them to additional sites with a Phase 2 (FFP2). One of the systems included in the Free Flight Program is the User Request Evaluation Tool (URET) [FAA, www.ffp1.faa.gov]. URET provides automation tools to en route controllers that help them do strategic planning for flights entering their sectors, alert them to potential conflicts, and help them develop resolutions to conflicts. It is believed that URET can support the elimination of the AFD rule in some airspaces by helping controllers deal with the additional conflicts that would occur, and by providing reliable, timely warning of opposite-direction, same-altitude conflicts. The feasibility of eliminating the rule depends on the number and nature of additional conflicts experienced. One of the tasks of the study described here was to assess the number and nature of such conflicts.

Research: The AFD Rule and How It Might Be Relaxed

Figure 4 illustrates the effects of the current AFD rule and two proposed approaches to eliminating the AFD rule at higher altitudes. The column labeled Today illustrates the current rule. As indicated by the arrows, eastbound flights in cruise must fly at FL290, 330..., and westbound flights in cruise must fly at Flight Level 310, 350... No flights fly at even 1000-foot levels. The heavy black line indicates the altitude profile that an eastbound flight would fly, remaining at one flight level until the weight decreases to the point where it is efficient to climb 4000 feet to the next correct-for-direction level. The second column shows the situation if the AFD rule were suspended in such a way that a flight could fly any of the flight levels currently used, regardless of direction. This is indicated by the two-way arrows at each odd 1000-foot level. With the AFD rule removed, flights could fly with 2000-foot steps as

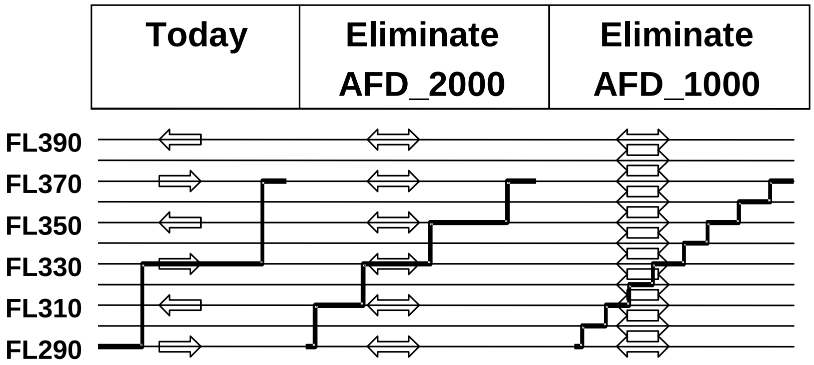

Figure 4. Alternatives for Eliminating the AFD Rule

indicated by the heavy line. In this case, flights at one cruise altitude would be automatically separated by the 2000-foot minimum from flights at an adjacent altitude. The final column shows the situation if the AFD rule were eliminated and flights were allowed to fly any 1000-foot level. In this case, flights could fly 1000-foot step profiles as shown by the heavy line. If the 2000-foot vertical separation minimum is preserved, a flight would be potentially in conflict with a flight at an adjacent 1000-foot level, but not with flights at levels 2000 feet or more away.

Today the AFD rule can be waived, if necessary, for the air traffic situation or for a pilot request during periods of light traffic. But during normal operations the rule is observed quite consistently. As an indication, of 42 flights in cruise at or above FL240 in Indianapolis Center (ZID) at 12:14PM local time on 4 August 1997, all but two were at the correct altitude for direction of flight. While the AFD rule may be waived today on an occasional basis, FAA has not yet announced procedures or designated airspace where the rule is suspended on a regular basis, as studied in this paper.

Figure 5 illustrates how smaller step sizes in the altitude profile contribute to fuel savings. Nearly all turbojet aircraft at higher altitudes have a curve for optimum altitude versus time after takeoff, similar to the optimum altitude curve shown here. As the weight of the aircraft is reduced through fuel burn-off, the optimum altitude increases. Aircraft operators take advantage of the fuel savings today by leveling at the highest correct-for-direction altitude that is less than the optimum for their top-of-climb weight. When the weight reduction is sufficient, they climb 4000 feet to the next altitude that is optimum for that weight. The process may be repeated as shown by the 4000-foot step profile. The altitude profile with 1000-foot steps has also been depicted to illustrate the effect of smaller step sizes. At all times on the 1000-foot step profile, the aircraft is at the same or higher altitude, and is always at the same or lesser distance from the optimum altitude curve. For aircraft with a substantial time at cruise altitude, the difference in fuel burn between the two profiles can be significant.

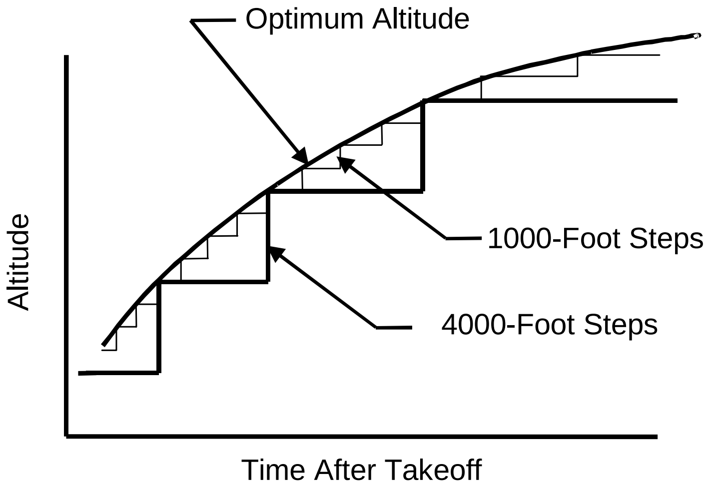

Figure 5. Effect of Smaller Altitude Steps

Details of Study

As a convenience, this study used the strategy where flights remained at one level until a climb of 1000, 2000, or 4000 feet would place them exactly on the optimum profile. A more efficient strategy is to fly for periods both above and below the optimum altitude. Use of the former strategy is not considered a serious limitation for the current study, where the primary measure of comparison between the 1000-, 2000-, and 4000-foot step climbs is the number of conflicts experienced.

The study reported here created three simulation scenarios, all based on the same real-world scenario from Indianapolis Center. The scenarios are named Today, Eliminate AFD-2000, and Eliminate AFD-1000, representing Columns 1, 2 and 3 respectively in Figure 4. Data was collected for the period 11:45AM to 3:00PM on 14 June 1999. This day was selected from a number of days for which reduced data sets were available to support this analysis. A review of archived weather and ATC performance information indicated that there were no unusual delays in the ATC system, and no disruptive weather that would influence the paths or altitude profiles of the flights. Only flights with a flight plan altitude of FL270 or greater were retained, since the study only considered eliminating the AFD rule above FL290. The ATC environment is considered simpler above than below FL290 because fewer aircraft are climbing and descending. While eliminating the rule at lower altitudes could be considered, an initial introduction at higher altitudes seems likely and was used for the North American Route Program, that allowed flights to fly more efficient routes at higher altitudes.

Twenty-four hours of Enhanced Traffic Management System (ETMS) data for 14 June 1999 was used to derive observed altitude-versus-time curves for nine classes of high altitude aircraft. Each aircraft observed was assigned to one of these classes. These curves are not strictly the optimum altitude-versus-time curves, because not all flights in the real world were following the optimum 4000-foot step strategy. But they produce realistic dispersion of flights within a class, place flights at roughly appropriate altitudes for their class, provide a basis for generating plausible step profiles, and to some degree account for the fact that some flights would not follow the optimum step profiles, even if the AFD rule were eliminated.

Software created the scenario data for a subject flight in each of the three scenarios. The software designated the flight as an eastbound or westbound flight based on the longitudes of the departure and destination airports. The software then generated 1000-, 2000-, or 4000-foot step climbs as appropriate for the direction of flight and the scenario. The scenario data for a single flight in one of the scenarios consisted of one initial flight plan followed by several flight plan amendments with higher altitudes. The times of the altitude amendments were calculated from the observed altitude-versus-time curves for the aircraft class to which the subject flight belonged. The initial flight plan containing the initial cruise altitude, and the flight plan amendments containing the altitudes for subsequent climbs represented the step climb profile. The horizontal routes for a flight were identical in all scenarios generated. The set of flight plans for all flights and a given scenario were merged and sorted in time order to create the simulation scenario file for that scenario. Each of the three derived scenarios was run through CAASD's URET prototype without any human involvement. These runs are called the Unattended Runs. Records of the conflicts notified by the URET prototype to each sector were collected. Throughout this study, any conflict notified by URET, regardless of alert level (red alert if closest predicted lateral separation is less than 5 nmi, and yellow alert if greater than 5 nmi), was counted as a conflict. In the simulations, all aircraft fly the flight plans perfectly. There is no navigation or pilot error represented. In Unattended Runs, aircraft do not carry out maneuvers to resolve the conflicts and maintain the required separation minima. Since only the first flight plan for each flight, and none of the real-world track data, is used in these scenarios, the maneuvers used by controllers to resolve conflicts in the real world are not reflected in our scenarios.

Output from the Unattended Runs was analyzed to identify time periods and sectors where the greatest number of conflicts were notified in a 15-minute interval. The simulation runs were repeated in an interactive mode called Interactive Runs. An analyst at a URET sector workstation interactively resolved all conflicts that were notified to that sector. The analyst stopped the simulation clock when an alert was notified, and constructed URET trial plans of proposed resolutions. (The analyst was not skilled enough to generate efficient resolutions in real time at the busiest sectors.) When a suitable resolution was found, it was implemented. The resolutions chosen followed general rules that field controllers have articulated over many years while helping analysts generate resolution rules for problem resolution aids [Kirk et al., 2000]. These rules avoid as many undesirable aspects of resolutions as possible. A list of some of the undesirable aspects is presented in an earlier paper [McFarland et al., 2001]. The flight would follow the resolution maneuver, and the alert would disappear. The analyst restarted the simulation clock and proceeded to the next notified conflict. This exercise was completed for the Today scenario, and then repeated for one of the scenarios with modified rules. The analyst noted observations about the impact of the modified rules on ATC operations.

Results

Table 1 presents results of Unattended Runs with the AFD rule eliminated at all altitudes above FL290. The current 2000-foot vertical separation minimum was retained for these results. Results are provided for both the Eliminate AFD-1000, and the Eliminate AFD-2000 scenarios. The entries in the table are the total number of conflicts notified to each sector during the 3¼-

hour data collection interval. Results are grouped for the high sectors and the super-high sectors. Subtotals are presented for the two groups and a total is presented for the groups combined. The ZID super-high sector map shown in Figure 6 helps interpret these results. Interactive Runs were made for sectors with shaded entries in Table 1 and shaded areas in Figure 6. The ZID high sectors, except for Sector 79, have the same lateral boundaries as the super-high sectors. High and super-high sectors with the same lateral boundaries have the same last digit in the sector number.

The total number of conflicts increases from 672 in the Today scenario to 823 in the Eliminate AFD-2000 scenario. The number increases further to 945 in the Eliminate AFD-1000 scenario. For the high sectors combined, and for some high sectors individually, the number of conflicts decreased between the Today scenario and the Eliminate AFD-2000 scenario. Investigation showed that some conflicts notified to a high sector in the Today scenario were notified to a super-high sector in the Eliminate AFD-2000 scenario. This is because, on average, aircraft are at a higher altitude with the AFD rule eliminated than in the Today scenario. Correspondingly, the increases in conflicts for the super-high sectors are due, not only to eliminating AFD, but also to the transfer of some conflicts from the high to the super-high sectors. One might expect that the number of conflicts in the Eliminate AFD-1000 scenario would be less than in the Eliminate AFD-2000 scenario, since the flights could be dispersed over a greater number of flight levels. However, the majority of turbojet aircraft have similar optimum altitudes and they tend to cluster in the FL330 to FL350 band in ZID. Permitting flight at any 1000-foot level allows more flights to cluster closer to the optimum altitudes and come within the 2000-foot separation minimum of other flights, thereby creating more conflicts in the Eliminate AFD-1000 scenario.

High Sectors				Super High Sectors			
Sector	Today	Elim. AFD 2000	Elim. AFD 1000	Sector	Today	Elim. AFD 2000	Elim. AFD 1000
I-79	43	59	62	I-90	8	16	19
I-80	39	36	40	I-91	9	33	42
I-81	53	50	54	I-92	20	38	46
I-82	42	38	38	I-93	19	31	41
I-83	18	19	20	I-94	10	12	17
I-84	34	30	37	I-95	28	41	60
I-85	21	14	12	I-96	12	16	25
I-86	27	19	25	I-97	17	29	31
I-87	68	81	86	I-98	71	116	128
I-88	72	63	64	I-99	13	28	38
I-89	48	54	60				
Total High	465	463	498	Total Super High	207	360	447
Total High and Super High					672	823	945

Table 1. Number of Conflicts If AFD Eliminated (Indianapolis Center)

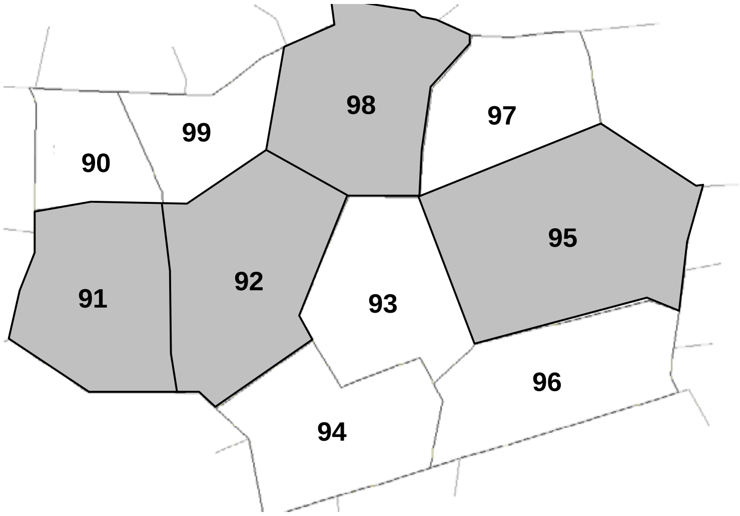

Figure 6. Indianapolis Super-High Sectors

Sector 98 had by far the most conflicts among the super-high sectors in the Today scenario, and the number increased significantly with AFD eliminated. On the other hand, with AFD eliminated, a number of other super-high sectors had no more conflicts than the high sectors had in the Today scenario, suggesting that the increase in conflicts would be manageable for some of these sectors.

Interactive Runs for Sectors 98, 95, 92, and 91 were also conducted on the Eliminate AFD-1000 scenario. At the peak time in the Interactive Runs with Sector 98, it was challenging to resolve all conflicts. Traffic was converging from all directions in the center of Sector 98, and all altitudes in the FL330 to FL350 band were occupied with multiple flights. An Interactive Run was also carried out on the Today scenario, where it was also difficult to resolve all conflicts in Sector 98 with reasonable maneuvers. Despite limitations of the simulation, the conclusion that eliminating the AFD rule in Sector 98 would be difficult seems justified. The number of conflicts with the AFD rule in place is high enough that the controller team should not be expected to handle the significant increase arising from eliminating the rule. In the Interactive Runs for Sectors 95 and 92, it was challenging to resolve the conflicts but not to the degree it was for Sector 98. Air traffic controllers would need to assess operations in these sectors with the Eliminate AFD-1000 or Eliminate AFD-2000 scenario to establish the feasibility of eliminating the rule in these sectors. In Sector 91 there was no difficulty in resolving all conflicts in the Eliminate AFD-1000 scenario. The experience at Sector 91 during the peak in this scenario, and the experience at the other three sectors during non-peak periods, suggests that controllers in the sectors other than these four would have no difficulty handling traffic with the AFD rule eliminated.

Numerous studies of the nationwide airspace system have shown that the areas of greatest air traffic congestion in the US lie in the northeast. Views of the ETMS Traffic Situation Display typically show high concentrations of traffic in this area. An on-line version of a previous paper presents a color figure of the density of air traffic in the US [McFarland et al., 2001]. This figure shows that the most wide spread areas of high density traffic in the country occur in Indianapolis

Center and areas to the north and east. Also, all of the choke point areas recently identified for early attention as part of the National Airspace Redesign fall in the Great Lakes and northeastern areas of the country [FAA, www.faa.gov/ats/nar].

The results of Unattended Runs, the experiences from Interactive Runs, and the observations about the locations of traffic congestion and choke points lead to the conclusion that elimination of the AFD rule is promising in some sizable areas of the US, but not in the most congested areas. This idea is represented notionally by the solid line drawn on Figure 7. Elimination of AFD is uncertain to the north and east of that line. But results of this study suggest that elimination is promising within areas of URET coverage (indicated by the centers with shading in Figure 7) to the south and west of the line. The existence of such a line seems likely, but this study was not comprehensive enough to firmly establish its location. Even though it does not appear feasible to eliminate AFD everywhere, a flight does not have to fly entirely in airspace where the rule was eliminated to derive a benefit. For example, a flight from Boston to Dallas Fort Worth would fly correct altitudes for direction until crossing this line and then would be able to fly 2000-foot steps until ready to descend.

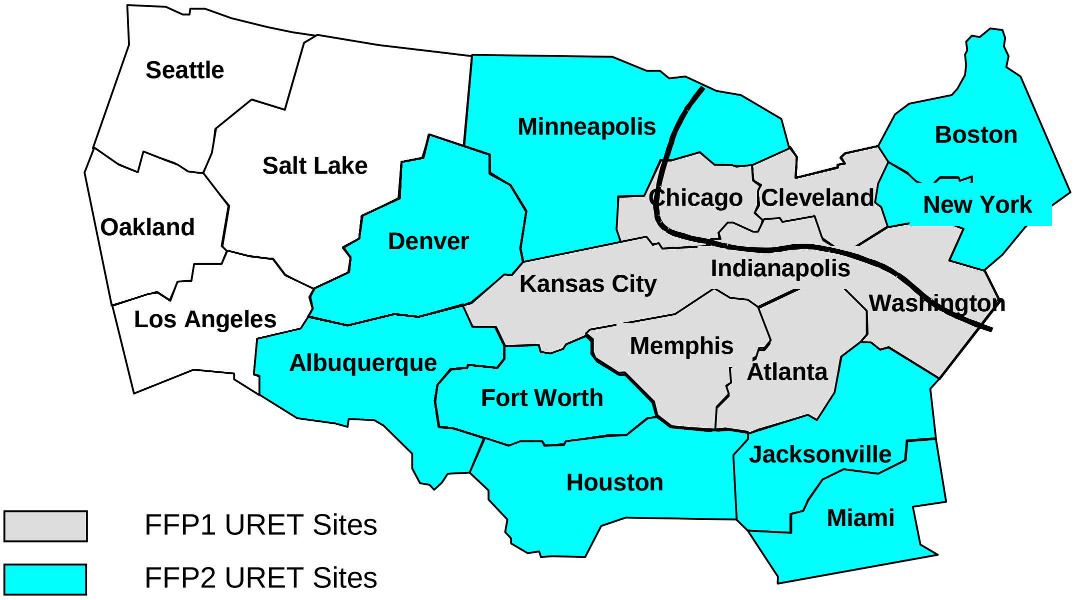

Figure 7. FFP1 and FFP2 URET Sites

The scope of this study did not allow us to calculate the fuel savings from relaxing the AFD rule for each flight. However, individual calculations were carried out for four B757 flights from the Eliminate AFD-2000 scenario. These flights were selected to show a variety of conditions with respect to trip length, initial weight, time of first step, and proportion of trip north and east of the line where the AFD rule was retained. All four flights flew correct altitudes for direction (4000-foot steps) for portions of the flight north and east of the line in Figure 7, and flew the optimum odd thousand-foot level (2000-foot steps) for the other portions of the flight. Eleven repetitions were calculated for each of the four flights with take-off weight varying by 1000 pounds between repetitions. Results of the eleven repetitions were averaged. A fuel cost of 80 cents per gallon was used for these calculations. The flight from Newark to Phoenix saved 1 percent of the total trip fuel ($51 for the flight). The flight from Cincinnati to Salt Lake City saved 1.5 percent of the total trip fuel ($54 for the flight). For the flight from Miami to Chicago, the altitude profiles with and without the AFD rule were the same for six of the eleven repetitions. The savings from the five repetitions where the profiles differed, when averaged over the eleven repetitions, were 0.1 percent or $3. The altitude profiles for all eleven repetitions for a flight from Detroit to Orlando were the same and there were no savings for that flight. While these savings seem small in an absolute sense, they represent a pure increase in profits to the airlines, because there is no associated cost to the airlines to realize them. To place the savings in perspective, we calculated an average profit per flight for the ten major US airlines for 1999 [Lampl, 2000]. The average profit per flight was $772.

Summary

This study produced the following conclusions:

- AFD rules reduce the potential for conflicts that controllers otherwise would have to resolve.
- Some congested areas such as Indianapolis Sector 98 receive significantly more conflicts when AFD is eliminated, so elimination is uncertain in these areas.
- There is a large contiguous area served by FFP1 and FFP2 URET where it appears promising to eliminate AFD in super-high sectors.
- There are no aircraft equipage or additional maintenance or certification requirements levied on operators to support elimination of the AFD rule.
- FFP1 and FFP2 URET capabilities support elimination of the AFD rule and may be required in some areas. These capabilities include continuous and reliable detection and notification of conflicts; convenient trial planning to support early, gradual maneuvers to resolve conflicts; quick and convenient trial planning of requested altitude changes; capability to amend the flight plan to a trial planned altitude with a single click; and inter-facility capability that provides the same alerting and trial planning capability for traffic inbound to the center, as for traffic within the center.
- The benefits and feasibility of eliminating AFD in some super-high sectors are sufficiently promising that this subject should be discussed with ATC controllers and managers. Field evaluations should be conducted, if possible, to confirm the feasibility.

TRANSITION AIRSPACE CONTROLLER TOOLS

The research effort on Transition Airspace Controller Tools (TACT) was initiated to investigate the potential for visualization tools to assist the en route radar controller with elements of complexity they face in transition airspace.[9] The initial stages of this research included visits to six ATC facilities to observe, survey, and identify areas of complexity. Input from controllers regarding tasks performed, and means of approaching them, was instrumental in this early stage of problem identification. One of several issues identified in this effort was that of time-based metering.

Problem Being Addressed

FAA plans include the increased use of time-based metering by the radar controller. Historically, time-based metering has met with some opposition from the operational community. Early metering software, known as the Arrival Sequencing Program, had inherent limitations which contributed to that opposition. Many of these limitations were to be addressed by the FFP1 implementation of Traffic Management Advisor (TMA) as a replacement for the outdated metering software.

Several factors contribute to the complexity of time-based metering for the radar controller, two of which relate to the manner in which information is displayed: the location where metering information is displayed, and the time-oriented method of displaying metering information.

In FFP1 metering information is not integrated into the controller’s traffic. Instead, it is displayed in lists that require the controller to switch focus away from traffic to scan for metering information, then correlate that information back into the traffic. Air traffic operations in transition airspace can be quite dynamic, occasionally requiring the controller’s planning horizon to span the entire range from strategic to tactical. For some, the requirement to constantly switch focus between a Metering List and their traffic represents a significant element of added complexity and workload.

Air traffic control, using a radar display, is fundamentally a spatial activity. This is evident in the fact that the radar controller’s environment is heavily oriented toward spatial perception. The majority of their radar separation standards are distance-based, and their two-dimensional video map display is measured and scaled to facilitate the use of mileage increments. With daily reinforcement, controllers become very adept at working within this environment and develop a keen ability to anticipate and achieve specified amounts of spacing between aircraft. The time-oriented display of metering information is non-intuitive in the context of this spatial environment. For example, consider two aircraft 10 miles-in-trail. A controller, if asked to increase this spacing to 20 miles, would know almost instantaneously what control action will achieve the spacing. This action would include speed control, and the direction and estimated magnitude of vector maneuvers. If instead you ask the controller to increase the spacing by 4 minutes, this too can be achieved. But for the radar controller, the cognitive translation of the time-oriented information into control actions is generally neither as intuitive, nor as seamless, as with spatial information. Elements that before were relatively instantaneous, such as

---

[9] Airspace that generally entails the transition from the en route-cruise portion of flight to the TRACON area.

approximating how far to the left or right the trailing aircraft will need to be vectored, now require conscious effort, or a trial and error approach. Although skills with time-oriented information can be developed, air traffic control with a radar display remains fundamentally a spatial activity. Infusing time-oriented information into the spatial environment, as it is today, imposes a possibly unnecessary cognitive difficulty for the controller.

Research

With consideration for the above issues, as well as the characteristics of the operational environment, the focus of TACT research for metering assistance took on three basic tenets:

1. Make the information intuitive, visually apparent, and spatial to the degree possible.
2. Integrate the information into the controller's traffic display, mitigating their need to look away from the traffic to scan additional lists or supplemental displays of information.
3. Make the proposed tools so that they are usable at the radar position, meaning flexible enough to permit the controller to minimize or eliminate any potential distraction the display may cause when priorities dictate that the controller's attention be elsewhere.

The two primary tool concepts that TACT explored are the Mileage Distance Marker (MDM) and Mileage in the Data Block (MDB). Both are intended to enhance the effectiveness of the time-based schedule by rendering the metering task more achievable and acceptable to the controller. Each visualization tool displays aircraft delay information that spatially indicates the magnitude of the adjustment needed for that aircraft to meet a time-based restriction.

Whether the metering restrictions are scheduled by TMA software or some other scheduling algorithm, the TACT visualization tools work in partnership with them to assist the controller by providing information in a manner that is intuitive. The MDM allows the controller to display a marker that provides an indication of the delay magnitude that is visually apparent, spatial, and integrated into the controller's traffic (see Figure 8). The delay magnitude is represented by the spatial gap between the marker and the aircraft. The controller's goal is to merge the aircraft and the MDM marker.

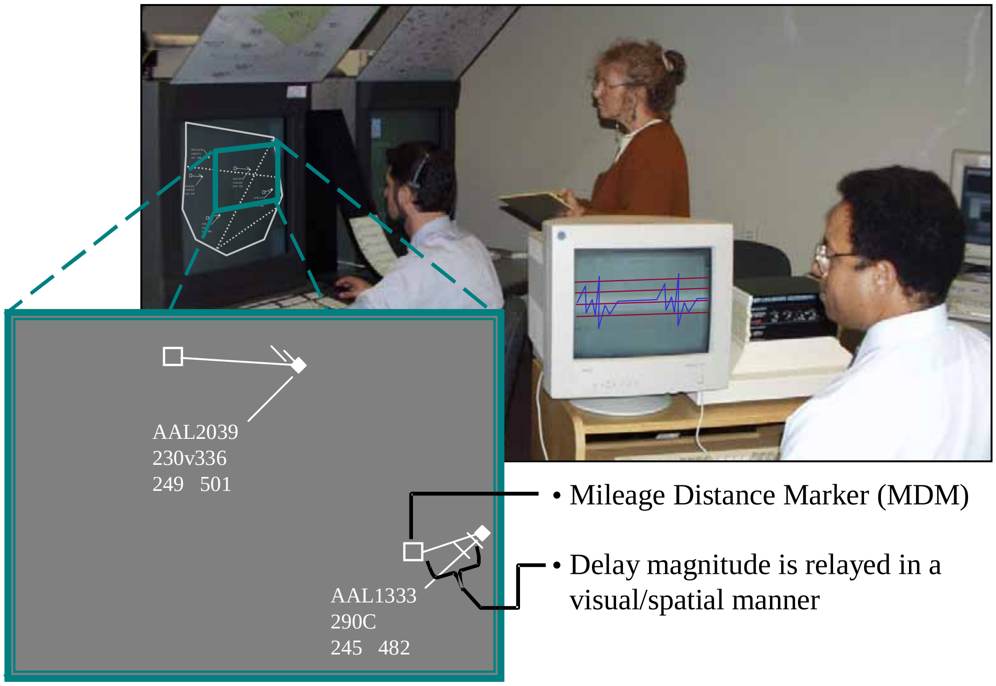

- Mileage Distance Marker (MDM)
- Delay magnitude is relayed in a visual/spatial manner

Figure 8. Mileage Distance Marker

Because the time-based restriction is depicted via the spatial gap between the aircraft and the marker, controllers can relate to the MDM display in much the same way as the spacing between two aircraft with a Miles-In-Trail restriction. The MDB displays a positive or negative number in the aircraft's data block that reflects the mileage adjustment needed to meet the required restriction (see Figure 9). The number will count down or up, displaying zero when desired spacing is achieved.

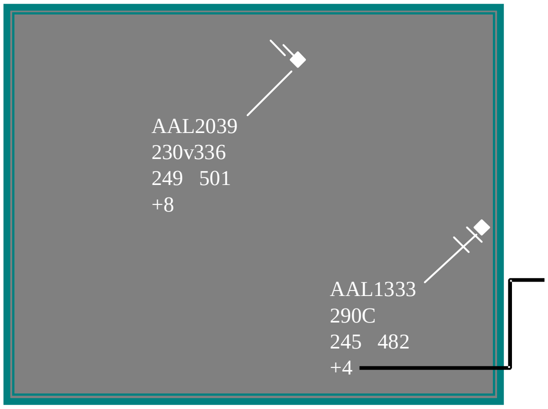

- Mileage in the Datablock (MDB)
- Delay magnitude relayed via numeric mileage increment in datablock

Figure 9. Mileage in the Data Block

The underlying functionality of the MDB is identical to that of the MDM. The key difference between the MDM and MDB is the method of displaying information to the controller.

Details of Study

To assess the TACT tools researchers used an interactive air traffic control environment in a CAASD laboratory. CAASD contracted access to appropriate participants through the National Aviation Research Institute. Active controllers from two ARTCCs participated. The interactive lab environment allowed controller participants to formulate informed opinions concerning operational utility and acceptability. This environment also provided a medium for capturing objective data on controller performance.

The experiment included a radar position with a Display System Replacement (DSR)-like display. The Host Computer System (HCS) display interface allowed access to most normal HCS functions plus access to new functions and displays representing the capabilities to be tested. The experiment used airspace modeled from sector 49,“Logen,” of the Atlanta ARTCC. This airspace was chosen based solely on the in-house availability of HCS adaptation for Atlanta, to minimize the impact on project resources.

The basis for scenario traffic was recorded real-world Atlanta ARTCC traffic. This was done to ensure that the general characteristics of the flows retained a measure of realism with respect to the airspace being used. To construct scenarios suitable for a laboratory evaluation, the traffic was tailored for specific purposes. For example, adjustments were made to help control the type and degree of complexity, and the overall changes in volume that were needed to meet evaluation goals. Overflight traffic flows were filtered to include only a subset of the many possible flight paths that transition the Logen sector. This allowed controllers, having no prior experience with this particular airspace, to gain a level of proficiency and comfort with the operations in a short period of time, facilitating their ability to effectively address the evaluation tasks.

This experimentation took place over two years and sought to gain insight into the effects of the TACT tools in several areas: workload, performance, acceptability of the tools, and acceptability of time-based metering.

Controller participants were given broad procedures to follow during each evaluated scenario run, as well as procedures specific to each tool evaluated. In general, these procedures were as follows:

1. The number one priority is separation and safety.
2. Attempt to meet meter fix times whenever possible without compromising separation and safety.
3. Use the (evaluated) tool when appropriate.
4. Use the tool display modes (or options) available, as appropriate.

Data was collected to indicate affects on controller performance and workload, as well as to capture subjective feedback about tool performance and acceptability. Methods used to capture data included: a modified NASA Task Load Index (TLX),[10] the Controller Acceptance Rating Scale (CARS), and structured post-run interviews. An automated data collection application captured run-time performance, and traffic information and stored the data for off-line post-processing and analysis.

As shown in Table 2, the experiment design utilized a factorial repeated-measures design, with a 2 x 3 (Scenario x TACT Tool) structure in order to control for most of the common threats to internal validity such as: history, maturation, testing, instrumentation, statistical regression, differential selection of respondents, and experimental mortality [Campbell et al., 1963]. Controlling for these classes of extraneous variables reduced the chance that they could produce effects that might be confounded with the effect of the TACT tools. Controllers were assigned randomly to one of three groups. Independent variables in the experiment were Scenarios A and B, the Baseline Metering List, and the TACT tools: MDB and MDM. Dependent variables included performance measures, subjective assessments using the TLX and the post-run questionnaire, and the CARS results [DeSenti et al., 2001]. The repeated measures design also increased the power of the effects tests for the TACT tools by removing the variation of each controller from the residual variance [Kenny, 1994].

Group	**Participant**	**Tool/Scenario**					
		Run 1	**Run 2**	**Run 3**	**Run 4**	**Run 5**	**Run 6**
1	1, 4, 7	BSN/A	MDM/B	MDB/A	BSN/B	MDM/A	MDB/B
2	2, 5, 8	MDM/A	MDB/B	BSN/A	MDM/B	MDB/A	BSN/B
3	3, 6, 9	MDB/A	BSN/B	MDM/A	MDB/B	BSN/A	MDM/B
BSN refers to the Baseline Metering List							

Table 2. Order of Treatments and Scenarios by Participant

Results

The laboratory research revealed that TACT visualization tools provide a clear advantage where controllers perform time-based metering. Indicators reported here include workload, performance, and acceptance. Among the methods used to gain insight with regard to workload was data collected using the NASA TLX. The TLX results showed the MDM tool as requiring the least mental demand, followed by the MDB tool. The baseline tool, or Metering List alone, was ranked by controllers as the most mentally demanding of the three conditions. The average controller responses rated the TACT tools with lower workload than the Baseline for all categories of the modified TLX including: mental demand, physical demand, time pressure, effort (i.e., overall physical and mental effort put forth), frustration, and performance.

The TLX data showed the TACT tool effects on Mental Demand ($F(2,16) = 4.94$, $p = 0.02$), Effort ($F(2,16) = 4.18$, $p = 0.03$), and Overall Workload ($F(2,16) = 4.56$, $p = 0.03$) were

---

[10] TLX was scored on a scale of 1-100, instead of the original 1-10. Wording of the questions was tailored so as to seem sensible with respect to the particular tools and tasks being evaluated.

statistically significant, and the MDM tool was favored specifically. Figure 10 shows the average TLX response for Mental Demand [DeSenti et al., 2001]

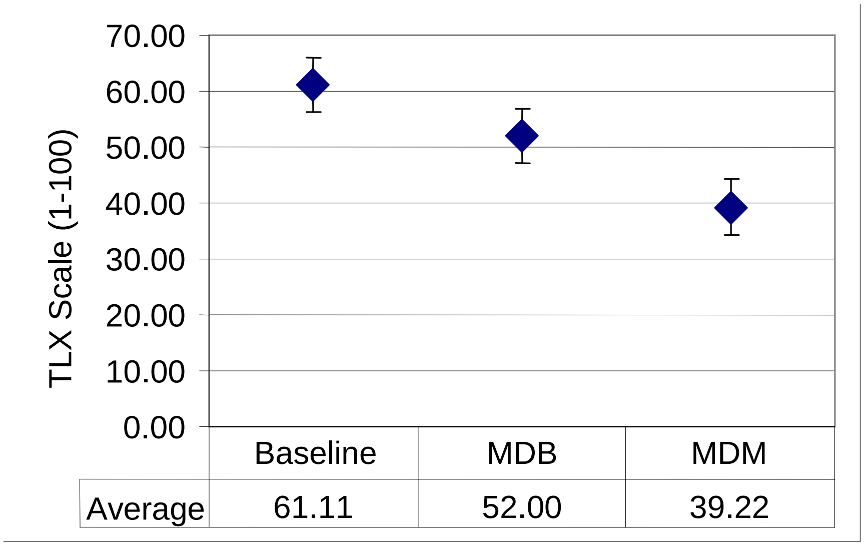

Figure 10. TLX Response for Mental Demand

Controller performance measures utilized in the experiment showed that controller metering performance was increased with the TACT tools. For one such measure, delivery accuracy with respect to the assigned metering fix time was tracked, and showed that controllers delivered aircraft on average more accurately with TACT tools than without. This is illustrated in Figure 11 below, where the Y-axis is in minutes and zero represents accurate delivery. Accuracy was greatest with the MDM tool, followed by the MDB and Baseline tools, respectively.

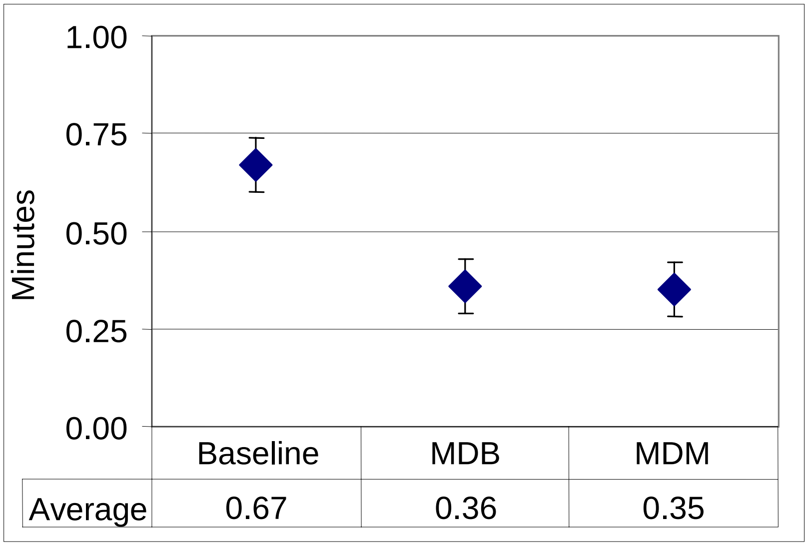

Figure 11. Metering Aircraft Delivery Accuracy

A repeated measures analysis of variance test showed the effect of the TACT and Baseline tools on delivery accuracy to be statistically significant, F(2, 16) = 9.65, p=0.002. Significant differences among the tools were obtained through a Tukey Honestly Significant Difference (HSD) test. The Tukey HSD test showed that there were significant differences between the MDM tool and the Baseline tool, as well as between the MDB tool and the Baseline tool, q=2.58, alpha=0.05. The power statistic for this test was 0.7906. Furthermore, a contrast between the MDM tool and the Baseline tool yielded a statistically significant difference, F=(1,16) = 15.26, p=0.001. Finally, a contrast between the MDB tool and the Baseline tool yielded a statistically significant difference, F=(1,16) = 13.63, p=0.002. No statistically significant differences were noted between the MDM and MDB tools [DeSenti et al., 2001].

Subjective feedback from controller participants indicated that they also perceived themselves as delivering more accurately with the MDM and MDB tools in the same rank order.

As an indicator of operational acceptability a CARS was completed by each controller subject, after each evaluation run. The CARS instrument permitted the controller to rate the usability and acceptability of the TACT tools and the Metering List in an operational setting. The acceptance scale ranged from 1 (unacceptable, improvement is mandatory) to 10 (desirable with no improvements needed). One way to gauge acceptance is to refer to the CARS results for each tool individually, as a specified level of acceptance corresponds to the answers provided. The average response was in the acceptable range for the MDM, MDB and Baseline. Another, perhaps more compelling, way to gauge acceptance is to make a comparison with the Baseline. The Baseline Metering List was designed to emulate the look, feel, and general functionality of the Metering List employed with TMA. The TACT tools ranked above the Baseline in terms of operational acceptance [DeSenti et al., 2001].

Another perspective on acceptance relates to tool effect on the general acceptance of the metering task itself. A question on the post-run feedback interview form related to acceptability of time-based metering. On average, the controllers felt that the MDM and MDB tools make the task of time-based metering more acceptable. On a 100-point scale where 0 represents "Very Hard to Accept" and 100 represents "Very Easy to Accept," the average ratings for the MDM was 84.3, the MDB was 78.4, and the baseline was 68.3.

Summary

Several years of laboratory research has revealed that TACT visualization tools provide a clear advantage where controllers perform time-based metering. CAASD continues to study how the TACT tools can work in partnership with given time-based metering schedules.

A CAPACITY ENHANCEMENT PROCEDURE FOR CONVERGING RUNWAY CONFIGURATIONS

Problem Being Addressed

Chicago O'Hare (ORD) can use many runway configurations. One of the preferred configurations uses runways 9L, 9R, and 4R for arrivals. This configuration is shown in Figure 12. During visual meteorological conditions (VMC), this configuration can be used to support over 100 arrivals per hour. When weather conditions fall below VMC, the triple runway configuration cannot be used because protection must be provided in case of simultaneous missed approaches on the converging runway and one of the parallel runways. Below VMC, therefore, the facility reverts to a two runway operation. According to arrival rates used by ATC System Command Center, Runways 9L and 4R may be used down to a ceiling of 700 feet and visibility of 2 miles with a rate of 80 arrivals per hour [FAA, www.fly.faa.gov]. However, below 700 feet and 2 miles, the only available configuration is the parallel runway option: in this case runways 9L and 9R. When such a simultaneous parallel runway operation is used, the facility can support an arrival rate of 68. It also requires the use of two more controllers to monitor the operation. Since the airlines generally schedule flights assuming that the conditions will be VMC, a reduction in the arrival rate from over 100 to 68 will cause significant delays.

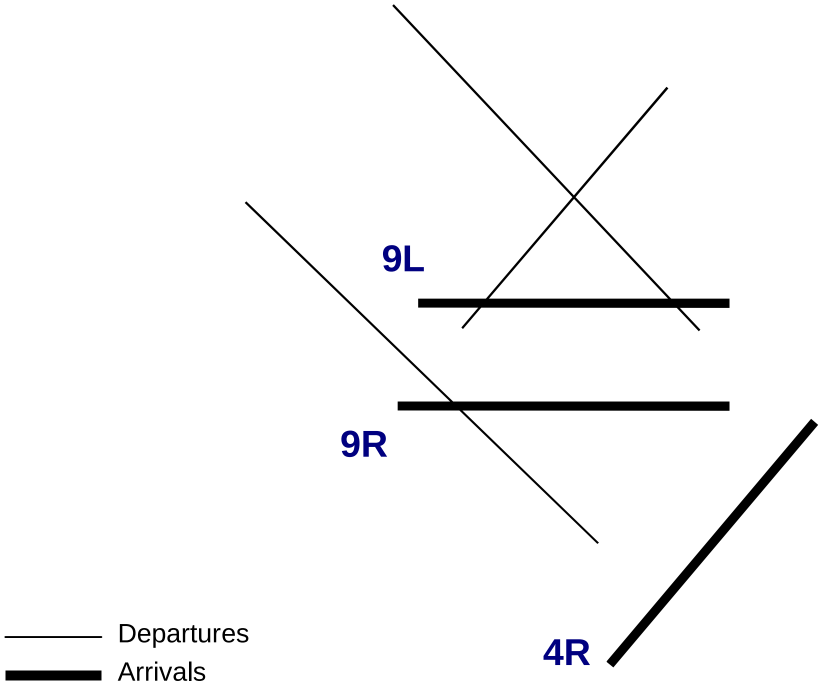

Figure 12. The 9L/9R/4R Configuration at Chicago O'Hare

If all three runways could remain open in instrument meteorological conditions (IMC), the arrival rate could be increased over that of the 9L/9R pair.

Research

A procedure called Dependent Converging Instrument Approaches (DCIA) was authorized in the US in 1992 in the FAA Air Traffic Order 7110.110 [FAA, 1992]. It authorizes the use of approaches to converging or intersecting runways in IMC and establishes the conditions for conducting such operations. One key requirement that makes this operation possible is the staggering of arrivals on the converging approaches by specific minimum values. These values are specified in the order by groups of runway geometries. They are designed such that even in certain worst conditions, even if both aircraft on the converging approaches missed at the same time, full safety is maintained without any controller intervention or any special pilot techniques. The DCIA procedure is currently in use at a few airports, notably at St. Louis International (STL). In use since 1992, it enables STL to provide a rate of about 44 arrivals per hour in IMC, compared with a single stream arrival rate of 36 [FAA, www.fly.faa.gov]. Before the introduction of DCIAs, STL was forced to operate the airport on a single-stream arrival basis in IMC, causing its hub airline to cancel flights regularly.

In principle, any airport with converging or intersecting runways is eligible for DCIAs. However, the capacity benefit that DCIAs provide depends on the minimum stagger values required for the eligible runway configurations. The minimum stagger requirements are governed primarily by the lengths of the runways or their extended centerlines to the point of intersection (called the common point). The longer these runway lengths are to the point of intersection, the greater the required stagger to provide adequate safety. If an aircraft with a low final approach speed is followed by one with a high final approach speed, a larger stagger is also required. The greater the distance to the common point is, the greater the distance the fast trailing aircraft can make up, and therefore, the larger the required stagger. When the stagger value exceeds 2.5 nmi, most airports derive no significant capacity benefit from the DCIA operation, except perhaps the benefit of providing an option in serving traffic on different runways. Based on Order 7110.110, the runways in question at ORD fall into this category of requiring too great a stagger to be beneficial.

The minimum staggers required by Order 7110.110, however, protect against simultaneous dual missed approaches on the converging approaches regardless of the approach speeds of the aircraft. Suppose the minimum stagger value required for the case of a slow aircraft on one runway followed by a fast aircraft on the converging runway is 2 nmi. The minimum stagger value required for the reverse case, i.e., for a fast aircraft followed by a slow aircraft on the same converging runway may only be 0.5 nmi to provide the same degree of safety. The current DCIA order requires *all* aircraft pairs to be staggered by at least 2.0 nmi. This is reasonable in the current system since the final approach speeds of the aircraft are generally not known by the controller. The procedure proposed in this paper offers the potential for an increased operations rate in IMC for this configuration.

Since the stagger requirements to ORD runways 9L/4R dominate the stagger requirements, Figure 13 is included to show the stagger surface (i.e., the minimum stagger requirements) for runway pair 9L and 4R. This surface was developed using the same logic on which the stagger values in Order 7110.110 are based. A simulation of two aircraft performing missed approaches was run, assuming that the leading aircraft does not accelerate and the trailing aircraft does. Within limits, the wind is chosen to minimize the separation of the aircraft at the intersection.

The stagger is adjusted until this minimum separation is 1 nmi (for a non-heavy leading aircraft) and 76 seconds for a leading heavy aircraft.

Figure 13 shows that the stagger requirements for converging operations to runways 9L/4R can vary anywhere from 0.5 to 2.8 nmi, depending on the expected final approach speeds of the two aircraft. Figure 13 also shows that if the approach speeds are nearly equal, a stagger of only about 1.4 nmi is required.[11]

Of course, if the runway geometry were different (i.e., the included angle between the runways and the distances from the runway thresholds to the intersection point), then a different surface would result. The general shape of the surface will show a maximum at the low leader speed and high trailer speed sloping down to a minimum at the high leader speed, low trailer speed point. The structure in the surface is due to the missed approach climb out model that was used. With shorter distances of the runway thresholds to the intersection point, the surface would be smoother because the effect of the climb out model would not come into play. It turns out in the case of runways 4R and 9L at ORD, the distances to the intersection point from the respective runway thresholds is about the same so the surface for 9L leading 4R would be the same as shown in Figure 13.

---

[11] The maximum stagger value of 2.8 nmi for this configuration is less than that implied in Order 7110.110. This is because Order 7110.110 provides stagger values for groups of runways of different lengths and different included angles. Computation of the stagger required for any particular configuration can result in smaller values, especially when the geometry is favorable. The included angle for ORD's 9L/4R configuration is very favorable. Hence the lower stagger value 2.8 nmi in this chart, compared with >3.0 nmi in 7110.110. The stagger value for configuration 9R/4R is found to be 2.4 nmi. Capacity runs with these particular values, 2.4 nmi and 2.8 nmi, show that no significant capacity gain can be achieved with these fixed values for this triple configuration compared with the dual configuration 9L/9R.

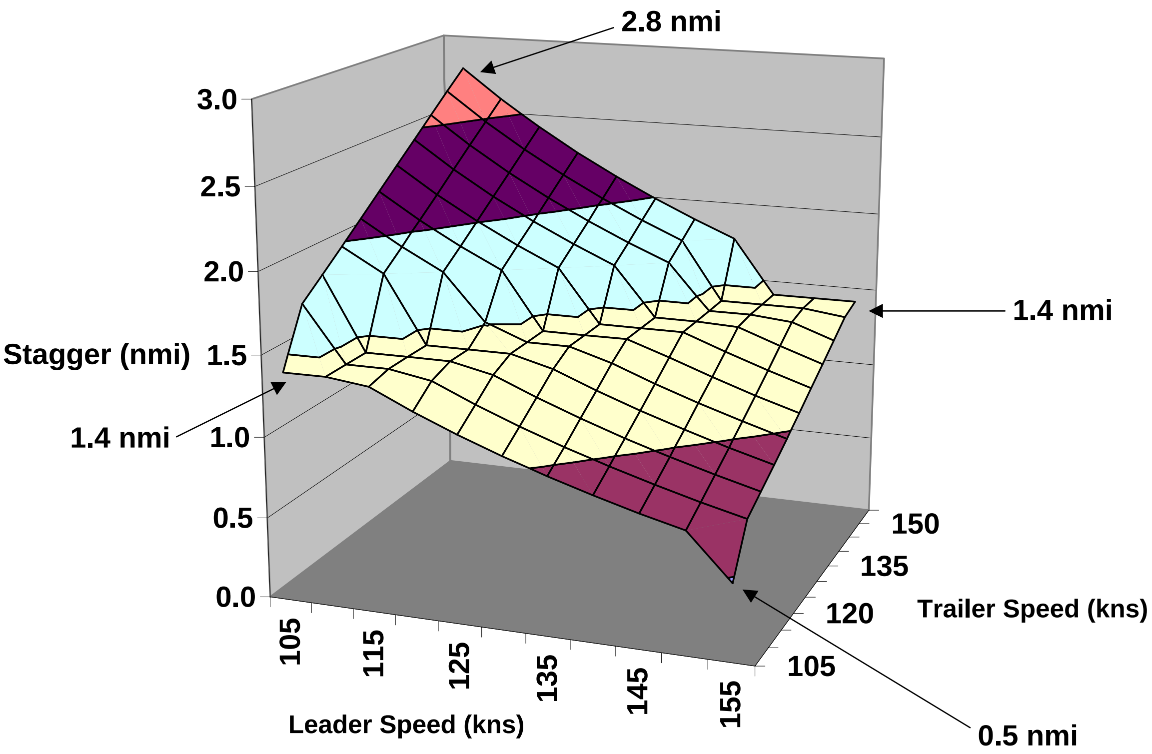

Figure 13. Required Stagger Surface for 4R Leading 9L at Chicago O'Hare

Thus, if the expected landing speeds of the leading and trailing aircraft were known to the ATC system (e.g., the speeds were down-linked from the aircraft) and the controller could set up the traffic streams to have the proper staggers using the FAA's Converging Runway Display Aid (CRDA), then this procedure could yield a significant benefit over that of the DCIA procedure now approved by the FAA.

Details of Study

Figure 14 shows potential arrival capacity values that may be possible with such an operation. The arrival rate values displayed in Figure 14 were computed based on a monte carlo simulation of a string of arrival aircraft. The final approach landing speed of each aircraft was randomly chosen to have an indicated airspeed between 125 knots (kn) and 145 kn representative of the "large" class of aircraft. The distance that each aircraft would be behind another aircraft as the leading aircraft crossed the runway threshold would also be a random value. In the case of the single runway, it would be 3 nmi plus a value between 0 and the Trailing Precision value shown in Figure 14. For the converging runway cases, the distance between the aircraft would be the minimum stagger distance plus a value between 0 and the Trailing Precision. The arrival rate was determined by taking the cumulative time it takes 20 aircraft to fly the separation distances to the runway divided by the number of aircraft intervals (19). Twenty aircraft on each runway were chosen as representative of a typical arrival push. One thousand 20-aircraft strings were simulated to determine the average arrival rate.

Results

To model the independent triple and the independent pair, a single runway arrival rate was computed and then the rate was tripled or doubled, respectively. The dependent triple model used the fact that at ORD the runway pairs 9L and 9R are independent. Therefore, one considers the required staggers between runways 4R and 9R and between runways 4R and 9L. In the simulation there is, in fact, a dependency between 9R and 9L via the staggers with aircraft on 4R. The stagger values from the analysis that produced the stagger surface depicted in Figure 13 were used in this model.

This simulation was run several times, varying the Trailing Precision — the buffer that the controller would introduce for several reasons, such as wind uncertainty or lack of knowledge about how a particular aircraft might decelerate. From results in the laboratory, buffers on the order of ½ to 1 nmi behind ghosts were typical. Matching these curves to experience of landing rates at ORD implies a Trailing Precision between 1.0 and 2.0 nmi.

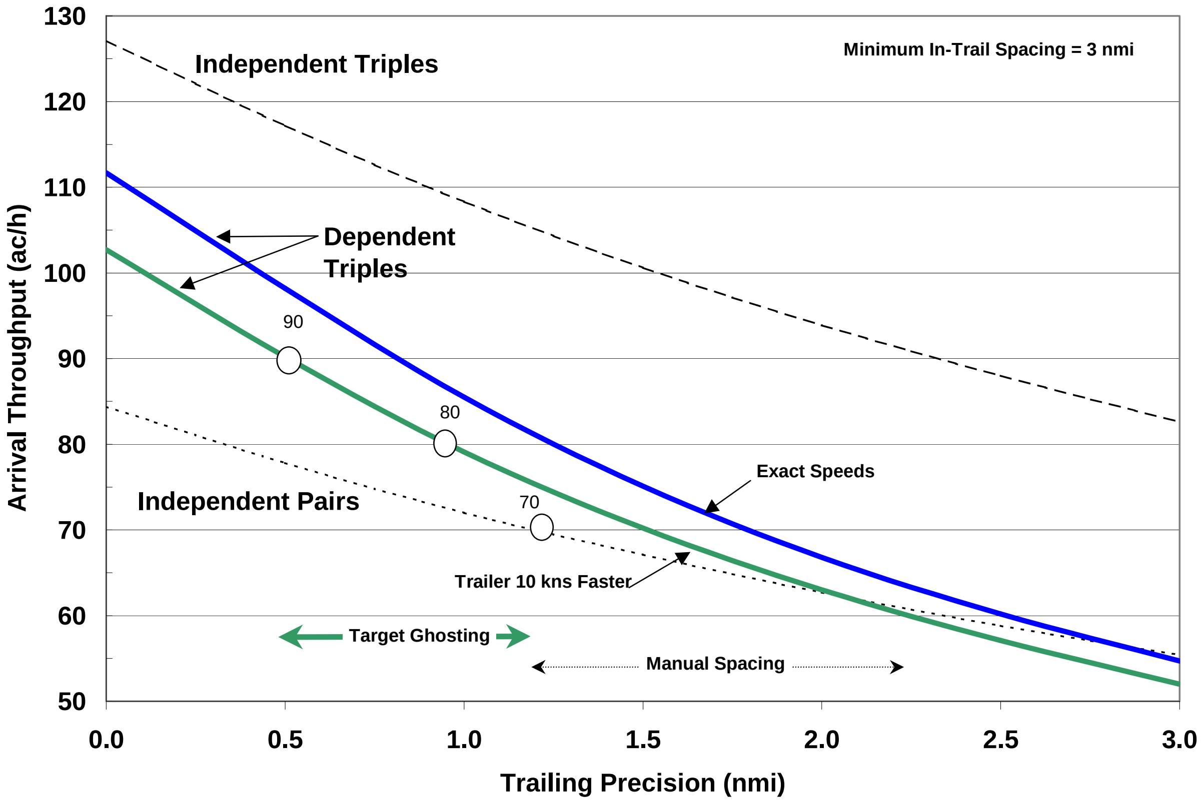

Figure 14. Potential Arrival Rates for Runways 4R, 9L and 9R at Chicago O'Hare

The line labeled "Trailer 10 kn Faster" in Figure 14 provides a buffer of 10 kn in case the aircraft is flown faster on final approach than the pilot initially indicated; i.e., it accounts for a 10 knot error in expected final approach speed values. This figure shows arrival throughput estimates for various configurations of interest. It shows that for operations to an independent pair of runways, assuming a manual spacing accuracy of about 1 to 2 nmi, a rate of about 68 ±5 aircraft per hour might be supported. For an independent triplet of runways, an operation conducted only in VMC, an arrival rate of about 100 ±10 aircraft per hour might accrue. For the operation being proposed here, there will be uncertainty between the speeds down-linked as expected landing speeds, and those actually flown, due to changes in wind values. The procedure must account for such uncertainty. This has been accounted for by assuming the trailing aircraft is really 10 knots faster. Thus, for the dependent triple operation being proposed, and assuming a delivery accuracy of 0.5 to 1.0 nmi (which is expected from experience with ghosting operations), Figure 14 shows that a rate of somewhere between 80 to 90 may be sustainable in IMC to Category I minima.

One of the issues of implementing this procedure is how to obtain and input accurate estimates of the landing speeds of the aircraft into the ATC automation system without causing the workload on the controller or the pilot to be unacceptable. This issue is currently being investigated.

There are at least three technologies currently available or being developed that provide a possible way to down-link planned landing speeds to the ground automation system: ADS-B (Automatic Dependent Surveillance-Broadcast), CPDLC (Controller Pilot Data Link Communications) and ACARS (Aircraft Communications Addressing and Reporting System). Research continues as CAASD explores each of these technologies—each is complex, and therefore is not addressed in any detail in this paper.

Summary

This particular example of the procedure focuses on Chicago O’Hare. However, if at any other airport, the DCIA staggers were to be tailored to that airport and computed as a function of the expected landing speed, similar benefits, yet airport-specific, could be achieved.

CLOSING

These concepts represent a broad cross section of potential improvements to the ATM system that are being investigated and developed at CAASD. These concepts currently reflect ideas that may lead to later implementation but that require further investigation and development to assess their operational feasibility, suitability, and effectiveness. However, they are only a few of the ideas that are being developed at CAASD and throughout the aviation community in consultation with the FAA, airlines, other airspace users, and other stakeholders.

The FAA has comprehensive programs of planned improvements to the ATM system, such as the Operational Evolution Plan, the Research and Development Plan, and the Inter-Agency Integrated Product Team Plan. These plans and others have been developed in cooperation with the entire aviation community and are expected to provide notable capacity and performance improvements. With the implementation of these plans and the continuation of research programs developing into fielded prototypes, such as these presented here in this paper, increased system performance can be achieved.

List of References

Anagnostakis, I., Idris, H.R., Clarke, J.P., Feron, E., Hansman, R.J., Odoni, A.R. and Hall, W.D. (2000), "A Conceptual Design of a Departure Planner Decision Aid," *Proceedings of the 3rd USA/Europe Air Traffic Management R&D Seminar*, Naples, Italy, 13-16 June.

Barrer, J., Swetnam, G.F., Weiss, W.E. (1989), *The Feasibility Study of Using Computer Optimization for Airport Surface Traffic Management*, MTR89W00010, The MITRE Corporation, McLean, VA, March.

Campbell, D. and J. Stanley (1963), *Experimental and Quasi-Experimental Designs for Research*, Houghton Mifflin Company, Boston, MA.

Cooper, W.W., Cherniavsky, E., DeArmon, J., Foster, J.G., Mills, M., Mohleji, S., and Zhu, F. (2001), "Determination of Minimum Push-Back Time Predictability Needed for Near-Term Departure Scheduling Using DEPARTS," *Proceedings of the Fourth International Air Traffic Management R&D Seminar (ATM-2001)*, Santa Fe, New Mexico, USA, 3-7 December.

Cooper, W.W., Mohleji, S., Burke, C., Foster, J.G., Mills, M. (2001), "DEPARTS: A Tool for Improving Airline Departure Scheduling and Reducing Flight Delays at Busy Airports," *Handbook of Airline Strategy: Public Policy, Regulatory Issues, Challenges, and Solutions*, Gail F. Butler and Martin R. Keller, Exec. Ed., First Edition, New York: McGraw-Hill Aviation Week, pp. 577-600.

DeSenti, C., Ball, C., Rowe, D. and Jacobs, G. (2001), "Transition Airspace Controller Tools (TACT) Visualization Aids for Radar Controllers," *Proceedings of the Fourth International Air Traffic Management R&D Seminar ATM-2001*, Santa Fe, NM, December.

FAA, Air Traffic Control Command Center, website, www.fly.faa.gov.

FAA, Free Flight Phase I, website, www.ffp1.faa.gov.

FAA, National Airspace Redesign, website, www.faa.gov/ats/nar.

FAA (1992), FAA Air Traffic Order 7110.110.

Idris, H.R. (1999*), Queuing Dynamics and Control of Departure Operations at Boston Logan*, PhD Thesis, Massachusetts Institute of Technology, August.

Idris, H.R., Delcaire, B., Anagnostakis, I., Hall, W.D., Pujet, N., Feron, E., Hansman, R.J., Clarke, J.P., Odoni, A.R. (1998), "Identification of Flow Constraints and Control Points in Departure Operations at Airport System," AIAA 98-4291, *Proceedings AIAA Guidance, Navigation and Control Conference*, Boston, MA, August.

Kenny, D. (1994), *Statistics for the Behavioral and Social Sciences*, Harper Collins, New York, NY.

Kirk, D., and Bowen, K. C. (2000), *Preliminary Observations About Providing Problem Resolution Advisories to Air Traffic Controllers*, MP-00W186, The MITRE Corporation, McLean, VA, June.

Lampl, R. (Executive Editor) (2000), "The Aviation & Aerospace Almanac – 2000 Edition," *Aviation Week*, Washington, DC.

McFarland, A., and Maroney, D. (2001), *Eliminating the Altitude-for-Direction Rule and Implementing Reduced Vertical Separation Minimum in the U.S.*, MP-01W167, The MITRE Corporation, McLean, VA, September.

Nemhauser, G. L., and Wolsey, L.A. (1999), *Integer and Combinatorial Optimization*, John Wiley and Sons, Inc., New York, New York.

Pujet, N. (1999), *Modeling and Control of the Departure Process of Congested Airports*, Massachusetts Institute of Technology, PhD Thesis, August.

RTCA (2000), *Government and Industry Guidelines and Concepts for NAS Analysis and Redesign*, RTCA/DO-266, Washington, DC, December.

RTCA Select Committee for Free Flight Implementation (2000), *Documentation of 2003-2005 Capabilities Working Group Deliberations and Recommendations, National Airspace System Concept of Operations Addendum 4: Free Flight Phase 2 Appendix B*, RTCA Inc., December 13.

Shumsky, R. A. (1997), "Real Time Forecasts of Aircraft Departure Queues," *Air Traffic Control Quarterly*, Vol. 5(4).

Williams, H.P. (1999), *Model Building in Mathematical Programming*, 4th Ed., John Wiley and Sons, Ltd., Chichester, England.

Acronyms

ACARS	Aircraft Communications Addressing and Reporting System
ADS-B	Automatic Dependent Surveillance - Broadcast
AFD	Altitude-for-Direction
AOC	Airline Operations Center
ARMT	Airport Resource Management Tool
ARTCC	Air Route Traffic Control Center
ASQP	Airline Service Quality Performance
ATC	Air Traffic Control
ATCSCC	Air Traffic Control System Command Center
ATCT	Air Traffic Control Tower
ATL	Atlanta Hartsfield International Airport
ATM	Air Traffic Management
CAASD	Center for Advanced Aviation System Development
CARS	Controller Acceptance Rating Scale
CDM-Net	Collaborative Decision Making - Network
CPDLC	Controller Pilot Data Link Communications
CRDA	Converging Runway Display Aid
DCIA	Dependent Converging Instrument Approaches
DDTC	DataLink Delivery of Taxi Clearance
DEPARTS	Departure Enhanced Planning and Runway/Taxiway-Assignment System
DSP	Departure Sequencing Program
ETMS	Enhanced Traffic Management System
FAA	Federal Aviation Administration
FFP1	Free Flight Phase I
FFP2	Free Flight Phase II
FFRDC	Federally Funded Research and Development Center
FL	flight level
GDP	Ground Delay Program
HCS	Host Computer System
IFR	Instrument Flight Rules
IMC	Instrument Meteorological Conditions
MDB	Mileage in the Data Block
MDM	Mileage Distance Marker
MIP	Mixed Integer Program
MIT	miles-in-trail
OOOI	Out, Off, On, In
ORD	Chicago O’Hare
SMA	Surface Movement Advisor
SMS	Surface Management System
STL	St. Louis Lambert Field
TACT	Transition Airspace Controller Tools
TFM	Traffic Flow Management
TLX	NASA Task Load Index
TMA	Traffic Management Advisor
TMC	Traffic Management Coordinator
TRACON	Terminal Radar Approach Control
URET	User Request Evaluation Tool
US	United States
VMC	Visual Meteorological Conditions
ZID	Indianapolis Center
ZTL	Atlanta Center

Biography

Michelle J. Blucher, Wayne W. Cooper, Christopher T. DeSenti, Dr. Alvin L. McFarland, Dr. Satish C. Mohleji, and Arthur P. Smith, III are members of the technical staff at The MITRE Corporation's Center for Advanced Aviation System Development (CAASD) in McLean, VA.