

A CASE FOR INTEGRATING THE CTAS TRAFFIC MANAGEMENT ADVISOR AND THE SURFACE MANAGEMENT SYSTEM

Stephen C. Atkins
NASA Ames Research Center
Mail Stop 210-6
Moffett Field, CA 94035-1000

William D. Hall
C. S. Draper Laboratory, MS #3F
555 Technology Square
Cambridge, MA 02139

Abstract

Arrival and departure capacities are interdependent at many high-traffic airports. At these airports, total capacity may be dynamically reallocated between arrivals and departures in response to the time-varying demands for both types of operations. Moreover, controllers appear to use a working understanding of the tradeoff between arrival and departure capacities in daily practice. This paper investigates how capacity allocation decisions are currently made, the degree to which and the mechanisms by which arrival and departure capacities are controllable, and whether automation could help controllers better match arrival and departure capacities to time-varying demands. Furthermore, this paper studies how the interdependence of capacities affects decision support tools such as the Traffic Management Advisor (TMA), which schedules arrivals subject to the arrival capacity, and the Surface Management System (SMS), which manages departures on the surface. To this point, TMA and SMS have been developed independently. However, at airports where capacities are interdependent, significant benefit may be achieved by integrating TMA and SMS to coordinate how limited airport resources will be shared. A decision aid to help controllers dynamically reallocate total airport capacity between arrivals and departures may achieve TMA-SMS interoperability. Two possible decision aids are suggested: one which provides arrival and departure demand information not currently available, and one which advises an efficient sequence of capacity allocations, possibly determined collaboratively by air traffic control and the air carriers.

Introduction

The number of departures that an airport can accommodate in a given period of time depends, in general, on the number of aircraft that arrive during that

period. This characteristic is widely known. For example, Gilbo [1] demonstrated empirically that a tradeoff between arrival capacity and departure capacity exists in certain configurations at certain airports, and the FAA's Engineered Performance Standards present the achievable hourly arrival rate for an airport as a function of the configuration and the hourly departure rate. At these airports, total capacity may be dynamically reallocated between arrivals and departures in response to the time-varying demands for both types of operations. Moreover, the authors' interviews of air traffic controllers have revealed that controllers use a working understanding of the arrival-departure tradeoff in daily practice. This paper presents an initial investigation of the efficiency of controllers' runway use decisions. To be efficient, interdependent arrival and departure capacities must be coordinated. Moreover, strategic traffic management (e.g., TMA and SMS) requires capacities to be planned in advance.

NASA Arrival and Departure ATM Tools

NASA Ames Research Center, in cooperation with the FAA, is developing a suite of decision support tools (DSTs), collectively known as the Center-TRACON Automation System (CTAS), to assist controllers in managing aircraft more efficiently [2]. In many terminal airspaces there is limited ability to substantially delay arrivals without unacceptably high controller workload. Therefore, large delays must be absorbed in en-route airspace (or through ground delays) prior to aircraft entering the terminal area. The Traffic Management Advisor (TMA), one element of CTAS that the FAA is fielding nationally under the Free Flight Phase 1 program, computes an efficient schedule for arrivals to a busy airport, for the purpose of metering the arrivals into the terminal area (starting when arrivals are approximately 20 minutes away). TMA schedules arrivals to not exceed capacity parameters set by the Traffic Management Coordinators (TMCs), such as the Airport Acceptance Rate (AAR), and to not violate aircraft separation requirements at the meter fixes, final approach fixes and thresholds. To achieve this planned flow, TMA provides information and advisories to TMCs in the Center and TRACON and to the sector controllers in the Center [3].

Copyright © 2000 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

As with arrivals, there is often limited controllability in the terminal airspace to substantially delay or re-sequence departures after they are airborne. Instead, the ability to sequence departures and absorb significant delays exists on the airport surface. The Surface Management System (SMS) is being developed by the NASA Ames Research Center, in cooperation with the FAA, to help controllers manage departures prior to takeoff (as well as arrivals on the surface). Although a departure capacity parameter similar to the AAR is not yet well defined, SMS will plan departures to maximize departure throughput for a given amount of resources, and incur necessary delays efficiently. To do so, SMS will provide information and advisories to FAA tower controllers as well as provide information to and collect information about demand and preferences from the air carriers, interacting with both the centralized Airline Operations Control (AOC) centers and the local ramp towers/stations [4].

TMA-SMS Interoperability

To this point, TMA and SMS have been developed independently. However, since TMA and SMS manage traffic to use the capacities available to them efficiently, at airports where arrival and departure capacities are interdependent the two tools must be interoperable (i.e., the arrival and departure capacities to which the tools schedule must be coordinated). In addition, each tool must know what capacity will be available to it far enough in advance to efficiently manage aircraft. A controller decision aid to help plan coordinated arrival and departure capacities could achieve TMA-SMS interoperability. One possible decision aid would provide arrival and departure demand information not currently available, enabling controllers to plan coordinated arrival and departure capacities. An alternative decision aid would advise an efficient sequence of capacity allocations, possibly determined collaboratively by air traffic control (ATC) and the air carriers.

Note that TMA-SMS interoperability represents a macroscopic approach to arrival-departure interoperability (i.e., aggregate arrival and departure rates are planned without considering the interactions between individual aircraft). Consequently, this approach assumes that the TRACON and tower controllers will be able to resolve timing miscues between individual arrivals and departures at the runway. For example, at New York's LaGuardia airport, arrivals and departures are frequently mixed on the same runway, requiring controllers to control arrivals to leave gaps for departures. An interoperable TMA-SMS would plan arrival and departure rates and deliver smooth traffic flow (similar to what TMA currently does for arrivals).

However, TRACON and tower controllers will still be required to manually sequence arrivals and departures. If the controllers are unable to construct gaps for departures, the departure rate will be less than the planned capacity. Future work will investigate whether interactions between individual aircraft must be considered when planning efficient arrival and departure capacities. Automation such as the Final Approach Spacing Tool (FAST), the Expedite Departure Path (EDP) tool (both elements of the CTAS suite), and tactical surface movement functionalities of SMS could help controllers coordinate the timing of individual aircraft (i.e., properly space arrivals to leave gaps for departures, and release departures into those gaps) to achieve the planned arrival and departure capacities [5].

By not considering the exact sequence of operations, an interoperable TMA-SMS may not accurately know the overall airport capacity. To be robust to uncertainties about capacity, an interoperable TMA-SMS could plan and manage arrivals and departures to maintain constant pressure on the constrained resources. Similarly, TMA does not absorb all of the delay which it expects the arriving aircraft to incur before landing. Rather, it meters aircraft to maintain constant pressure on the TRACON and airport [2].

Overview

The next section discusses difficulties in modeling the combinations of arrival and departure capacities that are feasible at an airport at a particular time, illustrating the limitations of an established method for describing the arrival-departure interdependence. The subsequent section studies how the arrival-departure mixture is currently determined, and demonstrates the effect of inefficient capacity allocation decisions through an analytic example. The paper then discusses the need for air carrier participation in setting the arrival and departure capacities, and outlines a possible algorithm for ATC-air carrier collaboration. This algorithm could serve as the foundation for a decision aid to enable planning coordinated arrival and departure capacities.

Arrival-Departure Interdependence

Arrivals and departures are interdependent when they compete for limited resources, which occurs, for example, when they share runways or the runway or airspace geometry is such that operations on different runways must be coordinated.

Figure 1 illustrates the flow of arrivals and departures at Boston's Logan Airport (BOS) in the 22L-27, 22L-22R configuration. When the departure demand does not fill

the capacity of runway 22R, arrivals can cross runway 22R in the naturally occurring gaps in the departure traffic. Otherwise, departures on 22R must be stopped while arrivals on runways 22L and 27 taxi across the departure runway to reach the terminal. Since there is limited space on the taxiways to hold arrivals in this configuration, when the airport is under pressure to operate both arrivals and departures, an increase in the number of arrivals requires a decrease in the number of departures, and vice versa. References [6] and [7] discuss other resources that constrain arrivals and departures at BOS.

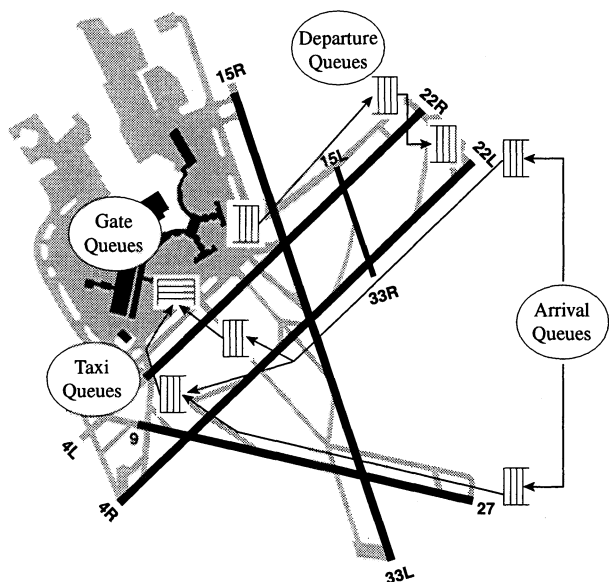


Figure 1. BOS arrival and departure flows for configuration 22L-27, 22L-22R.

The resulting relationship between arrival and departure capacities is illustrated by the *arrival-departure capacity curve*; Figure 2 shows a typical example. Each point in the arrival-departure plane represents a combination of an arrival rate and departure rate. Points on or underneath the curve represent feasible *operating points* for the airport, while points outside the curve represent user demand in excess of the airport's capacity.

The example curve is horizontal at point A where it intercepts the arrival axis, indicating that constraints restricting the arrival rate are locally independent of departure operations. The segments between points B and C and between C and D indicate that additional departures can be achieved by reducing the arrival rate. The slopes of the segments indicate the tradeoff that must be made – how large a reduction in the arrival rate is required to achieve a certain increase in the departure rate [1]. Beyond point D, further decreases in the

arrival rate do not allow any increase in the departure rate. The extent of the arrival-departure interdependence varies across airports and runway configurations. When the arrival and departure rates are independent, the arrival-departure capacity curve forms a rectangle with the axes. If point 3 represents the total demand for arrivals and departures, then any point on the curve between 1 and 2 would be efficient with respect to not wasting airport capacity. Which operating point is best depends on a variety of factors, such as the user's preferences, how demand varies in time, and how frequently the operating point can be adjusted. Note that the airport's arrival-departure capacity curve also varies in time.

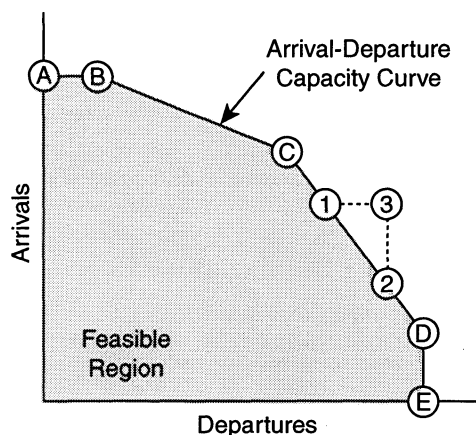


Figure 2. Typical Arrival-Departure Capacity Curve.

In complex airspaces (e.g., the New York metropolitan area), the capacities at proximate airports can be dependent. If operations at one airport affect the capacities available to another airport, the arrival and departure capacities at each of the dependent airports must be determined simultaneously. A higher dimensional arrival-departure capacity curve may be used to visualize the tradeoffs between dependent airports. Furthermore, capacity curves may be a useful mechanism for visualizing the necessary tradeoffs resulting from airborne capacity limitations (e.g., between two streams of flights through a congested sector).

Controllability of Arrival and Departure Capacities

This section identifies several issues concerning the controllability of arrival and departure capacities. Understanding what operating points are feasible at any given time is necessary both to evaluate how efficiently controllers currently allocate airport capacity and to design automation to advise a sequence of operating points.

Figure 3 plots the operating points that occurred at Chicago O'Hare (ORD) during December 1998 and January 1999. Each circle represents the numbers of arrivals and departures that occurred during a 15-minute interval, reported by Enhanced Traffic Management System (ETMS). Gray-scale is used to depict the number of times the airport operated at that combination of arrival and departure rates; darker points signify operating points that occurred more often.

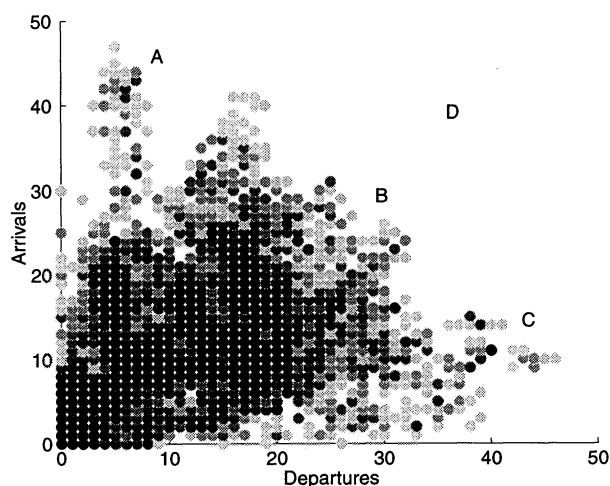


Figure 3. Numbers of arrivals and departures that occurred in 15-minute intervals at ORD during December 1998 and January 1999.

Notice that a smooth arrival-departure capacity curve such as that drawn in Figure 2 is a simplistic representation of the points at which an airport will operate. Although the theoretical arrival-departure capacity curve is a useful model, it does not completely describe the realizable operating points. Furthermore, the actual number of operations that are realized during a period of time depends on the sequence of the individual aircraft, as well as buffers controllers and pilots add to the minimum separation requirements. Since the airport might operate beyond a sustainable traffic rate for short periods of time, additional research is required to understand to what combinations of arrival and departure rates controllers or automation (e.g., TMA and SMS) should schedule.

Airports can operate in a finite number of configurations, where multiple distinct *operating modes* may be possible within each configuration.* Each

**Operating modes* are distinct ways in which runways are used within a configuration. For example, the departure split (i.e., the mapping of departure fixes to runways) is varied at Atlanta Hartsfield airport (ATL) to balance departure demand between the departure runways. Since different modes, like different configurations, can have distinct

configuration/mode will have a distinct arrival-departure capacity relationship. At any point in time, a subset of these configurations/modes will be feasible, depending on weather and noise-impact considerations. Therefore, the arrival and departure capacities are determined by first selecting a configuration/mode from those feasible at that time, and then choosing an operating point realizable under that configuration.

Arrival-departure capacity curves are typically drawn from historical observations using Airline Service Quality Performance (ASQP) or ETMS data sets. A limitation of this data is the absence of airport configuration information. Figure 3, for example, plots data for every 15-minute interval during the two-month period, combining operating points that occurred under different configurations. Therefore, the arrival and departure capacity pairs that are available at a specific time may be a subset of the figure. Although historical data for specific configurations/modes is less readily available, there are no fundamental reasons plots like Figure 3 cannot be drawn for individual configurations/modes.

The distribution of points in Figure 3 likely resulted from variations in both demand and capacity allocation. Did ORD not operate at point D (see Figure 3) because the demands for arrivals and departures were never simultaneously that high, or because the airport's overall capacity limited operations to points A, B, and C? Empirically derived arrival-departure capacity plots cannot distinguish between limited capacity and a lack of demand.

Figure 4 plots arrival-departure rates observed at Dallas/Fort Worth airport (DFW) during December 1998 and January 1999. With abundant runways and TRACON airspace, arrivals and departures at DFW are procedurally separated. For example, in a North flow configuration, aircraft depart using runways 36R, 35L, and 31L, while arrivals land on runways 36L, 35C, 35R, and 31R. Although arrivals must taxi across the departure runways (and departures on 31L must taxi across a departure and an arrival runway), the interruption to departures is minimized by queueing arrivals on the multiple taxiways that cross the departure runways and then crossing several aircraft in parallel during the gap between two departures. Consequently, arrival and departure capacities at DFW are largely independent, implying the arrival-departure capacity curve would be expected to be rectangular. Although the data in Figure 4 shows a tradeoff between the arrival and departure rates that were observed, it

arrival-departure relationships, this paper does not make a distinction between modes and configurations.

does not disprove independence of arrival and departure capacities. DFW may not experience periods of time during which the demands for arrivals and departures simultaneously stress the airport's resources. This explanation is feasible since the airport is dominated by a single air carrier operating a large hub. Since arrivals and departures are concentrated in banks, arrival and departure demands could be simultaneously large only if a delayed departure bank overlaps the next arrival bank. Future work will combine demand information with operation counts, to better understand decisions concerning the arrival-departure mix.

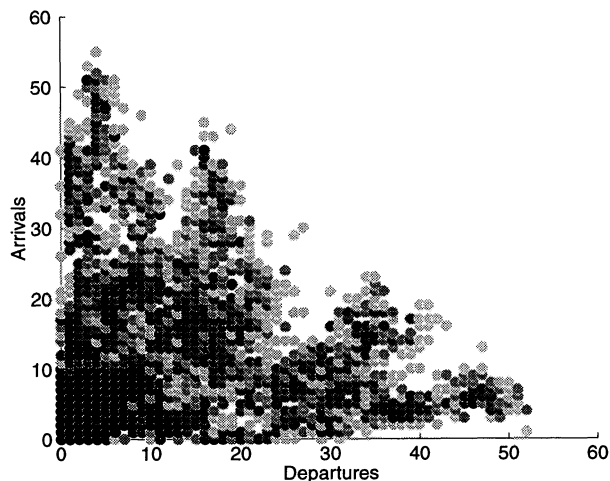


Figure 4. Numbers of arrivals and departures that occurred in 15-minute intervals at DFW during December 1998 and January 1999.

Figure 5 plots numbers of arrivals and departures that occurred during 15-minute intervals as functions of time (source ETMS, DFW, 1 December 1998), showing the bandwidth of the operating point. Note that if the discrete time bins are too large, significant details of the arrival-departure interactions may get hidden. Separate weather data shows that meteorological conditions were favorable, suggesting a high capacity configuration was likely being used and there may have been no configuration/mode changes during this time period. At approximately 13:00 the demands overlap; again, the data does not reveal when operations are limited by airport capacity versus demand. For example, where both the arrival and departure rates are moderate, the data does not show whether additional demand for one or both types of operations existed, and controllers chose to allocate capacity in this way, or whether the airport was accommodating the full demand for each. Future work will incorporate demand information to better illustrate how controllers mix arrivals and departures when demands are simultaneously high. Figure 5 also illustrates the variety in how the operating

point moves in the arrival-departure space. Around 13:00, the operating point transitions from the lower-right corner of the arrival-departure plane to the upper-left (see Figure 4) by moving along the envelope of observed operating points. The subsequent shift from arrival to departure emphasis at 13:30 goes through the interior of the feasible space.

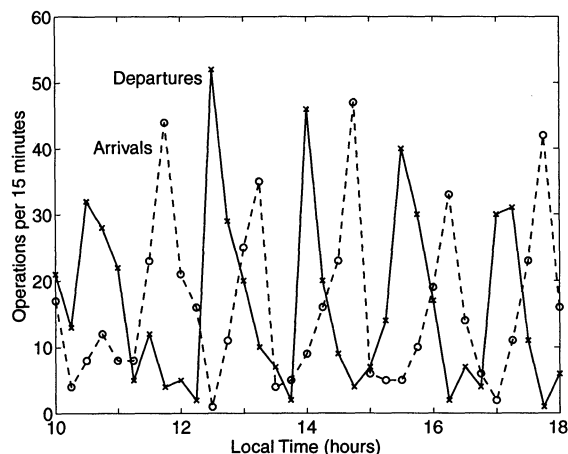


Figure 5. Numbers of arrivals and departures that occurred in 15-minute intervals at DFW on December 1, 1998.

How rapidly the airport may switch its configuration/mode must also be understood when planning arrival and departure schedules. When the end of an arrival rush overlaps the start of a departure push, for example, the configuration/mode may need to change quickly. Observations conducted at Boston's Logan airport (BOS) indicate that the configuration/mode may be changed rapidly under some conditions, whereas under other conditions changing the operating mode requires advanced planning [6]. Furthermore, configuration and mode changes can temporarily interrupt operations, reducing capacity, especially when not planned in advance. Radar data from DFW, from which the configuration being used can be identified, suggests that configuration changes (e.g., from North to South flow) are typically planned to occur during gaps in demand. Finally, a variety of factors (e.g., controller workload, the lead time required, or the operations lost during the transition) may limit how frequently the configuration/mode can be changed. Although arrival and departure rushes typically last an hour or more, runway balancing at ATL, for example, might benefit from more frequent runway split changes. Consequently, when planning a schedule for the airport configuration/mode, each of these issues – how far in advance configuration/mode changes must be planned, how much capacity will be lost during the transition, and how frequently changes may be made – must be considered.

Finally, the empirical data shows the achieved operating points after controllers have made decisions about arrival and departure capacities, not the operating points that were available prior to their decisions. In summary, plots of observed operating points do not by themselves reveal the realizable operating points. Consequently, we do not yet know how to draw theoretical arrival-departure capacity curves to model the set of operating points to which interoperable arrival-departure management can plan.

Capacity Allocation

Typically, airport capacity is allocated between arrivals and departures indirectly; an Airport Acceptance Rate (AAR) is specified, which implicitly defines the departure capacity. Although the tower supervisor or Traffic Management Unit (TMU) may consider historical departure demand when manually choosing the AAR, the accurate departure demand information needed to predict the consequence of the allocation decision on departure delays is not currently available. Consequently, this approach tends to favor arrivals, forcing departures to be opportunistic. Although being opportunistic with departures avoids wasting airport capacity, this approach does not necessarily efficiently balance arrivals and departures.

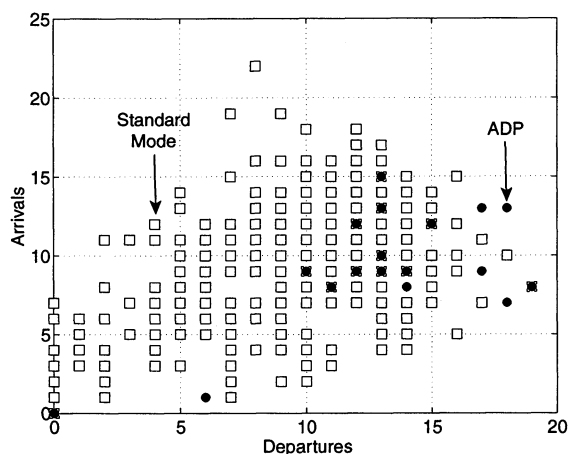


Figure 6. Numbers of arrivals and departures occurring in 15-minute intervals at BOS in configuration 22L-27, 22L-22R.

For example, within the BOS 22L-27, 22L-22R configuration, controllers may either mix arrivals and departures on runway 22L (standard mode) or use 22L exclusively for departures (accelerated departure procedure). Figure 6 plots counts of the arrivals and departures that occurred during 15-minute intervals when the airport was operating in one or the other of these modes. A square indicates the standard mode,

while a filled circle indicates the accelerated departure procedure (ADP). The data was derived from ETMS and air traffic control tower records over a 15-day period by Idris [4]. Under the standard mode, tower controllers cannot plan departures. Planning departures (or arrivals), whether done by a controller or automation, requires knowing the capacity that will be available in the future. From the tower controller's perspective, arrivals to 22L arrive whenever they do and have priority over departures. Consequently, the tower controllers have limited control over the arrival/departure mixture at which the airport operates. Almost any combination of arrival and departure rates on (or under) the arrival-departure curve for this configuration may occur, as seen in Figure 6, depending on the arrival demand and TRACON. Under the ADP, the TRACON must send all arrivals to runway 27, allowing the tower controllers to plan a higher departure rate (at the expense of arrival capacity). In Figure 6, the operating points observed under ADP are clustered at the higher departure rates.

Arrival and departure capacities should be coordinated and planned in advance. The AAR and operating modes such as the ADP are two of the existing mechanisms by which the arrival-departure mix can be controlled. However, planning arrival and departure capacities, whether done by controllers or automation, requires knowing the demands for both types of operations, as well as the overall constraints in the arrival-departure plane. Arrival demand can be predicted accurately from radar surveillance (CTAS currently displays arrival demand in the towers at DFW). However, departure demand forecasts are currently unreliable; the air carriers themselves do not have good predictions of the times at which their flights will be ready to push back. The Surface Management System (SMS) will generate improved departure demand information to support departure management and arrival-departure interoperability.

One approach to coordinating arrival and departure planning is to provide controllers with information about the future demands for both types of operations. Alternatively, automation could use this information to predict the delay impact of trial allocation plans, or to advise an efficient schedule for capacity allocation.

Inter-facility Coordination

The airport tower, TRACON, and Center may need to coordinate the arrival and departure capacities. Arrival capacity is typically set by the airport tower supervisor/TMU and communicated to the TRACON and Center TMUs, so that the Center and TRACON can

delay arrivals as necessary. However, either the TRACON or Center can reduce the arrival rate below the airport's capacity, for workload or other considerations. If this occurs, the airport tower supervisor/TMU may receive little or no advanced warning that the allocated arrival capacity will not be filled, possibly wasting capacity that might have been used for additional departures.

The airport runways are frequently the constrained resource that limits arrival throughput. Although the runways may also limit the departure throughput, downstream airborne issues (e.g., merging aircraft over the departure fixes and into en-route airspace) may be more restrictive. In this case, the arrival and departure capacities are set by controllers in different facilities. The Center TMU sets the departure capacity, typically in terms of miles-in-trail restrictions on consecutive departures to certain airports or to the same departure fix. These constraints are communicated to (i.e., imposed on) the TRACON TMU, and from there to the airport tower supervisor/TMU. Given departure restrictions, the tower supervisor/TMU must choose the appropriate airport operating mode. If the airport's departure capacity exceeds the downstream capacity set by the Center, airport capacity that might be used for additional arrivals will be wasted. Graphically, when airborne considerations limit departure throughput, a region of the arrival-departure capacity curve may not be achievable. An arrival-departure capacity curve may be drawn for the extended terminal area, which would lie under the curve representing the capacity of the airport alone.

Analytic Queueing Example

Arrival and departure capacities should also be adjusted dynamically, since the demands for both types of operations are not constant. Controllers currently change the airport operating mode, to reallocate capacity between arrivals and departures, based on information that they observe directly, such as the length of the departure queue. However, this is only effective when little planning is required to implement a change. Moreover, controllers may respond late, after a large queue has already formed. This section uses a single analytic example to illustrate the effect of the airport responding late to shifts in demand. Again, improved information about future arrival and departure demands is necessary to improve planning of operating mode changes.

The example considers an arrival rush of 80 aircraft, followed by a departure push of 80 aircraft, in a period of 3 hours. Although the next arrival rush may closely follow these departures, the example isolates a single

shift in demand. Clustering of arrivals and departures can be observed at most busy U.S. airports, due to "hub and spoke" flight schedules and time-of-day travel tendencies. At airports with dependent arrival and departure capacities, the airport operating point (i.e., the allocation of capacity between arrivals and departures) may be adjusted dynamically in response to shifts in the demands for both types of operations. If the demands for arrivals and departures are constant, then a constant airport operating mode may be appropriate.

Figure 7 plots the times at which aircraft join virtual arrival and departure queues (x's denote arrivals and o's denote departures). These times are random samples taken from the arrival and departure demand distributions (also plotted in the figure) assumed to be normally distributed with equal variance (20^2 minutes²) and means at 60 minutes and 120 minutes, respectively.

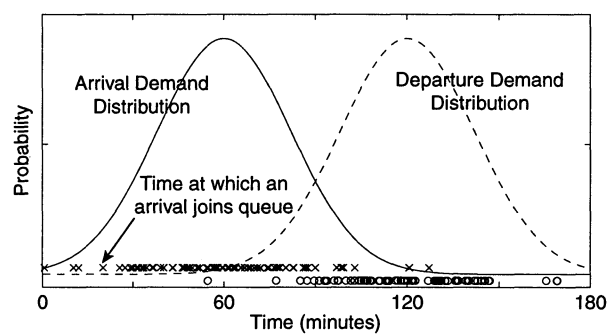


Figure 7. Arrival and departure demands.

Assume the airport can operate in two modes: one mode gives preference to arrivals, permitting 60 arrivals but only 30 departures per hour, while the other mode gives preference to departures, permitting only 30 arrivals but 60 departures per hour. These modes may be achieved through distinct runway configurations or through different schemes for using the runways within a configuration. By ignoring the interactions between arrivals and departures, this example studies arrival/departure interoperability in aggregate. Tactical planning of the interactions between individual arrivals and departures will be the focus of future work. Strategic planning of average capacities avoids the well-studied optimization problem of minimizing inter-operation times by sequencing operations according to aircraft class.

Assume the airport initially operates with the higher arrival capacity to accommodate the arrival rush. At some time during the overlap of arrival and departure demands, the airport switches to favor departures. The controller's task is to decide the time at which to make this switch. Two policies for determining the time to

reallocate capacity from arrivals to departures are compared: the first, labeled “late departure emphasis” in the figures, operates in the higher arrival capacity mode until the arrival queue is empty; the second, labeled “early departure emphasis,” switches to the mode with the higher departure capacity earlier, when the departure queue would otherwise begin to form.

Figure 8 plots the lengths of the arrival and departure queues as functions of time, for the two capacity allocation policies, illustrating the sensitivity of the queues to the policy. If the operating mode is switched too early or too late, capacity may be wasted if departure or arrival demand, respectively, is less than the allotted capacity for some period of time. If arrival and departure demands overlap sufficiently, then arrival and departure delays must be traded against one another. However, Figure 8 shows that a large decrease in departure delays (i.e., queue length) can be achieved at the cost of a much smaller increase in arrival delays.

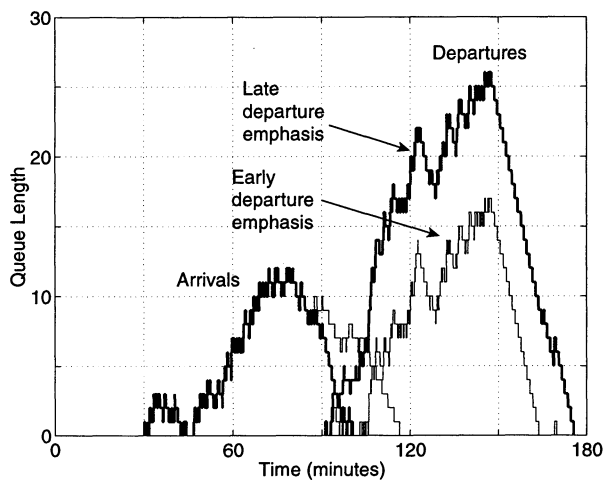


Figure 8. Arrival and departure queues for the two capacity allocation policies.

Figure 9 plots the times at which aircraft enter and leave the queues, for the two policies. Switching the operating mode to emphasize departures at an earlier time delays the formation of the departure queue. By serving a few departures at the start of the departure push earlier, every subsequent departure waits in the queue for a shorter time (Figure 9). However, the earlier reduction in the arrival rate affects fewer arrivals, since the arrival rush is near its end. The time at which an aircraft enters a queue represents the undelayed arrival or departure time; the time spent in the queue equals the delay. Table 1 compares the total delays (i.e., the sum of the times each aircraft spends in its queue) for the two policies. Although reacting late to the departure demand minimizes arrival delays, the

total resulting delays are considerably larger than when departures are emphasized earlier.

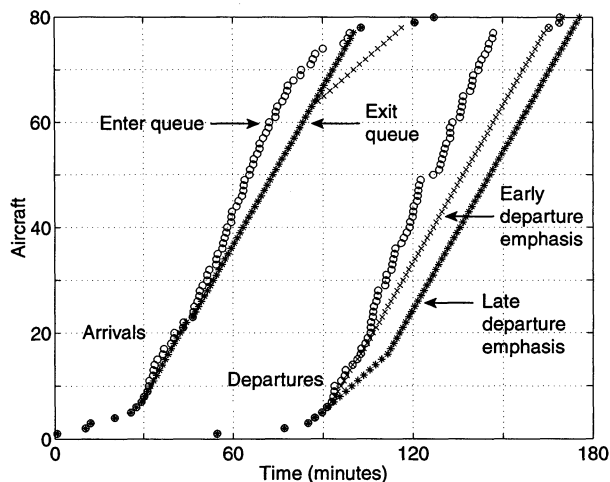


Figure 9. Times at which aircraft enter and exit the queues, for the two capacity allocation policies.

The purpose of this exercise is not to advocate either of these particular policies. Rather, the exercise demonstrates the impact that the decision of when to reallocate capacity from arrivals to departures can have on the arrival and departure queues. One of the challenges controllers face in making this decision is the absence of information about the future demands for both types of operations. Currently, at Dallas–Fort Worth, CTAS displays predicted arrival times in the FAA towers, derived from radar data. Although predicting departure demand is more problematic, since even the air carriers don’t currently know when their flights will be ready to push back, SMS is working to facilitate the availability of this information in a timely manner.

Table 1. Comparison of delays (in minutes).

Policy	Arrival	Departure	Total
Late Emphasis to Departures	380	1211	1591
Early Emphasis to Departures	498	597	1095

A decision aid which predicts what queues would result from various decisions, or which advises operating mode changes, may provide further benefit to controllers and users. Moreover, although a variety of cost functions could be proposed to justify various policies as being optimal, only the air carriers know how delaying arrivals versus departures will affect their network-wide schedule and business efficiency.

Therefore, a collaborative algorithm will be suggested as the foundation for a controller decision aid.

Air Carrier Response to Arrival Delays

The service provider does not know the user's relative (and time-varying) values for arrival and departure capacities. Consequently, air carriers would like the ability to dynamically adjust the airport operating point in the manner that best achieves their business objectives. The user-selected operating point is limited by the capacity curve and must not negatively affect overall capacity, fairness between all users, or controller workload. Involving the air carriers in the capacity allocation decision could also improve airport efficiency because demand and capacity are not necessarily independent – the air carriers determine demand, information about demand is necessary in allocating capacity, and capacity can affect demand.

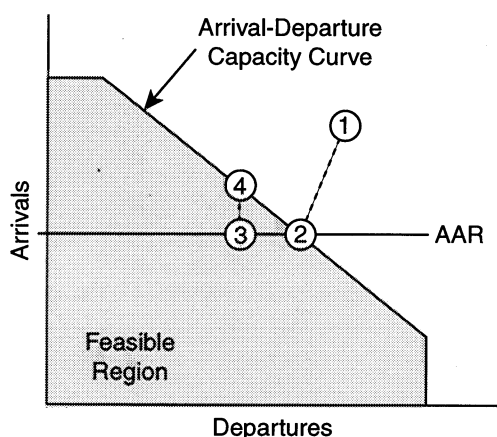


Figure 10. Air carrier response to arrival delays.

Assume an airport operates constrained by the hypothetical arrival-departure capacity curve shown in Figure 10, and that the total unconstrained demand for all of the users is represented by point 1. Since the combination of the arrival and departure demands lies outside the curve, strategic traffic management is required to delay some aircraft. By setting the Airport Acceptance Rate (AAR) as shown in the figure, arrivals and departures are limited to point 2. However, the air carriers may respond to limited arrival capacity by reducing their departure demand, since some of the aircraft and crews which make up the departures will be late and the air carrier may choose to hold departures to make passenger connections off the late arrivals. Therefore, by setting the AAR without considering the effect on departure demand, the airport will operate at point 3, rather than at 2, and the capacity between 3 and the capacity curve will be wasted. Given the airport cannot accommodate the unconstrained demand 1, the

air carriers may prefer to operate at 4 (or elsewhere on the arrival-departure capacity curve between 2 and 4) rather than at 3. However, the air traffic control system cannot currently know how the air carriers will respond to limited arrival capacity or, more generally, how they would want the airport to balance arrivals against departures.

All points on the arrival-departure capacity curve are equally efficient from the perspective that the airport is not wasting any capacity. Therefore, assuming all points on the curve are realizable without exceeding acceptable controller workload, how the airport operates should not matter to the ATC system. However, even if an air carrier's demand for arrivals exceeds capacity, the air carrier may prefer, for internal reasons such as gate availability, to further delay its arrivals to create additional departure capacity. Preferences such as this, which currently cannot be known a priori by the ATC system and will continuously change throughout the day, can significantly affect the air carrier's business efficiency. Therefore, collaboration between ATC and the air carriers could yield a more efficient allocation decision.

Collaborative Capacity Allocation

A variety of controller decision support tools could achieve TMA-SMS interoperability. The simplest would display arrival and departure demands, giving controllers the information needed to set efficient arrival and departure capacities, as well as plan capacity changes in advance. At DFW, CTAS currently provides information about arrival demand in the towers. A benefit of this approach is that the automation does not need to know the capacity limitations. An extension, which also relies on the controllers to coordinate arrival and departure capacities, could predict arrival and departure delays for proposed capacity allocations.

Alternatively, a decision support tool for interoperable arrival-departure management could advise controllers with an efficient sequence of coordinated arrival and departure capacities. This section proposes a possible algorithmic foundation for computing these capacity allocations. Previous approaches to "optimizing" capacity allocation have proposed minimizing various measures of the total cost for arrival and departure delays, often accounting only for direct operating costs. These approaches, which generally favor arrivals because the per-minute fuel and maintenance costs are higher for airborne aircraft, consider all flights to be equally important, which is not true from the perspective of the air carriers.

The Arrival-Departure Capacity Allocation Method (ADCAM), introduced by Hall [8], is a collaborative approach, between ATC and air carriers, to allocating airport capacity. ADCAM provides a method for calculating an efficient schedule of airport operating points that allows air carriers to trade arrivals for additional departures, or departures for additional arrivals, to best achieve their business objectives. More recently, Gilbo [9] has proposed an alternative optimization rule for collaboratively allocating arrival and departure capacities.

ADCAM is a generalization of the Collaborative Decision Making (CDM) Flight Schedule Monitor (FSM), which is currently used by the FAA and air carriers to manage ground delay programs [10]. If departure capacity is unlimited, ADCAM simplifies to the FSM algorithm. ADCAM consists of four steps, the first being to forecast the sequence of airport configurations (and modes) for the period of interest. Since the configuration determines the capacity curve, configuration planning is an integral part of arrival-departure interoperability. Improved knowledge of the arrival and departure demands, which ADCAM generates, may be used to improve the configuration plan, and the resulting new capacity forecast can be used in subsequent iterations of ADCAM. Note that this approach assumes the operating points that are feasible under each configuration are known.

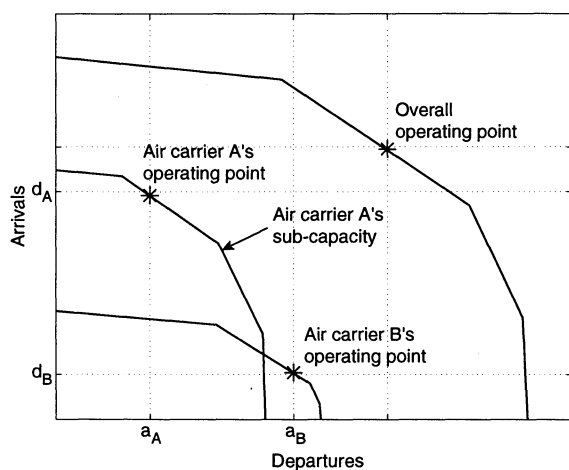


Figure 11. User sub-capacities calculated by the

The second step is to ration constrained airport and airspace resources among the air carriers. ADCAM employs ration-by-schedule, a paradigm that has been well accepted by all of the stakeholders under the CDM program. The set of resources allocated to an air carrier is the user's arrival-departure *sub-capacity*, illustrated in Figure 11. Hall [8] provides details concerning the

rationing algorithm and the construction of the sub-capacity curves. Graphically, a user's sub-capacity curve to the left, for example, of the user's operating point is a scaled copy of the overall airport capacity curve to the left of the overall operating point, where the scaling preserves the slopes of the capacity curve segments. Therefore, the sub-capacity curves look different because each air carrier's demand consists of a different ratio of arrivals and departures.

The third step of ADCAM is for each air carrier to re-plan its schedule, delaying and canceling flights, to best achieve its business objectives while adhering to its sub-capacity. An advantage of this approach is that it does not require that the ATC system know the arrival and departure demands explicitly. Rather, the method assumes the air carriers will know their demands, especially for departures. Each air carrier is allowed to operate any combination of arrivals and departures permitted by its sub-capacity curve. Hall [8] proves mathematically that the construction of the sub-capacity curves guarantees that the resulting total operations will be feasible with respect to the overall airport capacity constraints. However, the overall operating point may be under the airport's overall capacity curve, while a sub-capacity constraint is binding on every air carrier's solution. Furthermore, air carriers may decide to not use all of their allocated sub-capacity (e.g., if mechanical problems ground an aircraft). Therefore, the final step of ADCAM is to reallocate unused capacity by iterating the above steps. This step parallels the FSM compression step. Compression gives capacity that cannot be used by one air carrier to another air carrier that can use the capacity. In exchange, the releasing air carrier receives capacity at another time that it can use.

Conclusions

This paper described the complex interactions between strategic arrival and departure traffic management. When arrivals and departures at an airport are interdependent, the airport operating point (i.e., the mixture of arrivals and departures) may be adjusted to accommodate the time-varying demands for both types of operations. Currently, arrival and departure capacities are neither precisely coordinated nor planned in advance, resulting in inefficiencies. In fact, the knowledge of future arrival and departure demands needed to plan future capacity allocations is not presently available. Furthermore, the introduction of traffic management tools, such as TMA and SMS, will require that the arrival and departure capacities, to which the automation tools schedule, be coordinated and planned in advance.

Planning coordinated arrival and departure capacities requires knowing what operating points are (and will be) feasible. Plots of arrival-departure operation counts were examined. Empirical data does not show how the airport should or can operate at a specific time, since it does not show when operations are limited by demand rather than capacity. To more fully understand how the arrival-departure mixture can be controlled, information about the demands must be combined with the observed operations to identify when capacity limits operations. Furthermore, all regions of a theoretical arrival-departure capacity curve may not be feasible due to airborne capacity restrictions or ATC procedures that limit the airport's ability to set the arrival-departure mixture. An interoperable TMA-SMS could coordinate airborne and surface planning of arrival and departure rates, thereby improving the efficiency with which limited airport resources are used. Details of an interoperable TMA-SMS will be the subject of a follow-on investigation.

Two automation aids for macroscopic arrival-departure interoperability (i.e., managing aggregate arrival and departure rates without considering interactions between individual aircraft) were suggested. Displaying future arrival and departure demands would provide controllers with the information needed to plan an efficient schedule of coordinated arrival and departure capacities. Alternatively, automation could advise a schedule, possibly yielding one that is more optimal and enabling more frequent changes. An algorithm in which ATC and the air carriers collaboratively determine the plan for coordinated arrival and departure capacities was outlined.

References

- [1] E. Gilbo, "Airport Capacity: Representation, Estimation, Optimization," *IEEE Transactions on Control Systems Technology*, vol. 1, no. 3, 1993.
- [2] H. Erzberger, T. Davis, and S. Green, "Design of Center-TRACON Automation System," AGARD Meeting on Machine Intelligence in Air Traffic Management, Berlin, Germany, 1993.
- [3] H. Swenson, T. Hoang, S. Engelland, D. Vincent, T. Sanders, B. Sanford, and K. Heere, "Design and Operational Evaluation of the Traffic Management Advisor at the Fort Worth Air Route Traffic Control Center," 1'st USA/Europe ATM R&D Seminar, Saclay, France, 1997.
- [4] I. Anagnostakis, H. Idris, J.-P. Clarke, E. Feron, R. J. Hansman, A. Odoni, and W. Hall, "A Conceptual Design of a Departure Planner Decision Aid," 3'rd USA/Europe ATM R&D Seminar, Napoli, Italy, 2000
- [5] T. Davis, D. Isaacson, J. Robinson, W. den Braven, K. Lee, and B. Sanford, "Operational Field Test Results of the Passive Final Approach Spacing Tool," IFAC 8'th Symposium on Transportation Systems, Chania, Greece, June, 1997.
- [6] H. Idris, B. Delcaire, I. Anagnostakis, W. Hall, J.-P. Clarke, R. J. Hansman, E. Feron, and A. Odoni, "Observations of Departure Processes at Logan Airport to Support the Development of Departure Planning Tools," 2'nd International ATM R&D Seminar, Orlando, 1998.
- [7] H. Idris, B. Delcaire, I. Anagnostakis, W. Hall, N. Pujet, E. Feron, R. J. Hansman, J.-P. Clarke and A. Odoni, "Identification of Flow Constraints and Control Points in Departure Operations in Airport Systems," AIAA Guidance, Navigation & Control Conference, Boston, 1998
- [8] W. Hall, "Efficient Capacity Allocation in a Collaborative Air Transportation System," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1999.
- [9] E. Gilbo and K. Howard, "Collaborative Optimization of Airport Arrival and Departure Traffic Flow Management Strategies for CDM," 3'rd USA/Europe ATM R&D Seminar, Napoli, Italy, 2000
- [10] M. Wambsganss, "Collaborative Decision Making Through Dynamic Information Transfer," *Air Traffic Control Quarterly*, vol. 4, no. 2, pp. 107-123, 1997.