NASA/CR—97— 207142

NA62-1088

IN-04-CR
(C) UNITED
067295

Evolutionary Concepts for Decentralized Air Traffic Flow Management

Milton Adams, Stephan Kolitz, Joseph Milner, and Amedeo Odoni

Alternative concepts for modifying the policies and procedures under which the air traffic flow management system operates are described, and an approach to the evaluation of those concepts is discussed. Here, air traffic flow management includes all activities related to the management of the flow of aircraft and related system resources from "block to block." The alternative concepts represent stages in the evolution from the current system, in which air traffic management decision making is largely centralized within the FAA, to a more decentralized approach wherein the airlines and other airspace users collaborate in air traffic management decision making with the FAA. The emphasis in the discussion is on a viable medium-term partially decentralized scenario representing a phase of this evolution that is consistent with the decision-making approaches embodied in proposed Free Flight concepts for air traffic management. System-level metrics for analyzing and evaluating the various alternatives are defined, and a simulation testbed developed to generate values for those metrics is described. The fundamental issue of modeling airline behavior in decentralized environments is also raised, and an example of such a model, which deals with the preservation of flight bank integrity in hub airports, is presented.

INTRODUCTION

Air Traffic Flow Management (ATFM) is one of the two major components of Air Traffic Management (ATM), the other being ATC. ATFM includes all activities related to the management of the flow

© 1996 Air Traffic Control Association Institute, Inc.   CCC 1064-3818/95/030163-20

of aircraft and related system resources from "block to block," including strategic flow management of airport arrival and departure capacities, tactical en route flow management, near terminal area flow management, and ground traffic flow management.

ATFM has become increasingly critical to the successful operation of the air traffic system both in the United States and in Europe. Continuing growth in system traffic demand is not being met by corresponding increases in the physical capacity of the system (e.g., new airports), and therefore must be accommodated by increasing the system's effective[1] capacity through improved management and utilization of the existing system resources. The objective of both strategic and tactical traffic flow management is to match as best[2] as possible the projected demand for air transportation system resources (airports, terminal areas, en route sectors) with their anticipated capacity. ATFM system functions are most critical to system performance on precisely those days and at precisely those locations where the demand vs. capacity relationship is most unfavorable. The ATFM system relies on a combination of mechanisms, some of which are more global and strategic in nature with longer time horizons (e.g., ground holding of aircraft prior to departure, ground stop programs and traffic rerouting), while others are more local and tactical with shorter time horizons (e.g., miles-in-trail, airborne holding, arrival sequencing, and ground traffic management).

The ATFM system is on the verge of a transition that is likely to bring about dramatic changes. This transition is unavoidable in view of the confluence of several factors, including:

- The emergence of new technologies that offer the opportunity to correct some of the perceived deficiencies in today's ATFM system.
- The expressed preference by the airlines and other aircraft operators for a more decentralized system wherein they participate more broadly in ATFM decision making.
- A general recognition of the need for increased use of decision support tools and automation aids in order to cope more effectively with the highly dynamic environment in which the ATFM system operates, including substantial uncertainty in predictions of demand and available airport capacity when weather conditions deteriorate.

This paper is concerned with alternative concepts for modifying the policies and procedures under which the ATFM system operates

---

[1] *Effective* capacity refers to the capacity of a system resource that is realized as a result of the application of a set of policies and procedures for utilizing that resource.
[2] Here, the "best match" is really a multi-objective problem in that "best" is interpreted differently by the various system participants.

in the United States. These alternative concepts represent stages in an evolution from the current system, in which ATFM decision making is largely centralized within the FAA, to a more decentralized approach, wherein the airlines and other airspace users collaborate in decision making with the FAA. This evolution is consistent with the decision-making approaches embodied in the Free Flight approach to ATFM. In particular, with more substantial participation from the airlines, air traffic flow will be influenced more directly in ways that accommodate the business objectives of the airlines and the interests of their passengers. At the same time, the FAA will continue to be responsible for the safe operation of the U.S. air transportation system. Thus, a collaboration between the FAA and airlines will be required to ensure that system resources will simultaneously address both airline business and system safety objectives.

The next section elaborates on the themes of centralized and decentralized ATFM and discusses a spectrum of proposed alternative concepts for ATFM, ranging from highly centralized (nearly all decision making made by the FAA) to highly decentralized (nearly all decision making performed by the airlines and other airspace users). The third section further explores a viable medium-term partially decentralized scenario, summarizes metrics that can be employed in analyzing and evaluating the various alternatives, and describes briefly a simulation testbed that has been developed to generate values for those metrics. In addition to modeling activities of individual aircraft, the simulation testbed must contain behavioral models of the FAA, the airlines, and other airspace users. The fourth section raises the fundamental challenge of modeling airline behavior in decentralized ATFM environments and presents an example of such a model, which deals with the preservation of flight bank integrity in hub airports. The final section presents conclusions.

ATFM CONCEPTS

As long as there are times, such as during severe weather conditions, when the capacity of one or more air traffic system resources falls considerably below the scheduled demand, a coordinated approach will have to be employed to allocate fairly scarce resources during the period of excess demand. It is highly probable that periods when demand for the limited departure and arrival capacities at the busiest airports exceeds the supply will continue to exist for the foreseeable future.

In a centralized ATFM system during periods when demand exceeds capacity, the ATFM system operator (the FAA) makes largely unilateral decisions required to align demand with capacity, such as the assignment of delays to aircraft and the modification of their

routes, and monitors closely each aircraft's compliance. A "moderately" decentralized approach to addressing excess demand would be similar in many ways to today's ATFM system, but with increased FAA/airline cooperation and coordination in ATFM decision making. In a "fully" decentralized ATFM system, each aircraft and aircraft operator would be given accurate and timely information about existing and projected demand and capacity for each ATFM system resource, allowing each aircraft operator to determine independently its own preferred strategy with regard to its own set of flights, and to plan, execute, and monitor the flights. The ATFM system operator would be kept apprised of those plans and intervene only when needed for safety.

The airlines have generally indicated a preference for the relatively decentralized end of the spectrum. Indeed, under decentralized schemes, airlines, in theory, would have more freedom to optimize their individual operations, with the potential for providing passengers with more efficient service, as measured by shorter flight times and more reliable schedules and connections. The Free Flight concept is an expression of this point of view (RTCA, 1995, 1996). The Collaborative Decision Making (CDM) (formerly called FADE) program (Wambsganss, 1997), already under way, is an attempt to increase cooperation and collaborative decision making between the FAA and the airlines in the existing ATFM system.

Evolution of Alternative ATFM System Concepts

The ATFM system is an extremely complex, large-scale system that can be decomposed into three highly coupled physical segments—en route, near terminal area, and ground operations (see Figure 1). The decisions made for planning and controlling the traffic management activities within each of these segments impact those within the other two, and together they determine the overall flow of traffic through the air transportation system network. Ultimately, the objective of any modification made to the ATFM system is to increase the effective capacity of the overall system in ways that benefit all participants (FAA, airlines, general aviation, and passengers), while sustaining or improving the level of safety afforded by the system. Because the decisions and activities within one segment impact those in the others, some form of coordination of air traffic management across segments will be required (1) to ensure that solutions to air traffic management problems in one segment do not have a significant negative impact on other coupled segments, and (2) to increase opportunities for synergism across segments that will lead to improved system-wide performance. Metrics or figures of merit that are useful in evaluating alternative ATFM concepts are discussed in the next section.

Figure 1 represents a functional decomposition of decision making

A "moder-
l would be
increased
n making.
d aircraft
about ex-
M system
pendently
ghts, and
l operator
nly when

relatively
ntralized
optimize
g passen-
ght times
it concept
The Col-
program
increase
FAA and

tem that
ents—en
e 1). The
agement
the other
through
ective of
effective
ticipants
staining
ause the
the oth-
ross seg-
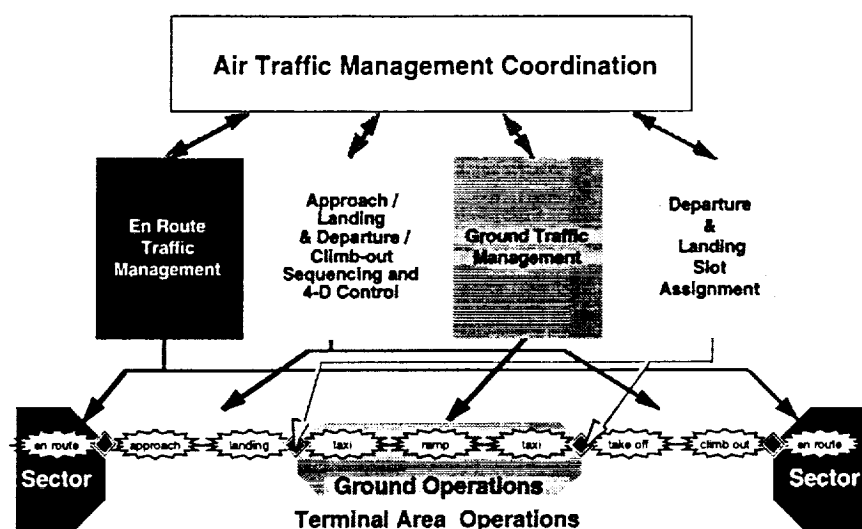fic man-
negative
tunities
system-
in eval-
section.
making

Figure 1. Coordination of traffic flow management across segments.

in ATFM, but does not reflect an allocation of those decision-making functions to the FAA, airlines, or any other potential participant in the system.

In this paper, the term *decentralization* refers to both the decomposition of the decision-making functions as illustrated in the middle tier of Figure 1, and the shared allocation of those functions among the FAA, the airlines, and other airspace users. In the latter case, decentralization connotes that the airlines and other airspace users play a more significant role in decision making than heretofore.

The evolution from the policies and procedures under which the current ATFM system operates to those of future, more decentralized systems should occur along a migration path of feasible, cost-effective changes. Each change along such a path should provide improvements in system performance. Specifically, new equipment and procedural changes must be phased into the system in such a way that there are benefits to the overall system as measured by its benefits to the individual participants: the FAA, the airlines and other airspace users, and the passengers.

Feasibility of implementation is a criterion that must be applied in defining the sequence of system modifications. In particular, the introduction of new technologies cannot occur "overnight," so that the system must operate under each modification with a mix of equipped aircraft, some of which will have new equipment and some of which will not.

Characterization of Alternative ATFM Concepts

Descriptions of ATFM Alternatives. Table 1 outlines a spectrum of alternative ATFM system concepts, with each row representing a

Table 1. Spectrum of Alternative ATFM System Concepts

Degree of Centralization	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized	1a ATFM system operator (FAA) allocates arrival slots to individual flights.	2a ATFM system operator assigns slots; each airline may concel and substitute flights.	3a ATFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; ATFM system operator controls.	
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); ATFM system operator approves or rejects.			
Partially Decentralized I	1b ATFM system operator allocates sets of arrival slots to individual airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assign departure slots to individual aircraft; ATFM system operator approves or rejects.	4b Airlines plan; ATFM system operator specifies a region in which airlines control their own aircraft; ATFM system operator controls other regions and monitors globally for feasibility, conflicts.	5a Airlines plan; ATFM system operator controls.
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own			

	(1)	(2)	(3)	(4)	(5)
(row above — cut off)	allocates sets of arrival slots to individual airlines.	sets of slots among their own flights.	individual aircraft; ATFM system operator approves or rejects.	individual aircraft; ATFM system operator controls other regions and monitors globally for feasibility, conflicts.	ATFM system operator controls.
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c ATFM system operator informs airlines of the guaranteed capacities at potentially congested airports.	2e Airlines may bargain among themselves for the guaranteed capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft		5b Airlines plan; they also specify each aircraft's heading directly after departure; ATFM system operator can approve or reject heading; ATFM system operator controls everything else.
Decentralized II	1d ATFM system operator informs airlines about anticipated availability of capacities at congested airports.	2f Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; ATFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; ATFM system operator monitors for feasibility, conflicts.

potential operational alternative. A dynamic, real-time decision-making environment is assumed throughout. The columns in the table decompose the strategies for implementing each concept into:

1. Initial allocation of arrival slots among aircraft operators.
2. Final assignment of arrival slots to individual flights.
3. Assignment (if any) of departure slots to individual flights.
4. En route flight planning and control.
5. Transitional area and terminal area flight planning and airport surface movement control.

The concepts in the table range from "highly centralized" in the top row (where decision making is centralized within the FAA) to "highly decentralized" in the bottom row (where decision making is decentralized across the airline operators). Moving down a row in the table represents an evolutionary step in the process of system decentralization and a consequent move in the direction of more Free Flight. Each alternative ATFM operational concept (each row) must be analyzed and evaluated to find the best operating point from the perspective of both system performance and system safety.

The second row, *Partially Centralized*, roughly describes the present state of the system. The third row, *Partially Decentralized I*—highlighted in gray—is the focus of the partially decentralized scenario discussed in the next section. The ATFM concept represented by this third row is a viable evolutionary step from the present system; indeed, most of the technologies to implement this concept exist, and consensus for it among the FAA and airlines is beginning to take shape.

All concepts outlined below the *Partially Decentralized I* row in Table 1 should be considered highly speculative at this time and are listed here only as a rough indication of the types of potential approaches that may emerge in the future. The brief descriptions of these concepts below are necessarily sketchy and incomplete. Moreover, to our knowledge, no analysis of the robustness of these concepts with respect to ensuring a fail-safe system operation has been performed to date. Thus, the feasibility of moving beyond a concept similar to *Partially Decentralized I* is an open question.

Discussion of Table Entries

Column 1: Allocation of Arrival Slots. One of the principal flow management strategies employed in the current system in reacting to reduced capacity at an individual destination airport is that of controlling the departure times of aircraft destined for that airport through the use of Ground Delay Programs. In effect, these controlled departure times implicitly represent a set of arrival slot allocations

decision-
ns in the
:ept into:

ors.

:hts.

nd airport

ed" in the
e FAA) to
making is
row in the
em decen-
nore Free
row) must
t from the

s the pre-
alized I—
ilized sce-
presented
esent sys-
:ept exist,
ng to take

I row in
e and are
:ntial ap-
iptions of
te. More-
concepts
:een per-
l concept

:ipal flow
reacting
s that of
t airport
:ntrolled
locations

at the destination airport of interest. The current strategy employed in arrival slot allocation is one wherein the FAA attempts to minimize airborne delays and maximize utilization of available airport take-off and landing capacity, with the objective of adhering as closely as possible to published schedules (i.e., the OAG). Recently, to ensure that airport arrival capacity would be fully utilized, a program referred to as Managed Arrival Reservoir (MAR) was also instituted. This program plans for up to 15 minutes of airborne delay for arrivals beyond the forecasted capacity (an airborne reserve).

In a more decentralized ATFM system, where Free Flight-like concepts are employed for en route traffic, the allocation of arrival slots will be the FAA's principal flow control mechanism, and thus is likely to play an even more significant role in air traffic management than it does in the current system.

In a first step toward decentralization (**1b** in the table), the FAA allocates a set of slots to each airline, e.g., based on first-come, first-served, or on the proportion of flights held, over predefined intervals without assigning specific flights, giving the airlines more flexibility to assign flights to the allocated slots. This provides an advantage from the airlines' perspective in that each individual airline can assign flights to its allocated slots based on its own business objectives.

In the next step, no specific landing slot assignments are made to individual airlines (**1c**). However, to ensure safe operation, there is a cap on the total number of slots for predefined time intervals. This opens the possibility of creating a "real-time market" within which the airlines "trade" for slots up to the specified limits. Finally, in a fully decentralized scenario (**1d**), the market of available slots is not constrained, but the FAA keeps the airlines informed of the expected safe limits on arrival capacities at individual airports.

Column 2: Assignment of Arrival Slots to Individual Flights. Under the approach taken today (**2b**), each airline, within limits, suggests alternative assignments of the slots allocated to it (its own flights); the ATFM system operator may approve or reject those alternative assignments. The airlines have the freedom, within established constraints, to cancel flights and substitute other flights. Substitutions are done one at a time, with each being subject to approval by the ATFM system operator.

Under **2c**, individual airlines freely assign their own flights to their own sets of slots. In this case, the ATFM system operator performs no assignment of flights to slots. Each airline is allocated slots and is given the freedom to assign any of its own flights to the allocated slots, with the assignments being subject to approval by the ATFM system operator. The ATFM system operator awards any slots left unused by a particular airline to other airlines on a first-scheduled, first-served, or other agreed upon, basis.

Under **2d**, airlines are allocated slots by the system operator and may trade slots among themselves. They are free to assign flights to their allocated/traded slots as they see fit.

Under **2e**, airlines negotiate among themselves for slots within guaranteed, minimum capacities. No slots are initially assigned; however, the ATFM system operator sets a limit on slots for each airport as a function of time, based on considerations of safety. Airlines negotiate among themselves for slots within this limit and may conceivably launch flights that exceed it, but such flights are subject to diversion if they cannot be handled safely at the destination airport.

Under **2f**, airlines decide freely and independently on their own slot allocations/assignments. The ATFM system operator disseminates information with respect to guaranteed, minimum capacities and expected aggregate demand based on the most up-to-date information regarding the posted assignment decisions of all the airlines. The system operator sets a limit on slots for each airport as a function of time, based on considerations of safety. An approach to sequencing landings when demand exceeds guaranteed supply must be developed.

Column 3: Departure Slot Assignment.

Given an assignment of flights to arrival slots at destination airports, departure slots for those flights from their origination airports can be assigned in a variety of ways. The most straightforward is simply to subtract the nominal flight time from the arrival slot assignment time and assign a departure slot at the origination airport for that time. Alternatively, in situations where the assigned arrival slot for a given flight represents a delay with respect to the scheduled arrival time for that flight, the airlines or the system operator may choose for that flight to depart earlier than the simple difference between the assigned arrival time and the nominal flight time (e.g., as in the MAR program) in anticipation that either (1) weather may improve, and the increased capacity will result in newly opened slots that a "flight already in the air" could take advantage of, or (2) a slot may open up at the arrival airport as a result of delay or cancellation of other flights slotted to arrive. Thus, a buffer of aircraft in the air representing demand that slightly exceeds anticipated capacity ensures that an unexpected increase in realized capacity will not go unused. Given the freedom to make this kind of decision, airlines may decide to take some of the scheduled delay on the ground (a ground hold) and some in the air, anticipating that capacity will improve with some nonzero probability.

Under **3b**, the airlines assign departure slots to individual aircraft, and the ATFM system operator approves or rejects those assignments. Thus, the airlines have the opportunity to anticipate potential

rator and
.. flights to

ots within
assigned;
s for each
afety. Air-
t and may
s are sub-
estination

their own
r dissemi-
capacities
late infor-
e airlines.
a function
sequenc-
must be

mment of
slots for
med in a
tract the
nd assign
natively,
ht repre-
at flight,
flight to
igned ar-
program)
l the in-
flight al-
open up
of other
ir repre-
ensures
unused.
y decide
nd hold)
ve with

aircraft,
assign-
otential
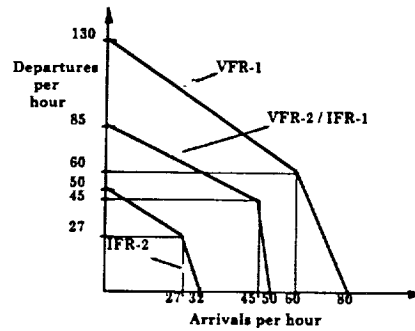
Figure 2. Typical trade-off of arrival vs. departures.

improvements in arrival airport capacity and can choose to "leave early" in order to take advantage of any realized improvements in arrival capacity. To ensure that the number of airborne aircraft is not so large as to cause potential safety problems due to unacceptably high levels of congestion near arrival airports, the chosen departure slots are subject to approval by the system operator.

Under **3c**, the airlines assign departure slots to individual aircraft. Here, the system operator is responsible only for controlling departures to ensure safety. There is, of course, no guarantee that a flight will be able to depart at the desired time, if the total requested number of departures from the origination airport exceeds the departure capacity of that airport during the period of interest.

For given weather conditions at many airports, the number of departures per unit time interval and the number of arrivals over that same interval are often coupled, as notionally illustrated in Figure 2. Thus, arrival and departure slot allocation cannot be performed independently when demand exceeds the available supply.

Column 4: En Route Planning and Control. For the en route segment, control strategies range from assignments of 4-D waypoints by the system operator to Free Flight.

Under **4a**, airlines (and other airspace users) plan; the ATFM system operator controls. Airlines file a flight plan for each flight, and the system operator suggests or mandates modifications to the plan if deemed necessary for safety purposes, and monitors the flight. Flight plans are expected to conform largely with established airways.

Under **4b**, airlines (and other airspace users) plan; the ATFM system operator specifies regions in which user-preferred trajectories are acceptable, controls flights in other regions, and monitors globally for feasibility and conflicts. This is a mix of the previous approach for certain regions (e.g., highly congested en route sectors and near terminal areas) and more freedom for the airlines to fly in Free Flight in other regions. The current National Route Program (NRP) repre-

sents the first step in this direction. In all cases, the system operator is required to monitor flight trajectories in order to predict conflicts, and to intercede and resolve conflicts when potential conflicts are detected.

Under **4c**, airlines (and other airspace users) plan and reevaluate continuously their flight trajectories; the ATFM system operator monitors for feasibility and conflicts. Here the airlines are afforded the flexibility to freely choose en route trajectories, with the system operator being responsible for conflict monitoring and, when necessary, resolution.

Column 5: Transition Area, Terminal Area, Ground Movement Planning and Control. To ensure safe operations under all of the alternatives discussed here, the ATFM system operator must monitor and manage traffic in the near terminal area for arrivals and, to a lesser extent, for departures. Under current approaches to ATFM, the system operator implicitly controls the demand on near terminal area airspace through both arrival slot allocation and en route traffic control. With the evolution toward increased freedom for the airlines to choose arrival and departure slots and en route flight plans, there comes the increased potential for substantial congestion in the airspace near airports. To avoid the associated potential for terminal area overload and attendant risks, measures that serve to coordinate ATFM between the en route and terminal areas must be applied by the system operator in order to influence traffic behavior in a manner that will reduce this potential.

Although the terminal area traffic may be tightly controlled by the system operator, this does not preclude collaboration by the airlines in creating the plans for the sequencing of that traffic. Currently, arrival sequence planning under even the most advanced systems (e.g., Center-TRACON Automation System [AS]; see Erzberger, 1993) does not directly include participation by the airlines. Increased decentralization may provide opportunities for the airlines' participation in formulating the objectives employed by the system operator in choosing arrival sequences.

Under **5a**, airlines and other airspace users plan their activities within the terminal area and on the ground; the ATFM system operator controls aircraft in this region (except when they are on the ramp) including movement on the taxiways, aircraft departures, climb-outs, and descents. Departure headings are limited to predefined fixes.

Under **5b**, airlines (and other airspace users) plan; they also specify each aircraft's heading directly after departure (here we assume that departures are not restricted to flying departure fixes); the ATFM system operator can approve or reject headings. Prior to departure, airlines and other airspace users inform the ATFM system

operator of their preferred departure heading. To ensure that the configuration of departing airborne aircraft does not cause potential safety problems, the preferred departure headings are subject to approval by the system operator.

Under **5c**, airlines and other airspace users have the freedom to operate within the "rules of the road." The ATFM system operator monitors for feasibility and conflicts, and has the capability to impose control.

Another component of terminal area ATFM is ground traffic management. In ground operations, there can be a strong coupling between the assignment of departure slots and the control of the ground movement of aircraft. For example, in poor weather with limited visibility, arrival and departure rates (airport capacities) can be limited by the ability to move ground aircraft traffic through taxiways and ramps in a timely fashion. Ground traffic planning, management, and control can benefit from timely and accurate traffic location information and digital data communications. Thus, improvements in ground traffic surveillance and planning can effectively improve the capacity of some airports in poor weather, reducing delay costs by apportionment of delay among gates, ramps, and taxiways. Again, as in the case of arrival sequencing, decentralized approaches to traffic flow management must afford the airlines the opportunity to influence ground traffic plans to the extent that those plans affect the sequence of departures at an airport.

EVALUATION OF NEW ATFM SYSTEM CONCEPTS

Since the ATFM system consists of several highly coupled segments, as described earlier, it would be misleading to evaluate the impact of a modification to ATFM operating procedures for any single segment without accounting for its impact on the rest of the system. In this section, an ATFM simulation testbed that embodies a variety of system-level modeling and analysis tools is described. This testbed has been developed to evaluate the system-wide impacts of candidate modifications to the existing ATFM environment.

Since each of the three principal stakeholders in the system—the FAA, the airlines, and the traveling public—may have different sets of priorities and objectives, each ATFM system concept should be evaluated using a variety of metrics that reflect the utility functions of each stakeholder. The use of a simulation testbed that contains a complete system description, appropriate evaluation metrics, models of each of the system entities, and an array of analytical capabilities will ensure that proper preliminary system-wide evaluations are performed.

A Viable Medium-Term Partially Decentralized Scenario

Our recent research has concentrated on the development of a set of tools and experiments that would make it possible to evaluate the *Partially Decentralized I* concept highlighted in Table 1. As noted earlier, this concept seems to be a viable evolutionary step from the present system. Indeed, some of its aspects will be implemented within the next 5 years under the *Free Flight Action Plan* recently announced by the FAA (1996) in response to the work of the RTCA Task Force on Free Flight, while several other aspects have been investigated, at least in a preliminary manner, in a number of recent studies (Milner, 1995; Wambsganss, 1997; DeArmon and Lacher, 1996).

The principal characteristics of the concept can be summarized as follows:

1. For slot allocation purposes, the busy hours of the day are subdivided into intervals that can accommodate several arrivals (e.g., intervals of 10 or 15 or 20 minutes). This increases flexibility for both the system operator and the airlines with regard to dynamic arrival scheduling, while providing protection from excessive arrival clustering (DeArmon and Lacher, 1996). (It should also be noted that none of the operational characteristics 3-8 listed below depends critically on this point; the proposed concept is also compatible with a slot allocation system that would allocate slots on a one-flight basis, i.e., by subdividing the time axis into intervals of the order of 1 minute.)

2. Allocation of arrival slots at congested airports is performed on a dynamic basis, according to predicted airport capacity over the next few hours. Whenever arrival capacity at one or more airports is predicted to be exceeded by demand, available slots for arrivals at these airports are allocated among the airlines on a first-scheduled, first-served basis to ensure fairness. For example, suppose that 15-minute slot allocation intervals are in use and that airline XYZ originally had six arrivals scheduled to arrive at A between 10:00 and 10:14. On a particular day when capacity is low, XYZ might then receive four slots on a first-scheduled, first-served basis, with the other two slots being moved to the 10:15–10:29 interval. Note that, while the number of slots that XYZ will receive in the interval is specified, the identity of those XYZ flights which will occupy these slots is not. This is a fundamental aspect of this partially decentralized concept and applies irrespective of whether the intervals into which slots are allocated are 1, 10, or 15 minutes long.

3. Each airline (or, more generally, each aircraft operator) is now free to utilize its slots in each interval in the way it deems best.

Thus, each airline may schedule any one of its flights into any one of its arrival slots. Each airline must also keep the ATFM operator informed as to which flight has been assigned to each slot and, most important, what slots, if any, will be left unused because of flight cancellations. Any slots left unused by a particular airline will be awarded by the ATFM operator to other airlines on a first-scheduled, first-served or other agreed-on basis.

4. The ATFM operator then estimates a controlled time of arrival (CTA) for each flight, (taking into account each airline's preferred sequencing of its own flights) and sends to each airline the list of its CTAs. The point of "arrival" is not necessarily the runway; in fact, in the presence of congestion, this point will usually be the boundary between en route airspace and the transitional airspace into the terminal area of each congested airport.

5. Little or no use is made of departure slot assignments (known currently as Expected Departure Clearance Times [EDCTs]). Thus, each aircraft operator is responsible for determining the time of take-off that is most appropriate for meeting the CTA of each of its flights. This means the aircraft operators will also decide how best to allocate a predicted delay between delay taken on the ground and that taken while airborne. In other words, the airlines will determine the size of their own "Managed Arrival Reservoirs."

6. Free Flight (user-preferred routing) is permitted in large portions of en route airspace and is used to travel from the airport of origin to the "arrival point" at which the CTA has been specified.

7. Air traffic management in the transitional area to the airport of arrival, in the airport's terminal area, and on the airport's surface is supported by advanced decision support and automation aids such as CTAS and Surface Management Automation (SMA). (An SMA-like concept is described in Böhme, 1994).

8. The ATFM operator checks for compliance with slot allocations and CTAs on the part of aircraft operators. The ATFM operator also continuously monitors traffic operations to ensure safety.

Numerous additional details can be specified with regard to the partially decentralized concept described by the above characteristics, but they are superfluous for the purposes of this paper. It should also be noted that many other plausible variations on the above themes merit investigation. The important point, however, is that a consensus is beginning to emerge among the FAA and the airlines about the desirability of moving toward a system that complies with the general framework outlined here.

Table 2. ATFM Performance Evaluation Metrics

Congestion-related	(1) Demand-to-capacity ratios at each airport and each defined airspace region by time of day
	(2) Number of aircraft held/delayed in the air on arrival for each airport by time of day
	(3) Number of aircraft held/delayed on the ground on departure because of ATFM intervention
	(4) Spatial density of airborne aircraft by defined airspace region over time
Delays and delay costs	(1) Minutes of delay incurred per operation (e.g., arrival, departure, taxi-in, taxi-out, transit portion of the flight plan)
	(2) Associated aircraft direct operating costs, according to a general, user-specified function of delay time
Schedule reliability	(1) Distribution of arrival times of flights relative to scheduled arrival times
	(2) Distribution of arrival times for flights defined to be members of a flight "bank" relative to scheduled arrival times for the bank
	(3) Distribution of the percentage of other flights in a bank with which each member of a bank connects
Aircraft utilization	(1) Number of aircraft of a given type needed to perform a specified part of a daily schedule of flights

Evaluation Metrics

A set of key metrics related to congestion, delay costs, schedule reliability, and utilization of aircraft and other resources has been identified for the evaluation of alternative ATFM systems. Each of these metrics quantifies an aspect of performance that is of particular interest to one or more of the three principal stakeholders identified previously. Fine-granularity metrics are listed in Table 2; aggregate metrics can be derived from these:
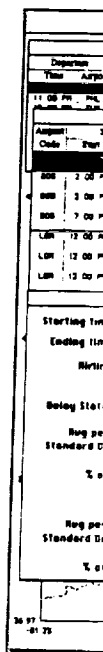
ASCENT—The Draper ATFM Simulation Testbed

Draper Laboratory and MIT have been working together to investigate ATFM concepts since 1989. Since 1991, as part of that collaboration, an ATFM simulation testbed, ASCENT (ATFM System Concept Evaluator for New Technologies), has been designed and implemented to evaluate the system-wide impact of new procedures, technologies, and improved infrastructure under existing or anticipated future approaches to ATFM. The current version of ASCENT contains the following:

- Models for a national network of capacitated[3] and noncapacitated[4] airports

---

[3] At a capacitated airport, capacity can be less than demand. ATFM planning deals with all flights into and out of capacitated airports.

[4] At a noncapacitated airport, demand is always less than capacity.

Figure 3. I

- Algorith delay be
- Algorith
- A syster space Sy
- Databas

The foll

- Models t
- A tool fo

Figure 3 available f ASCENT analyst, re defining ar pable of ev of scenario ities (suppl ing factors user intera Through erator; see
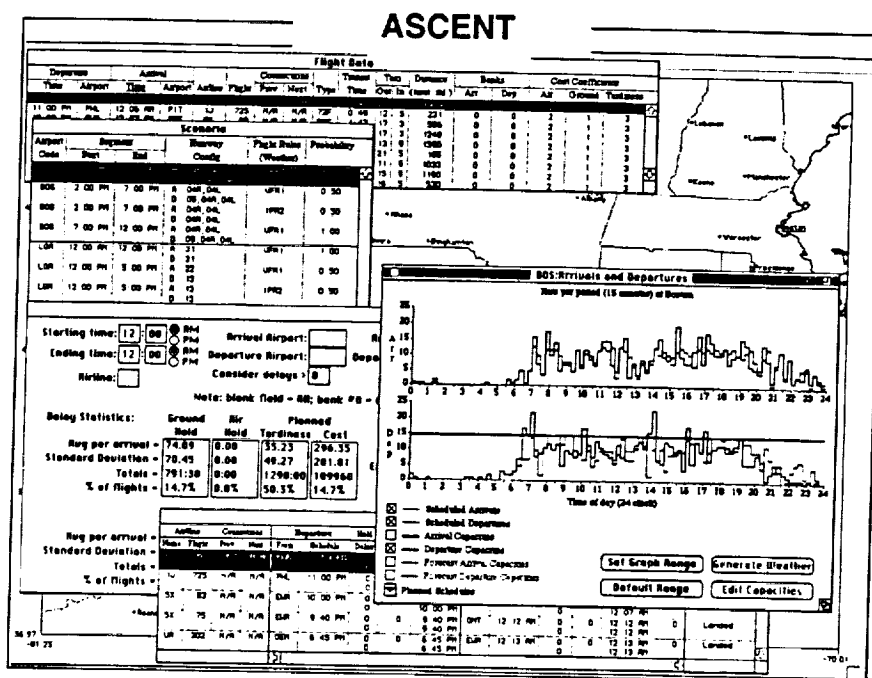
Figure 3. Representative ASCENT windows.

- Algorithms for planning ground holds and for allocating mandated delay between the ground and the air
- Algorithms for tactical planning of arrivals at airports
- A system-level simulation of a day's activities in the National Airspace System (NAS)
- Database and analysis capabilities

The following supporting utility programs are included:

- Models to simulate the evolution of airport weather and capacity
- A tool for generating realistic demand schedules at airports

Figure 3 illustrates a subset of the many output window formats available from ASCENT's graphical user interface.

ASCENT has been designed so that it can be used by a single analyst, requiring a minimum of overhead activity associated with defining and setting up scenarios and performing analyses. It is capable of evaluating candidate ATFM approaches across a spectrum of scenario variations. Flight schedules (demand) and airport capacities (supply) have been determined to be the most significant defining factors for any given scenario. Tools have been created to allow user interaction in the creation of each of these scenario components.

Through the demand-generation tool POAGG (Pseudo-OAG Generator; see Hocker, 1994) the user can easily generate OAG-like hy-

pothetical flight schedules for a network of airports. POAGG uses a combination of heuristics and mathematical programming to create reasonably realistic flight schedules that satisfy user-specified parameters. These input parameters include the following for each airport:

- The number and hourly distribution of arrivals
- The percentage of flights that connect to each of the other airports in the network
- Directional travel times between airport pairs
- The presence, if any, of shuttle flights and their characteristics
- The presence, if any, of airline flight "banks" and their characteristics

Once defined, a parameter set can be saved and edited to create new scenarios.

Airport capacities can be defined explicitly by the user through an interactive graphical input format and can be required to change dynamically over time. Alternatively, the capacity can be generated automatically by a Markov model of weather and correlated capacity at individual airports (Robinson, 1992). The transition probabilities for the Markov model are estimated using NOAA data. Also in the testbed environment is a sawtooth wave model of weather (Yu, 1996) that has realistic spatial and temporal correlation characteristics. The airport arrival and departure capacities are modeled using the FAA Engineered Performance Standards.

In setting up a simulated test case, the analyst selects a flight schedule and an airport capacity forecast as inputs. One of a set of ground-holding algorithms is selected to create planned aircraft ground holds for the day. Reductions in en route times due to Free Flight, reductions in airport ground delay times due to the improved ground traffic management, or increases in effective airport capacity due to improved arrival sequencing (resulting, e.g., from CTAS) can also be selected or specified by the analyst. Once the test case has been set up, the simulation of a day in the NAS is realized, and the resulting delays and other desired evaluation metrics are computed. Note that when weather/capacity are modeled probabilistically, their realizations may not exactly match forecasts. If at some point during the simulated day a (weather or capacity) forecast changes, the analyst can choose to invoke an algorithm that replans ground holds or can select an algorithm to resequence arrivals tactically at a given airport, on the basis of both the current state of the system and the new forecast.

The left margin contains fragments of text from an adjacent page:

)AGG uses a
ing to create
specified pa-
ng for each

ier airports

eristics
character-

:reate new

irough an
to change
generated
l capacity
babilities
so in the
Yu, 1996)
:teristics.
ising the

a flight
a set of
aircraft
to Free
nproved
:apacity
AS) can
ase has
and the
iputed.
y, their
during
he an-
olds or
given
nd the

Modeling Airline Behavior

Modeling airline behavior and quantifying the benefits of alternative ATFM concepts are the most challenging aspects of evaluating decentralized AFTM. This section outlines some particularly interesting work on quantifying the benefits that may accrue from the airlines' improved ability under decentralized systems to preserve the "integrity" of their flight schedules (Milner, 1995). Additional algorithms and approaches related to modeling the behavior of airlines are described elsewhere (see Bertsimas and Stock, 1994; Richetta and Odoni, 1993; Brunetta et al., 1995; and Adams et al., 1996). These are also being incorporated into ASCENT.

Airlines that operate in hub-and-spoke environments constantly face the problem of trying to maintain the integrity of their banks of flights, whenever major delays occur at these airports. By allowing airlines to utilize their own slots as they deem best, decentralized ATFM concepts may make it possible for the airlines to accommodate their banks better. For example, a bank may originally have been scheduled to arrive over the course of 30–40 minutes, but may be assigned a set of arrival slots that span several hours. The airline's response may be to choose a particular subset of flights of the bank and keep them together, delaying or canceling the rest. Another airline facing the same problem may choose differently. What is clear is that an ATFM system operator cannot be as effective as the airlines themselves in making these decisions, because the ATFM operator cannot know perfectly each airline's preferences and utility functions.

Milner has developed two models of how airlines might respond to arrival slot allocations. The Independent Flights (IF) model describes the problem airlines face in allocating their arrival slots under the assumption that the airline views its flights as being independent of each other. Such a model might be applicable at a spoke airport and is similar to others presented earlier (e.g., Vasquez-Marquez, 1991).

The second model takes into account the dependencies among flights in a bank. Specifically, it includes information regarding the total delay flights experience when they arrive in a bank of flights. This delay is greater than that experienced by individual flights because of the time flights will spend at the hub airport waiting for the completion of the bank. In this second model, referred to as the Cancellation/Delay (C/D) model, an airline attempting to allocate its arrival slots will either assign a flight to arrive as part of the bank in which it is scheduled to arrive, assign the flight to a slot after the bank is completed, or cancel the flight outright. If a flight is assigned to a slot after its bank's completion, passengers on that flight will likely miss their connections. Further, passengers at the hub airport connecting onto the next flight that this particular aircraft was ex-
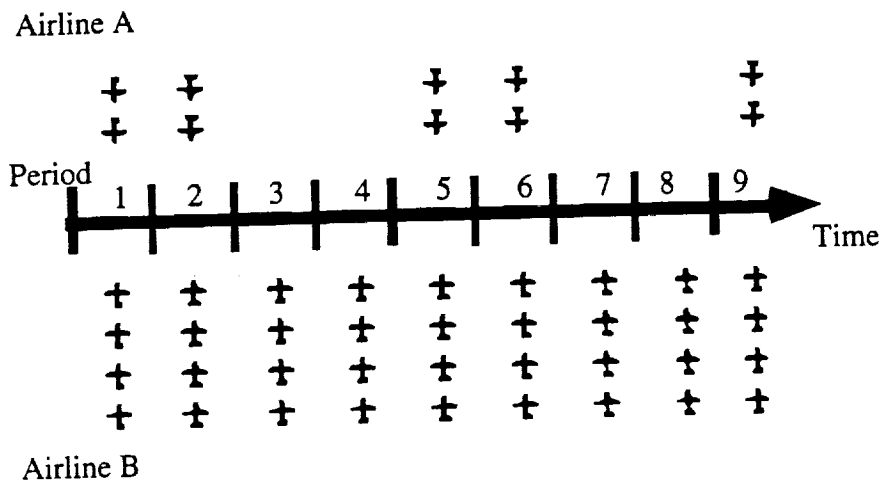
Airline A

Figure 4. Arrival pattern used in experiments.

pected to fly will also be delayed. The assumed objective of the airline in the model is to minimize a weighted combination of the delay incurred by flights that remain with their banks, the delay incurred by flights separated from their banks, and the cost of canceling flights.

Milner has shown that airlines can allocate their arrival slots in ways that are consistent with their preferences. In one set of experiments, he simulated the arrival of several banks of flights at a hub airport for an airline A. Each day was divided into 15-minute periods. Each bank of flights for the airline was scheduled to arrive within a 30-minute interval. A second airline B was also scheduled into the airport, the flights for that airline being a constant number per period. This second airline's schedule represented the total schedule for all other airlines operating at the airport. The simulated arrival pattern is displayed in Figure 4. In the figure, Airline A is seen to have scheduled banks of four flights arriving in each hour, in a 2-period or 30-minute interval. Airline B has four flights arriving in every period. In the actual experiments, the number of flights in Airline A's banks varied between 20 and 60, while Airline B had a constant 10 flights per 15-minute period.

In the experiments, a nominal capacity sufficient to satisfy the entire demand without delays was reduced systematically to 90, 80, 70, and 60 percent of its original amount. Airlines were allocated arrival slots based on the simulated schedule. Under these conditions, it was shown that an airline could benefit by allocating its arrival slots depending on its preferences. In particular, the amount of total delay encountered and the number of flights canceled varies depending on the cost an airline assigns to the cancellation of a flight.

Figure 5 displays how an airline might trade off the number of canceled flights and the delay cost (the total number of periods of
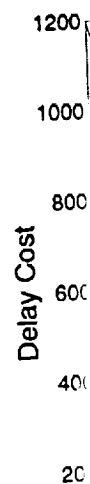
Time

'the airline
: the delay
ly incurred
: canceling

al slots in
t of exper-
s at a hub
te periods.
e within a
d into the
er per pe-
hedule for
rival pat-
n to have
-period or
ry period.
A's banks
10 flights

tisfy the
o 90, 80,
allocated
se condi-
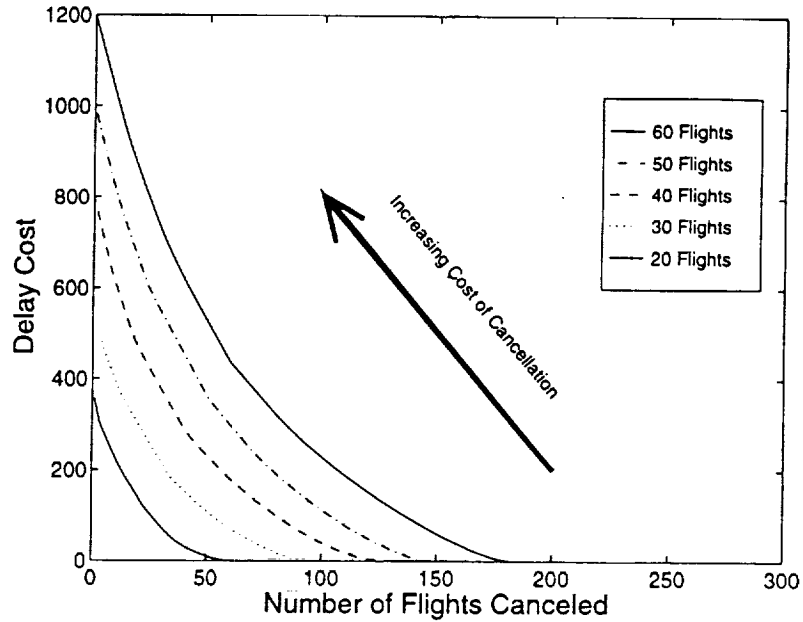ating its
amount
d varies
a flight.
mber of
riods of

Figure 5. Number of canceled flights vs. delay cost with changing cost of cancellation.

delay experienced by flights in the airline's schedule). The trade-off curves are given for banks of 20 through 60 flights. An airline that assigns a low value to canceling flights will cancel many flights; however for those flights which were allocated arrival slots, fewer delays will be incurred. An airline that assigns a high value to canceling flights will schedule many flights, resulting in high delay costs, but low total cost for the canceled flights. The conclusion from the figure is that different airlines can position themselves at varying places on the curve, a result achievable only with some form of decentralization.

These experiments also showed that airlines that consider their flights as being interdependent will allocate their arrival slots differently from airlines that view their flights as being independent. Figure 6 shows that the total amount of delay experienced by airlines that consider their flights independent is greater than the delay encountered by airlines that allocate arrival slots taking the true dependency of flights into account. The 'x' marks represent an airline using the IF model, whereas the 'o' marks represent the results of the C/D model. The curved line in the figure is a best fit line for the C/D model's results. The experiments showed that when few or many flights are canceled, the total delay is about the same for either model; however, when the total number of flights canceled is between 5 and 20 percent of all flights, airlines using the C/D model can reduce the delay for each flight flying by an average of 1 to 3 periods (15 to 45 minutes) as compared with the IF model. The IF model selects less
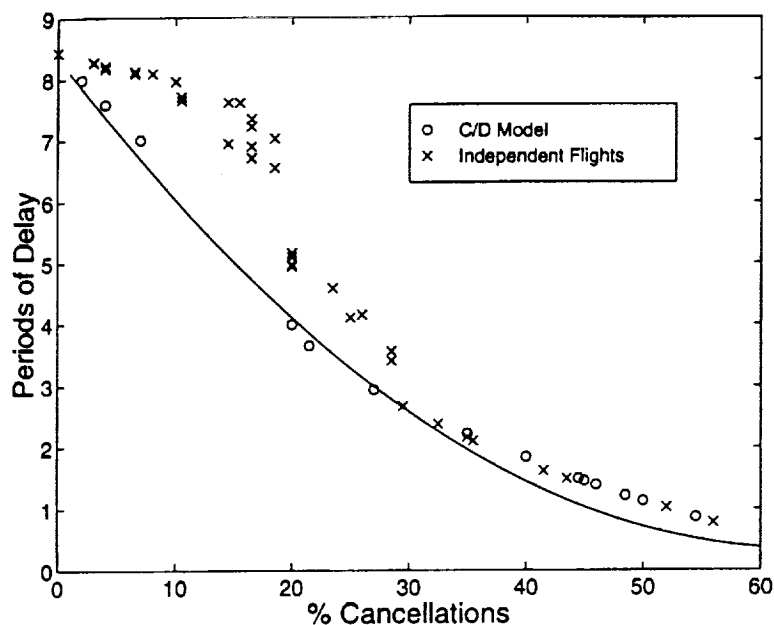
Figure 6. Comparison of IF and C/D models.

expensive flights to cancel, while the C/D model selects flights that
will reduce both the delay and cancellation costs. Again, a partially
decentralized system is required for airlines to achieve these results.

The schedules resulting from allowing airlines to allocate their
arrival slots can vary greatly, depending on how an airline views the
cost of canceling flights and the cost of separating them from their
banks. Figure 7 shows a result of applying the C/D model for an
example where an airline had five banks of flights scheduled to arrive
in five consecutive hours. In each 15-minute period, the airline was
assigned a varying number of arrival slots based on the original
schedule of all the flights scheduled into the airport. Each distinct
slot is depicted by a circle, square, or triangle. A circle with a number
in it indicates that the slot was assigned to a flight from the indicated
bank number, and that the flight arrived with the rest of the bank.
A square represents a slot not used by the airline because it canceled
some flight and did not substitute another into the available slot. A
triangle with a number in it indicates that some flight from the
indicated bank number was assigned to that slot, but the flight was
separated from its bank. The figure shows that four flights from bank
2 were separated from the rest of the bank. Even though those four
flights arrived immediately after all of the rest of bank 2, the other
flights would not wait for the connecting passengers from those four
flights. Similarly, two flights from bank 3 and one flight from bank 4
were scheduled to arrive after bank 5 was completed, with no
possibility for passengers on those flights to make their original
connections.

flights that
, a partially
lese results.
locate their
ie views the
i from their
iodel for an
led to arrive
airline was
the original
ach distinct
h a number
ie indicated
if the bank.
it canceled
able slot. A
it from the
e flight was
s from bank
i those four
, the other
i those four
rom bank 4
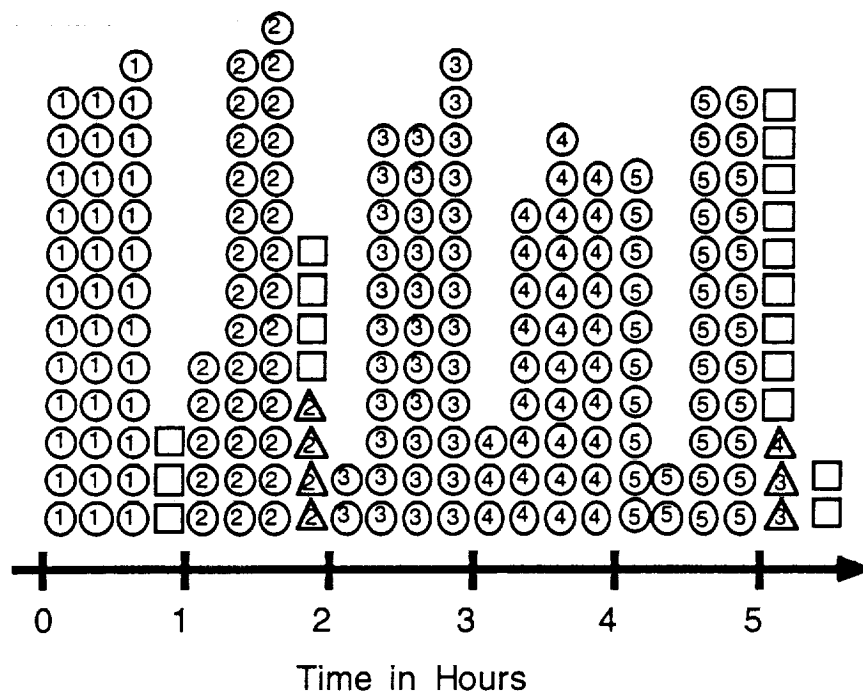d, with no
ir original

Figure 7. Assignment of flights to slots by C/D algorithm.

In the example of Figure 7, several slots were not used, in particular many of the slots occurring in the sixth hour. Under current practice, airlines have little incentive to reveal the presence of these open slots until it is too late for their competition to occupy them. While allowing a limited form of decentralization can aid airlines in allocating their arrival slots in accordance with their preferences, additional decentralization involving some exchange of slots between airlines as described in row 2d of Table 1 would be needed so that these unoccupied slots might be filled. Alternatively, either of the more decentralized plans indicated in row 2e of the table would likely ensure that open slots would not exist during these periods of reduced capacity.

CONCLUSIONS

This paper has presented a review of issues related to the evolution of the ATFM system in the United States toward more decentralized

decision-making environments. A range of conceivable alternative concepts has been broadly outlined. For the short- to medium-term future, it would seem that one of these alternatives, described in more detail, may be technically feasible, as well as consistent with expressed airline preferences and with the current emphasis on advancing Free Flight. It should be noted, however, that many open technical, procedural, and operational issues need to be addressed with regard to such a partially decentralized system. A flexible simulation environment to support the evaluation and assessment of the benefits and costs stemming from such a concept is also needed; we have described such an environment currently under development, along with a set of metrics being used to evaluate alternatives. It is also clear that it is very difficult at this time to predict exactly how major airspace users, such as the airlines, would behave in decentralized ATFM environments. An approach to modeling and understanding some aspects of this behavior has been described here, but this is a general area that will require extensive basic research over the next several years.

ACRONYMS

ATFM	air traffic flow management
ATM	air traffic management
C/D	Cancellation/Delay
CDM	collaborative decision making
CTA	controlled time of arrival
CTAS	Center-TRACON Automation System
IF	Independent Flights
MAR	Managed Arrival Reservoir
NAS	National Airspace System
NRP	National Route Program

REFERENCES

Adams, M., S. Kolitz, and A. Odoni (November 1996), *Evolution Toward a Decentralized Air Traffic Flow Management System*, CSDL-R-2767, C.S. Draper Laboratory, Cambridge, MA.

Andrews, J. W. (1993), "Impact of Weather Event Uncertainty upon an Optimum Ground-Holding Strategy, *ATC Quarterly*, 1, 59–84.

Bertsimas, D. and S. Stock (1994), "The Multi-Airport Flow Management Problem with En Route Capacities," Working Paper, Operations Research Center, MIT, Cambridge, MA (to appear in *Operations Research*).

Bohme, D. (1994), "Improved Airport Surface Traffic Management by Planning," in *Advanced Technologies for Air Traffic Flow Management*, H. Winter and H.-G. Nußer (eds.), Springer-Verlag, Berlin.

Brunetta, L., G. Guastalla, and L. Navazio (1995), "A New Approach for Solving the Multi-Airport Ground Holding Problem," Working Paper, University of Padua, Padua, Italy.

DeArmon, J. and A. R. Lacher (1996), "Aggregate Flow Directives as a Ground Delay Strategy: Concept Analysis Using Discrete Event Simulation," Working Paper, The MITRE Corporation, McLean, Virginia.

Erzberger, H., T. Davis, and S. Green (1993), "Design of Center-TRACON Automation System," *Proceedings AGARD Conference on Machine Intelligence in Air Traffic Management*, Berlin, Germany.

Federal Aviation Administration (April 1996), *FAA Free Flight Action Plan*, Washington, DC.

Hocker, G. A. (1994), *Airport Demand and Capacity Modeling for Flow Management Analysis*, Master's Thesis in Operations Research, MIT, Cambridge, MA.

Milner, J. M. (1995), *Dynamic Slot Allocation with Airline Participation*, Ph.D. Dissertation, MIT, Cambridge, MA.

Odoni, A. R. (1987), "The Flow Management Problem in Air Traffic Control," in *Flow Control of Congested Networks*, A. R. Odoni, L. Bianco, and G. Szego (eds.), Springer-Verlag, Berlin.

Richetta, O. and A. R. Odoni (1993), "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control," *Transportation Science*, 27, 228–238.

Robinson, J. D. (1992), *A Simulation Testbed for Flow Management Analysis in Air Traffic Control*, Master's Thesis in Operations Research, MIT, Cambridge, MA.

RTCA, Inc. (1995), *Final Report of RTCA Task Force 3, Free Flight Implementation*, Washington, D.C., November.

RTCA, Inc. (1996), *Free Flight Action Plan*, Washington, D.C., August 1996.

Vasquez-Marquez, A. (1991), "American Airlines Arrival Slot Allocation System (ASAS)," *Interfaces*, 21 (1).

Wambsganss, M. C. (1997), "Collaborative Traffic Flow Management," in *Air Traffic Control Quarterly*, 4, (2).

Yu, J. (1996), *Airport Capacity and Regional Weather Modeling*, Master's Thesis in Operations Research, MIT, Cambridge, MA.

BIOGRAPHIES

Milton Adams: Dr. Adams has been at the Charles Stark Draper Laboratory since 1972 and is currently the Associate Director of Applied Information and Automation Systems. He spent the academic year 1991–92 visiting at MIT, where he taught courses in multivariable control and large-scale systems. His technical interests include the design, development, and implementation of real-time, mission planning systems for unmanned vehicles and the analysis and design of traffic flow planning and control for large-scale rail and air traffic systems. Dr. Adams received the Sc.B. and Sc.M. degrees in Electrical Engineering from Brown University and the Sc.D. in Aeronautics and Astronautics from MIT.

Stephan Kolitz: Dr. Kolitz joined the Charles Stark Draper Laboratory in 1989 and is currently a Principal Member of the Technical Staff in the Simulation and Decision Technologies Division. His technical interests include planning and scheduling in the management and operation of large-scale transportation systems and command and control systems. Prior to joining Draper Laboratory, Dr. Kolitz taught at Loyola University of Chicago, the University of Illinois at Chicago, and the University of Massachusetts at Boston. Dr. Kolitz received the S.B. degree in Physical Science from MIT, the M.A. degree in Economics from Yale University, and a Ph. D. in Operations Research from Northwestern University.

Joseph Milner: Dr. Milner is currently an Assistant Professor of Operations and

Manufacturing Management at the Olin School of Business at Washington University in St. Louis, where he has been since 1995. His research interests include integer and network programming, air traffic flow management modeling, and applications of operations management to the service sector. Dr. Milner received the B.S. in Operations Research and Industrial Engineering from Cornell University in 1989 and the S.M. and Ph.D. in Operations Research from MIT in 1991 and 1995.

Amedeo Odoni: Dr. Odoni is the T. Wilson Professor of Aeronautics and Astronautics and a Professor of Civil and Environmental Engineering at MIT. He has done extensive research on airport and airspace modeling and simulations, on air traffic flow management, and on airport airside and landside operations. He has also acted as technical advisor to many airport authorities in the United States, Europe, and Australia and to several national and international civil aviation organizations. Dr. Odoni is the author or co-author of 5 books and more than 60 technical papers and has obtained S.B. and M.S. degrees in Electrical Engineering and a Ph.D in Operations Research at MIT. In 1991, he received the FAA Administrator's National Award for Excellence in Aviation Education.

Agg
as a
Con
Disc

James S.

When weat!
at an airpo
Currently, ï
specific del:
carriers car
would be de
ture schedu
to be addre
Flow Direct
port into Y-
flexibility, b
balance at ℕ
value for Y, t
impact of A
Flow Manag
erate possib
assess dema
demand/cap:

INTRODU:

The field of
two parts: s: