Collaborative Decision Making in Air Traffic Flow Management

ROBERT HOFFMAN WILLIAM HALL MICHAEL BALL
AMEDEO ODONI MICHAEL WAMBSGANSS

Key words and phrases: Air traffic flow management, collaborative decision making, compression, free flight, ground-holding problem, integer programming, ration-by-schedule.

ABSTRACT This paper describes the models and procedures underlying collaborative decision-making (CDM), a recently introduced approach to air traffic flow management. CDM is a first step toward free flight and promises to have a profound impact on research in traffic flow management. The initial focus of this joint FAA-industry venture has been the improvement of ground delay programs and their progression toward an environment in which the NAS users would gain more control and flexibility over their operations. CDM operational procedures for ground delay programs are now in place at all airports in the United States. CDM is extending to other areas of air traffic management, such as the routing of aircraft around overloaded sectors of airspace and inclement weather. This paper discusses the new roles of government and industry that have been forged by CDM philosophies as well as the algorithms, methods and technologies employed by CDM. We document the current efforts of CDM to improve the air transportation system and discuss the impact CDM will have on future research in air traffic management.

DRAFT: Saturday 9-12-98

1. INTRODUCTION

The traffic demand placed on the national air transportation system has increased steadily since the early days of flight. Until recently, the system capacity was adequate to handle this demand, but this is no longer the case. [**Bill, Amedeo: we need another stat here, perhaps, in addition to the following one.**] Indeed, the average gate-to-gate time for a flight between Chicago's O'hare and Boston's Logan airports is the same today as it was in 1950 [**Bill: do we have a ref for this?**]. Any potential reduction in travel time due to faster aircraft has been nullified by the increase in air traffic delays. According to several studies (see [19], [8] and [12], the delays incurred during routine operations will grow even larger in the next ten years.

Delays in the National Airspace System (NAS) in the United States are primarily due to limited capacity at the airports. Although there are restrictions

[Bill Hall will supply an appropriate graph]

Figure 1: Peaks in Demand

in en route capacity in the NAS, they are dwarfed by the restrictions imposed by airport capacity. Over 90% of the delays in air travel are caused by capacity shortfall and subsequent congestion at only 22 of the 18,224 airports in the country [**where did this statistic come from?**]. Not surprisingly, these 22 locations coincide with the major metropolitan areas, which are prime destinations and points of origin for passengers. Furthermore, air traffic demand at these airports is projected to grow very rapidly over the next several years. [**need statistic**] For these reasons, limited airport capacity is a major concern to the airlines, the Federal Aviation Administration (FAA), the aircraft manufacturing sector, and the public.

One aspect of the recurring capacity shortfall is that the demand at many of the major airports has grown far beyond what the airport was designed to handle. [**cite SFO as an example - call Zane to find out what they were designed to handle and are handling now.**] The mean level of demand is considerably less than the mean capacity but the uneven distribution of demand causes demand to exceed capacity regularly. Demand tends to peak during the morning and evening hours, when most passengers prefer to travel (see Figure 1). There are also peaks at the even hours and half-hours because the airlines tend to schedule flight departures at the times passengers are most likely to request.

Another cause of demand peaks is the operation of hub-and-spoke systems by many of the major carriers. This system provides an economically sound way for a large carrier to establish connectivity and more frequent service over the airport-pairs that it serves. Rather than scheduling a separate flight for each city-pair, the carrier schedules flights between many of the airports and a specified hub airport. Passengers can then take a flight to the hub airport, change aircraft, and depart for their destinations on the new aircraft. This practice has further aggravated crowded conditions at many of our major metropolitan airports.

Since there would be little support for any long-range plan that asks the airlines to curtail well-established operational paradigms or that asks the public to travel at odd times of the day to unpopular locations, it has proved to be more promising to focus efforts on increasing airport capacity rather than redistributing demand. Long range plans to increase airport capacity include the construction of new facilities and the expansion of existing ones. Unfortunately, such plans are often hindered by prohibitively expensive real estate and by the protests of surrounding communities concerned with the potential increase in noise and automotive traffic congestion. As a result, efforts in both research and daily air traffic management are concentrated on using current resources as efficiently as possible.

Many of the short-term demand-capacity inequities in the NAS are smoothed out by FAA traffic flow managers by tactics such as vectoring and miles-in-trail restrictions. Longer-term demand-capacity inequities (on the order of several hours) are more problematic. These usually occur when an airport acceptance rate is drastically reduced because of inclement weather, airport construction, or special runway operations. In these instances, air traffic flow managers at the FAA employ ground holding strategies in which aircraft bound for an afflicted airport are held at their points of origin in lieu of costly and hazardous airborne holding that would occur if they were allowed to depart on schedule. The most prominent of these strategies is the ground delay program (GDP), which is an initiative taken by the FAA to lower the arrival rate to a level that can be safely handled by airport controllers. GDP's have been traditionally run in lengths of 4-6 hours, although GDPs as long as 12 hours are becoming more commonplace.

For the last 50 years, the FAA has used a central-decision-maker paradigm to manipulate air traffic flow in the NAS in a way that maximizes safety and makes equitable use of NAS resources. But the growth in air traffic has pushed this paradigm to its operational limits. The network of the NAS has become so large and complex that many air traffic initiatives are reduced to a set of localized efforts. These efforts are, in turn, impaired by a lack of global information. The FAA does not have the communications infrastructure to support the exchange of information necessary for a more global management of air traffic. Much of this information is either unavailable to the FAA or there is no real-time data feed for their transmission.

The FAA often lacks the expertise in airline operations to make refined decisions regarding individual aircraft. In particular, there has been growing dissatisfaction with GDP's and the way that they are managed.

The idea of a GDP was initiated during the 1970's when fuel prices and availability were more of a national concern. Although fuel is the foremost airline expenditure, the waning of the fuel crisis of the 1970's has increased the willingness of airlines to forsake the safety of ground holding in favor of a chance to have their aircraft arrive sooner, although it means running a risk of airborne holding. In some flight instances, they even find predetermined airborne delay preferable over ground holding. The implementation of a GDP is now seen almost strictly as a government intervention necessary for the security of air traffic conditions at a terminal space. Combined with an airline inclination toward "free market", there has been growing resentment against FAA decisions that are more economic in nature than safety related.

There are many technological innovations with widespread applicability to aviation navigation and air traffic management that have not yet been put into place.[**Bill: can you fill out this paragraph?**]

All of these elements, growing demand, crowded airports, dissatisfaction with GDP's, and new technologies have set the stage for an environment in which the users of the NAS have more control over their economic destinies. For this reason, the central-authority paradigm has fallen under considerable scrutiny in recent years. NAS users would like to assume a wider range of movement within safety related constraints and to establish procedures for relaying to the FAA their preferences from a list of operational alternatives.

These trends have paved the way for *free flight*. In its strictest sense, this is a new concept for the operation of aircraft while en route. The FAA currently dictates that aircraft fly from one waypoint [**Is this the correct term?**] to another in the NAS so that their position can be monitored by ground-based equipment. Recent technologies such as GPS (global positioning system) may have made this operational paradigm obsolete. With GPS capabilities, it is now possible to track of the precise location of an aircraft anywhere in the world. Channeling a flight from one waystation to another is usually fuel-inefficient and can produce points of congestion. It has been proposed that aircraft equipped with GPS be allowed to fly the route of their choice. In some cases, this would be a flight path that minimizes fuel consumption while, in others, it might be a geodesic that minimizes en route time.

In a larger sense, free flight is the concept that the airlines should be able to operate more freely within the bounds of safety and overall traffic flow; that the primary function of the FAA would be as a NAS service-provider, whose job it is to maintain the big picture and to alert the users of the NAS of constrained situations, such as air traffic congestion and inclement weather. Collaborative Decision Making (CDM) is one of the first, and perhaps largest, step toward the realization of free flight in the NAS. The challenge for CDM and the FAA has been to find a way to combine information from the many competing interests to produce solutions that all NAS users can agree are good. The information from the competing interests must be used in a way that motivates them to provide accurate information. In the literature of game theory, this is called incentive compatibility (see, for example, [13]).

In a short period of time, the CDM effort has shown that government and industry can work together toward a common vision of improved conditions in the NAS, that the FAA is willing to relinquish some of its control, provided that it does not jeopardize overall safety, and that the airlines are willing to accept the responsibilities incurred by their increased degrees of freedom. The aviation industry and the FAA are now working together in an unprecedented fashion. New methodologies for air traffic management have been realized through algorithms developed by the CDM working group. The communications infrastructure established by CDM has supplied to both the FAA and participating airlines a common arrival demand picture at every major airport in the United States. On January 20, 1998, prototype operations of CDM practices were put into place at two U.S. airports (San Francisco-SFO and Newark-EWR). Later in 1998, prototype CDM operations went into effect at all major airports in the United States.

One of the essential components for the success of CDM is the concentration of efforts on well-defined, attainable goals. At first, these efforts were focused on key aspects of the development and implementation of ground delay programs. Spurred on by its initial success, CDM efforts are now ferreting out other problematic areas of air traffic control, such as the routing of aircraft around crowded or weather impacted airspace.

CDM and its parent concept, free flight, represent a revolutionary movement in air traffic management. Roger Beatty of American Airlines has declared that “CDM has the potential to be as important to airline operational centers as yield management has been to reservation systems.” The growing acceptance of CDM idealogies is rapidly dating research in air traffic; some of the most fundamental presumptions must now be reexamined in the light of CDM philosophies. As CDM philosophies forge new paradigms in government-industry working relations, opportunities are being created for innovative applications of modelling and optimization techniques.

The next section of this paper, section 2, provides necessary background information and a brief literature survey to put into context the CDM efforts. Section 3 describes the current status of CDM and its most prominent accomplishments. Section 4 details the procedures and operational paradigms developed by CDM that have been put into place and are bound to impact the future of air traffic management. Section 5 discusses past and present research topics that have spun off from CDM and speculates on the future direction of CDM. The paper concludes with a few closing remarks.

2. BACKGROUND, LITERATURE AND ISSUES IN GROUND HOLDING

The original focus of CDM efforts was the implementation and management of ground delay strategies. Many of the underlying philosophies of CDM and its operational paradigms have been refined in this setting. Also, the development of ground holding strategies have been a continued source of research. We now present relevant background information on ground holding strategies, a brief literature survey and a discussion of the some of the issues surrounding ground holding that have motivated CDM activities.

The Air Traffic Control Systems Command Center (ATCSCC) facility of the FAA monitors airports throughout the United States for capacity shortfall. Whenever it is predicted that the number of flights arriving at an airport within a 15-minute time interval will exceed the number of flights scheduled to land, the ATCSCC is required by FAA regulation to take some form of action.[1] Short-term periods of capacity shortfall are alleviated by airborne tactics such as re-routing and variations in airborne speed. Longer-term periods of capacity shortfall are met by the ATCSCC with ground holding strategies in which aircraft are held at their departure gates in lieu of costly and dangerous airborne delay.

In some cases, the ATCSCC will issue a ground stop in which all flights destined for an afflicted airport are held on the ground at their departure gates until airport arrival capacity rises above demand. These ground stops are reserved for extreme cases in which arrival capacity was severely underestimated or dropped suddenly without warning.

The primary tool of the ATCSCC for addressing arrival capacity shortfall is a ground delay program (GDP). In a GDP, each flight scheduled to arrive at an afflicted airport during a fixed time period is held at its departure gate long enough to ensure that it will be able to land without delay. For instance, if flight f is due to arrive at airport A at 12:00 and it is known that f will not be able to land until 12:30 due to limited arrival capacity at A, then f would be held at its departure gate for 30 minutes. The construction of a GDP requires the assignment of both a controlled time of departure (CTD) and a controlled time of arrival (CTA) to each incoming flight . Since en route travel times can be predicted with reasonable accuracy, the CTD of each flight is easily calculated once the flight has been assigned an arrival slot and its CTA is known.

The primary purpose of a GDP is to take control of the airspace surrounding an airport so that excess arrivals do not interfere with the safety and operation of the airport. While the airlines respect the need for such a program, the implementation of a GDP is generally met with trepidation by the airlines. This is partly for the obvious reason that it heralds large-scale delays

---

[1] Much of the information in this section concerning the operation of the ATCSCC was obtained through meetings with ATCSCC personnel.

and government intervention. But also, it exposes airline operations to miscalculations on the part of the ATCSCC. For instance, if the future capacity of an airport is overestimated, then planes will absorb delay in costly airborne holding patterns rather than on the ground at their departure gates. On the other hand, if future capacity is underestimated, then arrival slots become a wasted resource and flights absorb unnecessary ground delay. The airlines find it particularly disruptive to their operations when a GDP is aborted in mid-operation. This happens whenever the ATCSCC has clear evidence that the original capacity and weather forecasts were overly pessimistic. Estimates vary between the airlines and the FAA as to the percentage of GDP's that are aborted, but these estimates are as high as 60%.

2.1. Arrival Slot Allocation: the Grover-Jack Algorithm

Prior to CDM, the Grover-Jack algorithm was being used to implement ground delay programs at all U.S. airports. A grasp of the fundamental concept behind this algorithm is crucial to an understanding of the structure of ground delay programs and the proposals made by CDM.

In preparation for execution of the Grover-Jack algorithm, the time horizon is divided into hourly periods $t = 1, 2, ..., T$. The airport acceptance rate (AAR) of a time period t, denoted X_t, is defined as the number of aircraft that can be accepted during period t. (Strictly speaking, this is a capacity, not a rate. Nonetheless, it is the established terminology in the air traffic community.) A value of X_t is set for each time period t by a traffic flow manager at the ATCSCC. In spirit, at least, the Grover-Jack algorithm assigns controlled arrival times as follows. First, a list of incoming flights is formed, ordered by increasing estimated times of arrival. Then, the first X_1 flights on the list are assigned to the first time period, the next X_2 flights are assigned to the second time period, and so on, preserving order of the list. The net effect of the Grover-Jack algorithm is to stretch out the list of incoming flights over time so that arrival demand meets capacity in each time period.

In practice, there are several complications surrounding the execution of the Grover-Jack algorithm. Since a flight cannot be assigned to a time slot earlier than its ETA, some time slots will be passed over during the assignment process and have no flight assigned to them; international flights, general aviation, and flights airborne at the time of formulation of a GDP are exempt from the program, meaning that they cannot be issued a ground delay (see [16] or [17] for more on exemptions). In addition, the traffic flow manager may choose to exclude other categories of flights from ground delay, usually based on geographical location of point of origination. The arrival demand of these flights must be taken into account when assigning new arrival times to flights.

Once the controlled time of arrival (CTA) of each flight has been determined by the Grover-Jack algorithm, a controlled time of departure (CTD) must also be set. This can be back-computed from the CTA, once an en route time is estimated: the CTD is simply the CTA minus the en route time. The FAA uses its own estimates of en route time. This is, at times, a source of contention with the airlines. The estimation of en route time naturally involves two parameters: cruising speed, and the route flown. Both of these, in turn, are heavily related to fuel consumption, which is the primary expenditure for an airline operation. Thus, en route time estimation depends on airline intentions and the economics of its daily operations. It has been proposed that the airlines determine their own departure times, once arrival times have been assigned to them. See [reference: Airline Operational Control Overview, Sept. 1995. Airline Dispatchers federation and Seagull Technology, Inc., Sunnyvale, CA.] [**Bill: are these the correct references?**]

Grover-Jack is a first-come, first-served algorithm with respect to estimated time arrival for the allocation of arrival slots. Although the airlines agree that some sort of allocation scheme is required when arrival demand greatly exceeds capacity, they have objected to the use of ETA (estimated time of arrival) as the basis for the initial allocation of arrival slots. Their objection is called the *double-penalty issue* and is founded on the following dilemma. Suppose that airport capacity is cut in half (a frequently occurring scenario). Then in order to make demand equal to capacity, the ETA-based arrival sequence should be dilated to twice its current length. This means that every time a flight is moved x minutes down in the original arrival sequence, it is effectively moved down $2x$ minutes in the final arrival sequence. In particular, if an airline reports, say, a 30-minute delay of one of its flights prior to the implementation of the Grover-Jack algorithm, then the flight will receive 60 minutes of delay: 30 minutes of original delay plus another 30 minutes of imposed delay.

One might argue that if an airline experiences some sort of internal delay and cannot take off on time, then it should suffer the consequences. However, adding delay to an already-delayed flight is viewed as "kicking a guy while he is down" (consider this from the passengers point of view). Also, the airlines offer the argument that, since the ultimate goal of a GDP is to slow down the rate of arrivals, self-inflicted delay meets this need and should serve toward any government-issued delay.

Whether or not one agrees with the airline's position on the double-penalty issue, it has proven to be a strong disincentive for the airlines to submit updated ETA's to the FAA. The CDM working group has devised a scheme for the initial rationing of arrival slots that removes this disincentive. The scheme is called *ration-by-schedule* (RBS) because it rations slots according to scheduled arrival times instead of estimated arrival times. Scheduled times are taken to be those as published in the official airline guide (OAG). For instance, suppose that American Airlines flight AAL101 is fifth in the scheduled arrival sequence. Then when a GDP is imposed and the arrival sequence is stretched out over time, American Airlines will maintain ownership of the fifth arrival slot, independent of the arrival status of AAL101. In the event that AAL101 is delayed, American Airlines has the opportunity to move another one of its flights into the fifth slot and move AAL101 to a later arrival slot that it can use. This way, the scheduled arrival resources of American Airlines are retained by American Airlines and the natural delay of AAL101 serves as its ground delay.

This example presupposes, of course, that AAL has another flight occupying an arrival slot feasible for AAL101. There are many such complications associated with the practical implementation of RBS. See section 5 for complete details of the algorithm. Also, it assumes the existence of a substitution procedure in which the airlines can freely manage the arrival resources that have been allocated to them. This is the topic of our next subsection.

2.2. Flight Substitution-Cancellation

The initial allocation of arrival slots by the FAA based on Grover Jack (or any allocation scheme, for that matter) may be in the best interests of the overall traffic flow but it may not be in the best interests of an individual airline. For this reason, the airlines reserve the right to reallocate their arrival slots amongst their own flights. Under existing rules, an airline can perform any of the following routines:

- cancel a flight;
- move a flight into an arrival slot vacated by one of its flights;
- swap the arrival slots (arrival times) of two if its flights.

Example 1. For example, suppose that arrival slots have been allocated to the flights of airlines A, B and C, as in Table 1.

Airline	Flight	ETA	CTA	Delay (minutes)
A	1	1000	1000	0
B	2	1005	1010	5
A	3	1010	1020	10
C	4	1015	1030	15
A	5	1020	1040	20
B	6	1025	1050	25
A	7	1030	1100	30
A	8	1035	1110	35
C	9	1040	1120	40
A	10	1045	1130	45

Table 1. Initial Allocation of Arrival Slots to Flights

Let us adopt the vantage point of airline A and determine their best substitution/cancellation strategy, based on the following information: The cost of delaying A3 is relatively low (e.g., it has a small number of passengers and the aircraft has a long layover); the cost of delaying A5 is quite high; passengers of flight A7 can be easily rerouted onto flight A8 (recall that in a GDP, these flights have a common destination). Then one strategy airline A might use is to

1. swap flights A3 and A5;
2. cancel A7 (transfer passengers to A8);
3. move A8 into slot vacated by A7;
4. move A10 into slot vacated by A8;
5. hold onto the slot vacated by A10.

This substitution process reduces the overall delay for airline A from 135 minutes to 80 minutes, as seen in Table 2.

Airline	Flight	ETA	CTA	Delay (minutes)
A	1	1000	1000	0
B	2	1005	1010	5
A	5	1020	1020	15
C	4	1015	1030	20
A	3	1010	1040	30
B	6	1025	1050	25
A	8	1035	1100	25
A	10	1045	1110	25
C	9	1040	1120	40
A	hold	-	1130	-

Table 2. Allocation of Arrival Slots after Airline A subs/swaps

Similarly, airlines B and C would make substitutions and cancellations beneficial to their operations. Throughout this paper, we will refer to this reallocation procedure on the part of the airlines as *substitution-cancellation.* The substitution-cancellation process is vital for the mitigation of damages incurred by delays in a ground delay program.

The substitutions performed by airline A happened to minimize their overall delay but this is not necessarily their objective. Overall delay minimization assumes that delay costs are quantifiable entities that are the same for every flight and that they increase linearly with delay. Substitution and cancellations decisions for a given airline are highly peculiar to its flight schedule, flight connectivity, and operational methodology. ■

2.3. The Need for Airline Incentives and Improved Procedures

The following example demonstrates that circumstances can (and do) arise under the Grover-Jack system in which there is an incentive for an airline to withhold cancellation information and, worse yet, that overall efficiency can suffer as a result of these withholdings.

Example 2. Assume that airport conditions allow for the intake of only one flight every ten minutes (this is chosen to be unrealistically small for the purposes of exposition). Under the Grover-Jack algorithm, flights would be ordered by ETA, then assigned to arrival slots so as to preserve the ordering. Table 3 shows the hypothetical assignments of 11 flights to 11 arrival slots and their respective delays under Grover-Jack. For instance, flight 2, with an ETA of 0700, has been given a controlled time of arrival of 0710, thus leading to a delay of ten minutes. The delay over all flights is 320 minutes.

Airline	Flight	ETA	CTA	Delay (minutes)
A	1	0700	0700	0
A	2	0700	0710	10
B	3	0705	0720	15
B	4	0705	0730	25
B	5	0710	0740	30
B	6	0710	0750	40
A	7	0710	0800	50
C	8	0720	0810	50
B	9	0740	0820	40
C	10	0740	0830	50
A	11	0830	0840	10
Total A				70
Total B				150
Total C				100
Total				320

Table 3. Delays Generated by Grover-Jack Algorithm

Further assume that flight 1 is cancelled by airline A. If airline A were to report this cancellation prior to the implementation of the ground delay program, then each of the subsequent flights (from all airlines) in the sequence would be effectively shifted up one arrival slot when the Grover-Jack algorithm is run. This would save 10 minutes of delay for each of the 10 remaining flights, resulting in a reduction of cumulative delay from 320 minutes (Table 3) to 220 minutes (Table 4). In particular, the delay for airline A would be reduced from 70 minutes to 40 minutes.

Airline	Flight	ETA	CTA	Delay (minutes)
A	2	0700	0700	0
B	3	0705	0710	5
B	4	0705	0720	15
B	5	0710	0730	20
B	6	0710	0740	30
A	7	0710	0750	40
C	8	0720	0800	40
B	9	0740	0810	30
C	10	0740	0820	40
A	11	0830	0830	0
Total A				40
Total B				100
Total C				80
Total				220

Table 4. Delays Generated by Grover-Jack Algorithm After Cancellation of Flight 1

It appears advantageous to airline A (and the two other airlines) to have reported this cancellation. However, consider the added benefit to airline A, if it had withheld the cancellation until after the execution of the Grover-Jack algorithm. Specifically, starting with the allocations in Table 3, Airline A could execute the following set of substitutions:

- move flight 2 into the slot vacated by flight 1 (CTA 0700) to save 10 minutes of delay;
- move flight 7 into the slot vacated by flight 2 (CTA 0710) to save 50 minutes delay.

The result is shown in Table 5.

Airline	Flight	ETA	CTA	Delay (minutes)
A	2	0700	0700	0
A	7	0710	0710	0
B	3	0705	0720	15
B	4	0705	0730	25
B	5	0710	0740	30
B	6	0710	0750	40
void	void	-	0800	-
C	8	0720	0810	50
B	9	0740	0820	40
C	10	0740	0830	50
A	11	0830	0840	10
Total A				10
Total B				150
Total C				100
Total				260

Table 5. Flight 1 Cancellation Goes Unreported; Airline A Makes Substitutions.

Through these simple substitutions, airline A has reduced its overall delay from 70 minutes to only 10 minutes. Contrast this with the delay of 40 minutes for airline A in Table 4. Unfortunately, the savings for airline A are at the expense of overall efficiency. Total delay (for all three airlines) is now 260 minutes, which is 40 minutes more than the 220 minutes of total delay, if the cancellation were reported. The 40-minute loss in efficiency is caused by the blockage of the upward movements of flights 8, 9, 10 and 11. Under the rules of substitution, a flight can be moved into a slot by an airline only if the slot was vacated by one of its own flights. Since the 0800 slot was vacated by a flight from airline A, none of the Flights 8, 9 or 10 (with respective airlines C, B and C) can be moved up into the 0800 slot. Flight 11 is a candidate for movement but even this flight is blocked from movement because its ETA of 0830 is too late for the corresponding slot arrival time (CTA 0800). This leaves the 0800 slot unutilized. ■

One can see the need for a compression mechanism to moves flights upward in the arrival sequence to fill slots vacated by competing airlines. This cannot happen, of course, unless there is notification of these vacancies and airline A has agreed to relinquish the corresponding slots. An incentive for notification can be provided by giving preference in this upward movement to the owner of a vacated slot and, more importantly, by giving appropriate compensation to that owner when one of its slots is rewarded to a competitor.

To see how a compression algorithm would work in this example, let us return to Table 5. The algorithm first recognizes the vacancy of slot 0800 created by the substitution of flight 7 into an earlier arrival slot. Since this slot is owned (initially assigned to) airline A, the algorithm searches the flights below this slot (flights 8, 9, 10, 11) for a flight from airline A that can moved into this slot. The only candidate is flight 11. We have already observed that the 0800 slot is infeasible for this flight because it will not arrive until (approximately) 0830. So, the algorithm instead selects the earliest flight feasible for that slot, independent of its owner. In this case, it happens to be flight 8, from airline C. Flight 8 is moved up one position, thus leaving a vacancy in the 0810 arrival slot. When trying to fill this slot, priority is given to airline A because it already lost a slot. The algorithm searches the remaining flights (9, 10, 11) for a feasible flight from airline A. Again, the only candidate is flight 11 and, again, its ETA of 0830 is too late for this 0810 slot. So, the 0810 slot is filled with flight 9, from airline B. This creates a vacancy in the 0820 slot, which, again, proves to be too early for flight 11, so flight 10 is moved into this slot. This time, however, when the algorithm searches for a flight from airline A to move into the 0830 slot vacated by flight 10, it finds that flight 11 is feasible. The final slot assignments are shown in Table 6.

Airline	Flight	ETA	CTA	Delay (minutes)
A	2	0700	0700	0
A	7	0710	0710	0
B	3	0705	0720	15
B	4	0705	0730	25
B	5	0710	0740	30
B	6	0710	0750	40
C	9	0720	0800	40
B	9	0740	0810	30
C	10	0740	0820	40
A	11	0830	0830	0
Total A				0
Total B				140
Total C				80
Total				220

Table 6. Slot assignments after the Compression Algorithm.

In effect, the vacancy of slot 0800 by Flight 7 created a stream of upward substitutions of the flights below it. The ownership of slot 0800 by airline A was transferred downward through this stream until it came to rest at the first slot in which it could benefit airline A. Flights 8, 9 and 10 created what is known as a “bridge” for this transfer. The compression algorithm was not able to fill the 0800 slot with a flight from airline A but it got the best candidate “as close as it could”.

Since Flight 11 could not possibly be moved up any farther, airline A is satisfied. Because all slots are utilized, total delay (for all three airlines) has been restored to the optimal (lowest) value, 220 minutes. Airlines B and C were able to make use of vacancy created by Flight 7. Airline A delay has dropped from 10 to 0 minutes, airline B delay has dropped from 150 to 140 minutes and airline C delay has dropped from 100 to 80 minutes. This is a win-win situation: each airline is better off than if airline A had not reported the cancellation.

In actuality, there are many complications to be addressed in the formulation of such an algorithm. For instance, should a flight be denied placement into an arrival slot just because it is estimated to arrive one minute later than the end of the arrival slot? Or, what if a slot is vacated by an airline A but there are no flights from airline A below that slot to be moved up, thus denying compensation to airline A - should the slot be left unutilized or given away to its competitors?

The CDM working group has developed a compression algorithm, which we will henceforth refer to as “the” compression algorithm. The many facets of compressing arrival slots have been captured in this algorithm through input from CDM participating airlines and through practical experience during prototype operations. The details of the algorithm are given in the section 5.

3. Overall CDM Architecture

3.1. The Cycle of Decision Making

Through ground delay program enhancements (GDP-E), CDM has established a cycle of exchange between the airlines and the ATCSCC that can effectively manipulate the arrival and departure status of flights before, and during, the execution of a ground delay program. See Figure 2. The ATCSCC initializes a ground delay program by executing the RBS (ration-by-schedule) algorithm. This forms an initial allocation of the arrival slots at the afflicted airport. Next, the airlines have the opportunity to reallocate the slots assigned to them by RBS by making cancellations and substitutions. This often leaves some slots with no flight assigned to them. The compression algorithm attempts to move flights up in the arrival heirarchy into these vacated slots, honoring original ownership, whenever possible. This results in a controlled time of arrival and controlled time of departure for each flight in the GDP.

One of the objectives of CDM has been to develop a set of simplified substitution rules (SSR) to provide a higher level of flexibility in the substitution-cancellation process. In fact, for many airlines, this has been a prime incentive

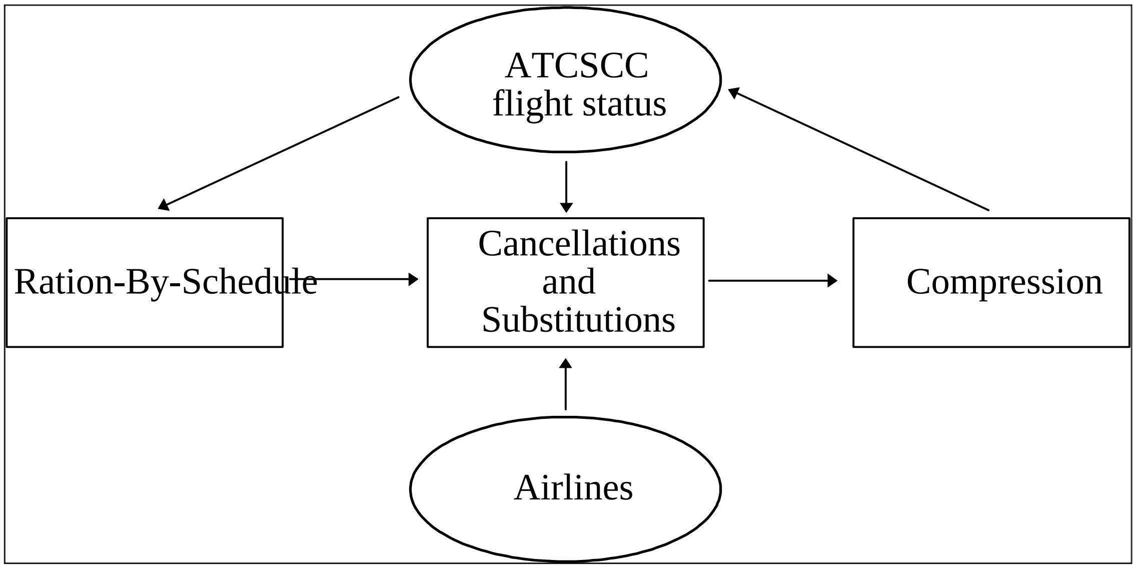

Figure 2: The Cycle of Decision Making During a GDP

for participation in CDM.

Both the initialial allocation of arrival slots and the substiution-cancellation were in place prior to CDM (RBS has replaced Grover-Jack). So, the truly new component of this cycle provided by CDM is the compression algorithm that makes efficient use of airport resources. One of the main strengths of the cycle is the fact that it can be executed several times during a GDP (usually, once per hour). This allows the airlines to mitigate delay damages incurred during a GDP, as events unfold.

The execution of the RBS algorithm can be done once, then by-passed in subsequent iterations of the cycle. However, the ability to repeat this initial allocation procedure allows the ATCSCC to periodically revise their estimations of airport capacity, as weather conditons become known. Prior to CDM, there was no allowance for revisions during a GDP, with the consequence that there was no mechanism for adapting to changing airport capacitites and inaccurate forecasts.

The combination of revisions and compressions in this cycle is projected to save [**millions of dollars over the next blank years - what source should we cite here? Mike W, can you help here?**].

3.2. Databases and Processing Strings

The implementation of CDM required a significant enhancement to the data management and distribution capabilities that support traffic flow management. Figure 3 illustrates that enhanced system. The key elements are

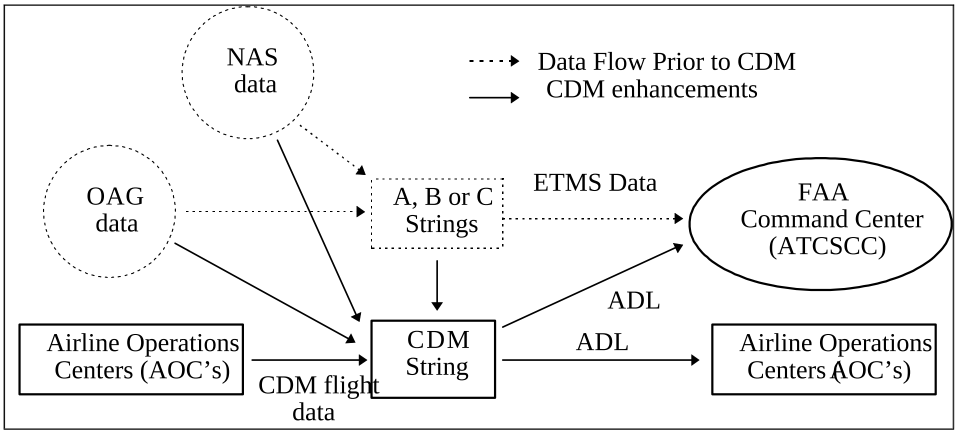

Figure 3: Databases and Processing Strings

described below.

ETMS (Enhanced Traffic Management System): Prior to CDM, this was the only flight information database available to the ATCSCC. Ground delay actions taken by the ATCSCC have been traditionally based upon data from ETMS.

ATMS (Advanced Traffic Management System): ATMS is the R&D version of ETMS. ATMS has all the ETMS data and functions plus additional data and functions that are being prototyped. In particular, ATMS merges into the ETMS flight database real-time schedule updates flowing from the airline operations centers (AOCs). At the time of this writing, both ATMS and ETMS feed flight information to the ATCSCC

A, B, and C strings: A "string" is a group of about ten computers that perform all the ETMS hub site processing; a string receives input data from many different sources, e.g., NAS, OAG, weather providers, specialists, and airlines, and then integrates this data in a database. This data is then made available to the various ETMS sites. There are two operational strings called A string and B string; there is also a test string called C string; these three strings should have the same data except for alterations to C string made as part of system tests. All three of these strings transmit ETMS data to the ATCSCC.

CDM string: This is the research and development string that receives all the input that is received by A, B, and C strings, but that in addition receives the data feed of real-time airline schedule updates developed by the CDM working group.

3.3. The AOCnet

The participating airline operations centers (AOCs), the ATCSCC, the Volpe National Transportation Systems Center (the hub site of the Enhanced Traffic Management System – ETMS), Metron (the developer of FSM) as well as certain other parties are all interconnected via a private intra-net, the AOCnet. This network is used to exchange CDM operational information.

Figure 4 describes the flow of information through the AOCnet. The Volpe hub site (shown on the left) forms ADLs (aggregate demand lists) based on three sources of flight information:

(1) the Airline Operations Centers (AOCs);
(2) the National Airspace System (NAS);
(3) the Official Airline Guide (OAG).

The ADLs are then pushed out over the AOCnet to the Air Traffic Control System Command Center (ATCSCC) and the AOCs. The ADLs are used by both the ATCSCC and the AOCs to generate the information displayed through the CDM decision support tool, FSM.

4. New Roles and Procedures

The responsibility of the FAA is to manage air traffic flow in a manner that makes safe, efficient and equitable use of resources within the NAS. The pre-CDM paradigm that guided the FAA was based on the belief that the FAA and its agencies should act as a centralized authority and decision-maker. This paradigm, as well as many of the procedures and standards for air traffic management, were developed just after World War II. There are several problems with maintaining the central decision-making paradigm. First, the National Airspace System has grown too large to be efficiently managed by one authority. Secondly, the airlines have been given little input into the decision making process.

A relationship in which decision making is shared between the scheduled carriers and the FAA was formalized in the document, "Roles and Responsibilities", written by the FADE program manager and airline representatives in early 1995. It specifies that the ATCSCC should remain a neutral service provider within the NAS and that its primary responsibility is to maintain situational awareness and alert the users of the NAS to situations that place constraints upon their operations. The users, on the other hand, are responsible for responding to those constraints with actions, intents and preferences that lie within the constraints specified by the service provider. In addition, they are responsible for supplying the service provider with accurate and timely information necessary for the monitoring and management of constrained situations.

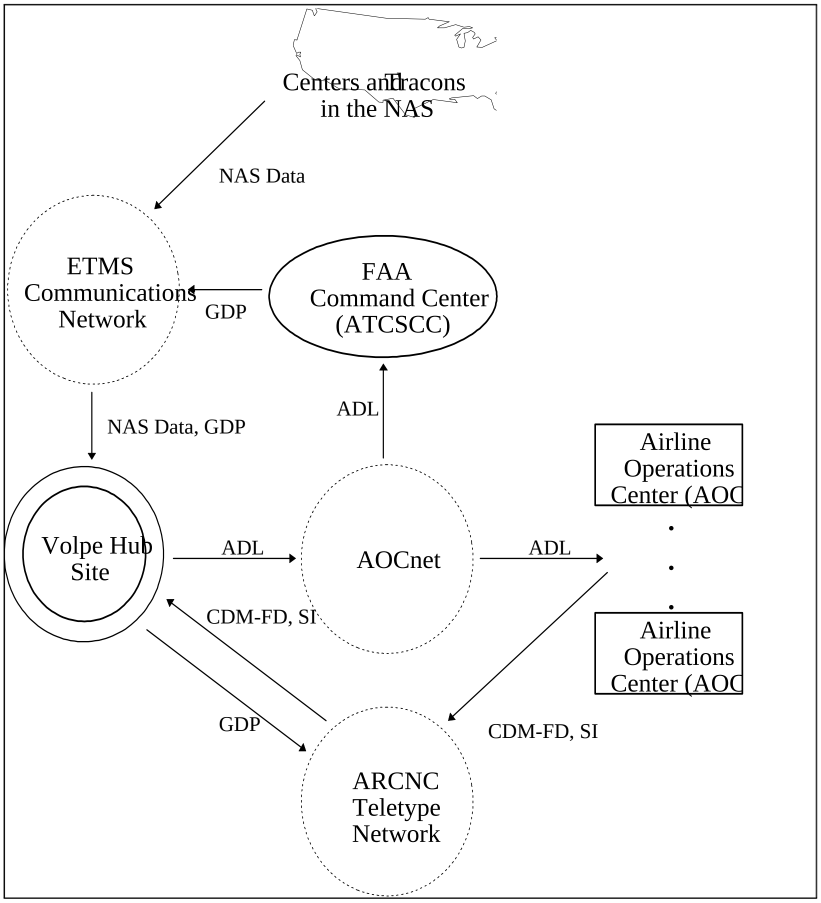

Figure 4: The AOCnet

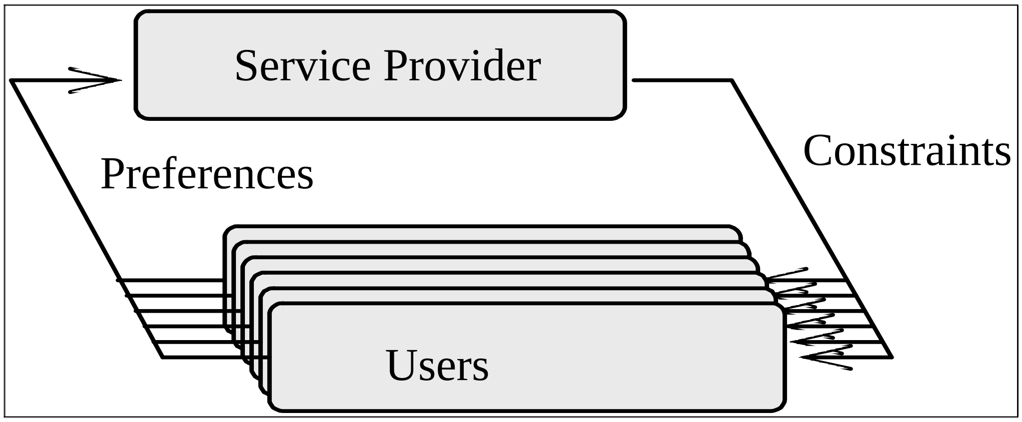

Figure 5: The constraint-dissemintation user-preference (CDUP) loop.

4.1. User Preference and Constraint Dissemination

The new roles and responsibilities specified for the service provider and users of the NAS have lead to a ubiquitous operational procedure for handling a constrained situation in the NAS, which we will call *constraint-dissemination and user-preference* (CDUP). This is a cycle of feedback between the users and service provider. A single iteration of the cycle begins with an analysis of the constrained situation and the dissemination of constraints from the service provider to the users. Next, the users asses their role in the constrained situation and the potential impact on their operations. Each user reacts with a set of preferences. Lastly, the service provider updates the information, honors or rejects user preferences, and the users take action. This completes one iteration of the cycle. The cycle repeats, perhaps at regular intervals, either until the constrained situation is played out or until it is alleviated through user reactions. See Figure 2.

In theory, CDUP allows for any degree of decentralized decision-making. At one extreme, the service provider provides only an aggregate picture of a traffic situation and the users shoulder the entire responsibility for alleviating the situation. For instance, suppose that a sector of airspace has been overloaded because a number of airlines have filed flight plans passing through it. Then the service provider would issue the projected demand and capacity of the sector, highlighting the inequity. Hopefully, enough flights would re-file their flight plan over another sector to bring demand back in line with capacity. This might require several iterations of the cycle. This type of collaboration is attractive because it gives the users maximum control over their operations. The problem with the maximally decentralized CDUP is that it assumes that user preferences and intentions are in accordance with overall traffic flow efficiency. This could fail when users are competing for highly valuable resources.

At the other extreme, the users are given a set of constraints by the service provider that effectively reduces their options to a single option. The user "preference" that is relayed back to the service provider is just an acknowledgment of compliance with their only option. In this case, CDUP reverts to centralized control. The centralized decision-making paradigm is exactly what CDM is trying to avoid. For most situations, the optimal amount of collaboration lies somewhere between these two extremes. A blend of centralized decision-making and user decision-making can be achieved by generating additional artificial constraints that severely limit user options. These options can be used to channel user actions toward a more global solution. The CDM GDP process outlined in the previous section represents such a blend and is now working successfully.

The practical application of this CDUP paradigm requires careful synchronization of user actions and constraint generation: there must be sufficient time between cycles for the service provider to correctly assess the situation and for the users to analyze their options and respond in their best interests. There needs to be a reliable mechanism for data exchange between the service provider and the users. Also, there needs to be clearly defined rules and procedures for the dissemination of constraints, user compliance, formatting of user preferences, and so on. These rules must be perceived as fair by all parties. A great deal of effort has been applied by the CDM working group to achieve these necessities for GDP's.

5. Ration-By-Schedule and Compression Algorithms: Fair Allocations

The two algorithms we are about to describe, RBS and Compression, embody the concepts of fairness for allocation of arrival resources embraced by the CDM community. These concepts, hence, these algorithms, are likely to serve as models for future notions of fairness and allocation of resources during other constrained situations in the NAS.

Input to algorithms (RBS and Compression): We assume that a GDP is about to be implemented from $start\_time = t_0$ to $end\_time = t_1$. Based on planned arrival acceptance rates for each future hour, the time horizon has been subdivided into contiguous time slots $S = t_0, t_0 + 1, ..., t_1$. The set of flights F that will be affected by the GDP form a partition Φ of airlines (general aviation comprises one "airline"). For each $f \in F$, the following data fields are required.

$AIR_f \in \Phi$ =airline of flight f
$ETA_f,\ ETD_f$ =estimated arrival/departure times for f
$CTA_f,\ CTD_f$ =controlled times of arrival/departure for f
$EDCT_f$ =estimated departure clearance time of f

$OGTA_f,\ OGTD$ =original gate time of arrival/departure of f
(in general, $O-$ prefix implies an original time e.g., $OETA_f$)

For each time slot $t \in S$, the fields $flight\,(t) \in F$ and $owner\,(t) \in \Phi$ are established to reflect the flight assigned to t and the controlling airline, respectively. In addition, each time slot t will maintain a status, $status\,(t)$, which can take on one of four states: *open* (to compression, to be filled with the next available flight), *released* (by the controlling airline to compression to be filled, hopefully, by one of its own flights), $filled$ (by some flight), or *held* (by an airline for later use).

5.1. The RBS (Ration-By-Schedule) Algorithm

The purpose of RBS is to make an initial allocation of arrival slots to flights bound for the airport that is subject to degraded arrival capacity. Essentially, the flights are placed in a list ordered by their arrival times as published in the OAG (Official Airline Guide), then the flights are assigned to the arrival slots, preserving order of the list.

1. Find the set of flights, $I \subseteq F$, to be included in (affected by) the ground delay program. Let I $= A \bigcup B$, where

$$A = \{f \in F : t_0 \le ETA_f \le t_1\}$$
$$B = \{f \in F : t_1 < ETA_f \quad and \quad (t_0 \le CTA_f \le t_1 \ \ or\ \ t_0 \le OGTA_f - Taxi \le t_1)\}$$

2. Find the set of flights, $E \subseteq I$, to be exempted from (delay within) the program.

   This set is user specified.

3. Split non-exempted (but included) flights $I - E$ into two disjoint sets, F_1 and F_2, defined

   via

$$F_1 := \{f \in I - E : f\ has\ been\ assigned\ a\ CTA\ at\ least\ once\}$$
$$F_2 := (I - E) - F_1$$

   Note: $I = E \smile F_1 \smile F_2$.

4. Compute *earliest*_CTA for each $f \in I$

   if $f \in E$,

$$earliest\_CTA := \begin{cases} ETA_f\ ,\ if\ f\ has\ a\ slot\ ID \\ min\,(ETA_f,\ OGTA - Taxi)\,,\ else \end{cases}$$

if $f \in F_1 \smile F_2$,

$$earliest\_CTA := \begin{cases} OGTA - Taxi, & if\ current\_time \leq OETD_f \\ current\_time + p + \left(OETA_f - OETD_f\right), & otherwise \end{cases}$$

Note: p is a positive parameter, user-specified

5. Create queues of flights.

   (a) $Q\left(E\right) := q_1^E, q_2^E, ..., q_{|E|}^E$, from E using ETA as a priority.

   (b) $Q\left(F_1\right) := q_1^{F_1}, q_2^{F_1}, ..., q_{|F_1|}^{F_1}$, from F_1 using CTA as a priority.

   (c) $Q\left(F_2\right) := q_1^{F_2}, q_2^{F_2}, ..., q_{|F_2|}^{F_2}$,from F_2 using $\left(OGTA - Taxi\right)$ as a priority.

   (d) $Q := q_1^E, q_2^E, ..., q_{|E|}^E, \quad q_1^{F_1}, q_2^{F_1}, ..., q_{|F_1|}^{F_1}, \quad q_1^{F_2}, q_2^{F_2}, ..., q_{|F_2|}^{F_2}$

   Re-index via $Q = q_1, q_2, ..., q_{|I|}$. Note that $\bigcup_{k=1}^{|I|} q_k = I$.

6. Create the set S of virtual slots according to user-specified AAR's.

   For instance, if AAR$=6$ $flights\ per\ hour,\ t_0 = 1801,\ t_1 = 2159,$

   then $S = \{1810, 1820, ..., 2150, 2200\}$. When AAR$> 60$, use suffixes to mark

   subdivisions of minute intervals, e.g., $1801A$, $1801B, ...$

7. Assign a CTA and $OCTA$ (if necessary) to each $f \in I$

   Let Q be as in Step 5(d) and let S be as in Step 6.

```
for $f = q_1, q_2, ..., q_{|I|}$,
    $Let\ t' := min\left(t \in S : t \geq earliest\_CTA_f\ and\ status\left(t\right) = open\right)$
    $CTA_f := t'$
    $CTD_f := CTA_f - OETA_f + OETD_f\ ,\ if \in F_1 \cup F_2$
    $EDCT_f := \begin{cases} OETD + 1,\ if\ f \in E \\ CTD_f\ ,\ if\ \ f \in F_1 \cup F_2 \end{cases}$
    $OCTA_f := CTA_f\ ,\ OCTD_f := CTD_f\ ,\ if\ \ f \in F_2$
    $ETA_f := CTA_f\ ,\ if\ f \in F_1 \cup F_2$
    $flight\left(t'\right) := f$
    $owner\left(t'\right) := AIR_f$
    $status\left(t'\right) := \begin{cases} open,\ if\ f \in E\ and\ CTA_f < ETA_f \\ filled,\ else \end{cases}$
    $S := S - \{t'\}$
end for
```

end RBS algorithm.

The RBS algorithm is surprisingly complex considering that it is intrinsically a fist-come, first-served algorithm. Steps 1 - 4 establish the status of each flight (included in the program or not, exempt or not) and the earliest time of arrival that can be assigned to it. It is arguable that steps 1-4 are preliminary to the entire GDP, and are not really peculiar to RBS.

Note that, in step 1, there are three conditions each of which is sufficient for a flight f to be included in the program: (1) that ETA_f falls within the program horizon; (2) that ETA_f is beyond the program horizon but CTA_f falls within the program horizon; (3) that ETA_f is beyond the program horizon but the original scheduled arrival time of f falls within the program horizon. Condition (1) simply says that the flight will require an arrival resource even if it is not subject to ground holding. Condition (2) allows for the fact that RBS may have been run before and that previously issued controlled times of arrival should be honored. Condition (3) handles the case in which a flight has effectively reserved an arrival slot within the program by virtue of its scheduled arrival time but has been delayed beyond the scheduled end of the program. This removes the incentive for an airline to withhold delayed flight information.

In step 2, the set of flights that are declared exempt from ground delay are specified by the user. Flights are exempted for any of a number of reasons: special permission from the traffic flow manager running the program, they have been de-iced at their departure airport and will get iced over again if delayed, they are far enough away (in time) that there is not sufficient confidence in arrival capacity forecasts to warrant their delay, and so on.

Step 5, the heart of the algorithm, rations slots by the following priority scheme: exempt flights first, flights already assigned a CTA second, and, lastly, non-exempt flights that require a CTA. Step 6 is a preparatory division of the time horizon into arrival slots; the division of one hour into more than 60 arrival slots is accomplished through suffixing one-minute arrival slots (e.g., 1810A, 1810B). Step 7 is a labeling scheme necessary for subsequent executions of the RBS algorithm.

Although RBS was originally intended to be run just once at the start of a GDP, a combined routine has been developed, called RBS++, in which RBS is run immediately prior to Compression.

5.2. The Compression Algorithm

The Compression Algorithm is a dynamic tool designed to move flights up in the arrival hierarchy during a GDP in order to fill slots vacated by canceled flights. This makes efficient use of airport arrival resources. Most of the direct airline cost savings forecasted for CDM prototype operations are a result of delay reductions from compression (see [5] for a benefits analysis of CDM prototype operations). Section 2.3 of this paper contains a working example of the core compression idea. Also, see [31] or [16] for more examples of compression.

```
main Algorithm (Compression)

1. for each slot $t = t_1, t_2, ..., t_{|S|}$
        if status($t_k$) = $filled$ or $hold$
          go to $t_{k+1}$
        else      [$status = open$ or $released$]
          call subroutine $fill\_slot\,(t)$
  end main algorithm


subroutine $fill\_slot\,(t)$      [try to fill slot $t$ with a flight]

1. if $status\,(t) = open$ then

     (a) form one queue $Q$ of all flights ordered by CTA
     (b) call subroutine $search\,(Q,\, t)$

1. if $status\,(t) = released$, then

2. let $A := owner\,(t)$

     (a) form two queues
           $Q^A$ := flights from airline $A$ ordered by CTA
           $Q^{\overline{A}}$ := flights not in airline $A$ ordered by CTA
     (b) call subroutine $search\left(Q^A, t\right)$
     (c) if $search\left(Q^A, t\right)$ returns $no\_flight\_found$, then
           call $search\left(Q^{\overline{A}}, t\right)$
     (d) if $search\left(Q^A, t\right)$ returns $hold\_slot$, then end subroutine
end subroutine $fill\_slot\,()$

subroutine $search\,(Q,\, t)$      [find a flight in $Q$ to assign to $t$]

1. input queue $Q$ and time slot $t$

2. if $Q = \phi$ then return $hold\_slot$ and end subroutine
```

```
3. repeat
       (a) $f$ := next flight in $Q$
       (b) if $feasible(f, t) = true$ then
           $old\_owner := owner(t)$        [save info on $t$]
           $old\_status := status(t)$
           $t' = \mathrm{CTA}_f$                    [next slot to consider]
           $\mathrm{CTA}_f := t$                    [assign $f$ to $t$]
           $status(t) := filled$
           $owner(t) := Air_f$
           (c) if $old\_status = open$ then
               $status(t') := open$
           (d) if $old\_status = released$ then
               $status(t') := released$
               $owner(t') := old\_owner$
               $fill\_slot(t')$
       (e) $Q := Q - \{f\}$
   until $status\,(t) = filled$ or $Q = \phi$
4. if $status\,(t) = filled$ then return $flight\_found$ else return $no\_flight\_found$
   end subroutine $search\,(\ ,\ )$

   subroutine $feasible\,(f, t)$
5. input flight $f$, time slot $t$
6. check that $f$ meets feasibility criteria (user specified) for assigning $f$ to $t$
7. if $f$ feasible for $t$ then return $true$ else return $false$
   end subroutine $feasible\,(\ ,\ )$
```

The main algorithm of compression scrolls through the time slots in ascending order and attempts to fill available slots by calling the subroutine $fill\_slot(\,)$. This subroutine partitions slots according to their status, *open* or *released*. If a slot t is *released*, then it has a controlling airline, A, who would like to see one of its own flights assigned to t. In step 1(b), the subroutine *search* (,) attempts to fill t with the next available flight from A. If this is not possible, then in step 1(c), the subroutine *search* (,) attempts to fill t with the next available flight not from airline t. If a slot is *open*, then it has no controlling airline and is generally available. In step 2, *search* ( ) is called on the entire collection of flights. Note that no matter what the status of a slot is, flights are considered in order of increasing CTA - see steps 1(a) and 2(a).

It is quite possible for the subroutine *search* ( ) to fail to assign a flight to a time slot from the input queue. More globally, the entire algorithm may leave a slot with no flight assigned to it. This is a desirable feature, since a slot may be too early for any flight in the program.

One of the key aspects of this algorithm is that the status of these unfilled slots is unaltered. For instance, suppose that an *open* slot t cannot be filled with a flight from the controlling airline, A, because it has no flights below slot t. Then the subroutine *search* ( ) returns a *hold_slot* status in step 2, indicating that the input queue of flights from A is empty, and by step 1(d) of $fill\_slot()$, the attempt to fill slot t is halted. If the ground delay program is extended, A might suddenly have a flight that can be moved into t. By retaining the *open* (but not *released*) status of t, A can make future use of the resource.

The upward movement of flights from a single airline (substitution stream) is ignited by step 3(c). Suppose that a slot t has the *open* status and owner A. If a flight f is moved into t from a (lower) slot t', then ownership of slot t is transferred to t', regardless of the airline that owns f, and the subroutine $fill\_slot()$ is recursively called on t'. A flight from A is then sought for t'. If f belongs to some airline B $\neq$ A, then f is said to have acted as a *bridge* for the transfer of ownership of slot t to t'. More generally, the term *bridging* refers to the downward transfer of ownership of a slot over flights from competing airlines. Note that if airline B creates a bridge for airline A, then both A and B have profited: a flight from B is moved earlier in time and the slot resources of airline A are moved later in time preserved. moved to a place where that they can make use of it.

In the interest of NAS user compliance, a more refined version of Compression is used in which flights are designated as Class I or Class II. Class II flights are those created within 48 hours of the current time. If a Class II flight is cancelled from a slot t, then the status of t is declared to be *open*, so that the airline loses control of the slot. This is intended to discourage the airlines from packing the schedule with dummy flights just prior to the program.

The feasibility criteria in the subroutine $feasible()$ has deliberately been left vague. There are several possibilities for this. The most direct is that a flight f can be moved into a slot t only if $ETA_f \geq t$. More generally, the requirement is that $ETA_f \geq t+k$, where $k \geq 0$ is a fixed parameter to prevent minor movements of flights. Also, feasibility of assignment should vary with the current time. The airlines require a minimum notification if a flight is to be moved earlier than its current CTA. This minimum requirement is currently set at 30 minutes for all airlines.

We emphasize that there are many potential variations on compression; we have presented the one that has been in use at the inception of prototype operations of CDM. The details of compression are still being worked out through experimentation. For instance, the need for a 20-minute window is under question. Such a time window would allow an airline to substitute a flight into a time slot that is no more than 20 minutes earlier than its earliest possible time of arrival.

6. Research and Extensions of CDM

6.1. Past Research on Ground Holding Problem

The problem of assigning ground delay to flights bound for a single airport can be mathematically modeled as a transportation problem known as the *ground holding problem* (GH). The model requires the following assumptions.

Assumption 1: (Discrete time horizon) There is a fixed time horizon which has been discretized into T equally-sized contiguous time periods, $t = 1, 2, ..., T$.

Assumption 2: (Deterministic demand) The number of incoming flights is known in advance; for each flight f, there is a scheduled time (period) of arrival, denoted a_f This is the earliest arrival time that can be assigned to the flight.

Assumption 3: (Deterministic capacity) For each time period, t, let b_t be the arrival acceptance rate (AAR) of the airport, meaning the maximum number of flights that can be accepted by the airport during that time interval. Then we assume that b_t is known in advance for each time period t. Strictly speaking, this does not hold in practice because the AAR's are dependent upon weather conditions and runway configurations, which are stochastic in nature. However, the specialist who formulates the GDP fixes these numbers in accordance with the current best estimate, so, for purposes of this formulation, we will assume that they are deterministic and known in advance.

Let F be the set of incoming flights that require arrival slots. We define for each f and each t, a binary variable, $X_{f\,t}$, such that

$$X_{f\,t} = \begin{cases} 1, \text{if flight } f \text{ is assigned to time interval } t \\ 0, \text{otherwise.} \end{cases}$$

We have the following integer program.

(GH)

$$Minimize \sum_{f \in F} \sum_{t=1}^{T} X_{f\,t} C_{f\,t}$$

subject to

$$\sum_{t=a_f}^{T} X_{f\,t} = 1$$

$$\sum_{f} X_{f\,t} \leq b_t$$

$$0 \leq X_{f\,t} \leq 1$$

$$X_{f\,t} \in \{0, 1\}$$

Let $C_{f\,t}$ be the cost associated with assigning flight f to time interval t. A commonly accepted form is $C_{f\,t} = (t - a_f)^{\sigma}$. The parameter σ yields super-linear growth in the tardiness of a flight as t increases. This favors the assignment of a moderate amount of delay to each of two flights rather than the assignment of a small amount of delay to one and a large amount to the other.

Since GH is a transportation problem, large instances can be efficiently solved. GH was first systematically described by Odoni in [20]. Andreatta and Romanin-Jacur [1] treated the stochastic version of GH for an airport with constrained arrival capacity in (at most) one time period. In [26], Terrab and Odoni developed a dynamic programming formulation for the stochastic version of GH as well as heuristics to handle the larger cases. [**Amedeo: what should we say about ref [2] ?**] Using stochastic linear programming with recourse, Richetta and Odoni expanded this work to include the dynamic case, in which ground holdings are updated as time progresses (see [22]). Although their dynamic solution yielded considerable savings over the static solution, the speed of solution proved to be too slow for realistic cases. See [4].

Ideally, the ground holding problem should be solved on a network-wide level, taking into account the connectivity of flights. Flights can be connected in one of three ways: by passenger, crew or aircraft. In the former sense, passengers are scheduled to travel from airport A to C by taking a flight from A to B, then B to C. The arrival of the first flight should coincide (roughly) with the departure of the second flight. In the latter two senses, a single crew or aircraft may be scheduled to traverse many flight legs, e.g., from city A to city B to city C, and so on. The delay of even a single flight can propagate throughout the entire system.

Both GH and air traffic flow management in general have been treated on a network-wide level (taking multiple ‘airports and flight connectivity into account) in Attwool [3], Sokkapia [25], Andreatta and Romanin-Jacur [1], Wang [32] and by Vranas, et. al., in [29] and [30], and, more recently, by Bertsimas and Stock [6]. These models are generally difficult to solve (NP-hard) integer programs, although much progress has been made toward solving these models in real time.

So far, none of the optimization methods designed to solve the FAA's problems has been accepted by the air transportation community. There are several problems associated with the practical application of these models.

First, is that the models assume the existence of a centralized data source providing an accurate, up-to-the-minute status on all flights and connectivity at every airport and in each sector of airspace over the United States. No such source exists at this time; moreover, it would be an enormous undertaking to set up the required real-time data feeds and to solicit the necessary cooperation of the appropriate businesses, local governments, and private concerns.

Second, is that the models tend to adopt the economically neutral viewpoint of a central controller (presumably, the FAA) whose sole interest is the optimization of overall system efficiency. To the contrary, the FAA serves the needs of a number of competing airlines, each of whom would like to have a performance edge over their competitors or, perhaps, to improve the efficiency of its own operation at the expense of overall efficiency. Worse yet, many of the optimization-based models assume that the FAA has precise control over each and every flight in the NAS. This is out of step with the latest trend in air traffic management to relinquish to the airlines some of the control possessed by the FAA. In fact, these are the issues that motivated CDM.

In addition to the technical difficulties involved in gathering the data from the disparate sources, there are fundamental game-theoretic issues that must be addressed to ensure the accuracy of the data. It may be to a party's advantage to submit skewed or inaccurate data. Indeed, this is a problem with the present (non-CDM) system used by the FAA to schedule Ground Delay Programs (GDP's), as is thoroughly discussed in [31]. The research topics discussed in this section define problems that fit within this context.

6.2. Current and Future Research Stemming From CDM

CDM methodologies have had a major impact on air traffic flow management, particularly in the context of ground holding strategies. Given the current momentum of CDM and its growing support, it seems inevitable that CDM will continue to impact air traffic management. It is appropriate that the current body of research in air traffic flow management should reflect, if not be tailored to, the CDM philosophy and methodologies. In this section, we examine the research issues that have arisen as a result of CDM and those that we foresee.

(1) The Stochastic Ground Holding Problem and Its Extensions

The primary purpose of a ground delay programs is to slow down rate of arrivals into an airport suffering from degraded arrival capacity. Flights are issued ground delays at the respective points of origin. Since most flights in the United States are scheduled for an en route time in the range of 1-6 hours, a ground delay program must be planned several hours in advance of anticipated degraded arrival capacity. Otherwise, the ability to take control of the arrival rate is lost as flights become airborne.

One must be careful to distinguish between the airport acceptance rate (AAR) that will actually occur and the planned arrival acceptance rate (PAAR) that is set by the specialist at the ATCSCC. The former is the number of flights that will actually be able to land at the airport while the latter is the number of flights that will attempt to land, based on controlled times of arrival assigned during a GDP.

In a perfectly planned and executed ground delay program, AAR = PAAR for each hour of the time horizon. Due to the unpredictably of acceptance rates, this is rarely the case. If AAR $>$ PAAR in any given hour, then airport capacity is greater than expected and some flights will have suffered needless delays, and the airport arrival resources will be wasted. On the other hand, if AAR $<$ PAAR in any given hour, then the arrival capacity is less than expected and flights will incur airborne delay that could have been absorbed on the ground, had a more aggressive GDP been enforced. One can see that a balance must be achieved between airborne holding and ground holding that maintains throughput at the airport.

AARs are dependent upon airport configurations which are, in turn, dependent upon meteorological conditions such as visibility, wind velocity/direction, and precipitation. Thus, AARs are stochastic in nature, rather than deterministic.

In [17], an integer programming model is proposed for determining a set of PAARs that minimizes the sum of the cost of ground delay and the expected cost of airborne delay. The approach they have used assumes a discrete number of scenarios, or profiles, of arrival capacity, exactly one of which will come to pass. Each scenario consists of a sequence of T arrival capacities $A_1, A_2, ..., A_T$, where each t corresponds to a time interval in the horizon. Thus, A_t is the arrival capacity that will occur in the t^{th} time period if that scenario should be realized. Figure 6 displays multiple AAR scenario forecasts for a fictitious airport whose normal AAR is 70 flights per hour. Their use of discrete scenarios is similar to the technique applied by Richetta and Odoni in [23] and Chapter 3 of [24]. There are many approaches to stochastic programming. See [7], [18], or [21] for background.

The model produces excellent computational results. The network structure of the problem allows integer solutions to be obtained directly from the linear programming relaxation. For several large-scale, realistic test cases, optimal results were obtained in fractions of a second with modest computing power .

(2) Weather Forecasting Techniques

The integer programming models of the stochastic ground holding problem such as those proposed in [17] and [23] assume that weather can be turned into runway configurations and then into airport capacity. This is already being done with acceptable accuracy at airports across the country by experienced airport controllers [**is this the correct terminology?**]. But these scenario-based models also assume the availability of a discrete set of mutually exclusive weather forecasts, each of which has an associated probability. To our knowledge, there is no research that addresses the question of whether or not this type of forecasting is possible.

We foresee two possible approaches to multiple-scenario weather forecasting. One would be to look at the fundamental process of weather prediction and to modify it to produce scenarios and their associated probabilities. The second would be to analyze historical weather patterns and AARs at individual airports and to determine scenarios and probabilities based on these.

Since most forecasts are based on an amalgamation of weather models, it seems reasonable that, indeed, multiple forecasts could be generated. However, the rub is how to assign accurate probabilities to them. It's possible that these numbers could be based on historical data peculiar to that airport. Moreover, one needs to question the assumption that weather scenarios can be treated as mutually exclusive. The intuition for this assumption is based on experience at an airport such as San Francisco where fog is the main problem and a set of scenarios can be generated such that each one reflects a time at which the fog might roll in and a time at which it might roll out. Similarly, scenarios can be generated for an inevitable winter storm at an airport in the eastern United States. In this case, the scenarios would reflect the duration and severity of the storm. However, weather patterns at other airports may not behave this way.

(3) Stochastic Demand Prediction

In addition to uncertainty about airport capacity, GDPs must contend with the randomness associated with the demand at the subject airport. Experience with GDPs during the first six months of CDM prototype operations in 1998 has underscored the fact that, even with the improved information provided by CDM, there are factors that lead to inaccuracies in demand forecasts. These factors include cancellations, pop-up (unforeseen) flights, delays in departure times and variability of en route times. The demand profile that actually materializes during GDPs, more often than not, is significantly different from the profile that the GDPs had anticipated.

There is a clear need for a stochastic model of arrival demand. Since it is difficult to predict the arrival time of an individual flight, it would probably be best if the demand profile were not based on the traditional accumulation of specific flight-by-flight arrival time estimates but rather on a summing of flight-by-flight arrival time probability distributions. In other words, the demand profile could take into account that there is, for example, a 5% chance that a certain flight will be cancelled, a 30% chance it will be 30 minutes late, a 10% chance it will be 60 minutes late, etc. Such a model should be based on statistical analysis of historical data and take into account that fact that probability distributions will vary with airport, type of flight, time of day and weather conditions.

(4) Balancing Arrivals and Departures

At most airports, runways are used for both departing and arriving aircraft. Research attention to date has been focused on arrivals. Presumably, this is because the spacing of departing aircraft is not as sensitive as the spacing of arriving aircraft. Also, having planes queued up on the runways waiting for take off is not as much of a safety issue as having planes queued up in the air waiting for arrival. But when an airport is impacted by adverse weather conditions or special runway operations, departures are afflicted as well as arrivals. A balance must be achieved in the allocation of airport resources toward both arrivals and departures.

The general reserach question to consider is the development of models that aid in the control of the balance between arrivals and departures. Ideally, such models should provide airlines with the ability to control arrival-departure trades relative to their own flights.

(5) Determining the Optimum Departure time for Aircraft

Currently, controlled departure times for aircraft subject to a ground delay program are back-computed from their controlled arrival times, using ETMS-modeled en route times. For instance, if a flight with a target arrival time of 0900 were predicted to have a three-hour en route time, then it would be given a departure time of 1200. Since en route time is a function of aircraft speed, altitude and the path flown, many of the en route alternatives for a given flight are predetermined implicitly by the CTD-assignment process. However, the selection of en route alternatives is more of an economic issue than an air traffic control issue. In keeping with CDM philosophy that purely economic issues be settled by the airlines, the RTCA Select Committee has prescribed [**do we have a ref for this?**] that, in the future flow management environment, flights subject to a GDP will be assigned only a controlled time of arrival (CTA). It will then be up to the airlines to determine the times at which their respective flights should leave their airports of origin in order to meet their assigned CTAs at the airport of destination (the GDP airport).

This raises the question of how a given airline should determine the optimal departure time in the absence of CTDs. This time must be determined in the face of two principal sources of uncertainty: (a) the uncertainty about how much delay the flight will experience in departing from the airport of origin due to taxi-out time to the departure runway and queueing time for take-off, and (b) the uncertainty about how much delay the flight will experience after arriving at the terminal airspace of the airport of destination due to congestion. There are other sources of uncertainty, such as en route travel times, but they are typically dominated by the uncertainty about the delays incurred at take-off and landing.

Some research was done at MIT Lincoln Laboratories in 1993 [**Bill: do we have a ref?**] that dealt only with uncertainty about delay on arrival. This work should be extended to a model that the airlines could use to determine optimal departure times. This research would not only be helpful to the airlines but could also lead to some policy conclusions regarding the RTCA proposal to conduct flow management by assigning CTAs but not CTDs.

(6) The Integration of CDM with Existing and Developing Systems (CTAS, SMA)

There is currently a lot of activity in research and development of terminal area (TMA) decision support systems for air traffic management. NASA is playing a leadership role in this respect, through both in-house work and funding of related work elsewhere. The principal product to date has been CTAS, which was prototyped several years ago, but is still undergoing significant evolution at the Dallas-Fort Worth test site. Other products include the Surface Management Automation (SMA) program in Atlanta, various tools for Departure Planning (CTAS is currently limited to arrival air traffic management) and a host of more specialized tools under the Terminal Automation Program (TAP).

CDM could influence and, in turn, be influenced in fundamental ways by the development of such an integrated TMA decision-support system(s). For example, and at the most obvious level, some of the data available to, or through, CDM, such as timely information about cancellations, flight substitutions, compression, projected flight departure times, etc., are of vital importance to a TMA decision-support system. As another example, information that could be made available in the future through CDM could be very valuable to an integrated TMA decision-support system in scheduling and sequencing aircraft, arrivals and departures at a busy airport. This would be the case, for instance, if real-time information were made available by the airlines on flights within a tightly connected "bank" at a hub. An "enhanced CTAS" could then schedule arrivals and departures and assign them to runways in a way that preserves these flight connectivities as much as possible. Conversely, CDM and flow management, in general, would greatly benefit from more accurate estimates about projected arrival and departure airport capacities, taxi-out times and landing and take-off

(7) Formal Models of Fair Allocation and Competition Among Airlines

The key algorithmic components for the current CDM GDP procedures are ration-by-schedule and compression. These procedures carry out a resource allocation that is based on certain implicit fairness criteria. A second key ingredient that determines GDP formulation is the manner in which the airlines provide their inputs through the cancellation and substitution process.

These procedures have been developed and refined in a truly collaborative setting through discussions and mock testing and real-world implementation. Provided that future revisions to these algorithms are minor, these heuristics should serve CDM well. But it is hard to know in advance what affects major, or even minor, revisions will have. Formal models should be developed that capture each of these two components. The purpose of such a model would be to

- provide a more formal presentation of the current allocation and reallocation schemes;
- provide flexibility in incorporating new rules or concepts of fairness;
- provide a model whose behavior can be better tested.

An optimization-based approach was developed by Butler in [10] to implement the ration-by-schedule and compression functionality.

(8) Gaming Issues

Much of the collaborative process in CDM has gone toward devising procedures that are deemed fair by the competing parties. However, there may be ways in which the current procedures established by CDM for the implementation and maintenance of a ground delay program allow for gaming by the participating airlines. For instance, under the ration-by-schedule algorithm, an airline has effectively reserved an arrival slot when it has scheduled a flight in the OAG (official airline guide). One possibility that should be guarded against is that an airline creates strategically scheduled dummy flights in order to gain an advantage in the substitution-cancellation process that occurs after the initial rationing of slots. This potential problem has been alleviated by the institution of a policy that disallows an airline from substituting a flight into a slot vacated by a flight that has been created within 48-hours of the GDP (such a flight is considered a class-2 flight, all others are class-1). But there could be more subtle ways of gaming the system that are, as yet, not revealed.

Another example of potential gaming is the bridge-only status that some airlines have received at their hubbing airport. Under this agreement, the compression algorithm is prevented from filling the vacant slots of a bridge-only airline. This is requested so that the bridge-only airline can postpone substitutions into its vacated slots until such a time that more information is available and the in-house substitution software can act through its own logic.

(This status is reserved for airlines that, presumably, have enough flights to eventually fill all of their slots.) See section 5 for a discussion of brdiging.

In theory, a bridge-only airline could allow compression to fill its slots then reallocate the slots to its flights at a later time through the substitution-cancellation process. However, this is (currently) a cumbersome process that involves many pair-wise swaps. For now, bridge-only carriers will be monitored to see to it that their slots are eventually used. Nonetheless, this raises the question of whether or not the bridge-only status is a hinderance to the overall goals of compression or if it could be used to unfairly game the situation.

(9) Collaborative Routing

When a sector becomes overloaded, the controllers in that sector usually issue miles-in-trail restrictions on the aircraft that have yet to enter the sector. This spaces the aircraft out to maintain safety. However, the miles-in-trail restrictions force the aircraft that have not yet entered the sector to sustain delay in the queue of aircraft waiting to enter. Anecdotal evidence indicates that, in hindsight, the aircraft waiting to enter the sector could in many cases have reduced their overall delay had they chosen a less congested route [11]. Unfortunately, in the current system, there is little information to help an airline decide which routes are likely to be congested a priori.

The CDM community is currently considering benefits that could be produced by applying the underlying CDM procedure, CDUP, to the flight routing problem. The envisioned system would use data from the Enhanced Traffic Monitoring System (ETMS), a tool developed by the Volpe National Transportation Systems Center that predicts aircraft trajectories, to predict when and where the congestion problems will occur. These predictions of sector loading would be distributed to the airlines over the AOCnet or over the world wide web. The airlines then have the option to alter their flight plans in response to the predicted congestion. This results in an update to the predictions, at which point, the airlines have the option to again alter their flight plans. This completes the CDUP loop. Thus far, the efforts have been geared toward the identification of useful and enabling technology, mainly "white-boarding", and teleconferencing. There is a need to identify areas for the application of CDM technology to aircraft routing and to develop operational concepts related to the identified areas and to explore how previous work on large-scale routing models can be adapted to the CDM setting.

Some of the large body of research that addresses large-scale routing problems was developed specifically for the air traffic management context and some was developed for use in other areas. This work addresses the difficult task of finding a set of routes that collectively meet system-wide objectives from the combinatorially explosive number of possibilities that can arise in large problem settings. To our knowledge, no work to date has addressed the special environment that arises in the context of CDM. We believe that there is a need within CDM for such large scale route identification schemes and that effective procedures can be derived by modifying work developed for other (non-CDM) contexts.

7. Closing Remarks

The tremendous growth in air traffic in the last quarter of this century has induced a vast infrastructure for serving the needs of NAS users, mainly by exercising centralized control. It is ironic that this same growth has incited a movement away from centralized control toward a more collaborative setting. The infrastructure is now being recast in a user-service provider paradigm in which air carriers gain more control over their operations and the government is more responsive to their needs. Much of this has been enabled by the voluntary exchange of timely data and the dissemination of situational displays.

Just as the air traffic community struggles with the collaborative paradigm, so must researchers struggle with the implications of it. Many of the guiding maxims of air traffic control are in jeopardy. CDM has presented the research community with a host of new problems revolving around the technological and logical integration of its ideas.

8. Acronyms and terms:

AAL - American Airlines

AAR - Airport Acceptance Rate

AOCnet - The private intranet that links the airline operation centers, the FAA and other organizations

ATMS - Advanced Traffic Management System: an air traffic database maintained by CDM; the counterpart to the more established database, ETMS.

CDM - Collaborative Decision Making: a project aimed at the improvement of air traffic flow in the United States by combining the expertise and resources of both government and industry to improve information exchange, create situational awareness, and disseminate decision-making to appropriate parties.

CDUP - Constraint-Dissemination User-Preference: a cycle of feedback between users of the NAS and the service provider in which constraints are issued and user preferences are expressed.

Compression (algorithm) - An algorithm for moving flights

CTA - Controlled Time of Arrival: issued to a flight during a GDP as a means of delaying arrival

CTD - Controlled Time of Departure: issued to a flight during a GDP as a means of delaying arrival

ETA - Estimated Time of Arrival

ETD - Estimated Time of Departure

ETMS - Enhanced Traffic Management System: the operational database that supplies air traffic information to the ATCSCC.

FAA - Federal Aviation Administration.

FADE - FAA/Airline Data Exchange: the forerunner of CDM; a government-industry project aimed at determining the impact of improved data exchange on decision-making in air traffic management.
FSM - Flight Schedule Monitor: a decision support/situation display tool developed by CDM.
GH - Ground Holding (Problem)
GDP - Ground Delay Program: an extended air traffic initiative by the ATCSCC to regulate air traffic flow into an airport with degraded arrival capacity.
Grover-Jack - An algorithm used by the ATCSCC to allocate arrival resources during a GDP.
LP - Linear Programming.
NAS - National Airspace System the FAA components (traffic control centers, radar, etc.) that make up the national air traffic network.
NEXTOR - National Center of Excellence for Aviation Operations Research
OAG - The Official Airline Guide: A publication of scheduled airline activities (arrival, departure, etc.).
Prototype Operations - The implementation of CDM-GDP enhancements on a trail basis at several airports in the United States.
Service Provider (to the NAS) - another name for the FAA, used especially particularly in the context of its responsibilities toward users of the NAS.
User (of the NAS) - any person or body that operates aircraft in the NAS
RBS - Ration-by-Schedule: an algorithm developed by CDM for the allocation of arrival slots during a GDP according to OAG-scheduled arrival times. This is the CDM counterpart to Grover-Jack.

REFERENCES

[1] Andreatta, G., and G. Romanin-Jacur. 1987, "Aircraft Flow Management Under Congestion," *Transportation Science*, **21**, 249-253.

[2] Andreatta, G., A. Odoni, and O. Richetta, O. 1993, "Models for the Ground-Holding Problem," in *Large-Scale Computation and Information Processing in Air Traffic Control*, L. Bianco

and A.R. Odoni, eds., Springer-Verlag, Berlin, 125-168.

[3] Attwool, V.W. 1977, "Some Mathematical Aspects of Air Traffic Systems," J. Inst. Navigation., **30**, 394-411.

[4] Ball, M., R. Dahl, L. Stone, and T. Thompson. 1993, "OPTIFLOW Build-I Design Document, Version 1.5," *Optiflow Project Report to FAA*, December. [**Mike Ball: should this be referenced?**]

[5] Ball, M. and R. Hoffman. 1998, "Collaborative Decision Making in Air Traffic Management: A Preliminary Assessment;" NEXTOR Project Report to FAA, August.

[6] Bertsimas, D. and S. Stock. 1996, "Air Traffic Flow Management Problem with Enroute Capacities," to appear in *Operations Research.*

[7] Birge, J. and F. Louveaux. 1997, *Introduction to Stochastic Programming*, Springer-Verlag, New York, New York.

[8] Boeing Corporation. 1996, "Current Market Outlook"[**Bill: is this a tech report, or what?**]

[9] Boswek, S.B. and J.E. Evans. 1997, "Analysis of Downstream Impacts of Air Traffic Delay," Project report, MIT Lincoln Lab, March.[**Bill: did we cite this in the paper? If so, where?**]

[10] Butler, T. 1998, "Optimization Model with Fairness Objective for Air Traffic Management," Masters thesis, University of Maryland.

[11] CDM Working Group. 1997, *CDM Meeting Minutes,* July.

[12] FAA. [**Bill:year?**], "FAA Aviation Forecasts: 1996-2007", Technical Report.

[13] Fudenberg, Drew and Tirole [**Bill: first initials?**] 1991, *Game Theory,* MIT Press.

[14] Glover, F., D. Klingman, and N. Phillips. 1992, *Network Models in Optimization and Their Applications in Practice*, John Wiley and Sons, Inc., New York, New York.

[15] Hoffman, R. and M. Ball. 1997, "A Comparison of Formulations for the Single-Airport Ground Holding Problem with Banking Constraints," Submitted for publication.

[16] Hoffman, R. 1997, "Integer Programming Models For Ground-Holding in Air Traffic Flow Management," Ph.D. dissertation, University of Maryland.

[17] Hoffman, R., R. Rifkin, M. Ball and A. Odoni. 1998, "The Static Stochastic Ground Holding Problem," submitted for publication.

[18] Infanger, G. 1994, *Planning Under Uncertainty: Solving Large-Scale Stochastic Linear Programs.* Boyd and Fraser publising company, Danvers, Massachusetts.

[19] Kostiuk, P., D. Lee, D. Long and G. Shapiro. 1997, "Aviation Studies Program," Presentation given at MIT.

[20] Odoni, A. 1987, "The Flow Management Problem in Air Traffic Control," In *Flow Control of Congested Networks*, A. Odoni, L. Bianco and G. Szego, eds., 269-288, Springer-Verlag, Berlin.

[21] Powell, W., P. Jaillet and A. Odoni. 1995, "Stochastic and Dynamic Networks and Routing," *Handbook in Operations Research and Management Science*, **8**, M. O. Ball, et. al., editors, Elsevier Science, Berlin.

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

[23] Richetta, O. 1995, "Optimal Algorithms and a Remarkably Efficient Heuristic for the Ground-Holding Problem in Air Traffic Control," *Operations Research*, **43**, 758-770.

[24] Richetta, O. 1991, "Ground Holding Strategies for Air Traffic Control under Uncertainty," Ph.D. dissertation, Massachusetts Institute of Technology.

[25] Sokkapia, B. G. 1985, "Arrival Flow Management as a Feedback Control System," The Mitre Corporation, Washington, D.C.

[26] Terrab, M. and A. Odoni. 1993, "Strategic Flow Management for Air Traffic Control," *Operations Research*, **41**, 138-152.

[27] Vanchieri, Anthony, R. 1997, "Aviation Dream Team," *ORMS Today*, April, 39-41.[**This is a good ref on COE - can we fit it in somehow?**]

[28] Venkatakrishnan, C., A. Barnett, and A. Odoni. 1993, "Landings at Logan Airport: Describing and Increasing Airport Capacity," *Transportation Science,* **27**, 211 - 227.

[29] Vranas, P., D. Bertsimas and A. Odoni. 1994, "The Multi-Airport Ground-Holding Problem in Air Traffic Control," *Operations Research*, **42**, 249-261.

[30] Vranas, P., D. Bertsimas and A. Odoni. 1994, "Dynamic Ground-Holding Policies for a Network of Airports," *Transportation Science,* **28**, 275-291.

[31] Wambsganss, M. 1996, "Collaborative Decision Making Through Dynamic Information Transfer," *Air Traffic Control Quarterly*, **4**, 107-123.

[32] Wang, H. 1991, "A Dynamic Programming Framework for the Global Flow Control Problem in Air Traffic Management," Transportation Science,} **25**, 308-313.