Auctions for the Safe, Efficient and Equitable Allocation of Airspace System Resources

Michael Ball
Robert H. Smith School of Business & Institute for Systems Research
University of Maryland
College Park, MD 20742

George L. Donohue and Karla Hoffman
Systems Engineering and Operations Research Department
George Mason University
Fairfax, VA 22030

ABSTRACT:

In this paper, we present evidence of the need for the specification of property rights associated with airport departure and arrival time slots and then provide the rationale and means for using auctions to allocate such property rights. We provide a historical look at slot allocation in the U.S. and the consequences of prior and current practices on both safety and competition. It is shown that the particular characteristics of aviation slot allocation suggest the need for three types of market mechanisms: an auction of long-term leases of arrival and/or departure slots, a market that supports inter-airline exchange of long-term leases and a near-real-time market that allows for the exchange of slots on a particular day of operation. We describe how certain concepts and tools underlying Collaborative Decision Making (CDM) provide a natural basis for developing the near-real-time market and present a broad overview of the special characteristics of aviation slot auctions and appropriate research topics.

1.0 Introduction/Background

Most countries attempt to design their air transportation system so that it is economically viable, safe and efficient. As the system evolves, changes are necessary to assure that these goals continue to be met. Although air transportation in the United States has a comparable safety record to that of automobile travel (on an exposure to risk time basis, (The Royal Society 1992)), the margin of safety is slowly eroding under the demands for more enplanement opportunities. The 1978 deregulation of the US route structure was intended to increase competition within the airline industry and thereby improve efficiency, decrease cost to travelers and expand the overall flying opportunities. This policy initially provided increased enplanement opportunities at reduced prices because there was sufficient capacity in the system to allow such growth. However, the current policies and procedures do not produce a similar effect in a capacitated system. In fact, these policies impede the need to build additional airports and overhaul the technology both within air traffic control and on airplanes. Without such expansion, more system elements are likely to become capacity-limited. In such a system it is essential to use system resources efficiently. We therefore provide suggestions for mechanisms to both expand the capacity and to assure that the current, limited capacity is used both safely and efficiently.

As the U.S. National Air Transportation System (NATS) becomes highly capacity constrained along multiple dimensions, it requires feedback mechanisms that can react along multiple time scales to adjust system behavior (Fan and Odoni 2002). Today, almost a quarter of a century after airline deregulation in 1978, strategic airspace management exercises little or no control over the number of aircraft that are scheduled to land and depart from various airports. It can only react when the system is overloaded. Thus, the U.S. airlines implicitly are responsible for setting constraints on airport operations as part of their scheduling process. The policies of the U.S. Department of Transportation (DoT) and the Federal Aviation Administration (FAA) effectively encourage these airlines to overbook and then cancel or delay flights leaving the system regularly in crisis mode, with re-scheduling the norm rather than the exception. Similarly, regional governments that wish to determine their demographic growth patterns are powerless to shape, or even to suggest, how the airspace in their region is used. The driving forces behind this paper are the questions: what forces led to this situation and what policy changes might be made to improve the U.S. national air transportation system crisis.

We begin by providing a description of the history of the U.S. Aviation System and then proceed, thereafter, to explain how market-clearing mechanisms might be able to rectify many of the shortcomings of the current system (see (Mineta 1997; Commision on the Future of the US Aerospace Industry, 2002) for background.)

.

History of US Aviation:

From 1938 to 1978 the Civil Aviation Board (CAB) managed the nation's air transportation route (and industry) structure (Gleimer 1996). Many, including most economists, felt this administrative process did an inefficient job of providing transportation services (Rochester 1976; Preston 1987). Figure 1 shows how the growth rate of Revenue Passenger Miles (RPM), normalized by GDP, was stagnating just prior to 1978 (data taken from (DOT/BTS 2001)). Prior to 1978, air travel was relatively expensive and considered by many to be only for the upper echelons of society. Immediately after the deregulation act, prices fell and industry productivity and frequency of service increased dramatically. Fig. 2 illustrates how deregulation in both the U.S. and Europe initially increased airline productivity (Alamdari and Morrel 1997), even though FAA productivity did not change (with the exception of the effect caused by the air traffic controllers strike in 1981). However, after 1990, there was a leveling off of airline productivity (Donohue 2002). The lack of incentives to adopt new technology and the political inability to add new airport infrastructure began to limit growth, thereby creating the inevitable rise in queuing delays as the system approached the maximum demand to capacity ratio (Donohue and Shaver 2000).

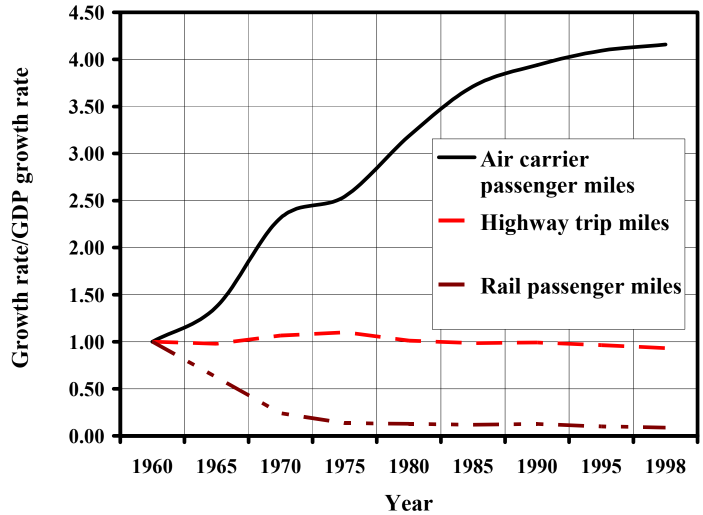

Figure 1. US Transportation Growth for Major Modes normalized to 1960 levels and GDP growth. Source US Department of Transportation Statistics.

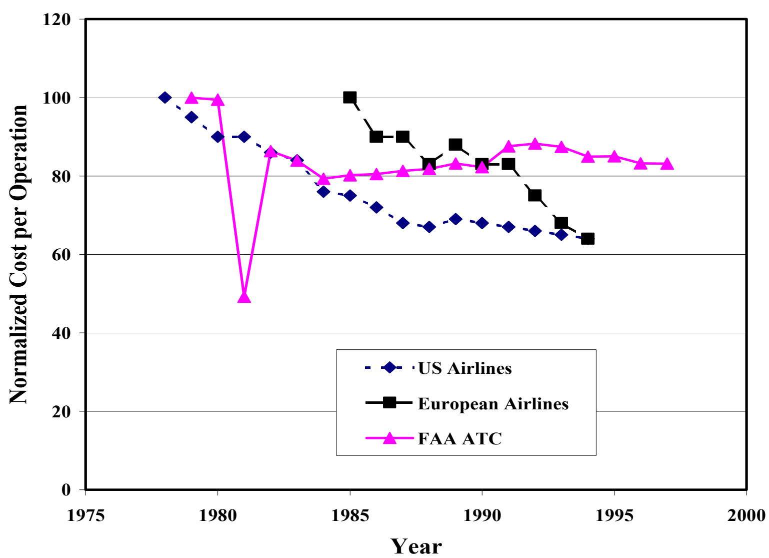

Figure 2. A relative comparison of Productivity Trends for the US Air Traffic Control System and both US and European Airlines Post Deregulation (industry data taken from (Alamdari and Morrell 1997)

Even prior to 1978, however, some airports were already congested. Four airports had been arrival slot controlled since 1968 under the High Density Rule (HDR): New York's Kennedy (JFK) and LaGuardia (LGA), Chicago O'Hare (ORD) and Washington's (Reagan) National (DCA). Today, the air transportation situation looks very different than it did in 1978. Many U.S. airports are becoming scheduled at levels that exceed the FAA's estimate of a maximum safe operational rate (DOT Benchmark report 2001) (Haynie 2002). The major domestic U.S. air carriers use the hub-and-spoke system that brings passengers from smaller cities to hubs that will then transport them in an economical way to their final destination. Hub operations tend to concentrate very large numbers of flight arrivals and departures over short time periods. In some cases, airlines maintain near-monopoly control over hub airports so that newer airlines face significant barriers to entry into these airports.

On April 5, 2000, the semi-deregulation of the slot controls went into effect with the enactment of the AIR-21 bill (Federal Register 2000), which among other things, directed the DoT to totally eliminate slot controls at the four US HDR airports by 2007, and to increase immediately the number of slots allocated for regional service at LGA. This act led to the immediate and extreme congestion of air traffic activity at LGA (Fan and Odoni 2002). Strong "network effects" meant that the LGA delays induced additional delays throughout the NATS.

LGA has been arrival slot controlled (approximately 32 arrivals per runway per hour) since 1968, due to concerns about congestion and community noise at that airport. These slot controls were maintained even after the CAB was abolished in 1978. Figure 3 shows the scheduled number of flights at LGA in 2000. The schedule consists of both arrivals and departures in 15-minute intervals from 7 am in the morning to 10 pm at night. LGA has one arrival runway and one orthogonal crossing departure runway. The FAA officially considers the maximum safe level of operations under favorable weather conditions to be 40 arrivals and 40 departures per hour (i.e. 10 arrivals per 15 minute epoch) under visual conditions. Under reduced visual conditions (Instrument Flight Rules or IFR), this airport is supposed to be reduced to 32 arrivals and 32 departures per hour (i.e. 8 arrivals per 15 minute epoch). Figure 4 shows that the actual operational rates under the more restrictive, and slightly more hazardous IFR conditions frequently exceed the (32,32) rate (DOT Benchmark 2001). It will be shown later that this rate was set by Runway Occupancy Time (ROT) considerations and not aircraft wake vortex separation standards, which are more restrictive. The wake vortex problem was unknown in 1968 when most commercial aircraft were of medium size. With the introduction of both wide-body aircraft (heavy) and small regional jets (RJ's) in a highly dynamic mixture, this safety problem is of growing concern (Haynie, 2002).

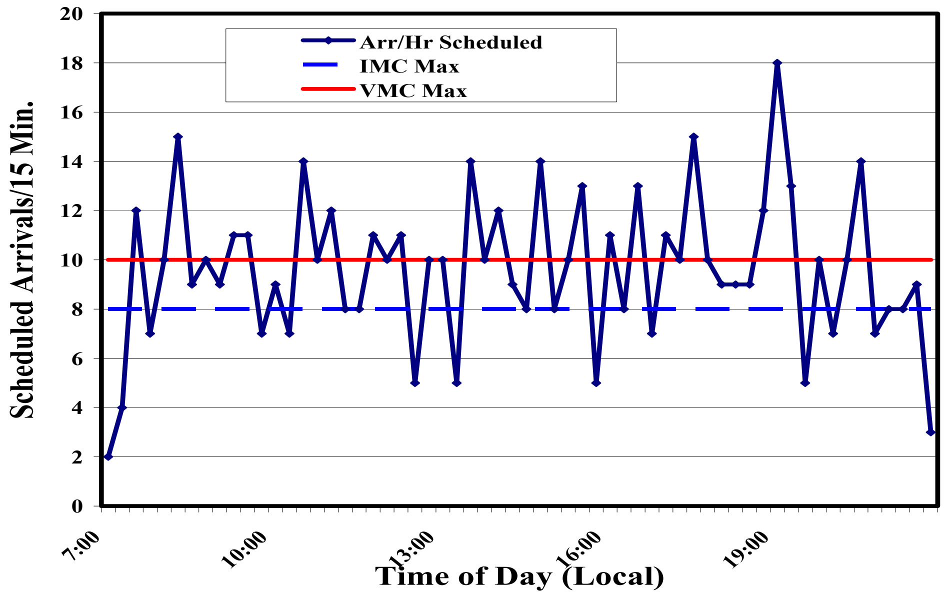

Figure 3. LGA Scheduled Number of Arrivals in 15 minute increments compared to the FAA estimate of safe aircraft operational separation of 8 (32 Arrivals / Hr.) under Instrument Meteorological Conditions (IMC) (FAA 2001)

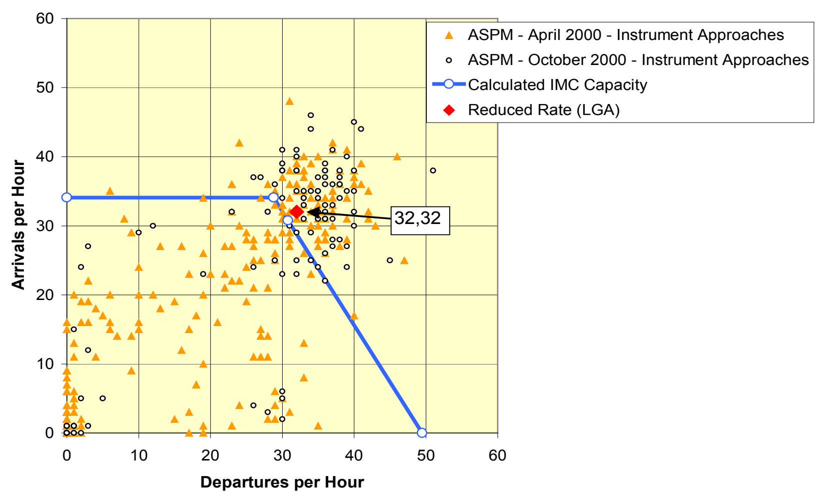

Figure 4. LGA one-hour IMC Arrival-Departure Operation Rates for April and October 2000 (FAA 2001). Notice that Arrivals in excess of 33 Arrivals per Hour imply average aircraft separations of less than 3 n mi or 80 second separation**.**

The fact that there are two different capacity levels: one for good weather conditions and another for inclement conditions, further complicates the process of scheduling. Also, FAA regulated separation rates change depending upon whether a small aircraft follows a large aircraft (in which case the separation must be larger) due to aircraft wake vortex encounter concerns. These alternative landing and takeoff separation rules are quite complex and not considered in the overall determination of how many flights should be allowed to be scheduled during any given time period.

A question naturally arises: why do the airlines schedule operations that exceed the safe departure/arrival rate that an airport can support, thus generating excessive flight delays, cancellations and loss-of-separation violations? The answer is competition. If airline A acts responsibly and does not increase its schedule at a congested airport, it will have voluntarily provided another airline with the opportunity to schedule more flights at that airport. Conversely, if airline B decides to increase its schedule in an attempt to increase city pair options and flight frequency (but in reality only increasing congestion and delay for all airlines), airline A may lose market share. Under policies currently in effect, if scheduled flights at that airport are, at some time in the future, legislatively or procedurally reduced through re-regulation, then the airline with the greatest scheduled flights is likely to argue and receive more of the available flights. Thus, the risk is not only the reduction of market share currently, but also the risk of a permanent reduction in market share. In addition, at an airport not dominated by a single carrier (e.g. LGA), when an airline adds another flight the delay experienced by that flight may be small relative to the total delay to other flights caused by that additional flight. Thus, only a small portion of the "delay cost" of adding the flight may be internalized by the initiating airline.

LGA may be the extreme case, but many US airports and airlines are experiencing similar situations. Table 1 shows the demand to capacity ratio of 20 major US airports. The Demand/Capacity (D/C) ratio is based upon FAA computed good weather capacity calculations and measured average operational rates. When the D/C ratio approaches one, queuing theory predicts that the delay will grow exponentially.

Another way of looking at the congestion issue is to look at a measure of competition at a given airport. The Herfindahl-Hirschman Index (HHI) is a measure of industry competition. According to (Cooper 2000), studies of the hub system within the US have shown that fare revenues are higher on average for trips to and from major hub airports, with a few concentrated hub airports showing significant premiums over a decade; also, the higher premiums are realized by the dominant carrier in that hub, while fares charged by the other carriers in that hub are similar to the fares charged in less concentrated airports. The hub airports that have a dominant carrier typically have a high HHI.

One could argue, on a theoretical basis, that the more competition at an airport, the more tendencies there would be to over-schedule (i.e. monopolists will totally internalize the cost penalty of the delays they cause themselves and therefore not over-schedule). Thus, under current procedures, from the perspective of driving down prices, one would like competition, but from the perspective of keeping congestion at a reasonable level one would like to discourage competition but only if the monopolist would internalize all delay costs and therefore operate at the optimum delay value. Unfortunately, the data in Table 1 does not support this hypothesis.

Table 1. Major Airport Slot Control Status, Maximum Arrival Rate and Aircraft Size Mixture.

City	Airport	Operations per yr X 1000	Delays >15 min per 1000	HHI Index	Demand to Capacity	Slot Controlled
NY LaGuardia	LGA	392	156	<1800	0.68	yes
NY Newark	EWR	457	81	3600	0.66	
Chicago	ORD	909	63	3200	0.72	yes
San Fransisco	SFO	431	57	3500	0.67	
Philladelphia	PHL	484	45	3100	0.67	
NY Kennedy	JFK	359	39	<1800	0.42	yes
Atlanta	ATL	913	31	5500	0.78	
Dallas FW	DFW	866	24	4500	0.52	
Los Angles	LAX	784	22	<1000	0.78	
Washington	IAD	480	20	2700	0.54	
St. Louis	STL	484	18	3800	0.72	
Detroit	DTW	554	18	3600	0.53	
Minneapolis	MSP	522	13	4500	0.68	
Seatle	SEA	445	11	<1800	0.71	
Baltimore	BWI	315	7	1900	0.42	
Charlotte, NC	CLT	460	6	5200	0.52	
Pittsburg	PIT	448	4	5300	0.47	
Denver	DEN	528	2	4200	0.41	

Note in Table 1 that non-slot controlled airports like Atlanta Hartsfield (ATL) have both high market concentration (high HHI) and high delay values. Also, even the current slot controlled airports are experiencing high delays because they are operating near the maximum capacity level. If the current slot controls on these high demand airports expire in 2007, they will undoubtedly become even more congested, as seen at LGA in 2001. This argues for the need for slot allocation measures that both encourage competition and avoid congestion.

Perhaps the most important argument for a new slot control system is the effect of operating overscheduled airports at high demand to maximum capacity ratios. The recent data and analysis in (Haynie 2002) indicate that loss of the regulated safe separation distance of aircraft is positively correlated to the aircraft arrival demand to maximum runway capacity ratio (Figure 5). This will be discussed in more detail in the next section.

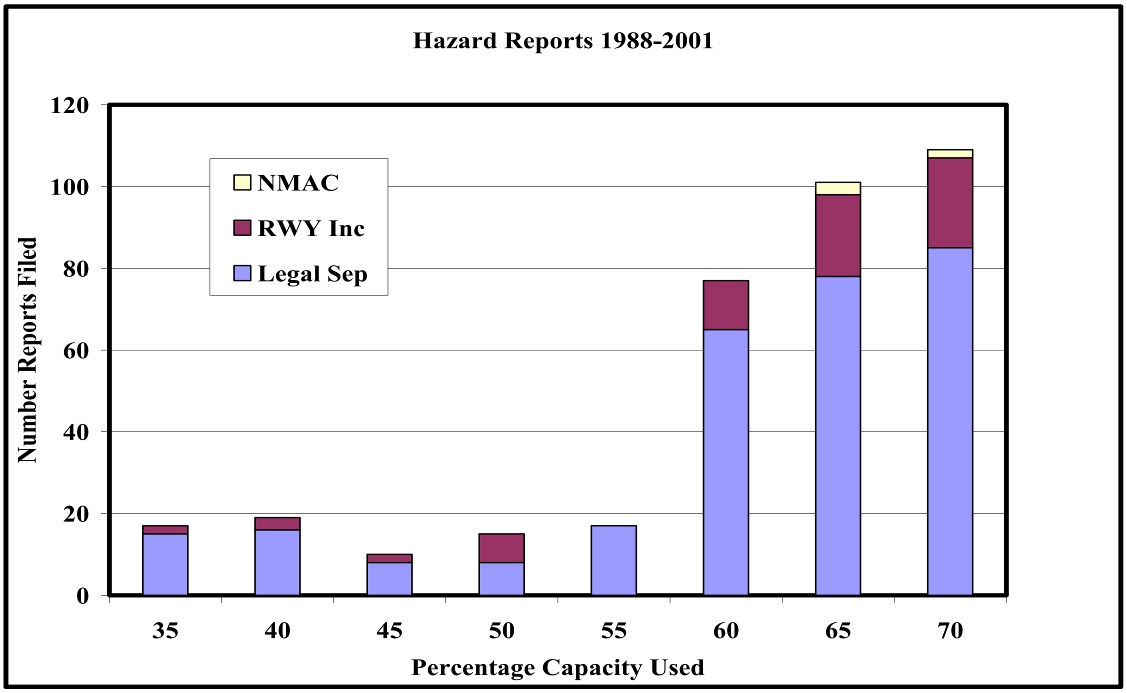

Figure 5. Number of Near Midair Collisions (NMAC), Runway Incursions (RWY Inc) and Loss of Legal Separation (Legal Sep) Reports filed at 4 airports over the last 13 years correlates with the capacity fraction the airport was operating at the time the incident Occurred (Haynie 2002)

2.0 Current Procedures for Allocating Landing Time Slots

2.1 Technical Procedures

A brief description of how aircraft spacing is determined is in order, as this will become a central issue in any demand regulation scheme. There is a long-standing, internationally recognized, safety principle that two aircraft should never be on an active runway at any one time. Thus aircraft must be separated by a time interval that achieves this fundamental safety objective. The concern is that the leading aircraft may not be able to exit the active runway for a variety of reasons and the following aircraft must not land until the active runway is clear in order to avoid a potential high-speed collision. Aircraft deceleration times vary and Runway Occupancy Time (ROT) can be represented as a Gaussian distribution with a mean of 40 seconds and a standard deviation of 8 seconds. Thus, there is a 99% probability that a preceding aircraft has departed the runway after 64 seconds and a 99.999% probability after 88 seconds. Figure 6 shows the relationship between arrival rates and mean aircraft spacing for a representative range of aircraft landing speeds. In general, larger aircraft approach a runway at a higher speed than do small aircraft because of the extra aerodynamic lift required to keep them aloft. Table 2 shows a representative sample of arrival rates set at a number of major international airports. Note that many are set at a maximum arrival rate of 40 arrivals/rw/hr. As can be seen in figure 6 this represents an AVERAGE aircraft spacing of 3 nautical miles or an AVERAGE time separation of 90 seconds.

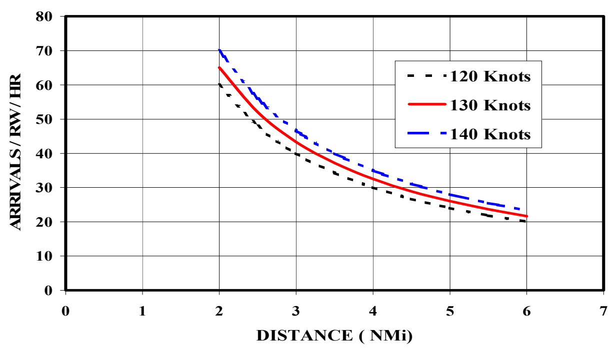

Figure 6. Relationship between Aircraft arrival rate and Average Aircraft spacing in distance and time for a representative range of approach speeds.

Table 2. Representative list of major international airports with published maximum runway arrival rates. Note that for this set of major airports the four non-slot controlled airports are all in the US.

Airport	Abr	No. R/Ws	Slot Control y/n	Approx Arrivals per Hr. *	% Wide Body	% Narrow Body	% Small Body
Atlanta	ATL	4	n	45	23	55	22
Auckland	AKL	2	y	42	16	17	67
Chicago	ORD	7	y	33			
Frankfurt	FRA	3	y	41	30	49	21
Hong Kong	HKG	2	y	25			
London	LHR	3	y	42	46	50	4
Los Angeles	LAX	4	n	40	25	43	32
New York	JFK	3	y	40	45	18	37
New York	LGA	2	y	38	7	63	30
Paris	CDG	2	y	42	18	58	23
San Diego	SAN	1	n	26	2	88	10
San Fransisco	SFO	4	n	30	29	52	19
Seoul	SEL	2	y	22	58	20	22
Singapore	SIN	2	y	33	70	30	0
Washington DC	DCA	3	y	42	4	57	39

Note: 40 Arrivals/Hr/Rw implies an average aircraft seperation of 90 seconds

Arrival per Hr. estimate is for an independent runway pair

With the introduction of the Boeing 747 wide-body aircraft into commercial service in the 1970's, the aviation community became aware of a new hazard for smaller, lighter-weight aircraft following very heavy aircraft. The lift required to support any aircraft ultimately gets left in it's wake in the form of both turbulence (a result of drag) and a set of counter-rotating vortex pairs. Aircraft are generally designed to withstand a significant amount of turbulence due to convective weather (i.e. thunder storms) but the coherent induced rolling encounter of a significantly smaller aircraft in the wing-tip vortex of a wide-body heavy aircraft can be fatal. In the 1970's and 1980's, conservative safe aircraft separation times were estimated (based upon wake vortex knowledge at the time) and established as separation times in excess of ROT during weather conditions that required air traffic control separation responsibility. In clear weather, the aircraft pilot was warned that he/she was following a heavy aircraft and the wake should be avoided (Unfortunately for the pilots, the wake vortex is invisible most of the time.) Table 3 shows aircraft separation time estimates taken from a study on the aircraft mixture and maximum arrival rate at LAX based upon wake vortex separation times (Hansen, 2001). This table illustrates that for many aircraft pair combinations observed at LAX, separation times in excess of 90 seconds should be maintained. This assumes that both the FAA air traffic control function and the pilot in the aircraft can assure planned separation to within accuracy of several seconds. Using today's technology, this is impossible.

Table 3. Representative aircraft separation time estimates for wake vortex separation (Hansen 2002). The shaded conditions represent aircraft sequencing whose separations are set by runway occupancy time safety margins (approx. 90 sec).

	Trailing Aircraft			
Leading	**Small**	**Large**	**B757**	**Heavy**
Small	80	68	66	64
Large	160	73	66	64
B757	200	120	100	100
Heavy	240	150	140	100

Figure 7 shows that almost half of the arriving aircraft at LGA have separation times less than the standards at an arrival rate of 39 arrivals/rw/hr (Haynie 2002). With today's technology, the safety limit is only 30 arrival/rw/hr. However, technology exists (Ballin, et. al. 2002) to reduce the variance of aircraft separation that would allow up to 40 or more arrivals/rw/hr. However, implementation of this technology requires that the airlines purchase and install new equipment on aircraft and that the FAA certify new procedures. In general, significant benefits to any airline will not be accrued until *all* airlines are equipped. No airline wants to be the first to equip and have this investment languish (Aerospace Industry Commission Report, 2002).

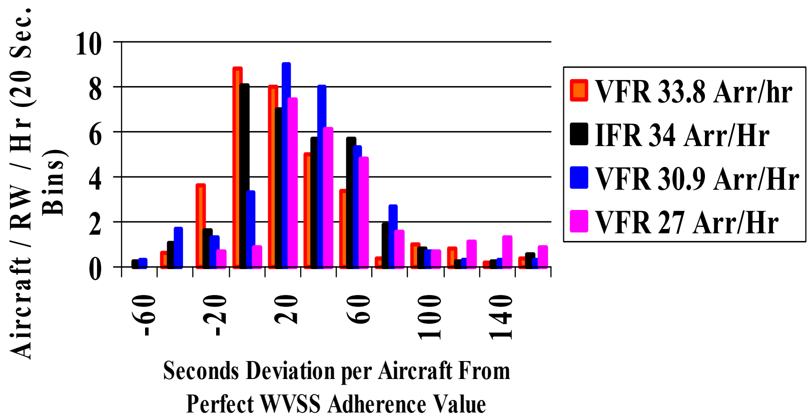

Figure 7. LGA Arrival Histograms by Period and Day, Normalized by Arrival Rate, Displaying Positive or Negative Deviation from safe Wake Vortex Separation (Haynie 2002).

2.2 Administrative Procedures and Property Rights

The definition of a slot is somewhat ambiguous. In general, the concept refers to the ability to access a resource over a particular time. Title 49 of the United States Code (U.S.C.) subtitle VII **(**49USC41714**)** defines the term **"**slot**"** to mean "a reservation for an instrument flight rule takeoff or landing by an air carrier of an aircraft in air transportation."

In the US, prior to 1969, the concept of an arrival or departure slot essentially did not exist. Rather airlines interested in providing service simply published schedules at the airports of interest and, on a given day of operations, requested access to runways as needed. There were other resources, of course, that had to be arranged to support such services. These included most notably gates, ticketing facilities, and baggage handling capabilities.

When the FAA instituted the high density rules (HDR) at JFK, LGA, ORD and DCA in 1969, it assumed responsibility for determining the appropriate number of slots and of overseeing their allocation. Slots were divided into three categories: air carriers, commuter airlines and general aviation. The separate allocation at each of the four airports was accomplished via airport-based scheduling committees, which were granted limited anti-trust immunity. The FAA intervened when committee deliberations became deadlocked. Slot ownership was ceded to the incumbent operators at the HDR airports in 1985. However, "use it or lose it" rules were established to ensure that each allocated slot was used. Slot exchange among owners has been allowed over the years.

On April 5, 2000, through the so-called "Air-21" Act, Congress made additional slots available at the HDR airport (Federal Register 2000). The restrictions placed on the usage of these slots effectively directed their use toward small aircraft providing access to small communities. By the end of September of 2000, there were 192 additional operations scheduled at LGA based on the Air-21 exemptions. Since the additional slots made available were not based on an actual increase in physical capacity, the increase in operations lead to substantial performance degradations in the form of extreme delays and numbers of flight cancellations (Fan and Odoni 2002). To address these performance problems, limitations on the number of Air-21 slots were instituted and these slots were allocated based on a lottery. The lottery is considered a temporary measure and replacement mechanisms are under investigation (Federal Register 2001).

Internationally, slot allocation activities are essentially similar to the US experience. Administrative procedures are used to allocate slots (see DotEcon Ltd 2001). Generally, there are provisions to encourage new entrants and access to small communities, however, grandfather rights predominate and limit the effect of such provisions. While there appears to be interest in exploring the use of market mechanisms there has been little or no use to date. In fact, in the UK, while the exchange of slots on a one-for-one basis is allowed, the buying and selling of slots is prohibited. Two IATA conferences, held each year in June and November, provide a forum for airlines to coordinate slot assignments relative to international schedules. At these conferences airport schedule coordinators and airlines confirm and possibly renegotiate slot assignments. Airlines also engage in slot exchanges. These activities are essential in order for airlines to properly coordinate schedules and to optimize ground, maintenance and other operations. They can continue after the conferences as airlines and airport coordinators continue to refine schedules.

When considering options relative to market mechanisms for allocating slots and, possibly related airport resources, a number of ambiguous legal questions arise. Regional airport authorities and municipalities typically own the airports in the US and Europe. Because these authorities are almost always public agencies, they are typically restricted in that their charges for services can only achieve cost recovery. Thus, the prospect of generating revenue streams from auctions that are well in excess of costs may be legally prohibited. In the US, virtually all of these airports have received federal aviation trust funds to partially fund infrastructure investments and have therefore agreed to abide by a number of federal regulations governing access to and the safe operation of the airport. Most major airports have also funded numerous land-side infrastructure investments with municipal bond debt financing. Many airlines have long-term leases with the municipal airport authorities on gates that give them exclusive rights to the gates. In fact, in many cases, the airlines have paid for the construction of these gates. The owner airlines can effectively use these long-term rights to exclude new entrants. These local gate ownership arrangements represent a major impediment to the prospect of auctioning gate access.

Currently, one might consider landing fees the charge most similar to a fee one might pay to lease a slot. However, landing fees represent a component of an airport's revenue stream that is used to recover overall airport costs. Conceptually, this fee pays for the cost required to maintain, and possibly construct, runways, taxiways and related infrastructure. On the other had, the FAA has exclusive and unambiguous rights to all aircraft use of the airspace. Title 49 of the United States Code (U.S.C.) subtitle VII

(49USC40103) states “The administrator of the Federal Aviation Administration shall develop plans and policy for the use of the navigable airspace and assign by regulation or order the use of the airspace necessary to ensure the safety of aircraft and the efficient use of airspace.” This authority is exercised daily with the routine practice of aircraft departure denial under Ground Delay Programs, the LGA lotteries and other activities have clearly established the precedent for network-wide government control of scheduled airline flights. Thus, the FAA “owns” the airspace immediately above airport runways and, consequently, has the authority to allocate the use of that airspace.

3.0 Slot Allocation and Reallocation on the Day-of-Operations

In an ideal world, airport capacity could be partitioned into well-defined slots and market or administrative mechanisms used to determine owners (or renters) of each slot. On any given day, slot ownership would afford an airline with the rights to carry out an operation within the precise time limits specified by the lease or ownership agreement. If the airline failed to exercise this right for any particular slot then no off-setting compensation would be granted. A variety of factors render this ideal impractical or impossible. Uncertainty on several levels necessitates a more flexible approach. We use the term *demand uncertainty* to categorize effects that can cause flights to fail to meet planned departure or arrival time slots and the term *capacity uncertainty* to categorize effects that cause changes to the number and/or timing of slots. Examples of factors contributing to demand uncertainty include problems in loading passengers onto an aircraft, mechanical problems, queues on the airport surface or in the air and en-route weather problems. Examples of factors contributing to capacity uncertainty include weather conditions at the airport and changes in flight sequences that cause the need to alter flight departure or arrival spacing.

In the US, under normal conditions, there are essentially no controls placed on an airline to adhere to “scheduled slot assignments”. Airlines have total control over when flights push back from their gates and, within the safety constraints imposed by ATC, control over when flights arrive. Within this framework, the stated FAA policy is to provide access to NAS resource on a first-come, first-served basis. That is, as aircraft push back from gates, they are placed in queues and given access to departure runways based on their place in the queue. Similarly, flights are placed in arrival sequences as they approach the airspace of their destination airports.

FAA traffic flow management (TFM) procedures in many cases make exceptions to this policy for safety and efficiency reasons. For example, flights routed toward congested airspace with flow (miles-in-trail) restrictions might not be given sequential access to a departure runway in order to delay their arrival to a portion of airspace. On the arrival side, flights can be sequenced to maximize airport arrival throughput using the capabilities of the Center Tracon Automation System (CTAS) (Erzberger 1995). It is noteworthy that IATA guidelines explicitly recognize a decoupling of schedules and the actual timing of operations (IATA 2000):

> "The Conferences deal with adjustments to planned schedules to fit in with the slots available at airports. This activity has nothing to do with adjustments to schedules on the day of operation for air traffic flow management. The two types of slot allocation are quite different and unrelated."

In the US, the most significant TFM adjustments to operations occur during ground delay programs (GDPs). FAA traffic flow managers institute a GDP whenever the anticipated arrival demand is significantly greater than the estimated arrival capacity at an airport. This most often occurs when degraded weather conditions at an airport cause a change from visual flight rules (VFR) to instrument flight rules (IFR).

GDPs have been used by the FAA for close to 20 years now. Recently however, the emergence of a new paradigm for TFM has led to significant changes in the implementation of GDPs. This paradigm, called *Collaborative Decision Making* (CDM), is based on the recognition that improved data exchange and communication between the FAA and airlines will lead to better decision-making (see Wambsganss 1996; Ball et al 2000). In particular, the CDM philosophy emphasizes that decisions with a potential economic impact on airlines should be decentralized and made in collaboration with the airlines whenever possible. The GDP enhancements introduced under CDM are numerous, and include improved data-exchange, better situational awareness tools, and increased flexibility for the airlines. All major US airlines participate in CDM. An extranet interconnects the airline operational control centers with the ATCSCC and all participants have a common decision support tool, the Flight Schedule Monitor (FSM). The most significant improvements from CDM are derived through the use of different procedures for allocating ground delays. Under CDM, arrival capacity is allocated to the airlines by a procedure called Ration-by-Schedule (RBS). RBS is based on the consensus recognition that airlines have claims on the arrival schedule based on the original flight schedules. In addition, CDM has introduced a new reallocation procedure called Compression. This procedure aims to ensure optimal capacity utilization in the presence of delays and cancellations.

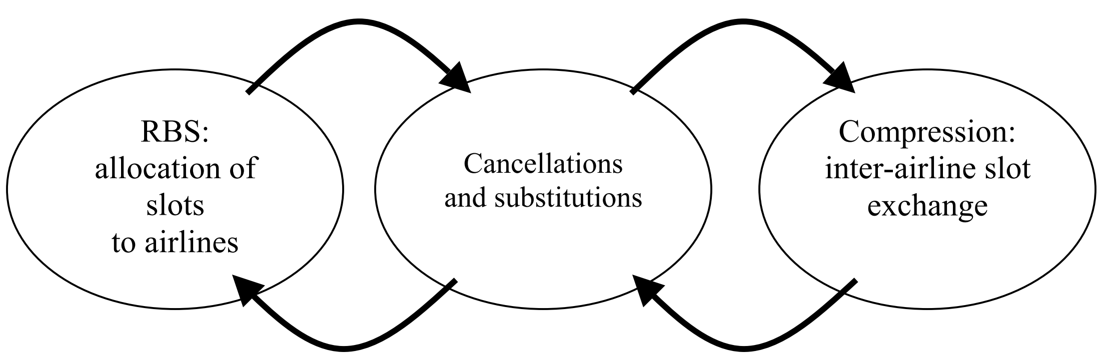

Figure 8: CDM Resource Allocation

Figure 8 illustrates the overall allocation process. Note that, in addition to RBS and compression, there is a third significant process, *cancellations and substitutions,* that is totally controlled by the airlines. Under this process, each airline may cancel flights and interchange slot-to-flight assignments for its own flights. Thus, although RBS in concept allocates slots to flights, the cancellation and substitution process effectively converts the slot-to-flight assignment into a slot-to-airline assignment.

The principal output of either RBS or Compression, is a controlled time of departure (CTD) for each flight in the GDP. Calculation of a CTD is accomplished by assigning a controlled time of arrival (CTA) to each flight and then computing a CTD by subtracting the estimated enroute time from the CTA. The assignment of CTAs by RBS can be viewed as a simple priority rule. A set of arrival slots consistent with the degraded capacity is created. Using the OAG arrival order as a priority order, each flight in the OAG is assigned the next available arrival slot. If this rule was applied to all flights and there were no cancellations or substitutions, then the flights would arrive in their original sequence but generally later in time. There are two groups of flights exempted from this basic allocation scheme: 1) flights that are currently airborne (clearly these cannot be assigned ground delay) and 2) a set of flights characterized by the distance of their departure airports from the GDP (arrival) airport (see Ball and Lulli 2002). The motivation for 2) is to include in the allocation scheme flights close to the airport and to exempt flights further away from the airport. Flights a greater distance away must be assigned ground delays well in advance of their actual arrival, e.g. 4 or 5 hours. At such a long time frame in advance of arrival, there tends to be a greater level of uncertainty regarding weather and, as a consequence, airport arrival capacity. Thus, if these flights are assigned ground delay, then there is a significant likelihood that the ground delay could be assigned unnecessarily. Thus, distance-based exemptions constitute a mechanism for improving expected airport throughput.

After a round of substitutions and cancellations the utilization of slots can usually be improved. The reason for this is that an airline's flight cancellations and delays may create “holes” in the current schedule; that is, there will be arrival slots that have no flights assigned to them. The purpose of the Compression algorithm is to move flights up in the schedule to fill these slots. The basic idea behind the compression algorithm is that airlines are “paid back” for the slots they release, so as to encourage airlines to report cancellations.

To illustrate the compression algorithm, let us consider the example shown in Figure 9. The leftmost figure represents the flight-slot assignment prior to the execution of the compression algorithm. Associated with each flight is an earliest time of arrival, and each slot has an associated slot time. Note that there is one cancelled flight. The rightmost figure shows the flight schedule after execution of the compression algorithm: as a first step, the algorithm attempts to fill AAL's open slot. Since, there is no flight from AAL that can use the slot, the slot is allocated to UAL, and the process is repeated with the next open slot, which using the same logic is assigned to USA. The process is repeated for the next open slot, which is assigned to AAL. The AAL receives the earliest slot that it can use.

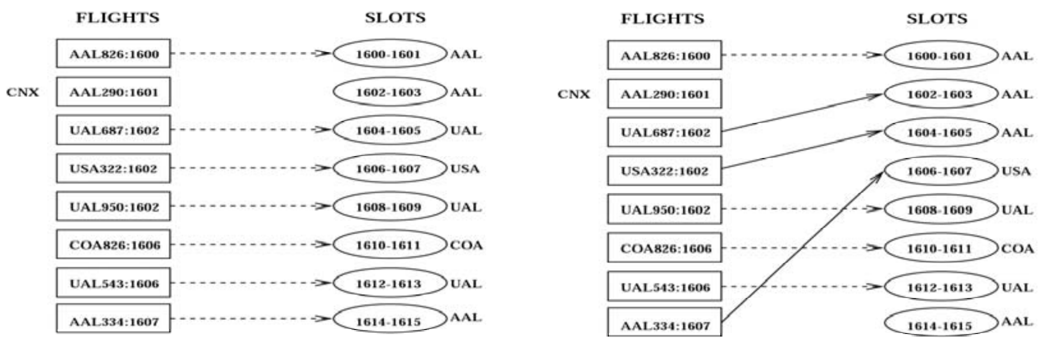

Figure 9: Execution of Compression Algorithm

The Compression algorithm results in an exchange among airlines of the initial RBS allocation. One could interpret this result as a reallocation. However, there is also a natural interpretation of Compression as an inter-airline trading or bartering process (see Vossen and Ball 2001). For example, in Figure 9, American Airlines "traded" the 1600 – 1601 slot, which it could not use for the 1607 – 1607 slot, which it could use, and United Airlines reduced its delay by trading the 1604 – 1605 slot for the earlier 1602 – 1603 slot. Vossen and Ball show that a bartering process can be structured so as to produce a result essentially equivalent to Compression. This view of Compression suggests many possible extensions. For example, (Vossen 2002) defines a more complex 2-for-2 bartering mechanism and shows that there is a substantial potential for improved economic performance from this using mechanism. Probably the most intriguing enhancement is to allow "side payments" with any exchange as well as the buying and selling of slots. Such a process, which can be viewed as a day-of-operations aftermarket, is discussed and analyzed in Section 8. It is noteworthy that experience with the current process can provide insights into the design of future market mechanisms and, also, that the CDM IT infrastructure can potentially serve as a basis for implementing such a market.

4.0 Objectives and Fundamental Issues Associated with the Design of Auctions within Aviation

The concept of creating a market-clearing mechanism for the allocation of takeoff and landing rights at airports is not new. (Rossenti, Smith and Bulfin 1982) provided the first general package-bidding framework for slot auctions. They suggested a combinatorial auction design that allowed bidders to be able to couple takeoff and landings in a single bid of the form "I want to purchase the right to take off from airport A at time X and also have the right to land at airport B at time Y for a total price of $Q." and they acknowledge that even more complicated expressions of uses and related values may be needed. For example, an airline may need to express a collection of arrivals and takeoffs in a single bid to ensure that arrivals at a given airport are coupled with a takeoff from that airport and each takeoff is coupled with a landing at some other airport. The need for alternative collections, describing different business plans, also argues for a bidding language that allows bidders to say that two packages are mutually exclusive. With this type of expressive language, a bidder does not experience the concern that he will win only part of what he needs to create a commercially viable schedule of flights or conversely, that he will win more than expected (called the “exposure” problem). A very detailed and excellent account of the issues and benefits associated with auctions for slot allocation can be found in a report (DotEcon, Ltd. 2001). We begin this section with a short summarization of some of these issues.

The current process for slot allocation whereby airlines have landing and takeoff rights grand-fathered for decades at airports that are over-capacitated leads to inefficiencies, inertia, and distorted incentives (e.g. use available slots so as to keep competition at bay). When slots are allocated through an administrative process, the cost of ownership is only the cost of acquiring the slot. When slot are scarce, the decision of assigning new slots must be somewhat subjective and vulnerable to legal challenge. Such decisions often limit entrance or expansion of new carriers into a given market because they have less ability to impact the outcome of the administrative process.

On the other hand, the grand-fathering of slots (a) allows a maximum degree of certainty over future slot holdings thereby helping an airline with its long-term investment decisions; (b) makes the job of determining the allocation less cumbersome, because it is done infrequently and (c) helps the airline do long-term planning since it can predict most of its competition over long periods of time. Finally, the management of the overall airspace is less complex because the activities at a given airport or over the entire airspace are predictable and change little over time. Each of these benefits, however, also highlights how such allocation impedes competition and change over long periods of time.

We believe that the use of a market-based allocation system, such as auctions, for the allocation of slots for use over a specified time period is more likely to provide a system that is efficient, i.e. produces an allocation that results in maximizing the benefits to the consumer and the economy by allocating them to those that can generate the greatest benefit from the use. The knowledge that the slots will be re-auctioned in the future assures that the industry must actively evaluate the market and the value of ownership of such slots. Another benefit of the use of auctions is that the process is both transparent and is less open to legal challenges.

One argument against auctions is that they are an easy mechanism for the government to raise revenues. Auctions do raise revenue, but the revenues raised are no more than a reflection of the market value associated with a scarce resource. How the revenue generated is used warrants discussion. It is our belief that the revenue can be used to help pay for the infrastructure necessary to safely administer and expand airspace use. Since such costs must be incurred somehow, one can argue that the revenues generated are in lieu of other taxes that would be required if they were not generated from the auctioning of slots. However, we note that careful auction design often works to *not* maximize the revenue from high value bidders, but rather chooses objectives that encourage new entries and discourages or disallows monopolistic control over markets. Thus, the revenue generated is a consequence of ensuring efficient outcomes but need not generate any more than the minimum necessary to do so.

There are a number of design issues that must be considered for an auction that allocates slot at airports. We discuss many of these issues prior to providing the details of a basic auction framework.

Capacity Issues: Since airport slots cannot be considered separately from the other resources that will be used with the slot, one must consider runway, gating, baggage and terminal capacity when determining the number of slots that might be available within a given time period. One must also know how the slots are to be used – i.e. What is the distribution of aircraft types that will use the slots? Since the distribution impacts the capacity of the system, the number of the items (e.g. arrival slots between 9:00 am and 9:15 am) to be auctioned may vary depending on the type of aircraft that will use these arrival slots. Since the *number* of slots may be dependent on demand characteristics that are unknown prior to the auction, the auctioneer – when determining winners – must assure that the physical limitations of the airport are satisfied when determining a feasible allocation set.

Property rights: Currently, airlines receive rights that they expect to be able to renew indefinitely if they adhere to conditions of usage specified by the FAA. In fact, these "use it or lose it" requirements have encouraged carriers to inefficiently use slots allocated to them in order to assure continued ownership and to preclude competition by other carriers. Although allocation exercises take place semi-annually, the turnover of slots is very small. When airports expand, new slots become available and an administrative process measures the value of providing these slots to new entrants versus expanding the slots of the major existing carriers at that airport.

Although carriers *currently* enjoy what they consider "perpetual right of usage", this usage does not confer any future property rights. Indeed, the law indicates that the DoT has exclusive rights to determine the takeoffs and landings at all airports in the continental U.S.

Property rights would need to be carefully specified. One approach would allow the buyer of a slot to have exclusive use of the slot during a given time window (e.g. within a fifteen minute period) every day. The owner of this slot would keep this right over a relatively long period of time, thereby providing the owner with the ability to create long-term plans and schedules. For example, if 20% of the slots were auctioned each year, then the duration of the property rights for any given slot could be a five-year period. Research to determine the optimal period for the right is necessary.

In addition to these yearly auctions, we perceive the need for a secondary auction that would allow trades of slots among carriers. Thus, the carrier acquires an assumable lease for the given lease period. That carrier has the right to trade or sell this slot for any subset of the leasing period to any other carrier contingent upon regulatory conditions specified by the DOT/FAA (e.g. proof that the buyer-carrier can operate in this market safely and that this carrier adheres to any limitation on market share.)

Finally, a third trading mechanism must exist for the day-of-operations. This market is a classic "trade" exchange among carriers having rights to takeoff and land at a given airport to accommodate issues of weather, maintenance, and other problems that cause planes to be delayed in departure or landing.

Thus, if the FAA decides to allocate long-term rights to departures and landings through a market clearing mechanism, then the FAA must also stipulate how carriers can trade these slots during the period of that lease. In Section 5-7, we will discuss the differences in auction design necessary for each of these three situations.

Valuation Issues: Currently, the value of a takeoff or landing slot is unknown. This is consistent with many other government administrative processes. For example, prior to the Federal Communication Commission auctioning off spectrum rights, the value of this spectrum was unknown but perceived to be quite valuable. When the first PCS auction was held (this was an auction for wireless communication), and brought over $7 billion into the US Treasury, the value of spectrum was no longer questioned. Similarly, since the values of airline slots are not well understood at present, we believe that one should choose an auction design that allows price discovery.

Market Power: Auctions allocate markets efficiently when there is sufficient competition within the market. In this case, competitors bid against each other and the market is allocated to the bidder that values the market the most and is, therefore, willing to pay the associated price of acquiring the market. However, for markets where the goods are scarce and where the goal is to assure competition within the industry, the rules of the auction must ensure that an airline does not bid based on the airlines' willingness to buy a monopoly. Without rules to preclude this situation, an auction can lead to monopolistic control of an airport or region. One can avoid this pitfall by placing ownership restrictions both on the entire airspace, within any given airport, and even within an airport by time of day. Carriers may need a certain amount of activity at given airports to optimize their hub/spoke system, but no airline needs to control a majority of the airport. Rules similar to those imposed by the FCC on spectrum, can control total ownership of slots within regions, airports and within time windows.

Implementation Issues: A transition period that moves the airline industry from an administrative process to a market-clearing process must be carefully thought through. As discussed earlier, we would expect only some small percentage of the slots to be auctioned in any given year; our intitial suggestion is 20% per year. The question of where to start must be studied (see (Rothkopf 2002) for a general approach to this problem). We propose that one starts with the most congested airports, and during the most congested time periods, thereby relieving some of the safety and delay issues that exist at these airports. Since peak-demand slots will command higher prices than off-peak hour schedules, auctioning these slots first will force airlines to reexamine the sizing of flights that take off during these high-demand times. Thus, although takeoff and landings might decrease (for safety reasons), the number of passengers being serviced during these time periods may remain the same or increase. Similarly, the less demanded flights, with smaller aircraft, are likely to move to less demanded time periods.

We emphasize that the definition of a "slot" is either a takeoff or a landing within a given time interval. This "slot right" provides the airline with the right to schedule a published arrival or departure during that time window. We believe that a "time window" of 15-minutes is reasonable given the physical limitations of aircraft movements. A larger window could allow airlines to "bunch" takeoffs and landings during a peak interval resulting in significant delays and safety concerns.

If one "wins" a slot in this auction, then one must also acquire commensurate rights to terminal space, e.g. ticketing, baggage and gating facilities. Thus, it is essential that the winner be able to obtain such rights, e.g. by paying the "going rate" for these facilities as computed based on the current-long-term contracts with the local airport authority. Dominant carriers have historically limited rivals from entering the market by limiting the available gates at the peak demand periods. When dominant carriers lose their ability to control gates and ticketing counters, competition during these periods is likely to increase at capacitated airports. The gradual transition over a five-year period will allow dominant carriers to continue to have their "hub" presence since no more than 20% of the flights at the airport would change during a given year.

In the next few sections, we highlight some of the specifics of auction design relevant to this application.

5.0 Design Principles and Research Questions for Long-Term Lease Auction

When considering the design principles for the long-term lease auction, one must first be certain that all property rights are well defined. Such rights must include the length of the right (e.g. for *x* years), the right itself (arriving at runway *y* at airport A between x and y time on an aircraft of size *b*), the transfer rights, and the obligations (safety, adhering to restrictions on market concentration, etc). One must also determine what happens if the airline cannot use the slot at the time allocated. We are presuming that the airline must either cancel the flight or bid in the same-day auction for a different time slot. We also note that rules must specify what happens if the airline wishes to substitute an alternative aircraft type into that slot. We only mention these "same day" issues here because any ambiguity in the overall process can create bad outcomes. Similarly, rules that specify the maximum amount of slot ownership at any airport, region or globally must be carefully considered. Once property rights have been completely and unambiguously specified, a one-sided combinatorial auction can be designed. This auction design is one-sided because the US Government is the seller and is selling a collection of slots. Thus, there is only one seller and multiple buyers.

Since the success of an auction design is dependent on all of the details "fitting", we next present a "straw man" for discussion and study.

General framework:

We believe that the overall framework for this auction should be a simultaneous multiple round ascending bid auction. Because the items being auctioned are scarce commodities with both private and common values (i.e. the value to a buyer is based partially on the value that others place on this item and partially based on the buyer's own business plan), there is a strong need for price discovery. Such auctions also have the added attribute that they close at slightly over second highest prices. Since all items are awarded simultaneously, buyers can alter their business plans as they collect information about prices and competition among slots. Because of the complexity of an airline's overall business plan, there is a strong need for an expressive language that allows bids to be treated as mutually exclusive.

Since there have been few combinatorial auctions, we believe that careful study – both empirically and computationally is necessary to assure its success. We present some of the components below:

1. *activity:* We argue that activity rules similar to those set by the FCC in simultaneous multi-round auctions be considered.

2. *bidding language:* The literature presents a number of expressive bidding languages. The general XOR language allows complete expressivity but at the cost of requiring an extraordinary number of bids to be placed by the bidder. Alternatively, other bidding languages are compact but only appropriate for specific kinds of bidders. We note that (Ledyard 1999) describes the need to express mutual exclusive bids based on his experience in transportation auctions. For more on bidding languages, see (Nasam Chapter 9).

3. *bidder's aid tools:* Providing the bidder with tools that help the bidder to compactly express her business plans make package bidding auctions easier. They include tools that (a) help the bidder estimate a competitive price for his package; (b) help the bidder determine his *best bids* based on his business plans and the current prices of the various slots and (c) help the bidder to identify "partnering" bids (i.e. bids that fit with his bid to create a winning set).

4. *pricing:* In a combinatorial auction, it is hard to infer the prices of the individual pieces from the winning prices of packages. Yet, this information is critical to a bidder being able to create new packages that have some hope of "winning" in future rounds. There are a number of papers that discuss both *linear* pricing and *nonlinear* pricing (see Hoffman,et al. Chapter 18 for details). Pricing information is critical to both the overall outcome (efficiency) and speed of the auction.

5. *winner determination calculations:* Since a slot allocation auction is likely to have many objects auctioned simultaneously, the size of the winner determination problem can get very large. The federal government must weigh the benefits of

speeding up these calculations against the issues associated with fairness and transparency. Indeed, the FCC chose to not only solve the winner determination problem to proven optimality but also to choose among *all tied solutions* randomly. These calculations are expensive, but may be necessary to assure that the results of the auction are not contested in court.

6. *length of time bids are kept:* In order to assure sincere bidding, we believe that bids should be kept active throughout the auction.

7. *setting the minimum bid increment:* The minimum bid increment determines the allowable new bids for the next round. If this increment is set too small, then the overall length of the auction will be extended significantly. If, on the other hand, the increment is too large, efficiency can be lost. Smoothing procedures that use a relatively large bid increment when there is much competition and a relatively small increment when there is little competition should be considered.

8 *proxies:* Many auctions employ either an optional or a mandatory computerized facility that bids for the bidder within an auction, called "proxies". Proxies allow bidders to participate in multiple rounds without the cost of continuing monitoring the system. It has been shown that forced proxy bidding can eliminate much signaling and gaming, because the proxy will always bid in a straightforward manner (see (Ausubel and Milgrom, Chapter 5) for details). We believe that for the slot-allocation auction, such designs would require modification to allow "stages" whereby bidders could re-adjust their maximum bid prices and set of bids. With stages, eligibility and activity rules would need to be imposed. Little experience with such an auction system exists, and careful study of these systems is warranted due to their ability to reduce disruptive bidding.

9 *stopping rule:* In a multiple round auction, the stopping rule must be coordinated with the eligibility and activity rules of the auction. The rule that the auction ends when no new bids are placed encourages activity. However, in package bidding, there can be many rounds with new bids but with little or no change. An alternative with promise is to consider merging an ascending bid auction with a "last proxy stage". In this stage, bidders provide the "maximum bid amount" for every bid they might wish to win and the proxy works with these bids to determine the winners.

Clearly, for an application as critical as slot allocation, one must be careful in choosing the auction design. The problem is sufficiently important to warrant extensive experimental and computational study of the alternatives.

6.0 Design Principles and Research Questions for Medium-Term Exchange of Slots

In order for a market-clearing strategy to work, airlines must have the right to reassign their lease to another airline for the remaining lease period. The only regulatory oversight of such transactions would be to insure that airlines do not, by such transactions, violate any of the rules of aggregation originally specified for the long-term auction. Thus, the FAA would specify rules restricting the overall leasing of slots at (a) a given airport, (b) within a given region and (c) throughout the US airspace. Trades must not violate these rules. Within this overall restriction, the industry can determine the market trading mechanism that they feel most suits their needs.

These trades can be performed by bi-lateral negotiation or through an organized market (auction) mechanism. If the airline industry chooses an auction mechanism, the choices are large: the mechanism can be a simple one-sided exchange mechanism whereby one lists slots for sale, and an auctioneer then sells these for the airlines. Alternatively, one can create a two-sided combinatorial exchange whereby one can create packages of items for sale where items are exchanged either (a) whenever a seller's ask price "matches" a buyer's bid price; or (b) at the end of a full combinatorial exchange whereby bidders provide complicated buy and sell bundles and trades take place simultaneously with the surplus from those trades being allocated based on a given set of rules. A summary of issues associated with combinatorial exchanges can be found in (Parkes 2001; Wurman 2002; Ledyard 2002). The auction design for this application is likely to evolve over time. We believe that as long as property rights are established and the industry is free to exchange these rights in the market place, the industry will determine the exchange that works best.

We now move onto the issues associated with what occurs on day of operations, given the fact that weather conditions, mechanical breakdowns and other operational conditions might significantly alter an airline's ability to adhere to its published schedule.

7.0 Day of Operations Slot Exchange

Section 1 of this paper argued the merits of the general use of market mechanisms within air traffic management. In the section, we first specifically consider whether a day-of-operations market mechanism may provide significant added value when compared with the current CDM procedures. We then discuss design issues related to such a market mechanism.

7.1 Value Proposition and Need

If one examines the two types of offers associated with the bartering model described at the end of Section 3 as well as Figure 10, it becomes apparent that the current slot exchange procedure addresses a very specific, limited scenario. That is, compression and/or the bartering process are driven by one or more slots that have become unusable by their owner-airlines. As illustrated in Figure 10-i, airline A owning an unusable slot, [x], places a value of $0 on [x] but say a value of $1000 on a slightly later slot [y]. On the other hand, airline B might place a value of $800 on [y] and a value of $900 on [x]. An exchange of ownership provides an added value of $1000 to A and an added value of $100 to B for a total economic surplus of $1100. It would seem clear that many other types of scenarios are not only possible but likely. In Figure 10-ii, airline A's flight has a relatively small number of passengers and airline B's has a much larger number, leading to A's valuations of [x]: $500, [y]: $450 and B's valuations of [x]: $2000, [y]: $1600 for a total economic surplus of $350. In Figure 10-iii, by switching from slot [y] to slot[x], airline B is able to avoid a timing out of its crew thus, saving a much larger incremental amount than A incurs. The second and third scenarios are quite realistic, but could not be addressed under the current CDM procedures.

Recent research on non-monetary extensions of the current CDM exchange mechanism provides further evidence of the potential value of a day-of-operations market. In (Vossen 2002), an extension from the current 1-for-1 bartering model, i.e. compression, to a 2-for-2 bartering model is described. Experimental results show that airlines can derive substantial additional value from using such a system.

Because of the high degree of uncertainty in daily NAS operations there is always a need to dynamically adjust schedules. The evidence just provided clearly indicates that during time of severe or even moderate disruptions, substantial value can be derived from a day-of-operations slot exchange market.

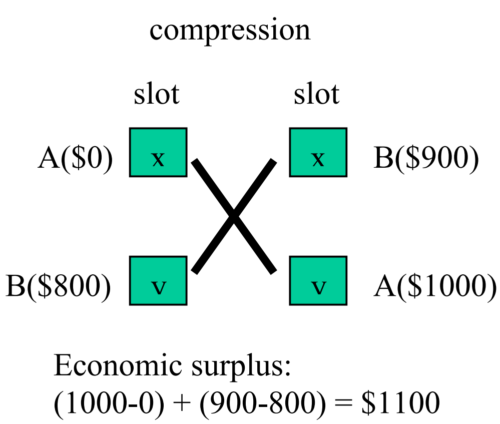

Figure 10-i

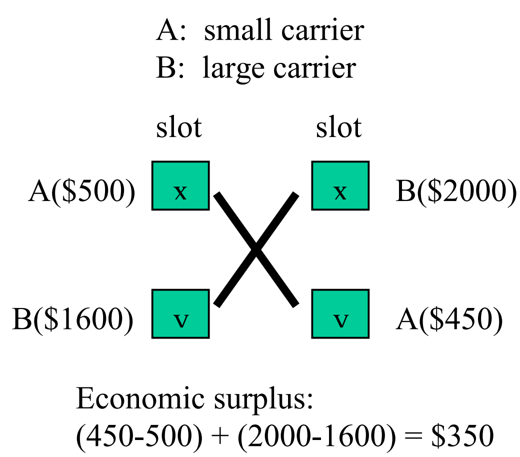

Figure 10-ii

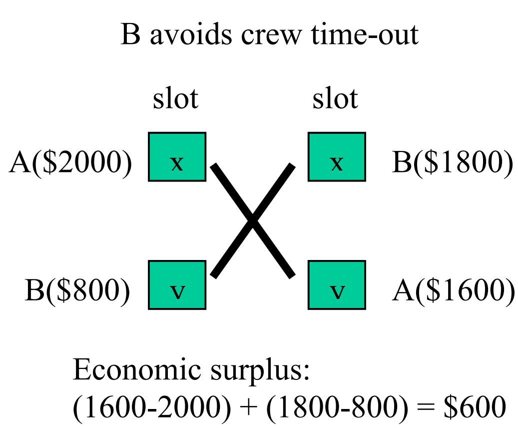

Figure 10-iii

7.2 Design Issues

The current CDM exchange mechanism, i.e. the Compression algorithm, together with the CDM data communications infrastructure provides a firm foundation for structuring and understanding a possible "day of operations" slot exchange market. Viewed most simply, one "only" needs to convert this non-monetary exchange into a monetarily-based exchange. Prior to pursuing this approach, it is worthwhile to step back and consider the basic premise for such a market so that more fundamental changes can be considered.

Currently slot rationing and slot exchange are employed for arrival slots during periods of time when the arrival capacity at an airport is reduced usually as a result of poor weather conditions. These operations occur before the departure of the impacted flights. The arrival slot allocated to each flight is converted into a departure slot by subtracting an estimate of the en-route time. The difference between the revised departure time and original departure time of a flight is the ground delay assigned to a flight. Thus, ultimately control is exercised through the assignment of ground delays while planning and allocation are carried out in terms of arrival slots. No constraints are placed on the revised departure times so there is an implicit assumption that departure capacity is unconstrained.

When considering more general market structures, some questions naturally come to mind.

1) *Should a day-of-operations slot exchange market only exist during conditions of reduced arrival capacity?*

Air carriers routinely enter into conditions of irregular operations. Under such conditions there can be large deviations from normal schedules so the set of required slots can substantially deviate from the set "owned" by the carrier. Clearly this suggests the usefulness of market mechanisms to achieve a reallocation under a broad set of circumstances. On the other hand, today it is only deemed necessary to carry out an explicit allocation during times of reduced capacity since relatively few delays result from exercising a basic first-come, first-served rule. It seems clear that some expansion over the current practice is warranted, however, the extent of such an expansion deserves additional consideration.

2) *Should both arrival and departure slots be exchanged in the market?*

At this time it is difficult to provide a definitive answer to question 2). It certainly is clear that there are times of significant congestion with respect to departure slots, which strongly suggests the need for allocation mechanisms. On the other hand, the notion of dynamically coordinating NAS-wide arrival and departure slot allocation, e.g. through a NAS-wide auction with package bidding would seem to be extremely onerous and perhaps overkill. Thus, we will proceed here with the development of a day-of-operations slot exchange for arrival slots. At the same time, we view 2) as an open question worthy of study.

3) *Should the basic control via departure times remain the same or should other control mechanisms be used?*

The current approach to controlling arrivals may seem rather convoluted, in that *arrival slots* are allocated and this allocation is implemented by controlling when flights *depart*. The logic of this strategy might seem especially dubious when one considers the high degree of uncertainty associated with departures and en-route traffic flows as described in Section 3. In fact, one of the goals of CDM development activities has been to migrate to the so-called "control by CTA". This concept states flights should be given a CTA (controlled time of arrival) and that the responsibility to meet the CTA rests with the airline/pilot. No restrictions would be placed on departure times. In fact, we do not believe this to be a desirable method of control for CDM nor do we believe it is desirable under a future day-of-operations auction scenario. The principal disadvantage of such an approach rests in the fundamental difference in how flights on the ground and flights in the air can be handled. Once a flight is airborne, by necessity, it must be granted certain priorities and privileges. For example, the fuel on-board would represent a hard constraint on the degree to which it could absorb airborne delay.. Of even greater concern, is that prospect that if significant penalties were attached to either early or late arrivals into the terminal area, then it is likely that the added pressures placed on pilots and air traffic controllers to meet certain time windows would lead to a degradation in safety. Thus, our answer to question 3) is that the current approach of allocating arrival slots and controlling access to those slots via departure time modifications should be continued.

To summarize, we have concluded that a basic "conversion" of the current system to a monetarily-based system is the best alternative to pursue at this time. We note that even when proceeding along these lines, many design issues and open problems remain. We mention three below.

Exchange markets: A key property of these slot-trading markets is that each airline is potentially both a buyer and seller. In fact, the natural extension of the current exchange system suggests simply adding the possibility of side payments to the current trades. While outright slot buying and selling should be supported, the dynamics of the situation would seem to mandate support exchanges with side payments, which represents a design challenge.

Package bidding: **T**he results of Vossen (2002) show that the use of complex, e.g. 2-for-2, trades can lead to substantially improved performance under certain assumptions. These facts point to the possible advantages of exchanges involving sets of slots. However, the merits of supporting general package exchanges is less obvious when side payments are supported.

Dynamic bidding vs discrete cycles: Currently the slot exchange mechanism, compression, is executed periodically, e.g. every ½ hour. A more dynamic, transaction oriented system (slot credit substitution) has recently been implemented. The choice between a dynamic (stock market-like) approach and the use of discrete cycles represents an open design question.

8.0 Conclusions

Airport arrival and departure slots at certain critical airports have become scarce commodities that have substantial intrinsic value. In spite of recent downturns in demand for air transportation, demand for the most desirable slots remains greater than supply. Further, there is little doubt that this condition will remain in the near term and most likely will worsen in the long term.

The effective provision of scheduled air transport services requires the acquisition of multiple coordinated slots as well as related airport resources, including access to gates and terminal facilities. Complementary slots are required both at a single airport and at multiple airports.

In the U.S. the FAA has statutory authority to control access to the airspace and consequently has authority over the allocation of slots. At many airports, certain airlines hold implicit or explicit grandfather rights to certain slots. Airport resources, most notably gates, are typically owned and managed by local public airport authorities. However, in many cases, airlines owe long-term leases to gates. Slots cannot be effectively used without corresponding gate resources so any approach to slot allocation must allow for the allocation or reallocation of corresponding gate resources.

At many airports, the airport capacity, i.e. number of available slots, depends on the weather. This dependency implies a high degree of “supply” variability and uncertainty. Variability in a variety of airline and air traffic control processes makes it difficult to achieve precise timing in the delivery of aircraft to departure and arrival slots. This “demand” uncertainty coupled with the supply uncertainty make it necessary to employ a flexible approach to the rights associated with slot ownership on a particular day of operations. Furthermore, a day-of-operations reallocation or exchange of slots can produce substantial improvement in overall system and market efficiency. Such a reallocation currently is used under the Collaborative Decision Making procedures used for the planning and control of ground delay programs.

The various circumstances of the aviation slot allocation setting suggest the possible use of three types of market mechanisms. First, an auction of long-term leases of arrival and/or departure slots at capacity constrained airports. Second, in order to refine the initial allocation and achieve market efficiency, a market that allows inter-airline exchange of such long-term leases. Third, a near-real-time market that allows inter-airline exchange of slots on a particular day of operation.

The absence of an effective slot allocation mechanism has led to the gradual degradation of safety levels as well as instances of excessive delays. The use of market mechanisms offers the prospect of addressing these issues as well as providing a number of other economic benefits including use of existing slots by air carriers who can most productively make use of them, easier entry into markets by new and emerging carriers, breaking up of monopolistic “fortress hubs” and disclosure of true market value of slots.

Of course, auctions of long-term slot leases will generate additional revenue. This most naturally should be used to invest in NAS infrastructure and aircraft equipage. Identifying funds or economic necessity for such investments has been difficult in the past.

The successful use of auctions for telecommunication spectrum, energy and other commodities provide valuable insight into how to design auctions for NAS resources. Furthermore, this experience provides ample evidence that effective auctions can be designed and implemented in the aviation setting. A number of issues and problems need to be addressed before effective auctions for NAS resources can be designed. These include determining the number of slots offered per time period at an airport; defining the property rights associated with slot ownership; transition issues related to current implicit and explicit airline property rights and handling the need for airlines to acquire groups of slots and related airport resources, through package bidding and/or after-markets. Auction design for both long-term slot leases and for slot exchange on the day-of-operations, while fundamentally similar to other cases previously tackled; present a number of challenges, some of which represent interesting research challenges.

We believe that market mechanisms should play a fundamental role in any comprehensive approach to the management of an air transportation system. Furthermore, such mechanisms show strong promise for improving the safety, delay performance and economic efficiency of today's NAS.

9.0 Acknowledgements

The work of the first author was supported in part by NSF Grant: ITR0205489. Many of the ideas for this paper grew out of the workshop, "National Airspace System Resource Allocation: Economics and Equity" organized by the authors. The workshop was supported by George Mason University, NEXTOR, the National Center of Excellence for Aviation Operations Research and the University of Maryland.

References:

Air Transport Action Group (1998), *Airport Capacity / Demand Profiles* , database on CDROM published by the International Air Transport Association, London.

Alamdari, F. E. and P. Morrell (1997), "Airline labor cost reduction: Post-Liberalism experience in the USA and Europe", *Journal of Air Transport Management*, 3, 2, 53-56

Aviation Week and Space Technology (2002), "Justice Dept. Backs Slot Auctions at LGA", July 15, 2002, pp.37-38.

Ball, M.O., R. L. Hoffman, D. Knorr, J. Wetherly and M. Wambsganss (2000) "Assessing The Benefits Of Collaborative Decision Making In Air Traffic Management", in *Proceedings of 3rd USA/Europe air Traffic Management R&D Seminar*.

Ball, M.O. and G. Lulli (2002) "Ground Delay Programs: Optimizing over the Included Set Based on Distance", submitted for publication, available at: http://bmgt2-notes.umd.edu:8080/Faculty/KM/papers.nsf

Ballin, M. G., J. M. Hoekstra, D. J. Wing and G. W. Lohr (2002), "NASA Langley and NLR Research of Distributed Air/Ground Traffic Management", American Institute of Aeronautics and Astronautics Paper AIAA-2002-5826, Los Angles Calif. Oct. 2002.

Commission on the Future of the US Aerospace Industry (2002), "Final Report of the Commision", USGPO, Washington DC.

Cooper, M. N. (2000), "The Proposed United Airlines-US Airways Merger", Testimony before the Antitrust Committee, Unites States Senate, June 14, 2000.

Donohue, G. L. and R. Shaver (2000), "United States Air Transportation Capacity: Limits to Growth Part I (Modeling) and Part II (Policy)" National Research Council Transportation Research Board 79th Annual Mtg., Washington DC, Papers Numbered 00-0582 and 00-0583.

Donohue, G. L. (2002) The U. S. Air Transportation System: A Bold Vision for Change: A White Paper prepared fro the Commission on the Future of the U.S. Aerospace Industry, George Mason University

DOT/BTS (2001) U.S. Department of Transportation, Bureau of Transportation Statistics web site (http://www.bts.gov/)

DotEcon Ltd, 2001, "Auctioning Airport Slots, A Report for H M Treasury and the Department of the Environment, Transport and the Regions", (http://www.dotecon.com).

Erzberger, H. (1995) "Design Principles and Algorithms for Automated Air Traffic Management", *AGARD Lecture Series 200 Presentation*; Madrid, Spain; Paris, France; Moffett Field, CA; available at http://ctas.arc.nasa.gov

Federal Aviation Administration (2001), Airport Capacity Benchmark Report (http://www.1faa.gov/events/benchmarks/)

Federal Register, 2000, Department of Transportation, Federal Aviation Administration, "High Density Airports; Notice of Lottery of Slot Exemptions at LaGuardia Airport", Docket Number [FAA2000-8278]vol. 65, no. 221.

Federal Register, 2001, Department of Transportation, Federal Aviation Administration, "Notice of Alternative Policy Options for Managing Capacity at LaGuardia Airport and Proposed Extension of the Lottery Allocation", Docket Numbers [FAA2001-9852], [FAA2001-9854] vol. 66, no. 113.

Fan, T. P. and A. R. Odoni (2002), "A Practical Perspective on Airport Demand Management", *Air Traffic Control Quarterly*, 10. 3, 285-306.

Gleimer, E. (1996), "Slot regulation at high density airports: how did we get here and where are we going?", *Journal of Air Law and Commerce*, 10, 4, May-June, 877-931.

Goetz, A. R.. (2002), "Deregulation, competition and antitrust implications in the US airline industry", *Journal of Transport Geography*, 10, 1-19.

Haynie, R. C. (2002), "An Investigation of Capacity and Safety in Near-Terminal Airspace for Guiding Information Technology Adoption", Ph.D. Dissertation, George Mason University.

Hansen, M. (2002), "Micro-level analysis of airport delay externalities using deterministic queuing models: a case study", *Journal of Air Transport Management*, 8, 73-87.

IATA, 2000*, Worldwide Scheduling Guidelines, 3rd Edition*, International Air Transport Association.

Mineta, N. Y., (1997) *Avoiding Aviation Gridlock & Reducing the Accident Rate: A Consensus for Change*, National Civil Aviation Review Commission, (http://www.faa.gov/narc/)

Preston, E. (1987), *Troubled Passage: The Federal Aviation Administration during the Nixon-Ford Term 1973-1977*, US DOT USGPO book

Rassenti, S.J. ,V. L. Smith and R. I. Bulfin "A Combinatorial Auction Mechanism for Airport Time Slot Allocation" *Journal of Economics* V13 pp 402-417

Rochester, S. I. (1976*), Takeoff at Mid-Century: Federal Civil Aviation Policy in the Eisenhower Years 1953-1961*, US DOT USGPO book

The Royal Society (1992), *Risk: Analysis, Perception and Management,* Report of a Royal Society Study Group, London. Fifth Printing 2002.

Vossen, T. (2002) *Fair Allocation Methods in Air Traffic Management*, PhD Dissertation, U. of Maryland.

Vossen, T. and M.O. Ball (2001) "Optimization and Mediated Bartering Models for Ground Delay Programs", submitted for publication, available at: http://bmgt2-notes.umd.edu:8080/Faculty/KM/papers.nsf

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

Welch, J. D. (2002), "Assessing Capacity Needs of Congested Airports", AIAA