Evaluation of Future National Airspace System Architectures

George Hunter,[*] Kris Ramamoorthy,[†] Paul Cobb,[‡] Alex Huang,[§] Matt Blake,[**]
Sensis Corporation, Campbell, CA, 95008

and

Alexander Klein[††]
George Mason University, Fairfax, VA, 22030

The National Airspace System is evolving as both service providers and operators migrate to new products and procedures and user demand continues to grow. Consequently policy makers need decision support tools and performance projections to help them prepare for potential evolution strategies. Here we describe an emerging suite of tools that evaluates hypothetical future National Airspace Systems. These tools model the future user demand, the resulting operations and delays in the National Airspace Systems, and the economic constraints and impacts. We simulate and analyze several example case studies of National Airspace System evolution.

I. Introduction

A number of factors have contributed to a significant increase in air transportation demand in recent decades. With continued competition, cost containment, and air transportation product diversification, this demand is projected to continue to increase in coming years. Planners are now assuming a doubling (2X) and tripling (3X) of demand in the coming decades. The current demand-to-capacity ratio in the National Airspace System (NAS) is rapidly approaching unity. This capacity limit will constrain future growth in the aviation sector. Adjustments to the current system can relieve this constraint to a limited extent. Examples of such adjustments include (i) flight schedule shifting to the few remaining periods of low demand, (ii) air traffic control (ATC) productivity enhancements, and (iii) new procedures and airspace redesigns. It is doubtful, however, that these adjustments can accommodate 2X and 3X scenarios. More significant system transformation is required. Planners are now investigating possible NAS architecture transformation strategies that can support the 3X scenario.

In this paper we use the ACES (Airspace Concepts Evaluation System), AvDemand and AvAnalyst suite of tools to evaluate the future NAS. We briefly describe these tools. Next we present our baseline model that represents the current NAS. We show that today's system is not capable of accommodating the increased demand levels expected in the future. This is also true given the system capacity improvements that are currently planned. We discuss where the congestion bottlenecks are likely to occur, and in particular how surface congestion and airspace congestion are both important.

Next we evaluate a number of candidate architecture transformation strategies. We note that in addition to increasing system capacity, these strategies must also meet a number of other system performance targets in areas such as noise, emissions, safety and security. The strategies may be used in various combinations to achieve the required performance targets. The strategies include new operator business models, new aircraft designs, and new air traffic management procedures and paradigms.

We evaluate two new operator business models. First, we use time-shifting to move flights away from time periods of heavy congestion. This strategy reduces congestion and delay but forces travelers to use potentially less

---

[*] Sr. Scientist, STC, 1700 Dell Avenue, Campbell CA 95008, AIAA Member
[†] System Engineer, STC, 1700 Dell Avenue, Campbell CA 95008, AIAA Member
[‡] Technical Director, STC, 1700 Dell Avenue, Campbell CA 95008, AIAA Member
[§] System Engineer, STC, 1700 Dell Avenue, Campbell CA 95008, AIAA Member
[**] Vice President, STC, 1700 Dell Avenue, Campbell CA 95008, AIAA Member
[††] Research Professor, Systems Engr & Op Research, 4400 University Drive, MSN 2F2, AIAA Member

desirable time slots. Second, we use spatial shifting to move flights from congested hub airports to less congested auxiliary airports within 30 miles of the major hub (Hub Saturation and Local Airport Growth). This strategy preserves the desirable time slots, but forces travelers to use alternate airports. We alleviate this penalty by not allocating connecting traffic to these flights.

We make an initial evaluation of one new aircraft design, the Extremely Short Takeoff and Landing aircraft (ESTOL). ESTOL aircraft promise a number of advantages including reduced noise and operating independently of an airport's major runways. The latter advantage offers reduced terminal area congestion. The disadvantage of ESTOL aircraft is potentially increased costs. Finally, we model increased airspace capacity assuming enhanced decision support tools for air traffic controllers and managers.

Our results so far suggest that there is no single strategy that can solve the 2X demand problem; however, different combinations of strategies are capable of accommodating future increased demand. Each combination of strategies cannot be fully evaluated without understanding the economic constraints and impacts. For instance, the NAS experiences congestion not because it lacks capacity, but because only a fraction of the total capacity is in demand. The demand for air transportation is highly focused on certain locations and time periods. Therefore any proposed strategy for resolving NAS congestion ultimately must account for these economic considerations.

It would be relatively simple, for example, to solve the NAS congestion problem by shifting flights to low-congestion time slots or airports. But such a solution ignores the important economic constraints. As traffic levels escalate in the future, congestion management will come at a cost. We discuss these economic aspects of the problem and our concept of how to integrate economic models with our suite of NAS modeling tools. The goal is a simulation that provides policy decision makers with both NAS performance and economic information for future planning.

Simulating the entire NAS at a reasonable level of detail – let alone the NAS at 2X or 3X demand levels – is a serious computational challenge for the software development and modeling community. A solution may be a medium-fidelity model derived from ACES and enriched with an economic modeling layer. Such a model would be positioned between the simpler queuing network models and the high-fidelity microscopic simulation tools. It would probably include a discrete-event simulation core and a 4D trajectory modeling component. We present high-level requirements and architecture considerations for this type of model, as well as software development strategies. In this context, we also discuss the issue of randomization and probability-based decision making in the model. Our goal is to contain the "parameter explosion" and keep the computation time within sensible limits while at the same time providing a range for each output value so as to increase confidence in the model.

II. NAS Demand Modeling

We use our NAS demand modeling tool, *AvDemand*, to create daily traffic schedules.[1] AvDemand uses a nominal traffic schedule as a starting point, such as that supplied by the FAA's Enhanced Traffic Management System ETMS. From there AvDemand both escalates the traffic level and modifies the schedule according to new business models or NAS transformation strategies.

Traffic can be escalated by either integer or fractional growth multipliers. Flights can be shifted in time to exploit valleys in the demand profile. They can also be shifted spatially, to outlying airports in order to reduce hub congestion. Increased point-to-point (PTP) scheduling can also be used, for instance to exploit less-used regional airports, or to obviate congested hub airports. Fleet mix shifts can be hypothesized, for instance toward larger aircraft to help handle increased demand levels, smaller regional jets, very light jets, and so forth.

A. Background and motivation

Generation of NAS-wide flight schedules and flight plans is currently a resource-intensive activity. The need for air traffic demand sets that represent the future NAS is expected to increase as researchers investigate NAS transformation strategies. Our research fulfills this need by modeling and automating this process.

Currently, generation of realistic and useful future air traffic demand sets for air traffic simulations can take upwards of several months. This effort typically requires subject matter experts who are familiar with the demand generation process (see Figure 1). The intensive resource requirements of demand generation can limit the quality, frequency, and depth of NAS simulation investigations.

Researchers who generate air traffic demand sets usually rely on subsets of historical traffic data to perform their analysis and technology evaluations. Yet any proposed NAS technology requires many years for deployment. By then traffic levels have grown and the scheduling structure may have changed. The future demand level has a first-order, and nonlinear, impact on the resulting NAS delays and the benefits from a given technology improvement.

Therefore, it is crucial for NAS investment decision makers to conduct assessments with future demand projections that are as realistic as possible.

Figure 1. The three-step demand generation process requires significant data and domain expertise.

Since the future is inherently uncertain, it is important for researchers to analyze multiple potential future demand scenarios. Because of the multidimensional nature of the demand problem, sensitivity analysis requires many demand sets. Without an automated demand generator, a significant number of demand sets becomes impossible. For instance, the different levels of traffic escalation and the different NAS transformation strategies would alone produce a great many demand sets.

AvDemand is targeted for use in fast-time and real-time NAS simulators such as the Airspace Concept Evaluation System (ACES) , Future ATM Concept Evaluation Tool (FACET) , Total Airport and Airspace Modeller (TAAM) , the Reorganized ATM Mathematical Simulator (RAMS) +, SIMMOD , AwSIM, and Pseudo Aircraft Systems (PAS).

B. Generating baseline future NAS demand scenarios

We have used AvDemand to generate future NAS demand scenarios using February 19, 2004 as our baseline day. We scaled up the seed day traffic level according to the growth forecast by the FAA in its Terminal Area Forecast (TAF),[2] which is updated and published annually. The TAF forecasts are based on past airport activity levels and on socio-economic and institutional factors, especially for the larger airport forecasts. We projected the baseline seed day to the years 2014 and 2025. Since the TAF only extends out to 2020 we extended the annual growth rates implied in the existing TAF forecasts through 2025.

The TAF provides forecasts for airport specific levels of operations, but it does not provide forecasts of the future frequency of flights between specific airports. To model this, AvDemand uses a standard transportation traffic distribution algorithm, known as the Fratar method. This tool extrapolates future trip distributions based on TAF forecasts of airport specific activity levels. AvDemand then uses an Airport Pair Demand Profiling (APDP) algorithm to determine schedule timing to adhere to the existing business model for scheduling departure and arrivals at each airport. A detailed account of the Fratar algorithm, the APDP, and other underlying future demand generation assumptions can be found in Reference 1.

In these 2014 and 2025 projections we assumed that the NAS fleet mix remains consistent. The NAS-wide aggregate demand escalation in the 2014 baseline scenario is approximately 1.2 times the baseline demand (referred to as “1.2X”). The 2025 demand is approximately 1.4X. We also generated 2X and 3X demand scenarios. Figure 2 lists the key metrics for each demand scenario.

Demand Scenario	2004	2014	2025	2X	3X
Total Ops	31,791	36,733	46,192	60,274	90,395
Available Seats	2,451,323	3,063,086	3,955,444	5,323,272	8,197,756
Available Seat Mile	1,910,429,491	2,467,611,567	3,243,880,242	4,431,223,069	6,939,473,461
Average Stage Length (miles)	601	646	663	679	695
Flight Growth	1.00	1.16	1.45	1.90	2.84
ASM Growth	1.00	1.29	1.70	2.32	3.63
Available Seat Growth	1.00	1.25	1.61	2.17	3.34
Average Airport Growth	1.00	1.18	1.51	2.03	3.18

Figure 2. AvDemand baseline demand metrics..

C. Generating off-baseline future NAS demand scenarios

There are many ways the aircraft operators of the future may choose to operate. For instance, they may shift to fleets of smaller or larger aircraft, operate with more or new hub airports, shift to under-utilized secondary airports near congested hubs, operate direct PTP routing to avoid congested hub airports, shift flight departure times to take advantage of brief periods of airport underutilization, or schedule flights earlier in the morning or later at night. The initial baseline scenarios for future activity described above rely on the assumptions that the distribution of airport utilization and the fleet composition used remain stable in future years or at higher levels of overall activity. This assumption provides a baseline from which to frame other scenarios for future NAS activity.

As examples of alternative scenarios we now discuss the spatially-shifted schedule and a temporally-shifted schedule strategies. In the spatially-shifted schedule we assume a migration toward greater use of secondary airports and somewhat smaller aircraft. This represents an increased use of PTP scheduling in addition to the hub and spoke scheduling strategy that the major airlines have traditionally used.

The principal rationale for considering this business shift scenario is that passenger and aircraft demand at the 34 major CONUS hub airports, as defined in the FAA Operational Evolution Plan (OEP),[3] may soon exceed the capacity of those airports. Therefore, passenger flights at neighboring regional airports becomes more attractive and feasible. This may especially be true during high traffic hours of the day. During these hours, markets with high frequency flights are identified as candidates for diversion or reassignment to other regional airports. This strategy may be attractive because a significant fraction of traffic at most hubs is actually originating or terminating at the hub, and therefore may potentially choose to use services to a regional airport if available.

In this particular AvDemand exercise, to be considered as a regional airport that might receive reassigned flights from over-utilized OEP airports, an airport must be a public use airport and within 30 miles of the OEP airport. It must have at least one runway that is in good condition and at least 5,000 feet in length. Furthermore, flights that are shifted to regional airports as part of this business shift alternative must have a stage length of less than 1,000 miles (in great circle distance). The alternative airport must also have arrival or departure slots available within the same time window as the diverted flight.

We use the smaller regional jet aircraft for the flights that service the alternative regional airports. Therefore, diverted flights that originally used 100 seat or larger aircraft will have a change of aircraft type. The diverted flight will be changed to a regional jet with seat capacity of 100, and additional flights will be added to ensure that the total seats provided in individual markets is retained. Therefore this scenario has more flights than the baseline scenario, and the average seat size of aircraft is lower.

Although this scenario results in higher numbers of NAS aircraft operations, it does not necessarily produce higher numbers of passenger miles flown or available seat miles. Differences in seat miles flown may occur because of changes in the relationships between direct trips between origins and destinations and connecting itineraries. Finally, we do not assume changes in the number of passenger trips, since passenger demand for trips is mostly driven by socio-economic factors, including passenger yields.

This scenario results in a decreased load at the 34 OEP airports in the 2014, 2025, 2X and 3X scenarios. Accomplishing this required one percent more flights in the 2014 scenario, three percent more flights in the 2025 scenario and 12 percent more flights in the 2X scenario.

In the temporally-shifted schedule we shift flight departure times to take advantage of temporally-proximate periods of airport underutilization. Flight peaking activities can easily exceed the airport capacity for a given time period. On the other hand, there are also times when the airport has runway slots available for departure and arrival. To take advantage of the available capacity, as airlines may do in the future, we use an algorithm to switch flights from excessive demand time windows to the nearby time windows with available runway slots. We refer to this process as “timeshifting.” Starting from the most congested airport, the timeline at each airport is divided into 15-

minute intervals. The algorithm starts at the beginning of the day and proceeds through each time period until the end of the day.

This timeshifting examines flights in 15-minute time windows and determines the airport departure and arrival throughput using the Pareto capacity envelope, which is shown in Figure 3. During a 15-minute time window, if there are more departure flights than arrival flights, then the departure rate is set to the 15-minute maximum departure rate, or the departure capacity. This hourly departure capacity is divided by four. The arrival rate at this 15-minute interval is set to the remaining capacity, which is the difference between the total capacity and the departure capacity. The same logic holds when the arrival flights dominate the traffic. When the number of departure flights equals to the arrival flights, both departure and arrival rate are set to half of the total capacity.

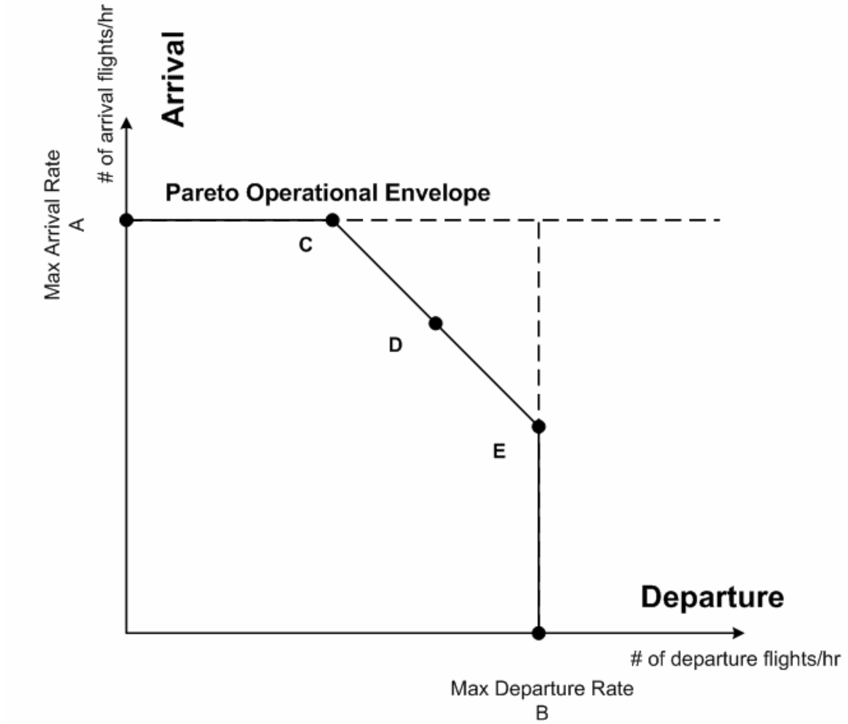

Figure 3. Illustration of airport capacity pareto envelope..

Once the rates for both arrival and departure are determined in the 15-minute time window, the amount of excessive departure/arrival flights at this time window can be determined. If there is any excessive demand, the algorithm searches for unused capacity in the hour before and hour after the 15-minute time window (eight 15-minute time windows in all, four before and four after). A flight can be re-assigned only if both the departure and arrival airports have runway slots available at the time to which the flight is intended to be re-assigned. In addition, when deciding which flights to shift to the other time window, flights from origin-destination airport pairs with a low daily flight frequency are given priority. The reason for preferring to move flights with low daily airport-pair flight frequency is that we want to minimize our timeshifting algorithm impact on the existing airline banking structure and the lower the airport-pair flight frequency, the lower the probability that the flight is involved in a schedule bank. Starting from the beginning of the day of the most-congested airport, the algorithm processes through the whole airports in the dataset.

Figure 4 compares the runway departure distribution based on augmented flight demand at Atlanta Hartsfield International Airport (ATL). Assuming that the traffic from ATL increases two-fold, the departures at ATL increase from 1,208 operations per day to 2,416 operations. The red curve represents the non-airport capacity constrained departure distribution and the blue curve shows the airport capacity constrained departure distribution. The dashed line indicates the departure throughput for the 15-minute time windows when VFR departure traffic dominates the arrival traffic.

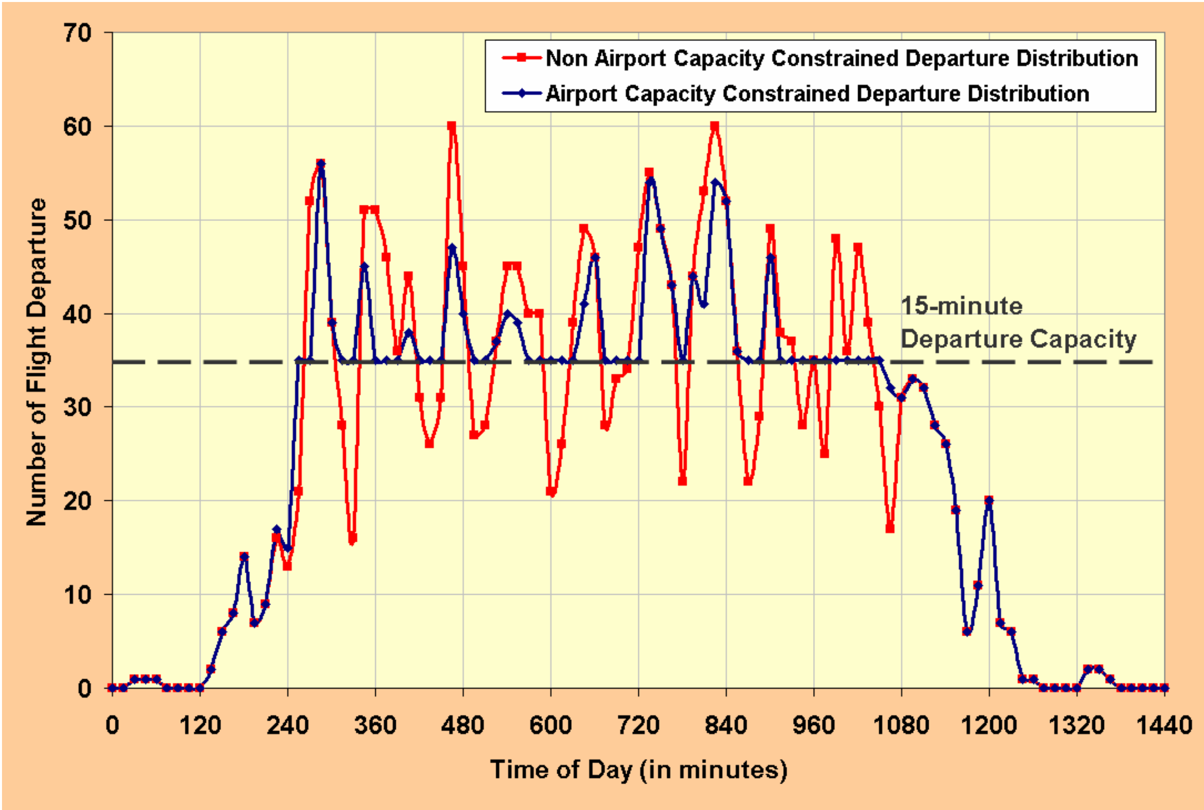

Figure 4. Airport Capacity Constrained versus Non-Airport Capacity Constrained Flight Departure Distribution by Time-of-Day using the timeshift algorithm..

The distribution logic, used in creating the airport capacity constrained departure distribution curve, takes advantage of time windows where departure slots are available. It shifts flights from congested time windows to those with excess departure slots. The non airport capacity constrained departure distribution curve, on the other hand, simply reflects the original demand. In this case, excess departure slots are not being used. As Figure 5 illustrates, we use both time-shifting and/or spatial-shifting (i.e., the PTP flight demand) logic independently or together.

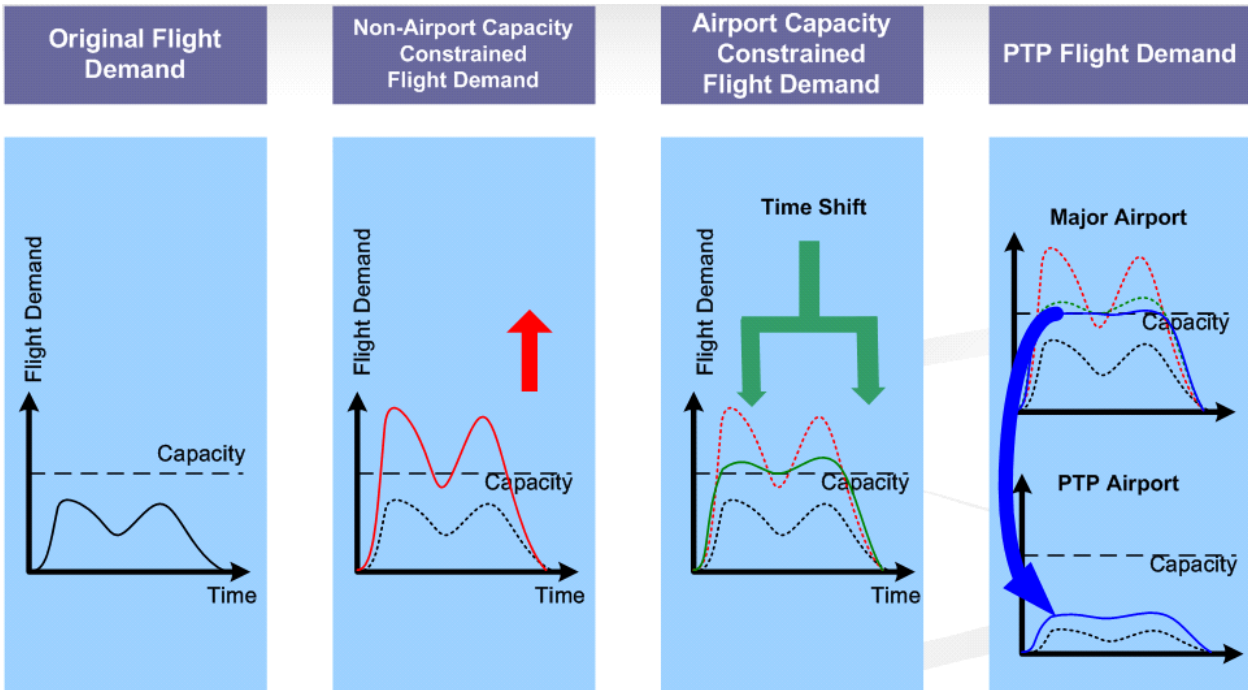

Figure 5. Flight demand types and combinations. The baseline demand profile for an airport is escalated according to projections of future growth. This may result in a congested airport with demand exceeding capacity. The timeshifting or regional airport strategies may then be used independently or together.

III. Airspace Concepts Evaluation System (ACES)

This section describes the ACES simulation tool that uses the AvDemand demand scenarios as input. ACES is a fast-time, computer simulation of local, regional and nationwide factors covering aircraft flight from gate departure to arrival.[4,5] ACES provides a flexible NAS simulation and modeling environment that can assess the impact of new NAS tools, concepts, and architectures, including those that represent a significant departure from the existing NAS operational paradigm. To meet this objective, ACES utilizes the High Level Architecture (HLA) and an agent-based modeling paradigm to create the large scale, distributed simulation framework necessary to support NAS-wide simulations. HLA is a set of processes, tools and middleware software, developed by the Department of Defense, to support "plug-and-play" assembly of independently developed simulations. For ease in integration and efficient runtime execution of the simulation, the ACES simulation framework employs agent-based modeling. The ACES agent-based processes simulate airspace and aircraft operations.

The ACES architecture is designed to accommodate models of each operational component of the NAS. ACES contains models for Air Traffic Management (ATM), encompassing Air Traffic Control (ATC) and Traffic Flow Management (TFM) operations; aircraft dynamics; and en route winds. The modeling accounts for airspace and airport designs and procedures, including airport visual flight rules (VFR) and instrument flight rules (IFR). Agents represent present-day NAS operations and include Air Traffic Control System Command Center (ATCSCC), the en route Air Route Traffic Control Center (ARTCC), Terminal Radar Approach Control (TRACON), Airport Traffic Control Tower (ATCT), and aircraft and pilot entities.

The agents are autonomous entities, each simulating operations according to algorithmic and data processing logic defined by its model based on information exchanged with other agents and supporting constructs. In the ACES modeling concept, the unifying factor is the aircraft trajectory. Each of the basic NAS agents performs modeling functions that operate on flight trajectories. AOC agents conduct pre-takeoff flight planning to define four-dimensional (4-D) desired/requested trajectories, post-takeoff flight following to coordinate trajectory flight plan revisions, and flight schedule revisions based on delay, diversion and cancellation assessments. The TFM agents conduct local, regional and nationwide flow constraint assessments and determinations based on flight plan, traffic surveillance, meteorological and airspace and airport status and constraint information. The ATC agents conduct trajectory intervention assessments and resolutions based on traffic situation surveillance, procedures, separation rules and TFM constraints, trajectory state and intent, meteorological and aircraft performance information. The aircraft trajectory simulations are based on trajectory state estimate and intent, planned trajectory, ATC clearance, aircraft performance and meteorological information.

Each agent has information describing a trajectory and performs action on the trajectory based on this information. But each agent does not necessarily have the same information as another, and none may know the true trajectory state of a flight. The ACES design enables each agent to maintain separate or different trajectory data and trajectory management logic, hence implementing a multi-trajectory modeling concept. ACES also maintains a model of trajectory truth for each flight. Hence ACES would have the ability to model effects of trajectory estimation errors concurrently with the modeling of agent operational processes. The degree of accuracy and fidelity with which the agents operate on trajectories depends on the technologies and functional capabilities of the NAS operational concept being simulated. In ACES these are represented by the logic encoded into each agent and associated modeling parameters.

Independent validation of ACES was performed using FAA Aviation System Performance Metrics (ASPM) data. Archived ASPM data contain a variety of NAS performance metrics useful in verifying and validating NAS-wide simulations. The V&V analyses demonstrate that ACES provides a good approximation of today's NAS.[6]

IV. NAS Transformation Strategies

This section establishes the need for NAS transformation. The forecasted traffic growth will cause significant congestion in all phases of air transportation, including the airspace, airport surface, and landside phases. Therefore congestion alleviation strategies are required in all phases. Furthermore, multiple strategies are likely to be required even for a given phase. This section presents initial results for several candidate airport surface strategies. While no single strategy appears to be capable of solving the predicted congestion, multiple strategies in combination probably can solve this problem.

A. NAS baseline performance

Plans are in place to increase the capacity of the NAS using OEP investments.[3] These capacity improvements will help the NAS accommodate higher levels of demand. These improvements, however, do not appear sufficient to accommodate the 2X and 3X levels of traffic escalation that are predicted.

These demand scenarios quickly outstrip current and anticipated NAS capacities, resulting in unacceptable levels of delay or flight cancellation. ACES simulations show that at higher levels of demand, system delays quickly rise over the course of the simulated day to untenable levels. Simply put, it is not possible to operate a scheduled air transportation system in such a congested and unreliable environment, nor are passengers likely to be willing to make use of such a system

Figure 6 shows the ACES Build 2 growth of flight delay by hour of day for a range of traffic levels leading up to the 2X scenario. Delay over the course of the 2004 seed day is modest and relatively stable; average delay per flight is just under three minutes per flight. As the figure illustrates, even with anticipated NAS capacity improvements from OEP investments in place, delay grows to unacceptable levels after a 20 percent increase in demand (to 1.2X). While average delay with 1.2X demand seems manageable at 11.2 minutes per operation, the impact of congestion is illustrated in the growth of aircraft delay over the course of the day. Flights become increasingly delayed as the day goes on, and some become so delayed that they cannot be completed by the end of the day. This situation worsens with demand at 1.4X, when average delay grows to 38.2 minutes and the end of day delays increasingly aggravated. At a demand level of 2X, average delay is nearly 100 minutes, and end of day delays are even worse.

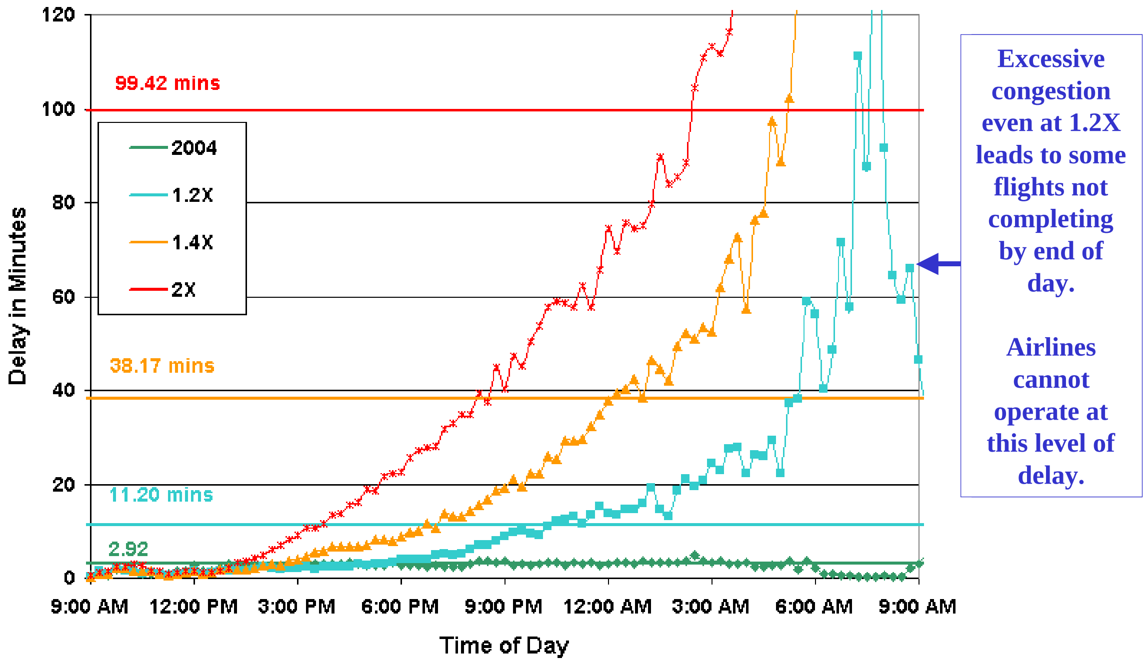

Figure 6. Hour-by-Hour Delay Profiles for Baseline Demand Scenarios. These simulation runs were all under good weather conditions

B. Temporally- and spatially-shifted schedules

In Section II above we discussed methods of shifting traffic schedules, both temporally and spatially. We now use these methods in our case study of the Chicago metropolitan area air traffic congestion problem. Figure 7 illustrates the number of flights at the Chicago area airports in the 1X, 2X and 3X demand scenarios, for the baseline, and the temporally- and spatially-shifted schedule scenarios.

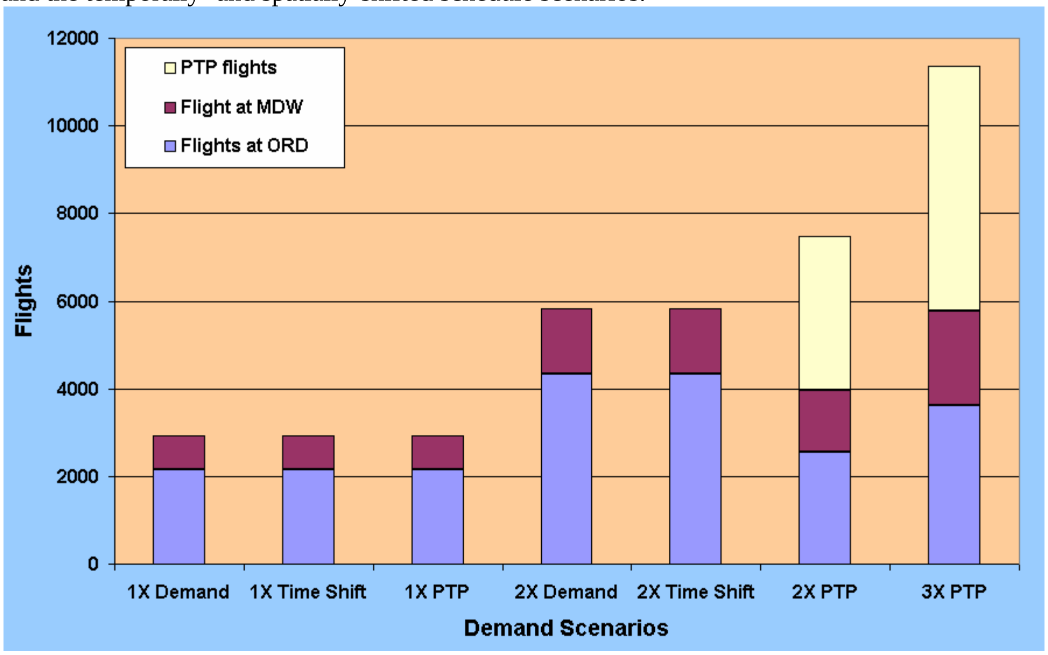

Figure 7. Demand Scenario Number of Flights Comparison..

Figure 7 shows that in the PTP scenarios a significant fraction of the flights are reassigned from ORD and MDW to regional airports in the area (the majority of these shifts are from ORD). The result is that the demand at the main hub, ORD, is relieved in both the 2X and 3X scenarios, although it remains greater than in today's baseline scenario. The total number of daily Chicago area flights in the 2X PTP scenario, for example, is about 7,500.

Figure 8 illustrates the predicted mean flight delay for these different scenarios (more data points are required to refine these curve fits). It shows that the 2X PTP scenario results in a mean flight delay of about 17 minutes across the entire day. In contrast, to preserve this level of delay in the baseline with timeshift scenario, the flight count would have to be reduced to 4,400.

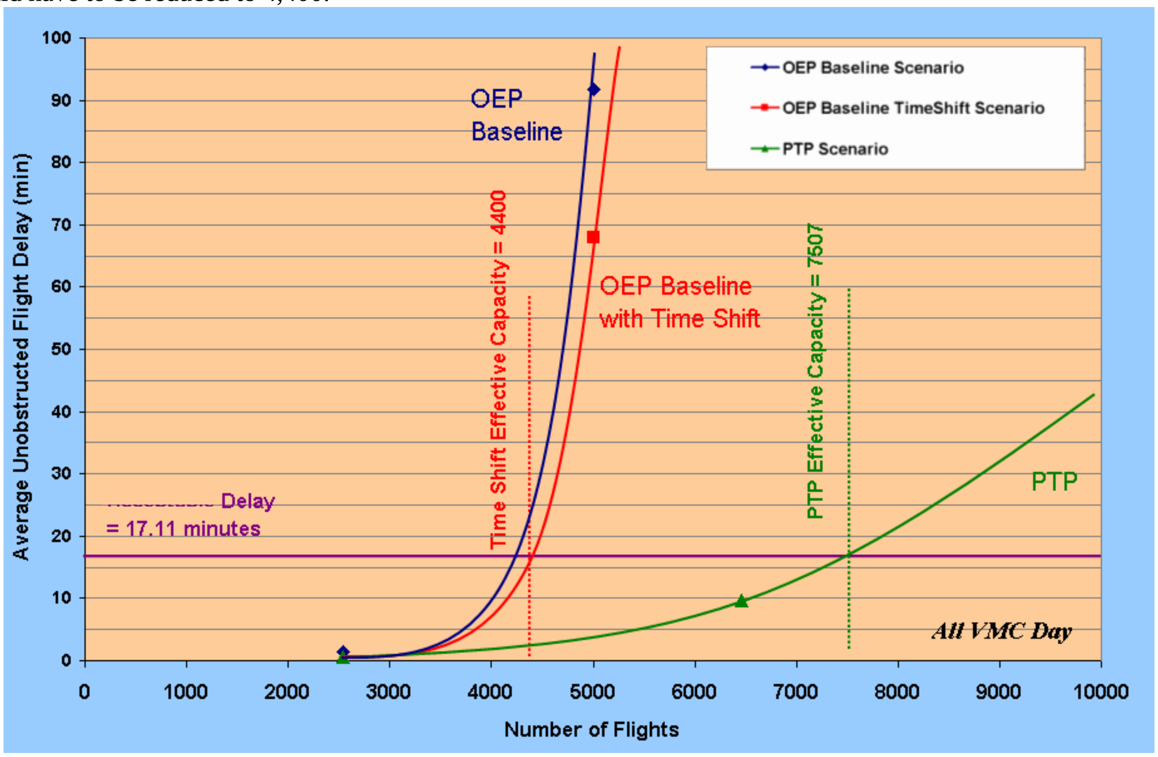

Figure 8. Effective capacity estimation assuming exponential demand-delay relationship. These effective capacity curves are notional. Additional data points are required to refine the curve fits.

The Figure 8 results do not account for airspace congestion and its associated delay. To account for this, these results are extrapolated NAS-wide using ACES Build 2 in Figure 9. Here we compare three different strategy combinations: timeshifting alone (red), PTP scheduling using regional airports (green), and PTP scheduling with increased airspace capacity (purple). In the increased airspace capacity case we assume the en route airspace capacity is tripled (3X airspace). These results illustrate the importance of the airspace bottleneck in the NAS-wide congestion problem. Without increasing the airspace capacity, the mean flight delay is untenable, even with substantial surface capacity increase. It is also true that increasing the airspace capacity alone results in untenable flight delay (results not shown). Therefore, both surface and airspace domains are predicted to have congestion bottlenecks.

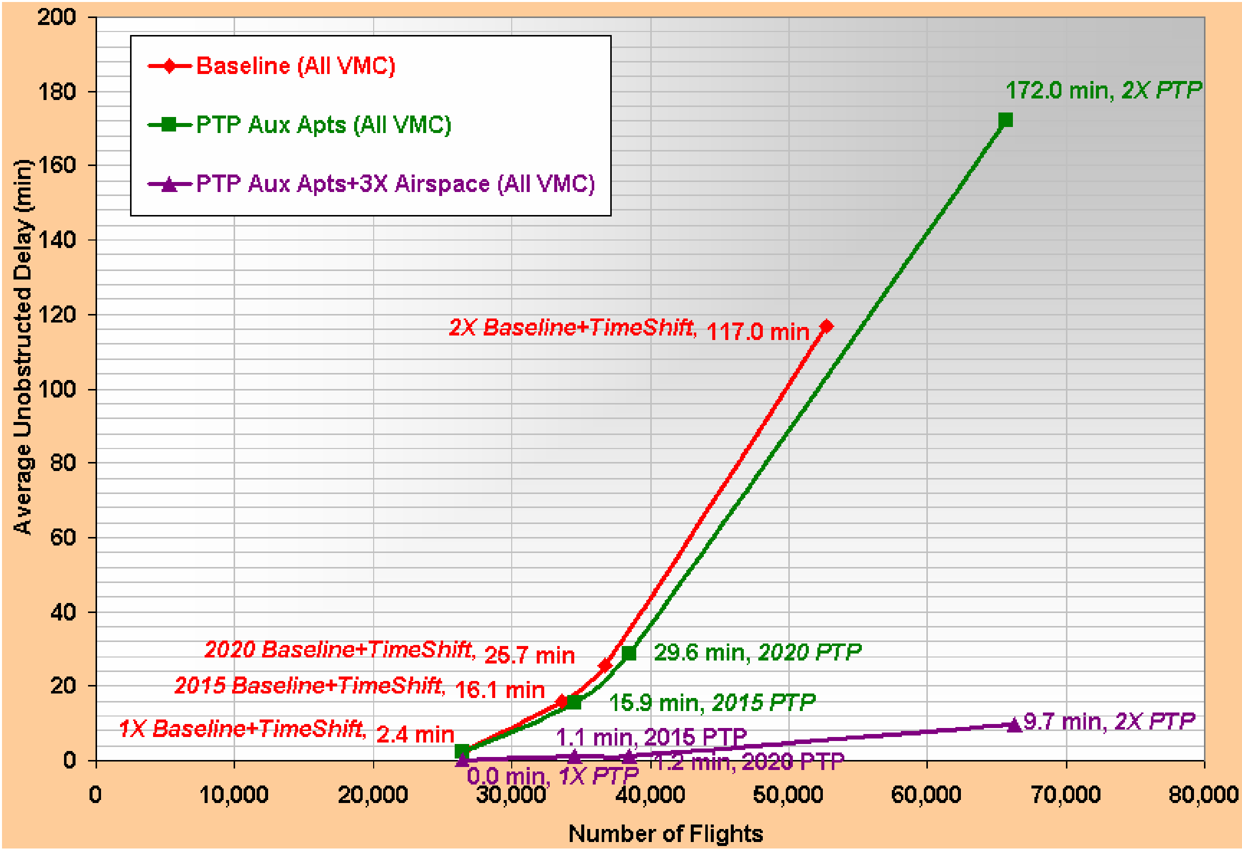

Figure 9. NAS-wide mean flight delay for entire day for three different capacity increasing strategies: Timeshifting, PTP auxiliary airports, and PTP auxiliary airports plus En Route ATM Automation. All VMC conditions using ACES Build 2.0.3.

C. Extremely Short Takeoff and Landing aircraft

Extremely Short Takeoff and Landing (ESTOL) aircraft may help the future NAS reduce both congestion and environmental impact. Congestion is reduced by moving the ESTOL operations off the main runway. By diverting a fraction of the traffic to shorter side runways, the throughput bottleneck effect at the main runway is reduced. Environmental impact is reduced by the ESTOL's reduced noise levels and fuel consumption.

The current state of the art in ESTOL air transport performance is approximately a 120 kt takeoff and landing speed with a 4,000 ft runway requirement. This is a 90 seat vehicle with a 1,000 nmi range and Mach 0.6 cruise speed. The ESTOL research and development community has set target performance goals to be reached within the next two decades. These include halving the takeoff and landing speed and the runway length, reducing the noise, and increasing the cruise speed and range to Mach 0.8 and 1,400 nmi, respectively.[7]

To meet these performance goals several technical problems are under investigation. A key challenge is to integrate, in one vehicle, (i) the powered lift capability which enables short takeoff and landing, and (ii) efficient cruise capability. Other challenges include: reduce aerodynamic drag, and delay the transonic drag rise to higher speeds; manage the noise effects of the ESTOL lift system while maintaining performance; reduce specific fuel consumption, and maintain acceptable handling qualities while improving performance.

Whether or not ESTOL aircraft are able to achieve substantial market penetration in future years is an open question. Assuming that they do we use AvDemand to make a preliminary investigation of the congestion reduction at the 2X traffic level. Figure 10 breaks down a 2X demand scenario, showing the flights for a day in the NAS in payload-range bins. The data reveal several distinct clusters of flights, including (i) 101 – 150 seat aircraft flying 100 – 1,500 nmi flights, (ii) 30 – 50 seat aircraft flying 100 – 750 nmi flights, and (iii) 1 – 10 seat aircraft flying 100 – 500 nmi flights. The green shaded box illustrates the potential ESTOL performance envelope.

Flt. Length \ Seats	0	1-10	11-20	21-30	31-40	41-50	51-60	61-70	71-80	81-90	91-100	101-150	151-200	201-250	251-300	301-350	>351	Total	%Total
0	1	4	1	0	0	0	0	0	0	0	0	0	1	0	0	0	0	7	0.01
1-100	99	310	57	119	198	170	0	30	0	27	0	12	16	0	0	0	0	1038	1.72
101-200	601	2047	557	467	1325	1498	0	253	6	77	0	1845	221	2	2	0	7	8908	14.78
201-300	627	1285	333	351	954	1885	1	236	0	31	82	3542	206	35	3	0	13	9584	15.90
301-400	580	709	80	93	612	1902	0	151	0	33	6	2880	226	149	5	9	0	7435	12.34
401-500	449	501	91	18	219	1662	0	139	0	0	18	1991	265	84	0	1	2	5440	9.03
501-600	391	433	48	0	148	1283	0	143	0	43	107	2820	398	113	2	0	9	5938	9.85
601-700	304	288	37	0	55	830	0	99	0	0	56	1602	288	47	9	8	18	3641	6.04
701-800	286	248	35	1	81	647	0	113	0	11	0	1662	404	52	0	26	13	3579	5.94
801-900	350	270	55	0	39	287	2	54	0	0	31	2006	621	84	5	10	7	3821	6.34
901-1000	197	210	34	0	6	58	0	7	0	0	14	1309	399	18	5	5	2	2264	3.76
1001-1200	206	168	52	0	0	18	0	64	0	0	0	1322	581	72	2	5	5	2495	4.14
1201-1400	174	77	29	0	2	19	0	11	0	0	0	1085	433	109	0	24	0	1963	3.26
1401-1600	186	37	22	0	0	5	0	0	0	0	0	707	461	89	5	20	1	1533	2.54
1601-1800	143	23	13	0	0	5	0	0	0	0	0	268	342	64	5	4	0	867	1.44
1801-2000	30	10	24	0	0	5	0	0	0	0	0	280	326	40	0	12	0	727	1.21
>2001	56	31	32	0	0	11	0	0	0	0	0	338	490	69	0	4	2	1033	1.71
Total	4680	6651	1500	1049	3639	10285	3	1300	6	222	314	23669	5678	1027	43	128	79	60273	
% Total	7.76	11.03	2.49	1.74	6.04	17.06	0.00	2.16	0.01	0.37	0.52	39.27	9.42	1.70	0.07	0.21	0.13		100

Figure 10. AvDemand 2X traffic scenario divided into payload-range bins. The green shaded box illustrates the potential ESTOL performance envelope of less than 100 seat aircraft and 1,400 nmi range.

ESTOL aircraft can effectively increase capacity at major airports by moving those operations off the main runway and onto shorter runways that otherwise are underutilized. A key question is: What clusters in the payload-range space will ESTOL target for market penetration? To be economically viable an aircraft must be used at reasonably high load factors, so that for example, a 50 passenger market would not be served by a 100-seat aircraft.

Charts such as Figure 10 help to identify economically viable design points for ESTOL aircraft. It indicates that in today's schedule structure, a 90 seat vehicle will find less utility than a 50 seat vehicle where there is a much larger demand cluster. Figure 11 illustrates this for the Chicago O'Hare airport (ORD). In this analysis, ORD demand levels are computed at 15-minute intervals throughout the day where ESTOL aircraft are assumed to be exclusively used for a range of aircraft sizes (seating capacity). In each case, the ESTOL aircraft are assumed to be independent of the main runway operations, thus reducing demand for the main runways. The results illustrate that the 41 – 60 seat case reduces demand the most.

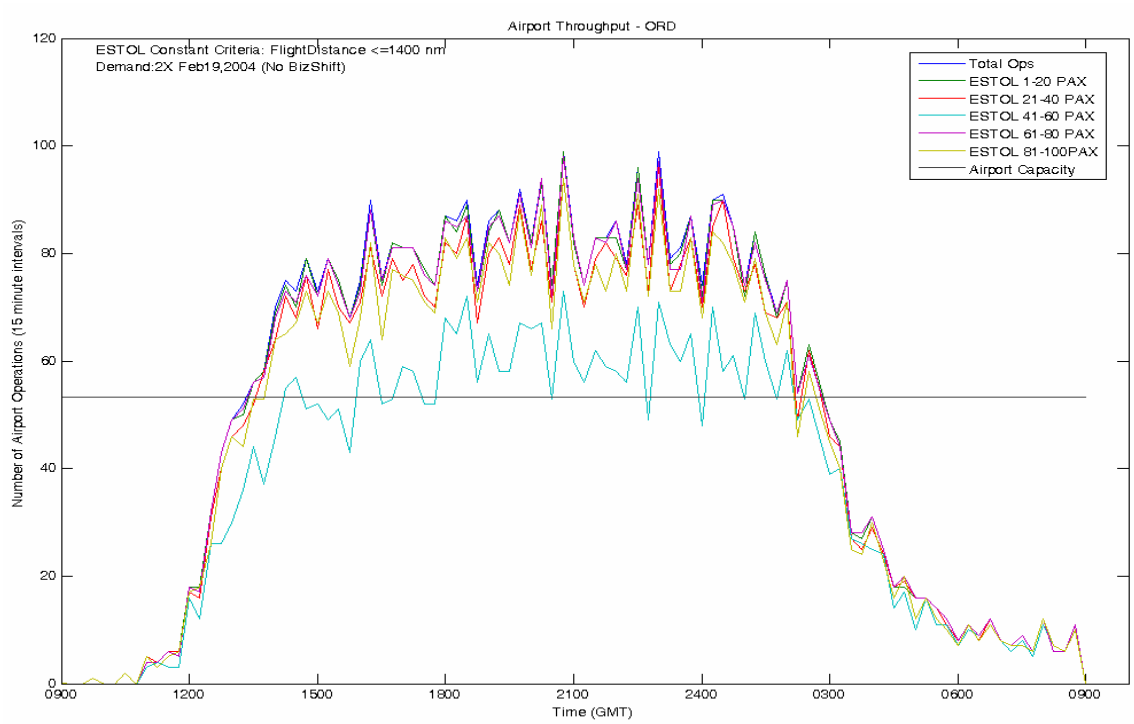

Figure 11. Demand time histories for Chicago O'Hare airport (ORD) main runways. Demand computed at 15-minute intervals throughout a day in the NAS with 2X traffic level. Demand is reduced most dramatically when 41 – 60 seat aircraft are moved off the main runways.

As with the time-shifting and the PTP strategy examined above, the introduction of ESTOL aircraft will not likely, by itself, solve future congestion. Combinations of different strategies, however, do have the potential to solve future congestion. Figure 12 illustrates how the NAS-wide delay is likely to be tenable with the combined strategies of PTP and ESTOL aircraft for surface capacity enhancement.

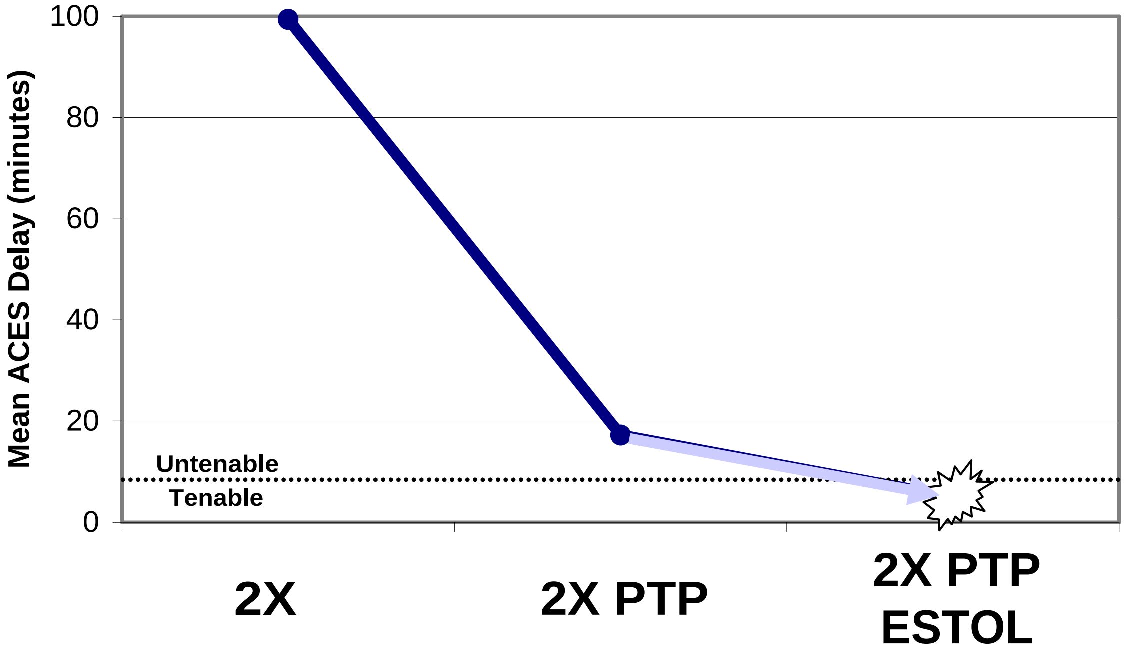

Figure 12. Illustration of NAS-wide delay reduction with PTP and PTP plus ESTOL strategies to enhance surface capacity. Other system capacity constraints, assumed to be solved..

In this section we have analyzed a small set of candidate strategies for accommodating future demand growth. There are several other strategies that are possible as well. The predicted 2X and 3X traffic levels create throughput bottlenecks in the en route airspace, terminal airspace, airport surface, and landside domains. Therefore there is no single strategy that alone accommodates these escalated traffic levels. Furthermore, we do not yet find any single strategy that accommodates these escalated traffic levels on the airport surface. Multiple strategies, however, do appear capable of solving the surface congestion problem. In the future we plan to analyze other strategies such as the use of reduced wake vortex separation minimums on final approach when and where appropriate, and the addition of closely spaced parallel runways using precision approach techniques.

In all phases, the details of how these strategies will work, either alone or in combination, ultimately depends on economic and policy factors as well as engineering factors. For instance we need to evaluate the assumptions imbedded in each of our hypothesized strategies in light of the relevant policy and economic factors. Are there limitations we have not accounted for? What fraction of market penetration can ESTOL aircraft reasonably be predicted to achieve? How willing will passengers be to fly at off hours or outlying airports? Ultimately, a policy and economic analysis must be considered to better prediction what strategies are feasible, and to what extent, in the future NAS.

V. Economic Modeling

Traditionally, “pure” ATM simulation models generate a series of metrics such as capacity, throughput, miles flown, and delays. Output from these models can be handed over to specialized economic analysis tools for post-processing. Such tools often use simplified airline modeling, with all airlines treated as a single group. Such economic modeling does not include important details such as: low-cost vs. legacy carriers, general and business aviation, the FAA cost/revenue on a daily basis, security-related investments, and other factors. NAS-scale macroeconomic analysis tools are powerful but do not model these details. For example, what is the economic impact of wider RJ use or, conversely, of increased use of larger aircraft at airports where congestion pricing may make RJs less viable? What will be the NAS-wide implications of the planned $6-10B Chicago O’Hare airport redesign? If the use of micro-jets becomes wide-spread, how will this affect airline economics, FAA’s costs and revenues, and state of transportation in selected US regions?

A medium-fidelity model with an “economic layer” and ability to address this type of questions fills the existing gap in NAS-scale simulation model line-up. High-level requirements for such a model include the following.

- Sufficient fidelity to "feel" the impact of significant enough – but not necessarily global – factors such as an extra runway added at a major airport; rapid growth of traffic at secondary/small airports in a specific region; changes in aircraft type mix in a specific region or market segment; difference in speed between mainline jets and micro-jets; or benefits from expedited introduction of a NAS-wide datalink capability.
- A degree of randomization and "fuzziness" so as to provide output as a range of values. Low-fidelity models require a large number of random iterations to compensate for lack of detail and, since they are usually very fast, this is feasible. High-fidelity models have a rich "fabric" of events which reduces the need for massively stochastic analyses. For a medium-fidelity model, a balance needs to be struck between desired (randomized) output, number of core scenarios to be studied, complexity of setup, and computation time. A possible approach may be to combine randomizing a very limited number of factors in a small series of simulation runs with probability-based decision making. High growth rate of uncertainty in, say, sector traffic load might cause the model to ignore this factor in decisions beyond a 2-3 hour time horizon.
- This will be a tactical (day-of-operation), not strategic (multi-year outlook) model. By simulating a variety of typical days of operation we will, however, be able to assess longer-term economic and ATM impact of the factors being analyzed.
- Dynamic visualization. At this level of fidelity, dynamic visualization becomes important. Therefore, the model should probably be a discrete-event fast-time simulator, coarser than microscopic simulators but finer-grain than the aggregate queuing models. This implies reasonably accurate 4D trajectory modeling, at least down to "inflection point" level (e.g., commence/complete a turn at waypoint, commence/complete climb or descent to cleared altitude etc). At closer range, each aircraft can be visualized as a separate object; at a NAS-level view, the model can be limited to showing flows.
- Ability to simulate impacts of inclement weather and the corresponding TFM and airline responses (including cancellations), for both severe weather en-route and local weather at major airports.
- Presence of the ATCSCC and individual facilities (Centers), as distinct objects; for selected portions of airspace it may be necessary to go down to ATC Sector level. Also, the model should include AOCs as objects.

In addition to modeling air traffic in a classic sense – i.e., movement of airplanes, factors like airspace and airport capacity, throughput, and delays – the model's economic layer would provide the ability to simulate economic and regulatory factors affecting daily operations and quantify their impact in direct simulation output, not just through post-processing.

- Operational economics would include airlines (pricing strategies, revenues, fuel costs, crew costs, maintenance, other operating costs, disruption management), the FAA (revenues – different models, operations, staff, facilities maintenance etc), airports (revenues from landing fees, congestion management, slots), and passenger perspective (pricing, schedule reliability / predictability, cost of delays, lost connections and cancellations);
- Shorter-term tactical aspects would be reflected in the model (competitive pricing and schedule adjustments, fuel hedging, schedule development, temporary surcharges etc). Additionally, the economic layer would include longer-term investment strategy (new airport infrastructure, airline fleets, aircraft equipage, facilities), political factors (labor issues, environmental impact), safety, security and other government regulations (e.g., cost of additional baggage screening).

Given the massive traffic loads and the need to run multiple sets of simulation scenarios, finding the right balance between model fidelity and performance will be a key challenge.

In terms of architecture, multi-threading should be part of the design but the requirement to run multiple instances of the model on separate processors or computers (in an HLA environment) is much less certain. HLA was initially designed for distributed real-time simulations with a limited number of players / objects. Running super-massive traffic scenarios in fast time may not be something HLA can support. The amount of overhead due to data shunting between HLA federates may well nullify any potential gains from splitting the simulation into a network of interconnected objects (Centers, airports etc).

Similar caution must be exercised regarding the ability to handle external plug-ins or interpreted scripts. In order to maximize performance under heavy object load, the simulation module's interactions with external modules ought to be kept to a minimum. On the other hand, using plug-in components in binary mode, as libraries, is entirely feasible.

In order to satisfy the above requirements, a spiral development approach could be proposed. An initial version of the model would not include randomization but would feature a small set of probability-based decision making simulation mechanisms. The initial economic layer could include the notion of airline pricing, load factors and yield; cost of delays, excess flying distance and cancellations; and FAA revenue calculation. Stress-tests would be conducted from early on in the process to ensure adequate performance for NAS-scale scenarios. The economic layer would be gradually expanded at subsequent levels of the development spiral.

VI. Conclusion

This paper presents preliminary results of the performance of the future NAS. The results indicate the type, depth and scope of performance issues that are likely to arise due to increased demand levels. The results also suggest NAS architecture strategies that may improve future performance and accommodate escalated traffic levels. These results need to be refined with more detailed investigations. In addition to investigating more NAS strategies, we recommend detailed economic modeling in the investigation. We also need to evaluate the NAS in a variety of futures in addition to the baseline future used in this research.

Acknowledgments

The ACES simulation system was developed for the NASA Virtual Airspace Modeling and Simulation Project under the Air Traffic Management System Development and Integration Contract. The authors would like to acknowledge Harry Swenson and Dr. Sherry Borener for helpful inputs and comments in performing this work.

References

1. Huang, A. S., Schleicher, D. and Hunter C.G. "Future Flight Demand Generation Tool," *AIAA Aviation Technology, Integration and Operations (ATIO) Forum*, Chicago, IL (September 2004).

2. Anon., "Terminal Area Forecasts," FAA, Washington, DC, various years. http://www.apo.data.faa.gov/faatafall.htm

3. Anon, "National Airspace System Operational Evolution Plan, 2004-2014," Version 6.0, FAA, Washington, DC, 2004. http://www.faa.gov/programs/oep/

4. Couluris G. J., Hunter, C. G., Blake, M., Roth, K., Sweet, D., Stassart, P., Phillips, J., Huang, A., "National Airspace System Simulation Capturing the Interactions of Air Traffic Management and Flight Trajectories", *AIAA Modeling and Simulation Technology Conference*, 2003-5597, Monterey, CA (August 2002).

5. Sweet, D. N., Manikonda, V., Aronson, J. S., Roth, K., Blake, M., "Fast-Time Simulation System For Analysis of Advanced Air Transportation Concepts," *AIAA Modeling and Simulation Technology Conference*, 2002-4593, Monterey, CA (August 2002).

6. Zelinski, S., "Validating The Airspace Concept Evaluation System Using Real World Data," *AIAA Modeling and Simulation Technology Conference*, San Francisco, CA (August 2005).

7. Zuk, J., "ESTOL Vehicle Sector," *Vehicle Systems Program 1st Annual Industry Meeting*, (May 2004).