

# Modeling the NAS: A Grand Challenge for the Simulation Community

Frederick Wieland, Craig Wanke, Bill Niedringhaus, Len Wojcik

The MITRE Corporation  
Center for Advanced Aviation Systems Development  
McLean, VA 22102

**Keywords:** aviation, grand challenge, national airspace system, decisionsupport system.

## ABSTRACT

The National Airspace System (NAS) is the air traffic management system that supports flight operations around the United States. It is an example of an engineered complex system, in that it is composed of mechanical components (i.e. flights), human decision making through organizations (air transportation service providers and flight controllers), and the information flow that supports these decisions. The thesis of this paper is that representation of the NAS with all its interrelated components—mechanics, human decision making, and information flow—is a large effort involving multidisciplinary and “out-of-the-box” thinking, and thus constitutes a Grand Challenge for the modeling and simulation community.

## INTRODUCTION

The purpose of this paper is to introduce a problem domain concerning the National Airspace System (NAS) and to propose that modeling it is a Grand Challenge for the modeling and simulation (M&S) community. Our purpose is not to propose a specific NAS model, nor even to outline such a model, but rather to discuss the critical characteristics of the NAS and the interrelationships that need to be represented in any NAS model. This paper will focus on a conceptual understanding of this characterization, at the highest level of abstraction. We will not discuss details of the implementation of an NAS model. There will be no discussion of the programming model, such as whether the object-oriented or agent-based paradigms would better fit the system. Similarly we will avoid discussing implementation strategies, such as whether a loosely distributed or tightly coupled parallel architecture would be better, or whether such a model should be composed of small models linked via CORBA or the High Level Architecture. These questions are vitally important and need to be addressed before the Grand Challenge can be met. We will, however, concentrate on more basic questions in this paper. Why is modeling the NAS important? What types of questions must such a model address? Is it possible to develop an abstract characterization of the NAS against which specific models of the NAS can be built? We shall focus on these questions because a

basic understanding of the problem space must necessarily precede any discussion of its architectural realization.

The NAS is a collection of physical and human assets that effectuate air travel throughout the United States. Its physical component includes airports, airways, navigation aids, aircraft, and waypoints, while its human component includes air transportation service providers (ATSPs), controllers, pilots, traffic management specialists, airline dispatchers, and so forth [Nolan 1999]. There are about 16,000 airfields in the United States, of which about 1,000 of them handle most of the commercial air traffic, while 500 of them handle 80% of the general aviation traffic. There are approximately 60 airports where congestion produces measurable delays, while a number of them have been identified as “pacing” airports. Pacing airports are those where unexpected delays can have a significant cascading effect throughout the system.

In terms of volume, there are approximately 40,000 commercial flights per day, and approximately 10,000-20,000 general aviation (GA) flights. Both commercial and GA flight volumes are influenced by the weather, with a greater effect upon GA traffic. The system handles approximately 100,000 passengers per hour on more than 4,000 aircraft, which equates to about 650 million passengers per year. Growth in air traffic is closely correlated with the growth of the economy, with growth rates averaging two to four percent per year in the recent past.

We propose that representing this system in an integrated manner constitutes a Grand Challenge for the M&S community. When most engineers ponder the NAS, they think solely in terms of modeling the physical processes involved in NAS operations—moving airframes from point to point governed by air traffic control (ATC) procedures. Although there are a few unresolved issues in the mechanical movement of aircraft, we believe that it is mostly a solved problem, and therefore this alone is not worthy of a Grand Challenge. Rather the challenge is not only to represent physical NAS dynamics, but also to incorporate the behavioral and reactionary components of NAS decision making that are an important part of the system. We argue that representing merely mechanical processes (moving airframes) captures

only a small part of the larger problem. A comprehensive capability that integrates the mechanics, economics, and information flow of the NAS is required to provide comprehensive answers.

But before we continue, the relevance of addressing the NAS needs to be addressed. Why should we model the NAS? Why is it an important problem? Why should it be elevated to that of a Grand Challenge? Operation of the aviation system is becoming an increasingly important role in society. Not only is passenger air travel growing, but freight traffic is increasing as well. The economy depends, in part, on the efficient and timely movement of people and goods around the country and the world. The safety of these operations is the most critical variable in controlling the system. The increasing growth of air traffic has led to increased congestion, and thereby delays, causing personal and economic hardship on the users of the system. The increasing growth has also led to increasing concerns about the impact of airport operations on the environment, particularly noise and air pollutant emissions. Behind all growth are wholesale changes in NAS operations, from point-to-point scheduling before deregulation to the hub-and-spoke system in wide use today. There are anticipated future changes with the introduction of new equipment, the segmentation of the market with fractional ownership of airplanes, and, in the very far future, personal airplanes as envisioned by the NASA/FAA Small Aircraft Transportation System (SATS) initiative. Finally, with the recent terrorist attacks involving the airlines, there are new, emerging, and very challenging aspects of the system to be considered.

The criteria for a model to be considered as a Grand Challenge are threefold: (1) the problem must be demonstrably hard to solve; (2) the problem must not be known to be unsolvable; and (3) the solution must have a significant economic and/or social impact. The first consideration will be the subject of much of the rest of this paper. The second consideration should be self-evident: there is nothing in the modeling of the NAS that is a known unsolvable problem from a computational viewpoint. The third consideration was argued above, that the NAS itself is a significant engine of economic activity. Modeling it so that it can be studied and made more efficient will have a significant impact on the economy, individuals, and the environment.

## TOWARD AN UNDERSTANDING OF THE NAS PROBLEM DOMAIN

In this section, we present several different views of the NAS. These views collectively outline the characteristics and nature of the NAS that need to be captured in any integrated representation. We start by outlining the three primary reasons why NAS modeling is needed, and then follow

with a description of the NAS problem domain designed to characterize the NAS as a complex heterogeneous modeling challenge.

### Why Model the NAS?

There are three primary reasons why a model of the NAS is needed: tactical decision making (predictive modeling), long term decision making (strategic modeling), and post-priori event analysis (also a type of strategic model). Tactical decision making involves understanding the NAS in enough detail and with enough fidelity to influence near-term decisions—sometimes decisions that must be made immediately to influence events an hour from now. Tactical decision making is used in an operational context. The FAA regularly issues directives that manage the flow of traffic in and around congested areas that may arise from either excessive traffic or poor weather conditions (or some other exogenous event). Additionally, airlines and the FAA are both collaborating about the imposition of such directives, and the airlines are reacting to them by changing flights and routes. All these decisions are greatly enhanced through the use of tactical modeling.

Strategic decision making differs from tactical decision making in that it is used in a planning context and thus involves decisions made over long time scales. Strategic decisions typically are on the order of weeks to years—everything from changing a procedure to building a new airport. Post-event analysis is conducted to understand what happened during some previous events so that useful information can be extracted and applied to future situations. This process involves evaluating decisions made in the past, as opposed to improving decisions to be made in the future. Both types are (in some sense) duals of each other in the temporal domain. Table 1 shows some of the characteristics of tactical and strategic decision making.

An interesting research question is whether one modeling capability can handle all three uses. Can a single model (or set of cooperating models) be built that handles near-term predictions as well as long-term analysis? On the one hand, we can argue that the underlying NAS is the same for all three uses; that a valid fast tactical model employing real data can be used with synthetic data for strategic decision making; and that such a model can have a wide focus that can be easily narrowed down for specific uses. On the other hand, these requirements lead to different implementation technologies, different detailed design decisions, different levels of resolution, and different modeling styles such that a “one size fits all” model may not fit any one of the needs well enough for practical use. We will leave a definitive look at this question for further study.

**Table 1. Requirements derived from the main drivers for NAS modeling.**

<b>Tactical Decision Making (near-term decisions support system)</b>	<b>Strategic Decision Making (long term decisions or post event analysis)</b>
<ul style="list-style-type: none"> <li>• Fast turnaround time</li> <li>• Uses live (real-time) data</li>   <li>• Tend to be deterministic, with some (but few) stochastic components</li>   <li>• Tolerance for error is very low (must have a high degree of validity)</li> <li>• Narrow focus on one decision</li> </ul>	<ul style="list-style-type: none"> <li>• Slow turnaround time</li> <li>• Typically uses synthetic data (i.e. future traffic profiles, future airport/runway configurations, etc.)</li> <li>• Based on information about the future that is highly uncertain, and hence uses some form of stochastic decision making</li> <li>• High tolerance for error, although model must be valid to a certain degree</li> <li>• Wide focus on numerous interrelated variables</li> </ul>

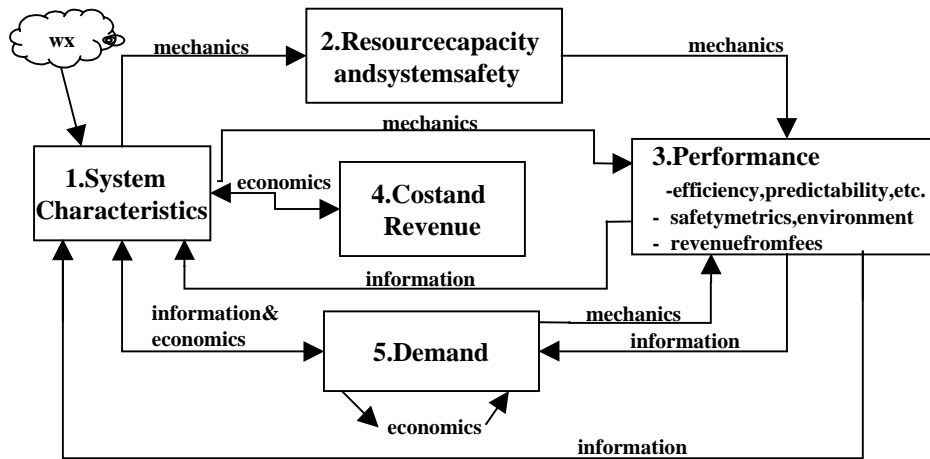
**Characterization of the NAS Problem Domain**

There are many questions that can be supported by a model of the NAS. They span the gamut from “how does a change in the system affect performance?” to “how will air transportation service providers (ATSPs) react to changes in demographics?” We have developed an architectural map—a superstructure of components of the NAS and the related question domains—that allows us to understand the various questions which a Grand Challenge NAS model must address. This characterization provides a framework against which specific models can be implemented. It allows us to understand how different models and components of the NAS might work together to address certain questions, as well as how to use seemingly disparate models to answer practical questions. Because there are hundreds of such questions, it may seem impractical to build models that answer them all. However, these questions can be generally classified as operating on a certain set of NAS characteristics, and this view (Figure 1) shows how these characteristics relate to each other.

The boxes in Figure 1 refer to measurable features of the NAS, while the arrows in the figure refer to mathematical models that use features as input parameters (or constraints) to predict other features. The mathematical models (the arrows) can be categorized into three basic types: mechanical, information, and economic models. *Mechanical models* represent the NAS in the physical domain: for example, the movement of airframes through the system, governed by procedural rules. When contemplating NAS modeling, most analysts limit themselves to such mechanical representations. Although they are an important component of NAS analysis, they should be supplemented with two other types of models. *Economic models* account for the behavior of the agents in the NAS. They typically account for the profit-

seeking and/or cost-reducing motivation that change the agents’ behavior, and therefore influence the mechanical models. We will return to this relationship shortly. *Information models* account for the effect of imperfect information, the effect of imperfect forecasts, or the effect of information flow on the NAS agents. Our claim is that a useful simulation of the NAS must contain all three types of system dimensions

The figure also contains five boxes that represent features of the NAS that are used as input or output for each of the models. Box 1, system characteristics, is the physical definition of the NAS: airports, runways, waypoints, navigaids, aircraft, and so forth. Also included in the box are the assets that are owned by the commercial operators: their fleet mix and their overall scheduling strategy. The weather is included here as a system characteristic, however, it could easily link to all other boxes on the diagram. Box 2 represents the capacity and safety characteristics of the NAS. Box 3 represents the overall performance of the system, using metrics that are tailored to a particular study. Performance metrics include efficiency, predictability, throughput, delay, safety, environmental, and even revenue from user fees. The latter is a result of NAS performance and therefore is included in box 3; revenues generated from changes in system characteristics are included in box 4. Box 4 represents the costs and benefits associated with specific system characteristics, i.e. the cost to users of installing new avionics equipment, or the benefit from a capacity for investing in airport expansion. Box 5 represents the NAS demand, in all of its manifestations. Demand can arise from users requesting service from the NAS (the usual definition of demand). It can also be viewed as passengers demanding service from users, or even as airframes demanding communications services from the NAS itself.



**Figure 1. A proposed framework for NAS modeling**

While this discussion of the models and characteristics in Figure 1 is necessarily brief, a deeper understanding can be gained through a use case that illustrates this framework. Consider one of the more common analytical questions: if we change the system through  $X$ , what is the performance impact as measured by  $Y$ ?  $X$  can be any system change: a new controller procedure, new avionics, a new runway, new airport, or a wholesale change (like privatization of air traffic management services).  $Y$  can be any resulting metric of interest, like throughput, delay, safety, controller workload, and so forth. To model this, we might begin with the system change in question, then run a mechanical model to determine what effect this change has on either capacity or safety of the system (boxes 1 to 2). For many questions, the analysis can stop here. But if we continue, we can use a mechanical model to predict performance given the change in system capacity (boxes 2 to 3). The new performance, then, can feed into an information model that influences passenger and/or airline demand (boxes 3 to 5). The new demand profile, via an economic model, will affect airline schedules (their demand on the system) (box 5 to itself). The new schedules can be fed into a mechanical model that then computes NAS performance (box 5 to 3). If the NAS performance is worse, this information will be used to change the system characteristic either short term (flow control programs) or long term (expand an airport) (box 3 to 1). Those system changes will have direct operating costs associated with them that can be modeled economically (box 1 to 4).

Virtually any NAS question can be modeled with this framework. One example is, what happens if the airlines adopt more of a point-to-point system instead of hub-and-spoke? This would be a procedural change by the airlines, which is a type of system change (box 1). This change modifies demand (box 5) that in turn influences performance

(box 3). Another example is, what happens if the population dynamics shift such that there is significant unmet demand for certain city pairs? This would be a change in underlying passenger demand (box 5) that feeds into a possible system change, for example a new airport (box 1).

This figure represents a structured way to characterize the NAS. It is intended to illustrate the aspects of the NAS that would need to be incorporated in a NAS model. But it is not specific enough to construct such a model. Consider the mechanical model that links a system change (box 1) to capacity and safety issues (box 2). A realization of that model would be entirely different if it worked on the airspace than if it worked on airports. An airport model would take into account the geometry of the runways, arrival and departure spacing thresholds, airport approach procedures, runway occupancy times, and so forth. An airspace model would take into account controller workload, the complexity of the route structure, letters of agreement between adjacent sectors and centers, as well as standard operating procedures. A model operating in the transition airspace might look entirely different as well.

If we step back and analyze this framework, we can characterize the NAS (and therefore model of it) in a relatively simple way. The NAS is an example of a multi-level distributed command and control system, with imperfect communication, where the various players (agents) have competing agendas in some circumstances, while they cooperate in others. The interaction of competition and cooperation among the players makes NAS modeling different from, for example, combat modeling, where the players are always antagonistic and do not cooperate. It is the interaction of the cooperating and competing agents with the NAS that changes the NAS itself, which in turn changes the environment in which the agents operate. Thus the system is truly

dynamic. Mechanical models alone cannot capture the richness of these dynamics, and thus we assert that a NAS model must incorporate the economic and information components as well.

## EXISTING NAS MODELS

There are dozens of existing NAS models. Most of them concentrate on modeling the mechanical nature of the NAS in different venues and different levels of fidelity. In this section we will review some of these models.

### Same Day Decision Support Models

Many NAS models have been developed for real-time decision support applications. As noted above, these models are primarily mechanical, in that they predict airframe movement through the NAS. They generally work on the individual flight level, including data such as flight schedules, flight plans, radar tracking reports, and wind forecasts. While the physics of atmospheric flight are well known and straightforward to simulate, the NAS itself has a strong effect on the trajectory a flight will take. Air traffic control procedures, for example, may restrict a flight to particular altitudes and speeds. Also, although flight plans are filed in order to establish the intended route of flight, speed, and altitude that an aircraft will fly, these plans are subject to change. Also, some of these changes may not be entered into the automation system, and as such must be heuristically inferred from observed flight behavior. All these problems become more prevalent as congestion increases, and that is when system management and improvement becomes more important.

The User Request Evaluation Tool (URET) provides a good example of such a model [Brudnick and McFarland, 1997]. URET is a "conflict probe" capability. By modeling future trajectories of flights, it predicts when aircraft on their current flight plan may come within 5 miles laterally and 1000 feet vertically of another aircraft (an aircraft-to-aircraft conflict) and when aircraft may penetrate restricted airspace (for example, a military operations area). Potential aircraft-to-aircraft conflicts are predicted up to 20 minutes in advance, alerting controllers to take action long before the situation will occur. URET also provides a "trial planning" capability, with which controllers can postulate solutions to the conflict, and probe candidate solutions to ensure that they are conflict-free. Current research and development efforts are focused on a capability to suggest conflict-free solutions that minimize the impact on the involved flights in terms of additional fuel burn and distance flown.

At the short time-scale used in predicting aircraft-to-aircraft conflicts, the impact of ATSP and air traffic management decision-making is mostly invisible. Once flights are airborne, the planning phase is complete; the job of the pilot,

h- ATC system, and ATSP dispatch office is to complete the flights safely and with a minimum of delay. A detailed, adaptive mechanical model is appropriate and sufficient.

This is not a true for the longer-ranged decision support systems used in traffic flow management operations. The Collaborative Routing Coordination Tools prototype is an example of such a system [Wanke, 2000]. It includes tools for flow problem recognition and strategy evaluation, designed to aid in minimizing delays due to airspace congestion and weather. For example, the system may identify areas of airspace where aircraft density is predicted to be greater than a specified safe capacity, or where hazardous weather is present (Figure 2). Some or all of the aircraft entering these "flow constrained areas" must be rerouted for safety. The traffic manager can then use the automation tools to trial-plan reroutes for groups of flights, evaluate the predicted changes in aircraft density and additional flying time required, and hopefully arrive at a solution that resolves the safety issues while minimizing the imposed flight delays.

i-  
e-  
d-

t-

s-

a-

r-

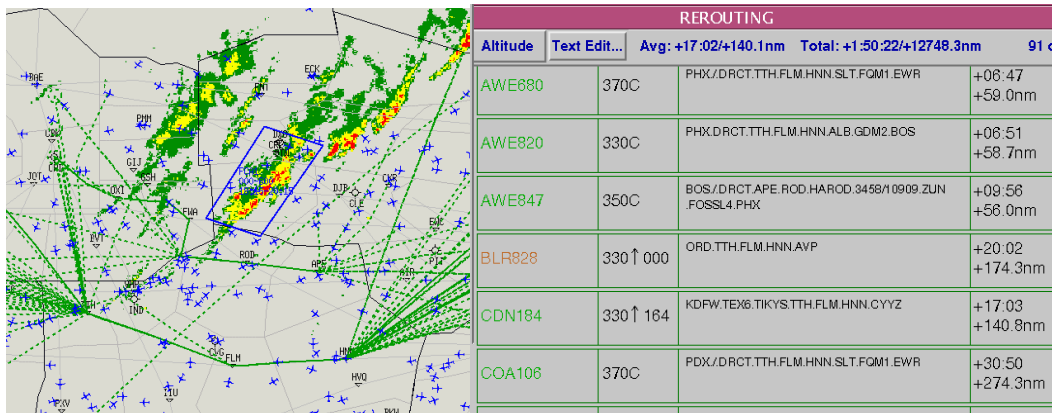
n-

d-

The basic modeling process used in the CRCT prototype—and in the currently operational Enhanced Traffic Management System (ETMS), upon which CRCT-derived tools are being deployed—is a similar mechanical model to that used by URET. However, this type of decision support requires predictions of aircraft movements several hours in advance, and therefore is subject to the interactions of the various NAS agents. For example, flight plans may not be filed until one hour before a flight is scheduled to depart. Until this occurs, predictions for that flight are based solely on published flight schedules and historical route and altitude departures for the flight, and therefore are subject to change by the airline's strategy on that day. Also, flights cannot be treated as wholly-independent entities. If the airport from which a flight is planning to depart is currently undergoing large departure delays, then the flight will likely depart late; but how late? Under these conditions, the most difficult element of modeling a flight's trajectory may be simply predicting the departure time.

A screenshot of CRCT is shown in Figure 2. The traffic flow manager can graphically specify possible reroutes around an area of severe weather. In this example, an area of severe weather has blocked a primary air route. Two reroutes have been specified; one for eastbound flights and one for westbound flights. The time and distance added to involved flights are calculated, and displayed individually and in total in the rerouting window at right. The system will also evaluate the impact of these reroutes on airspace sector loading, in order to maintain safe controller workload.

r-



**Figure 2. Collaborative Routing Coordination Tools (CRCT) Prototype.**

### Longer Term Information Models

Information models attempt to predict the behavior of agents as they process information about the NAS. As such, they represent a feedback component of NAS modeling: as the NAS changes, the agents then change their behavior, which in turn changes the NAS yet again. A good way to approach such models is to combine the mechanical models, which produce the state of the NAS, with information models that represent the decisions made by the agents as the state of the NAS changes. As often there is little known about the detailed behavior of the agents (for example, who has an algorithm that can predict exactly which flights an airline will cancel in poor weather?), such models usually contain a learning component that relies on genetic or evolutionary programming. We will review one such model here, called Jet:Wise.

Jet:Wise is an agent-based model (ABM) that explores the evolution of the airline industry and its interactions with NAS. Jet:Wise, developed at MITRE, models airline decisions such as markets (what airports to serve and where to establish hubs), fleet mix, schedules, fares, and airline responses to delays, congestion and missed connections. Jet:Wise also models decisions by the Federal Aviation Administration (FAA) or individual airports, related to airport hourly capacities and possible reductions imposed due to weather or congestion. Because these factors both affect and are affected by each other, understanding their complex interaction is beyond the ability of traditional modeling tools and has not been attempted until now. The problem is difficult enough thanks to the sheer numbers involved: thousands of aircraft and millions of possible flight connections. But the main source of difficulty is the tangled web of interactions with multiple feedback loops.

There are a number of interesting potential uses for such a model. It could be run to evaluate the impact of possible policy

changes, such as adoption of Free Flight [RTCA, 1995] concepts. It could help analyze airline's reaction to future changes such as a new runway or airport, invention of a new type of jet, a change in operational costs, an improvement in the NAS's ability to predict aircraft arrival times, removal of certain ATC restrictions (thereby shortening travel times), or changes in passenger demand.

The Jet:Wise software is a loop over three phases: MARKET, FLY, LEARN, discussed below.

- The MARKET phase is an economic model, with input (a) a model of passenger demand and (b) a set of airline schedules and fares (e.g. for a particular day over the entire US). The MARKET phase determines which passengers by ticket on which flights
- The FLY phase is a simple mechanical simulation of the aircraft moving along their routes, subject to queuing for takeoffs and landing, to determine averaged delay, operational costs, passenger connection times, and whether passengers that missed their connections were able to take a later flight.
- In the LEARN phase, the heart of Jet:Wise, airlines may (a) decide on values for decision parameters such as fares, aircraft size, flight departure times and (possibly padded) arrival times, percentage of seats reserved for business travel, etc, (b) buy and sell aircraft and alter itineraries, and (c) meta-learn—that is, evolve their learning mechanisms, perhaps by weighting market share or some other factors in addition to immediate profit, or emphasizing/de-emphasizing certain learning techniques of (a) and (b).

An ABM such as Jet:Wise is almost a philosophical opposite from an *expert system* (ES) [Durkin 1994], in which a series of rules intended to reflect the knowledge of human experts is programmed into the model. An ES alternative to Jet:Wise might, for instance, try to encode rules-of-thumb used by airline marketing and scheduling departments when establishing hubs and arrival/departure banks. Jet:Wise, in contrast, has no built-in knowledge of such phenomena as hubs and banks. It only a priori knowledge about transfer flights is that the scheduled arrival time of the incoming flight needs should sufficiently precede the scheduled departure time of the outgoing flight. In Jet:Wise, the idea is that hubs and banks emerge only if they are warranted by market forces. Complex second-order NAS phenomena such as hubs and banks need to be reinvented with each Jet:Wise run. In an ES they are inputs, in Jet:Wise they are outputs. As a result, Jet:Wise can provide surprises that might not be observed in an ES.

### CONCLUSION: THE NAS AS A GRAND CHALLENGE FOR MODELING & SIMULATION

We have been arguing that modeling the NAS is far more than merely modeling the mechanical movement of airplanes from departure to arrival gate. While the mechanical movement of flights is a vital component to NAS modeling, as an intellectual and engineering problem it has been solved, many times, with many different technologies for many years. The mechanical modeling component is not the Grand Challenge. Rather, it is the complex interactions that constitute the NAS that must be integrated to create a practical, realistic, and usable model. These interactions are important because the aviation industry is remarkably successful at adaptation. For example, when a new runway at a large airport is built, its capacity is often realized within a short time. Mechanical models that merely account for the increased capacity of the airport do not have the richness to predict the shifting use of that new capacity. If the mechanical model were supplemented by economic models of airline behavior as well as information models that describe the flow of data through the NAS, the predicted impact of the new runway could be determined more realistically. There are many such dynamic forces that operate on the NAS. Travel patterns shift, demographics changes, passenger and cargo demand increases,

seasonal effects vary. Against these shifting forces—some short term and some long term—the NAS must be modeled.

Our conclusion is that a comprehensive model of the NAS—however it is designed and built technologically—is incomplete and subject to first-order errors unless all such interactions are incorporated to some degree. Because modeling the mechanical, the human decision making, the procedural, and the economic together in one integrated representation is a difficult problem, but not unsolvable, we nominate it as a Grand Challenge.

### ACKNOWLEDGEMENTS AND REFERENCES

We would like to thank our reviewers, Dr. Judith Dahman and Joseph Sinnott, for their insightful comments that have been incorporated throughout this paper.

The contents of this paper reflect the views of the authors. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of the views expressed herein.

Brudnicki, Daniel J. and McFarland, Alvin L., 1997, "User Request Evaluation Tool (URET) Conflict Probe Results," *1st USA/Europe Air Traffic Management Research and Development Seminar*, Saclay, France, 17-20 June, 1997.

Durkin, John, 1994, *Expert Systems: Design and Development*, New York: Maxwell McMillan International.

Nolan, Michael, 1999, *Fundamentals of Air Traffic Control*, Wadsworth Publishing Corporation: 1999.

RTCA 1995, *Report of the RTCA Board of Directors' Select Committee on Free Flight*, Washington DC, January 1995.

Wanke, Craig, 2000, "Collaborative Routing Coordination Tools (CRCT): Developing Traffic Flow Management Decision Support Concepts", *Proceedings of the 2000 Air Traffic Control Association Conference*, November 2000.