

Simulating the National Airspace System

William E. Weiss
Elizabeth S. Lacher
The MITRE Corporation
McLean, Virginia

ABSTRACT

This paper describes an event-step simulation model being developed for the FAA that can be used to evaluate how a change in one component of the airspace system affects the flow of air traffic in the rest of the system. The primary measure of system performance in the model is delay. The model design, sample model applications, and future model enhancements are discussed.

1. INTRODUCTION

The Federal Aviation Administration (FAA) has identified the need for a capability to analyze high-level, system-oriented air traffic flow problems in the National Airspace System (NAS). (The NAS includes airports, terminal airspace, and en route airspace.) In response to that need, the National Airspace System Performance Analysis Capability (NASPAC) is being developed by the MITRE Corporation for the FAA. In order to meet FAA objectives, the NASPAC group has produced several tools including an event-step simulation model. That model is the focus of this paper.

The primary purpose of the model is to evaluate how a change in one component of the airspace system affects the rest of the system. It is well known that demand in excess of airport capacity causes delays. An increase in capacity at an airport can reduce delays; similarly a change in the total demand or the demand pattern can affect delays. However, delays change not only at the airport in question and in the surrounding airspace, but throughout the airspace system.

Until now, analysts have only been able to study delays at individual airports or in localized airport/airspace regions. The NASPAC simulation model is the first model to allow analysts to study how individual components of the entire airspace system interact, how delays ripple through the system, and how the entire system will react to projected demand or capacity changes.

The model is being developed in a multi-phased effort. The purpose of this paper is to describe the Phase 1 model and its design. This model was designed, coded, and partially validated in less than six months. The modeling effort during this phase concentrated on airports and nearby airspace, and used readily available data. The Phase 1 model was intended as a proof of concept and a quick response to modeling needs of the FAA. This paper does not describe model validation nor does it discuss results in detail.

The Phase 1 simulation model was written in the Simscript II.5 programming language and developed on a Digital Equipment Corporation VAX 11/780-5 computer. For its initial runs, the model used approximately three hours of central processing unit (CPU) time on the VAX 11/780-5 or 45 minutes of CPU time on a VAX 8700 computer. These runs simulated one 24-hour day.

The remainder of this paper is divided into five sections: Model Design, Model Inputs, Model Outputs, Model Applications, and Future Enhancements.

2. MODEL DESIGN

The model has been designed to be applied to a wide variety of analyses. To this end, it works in conjunction with a preprocessor and a postprocessor. The preprocessor generates inputs for the model, while the postprocessor generates detailed output reports and graphic displays. Both the preprocessor and the postprocessor off-load time-consuming data processing from the model itself and provide analysis flexibility.

Because the model is intended to be used for system-wide analyses, it was not necessary to model the system components in the high level of detail that might be necessary for a local-area analysis. The design had to maintain a balance such that the level of detail would be meaningful without being overly cumbersome.

The focus of the Phase 1 model was airports, and arrival and departure fixes. A fix is a location in the airspace used to route traffic; often traffic is sequenced and/or merged at a fix. Arrival and departure fixes are relatively close to the airport, and are near the boundary between terminal and en route airspace.

The components in the airport/fix network are defined as input to the model; thus, the network can be easily adapted to a particular analysis. The airports and fixes are represented as service nodes with queues. All airports in the NAS which are not defined explicitly in the simulation network are modeled as a single sink/source airport; there are no queues at the sink/source.

Although the Phase 1 model is completely deterministic at present, the model was designed so that stochastic elements may be readily added. Time constraints and difficulties in data collection did not allow the addition of these stochastic elements in Phase 1. Once appropriate distributions for stochastic elements are determined, they will be added to the model.

Key elements of the model design--simulation events, airport structure, and fix structure--are described below.

2.1 Simulation Events

The model is centered around the "movement" of aircraft through the airport/fix network during a simulated day. This movement of an aircraft is called the aircraft process, and defines the events "experienced" by the aircraft (see Figure 1). Each aircraft process may consist of several flight legs. A flight leg includes the following events:

- Pushback from gate
- Enter departure queue
- Takeoff
- Enter arrival or departure fix queue
- Cross arrival or departure fix
- Enter arrival queue
- Land

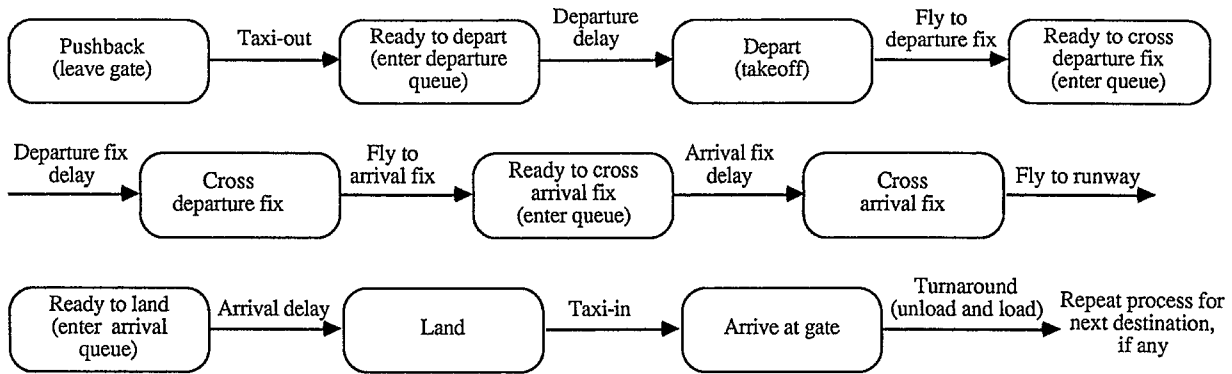


Figure 1: Aircraft Process

Each origin and destination airport, as well as any intervening fixes, for particular a aircraft process is determined by the preprocessor. Because only a subset of all arrival and departure fixes were modeled, not every flight leg includes an arrival or departure fix; also, depending on the origin and destination, a flight leg may include both an arrival fix and a departure fix.

A turnaround time, which represents the minimum time an aircraft spends at the gate, must elapse between flight legs; this time is constant for each aircraft type. An aircraft will not push back until its scheduled departure time. Thus the turnaround time can only serve to make the aircraft pushback late, not early.

Between the pushback-from-gate and enter-departure-queue events there is a taxi-out time, which is constant for each airport. Similarly, there is a taxi-in time, also constant for each airport, between the land and arrive-at-gate events.

"Flying" times between elements in the network (airports and fixes) are deterministic, but are different for each flight as generated by the preprocessor (see "Model Inputs" for detail).

2.2 Airport Structure

Although Phase 1 development concentrated on modeling airports, to model runway operations explicitly would have been extremely time consuming in terms of data collection and model development. Additionally, it was not necessary to model runway operations because the focus of the simulation was system-wide analysis rather than a single airport.

An airport was modeled as two servers with separate queues; one was dedicated to arrivals, the other to departures. Service time is the minimum time required between two arrivals or between two departures; runway occupancy is assumed to be instantaneous. The arrival queue is FIFO; however, the departure queue is more complex. The service time and departure queue algorithms are described in detail below.

Service Time Algorithm. The calculation of service times at airports is a key element in the modeling of airport operations and controller actions in the simulation. During the course of a day, controllers vary the service times for arrivals and departures. Although they may give priority to arrivals, the gaps between arrivals are often "stretched" to accommodate departures. When departure demand is far greater than arrival demand, arrivals may be delayed upstream while departures are released. Thus, arrival and departure service times vary inversely, depending on demand.

To capture this phenomenon, the following approach is taken in the simulation. Service time is assumed to be the inverse of hourly capacity; the relationship between arrival and departure capacities is assumed to be linear. (The linear assumption is based on capacities generated using the FAA Airfield Capacity Model.)

For each airport three parameters are determined: maximum arrival capacity, maximum departure capacity, and minimum departure capacity. (Minimum arrival capacity is assumed to be zero; that is, when the maximum number of departures are being released, no arrivals can be accepted). The relationship between arrival and departure capacities can be defined by these parameters (see Figure 2).




Figure 2: Arrival versus Departure Capacity

Queue length is the measure of demand used to determine arrival and departure service times. The lengths of the queues are not only readily available within the model, they are also obvious signs of demand visible to controllers. As each arrival or departure reaches the front of the queue in the model, the ratio of the length of the departure queue to the length of arrival queue is used to determine the appropriate point on the capacity curve for servicing the demand. The inverse of the identified capacity is then taken as the service time.

If the arrival queue is empty, the airport operates in a departure-priority mode, imitating the process of controllers releasing as many departures as possible until an arrival appears. Conversely, if the departure queue is empty, the airport operates in arrival-priority mode.

The algorithm accounts for the fact that it is unrealistic to create a very large interarrival gap in order to allow the release of

multiple departures between two arrivals. (Because arrivals are airborne, they are generally given priority.) To avoid this problem, the departure capacity is restricted to 75 percent of maximum departure capacity unless the arrival queue is empty.

Additionally, the algorithm can force arrival and departure service times to remain constant and independent for airports where arrival and departure capacities are, in fact, independent. (For example, this occurs at airports where there are runways dedicated specifically to arrivals and departures.) This is achieved by setting the maximum departure capacity equal to the minimum departure capacity.

Departure Queue Algorithm. The departure queue differs from a simple FIFO queue for two reasons.

First, if an aircraft in the departure queue has an Estimated Departure Clearance Time (EDCT), it may not be able to depart when it reaches the front of the queue. An EDCT is the scheduled departure time plus a ground delay assigned by the FAA due to congestion expected at the arrival airport.

Second, if the aircraft will be using a departure fix, it may be subject to a departure restriction associated with that fix. A departure restriction is a minimum spacing requirement between two aircraft departing from the same airport for the same departure fix. This spacing helps the controller who is merging traffic from several airports at the departure fix.

In the simulation, entries in the departure queue are ranked on the maximum of the time they entered the queue and their EDCT. Thus, aircraft with assigned EDCTs do not reach the front of the queue and hold up aircraft behind them with no EDCT. An aircraft is removed from the queue (that is, the aircraft "takes off") when it is the first aircraft in the queue with its departure restriction and EDCT satisfied. After the calculated departure service time elapses, another aircraft is removed from the queue (provided there is another aircraft in the queue that is qualified to depart).

2.3 Fix Structure

A fix is modeled as a single server with one first-in/first-out (FIFO) queue. The service time is constant and is an input parameter. The service time for a fix is the minimum time required between two aircraft crossing the fix. Aircraft are assumed to cross the fix instantaneously.

3. MODEL INPUTS

One of the greatest challenges in creating a simulation of this scope was finding the data to drive the model; these data are described below.

An on-line version of the Official Airline Guide (OAG) from October 1, 1987 was used to generate air carrier demand (the OAG contains the complete schedule of commercial passenger--air carrier--flights, which is used by travel agencies and other organizations for planning purposes). General aviation and military flights were distributed randomly throughout the day based on historical data. En route flight times were derived from the OAG.

Initially, 58 airports were modeled; airports not explicitly modeled were represented by a source/sink airport. The 50 busiest airports in the continental U.S. were included; approximately 70 percent of all air carrier operations occur at these airports (based on the October 1, 1987 OAG). Eight more airports were added because they shared terminal airspace and fixes with airports in the original set of 50.

Airport capacities were supplied, for the most part, by Air Traffic Operations Service of the FAA. Some capacities were generated using the FAA Airfield Capacity Model. Capacities were compared to those found in Airport Delay Task Force Reports and

FAA Engineered Performance Standards (EPS), for airports where these were available (EPS airport capacities are used in air traffic flow management programs for the busiest airports in the country).

A total of 48 arrival and departure fixes were included in the Phase 1 analyses. Fixes were chosen based on interaction with major airports. For arrival fixes, three terminal areas with a total of 13 arrival fixes were included; for departure fixes, seven terminal areas with 35 total departure fixes were included.

Data such as flight itineraries, taxi times, and turnaround times are presently being obtained from airlines. These data are being analyzed and will be incorporated into the input data set. Additionally, on-line reports of aircraft position are being obtained from FAA computers. These will be used to increase the accuracy of en route flight times, fix usage, and arrival and departure service time distributions. Traffic counts, runway configurations, and weather information from FAA Air Traffic Control Towers will also contribute to increasing the accuracy of arrival and departure service times.

4. MODEL OUTPUTS

Delays are the primary output measure in the model. Two types of delay are calculated: technical and effective. A technical delay is the amount of time an aircraft spends in a queue; that is, the amount of time it spends waiting for a resource (e.g., runway, fix). Effective delay is the difference between the simulated arrival time at the gate and the scheduled arrival time for a flight leg.

Effective delay is an extremely useful measure and is not found in other models. Effective delay is a measure of all the delays an aircraft experiences, including those due to the late arrival of that aircraft on a previous flight leg. It is the delay that a passenger on board the aircraft would perceive. Thus, the effect of delays "rippling" through the system can be observed as a flight accumulates delay during the simulated day.

The delays and throughput are recorded in a set of default reports produced by the model. In addition to the default reports, trace files containing one record for every event are produced. The trace files are used by the postprocessor to generate more detailed reports as required by a particular analysis. Both the default reports and the reports generated by the postprocessor from the trace files are used to produce graphic displays.

5. MODEL APPLICATIONS

The Phase 1 simulation model was designed to be used for a wide variety of system-level analyses. The model is best suited for analyses that ask "what if" questions about the NAS. A few sample applications are listed below.

- o The model can be used to study the sensitivity of delays to changes in airport capacity. Thus the model can be used to identify airports where even a small increase in capacity will have a large effect on delays, not only at the airport in question but at other airports as well. For example, a preliminary analysis found that a small increase in capacity at Los Angeles International reduced delays at airports such as Las Vegas, San Francisco, and San Diego.
- o As a follow-up to the analysis described above, the model can be used to evaluate the effects of specific procedural changes that would increase airport capacity by a small amount. Changes that involve a different use of arrival or departure fixes (for example, changing the number of aircraft that use each fix or adding a new fix) are especially suited for study in the model. Changes in the procedures for issuing EDCTs are another example of this type of application.

- o The model can be used to determine the impact on system-wide delays of changing capacity at a single airport. In a preliminary analysis, it was found that for a 10% increase in capacity at any given major airport, 50-60% of the decrease in delays occurred at that airport, while the remainder of the decrease was distributed among other airports in the system.
 - o The effects of proposed airport improvements can be analyzed using the model. First, the impact of the improvement on capacity can be estimated, or in many cases is included in the proposal. This information, as well as projected demand data, can then be used as inputs to the model. The model results would show the benefits of the improvement at that airport and throughout the system.
 - o The model can be used to identify the sensitivity of delays to changes in airport demand. Recently, a runway at Los Angeles International (LAX) was closed. In preparation, the airlines and FAA met to discuss and implement schedule changes. When the runway was closed and the new schedule activated, delays actually went down. The model could identify other airports where a cooperative effort between the FAA and airlines might reduce delays.
 - o Specific proposed schedule changes (even those as radical as the addition of a new airline hub) can be evaluated. While a schedule change might reduce delays at the airport in question, it could increase delays elsewhere. This type of effect can be predicted with the model.
 - o The model can be used to identify areas where the constraint on traffic flow is not the airport itself but the surrounding airspace. This happens in large metropolitan areas where several airports use the same or overlapping terminal airspace. An understanding of the interaction between the airspace and the airport is especially critical when evaluating proposed changes.
 - o All of the above applications can be expanded to include varying weather conditions. Many airports, theoretically, do not have a delay problem. However, adverse weather conditions can radically diminish capacity and thereby affect delays. Understanding weather effects is especially important at airports that are regularly affected by severe weather.
- o En route sectors will be added to the model. A sector is a volume of airspace generally handled by one air traffic controller. It will be modeled as a multi-server service node with service times (sector transit times) that are aircraft dependent.
 - o Other en route nodes will also be added to the model. These will represent points or areas in the en route airspace used for sequencing, merging, or spacing aircraft; they may or may not be associated with an en route fix in the airspace. Their representation in model will be similar to that of the arrival and departure fixes, regardless of whether the node is actually associated with an en route fix.
 - o Various algorithms will be enhanced or added to the model as new information becomes available. Air traffic flow control procedures and airport operations are examples of areas that might be improved after further observation of controllers in the field. Additionally, logic about flight cancellation and dynamic adjustment of aircraft schedules may be added when airline operations are better understood.
 - o Inputs to the model will also be enhanced as new data become available and are analyzed. Although these new data will not change the structure of the model, they will improve the credibility and validity of the results.

AUTHORS' BIOGRAPHIES

WILLIAM E. WEISS received a S.M. in Civil Engineering from the Massachusetts Institute of Technology (specializing in transportation systems) in 1980, and a B.S. in Civil Engineering from Michigan State University in 1978. He recently completed a simulation of surface traffic operations at La Guardia Airport in New York; this simulation, which runs on an IBM PC-AT, includes interactive, animated graphics. His previous experience includes terminal-area air traffic control systems analysis, real-time systems engineering, civil engineering in airports, and highway engineering. He has been a Member of the Technical Staff at The MITRE Corporation for four years.

William E. Weiss
The MITRE Corporation
7525 Colshire Drive, W335
McLean, Virginia 22102
(703) 883-6590

ELIZABETH S. LACHER received a B.S. in Operations Research from The George Washington University, Washington, D.C. in 1987. She has been a Member of the Technical Staff at The MITRE Corporation since 1987. She is a member of ORSA, TMS, Tau Beta Pi, and Omega Rho.

Elizabeth S. Lacher
The MITRE Corporation
7525 Colshire Drive, W335
McLean, Virginia 22102
(703) 883-5841

It should be noted that although the model is a very useful tool for evaluating the benefits of proposed changes in terms of delays, it is not adept at evaluating the safety implications of the changes. Moreover, the model is not intended to provide THE answer in an evaluation, but is one tool to be used in the decision-making process.

6. FUTURE ENHANCEMENTS

The Phase 1 simulation model, designed as a quick response to the modeling needs of the FAA, is still under development. Planned enhancements have been prioritized and will be implemented as time and other resources permit. Near-term enhancements are described below.

- o The addition of stochastic elements will begin shortly. Almost all aspects of the model can be randomized once appropriate distributions have been determined. En route times, taxi time, and turnaround times have been identified as candidates for randomization in the near term.