APPLYING SIMULATION TECHNIQUES TO AN AIR TRAFFIC CONTROL STUDY

Richard C. Baxter, Julian Reitman
and Donald Ingerman*

Norden Division
United Aircraft Corporation
Norwalk, Connecticut 06856

INTRODUCTION

Air traffic in the United States has expanded rapidly due to the rising demand of passengers for rapid, convenient, comfortable transportation, the increasing use of air freight, and the soaring growth of general aviation, which includes the recent proliferation of business jet aircraft. The primary purpose of the United States air traffic system is the safe, efficient movement of air traffic; it is required to accommodate all aircraft which place a demand upon it, from the large, fast commercial jet airliner to the small, slow privately owned airplane, and in so doing must accommodate the widely varying navigational and communications capabilities and performance characteristics of these aircraft.

Aircraft types can change or be modified in a short time to accommodate a market--witness the increase in instrument flight capabilities among general aviation aircraft, or the experimental introduction of STOL aircraft into commercial service. Changes in the air traffic system, with its airports, air routes, navigation, and communication facilities, and geographic subdivisions, require much more time, partly because the system has to maintain service as it is being changed. Demands to modernize the system and to provide more airports have come from users and from Congress, while people living in neighborhoods of existing airports have demanded decreases in the aircraft noise levels emanating from these airports.

Air traffic in the United States represents a very complex system of very many elements (aircraft, navigational aids, traffic controllers, regulations, weather factors, etc.) which interact with each other in complex ways. No simple description or analytical expression is available to describe the behavior of this system, nor is one likely to be available in the near future. Therefore, deductions drawn covering the effect of changes in the system parameters such as introduction of new airports, or changes in control procedures, etc. on the system performance as measured by criteria such as traffic delay, number of near misses, etc. depend on computer simulation. Such highly complex, multiple interconnected systems operating with constraints tend to behave in a counter-intuitive manner. Similarly, analyses of a portion of the system without proper regard for the rest of the system and further constraints operating on it can yield highly misleading answers. For instance, installation of a second runway at an airport may attract more traffic to the airport and thus create greater rather than less congestion.

On the other hand, experimentation with the entire actual system to determine the effect of changes in its characteristics or its performance is expensive and time consuming to the point where this approach is completely impractical and possibly dangerous.

Modeling or simulation of a complex system using digital computers has been used to an increasing extent in recent years, and the possibility of modeling the air traffic system has been considered many times. Because of its complexity, the conclusion in the past has been that a modeling of the entire system is impractical with available computing equipment and techniques. However, there has been continuous improvement in both, and there is reason to believe that a meaningful simulation of the air traffic in the United States may now be practical, i.e., technically feasible at a cost in development and computer running time commensurate with the benefits obtainable for such a simulation.

As such, we undertook to develop an exploratory model to determine if a simulation of such an extensive and complicated system is feasible.

MODELING APPROACH

Our approach in modeling complex problems of this type is one of successive approximations, each covering a greater portion of the problem. These successive approximations are started at a very simple level and expanded as more is learned about the problem. Therefore, the degree of success in simulation is predicated on the following: flexibility to make changes, coping with data banks of this size, establishing data structures, and ease of debugging.

Norden's proprietary version of GPSS more than adequately meets these criteria. Specifically, GPSS/360-NORDEN has a data library, core overlay ability, and man-machine interactive features. The data library feature

---

*Formerly with Norden, now with Bell Laboratories, New Brunswick, New Jersey permits the permanent storage of information as matrices on disk. This information can be accessed, manipulated, and changed by a normal run of the model; the changes in the information can be temporary or permanent. Load sections allow core overlay by allowing the model to use an unlimited number of matrices and blocks by storing them on disk and calling in only a limited predesignated number at one time to be core resident. Using an IBM 2250 Display Unit the analyst can change his model, insert or delete lines, execute, and interact with a dynamic display during execution. Assembly and input errors are displayed on the CRT enabling the analyst to rapidly debug the model.

INVESTIGATION OF SIMULATION METHODS

A major concern in the attempt to simulate the entire Air Traffic Control problem was the amount of computer time required to perform the simulation. Therefore, various methods of developing a more efficient simulation were investigated and utilized in structuring the model.

One basic consideration in a fast time simulation is computer operating time per run, and alternative methods must be evaluated so that the time required to perform functions in the computer will be minimum. A coarse model was generated to investigate the time required to execute various equivalent logic structures.

To enable actual timing of computer use, GPSS/360-NORDEN was modified to note time of entry and completion of each phase of GPSS. One timing run evaluated a method of transferring information (speed) between transactions (aircraft) in the program, which can be done using either an ALTER-SCAN technique, or SAVEVALUES. In the ALTER-SCAN, the parameter conveying speed in each of a group of transactions, representing for example, all aircraft in Colt's Neck holding pattern, is examined to see if it matches a desired parameter, the limiting speed, in a control transaction. The value of this desired parameter can then be placed in the corresponding parameter positions of the group transactions. Alternatively, the value of the limiting speed parameter can be placed in a SAVEVALUE position, which is then sampled by each transaction in turn. The ALTER-SCAN method provides a 15% faster transfer of information between transactions.

Another timing run investigated the quickest means of controlling the sequence of transactions moving through points in the model, using either GATE-LOGIC handling or PRIORITY-BUFFER. The GATE-LOGIC uses a gate block to check the priority or sequence number of a transaction attempting to enter and refuses entry if the priority is not the desired one; if it is, the transaction proceeds through, and passes a logic switch which makes the gate available for the next transaction. An example might be a runway with mixed operations arrivals having priority over departures. The PRIORITY-BUFFER mode is used to reset the priority order on the current events chain (actions that are waiting to happen at the particular instant of computer clock time), and thereby cause different sequencing. The GATE-LOGIC mode is 15% faster.

- when sending transactions through the model it is faster to LOOP, when possible, than to GENERATE new transactions.

- when trying to use a limited capacity item, a FACILITY preceded by a LINK with alternate exit is much faster than the FACILITY alone. The FACILITY preceded by a LINK with alternate exit takes essentially the same time as a TEST based on the contents of a SAVEVALUE.

- saturating the Current Events Scan has little effect on execution time.

DISCUSSION OF COMPUTER MODEL STRUCTURE

Air Traffic Control consists of three phases:

- activities on the ground (airport phase)

- activities in the airspace immediately surrounding the airport--that is, arrivals and departures (terminal area phase)

- activities between terminal areas (enroute phase)

The computer logic model was also structured into ground, terminal, and enroute phases. Each section can stand alone and, by adding further elaborations, could be used for a detailed study of that phase by itself, or they can be interconnected to allow for the full interaction of the entire system.

It should be mentioned that the logic structure of the model is general in nature and is independent of the number of airports, gates, flights, and route structure. Thus to model a completely new situation requires only that the data specifying these items be inputted to the model.

The model is structured to accept input information in the form of matrix savevalues. These matrix savevalues are used to define the airports, airline rosters, routes, and schedules. This type of structuring allows the inputs to be changed, increased, or decreased with ease.

The advantages of this approach are many. GPSS/360-NORDEN permits these data bank inputs to be part of a permanent library of matrices. Thus, the user can readily recall a given configuration. The input data can be generated or revised through the use of the 2250 Display Unit. He merely detects on the row and column of interest and keys in the desired data. Facilities are also provided for inputting in entire rows or columns of data at one time. The use of awkward INITIAL cards is eliminated. (Similarly selected results can be stored in matrix savevalues and stored in the library for future reference.)

AIRPORT PHASE

The scheduling program scans the input data matrix that defines each flight in terms of its airport pairs, airline, aircraft type and tail number and is ordered chronologically by departure times (E.S.T.). When the simulated time for each flight occurs, the scheduler creates a GPSS transaction representing that flight and places it on the current event chain for processing by the model.

It should be noted that airports beyond the areas of interest can be used to generate traffic into and out of that area. In these areas much of the detail of the activities at these outlying airports and terminal areas are omitted by the model.

The processing of each scheduled flight occurs in this section of the model and is described below.

The scheduled flight searches the available aircraft from its airline at its airport of origin by tail number. If the tail number is located among the available aircraft the flight proceeds, otherwise it waits for the tail number to become available. A check is made to determine whether the destination airport can accept the flight at its scheduled arrival time, then the aircraft transaction proceeds via taxiways a specified distance from gate to the selected runway at an input taxiing velocity. If the runway is occupied, the aircraft transaction enters a queue and waits its turn, first-in-first-out, to acquire the runway. Once the aircraft is on the runway, a check is made to determine if the enroute phase can accept the flight, another to determine if a sufficient time has elapsed since last runway use, and another to determine if a wind shift has closed or changed the runway assignment. Then the flight departs, and frees the runway for the next aircraft transaction. The present airport configuration has parallel runways, one for takeoffs and the other for landings.

Upon landing, the flight occupies a usage time on the runway, then taxis to the terminal and awaits availability of a gate there. A service time is incorporated, depending on whether the flight terminates or continues, and the tail number is then made available at that airport for future scheduling.

TERMINAL AREA PHASE

Upon takeoff, a flight proceeds to the first point of the standard instrument departure (SID) for the runway used and direction of flight. There is no provision for geographically locating the SID data or for conflict concern enroute to the first SID point because traffic in the terminal areas are restricted to blocked airspace.

Upon reaching the first SID checkpoint, the flight enters the airway and attempts to climb to its enroute altitude. Should a potential conflict occur in crossing traffic using the same segment of the airway, the aircraft's speed is reduced by 25 knots and the test for possible conflicts is repeated. This procedure is repeated until a speed is reached for which there are no conflicts.

After gaining the enroute altitude, the flight transaction proceeds to the enroute phase of the model.

Upon arrival in the designated terminal area, the flight transaction enters the airport holding area from which it attempts to enter the approach area. It should be noted that, if the destination airport is not under consideration in the particular study, no further consideration is given to this flight. The approach area has a multiple aircraft capacity which depends upon the spacing required between aircraft and on the length of the approach. This spacing is modified by the weather at the particular airport.

The flight then proceeds through the approach area and then into the final approach where it proceeds to touch down. Again, this is blocked airspace and crossing conflicts are not considered.

ENROUTE PHASE

Matrix savevalues have been incorporated in the model to handle the routings. Rather than use the standard FAA nomenclature, such as J80 for jet route number 80, each point of departure and point of termination in the network is connected by arbitrarily numbered routes. Thus, for example, Route Number 1 connects JFK undirectionally with Los Angeles International Airport. Along Route 1, there are segments of the route between VOR checkpoints, which are part of the jet airway system. These segments are independently numbered and can be used by any route. Each segment, starting at a VOR, is characterized by a distance in miles and an inbound course in degrees to the next VOR. In addition, airways may cross these segments at points along their lengths, and the distances to these crossing points for each segment are tabulated in matrix savevalues.

These segments are numbered serially in the reverse order to the progress of the flight, so that as one segment is completed, the GPSS program decrements to the next segment. Decrementation can be done readily in the GPSS language while incrementation (i.e., going from segment 1 to 2) is more involved and consumes more running time on the computer.

As each airway segment is traversed by a flight, the data on crossing airways are extracted from matrix savevalue storages to identify areas of potential conflict. That is, if aircraft are on the crossing airways at the same altitude, it is necessary to insure that safe separation standards are met at the points where their paths cross that of the flight under consideration. Initially, a search is made of the segment to determine if this aircraft will overtake the preceding aircraft, traveling along the same segment. If so, the speed is reduced to the value required to bring it in trail with the proper separation. Next, a search is made of the airways segments crossing the segment to determine if any traffic is flowing on them at the same altitude. If so, the matrix savevalue storages are called upon to obtain data on distances of each crossing airway along the airway segments from the VOR to checkpoint. With these distances available, the time that the aircraft under investigation would arrive at these crossings points can be computed, as well as the time the crossing aircraft is due to arrive, and these values computed. Safety criteria, in terms of time separation, is used to determine whether safe separation exists under the compared time values. If not, the aircraft that has been previously cleared to the crossing point is given the right of way, and the other aircraft delayed to compensate for the safe separation value.

Aircraft in ascent or descent are required to have safe separation through a multiple layer of altitudes during these phases of their operations. Since the airways are tiered, with traffic flowing in opposite directions on adjacent tiers, it is necessary for an aircraft changing altitude to verify safe separation at least at the next two altitude levels in its direction of change, to ensure that it can change to a level wherein aircraft are flowing in its direction of flight. That is, the aircraft under investigation could not simply check the first level through which it will pass and proceed there if safe separation obtains because it would be moving in a direction opposite to the flow of traffic at that altitude and could not continue flight at that altitude if safe separation could not be obtained at the succeeding level. The altitude difference to be traversed between each level is known (2000 ft. above flight level 290, 1000 feet below this flight level) and the aircraft rate of change of altitude in feet per minute during the ascent or descent maneuver is known so the time of crossing the altitudes under consideration may be computed. Correspondingly, since the aircraft's forward velocity is known during the ascent or descent maneuvers, the point at which the investigated aircraft will cross an altitude level on an airway segment may be calculated using the forward velocity and the time between altitudes. Here again, the existence of safe separation may be determined with aircraft traffic flowing at the intersected altitudes. If safe separation does not obtain, the aircraft changing altitude is required to delay until a safe passage can be made.

Delays to ensure safe separation and actions to avoid conflict can be accomplished in several ways. Variations in the delayed aircraft's speed can be made. The aircraft can be required to make a 360° standard rate turn, thus delaying it two minutes. It can be put in a holding pattern. It can be required to change altitude to avoid conflict, or can be made to go off the airway segment a safe distance, then return to the airway when the conflict situation has been resolved. At the present time, the program incorporates speed variations to institute delays for aircraft requiring them.

Three approaches were taken to develop conflict evaluation and avoidance logic for the enroute phase of the model. The first was to generate trigonometric functions describing the position of an aircraft along its flight path in x, y, and z coordinates so that aircraft positions and times would be calculable to analyze the possibility of conflicts. In this approach, the values of the trigonometric functions would be inputs to the program in tabular form, and the GPSS language would handle this process (in a time-consuming manner) by searching out the required value from those tabulated. In the second approach, a FORTRAN subroutine was developed to handle aircraft position changes and establish points of closest approach between aircraft. In the absence of the recently received IBM Help Routine, a means was developed to put the GPSS program on external disk storage when this subroutine was called into core for computation, with the computed result being left in a selected core area to be used by the GPSS program on its return to core after the FORTRAN program has run. While feasible, this approach is considered very time consuming. The third approach was to consider the aircraft to be operating solely on the controlled airway segments, such as high altitude routes, in which case conflicts would occur only with aircraft on the same or crossing airway segments. The problem can then be treated in terms of time and distance along the airway segment, and the times for aircraft to reach points of confluence compared. In the ascending or descending phases of aircraft operation, the vertical rate is used to determine times of intercepting transisted altitudes, and the aircraft's horizontal velocity used to determine its position along the airways segment. Since the data on the segmentation of the air routes and on their crossing segments are stored in matrix savevalues, they are readily obtainable and, under the assumption of controlled airway segment operation only, the need for the trigonometric solutions is eliminated.

ENVIRONMENTAL CONSIDERATIONS

The wind and weather submodel randomly generates the wind direction and weather conditions at each airport. The wind and weather are updated every 30 minutes. The weather factor is used in the spacing of aircraft during arrivals and departures while the wind direction is used to determine which runway to use. When the wind changes sufficiently to cause a runway change, the wind direction is forecasted for the next 30 minutes and checked, if the wind direction is to return, the runway is closed down, otherwise the queue of departing aircraft is sent to the new runway.

FLOW CONTROL

Flow control is the constraint imposed by the FAA to limit holding pattern delays at a destination airport by keeping aircraft on the ground at their origins when delays are expected to exceed one hour at their expected arrival time. A random number in the range from 0 through 9 is selected every 15 minutes of program time; selection of the digits 8 or 9 represents flow control and delays a scheduled departure until the next digit selection. The digits 0 through 7 represent no flow control effect. In a more realistic model, the simulated flow control should be a continuously updated feedback loop between destination and origin airports.

ELEMENTAL ATC SIMULATION

In keeping with the feasibility concept the model was formulated to determine the problem areas. Ten of the nation's busiest airports (Boston, JFK, Washington National, Atlanta, Cleveland, Indianapolis, O'Hare, St. Louis, Los Angeles, and San Francisco) were selected as origin and destination points. This was considered the minimum number to cause significant interactions. Ninety FAA-established high-altitude routes were used to connect these ten airports. High-altitude routings were selected since they are in positive controlled air space (above 18,000 ft. MSL in the Northeast U.S., and above 24,000 ft. MSL elsewhere). These were grouped by distance into short (<400 mi), medium (400-1500 mi), and long (>1500 mi) categories for aircraft altitude assignment purposes. The routes were segmented into 174 parts corresponding to segments on the routes between VHF omnirange (VOR) stations, with many of the segments being common to more than one route. Data on the routes, the segment distances, and the inbound magnetic headings to the VOR stations were stored in matrix savevalues. The FAA rules were incorporated for aircraft altitude assignment on these routes according to direction of heading with the altitude band of 29,000 ft. MSL to 39,000 ft. MSL being used. One hundred discrete aircraft, each with a unique tail number, were chosen as the population of the model, and these were apportioned among five airlines (UA, TWA, EAL, AA, and a fictitious airline) at 20 per airline. The aircraft differed in cruise speeds and rates of change in altitude, and had a common taxi speed.

Selected flight departures, scheduled in the Official Airline Guide in the time span 7:00 a.m. to 7:00 p.m., were used between these ten airports where actually assigned by the carriers, otherwise flights were assigned to the fictitious airline. Diurnal peaking of the flights was incorporated by means of the scheduling, with periods 7:00 a.m. to 9:00 a.m., and 4:00 p.m. to 6:00 p.m. being assigned the most flights. Since coast-to-coast flights were involved covering four time zones, all departure times in the program were converted to EST. It was decided not to generate any new flight departures after 7:00 p.m., but to continue the airborne portion until 10:00 p.m., to allow time for ATC interactions by those flights departing near the 7:00 p.m. departure time.

One hundred forty-four flights were scheduled for the 100 aircraft in this ten-airport network, some being one way, some one or more round trips, and one was a continued flight from JFK to LAX to SFO. Since each aircraft was uniquely identified by tail number, delays to flights encountered in the simulation then appear in a cumulative fashion, with the delay to one flight possibly causing delay to the next scheduled flight requiring that aircraft.

The items selected for parametric variation in the program were the number of airports in the system, the number of scheduled flights, and the weather. The effects on computer running time of these parametric variations were sampled in four runs of the program, and running time varied from about 5 to 11 minutes. The IBM 2250 Dynamic Display unit with light pen was used as an operator interface for the program in origination, debugging, making parametric changes, and displaying the results. This method of working with the computer is economical of programmer's time.

RESULTS

Data collected in the four computer runs were grouped under two headings, airport operations and en route. These results are to be considered as indicative of the types of data obtainable, rather than definitive in a quantitative sense, since this is an exploratory model. The data selected for collection on operations at each airport included delays and number of aircraft waiting at runway and pattern queues, and for takeoffs and landings. A sample of this data is shown in Table I. Data collected for en route phase included: flights waiting for aircraft, numbers of potential conflicts, maximum number of en route aircraft, and flow control effect on flights, as illustrated in Table II. It should be emphasized that any desired data, extractable by the program, can be accumulated and presented by the GPSS program numerically or graphically, either in hard copy or on the 2250 display.

CONCLUSION

The development effort showed that a GPSS program has provided a flexible and expandable model independent of the number of airports, aircraft, gates, flights, routes, or ATC constraints entered. It is portable, not tied to a particular computer or location. Increased complexity and sophistication in the program would lead to use of a larger faster computer. Some potential uses of this program, made more sophisticated and expanded into an operational model for an evaluation tool, include studies of airport and terminal congestion, airport use restrictions, flow control policies, aircraft spacing constraints, proposed new airports, variations in route segments, and area navigation concepts.

In conclusion, a simulation of the flow of air traffic in the United States for the purpose of evaluating alternative ATC concepts was implemented using the GPSS/360-NORDEN programming language. Four computer runs were made, and sample data were collected. It is felt that this model provides a strong base for straightforward extensions so that it can be used as an evaluation tool in a number of air traffic problem areas.

ACKNOWLEDGMENT

Inasmuch as this effort resulted from a working agreement between the FAA and The United Aircraft Research Laboratories, the authors would like to acknowledge the efforts of C. W. Burlin, United Aircraft Research Laboratories, as well as Stanley Price and Walter Felton of The Systems Research and Development Service, FAA.

		A/C WAITING AT RUNWAY			TOTAL NO.	A/C WAITING IN HOLDING PATTERN			TOTAL NO.
AIRPORT	RUN NO.	MAX. NO.	TOTAL NO.	AVER. WAIT (MIN)	TAKEOFFS	MAX. NO.	TOTAL NO.	AVER. WAIT (MIN)	LANDINGS
BOS	1	3	12	21.25	19	2	3	5.34	13
	2	3	11	22.62	14	-	-	0.00	6
	3	2	11	7.85	19	1	1	2.00	13
	4	3	12	21.96	19	1	2	12.04	13
JFK	1	4	13	27.46	20	1	1	2.25	13
	2	4	10	31.09	14	1	1	2.25	7
	3	3	13	13.98	20	-	-	0.00	13
	4	4	12	24.48	20	1	2	2.90	14
DCA	1	4	10	36.80	18	1	1	5.67	11
	2	4	10	36.80	12	1	1	5.67	7
	3	3	11	17.41	18	1	1	2.00	11
	4	4	10	34.89	18	1	1	5.67	12

TABLE I - AIRPORT OPERATIONS DATA

RUN	MAX. NO.	TOTAL NO.	AVERAGE WAIT (MIN)	TOTAL NO. POTENTIAL CONFLICTS	MAX. NO. AIRCRAFT EN ROUTE	MAX. DELAYED	TOTAL DELAYED	AVER. DELAY (MIN)
1	9	32	83.16	54	48	4	42	6.44
2	5	5	35.23	32	43	3	17	7.91
3	7	28	73.10	52	48	4	41	6.83
4	5	20	77.38	29	39	3	27	5.79

TABLE II - EN ROUTE DATA

REFERENCES

(a) Final Report on Air Traffic Control System Simulation Study, Norden Report 3818-R-0003.

(b) Exploratory Study Into The Development of An Air Traffic Control Simulation Model, C. W. Burlin, J. Reitman, D. Ingerman, AIAA Conference, Houston, October 1970.