

Airline performance modelling to support schedule development: An application case study

R. A. Haeme
Booz-Allen & Hamilton Inc.
1725 Jefferson Davis Highway
Arlington, Virginia 22202-4158

and

J. L. Huttinger
Booz-Allen & Hamilton Inc.
1953 Gallows Road
Vienna, Virginia 22180

and

R. W. Shore
Booz-Allen & Hamilton Inc.
1725 Jefferson Davis Highway
Arlington, Virginia 22202-4158

ABSTRACT

Flight schedules are the key to planning and executing an airline's operation. With each schedule, the airline defines its day to day operations and commits its resources to satisfying its customers' air travel needs. When developing a schedule, the airline aims to provide efficient service to its customers while maximizing its own profit. Key to both of these goals are keeping three concurrent processes on schedule and synchronized with one another. These processes are: the movement of each of aircraft along its assigned route, the movement of each crew along its assigned route and the movement of passengers from their origin to destination. This paper addresses a monte carlo simulation model which was developed and used to help one airline evaluate its on-time arrival performance. Using the model, the authors and airline operations planners were able to examine alternative strategies for maintaining high on-time performance without increasing costs.

INTRODUCTION

The dramatic growth in air travel which has occurred over the past several years has placed a severe burden on air traffic control and airport facilities. While the Federal Aviation Administration (FAA) is working hard to upgrade its facilities and train more air traffic controllers, air travelers continue to face increasing delays in both number and duration. The Department of Transportation (DOT) recently initiated action to encourage the airlines to change operational practices that contribute to air traffic congestion at high density airports during peak travel hours. DOT now publishes monthly reports comparing airline on-time arrival performance overall. In addition, the Official Airline Guide (OAG) now contains a numeric indicator which gives the percentage of time each scheduled flight has historically arrived on time. For example, a "5" associated with an airline's 2:00 PM daily flight from Boston to Washington National means that it arrives on-time 50-60 percent of the time. Officially, on-time is defined as within 14 minutes of scheduled arrival; only delays due to mechanical problems are not counted. DOT's actions have caused the airlines to take a careful look at their schedules.

AIRLINE SCHEDULING

Understanding the Problem

The problem for the airline is not trivial. The scheduler must serve several masters. Passengers, particularly business travelers, want to arrive at their destination by a particular time, e.g. to make that 10:00 AM meeting. Given nominal travel times between cities and passenger preferences, the airline's choice of departure times becomes quite restricted. Satisfying passenger demands leads to higher load factors which translate to profit for the airline. For

safety reasons, each aircraft must be inspected and maintained at specific intervals. To accomplish this, aircraft must be scheduled and routed to arrive at maintenance facilities when the aircraft is due for an inspection or maintenance. Crews may only fly for a certain number of hours per day, week and month. They must also have adequate time to transfer from aircraft to aircraft, eat meals and rest overnight. Ground personnel must have adequate time to service aircraft and transfer baggage. Customer service personnel must have time to deplane, transfer and board passengers. The scheduler faces a delicate balancing problem which has fewer feasible options than one might think.

External events create uncertainty for the scheduler. Competition generates traffic. To get their share of the market, the airline's competitors will also offer arrival times in the desired time window at popular destinations. Passengers from other originating cities also want to arrive at the popular destination city at the same time. Air traffic congestion is inevitable. Other random events such as weather, mechanical problems, runway repairs and air controller decisions make it impossible for a scheduler to predict the exact arrival time for a flight segment on a given day.

Considering the number of random events impacting on-time arrival performance of the airline, it is surprising that the airlines can run on schedule at all. Recent statistics show that the 14 major airlines, in fact, arrive on time between 75 and 90 percent of the time. That is an improvement over the first set of statistics published in October 1987 when on-time arrivals ranged from approximately 60 to 80 percent.

Solving the Problem

The airlines have taken varying approaches to solving this problem. The most common and most costly is to add time to the schedule. This approach was envisioned by DOT when they initiated reporting on-time arrival performance. It is a "truth-in scheduling" philosophy.

The basic element of the schedule is the flight segment. The scheduled time for any flight segment is called block time, or the time from pushing away from the gate at city A until docking at the gate at city B. Actual block time is the sum of three random variables. They are the taxi-out time, air enroute time and taxi-in time. Figure 1 provides a notional example, showing also the relationship of ground time to the segment:

turn - deplane, service, handle baggage, clean, cater, enplane

taxi-out - from push off gate to lift off runway

air enroute - from lift off runway to touch down on runway

taxi-in - from touch down on runway to docked at gate

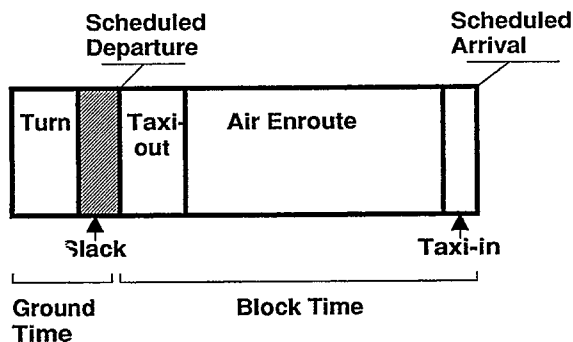


Figure 1: Schematic Representation of a Flight Segment.

The scheduler calculates the block time using historical data for those variables, the equipment type and cities involved in that flight segment.

Suppose that an airline flies from city A to city B and its schedule says that block time is 90 minutes. Figure 2 shows the cumulative distribution of time it has historically taken to fly the segment. In fact, the flight segment takes more than 105 minutes 41 percent of the time it is flown. According to DOT, the airline rates a "5" for that flight segment. To increase its rating the airline should increase its scheduled flight time. While it seems reasonable to add some time to the schedule, it is difficult to determine the amount of time to add. Adding 15 minutes will improve the rating to a "7". The aircraft will arrive in less than 120 minutes between 70 and 80 percent of times it flies the segment. It also means that the flight will arrive early 59 percent of the time. For 45 percent of the time, it will arrive more than 15 minutes earlier than scheduled. That creates its own set of problems such as meeting arrival slot times at slot controlled airports, scheduling gates and ground crews, and family members meeting passengers.

Moreover, adding time to the schedule adds to the airline's costs very quickly. Airline contracts usually call for paying the crews for scheduled or actual time in the air. In this case, the airline would be paying for more than actual time 60 percent of the times this segment is flown. Over half of those flights are more than 15 minutes less than scheduled. For this segment alone, costs can add up to over \$12,000. If an average airline, flying about 1000 flight

segments per day, added just one minute to every scheduled flight, it would add over 16 hours of flying. That is at least 2 additional crews and an additional aircraft. Neither are trivial costs.

No single flight segment is an independent event in terms of on-time arrival performance. Other events impact on each segment and segments impact on one another. Individual flight segments interact in three ways:

Flight segments interact with external events. Weather, mechanical problems, crew or passenger delays may all delay a flight's departure.

Flight segments interact with the next segment in its "line of flying". Each aircraft is scheduled in a line of flying, beginning at its originating station, flying a series of flight segments along a prescribed route to its final station for the day. When one aircraft arrives late, it may delay its own departure on the next segment.

Flight segments interact with segments in other lines of flying at connecting hubs. When inbound flights arrive late at a connecting hub, some outbound flights must delay their departure to wait for passengers and crews.

If there is insufficient slack time in a single aircraft's line of flying, one delay early in the day will cause the aircraft to be late all day. If that aircraft passes through the hub, it may transfer delay to several other aircraft.

Historical data exists to describe nearly all of the random events and variables involved. However, no simple formula exists that allows the scheduler to measure their complex interaction. The key to solving the airline scheduling problem is to recognize the random processes involved and make scheduling and policy decisions that minimize the risk of delays.

SIMULATION MODEL

To allow the scheduler to test a variety of scheduling strategies and operations policies which might impact schedule performance, we constructed a stochastic simulation of the airline's operation. The simulation model was built to represent the airline's entire "hub and spoke" operation. While simulating the airline's daily operation, the model collects and stores data that may be used to

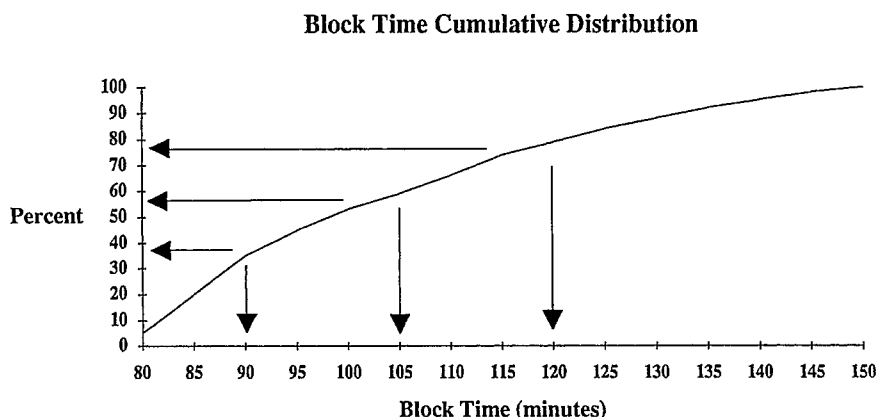


Figure 2: Hypothetical Distribution of Block Time between Two Cities.

evaluate the airline's schedule performance. It also identifies delay causes and duration. The model was designed with the flexibility to:

- fly any flight schedule
- fly any crew schedule
- fly to any number of cities
- fly any number of flight segments
- fly any number of aircraft
- fly any number of aircraft types.

In addition, the model incorporates the interaction with ground operations to include delays caused by:

- late originating crews
- late connecting crews
- late connecting passengers
- excess ground "turn" time
(ground and customer service)
- weather and air traffic control (ATC) hold.

The simulation model was programmed in FORTRAN on an IBM PC AT. Obviously, some limits apply to the "any number" statements above based on internal memory and disk storage capacity. The PC AT version can handle up to approximately 150 aircraft flying 1500 flight segments to 100 airports. The program should be easily portable to a larger machine if larger dimensions are required.

Conceptual Design

The model design stems from the concept of the airline flight schedule itself. The airline flight schedule is a sequenced list of flight segments, usually grouped by equipment type and line of flying. The route periodically passes through a hub, allowing inter-connections of passengers and crews. Each flight segment is a unique entity, flying a specific aircraft from an originating city to a terminating city at a specified departure and arrival time with a designated crew. Figure 3 lays out the conceptual flow of the model.

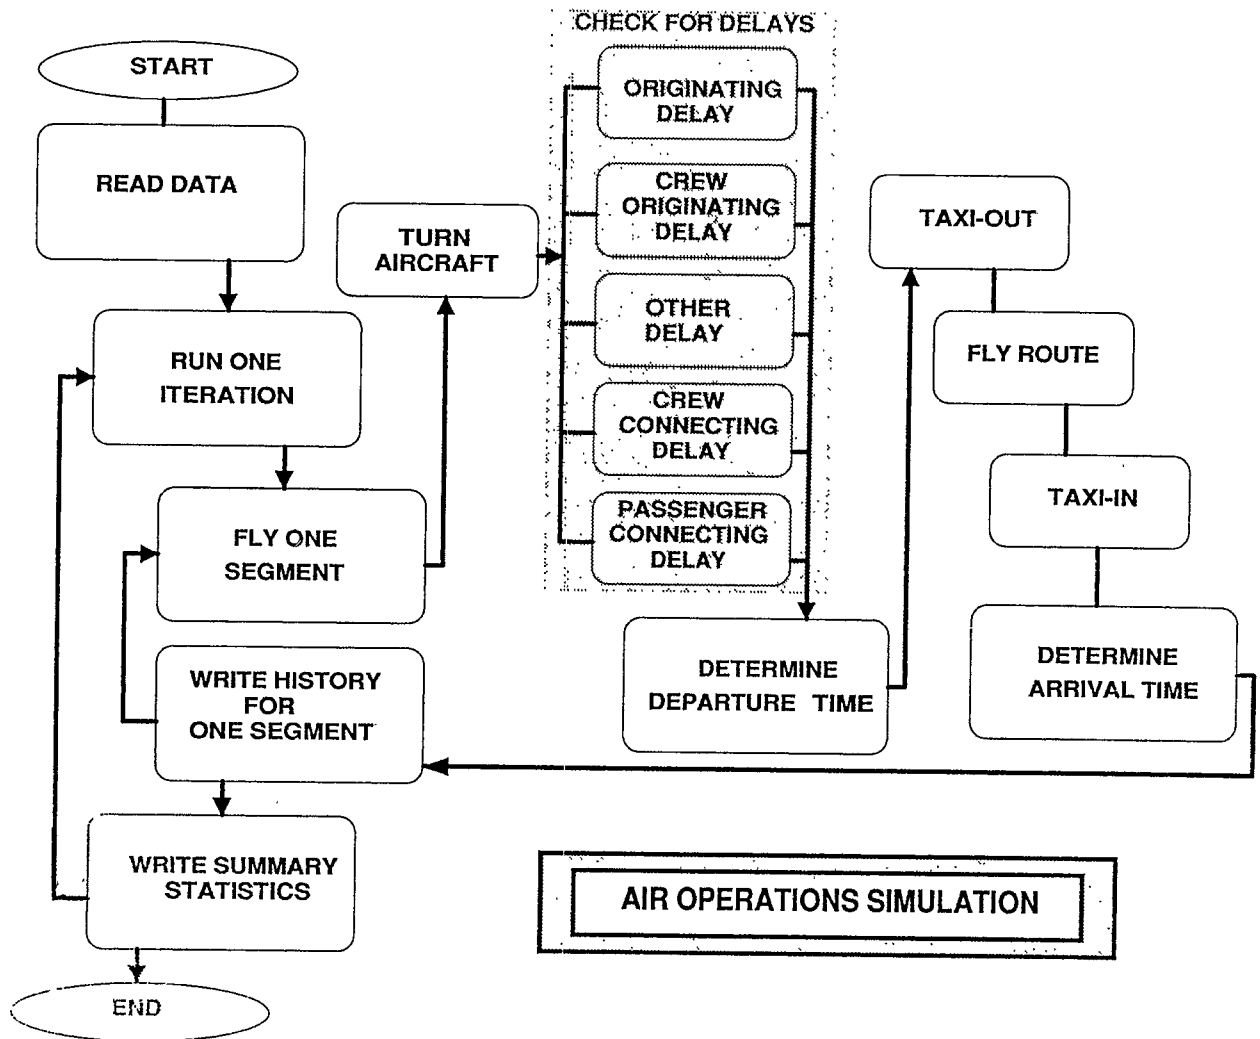


Figure 3: Simulation Model Logic Flow.

The deliberate synchronization and sequential nature of flight segments in the airlines flight schedule allowed considerable simplification in the simulation model design. Arranging flight segments in order of their departure time from earliest to latest obviated the need for any mechanism to keep track of event times. The model is able to fly one segment at a time, beginning with the earliest and ending with the latest. Crews and passengers are related to aircraft via the common segments on which they fly. The inter-relationship of segments occurs when the model calculates the actual departure time for a segment. To calculate the actual departure time for a segment, the model must first check for the actual arrival time of the aircraft, the crew and any connecting passengers from previously flown segments. Since the model flies segments to completion in scheduled departure time order, all information about that flight are available to later scheduled flights even when a flight actually flies much later than scheduled.

Logic Flow

As the model encounters each flight segment, it first determines when the aircraft was available to begin processing for this flight. That is, its last actual arrival time. It then draws from a historical distribution of turn times for the originating station and aircraft type to establish an actual departure time. The model then checks for other times which may contribute to delay of the aircraft past its scheduled departure. These include originating station, originating crew, connecting crew, connecting passenger, weather, and ATC delays as appropriate for the current departure station. The actual departure time is set at the actual arrival time plus the longest delay or the scheduled departure time, whichever is later.

Next the model draws a taxi-out time from a historical distribution for the departure city, and for that time of day if the city is a high traffic city. The time of day will determine whether the time is drawn from a peak time distribution or a normal time distribution at high traffic cities. Then the model draws an air enroute time from a historical distribution for the city pair defined by the segment. If the destination city is high traffic city and the time of day it will arrive is a peak period, the model will draw from a peak time distribution as before. Finally, the model will draw a taxi-in time from the destination city's historical distribution. Once again, if the destination city is a high traffic city and the time of day it will arrive is a peak period, the model will draw from a peak time distribution. The sum of the three times drawn added to the actual departure time for the segment will yield its actual arrival time.

Data

Table 1 describes the nature of data read and used by the model. In general, data files containing historical time distributions were formatted with times in deciles. Parameter files allow this to be changed to any set of desired percentile intervals.

The crew data file presented an interesting challenge. Crew "trips" are scheduled to coincide with each flight segment, obviously, but not with each aircraft's line of flying. Crew trips are scheduled to maximize crew utilization while conforming to numerous constraints which result from both regulation and contracts. A crew trip consists of a set of flight segments over several days. In some cases, a crew trip requires the crew to "deadhead", that is, travel as passengers on an aircraft, to get to its next segment's originating station. Therefore, an aircraft may have to wait for more than one crew before departing. In any case, it was necessary to cull out crew trip data which corresponded to the same day as the line of flying schedule.

Passenger connection data was derived from data used at the hubs to control minute by minute operations. This data is not normally saved, but, for this effort, a brief history was accumulated. This data file defines the connecting bank times for each hub city. A connecting bank is a time interval at a hub city when several of the airlines flights arrive and depart allowing passengers to connect to as many different destinations. This data also provides a profile of the average number of passengers connecting from each flight inbound to each flight outbound from the connecting bank. In addition, the file incorporates a decision rule for each connecting bank which determines how long the model will wait for passengers. The decision rule considers number of passengers, the difference between arrival and anticipated departure time, and the wait time allowed. Initially, the rules were set to represent the airlines operational practices. To test the impact of the rules on schedule performance, they were modified in various cases.

RESULTS

The analysis of an airline's schedule performance using this model provides some useful insights. The analysis should focus on the the high leverage drivers of schedule performance. First, uncontrollable factors are eliminated from consideration. The airline can't control the weather, a temporary runway closing at an airport in their network, or predict when maintenance problems may occur. Factors the airline can control fall into two major categories; operating procedures and schedule design.

Operating Procedures

Analysis of the airline's performance data will show that some delays can be attributed to operating procedures. An example is originating delay. That is, the first flight segment of the day for some lines of flying typically depart late. Originating delay concerns the airlines because it can impact the entire line of flying. While it was possible to calculate the immediate impact of originating delays, it is not possible to calculate their impact on cumulative delay. If a late originating aircraft has no slack in its down line schedule, it will continue to be late. If that aircraft enters a connecting bank, it can pass its lateness on to other aircraft. The simulation model makes it possible to see the total impact of late originating aircraft. Of course, the result can range from insignificant to significant depending on the status of other control factors.

Many of the operating procedures that impact an airline's schedule occur on the ground. As previously discussed, ground service crews fuel the aircraft and transfer baggage; other service teams clean and cater the cabin after passengers have deplaned; and customer service personnel help passengers connect to other flights and board new passengers. Because this turn time falls within the scheduled ground time at a station, which may include slack, its impact is hard to predict. Using the simulation model, it is possible to see the impact of adding or reducing a few minutes from turn times. Under certain schedule circumstances, it appears that reducing turn time can pay big dividends in overall schedule performance. A few minutes at each stop adds up to significant slack time buffers, allowing the airline to catch up when other delays occur. Better still, reducing time from aircraft servicing and customer service operations ought to be easily managed and very low cost to achieve. The simulation model provides a means to measure the performance improvement gained per minute of turn time reduced. The airline must weigh the performance gains against real-world ability to reduce their turn time standards.

Table 1: Sample Data Used in Air Operations Simulation Model

AIR OPERATIONS SIMULATION DATA

DATA NAME	FORMAT	PURPOSE
Line of Flight	LOF FLT ORIG DEP TERM ARV 7303 600 AAA 50 BBB 134 7306 620 AAA 150 CCC 615	Describe schedule by segment, in order of departures.
Crew Trip	BASE/TRIP FLT ORG TRM DEP ARR CCC73 4552 600 AAA BBB 50 134 CCC73 6443 620 AAA CCC 150 615	Describe crew trips by segment for same LOF schedule.
Enroute Time	DEP ARR EQ P/N 10% 20% . . . 100% DDD EEE 73 N 12 13 28 DDD CCC 73 N 42 45 98	Discrete distribution of flight times in deciles for each city pair by type equipment peak or non peak arrival time.
Taxi in Time	ARR EQ P/N 10% 20% . . . 100% DDD 73 N 2 3 8 CCC 73 N 2 5 18	Discrete distribution of taxi in times in deciles for each city by type equipment peak or non peak arrival time.
Taxi out Time	DEP EQ P/N 10% 20% . . . 100% DDD 73 N 5 8 12 CCC 73 N 8 10 18	Discrete distribution of taxi out times in deciles for each city by type equipment peak or non peak departure time.
Turn Time	DEP EQ 10% 20% . . . 100% CCC 73 26 28 68 CCC 73 25 26 63	Discrete distribution of turn times in deciles for each city by type equipment.
Peak Times	CTY STRT END FFF 0600 1000 FFF 1600 2000	Designation of heavy traffic cities and definition of peak times at each.
Passenger Connections	CCC 0600 0800 0830 0855 1 TO FROM FROM . . . DEST FFF GGG HHH NO. PAX 6 12	Definition of connecting bankcities arrival and departure windows. For each outbound, a profile of average connections from inbounds.

Similarly, discrete distributions for delays due to originating crews, originating flights, crew and passenger transfer between gates, and other gate delays have been placed separately in appropriate data files .

Performance Improvement vs Added Block Time

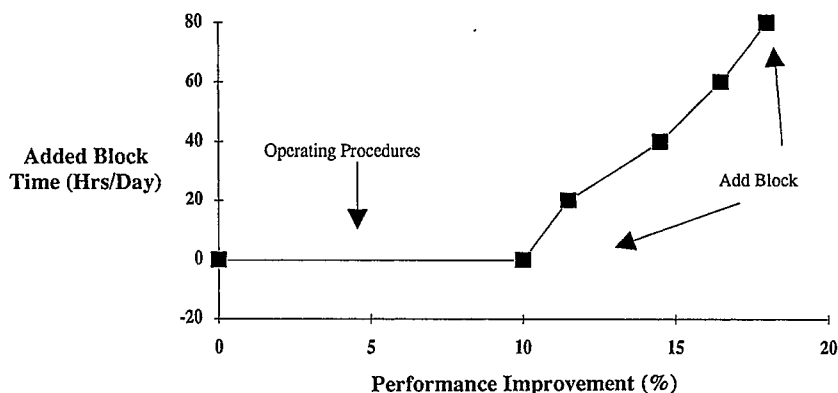


Figure 4: Performance Improvement vs. Added Block Time.

Schedule Design

Schedule design involves establishing a consistent rule for selecting the correct amount of time to allocate to each flight segment. Given the historical data the airlines have available for flight and taxi times, the scheduler must determine which parameter of the block time distribution will minimize late arrivals without adding time to the total schedule. For example, the scheduler might use the mean, median, mode or a specified percentile of the distribution. Using the simulation model, it is possible to test sensitivity of overall schedule performance to the schedule time parameter, by setting the parameter at several levels to observe its impact.

Figure 4 shows the relative improvement in schedule performance obtained in one schedule by first reducing turn time then adding increments of block time as shown. It is possible that a substantial improvement can be attained for virtually no cost. It is also possible to see that adding block time provides a less significant improvement in performance. That is not to say that an airline should not or would not add block time, particularly in instances where customer satisfaction and or competitive posture may be improved.

What is not obvious is the fact that adding block time here and there may not improve the schedule performance at all. This approach will certainly improve performance of the segment involved, but it will not necessarily improve the overall schedule performance. The only certainty is that it will cost more. In fact, adding block time to a few problem segments may make the overall schedule performance worse. It is difficult to predict the interactions that occur down line particularly at connecting banks. The simulation model provides a very effective means to test the schedule changes and to predict the overall performance outcome.

CONCLUSION

The simulation model proved to be an effective tool for evaluating airline schedule performance. The complex interactions of airline operations defy prescriptive modeling. Airlines must make frequent changes to their schedules. Replacing their fleet with more modern equipment, adding routes and changing schedules to meet customer demands, and adjusting to changes which occur industry-wide can all impact an airline's schedule. In this environment, a descriptive, stochastic simulation model of the events and interactions that take place, provides a useful laboratory for an airline operations planner and scheduler.

REFERENCES

- Fishman, G.S. (1978). *Principles of Discrete Event Simulation*. John Wiley & Sons, New York.
- Hammond, R.H., Rogers, W.B., and Crittenden, J.B. (1987). *Introduction to FORTRAN 77 and the Personal Computer*. McGraw-Hill Book Company, New York.
- Hillier, F.S. and Lieberman, G.J. (1986). *Introduction to Operations Research*. Fourth Edition. Holden-Day, Inc., Oakland, California.

AUTHORS' BIOGRAPHIES

R. A. HAEME is an associate in the Analysis and Computational Technology Practice at Booz-Allen & Hamilton Inc. in Arlington, Virginia. He received a B.S. in Chemistry from Bowling Green State University, and an M.S. in Industrial Engineering from Ohio State University. His current interests include applying modeling, gaming, and expert systems techniques to problems in logistics. He is a senior member of IIE and a member of ORSA.

Booz-Allen & Hamilton Inc., Crystal Square 2, Suite 1100,
1725 Jefferson Davis Hwy, Arlington, Virginia, 22202-4158,
(703)769-7700.

J. L. HUTTINGER is an associate in the Decision Support Systems Practice at Booz-Allen & Hamilton Inc., Vienna, Virginia. He holds a B.S. in Computer Science from the University of Dayton and an M.S. from American University and an M.S. from George Washington University. His current interests include embedded operations research techniques in expert systems and executive information systems.

Booz-Allen & Hamilton Inc., Suite 600, 1953 Gallows Rd.,
Vienna, Virginia, 22180, (703)893-0040.

R. W. SHORE is a senior associate in the Analysis and Computational Technology Practice at Booz-Allen & Hamilton Inc. in Arlington, Virginia. He received a B.S. in Mathematics from Harvey Mudd College and a Ph. D. in Statistics from Carnegie-Mellon University. His interests include symbolic processing and applying advanced computer technology to intelligence analysis. He is a member of ACM, IEEE and the American Statistical Association.

Booz-Allen & Hamilton Inc., Crystal Square 2, Suite 1100,
1725 Jefferson Davis Hwy, Arlington, Virginia, 22202-4158,
(703)769-7700.