THE INTERNATIONAL JOURNAL OF AVIATION PSYCHOLOGY, 3(4), 327–343

En Route Air Traffic Controllers' Use of Flight Progress Strips: A Graph-Theoretic Analysis

O. U. Vortac, Mark B. Edwards, and Judi P. Jones
University of Oklahoma, Norman

Carol A. Manning
FAA Civil Aeromedical Institute
Oklahoma City

Allen J. Rotter
Federal Aviation Administration Academy, Oklahoma City

In the United States, flight data are represented on a paper flight progress strip (FPS). The role of the FPS has recently attracted attention because of plans to automate this aspect of air traffic control (ATC). The communication and FPS activities of controllers were classified as they controlled air traffic of varying complexity. Transition networks were derived from the empirical transitions. These networks indicated that several aspects of ATC generalize across complexity including the centrality of writing on the FPSs and the control of traffic. Complexity affected when FPSs were used with high-complexity traffic situations that required the controller to devote uninterrupted periods of time to the management of the FPSs rather than integrate these board management responsibilities with the responsibilities of separating aircraft.

In the United States, air traffic control (ATC) of high-speed and high-altitude aircraft en route between takeoff and landing is currently accomplished with three tools: radar, communication devices (radio and telephone), and a representation of flight information. Flight information is displayed on a flight progress strip (FPS), a rectangular piece of paper divided into 31 logical fields, each of which displays

---

Requests for reprints should be sent to Francis T. Durso, Department of Psychology, University of Oklahoma, Norman, OK 73019–0535.

particular information available from the flight plan. When a flight enters the volume of airspace (sector) for which a controller is responsible, the flight is considered active, and the controller moves the corresponding strip from a suspense bay where the imminent entries to the sector are held to the active bay; thereafter the controller interacts with the corresponding FPSs by (a) moving the strips within the active bay, (b) writing on the strip itself, and (c) looking at the strip to acquire or confirm flight information. The organization and upkeep of the suspense bay and the active bay is referred to as *board management.*

The responsibility for board management depends on the number of controllers assigned to control a sector. When a sector is handled by a team of two, the radar controller is primarily responsible for observing the radar screen and for talking to the pilot, whereas the data controller (sometimes called the *radar associate*), seated next to the radar controller and in front of the strip bays, is usually responsible for the FPSs. On several occasions, traffic loads permitting, both functions are assigned to one controller, as they were in the study described herein. This requires the controller to integrate board management with aircraft separation responsibilities.

This research focused on the FPS for two major reasons. First, although controllers must use FPSs to maintain a legal record of control actions as dictated by Air Traffic Control Publication 7110.65 (U.S. Department of Transportation, 1989), it is otherwise unclear why or when the controller uses the FPSs. Anecdotes from controllers often support the notion that the FPS is viewed as unimportant. This perceived lack of importance was reinforced by the results of a structured interview conducted by Human Technology, Inc. (1990). Controllers were asked about their priorities of activities under normal workloads. On a scale ranging from *lowest priority* (1) to *highest priority* (9), experts gave reviewing FPSs a 6.0 and writing on FPSs an 8.8.

On the other hand, observers of the ATC situation have advanced compelling arguments that the FPS, even if used only because of legal requirements, may provide coincidental but substantial benefits to the controllers (e.g., Hopkin, 1988; Means et al., 1988; for a review, see Vortac & Gettys, 1990). Hopkin (1988) believed that the relatively effortless incidental encoding of flight data that results from following the legal requirements helps to build understanding and memory. For example, when an aircraft is given a new speed, the old speed is crossed out and the new speed written on the strip (Weston, 1983), making the strip more distinctive and presumably easier to locate and remember. However, it is clear that not all the interaction with the strips is legally required. For example, the controllers may offset a strip from the bay as a reminder to take a future action.

A second reason to focus on the FPS is that the nation's ATC system will undergo a period of radical and unprecedented change in the next decade. The Advanced Automation System (AAS), to be phased in during the 1990s, will have substantial ramifications for all aspects of ATC (e.g., Ammerman & Jones, 1988). The first stage of AAS in the en route environment involves introduction of the Initial Sector Suite System (ISSS). The area of greatest change due to ISSS involves the way in which flight information is displayed and manipulated. The paper FPS will be replaced by electronic flight data entries. It is the automation of the FPS that has attracted the attention of aviation psychologists. If the introspections of controllers are correct, then the automation of the strips will make little or no difference. If these introspections are incomplete or incorrect, and the observations offered by human factors experts turn out to be true, automation could have a substantial impact.

Unfortunately, little empirical evidence is available that describes when FPSs are typically used. Work by Standard Technology, Inc. (1990) suggests that FPS usage does not bear a simple relation to air traffic complexity. Although strip activity was weakly predicted by the time aircraft spent in a sector for the two-controller situation, no scenario characteristic was predictive of strip activity for the single-controller situations. Nevertheless, other actions taken by the air traffic controller (e.g., the commands he or she issues) and other types of information received by the air traffic controller (e.g., communications from aircraft and other ATC facilities) may illuminate ways in which the FPSs are used.

For example, Buckley, DeBaryshe, Hitchner, and Kohn (1983) measured 28 potential indices of ATC system performance, including a variety of behavioral measures as well as various radar-related (plan-view display) indices (e.g., aircraft separation, time under control for each aircraft, etc.). A factor analysis revealed that a set of four factors could account for most of the variance between scenarios of different complexity. Three of these factors (Confliction, Occupancy, and Delay) summarized the plan-view-display–related technical measures and were difficult to observe, whereas the fourth factor (Communication) subsumed behaviors of the controller. Thus, looking at the relation between communication events and FPS activities seems an appropriate first step.

The goal of this study was to improve understanding of when controllers currently use FPSs to control traffic in an en route environment. Potential implications of automation (Wise, Hopkin, & Smith, 1991) can only be assessed in comparison to a well-understood existing system. The present study utilizes an observational methodology and focuses on data obtained from a sample of controllers who each worked individually, fulfilling both radar- and data-control functions. The data involve the online classification of communication events (i.e., controller commands, controller queries, pilot requests, and sector transitions) and FPS activities (i.e., looking at FPSs, writing on FPSs, and manipulating FPSs). These data were used to construct transition matrices that summarized, for example, how often a look at the FPSs was followed by a controller command.

However, attempts to determine the behavioral structure that underlies ATC from the raw transition matrices is unsatisfactory for at least two reasons. First, the transition matrices are quite complex. Second—and more important—the matrices do not distinguish between those transitions present in the data that reflect the latent structure of controlling air traffic and those present in the data because of random noise. What is needed is a way to reduce the complexity by eliminating those transitions that merely reflect noise in the data, so that the remaining transitions will provide structural insights. By assuming that the underlying structure is a graph (in the mathematical sense), we can distinguish between those transitions that are necessary to the structure and those that are not.

A *graph* is a formalism in which the concepts (e.g., controller command) are represented by a *node* and the transitions (e.g., controller command to write on FPS) are represented by *arcs* connecting nodes. The graphs for these data should be directed (i.e., event A leads to event B) because the transition matrix is asymmetric. Thus, the proportion of transitions from "write on FPS" to "look at FPS" need not be the same as the proportion of transitions from "look at FPS" to "write on FPS." If the arcs are weighted so that some arcs are traversed more frequently than others, the resultant formalism is a *network*.

One node can be connected to another in a network either directly—by an arc from that node to the other—or indirectly—by a *path* through other nodes. If the path between two nodes is the most efficient (shortest) path, then that path is the *geodesic* path. The representation of all such geodesic paths yields a geodesic network that represents the shortest distance between all pairs of nodes (see Schvaneveldt, Dearholt, & Durso, 1985). This is important because the geodesics of a graph meet several assumptions—for example, the distance from A to B plus the distance from B to C cannot be less than the distance from A to C. The geodesics of a network are unchanged if paths that violate the triangle inequality are omitted. In most modern scaling procedures, violations of the triangle inequality are assumed to be due to random noise and distortions, not to valid indicators of the underlying structure.

In an effort to reveal the underlying structure in our transition matrices, the Pathfinder scaling algorithm was selected (Schvaneveldt, 1990; Schvaneveldt & Durso, 1981; Schvaneveldt, Durso, & Dearholt, 1989). The Pathfinder analysis represents the relations among events graphically so that the underlying structure in the transition matrix can be more readily interpreted. The algorithm reduces a matrix of proximity data (e.g., transitions) by eliminating those connections that do not satisfy the metric properties of a network. Thus, the connections remaining are those that are ordinally necessary (Hutchinson, 1989).

The algorithm has been successfully employed in a variety of domains within cognitive psychology, engineering, and artificial intelligence (see Schvaneveldt, 1990). For example, Pathfinder has been used to articulate the structure of natural categories (Durso & Coggins, 1990), to distinguish between expert and novice fighter pilots (Schvaneveldt et al., 1985), to predict free recall (Cooke, Durso, & Schvaneveldt, 1986), to develop menus for automated cockpits (Roske-Hofstrand & Paap, 1986), and to establish connections in hypertext (McDonald, Paap, & McDonald, 1990). The mathematical foundations can be found in Schvaneveldt et al. (1985).

METHOD

Subjects

Nine Full Performance Level (FPL) controllers participated. They had served as en route FPLs from 3.5 to 9.3 years ($M = 5.7$ years) and had last been in the field

2 to 16 months before their participation (M = 7.4 months). The study was conducted at the Radar Training Facility at the Federal Aviation Administration (FAA) Mike Monroney Aeronautical Center in Oklahoma City, which can provide high-fidelity en route traffic simulations using the fictitious AERO Center airspace used in training. Because subjects had to be familiar with AERO Center, but naive to the particular selection of scenarios, we used FAA Academy instructors in the nonradar screen program.

Scenarios

All subjects were observed under low, medium, and high levels of complexity, according to a randomized counterbalancing schedule. Across subjects, numerous different scenarios were used. Therefore, conclusions drawn for a level of complexity are unlikely to be due to a particular scenario. We used existing scenarios as outlined in the *FAA Academy Scenario Guide* (U.S. Department of Transportation, 1990). The scenarios in the guide varied in complexity from 25% to 95%. Complexity was measured using the complexity worksheet found in the *Instructional Program Guide* (Appendix B, Section 3, Phase 8A for nonradar). Complexity is computed in the guide in the following way: (a) Departures received 5 points; (b) arrivals, en route aircraft needing a control action, emergencies, and radio failures each received 4 points; (c) special flights received 3 points; (d) en route flights not needing a control action received 2 points; and (e) each additional coordination action (e.g., a point-out) received 1 point.

All but the most complex scenarios represented a level of traffic density that, in the field, could be handled by a single FPL controller. The high-complexity scenario was comparable to a situation in the field in which a supervisor would provide or a controller might request a data-side controller to assist. Table 1 summarizes the scenarios used in this study.

Behavioral Categories

The onset of four types of communication events and three types of FPS activities were coded. Communication events were categorized into controller commands (CCOM), controller queries (CQUERY), pilot requests (PREQ), and sector transitions (SECTOR). FPS activities were categorized into looking at FPSs (LOOK), writing on an FPS (WRITE), and manipulating FPSs (MANIP).

TABLE 1
Summary of Scenarios for Different Complexities

Scenario	*Complexity (Percentage)*	*Departures*	*Arrivals*	*Overflights*	*Length (Min)*
Low	50	2.4	3.8	4.0	30
Medium	75	3.0	5.6	5.2	30
High	95	4.2	8.6	18.6	60

Communication Events

CCOM. CCOM involved one or several of the following six subcomponents: (a) change route, (b) change speed, (c) change altitude, (d) dispense information, (e) issue clearance, and (f) other. The "other" category was rare and included instances such as when the controller called an adjacent facility to request notification when a military flight was on the ground.

CQUERY. CQUERY was a controller-initiated request for information from the pilot. A CQUERY occurred whenever the controller asked a pilot to report (a) aircraft speed, (b) altitude, (c) route, or (d) other. The "other" subcategory tended to be requests for a verification of aircraft callsign or characteristics (e.g., "heavy").

PREQ. PREQ was a pilot-initiated request of the controller. We coded six possible requests from the pilot: (a) a change in route, (b) speed, (c) altitude, (d) information, (e) clearance (direct clearance to destination), and (f) other. Some of the "other" requests concerned air refueling, among other things.

SECTOR. Another event of interest involved interactions between the controller and adjacent centers and other ATC facilities, in particular when aircraft are entering or exiting a sector. Three components were coded: (a) departures, (b) handoffs, and (c) initial contact. A *departure* was a communication initiated by another ATC facility, a *handoff* was initiated by the controller, and *initial contact* was initiated by the pilot. The distinction between the PREQ category and the initial contact subcategory of SECTOR was evident in the content of the pilot's communication.

FPS Activities

To examine the relations between the events just described and controller activities relating directly to the FPS, the following FPS activities were coded.

LOOK. We coded LOOK whenever the controller looked at the suspense or the active bay. Because LOOK obviously preceded other FPS activities, this category only included those looks at the FPSs that were not immediately followed by writing or manipulating—those events were instead coded as WRITE or MANIP. In addition, multiple looks at the FPSs implied that the controller looked away (presumably to the radar display) and returned to look at the FPSs. Thus, a long, single look at the FPSs, or a search through the strips, was coded as one LOOK.

WRITE. We coded WRITE whenever the controller verified or changed an entry on an FPS. Verification involved one or more of the following four symbols: (a) "D" for departure, (b) "R" for radar contact, (c) a check mark comment when a new altitude was achieved, and (d) "C" for handoff to the next sector. Change involved altering markings on the FPS, such as a revision of altitude or change in route.

MANIP. MANIP of an FPS was any physical contact with the strip that did not involve writing. Five categories of MANIP were coded: (a) moving a strip from the suspense to the active bay when a flight entered the airspace, (b) sequencing the strips within the active bay, (c) offsetting or "cocking" a strip as a reminder, (d) flattening or removing the offset, and (e) tearing down or removing the strip from the active bay when a flight leaves the airspace.

Procedure

Subjects first completed a brief background sketch. They were then given the opportunity to organize the strip bay in preparation for the scenario. Subjects were provided with all the strips for the problem at this time. In our situation, the strip bay was located to the right of the radar screen. The two observers sat behind and to the right and left of the controller, with notebook computers on their laps. The computers were used for online data collection. Both computers were synchronized with the clock on the radar screen, and observers coded behaviors by pressing different keystroke combinations as they occurred, yielding a time-indexed behavioral record of each scenario. Simultaneous events were codable, but appeared as sequences in the data trace. Coding errors could be corrected online by escaping from the submenus.

In addition to the observers, two ghost pilots were required to control the planes, and another assumed the communication functions of adjacent centers and other ATC facilities. Radio and audio communications were recorded using a multitrack cassette recorder. One input channel included communications of the controller, the pilots, and the center. Each observer wore a lapel microphone and was recorded on their own input channel, which allowed them to annotate orally their event-recording.

Subjects were not informed about the emphasis placed on FPSs in this study, but were told to control traffic normally. Each experimental session lasted approximately 3 hr (two 30-min scenarios, plus 1 hr for the most complex scenario). Break periods between scenarios were approximately 20 min.

RESULTS

Any coding errors committed by the observers that were noted on the audiotape were corrected prior to analysis.

Event Rates

Events per minute were computed for each of the seven classes of events. (Calculating events per minute took into account that, although the high-complexity scenario had twice as many aircraft, it was also twice as long; see Figure 1.) An analysis of variance (ANOVA) was conducted for each event class, with a test-wise

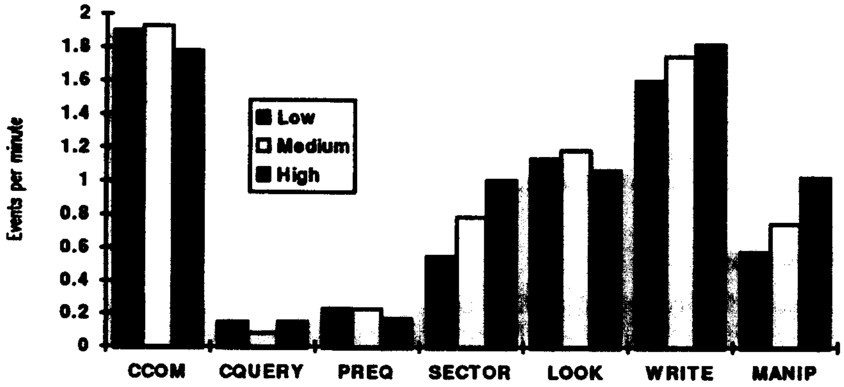

FIGURE 1 Events per minute as a function of scenario complexity for all classes of events.

alpha of .007 to yield an experiment-wise alpha of .05. As would be expected, the rate of SECTOR increased with complexity, $F(2, 16) = 23.88$, $MS_e = .02$, $p < .0001$. The most frequent events, CCOM and WRITE, tended to occur at the same rate regardless of complexity. In fact, the only other event to vary significantly with complexity was MANIP, $F(2, 16) = 15.51$, $MS_e = .03$, $p < .0002$, which increased until the high-complexity scenario, in which manipulations occurred once per minute. At the other extreme, PREQ and CQUERY were relatively rare, never reaching >4% of the recorded events, and not revealing any significant changes with complexity.

Transition Analyses

A 7 × 7 matrix of transitions from one event to another was computed for each subject, normalized by dividing by the total number of events, and weighted by the temporal separation between events. The temporal weighting was accomplished in the following way. When compiling the frequency of transitions, the increment to the frequency count was a negative exponential function of the elapsed time. Specifically, the time-weighting function was:

$$I = e^{-\lambda s} \tag{1}$$

where s refers to the number of sec between the two events and λ is a decay parameter, with larger values of λ corresponding to more rapid decay, or increasingly less impact from more time-distant events. The increment I was used as the increment to the transition count. All analyses used a value of $\lambda = 0.1$, which yielded a time-weighting function of the desired characteristics. Thus, two events that occurred simultaneously ($s = 0$) would produce an increment of unity ($I = 1$), two that were separated by 1 sec would be given an increment of .90, and so on. Sequential events separated by 20 sec or more contributed little to the transition count between those two events ($I = .13$). All analyses to be reported were conducted on these normalized and time-weighted transition matrices. Interobserver reliability for these transition matrices ranged from .75 to .90.

Individual Differences

Before creating a description of the average ATC, the presence of individual differences between controllers was addressed (with regard to training issues, see discussion by Smith, 1991). In the worst case, if variability between controllers exceeds variability between different levels of the situational variables, little meaningful generalization of results is possible. Transition matrices were averaged across subjects for each level of scenario complexity, and individual matrices were then correlated with the average matrices. All the correlations between the individual matrices and the average matrices were moderate to strong ($r = .62$ to .97). The number of correct classifications appears in Table 2. For high- and low-complexity scenarios, classification was quite good, with no individual being classified at the other extreme. On the other hand, the individual medium-complexity scenarios did not classify well, suggesting that the medium-complexity scenarios shared properties of both the low- and high-complexity scenarios. This lack of uniqueness for the medium-complexity scenario manifested itself in other analyses.

Pathfinder

We submitted the normalized and time-weighted transition matrices to the Pathfinder algorithm. The algorithm produces a network in which the communication events and the FPS events are represented by nodes and the transitions are represented by arcs between the nodes. The weight of the arc reflects the frequency with which that transition occurred.

Pathfinder is capable of producing a family of networks depending on the metric used to determine path distance (Minkowski r parameter) and the maximum path

TABLE 2
Number of Individual Transition Matrices Most Highly Correlated With Each Group Transition Matrix

	Group		
Individual	*Low*	*Medium*	*High*
Low	7	2	0
Medium	2	4	3
High	0	2	7

length for which the triangle inequality must be satisfied (q parameter). The Minkowski distance measure is presented here:

$$w(P) = \left[\sum w^r\right]^{1/r} \tag{2}$$

where $w(P)$ is the weight of the path, w is the weight of an arc on the path, and r is the Minkowski exponent. For example, when $r = 2$, the distance is the common Euclidean distance, for which the distance between two points is given by the Pythagorean theorem.

The Minkowski r distance can be used as a generalized distance function for computing distances in a graph or network. For example, when $r = 1$, the distance between two nodes is the sum of the distances along the existing paths (the city block metric). When r is infinity, in the limit the distance becomes the maximum (the dominance metric).

The particular application of Pathfinder that we used was one guaranteed to produce the simplest network, the *minimal cost network (MCN)*. This involved setting the r parameter to infinity, thus causing Pathfinder to compute distance using the dominance metric; and setting q to $k - 1$, where k is the number of nodes. This forced Pathfinder to eliminate violations of the triangle inequality in paths of any length (cf. Hutchinson, 1989). Thus, q was set to 6.

The mean 7 × 7 transition matrices were submitted to the Pathfinder algorithm for each of the three levels of complexity. The resultant MCNs appear in Figure 2. The arcs were supplied by Pathfinder and the width of the arcs reflects the weight of the transition, with thicker arcs having occurred proportionately more often than thin arcs. There were six widths corresponding to six proportion categories. The figures are further augmented by the proportional frequencies of occurrence of the events as represented by the size of the nodes: Events that occurred proportionately more often are represented by larger nodes. Note that in contrast to multidimensional scaling solutions, in Pathfinder networks the physical distances between nodes in the depiction is meaningless. All that matters is whether the nodes are linked or not, the direction of the arc, and how "strong" that arc is.

Visual inspection of the MCNs revealed that all of the networks are greatly simplified as compared with the original input matrices, which contained 49 transitions. The graphs also look remarkably similar, both qualitatively and quantitatively. Qualitatively, a number of arcs appear in all three networks. Quantitatively, often traveled arcs (thick lines) in one network tend to appear in the others. In fact, the medium-complexity MCN contained only one arc (a pilot request loop) not present in either the high- or low-complexity MCN. As in the individual-to-group correlations, we see that the graph representations of the medium-level scenarios were composites of both the high- and low-complexity scenarios.

Computation of various graph-theoretic measures (see Table 3) confirm that—at least at a macroscopic level of analysis—the networks are quite similar. The networks have a comparable number of arcs, and the most distant path (the

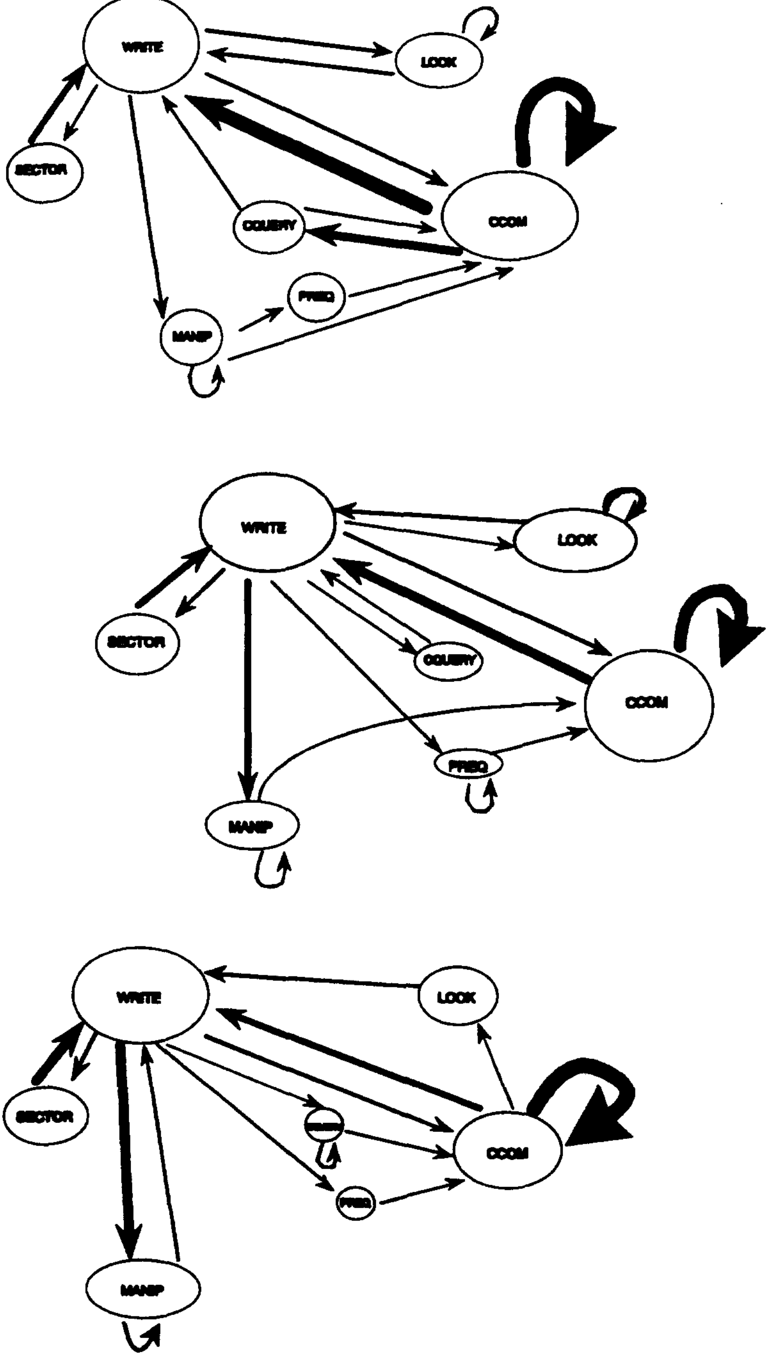

FIGURE 2 MCNs for low (top), medium (middle), and high (bottom) complexity. Size of nodes reflects the proportion of times that that event occurred. Transition is represented by an arrow (arc) from one node to another. Thickness of the arcs reflects the proportion of times that that transition occurred.

TABLE 3
Global MCN Characteristics

	Scenario Complexity		
Global MCN Characteristic	*Low*	*Medium*	*High*
Number of links	16	16	15
Diameter	3.80	2.88	2.90
Prestige			
In-degree	CCOM	CCOM and WRITE	CCOM and WRITE
Number of arcs	5	4	4
In-center	WRITE	WRITE	WRITE
Distance	1.86	1.88	1.91
In-median	WRITE	WRITE	WRITE
Distance	1.33	1.08	1.35
Influence			
Out-degree	WRITE	WRITE	WRITE
Number of arcs	4	6	5
Out-center	WRITE	WRITE	CCOM
Distance	1.93	1.84	1.91
Out-median	WRITE	WRITE	WRITE
Distance	1.36	1.34	1.22

diameter) of each network is similar, with perhaps the low-complexity network being slightly less compact.

Centrality

The nodes also seem to play similar graph-theoretic roles in the three networks. For example, WRITE plays a central role in all three networks. Table 3 also reports three different measures of centrality for both the incoming transitions (prestige) and the outgoing transitions (influence).

Prestige. A node that receives a large number of arcs, or that serves as a central location for incoming arcs, is said to have *prestige.* One measure of prestige is the *in-degree,* a simple count of the number of arcs terminating on a node. Another measure of this type is the *in-center,* the node that minimizes the distance to the farthest node. A final measure of prestige is the *in-median,* which minimizes the distance to all the nodes. Interestingly, the WRITE node appears to be the most prestigious based on the in-center and in-median, regardless of the complexity of the scenario. For in-degree, it shares prestige with CCOM. The prestige of WRITE may be the result of the mandatory strip-marking done for legal purposes.

Influence. *Influence* measures reflect the centrality of the node in terms of the number of outgoing transitions. Thus, the *out-degree* is a simple count of the number of arcs leaving a node, the *out-center* is the node that minimizes travel from that node to the most distant node, and the *out-median* is the node that minimizes travel to all the nodes. Again, WRITE plays a central role in the three scenarios. It is the most influential in all cases except for the out-center for high complexity.

The centrality of WRITE in all three networks strongly supports the importance of studying strip activity. We had expected, perhaps naïvely, that CCOM would have been at the center, at least in terms of prestige. The fact that WRITE appears to be a hub of both incoming and outgoing arcs suggests that much of ATC is organized around this strip activity. We should note that this centrality has two potential consequences for automation. One possibility is that ATC will be greatly facilitated by allowing the computer to control more strip management. The other possibility is that such automation will have negative side effects on controlling traffic (e.g., Hopkin, 1991). Of course, this is the question that originated this project. However, it is now clear that the debate focuses on an issue of true importance. When the center of a transition network is modified, it will have definite effects on the structural whole of the network. Whether these effects are ultimately positive or negative requires additional investigations, but the effects will be there and they may be substantial.

Comparing Across Complexity

Consideration of the intersection of the three graphs revealed a number of transitions that were fundamental to the way in which controllers operated. Figure 3 shows the arcs common to the three networks. Nine arcs are shared by all three; thus 56% of the arcs in the low-complexity and medium-complexity networks are fundamental, and 60% of the arcs in the high-complexity network are fundamental.

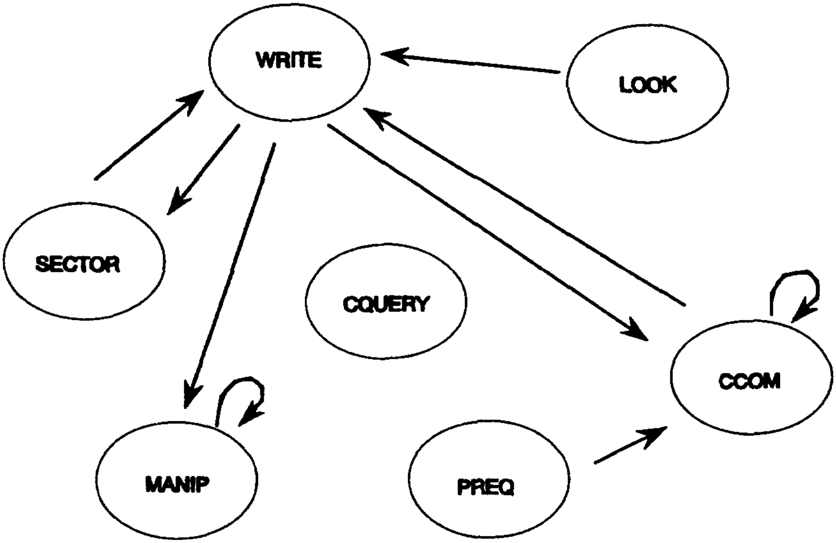

FIGURE 3 Fundamental arcs—subgraphs of the arcs shared by all networks, regardless of complexity.

There were two loops fundamental to the scenarios: a CCOM loop and a MANIP loop, with the former being much more substantial than the latter in all three scenarios. The loops indicate that the controllers often followed one controller command with another and one manipulation with another. Controller commands follow controller commands about as often in the low-complexity scenario as in the high-complexity one. Manipulations, on the other hand, not only occured at a higher rate overall in the more complex scenarios (see Figure 1), but they also tended to follow other manipulations slightly more often in the high-complexity scenarios. In addition to agreement on the two loops, the three networks agree on bidirectional transitions between CCOM and WRITE and between SECTOR and WRITE. Each of the networks also had transitions between WRITE and MANIP, between WRITE and LOOK, and between PREQ and CCOM. Thus, across complexity, several important transitions appear to be fundamental to the control of air traffic.

The MCNs of the high- and low-complexity conditions also differed in a number of interesting ways. These differences can supply insight into how controllers interact with strips differently depending on the complexity of the air traffic situation. We turn now to these differences.

One of the most frequent interactions between a communication event and FPS activities is captured in the arcs between CCOM and WRITE. This likely reflects in part the role of the FPS as a legal record. That complexity affects this interaction is evident: As complexity increased, the frequency of the transition from CCOM to WRITE decreased. It is important to note that controllers do not write less in the complex scenarios. On the contrary, the sizes of the nodes for WRITE are virtually identical across the three scenarios (see Figure 1). What does change as a function of complexity, however, is when writing takes place. As complexity increases, the controllers no longer use a command as a departure point for updating the strips. An implication of this phenomenon is that, in a high-complexity traffic situation, the controller will begin to fall behind in updating strip information. In fact, whereas the CCOM-to-WRITE arc became less prominent in the high-complexity scenario, a CCOM-to-LOOK arc emerged, suggesting that controllers had to settle for a look toward the bay. Being unable to record information on the strips as it is acquired may have other consequences. For example, it appears that controllers engaged in a series of controller queries (the CQUERY loop) in the high-complexity scenario. There may have been too many aircraft for the controller to remember all the altitudes, speeds, and headings, and because the current information had not been written on the strips, the controller might have needed to ask the pilot to provide that information, which added further to the workload.

Although each network showed a WRITE-to-MANIP arc, it is apparent that the frequency of this transition increased with complexity, mirroring the decrease with complexity shown by the CCOM-to-WRITE arc. Not only did the transition from WRITE to MANIP become more prominent with increasing complexity, an arc from MANIP to WRITE emerged in the high-complexity scenario, presumably replacing the activities captured by the arc from MANIP to CCOM found in the low-complexity one. Apparently, in high-complexity situations, the controller engages in periods of uninterrupted interaction with the FPSs, whereas in low-complexity situations, the management of the strips is more integrated into the control of traffic. This tendency to engage in uninterrupted board management as complexity increases also helps to explain why controllers were much less likely to write on a strip immediately after issuing a command in the high-complexity scenario than they were in the low-complexity one.

We find it intriguing that in the low-complexity scenario, the LOOK node is structurally similar to the MANIP node from the high-complexity scenario; that is, in the low scenario, LOOK has a loop and connects to the network only through WRITE. It is tempting to speculate that looking serves the board management function in less dense traffic situations, whereas looking is replaced—or at least augmented—by manipulating (e.g., moving or offsetting) in the more dense traffic situations.

If controllers were more likely to engage in uninterrupted board management, as the Pathfinder analysis suggests, then the frequency of sequences of board management activity should be greater in the high-complexity scenarios than in the low-complexity scenarios. We tabulated the frequency of *all* triple and quadruple sequences terminating in an FPS activity that occurred together within a 10-sec window. None of the low- or medium-complexity triples exclusively involved writing and manipulating, but 19% of the high-complexity triples did. For the quadruples, the five occurring most frequently in the high-complexity scenarios did not involve CCOM, suggesting again that controllers devoted more uninterrupted clusters of time to FPS-board management in the high-complexity scenario, and relatively less uninterrupted time to communication activity. This conclusion is also supported by a time-series analysis of these and related data conducted by Edwards, Fuller, Vortac, and Manning (1992).

DISCUSSION

By observing controllers in situations of varying traffic complexity, it was possible to discover the relations between communication events and strip activity. This was done by applying the Pathfinder scaling algorithm to matrices of the average transitions between events. Several aspects of ATC were shown to generalize to all levels of complexity, including the central role that writing on the strips plays in the current system.

It was also possible to delineate those aspects of ATC that tended to change as a function of the traffic complexity. Changes due to complexity applied to all controllers, indicating that individual differences were minor relative to situational factors. Moderate levels of complexity proved to be a composite of low complexity and high complexity.

The clearest difference between low-complexity and high-complexity situations was in the manipulation of the strips. At higher complexity, manipulations occurred at a higher rate, tended to occur with other board management functions, and were less integrated with communication events. The pattern of findings implies that board management was affected by the complexity of the scenario. In the more complex scenarios, the controller was forced to find time to keep the board configured and updated, and he or she did this in concentrated time segments. During these times, controllers may feel that they have "been taken away" from the plan-view display and the traffic situation. In this sense, the individual controller temporarily divorces from the radar to perform board management duties, essentially serving as his or her own data controller. Fortunately, with a traffic situation as complex as that found in our complex scenario, these data-side board management duties would normally be performed by a different individual, sparing the controller from the schism of switching from radar side to data side. In fact, several of controllers indicated the need for a data-side controller during the high-complexity scenario.

By contrast, while controlling simpler traffic situations, the controller could manage the board as an integral part of controlling air traffic. That is, he or she more effectively time-shared between controlling traffic and managing the board: Strips were updated immediately after issuing a command and board management was integrated with other controller activities.

This finding implies that automation may in fact facilitate control in higher-complexity situations if it allows the computer to take responsibility for the (now apparently segregated) board management duties. This is a very viable outcome of automation. However, an unwanted byproduct of automation may be an increased workload in other areas (e.g., more keyboard use), which may offset any advantage gained from having the computer take over board management duties. Nevertheless, it is likely that any facilitating effects of automation will be more pronounced in more complex situations. In those situations the controllers would be better able to control traffic without the assistance of a data-side controller if the automation allowed for the integration of board management duties with separating aircraft.

ACKNOWLEDGMENTS

We thank the Federal Aviation Administration Academy in Oklahoma City for its help; Cornelia Rea of the University of Oklahoma for her assistance; and Gwen Sawyer, Manny Torres, Claude J. Schuldt, and Tom Lynch. We also thank Bob Blanchard, William Collins, Pam Della Rocco, Ron Lofaro, Richard Mogford, Betty Murphy, Dave Schroeder, Michael Wayda, and Hilda Wing for their insightful comments.

"O. U. Vortac" represents the collaborative research efforts of Francis T. Durso, Scott D. Gronlund, and Stephan Lewandowsky.

REFERENCES

Ammerman, H. L., & Jones, G. W. (1988). *ISSS impact on ATC procedures and training* (FAA Rep. No. DTF-A01-85-Y-0101304). Washington, DC: Federal Aviation Administration.

Buckley, E. P., DeBaryshe, B. D., Hitchner, N., & Kohn, P. (1983). *Methods and measurements in real-time air traffic control system simulation* (Rep No. DOT/FAA/CT-83/26). U.S. Department of Transportation, Technical Center.

Cooke, N. M., Durso, F. T., & Schvaneveldt, R. W. (1986). Recall and measures of memory organization.

Journal of Experimental Psychology: Learning, Memory, and Cognition, 12, 538–549.

Durso, F. T., & Coggins, K. A. (1990). Graphs in the social and psychological sciences: Empirical contributions of Pathfinder. In R. W. Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization* (pp. 31–51), Norwood, NJ: Ablex.

Edwards, M. B., Fuller, D. K., Vortac, O. U., & Manning, C. A. (1992). *The role of flight progress strips in en route air traffic control: A time series analysis.* Manuscript submitted for publication.

Hopkin, V. D. (1988). *Human factors aspects of the AERA 2 program.* Hampshire, UK: Royal Air Force Institute of Aviation Medicine.

Hopkin, V. D. (1991). The impact of automation on air traffic control systems. In J. A. Wise, V. D, Hopkin, & M. L. Smith (Eds.), *Automation and systems issues in air traffic control* (pp. 3–19). Berlin: Springer-Verlag.

Human Technology, Inc. (1990). *Cognitive task analysis of prioritization in air traffic control.* McLean, VA: U.S. Office of Personnel Management, Training Management Assistance Division.

Hutchinson, J. W. (1989). NETSCAL: A network scaling algorithm for nonsymmetric proximity data. *Psychometrika, 54,* 25–51.

McDonald, J., Paap, K., & McDonald, D. (1990). Hypertext perspectives: Using Pathfinder to build hypertext systems. In R. W. Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization* (pp. 197–212), Norwood, NJ: Ablex.

Means, B., Mumaw, R., Roth, C., Schlager, M., McWilliam, E., Gagne, V. R., Rosenthal, D., & Heon, S. (1988). *ATC training analysis study: Design of the next-generation ATC training system.* Washington, DC: Federal Aviation Administration.

Roske-Hofstrand, R., & Paap, K. (1986). Cognitive networks as a guide to menu organization: An application in the automated cockpit. *Ergonomics, 29,* 1301–1312.

Schvaneveldt, R. W. (Ed). (1990). *Pathfinder associative networks: Studies in knowledge organization.* Norwood, NJ: Ablex.

Schvaneveldt, R. W., Dearholt, D. W., & Durso, F. T. (1985). Graph theoretic foundations of Pathfinder networks. *Computers and Mathematics With Applications, 15,* 337–345.

Schvaneveldt, R. W., & Durso, F. T. (1981, November). *General semantic networks.* Paper presented at the meeting of the Psychonomic Society, Philadelphia.

Schvaneveldt, R. W., Durso, F. T., & Dearholt, D. W. (1989). Network structures in proximity data. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 24, pp. 249–284). New York: Academic.

Schvaneveldt, R. W., Durso, F. T., Goldsmith, T. E., Breen, T. J., Cooke, N. M., Tucker, R. G., & DeMaio, J. C. (1985). Measuring the structure of expertise. *International Journal of Man–Machine Studies, 23,* 699–728.

Smith, M. L. (1991). Adaptive training to accommodate automation in the air traffic control system. In J. A. Wise, V. D. Hopkin, & M. L. Smith (Eds.). *Automation and systems issues in air traffic control* (pp. 481–495). Berlin: Springer-Verlag.

Standard Technology, Inc. (1990). *FAA staffing standards for air route traffic control centers.* Washington, DC: Author.

U.S. Department of Transportation, Federal Aviation Adminstration. (1989). *Air traffic control* (Publication No. 7110.65). Washington, DC: Author.

U.S. Department of Transportation, Federal Aviation Administration. (1990). *FAA Academy scenario guide* (Publication No. EIL-10a-2). Oklahoma City: Author.

U.S. Department of Transportation, Federal Aviation Administration. (1988). *Instructional program guide* (Pub. No. EP 12-0-1E). Washington, DC: Author.

Vortac, O. U., & Gettys, C. F. (1990). *Cognitive factors in the use of flight progress strips: Implications for automation.* Norman: University of Oklahoma, Cognitive Processes Laboratory.

Weston, R. C. W. (1983). Human factors in air traffic control. *Journal of Aviation Safety, 1,* 94–104.

Wise, J. A., Hopkin, V. D., & Smith, M. L. (Eds.). (1991). *Automation and systems issues in air traffic control.* Berlin: Springer-Verlag.