Automation of Flight Data in Air Traffic Control

O. U. Vortac
Ami L. Barile
Chris A. Albright
Todd R. Truitt
University of Oklahoma

Carol A. Manning
Federal Aviation Administration Civil Aeromedical Institute

Dana Bain
Federal Aviation Administration Academy

Every day, the airspace of the United States accommodates hundreds of thousands of aircraft movements. Some 70% of all flights occur under instrument flight rules and are thus handled by the nation's air traffic control (ATC) system. Currently, a major ATC center may handle as many as 7,000 flights a day, the majority of them jetliners traveling at airspeeds in excess of 450 knots. The responsibility for the separation of these flights rests with individuals or small teams of controllers, each assigned a volume of airspace, a sector. The controller manages the sector using technologies that, given the traffic loads and the rapid evolution of in-flight avionics, appear anachronistic. For example, some important data pertaining to a flight are currently recorded on small strips of paper, known as flight progress strips (FPSs), whose format and function have changed little since they were introduced in the 1930s and 1940s.

The FPS augments the primary flight information displayed on the computer-enhanced radar: The radar screen shows the position of each aircraft in the sector together with a representation of sector boundaries and other landmarks. Each aircraft is shown together with its call sign, altitude, and a limited choice of other information. Additional flight data, such as planned route, time of arrival, assigned altitude, type of aircraft, and so on, are printed on the FPS that is associated with each flight and that is posted in a "bay" next to the radar display. The FPSs are printed prior to departure of an aircraft and are thus primarily based on projected altitudes and times.

Several minutes prior to entering the sector, forthcoming flights are posted in a "suspense" bay, and once a flight becomes active by entering the sector, its FPS is moved by the controller to the "active" bay. While a flight is active, the controller frequently interacts with the corresponding FPS, noting activities entered into the computer, recording instructions to the pilot, and indicating the controller's plans for that flight. Estimated times are updated and turned into actual times, reroutings are noted, changes in altitude are recorded, and so on. All these changes are recorded on the strip by crossing out the previous data and writing the new information in its place. In addition, controllers often "offset" (move horizontally beyond the edge of the bay) an FPS to remind themselves that an aircraft requires some activity (e.g., clearance to descend) in the near future. In short, the FPS contains a record of the flight's progress through the sector, and, in fact, using the FPS to retain a legal record of the flight (at the time of this writing) is required by Federal Aviation Administration (FAA) procedures (U.S. Department of Transportation, 1989).

It should be clear, even from this brief description, that the FPS appears to form an integral component of the ATC task with potentially important and unique cognitive functionality. For example, the seemingly simple act of moving the strip of an incoming flight from the suspense to the active bay entails scanning other FPSs in order to find the appropriate place for the new flight, perhaps consolidating the controller's memory for the traffic pattern (Hopkin, 1988; see also Slamecka & Graf, 1978). Similarly, the fact that an FPS maintains the history of a flight, because previous information remains visible even when crossed through, has been listed as a unique and important advantage of paper strips (e.g., Hopkin, 1991; Weston, 1983). Finally, the frequent use of FPSs as memory aids (e.g., their offsetting when a flight requires some future control action) may reveal another important cognitive role of this seemingly simple strip of paper.

The gap between advanced in-flight technology and the more traditional ATC equipment is expected to close during the next decade. Because the most dramatic change in en route[1] control is likely to be the replacement of the paper FPSs with electronic versions, there has been some concern that the cognitive functionality provided by the current paper strip may be lost: Incoming flights will (by default) appear automatically in the active portion of the electronic display, the previous history of a flight will not be readily available, and highlighting an electronic entry to serve as a reminder cue will require keystrokes as opposed to a simple manual offset (see Ammerman & Jones, 1988, for more details). It follows that the introduction of

[1]En route control handles flights outside the immediate perimeter of major airports; this chapter specifically addresses issues in en route control.

automation, as feared by Hopkin (e.g., 1988), may entail diverse negative consequences for controller performance because the cognitively beneficial interaction with FPSs is eliminated. We call this the *interaction hypothesis.*

On the other hand, it is possible that removal of the FPSs will merely relieve the controllers of tedious, but required, record keeping that has little to do with the primary task of separating aircraft. That reduction in workload may improve performance, because the controller has more cognitive resources available for resolving traffic conflicts. We refer to this as the *workload hypothesis.*

The available data tentatively support the workload hypothesis. Consider a relevant study by Vortac, Edwards, Fuller, and Manning (1993), in which the availability of FPSs was manipulated between groups of controllers. In one condition, air traffic controllers controlled traffic normally, using paper FPSs as they would in the field. A second group of controllers controlled the same scenarios, but received skeleton FPSs that contained only a subset of the typically available information. In addition, as an extreme analog of automation, controllers in that condition were prohibited from writing, touching, or manipulating the skeleton strips. The study included a set of cognitive measures (attention, visual search, spatial recall, flight data recall, prospective memory, planning) as well as a battery of possible performance indices (over-the-shoulder evaluation by a subject-matter expert, completeness of relief briefings, overall efficiency of traffic management).

Contrary to the interaction hypothesis, Vortac et al. (1993) found few differences between the two conditions. Even when controllers were prevented from interacting with (restricted) flight data, they controlled traffic as efficiently and safely as their counterparts in the control condition. Similarly, and to our surprise, the only cognitive effects favored the group that had less interaction; of relevance here was the result that controllers granted more requests and granted them sooner when they did not interact with the FPSs.

Specifically, the study included several requests by pilots that the controller could not grant without some delay, because the aircraft was still outside the sector. In those cases, the controller must rely on prospective memory, memory for activities to be performed in the future (e.g., Einstein & McDaniel, 1990; Meacham & Singer, 1977). Given that FPSs function as external memory aids, for example by "offsetting" when future actions are required, a prospective memory advantage (Meacham & Leiman, 1982) constitutes the a priori most likely prediction of the interaction hypothesis (cf. Lansdale, Simpson, & Stroud, 1990). The fact that the opposite result was obtained by Vortac et al. (1993) suggests that workload is considerably reduced when FPS updating tasks are eliminated.

However, any conclusions based on that study must remain tentative because the experimental manipulation was not a complete analog to automation. Any automation of a data display likely entails two major consequences: greater operator passivity and enhanced display dynamics. Vortac et al. (1993) captured the former by restricting access to the strips, but they ignored the latter by presenting a single static set of FPSs that remained unchanged throughout the experiment.

The present study, then, was designed to provide a further contrast between the interaction and workload hypotheses using a more sophisticated experimental instantiation of automation. Each controller managed traffic under three different conditions: With the current paper FPSs (normal), with a completely automated electronic display (full), and with a partially automated display (partial). Under full automation, the electronic strips were automatically moved between bays, updated as necessary, and removed once a flight was handed off to the next controller. Under partial automation, updating of the electronic strips was automated, but it remained the controller's responsibility to move strips from the suspense to the active bay, to remove them when no longer needed, and to resequence and highlight them if so desired. Both instantiations of automation used a one-line electronic strip, modeled after the current plans for ISSS.

Turning to the predictions for this experiment, the interaction hypothesis would expect poorest performance under full automation, owing to the absence of controller interaction with flight data. Even under partial automation, performance should be impaired because interaction with the flight data is indirect, via some input device, rather than through direct physical manipulation. Conversely, the workload hypothesis would predict best performance under full automation, and a simple version of this hypothesis would also expect partial automation to be superior to the normal condition.

A second purpose of the current experiment was to investigate the extent to which flight-strip or board management responsibilities can be divided between the controller and the automation. Inclusion of the partial condition allowed assessment of the likely ease with which board management duties can be relegated to automation. We argued that partitioning of subtasks across distributed intelligences may be an important aspect of ATC automation (see Vortac, Edwards, & Manning, 1994, for details).

METHOD

Subjects

A total of 12 full-performance level controllers participated. All were instructors at the FAA Academy, had been controllers for an average of 6.4 years, and last served in the field 26 months prior to the study. All subjects participated in all conditions (normal, full, and partial) during a single 3-hour session.

Materials

Facility and Apparatus. The experiment was conducted at the Radar Training Facility (RTF) at the FAA's Mike Monroney Aeronautical Center in Oklahoma City. The RTF provides high-fidelity air traffic simulation using the fictitious Aero Center airspace. Subjects were familiar with the airspace and with the standard en route equipment. Automation was provided by a laptop computer with an external 17-inch color monitor and a trackball. The monitor was mounted at eye level at a comfortable viewing distance from the subject, and the trackball was used by the subject to move and sequence the electronic "strips" in the partial condition (see Fig. 25.1). Input not provided by the subject (e.g., updating of altitudes, etc.) was entered on the keyboard of the laptop by an experimenter who had received extensive training on that particular scenario. Thus, the experimenter was able to automate the flight data display because he or she could anticipate probable commands, knew likely flight paths, and was familiar with all relevant call signs and flight data.

In the automated conditions, flight data were displayed in the form of one-line entries (see Fig. 25.2). The top two thirds of the display corresponded to the active bay, the bottom one third to the suspense bay. Each strip contained the aircraft call sign and aircraft type, assigned altitude, route of flight, and the flight's computer identification number.

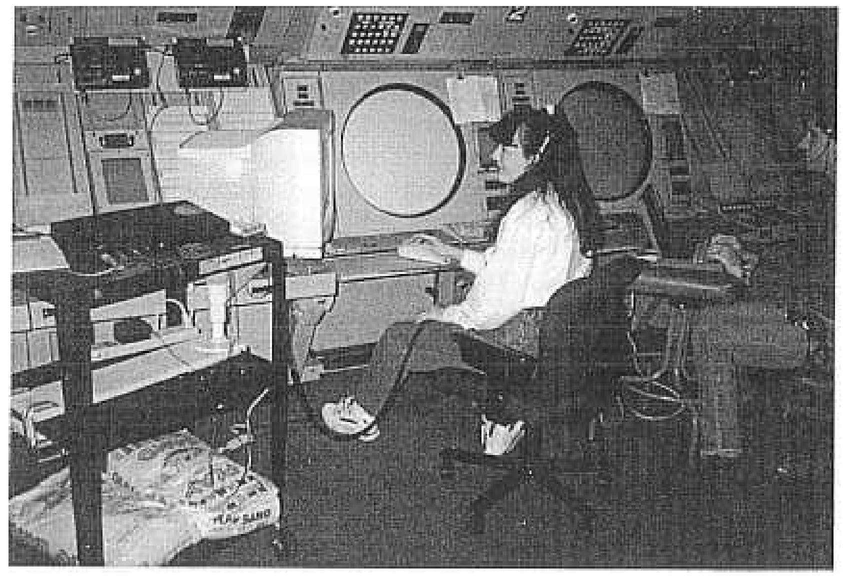

FIG. 25.1. Subject in the partial automation condition seated in front of the radar display, manipulating electronic strips through the trackball. The experimenter is using a laptop to control the automated aspects of the display.

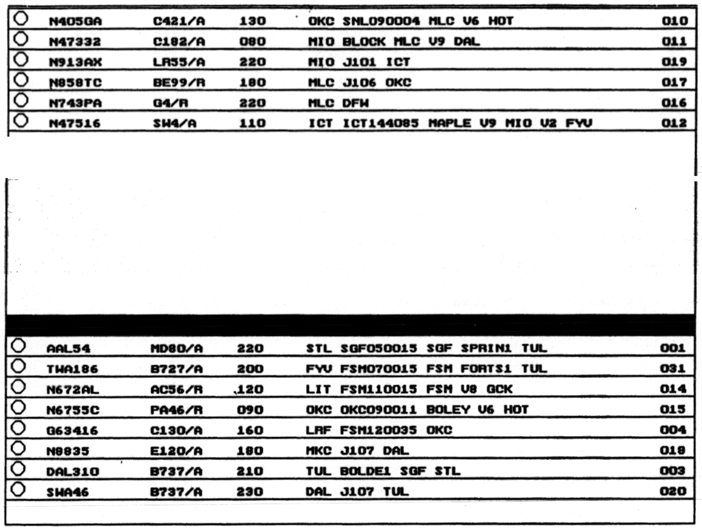

FIG. 25.2. The electronic display of flight data used in the automated conditions. Each strip had a control point, the call sign, the aircraft type and transponder, altitude, route, and computer identification number.

Scenarios. Three scenarios were constructed, each approximately 25 minutes in length and consisting of an average of 21 planes (8.7 departures, 3.7 arrivals, and 8.7 overflights). They were judged by a subject-matter expert to represent a traffic density that an individual controller could handle in the field. None of the subjects were familiar with the scenarios. Assignment of scenario to condition was counterbalanced such that each scenario appeared four times in each condition. All subjects first controlled traffic normally; then half of the subjects were exposed to partial automation, and the remaining half to full automation.

Prospective Memory Events. Each scenario contained three prospective requests spaced on average 7.5 minutes apart. These were pilot requests for an altitude change or rerouting made prior to entering the controller's airspace, when the controller cannot issue commands without permission from the adjacent facility.[2] Hence, we could be certain that the controller had to remember to grant these requests at some time in the future. We measured whether or not the request was granted once the aircraft entered the sector, as well as the latency to grant the request.

[2]Unlike in the field, any attempts to achieve control were denied by the (simulated) adjacent facility. Only one subject commented on the inability to obtain control of out-of-sector flights as unrealistic. Analyses were conducted excluding this subject, but did not alter any of the patterns reported in this chapter.

Conditions

Normal. In the normal condition subjects received their flight data on paper FPSs, just as they do in the field. Subjects were instructed to interact with the FPSs as they normally did. This included a certain amount of legally required strip marking, as well as movement of the strips from the suspense to the active bay, offsetting the strips when desired, and removing a strip after a flight was handed off to the next controller.

Full Automation. The experimenter was responsible for updating altitude and route changes on the electronic strips. By retyping what the controller entered on the en route keyboard, the experimenter provided the necessary information to update the electronic strips on the monitor. Subjects were told that the electronic strips would automatically appear in the suspense bay. The strips were then automatically moved to the active bay when a hand-off was taken; they would turn yellow for 10 seconds before reverting to the background green. The strips were sequenced in ascending order by time of arrival. The automation also removed strips from the active bay 1 minute after the aircraft had been handed off to the next sector. Similarly, updates for route and altitude replaced the existing route and altitude.

Partial Automation. The partial automation was similar to full automation with the following exceptions. Subjects were required to use the trackball to click a strip in the suspense bay to move it into the active bay. A method analogous to offsetting a strip was provided by allowing the controller to click a strip in the active bay, whereupon it was highlighted in red. Subjects were also told how to resequence and remove strips from the active bay by using drag and drop operations with the trackball. Finally, subjects were told that they had to keep their electronic strip "bay" managed in order to ensure that new flight data would become available at appropriate times. However, offsetting and resequencing were optional, as they are in the field.

Procedure

Subjects first completed a background questionnaire and were then familiarized with the particulars of controlling traffic in Aero Center (letters of agreement, hand-off procedures, radio frequencies, etc.). The paper strips for the initial normal condition were prepared and the controllers were given several minutes to adjust the FPSs to their preferences. Subjects were given minimal instructions ("control traffic as you normally would in the field") for the normal condition.

At a predetermined stopping point, the scenario was frozen and the controller was dismissed for a 15-minute break. The subject-matter expert then conducted a postscenario performance analysis (Vortac et al., 1993): For each aircraft remaining in the airspace, the subject-matter expert decided how many speed, route, altitude changes, and so on remained to get the aircraft out of the controller's airspace safely. Because the scenario was stopped at the same point and had the same starting conditions for all subjects, we reasoned that a more efficient controller should have fewer control actions remaining at the end of the scenario than would a less efficient controller. A postscenario analysis was conducted after each scenario.

Subjects next participated in the automated conditions, in an order determined by counterbalancing. For familiarization, subjects first viewed electronic versions of the FPSs from the immediately preceding normal scenario. In the case of partial automation, subjects also practiced the allowable trackball operations using electronic versions of the strips from the preceding scenario. Subjects were given another 15-minute break between the automated scenarios. After completion of the third scenario, subjects were given a questionnaire assessing their views on the scenarios, the automation, and the format of the electronic strips. Subjects also indicated whether or not they controlled traffic any differently than they would in the field.

RESULTS

Postscenario Analysis

The mean number of control actions remaining at the stopping point of the scenarios is shown in Table 25.1. The results provide no support for the contention that automation impairs efficient management of traffic. If anything, full automation appears superior to normal, although this difference did not reach significance [$t(11) = 1.69, p < .12$]. Moreover, although decomposing the total number of remaining actions into different classes of actions further increased variability, the table indicates a general superiority of the full automation condition. These aspects of the data cast doubt on the interaction hypothesis stated at the outset. Finally, note that partial automation yields performance between the normal and full automation conditions.

TABLE 25.1
Mean Number of Control Actions

	Normal	*Partial*	*Full*
Total remaining actions	27.0	25.6	23.3
Communications/hand-offs	24.1	22.8	20.8
Altitude changes	1.6	1.9	1.3
Speed changes	.7	.6	.8
Point-outs	.4	.1	.3
Route amendments	.2	.2	.2

Controller Opinion

Mean Likert judgments were used to help quantify controllers' opinions. Table 25.2 compares the three conditions, with lower scores indicating more favorable opinions. There was no apparent difference between partial and full on these scales. However, there seemed to be some interesting differences between the paper FPSs (normal) and the electronic conditions. Electronic strips were rated as slightly less useful than FPSs; however, electronic strips were rated as slightly easier to use and were liked slightly more than were the paper FPSs. For ease of use, six controllers preferred electronic strips, two preferred FPSs, and three did not indicate a preference. For satisfaction, five preferred the electronic strips, two preferred paper, and five did not indicate a preference. These leanings toward electronic strips were reversed, however, when usefulness was judged: Seven controllers thought the traditional paper strip more useful, two preferred the electronic strips, and three thought the two formats equally useful.

Protocols helped illuminate the Likert values. The most mentioned likable feature of the electronic strip was the automatic update ($n = 5$) of altitude and route information. The ability to highlight a strip ($n = 4$) in partial automation was also well liked. Subjects thought the electronic strip required less physical manipulation due to the ease with which they could be moved, the ease of their removal, and the automatic update. Overall, controllers felt as if they spent less time on board management.

The electronic strips were apparently judged less useful than the FPSs because they lacked information. Subjects recommended the addition of a beacon code ($n = 10$) as the single most agreed-on addition. Others suggested a space for writing on the strip or a small notepad ($n = 8$), that both interim and final requested altitudes should be visible along with filed true airspeed ($n = 4$), and that more than one active bay would be helpful ($n = 3$). It is interesting that despite the feeling that more information was needed on the electronic strips, the controllers were at least as efficient (postscenario

TABLE 25.2
Mean Likert Judgments

	Normal	*Partial*	*Full*
Usefulness	2.1	2.4	2.5
Ease of use	2.5	1.9	1.8
Satisfaction	2.6	2.0	2.2

TABLE 25.3
Prospective Memory Measures

	Normal	*Partial*	*Full*
Time from boundary to request granted (seconds)	77.4	79.4	70.4
Proportion of requests correctly granted	.83	.78	.90

analysis) in the full automation condition as they were when controlling traffic normally.

Prospective Memory

We looked at both the proportion of times the controller remembered to grant a request, and how long the controller waited to grant the request once he or she had control of the aircraft. As shown in Table 25.3, under full automation, controllers granted slightly more requests and granted them sooner than in the normal condition. This replicates our earlier finding (Vortac et al., 1993), although here the trends were not significant. Despite this null effect, it is important to note that again the pattern is opposite to that predicted by the interaction hypothesis. It is also worth noting that any advantage over normal conditions gained with full automation is completely lost under partial automation.

Other features of the data were interesting, although they do not distinguish among types of automation. Only one request was granted incorrectly: A controller cleared an aircraft direct to Wichita rather than to Kansas City. There were eight additional occurrences in which the controller remembered that he or she had received a request but did not remember the content of the request. The frequency with which this partial failure occurred supports Einstein and McDaniel's (1990) componential analysis of prospective memory: The component that triggers an action was intact, but the content of the required action was lost. In these cases, the controller contacted the aircraft and asked the pilot to repeat the request.

DISCUSSION

Notwithstanding meager statistical support, integrating across performance, opinion, and cognitive variables suggests that full automation has an advantage over the way traffic is controlled currently: Subjects were slightly more efficient, having fewer actions left to perform at an arbitrary stopping point; they liked the ease of use of the automation and were more satisfied with full automation than the standard system; finally, delayed pilot requests were granted 7% more often and a few seconds earlier under full automation. By contrast, there is no evidence supporting the competing interaction hypothesis, namely that automation will inhibit performance or cognition.

On the other hand, it is more difficult to marshal a similar argument for partial automation. First, partial automation was not always superior to the normal condition (e.g., prospective measures), and even when it was, the size of the difference did not invite lack-of-power excuses. Because partial and full automation used the same one-line electronic strips, it is unlikely that this particular display format explains the patterns of results. However, partial automation and normal operations both involved the controller engaging in board management. Apparently, when the controller must take time and resources to keep the flight data bay organized, we cannot expect automation of other functions, such as updating altitude and route, to yield an advantage. Thus, the current data suggest that it may be difficult to divide responsibility for board management between the computer and controller. In our simulation of partitioned board management between computer and controller (partial condition), the advantage of automation was completely (prospective memory) or partially (performance) lost.

We have argued elsewhere that board management is a behavioral module to experienced controllers, and that if automation replaces part of, rather than the entire, module, then no advantage (possibly even an inhibition) can be caused by the automation. This modular automation hypothesis (Vortac, 1993; Vortac et al., 1994) is based on the argument (Hayes-Roth, 1977) that cognitive processes that frequently occur together will eventually become unitized. When a unitized assembly of processes exists for the task at hand, cognitive processing should be most efficient. If, instead, processes must be combined to deal with the task, then additional time and resources would be needed. Or, if a unitized module must be disassembled or fractionated in Hayes-Roth's terms, then again additional time and resources would be needed.

Modularity may prove to be an important part of the success of automation attempts and should be considered in the design of automated systems. Fractionation could occur if uninformed attempts to automate a complex dynamic system places part of an existing cognitive or behavioral module under automation. This argument, we believe, is especially applicable to a situation like that of the current ATC system where expert, highly skilled individuals will be placed in a similar environment with automated aids designed to replace existing skilled behaviors. In such a situation, existing modules will be brought to bear; the new situation will force their fractionation; performance will be comparable to, or in extreme cases worse than, no automation. Whether a particular module should be automated or left under manual control depends on other factors (Vortac et al., 1994), but the module in its entirety should either be automated or preserved as a manual subtask.

In a related laboratory experiment using undergraduate volunteers (Vortac & Manning, 1994), we found support for the disruptive impact of automation that fractionates existing modules. Subjects learned to enter sequences of commands to control various aspects of a fictional complex process. Each task required that the subject enter eight keystrokes. By requiring subjects to make the same four-stroke entries frequently and in various situations we forced the subjects to create modules. Thus, the training ensured that the eight keystroke tasks comprised two four-keystroke modules, and it also ensured that the experimenters knew which sequences of keystrokes were unitized and which were not. We then automated some of the functions. In the modular automation condition, entire modules were automated: The computer performed one four-keystroke module and left the subject to perform the other module. In the fractionated condition, the same number of keystrokes were automated, but they fractionated the preexisting modules: The computer would perform two keystrokes from each of the two modules constituting the task. Fractionation caused a disruption in performance that required many subsequent retraining sessions to reduce.

Thus, we have some reason to believe that the fractionation tenet of modular automation holds. In addition, earlier work with air traffic controllers working singly and in teams supplied data that bear on the modules that exist in experienced air traffic controllers. Using the Pathfinder scaling algorithm (Vortac, Edwards, Jones, Manning, & Rotter, 1993; Vortac et al., 1994) and time-series analysis (Edwards, Fuller, Vortac, & Manning, in press) we investigated the structural characteristics of transition frequencies between various behaviors emitted by controllers in different traffic situations. Of relevance here is the finding that manipulating strips and writing on strips appeared to represent a module to controllers. We referred to this complex of writing and manipulating FPSs as the *board management module.* Board management meets several criteria suggestive of a module. Writing and manipulating strips co-occur frequently and in a variety of situations and were strongly linked in the Pathfinder network. When the task of ATC was divided between two controllers, all of board management becomes the responsibility of one controller, whereas the events that trigger board management are the actions of the other controller. Thus, teams of controllers do not fractionate the board management module.

Given that board management appears to be a module in ATC, it is likely that the partial automation used in the current experiment fractionated that module. Thus, we should not be surprised that there was no real benefit to partial automation, and in some cases a slight inhibition. We mention this because it may be the case that other partial automation configurations can be developed that would not fractionate board management, although it is difficult to imagine a more natural partition of responsibilities than the one implemented here.

Overall, we suspect that full automation will surpass the normal ATC configuration because it should relieve the controller of all board management responsibilities, allowing him or her to focus on the control of active traffic. We believe that to the extent that automation allows controllers to be relieved of all board management responsibilities it will prove to be a successful system. We warn, however, of the problems associated with automating only part of the board management module.

ACKNOWLEDGMENTS

O. U. Vortac was the collaborative research effort of Francis T. Durso, Scott D. Gronlund, and Stephan Lewandowsky. Lewandowsky is now at the University of Western Australia. This work was supported by contract #DTFA-02-91-C-91089 from the FAA. We are grateful to the instructors and management of the FAA Academy at Mike Monroney Aeronautical Center.

REFERENCES

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

Edwards, M. B., Fuller, D. K., Vortac, O. U., & Manning, C. A. (in press). The role of flight progress strips in en route air traffic control: A time series analysis. *International Journal of Human-Computer Studies.*

Einstein, G. O., & McDaniel, M. A. (1990). Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 717–726.

Hayes-Roth, B. (1977). Evolution of cognitive structures and processes. *Psychological Review, 84,* 260–278.

Hopkin, V. D. (1988). *Human factors aspects of the AERA 2 program.* Farnborough, 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.

Lansdale, M. W., Simpson, M., & Stroud, R. M. (1990). A comparison of words and icons as external memory aids in an information retrieval task. *Behavior and Information Technology, 9,* 111–131.

Meacham, J. A., & Leiman, B. (1982). Remembering to perform future actions. In U. Neisser (Ed.), *Memory observed: Remembering in natural context* (pp. 327–336). San Francisco, CA: Freeman.

Meacham, J. A., & Singer, J. (1977). Incentive effects in prospective remembering. *The Journal of Psychology, 97,* 191–197.

Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology: Human Learning and Memory, 4,* 592–604.

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

Vortac, O. U. (1993). Should Hal open the pod bay doors? An argument for modular automation. In D. J. Garland & J. A. Wise (Eds.), *Human factors and advanced aviation technologies* (pp. 159–163). Daytona Beach, FL: Embry-Riddle Aeronautical University Press.

Vortac, O. U., Edwards, M. B., Fuller, D. K., & Manning, C. A. (1993). Automation and cognition in air traffic control: An empirical investigation. *Applied Cognitive Psychology, 7*, 631–651.

Vortac, O. U., Edwards, M. B., Jones, J. P., Manning, C. A., & Rotter, A. J. (1993). En route air traffic controllers' use of flight progress strips: A graph-theoretic analysis. *The International Journal of Aviation Psychology, 3*, 327–343.

Vortac, O. U., Edwards, M. B., & Manning, C. A. (1994). Sequences of actions for individual and teams of air traffic controllers. *Human Computer Interaction, 9*, 319–343.

Vortac, O. U., & Manning, C. A. (1994, April). *Modular automation: Automating sub-tasks without disrupting task flow.* Paper presented at the First Automation Technology and Human Performance Conference, Washington, DC.

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