

DOT/FAA/AM-02/22

Office of Aerospace Medicine
Washington, DC 20591

Aiding Planning in Air Traffic Control: An Experimental Investigation of the Effects of Perceptual Information Integration

Peter M. Moertl
University of Oklahoma
John M. Canning
White Oak Technologies
Scott D. Gronlund
University of Oklahoma
Michael R.P. Dougherty
University of Maryland
Joakim Johansson
Ericsson, Inc.
Scott H. Mills
SBC Technology Resources, Inc.

December 2002

Final Report

This document is available to the public
through the National Technical Information
Service, Springfield, VA 22161.



U.S. Department
of Transportation

**Federal Aviation
Administration**

N O T I C E

This document is disseminated under the sponsorship of the U.S. Department of Transportation in the interest of information exchange. The United States Government assumes no liability for the contents thereof.

Technical Report Documentation Page

1. Report No. DOT/FAA/AM-02/22		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle Aiding Planning in Air Traffic Control: An Experimental Investigation of the Effects of Perceptual Information Integration				5. Report Date December 2002	
				6. Performing Organization Code	
7. Author(s) Moertl PM ¹ , Canning JM ² , Gronlund SD ¹ , Dougherty MRP ³ , Johansson J ⁴ , and Mills SM ⁵				8. Performing Organization Report No.	
9. Performing Organization Name and Address					
¹ University of Oklahoma Norman, OK 73109		² White Oak Technologies, Inc. Silver Springs, MD 20910		³ University of Maryland College Park, MD 20742	
⁴ Ericsson, Inc. Stockholm, Sweden		⁵ SBC Technology Resources, Inc. Austin, TX 78759		10. Work Unit No. (TRAIS)	
12. Sponsoring Agency name and Address Office of Aerospace Medicine Federal Aviation Administration 800 Independence Ave., S.W. Washington, DC 20591				11. Contract or Grant No. Grant #97-G-037	
				13. Type of Report and Period Covered	
14. Sponsoring Agency Code					
15. Supplemental Notes Work was accomplished under approved subtask AM-A-00-HRR-522. Carol Manning (AAM-500) served as the COTR.					
16. Abstract Prior research examined how controllers plan in their traditional environment and identified various information uncertainty by perceptually representing important constraints. This included integrating spatial information on the radar screen with discrete information (planned sequences of air traffic). Canning et al. (1999) and Moertl et al. (2000) reported improved planning performance and decreased workload in the planning aid condition. The purpose of this paper was to determine the source of these performance improvements. Analysis of computer interactions using loglinear modeling showed that the planning interface led to less repetitive, but more integrated, information retrieval gave rise to the performance improvements. Potential applications of this research include the design and evaluation of interface automation that keeps users in active control by modification of perceptual task characteristics.					
17. Key Words Information Integration, Air Traffic Control, Planning Aids, Strip Marking			18. Distribution Statement Document is available to the public through the National Technical Information Service, Springfield, VA 22161		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 12	22. Price

ACKNOWLEDGMENTS

The support of the Federal Aviation Administration under agreement 97-G-037 is gratefully acknowledged. We thank Frank Durso (Texas Tech University) and Robert Terry (University of Oklahoma) for their methodological expertise, and Carol Manning and Scott Shappell (FAA-Civil Aerospace Medical Institute) for their comments on this manuscript. We also extend our thanks to Ron Andrei and Paul DeBenedittis (of the FAA Academy) and Duane Duke and Rob Enos (FAA Office of Information Services) for their expertise.

AIDING PLANNING IN AIR TRAFFIC CONTROL: AN EXPERIMENTAL INVESTIGATION OF THE EFFECTS OF PERCEPTUAL INFORMATION INTEGRATION

Air traffic control equipment has changed in recent years as the Federal Aviation Administration (FAA) adapts its procedures to the growing volume of air traffic across the country. However, two major components of control equipment have stayed constant over the years. Specifically, generations of air traffic controllers have utilized a radar screen and flight progress strips as separate representations of aircraft entering their controlled sector and cognitively integrated those representations. This equipment has proven to be highly beneficial and, therefore, forms the foundation from which any innovation to the air traffic control system should begin.

If controllers are to manage the increasing volumes of air traffic, planning will be of increasing importance. This is evidenced by recent efforts to provide controllers with plan-aiding technology (e.g., URET, User Request Evaluation Tool, Arthur & McLaughlin, 1998; CTAS, Center-Tracon Automation System, Denery, & Erzberger, 1995; ERATO, En Route Air Traffic Organizer, Bressolle, Benhacene, Boudes, & Parise, 2000). Such interfaces offer additional functions to the controller like conflict detection algorithms (URET), automated traffic advisory functions for descending, sequencing, and spacing aircraft (CTAS), and decision aid tools like filtering options and problem reminders (ERATO). The approach we took in this study can enhance plan-aiding technology by identifying essential informational elements that support air traffic planning and determining the extent to which air traffic planning could be improved by optimizing the representation of that information.

Air traffic controllers manage a complex flow of aircraft through their airspace. They maintain strict rules of separation between the aircraft while allowing all aircraft to reach their destinations as safely and expeditiously as possible. In planning the routes for the aircraft, two forms of planning can be distinguished. Controllers make tactical plans when they make decisions that relate to the current moment and involve the separation of (usually) pairs of aircraft that could soon violate the separation rules and hence, need immediate action. They make strategic plans when their plans span longer periods of time (about 10 minutes or longer) and typically involve multiple aircraft. An examination of strategic planning in air traffic control is timely, given future concepts being

proposed. For instance, there have been discussions regarding the creation of a strategic controller position (N. Lawson & K. Thompson, personal communication, Dec. 15, 1997; see also Vivona, Ballin, Green, Bach, & McNally, 1996). The proposal provides for one person who would be responsible for a multiple-sector airspace, making decisions about traffic in that airspace, and delegating responsibility for tactical decisions to sector-level controllers. A goal of our project was to develop interface tools for a strategic controller position.

Dougherty, Gronlund, Durso, Canning, and Mills (1999) studied how air traffic controllers make strategic plans for en route traffic (high altitude, high speed traffic between destinations) using the radar screen and paper flight progress strips. They identified aircraft sequencing for approach to a common destination as a strategic planning task by analyzing controller verbalizations and use of flight progress strips. The specific sequence of a group of aircraft is determined by many factors— aircraft speed, altitude, destination, and airspace restrictions. Therefore, sequencing aircraft was a complex cognitive task that involved the consideration of many aircraft over an extended period of time. Dougherty et al. (1999) argued that controllers could benefit from an interface that supported planning the sequences of aircraft.

We begin by outlining the relevant aspects of the traditional air traffic control environment that guided our interface design. Following that, we describe the electronic planning aid and outline its design principles. Finally, we report the evaluation of the planning aid by comparing participants' planning performance using the interface to their performance in the traditional air traffic control environment.

Traditional En Route Air Traffic Control Environment

Air traffic controllers primarily use information from two different sources, the radar screen and the flight progress strips. The radar screen shows the spatial position and progress of aircraft, together with some characteristics of the controlled sector (e.g., boundaries and airways). The radar screen displays the spatial location of aircraft as well as the most vital flight information (identifiers, altitude, speed, and sometimes, flight destination). Discrete information

about an aircraft (e.g., its destination, flight origin, planned altitude) can be read from the strips (see Federal Aviation Administration, 1995; this information also is available from the Computer Readout Display). Flight progress strips are small paper strips that are computer-printed and posted next to the controller about 20 minutes prior to arrival of that flight in the sector. Strips also are used to indicate the progress of traffic, as controllers manually update the strips when information about an aircraft changes. In addition, strips carry identifiers that link them to their representations on the radar screen. However, the linkage is not direct because controllers have to search for the aircraft on the radar screen to coordinate the representations. The radar is of primary importance because it tells the controllers where the aircraft currently are. The strips can be useful for determining where the aircraft are going to be.

Dougherty et al. (1999) described how controllers use strips for strategic planning purposes. Controllers often classified aircraft into logical groups. These groupings were based on flight plan characteristics (e.g., destination, route, or altitude), and/or flow restrictions. Then strips were ordered to show the planned sequence for these aircraft. In other words, controllers used strips as planning aids with which they externally represented their planned sequences. For example, if a controller wanted to plan a sequence of aircraft A, B, and C, the controller would arrange the three strips in that order. In this way, the strips physically hold the planned sequence of aircraft. In contrast, the radar screen represents the actual air traffic situation. It is important to note that the planned and current sequences of aircraft are not linked to one another, and changes on the radar screen are not reflected by changes to the strips.

The discrepancy between the planned sequence represented by the strips and the positions of the aircraft on the radar screen is essential information for a controller. It makes apparent that something must be done to update the intended path of the aircraft. In order to assess the discrepancy between the planned and current sequence, controllers have to locate the strip and the corresponding aircraft representation on the radar screen and mentally link those representations to one another. In this way, they determine the mismatch between the planned sequence and the current air traffic situation. In other words, the controller is responsible for linking spatial information (location of the aircraft in the sector, distance to other aircraft, distance to specific points in the sector) to discrete information (flight identifier, optimal sequence of aircraft, flight plan, and aircraft type).

In order to help link the spatial and discrete aircraft representations, controllers have to “declutter” the radar screen. A block of information (the datablock) is connected to each aircraft on the radar containing important flight information about that aircraft. As more and more aircraft come into a sector, datablocks overlap and become illegible. Therefore, controllers often adjust datablocks into one of eight different positions around the aircraft target and also adjust the distance between the datablock and its aircraft. However, some information on the datablock is redundant; controllers can get the same information from strips by matching a unique identifier on the strip to the radar. Therefore, decluttering the radar screen to make aircraft distinguishable could be viewed as a hindrance to more important tasks like strategic planning. Though controllers can utilize radar decluttering to create memory cues, a perceptually-improved interface might supplant the memory demands. The necessity for decluttering seems to be caused by the separation of the different information sources.

We addressed this problem by designing a new interface that took over the task of linking spatial and discrete information. The interface should allow controllers to perceptually assess the difference between the current air traffic sequence and the planned sequence without having to frequently declutter the radar screen. We discuss this electronic planning aid next.

Electronic Planning Aid

The electronic planning aid changed the task of planning air traffic sequences by perceptually representing a major constraint in the controllers’ environment that was not visible in the traditional environment (see <http://www.ou.edu/cas/psychology/HTIC/FOPA/> for screenshots of the interface). It linked discrete information to its spatial representation on the radar screen. This dynamic linkage between the planning aid and the radar screen was implemented by presenting the participants with two computer screens: Figure 1 shows a snapshot of the two screens.

The planning aid was linked to the radar screen and the controller could move the cursor between the two screens. Aircraft in a sector were represented on the planning aid by rectangular tokens as shown in Figure 1. The tokens contained the aircraft’s identification number, destination airport, equipment type, and altitude. Aircraft tokens were marked and unmarked by clicking on them, in which case a check mark symbol appeared or disappeared on the token. All aircraft were represented on both screens and could be perceptually located on one screen by moving the cursor over the token on the planning aid or by clicking on its target on the radar screen.

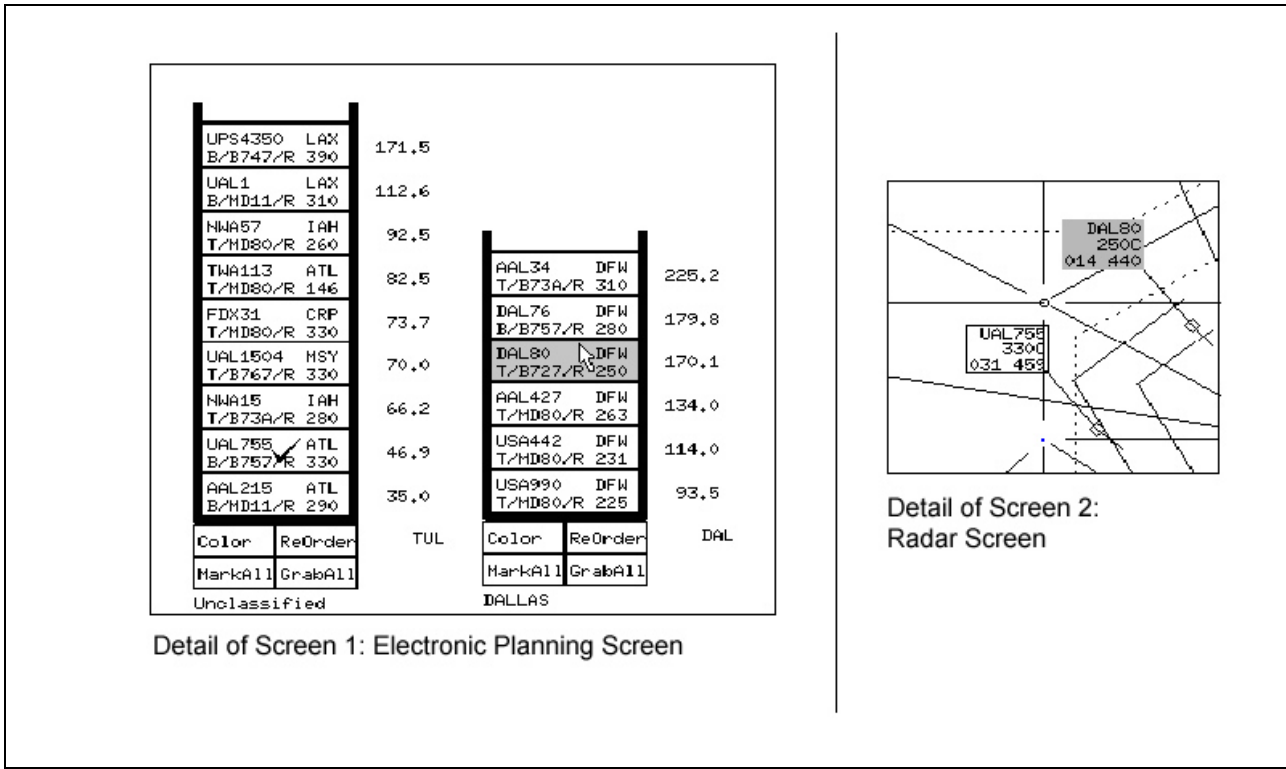


Figure 1. The dynamic linkage between the planner screen and the radar screen. On the radar, every diamond-shaped aircraft representation is linked to a datablock of flight information (in order from top, left to right: aircraft identification, altitude in 100 feet, computer identification number, and speed). The two aircraft on the radar screen (UAL755 and DAL80) are grouped into different categories of traffic (i.e., final destination Dallas and Unclassified). Different categories of traffic are represented in different colors. Aircraft are selected by clicking on them (UAL755). Moving the cursor over any token on the planning screen or a target on the radar screen puts a rectangle around the two representations (DAL80).

The planning aid allowed participants to categorize aircraft and updated that representation on the radar screen in the color corresponding to their category. Figure 1 shows two aircraft (UAL 755 and DAL80) that a controller placed into different categories. Because each column of aircraft was in different colors, participants do not have to encode and remember the categories of each aircraft after they were categorized because they were represented perceptually on the radar screen.

The planning aid also allowed the controller to automatically sort categorized air traffic according to time or distance to specific points (fixes) along the route. This information was displayed adjacent to each aircraft token. Note that this automatic sorting did not include higher-level conflict information; it simply was based on distance/time measures and therefore provided only an initial approximation of aircraft order. These initial sequences needed manual updating and checking. Participants also could get distance

information between points on the radar screen by using a distance-measuring tool. This tool was similar to how controllers measure distance on traditional radar screens.

One essential consequence of this design was that a planned sequence could be perceptually compared with the current sequence by sliding the cursor across the sequence on the planning screen. This allowed the controller to observe how the sequenced position of each aircraft corresponded to its current position on the radar screen. Any discrepancy between the planned sequence and the current sequence was therefore made perceptually salient to the controller. If the planned sequence differed from the current sequence, this discrepancy signaled the need for modifications to either the planned sequence or to an aircraft's path. The discrepancy represented an important constraint as it guided the controller toward the aspects of the situation where control interventions were needed.

Do participants rely to the same extent on the radar screen when using the planning aid as when using the strips? The answer to this question should provide evidence of the extent to which the participants in our study made differential usage of the same radar functions in each condition. For example, they should not have to declutter the radar as much. From observed differences, we can infer how participants used and accessed information differently in the two conditions.

METHOD

Twelve en route air traffic controllers participated in the experiment. All participants were full-performance-level controllers who served as instructors at the FAA Academy. Full-performance-level controllers are qualified to control traffic in a sector by themselves. The experiment was conducted at the FAA Civil Aerospace Medical Institute in Oklahoma City. All participants were familiar with the airspace (i.e., sector) used in the experiment.

Each of the 12 participants was given the task of planning and maintaining aircraft sequences for a generic en route sector and to communicate the sequence orally to a tactical controller (tactician) who implemented it. The tactician, another en route air traffic controller, was the subject-matter expert. All participants performed the same task under two conditions: 1) using the electronic planning aid without strips (planner condition) and 2) using paper flight strips without the electronic planning aid (strips condition). Participants used the same radar screen in both conditions, although some of the radar functionality was absent without the planning aid (i.e., the color coding of aircraft). The two scenarios were busy; up to 44 aircraft were in the sector during 20 minutes and included arrivals, departures, and overflights. The order of the scenarios and the order of the two conditions (planner or strips) were counterbalanced across participants.

Each participant completed a 45-minute training session on the use of the planning aid and radar screen. One or two days later, each participant completed a 1-hour practice session. No training was given on the use of the flight progress strips, as all participants had extensive experience with them in the field. The experiment started a day or two following the practice session. Immediately prior to testing, participants practiced with the planning aid for a few minutes to re-familiarize themselves with its operation.

For each scenario, the participant was told about any special airspace restrictions in effect and was given sufficient time to set up either the planning aid or paper strips while the aircraft in the scenario remained paused. During this time, they verbally conveyed their plan to the tactician. The simulation was started after the plan was conveyed. The simulation ran for 20 minutes, with a brief pause after 10 minutes.

RESULTS AND DISCUSSION

We have previously described the details of the participants' superior performance and reduced workload when using the planning aid (see Canning, Johansson, Gronlund, Dougherty, & Mills, 1999; Moertl et al., 2000). The advantage for the planning aid condition was especially noteworthy because participants had, on average, about 10 years experience with paper strips and less than 2 hours practice with the planning aid. In the present study, we were interested in the differential usage of the radar screen between the conditions and what that can tell us about *why* the planning aid led to superior planning performance and reduced workload.

Loglinear Modeling of User Interactions

Olson, Herbsler and Rueter (1994) presented loglinear modeling as a technique for analyzing human-computer interactions. Loglinear models allow the test of main effects and interactions of models that fit empirical human-computer interaction frequencies and thereby allow one to determine sequential structure and structure changes between conditions. Olson et al. (1994) used loglinear modeling on a set of interaction data that were semantically categorized by human judges. Such a procedure can prove highly useful but has the disadvantage of costing considerable time resources for the manual categorizations of actions. In contrast, we used this approach to analyze raw data-files that were not categorized by human judges. Therefore, we sidestepped the objection that the analysis of categorized data might reflect the judgments of human judges, in addition to actual empirical patterns in the computer interactions. In contrast to Olson et al. (1994), our design included repeated measurements for which we accounted in our models.

Although participants interacted with the radar screen more frequently in the strip condition (significant main effect of experimental condition, Wald

$\chi^2(1) = 17.07, p < 0.01$), of primary interest was determining the source of these differences. We compared the observed frequencies with the expected frequencies and determined how participants interacted with the radar screen differently in the two conditions.

Participants performed nine different types of actions on the radar screen (listed in Table 1). We compared observed frequencies for each action with expected frequencies assuming no differences between conditions (i.e., standardized residuals for each action and condition). The standardized residuals were calculated as the difference between the predicted and observed frequency divided by the square root of the predicted frequency. Table 1 displays the results of this analysis; the last column shows the standardized residuals. The model predicted the observed frequencies satisfactorily (within a 95% confidence interval) for all but three of the actions. This meant that these three actions occurred with differing frequencies in the two conditions. We discuss these three user actions in turn.

Select token. Participants selected significantly more aircraft on the radar screen in the strip condition than in the planning aid condition. Selecting an aircraft creates a border around its datablock that enhances its visibility. Participants in the planning aid condition did not need this perceptual aid as frequently, presumably because they could rely on the dynamic linkage between the two screens to perceptually locate aircraft on the radar screen.

Distance measurement. Participants measured the spatial distance between points on the radar screen more frequently in the strip condition. Distance information was crucial for planning, as it allowed estimation of when aircraft would reach specific points in the sector. Participants in the planning aid condition did not measure the distances on the radar screen as frequently, presumably because they could rely on the time/distance information that was presented to them next to each aircraft on the planning aid.

Datablock adjustment. Datablock adjustments included changing and adjusting datablock position. Participants adjusted datablocks more frequently in the strip condition than in the planning aid condition. As mentioned above, controllers declutter their radar screen to make datablock information visible. These adjustments are an important index of the usage of the radar screen. In the strip condition, participants had to get their flight information from the radar screen and had to declutter the radar screen to get to this information. However, when using the planning aid, participants adjusted datablock position less frequently and instead relied on the planning aid to review flight information. This was consistent with the greater ease of information access in the planning aid condition and more time spent in the strips condition on “housekeeping” functions.

Participants interacted with the radar screen less when they worked with the planning aid. They manipulated aircraft less (visually highlighted or marked aircraft less) and adjusted datablocks less often. They

Table 1.

Observed and Expected Frequencies for a Loglinear Model Assuming no Difference Between Experimental Conditions

User action	Frequency in strip condition	Frequency in planner condition	Predicted Frequency	Standardized residual for strip condition
Adjust Vector Length	3	6	4.5	-0.71
Invalid command (error)	72	83	77.5	-0.62
Zoom in/out	10	11	10.5	-0.15
Move information table	2	2	2	0
Altitude filter	12	10	11	0.30
Invalid track (error)	3	2	2.5	0.32
Select Token	482	408	445	1.75*
Distance measurement	108	56	82	2.87*
Datablock adjustment	661	520	590.5	2.90*

Note. * Observed frequency is outside a 95% confidence interval around the expected frequency.

also assessed spatial distance less often. The greater ease of accessing information was the result of replacing radar functionality with the functionality of the planning aid. The prior results examined interaction frequencies. In what follows, sequences of user interactions were examined to provide additional insights into how participants used the radar screen differently between the two conditions.

User Interactions Sequence Analysis

Differences in sequences of interactions between conditions were also testable using loglinear modeling. For each condition, the empirical probabilities of an interaction occurring after another interaction were compared with the estimated probabilities, assuming that there were no differences between conditions. In a loglinear model, the difference between conditions was tested by a three-way interaction of the main factors (antecedent action, consequent action, and experimental condition, see Olson et al., 1994). We found a significant three-way interaction (Wald $\chi^2(13) = 610.97, p < 0.01$), which indicates different, non-random sequences between the conditions. Accordingly, we found significant differences by comparing specific action sequences between conditions. Again, we calculated the standardized residuals using antecedent and consequent actions, and their interaction. Because there were 115 empirically occurring sequences, only the four sequences that were not predicted by the model are reported in Table 2. These differences reinforced the usage differences between the conditions noted above.

Participants had more frequent repetitions of distance measurement and datablock adjustment in the strip condition. They also cycled between adjusting datablocks and selecting aircraft, and vice versa, more frequently in the strip condition. Information retrieval

in the strip condition appeared more repetitive, compared with the planning aid condition. Participants needed to declutter the radar screen more because they had to integrate the strips with the aircraft on the radar screen. Furthermore, the integration effort seemed not to occur as an isolated event but happened in bursts of repeated decluttering activity. It is hard to see how these bursts of decluttering activity could entail any memory-enhancing function. This manual matching was much less prevalent in the planning condition, where participants used fewer interactions in absolute numbers and fewer self-repeating interactions to retrieve the necessary information.

CONCLUSIONS AND IMPLICATIONS

Complex cognitive tasks such as planning the sequencing of air traffic can be supported by integrating different information sources. The planning aid resulted in less repetitive, but more integrated, information retrieval compared with the traditional planning environment. Less repetitive interactions were possible because the interface itself provided the physical integration of the information. This provided one reason why Moertl et al. (2000) found that planner workload was reduced. The integration of the discrete flight information with the radar information allowed participants to develop their plan 6.3 minutes faster in the planning aid condition than in the strip condition.

Not only did the planning aid make planning easier, it also improved the quality of those plans. Moertl et al. (2000) found no difference in plan quality during the first half of the scenario (although remember, those plans were developed much faster). However, in the second half of the scenario, when participants were changing and updating their plans, planning performance was superior in the planning

Table 2.

Sequences of User Actions. The Positive Standardized Residuals Indicate That User Action Sequences Occurred in the Strip Condition More Frequently Than Expected by the Model

Sequence of user actions	Frequency in strip condition	Frequency in planner condition	Frequency predicted	Standardized residual for strip condition
Repeated distance measurement	60	14	37	3.78*
Repeated datablock adjustment	202	129	165.5	2.83*
Aircraft selection – Datablock adjustment	406	331	368.5	1.95*
Datablock adjustment – Aircraft selection	430	361	395.5	1.73*

Note. * Observed frequency is outside a 95% confidence interval around the expected frequency.

aid condition. The planning aid especially was beneficial in this situation because adapting plans to the current traffic situations depended strongly on the integration of planned sequence information with the current air traffic situation. The planning aid was designed to do exactly that. Participants could see on the radar the sequence of aircraft that they had proposed on the planning aid. By sliding the cursor over the corresponding aircraft on the planner, they could see how the plan was progressing. This visual display of a planned sequence on the radar gave the controller an important indication of where changes were required. Aircraft that were out of sequence and did not “light-up” where they should have would focus the participant on relevant decision points.

The current interface has many characteristics of an ecological interface (e.g., Effken, Kim, & Shaw, 1997; Lintern, 2000; Pawlak & Vicente, 1996; Rasmussen & Vicente, 1990; Vicente & Rasmussen, 1990). Ecological interface design argues for a perceptual formulation of user goals within the interface. The interface then facilitates actions as the user perceives his or her goals mirrored in the affordances of the interface. In this way, the interface guides users’ interactions without major intrusions or the need for automation. It replaces effortful cognitive processes with parallel, perceptual processes.

Perceptual information integration proved a successful design principle when we examined the cognitive task of planning air traffic in isolation. Future experiments should be directed at integrating strategic planning with other controlling tasks (e.g., tactical planning). Only then can it be determined if the planning aid can replace strips, or if other controller tasks still require the availability of paper flight progress strips. Also, it is important to keep in mind our focus on strategic planning and the accompanying decision to isolate the strategic planning tasks from the tactical planning tasks by assigning these responsibilities to two different individuals. It is possible that a single controller responsible for both tactical and strategic planning would not find the planning aid useful. However, that was not the goal of our project; the goal was to develop an interface for a possible future strategic controller position. A different interface may have resulted if our goal had been to develop an interface to enhance the strategic planning capabilities of a controller working a sector alone. An important next step will be to compare the planning aid with conflict probe software to determine what aspects of the air traffic control task can best be accomplished through information re-organization and what can be best handled by automation.

Recent research suggests the advantage of active control over passive monitoring in air traffic management (e.g., Metzger & Parasuraman, 2002). Our findings can be applied to the design and evaluation of interfaces that keep the controller in-the-loop. Specifically, reliance on perceptual processes could serve as an alternative to outsourcing plan development to a piece of software. This allows the controller to do what humans are good at (parallel perceptual processing) while allowing the computer to do what it is good at (organization and linking information). Simple modifications to the perceptual properties of an interface will decrease task difficulty and increase human performance without impinging on higher-level cognition. Furthermore, the perceptual optimization of interfaces should be accompanied by an empirical analysis of behavioral differences between the new and old interface. As a result, tasks that are best accomplished by a human operator could be delineated from those more appropriately left to a computer. The design and evaluation of interfaces would benefit from this process.

REFERENCES

- Arthur, W. C., & McLaughlin, M. P. (1998). User request evaluation tool (URET) interfacility conflict probe performance assessment, *MP 98W0000204*, The Mitre Corporation, McLean, VA.
- Bressolle, M. C., Benhacene, R., Boudes, N., & Parise, R. (2000). *Advanced decision aids for air traffic controllers: Understanding different working methods from a cognitive point of view*. Paper presented at the 3rd International Air Traffic Management Seminar, Napoli, Italy.
- Canning, J. M., Johansson J., Gronlund, S. D., Dougherty, M. R. P. & Mills, S. H. (1999). Effects of a novel interface design on strategic planning by en route controllers. *Proceedings of the International Symposium on Aviation Psychology, 10*, 528-33.
- Denery, D. G., & Erzberger, H. (1995). The center-TRACON automation system: Simulation and field testing, *NASA Technical Memorandum 110366*, Moffett Field, CA, NASA-Ames.
- Dougherty, M. R. P., Gronlund, S. D., Durso, F. T., Canning, J. M., & Mills, S. H. (1999). Plan generation in air traffic control. *Proceedings of the International Symposium on Aviation Psychology, 10*, 522-7.

- Effken, J. A., Kim, N. G., & Shaw, R. E. (1997). Making the constraints visible: Testing the ecological approach to interface design. *Ergonomics*, *40*, 1-27.
- Federal Aviation Administration. (1995). *Air traffic control handbook (7110.65J)*. Washington DC: Federal Aviation Administration Air Traffic Operations Program.
- Lintern, G. (2000). An affordance-based perspective on human-machine interface design. *Ecological Psychology*, *12*, 65-9.
- Metzger, U. & Parasuraman, R. (2001). The role of the air traffic controller in future air traffic management: an empirical study of active control versus passive monitoring. *Human Factors*, *43*, 519-28.
- Moertl, P. M., Canning, J. M., Johansson, J., Gronlund S. D., Dougherty, M. R. P., & Mills, S. H. (2000). Workload and Performance in FOPA: A Strategic Planning Interface for Air Traffic Control. *Proceedings of the Human Factors and Ergonomics Society*, *44*, 65-8.
- Olson, G. M., Herbsleb, J. D., & Rueter, H. H. (1994). Characterizing the sequential structure of interactive behaviors through statistical and grammatical techniques. *Human-Computer Interaction*, *9*, 427-72.
- Pawlak, W. S., & Vicente, K. (1996). Inducing effective operator control through ecological interface design. *International Journal of Human-Computer Studies*, *44*, 653-88.
- Rasmussen, J., & Vicente, K. J. (1990). Ecological interfaces: A technological imperative in high-tech systems. *International Journal of Human-Computer Interaction*, *2*, 93-110.
- Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems II: mediating "direct perception" in complex work domains. *Ecological Psychology*, *2*, 207-49.
- Vivona, R.A., Ballin, M.G., Green, S.M., Bach, R.E, & McNally, B.D. (1996). A system concept for facilitating user preferences in en route airspace, *NASA Technical Memorandum 4763*, Ames Research Center, Moffett Field, CA.