

## ATC-lab: An air traffic control simulator for the laboratory

SHAYNE LOFT, ANDREW HILL, ANDREW NEAL, MICHAEL HUMPHREYS, and GILLIAN YEO  
*University of Queensland, Brisbane, Australia*

Air Traffic Control Laboratory Simulator (ATC-lab) is a new low- and medium-fidelity task environment that simulates air traffic control. ATC-lab allows the researcher to study human performance of tasks under tightly controlled experimental conditions in a dynamic, spatial environment. The researcher can create standardized air traffic scenarios by manipulating a wide variety of parameters. These include temporal and spatial variables. There are two main versions of ATC-lab. The medium-fidelity simulator provides a simplified version of en route air traffic control, requiring participants to visually search a screen and both recognize and resolve conflicts so that adequate separation is maintained between all aircraft. The low-fidelity simulator presents pairs of aircraft in isolation, controlling the participant's focus of attention, which provides a more systematic measurement of conflict recognition and resolution performance. Preliminary studies have demonstrated that ATC-lab is a flexible tool for applied cognition research.

Air Traffic Control Laboratory Simulator (ATC-lab) is a simulated air traffic control (ATC) environment that supports the investigation of human performance of complex tasks. ATC-lab was specifically developed as a tool for which formal models of operator performance, including error, could be constructed and empirically tested. The simulator is dynamic, placing competing demands on operator attention, and it allows tight control over experimental conditions. These task characteristics allow a controlled examination of variability in performance under different air traffic scenarios, as well as the measurement of the conditions that lead to human error.

ATC simulation is currently a popular vehicle for human factors research concerning performance in complex real-world systems. The aim of ATC is to ensure the safe, orderly, and expeditious flow of air traffic. Air traffic control involves a range of activities, including conflict detection and resolution, traffic sequencing, workload management, and coordination with other controllers (Wickens, Mavor, & McGee, 1997). ATC-lab provides a simulation of the conflict detection and resolution components of the job. Aircraft are considered to be in *conflict* if they will, given their current respective speeds and bearing, violate the minimum requirements for separation at some time in the future. A number of different separation standards are used in ATC, some of them expressed in distance, and some in time. In ATC-lab, the participants are given a distance standard and are instructed to either (1) identify

the aircraft that will violate this distance standard, or (2) prevent the loss of separation by changing the speeds of the aircraft.

In this article, we describe two general versions of ATC-lab. The first represents a medium-fidelity simulation of en route ATC. In the medium-fidelity version, participants monitor a radar-like screen. There are typically several aircraft on the screen at once, some of them in conflict. These conflict events have variable onsets and durations and may overlap in time. As is typically the case in real ATC, however, the majority of aircraft are not in conflict. The task, therefore, has a visual search requirement, and it requires selective attention and prioritization, as well as decision making.

The second version of ATC-lab is a lower fidelity simulation that allows the experimenter to isolate specific components of the task, in the absence of a visual search requirement or competing demands on attention. The task requires participants to make decisions regarding pairs of aircraft presented in isolation. These decisions can include identifying whether a given pair of aircraft are in conflict or not, and deciding how to resolve such conflicts. The second version of ATC-lab, therefore, allows the experimenter to examine in detail the factors that influence judgment and decision making in real time.

The aim of this article is not only to introduce ATC-lab, but to indicate how the simulator can be used in applied cognition research. First, we will describe the types of ATC simulators currently available and discuss how ATC-lab is related to these. Next, we will outline features of the different versions of ATC-lab and the cognitive functions that it can be used to investigate. We will then detail the tools that are available to provide the experimenter with technical support. In general, these tools are used to script air-

---

The development of ATC-lab was supported in part by Large Grant A00104975 from the Australian Research Council. Correspondence concerning this article should be addressed to Shayne Loft, who is at the School of Psychology, University of Queensland, St. Lucia, Brisbane, QLD 4072, Australia (e-mail: s.loft@humanfactors.uq.edu.au).

craft traffic scenarios, select scoring and feedback options, and control how the data are recorded. Finally, we will describe the ATC-lab's system requirements.

### ATC Simulated Task Environments

A number of different ATC simulators are currently used in research. These differ in the kinds of activities that they simulate and in the realism of the air traffic scenarios that they present.

**High-fidelity ATC simulators.** In high-fidelity simulations, participants are required to handle simulated air traffic similarly to real air traffic controllers. Examples of high-fidelity simulators used in human factors research include the FAA Academy Training Simulator (Jones & Endsley, 2000) and TRACON (Ackerman, 1992). These simulators typically require participants to coordinate the safe and expeditious movement of air traffic by issuing commands (e.g., turn, ascend) to aircraft in their sector. Various types of aircraft are used, flying under either instrument or visual flight path rules, and the simulations include arrivals, departures, and overflights. In some instances, participants are also required to handle additional aspects of ATC, such as hand-offs and weather updates.

High-fidelity ATC simulators are particularly useful for assessing global controller performance and for training controllers and control teams because they mimic the complexity of real ATC, simultaneously overcoming the ethical and logistical constraints associated with field studies (DiFonzo, Hantula, & Bordia, 1998; Gray, 2002). The task conditions that the controller must face, such as the number of aircraft, the type and mix of aircraft, the number of flight paths and intersections, and the number and types of altitudes used, can be varied systematically. In many research situations, however, it is important to know how these factors and their combinations make one ATC situation more difficult than another, and how they contribute to performance on specific ATC task components such as conflict detection and conflict resolution. Researchers may also wish to test theoretical constructs developed from simple laboratory experimentation. High-fidelity simulations are not suited to such investigations because their complexity does not permit the researcher to isolate the effects of independent variables (Gray, 2002).

**Medium-fidelity ATC simulators.** Researchers assessing performance on individual task components or investigating specific theoretical questions use computer-generated ATC environments that present medium-fidelity simulations of conditions encountered in the field because these provide higher levels of experimental control (Brehmer & Dorner, 1993; DiFonzo et al., 1998; Omodei & Wearing, 1995). Medium-fidelity simulations allow experimenters to isolate one or several independent variables and measure their effects on a chosen dependent variable, while systematically controlling for confounds and other extraneous variables. Another advantage is that participants with little experience of the task environment can quickly learn how to perform medium-fidelity tasks (Gray, 2002).

Medium-fidelity ATC simulators have been developed for theoretically based research concerning variations in human performance in complex task environments. Examples include investigations of prospective memory (Stone, Dismukes, & Remington, 2001) and of the effects of time demands and attentional demands on operator information processing load (Hendy, Liao, & Milgram, 1997). Other researchers have used medium-level ATC simulators to isolate key variables needed in the study of applied problems—such as controller performance in free flight (Metzger & Parasuraman, 2001; Remington, Johnston, Ruthruff, Gold, & Romera, 2000). Most of these ATC simulators are generally flexible in the number and type of aircraft presented, what is displayed on aircraft data blocks, the temporal (e.g., onset time) and spatial (e.g., angle and speed) properties of aircraft events, and the types of responses that participants are required to make (e.g., detect conflicts, control traffic/avoid separation loss, hand-off aircraft).

Nevertheless, ATC-lab offers several advantages over the existing medium-fidelity ATC simulators. First, ATC-lab was specifically designed to make it easy for researchers to pursue a wide range of applied and theoretical research questions. To this end, the software facilitates manipulation of a variety of task parameters and the creation of standardized air traffic scenarios, using precise levels of experimental control. In contrast, the systems mentioned above were designed to answer more specific research questions, and it would be difficult for researchers to adapt such simulators for use in their own experiments. Second, ATC-lab provides the experimenter with flexibility in the fidelity of simulation. ATC-lab runs lower fidelity ATC simulations that can be used in the investigation of simpler phenomena that are hypothesized to contribute to the effects found in medium-fidelity ATC-lab simulations. In addition, the two-tiered approach to fidelity offers researchers the means to examine variables of interest in the more complex environment and immediately port them into a lower level simulator where greater experimental control is available.<sup>1</sup> The following sections outline the medium- and low-fidelity versions of ATC-lab and the types of research questions that can be investigated with them.

### The Medium-Fidelity Version

Figure 1 illustrates an example of a sector display generated by the medium-fidelity version of ATC-lab. Participants assume the role of a controller responsible for the safety of air traffic within the sector. The sector is made up of a series of flight paths; the flight paths contain a number of waypoints and airports. The user interface loosely resembles that currently used by air traffic controllers in Australia. The flight paths, waypoints, and airports are presented in dark gray font against a lighter background. Small green circles symbolize aircraft. Each aircraft on the sector has a green data block attached to it that displays its call sign, type, speed, and the designated route. The route information on the data block contains the first two letters of each waypoint

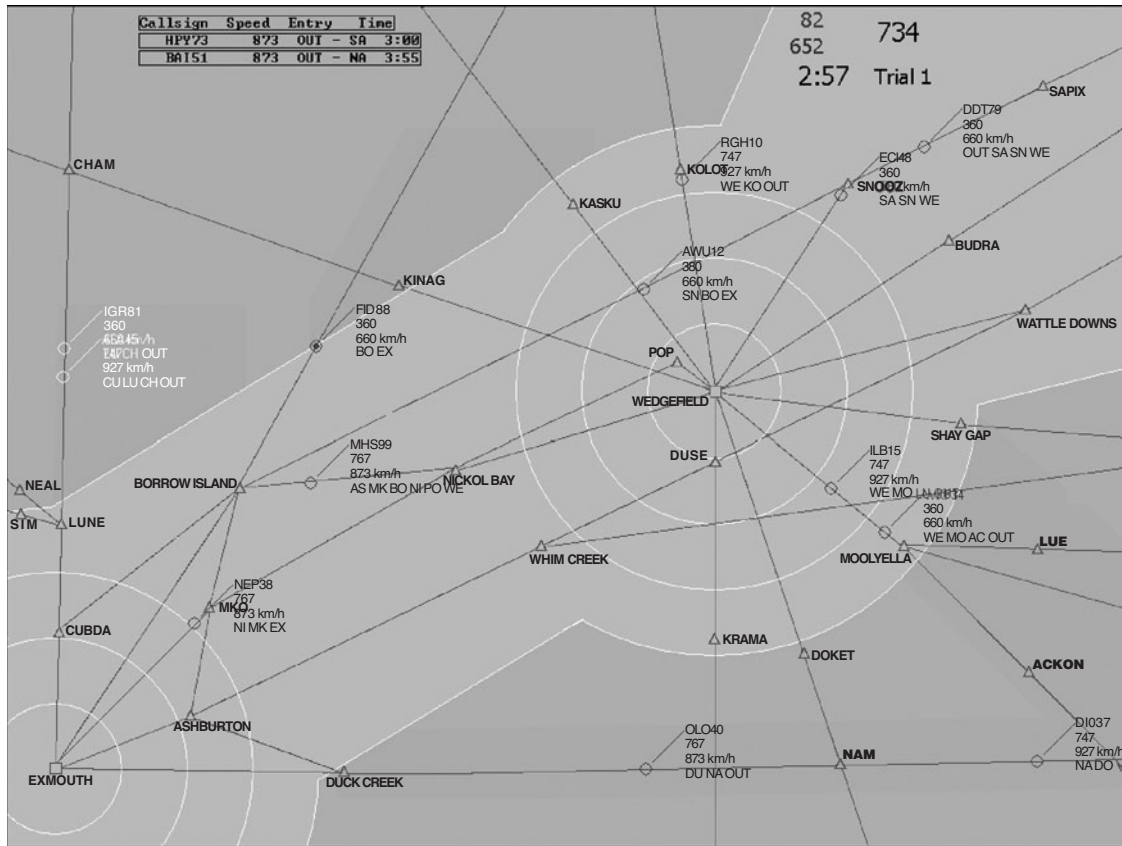


Figure 1. The display in the medium-fidelity simulations of ATC-lab. Triangles represent waypoints. The two squares denote airports. Typically, from 10 to 14 aircraft are active on the screen at any one time. The trial number, time in trial, error penalty score, efficiency penalty score, and total running score are displayed to participants continuously in each trial.

and/or airport that the aircraft will pass in its sequence of travel. All aircraft proceed through their sequence of waypoints and either land at one of the two airports or exit the sector on one of the outbound flight paths. At the beginning of any trial, aircraft are already on the screen at various points in the sector. Throughout the duration of a trial, new aircraft enter inbound from off-screen origins or take off from airports. Although the simulator can also represent the altitude of the aircraft, we have not included this in any current experiments. Incorporating altitude substantially increases the complexity of the task for participants.

In order to guarantee the safety of aircraft in their sector, participants have to ensure that all aircraft conform to specific separation standards. The separation standard used in ATC-lab is 5 km.<sup>2</sup> If the distance between two aircraft ever falls below 5 km, the aircraft on the display turn yellow and an alert is sounded. When the 5 km separation standard is reestablished, the aircraft turn green again and the alert sound stops.

The experimenter has control over a wide range of simulator features. These include (1) the sector map flight paths and intersection points, (2) the responses that the participants are required to make, (3) the timing of

the responses, and (4) the temporal and spatial properties of the events that are scripted within the scenarios. It is important to control the temporal and spatial properties of aircraft events so that independent variables of interest can be systematically manipulated without the presence of confounding factors. The combination of simulator features used depends upon the types of research questions being investigated. Two configurations of the medium-fidelity simulator are presented below. These tasks differ primarily in the behavioral responses that they require of participants.

**Conflict detection.** In the conflict detection task, participants watch sets of scripted air traffic scenarios evolve over time and are asked to identify conflicts between aircraft as quickly and as accurately as possible. The participants register pairs of aircraft as conflicts by clicking a small rectangular box displayed in a corner of the sector screen and then clicking consecutively in the green circles representing the two offending aircraft. This provides an accurate measure of conflict detection time to the nearest second. A running score is displayed in the upper right corner of the screen. Points are awarded on the basis of how quickly conflicts are detected (e.g., 40, 30, 20, or 10 points, with more points awarded the faster

the response); and points are deducted for misses (e.g., 10 points) and false alarms (e.g., 10 points). In all ATC-lab simulations, the experimenter sets the scoring formulae and can change them. The presented aircraft proceed in real time, but their progress remains unchanged by the participant, who can only observe it. The experimenter can control factors such as the number and characteristics of conflicts and near misses, the location and timing of events, and the number of aircraft on the screen. Another feature of the simulator is that the task can be frozen at specific moments. Self-report scales can then be presented to the participants in order to obtain subjective ratings of relevant psychological constructs. The experimenter has control over when the task is frozen, the number of questions presented, the wording of the questions, and the number of points on the self-report scale. Examples of constructs that can be assessed include perceptions of effort, difficulty, workload, situation awareness, and self-efficacy.

A number of different types of theoretical issues or questions can be examined with this version of the simulator. Examples include the influence of memory for prior instances on conflict detection performance (Loft, Humphreys, & Neal, in press), and the effects of time pressure on performance in a dynamic task environment (Neal, Kwantes, & Loft, 2002). For example, when studying the effects of episodic memory on conflict recognition performance, an experimenter might examine factors such as the similarity between the current event and previously seen events, the similarity between the contexts associated with events, the recency of specific events, the frequency of different events, and workload.

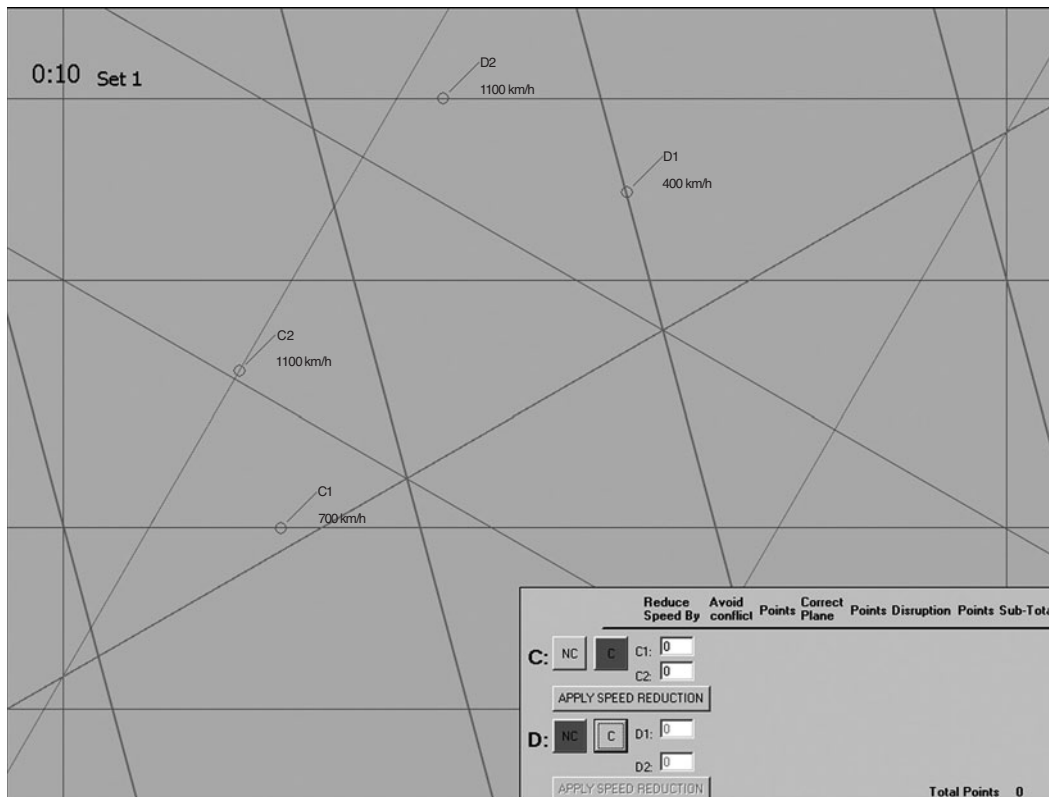
**Conflict resolution.** In the conflict resolution task, participants are given control over the speeds of aircraft. Figure 1 illustrates the display. To change the speed of an aircraft, participants select the aircraft, right-click to activate the speed drop-down menu, and left-click to select the desired speed. The objective is to guide aircraft to their destination as safely and as efficiently as possible. To do this effectively, participants need to identify conflicts between aircraft in sufficient time to avoid loss of separation, while ensuring that they do not jeopardize efficiency by slowing the aircraft unnecessarily. To facilitate this process, an entry box is provided that displays the time(s) at which new aircraft will enter the sector, and the position at which they will enter. The scoring system is based on safety and efficiency. Error penalty scores are based on the number of conflicts that occur and the length of time when aircraft are in conflict. Efficiency penalty scores, which accrue whenever an aircraft is not traveling at its maximum speed, are directly proportional to the speed reduction that has been applied to the aircraft. As with the conflict detection version of this task, the task can be frozen at specific moments, and self-report scales can be presented to the participants. This version of ATC-lab has been used to investigate the effects of ability and goal orientation on the relationship between effort and performance during skill acquisition (Yeo & Neal, 2004).

### The Low-Fidelity Version

Figure 2 illustrates an example of a display used in the low-fidelity version of ATC-lab. For low-fidelity simulations, ATC-lab presents one or more pairs of aircraft to participants in isolation. The participants are required to make conflict status classifications and/or decide on the most efficient speed changes for avoiding conflicts. The sector map is less realistic than the version used in the medium-fidelity simulator. However, this map provides the experimenter with a much wider range of options for generating conflicts and near misses. For example, the map provides a set of angles that are replicated at different points on the screen and at different degrees of orientation. The experimenter is able to select the waypoint toward which the aircraft are heading, and the starting positions and speeds of the aircraft. This gives the experimenter control over variables such as the angle of convergence, time until minimum separation, distance between the aircraft at the start of the event, distance between the aircraft at their point of minimum separation, order of passing (whether the slower aircraft passes in front of the faster aircraft, or vice versa), and the position and orientation of the event. It is difficult to develop scripts for the medium-fidelity simulator that systematically manipulate these types of variables. The structure of the flight paths and the presence of distractor aircraft place constraints on the types of events that can be scripted, and substantially increase the complexity of the scripting process. These constraints are removed in the low-fidelity version. Furthermore, the removal of the distractor aircraft eliminates the visual search requirement of the task, effectively controlling the participant's focus of attention. Two versions of the low-fidelity simulator, with different response tasks, are presented below.

**Conflict detection.** In the low-fidelity conflict detection task, participants are presented with pairs of aircraft converging on a common point of intersection. The participants decide whether or not the two aircraft will conflict, as quickly and as accurately as possible. These conflict status decisions are registered via a response box. Points are awarded on the basis of accuracy and speed of classification. For transfer trials, single pairs of aircraft are presented one at a time and participants are given a set time to make a classification (e.g., 30 sec). The experimenter has the option of allowing participants to make a response at any time during the exposure interval, or of asking them to respond at the end of the exposure interval.

The key differences between the low- and medium-fidelity versions of the conflict detection task are that the low-fidelity version requires participants to (1) make decisions in response to designated pairs of aircraft, and (2) indicate whether each pair of aircraft is, or is not, a conflict. In contrast, the medium-fidelity simulations contain filler aircraft (i.e., aircraft not involved in conflicts or near misses) and only require participants to make a response if they believe any pair to be in conflict. These differences have a number of implications for the way in which the data are analyzed and interpreted. First,



**Figure 2.** The display in the low-fidelity simulations of ATC-lab. Participants enter their classifications and make speed reductions on the response panel provided. Running scores are displayed to participants continuously in the response panel. The time in the trial and training set number are presented in the top corner of the screen.

signal detection theory (Stanislaw & Todorov, 1999) can be used to analyze participants' sensitivity (i.e., ability to discriminate between conflicts and nonconflicts) and response bias (i.e., tendency to favor one of these responses over another) under different experimental conditions in low-fidelity simulations. Owing to the inclusion of filler aircraft in the medium-fidelity simulations, the experimenter does not know the total number of opportunities that a participant has to false alarm, and false alarm rates cannot be calculated. If the hit rate varies between two conditions, it is not clear whether the conditions differ in sensitivity, response bias, or both. Second, given that medium-fidelity simulations do not require responses for nonconflicts, it is unclear whether the participant has made an incorrect conflict status decision or has not attended to the aircraft at all when a conflict is missed. Finally, in medium-fidelity simulations, reaction times cannot be measured for correct rejections, making it difficult to test the assumptions of some types of models (e.g., random walk models: Ratcliff, Van Zandt, & McKoon, 1999).

To date, experiments using the conflict detection version of the low-fidelity simulator have focused on the way in which participants learn from individual examples (Loft, Neal, & Humphreys, 2002). Experimental conditions have been designed for investigating how dif-

ferent types of training examples, and their schedules of presentation, affect subsequent transfer and retention. The task allows the relationship between training and transfer stimuli to be systematically manipulated, in order to measure the transfer of skill from trained examples of events to new examples that differ on any number of structural and contextual dimensions.

**Conflict resolution.** The dependent variable here is conflict resolution performance. Decisions involve selecting either of the two aircraft involved in an upcoming conflict and nominating a speed reduction value in order to avoid it. The experimenter has the option of highlighting the pairs that are in conflict, so that performance is purely a function of decision accuracy and is not influenced by participants' ability to detect the conflict. Alternatively, the experimenter may choose not to highlight the pairs that are in conflict. In this case, performance will be influenced by both the ability to detect conflicts and the ability to resolve them.

Figure 2 presents the display for the low-fidelity conflict resolution task. The experimenter has control over whether one pair or two pairs of aircraft are presented. The participants change the speed of the aircraft with the response box. Their task is to select the aircraft that they wish to slow down, type in the amount that they wish to reduce the speed by, decide when to apply this speed re-

duction, and press a button to apply it at the appropriate time. Conflict resolution decisions are assessed according to a number of criteria. These include selection efficiency, adjustment efficiency, accuracy, and response latency. Selection efficiency is a dichotomous variable that refers to whether or not the participant has selected the aircraft that requires the smaller speed reduction. Adjustment efficiency refers to the extent to which the disruption of the flow of air traffic is minimized as a result of the speed change. Accuracy refers to whether or not the conflict is resolved, and the response latency refers to how long the participant takes to make a speed reduction decision.

### Experimenter Interface

ATC-lab incorporates extensive support for the experimenter. Application tools include script files that determine what happens on the screen, a script developer and scenario tester, scoring algorithms, and data log files. These tools are described below.

**Air traffic scenario scripts.** Each ATC-lab experiment comprises a series of trials. The experimenter has control over the number and length of the trials. Air traffic scenario scripts are files that specify the events that occur during each of these trials. The script allows the experimenter to specify details for each aircraft. Specifically, the script contains fields listing (1) the trial number and time in the trial that the aircraft enters, (2) the call-sign of the aircraft, (3) the type of aircraft, (4) the path that the aircraft will travel, represented by a series of  $x, y$  coordinates, (5) the minimum and maximum speeds of the aircraft, (6) the starting speed of the aircraft, and (7) route information consisting of a series of waypoints. The values for each aircraft's call sign, type, current speed, and route are derived from these scripts and are displayed to the participant on the data block adjacent to the aircraft. Speed values can be changed during a trial in response to input from participants.

Taken together, the details contained in each air traffic scenario script determine (1) the number and type of aircraft on the screen at each point in time, (2) the speeds and times at which aircraft enter the sector screen, (3) the direction in which the aircraft travel, and (4) the number and specific characteristics of conflicts, near misses, and filler aircraft.

**Script developer.** The primary function of the script developer software is to precisely control the spatial and temporal characteristics of aircraft events scripted for ATC-lab trials, given parameters entered by the experimenter. Figure 3 presents the display of the script developer.

The most common order in which parameters are entered is as follows.<sup>3</sup> First, the experimenter chooses a map from the map options menu and the script developer displays the map in the right-hand corner of the interface. The map option menu also allows the experimenter to zoom in, show the  $x, y$  coordinates, and edit maps. The next parameters to be entered are contained in the plane (i.e., aircraft) options menu. The experimenter enters the details of the pair of aircraft that will be involved in the

event. This information includes the aircraft type, speeds, and call sign.

The information entered into the conflict options menu varies on a number of parameters. These include the point of intersection (i.e., the  $x, y$  coordinate where the aircrafts' paths intersect) and the DOMS (distance of minimum separation), which represents how close in kilometers the two aircraft will get to one another during the scripted event. The experimenter has the option of specifying either the time until conflict (if the event is a conflict), or the time until minimum separation. These two variables refer to the number of seconds that elapse between the time when the aircraft first appear together on the screen and the time when the 5-km separation standard is first violated or separation is at a minimum (i.e., when DOMS is reached), respectively. The user must also specify which of the two aircraft will reach the intersection first.<sup>4</sup> After these parameters have been entered, the experimenter returns to the plane options menu to enter the pre- and postintersection  $x, y$  coordinate points through which the aircraft will travel. The experimenter can view the proposed travel route of each aircraft by clicking on the view route button. To make specifying parameters easier, the program is also capable of suggesting and plotting, for any given desired point of intersection, every possible combination of pre- and postintersection routes, using the auto fill options menu.

Once all the details of the event have been entered, the script developer produces precise starting  $x$  and  $y$  coordinates (i.e., starting positions) for the two aircraft. These are based on the aircraft speeds, the intersection point, DOMS, the seconds until conflict or minimum separation, and the aircraft that was chosen to reach the intersection first. The script developer automates the process of scripting conflicts and nonconflicts involving pairs of aircraft. Without this tool, scripting aircraft events would be a laborious, time-consuming task, involving manual calculations or trial and error scripting. The ability to control the distance of minimum separation and the time until minimum separation is a valuable standardization feature for all ATC-lab experiments, because both variables exert a strong influence on performance. Existing simulation packages, such as TRACON, do not give the experimenter this level of control.

**Scenario tester.** Scripting trials for medium-fidelity simulations is an incremental process. That is, because the task is dynamic, it is difficult to imagine what the sector screen will look like until the trial is run and visual verification is made. The scenario tester is capable of running a trial at up to 50 times its normal speed. This allows the experimenter to quickly check details such as aircraft entry times, event types and timings, filler aircraft spacing, the number of aircraft on the screen, and the general spread and pattern of air traffic. The experimenter also has the option of using an automatic freeze function that pauses the screen each time a conflict occurs, allowing the verification of conflict timings through reference to the on-screen clock. Other functions include the ability to increase or reduce the speed of the trial during

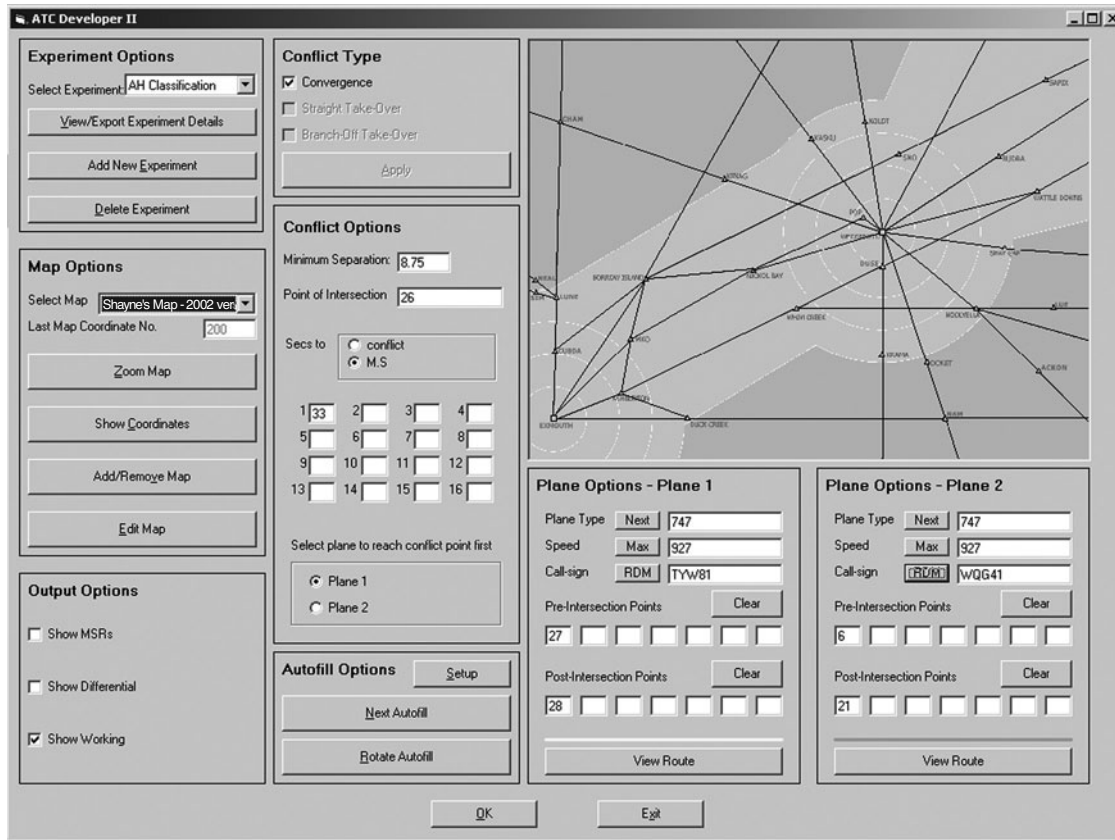


Figure 3. The script developer interface.

playback, and to accurately verify the DOMS and the time at which it is reached.

**Scoring.** ATC-lab simulations display scores to participants. Typically, scoring is based on the safety and efficiency of air traffic control performance. However, the type of scoring system used can vary greatly between experiments. This flexibility is achieved by varying the scoring formulae. For example, in low-fidelity conflict detection experiments, participants who classified an aircraft pair correctly were awarded 40 points if they responded during the first quarter of the available time, 30 points for responding in the second quarter, 20 points for the third quarter, and 10 points for the last quarter. If the classification was incorrect, these same points were deducted (Loft et al., 2002).

**Data log files.** The specific contents of the data log file recorded for each participant vary markedly according to the type of experiment being run. Nevertheless, these text files collect two types of general data. The first type of data are the system parameters. These data comprise details of the air traffic scenarios executed in each trial, such as the types, timings, and durations of aircraft events. The participants' actions are the second general source of data. Such data include participant numbers, reaction times, accuracy, subjective ratings, speed changes,

and errors. The log files generated for each participant can be imported into statistical packages such as SPSS.

**System requirements.** The ATC-lab task runs on a Pentium 450 with 128MB of RAM running Microsoft Windows 2000. The source language for ATC-lab is Visual Basic (Version 6). The task does not require any additional software or hardware for running its various applications. The screen shows a map representation of airspace 160 km east to west by 120 km north to south. The program updates and displays each aircraft's position in the sector once every second on the basis of the aircraft's current speed and direction values. These direction and speed values are preset by a simulation script that dictates the series of *x, y* coordinates through which the aircraft will travel and the speed at which the aircraft will be traveling. Of course, in medium-fidelity simulations participants may change an aircraft's speed during a trial.

**Conclusions**

ATC-lab is a low- and medium-fidelity ATC simulation laboratory package.<sup>5</sup> ATC-lab simulations present complex cognitive tasks to participants and share many of the practical characteristics of real ATC and of complex real-world task environments in general. The air

traffic scenarios used in the medium-fidelity simulator are designed to assess performance constructs central to the controller's job, such as detecting aircraft conflicts, maintaining aircraft separation, and maintaining situational awareness. The low-fidelity simulator is designed to control for operator attention, breaking down the task into discrete components, and allowing the direct measurement of conflict detection and conflict resolution performance. The ATC-lab developer program provides the experimenter with the capability of building and running scripts that present simulated ATC task environments to participants. This paper has described different types of research questions that can be investigated with ATC-lab, illustrating its flexibility as a tool for applied cognition research.

#### REFERENCES

- ACKERMAN, P. L. (1992). Predicting individual differences in complex skill acquisition: Dynamics of ability determinants. *Journal of Applied Psychology*, *77*, 598-614.
- BREHMER, B., & DORNER, D. (1993). Experiments with micro-simulated worlds: Escaping both the narrow straights of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior*, *9*, 171-184.
- DiFONZO, N., HANTULA, D. A., & BORDIA, P. (1998). Microworlds for experimental research: Having your (control and collection) cake, and realism too. *Behavior Research Methods, Instruments, & Computers*, *30*, 278-286.
- GRAY, W. D. (2002). Simulated task environments: The role of high-fidelity simulations, scaled worlds, synthetic environments, and microworlds in basic and applied cognitive research. *Cognitive Science Quarterly*, *2*, 205-227.
- HENDY, K. C., LIAO, J., & MILGRAM, P. (1997). Combining time and intensity effects in assessing operator information-processing load. *Human Factors*, *39*, 30-47.
- JONES, D. G., & ENDSLEY, M. R. (2000). Overcoming representational errors in complex environments. *Human Factors*, *42*, 367-378.
- LOFT, S., HUMPHREYS, M., & NEAL, A. (in press). The influence of memory for prior instances on performance in a conflict detection task. *Journal of Experimental Psychology: Applied*.
- LOFT, S., NEAL, A., & HUMPHREYS, M. (2002). *Learning and transfer in an applied visual spatial task*. Paper presented at the HF2002 Human Factors Conference, Melbourne.
- METZGER, U., & PARASURAMAN, R. (2001). The role of the air traffic controller in future air traffic management. *Human Factors*, *43*, 519-528.
- NEAL, A., KWANTES, P., & LOFT, S. (2002, November). *Development of a model for prediction of conflict detection performance in a simulated air traffic control task*. Paper presented at the meeting of the Defence Human Factors Special Interest Group (DHFSIG), Melbourne.
- OMODEI, M. M., & WEARING, A. J. (1995). The Fire Chief microworld generating program: An illustration of computer-based microworlds as an experimental paradigm for studying complex decision-making behavior. *Behavior Research Methods, Instruments, & Computers*, *27*, 303-316.
- RATCLIFF, R., VAN ZANDT, T., & MCKOON, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261-300.
- REMINGTON, R. W., JOHNSTON, J. C., RUTHRUFF, E., GOLD, M., & ROMERA, M. (2000). Visual search in complex displays: Factors affecting conflict recognition by air traffic controllers. *Human Factors*, *42*, 349-368.
- STANISLOW, H., & TODOROV, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, *31*, 137-149.
- STONE, M., DISMUKES, K., & REMINGTON, R. (2001). Prospective memory in dynamic environments: Effects of load, delay, and phonological rehearsal. *Memory*, *9*, 165-176.
- WICKENS, C. D., MAVOR, A. S., & MCGEE, J. P. (1997). *Flight to the future: Human factors in air traffic control*. Washington, DC: National Academy Press.
- YEO, G., & NEAL, A. (2004). Multilevel analysis of effort, practice and performance: Effects of ability, conscientiousness and goal orientation. *Journal of Applied Psychology*, *89*, 231-247.

#### NOTES

1. We would like to acknowledge Deborah Boehm-Davis for this making this point to us in her review of the manuscript.
2. ATC-lab was designed in Australia and uses the metric system for units of measurement. Changing these units to nautical miles or miles per hour would require re-programming by the researcher.
3. The order in which the parameters are entered into the script developer can vary, depending on the experimenter's personal preference.
4. In situations where DOMS = 0, aircraft will reach the intersection point at exactly the same time.
5. Research groups interested in using ATC-lab for noncommercial purposes should contact Shayne Loft at the Key Centre for Human Factors and Applied Cognitive Psychology. Upon receipt of your request, you will be sent a copy of the script developer and other experimenter tools as well as the medium- and low-fidelity programs and samples.

(Manuscript received January 7, 2003;  
revision accepted for publication January 4, 2004.)