


NASA/TM—2003–212261



A Cockpit Display Designed to Enable Limited Flight Deck Separation Responsibility

*Walter W. Johnson and Vernol Battiste
Ames Research Center, Moffett Field, California*

*Sheila Holland Bochow
San Jose State University Foundation, San Jose, California*

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

March 2003

Available from:

NASA Center for Aerospace Information
7121 Standard Drive
Hanover, MD 21076-1320
301-621-0390

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
703-605-6000

Summary

Cockpit displays need to be substantially improved to serve the goals of situational awareness, conflict detection, and path replanning, in Free Flight. This paper describes the design of such an advanced cockpit display, along with an initial simulation based usability evaluation. Flight crews were particularly enthusiastic about color coding for relative altitude, dynamically pulsing predictors, and the use of 3-D flight plans for alerting and situational awareness.

Introduction

Research and development on Cockpit Displays of Traffic Information (CDTIs) has received a boost by the FAA's Free Flight Initiative and NASA's Advanced Air Transportation Technologies (AATT) program. Both of these programs seek to advance the understanding of Free Flight issues, and most (although not all) existing descriptions of Free Flight give the flight crew a role as primary players in real-time navigation and conflict avoidance. As a result, effective CDTI design is a fundamental issue. It is fundamental because the flight crew's effectiveness will depend upon good awareness of the traffic situation, and this, in turn, will depend upon a good display of the traffic situation.

At the present time the Traffic Collision Avoidance System (TCAS) provides the most prevalent cockpit display of the traffic situation. During the research and development of the TCAS CDTI, the goal was to design a display which would generate confidence in the automated resolution advisories produced by TCAS. Thus, while apparently effective (although see Pritchett and Hansman, 1997), these displays were not designed to support the range of flightdeck activities being contemplated within Free Flight. For example, TCAS was designed as an emergency collision avoidance system, while

the operational concept behind Free Flight systems typically requires earlier maneuvering to ensure that minimal separation (e.g. 5 NM lateral, or 1000 ft vertical) is maintained. This is called conflict avoidance.

Comparisons of the two systems can be made at a more detailed level. On the one hand the TCAS system was designed to utilize transponder-based surveillance data of relatively limited bandwidth (azimuth and range derived from a beacon system, plus some additional broadcast information for altitude and aircraft ID), relatively limited range (~40-50 NM), and relatively limited accuracy (the accuracy of beacon derived information drops with range). On the other hand, the proposed newer systems would directly broadcast and receive more information (ground position, altitude, track vector/heading, ground speed, vertical speed, aircraft ID, and the location of 2 or more future flight plan waypoints). Furthermore, the range of these newer systems will extend to at least 120 NM. Finally, since these newer systems use entirely direct broadcast, and not beacon-based, information, accuracy will not depend on range; and because these systems use the satellite based Global Positioning System (GPS), information for ground position is precise.

Therefore, several new issues must be addressed in designing CDTIs appropriate for Free Flight. First, the potential number of aircraft on a display has increased dramatically. If a typical TCAS range is 40 NM, then a range of 120 NM increases the potential traffic, and potential clutter, nine-fold. And while selective filtering of aircraft from displays is always an option, the desire for increased situational awareness may run counter to this solution. Thus the management of clutter via other means than filtering out traffic is one issue.

Second many versions of Free Flight envision the flight crew as helping to specify, or design,

conflict resolutions, not just verify/evaluate automated resolutions. Combined with this, the flight crew may also manage or oversee how these resolutions are coordinated with other involved aircraft and with the air traffic controller (controller). Therefore the flight crew must have avionics that are adequate for displaying traffic, intent, and alert status information, devising resolutions, sharing these resolutions with other involved parties, and then implementing these resolutions. The most important goal of the present work is to develop a display that will enhance traffic awareness, and makes possible many of the proposed roles for the flight crew in Free Flight. In the short term this work could provide other researchers with displays, or display concepts, that they may use when simulating Free Flight scenarios. In the long term this work will hopefully contribute concepts to the design of future CDTIs.

The outline for the rest of the paper is as follows. First, the display will be described, along with its functionalities. This will comprise the main body of this report. Next, a simulation study in which the display was evaluated will be described. Finally, subjective reports from the pilots participating in the simulation will be examined. Performance data from this simulation will be analyzed in a subsequent report.

Cockpit Situational Display

This section describes a combined CDTI, conflict alerting, and flight path replanning system that was designed to support an integrated depiction and resolution activity. This system, which is an outgrowth of a display described in a previous paper (Johnson, Battiste, Delzell, Holland, Belcher, &

Jordan, 1997), will be referred to as the Cockpit Situational Display (CSD). Since depiction of the traffic and the resolution planning requires the same spatial framework, it is simply good human factors design to integrate these components within the same spatial display.

The design of the CSD utilized the basic format and size of the Navigation Display Map Mode in the Boeing 747-400 (i.e. a square 10 inch diagonal display with 1024 x 1024 resolution). The changes/additions to this basic display, seen in Figure 1, were as follows: (1) ownship was presented as a filled chevron, (2) closed chevrons were used to depict traffic aircraft, and (3) a toolbar to control display functions was included at the bottom of the navigation display. Significant features and design elements of the display included: (1) the presentation of traffic and ownship intent information in the form of flight plans, (2) formats for the presentation of aircraft and flight plan altitude information, (3) the presentation of traffic and ownship predictive information, (4) the inclusion of clutter management features, (5) the presentation of conflict alerts that would function in an integrated fashion with the path replanning tool, and (6) the path replanning tool itself. In addition, controls for the display functions were provided using two redundant input systems. First, two control panels were fabricated (see Figure 2). The flight crew could use these to control almost all display functions. Second, touch pads were mounted on the arms of the seats of the Captain and First Officer. These touch pads had two buttons mounted on them, and, together with display symbology, could be used to control all display functions. The details of the display design are given below.

Traffic Information

The display was designed to provide a maximum of information concerning surrounding traffic, and to be presented in a

picture that could be easily and quickly comprehended "at-a-glance". The goal was to keep heads down time to a minimum. While most of the additions to the basic 747-400 navigation display provided the



Figure 1. The CSD display was the same basic format as the Navigation Display Map Mode in the Boeing 747-400, with the following modifications/additions. ownship, in the lower middle portion of the display, was presented with a filed chevron. Closed chevrons depicted the position of traffic. The toolbar at the bottom of the display was used by touch pad input and controlled the presentation of information and the ARAT.

information as standard display elements, to reduce clutter, much of the information was provided only upon the request of the pilot. Pilots were provided with a touch pad for obtaining the traffic information on individual aircraft (i.e., data blocks and flight plans). When the crews wanted information to be presented on all of the aircraft, such as predictors, they had the options of using the control panels located on the console, or they could use the touch pad to operate a toolbar located on the navigation display. Both the touch pad and control panels provided the same functionalities, and were provided to accommodate user preference.

Intent Information

Information about traffic intent was provided by 3-D flight plans. These plans contained the intended 3-D routes of ownship and of traffic. In order to display the flight plan for an aircraft the crew had to use the touch pad to place a cursor over that aircraft (either traffic or ownship), and then use the right touch pad button to toggle the flight plan on. The same operation was used to toggle the flight plan off.

Depiction of flight plans corresponded to the typical flight plan combination of waypoints and legs, with two exceptions. First, periods of time during which aircraft were planning to climb or descend were included as separate legs with separate symbology indicating the start and end of the period of altitude change (this is further described in the following section on Altitude Information). Second, there were non-standard names for waypoints at non-standard locations. A non-standard location was a location not associated with a known VOR or waypoint. These non-standard locations were created by the crew when conducting flight replanning, or were waypoints inserted automatically by the system on legs greater than 100 miles. (These latter waypoints were necessary in order to make the flat earth model of the CSD handle the great circle routings generated by the 747-400

simulator's flight management system.) It was necessary to have names for all waypoints in order to facilitate communications with the controller during times when the controller was reviewing proposed flight plan changes (this process will be described in a later section). The names for non-standard waypoints always took the form of the name of the last standard waypoint passed on the route, plus the number of nautical miles between the referenced waypoint and the newly inserted point. Thus, if a waypoint was created and inserted 99 miles beyond OAL and the aircraft was flying direct between those two points, the name for that waypoint would be 'OAL-099' (see Figure 1 for example).

Altitude Information

A significant challenge in developing at-a-glance symbology is conveying altitude information. The standard two-dimensional map displays used by air traffic controllers, and the similar TCAS displays utilized by pilots, employ textual tags to show the altitude of aircraft. This text has drawbacks. First, it requires focal attention to read each tag, and this increases workload and the time needed to effectively scan a display. Second, since the altitude information within individual aircraft tags are not graphically rendered, it is difficult to extract relationships among the separate elements within the display. Attempts to provide graphical indicants of altitude on a CDTI have been attempted by Ellis and McGreevy (1987) with a perspective display, and by Merwin and Wickens (1996) with a co-planar display that showed both a standard horizontal view and a profile view. Both of these displays have strong and weak points. Perspective displays graphically integrate all of the information on the display, but they have been shown to be prone to perspective ambiguity with respect to range to a target. This ambiguity is eliminated in the co-planar presentation of Wickens, but this concept requires the viewer to cognitively integrate the information across the two separate displays.

The present design addressed these issues using a different method of displaying altitude information. Color was used to display basic categorical information for relative altitude, while text was used to provide more detailed altitude information (see Figure 3). Since humans extract color information pre-attentively (i.e. process it nearly effortlessly and extremely fast), color is a powerful graphical technique for providing at-a-glance information. Color was used to segregate the traffic into groups which were flying above, below and co-altitude with ownship: co-

altitude traffic was white, traffic above ownship was blue, and traffic below ownship was green (see Johnston, Horlitz, and Edmiston (1993) for an application of aircraft symbol color coding to ground air traffic displays). Debriefings of pilots in an earlier study (Johnson, et al, 1997) also suggested pilot preference for this type of categorical approach. This color scheme was also extended to depictions of the flight plans (see Figure 4). The legs of the flight plans were similarly coded, white if they were at ownship's current altitude, and blue or green if these legs

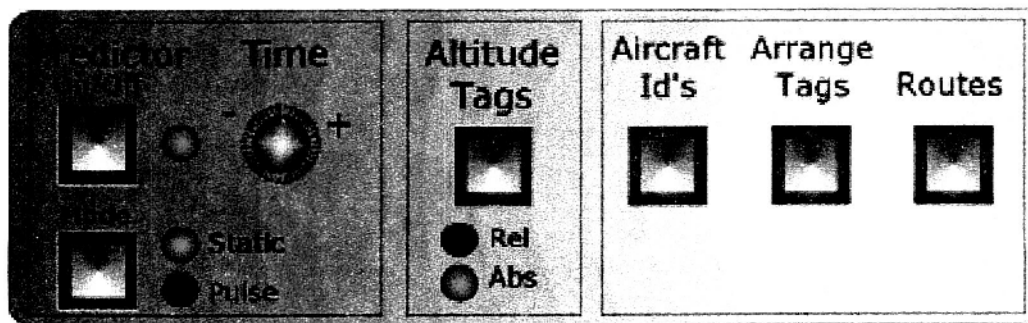
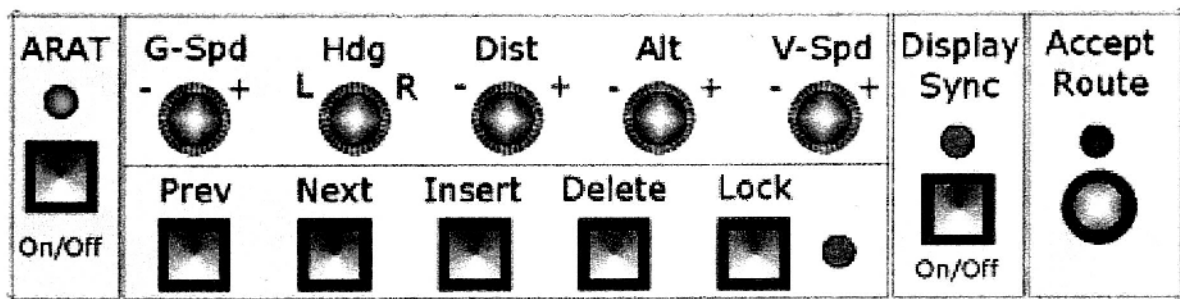


Figure 2. The top figure shows the control panel used to control the Advanced Route Analysis Tool. The top portion of the panel allowed the positioning of horizontal and vertical waypoint. The bottom portion of the panel would be used to scroll through or add/delete waypoints. The Lock and Accept buttons datalinked the altered route to ATC and then to the Flight Management System. The bottom figure shows the control panel used to control the amount of information presented on the display. This panel could be used to control the predictors, the mode for altitude information, the presentation of all Ids and the removal of individually selected routes. On both panels, green lights would be illuminated if that setting was currently selected or if that function was currently on. Additionally, the Accept Route light flashed when an ARAT modified flight plan was ready to be loaded in the Flight Management System.



Figure 3. Color was used to display basic categorical information for relative altitude.* Aircraft above ownship current flight level (38,500) are color coded blue. Aircraft below ownships current flight level (37,500) are color coded green. The aircraft with tags are climbing and descending, as indicated by the arrow on the tail tag of the aircraft. Also, as shown, the arrow is placed next to the current flight level in the extended tag.

Note. Due to the limitation of black and white print for this presentation, colors will be described as best as possible. For a color version of the display, please contact Walter W. Johnson at NASA Ames Research Center, MS 262-2, Moffett Field, CA 94035. Additionally, values for distance and time to next waypoint (upper right hand corner) were fed to the system via the flight simulator.

were above or below ownship's current altitude. The only exception was for the ownship route, where the legs at ownship's current altitude were maintained as magenta, the color convention for ownship route.

Periods of altitude change were also indicated graphically (see Figure 3). For the individual aircraft an up or down arrow, similar to the one used in TCAS displays, was displayed whenever an aircraft began climbing or descending at more than 100 feet per minute (this arrow was associated with an altitude tail tag described below). In addition, periods of planned altitude changes were depicted in the flight plan with a dashed line (see Figure 4). This dashed line was flanked by symbols indicating the beginning and end of the altitude change segment. This depiction provided additional important information since the flight plan color-coding could not indicate altitude changes that occurred entirely above or entirely below ownship's current altitude (where they would be entirely blue or green, respectively).

While the color-coding was designed to provide good at-a-glance awareness of important relative altitude information, text was used to provide more detailed altitude information should the pilot want or need it. Altitude tail tags were provided for all aircraft within the Free Flight SuperSector (described in a later section), and these tags were aligned with the bottom of each aircraft symbol. Pilots could select either a relative or an absolute (barometric) altitude format for these tags. Similar altitude information was also included within selectable aircraft ID blocks (also described later). Finally, when displayed, all flight plans also included text showing the destination altitude at the end of any altitude change segment.

Predictive Information

Another challenge in developing at-a-glance symbology was to provide predictive information. Predictive information, which extrapolates a traffic situation into the future,

is a critical element in any display. Predictive information could be displayed using either static, or dynamically pulsing, predictors.

When either type of predictor was turned on (using either the toolbar or a control panel button), lines would extend along the expected flight paths of all primary aircraft on the display (i.e. predictors were not selectable for individual aircraft). The length of these lines corresponded to a length of time in the future, and was adjusted by the crew with a control panel dial or a touch pad operation. The dial allowed the predictor to be adjusted in 20 second increments, while using the touch pad to adjust the predictor time button on the display toolbar allowed the predictor to be adjusted in two minute increments (the current predictor time length was always shown on the toolbar). When the flight plan for an aircraft was being displayed, the predictor line was shown by increasing the brightness of the corresponding segments of the flight plans. When the flight plan was not on, the predictor took on the corresponding flight plan color, and also showed the dashed lines associated with altitude change segments. However, the text and symbols associated with the altitude change segments and waypoints, were not shown. Thus displaying predictors was another method of displaying flight plans, but without the clutter (or information) associated with the textual information.

When static predictors were used, the crew could examine the expected relative positions of the aircraft at a future time by manipulating the prediction interval. As an alternative, dynamically pulsing predictors were also available. These predictors were designed to provide a more hands-off type of tool. The pulsing predictors were similar to the static predictors but also included small bright bullets that were repetitively emitted from the individual aircraft (see Figure 5). These bullets traveled at a speed proportional to the aircraft speed. In this way, the crew could see a projected close horizontal encounter when a bullet from ownship approached the bullet



Figure 4. Three dimensional flight plan information was available on all aircraft within the supersector. A flight plan is presented for UAL201. This aircraft is currently co-altitude with ownship and is color coded white. The flight plan shows that it will start descending after crossing OAL and it will level off below ownships current altitude. Therefore, the portion of the flight plan behind OAL is white (co-altitude) and the portion after OAL246042 is green (below ownship altitude). The dashed portion between the top of descent (T/D) and OAL246042 represent the altitude change segment and the destination flight level of the aircraft is presented at the latter waypoint (FL 350).

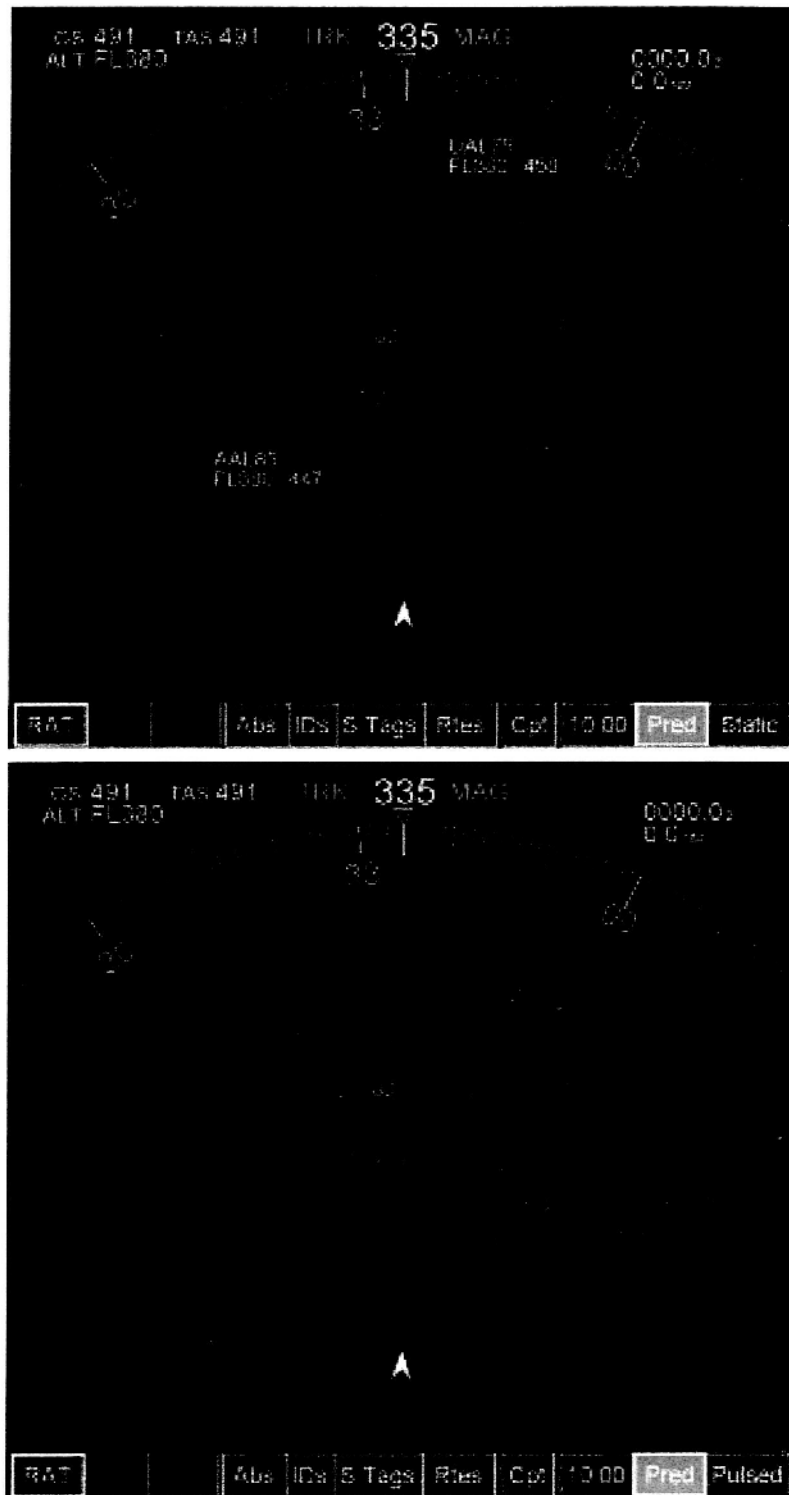


Figure 5. Static predictors (top) are presented at a range of ten minutes. The dashed portions of AAL85 and UAL75 represent the areas where these aircraft will be changing altitude. Pulsed predictors (bottom) were similar to static, but they also included small bright bullets that were repetitively emitted from the individual aircraft.

from any other aircraft. They would then only have to evaluate the projected altitude separation at this location to determine if this specified a potential conflict. In this manner the crew could evaluate the future path without having to continually manipulate the prediction time - the display showed the entire four-dimensional path up to the length of the predictor. The crews were allowed the optional use of either predictor because it was not known if the value of the pulsing predictors would be offset by annoyance with a constantly active display.

Clutter Management Features

A third challenge was to manage display clutter. Clutter has many causes. On the one hand aircraft density is obviously the main culprit in generating clutter, but on the other hand it is the textual information associated with the aircraft that truly drives the clutter problem. To the extent that the above graphical (non-textual) techniques for providing altitude and predictive information were successful, these must be considered methods for clutter management. However, many additional design features were incorporated to provide clutter management.

Since a goal of this study was to test the display's capabilities (or limitations) in rendering high traffic densities, no tools or techniques were included to decrease clutter by filtering out less important aircraft. Instead, it was managed in the following ways. First, only aircraft within the hypothetical high altitude SuperSector, or within 3000 feet of ownship, were presented at full brightness. All others were dimmed to reduce their salience and their impact on clutter.

Second, excluding alerts, the only standard (non-optional) traffic text on the display was the altitude text in the tail tags, and this was only shown for the brightened aircraft. Furthermore, aligning these tags with the tail of the aircraft further reduced apparent clutter, and made the aircraft plus altitude tag appear as a more integrated unit.

Third, aircraft ID and speed were placed within aircraft ID blocks that were easily viewed, but which were not standard display elements. The aircraft ID blocks could be brought onto the display in three ways. First, an aircraft ID block could be displayed by using the touch pad to simply point at an aircraft symbol (putting the cursor over the symbol). Once the cursor was removed from the symbol the ID block would go away. Second, if a crew member used the touch pad to tap on an aircraft symbol then the ID block would not disappear until the user tapped again on the aircraft symbol (or tapped on the ID block itself). This allowed an aircraft to be highlighted and remain highlighted for as long as the pilot desired. Third, the ID blocks for all aircraft could be turned on or off by pressing/clicking on the ID button on either the control panel or the toolbar (see Figure 6). Turning all of the ID blocks on allowed the pilot to search for a particular aircraft by its ID (a task not required in the present study), while the ability to turn all of the ID blocks off, including those individually selected, was a fast declutter option.

Other methods of managing clutter due to the effect of ID blocks were also made available. One method was to use the touch pad to select and reposition ID blocks by dragging them to a new location. A second method was the automated repositioning of the ID blocks to keep them from overlapping aircraft symbols or other ID blocks (when this was not possible the algorithm used a prioritization scheme to determine what should be overlapped). The automated positioning was applied to each ID block whenever the ID was turned on. However, in order to avoid the distracting effect of continuously jumping ID blocks, automated repositioning was not applied again to a tag unless the pilot pressed/clicked on a button (marked STags) which caused all visible ID blocks to once again be automatically positioned.



Figure 6. All Ids on was an option available to the crews. By clicking on the Ids button on the on screen toolbar or by pressing the ID button on the control panel, all Ids can be turned on/off, with the exception of any aircraft that is currently at any SA level (e.g., UAL201 is currently in an SA1 and is presented in yellow).

A final method fast declutter option was a button (RTES) that allowed the pilot to remove all of the flight plans from the display.

Conflict Alerting

Conflict alerting was incorporated within two parts of the system. The Primary Alerting System, which co-exists with the TCAS system,

was designed to advise the flight crew of a potential Loss Of Separation (LOS) along their currently active route. The operation and depictions of this system were similar, but not identical, to the operation and depictions used to warn of potential conflicts by the Advanced Route Analysis Tool (ARAT). The ARAT Alerting System was a version of the Primary

Alerting System that had been integrated with the ARAT. The ARAT (which will be more fully described in a later section) was used to create revised flight plans that resolved conflicts. The ARAT Alerting System was used to advise the flight crew if a potential route that they were designing had any conflicts associated with it. Both alerting systems were based upon an alerting algorithm designed by Yang and Kuchar (1998). This algorithm takes advantage of the intent information provided by flight plans, when these plans are available, and otherwise operates using only basic aircraft state information (position and velocity).

Primary Alerting System

The alerting logic within the Primary Alerting System was configured to provide three levels of alert (see Figure 7). When any primary alert was in progress the ownship symbol was turned into a filled yellow symbol, and a yellow ID block was brought up for the alerted aircraft (this ID block could not be turned off while the aircraft was in an alerted state). The three levels of alert were distinguished from each other in the following manner:

- A Level 1 Conflict Alert specified a moderately probable conflict that was well in the future, or a lower probability conflict that was closer in time. When it occurred the conflicting aircraft was depicted as an unfilled yellow chevron, a yellow ID Block for the conflicting aircraft came on, and a single chime sounded (no chime sounded when a Level 2 decreased to a Level 1 alert).
- A Level 2 Conflict Alert specified a moderately probable conflict that would take place moderately soon, or a highly probable conflict that would take place well in the future. When it occurred the conflicting aircraft was depicted as a solid (filled) yellow chevron, a yellow ID Block for the conflicting aircraft came on, and a double chime sounded (no chimes

sounded when a Level 3 decreased to a Level 2 alert).

- A Level 3 Conflict Alert specified a highly probable conflict that would occur in the moderate to near future. When it occurred yellow lines extended from the alerted aircraft, and from ownship, to their positions at the predicted point of initial LOS, and both lines terminated with a circle that was 5 NM in diameter. In addition, at this time a yellow ID Block for the conflicting aircraft came on, and an audible warning "Alert, Alert" sounded.

In addition to the conflict alerting, the Primary Alerting System was designed to work with TCAS. However, only the two severest TCAS advisories, the Traffic Alert and Resolution Advisory, were retained. These advisories were designed as a collision avoidance system and signaled a situation potentially more severe than a projected LOS.

ARAT Alerting System

The ARAT Alerting System operated in a manner similar to the Primary Alerting System. The purpose of the ARAT Alerting System was to show when a proposed new path would clear an existing primary alert, and to show when a proposed new path would create a new conflict. In general, this was done by changing the indicated alert levels of the displayed aircraft and ownship. For example, if there was a Level 2 primary alert, and the ARAT path would decrease it to a Level 1, then the Level 2 conflict symbology (a filled yellow chevron) would be replaced with the Level 1 conflict symbology (an unfilled yellow chevron). Or if the ARAT path would lead to a Level 1 conflict with a previously unalerted aircraft, the unfilled white chevron would be replaced with an unfilled yellow chevron.

However, a problem with such a system is that it can be easy to confuse it with the output of the Primary Alerting System. It was important

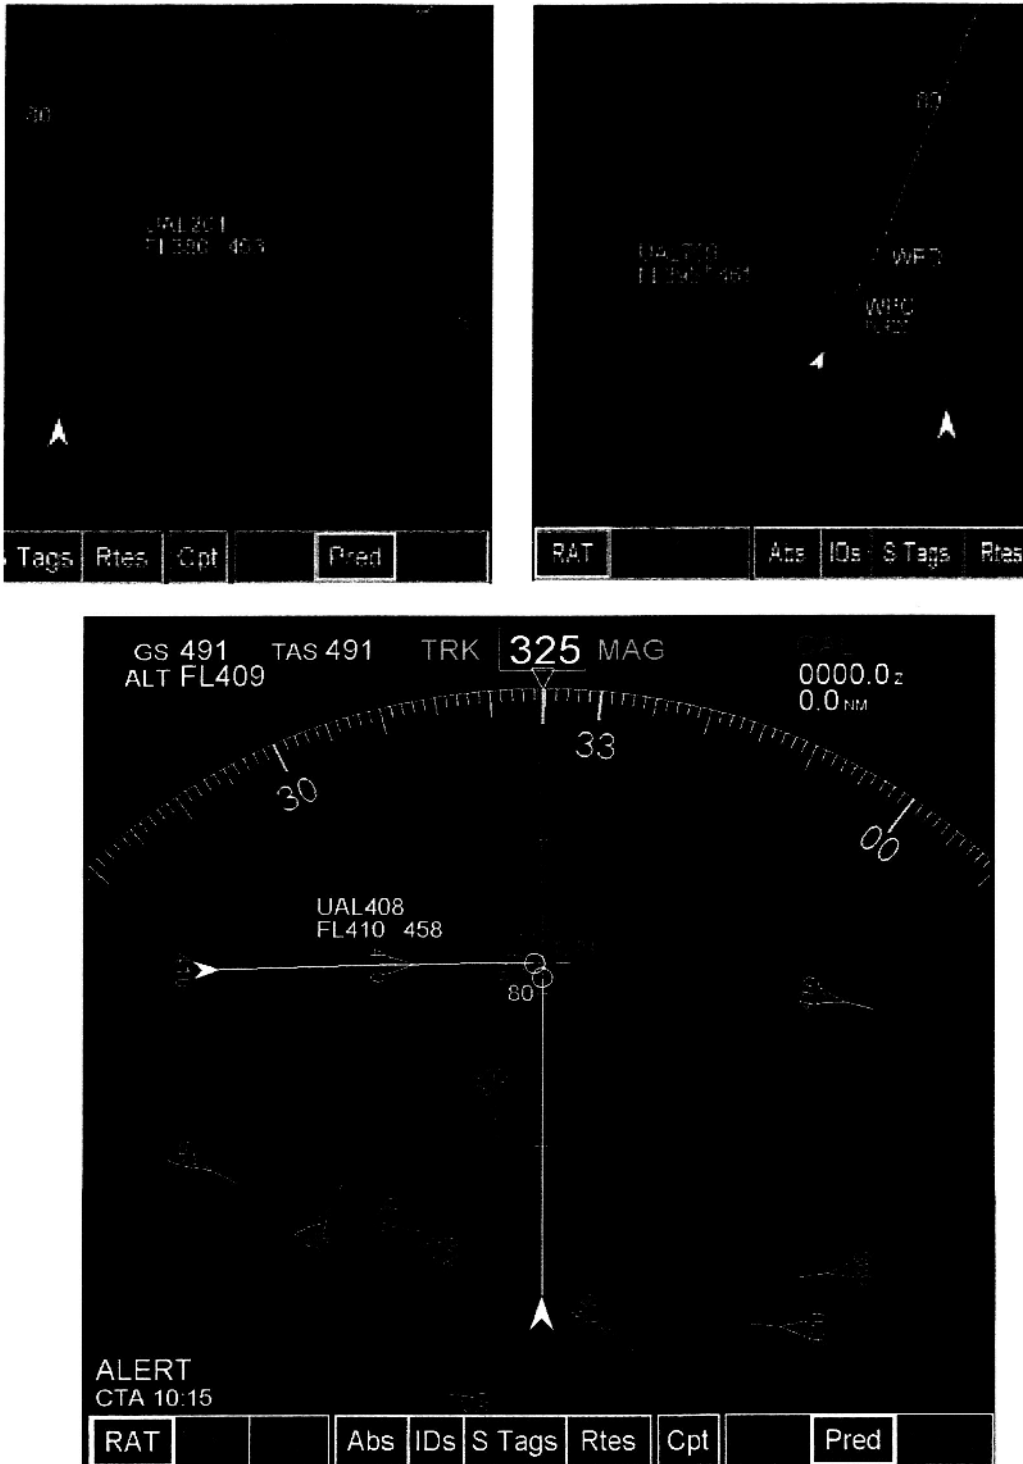


Figure 7. The upper left picture shows an SA1 (open yellow chevron), upper right an SA2 (closed yellow chevron) on which the flight plan has also been selected, and the lower picture shows an SA3 with lines and circles indicating the point of Loss of Separation. For all SA levels, the flight Id automatically comes on and can not be deselected. In addition to the changes in symbology with the increase in SA levels, the time until Loss of Separation is also reported in the bottom left with a SA3.

to make it clear when an alert came from the ARAT Alerting System and when an alert came from the Primary Alerting System (which could generate alerts while the ARAT was being used to plan a new flight path). In order to do this the yellow ID blocks triggered by a primary alert remained on even when the ARAT Alerting System showed that a proposed path would clear a conflict, and the ID blocks for ARAT alerted aircraft never came on automatically (see Figure 8). Furthermore, if ID blocks of ARAT alerted aircraft were turned on by the pilot, they would remain their original color (blue, white, or green, but not yellow). ARAT alerts also never contained the auditory signals associated with the Primary Alerting system.

In addition to the above, the outputs of the two systems differed in the two following ways. First, the ARAT Alerting System used the same Level 2 symbology for Level 2 and Level 3 ARAT alerts. That is, no lines extending to the point of LOS were shown for Level 3 ARAT alerts. These were not shown because the presence of multiple lines signifying points of closest approach was deemed to be confusing. Second, if an ARAT flight path would lead to a decrease in a Level 3 primary alert, the lines showing the point of closest approach would not disappear, but would turn white. This was done in order to avoid removing important and useful information from the display while this Level 3 primary alert was still in progress.

Advanced Route Analysis Tool (ARAT)

The ARAT was designed to allow the flight crew to 1) design changes to their three-dimensional flight plan, 2) submit this revision to the air traffic controller for approval, and 3) implement the revised plan if approved. For the purposes of the present study the ARAT, and only the ARAT, was used to control the aircraft (e.g., the pilots were instructed not to use the MCP). The ARAT was designed as a strategic tool, and therefore was most effective when used to make changes designed to take effect several minutes in the future. Procedures and features for dealing with more

tactical (near-term) problems were not included in this version of the ARAT.

General Procedure

The ARAT could be controlled using either the touch pad or the control panels. Both systems, replicating each other, could be used to perform all actions and were included to accommodate user preferences. The crews could also switch between input devices if one device was deemed easier for any particular task.

The ARAT tool could be turned on at any time and would display an alternative second flight path, called the ARAT Path. When the ARAT Path was initially displayed it would fully clone and overlap all legs and waypoints on ownship's current route. The current flight path would remain on the display, but would be hidden by the overlapping (orange) ARAT Path. If the position of the horizontal legs of the ARAT Path were modified during the use of the ARAT, then the underlying current path would be revealed. In no case was the current path affected by the use of the ARAT until the crew implemented the ARAT Path. At any time the crew could turn the ARAT off and on to clear all ARAT Path modifications, and start over. In addition, the crew could use the "UNDO" button to remove the last modification. However, the UNDO function had no memory, and thus only the last modifications could be removed using it. Once the crew was satisfied with the form of a modified ARAT Path, they would push a button marked "ENTER", and then send it to the controller by pressing a button marked "LOCK". This would cause the revised ARAT Path to appear in yellow on the ATC display, and to change to gray on the cockpit display. While in this "LOCKED" state no modifications to the ARAT Path were possible (see Figure 9).

Once received, the controller would review the proposed plan and verbally accept or reject the modification. If the controller rejected the modification, the crew would then press a



Figure 8. To prevent confusion between Primary alerts and ARAT alerts, primary alerts always had yellow tags up and the point of loss of separation was presented until the alert was solved. ARAT alerts were only identified by the alert symbols (e.g., filled chevron to left of ownship). ownship's symbol would only return to white if the proposed path was clear of the primary alert and all other aircraft.

button marked "REJECT" (on the display toolbar) or "UNLOCK" (on the control panel). This would cause the revised plan to be removed from the ATC display, while changing back to a modifiable orange path on the cockpit displays. If the controller accepted

the modification the crew would then press a button marked "ACCEPT" and the flight plan would be sent to the Flight Management System for implementation. Once implemented the new flight plan would be datalinked to the ground to update the ATC

flight data, and would also be broadcast to all other aircraft via ADS-B.

Finally, there were critical timing issues that had to be addressed in the procedures underlying the use of the ARAT. The first issue concerned how to handle the case where ownship arrived at an ARAT modified waypoint before the ARAT plan had been executed in the FMS. Once past this point it is not possible for ownship to fly the proposed path (not recognizing that the point had already been passed and that the aircraft should go to the next point, the aircraft would circle to recapture the last waypoint). There are two possible methods for resolving this dilemma. First, the proposed plan could be automatically, and progressively, modified to keep the first modified waypoint in front of ownship (e.g., push the first modification point on the path). A problem with this approach is that the path, which would be continuously modified, would need to be reevaluated by the alerting system. Since the system requires a minimum amount of time to update the flight plan in the FMS, this leaves open the possibility that a clear path upon "ACCEPT" would change to a conflicting path upon "EXECUTE." Furthermore, the degree of difficulty in devising an automated procedure for modifying the ARAT flight plan was determined to be beyond the scope of the present effort. Based upon these considerations, and the assumption that the automation should not try to anticipate what the crew would desire or find acceptable, the present study utilized a second method.

The second method was to cancel the proposed plan once the first modification point came within 40 seconds of ownship, the standard amount of time required to complete the procedure after an ARAT Path modification had been "LOCKED" and sent for ATC approval. During this interval the controller must evaluate and approve the plan, the crew must press the "ACCEPT" button, and then the crew must finish hand loading the modified flight plan into the FMS. This

location of this critical point in time was visually depicted on the CSD as a small yellow horizontal line segment called the Cancel Point. The Cancel Point was typically located about four to six miles in front of ownship (depending on ownship speed), and if the first modified/inserted waypoint point reached the Cancel point, the ARAT was reset and the crew was required to start over. The crew, of course, could place their initial modification further down the current path, and thus give themselves more time to create, submit, and implement the new flight plan before the initial modification passed the Cancel Point.

There are negative side effects to deleting an ARAT flight plan, e.g. lost work, pilot irritation, and increased workload. In order to avoid this it was clear that modifications (changes in planned heading or altitude) should not be initially located too near the Ownship, since this would result in overrunning the first modification point before the new flight plan can be fully evaluated by the crew, and then by ATC, and then implemented. The system required some planning aid that would ensure a minimum time (and distance) between ownship and the initial modified ARAT waypoint, and would therefore minimize the number of canceled flight plans

This was accomplished by the inclusion of a point, called the "Null Point," 2 minutes in front of ownship. The Null point was depicted by a solid orange circle, and changes to heading or altitude could not be introduced prior to this point. If the crew created an ARAT path modification using the Null Point, the point changed into a newly inserted waypoint that would be captured by ownship in 2 minutes. Thus, there would be 2 minutes for 1) the crew to finish creating the ARAT Path, 2) the crew to lock and send it to ATC, 3) the controller to issue a verbal approval, 4) the crew to press the ACCEPT button, and 5) the crew to load and execute the new plan in the FMS. If the crew was unable to implement all of the steps in the procedure, they could

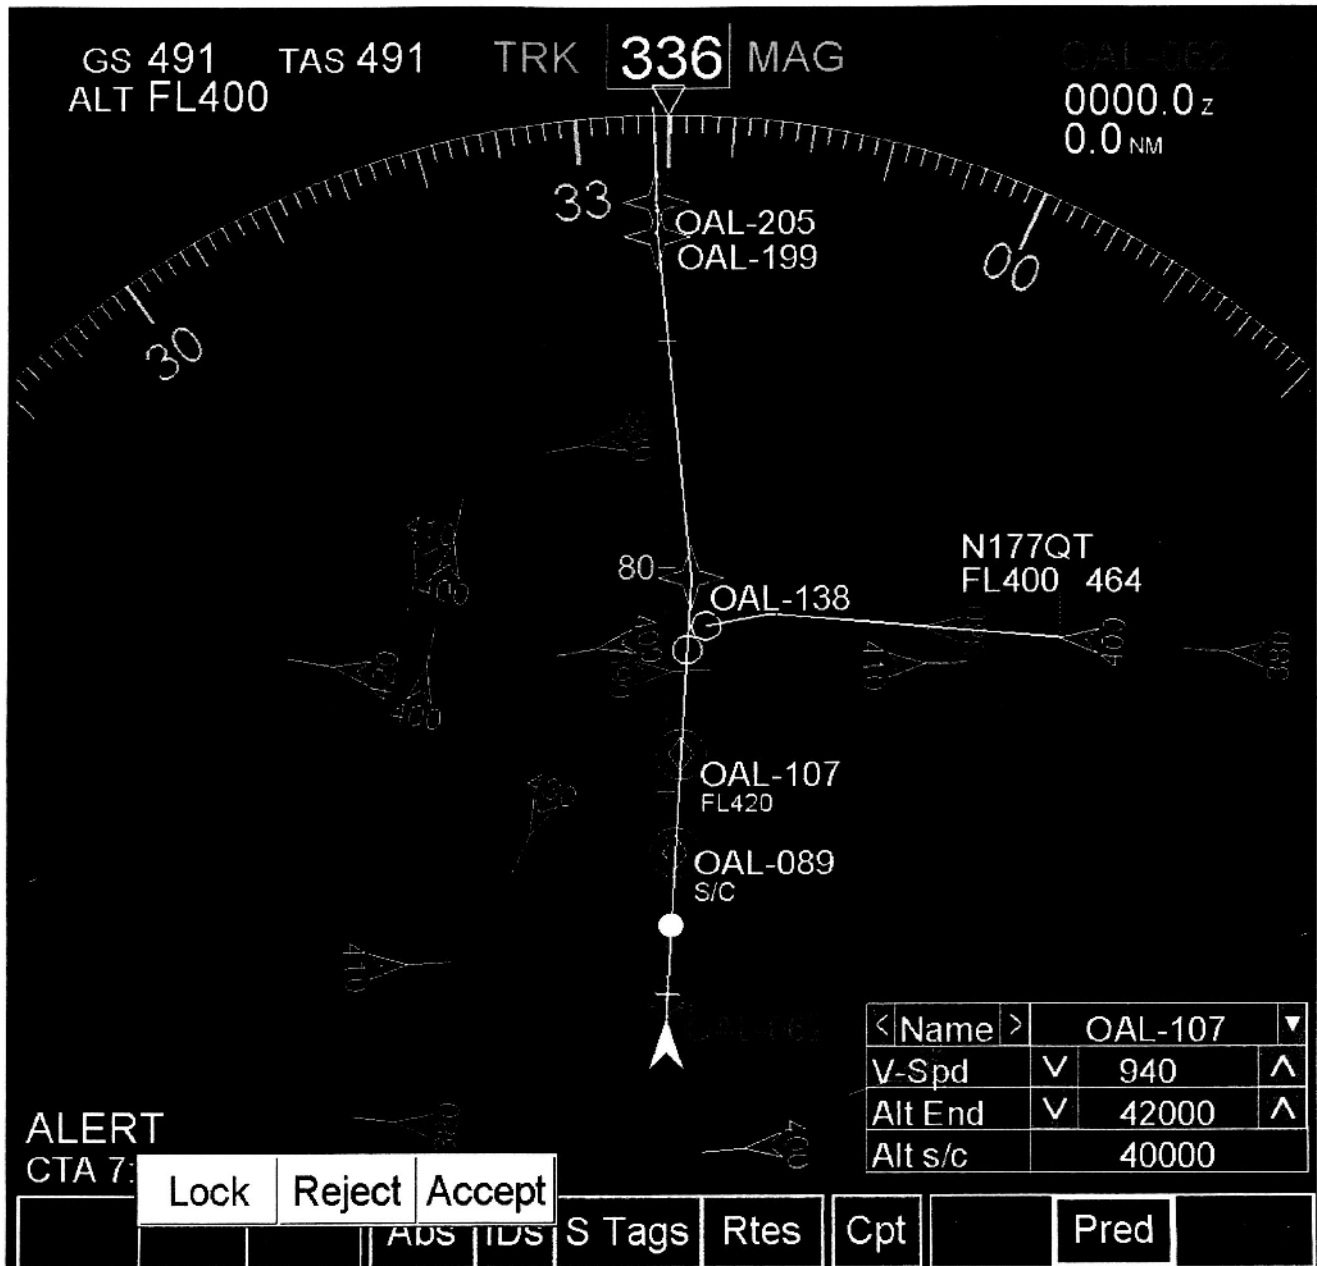


Figure 9. Once a crew finds a clear path with the ARAT (the color of all aircraft, including ownship, return to their original altitude color), the crew locks the new flight plan. Once locked, the crew can no longer edit the path until they unlock it. Lock transmits the new flight plan to ATC, and if approved, the crew would then “accept” the flight and “load and execute” it in the Flight Management System.

either let the route expire (cancel) or they could turn the ARAT off and back on and begin the planning process anew. Once again, if the crew introduced their initial modified waypoints further out on the route, they could ease these time constraints.

ARAT Flight Path Editing

The ARAT was designed to allow the crew to insert new waypoints, delete old waypoints, and to move waypoint locations. Using these two operations the crew could manipulate the horizontal flight path into any form they

wished. The touch pad and the panel controls operated in slightly different manners, but embodied the same capabilities. Essentially, both types of input device allowed insertion and positioning of waypoints. The main difference between the two was that the touch pad required fewer steps and worked in a more intuitive manner, but required greater fine motor control to operate effectively. For example, with the touch pad, a crew member could "pick up" a point and "drag" it to a desired point, never taking their eyes off the display. To do the same operation using the control panel, the crew member would have to manipulate both the heading and distance dials, after locating these dials on the control panel. Additionally, the touch pad allowed a simple method of selecting the point that crew members wished to edit (i.e. simply tapping the location of the waypoint), while selection using the control panel required them to scroll through the waypoints using the "Prev" (previous) and "Next" buttons. The benefit of the control panel was that the crew members did not need the fine motor controls sometimes necessary to use the touch pad. For example, using the touch pad, points could easily be inserted unintentionally by inadvertently tapping on the ARAT path (this was the method for inserting a point using the touch pad); while at other times it might be hard to keep the cursor over a point, which was a necessity when trying to do many of the manipulations (e.g., waypoint selecting, dragging, deleting).

The ARAT was also designed to allow the crew to insert/delete altitude control segments and to position these segments. An altitude control segment was created by adding a climb or descent to a previously existing waypoint or to an inserted point. Crew members did this by using the "Alt" dial on the control panel or clicking on the up and down arrows on the Waypoint Table (see Waypoint Table description below). When this was done that waypoint became the end of climb or descent, and a dashed line plus a start of climb or descent waypoint was inserted. The length of

the altitude control segment was determined by the size of the altitude change, together with the ownship's economical climb or descent. The location of any altitude control segment could be changed by sliding the segment (i.e., selecting and dragging with the touch pad or using the distance dial on the control panel) along the ARAT Path to a new location. The only constraint on this operation was that an altitude control segment could not slide past or over another altitude control segment. Rates of climb or descent could be set by the crew, but there was no insurance that these rates could be attained by ownship.

All of the modifications that could be performed using the ARAT were depicted within the visual picture of the display (i.e., waypoints, waypoint labels, graphical depiction of route) and in the Waypoint Table that appeared in the bottom right hand corner of the display when the ARAT was used (see Figure 9). The Waypoint Table reported the name of the currently selected waypoint, the headings to and from the waypoint, and the altitude associated with the point. If the point was part of an altitude change segment, you would also see the destination altitude and the vertical rate at which the aircraft would climb or descend to achieve that altitude. The waypoint table also had arrow controls which allowed the user to select the previous or next waypoint, allowed the user to modify altitude by thousand of feet, and allowed the user to modify the rate of climb or descent.

The next section describes the simulation that was used to evaluate the utility of the CSD.

CSD Simulation Evaluation

A simulation examining these display concepts during enroute flight was conducted using the Boeing 747-400 Level D flight simulator at the NASA Ames Research Center.

The simulation examined the utility of the CSD in the context of an Air Traffic

Management system (ATM) that supports limited enroute Free Flight. Specifically, the simulation examined how well the CSD supported crews who were given the task of maintaining separation and resolving conflicts. Furthermore, the context of this evaluation was varied in two ways: 1) whether or not intent information, in the form of 3-D flight plans, were globally shared between flight decks and the controller; and 2) the degree of flight deck responsibility for maintaining separation.

- **Intent Information:** In the Unshared Intent condition the flight deck only could view the location, altitude, speed, and direction of other aircraft, while the controller had flight plans in addition to all of this information. In the Shared Intent condition, the flight deck CSD, including the alerting system, also had the flight plans of the other aircraft.
- **Flight Deck Responsibility:** In the High Responsibility condition the flight deck was expected to resolve any potential LOS, while the controller was instructed to intervene only if the flight deck failed to accomplish this, and there was an imminent LOS (the definition of "imminent" was left to the discretion of the controller). In the Low Responsibility condition the flight deck was expected to resolve any potential LOS, but the controller was instructed to resolve conflicts at any time they deemed needed or appropriate. In this manner the controller could allow the flight deck to exercise whatever degree of separation responsibility seemed acceptable, but otherwise controlled the airspace based on current rules and practices.

The general procedural concept for engaging in Free Flight route modifications was based on current practices of requesting and granting of permission by the controller, and thus was a limited version of Free Flight. This procedure required a flight crew to use the cockpit avionics to design a proposed flight

plan modification, and then to datalink it to the controller for approval. Once this approval was granted (via voice link), then and only then could the flight crew implement the change. This procedure was consistent with current practices that are now conducted via voice transmissions.

In most other ways the simulation reflected optimistic assumptions about the quality of surveillance information and the alerting system. Specifically, the simulation incorporated the following features:

1. Error-free surveillance information (traffic position, velocities, headings, vertical speeds, and altitude.)
2. No weather, no winds.
3. No secondary tasks (very low workload.)
4. All aircraft in the SuperSector (described below) were assumed to have similar equipage, although the crews were told that it was up to them to resolve all conflicts.
5. While the controller had many operations to perform, ownship was the only aircraft requesting maneuvering, and received responses to requests quickly.
6. Ownship weight was very light for all scenarios, making all altitude changes possible as well as "economical."

Free Flight SuperSector Description

In order to evaluate the flight crews ability to utilize the CSD to support en route Free Flight, a new air traffic management entity, called the Free Flight SuperSector (FFS), was postulated. The initial design of the FFS was based on current Air Route Traffic Control Centers' (ARTCC) super high altitude sectors. However, the FFS concept had several additional defining characteristics. First, FFS's were assumed to be privileged airspace, requiring aircraft flying within a FFS to be specially equipped. This requirement for special

equipment is consistent with current special equipment requirements for operations in many terminal airspaces. For the present simulation, all aircraft flying within FFS were required to be minimally equipped with onboard surveillance systems, conflict alerting systems, and CDTIs. In addition, they were required to be able to datalink flight plan modifications to all aircraft within sensor range and to the controller for his or her approval.

For the present implementation the FFS was defined as airspace between FL370 and FL450, and covering the areas of two normal high altitude sectors (sectors 34 and 44 within the Oakland Center Airspace). Within the FFS the minimum vertical separation was reduced to 1000 feet, while the lateral separation minimum was maintained at 5 NM. Within the FFS cardinality rules continued to apply, with eastbound traffic located at odd number flight levels (i.e. FL370, FL390, FL410, FL430, and FL450), while westbound traffic was located at even numbered flight levels (i.e. FL380, FL400, FL420, and FL440).

Surveillance System Description

The simulated onboard surveillance system was loosely based on requirements for the Automatic Dependent Surveillance – Broadcast (ADS-B) system, and assumed all equipped aircraft would be broadcasting information derived from onboard sensors. The simulated system broadcasts/reception range was assumed to be 120 NM. All aircraft within this range were assumed to be exchanging information about their current position, altitude, ground speed, vertical speed, ground track, and aircraft ID at a rate of once per second. In addition, they exchanged information specifying their current approved flight plan. The flight plan specified the three-dimensional paths that the aircraft would follow from departure to destination. However, these flight plans did not include ground speed, and thus there was the possibility of conflicts (loss of separation), or collisions, at points where the flight paths of individual aircraft intersect or overlap.

Method

The following is a description of the experimental designs and procedures used in the simulation. This section is included to indicate the range of conditions in which the pilots flew. This is needed to provide an adequate context for the opinion ratings reported in the results section.

Design

The study was a 2 (Flight Deck Responsibility) by 2 (Intent) by 2 (Conflict Type) by 2 (Initial Altitude Separation) mixed design. The previously described Flight Deck Responsibility factor (High and Low responsibility) was the between subject factor, with four crews from a major airline assigned to each of the two levels.

The remaining three factors were within subject factors, and were used to create the eight pre-programmed scenarios presented to all eight flight crews. The previously described Intent factor (Shared and Unshared), presented each crew with four scenarios in which flight plans were part of the ADS-B broadcast (Shared condition), and four scenarios in which flight plans were not broadcast (Unshared condition). When flight plans were not broadcast the CSD system simply used aircraft state information to project the future paths of traffic.

Conflict Type (True and Apparent) and Initial Altitude Separation (Separation, No Separation) were included to reduce the scenario-to-scenario predictability of conflicts. Conflict Type refers to whether or not ownship really was on a conflicting course with another aircraft. In four cases it was on a conflicting course, while in the other four cases it was not on a conflicting course. However, for the four non-conflict cases, the alerting logic predicted a conflict when Intent was Unshared. That is, when the flight plans for the traffic aircraft were lacking, simple extrapolation of the present track of one of the traffic aircraft initially indicated a conflict. However, these

scenarios never showed a conflict when Intent was Shared, since there were planned changes in the future path of the apparent intruder. Initial Altitude Separation (Same Altitude, Different Altitude) refers to whether or not the intruder (Apparent or True) was initially separated in altitude from ownship.

Scenario Construction

For all of the scenarios, ownship was assigned to one of three altitudes (i.e., 380, 400, 420) within the SuperSector. Ownship was always assigned a ground speed of 491 knots. Ownship could either have a direct path to its destination or have one to two pre-planned lateral changes on their flight plan during the flown segment of the 20 minute scenario. No altitude changes or speed changes were included in ownship's flight plan.

The traffic for each of the eight scenarios was based upon traffic taken from radar data collected at the Oakland center. Most of this traffic was taken from altitudes 310 to 350, and placed in the SuperSector. Slight lateral modifications were also made to the flight plans for these aircraft. Finally, additional aircraft were inserted into the scenario so that approximately 10 to 16 aircraft were within the SuperSector within sensor range of ownship at all times. Within the SuperSector, a modified version of cardinal altitude rules was applied, with a requirement of 1000 ft vertical separation. Aircraft traveling east maintaining odd flight levels (e.g., 37000, 39000,..45000) and west bound aircraft maintaining even flight levels (e.g., 38000,40000, ..44000). Optimal ground speeds ranging between 400 and 500 knots were assigned to the SuperSector aircraft. Three to five context aircraft climbed into the SuperSector during the scenario, while a majority of the aircraft had a planned descent out of the SuperSector on arrival to the San Francisco area.

The time until the conflict aircraft was visible to the crew was different in each scenario, but the time to loss of separation varied between 7.5 and 12.5 minutes. Crews always had more

than 6 minute from when the SA3 was triggered to resolve the conflict. These conflicts were generated by the deliberate placement of Intruders within the traffic flow.

For each of the scenarios, the crew was told that they were responsible for maintaining separation, regardless of the encounter geometry. Additionally, since the system was designed for strategic conflict resolution, participants were told that the scenario would end if the aircraft came within the range of a TCAS advisory or resolution.

Task

Crews were required to fly eight enroute scenario segments. They were informed that during these scenarios there would be cardinal altitudes but that when maneuvering to avoid a conflict, the approval of the maneuver was ultimately the controller's responsibility. The vertical separation minimum would be 1000 ft (reduced from the standard 2000 ft high altitude minimum), and the lateral separation minimum was the standard 5 NM. To maintain separation, the crews were allowed to design and propose any maneuver to ATC for approval. Two experimental confederates who were former air traffic controllers in the Oakland Center played the role of the ATC in the high flight deck responsibility condition. The confederates were told to give the crews as much leeway as possible, only stepping in and proposing a maneuver when they felt separation minimums would be violated. In the low flight deck responsibility condition, current full performance level controllers (FPLs) from the Oakland Center controlled the sector. These controllers received a description of the surveillance equipment and tools on board the aircraft, and they were asked to control the sector as they felt necessary keeping the flight deck capabilities in mind. In both conditions, the crews were told to address the center with a call sign and a description of the deviation from the current course that they were requesting. The air traffic controllers could then view the modification of the course on their displays

and accept or reject the modification. The crews were not allowed to enter the new flight plan into the FMS until it had been approved by ATC. Two additional confederates played the roles of the Air Traffic controllers in adjacent sectors and the crews of the other aircraft.

Scenarios ended after fifteen minutes of flight in the scenario unless the aircraft was continuing the resolution with the primary threat or had created a secondary threat. If the initial avoidance maneuver did not result in any additional conflicting situations, this amount of time allowed the crew time for the avoidance of the designated Intruder and then some additional time to return to course or make additional modifications to the flight plan. In the case that the crew were unable to solve the conflict, the scenario automatically ended at the point that a TCAS advisory was issued.

Procedures

An initial briefing was held for each crew in which the experimenters described the flight environment (SuperSector) and the task. Additionally, the display tools were described. Following this, each crew received 2-3 hours of initial training in which they became familiar with the display and the display tools. The crews then flew the eight experimental scenarios, which took 4 to 5 hours. The crews then completed a debriefing questionnaire followed by a verbal question and answer session.

Equipment

The simulation was conducted using the Boeing 747-400 Level D flight simulator located at the NASA Ames Research Center. The basic traffic display was built upon the functionality present in the Map Mode of the Visual Navigation Display in the Boeing 747-400. Due to the complexity of the CSD, the navigation display was produced and presented by two stand alone Dual Pentium 2-400 computers running an NT Operating System. Graphics were produced by Leadtek

3100 graphics adapters. These two computers also served as the communication link between the 747 simulator and the Pseudo Aircraft Simulator (PAS), which generated the traffic in the scenarios and served as the interface for the controllers. The CSD computers transferred all information, communicating with the simulator and PAS via a TCP/IP datalink. The CVSRF system had no interconnection with the PAS system and vice versa. When the MAP mode was selected on the EFIS control panel in the simulator, all information displayed on the NAV displays were provided by these NT computers.

As far as the flight simulator was concerned, all user controls were in effect and were operating normally. Once the alternate route had been designed using the ARAT, new routes were implemented by the following steps: (1) the crew had to press the Accept Route pushbutton to engage the new route; (2) the new flight plan was automatically datalinked to the FMC and sent to the PAS system; 3) when the datalink was completed to the FMC, a message was displayed on the FMC stating the uplink information was ready; 4) the flight crew had to "Load" the uplink on the "Rte 1" page of the FMC; and 5) the flight crew had to then "Execute" the flight plan on the FMC, keeping the aircraft in the LNAV, and when applicable, the VNAV modes. Since the purpose of the experiment was to determine how well the crews could solve problems strategically using the tools provided, crews were asked not to interact with the Mode Control Panel (MCP). The MCP was controlled using CSD messaging, and was used to make altitude changes planned using the ARAT. Therefore, the crews never had to modify the positions of the dials on the MCP, and were asked to ignore the numbers that were present on that display.

Results

All of the following pilot opinion ratings were gathered at the end of the eight scenarios. As such, they reflect overall opinions, and are not

specifically related to individual scenario manipulations.

Unless otherwise stated, all items were scaled from 1 to 7 with the lower end of the scale having a negative statement (excessive, not desirable, poor, difficult, not effective, easily confused, insignificant, not important), the upper end being positive (not a problem, very desirable, good, easy, very effective, not easily confused, significant, very important), and the middle of the scale (4.0) having a neutral statement (tolerable, adequate, moderately desirable, moderately easy to use, somewhat effective, somewhat confusable, somewhat important). The nominal neutral point of 4.0 was used as an anchor to evaluate all 1 to 7 ratings, and those that were significantly ($p < .05$) above or below 4.0 on a two-tailed t-test are indicated as being significantly positive or negative. When the mean rating did not differ significantly from 4.0, it was indicated to be a neutral rating. In addition to the rated items, participants were also asked to make additional comments regarding the items scaled, specifying things they felt were particularly effective or particularly problematic.

General display clutter

The mean ratings for display clutter when the ARAT was engaged ($M = 4.13$, $SD = 1.34$) and when the ARAT was not engaged ($M = 4.53$, $SD = 1.07$) were both neutral. Pilot comments indicate that the number of aircraft was the main cause for the clutter. They suggested that the ability to remove aircraft that were flying away, at a different altitude, or which could in no way become a threat as means to resolve this problem.

Text

Finding an acceptable text size to accommodate most users while not adding too much clutter to the display is necessary to keep display clutter at a minimum. Text size and readability was looked at for tail tags, aircraft ID blocks, flight plan waypoint names, and ARAT waypoint names. The participants positively rated the first three of these items:

aircraft tail tag altitudes ($M = 5.69$, $SD = 1.03$), aircraft ID blocks ($M = 6.00$, $SD = .91$), and flight plan waypoint names ($M = 5.78$, $SD = 1.28$). However, the rating for the text in ARAT waypoint names was neutral ($M = 4.44$, $SD = 1.70$). Further comments indicated that the text sizes for waypoints may have been adequate, but problems with overlapping waypoints and aircraft symbols, in addition to some problems with the brightness of the waypoint names themselves, limited their readability.

Aircraft ID Blocks

The aircraft ID block contained the flight ID, the current altitude, and the speed of the traffic aircraft, and nearly all participants (15 of 16) felt that aircraft ID blocks were needed display elements. The crew members did not identify any other information that they would want included in the ID block, with the exception of a destination altitude of a climbing or descending aircraft (if that information was available). The ratings for the touch pad controls used to position and display ID blocks, and for the corresponding panel mounted controls, were both neutral ($M = 4.63$, $SD = 1.93$ for the touch pad controls, and $M = 4.88$, $SD = 1.70$ for the panel mounted controls). The comments suggested that some crew members experienced difficulties using the touch pad, resulting in a lower rating, and that either a more sensitive touch pad or more experience using a touch pad would increase their use and improve their opinion of this function.

Aircraft Symbolology

The effectiveness of the aircraft symbolology (size, shape, altitude format, 1 minute predictors to indicate supersector aircraft, and brightness) was rated positively ($M = 5.97$, $SD = .87$), as was the use of color to depict relative altitude ($M = 6.37$, $SD = .90$). There was some concern by the researchers that the non-horizontal format of the tail tags may have proved difficult to read, but only one of the sixteen crew members commented that they would prefer the tail tags to be presented

horizontally. Finally, some crew members commented that the color coding allowed an instant recognition of the traffic situation, although some problems with the brightness of these aircraft was also noted.

Flight Plans

Flight plans received a very positive response. Fifteen of the sixteen participants indicated that being able to view the flight plans of other aircraft was useful. The effectiveness of the flight plan symbology size, shape, waypoints, and altitude segments was rated positively ($M = 5.75$, $SD = 1.13$) as was the use of color in the flight plans to depict relative altitude ($M = 5.84$, $SD = 1.51$). Several participants commented that the use of a broken line to indicate where aircraft were going to change altitude was very useful, and only one participant commented that he didn't quite understand where the aircraft would start and end their altitude maneuver.

Non-ARAT Controls

Crews rated both the desirability ($M = 5.47$, $SD = 1.04$) and ease of use ($M = 5.61$, $SD = 1.43$) of the panel mounted non-ARAT controls positively. In contrast, they gave the touch pad controls neutral ratings ($M = 4.69$, $SD = 1.24$, and $M = 4.47$, $SD = 1.52$, respectively). The non-ARAT controls included all the controls for the predictors (i.e., on and off, time, and mode), a toggle to switch between a relative and an absolute altitude format, a toggle to turn all aircraft ID blocks on and off, a push button to rearrange all tags, and a push button to remove flight plans. The only difference in functionality between the panel controls and the touch pad controls was the increments of time allowed for the predictor control. Since time was modified by tapping or clicking the touch pad on the predictor time display, it was determined in the design that the control should use 2 minute increments. On the control panel, since the device was not as cumbersome, the scale was provided in 20 second increments. Some participants indicated that this difference in the tool was

the reason for their preference of the panel mounted controls. Additionally, crews commented that the control panels had a more direct one-to-one action. With the touch pad, the crews had to move the cursor to the toolbar and tap on the item, which was considered time consuming and inconsistent.

Aircraft Predictors

The participants rated the design of the pulsing predictor very positively ($M = 6.66$, $SD = .60$), while giving neutral ratings for the static predictor ($M = 4.44$, $SD = 2.22$). When asked to rate which predictor they would choose (1 - always static, 4 - equal preference, 7 - always pulsed), the mean rating was significantly favored the pulsing predictor ($M = 6.0312$, $SD = 1.49$). The participants also rated the ease of use of the two input devices when controlling the predictor. The ratings for the control panel were positive ($M = 5.56$, $SD = 1.64$), but neutral for the touch pad ($M = 4.22$, $SD = 1.78$). However, when asked to rate which input device they preferred (1 - always the control panel, 4 - no preference, 7 - always the touch pad), their mean rating was neutral ($M = 3.28$, $SD = 1.84$).

Primary Alert Timing

On average, crews felt that they needed the level 3 alert to trigger a minimum of 5.07 minutes ($SD = 1.62$) before a projected loss of separation, and that this would need to be increased to 6.13 minutes ($SD = 2.03$) for them to have a comfortable period in which to respond.

The crews gave positive ratings to the value of flight plans in increasing the utility of the Primary and ARAT alerting logic, (5.47 , $SD = 1.23$ and $M = 5.37$, $SD = 1.65$, respectively). Several crew members commented that having the flight plans helped them make decisions and was helpful for planning resolutions. The crews gave a neutral rating for the effectiveness of sharing of responsibility with ATC ($M = 4.63$, $SD = 2.00$), although their ratings of the importance of flight plan based alerting logic to "the successful

implementation of air-ground sharing of conflict detection and resolution responsibility” was positive ($M = 5.87$, $SD = 1.52$).

Alert Symbology

For cases where the ARAT was not engaged, the crews rated the effectiveness of the alerting symbology (shape, sounds, function) positively ($M = 5.63$, $SD = 1.20$). Some of the crew members commented that the chimes did cause some confusion for them, but most responded that the alerting symbology was very clear, easy to read, with effective color coding.

The effectiveness of the alert symbology when the ARAT was engaged was also rated positively ($M = 5.82$, $SD = 1.17$). The confusability of ARAT alert symbologies and Primary Alert symbology was rated positively (not confusing) ($M = 5.27$, $SD = 1.24$). Two problems were noted with the symbology. First, when the solutions were entered, some crews could not identify the aircraft that they originally had a problem, and felt that this loss of situational awareness was problematic. The other problem was that the symbology could be confusing when more than one alert occurred along the same path. Here the crew members were referring to a situation where two aircraft were simultaneously on a head-on course toward ownship. This situation, which was the result of a particular maneuver away from a conflict in one scenario, was only experienced by two crews, and generated two sets of level 3 alert symbology.

ARAT Symbology

The design of the ARAT flight plan symbology (Size, shape, waypoints, altitude segments, waypoint table) was considered effective ($M = 5.50$, $SD = 1.14$). The use of color in the flight plan to depict altitude was again considered effective (5.78 , $SD = 1.34$).

Alert Resolution

When crews were asked to rate their preference for maneuvers when resolving conflicts (1 –

vertical maneuvers, 4 – no preference, 7 – lateral maneuvers) they gave a neutral rating ($M = 4.47$, $SD = 2.13$). Furthermore, when asked where they liked to place their initial maneuver (1 – nearer ownship, 4 – no preference, 7 – away from ownship), they were again neutral ($M = 4.38$, $SD = 2.09$). Crew comments indicated that, in most cases, they would evaluate the fuel efficiency of an altitude maneuver, and where it would be more efficient, the simplicity of the maneuver, and the minimal effect on the total distance to travel, would make this their preference. Conversely, if their current course was not direct, they would look for modifications that would minimize distance.

When planning the resolutions, the crews would have placed the Null Point 2.43 minutes ($SD = .86$) ahead of the cancel point. This is in contrast to the 1.33 minutes between the Null point and the Cancel Point in the present design. The crew comments indicate that the additional time was needed in order to complete the communications with the air traffic controller in the real world. When possible, though, they would want the maneuver to occur as soon as possible since maneuvering earlier usually meant less modification to the route was necessary.

Finally, one crew member suggested that the pulsed predictor mode extend along the new ARAT path while the crew members planned a solution.

ARAT Controls

The ratings of the desirability and ease of use of the panel mounted controls for controlling the ARAT were both neutral ($M = 4.94$, $SD = 1.48$ and $M = 5.33$, $SD = 1.68$, respectively). The corresponding ratings for the touch pad controls were also neutral ($M = 4.5$, $SD = 1.57$ for desirability, and $M = 4.18$, $SD = 1.71$, for ease of use). The comments indicated that one problem with the current touch pad was the limitation of its surface area. Benefits of the touch pad, such as the ability to keep eyes on the display, were also noted. The positioning

of the control box was also noted as a problem. In either case, some crew members noted that they liked having both controls present, and their opinions and usage of the tools may have changed if they had more time to become experienced users.

For inserting and modifying the position of horizontal waypoints, crews gave neutral ratings for the desirability ($M = 5.03$, $SD = 1.63$) and ease of use ($M = 4.70$, $SD = 1.60$) of the touch pad. They gave similar neutral ratings for the desirability ($M = 4.4$, $SD = 1.50$) and ease of use ($M = 4.64$, $SD = 1.83$) of the panel mounted controls. The pattern was the same when the crews rated the desirability and ease of use of the two input devices when inserting and modifying the position of altitude change segments. Again the crews gave neutral ratings to the desirability ($M = 4.94$, $SD = 1.53$) and ease of use ($M = 4.63$, $SD = 1.61$) of the touch pad, and to the desirability ($M = 4.43$, $SD = 1.76$) and ease of use ($M = 4.61$, $SD = 1.96$) of the control panel.

Discussion

The main goals of the cockpit display design were to provide good 3-D and 4-D traffic situation awareness, control clutter, and provide an integrated flight-replanning tool. The first two of these goals appear to have been well met. The crews were very enthusiastic about the inclusion of 3-D flight plans. Cashion and Lozito (1999) have also found that pilots desire these types of flight plans. The main difference between the work of Cashion and Lozito and this study, is the enhanced 3-D flight plan format, and the methods for turning the flight plans on an off. The crews were also uniformly happy with the pulsing predictors. Finally, the use of altitude color-coding within both the traffic symbology and the flight plans was uniformly liked.

Given the neutral ratings for the questions about clutter, the goal of clutter management

was only partially achieved. The display concepts and tools dedicated to clutter management did work, however. The crew opinion data and other additional comments made during the debriefings indicated that the crews felt that they were able to maintain the amount of clutter on the display using the tools provided. However these comments also showed that the crews sometimes felt that there were just too many aircraft, and that this obscured important information. In these cases some crews felt that this clutter could have been alleviated if they were allowed to filter out less important aircraft. However, it was also clear from crew debriefing comments that it was when the ARAT was engaged that these aircraft become problematic, with aircraft symbology often overlapping and obscuring more important ARAT symbology.

Some of the suggestions offered during the debriefings to eliminate the clutter problem were: (1) a temporary button that would dim all aircraft except currently alerted aircraft when creating an alternate route with the ARAT; (2) filtering tools that would allow the crews to remove aircraft by importance, (e.g., flying away, behind current position at a slower rate, far to the sides), by altitude, or by some combination of these; and (3) to increase the brightness of the alternate route and dim every thing else on the display when the ARAT was on. Some crews desired tools (e.g., dials to filter on altitude or importance, or temporary suppression buttons) to allow them to select the amount of information presented; others wanted the filtering of information to be preprocessed and automatically presented. Filters of these types would mean that aircraft could "pop" on and off the display as they enter and leave the "importance" filter (i.e., climb, descend, or change their lateral course). Therefore distraction and workload effects, as pilots try to make sense of what is and is not on the display, must be evaluated. Manual or automatic importance filtering must also be consistent with the pilots understanding of importance, and also must ensure safety.

The ARAT also received mixed reviews. The design of the symbology was rated positively, but the crews wanted the Null Point located over a minute further away from the Cancel Point. Their comments indicated that they thought this would be needed in the real world, where additional time would be needed for communications with the air traffic controller. They also expressed no apparent preference when rating if they liked initial maneuvers to be closer or farther from ownship, although some noted that earlier maneuvers often require less modification to the flight plan. The degree of time stress brought about by the Cancel Point was obviously an issue, although the relatively minimal training with the system may have exacerbated this.

The ratings show that the crews' main difficulty in using the ARAT was with the controls. Neither the touch pad operated controls, nor the panel mounted controls, received better than a neutral rating. Comments by the crews indicated that the touch pad had certain advantages over the panel-mounted controls, such as the ability to

keep one's eyes on the display, but the physical operation of the touch pad was difficult. On the other hand some crew members cited the physical location of the control panel as a problem (they had to lean forward and reach for it). Improvements to both systems are possible. The panel-mounted controls could be better integrated into the cockpit. In the present system it was necessary to find an open space in a 747-400 cockpit, and that severely limited the available places. The touch pad, on the other hand, was an off-the-shelf piece of equipment, and was probably smaller than optimal (usable area of 2.5" wide x 1.5" high). Furthermore, the required methods of operation of the touch pad and its associated buttons were also less than optimal, and crew members unfamiliar with the operation of a touch pad had very little training time to become comfortable with it. Much work must go into developing a better touch pad interface for this system before it can be fairly evaluated.

References

- Cashion, P.A. & Lozito, S. (1999). The effects of different levels of intent information on pilot self separation performance. Proceedings of the Tenth International Symposium on Aviation Psychology, Columbus, Ohio.
- Ellis, S.R., McGreevy, M.W., & Hitchcock, R.J. (1987). Perspective traffic display format and airline pilot traffic avoidance. Human Factors, 29, 371-382.
- Johnson, W., Battiste, V., Delzell, S., Holland, S., Belcher, S., & Jordan, K. (1997). Development and demonstration of a prototype free flight CDTI, Proceedings of the SAE/AIAA World Aviation Congress, Anaheim, CA
- Johnston, J, Horlitz, K. and Edmiston, W. (1993). Improving situation awareness displays for air traffic controllers. Proceedings of the Seventh International Symposium on Aviation Psychology, Columbus, Ohio.
- Merwin, D. & Wickens, C. (1996). Evaluation of perspective and coplanar cockpit displays of traffic information to support hazard awareness in free flight. University of Illinois Aviation Technical Report (ARL-96-5/NASA-96-1). Savoy, IL: Aviation Res. Lab.
- Pritchett, A. & Hansman, R. J. (1997). Pilot non-conformance to alerting system commands during closely spaced parallel approaches", MIT Aeronautical Systems Laboratory Report, ASL-97-2, January 1997
- Yang, L. & Kuchar, J. (1998) Using intent information in probabilistic conflict analysis, Proceedings of the AIAA Guidance, Navigation, and Control Conference, Boston, MA.

