ALLOCATION OF FUNCTIONS IN A FAR-TERM
AIR TRAFFIC CONTROL ENVIRONMENT

Jeffrey Homola, Lynne Martin, Joey Mercer, and Christopher Cabrall
San Jose State University/NASA Ames Research Center, Moffett Field, CA
Thomas Prevot
NASA Ames Research Center, Moffett Field, CA

Abstract

A human-in-the-loop exploration of a ground-based automated separation assurance concept was conducted that involved the allocation of certain functions between humans and automation. This exploration included operations that were sustained for prolonged periods of time with high levels of traffic in the presence of convective weather and scheduling constraints. In this environment, the automation was allocated the functions of detecting separation conflicts, resolving strategic and tactical conflicts, providing trajectory trial planning assistance, and alerting the controller to urgent problems. The controller was responsible for supervising the automation, resolving conflicts deferred by the automation, resolving convective weather conflicts, monitoring and maintaining schedule compliance, and placing free track aircraft back onto their trajectory. An investigation into the acceptability of these roles and performance of tasks was conducted where it was found that the participants rated the concept and allocation of functions with a high level of acceptability. However, issues were encountered with the automation related to the detection of and response to tactical conflicts. Lower ratings were given on account of these concerns, and it was found that a key contributor to the underlying problems was transitioning aircraft and the uncertainty of their trajectories. Stemming from those results, participants responded that they would rather have direct control over aircraft transitions as well as more control over the tactical conflict resolution automation. In contrast, participants responded that they would rather have the automation place aircraft back on trajectory, and perform weather avoidance and scheduling tasks.

Introduction

The Next Generation Air Transportation System (NextGen) is expected to provide greater levels of capacity and efficiency while maintaining or improving upon current levels of safety in the National Airspace System (NAS). To do so in the face of a predicted increase in airspace demand requires a departure from the current paradigm of clearance-based air traffic control (ATC) to one that incorporates greater levels of automation and decision support. This paradigm shift, however, raises the important issue regarding the proper balance that must be struck between the controller and automation, particularly as it relates to the critical area of separation assurance. A human-in-the-loop (HITL) simulation was conducted that examined an approach to function allocation in a far-term air traffic control environment, and provided an opportunity to investigate controllers' perceived acceptability of their roles and preferred distribution of functions. This paper focuses on the results of that investigation and expands upon the areas identified as unacceptable.

Function Allocation

Any operational domain that involves a highly complex mixture of information gathering and integration, the need for timely and precise decision-making, and an extremely high level of safety criticality may benefit from, or even require, the introduction of automation. A subsequent allocation of certain functions to the automation previously limited to human operators can then occur, but not without substantial forethought. A great deal of research has gone into the area of function allocation between humans and automation that has spanned a number of domains [1-4]. Air traffic control is certainly no exception [5-8].

The nature of the ATC task coupled with the historical and projected increases in air traffic demand in the NAS has resulted in a variety of research and concept development efforts that incorporate varying levels of automation and allocate functions according to particular focus areas. Two primary focal points that define this area of research are the air- and ground-based approaches to separation assurance and air traffic management as a whole. Broadly speaking, in the air-based approach [9], certain tasks and functions (e.g., spacing, metering, providing safe separation) currently performed by controllers are re-allocated to the flight decks of aircraft.

In the ground-based approach, tasks and functions continue to reside with the control room and related infrastructure. However, function allocation is not directed between controller and pilot, but rather between controller and automation. The Advanced Airspace Concept (AAC) [10] developed by Dr. Heinz Erzberger is one approach that proposes the application of automation to such duties as conflict detection and resolution in both strategic and tactical timeframes, weather avoidance, and metering.

Despite its promises and underlying intent, the efficacy and benefits of automation are not necessarily guaranteed. Within the context of ATC operations, the integration of automation is often thought of as a means of reducing controller workload and summarily providing a path toward accommodating higher levels of air traffic. However, with saturated airspace and automation performing functions that today aid in the construction of situation awareness and maintain controller engagement in the sector's status, it is not clear whether such a path is feasible or acceptable.

Previous studies have shown positive support for this path [11-13], but few, if any, have examined an environment in which operations involving high levels of automation have been carried out in any sustained manner. Performing such an examination allows for a more relevant perspective and the identification of vital issues for a concept and its components than can be gained through brief exposures in isolation.

HITL Simulation

To partly address this issue, a HITL simulation was conducted in the Airspace Operations Laboratory at the NASA Ames Research Center [14] that explored control room operations in a far-term environment. This exploration consisted of four, three-hour simulation runs that tested the limits of airspace capacity and complexity. Such an environment required high levels of automation and humans interacting with it at various levels. This architecture allowed for an investigation into the function allocation issues that arise when tested in such an environment for prolonged periods of time as well as the acceptability of such operations.

Method

Participants

Six current Federal Aviation Administration (FAA) front line managers from different en route centers who were current on radar served as test participants. Each rotated through radar and area supervisor positions in order to operate and evaluate the concept from the two perspectives. Four recently retired controllers staffed the remaining test positions. Four additional retired controllers acted as confederate "Ghost" controllers that managed the air traffic outside of the test area. Ten general aviation pilots participated as pseudopilots and managed the flight decks of the simulated aircraft.

Experimental Design

This study was a 2x2 within-subjects design with Metering and Weather serving as the independent variables (Figure 1). The levels of Metering were Light and Heavy, which related to relative numbers of aircraft scheduled to area airports that also required adjustments and active management to meet scheduled times of arrivals. The distinction between Light and Heavy also translated to overall traffic load in that participants experienced higher levels of traffic in the Heavy Metering conditions than in the Light Metering conditions. As a comparison, the Light-Metering scenarios consisted of 2,216 simulation aircraft whereas the Heavy-Metering scenarios contained 3,060 aircraft.

The levels of the Weather variable were Growing and Decaying. This related to the change in

convective weather patterns over time as they transited the test airspace. In the Growing Weather conditions, convective weather began developing at approximately the midpoint of the simulation run and grew large enough to impact a significant portion of the test airspace. In the Decaying Weather conditions, the run started with weather already impacting the test airspace. As the run progressed, the weather dissipated as it exited the test airspace.

Figure 1. Experiment matrix for the Metering and Weather variables

Airspace

A total of eight, en route test sectors were staffed. As shown in Figure 2, four of the sectors were from the eastern portion of Kansas City Center (ZKC) and the other four were from the western portion of Indianapolis Center (ZID). This area was selected based on prior work and for its range of complexity with some sectors consisting of predominantly level overflights, and others having a mixture of overflights and transitioning aircraft. The altitude strata of the test area were set at flight level (FL) 290 and above.

Operational Environment

The environment simulated in this study was designed to incorporate a number of elements envisioned to be part of NextGen. In terms of aircraft and flight characteristics, full air-ground data communication (Data Comm) and Automatic Dependent Surveillance-Broadcast (ADS-B) equipage was assumed for all aircraft occupying the airspace. Trajectory Based Operations were also assumed to be in effect in which all flights were on

their user-preferred, 4-D trajectories and pre-cleared for their departure and arrival transitions.
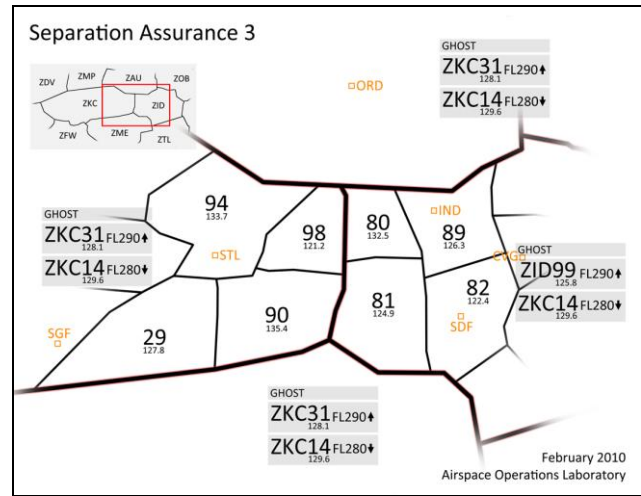

Figure 2. Test airspace centered on the adjacent areas of ZKC and ZID centers

The environment experienced in the control room included sustained levels of high traffic density (with peaks sometimes in excess of 60 aircraft within a sector) in the presence of convective weather and scheduling constraints to area airports. Such an environment required a significant departure from the standard mode of operations conducted today. As shown in Table 1, many of the tasks and functions currently assigned to controllers were instead assigned to ground-based automation.

The automation in this environment was responsible for handoffs and transfers of communication, detection of aircraft and weather conflicts, and the resolution of strategic aircraft conflicts (greater than three minutes to loss of separation) within defined parameters. Tactical aircraft conflicts (less than three minutes to loss of separation) were also handled by the automation through transmission of vectors to one or both involved aircraft via air-ground Data Comm. While controllers were responsible for weather avoidance, schedule conformance, and ensuring aircraft rejoin their trajectories following tactical vectors, their overall role in the separation task shifted from active control to management by exception [15].

Table 1 Allocation of functions between the automation and controller

Automation	Controller
Detect Separation Conflicts	Supervise the automation
Resolve trajectory-based conflicts (if within tolerances)	Resolve trajectory conflicts flagged by the automation
Resolve all time-critical traffic conflicts	Resolve convective weather conflicts
Alert controller to urgent problems	Monitor and maintain schedule compliance
Provide trajectory planning assistance	Place aircraft back on trajectory following automated tactical maneuvers

Just as the environment required a change in the task distribution and roles of the automation and controller, it also required a change in the way that information was displayed (Figure 3). In current en route operations, all aircraft within a sector of ownership are displayed with full data blocks. This is an important requirement as the information contained in the data block and the act of engaging with it helps controllers build and maintain situation awareness of the sector. However, with traffic levels ranging from twice to three times current levels as they did in this simulation, data blocks would quickly saturate the display, forcing the controller to spend an inordinate amount of time trying to separate them. As a result, by default all aircraft were displayed with a limited data block that denoted current altitude and transitioning status when applicable.

With limited data blocks on all aircraft, and the automation handling routine and most separation tasks, the controller was free to focus on weather avoidance and scheduling constraints. The only need to divert from those tasks was when the automation alerted to conflict situations outside its operating bounds or to follow up on a tactical conflict resolution.

Setup and Apparatus

A total of eight R-side sector positions were staffed. The positions were divided into two separate control rooms according to center such that the four ZKC positions were in one room, and the remaining four ZID positions in another (see Figure 4 for the actual layout of sector positions in one room). Each room was also staffed by an area supervisor that had two workstations available for use. One station was used primarily for traffic assessment and display, which controlled a wall projection of a Traffic Situation Display and traffic load graphs for each test sector. The other station was used for schedule monitoring and management when necessary.
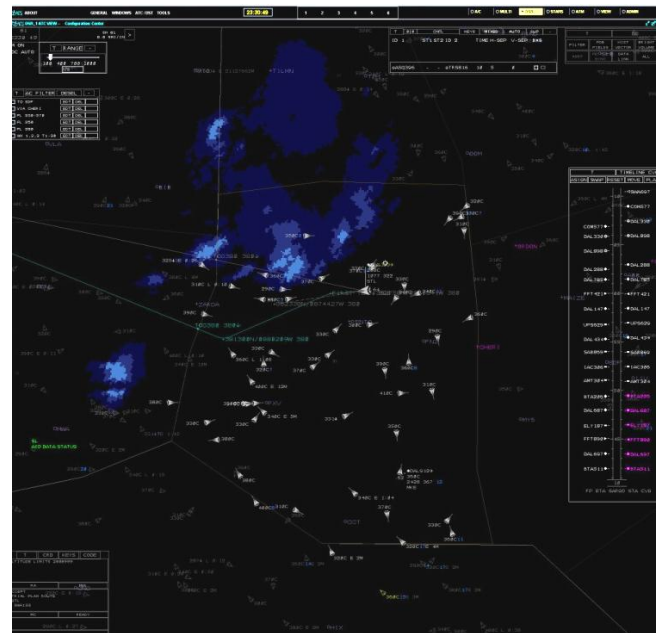

Figure 3. Advanced display with convective weather, interactive timeline and conflict list, with over 50 aircraft in the sector

Three confederate "ghost" stations were staffed that controlled aircraft outside of the test airspace. ZKC Ghost High controlled aircraft to the west, ZID Ghost High controlled aircraft to the east, and Ghost Low controlled aircraft below FL 290. A total of 10 pseudopilot stations were also staffed. Each station was assigned the aircraft owned by each associated test and "ghost" sector.

The common thread that linked all of the stations described was the Multi Aircraft Control System (MACS) software [16]. MACS allows for a wide range of simulation and rapid prototyping capabilities. The scalability of the software provides the ability to produce not only high-fidelity emulations of current displays and interfaces, but extends to the development of conceptual displays, systems, and environments such as those used in this study. All of the stations used in this simulation were networked, and data was communicated between them via the Aeronautical Data Link Radar Simulator. MACS' data collection function was used to record and output a range of data for later analysis. Separate screen recordings of each station were also captured to support the analysis.

Figure 4. Layout of the four R-side positions in the ZID control room during the study

Procedure

After receiving informed consent from all participants, three days of training were conducted in which they were initially briefed on the concept then stepped through the tools and procedures. This portion involved verbal instruction combined with hands-on interaction with the tools in the laboratory.

Following the training period, formal data collection runs were conducted over the course of four afternoons. The runs were three hours in length and tested the four conditions outlined in the Experimental Design subsection. The first condition tested was Light-Metering/Growing Weather followed by Heavy-Metering/Decaying Weather,

Light-Metering/Decaying Weather, and finally Heavy-Metering/Growing Weather.

Controller participants were first assigned to one of four, three-person teams. Throughout the course of each run, these teams rotated through two adjacent sectors such that each sector in the pairing was worked by each team member for one hour. The participants rotated through the positions according to a published schedule. Shift changes consisted of a relief briefing by the outgoing controller, followed by a sign-in entry by the oncoming controller. The relieved controller was then on break for 30 minutes. The first portion of this break was allotted for a short online questionnaire taken in the break room, and the rest of the time was free.

During the runs, the participant controllers operated according to the function allocation outlined in Table 1. Ground-based automation detected strategic aircraft conflicts and displayed related information to the controller in an interactive conflict list (shown in the upper right hand corner of Figure 3). With less than 10 minutes and greater than three minutes to predicted loss of separation (LOS), the automation attempted to find a resolution according to the algorithm developed as part of AAC. If the resulting resolution involved a heading change that was 60 degrees or less, an altitude change that was less than 2200 feet, and a speed change that was 50 knots or less, the resolution clearance was uplinked directly to the appropriate aircraft. The clearance was then automatically loaded and executed through the aircraft's flight management system (FMS). While the controller was not involved in this process, they were notified of the status through the conflict list: a status box on the associated conflict pair's row turned cyan to denote that the resolution was being developed, and, if successful within the stated parameters, the status color changed from cyan to green as the uplink occurred. An additional form of feedback regarding the uplink was available through the datalink status list.

If the resulting resolution exceeded the stated parameters, the status box turned yellow, which signified that the automation had deferred the resolution of the conflict to the controller. A number of options were available for the resolution of such conflict situations. An automated resolution could be requested through the conflict list or directly through the aircraft data blocks, modified if necessary, and

uplinked manually. A specifically tailored resolution could also be developed by the controller through the manual trial planning of trajectories. The trial plans were probed for conflicts in real time and feedback was provided on the display regarding the predicted conflict status. When deemed a suitable resolution, the trial plan could be manually uplinked to the aircraft. In addition to the options just described, the controllers always had the ability to issue verbal clearances as they do today.

In the case of urgent conflicts (less than five minutes to LOS) aircraft symbols were highlighted in yellow on the display. If a conflict progressed to, or was detected with less than three minutes to LOS, tactical vectors were developed by a separate algorithm module for one or both aircraft in the conflict pair. On the display, the involved aircraft were highlighted in red and their data blocks were pushed up to be full. Included in the data block was additional information regarding which aircraft would receive a tactical vector, and what the direction of turn and new heading would be. Once stable, these vectors were sent directly to the aircraft in an attempt to avoid a LOS. Once executed from the flight deck, the aircraft began their turns and the display indicated that they were in free track status. It was then the controller's task to rejoin the aircraft to their original trajectories.

In addition to separation functions, controllers were expected to manage timelines and scheduling constraints to area airports as well as initiate weather avoidance maneuvers. In terms of scheduling, the airports with meter fix constraints were BNA, CVG, MSP, ORD, SDF, and STL. Controllers were tasked with ensuring arrivals to these airports met their Scheduled Time of Arrival (STA). To aid in this task, sector positions were augmented with interactive timelines (shown to the right of Figure 3) that presented each aircraft's Estimated Time of Arrival (ETA) relative to its STA. If the two were out of sync, the controller could request an auto-generated clearance solution to align the ETA with the STA through the timeline. Additionally, aircraft that were early or late relative to their STA had an indication of the time difference included in the data block information. A solution could also be requested through this portal, or a trial plan trajectory could be initiated and manipulated. The impact of the proposed trajectory on the time difference was updated in real time. Once a satisfactory solution was reached, it was manually uplinked to the aircraft.

For weather avoidance, aircraft predicted to penetrate convective weather had data block information that included a blue number, which denoted the predicted time to penetration. In this case, controllers initiated a trial plan trajectory and manipulated it in such a way as to avoid the weather. As the trial plan was manipulated, it was also conflict probed and the time to weather impact updated in real time. Once a clear trajectory was found, the trial plan was manually uplinked to the aircraft.

At the conclusion of each run, participants completed an online questionnaire. Following the final data collection run, a comprehensive post-simulation questionnaire was presented to each participant that covered a range of topics related to the concept and operations as they were presented in the simulation. This was followed by a debrief discussion in which the participants were invited to interact with the researchers to discuss their experience and impressions of the overall simulation and to expand upon topics touched upon in the questionnaires.

Results

The primary results from this simulation were previously reported in [17]. Conclusions reported there were that the concept of ground-based automated separation assurance, as tested, was able to provide consistent benefits in terms of throughput. The automation performed well, overall, particularly considering the density and complexity of the environment. Participants reported general acceptability of the concept, and also provided valuable insight into some of the issues related to their role within the concept and the allocation and performance of functions as they were defined and tested.

The following results and ensuing discussion explore these insights through a combination of subjective data provided by the post-run and post-simulation questionnaires, and objective system data used as a vehicle for supporting and quantifying the subjective data.

Role Acceptability

A number of questions were asked of the participants regarding the functioning and use of the automation, as well as the distribution of roles in the simulation. Perhaps the most direct question posed was simply, "How acceptable was your role in this automated separation concept?" On an interval scale from one to five, with five representing highest acceptability, the mean response rating was 4.33 (SD= 0.69) suggesting a high level of acceptability for the distribution of roles as they were presented and its operational feasibility.

Automation Support and Performance

In terms of the level of support provided by the automation, participants rated that it provided "reasonable" support (M= 3.89, SD= 0.58). As the level of traffic increased throughout a run, participants also felt that they started to rely more on the automation (M= 4.67, SD= 0.49) and that, overall, they needed to place trust in the automation "often" (M= 4.33, SD= 0.59). An important part of this necessary trust and reliance was related to the automation's performance of separation assurance tasks. In terms of conflict detection, participants were asked, "How easy would it have been to determine whether a conflict between aircraft was imminent if you had not had the automation to assist you?" Participants responded that it would have been "almost impossible" without the automation (M= 1.44, SD= 0.70). When asked whether the automation solved an acceptable number of strategic conflicts, participants responded that it resolved a "reasonable" number (M= 3.50, SD= 1.20). A number of conflicts could not be solved strategically however. In such cases, a tactical conflict was presented to the automation for resolution. Acceptability ratings on the resulting tactical resolutions were less than ideal (M= 2.89, SD= 1.08).

To gain a better understanding of these subjective responses, analyses were performed on related system data. While there were four separate conditions, for the purposes of this analysis the data was aggregated across conditions. This was in keeping with the presentation of subjective data at the concept level and allowed for the establishment of a relationship between the two data components. In an effort to show the differences in operational and system characteristics across the test airspace and to help explain some of the rather large standard deviations observed in subjective responses, much of the objective data that follows is separated and presented according to test sector.

The following subsections are divided into two main categories: Separation Assurance, and Function Allocation and Performance. Presenting Separation Assurance first serves as a prelude and provides some background to the latter category of results.

Separation Assurance
Strategic and Tactical Conflict Detections

Following from the questionnaire response regarding the "reasonable" number of strategic conflict resolutions, the number of strategic and tactical conflicts detected was examined. After all, for a conflict to be resolved strategically, it must first be detected as such. In this analysis, for a conflict to be counted it must have met the following criteria: predicted LOS within one of the test sectors, greater than 12-second duration, and more than two consecutive conflict probe reports with less than a 90-second gap between reports. According to this classification scheme, conflicts between the same aircraft pair that were separated by more than 90 seconds and lasted for more than 12 seconds each were counted as two separate conflicts. In addition to these criteria, for a conflict to be classified as strategic, its initial predicted time to LOS needed to be three minutes or greater. In turn, less than three minutes to LOS was classified as tactical.

In this simulation, a total of 2,819 conflict detection events were recorded over the course of 96 sector hours (12 simulations hours multiplied by eight- the number of test sectors). Of this total, 2,458 (87%) were classified as strategic detections, and 361 (13%) as tactical. Figure 5 presents the distribution of conflict detections across the test sectors where it can be seen that the overall number and relative contributions of detection types varied quite a bit between sectors. At the extremes, sector ZKC 29 had relative contributions of strategic and tactical conflict detections at 92% and 8% respectively, whereas ZKC 98 had 81% of conflicts detected strategically and 19% tactically. Differences such as these were likely the drivers behind the large standard deviation observed in the responses to the strategic conflict resolution item from the questionnaire.
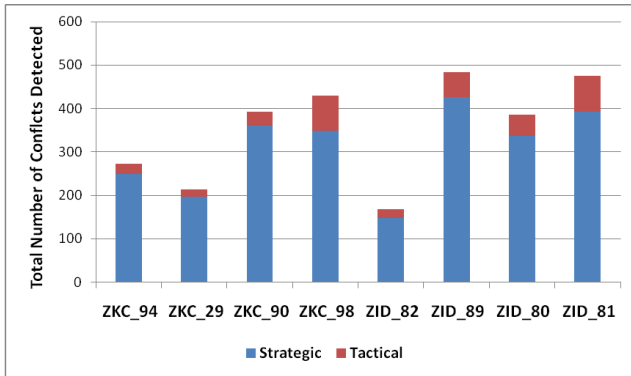
Figure 5. Total number of strategic and tactical conflict detections per sector

Conflict Characteristics

Further analyses were conducted on the conflict detection data to gain a better understanding of the characteristics and potential contributors. Of specific interest was the role that transitioning aircraft may have played in the detection of strategic and tactical conflicts. This was an area of interest because transitioning aircraft have a greater amount of uncertainty associated with their trajectories relative to aircraft at level flight. It is therefore likely that automation (e.g., conflict probe and resolution automation dependent on its information) performance is impacted by such uncertainty.

To examine this impact and the relative contributions of transitioning aircraft to strategic and tactical conflict detections, the two categories were further divided according to conflicts involving level and transitioning aircraft. The classification scheme used for the two categories was Level-Level in which both aircraft were predicted to be at level flight at time of LOS, and Transitioning in which at least one aircraft was predicted to be climbing or descending at time of LOS.

With respect to strategic conflict detections, of the 2,458 total, 1849 (75%) were classified as Level-Level and the remaining 609 (25%) involved transitioning aircraft. Figure 6 presents these results separated out by sector where it can be seen that some sectors (e.g., ZKC 98) had greater numbers of conflicts involving transitioning aircraft than others. This is due to the fact that some sectors had arrivals and departures to nearby airports as part of their traffic mix.

In contrast to the overall 75% Level-Level and 25% Transitioning composition observed for strategic

conflict detections, the makeup of the tactical conflict detections was a near reversal. Of the 361 tactical detections, 134 (37%) were classified as Level-Level and 227 (63%) were Transitioning. This reversal is most evident when comparing the two panels of Figure 6.

From these results, it is clear that there are inherent difficulties in detecting conflicts strategically that involve transitioning aircraft. This is particularly important in the type of environment tested in this simulation in which aircraft were cleared for their arrival and departure transitions without controller involvement or awareness.
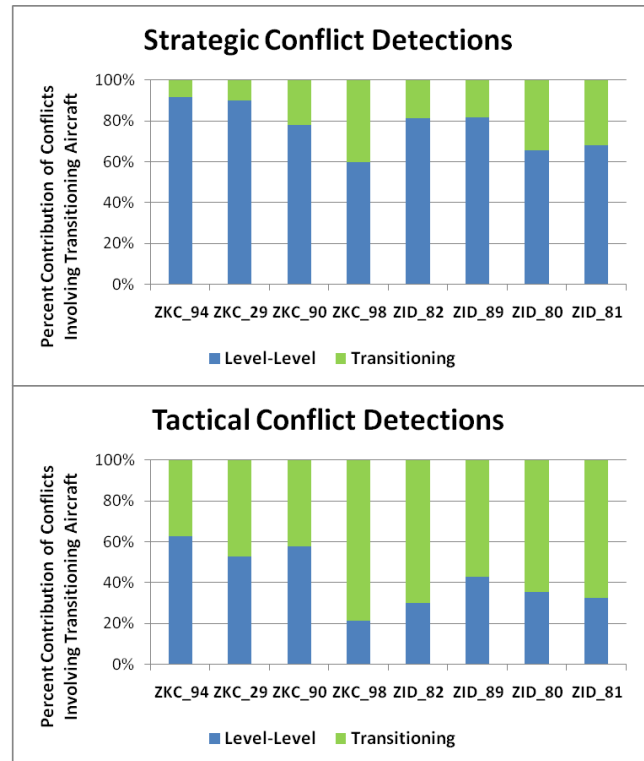

Figure 6. Relative contributions of level and transitioning aircraft to strategic and tactical conflict detections

Tactical Conflict Resolutions

When a conflict was either detected with or progressed to less than three minutes to LOS, a tactical vector was generated and uplinked via Data Comm to one or both aircraft in the pair. Throughout the study, a total of 508 tactical vectors were issued. Figure 7 presents a breakdown of the relative numbers between the test sectors with further categorization according to the magnitude of heading changes involved. This was included simply to get a

sense of the types of clearances that were generated by the tactical resolution algorithm. Of the total number, 51% of the vectors were between 30 and 60 degrees, 30% were less than 30 degrees, 15% were between 60 and 90 degrees, and 4% involved turns of 90 degrees or greater. It is worth pointing out that the issuance of these tactical vectors is heavily dependent on the information supplied by the conflict probe. Based on that, the closer an aircraft pair was to time of LOS (or the later a conflict was detected), the more extreme the required heading change was to avoid the LOS.
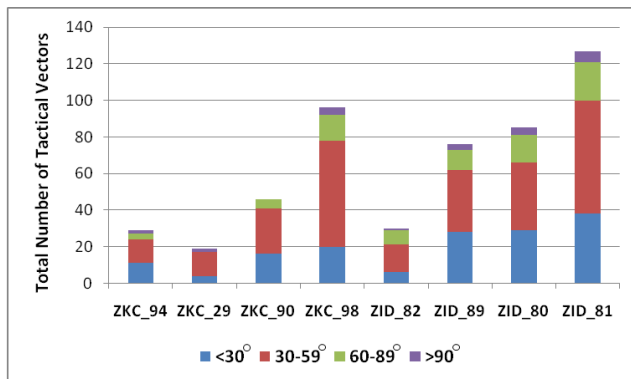

Figure 7. Total number of tactical vectors issued per sector with a breakdown of heading changes

From the post-run questionnaire, the acceptability ratings for the tactical resolutions were rated at (M= 2.89, SD= 1.08), which translated to "sometimes acceptable." Comments provided by participants later in the questionnaire shed some light on why the ratings were fairly low. One participant mentioned that they "got frustrated with the [automation] turning aircraft right when the aircraft is in a left turn to get back on course and avoid traffic." Another noted that, "when an aircraft was stopped in its descent due to traffic, [the automation] would still turn the level crossing traffic. This would compound the problem because the two would not cross but go off on a heading that would need to be corrected. They would also alert with more aircraft." Echoing this statement, another participant remarked that, "some of the [tactical vectors] turned the aircraft into each other with LOS. Then, as they were in free track, they alerted with other aircraft and it just continued."

These comments and others like them allude to the idea that the tactical resolutions themselves were not necessarily the issue, but that the timing, complexity of environment at time of clearance, and lack of control over the situation and automation by the controller were. These factors have implications for controller/automation responsibility and requirements that will be discussed later in the paper.

Losses of Separation

Despite the strategic and tactical auto resolvers working in conjunction with the controllers' efforts, a number of losses of separation still occurred. Over the course of the simulation, a total of 42 LOS events were recorded. Of those, 34 were classified as Operational Errors (closest point of approach (CPA) less than 4.5 nautical miles (NM) laterally and 800 feet (ft) vertically), and the remaining eight as Proximity Events (CPA between 4.5 and 5.0 NM laterally and less than 800 ft vertically).

A more detailed description of the LOS events can be found in [17]. Data presented here are the result of examining these events in terms of issues related to conflict detection, conflict characteristics, and tactical resolutions.

The first concern addressed was the type of conflicts that preceded the eventual LOS with the intent of being able to categorize the LOSs according to those that began as a strategically or tactically detected conflict. This involved a comparison of the LOS events with the conflict detection data presented earlier in the Strategic and Tactical Conflict Detections subsection. Each aircraft pair involved in a LOS was searched for in the conflict detection data set for a match. Results from this effort were that no matches were found that progressed directly from a strategically detected conflict to a LOS; all events classified as a LOS were either initially detected with less than three minutes remaining or not detected at all (at least according to the criteria set out for categorization as a detected conflict).

Similar to the conflict detections reported earlier, further analysis was performed on each LOS event to understand the relative contributions that level and transitioning aircraft might have made. As shown in Figure 8, the majority of LOS events (69%) involved at least one aircraft in transition. This is not necessarily surprising given that it is close to the proportion of transitioning aircraft in tactically detected conflicts (63%). However, it does reinforce the importance of trajectory uncertainty and the need to account for it in some way.

The final analysis performed on LOS data investigated how many of the events had an association with the automated tactical conflict resolver. This was done because the purpose of that safety layer is to avoid a LOS, so it was a way of examining its performance. What was found through this analysis was that of the 42 LOS events, 22 (52%) involved a tactical vector in some way. This meant that in half of the cases, the issuance of an automated tactical vector either caused or failed to avoid an imminent LOS. This is not, however, an indictment of the underlying algorithm or its functioning because there are a number of variables that affected its performance (e.g., late conflict detections, controller interactions, etc.). This result simply points to additional areas that need to be considered as the concept matures.
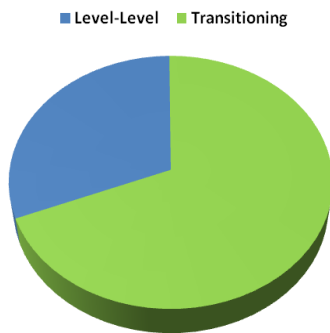

Figure 8. Contribution of level and transitioning aircraft to LOS events

Responsibility

With the automation allocated the strategic and tactical separation assurance functions and the controller managing by exception, the question of responsibility arises. One questionnaire item asked the participants, "Should you become the responsible party when you intervene in a conflict?" Of the 18 response opportunities, 16 were "Yes" responses and the remaining two were non-responses. However, later comments provided by the participants revealed that this was not such a clear-cut issue.

Some comments were in firm support of the responses. For example, regarding controller responsibility, one comment read, "Absolutely. This is why you are there: to monitor the automation and to solve conflicts that it can't." Other comments such as, "Yes, the controller should be responsible for their actions," and, "The controller should shoulder

the ultimate responsibility for the sector," provided further support.

However, comments such as, "I believe if you have enough time to fix the problem, the controller is responsible," begin to show that there are perhaps conditions on the responsibility. Following from this comment are others that are more explicit. Examples are, "If the computer can't come up with a resolution and asks you to but there isn't an out then I don't see how you can place blame/responsibility on the controller," and, "If you issue a control instruction, you are responsible. If the computer issues a control instruction, it is responsible. So, if you are responsible to solve the confliction when the computer cannot, then you are responsible. However, if the computer only gives you a minute or two to resolve the situation, that is a VERY grey area."

These valuable comments point out that the issue of responsibility is an important one that will likely need to be considered very carefully. Vital to this area is the performance of the automation in its assigned tasks. While the automation did perform well, in general, a number of areas for improvement were discovered. These will be discussed following the completion of the results.

Function Allocation and Performance
Task Distribution

Table 1 outlines the function allocation as it was defined and tested in this simulation. To understand how that translated to actual performance of those functions and associated tasks, the number of clearances sent via Data Comm were counted and categorized according to their source. Few, if any, voice clearances were transmitted and therefore not pursued.

Over the course of the simulation, a total of 7,296 clearances were uplinked to aircraft. The initial approach to categorizing these clearances was to separate them according to those initiated by the tactical vector automation, conflict resolution automation, and the controller. Figure 9 presents the overall distribution of these clearances according to test sector. Of the total, 508 tactical vectors were issued by the automation, which accounted for 7% of clearances. The automated conflict resolver was responsible for 1,729 clearances, which accounted for 24% of the total number. By far the largest source of uplinks was the controllers. A total of 5,059

clearances were uplinked by the controllers, which accounts for 69% of the total number and translated to approximately one uplink per minute.
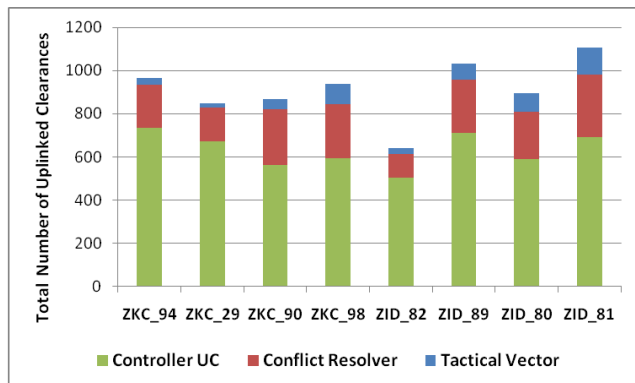

Figure 9. Total number of clearances issued by controllers, conflict resolution automation, and tactical vector automation

Controller Clearances

The observation that the controllers were responsible for 69% of uplinks in the simulation initially came as a surprise given the levels of automation in place. To examine this result more closely, the clearances issued by the controllers were further broken down in order to account for the tasks that that they were in support of.

The task categories were defined according to those assigned to controllers and presented in Table 1. In addition to the conflicts, weather, scheduling, and trajectory management categories, an additional multi-purpose category was added for clearances that met the assignment criteria for more than one category (e.g., a clearance sent to an aircraft with both a trajectory conflict and predicted weather penetration). Assignment to each of these categories was mutually exclusive such that each individual clearance only existed in a single category.

Assignment was based on basic indicators from the output. Assignment to the conflict category required that the aircraft had an associated conflict aircraft at the time of uplink. For weather assignment, the aircraft required an associated weather penetration prediction at time of uplink. For assignment to the scheduling and trajectory management categories, defining the relevant criteria to use was more difficult. For scheduling, a combination of data was used to base assignment (e.g., interaction with timeline for an aircraft just prior to uplink, associated speed assignments,

improved ETA and STA alignment, etc.). The trajectory management category was basically for uplinks that did not meet any of the criteria for the other categories. Upon review of screen recordings, it appears as though many of these were for placing off track aircraft back onto their routes and for route "fixing."

A great deal of time was spent trying to find the most relevant and completely accurate set of criteria for assignment to the scheduling category. Based on the difficulty encountered with that process, a caveat must be included here that there are likely a few clearances that appear in the scheduling category that belong in the trajectory management category and vice versa. This should not affect the overall trend however.

Figure 10 presents a breakdown of the clearances based on the categorizations just described on a per sector basis. Overall, out of the 5,059 clearances uplinked by the controllers, 393 (8%) were for conflict avoidance, 1102 (22%) were for weather avoidance, 2031 (40%) were for scheduling, 1101 (22%) were for trajectory management, and the remaining 432 (9%) were categorized as multi-purpose.
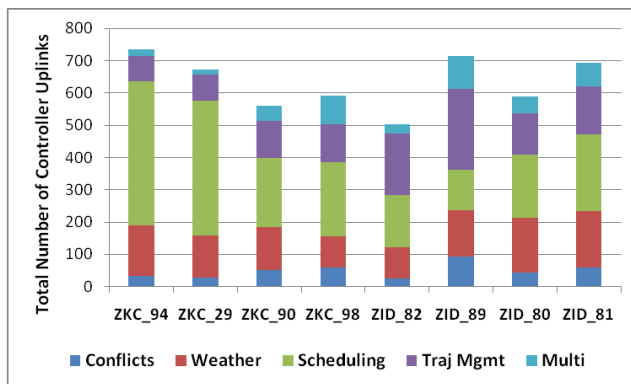

Figure 10. Controller-issued clearances per sector based on task

Based on these results, it is apparent that scheduling made up the bulk of the controllers' taskload followed by weather avoidance and trajectory management. In this context then, the fact that 69% of uplinks in the simulation were from the controllers is not surprising given that the majority of their uplinks were in support of routine tasks that they were assigned.

Desired Allocation

Based on the observations and results regarding the clearances issued both by the automation and controllers, an additional item of interest was how the controllers felt about the taskload that they experienced and what they would change regarding the allocation of functions. This was addressed through questionnaire items that first asked the participants to rate their taskload from the run, followed by questions regarding whether there were tasks that they would have rather performed themselves or tasks they would rather have had the automation perform.

A total of 46 responses were obtained regarding taskload. If a participant responded that their taskload was "low" in the previous run, they were then asked whether there were tasks that they would have liked to perform themselves. Of the 46 ratings, 33 (72%) were low taskload responses. Of those 33, 16 (49%) did not respond that there were tasks they would have rather performed, which suggests that they were satisfied with the allocation of functions. Figure 11 presents the remaining 51% of responses (in blue). Clearly, the task that participants wanted the most control over was aircraft climbs and descents. This relates very closely to the previous findings on the difficulties that resulted from aircraft transitions. The next highest response for function reallocation was for the solving of short-term conflicts. This is, in some way, related to the previous response in that many of the tactical conflicts were the result of transitioning aircraft. By being able to control transitions, they would likely reduce the number of tactical conflicts. Additionally, some participants commented that they would have liked the ability to selectively turn the tactical conflict automation off in order to maintain firm control of the situation.

Of the 46 total ratings, 13 were high taskload responses for the previous run. These respondents were then asked which tasks they would have liked the automation to perform. The most popular response for this item was for the automation to put free track (out of conformance) aircraft back onto their routes. Recall that after every tactical vector, the participant was responsible for this (in addition to other times when an aircraft managed to deviate from its trajectory). When an aircraft was free track, it entered a state in which the conflict probe was limited to five minutes directly ahead. As touched upon in some of the previous comments, there were a

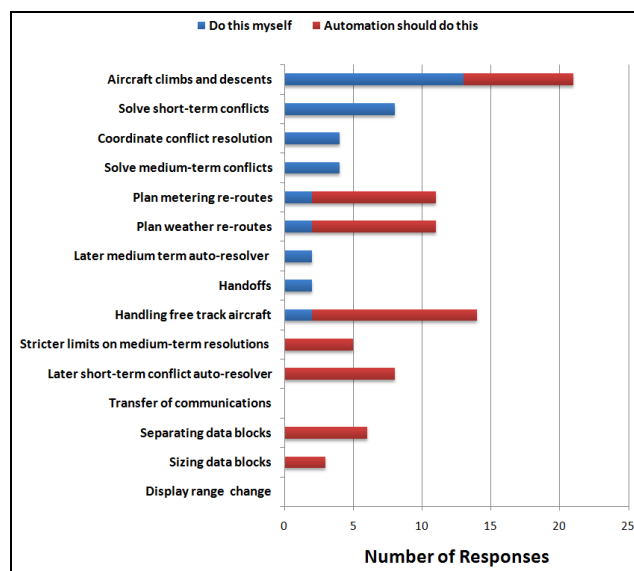number of times that an aircraft entering free track compounded an already difficult situation.

Figure 11. Questionnaire responses for items related to the reallocation of functions

The next most popular tasks that the participants would rather have had the automation perform were the trial planning and uplinking of weather and metering reroutes. This is not entirely surprising given that these were routine tasks that made up the majority of controllers' overall taskload.

For items in Figure 11 that have no response rating, the interpretation is that all participants asked either of the function allocation questions opted to keep the allocation consistent with what they experienced. This meant that participants were satisfied with the automation performing transfers of communication, and that there should be no automation involvement with the participants' display's range settings.

Discussion

Through this simulation, an investigation was conducted into the function allocation between controllers and automation in a far-term environment for prolonged periods of time. This environment was constructed based on the concept of ground-based automated separation assurance, and included high levels of traffic in the presence of convective weather and scheduling constraints. The focus of this paper was to characterize the defined allocation of functions according to the system environment in

which they were performed in addition to the performance and acceptability of assigned roles.

In this simulation, ground-based automation was responsible for the detection and resolution of strategic and tactical conflicts, detection of convective weather conflicts, trajectory planning assistance, and alerting of urgent problems. In contrast, the controller was responsible for monitoring the automation, resolving conflicts that the automation was unable to resolve, performing weather avoidance maneuvers, monitoring and maintaining schedule conformance, and placing free track aircraft back on track.

Overall, participants reported a fairly high level of acceptability for the function allocation as it was presented and tested. These positive ratings corroborate the operational feasibility of this paradigm, which was further supported by the findings in [17].

A significant departure from today's environment was that the automation was performing most of the separation assurance tasks. Given the complex environment in which the controllers were asked to operate, they reported that it would have been "almost impossible" to detect imminent conflicts between aircraft, and that they came to rely on the automation, placing trust in it "often."

In terms of conflicts, a total of 2,819 were detected by the automation: 87% were detected strategically (three minutes or greater to LOS) and 13% tactically (less than three minutes to LOS). When asked if the automation resolved an acceptable number of strategic conflicts, participants responded that a "reasonable" number were resolved. This suggests that there was room for improvement, particularly when 19% of the conflict detections in one sector were tactical.

To understand what contributed to the number of and responses to tactical detections, an examination of conflict cases was conducted. The focus of this examination was on how transitioning aircraft may have impacted the performance of the conflict detection automation. From this examination, it was found that, overall, 63% of conflicts classified as a tactical detection involved transitioning aircraft. Results for one sector in particular (ZKC 98) showed that 79% of its tactical conflict detections involved transitioning aircraft.

For those conflicts that had less than three minutes to LOS, an automated tactical vector was issued via Data Comm to one or both involved aircraft. When asked to rate the acceptability of the resulting resolutions, participants rated them as "sometimes acceptable." Subsequent comments revealed that they were sometimes "frustrated" with this automation over the timing of its resolutions, their lack of control over its functioning, and the additional problems that free track aircraft would occasionally cause following a tactical vector. Taking these issues into account, one participant perhaps summed it up best when they said simply, "The [tactical conflict] automation is good, it just needs to be tweaked a bit to be more practical to work with."

Despite a tactical safety layer being in place, 42 LOS events still occurred. An examination of these 42 cases revealed that they all involved late detections, 69% of the cases involved transitioning aircraft, and in half of the cases, the tactical conflict automation either contributed to or failed to prevent the LOS.

Aside from their contributions to LOS events, late conflict detections and the problems stemming from them provoked a range of responses to the question of responsibility. While all of those that responded felt that the controller should be responsible after intervening in a conflict situation, some provided qualifications to those responses through their follow-on comments. These comments tended to question the correctness of assigning responsibility to the controller if the automation could not resolve a situation, leaving little to no time for them to react. This is certainly an important issue, and one that is outside the scope of this paper.

Taken together, these results show that the uncertainty inherent in transitioning aircraft is a critical issue for the safety of an automated system. This provides firm support for the findings in [18], where it was reported that late conflict detections resulting from aircraft in transition (climbing in particular) were the "largest contributor" to LOS events. Addressing this concern would serve to reduce the number of tactical conflicts and likely improve the performance of the tactical conflict automation by providing it with more accurate information with which to operate. This would improve safety and possibly provide a means of being able to more concretely define the roles and

responsibilities that the automation and controller should have.

Regarding the role of the controller as it was defined in this study, aside from monitoring the automation, they were allocated the tasks of resolving conflicts deferred by the automation, avoiding convective weather, monitoring and maintaining scheduling conformance, and placing free track aircraft back onto their routes. Through the performance of these tasks, the controllers were responsible for 69% of the 7,296 clearances uplinked over the course of 96 sector hours.

While the number of controller issued clearances may seem like a sizeable proportion, when asked about their perceived taskload, the majority (72%) of participants responded that they felt it was low. Thirty-five percent of the time, participants did not offer alternatives to the allocation of functions that were set forth, suggesting that they were satisfied with its structure. For those that did respond with changes to the allocation, the function that participants most wanted control over was that of aircraft climbs and descents. This is closely aligned with the results discussed previously regarding tactical conflicts and losses of separation. Transitioning aircraft played a significant role in the creation of each. It follows, then, that the controllers wanted to remove this concern by being able to maintain control of when and where aircraft could transition. The next function that respondents wanted to control, rather than the automation, was the resolution of tactical conflicts. From some of the earlier comments on their experience with the tactical conflict automation, an emergent theme from the participants was that they would have liked the ability to suspend its actions preemptively or have it turned off upon manual intervention. These capabilities are, in actuality, part of the Advanced Airspace Concept. However, they were not implemented in this study as an intentional experimental design decision.

Over the course of the study, it was found that the total number of controller issued clearances was made up largely of scheduling related uplinks followed by weather avoidance and trajectory management. For those participants that rated their taskload as high, they were asked which tasks they would rather have the automation perform. Not surprisingly, the tasks selected most frequently for automation's responsibility were those three. Placing free track aircraft back onto their trajectories was the most commonly selected task for automation to perform. This echoes some of the earlier comments from participants in that free track aircraft could cause a cascade of subsequent problems and complexity that were difficult to recover from. This was due to the limited conflict probing and uncertainty associated with their trajectories in the presence of high-density traffic and convective weather. Having the automation take care of this task would likely prevent these types of cascading situations from occurring.

Next Steps

Since the completion of this study, efforts have been underway to address some of the issues previously highlighted. Work has been performed on refining the conflict resolution algorithms from both the strategic and tactical domains, and performance testing is currently being performed. A function has also been implemented within the MACS software for free track aircraft to automatically regain their tracks without the need for controller involvement. Additional work has also gone into the convective weather probing and resolution capabilities in an effort to improve upon earlier versions.

Conclusion

The concept of ground-based automated separation assurance and the allocation of functions as tested in this simulation had a high level of acceptability among the participants. However, if the automation is expected to perform the safety critical task of separation assurance impeccably, then improvements are necessary in the areas of trajectory prediction and tactical responses. Work is currently underway to address these and other issues relevant to the progression of the overall concept. Although there are important issues that will need to be worked out along the way, based on the results presented here and the work being performed in response, the concept holds promise for a safe and effective air transportation system of the future.

References

[1] Inagaki, T., 2003, Adaptive Automation: Sharing and Trading of Control, Handbook of Cognitive Task Design, Chapter 8, pp. 147-169.

[2] Billings, C., 1997, Aviation Automation: The Search for a Human-Centered Approach, Human Factors in Transportation, Lawrence Erlbaum Associates, Mahwah, NJ.

[3] Parasuraman, R. & C. Wickens, 2008, Humans: Still Vital After All These Years of Automation, Human Factors, Vol. 50, 3, June 2008, pp. 511–520.

[4] Desmond, P., P. Hancock, & J. Monette, 1998, Fatigue and Auto-Induced Impairments in Simulated Driving Performance, Transportation Research Record, 1628, pp. 8-14.

[5] Parasuraman, R., T. Sheridan, & C. Wickens, 2000, A Model for Types and Levels of Human Interaction with Automation, IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans, Vol. 30, 3, May 2000.

[6] Fitts, P., 1951, Human Engineering for an Effective Air-Navigation and Traffic-Control System, Columbus, OH, Ohio State University Foundation.

[7] Sheridan, T., 2006, Next Generation Air Transportation System: Human-Automation Interaction and Organizational Risks, Proceedings of the 2nd Symposium on Resilience Engineering, Juan les Pins, FR, Available: www.resilience-engineering.org/REPapers/Sheridan_R.pdf.

[8] Metzger, U. & R. Parasuraman, 2005, Automation in Future Air Traffic Management: Effects of Decision Aid Reliability on Controller Performance and Mental Workload, Human Factors, 47, pp. 35–49.

[9] Wing, D., 2005, A Potentially Useful Role for Airborne Separation in 4D-Trajectory ATM Operations, 5th AIAA ATIO Conference, Arlington, VA.

[10] Erzberger, H., T. Lauderdale, & Y. Chu, 2010, Automated Conflict Resolution, Arrival Management and Weather Avoidance for ATM, 27th ICAS, Nice, France.

[11] Farley, T. & H. Erzberger, 2007, Fast Time Air Traffic Simulation of a Conflict Resolution Algorithm Under High Air Traffic Demand, 7[th] USA/Europe ATM R&D Seminar, Barcelona, Spain.

[12] McNally, D. & C. Gong, 2006, Concept and Laboratory Analysis of Trajectory-Based Automation for Separation Assurance, AIAA GNC Conference and Exhibit, Keystone, CO.

[13] Prevot, T., J. Homola, J. Mercer, M. Mainini, & C. Cabrall, 2009, Initial Evaluation of Air/Ground Operations with Ground-Based Automated Separation Assurance, 8th USA/Europe ATM R&D Seminar, Napa, CA.

[14] Prevot, T., P. Lee, T. Callantine, J. Mercer, J. Homola, N. Smith, et al., 2010, Human-In-The-loop Evaluation of NextGen Concepts in the Airspace Operations Laboratory, AIAA MST Conference, Toronto, Canada.

[15] Dekker, S. & D. Woods, 1999, To Intervene or Not to Intervene: The Dilemma of Management by Exception, Cognition Technology and Work, 1, pp. 86-96.

[16] Prevot, T., 2002, Exploring the Many Perspectives of Distributed Air Traffic Management: The Multi-Aircraft Control System MACS, In S. Chatty, J. Hansman, & G. Boy. (Eds). HCI-Aero 2002, AIAA Press, Menlo Park, CA, pp. 149-154.

[17] Prevot, T., J. Homola, L. Martin, J. Mercer, & C. Cabrall, 2011, Automated Air Traffic Control Operations with Weather and Time-Constraints, 9[th] USA/Europe ATM R&D Seminar, Berlin, Germany.

[18] McNally, D. & D. Thipphavong, 2008, Automated Separation Assurance in the Presence of Uncertainty, 26[th] ICAS, Anchorage, AK.

Acknowledgements

30th Digital Avionics Systems Conference
October 16-20, 2011