Analysis of Interactive Conflict Resolution Tool Usage in a Mixed Equipage Environment

Jeffrey Homola[1], Susan Morey[2], Christopher Cabrall[3], Lynne Martin[1], and Joey Mercer[3]
San Jose State University/NASA Ames, Moffett Field, CA, 94035

Thomas Prevot[4]
NASA Ames, Moffett Field, CA, 94035

A human-in-the-loop simulation was conducted that examined separation assurance concepts in varying levels of traffic density with mixtures of aircraft equipage and automation. This paper's analysis focuses on one of the experimental conditions in which traffic levels were approximately fifty percent higher than today, and approximately fifty percent of the traffic within the test area were equipped with data communications (data comm) capabilities. The other fifty percent of the aircraft required control by voice much like today. Within this environment, the air traffic controller participants were provided access to tools and automation designed to support the primary task of separation assurance that are currently unavailable. Two tools were selected for analysis in this paper: 1) a pre-probed altitude fly-out menu that provided instant feedback of conflict probe results for a range of altitudes, and 2) an interactive auto resolver that provided on-demand access to an automation-generated conflict resolution trajectory. Although encouraged, use of the support tools was not required; the participants were free to use the tools as they saw fit, and they were also free to accept, reject, or modify the resolutions offered by the automation. This mode of interaction provided a unique opportunity to examine exactly when and how these tools were used, as well as how acceptable the resolutions were. Results showed that the participants used the pre-probed altitude fly-out menu in 14% of conflict cases and preferred to use it in a strategic timeframe on data comm equipped and level flight aircraft. The interactive auto resolver was also used in a primarily strategic timeframe on 22% of conflicts and that their preference was to use it on conflicts involving data comm equipped aircraft as well. Of the 258 resolutions displayed, 46% were implemented and 54% were not. The auto resolver was rated highly by participants in terms of confidence and preference. Factors such as aircraft equipage, ownership, and location of predicted separation loss appeared to play a role in the decision of controllers to accept or reject the auto resolver's resolutions.

I. Introduction

THE National Airspace System (NAS) is a highly complex and dynamic system that continues to evolve in response to the many different forces and demands acting upon it. An important driver of this process is the traffic demand placed on the NAS and the capacity available to accommodate it. The Joint Planning and Development Office (JPDO)[1] continues to project a significant increase in demand from both passengers and flights in the coming years. Incorporated into this demand increase is a concurrent increase in the mixture of aircraft types as well as levels of avionics equipage operating within the same airspace. Given that today's operations can provide significant challenges to the current supporting infrastructure, changes to infrastructure and the current paradigm of air traffic control (ATC) are being examined, developed, and implemented as part of the Next Generation Air Transportation System (NextGen)[2] modernization program.

---

[1] Research Psychologist, Human-Systems Integration Division, SJSU/NASA ARC, MS 262-4.
[2] Research Engineer, Human-Systems Integration Division, SJSU/NASA ARC, MS 262-4, AIAA Member.
[3] Research Psychologist, Human-Systems Integration Division, SJSU/NASA ARC, MS 262-4, AIAA Member.
[4] Research General Engineer, Human-Systems Integration Division, NASA ARC, MS 262-4, AIAA Member.

A key issue to be addressed in the transition to NextGen is that of controller workload. While the air traffic controller of today manages to provide safe separation of traffic, the ability to maintain such high levels of performance in the face of projected demand increases is doubtful. To that end, there have been efforts to research and develop concepts and accompanying decision support tools that aim to safely reallocate certain functions or tasks from the controller to automation. Such a reallocation would, in theory, reduce the associated workload and enable the air traffic controller to provide separation and service to greater or more varied levels of traffic.

Given that providing safe separation of aircraft is the primary and perhaps most workload intensive task of the controller, the area of separation assurance has been one of particular interest for research. The FAA's NextGen implementation plan calls for the inclusion of automated support for conflict resolution, specifically including trial planning using intent data, and conflict detection and resolution capabilities[3]. The Advanced Airspace Concept (AAC)[4,5] proposed by Dr. Heinz Erzberger is a particular example of a concept developed within this framework. In this concept, reliable ground-based automation is able to offload the task of conflict detection and resolution for appropriately equipped aircraft from the controller. This, in turn, enables the controller to focus on separation for unequipped aircraft as well as provide service on a more strategic, flow-level scale. At its furthest extent, this concept envisions a NAS in which ground-based automation would be able to perform the task of separation, from detection to resolution, for all aircraft without the requirement for direct controller involvement.

An important enabler for such operations is the conflict resolution algorithm developed as part of the AAC. This algorithm takes as input a number of state and predictive input parameters, and iteratively develops a resolution to a given conflict that seeks to maximize efficiency while ensuring safe separation. The Airspace Operations Laboratory (AOL) at the NASA Ames Research Center has served as a test bed for examining various aspects and components of the AAC with controllers-in-the-loop. From August 1-9, 2012, a simulation was conducted that investigated function allocation of separation assurance between controller and automation as well as flight deck and ground[6,7]. Four separate NextGen time frames were tested that ranged from a current day, completely voice and manual control environment, to a far term vision in which separation functions were performed almost exclusively by the automation with the controllers acting as supervisors of the automation.

This paper will focus on the third timeframe from the larger study in which controllers still maintained responsibility for separation, but had access to supporting functions from the automation. This provided a unique opportunity to examine the interaction characteristics of the controller participants with tools that provided on-demand automation-generated conflict resolutions. The results of the ensuing analyses are intended to extend prior analyses[8] on the subject with the intention to provide an understanding of the ways that the controller of a transitional NextGen system might use conflict resolution tools in an operational environment with increased demand and mixed levels of data communications (data comm) equipage. The remainder of this paper will first describe the tools and their modes of interaction, followed by a description of the airspace and traffic, a brief description of the overall procedure, then the results and discussion of those results.

II. Conflict Resolution Decision Support Tools

The current analysis focuses on two primary decision support tools that were available to each of the controller participants in the simulated NextGen environment: a pre-probed altitude fly-out menu and an interactive auto resolver. A description of each of these tools follows.

A. Pre-probed Altitude Fly-out Menu

The pre-probed altitude fly-out menu (see McNally, Erzberger, Bach, and Chan[9] for an earlier description) supported vertical resolutions by providing the controllers with instant feedback regarding the results of a conflict probe being applied to a range of altitude strata. This allowed the controller to quickly see which flight levels were clear of potential conflicts and the time to potential loss of separation (LOS) for the altitudes that were not clear. In contrast to the interactive auto resolver, the pre-probed altitude fly-out menu did not require an active conflict for access. Instead, the controller was able to access the menu via the altitude indicator on the second line of the data block (Fig. 1) at any time. In doing so, the menu displayed a range of altitudes with indications of whether the displayed options were clear or predicted to have a conflict if selected (with an associated time to predicted LOS). This function was important because it provided the controller with pre-probed conflict resolution options in the vertical dimension. As a result, the controller was able to quickly identify available altitudes and/or safe times for arrivals to begin their descents as well as identify altitudes for accommodating aircraft climbing into a sector.
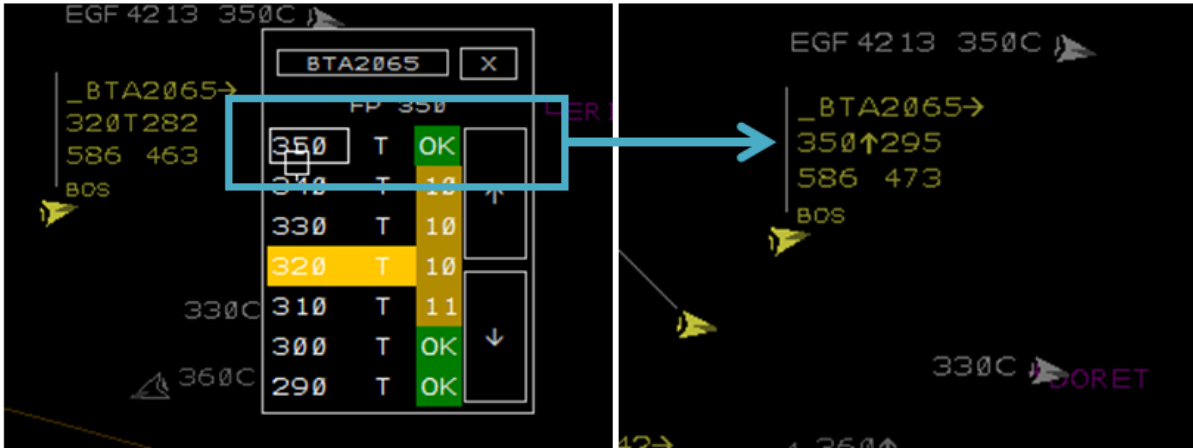
American Institute of Aeronautics and Astronautics

Figure 1. Altitude fly-out menu with conflict probed altitudes (left) and subsequent assignment by the controller.

B. Interactive Auto Resolver

The interactive auto resolver was designed to provide the controller with on-demand access to an automation-generated resolution for a given conflict in both the lateral and vertical dimensions. Access to the auto resolver was available to the controller via multiple control points. Figure 2 presents a graphical representation of these access points. The preliminary requirement for access was an active conflict, displayed to the controller either as a number representing time to loss of separation in the first line of an aircraft's data block or as a distinct row in the conflict list. As presented in Fig. 2, there were a total of four options in which the controller could access the auto resolver. The purpose of having this option set was to be able to consider the controller's maneuver preference for a particular aircraft in the algorithm's resolution development.

Figure 2. Points of access available to invoke the auto resolver and associated resolution preferences.

If, for example, a controller preferred to use altitude to resolve a given conflict, they could select the altitude indicator in the second line of an aircraft's data block. This provided as input to the auto resolver algorithm the soft constraint of searching first for a clear altitude for the aircraft from which the auto resolver was invoked. Likewise, if the controller preferred a lateral resolution for a particular aircraft, they selected the trial plan portal in the first line of the aircraft's data block. If there was no vertical or lateral preference for a resolution, but there was one for which aircraft to maneuver, the controller entered on the conflict probe number in the first line of the data block. Lastly, if there simply was no preference for which aircraft would be maneuvered and in what dimension, the controller selected the conflict pair's associated row in the conflict list to generate a resolution.

It is important to note that while each of the access methods was weighted with a particular preference for resolution type, the algorithm could not always honor that preference due to the requirements in place to find a conflict-free resolution. If the algorithm was simply unable to provide a resolution that fell within the constraints, often due to a dense and complex airspace, the controller was notified that the resolution attempt failed.

Figure 3 outlines a straightforward process for how the auto resolver was often used. The first step was to simply identify the conflict and assess the situation. The controllers often already understood the traffic situation, but could display the conflict pair's routes as well as the confliction point to gain a better awareness of the overall picture. At this point, the auto resolver was invoked, and the resulting resolution was reviewed by the controller for satisfaction, and then modified, canceled, or implemented by voice to unequipped aircraft or via air-ground data comm for equipped aircraft.
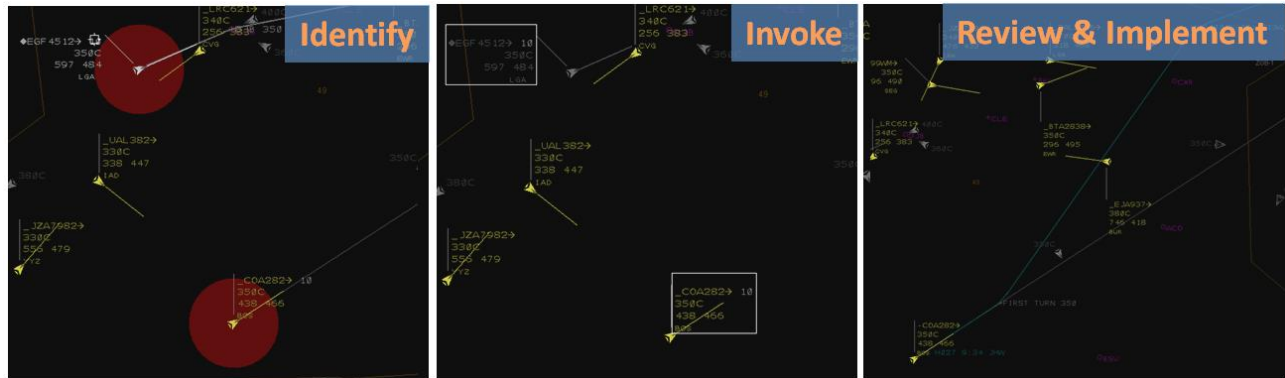

Figure 3. Typical process for interacting with the auto resolver and implementing its solution.

III. Operational Environment and Procedures

C. Operational Environment

The operational environment tested in this particular condition of the overall study was complex and demanding with high levels of traffic and a mix of aircraft data comm equipage. As detailed in Wing et. al., this environment proved to be difficult in terms of safety and acceptability. These results provided impetus to evaluate, at a finer level, the ways in which the controllers performed and interacted with automation within the simulated environment. This section will describe the airspace and traffic composition to provide context, and will be followed by a section that outlines the tasks and procedures followed during the experiment.

1. Airspace

The simulated airspace consisted of five adjacent high altitude, en route test sectors (see Fig. 4). These sectors were assigned to two areas of specialization with sectors 26, 38, and 79 assigned to the North area and the remaining 49 and 59 to the South area. The floor of the overall test airspace was set at flight level (FL) 330. One participant, working as the radar controller, and one supporting confederate controller, working as the radar associate, were assigned to each of the five sectors. Confederate "Ghost" controllers were responsible for the airspace surrounding the test area. The geometries and traffic characteristics of each test sector had differences such that there was variability between the sectors in terms of complexity, flow, and task requirements.

2. Traffic

The traffic scenarios were based on actual traffic from the Cleveland Air Route Traffic Control Center (ARTCC) area, but modified to increase demand levels from current day by approximately fifty percent (Fig. 5). Additionally, the mixture of aircraft was designed to reflect a transitional environment in which fifty percent of the traffic was data comm equipped and the other fifty percent was not data comm equipped; all were Automatic Dependent Surveillance – Broadcast (ADS-B) out equipped. The overall traffic included a mix of level overflights as well as a significant number of arrivals and departures to and from area airports.
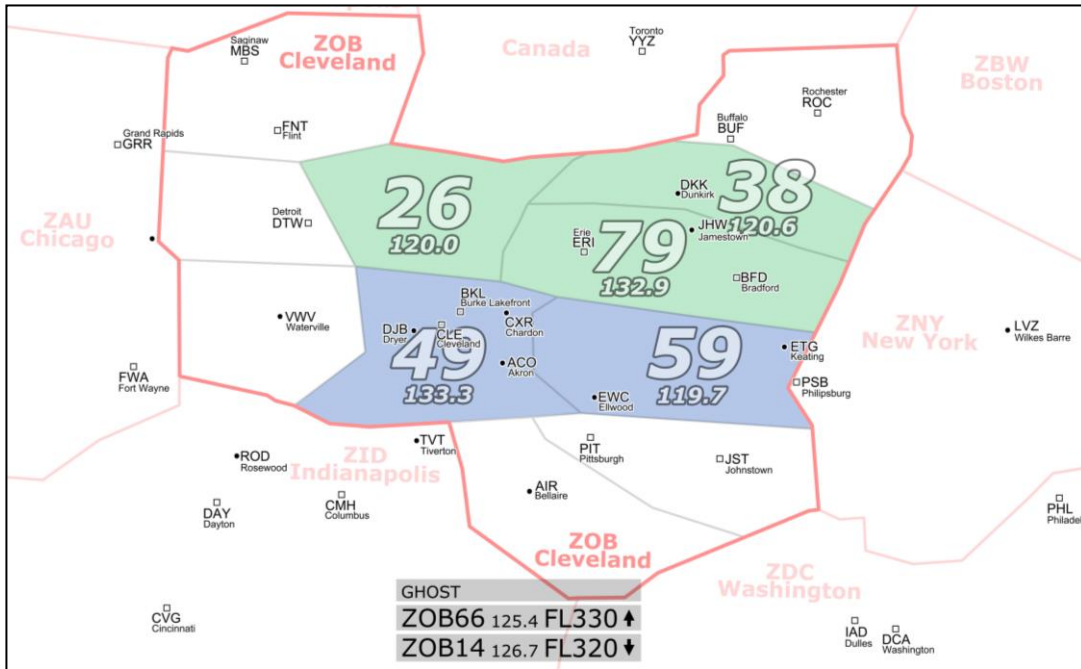
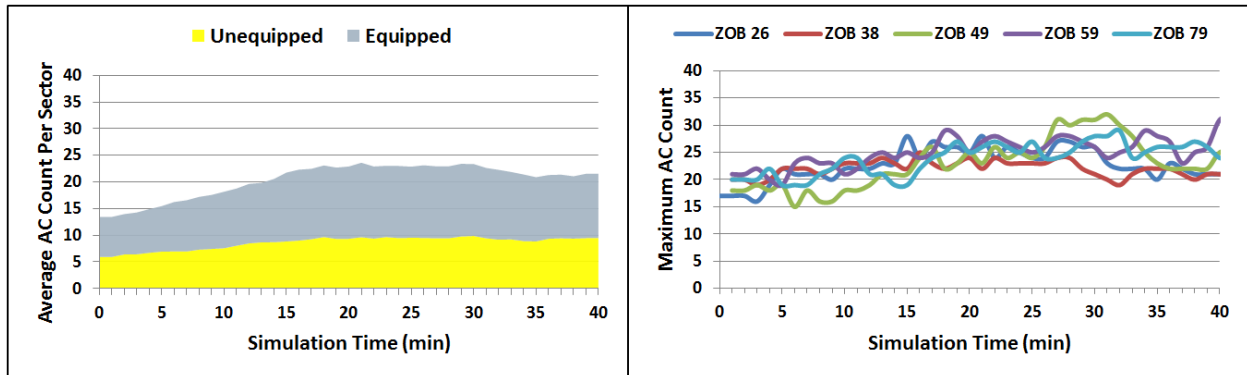**Figure 4. Test airspace within Cleveland ARTCC (areas of specialization denoted by color).**

Figure 5. Average and peak traffic composition over time by data comm equipage (left) and per sector (right).

D. Apparatus

The simulation platform used for the simulation was the Multi Aircraft Control System (MACS)[10], a Java-based software platform developed in the AOL. Each controller workstation was equipped with a Barco display and Display System Replacement (DSR) trackball and keyboard. Voice communications were enabled through a custom, stand-alone voice system, meant to emulate the fielded Voice Switching and Control System (VSCS). Simulation data was collected via MACS's internal collection processes as well as through Camtasia screen recordings and the voice communication application's data recording processes.

E. Participants

A total of seven individuals served as test participants for this study. Six were current front line managers from various US ARTCCs, and one was a recently retired front line manager. Five of the test participants served as radar controllers and two as area supervisors. In support of the test participants, five retired controllers staffed radar associate positions. Three retired controllers acted as confederate "ghost" controllers responsible for traffic outside

the test airspace. Ten airline pilots operated eight mid-fidelity, single-aircraft flight simulators, and ten general aviation/corporate pilots operated multi-aircraft stations.

F. Procedures

There were a total of six runs analyzed for this paper. Each run was 40 minutes in length with the traffic building up gradually to peaks around the midpoint (as seen in Fig. 5). Noted in the airspace description was the fact that the five test sectors were divided into two areas of specialization. The North and South areas of specialization were staffed in physically separate rooms, each with an assigned area supervisor that monitored the traffic situation as well as the workload of the participant radar controllers. It was the decision of the supervisor regarding when to provide radar associate support to the radar controller.

The role of the controllers was much like today: to provide the safe separation of aircraft. In terms of the unequipped aircraft, they were displayed using a yellow chevron as the target symbol and full data block while inside the sector. The controllers were responsible for making and taking hand-offs for these aircraft as well as the transfer of communication. Equipped aircraft were displayed with grey chevrons as target symbols with limited data blocks by default. Hand-offs were made and accepted via automation with an accompanying automatic transfer of communication message sent via data comm. Arrivals to area airports were also handled differently according to equipage. Equipped aircraft were cleared to descend at their top of descent point. The only requirement was that the pilot of the aircraft needed to notify the owning controller when leaving their current altitude. To aid in anticipation and awareness of this event, the limited data block popped up to full when the aircraft was within 150 nautical miles of the destination airport. Pilots of the unequipped aircraft, on the other hand, were required to request a lower altitude upon reaching their top of descent. Departures followed similar procedures in that equipped aircraft were cleared for their climb without the need for controller approval. Departures of unequipped aircraft that would not reach their top of climb prior to entering the test area were pre-assigned a temporary altitude limit of FL 320. It was then the controller's decision to allow those aircraft to climb into the sector to their filed or an amended altitude.

Through the design of the scenarios and the interactions of the controllers with the traffic, a number of conflicts occurred that required some level of controller involvement. Conflicts were detected and displayed automatically to the controller as a time to LOS in the first line of the conflict pair's data blocks and as an added row on the conflict list (see Fig. 2). There was a varied mix of conflict types in terms of level and transitioning aircraft as well as the types of aircraft involved (i.e., data comm equipage).

Controllers had at their disposal tools to aid in the process of conflict resolution. This paper will focus on the pre-probed altitude fly-out menu and interactive auto resolver. The participants were not required to use any of the available tools; they were presented as useful decision support tools available for them to use as they saw fit. This usage paradigm within the context of the airspace and operating environment provided a unique opportunity to observe and analyze the methods of interaction employed by the controllers. The following section will focus specifically on the controllers' interactions with the pre-probed altitude fly-out menu and interactive auto resolver as they involved greater levels of automation. Detailed analyses of these interactions will be presented, followed by a discussion of the results.

IV. Results

To gain an understanding of how the controller participants interacted with the conflict resolution automation, a number of analyses were performed. To provide context, this section will first present results for the analyses of overall conflict data and basic usage data of the automation in that regard. The remainder of this section will be presented in two parts with the first set of results relating specifically to the use of the pre-probed altitude fly-out menu. This will then be followed by analyses of auto resolver requests by the controllers and the subsequent actions by the controllers on the automation's suggested resolutions. This section will conclude with a description of relevant cases that resulted in a LOS and subjective feedback from the participants regarding the auto resolver and its integration into an environment like that simulated.

A. Conflicts

Across the six analyzed simulation runs, there were a total of 727 unique conflict pairs displayed to the controllers. However, because the display of conflicts was tied to aircraft ownership as well as predicted LOS location, there were cases in which a unique conflict was displayed to more than one controller. Therefore, when accounting for the total number of conflicts displayed to all controllers, the resulting aggregate was 1073 conflicts displayed. This was a more appropriate number to use for further analyses because each conflict displayed to each controller was an opportunity for them to use the conflict resolution automation.

The left panel of Fig. 6 shows a histogram of times to LOS at initial conflict detection. This shows that although there were some late conflict detections (i.e., less than three minutes to LOS), the vast majority were in a more strategic time horizon, with most conflicts detected between nine and ten minutes to LOS. The average time to LOS for all conflicts was six minutes fifty seconds. The right panel of Fig. 6 presents a breakdown by percentage of the overall conflict data by data comm equipage level of aircraft in the conflict pair. From this figure, it is clear that the majority (54%) of conflict pairings were mixed with one of the aircraft being unequipped and the other equipped. The next highest composition was conflict pairs involving both equipped aircraft (29%), followed by unequipped aircraft in the pairing at 17%.
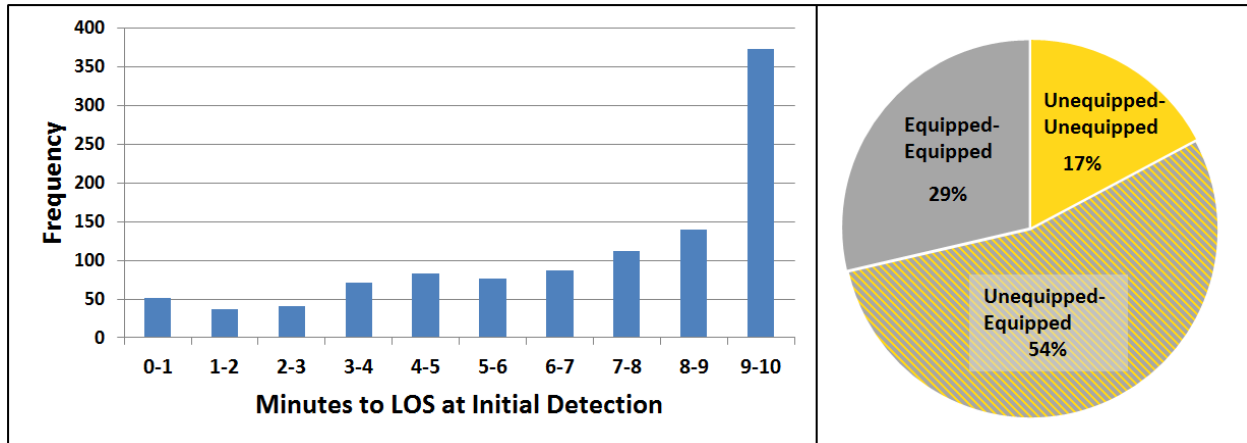

Figure 6. Distribution of times to LOS at initial detection (left) and composition of conflict pairs by equipage level (right).

B. Conflict Resolution Automation Usage

A total of 1073 conflicts were displayed to the controllers throughout the course of the runs of interest. Of that total, controllers used the conflict resolution automation in 354 of the cases. This translates to 33% of the overall number of conflicts in which the pre-probed altitude fly-out menu or interactive auto resolver was used. Between the two tools, the auto resolver was used in 22% of the overall number of conflict cases, whereas the altitude fly-out menu was consulted in 14% of cases. Note that the combined contribution of each tool is greater than the overall 33%. This difference is due to the fact that there were a small number of cases in which both tools were used for the same conflict. Rather than counting each of these cases twice and potentially inflating the overall usage rate, conflicts in which both tools were used counted as a single occurrence in the overall calculation. However, such cases featured equally in the calculation of each tool's relative usage rate.

While not the intention of the current analysis, it may be tempting to make a direct comparison between the usage rates of the pre-probed altitude fly-out menu and the interactive auto resolver and simply conclude that the auto resolver was used more, and, thus, more preferred. However, as noted in the tools' descriptions, the pre-probed altitude fly-out menu was accessible at all times and able to be used for conflict avoidance as well as resolution whereas the auto resolver was only accessible for active conflicts in support of their resolution. Because of these differences in functionality, results for each tool will be presented separately. It is intended that these results be viewed as complementary rather than comparative.

C. Pre-probed Altitude Fly-out Menu

1. Characteristics of Usage

With respect to the pre-probed altitude fly-out menu, the initial interest in its use was simply how and when the controllers decided to use it. In total, the pre-probed altitude fly-out menu was selected for display 519 times. Since the menu was available in both conflict and non-conflict situations, however, the first distinction to draw was simply how the total number of menu displays was divided between the two situations. The left panel of Fig. 7 presents the breakdown of the total count according to the conflict status of the aircraft at time of display. From this chart, it can be seen that the number of times the altitude menu was used on an aircraft in conflict was 235 times (45%), which was slightly less than the 284 times (55%) when the aircraft was not in conflict.

Given that the focus of this analysis is on the tool's usage with respect to conflict situations, the next item of interest examined when the controllers decided to first select the altitude fly-out menu in search of a potential resolution. The right panel of Fig. 7 presents a histogram of the times to predicted LOS at which the controller chose to initially display the menu. From the histogram, it is apparent that there was a tendency to use the menu in a more strategic manner such that the most common timeframe of display was when a given conflict had between eight and nine minutes until predicted LOS.
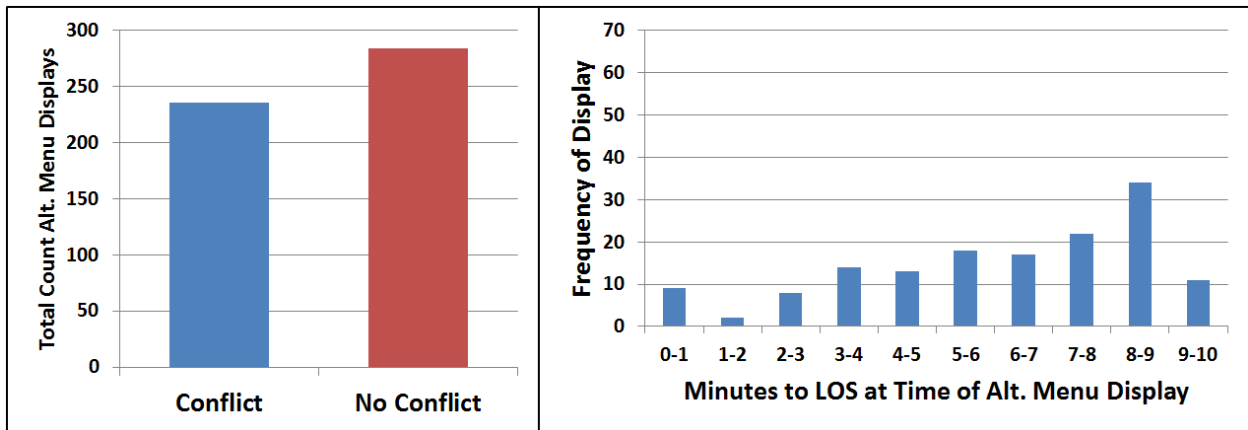

Figure 7. Conflict status (left) and time to LOS (right) at the time of the pre-probed altitude fly-out menu's display.

To further examine how the pre-probed altitude fly-out menu was used in conflict situations, analyses were conducted to assess whether equipage of the aircraft or whether it was maneuvering vertically played a role in the controllers' decision to use the menu. Figure 8 presents the results of this analysis where it can be seen on the left panel that there was a preference for using the altitude fly-out menu for aircraft in conflict that were data comm equipped. This preference may have been due to the relative ease in which clearances were sent to aircraft without the need for the voice communications and accompanying system entries associated with unequipped aircraft.

With respect to whether aircraft were maneuvering vertically and how it may have affected the usage of the pre-probed altitude menu, the right panel of Fig. 8 presents the relative counts of menu display according to whether the aircraft were level, climbing, or descending at the time. From these results, it is clear that the controllers preferred to use the menu on conflict aircraft that were in level flight.
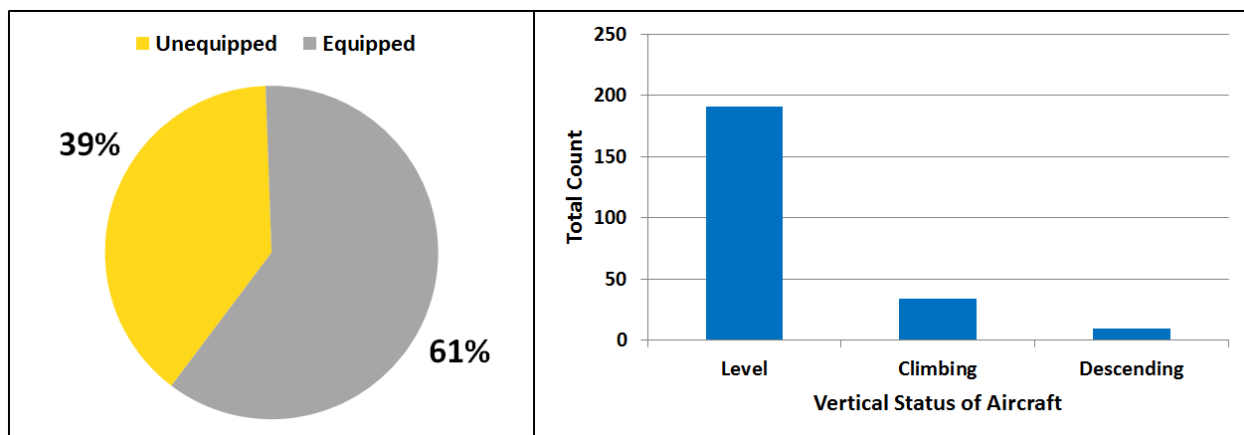

Figure 8. Comparison of altitude menu usage according to aircraft equipage (left) and vertical status of aircraft at the time of display (right).

2. Differences in Usage Between Controllers

In the description of the airspace and accompanying Figs. 4 and 5, it is clear that the sector geometries and associated traffic were not all equal. This prompted further analysis of conflicts on a per sector basis with the altitude menu's usage characteristics framed in that regard in order to see, at a basic level, how the tool may have been used differently.

It was stated previously that the altitude menu was displayed in 14% of conflict cases. Figure 9 presents a breakdown of that 14% in order to show each sector's contribution. The left panel of Fig. 9 presents a comparison of the raw count of unique conflict cases per sector and the number of those conflicts in which the altitude menu was consulted. In terms of conflict distribution, it can be seen that sectors ZOB 59 and 79 experienced the greatest number of conflicts followed by ZOB 26, 38, and 49 with the least number. The right panel of Fig. 9 presents the conversion of the conflict counts and altitude menu displays into associated percentages for each sector. From this panel it can be seen that sector ZOB 59 used the pre-probed altitude fly-out menu in the greatest number of conflict cases at 26%, followed by ZOB 49 at 19% of conflicts. ZOB 26 used the menu in 4% of conflict cases, the fewest of all the sectors, despite having the third highest number of conflicts. From the differences in usage rates between the sectors, it is clear that there is not necessarily a direct relationship between the number of conflicts and the use of the altitude fly-out menu. This is likely due to the fact that the complexity and composition of traffic was different for each sector, and each controller had their own set of preferences and strategies to apply, affecting their interactions not just with the pre-probed altitude fly-out menu but with the interactive auto resolver as well.
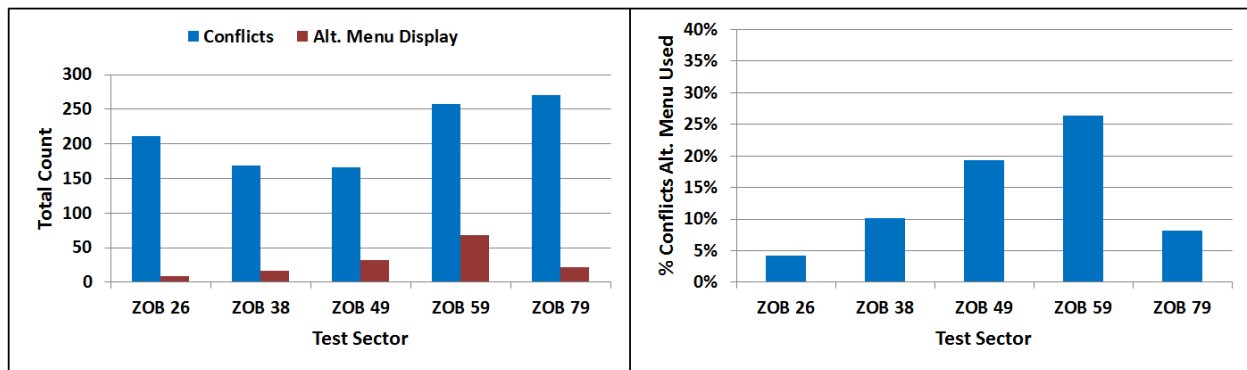

Figure 9. Number of conflicts and times that the pre-probed altitude menu was displayed (left) and percentages of conflicts in which the altitude menu was used (right) on a per sector basis.

D. Interactive Auto Resolver

1. Characteristics of Usage

Figure 1 presents the multiple methods that were available for invoking the auto resolver. The methods available, again, were to select the data block's conflict probe, select a given row in the conflict list, select the data block's trial plan portal, and select the data block's altitude line. Of interest in this regard was the controllers' preferences for how and when to use the auto resolver. The table in Fig. 10 presents results for how often each of the methods was employed in terms of percentages. From this table, it is clear that the preferred method, by far, was to use the auto resolver through the conflict probe number on the first line of the data block. This suggests that, for most cases, the controller did not have a preference for a lateral or vertical resolution. It is likely that the preference was for the resulting resolution to be for the particular aircraft that was associated with the data block used to invoke the auto resolver.

In addition to how the auto resolver was invoked, there was an interest in when the controllers requested a resolution in terms of time remaining until predicted LOS. The question addressed here was whether the controllers used the auto resolver within a more strategic timeframe or used it as a "last minute" support tool. Recall from Fig. 6 that the majority of conflicts were detected with between nine and ten minutes (9-10) to LOS. From the right panel of Fig. 10, it can be seen that most resolution requests occurred within the 8-9 minute to LOS timeframe and with a sharp drop at less than four minutes to LOS. It is important to note that although there were auto resolver requests within four minutes to LOS, it was not designed nor intended for use within such a tactical timeframe. However, overall, it appears as though the controllers tended to use the auto resolver as a more strategic, rather than close-in, support tool.
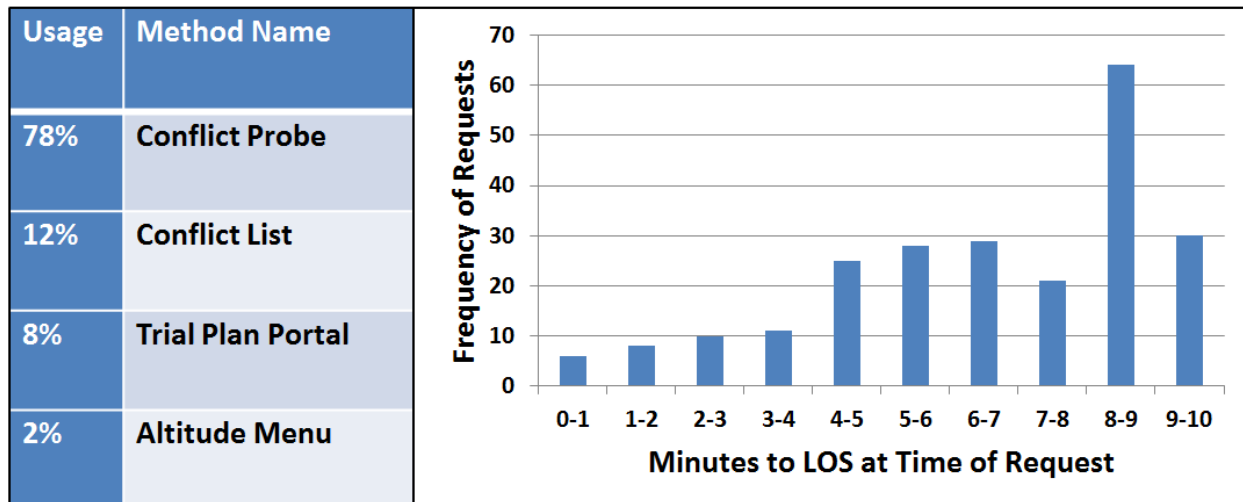
Usage	Method Name
78%	Conflict Probe
12%	Conflict List
8%	Trial Plan Portal
2%	Altitude Menu

Figure 10.   Method (left) and timing (right) of auto resolver resolution requests.

2.   Differences in Usage by Conflict Type and Between Controllers

As previously stated, there were a total number of 1073 conflicts displayed, and the auto resolver was used in 22% of those cases. One item of interest was how the 22% of auto resolver cases was distributed among the types of conflicts in terms of data comm equipage (i.e., was the usage of the auto resolver driven in any way by the equipage types of the aircraft making up the conflict pair?). Results of this analysis showed that 47% of auto resolver usage was on unequipped-equipped conflicts, making it the most common use case, followed by equipped-equipped conflicts at 38%, and unequipped-unequipped conflicts making up the smallest percentage of use cases at 15%.

Although it was shown that the interactive auto resolver was used the most on unequipped-equipped conflicts, that particular type of conflict also made up the majority of conflicts when categorized according to equipage mix. This meant that there was a possibility that greater numbers of auto resolver usage on unequipped-equipped conflicts may have simply been due to there being more of those types of conflicts displayed to the controller. To account for this possibility, further comparisons were made between how many times the auto resolver was used and how often a given conflict type occurred. Results from this approach differed from the previous results in that the auto resolver was used equally for the unequipped-unequipped and unequipped-equipped conflicts (19% in each case), and that there appeared, to some extent, to be a preference for using it in the cases of equipped-equipped conflicts (29%). This was likely driven by the ability of the controller to uplink resolutions to the aircraft without the need for a voice clearance.

In terms of the number of conflicts, the left panel of Fig. 11 presents a breakdown of conflict counts by sector where it can be seen that sectors ZOB 59 and 79 experienced the greatest numbers of conflicts relative to the other sectors with 258 and 270 conflicts displayed respectively. However, plotted next to the total conflict counts is the number of times that an auto resolution was requested. This contrast shows that the number of conflicts in a sector did not necessarily correlate directly to auto resolver usage. The right panel of Fig. 11 presents a clearer depiction of the difference in which the use of the auto resolver is plotted as a percentage of the number of conflicts per sector. Most notably, although sectors 59 and 79 had the highest number of conflicts, the controllers chose to use the auto resolver at the lowest rate compared to the other sectors. Sector 59, for example, used the auto resolver for only 6% of the 258 conflicts displayed choosing instead to manage them with the altitude fly-out menu or manually. In contrast, sector 49, which experienced the lowest number of conflicts, used the auto resolver in 34% of the cases-the highest of all sectors.

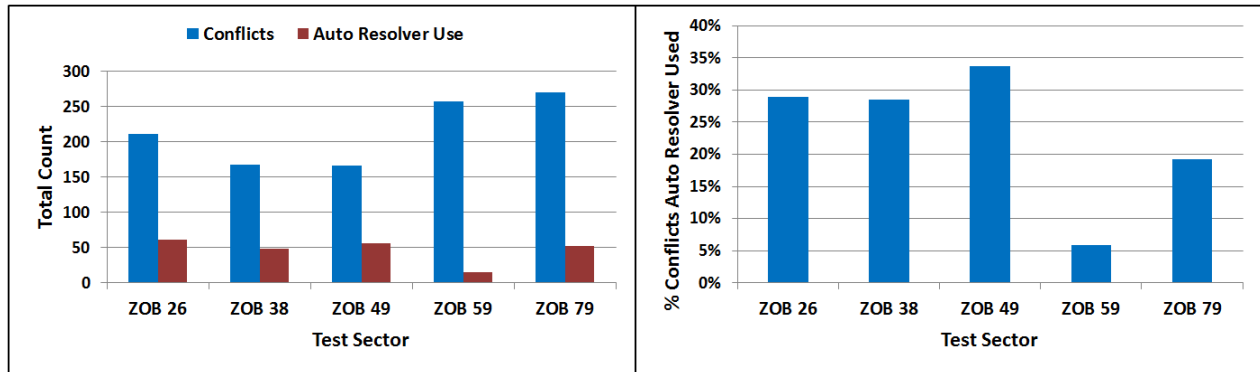American Institute of Aeronautics and Astronautics

Figure 11. Number of conflicts and auto resolution requests (left) and percentages of conflicts displayed for which the auto resolver was used (right), on a per sector basis.

3. Auto Resolver Resolution Implementation

Across the six runs analyzed, the auto resolver was invoked a total of 354 times. In 72 of these cases (20%), the auto resolver was unable to find a resolution within the defined parameters. Seventeen cases were also excluded due to reasons such as the controller re-requesting a resolution before the auto resolver had a chance to generate the first one, or the conflict became inactive during the auto resolver's iteration through the possible resolutions. Of the overall number, a successful resolution was developed within the defined parameters 265 times (75%). Of the 265 successful resolutions, a further seven were excluded from the subsequent analyses (four cases were inadvertent calls to the auto resolver and three were re-requests following shortly after the first one was displayed). This left a final number of 258 resolutions to analyze. The remainder of this section will focus on the number of successful resolutions that were displayed.

Once the auto resolver was invoked, if a resolution was successfully generated, the controller was presented with a trial plan resolution on the display. They then had to decide whether to implement the given resolution or try an alternative. Having examined some of the usage characteristics of the participants when invoking the auto resolver, the analysis that followed was on what was done with the resolution once it was presented. At a basic level, the way that this topic was addressed was by simply categorizing the controllers' actions with the resolution as either "accepted" or "rejected." This was done by creating and examining action sequences from the data output starting with the display of a resolution and continuing until an associated sequence termination entry was encountered. To be categorized as "accepted," this final entry was typically an "uplink clearance" event for trial plans sent to equipped aircraft via air-ground data comm, a route amendment for unequipped aircraft, or a clearance coordination message sent to adjacent controllers for proposed action via ground-ground data comm. To be categorized as "rejected," the final action in the sequence was typically the cancellation of all trial plans for that particular conflict.

Through the methods described previously, of the 258 successful resolutions displayed to the controller, 118 (46%) were categorized as "accepted," while 140 (54%) were "rejected." Figure 12 presents these results both as a whole (left panel) and further broken down by sector (right panel). It is interesting to note the differences in accept/reject patterns between the sectors, particularly in terms of the ZOB 26 controller's tendency to accept resolutions while ZOB 38 showed the opposite tendency.

The high rate of rejection prompted further analysis in an attempt to gain a better understanding of the factors that contributed to this observation. A number of factors were examined that were hypothesized to have an impact on the controllers' decision-making process. These included the flight rule of the aircraft selected for maneuver by the auto resolver, the location of the aircraft at time of display (inside/outside of sector), whether it was owned, the transitioning status of the selected aircraft, the transitioning status of aircraft in the conflict pair, and location where the loss of separation was predicted to occur. Given the differences observed between the sectors as shown in Fig. 12, it was initially thought that a noticeable difference in the form of trend reversals would be likewise observable when examining the results for the factors just listed. This, however, did not turn out to be the case as similar trends were observed between the sectors, albeit at different magnitudes.

The next step was to analyze the factors as aggregates in order to enable comparisons between the contributions of each to the acceptance/rejection rates. After analyzing the results for each of the factors, the only ones that drew any noticeable differences were those regarding the equipage and ownership of the selected aircraft, as well as the

location of the predicted LOS. Given the fact that the number of rejected resolutions came as a surprise, the following results will be framed with respect to the rate at which resolutions were rejected.
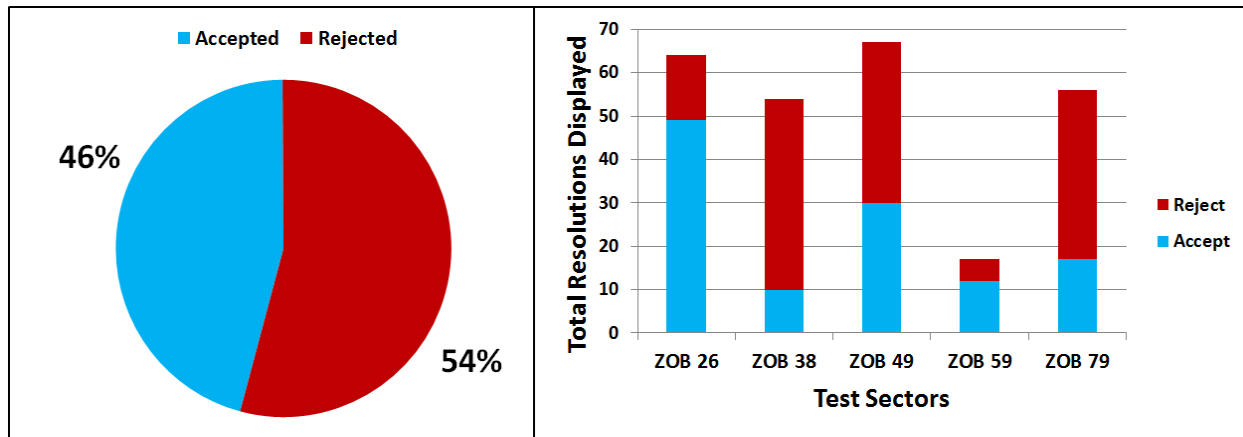

Figure 12. Percentage of resolutions accepted or rejected (left) further broken down by sector (right).

In terms of the equipage of the selected aircraft, it was hypothesized that resolutions for unequipped aircraft would be rejected at higher rates than for equipped aircraft given the extra steps involved in issuing the clearance by voice followed by a route amendment. Of the 258 resolutions displayed, unequipped aircraft were selected for maneuver by the auto resolver 85 times (33%), whereas equipped aircraft were selected 173 times (67%). Figure 13 presents a side-by-side comparison of the resulting rejection rates where it can be seen that resolutions for the unequipped aircraft were rejected 74% of the time whereas resolutions for equipped aircraft were rejected 47% of the time.
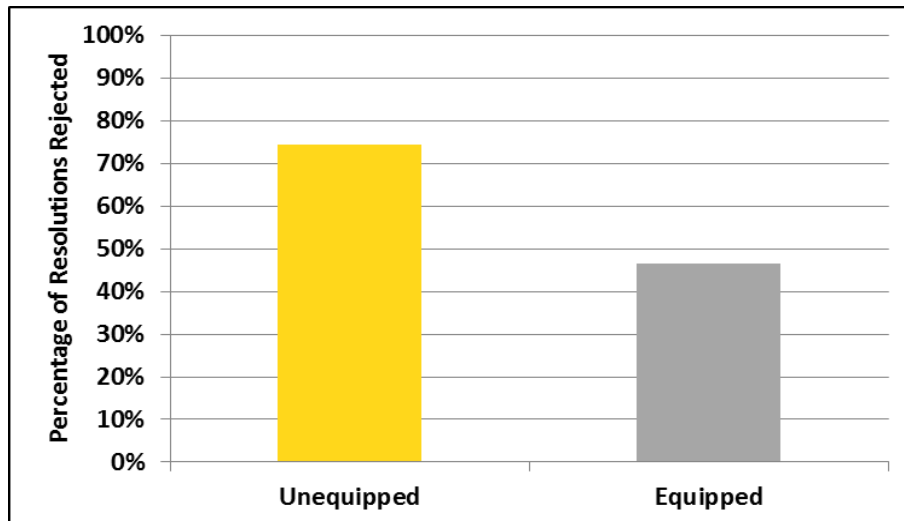

Figure 13. Comparison of rejection rates based on flight rule of aircraft selected for maneuver.

The next factor that exhibited differential rejection rates was that of aircraft ownership. Although the participants predominantly used methods to invoke the auto resolver that identified a preference for a specific aircraft to receive a resolution, they still used the conflict list 12% of the time. Through the conflict list, a preference is not considered and either aircraft in the conflict pair is eligible for a resolution maneuver. Additionally, there was the possibility that a resolution could not be found for a requested aircraft but could be found for the other, which prompted the hypothesis that, given a choice, there would be a tendency to accept resolutions for aircraft that were owned over those that were not owned. This is because if a resolution was suggested for an aircraft that was not owned, the

American Institute of Aeronautics and Astronautics

controller had to engage in some level of extra coordination to resolve the conflict rather than handle it internally. Figure 14 shows that 75% of resolutions for aircraft that were not owned by the requesting controller were rejected, whereas 51% for owned aircraft were rejected.
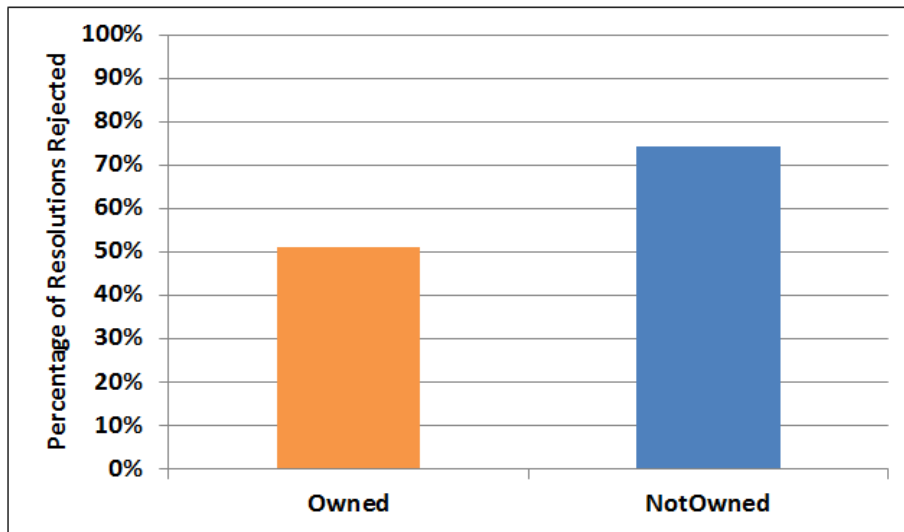

Figure 14. Comparison of rejection rates based on ownership of aircraft selected for maneuver.

The last factor from the current analysis relates to the location of predicted LOS with respect to controller actions in response to displayed resolutions. In this environment, controllers were alerted to a conflict if the LOS was predicted to occur in their sector and/or if they owned at least one of the involved aircraft. This resulted in situations where a controller had an opportunity and a choice either to resolve a conflict that was predicted to lose separation in a downstream sector or to allow the conflict to progress and be handled by the downstream sector. Going into this portion of the analysis, the hypothesis was that there would be a greater tendency to reject resolutions for conflicts with LOS locations outside of a given sector than for resolutions for conflicts with LOS locations inside due to there being, in theory, a somewhat lesser vested interest in resolving such a conflict. Figure 15 presents the overall results where it can be seen that there was a greater tendency to reject resolutions for LOS situations predicted to occur outside the sector with 67% of resolutions rejected as opposed to 45% of resolutions rejected for conflicts predicted to lose separation within the requesting controller's sector.
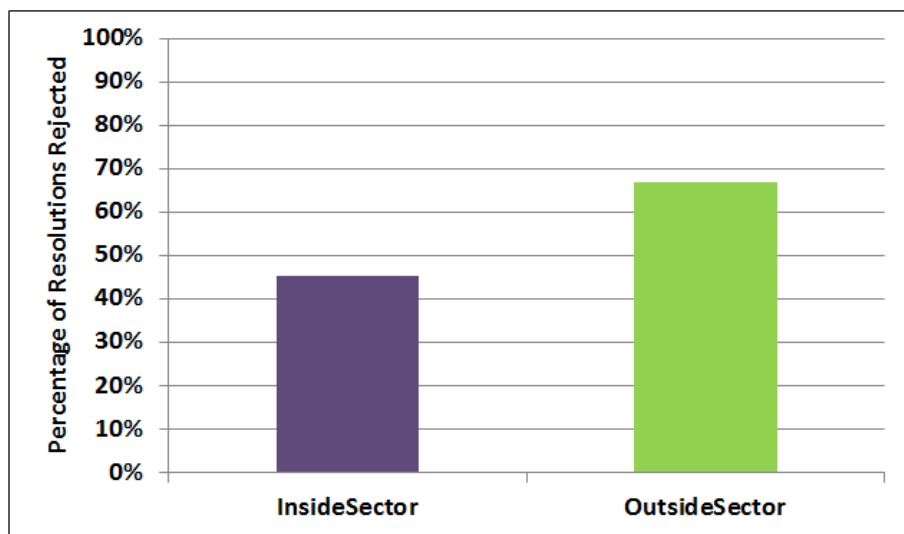

Figure 15. Comparison of resolution actions based on predicted LOS location of the conflict.

4. Review of Reject Cases

Following the factors influencing participants' decision to accept or reject auto resolver resolutions, the question of what was done after rejecting a resolution emerged as a topic of interest. This involved reviewing each of the 140 rejection cases through video display, voice recordings, and trajectory plots in order to assess what, at a general level, the controller's action was following the rejection. From this assessment, eight categories were established to assign actions: Handoff, Other AC, Same Type, Other Type, Retry, External Maneuver, Conflict Drop, and Misc. Handoff referred to cases in which a resolution was rejected and the controller let the conflict continue its progression until the aircraft were handed off either manually or automatically. Other AC refers to when the controller rejected a resolution for a given aircraft in a pair and subsequently moved on to attempt a resolution on the other aircraft in the pair. Same Type refers to situations in which the controller issued a clearance that was of the same dimension as the resolution (e.g., a heading was issued following a suggested path stretch by the auto resolver), but that the clearance was different from the original resolution (e.g., a left heading vs. a path stretch to the right from the auto resolver). Subsequent clearances that were identical to the rejected resolution were re-categorized as an acceptance. Other Type referred to cases in which a resolution in one dimension (e.g., lateral) was rejected and the follow up clearance was issued in the other dimension (e.g., vertical clearance). Retry simply covered cases in which the initial resolution was rejected, but the controller tried a resolution for the same aircraft again. External Maneuver referred to cases in which one aircraft in the conflict pair was maneuvered by another controller that ended up resolving a conflict for which an auto resolution had been developed and rejected. Conflict Drop involved cases, typically with transitioning aircraft, in which a previously detected conflict for which an auto resolution was rejected dropped out without any controller intervention. Finally, the Misc. category was necessary to cover the few cases that could not fit into any of those just listed.

Figure 16 presents the results of this assessment where it can be seen that the most frequent action following the rejection of an auto resolver resolution was to try the other aircraft in the conflict pair (Other AC). This was followed by Other Type and Handoff, with the other action categories following in decreasing order.
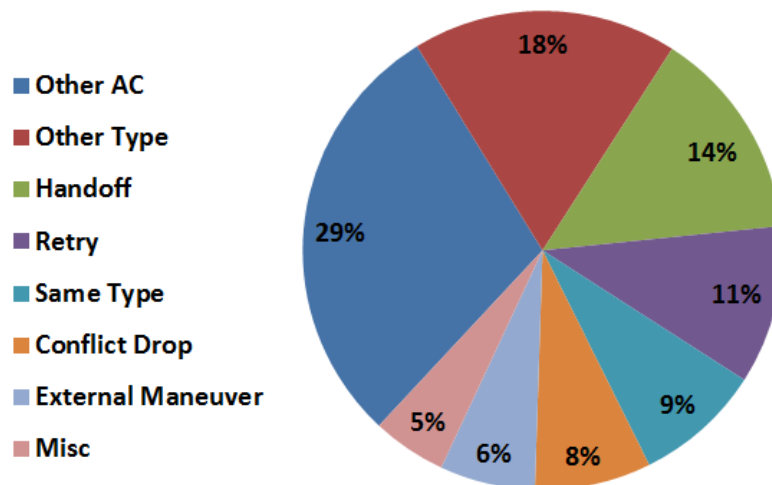

Figure 16. Categories and associated percentages of controller actions following an auto resolver resolution reject.

5. Subjective Feedback

At the end of each run and at the end of the simulation, the participants were presented with questionnaires that included a number of items related to a variety of topics. Of interest are their responses to those items regarding the auto resolver. One question asked them to rate, on a scale from 1 to 7, their confidence in the auto resolver's resolutions. Despite the fact that the participants rejected more than half of the auto resolver's resolutions, the mean confidence rating was 5.50 with ratings ranging from 4 (somewhat confident) to 7 (very confident). Additionally, when asked to rank the provided set of tools in order of preference for using them, two of the five controllers chose the interactive Auto Resolver as their most-preferred tool, and the other three controllers chose the interactive auto

resolver as their second-most preferred tool. When asked to identify the value added by the auto resolver, two chose workload reduction, two chose increased awareness, and one selected safer operations.

Although not subjective in nature, one observed auto resolver usage characteristic that supports the participants' positive feedback involves usage over time. Given that there was a fairly high rate of rejection for resolutions displayed to the controller, it might be expected that usage of the interactive auto resolver would decrease over time due to the fact that the resolutions were in some way not in line with the controllers' plans and strategies. This was not the case, however. The number of invocations of the auto resolver remained steady across the six simulation runs with an average of 43 resolutions displayed per run.

6. Safety

A final item of interest regarding the auto resolver was its role in supporting the controller in their task of maintaining safe separation. Throughout the six runs of this condition, a total of 15 losses of separation occurred. Two of these losses involved conflicts where the auto resolver was used in some way. In the first case, the auto resolver was used on a conflict that alerted less than one minute before the LOS occurred. As stated earlier, the auto resolver was designed for use in more strategic conflicts (greater than three minutes to LOS). In the second LOS, a departure aircraft climbed into the flight path of an overflight. A controller in an adjacent sector (who owned one of the aircraft at initial conflict alert) attempted to use the auto resolver three times successively, but it was unable to find a resolution within parameters. The controller owning the airspace where the LOS was predicted to occur did not use the auto resolver and the LOS occurred. Although the auto resolver was used at some point in both LOS cases, its use cannot reasonably be implicated in the event.

V.   Discussion

The current analysis focused on the interactions of controller participants with an instantiation of a pre-probed altitude fly-out menu and interactive auto resolver in an environment of high traffic density and mixed data comm equipage. Across the six runs analyzed, there were a total of 1073 conflicts that were displayed to the controllers. The majority of these conflicts had nine minutes or greater to LOS at initial detection, allowing for ample time to manage separation. The majority of these conflicts (54%) were composed of mixed equipage aircraft pairs while the others were either unequipped-unequipped or equipped-equipped pairs.

Of the 1073 conflicts, the conflict resolution tools were used in 33% of those cases. With respect to the altitude fly-out menu, the participants used the tool in 14% of the conflict cases and tended to use it most often in a more strategic timeframe with between eight and nine minutes remaining until predicted LOS. The fly-out menu was also used most often on aircraft that were data comm equipped and at level flight. At the sector level, it was apparent that each controller had their own preference for using the altitude fly-out menu during conflict situations that did not necessarily map onto the number of conflicts experienced. For example, the sector with the second highest altitude fly-out menu usage rate experienced the fewest number of conflicts relative to the other sectors. Conversely, the sector with the highest number of conflicts used the altitude menu the second fewest number of times.

The auto resolver was used in 22% of the overall number of conflict cases with some sectors showing a greater usage rate than others. Although there were a variety of methods available to invoke the auto resolver, the participants overwhelmingly chose the conflict probe's time to LOS number in the first line of an aircraft's data block. In doing so, the auto resolver attempted to take into account a stated preference for developing a resolution for the aircraft through which the algorithm was invoked. The timing of the auto resolver's use was also examined where it was found that the majority of use cases were for conflicts with between eight and nine minutes to LOS, suggesting that the participants were using it mostly as a strategic tool. In addition to the time, the participants showed a preference for the types of aircraft involved in a given conflict such that the auto resolver was used in 29% of the equipped-equipped conflict cases whereas it was used in 19% of the unequipped-unequipped and mixed equipage conflict cases. This trend is likely due to the fact that the controllers may have felt that uplinking trajectories to aircraft via air-ground data comm was easier and required less work than to amend a trial planned auto resolution and issue the associated clearance via voice. It should be noted that the researchers did not ask the participants to consider providing service for equipage.

For the conflicts displayed, a total of 258 auto resolutions were successfully developed and displayed to the controller. It was their decision to issue the suggested resolution, modify, or reject it according to their needs and preferences at the time. Of the 258 resolutions, 46% were accepted and issued while 54% were rejected in favor of an alternative. In an attempt to understand the factors that may have contributed to the observed rejection rate, a number of analyses were conducted that looked at relative contributions and trend behaviors between those resolutions that were accepted and those that were rejected.

Stemming from this analysis, the three factors that exhibited any discernible differences were those related to the equipage of the aircraft selected for the auto resolution, the ownership of the aircraft, and the predicted location of the separation loss. In terms of aircraft equipage, the tendency to use the auto resolver for equipped-equipped conflicts extended to include a greater tendency to accept resolutions for data comm equipped aircraft. The opposite was true for unequipped aircraft in which 74% of resolutions displayed for unequipped conflict aircraft were rejected. With respect to aircraft ownership, the majority of resolutions displayed were for aircraft that were owned by the controller. Of those, 51% were rejected whereas 75% of resolutions displayed for unowned aircraft were rejected. A similar trend, albeit more subtle, was observed when comparing resolutions for conflicts with LOS locations inside versus outside of a given sector. Here it was observed that resolutions for conflicts with LOS locations outside of the current sector had a rejection rate of 67% compared with 45% for LOS locations within the current sector.

To gain further insight into the nature of the controllers' interactions with the auto resolver and to get a better understanding of what they were looking for when a resolution was rejected, a review of the rejection cases was performed. The focus of this examination was on the participants' subsequent actions following the rejection of a resolution. Categorizations were constructed to encompass a number of actions, and the most consistent case observed following a rejection was for the controller to attempt a resolution on the other aircraft in the conflict pair. The next most consistent action following a rejection was for the controller to issue a resolution that was in the opposite dimension than the original resolution. Overall, the majority of resolutions that were initially suggested involved lateral path changes. In this regard, it appeared as though the participants often followed a rejection of a lateral resolution with a vertical clearance for the same aircraft.

The interactive auto resolver is a decision support tool designed to offload some of the cognitive workload needed to develop conflict resolutions. Despite the high rejection rates, it appears as though the tool fulfilled its intended role in that the controllers rated it highly in terms of the confidence they had in its performance as well as their preference for it relative to other tools. Initially, given the number of rejections, it was thought that there would be a resulting reduction in the auto resolver's usage as the number of runs progressed, particularly within certain sectors. This, however, was not the case. The tool usage remained consistent overall, even for sectors with the highest rates of rejection.

The disparity between the resolution rejection rate and the positive feedback coupled with the continued use of the interactive auto resolver offers the opportunity for various interpretations. Based on the review process of the rejection cases, one interpretation came to the fore, which was that the controllers were not necessarily using the auto resolver in a strictly straightforward manner (e.g., generate resolution then issue). This also applies to how the controllers used the pre-probed altitude fly-out menu. As described earlier, the airspace simulated in this study was highly complex. In many of the rejection cases reviewed, a pattern emerged where a resolution was invoked and displayed for a particular conflict, subsequently rejected, and typically followed by a different resolution. It may have been the case that the controllers were at times using the altitude menu and the auto resolver as a means of "testing the waters" to see how proposed trajectories would interact with the environment in order to get a sense or heightened awareness of the traffic situation and to inform subsequent decisions and actions. In this sense, the functionality of the altitude fly-out menu and auto resolver, or tools similar to those examined in this analysis, extends beyond conflict resolution to that of providing support for gaining situation awareness. Such an extension could be a useful and necessary addition to the resources available to controllers as the NAS continues its transition to the future environments of NextGen.

VI. Conclusion

The current analysis focused on the ways in which air traffic controllers interacted with two conflict resolution decision support tools in a highly complex future environment of high traffic demand and mixed equipage. The controllers showed a tendency for using the pre-probed altitude fly-out menu for conflicts involving data comm equipped aircraft as well as aircraft at level flight within a strategic timeframe. The controllers also used the auto resolver in conflict cases involving data comm equipped aircraft, aircraft under their ownership, and for conflicts with predicted loss of separation within their own sector. Although 54% of resolutions were rejected, the auto resolver was rated highly by participants. Subsequent review of rejection cases and use of the fly-out menu suggested the possibility that some participants used the tools not solely for conflict resolution but also as an aid to gain awareness of the larger traffic situation and inform subsequent decisions.

Acknowledgments

References

[1]Joint Planning and Development Office, "Concept of Operations for the Next Generation Air Transportation System (Version 3.2)," http://jpe.jpdo.gov/ee/docs/conops/NextGen_ConOps_v3_2.pdf.

[2]Federal Aviation Administration, "NextGen Implementation Plan," http://www.faa.gov/nextgen/implementation_/media/NextGen_Implementation_Plan_2013.pdf [cited July 30 2013].

[3]Federal Aviation Administration, "NAS EA OI 104127- Automated Support for Conflict Resolution," https://nasea.faa.gov/products/oi/main/display/133 [cited July 31 2013].

[4]Erzberger, H. "The Automated Airspace Concept." *4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM, USA*, *2001*.

[5]Erzberger, H. "Transforming the NAS: The next generation air traffic control system." *Proceedings of the international congress of the aeronautical sciences (ICAS)*, Yokohama, Japan, August 30, 2004.

[6]Wing, D., Prevot, T., Lewis, T., Martin, L., Johnson, S., Cabrall, C., et. al., "Pilot and Controller Evaluations of Separation Function Allocation in Air Traffic Management," *submission to Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)*.

[7]Cabrall, C., Gomez, A., Homola, J., Hunt, S., Martin, L., "Transitioning Resolution Responsibility between the Controller and Automation Team in Simulated NextGen Separation Assurance," *submission to ENRI International Workshop on ATM/CNS 2013*.

[8]Mercer, J., Prevot, T., Brasil, C., Mainini, M., Kupfer, M., and Smith, M. "An Integrated Tool Suite for En Route Radar Controllers in NextGen," *Proceedings of the international congress of the aeronautical sciences (ICAS)*, Nice, France, September 19, 2010.

[9]McNally D., Bach, R., & Chan, W., "Field Test Evaluation of the CTAS Conflict Prediction and Trial Planning Capability," *AIAA Paper No. 98-4480*, Reston, VA: American Institute of Aeronautics and Astronautics.

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