https://ntrs.nasa.gov/search.jsp?R=20190001984 2020-03-10T19:29:15+00:00Z

A CLOSER LOOK AT AUTOMATION BEHAVIOR DURING A HUMAN-IN-THE-LOOP SIMULATION

Joey Mercer, San Jose State University / NASA Ames Research Center, Moffett Field, CA

Ashley Gomez, Jeffrey Homola, San Jose State University / NASA Ames Research Center, Moffett Field, CA

Thomas Prevôt, NASA Ames Research Center, Moffett Field, CA

Abstract

A 2012 Human-In-The-Loop air traffic control simulation investigated a gradual paradigm-shift in the allocation of functions between operators and automation. Air traffic controllers staffed five adjacent high-altitude en route sectors, and during the course of a two-week experiment, worked traffic under four different function allocation concepts aligned with increasingly mature NextGen operational environments. These NextGen 'time-frames' ranged from near current-day operations to nearly fully-automated control, in which the ground system's automation was responsible for detecting conflicts, issuing strategic and tactical resolutions, and alerting controllers to exceptional circumstances. This paper continues the investigations reported in previous publications. Analyses of data surrounding the conflict-resolution task serve as the context in which we investigate the interactions between controllers and the automation.

Introduction

The transition to NextGen will include the introduction of conflict detection and resolution tools to help controllers perform their duties. Progress is expected to then continue towards higher levels of automation, enabling controllers' working environment to move from tactical separation management to strategic decision-making. Such automation is envisioned to expand performance beyond today's limits by off-loading workload from controllers onto automated functions for the majority of routine operations [1]. This introduces a fundamental paradigm-shift in which automation is allowed to perform safety-critical tasks that today are strictly the air traffic controllers' domain. Careful and thorough investigation is needed to better understand how the automation performs in such environments, as well as the associated human-automation cooperation issues.

Background

The Airspace Operations Laboratory at NASA's Ames Research Center conducted a Human-In-The-Loop (HITL) simulation, called SA5, which examined separation assurance and function allocation concepts in varying levels of traffic density and mixtures of aircraft equipage [2, 3]. In addition, decision-support tools and ground system automation capabilities varied across four operational environments (i.e., NextGen time-frames), which 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 [4].

The issue of how operators interact with automation is complicated at best. In complex working environments such as air traffic control, could you recognize an effective relationship between operators and automation? Perhaps. Could you recognize an unsuccessful relationship between operators and automation? Hopefully before it's too late. A further understanding of this issue is needed. In the context of the SA5 simulation, this paper builds upon previous analyses to explore the nature of the interactions between air traffic controllers and automation. Specifically, analyses examined how air traffic controllers used automation while resolving traffic conflicts. Comparisons are made across four conditions, distinguished by (among other things) the level of automation available.

The SA5 Simulation

Although documented in [3-9], a brief description of the study's conditions follows.

Baseline 'Current-Day' Time-Frame

The Baseline time-frame approximated a near-term NextGen system, by adding only few differences to current-day, fielded operations. This time-frame assumed that all aircraft had Flight Management System (FMS) capabilities and enhanced surveillance equipage to broadcast their position and state information. Roles and responsibilities in the Baseline time-frame were identical to those in today's operations, in which controllers were responsible for maintaining safe separation between aircraft, and did so with the available support of a flight-plan-aided conflict probe similar to the User Request Evaluation Tool (URET). Controllers issued all clearances via voice.

Minimum NextGen Time-Frame

The second time-frame, referred to as 'Minimum NextGen', added to the Baseline time-frame a limited Data Comm implementation. The Minimum NextGen time-frame assumed 25% of the aircraft were Data Comm equipped, which enabled two primary changes in handling for those 'equipped' aircraft: the automation would automatically perform the hand-offs and transfers of communication for equipped aircraft, and all equipped aircraft had the clearance to follow their FMS-computed vertical profile, unless otherwise instructed by the controller. Decision-support tool enhancements were two-fold: the conflict probe's information was integrated directly in the data block, shown as minutes-to-go until the predicted Loss of Separation (LOS); and trial-planning functions were available to help the controller craft provisional trajectory changes. The trial-planning functions benefited from instantaneous what-if feedback regarding potential conflicts, also a result of conflict-probe integration. Despite the presence of trial-planning capabilities, controllers still issued all trajectory-related clearances via voice, as Data Comm was only available for transfer of communication messages.

Moderate NextGen Time-Frame

The Moderate NextGen time-frame not only assumed that 50% of all aircraft were equipped for Data Comm, but also expanded the Data Comm capabilities to allow for the sending of trajectory changes. This gave the controller the ability to send any route modifications and/or altitude changes created through the trial-planning functions directly to the aircraft. The Moderate NextGen time-frame also added two specific decision-support tools to address the strategic and tactical aspects of conflict resolution, based on key elements of Erzberger's Advanced Airspace Concept [10]. In the presence of a detected conflict, controllers could invoke an Auto-Resolver algorithm to request a new trajectory from the automation that resolved the conflict. Presented in the form of a trial-plan, the controller could then send the trajectory change to the aircraft via Data Comm, just as if they had manually created the trial-plan themselves, or modify it, making any desired adjustments before sending it to the aircraft. Several access points were available to the Auto-Resolver, offering a means for the controller to communicate vertical, lateral, and/or aircraft-specific preferences to the automation [5]. While the Auto-Resolver addressed conflicts with medium-term or 'strategic' time-horizons (typically less than eight minutes until LOS), the Tactical Separation Assured Flight Environment (TSAFE) algorithm was available to provide advisories in response to any short-term, 'tactical' conflicts (less than three minutes until LOS) [10]. In the event of a short-term conflict, TSAFE would compute the aircraft heading changes needed to avoid the pending LOS. TSAFE then displayed the suggested headings to the controllers in the 5th line of the aircraft's data block. The advisories were informational only – action was still required on behalf of the controller to resolve the conflict, and whether or not they incorporated the TSAFE advisories while doing so, was at their discretion.

Maximum NextGen Time-Frame

The Maximum NextGen time-frame assumed all aircraft had Data Comm equipage, but more importantly, it represented a significant change to the air traffic operations. With full Data Comm equipage, it was technically possible to send all instructions electronically, enabling a new distribution of tasks between the controller and automation- with both members of the sector team using Data Comm to take care of certain responsibilities.

The automation's responsibilities included six tasks: 1) detecting conflicts, 2) using TSAFE

advisories to avoid separation violations, 3) using the Auto-Resolver to resolve medium-term conflicts (within certain pre-defined limits), 4) for aircraft receiving a TSAFE instruction, also sending a new trajectory to put the aircraft back on course to rejoin its original route, 5) all hand-offs and transfers of communication, and 4) alerting the controller to any problems or exceptional situations. Meanwhile, the controller was responsible for supervising the automation and for resolving any situations flagged to them by the automation. Controllers also had the option to inhibit the automation for selected aircraft, similar to a manual over-ride, allowing them to address a particular situation in their own manner.

Simulated Airspace and Traffic Scenarios

The simulated airspace used for this study, shown in Figure 1, consisted of five adjacent test sectors, all in the high-altitude (flight-level 330 and above) en route airspace of Cleveland Air Route Traffic Control Center (ZOB). The test airspace spanned two areas of specialization, which grouped three of the sectors in a 'North' area, and two sectors in a 'South' area. One participant, working as the radar controller, and one supporting confederate controller working as the radar associate (D-side), staffed each of the five sectors. However, the D-sides were not necessarily always staffed. Instead, each area of specialization (housed in separate rooms) had an area supervisor monitoring the traffic situation and controller workload, who decided when to staff the D-side position, based on their determination of when the radar controller needed additional support. Confederate "ghost" controllers were responsible for the airspace surrounding the test area.

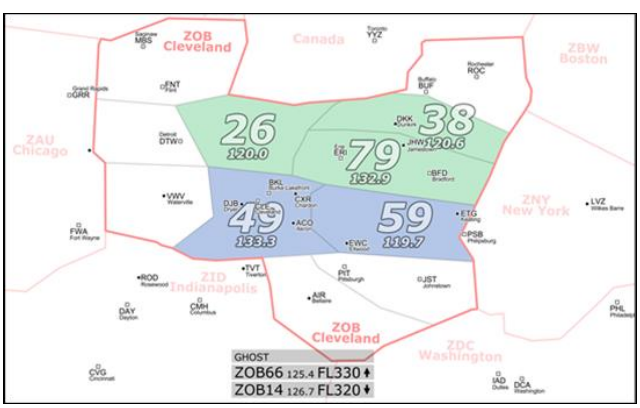

Figure 1. Test airspace used for the simulation.

Originally based on actual traffic flows from the ZOB area, the traffic scenarios included a mix of arrivals and departures from nearby airports, as well as overflights. Combined with the different sector geometries, the scenarios provided natural variations in complexity across the sectors. Traffic levels were representative of today's operations, with a Monitor Alert Parameter (MAP) value of 18 aircraft per sector. However, as the study progressed through the different NextGen time-frames, each new condition saw an increased level of traffic. There was a concern that as the technologies associated with each time-frame were introduced, the operations would require little-to-no workload and would not fully exercise the concepts. In response, the increasing traffic levels sought to expose the participants to varying degrees of workload, thereby more rigorously testing the different operations. As a result, the traffic levels increased by 20%, 50% and 100%, translating to MAP values of 22, 27, and 36 aircraft per sector in the Minimum, Moderate, and Maximum NextGen time-frames, respectively.

Participants and Equipment

Six current Federal Aviation Administration (FAA) front-line managers, and one recently-retired front-line manager, served as primary participants; five as radar controllers and two as area supervisors. Additionally, eight retired controllers supported the test participants by assuming the D-side and 'ghost' controller roles. Twenty pilots, with general aviation, corporate, and type-rated qualifications, operated the aircraft workstations. All participants received payment for their participation in the study.

The primary simulation platform used for the study was the Multi Aircraft Control System (MACS) [2], which, for each controller workstation, hosted a Display System Replacement (DSR) emulation on a large-format monitor. The controller workstation also included a specialized keyboard and trackball, similar to those used in current air traffic control facilities, as well as a custom, stand-alone voice application emulating the fielded communication system. Data recorded and collected at each workstation included aircraft flight states, operator task data and workload, automation states, voice communications, etc.

Previous Findings

The simulation generated a rich set of data, providing many valuable insights. Prior work has reported results on the impact of traffic density and sharing of intent information on flight-path efficiency [3]. Other analyses explored the relationship between recorded separation violations, and the nature of the conflicts that led to their occurrence, as well as the controllers' interactions with the conflict detection and resolution tools [5]. Analyses also showed differing tool usefulness ratings between strategic and tactical decision-support tools [6]. Additionally, an investigation of the factors contributing to workload across the four time-frames identified counts of conflicts, transitioning aircraft, and non-Data Comm-equipped aircraft, as well-correlated with the controllers' reported workload [7]. Another analyses examined how the controllers interacted with the automation, and how the traffic situations unique to each sector influenced those interactions [8], while yet another analyses focused just on the controllers' decisions to inhibit the automation [9].

Method

While extensive, the analyses to-date have not closely examined how the controllers resolved conflicts under the different NextGen time-frames. While [3] and [4] offer complete overviews of the study, their detail is very general. Results from [5] and [6] provide initial insights on the conflict-resolution task, but do so from a limited scope: both focused specifically on the Moderate NextGen time-frame, as it exhibited several unique traits that warranted targeted investigations. The findings in [8] and [9] also hint at this subject, but only in the context of the Maximum NextGen time-frame: [8] was more concerned with the link between general automation usage and individual sector characteristics, while [9] was even more specific, looking at trends in the automation-inhibition actions.

This paper continues the analysis of the SA5 data, exploring two areas. First, characterizations of all the addressed conflict situations are made to provide context for the controllers' actions, and to describe how the different NextGen time-frames impacted the traffic picture. Second, the sector teams' specific actions in response to conflicts are studied to identify any trends and/or strategies in their resolutions, and to see if the various NextGen time-frames showed any effect on the manner in which conflicts were addressed.

The analyses in this paper all share the same starting point: the set of selected conflicts. An initial look at the data revealed that some detected conflicts were never addressed by any action on behalf of the controller or the automation. False alerts are the most likely cause here, which contribute little to this paper's main focus: how the controllers and the automation resolved conflicts. For this reason, the analyses herein are not in relation to *all detected conflicts*. Instead, the set of conflicts comes from the sum of the following: all conflicts that ended in a LOS event, and all conflicts for which actions were recorded during the conflict's duration. For conflicts meeting either selection criteria, the first and last times at which a given conflict was 'active' (i.e., logged in the data), defined the conflict's window of duration. Further examination of the data files then produced a compilation of all related actions logged during the conflict's window of duration. A final filtering of the data removed cases in which the recorded action couldn't be traced back to the conflict-detection task (e.g., the only action logged during the conflict's window of duration was a transfer of communication message).

Organizing the data on the sector teams' actions benefited from 'action trees', which mapped the possible actions into relevant maneuver categories (i.e., altitude, speed, etc.). As an example, Figure 2 shows the Baseline condition's action tree. Other parameters associated with the actions afforded further investigations. The time stamps of each action, and their relation to the conflict's time of first

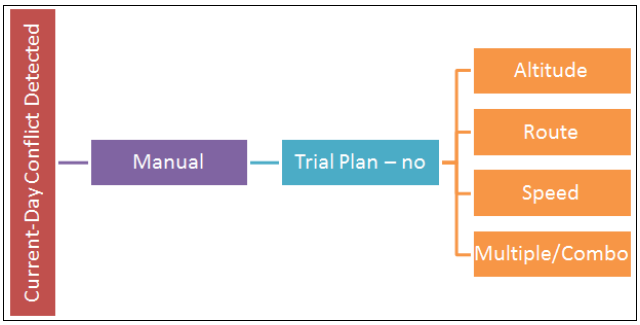

Figure 2. Action tree representing possible ATC actions within the Baseline condition.

detection, as well as which aircraft the controllers maneuvered in order to solve the conflict, help explore how controllers chose to resolve traffic conflicts.

Results

The analyses presented in this paper examine how controllers chose to solve conflicts, and in particular, search for any sensitivities in the data to the different NextGen time-frames. Breakdowns of the analyzed set of conflicts are followed by an examination of the sector teams' actions.

Conflicts Acted Upon

To help understand the actions in response to the analyzed conflicts, an initial examination of the data first describes the conflicts themselves. The set of analyzed conflicts - that is, conflicts ending in separation violations or for which existed logs of relevant actions, consisted of 1076 total conflicts. As expected, the number of conflicts worked by the controllers increased across the study's different time-frames, in correspondence with the increasing levels of traffic associated with each time-frame. Figure 3 illustrates these counts, which also show the number of conflicts worked by the different sectors. Sectors ZOB59 and ZOB79 experienced similarly high numbers of conflicts: working 291 (27%) and 265 (25%) of all analyzed conflicts, respectively. By comparison, sectors ZOB38 and ZOB49 worked 166 (15%) and 141 (13%) of the analyzed conflicts.

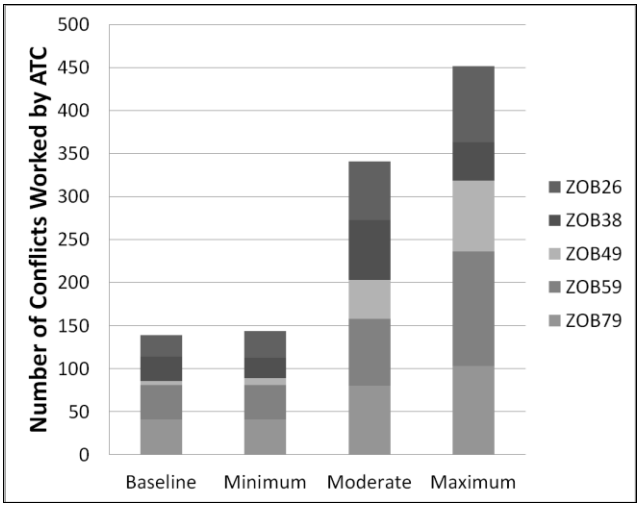

Figure 3. Conflict counts across the study's four conditions.

Time until predicted LOS

At the moment a conflict was first detected, the amount of time remaining until the automation predicted the conflicting aircraft would violate the separation minima serves as a measure of conflict urgency. A conflict initially detected with little time remaining would be more urgent than a conflict first detected with still many minutes-to-go. For the set of analyzed conflicts, the distribution of their initial time until predicted LOS is shown in Figure 4. The top two portions of the figure show the distributions for the Baseline and Minimum NextGen time-frames, and share a common y-axis scale for comparison. The bottom two portions of the figure have a different y-axis scale, showing the distributions of the Moderate and Maximum NextGen time-frames. The data clearly shows conflicts occurring across the full range of detection look-ahead times and, as expected, a strong skew towards conflicts detected in the earlier, strategic time-horizon. For each of the respective conditions, tactical conflicts (detected at three or less minutes-to-go) made up only 8%, 13%, 13%, and 4% of all conflicts, strategic conflicts (detected between three and eight minutes-to-go) accounted for 32%, 40%, 28%, and 31% of all conflicts, and conflicts detected with more than eight minutes-to-go made up 60%, 47%, 59%, and 66% (the majority) of all conflicts. The Moderate NextGen time-frame saw the highest number of short-term conflicts, offering support to the challenges and difficulties unique to that condition, theorized in [5] and [6].

Predicted Vertical State

Another way to describe a conflict is by examining the aircraft involved. Their predicted vertical states, to some extent, provide a simple proxy for conflict complexity. For an aircraft in conflict, its predicted vertical state represents the automation's belief in what the aircraft will be doing at the moment the loss of separation occurs. An analysis of this data characterized the vertical states as either: level flight, climbing, or descending. Since each conflict involves two aircraft, the analysis categorized the data according to the pairing of the aircrafts' respective vertical states. Illustrated in Figure 5, 1029 (96%) of the analyzed conflicts

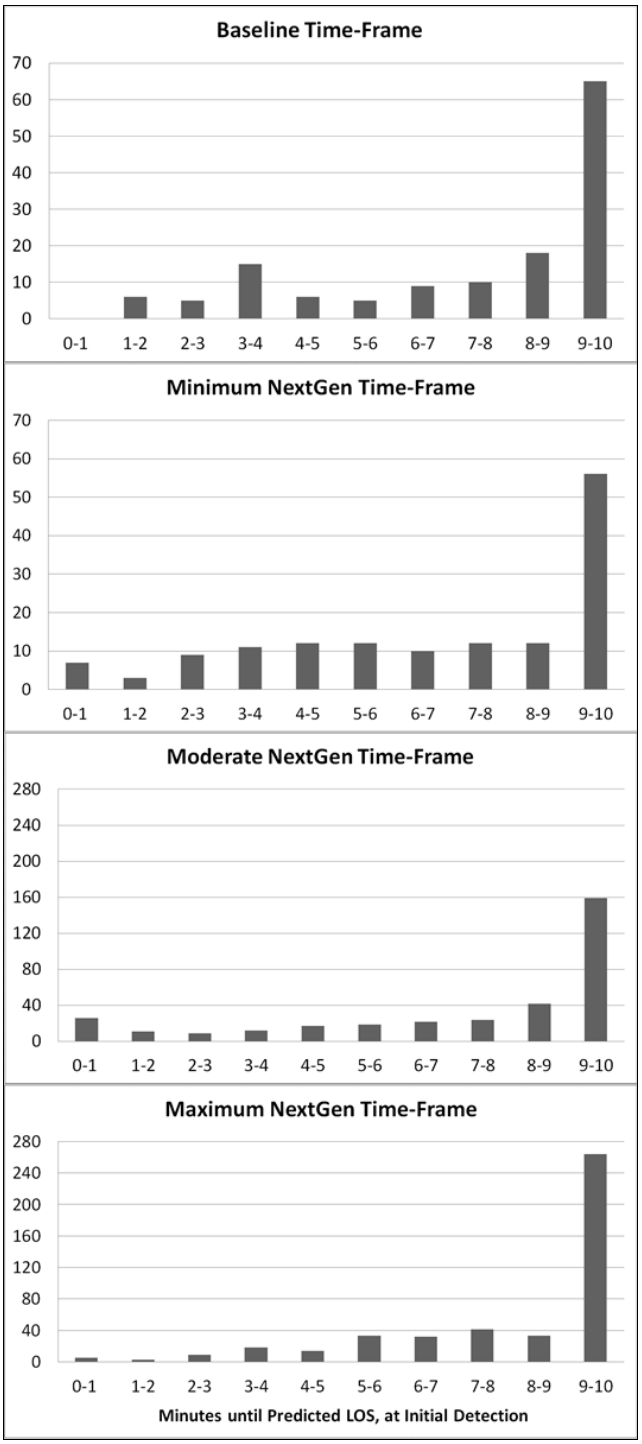

Figure 4. Histograms of conflict time-to-LOS, as measured at time of initial detection.

involved at least one level-flight aircraft, and in 625 (58%) of the conflicts, both aircraft were in level flight. Also, in the Baseline and Minimum NextGen

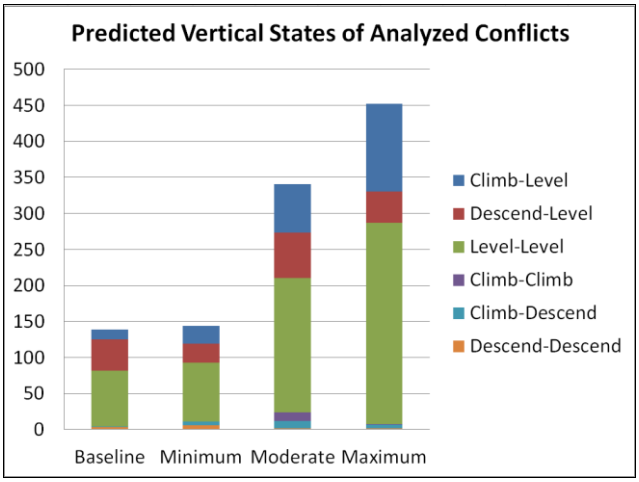

Figure 5. Counts of conflict vertical-state pairings.

time-frames, the data shows more conflicts involving descending aircraft than climbing aircraft, while the opposite is seen in the Moderate and Maximum NextGen time-frames.

Equipage Mix

Understanding which aircraft were involved in the analyzed conflicts also included an examination of the aircraft equipage levels. In the Baseline and Maximum NextGen time-frames, aircraft equipage level was not a factor since the traffic scenarios in those conditions were homogenous: all aircraft were unequipped in the Baseline time-frame, and all aircraft were equipped in the Maximum NextGen time-frame. However, in the Minimum and Moderate NextGen time-frames there was a mixture of aircraft equipage levels, and the implications of being equipped were different between the two time-frames. Equipped aircraft in the Minimum NextGen time-frame were automatically cleared to follow their FMS-computed vertical profile, and received transfer of communication instructions via Data Comm. In the Moderate NextGen time-frame, equipped aircraft could receive trajectory-change instructions via Data Comm as well.

The data shows that in conditions where mixed-equipage conflicts were possible, such conflicts played the biggest role in the overall conflict picture experienced by controllers. In the Minimum NextGen time-fame, 68 conflicts (47%) were between an unequipped and equipped aircraft, while 170 such encounters represented 50% of all conflicts in the Moderate NextGen time-frame (see Figure 6).

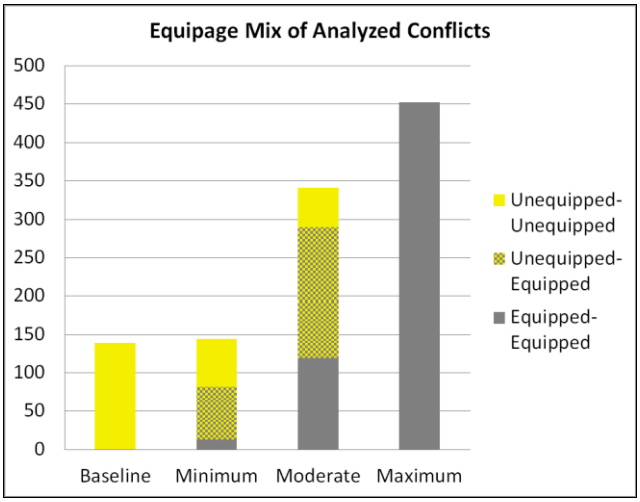

Figure 6. Counts of aircraft equipage pairings.

Conflict-Resolution Actions

Having described the number and types of conflicts analyzed, we now have the context in which to examine the actions taken by the sector team (i.e., the controllers and the automation). Data regarding how the sector team addressed the conflicts can provide another perspective on the interactions between the controllers and the automation. The four NextGen time-frames allow for comparisons between how the controller (with or without tools) and the automation (when working on its own) resolved conflicts.

When Did the Resolution Occur?

The first relevant conflict-resolution action logged during the conflict window helped to define when the sector team began to address the conflict. The time of that first action, compared to the conflict's initial detection time, describe one element of resolution implementation: timing. Submitted as a proxy for response time, Figure 7 shows the distributions of when the sector teams started working a conflict. Presented as the number of elapsed seconds between 'initial detection' and 'first action', the data's leftward skew suggests a preference by the controllers to work the conflicts as early as possible, perhaps helping them to more easily avoid the need for larger maneuvers. Interestingly, the skew appears increasingly pronounced over the progression of the different conditions. One possible explanation is that as the level of automation increased, and more tasks were delegated to the automation, controllers were more able to focus on the task of resolving conflicts.

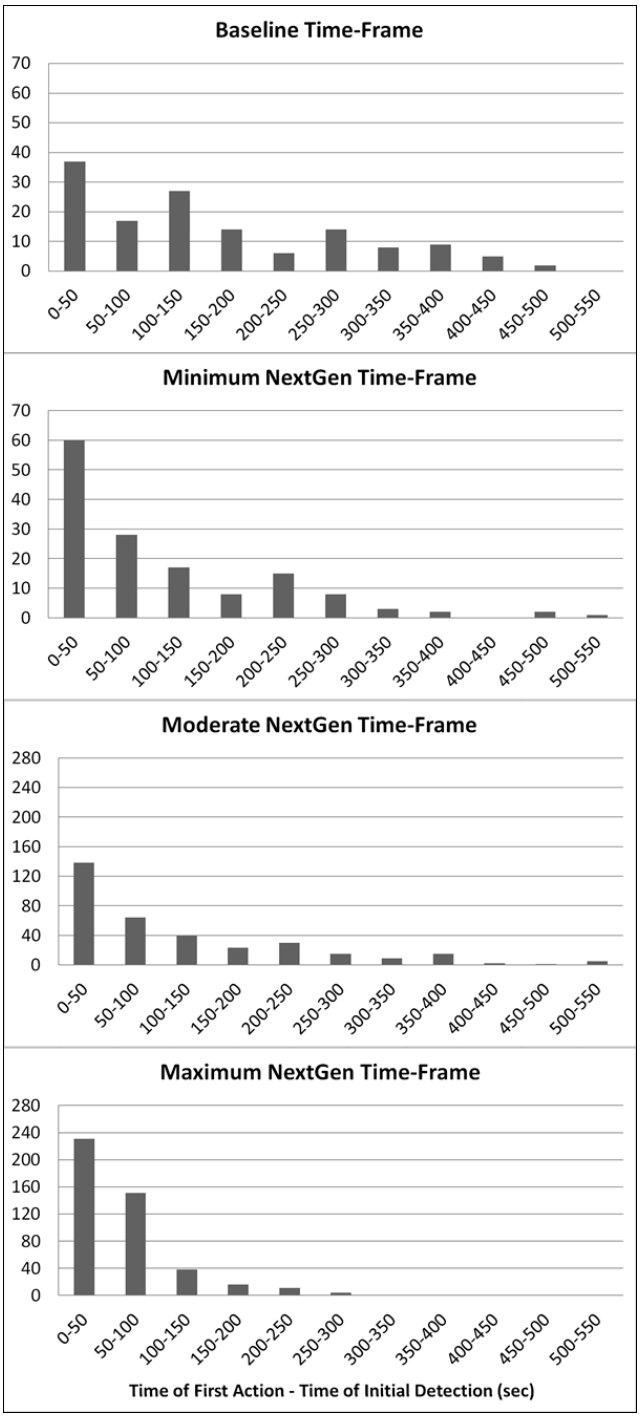

Figure 7. Histograms of conflict resolution initiation, measured as seconds since the initial conflict detection.

At this point, it is important to note that when describing conflict-resolution action data, the Maximum NextGen time-frame requires special consideration. This is because during that condition, independent of the controller, the automation was also resolving conflicts. The bottom portion of Figure 7, in other words, depicts controller- and automation-generated data together. This makes identifying any trends in the controllers' data difficult when comparing across conditions. To that end, Figure 8 separates the controller-generated data from the automation's data, where we see a curious discrepancy: the controllers' data continues with the strong skew towards shorter response times, while the automation appears content to often wait twice as long after a conflict's initial detection before addressing it.

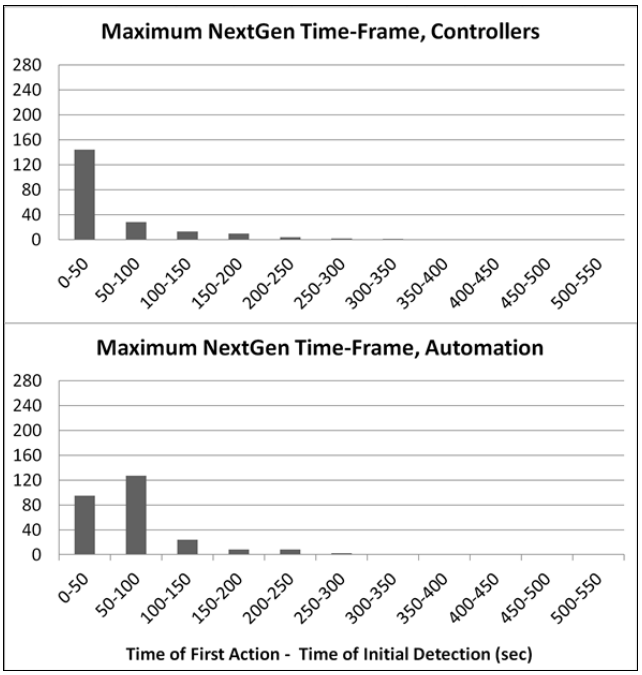

Figure 8. Timing data for conflict-resolution initiation, separated by agent.

A third aspect of timing is easily examined by simply taking the product of Figures 4 and 7. The result characterizes the first conflict-resolution action in terms of the conflict's urgency at the time of that action. Together with Figures 4 and 7, the data seen in Figure 9 provides a more complete understanding of when the resolutions occurred. For example, data from the Maximum NextGen time-frame show a peak

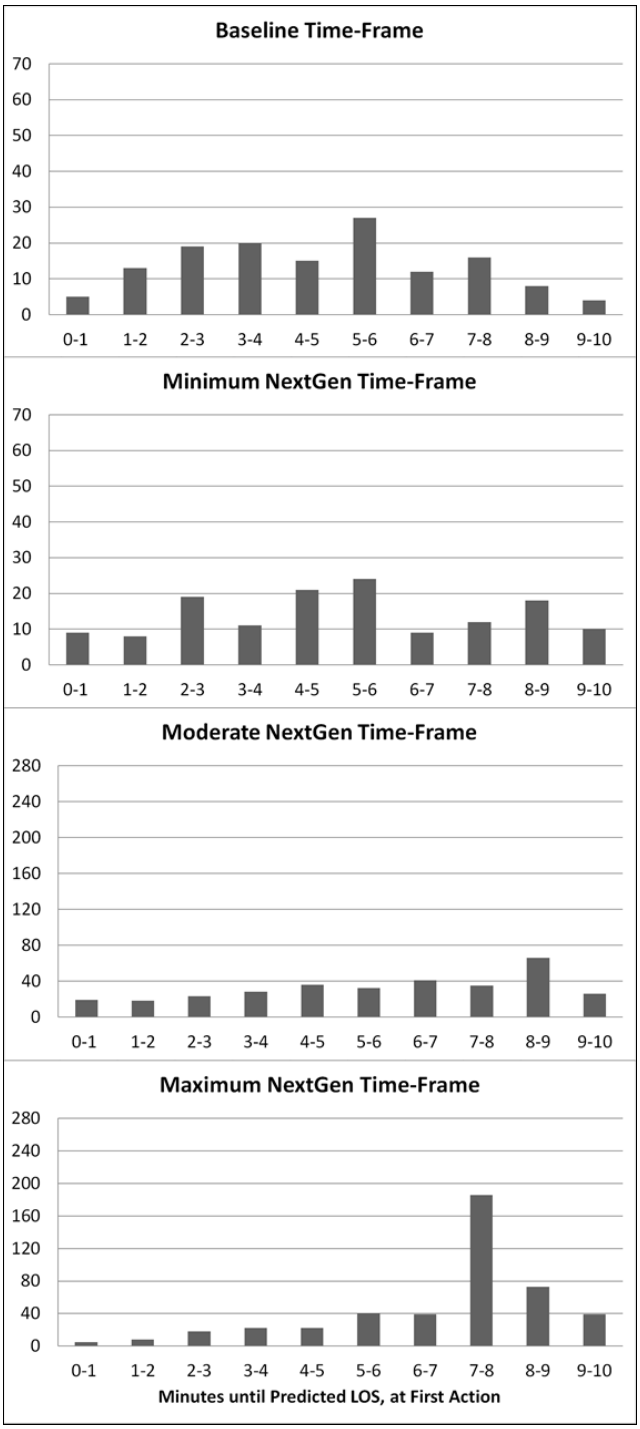

Figure 9. Histograms of conflict time-to-LOS, as measured at time of first action.

in the 7-8 minute bin, which could explain the slightly longer response times in Figure 7. However, just as with Figure 7, Figure 9's Maximum NextGen

data benefits from being split between controller and automation-generated data. A clearer picture then is shown in Figure 10, offering an explanation regarding the condition's longer conflict response times. The data confirms it is not the controllers who contribute to the issue; indeed, the upper portion of Figure 10, when compared to the Moderate NextGen data in Figure 9, suggests they addressed conflicts similarly, if not earlier. The peak seen in the lower portion of Figure 10 clearly attributes the longer conflict response times in the overall picture of the Maximum NextGen time-frame to the manner in which the automation responded to conflicts.

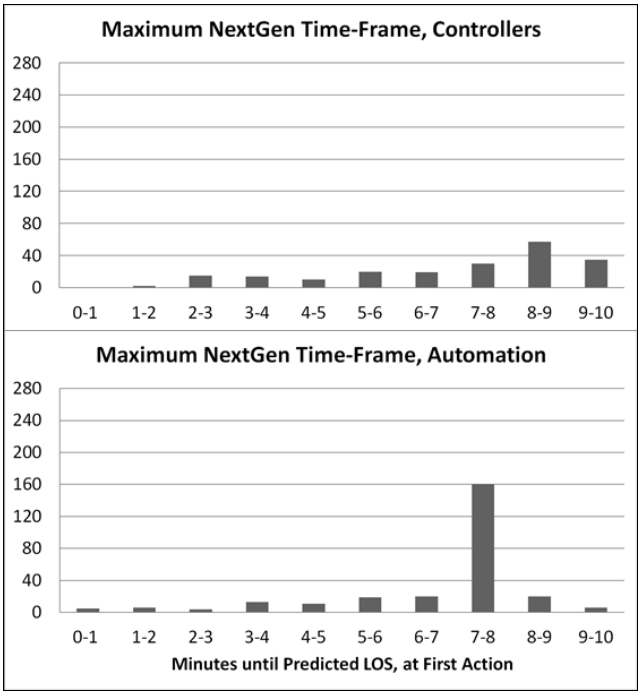

Figure 10. Distributions of time-to-LOS data measured at first action, separated by agent.

Which Aircraft Resolved the Conflict?

For a given pair of conflicting aircraft, the situation's resolution can involve one aircraft, the other aircraft, or both. In determining which aircraft to maneuver, controllers consider many factors, such as the route of flight, or an aircraft's 'maneuverability': where within its performance envelope is an aircraft currently operating, if the aircraft is or expects to soon be transitioning to a new altitude, etc. Analyses of the SA5 data examined two such factors in order to gain insights into the conflict-resolution strategies employed by the controllers: the predicted vertical states of conflicting aircraft, and their level of equipage. These analyses are described in the meaningful context of the aircraft pair rather than globally across an entire condition, helping interpretations from becoming too abstract.

For the six possible pairings of vertical-state encounters, there existed four outcomes identifying the maneuvering aircraft: a climbing aircraft, a descending aircraft, a level-flight aircraft, or both aircraft. Figure 11 ambitiously attempts to show this data across each of the NextGen time-frames, while splitting the Maximum condition's results between controller- and automation-generated data.

As highlighted earlier in Figure 5, Figure 11 reiterates that very few conflicts were without an aircraft predicted to be in level-flight. The sparse data in the right half of Figure 11 provides just one insight: in the few conflicts involving a climbing aircraft vs. a descending aircraft, the controller preferred to move the descending aircraft. It was the controller who addressed 18 of 21 such conflicts; on 12 occasions (57%) by moving the descending aircraft, while electing to move the climbing aircraft three times (14%). For the three remaining occasions, the controller moved both aircraft. On its own, the automation only addressed three of these conflicts, and although not nearly enough data to draw any conclusions from, always chose to move the climbing aircraft.

When a conflict involved two level-flight aircraft, maneuvers to both aircraft were infrequent. In the Baseline, Minimum, and Moderate conditions, such resolutions occurred less than 5% of the time. In the Maximum NextGen time-frame, controllers moved both aircraft in 16% of such conflicts, while the automation did so 11% of the time. For conflicts involving a level-flight aircraft and either a descending or climbing aircraft (the two left-most portions of Figure 11), maneuvers involving both aircraft were again infrequent, but more importantly, one trend emerges: the sector team more often moved the transitioning aircraft. The only minor exception to this is seen in the automation's data for climbing vs. level conflicts, where the automation moved the climbing aircraft equally as often as the level-flight aircraft. In all other cases, the resolution to maneuver only the transitioning aircraft represented the majority of the data.

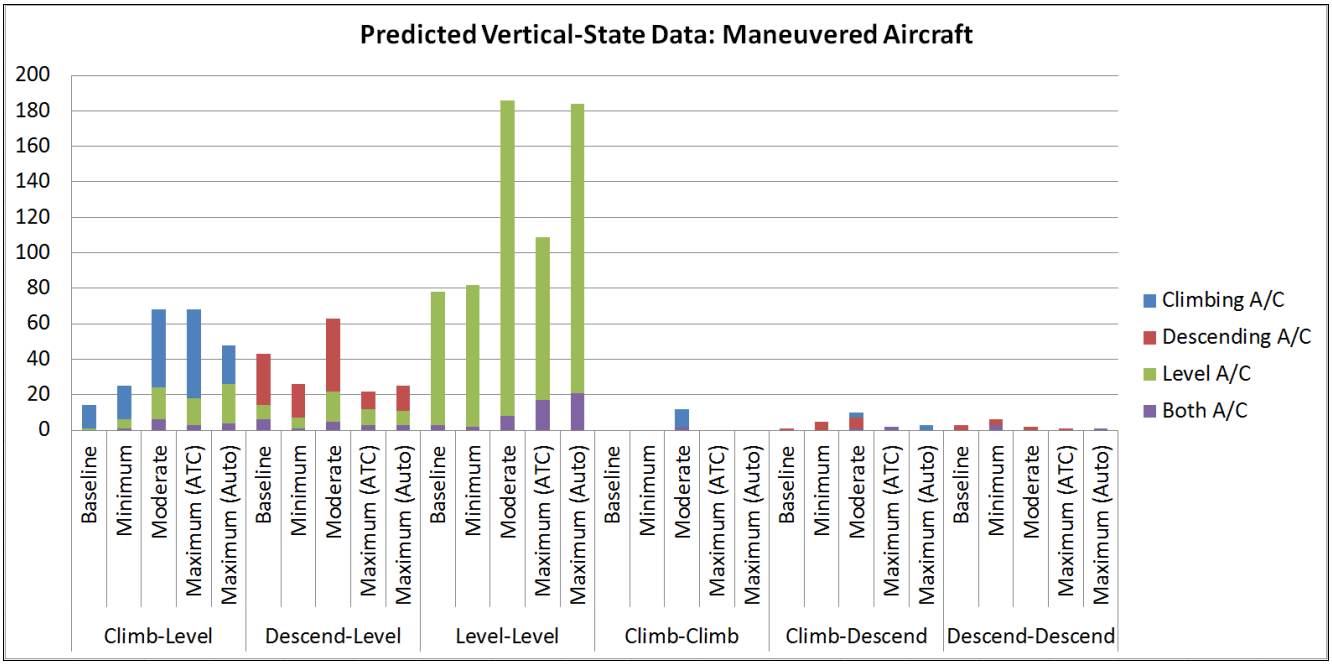

Figure 11. Analysis of which aircraft maneuvered as part of the conflict's resolution, organized across all combinations of vertical-state pairings and study conditions.

Another factor in the decision of which aircraft to maneuver is the aircraft's equipage level. An analysis of the data explored this issue, but did so only for the Minimum and Moderate NextGen time-frames. One reason for the more focused scope is that in the Baseline and Maximum NextGen time-frames, the traffic was completely homogenous, leaving little mystery in identifying the equipage of the maneuvered aircraft: the Baseline condition saw only unequipped-vs.-unequipped conflicts, while the Maximum condition saw only equipped-vs.-equipped conflicts. Secondly, the area in which data from these two conditions could prove insightful, whether the sector team maneuvered one or both aircraft, again speaks only to a small minority of the data. Thus, for all unequipped-vs.-equipped conflicts occurring in the Minimum and Moderate NextGen time-frames, Figure 12 illustrates the results regarding which aircraft received the maneuver.

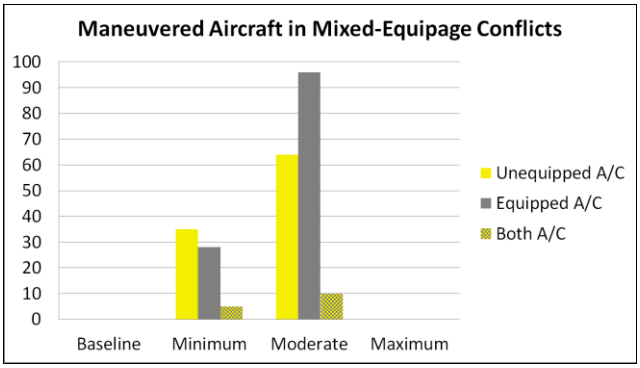

Figure 12. Counts of which aircraft controllers maneuvered to resolve mixed-equipage conflicts.

In the context of aircraft equipage level, it was possible for the sector team to reach one of three maneuver outcomes: resolving the conflict using the unequipped aircraft, the equipped aircraft, or both. In the Minimum NextGen time-frame, the sector team resolved 51% of these conflicts by moving the unequipped aircraft, while 41% of their resolutions involved the equipped aircraft. Interestingly, the preference to maneuver the unequipped aircraft is contradicted in the Moderate NextGen time-frame, where the majority of resolutions (56%) involved the equipped aircraft, and the sector team used the unequipped aircraft in only 38% of their resolutions.

Which Type of Maneuver Resolved the Conflict?

In resolving a conflict, the sector team must issue a maneuver instruction to one or both aircraft (the latter, we know, rarely occurred). In simplified terms, such resolution clearances can move the

aircraft in vertical, lateral, and/or longitudinal directions, as accomplished, respectively, through altitude, heading or route, and speed clearances. An analysis of the conflict-resolution actions taken by the sector team categorized the data according to clearance type, yielding four groupings: Altitude, Route (which included both heading and route instructions), Speed, and Multi. This last category helped to capture instances in which the sector team resolved a conflict using more than one type of clearance.

Figure 13 provides a high-level look at the distribution of clearance types across the four NextGen time-frames. In large part, the data reflects the increase in the number of conflicts across the different conditions, but indicates that the proportion of clearances consisting of route instructions did increase over the progression of the four NextGen time-frames. Of more importance are the altitude and route categories, whose contributions across the four conditions were as follows: Altitude – 83%, 82%, 61%, and 19%; Route – 10%, 7%, 32%, 75%. Presenting the data in relation to pair-wise encounters, should improve our understanding of this seemingly inverse relationship. An analysis of the maneuver type data examined conflict resolution actions according to the two pairings used earlier: predicted vertical state, and aircraft equipage.

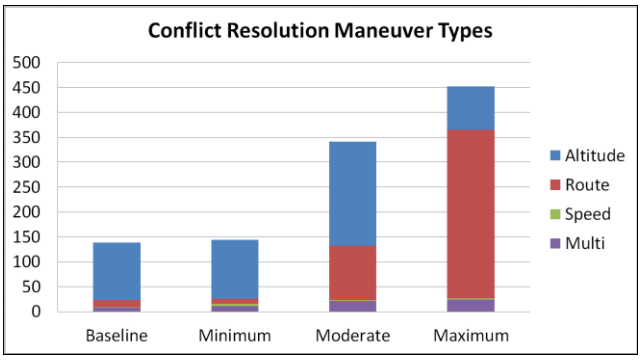

Figure 13. Global counts of maneuver types implemented during conflict resolution.

For each of the six possible vertical-state conflict pairings, Figure 14 illustrates the type of maneuver chosen by the sector team to resolve the conflict. The y-axis is scaled to 100%, but the actual counts of altitude and route clearances are shown as overlays. In addition to minimizing clutter, counts for the other maneuver types are not shown because of their relatively low frequency and minimal influence on this analysis.

Again, most conflicts involved at least one level-flight aircraft, as evidenced by the low numbers in the bottom half of the Figure. Focusing then, on the upper half of Figure 14, the data appears to display three principle trends. First, speed clearances, as well as resolutions consisting of multiple maneuver types, contributed only small amounts to the overall activity and manner in which the sector team resolved conflicts. For example, the 'high points' in this data come from the 13 multi-type clearances issued in response to Level-Level conflicts during the Moderate NextGen time-frame, and the controllers' four multi-type clearances issued in response to Descend-Level conflicts during the Maximum NextGen time-frame. Those data points translate to 7% and 13% of the total number of resolutions to the conflicts with those particular vertical-state pairings. Secondly, when comparing only the controllers' actions between the four NextGen time-frames, they seem to gradually trade an increasing number of altitude maneuvers for route maneuvers. By the Maximum NextGen time-frame, route maneuvers become the majority of their resolutions, whereas in the first three conditions, they issued more altitude maneuvers than any other type of maneuver. An imaginary diagonal line, from bottom-left to upper-right, closely approximates this observation. While similar to the data presented in Figure 13, it is encouraging that the trend is still observable after removing two extraneous factors: the automation's actions and all vertical-state pairings without level-flight aircraft.

The data representing the automation's own conflict-resolution actions helps uncover the third finding. While it appears as if the automation's resolutions align with the controller's trend to use increasingly larger proportions of route maneuvers and increasingly smaller proportions of altitude clearances, automation-generated data is not available from the other conditions. Without knowing how the automation would have handled the other conditions, such conclusions can be misleading.

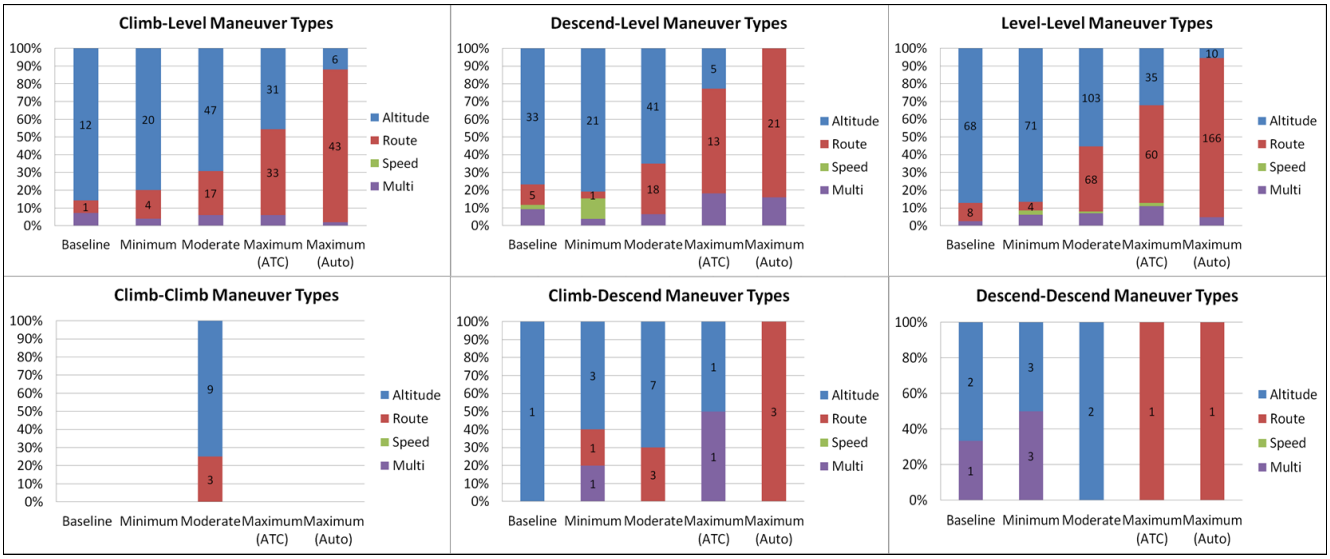

Figure 14. Analysis of which maneuver type the sector teams used to resolve conflicts, organized across all combinations of vertical-state pairings and study conditions.

What is interesting however, is comparing the automation's and controllers' approaches to resolving conflicts in the maximum NextGen time-frame. For Climb-Level, Descend-Level, and Level-Level conflicts, the controllers maintained a sizeable use of altitude maneuvers, nearly 25% or more. On the other hand, the automation's data suggest a different strategy; one that (seemingly in all vertical-state pairings), simply uses route maneuvers as much as possible.

Another analysis of the resolution maneuvers issued by the sector team examined the data according to the three possible equipage pairings, investigating whether an aircraft's equipage had any bearing on how the sector team addressed the conflict. Figure 15 depicts the distributions of the employed maneuver types, for each of the NextGen time-frames. As with the previous analysis, we first examine the controller-generated data across the four conditions, and then compare the controllers' and automation's data in the Maximum NextGen time-frame.

It is important to note that the actual counts of each maneuver type, when compared across equipage pairing, are subject to the traffic mixes of the scenarios in the different NextGen time-frames. Observations of the relative proportions however, highlight several elements. For conflicts involving two unequipped aircraft (seen in the upper portion of Figure 15), the controllers' approach remained fairly consistent, primarily relying on altitude clearances. The middle portion of the Figure shows the data for Mixed-equipage conflicts, those between an unequipped aircraft and an equipped aircraft. The data for such conflicts during the Minimum NextGen time-frame show a similar pattern to the conflicts involving two unequipped aircraft. Mixed-equipage conflicts in the Moderate condition bring an apparent change in strategy, where the controllers issued a larger proportion of route maneuvers (in this case, 26%). The trend to more often issue route maneuvers is also seen in conflicts involving two equipped aircraft. The bottom portion of Figure 15 shows an increase in the controllers' proportion of route maneuvers, continuing across the Minimum, Moderate, and Maximum NextGen time-frames, similar to that seen in Figure 13 and the upper half of Figure 14. Focusing on just the Maximum NextGen time-frame, when comparing the controller- and automation-generated data, the difference between the two distributions point to the same observations made with the vertical-state data: the automation issued relatively few altitude maneuvers.

Discussion

Results presented here have identified several issues regarding the sector team's strategy for resolving conflicts. Timing data suggests a

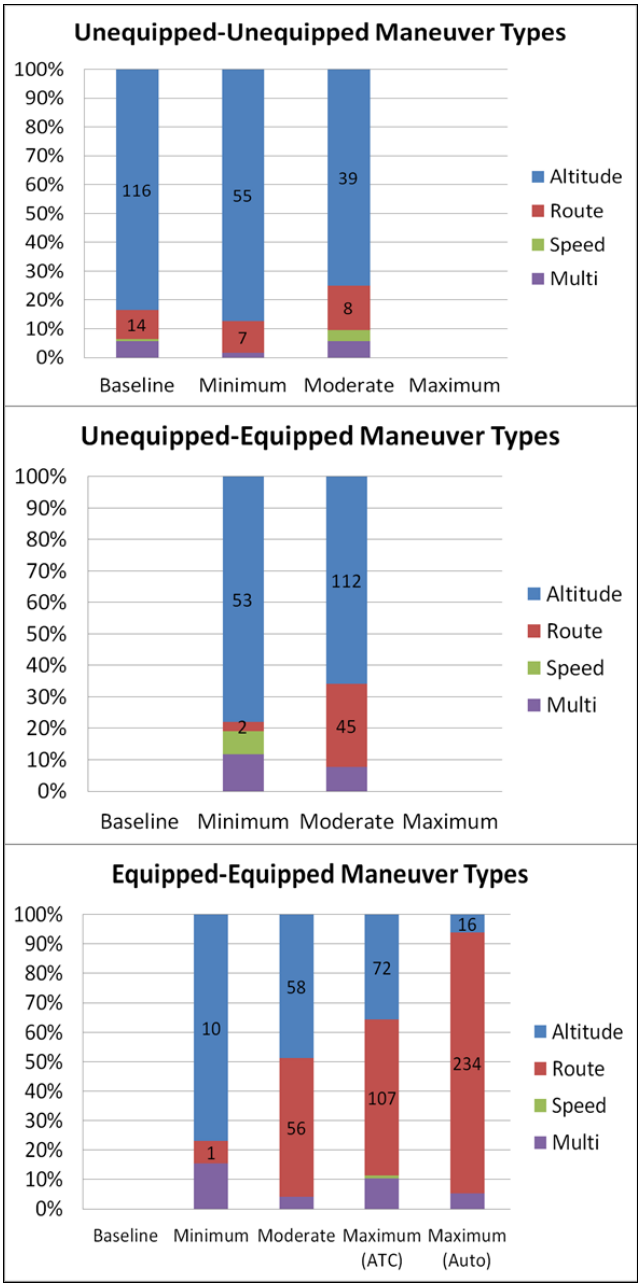

Figure 15. Proportions of conflict-resolution maneuver types, per aircraft-equipage pairing.

relationship between time of first action and NextGen time-frame. The plausibility of this relationship is supported by two theories. The progression of NextGen time-frames included an increase in the level of available automation and amount of tasks delegated to the automation. This perhaps had a positive effect on the controller's availability, resulting in quicker responses to the conflicts. The different NextGen time-frames also brought increasing levels of traffic. The controllers may have anticipated an increase in workload accompanying these traffic levels, and adapted their working strategy as a result. The controllers' theorized efforts to 'stay ahead of the problem' require careful consideration, since it is unknown if their self-reported workload data (see references [3-4]), would capture the anticipated workload, or the workload resulting after their coping strategy took effect. Another timing issue discovered in the data is the difference between the controllers' time of first action, and the automation's. While controllers acted more quickly than the automation, there is no evidence that this discrepancy is bad; it just exists. The automation simply followed settings which specified that it should not respond to conflicts any earlier than eight minutes until the predicted LOS. Further research is needed to explore the issues surrounding this form of implicit coordination.

Data identifying the aircraft used by the sector team to resolve a conflict provides two clear results: both the controllers and the automation rarely maneuvered both aircraft involved in a conflict, and in conflicts between level and transitioning (climbing or descending) aircraft, the controllers more often moved the transitioning aircraft. The latter finding can help improve the conflict-resolution algorithm, so that when used interactively by the controller, it can provide resolution suggestions that the controllers may more likely accept. Aircraft equipage played another role in determining which aircraft the controllers maneuvered in order to resolve a conflict, but did so in combination with the present NextGen time-frame. More specifically, the Moderate NextGen time-frame, whose traffic scenarios consisted of a 50-50 split between unequipped and equipped aircraft, also introduced the ability to send trajectory changes via Data Comm; an aspect which has significant implications for the conflict-resolution task. For mixed-equipage conflicts in the Moderate NextGen time-frame, controllers moved the equipped aircraft 30% more often than the unequipped aircraft, likely due to this interaction between aircraft equipage and the ability to send trajectory changes via Data Comm.

The type of maneuver chosen by the sector team to resolve a conflict also shows strong ties to the ability to send trajectory changes via Data Comm. Controllers issued increasing proportions of route

clearances, eventually over-taking altitude clearances in the Maximum NextGen time-frame. In a voice environment, route clearances can be quite laborious, perhaps explaining why controllers prefer to use altitude maneuvers. Just between the Minimum and Moderate conditions, route maneuver usage jumped from just under 10% to just over 30%, and by the Maximum NextGen time-frame it was 75%. In the Maximum condition, the preference for route maneuvers over altitude is perhaps explained by the high traffic levels associated with that condition. One possibility is that reducing their use of altitude maneuvers helped the controllers to keep aircraft on their 'right-for-direction' altitudes, thereby minimizing the complexity of their sector. In general, the automation rarely issued an altitude maneuver. While this may be less of an issue in Level-Level conflicts, improving the automation so that it can better utilize altitude, for example to start down early an aircraft nearing its Top of Descent, could prove beneficial. It is understood that 'teaching' automation to effectively use altitude is not without its challenges. Controllers can easily issue temporary altitudes, with the plan of later clearing aircraft to their requested altitude. An example of this would be to stop an aircraft's climb until a conflicting aircraft overhead passes by, leaving a clear space for the climbing aircraft to then reach its desired altitude. However, such logic can be difficult for automation, since step-climbs/descents are not easily sent via Data Comm (i.e., not easily read by an aircraft's FMS). Additionally, helping aircraft (who are not in conflict) reach their requested altitude is currently not something the automation addresses, but perhaps could. A simpler way to increase the automation's use of altitude maneuvers is to expand the TSAFE algorithm to include tactical vertical maneuvers.

Lastly, while a conflict's vertical-state pairing had little impact on the type of issued maneuver, aircraft equipage did. For conflicts between two unequipped aircraft, their lack of equipage kept the proportions of altitude and route maneuvers fairly consistent. In contrast, conflicts involving equipped aircraft are where we see large increases in the use of route maneuvers. Again, this is likely correlated with the ability to send trajectory changes via Data Comm in the Moderate and Maximum NextGen time-frames.

Conclusion

The results presented in this paper uncovered two discrepancies between the controllers' and the automation's conflict-resolution strategy: timing of first action, and the use of altitude maneuvers. Understanding the implications of the former requires more investigation, but the latter issue is something that has the potential to cause disruptive interactions and ineffective team work between the controller and the automation.

The analyses also revealed that of the examined aspects of a conflict, two factors most clearly impacted the controllers' resolution strategies, expressed in the data as the following observed patterns: when transitioning aircraft were involved, the controllers more often maneuvered the transitioning aircraft, and in NextGen time-frames where it was possible to send trajectory changes via Data Comm, the controllers more often maneuvered the equipped aircraft.

Further analyses of this data could provide additional insights by investigating, as opposed to a conflict resolution's end state, the process by which the sector team came to that action. Did they start down a particular path, but then switch to a different approach? Another more challenging exercise would be to trace the sector team's actions down many more 'and-then-what' layers. For example, analyzing the three aircraft-equipage pairings, for each of the six vertical-state pairings, for each of the 10 time-until-predicted-LOS bins, would provide a lucky individual with 180 cases to investigate.

Clearly, more research is still needed to better understand the conflict-resolution task's intricacies and contextual factors. If one assumes a future in which controllers resolve conflicts, understanding their strategies will go a long way towards building useful decision-support tools. If one assumes a future in which conflict resolution is the automation's responsibility, an understanding of controllers' strategies could help the automation resolve conflicts according to detailed 'context-dependent' criteria, thereby approximating the heuristics today's controllers typically develop.

References

[1] Joint Planning and Development Office. (2010). Concept of Operations for the Next Generation Air Transportation System, version 3.2. retrieved March 12 , 2014 from http://jpe.jpdo.gov/ee/request/folder?id=35434.

[2] Prevôt, T., Smith, N., Palmer, E., Callantine, T., Lee, P., Mercer, J., et al. (2014). *An Overview of Current Capabilities and Research Activities in the Airspace Operations Laboratory at NASA Ames Research Center*. AIAA-2014-2860, American Institute of Aeronautics and Astronautics, Reston, VA.

[3] Wing, D., Prevôt, T., Lewis, T., Martin, L., Johnson, S., Cabrall, C., et al. (2013). Pilot and Controller Evaluations of Separation Function Allocation in Air Traffic Management. *Proceedings of the USA/Europe Air Traffic Management Research and Development Seminar*, Chicago, IL.

[4] Mercer, J., Homola, J., Cabrall, C., Martin, L., Morey, S., Gomez, A., et al. (2014) Human-Automation Cooperation for Separation Assurance in Future NextGen Environments. *Proceedings of the HCI-Aero Conference*, Santa Clara, CA.

[5] Homola, J., Morey, S., Cabrall, C., Martin, L., Mercer, J., and Prevôt, T. (2013). *Analysis of Interactive Conflict Resolution Tool Usage in a Mixed Equipage Environment*. AIAA-2013-4700, American Institute of Aeronautics and Astronautics, Reston, VA.

[6] Gomez, A., Martin, L., Homola, J., Morey, S., Cabrall, C., Mercer, J., et al. (2013).Investigating the Complexity of Transitioning Separation Assurance Tools in NextGen Air Traffic Control. *Proceedings of the Digital Avionics Systems Conference*, Syracuse, N.Y.

[7] Homola, J., Martin, L., Mercer, J., and Prevôt, T., (2014). Exploring Workload Factors Across Future Environments. *Proceedings of the HCI-Aero Conference*, Santa Clara, CA.

[8] Cabrall, C., Gomez, A., Homola, J., Hunt, S., Martin, L., Mercer, J., et al. (2013). Transitioning Resolution Responsibility between the Controller and Automation Team in Simulated NextGen Separation Assurance. *Proceedings of the ENRI International Workshop on ATM/CNS*, Tokyo, Japan.

[9] Hunt, S., Homola, J., Mercer, J., and Cabrall, C. (2014, in press). Controller Inhibition of Automated Conflict Resolutions in a Maximum NextGen Condition. *Proceedings of the Digital Avionics Systems Conference*, Colorado Springs, CO.

[10] Erzberger, H. (2001). The Automated Airspace Concept. *Proceedings of the USA/Europe Air Traffic Management R&D Seminar*, Santa Fe, N.M.