Transitioning Resolution Responsibility Between the Controller and Automation Team in Simulated NextGen Separation Assurance

C.D. Cabrall, A.N. Gomez, J.R. Homola, S.M. Hunt, L. Martin, J. Mercer, and T. Prevot

Abstract As part of an ongoing research effort on separation assurance and functional allocation in NextGen, a controller-in-the-loop study with ground-based automation was conducted at NASA Ames' Airspace Operations Laboratory in August 2012 to investigate the potential impact of introducing self-separating aircraft in progressively advanced NextGen time-frames. From this larger study, the current exploratory analysis of controller–automation interaction styles focuses on the last and most far-term time frame. Measurements were recorded that firstly verified the continued operational validity of this iteration of the ground-based functional allocation automation concept in forecast traffic densities up to two times that of current day high altitude en-route sectors. Additionally, with greater levels of fully automated conflict detection and resolution as well as the introduction of intervention functionality, objective and subjective analyses showed a range of passive to active controller–automation interaction styles between the participants. Not only did the controllers work with the automation to meet their safety and capacity goals in the simulated future NextGen timeframe, they did so in different ways and with different attitudes of trust/use of the automation. Taken as a whole, the results showed that the prototyped controller–automation functional allocation framework was very flexible and successful overall.

C.D. Cabrall (✉) • A.N. Gomez • J.R. Homola • S.M. Hunt • L. Martin • J. Mercer
San Jose State University, NASA Ames Research Center, Moffett Field, CA, USA
e-mail: christopher.d.cabrall@nasa.gov; ashley.n.gomez@nasa.gov; jeffrey.r.homola@nasa.gov; sarah.m.hunt@nasa.gov; lynne.martin@nasa.gov; joey.mercer@nasa.gov

T. Prevot
NASA Ames Research Center, Moffett Field, CA, USA
e-mail: thomas.prevot@nasa.gov

1 Introduction

In the present day high altitude en-route environment of the United States National Airspace System, aircraft separation assurance is achieved by a highly labor intensive process of dutiful air traffic controllers on the ground. Monitoring the progress of aircraft across their display, the controllers scan their sector by watching each and every aircraft in order to identify potential separation risks and to mentally calculate conflict avoidance possibilities. From the time any aircraft checks in with one controller until it is handed off to the next, all clearances are devised mentally, manually, and individually by the controller and issued verbally over a radio frequency. With great scrutiny, attention, and positive personal control over each of the aircraft in their sector, en-route controllers have maintained a commendable safety record and contribute greatly to the overall US air traffic control system being the safest in the world [1].

Current forecasts by the Federal Aviation Administration (FAA) show continued growth in demand, in particular for the en-route centers, due to a faster growing commercial sector. The number of commercial aircraft is projected to grow from 2011 to 2032 at an average growth rate of about 1.5% or 127 aircraft annually. Similarly up to 2032, commercial IFR aircraft handled at FAA en-route centers has been projected to increase 2.4% annually [2]. These forecasts pose a problem for en-route controllers because they exceed monitor alert parameters (MAP) values which have been defined to limit the number of aircraft permitted in a sector as a safeguard prior to which performance is expected to decline. Natural cognitive processing limits of air traffic controllers have been accepted as potential bottlenecks against rising air traffic demand on account of the number of planes any person could reasonably be expected to track. Complementary to the FAA forecasts, the US Congress established the Joint Planning and Development Office (JPDO) to develop the Next Generation Air Transportation System (NextGen) which among its many visions, explicitly aims to overcome the capacity limits imposed by individually attended aircraft separation procedures of today and requires a restructuring of the roles of humans and automation and how they perform their respective functions to synergize human and automation performance [3].

To meet the forecast demand increase, separation management components of en-route NextGen environments are envisioned to rely on automation to augment human performance beyond today's limits by offloading workload from the human controllers onto automated functions for the majority of routine operations. Use of automated conflict detection and resolution decision aides that are seamlessly integrated within ground automation systems is planned to allow separation management tasks to move away from fixed human-based standards while always maintaining an unambiguous delegation of responsibility. Automation is anticipated to support a migration from tactical to strategic decision making as well as perform many routine tasks. With layers of protection that allow for graceful degradation of situations, automation reliance is planned to be coupled with modes that do not require full reliance on humans as backup. Building from today's current roles, the

corresponding NextGen roles for air traffic controllers that stand to benefit from use of automation include: identifying complex future conflicts, management of individual aircraft trajectories, and detecting and resolving conflicts via automation while eliminating residual conflicts [3]. Use of data communications that are integrated with ground automation is envisioned to reduce the number of voice communications and controller workload, and hence increase the controller's efficiency and ability to manage more traffic [4]. While providing tactical and strategic separation management, en-route trajectory based operations (TBO) automation is planned to provide the ability to request modifications of trajectories and support trajectory negotiation [5].

However, relying on automation to fully or partially replace a function previously carried out by a human operator means that automation need not be all or none, but can vary across a continuum of levels, from the lowest level of fully manual performance to the highest level of full automation. Furthermore, the specific function with its variant level of manual/automatic control, itself can range along a variety of human information processing sub-tasks or stages [6]. As a simple example, the function of detecting a conflict could be fully automatic, fully manual, or somewhere in between and this could exist along with different levels of automation for the subsequent separate function of conflict resolution, which could itself be fully automatic, fully manual, or somewhere in between. It has been shown that the flexibility of an automation system contributes to its use case as task load and complexity increase [7].

In addition to being flexible and multi-layered, other beneficial design factors can encourage effective trust and use within human–automation systems. Recent research suggests that humans respond socially to technology and reactions to computers can be similar to reactions to human collaborators [8]. In commonly observed effective human teamwork and collaboration, both parties walk a line to balance what they perceive the other is capable of while they, themselves, exhibit evidence of their own reliability in handling certain tasks. In general, someone's capability with simpler tasks is commonly held to reflect their capability with more complex tasks. Research on system credibility established within the context of simple decision tasks has been conducted and has shown that operators who experience an automated system's failures in easy tasks are less likely to comply with the automation's recommendations during a more difficult task [9]. Rather than taken on immediate face value, automation is scrutinized for its credibility through operator experiences with that automation, i.e. trust is learned. Considering that entire schools of cognitive theory have posited active participation as preferable over passive reception or observation, at least some sense of control is assumed to be of crucial importance for an operator to work with and appropriately trust automation. Automation surprises occur when technology autonomously performs tasks that cause a system to behave in a manner that the operator had not anticipated and it has been assumed that such a decrease in situation awareness arises from a non-satisfaction of a self-agency mechanism [10]. In other words, allowing the operator some form of control over the automation is expected to enhance the human–automation work dynamic.

2 Operational Concept

2.1 *Ground-Based Automated Separation Assurance*

Informed from the guidelines of the JPDO and human–automation functional allocation literature referenced above, our NextGen prototype instantiation of a ground-based automated separation assurance concept is next briefly described in this section. More detailed accounts of the precise characteristics and evolution of the concept can be found in the prior separation assurance (SA) research conducted at NASA Ames [11–18]. Additionally, complementary and collaborative airborne-based separation assurance concepts are detailed in research conducted at and with NASA Langley [19, 20].

Ground-based automated separation assurance involves automation components that monitor and/or manage nominal TBO equipped aircraft, while the controller handles off-nominal operations, provides additional services, and makes decisions when human involvement is needed. The primary difference from today's system is automated conflict detection and automated conflict resolution via data link. Controller involvement in routine conflicts is only required when an automatic trajectory change would exceed defined thresholds.

2.1.1 Enabling Environment

Each aircraft was assumed to be equipped with integrated data communications capabilities for route modifications, frequency changes, cruise altitude changes, and climb, cruise and descent speed modifications as well as high accuracy surveillance data provided via Automatic Dependent Surveillance Broadcast (ADS-B). Automated trajectory-based conflict resolutions were generated for conflicts with more than 3 min to initial loss of separation (LOS). For those with less time before LOS, a separate automated tactical conflict avoidance function (TSAFE) could generate a resolution and send heading changes to the aircraft directly.

2.1.2 Roles and Responsibilities

The automation detected conflicts, computed resolutions and/or alerted controllers. It nominally functioned by automatically sending instructions to aircraft via data link unless they exceeded a priori defined thresholds. The automation also augmented controller awareness and provided conflict status and probing tools. Primarily, the controller managed the automation, handled off-nominal situations and made decisions on situations when presented.
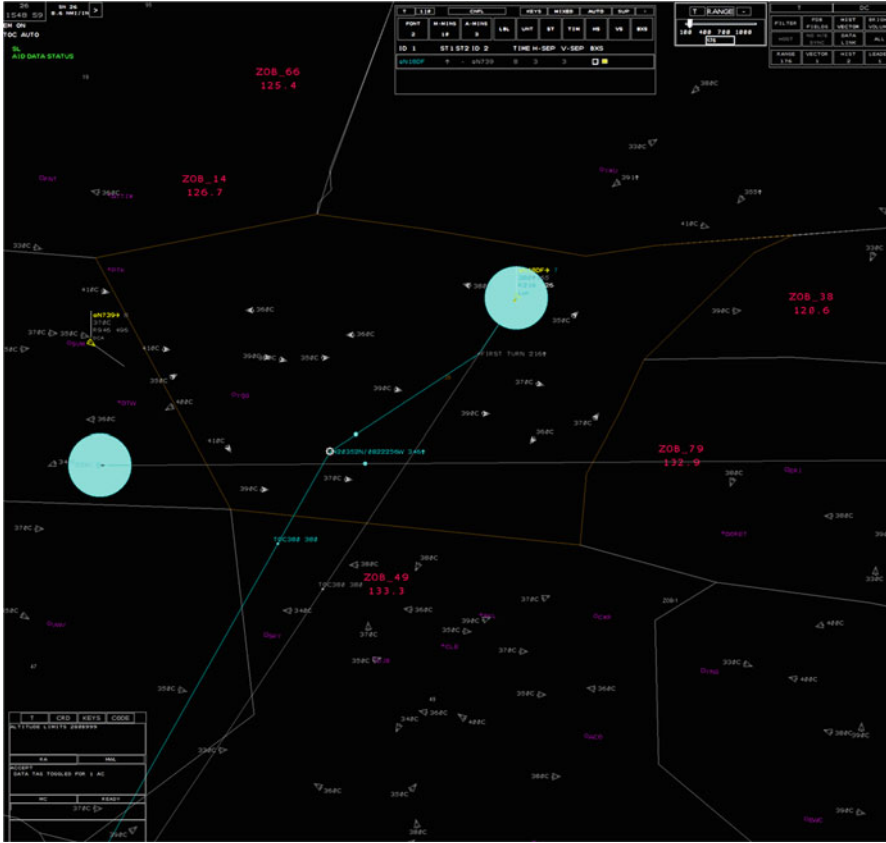
Fig. 1 Controller display with conflict list, an active conflict deferred by the automation to the controller (*yellow*), and a provisional resolution trajectory that currently conflicts with a third aircraft (*cyan*)

2.1.3 Air Traffic Controller Workstation

Figure 1 depicts the air traffic controller workstation prototype designed for the above distribution of roles and responsibilities. Aircraft that were managed by the automation and within the controller's sector are displayed in a brighter gray than low-lighted exterior aircraft. Additional information in data tags and colors were used to draw the controller's attention to a specific problem. The display was designed for general situation awareness and management by exception. The following figures present more detailed depictions of the various aspects of the interface controller's used to interact with the automation tools.

Nominally aircraft data tags were collapsed and appeared only as a chevron target with an altitude tag because routine aircraft operations such as frequency changes, hand-offs, climbs and descents, etc. were conducted automatically without controller involvement (Fig. 2).

Fig. 2 Collapsed aircraft
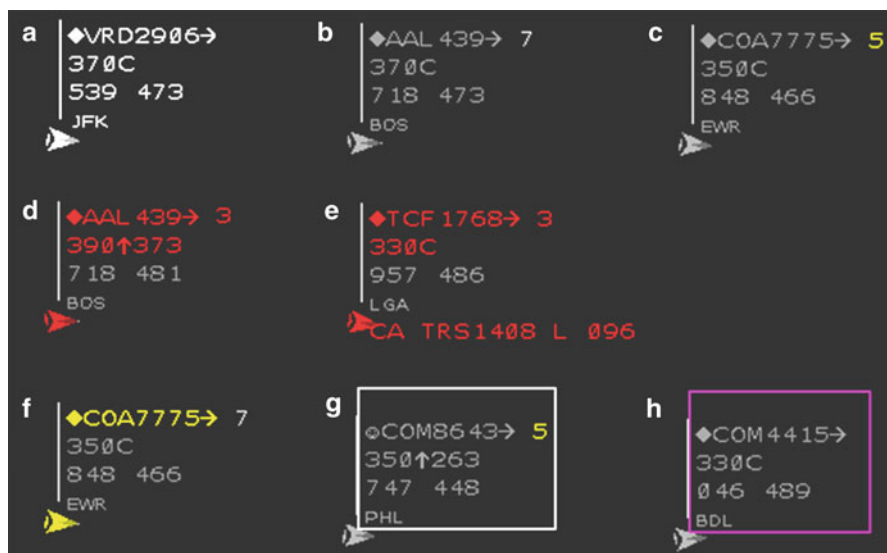data tags as chevrons with
altitude tags

Fig. 3 Expanded data tag examples

Expanded data tags were used only in conjunction with situations requiring human attention. Figure 3 provides artificially arranged and ordered examples of what these looked like: (a) highlighted when manually expanded by the controller, (b) a "long-term" 7 min to LOS conflict number in gray, (c) a "medium-term" 5 min to LOS conflict number in yellow, (d) a "short term" 3 min to LOS with target symbol, data tag, and conflict number in red, (e) an auto-generated short term conflict resolution advisory in red, (f) a conflict deferred by the automation to the controller in yellow, (g) a conflict that the automation is still "thinking" about, and (h) an aircraft placed in an auto-uplink inhibited status by a controller.

Data tags contained items (Fig. 4) that controllers could left/right click on to initiate different trial plans/auto-resolution requests. For example, left clicking on the arrow opened a lateral trial plan; right clicking on the time to LOS requested an auto-resolution from the automation along the lateral dimension; right clicking on the altitude requested an auto-resolution along the vertical dimension. Also, controllers could click on the diamond to access a data communications menu.

Figure 5 depicts the conflict detection alert and automation status table. Each row represented a conflict and provided the callsigns of the involved aircraft, their datalink eligibility and their vertical status (climbing, descending, or level).
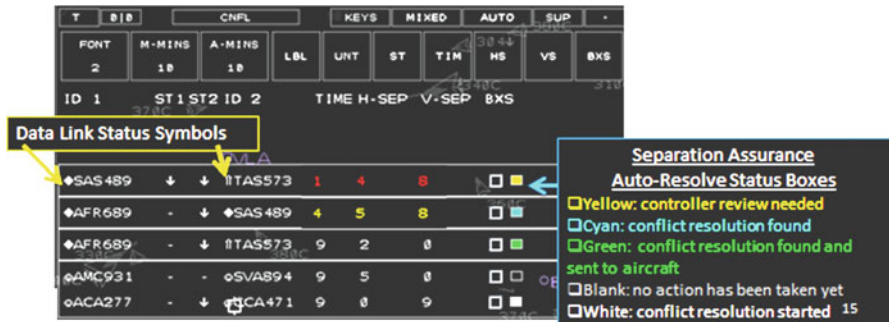
Fig. 4 Clickable data tag items

Fig. 5 Conflict detection alert and automation status table

The count down time to initial LOS was displayed in minutes, the predicted horizontal separation in nautical miles, and the predicted vertical separation in hundreds of feet. Last in the row, a dynamically color-coded box was used to indicate the current action-state of the automation in regard to that specific conflict.

2.1.4 Current Additions for Present Analysis

In line with the separation management standards and visions of the JPDO, the broader human factors research in human–automation functional allocation and trust introduced above, as well as participant comments from prior SA research, new adjustments and additions were made in the current study's human–automation interaction environment. Criteria thresholds for when the automation acted independently of the controller were changed and new intervention functionalities were introduced.

The thresholds were changed to provide a wider range of instances where the automation could uplink resolutions directly to aircraft without controller involvement (i.e. full-auto resolutions) while simultaneously increasing the amount of time available for a controller to observe or act prior to those uplinks by the automation. Specifically, full-auto resolution limits were increased to impositions on aircraft of up to 90 s or more of delay, 60 or more degrees of heading change, 2,200 or more feet of altitude change, and/or 200 or more knots of speed change.

Within these limits, the automation was permitted to directly issue an uplink resolution without involving the controller. Additionally, these criteria-bounded full-auto resolutions could only take place on conflicts that had no more than 8 min until LOS while the conflicts themselves could be displayed as early as 10 min until LOS. Furthermore, auto-generated TSAFE resolution advisories for short-term conflicts were eligible for direct uplink without controller involvement within 2 min to LOS. Suggested resolutions could be displayed as early as 3 min to LOS.

Based on prior feedback regarding the desire to maintain a certain level of control over the automation, new intervention functionalities were introduced that provided the controllers with an ability to inhibit/allow the automatic uplink aspect of the automation. At any point, a controller could input an "NU" (i.e., no uplinks) command and select one or more aircraft to put into a status where the automation was prohibited from uplinking conflict resolutions to the aircraft without their involvement. This status persisted for the aircraft until the same controller entered an "AU" (i.e., allow uplinks) command or the aircraft was handed off to the next controller.

2.2   Problem Statement of Current Analysis

The full-auto criteria adjustments described above combine to provide more opportunities for controllers to observe the automation successfully accomplish its work in handling simpler or "easy" conflicts. This complemented its already apparent proficiency with the routine hand-off and frequency changes. These opportunities are expected to support and engender actions from the controllers consistent with a perspective of reliability or trust in the automation. Such positive experience is assumedly essential as a precursor to effective interactions with the automation in more complex or critical situations. Furthermore, the addition of intervention functionality is expected to foster a sense of engagement, participation and control that should facilitate the controllers' confidence with and effective use of the automation.

First, verification that the prototyped human–automation functional allocation operational concept of this iteration of SA research continues to support the controllers in the NextGen envisioned environment by maintaining the FAA's safe separation standards and forecast levels of increased traffic densities is of principal interest to the current analysis. Next, the present analysis aims to provide a characterization of the transitioning separation assurance responsibilities between the controllers and the automation to explore the different interaction styles of controller trust and use of the automation, and lend insight towards possible factors that contribute to those shifting human–automation interaction styles.

3 Method

3.1 Apparatus

The entire operational environment was simulated using the Multi Aircraft Control System (MACS) software package [21] developed and maintained by the Airspace Operations Laboratory (AOL) software team. MACS is a java based scalable platform used for the prototyping of air traffic management displays and concepts that range from the current day and up through exploratory far term time frames. For each sector presently analyzed a radar controller (R-side) workstation consisted of a standard desktop PC with 75 cm Barco monitor and Display System Replacement (DSR) keyboard and trackball as input devices. These workstations were also equipped with tablet PCs that were used for voice communications similar to the presently fielded Voice Communications System (VCS). Seven pseudopilot stations with standard desktop PC setups were used for the management of flights within the simulation.

3.2 Design

The present analysis focuses on the last six runs of a larger human-in-the-loop SA study aimed to investigate the potential impact of introducing self-separating aircraft in progressively futuristic NextGen time-frames. The full study simulated four different time-frame environments and the last block of six runs were dedicated to representing the environment furthest into the future and with the most advanced human–automation operational paradigm. This portion of the study consisted of 2 days: one full day of training with a morning classroom briefing on the new environment assumptions and automation capabilities, hands-on learning activities, two training runs and discussion sessions followed by a second day of six different 40-min data collection runs.

Traffic scenarios were developed to present each controller participant with a varying range of aircraft densities for their sector over the course of a run to represent an approximate FAA NextGen forecast level of approximately twice that of current day levels (approx. 13–17 aircraft in a sector at any given point) resulting in peak instantaneous traffic counts of well over 30 aircraft in a sector. Scripted conflicts between aircraft trajectories were included in the density mix in addition to those that would naturally occur on account of the increased traffic levels.

Fig. 6 Simulated airspace

3.3 Airspace

The airspace simulated five high altitude sectors from Cleveland Center (ZOB) in the central region of the United States: ZOB 26, ZOB 38, ZOB 79, ZOB 49 and ZOB 59. The floor of each sector was set at flight level (FL) 330 with confederate controllers handling the traffic outside of the five test sectors as well as the aircraft below. As seen in Fig. 6, each sector has unique geographic boundaries and different characteristics of aircraft density, traffic flows and complexity. Arrivals and departures from local area airports (e.g., Toronto-YYZ) contributed to these individual sector differences.

3.4 Participants

The participants consisted of seven current FAA front line managers, each from different enroute centers and current on radar rating and certification. Of these, five served as radar (R-side) controllers and two served as area supervisors that had five different recently retired confederate controllers available for on-call data (D-side)

control positions to support the R-sides. In addition to the D-sides, two other recently retired controllers served as confederate "Ghost" positions that managed the air traffic outside of the test area. Additional confederates included seven general aviation and student pilots that acted as pseudopilots and were assigned to each of the test sectors and surrounding areas.

3.5 Procedure

During the runs, the tasks of the control team were different in many respects from what they are today. As this was a functional allocation study of ground-based automated separation assurance, the main departures were along such lines: the automation was responsible for handoffs, transfers of communication, conflict detection, and conflict resolutions within defined parameters; the controllers were responsible for monitoring the automation's performance, handling conflict situations deferred by the automation, and exercising control of the automation to ensure an efficient and effective flow of traffic through the sector. Additional irregular events of varying situational awareness demand and potential consequence to automation acceptance, e.g. descent/climb requests from pilots due to turbulence however were not yet explicitly controlled for investigation within the current experimental conditions.

Data were collected on the performance of these tasks from a variety of sources throughout the study for later consolidation and analyses. During each run, screen recordings were taken on each of the workstations. Actions performed by participants within MACS and the various states and aspects of the traffic were recorded in real-time by MACS data collection processes. Participants completed post-run questionnaires after the conclusion of each data collection run as well as one post-simulation questionnaire administered at the end of the entire study.

4 Analysis Results

Because of the open-ended nature of human-in-the-loop simulations whereby actions of the controller participants can change the nature of the simulated traffic flows, the next section will first cover descriptive analyses of throughout and safety both on the whole and per the individual sectors (Sects. 4.1 and 4.2) before moving into the analyses of human–automation interaction styles (Sect. 4.3).

NextGen Traffic Densities

Fig. 7 Average aircraft counts across all six runs per sector

4.1 Capacity and Safety

Aircraft counts were calculated and recorded in real time at 1-min intervals for each sector during each run. Figure 7 shows the average aircraft counts collapsed across all six data collection runs. After an initial ramping up of traffic in the first quarter of a run, traffic densities increased to sustained average levels of approximately 23 aircraft for the narrowest and most local flow constrained sector (38) and approximately 29 aircraft for the larger and less local flow constrained sectors (49 and 59). These results verify that the targeted levels of aircraft counts were met and maintained by the test sector controllers across the simulated runs.

To assess the basic level of operational safety in the test airspace, LOS events were examined. A LOS was recorded anytime two aircraft were simultaneously closer than 5 nautical miles (nmi) laterally and less than 800 ft apart vertically. To be included in the following analysis, a LOS had to occur within one of the test sectors after the first 5 min of a run and last for more than 12 consecutive seconds (one full, simulated radar position update). LOS events were further categorized into Operational Errors (OE) and Proximity Events (PE) based upon the lateral separation at the closet point of approach measured along the diagonal between the aircraft. If that lateral separation distance was between 4.5 and 5.0 nmi horizontally, the LOS was counted as a PE; whereas if that distance was less than 4.5 nmi, the LOS was counted as an OE.

Across the 240 min of the six runs multiple LOS events were scripted to occur inside the test airspace. Only two LOS events actually occurred: both classified as PE. However, both LOS events were found to be attributable to simulation artifacts. Specifically, the first PE was due to a confederate pseudo pilot failing to comply with a controller's issued clearance to maintain a specified flight level. The other was due to a traffic scenario design error that unrealistically stacked two departure aircraft

Fig. 8 Number of aircraft handled on average by a sector controller in a single run

together and did not provide the confederate ghost controller a fair amount of time to resolve prior to their entry into the test airspace. In sum, these results verify that appropriate levels of separation safety were maintained despite the increased levels of traffic and built in conflicts.

4.2   Individual Sector Differences

A priori differences in sectors in terms of a sector's demand for climbing and descending aircraft, average time and distances for aircraft to cross a sector, and the nature of the conflicts common to a sector were analyzed as potential contributing factors to a controller's interaction style with the automation. A characterization of each of these differences follows.

4.2.1   Transitioning Arrival and Departure Aircraft

Unique aircraft handled by each controller over the course of a run were counted and classified as either a transitioning aircraft or an overflight. These were averaged per run and the results can be seen in Fig. 8. Transitioning aircraft included those descending towards or climbing out of airports in the local vicinity of the test sectors (e.g., DTW, YYZ, BUF, etc.). These flights created additional complexity for controllers on account of the associated uncertainty and additional constraints and demands not attributed to overflights, which could nominally be left at the same altitude across a sector.

Fig. 9 Average aircraft sector crossing times (s) and distances (nm)

A single factor ANOVA was conducted to examine these differences and found a significantly higher proportion of transitioning aircraft for sector 79 over all the other sectors $F_{(4,25)} = 39.35$, $p < 0.001$. Sectors 26 and 38 had the next highest proportion, which were in turn significantly higher than the proportions of transitioning aircraft for sectors 59 and 49.

4.2.2 Sector Crossing Time and Distance

One of the most visibly apparent individual differences between the controllers is the shape and size of the sector they controlled (Fig. 6). These aspects combine to affect how much time and what kind of space controllers' have to work with for aircraft in their sector before the aircraft is handed off to the next sector. Sector crossing data were recorded for each aircraft that transited a sector to capture how many seconds an aircraft spent in a sector and how far it flew within that sector.

A separate single-factor ANOVA was run to test for differences in both the transit times and transit distances of aircraft for each sector. In both cases, statistical significance was found indicative of more time and space for sector 49 when compared to any other sector; time: $F_{(4,1410)} = 2.98$, $p < 0.05$, distance: $F_{(4,1410)} = 3.49$, $p < 0.01$ with other comparisons failing to obtain significance (Fig. 9).

Fig. 10 Average number of conflicts per sector per run

4.2.3 Conflicts

During the simulation, a conflict event was logged at each track update where two aircraft were predicted to come into LOS at a future point in time in one of the test sectors. Figure 10 shows the average number of conflicts predicted for each sector across all six runs. A single-factor ANOVA was used to test for average conflict frequency differences among the five controllers. Average occurrences of unique conflict pairs differed significantly across the controllers, $F_{(4,25)} = 12.43$, $p < 0.001$ with sectors 59 and 79 having significantly higher average number of conflicts per run than sectors 26 and 49 who in turn had a significantly higher average number of conflicts per run than sector 38 (Fig. 10).

Lastly, for each conflict pair the vertical state for each involved aircraft was recorded at that point in time. Conflict pairs were categorized as level conflicts if both aircraft were level, or transitioning conflicts if either aircraft in the pair was in a climb or descent. Figure 11 shows the average distributions of level conflicts on top of transitioning conflicts for each sector. A single factor ANOVA was used to test for differences in the average percentage of transitioning conflicts between the sectors. Sector 79 had a significantly higher proportion of conflicts that involved transitioning aircraft, $F_{(4,25)} = 4.55$, $p < 0.01$. Comparisons between the other sectors did not obtain significant differences.

4.3 Human–Automation Interaction Styles

For the present analysis of human automation interaction styles, four major sources of information were investigated. These included route, altitude, and/or speed amendments uplinked by the automation without any controller involvement;

Conflict Data, Level vs Transitioning

Fig. 11 Proportional number of conflicts that involved level versus transitioning aircraft

amendments uplinked by a controller with little to no automation involvement; interventions issued by a controller to inhibit the automation's ability to uplink to an aircraft; and subjective workload ratings and responses from questionnaires pertaining to participants' trust and use of the automation.

4.3.1 Uplinks

A total of 709 uplinks were counted across all five controllers and all six runs. 151 of these uplinks occurred without the presence of a conflict for the involved aircraft, whereas the remaining 558 uplinks concerned conflicts. From Fig. 12, it can be seen that sector 38 had the greatest percentage of non-conflict uplinks and sector 49 the least.

For uplinks where the automation detected conflicts, the status of the automation in resolving that conflict (Fig. 5) was recorded. Figure 13 shows the average proportions of different resolution automation states for each test sector.

4.3.2 Full-Auto Resolution Uplinks

288 of the conflict uplinks were full-auto resolutions not involving a controller and of these, 12.5% were tactical avoidance TSAFE resolutions and the remaining 87.5% were sent strategically with more than 3 min until LOS. A single-factor ANOVA was used to test for proportional full-auto resolution uplink differences among the five controllers. The average percentage of uplinks that were full-auto resolutions per run differed significantly across the controllers, $F_{(4, 25)} = 18.63$,

Total Uplinks

Fig. 12 Total uplinks categorized by conflict presence and sector
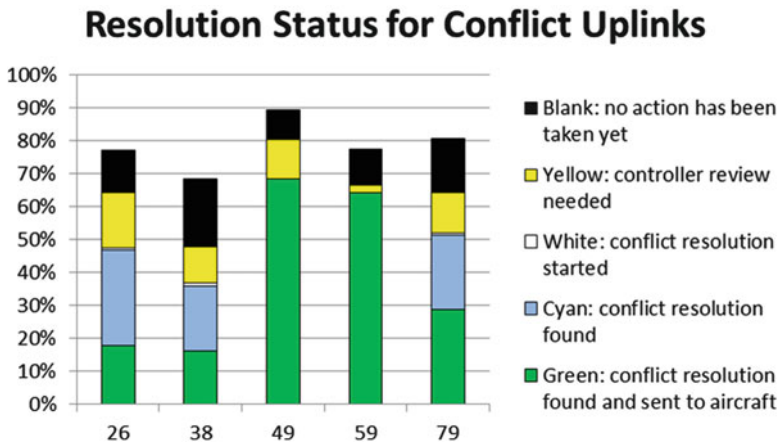
Resolution Status for Conflict Uplinks

Fig. 13 Averaged proportions of status of resolution automation for uplinks involving conflicts

$p < 0.01$ with sectors 49 and 59 having significantly higher average percentages of full-auto uplinks than sector 79 who in turn had a significantly higher full-auto percentage than sectors 26 and 38 (Fig. 14).

4.3.3 Pro-active Controller Resolution Uplinks

Prior to the automation getting involved in the resolution of a conflict, a controller could issue a resolution on his/her own in response to a conflict alert or even, as mentioned above, without a conflict alert at all. In addition to the 151 non-conflict uplinks (Fig. 12), a total of 97 uplinks were issued by controllers across the runs

Fig. 14 Average percentage of uplinks in a run that were full-auto, i.e. the *green bars* from Fig. 13

while the conflict automation status box was still black/blank; i.e. indicative that the automation had not yet begun to work on resolving that conflict (Fig 13). Taking these two numbers together provides a measurement of how pro-active/preemptive a sector controller was in issuing resolution clearances.

A single-factor ANOVA was used to test for differences in the pro-activeness of sector controllers in resolution clearance uplinks. Average percentages of uplinks that were executed by controllers preemptive of automation differed significantly across the controllers, $F_{(4,25)} = 3.95$, $p < 0.05$ with sector 38 showing significantly higher levels of pro-activeness than all the other sectors, and sector 49 the lowest (Fig. 15).

4.3.4 NU Intervention Frequency and Duration

An auto-uplink inhibit event "NU" was counted on a per plane basis and a total of 100 NU's were found issued by all test controllers across the six different runs. In 87% of these cases, controllers inhibited both aircraft involved in the conflict as opposed to just one. Figure 16 shows the total number of NU's for each sector as well as the proportionality of NU's that were issued on top of an active TSAFE advisory.

A single-factor ANOVA was run to test for differences in the number of NU's issued by controllers where the aircraft involved did not have an active T-SAFE advisory. Non-TSAFE NU's differed significantly across the controllers, $F_{(4,25)} = 99.89$, $p < 0.001$ with sector 38 issuing significantly more NU's outside of TSAFE status on average per run than any other sector (Fig. 17).

Avg % Pro-active Uplinks by ATC

Fig. 15 Average percentage of uplinks in a run that were preemptive

NU's Issued by ATC

Fig. 16 Total NU's issued by a controller in the presence and absence of TSAFE advisories

Controllers could revert an aircraft from NU status back to automatic uplink eligibility status at any point or prolong their NU status indefinitely while under their ownership. Length of time in NU status was measured as the time between an NU and a subsequent AU for the same aircraft by the same controller. Initial results indicate that elapsed time in NU status ranged from as short as 20 s to as long as 379 s with an average duration of 115 s across all runs and controllers. Sectors 38 and 79 had the longest NU duration average at 145 s. Trends in the results indicate a positive relationship between number of NU's issued and length of time aircraft were kept in NU status (Fig. 18). In other words, controllers with more frequent use of NU left aircraft in NU status for longer durations on average versus less frequent users of NU who more quickly transitioned aircraft back out of NU status.

Non-TSAFE NU's

Fig. 17 Average number of non-TSAFE NU's issued per run

Avg Seconds of NU Duration

Fig. 18 Average number of seconds controller left aircraft in NU status before reverting with AU command

4.3.5 Subjectives: Workload

Throughout each 40-min run, self-assessment workload prompts appeared in the margin at the top of the controllers' display and lasted for 40 s for each prompt. Workload ratings were made on a "1" to "6" scale (1 = "Very Low Workload" to 6 = "Very High Workload") with averages computed per controller for each run. Notably, in spite of working about twice the level of present day traffic volume, participants still recorded favorable workload ratings. Furthermore, while average ratings for *all* participants fell on the lower end of the scale, statistical analyses still indicated significant differences between their ratings (Fig. 19), F (4,360) = 14.02,

Workload Ratings

Fig. 19 Self-assessed workload ratings: 1–2 "Time on hands," 3–4 "In the groove," 5–6 "Overloaded"

$p < 0.001$. Sectors 49 and 59 were the only controllers to never rate their workload higher than a "2". Sector 49 provided a significantly lower rating than everyone else except 26, while sector 38's higher workload ratings obtained statistical significance as well.

4.3.6 Subjectives: Questionnaires

From their questionnaire responses on the topic of human–automation interaction styles, controllers showed some general consensus both as to what they liked and did not like about the automation tools. They also provided answers indicative of very different personal opinions on particular aspects.

For each of three different questions asking who should be responsible for the detection of conflicts, the generation of resolutions and the execution of resolutions, all controllers selected the answer "controller and ground automation should share." While the current analysis focuses only around the furthest "maximum" NextGen timeframe from the larger multi-timeframe experiment, the nature of the sequential design of the larger study lends itself to potential insights on controllers' developing attitudes over time and growing experience with the automation. While further analyses are planned to provide more detailed investigations, some relevant insight can still be seen at present that shed light on their "maximum" NextGen responses. Growth in controller confidence and trust of the automation can be seen from their increasing experience with it over time. Controllers' confidence grew in the trial planning tools as they used them. They were only "somewhat confident" when they used the tools in the minimum conditions ($m = 4$) but this confidence increased in the maximum conditions when they said they were "very confident" ($m = 6$).

Their confidence grew in a similar way when using the strategic conflict advisories: controllers were "quite confident" when they used the strategic advisories in the moderate conditions (m = 5.6) and this confidence increased to "confident" in the maximum conditions (m = 6). In the moderate condition, controllers' overall averaged opinion of the accuracy of the TSAFE advisories was that it was accurate (m = 4) and this increased slightly as they rated them as quite accurate (m = 4.66) in the maximum conditions.

In spite of the overall positive rating averages of the TSAFE tool, controllers did share some common reservations about its present implementation that limited its resolutions to using vectors without the possibility of using altitude resolutions. Example comments from controllers on sectors 38, 79, and 49 spoke directly to this: "I did not allow the computer to get to the point of needing a TSAFE resolution. The result of the computer applying a TSAFE was not acceptable to me."—38; "The better resolution was to stop the climb of one aircraft versus a turn"—79; "With climbing aircraft we had to be more aware to intervene before the computer vectored aircraft, when stopping at a lower altitude was a much easier, more efficient resolution"—49.

Despite where controller responses agreed with each other, other questionnaire answers alluded to striking differences in their overall experiences and dispositions. On one side of the spectrum, the sector 49 controller had a very easy time working with and trusting the automation. In a series of post simulation questions referencing a list of automation tools that asked what, if any, value was provided by that tool, this controller exclusively responded either "reduced my workload" or "increased my awareness." All other controllers selected answers that stated value was added to the operations (i.e., safer, more efficient) rather than to themselves or that a tool "had no added value." Additionally, for a question asked at the end of each run: "Did you feel rushed and that you did not have enough time to complete tasks? Or, did you feel that you did not have enough to do?" sector 49 marked the minimum value of "1 – very low time pressure" on the 7 point scale every time.

On the other hand, sector 38 indicated a personal preference and comfort for human control rather than trust of automation control in some areas. For example, "I don't always trust the solutions the computer comes up with, and never like the TSAFE resolutions" and "I think things will get easier as my comfort level increases. I do not always trust the solution or believe that they are in the best interest of the aircraft." Sector 38 answered "had no added value" to each of the three different value questions regarding TSAFE automation. At the end of a run, 38 was the only controller to answer "moderate compensation required to maintain adequate performance" to the question "how much did you have to compensate for the automation to make the tools and concept work?" all others selected either "minimal compensation" or "no controller correction." Another example of sector 38's confidence in himself over the automation comes from a question that asked at the end of a run for the controller to comment on whether or not they had enough time to resolve their most complex conflict, to which 38 responded "yes, only because I saw the potential loss of separation before the computer, put an NU on the involved aircraft, and separated them my way when the red fifth line appeared."

The questionnaire responses from sectors 26, 59 and 79 generally fell in between 49 and 38 with more moderate ratings and/or comments.

5 Discussion

Working within the human–automation interaction paradigm examined in the present analysis, controllers were able to maintain safe separation standards in spite of future levels of increased air traffic demand. From an absolute perspective of taking the group of participants on the whole, all the controllers worked well with and liked the automated tools. This can be seen from meeting the above goals along with low workload ratings and questionnaire responses that revealed they preferred sharing separation assurance responsibilities with a set of automated tools that they increasingly trusted as time and experience with them went on.

Exploring a relative comparison perspective between the controllers, the human automation interaction style measurements above observably divided the controllers along a spectrum with sector 38 placed towards a more manual end, sector 49 towards a more automated end, and the others falling somewhere in between. With the greatest proportion of non-conflict and pro-active trajectory uplinks, as well as the lowest proportion of full-auto uplinks, the highest number of NU's and non-TSAFE NU's, and the greatest average NU status durations, sector 38's objective data combine to stand out as a characterization of a more active approach to the human–automation team working dynamic. This higher level of engagement and activity is also reflected in the higher average workload ratings of sector 38 compared to the other sector controllers, though notably still well within the acceptable range of the workload scale. Sector 49 on the other hand, assumed a much more passive approach in the controller-automation dynamic, with the highest percentage of full-auto uplinks, lowest percentage of uplinks without automation involvement, and relatively low number and duration of automatic uplink interventions. Assuming such an approach, sector 49's peak workload ratings never exceeded a "2." The subjective questionnaire responses from 38 to 49 substantiate their differing styles of action, as their own words and ratings exhibit contrasting opinions of automation trust and use.

While individual differences in how much people trust and use automation will surely always exist based from their own personal experiences and attitudes, task characteristics like demand, pressure and complexity might reasonably be expected to influence a person's behaviors with automated tools. Individual sector differences from the pre-scripted traffic flows and sector geographic dimensions presented the controllers with very different and highly contextualized local work environments. Some of these factors exist on a level completely independent of a controller. For example, sectors 49 and 59 clearly had the most time and space to work with aircraft, as well as having to serve the lowest demand of transitioning aircraft. Other factors also reflect a local work environment that dynamically changes based on the actions taken from within that environment, as this is the nature of "human-in-the-loop."

For example, based on the pro-active resolution approach of the sector 38 controller, the lower number of conflicts certainly also reflect his solving of some conflicts early enough that they weren't recorded as conflicts. In contrast, sectors 59 (most likely by choice/comfort) and 79 (most likely by transitioning demand) had much higher levels of recorded conflict events.

Interestingly, sector 59 called out several instances in his questionnaire responses where he disapproved of the automation's handling of a situation. He also had approximately the same levels of sector crossing time/distances, and transitioning aircraft conflict demand as sector 38. However, unlike 38, we observed in his questionnaire comments a more passive approach like that of 49, i.e., "the hardest part will be to keep the controllers engaged"—59. Additionally, the arrangement of the simulation which had 59 co-located in the south area alone with 49 and separate from the other controllers, provided more opportunity for 59 to observe and be influenced by a functional passive approach than perhaps would have been afforded to him alone.

From the present analysis, the most clear and single mapping between individual sector characteristics and resultant human interaction style appears to be between lower levels of transitioning aircraft demand and lower levels of pro-active controller resolutions. Less transitioning aircraft have been observed to lead to fewer short-term conflicts and TSAFE advisories. The resultant trend in interaction with automation is underscored by the controller's expressed dislike that TSAFE resolutions were limited to the lateral dimension alone, which encouraged them to be more pro-active in assigning altitude stops themselves.

Several areas of future research are encouraged from the current analysis. While one can get some preliminary ideas of controller differences at present, more can be learned from subsequent tests in more precisely targeted and controlled studies. Most relevant to continuing from this exploratory vein of characterizing individually different controller–automation interaction styles would be a between-subjects designed study with either controller participants randomly rotated between or experimentally paired in specific sectors to ascertain relative effects of localized traffic and sector demands on a priori attitude towards trust/use of the automation. Additionally, further analysis of metrics to independently characterize traffic conflicts in open-loop runs would help to more accurately identify the variance in task or problem posed to each sector and speak towards levels of controller reliance on automation. Lastly, while flexible and accommodating to multiple styles of real-time usage, all the controllers in the present analysis shared the same underlying automation configuration parameters. In the future, this might not need to be the case. Individually tailored automation settings per the various localized sector environments and controller preferences for automation task sharing styles might be set ahead of time or flexibly adapted in real-time based on performance.

6 Conclusions

The automation's design was very flexible, with multiple interaction points for different stages of manual and automated control and so accommodated a variety of individually different passive to active work styles of the controller participants. The provision of increased ranges of opportunities for the automation to act independently and be previewed in doing so were well received by some sectors (49 and 59) while others felt much more comfortable with exercising the auto uplink intervention NU functionalities (38 and 79). Not only did the controllers work with the automation to meet their safety and traffic level goals in this simulated future NextGen timeframe, they also did so in different ways and with different attitudes of trust/use of the automation. The prototyped controller-automation functional allocation framework was on the whole very flexible and very successful.

References

1. Federal Aviation Administration (2012) NextGen Briefing. Retrieved online November 2012 from http://www.faa.gov/air_traffic/briefing/
2. Federal Aviation Administration (2012) FAA Aerospace Forecasts Fiscal Years 2012–2032
3. Joint Planning and Development Office (2010) Concept of Operations for the Next Generation Air Transportation System [Draft version 3.2]
4. Joint Planning and Development Office (2011) Targeted NextGen Capabilities for 2025
5. Joint Planning and Development Office (2011) JPDO Trajectory-Based Operations (TBO) Study Team Report
6. Parasuraman R, Sheridan T, Wickens C (2000) A model for types and levels of human interaction with automation. IEEE Trans Syst Man Cybern A Syst Hum 30(3):286–297
7. Hou M, Kobierski R, Brown M (2007) Intelligent adaptive interfaces for the control of multiple UAVs. J Cogn Eng Decis Mak 1:327–362
8. Lee J, See K (2004) Trust in automation: designing for appropriate reliance. Hum Factors J Hum Factors Ergon Soc 46:50–80
9. Madhavan P, Wiegmann D, Lacson F (2006) Automation failures on tasks easily performed by operators undermine trust in automated aids. Hum Factors J Hum Factors Ergon Soc 48(2):241–256
10. Maille N, Sarrazin J (2012) Sense of control in supervision tasks of automated systems. Aerospace Lab Journal, AL04-09. Retrieved online November 2012 from http://www.aerospacelab-journal.org/
11. Erzberger H (2001, December) The automated airspace concept. In: Proceedings of the fourth USA/Europe air traffic management R&D seminar
12. Erzberger H (2009) Separation assurance in the future air traffic system. In: Proceedings of the ENRI international workshop on ATM/CNS (EIWAC 2009)
13. Homola J (2008) Analysis of human and automated conflict resolution capabilities at varying levels of traffic density. Master's thesis, San Jose State University, San Jose, CA

14. Prevot T, Homola J, Mercer J, Mainini M, Cabrall C (2009) Initial evaluation of NextGen air/ground operations with ground-based automated separation assurance. In: Proceedings of the eighth FAA/Eurocontrol R&D seminar
15. Homola J, Prevot T, Mercer J, Brasil C, Martin L, Cabrall C (2010) A controller-in-the-loop simulation of ground-based automated separation assurance in a NextGen environment. In: Proceedings of the ENRI international workshop on ATM/CNS (EIWAC 2010)
16. Prevot T, Mercer J, Martin L, Homola J, Cabrall C (2010) Function allocation for ground-based automated separation assurance in NextGen. In: Proceedings of the international conference on human-computer interaction in aerospace
17. Prevot T, Homola J, Martin L, Mercer J, Cabrall C (2011) Automated air traffic control operations with weather and time-constraints. In: Proceedings of the ninth FAA/Eurocontrol R&D seminar
18. Prevot T, Homola J, Martin L, Mercer J, Cabrall C (2012) Toward automated air traffic control – investigating a fundamental paradigm shift in human/systems interaction. Int J Hum Comput Interact 28(2):77–98
19. Wing D (2008) Performance basis for airborne separation. In: Proceedings of the 26th International Congress of the aeronautical sciences
20. Wing D, Prevot T, Murdoch J, Cabrall C, Homola J, Martin L, … Palmer M (2010) Comparison of ground-based and airborne function allocation concepts for NextGen using human-in-the-loop simulations. Paper presented at the 10th AIAA Aviation Technology, Integration, and Operations Conference, September, Fort Worth, TX
21. Prevot T (2002) Exploring the many perspectives of distributed air traffic management: the multi-aircraft control system (MACS). In: Chatty S, Hansman J, Boy G (eds) HCI-Aero 2002. AIAA Press, Menlo Park, pp 149–154