Human-Automation Cooperation for Separation Assurance in Future NextGen Environments

Joey Mercer
Jeffrey Homola
Christopher Cabrall
Lynne Martin
Susan Morey
Ashley Gomez
San Jose State University
NASA Ames Research Center
Moffett Field, CA, 94035, USA
joey.mercer@nasa.gov

Thomas Prevôt
NASA Ames Research Center
Moffett Field, CA, 94035, USA

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 different function-allocation approaches aligned with four 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 the controller to exceptional circumstances. Results indicate that overall performance was best in the most automated NextGen environment. Safe operations were achieved in this environment for twice today's peak airspace capacity, while being rated by the controllers as highly acceptable. However, results show that sector operations were not always safe; separation violations did in fact occur. This paper will describe in detail the simulation conducted, as well discuss important results and their implications.

Keywords

Human-in-the-loop simulation, air traffic control, function allocation, automation, ground-based separation assurance

BACKGROUND

Predicted air traffic increases over the next 25 years may create a significant capacity problem that the United States' National Airspace System will be unable to accommodate [1]. An automated separation assurance concept was proposed to help solve this problem [2]. However, the adoption of such an approach involves a fundamental paradigm shift in which automation is allowed to perform safety-critical tasks that today are strictly the air traffic controllers' domain. Moving toward automated air traffic control, therefore, requires a careful and thorough investigation.

GROUND-BASED SEPARATION ASSURANCE

Separation management elements of en-route NextGen environments are envisioned to rely on automation to expand performance beyond today's limits by off-loading workload from the controller onto automated functions for the majority of routine operations [3]. Use of automated conflict detection and resolution tools integrated within ground automation systems can therefore enable the controller's working environment to move from tactical separation management to strategic decision-making. In this ground-based approach, air traffic control automation, not the air traffic controller, monitors traffic for potential conflicts. The automation additionally conducts several workload-intensive routine tasks such as transferring ownership and communication frequencies between sectors. Relieved of these tasks, the controller can concentrate on managing non-routine operations that often require human intelligence, ingenuity, and experience. The controller and the automation work cooperatively to enable the higher traffic demands of the future while ensuring the equivalent (or better) safety levels of today [4].

Key Technologies

For automation to successfully perform these new roles, technological enhancements over today's operations are needed. As such, effective integration of ground automation systems, controller work-stations, airborne automation systems, and flight-deck interfaces is an important enabler of this concept. To achieve this integration, connections between these elements are made with the following: Automatic Dependent Surveillance-Broadcast (ADS-B) surveillance providing more precise position, state, and intent data; Data Comm as the primary means of communication and clearance delivery for equipped aircraft; and Flight Management Systems that, in addition to allowing aircraft to fly along a 3D trajectory, can receive Data Comm messages such as frequency changes, cruise altitudes, route modifications, and speed changes. Together, these elements create an investigative environment targeting automation's role in separation

assurance, and the associated human-automation interaction issues.

Included in this simulation was a version of the ground-based separation assurance approach envisioned in a far-term environment, based on the Advanced Airspace Concept (AAC), originally developed by Erzberger [2]. The basic premise is to utilize two independent layers of separation assurance, with each technical sub-system designed to detect and resolve conflicts, but at different time-frames. The Auto-Resolver algorithm works within a 'strategic' time-horizon to ensure separation between aircraft for conflicts detected with more than three minutes until the predicted Loss of Separation (LOS). It computes complete (non-open-ended) trajectories to resolve these less urgent conflicts, which rejoin the aircraft's original route. Even with improved surveillance capabilities, automation systems of the future will still produce trajectory predictions errors. Using climb profiles as an example, some conflicts may not be detected early enough for the Auto-Resolver to address. The Tactical Separation Assured Flight Environment (TSAFE) algorithm focuses on these short-term, 'tactical' conflicts detected with less than three minutes until a predicted LOS. In these urgent situations, TSAFE computes an open-ended heading change to avoid the LOS and keep the aircraft clear a few minutes, and by doing so, provides time for the Auto-Resolver to find a complete trajectory-based solution.

HUMAN-IN-THE-LOOP SIMULATION

In addition to a roll-out of technological advancements over time, NextGen's maturation process will likely also include a growing number of equipped aircraft, enabling higher utilization of automation for monitoring trajectories and managing separation assurance. In August of 2012, the Airspace Operations Laboratory (AOL) at the NASA Ames Research Center conducted a human-in-the-loop simulation to investigate the issue of separation assurance function-allocation between human operators and the automation [5-7]. The simulation tested four separate NextGen time-frames, each representing a candidate level of aircraft equipage, and a potential stage of ground system capabilities in NextGen's evolution. The stages ranged from a near 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.

The simulation also investigated allocating separation functions between the air and ground, by incorporating eight aircraft equipped to manage their own separation. Results showed the presence of these self-separating aircraft had no impact on the ground system's performance [5]. Therefore, no distinctions between these two categories of aircraft are made in this paper; the analyses reported here treated all aircraft in the same manner.

Study Design

The study examined four conditions, one for each of the NextGen time-frames investigated. Differences in the available automation capabilities, as well as different levels of aircraft equipage characterized each condition. Keeping track of the differences between the complexities associated with each condition proved difficult for the participants, as evidenced during preparatory 'shake-down' simulations. As a result, a randomized or counter-balanced run schedule was not pursued, and the four conditions were tested in order of their maturation/automation level: first the Baseline condition, followed by the Minimum NextGen condition, then Moderate, and lastly the Maximum NextGen condition. Each stage took place over two consecutive days, during which the participants were briefed and trained on the operational environment, tools, and procedures on the first day. The data collection period for each condition entailed six 40-minute runs, occurring on the second day. The following sub-sections describe the four conditions in more detail.

Baseline 'Current-Day' Time-Frame

The Baseline time-frame served as a representation of a near-term NextGen with only few differences from current-day, fielded operations. The simulated operations for this time-frame assumed that all aircraft were equipped for ADS-B out; broadcasting their position and state information. In using the more precise surveillance data, controller workstation displays showed a given aircraft's target position, current altitude, and current ground-speed with a 1-second update rate. The improved surveillance data also allowed for a change to aircraft target symbols, shown as directional chevrons. Ownership and tracking status information were integrated into the chevron symbology, shown in Figure 1: solid and hollow chevrons represented owned and un-owned aircraft, respectively, while small and large chevron symbols indicated flat-track (on flight plan) and free-track (off flight plan) statuses, respectively.
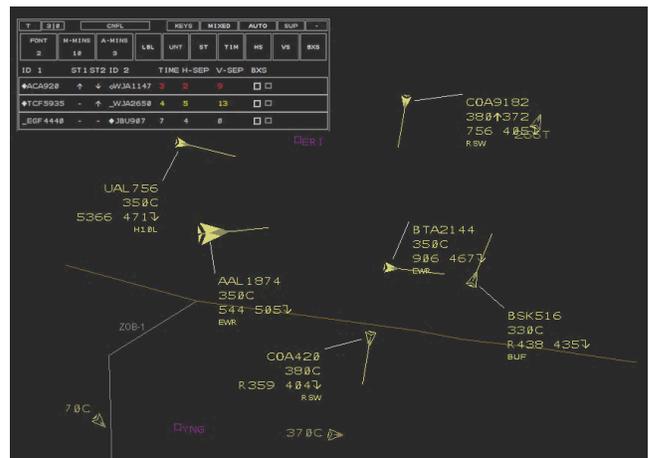

Figure 1. Chevron symbology and conflict list, as seen in the Baseline 'Current-Day' time-frame.

Roles and responsibilities in the Baseline time-frame were unchanged from today's operations. Being responsible for maintaining safe separation between aircraft, controllers performed all conflict detection, evaluation, and resolution tasks. Supporting the controllers in this task was a flight-plan-aided conflict probe, which displayed detected conflicts in a conflict list window, shown in Figure 1, much like the User Request Evaluation Tool (URET).

Controllers issued all clearances via voice, as no aircraft in the Baseline time-frame were Data Comm equipped. Controllers were also responsible for routine book-keeping tasks, just as they are in today's operations. These included hand-off initiation, hand-off accept, and transfer of communication (frequency changes).

Minimum NextGen Time-Frame
Envisioning a possible early stage of NextGen's evolution, the simulation's Minimum NextGen time-frame introduced a limited Data Comm implementation. 25% of the aircraft were Data Comm equipped, which received slightly different handling than the unequipped aircraft. For Data Comm equipped aircraft, the controller's workstation automatically initiated out-going sector-handoffs and automatically accepted incoming sector-handoffs. Once the receiving sectors accepted the automated hand-offs, the sending sector's workstation automatically prepared and sent a Data Comm transfer-of-communication message to the aircraft, instructing them to switch to the next sector's frequency. Additionally, Data Comm equipped aircraft did not need to verbally check-in with the controller upon receiving the frequency change instruction. Therefore, it was the automation that performed the hand-offs and transfers-of-communication for the equipped aircraft, easing any controller workload otherwise associated with such tasks. Shown in Figure 2, Data Comm messages sent by the ground system's automation, and the 'ROGER' or 'UNABLE' response message received from the flight-deck, were displayed to the controller in an on-screen Data Comm message status list.
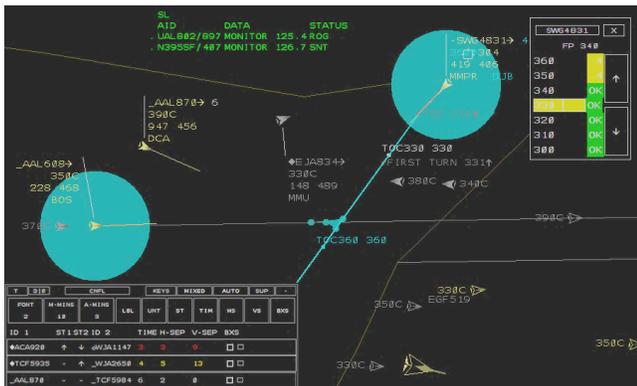

Figure 2. Screen image from the Minimum NextGen time-frame, showing the Data Comm status list, equipped (grey) aircraft, trial-planning capability, and conflict probe information integrated into the data blocks and altitude fly-out menu.

Serving as another potential early step in the evolution towards trajectory-based operations and more automated NextGen concepts, the Minimum NextGen time-frame included an assumption that allowed Data Comm equipped aircraft to follow their Flight Management System's (FMS-) computed vertical profile. Specifically, unless instructed otherwise by the controller, aircraft were 'pre-cleared' to climb to their cruise altitude, and allowed to descend at their Top of Descent (TOD) point. Because the controllers were still responsible for the separation of aircraft, they could issue any clearance deemed important for their sector's safe operation; at which point an equipped aircraft receiving such an instruction would behave exactly as an unequipped aircraft.

The automated hand-offs and frequency changes, together with the notion of being 'pre-cleared' to follow the FMS-computed vertical profile, increased the likelihood that, nominally, the controller would not need to interact with an equipped aircraft at all. This led to a change in the design of the equipped aircrafts' data blocks, illustrated in Figure 2. In contrast to unequipped aircraft's yellow data blocks, which cannot be collapsed while owned by the controller or geographically inside their sector, the data blocks of equipped aircraft were displayed in muted grey color, and were only shown as full data blocks if the aircraft was within 150 nmi. of the destination airport, or if the automation detected a traffic conflict for the aircraft. Intended to roughly coincide with an aircraft's being just before their TOD point, popping up the data blocks relative to the destination airport was thought to help controllers maintain a level of awareness of the 'pre-cleared' vertical changes expected of the equipped aircraft.

Decision-support tool enhancements for the controllers came in two other areas. First, the conflict probe's information was better integrated with the controller's display: in addition to the conflict list, the remaining time until predicted LOS for a detected conflict was shown (in minutes-to-go) directly in the data block, supporting the controller's traffic scan (see Figure 2). Secondly, trial-planning functions were available to help the controller plan aircraft trajectory changes. The trial-planning tools allowed the controllers to craft a provisional trajectory for an aircraft, which was fully integrated with the automation's conflict probe, as illustrated in Figure 2. Such integration provided what-if feedback, informing the controller as to the possible outcome of issuing the planned clearance, before actually doing so.

Additionally, the trial-planning capabilities helped to enhance the data block altitude fly-out menu. By incorporating data available from the conflict-probe-informed trial-planner, the altitude fly-out menu examined under-the-hood trial-plans for each flight level. This in turn, was displayed to the controller such that they could, over a range of altitudes, quickly see which were clear of potential conflicts, and, for the flight levels not clear, see the time-to-predicted-LOS. Originally based on the design

from [8], an example of the altitude fly-out menu with pre-probed altitudes is shown in Figure 2. It is important to note that, as part of the Minimum NextGen time-frame, none of the trial-planning capabilities were integrated with Data Comm; all trajectory-related clearances were issued via voice for all aircraft.

Moderate NextGen Time-Frame

The continuation to the Moderate NextGen time-frame provided controllers with additional decision-support tools, while keeping the roles and responsibilities identical to those from the Minimum NextGen time-frame. The Moderate time-frame included two additional assumptions, representing a further evolution of NextGen operations. First, the Data Comm capabilities were expanded to allow the sending of trajectory changes. This meant that, in addition to transfer-of-communication messages, route modifications and/or altitude changes created through the trial-planning functions could be sent directly to the aircraft. Also, a larger population of the traffic was assumed to be equipped for Data Comm; simulated here as 50% of all aircraft.

The Moderate NextGen time-frame added two decision-support tools to address the strategic and tactical safety layers. A version of the Auto-Resolver algorithm was available to help the controller with the task of conflict resolution. In the presence of a detected conflict, controllers could invoke the Auto-Resolver algorithm, which would search for a new trajectory that would resolve the conflict, and present that solution to controller in the form of a trial-plan. The controller could then send the displayed trial-plan to the aircraft via Data Comm, just as if they had manually created the trial-plan themselves. Moreover, a trial-plan generated by the Auto-Resolver could be manually modified by the controllers, allowing them to use it as a starting point to then make adjustments before sending the clearance 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 [6].

Additionally, the Moderate time-frame included TSAFE, marking the first time the automation provided support to

Figure 3. A TSAFE advisory from the Moderate NextGen condition, suggesting the predicted LOS between WJA2652 and JBU1119 can be avoided by turning the WJA right to a heading of 202 degrees, and turning the JBU left to heading of 156 degrees.

the controllers during urgent, tactical conflicts. In the event of a short-term conflict, the ground system's automation would calculate the heading changes needed to avoid the pending LOS, and, if possible, keep the aircraft free of other conflicts. Displayed to the controllers in the data block's 5th line (see Figure 3), the resulting headings served as reference information only: the controller could issue the suggested heading change via voice or Data Comm, or could issue something of their own choosing; whether it be a slight modification to the TSAFE advisory, or something completely different.

Maximum NextGen Time-Frame

Over the progression of the first three NextGen time-frames, the controllers received increasing support from the automation to assist them in their responsibilities. However, this trend changed slightly in the fourth condition. In the Maximum NextGen time-frame, all aircraft were Data Comm equipped, with an established electronic communication link to the ground. Consequently, Data Comm messages could be sent to the aircraft by the controller, and theoretically also directly from the automation, creating opportunity for a new distribution of tasks between the controller and the automation. In comparison to the Moderate time-frame, the capabilities of the tools themselves changed only slightly, but *how* those capabilities were used was quite different, resulting in a true paradigm shift of air traffic operations.

The Maximum NextGen time-frame investigated an allocation of air traffic functions between the controller and the automation, but it was the allocation of air traffic *responsibilities* that served as the foundation of the human-automation cooperation scheme, and is what distinguished the Maximum NextGen time-frame from the other conditions. Here, the automation's responsibilities included three critical tasks: 1) conflict detection, 2) tactical LOS avoidance by means of automatically sending any pending TSAFE advisories for detected short-term conflicts when two minutes or less remained until the predicted LOS, and 3) alerting the controller to any problems or exceptional situations. In addition, the automation worked within pre-defined limits to resolve conflicts detected within the strategic time-horizon, automatically sending the resolution trajectory to the aircraft. Also, when an aircraft received an automated TSAFE instruction, the automation would then later follow-up with a new trajectory that put the aircraft back on course, rejoining its original route. Also, since all aircraft were Data Comm equipped, all hand-offs and transfers of communication were performed by the automation. By allocating these responsibilities and tasks to the automation, the controller's role changed significantly, to one more focused on supervisory and management-by-exception duties.

It is important to note that the air traffic controller and the automation were jointly responsible for maintaining safe separation; however, under these operations, the controller
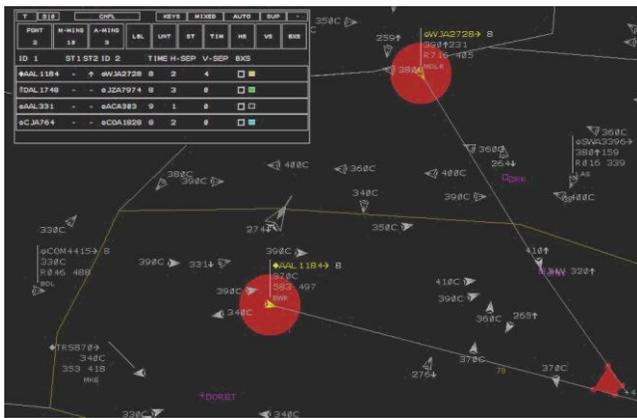
Figure 4. Screen image from the Maximum NextGen time-frame, showing the automation status indicators in their various colors. The small yellow box in the conflict list, signifying 'needs controller attention,' is coordinated with the conflicting aircraft highlighted with yellow full data blocks.

would not be held accountable for a LOS in which the automation never detected the conflict, and/or never alerted the controller, and/or never issued a TSAFE advisory.

In cooperation with the automation, the controller's primary responsibilities were to supervise the automation and to resolve any situations flagged to them by the automation. The automation displayed to the controllers the status of its efforts to resolve conflicts detected within the strategic time-horizon through colored boxes shown in the conflict list, examples of which are shown in Figure 4. These status indicators used five different colors to represent the five possible statuses of the automation's conflict resolution process: 1) empty boxes signifying the automation was not working on that conflict, typically when a conflict was detected more than eight minutes away from the predicted LOS; 2) white-filled boxes signifying the automation was actively looking for a resolution to the conflict; 3) green-filled boxes signifying the automation successfully found a resolution to the conflict; 4) cyan-filled boxes signifying the automation successfully uplinked the resolution to the aircraft; and 5) yellow-filled boxes signifying the automation either could not find a resolution to the conflict at all, or could not find a resolution within its pre-defined limits (i.e., one that imposed less than 90 seconds of delay, 60 degrees of heading change, 2200 feet of altitude change, or 50 knots of speed change).

Tactical conflicts were generally handled by the automation with the automatic uplink of TSAFE advisories, but exceptions in this domain were of a different sort. Rather than the automation alerting the controller to situations in which it needed help, here the controller could alert the automation to situations in which they wanted to assume a more hands-on role. Similar to a manual override, this meant that controllers could inhibit the automation's uplink of TSAFE advisories, so that they could address the situation in their own way.

Airspace and Traffic

Figure 5 illustrates the simulated airspace used for this study, consisting of five adjacent test sectors, all in the high-altitude en route airspace of Cleveland Air Route Traffic Control Center (ZOB). These sectors were divided across two areas of specialization: sectors 26, 38, and 79 pertained to the North area, while sectors 49 and 59 belonged to the South area. The floor of the overall test airspace was set at flight level (FL) 330. One participant, working as the radar controller, and one supporting confederate controller working as the radar associate, were assigned to each of the five sectors. Confederate "Ghost" controllers were responsible for the airspace surrounding the test area. Sector geometries and traffic flows combined to create natural variations in complexity between the sectors.
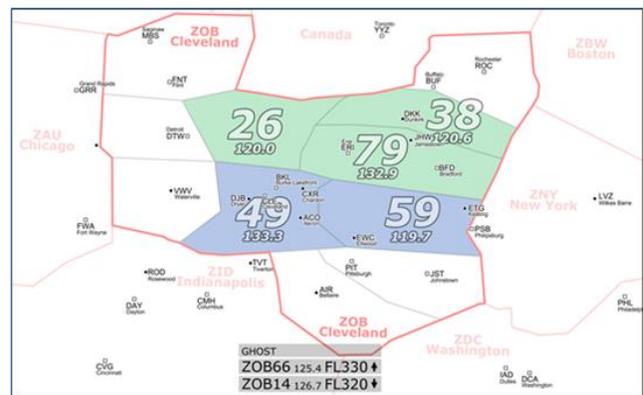

Figure 5. Test airspace used for the simulation.

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. In keeping with the availability of new technologies and decision support tools, as well as an increasing number of aircraft equipped for Data Comm, traffic densities in the scenarios increased from one condition to the next, illustrated in Figure 6. The Baseline condition used traffic levels representative of today's operations, with a Monitor Alert Parameter (MAP) of 18 aircraft per sector. The Minimum NextGen time-frame saw traffic levels increase by 20% to a MAP value of 22 aircraft per sector, while the Moderate NextGen time-frame simulated a 50% increase in traffic, with a MAP value of 27 aircraft. In the Maximum NextGen time-frame, traffic levels were double those of the Baseline condition, using a MAP value of 36 aircraft per sector.

Participants

Seven Federal Aviation Administration (FAA) front line managers (six current and one recently retired) served as primary participants; five as radar controllers and two as area supervisors. Eight additional retired controllers supported the test participants, five working as D-sides for the test sectors, and three handling the traffic in the 'ghost' sector outside the test airspace. Ten type-rated airline pilots operated eight mid-fidelity, single-aircraft flight
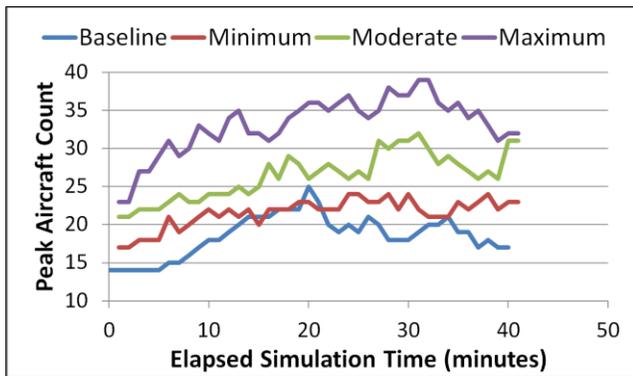
Figure 6. Peak traffic levels for each of the four conditions.

simulators, and ten general aviation/corporate pilots operated multi-aircraft stations.

Separate rooms housed the North and South areas (Figure 5). The configuration of each test sector included a primary radar display and a nearly identical radar associate (D-side) display. Area supervisors assigned to each room monitored the traffic situation as well as the workload of the participant radar controllers, and judged whether/when D-side support was needed.

Equipment

Other than the eight single-aircraft flight simulators using the Aircraft Simulation for Traffic Operations Research (ASTOR) software developed at NASA's Langley Research Center [9], the primary simulation platform used for the study was the Multi Aircraft Control System (MACS). Hosting a Display System Replacement (DSR) emulation, each controller workstation utilized a large-format monitor and a specialized keyboard and trackball, similar to those used in current air traffic control facilities. Voice communications occurred through a custom, stand-alone voice application, meant to emulate the fielded system. Data recorded and collected at each workstation included aircraft flight states, operator task data, automation states, voice communications, etc. Screen recordings captured as movie files were also saved. Workload Assessment Keypads (WAKs) probed controller workload at three-minute intervals during simulation trials using Air Traffic Workload Input Technique (ATWIT) [10] ratings on a modified six-point scale (1 as low workload, 6 as high workload). The controllers completed questionnaires at the end of each run, as well as a post-simulation questionnaire. Debrief discussions provided an additional opportunity for controllers to offer feedback.

RESULTS

Select results from the simulation are described in the following section. Other publications are available that provide additional descriptions and explanations of the data produced during the simulation [5-7].

Detected Conflicts

Predictions by the automation that an aircraft would come too close to another aircraft resulted in a detected conflict. The automation used 'detection buffers', thereby including additional safety margins in its search for conflicts. In most cases, the automation probed for conflicts using separation minima of 5.9 nmi. laterally and 1,000 feet vertically. However, uncertainties associated with climbing and descending aircraft, as well as aircraft off of their expected trajectory, required minor changes to system settings in order to increase stability and minimize false and late alerts. For climbing and descending aircraft, the automation utilized expanded buffers in the vertical dimension of 1,500 feet. For off-trajectory aircraft, the automation used a reduced look-ahead time of five minutes (compared to 10 minutes nominally).

Throughout the simulation, the automation detected a total of 2,323 unique pairs of aircraft in conflict. Figure 7 shows the distribution of these conflicts across the four conditions. The quantity of detected conflicts generally increased with the different NextGen time-frames; an expected result given their increased traffic levels.
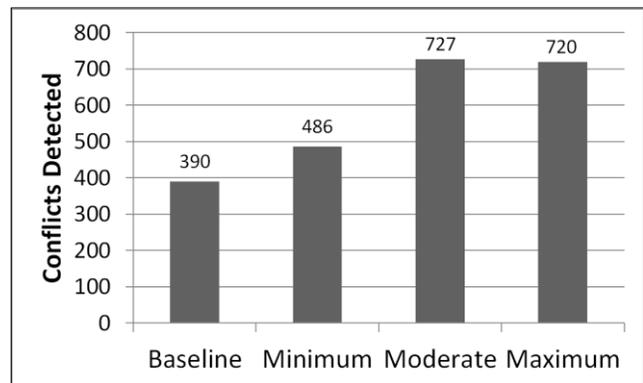

Figure 7. Number of conflicts detected by the automation in each of the four conditions.

Safety

A LOS event occurred when two aircraft in the simulation were simultaneously closer than 5 nmi. laterally and 800 feet vertically. Additionally, a LOS had to occur within the test sectors after the first 5 minutes of a run and persist for at least 12 consecutive seconds. From the 2,323 detected conflicts, a total of 25 resulted in LOS events, the distribution of which across the NextGen time-frames is shown in Figure 8. Interestingly, the Baseline and Maximum conditions showed equal safety performance, despite the latter having twice the amount of traffic. A clear majority of the LOS events (60%) occurred in Moderate condition, a surprising result, given that the Maximum condition had comparable amounts of detected conflicts (Figure. 7).

Workload

During the simulation, controller responses to real-time workload queries utilized the scale's entire range. In the Maximum condition however, controllers never rated their workload above 3. Illustrated in Figure 9, mean workload ratings were similar for the Baseline, Minimum, and Moderate conditions, while workload reported in the Maximum condition was lower. On average, the Maximum
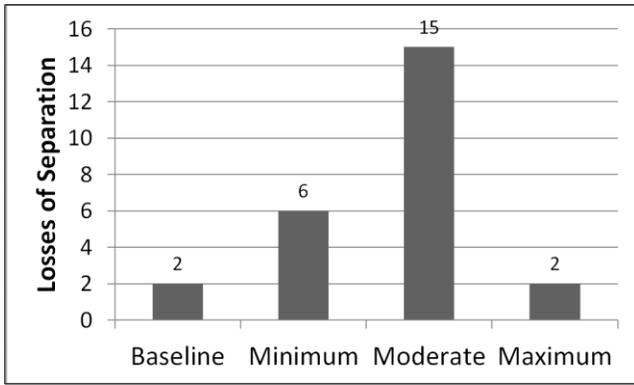
Figure 8. Number of losses of separation occurring in each of the four conditions.

condition's workload across the test sectors was just below 2, considered to be 'low' workload. Statistical testing confirmed this to be significantly different than the workload ratings from the other conditions ($z = 55.08$, df = 3, $p < .000$), all with averages of nearly 3, or 'moderate.' An analysis of workload data from post-run questionnaires verified these findings, showing similar trends [7].
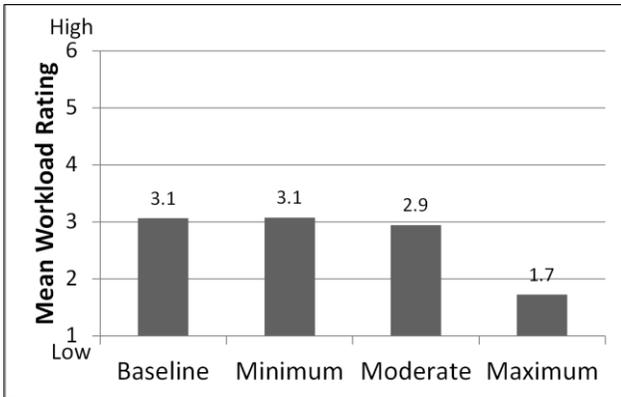

Figure 9. Mean real-time workload ratings across all participant radar controllers for each of the four conditions.

Situation Awareness
Post-run questionnaires asked the controllers to judge their own situation awareness by rating their level of 'situational understanding,' their 'capacity to take in more information,' and the 'demand on their attention.' Ratings from these three scales formed an overall situation awareness score, or Situation Awareness Rating Technique (SART) score, which ranged from -5 (low situation awareness) to 13 (high situation awareness) [11]. The mean SART scores computed for each condition are depicted in Figure 10.

Overall, the controllers gave higher situation awareness ratings in the three NextGen time-frames, as compared to the Baseline condition. The SART scores in both the Baseline and Moderate conditions, although slightly higher in the Moderate condition, signal 'moderate awareness' on behalf of the controllers. Perceived situation awareness increased in the Minimum and Maximum conditions, the
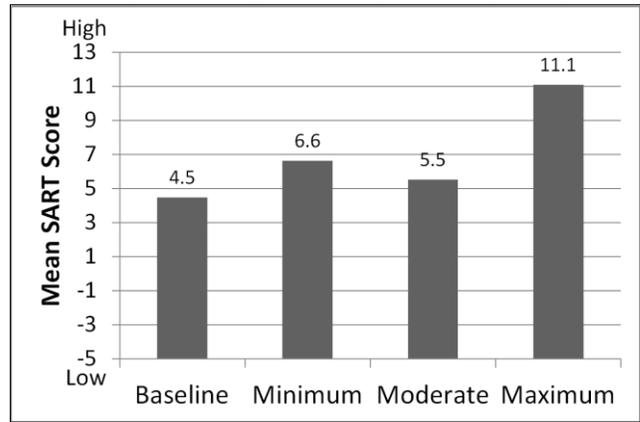

Figure 10. Mean SART scores across all participant radar controllers for each of the four conditions.

latter with a SART score just above 10, considered to be 'high awareness.' Statistical testing confirmed this to be a significant difference ($z = 46.14$, df = 3, $p < .000$).

Acceptability
Another subjective measure obtained through questionnaires was the Controller Acceptance Rating Scale (CARS) [12]. Derived from controllers' answers to several yes-no questions regarding how safe, adequate, satisfactory, and acceptable the operations were, the overall CARS scores served as an assessment of the controllers' comfort level with the simulated operations.

The mean CARS scores indicate that, for the most part, the controllers rated the operations as acceptable, although acceptability was lower for the Moderate condition. Statistical testing verified this as significantly different ($z = 12.22$, df = 3, $p < .007$), driven in large part by 'unsafe' ratings occurring in 40% of the controllers' scores in the Moderate condition.
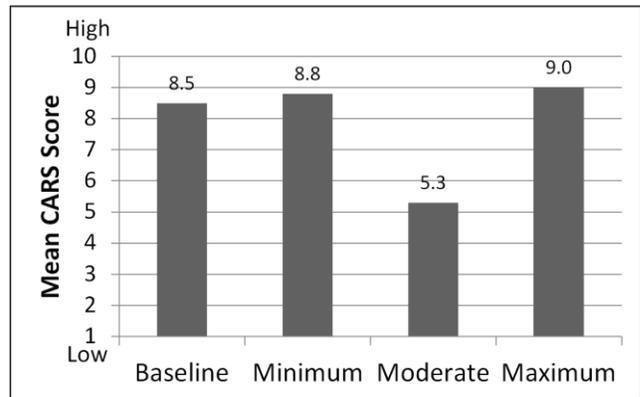

Figure 11. Mean CARS scores across all participant radar controllers for each of the four conditions.

CONCLUSIONS
With regard to the question: "How effectively did the controllers and automation cooperate?", promising results from the Maximum condition appear to suggest a successful cooperation: compared to the other conditions, it

had the highest levels of traffic, nearly the highest number of detected conflicts, safety performance equal to that of the Baseline condition; but recorded the lowest workload, highest SART, and highest CARS scores.

Even though the Moderate condition had substantially fewer aircraft than the Maximum condition, the number of detected conflicts for both conditions was quite comparable, suggesting that additional (i.e., secondary) conflicts were possibly created by the controllers' own actions in the Moderate condition. This theory may also help explain the increase in LOS events in the Moderate condition, but further analyses are needed. Indeed, the Moderate (and also the Minimum) condition, at times, resulted in unacceptable safety measures. Furthermore, the situation awareness and acceptability data from the Moderate condition provide additional evidence that may point to ineffective controller-automation cooperation. These results create a need to better understand the complexities of the simulated operations, and how they interacted with the studied function allocation schemes.

From this perspective, the question becomes: "is the effectiveness of a given controller-automation cooperation scheme context-dependent?" The progression of the four conditions saw an essentially linear increase in overall traffic, increase in Data Comm equipped aircraft, and increase in decision-support tool capabilities. Changes to task responsibility however, progressed in a discrete, step-function-like manner, affecting only the Maximum condition, in which the nature of the controllers' job was very different. Perhaps changes to the allocation of responsibilities occurring in earlier NextGen time-frames, such as in the Moderate condition, would have resulted in better system performance. What if, for example, the Moderate condition had allocated separation responsibility for Data Comm equipped aircraft to the automation, while the controller managed the unequipped aircraft? Follow-on investigations are needed to explore this trade-space.

The safety data indicates that during the Minimum and Moderate NextGen time-frames, the traffic situations may have been beyond the limits of the 'system's' capability, with 'system' represented by the function allocation schemes used during the simulation. One could conclude that the manner in which controllers and automation cooperate (i.e., the cooperation's effectiveness), is certainly a direct result of the function allocation schemes themselves, but perhaps is also impacted by the context, subtleties, and intricacies of the working environment (traffic density, etc.).

Investigations into the uniqueness of the Moderate condition have already begun [6-7]. Forthcoming analyses will focus on quantifying the complexities of the traffic situations and look to identify factors that correlate to when the operations started to break down. Additional analyses will examine other measures of performance for the four conditions, such as flight path efficiency.

REFERENCES

1. Federal Aviation Administration. *FAA Aerospace Forecast, Fiscal Years 2013-2033*. FAA, Washington, D.C., 2013.

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

3. Joint Planning and Development Office. *Concept of Operations for the Next Generation Air Transportation System. Version 3.2*, JPDO, Washington, D.C., 2010.

4. Prevôt, T., Homola, J., Martin, L., Mercer, J. and Cabrall, C. Toward Automated Air Traffic Control – Investigating a Fundamental Paradigm Shift in Human/System Interaction. *International Journal of Human-Computer Interaction*, 28 (2), 77-98, 2012.

5. Wing, D., Prevôt, T., Lewis, T., Martin, L., Johnson, S., Cabrall, C., et al. Pilot and Controller Evaluations of Separation Function Allocation in Air Traffic Management. *Proceedings of the USA/Europe Air Traffic Management R&D Seminar*, Chicago, IL., 2013.

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

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

8. McNally, D., Erzberger, H., Bach, R., and Chan, W. *A Controller Tool for Transition Airspace*. AIAA-99-4298, American Institute of Aeronautics and Astronautics, Reston, VA., 1999.

9. Liu, D., and Chung, W. *ASTOR: An Avionics Concept Test Bed in a Distributed Networked Synthetic Airspace Environment*. AIAA-204-5259 American Institute of Aeronautics and Astronautics, Reston, VA., 2004.

10. Stein, E. *Air Traffic Controller Workload: An Examination of Workload Probe*. DOT/FAA/CT-TN84/24, Federal Aviation Administration, Washington, D.C., 1985.

11. Selcon, S., Taylor, R., and Koritsas, E. Workload or Situational Awareness? TLX vs. SART for aerospace systems design evaluation. *Proceedings of the Human Factors and Ergonomic Society Annual Meeting,* 35 (2), 62-66, 2001.

12. Lee, K., Kerns, K., Bone, R., and Nickelson, M. Development and Validation of the Controller Acceptance Rating Scale (CARS): Results of Empirical Research. *Proceedings of the USA/Europe Air Traffic Management R&D Seminar*, Santa Fe, N.M., 2001.