THE EFFECTS OF LIMITED INTENT INFORMATION AVAILABILITY ON SELF-SEPARATION IN MIXED OPERATIONS

Timothy A. Lewis, Nipa Phojanamongkolkij, and David. J. Wing
NASA Langley Research Center, Hampton, VA

Abstract

This paper presents the results of a computer simulation of the NASA Autonomous Flight Rules (AFR) concept for airborne self-separation in airspace shared with conventional Instrument Flight Rules (IFR) traffic. This study was designed to determine the impact of varying levels of intent information from IFR aircraft on the performance of AFR conflict detection and resolution. The study used Automatic Dependent Surveillance–Broadcast (ADS-B) to supply IFR intent, but other methods such as an uplink from a ground-based System Wide Information Management (SWIM) network could alternatively supply this information. The independent variables of the study consist of the number of ADS-B trajectory change reports broadcast by IFR aircraft and the time interval between those reports. The conflict detection and resolution metrics include: the number of conflicts and losses of separation, the average conflict warning time, and the amount of time spent in strategic vs. tactical flight modes (i.e., whether the autoflight system was decoupled from the planned route in the Flight Management System in order to respond to a short-notice traffic conflict). The results show a measurable benefit of broadcasting IFR intent vs. relying on state-only broadcasts. The results of this study will inform ongoing separation assurance research and FAA NextGen design decisions for the sharing of trajectory intent information in the National Airspace System.

Introduction

The separation function performed by the ground-based air navigation service provider (ANSP) contributes to the extraordinarily low risk of mid-air collision between commercial aircraft in the sky today. Human controllers provide essential separation services to IFR aircraft, relying on a combination of procedural and surveillance-based control approaches to ensure aircraft maintain a safe distance from one another at all times. The aircrew on the IFR flight deck has limited knowledge of surrounding traffic and thus only a minimal ability to separate from other aircraft without support from the ground.

However, the emergence of ADS-B technology has the potential to change this reality. In support of NextGen, the FAA has mandated that aircraft equip by 2020 to frequently broadcast a state vector via ADS-B OUT, derived from GPS and other on-board sources of state data. With a sky full of aircraft frequently broadcasting their current position and velocity, an aircraft equipped to receive these reports via ADS-B IN would have access to an accurate and up-to-date picture of the traffic within its vicinity.

The availability of this surveillance information on the flight deck is an enabler for a concept previously known as *free flight* [1], more generically known as *self-separation*. This idea posits that, with detailed traffic information on the flight deck, the aircrew can identify potential trajectory conflicts with other aircraft and make course changes to resolve them autonomously, doing so without active supervision or control by a ground service. Such a mode of flight is a significant departure from conventional IFR operations that require ground-based separation services. Accordingly, the NASA concept has been formulated as a new set of *Autonomous Flight Rules* (AFR), recently described by Wing and Cotton in [2,3].

In contrast to other concepts which place self-separating aircraft in segregated regions of airspace, AFR is by necessity a mixed-operations concept. Newly-introduced AFR aircraft must be capable of coexisting with conventional IFR traffic in the same airspace. This integration of self-separating and ground-separated traffic presents unique challenges for the design of the AFR concept.

One such challenge is related to the availability of intent information from IFR aircraft. While IFR aircraft are mandated to broadcast their current position and velocity, this broadcast does not include information about the aircraft's intent. By intent, this paper refers to the aircraft's planned trajectory changes in the near future, consisting of upcoming

turns and altitude changes within the next 10 minutes. Without this information, IFR trajectory changes are less predictable to AFR aircraft, creating the potential for more short-term AFR-IFR conflict situations. The availability of this intent information, and the consequences to AFR operations given varying levels of intent, are the focus of this paper.

This paper is organized as follows. First, background on the AFR concept and the problem of IFR intent information availability is presented. Then, we describe an unpiloted computer simulation experiment conducted in the NASA Langley Air Traffic Operations Laboratory to study the effects of varying levels of intent broadcast on the performance of AFR conflict detection and resolution. Finally, the results of this study are discussed along with conclusions and recommendations that can be made for sharing trajectory intent information in future NextGen designs for the National Airspace System.

Background

Definitions

Loss of separation. A loss of separation (LOS) has occurred between two aircraft when they are closer than a given separation standard—in this paper, 5 nmi laterally and 1000 ft vertically.

Conflict. Two aircraft are said to be in conflict when a loss of separation is predicted to occur between them in the future.

Conflict detection. An on-board airborne separation assurance system (ASAS) continually scans surveillance information from surrounding aircraft for conflicts within a given time horizon. This function depends on an accurate prediction of an aircraft's trajectory into the future.

Conflict resolution. Once a conflict is detected, the ASAS must guide the pilot to alter the ownship's trajectory so that a loss of separation is no longer predicted to occur. Conflict detection and resolution are often referred to together as "CD&R."

Strategic. When the autoflight system is fully coupled with the flight management system (FMS) (i.e., in LNAV and VNAV modes), this is referred to as strategic flight. Conflict resolutions accomplished via FMS route modifications are referred to as strategic resolutions.

Tactical. Complementary to strategic operations, tactical flight refers to modes were the autoflight system is decoupled with the FMS, such as track select and altitude hold modes. Tactical conflict resolutions are accomplished via mode control panel (MCP) altitude and track changes rather than FMS route modifications.

Autonomous Flight Rules

The reader is directed to [2,3] for a thorough history and description of the AFR concept. The fundamental elements of AFR are summarized here.

An operator files a flight plan under the rules of AFR as an alternative to an IFR or VFR flight plan. By filing AFR, the pilot agrees to take on the responsibility for maintaining the separation of the aircraft from all other traffic.

The AFR pilot relies on an on-board ASAS (decision support tool) to detect conflicts with other aircraft, and must follow the tool's guidance to alter course and promptly resolve those conflicts. The pilot must also use the tool to probe planned maneuvers and route changes for potential conflicts, ensuring that any move is sufficiently conflict-free before executing it.

In similar fashion, the pilot uses the on-board tool to avoid weather and restricted airspace and to fit into the ANSP's traffic flow management plans (e.g., to meet a required time of arrival constraint at a metering fix of the terminal airspace).

For conflicts between a pair of AFR aircraft, a set of right-of-way rules determines which aircraft acts first to resolve the conflict. For AFR-IFR conflicts, the AFR aircraft always gives way to the IFR aircraft, relieving the ANSP of the burden to monitor for and resolve conflicts with AFR aircraft.

The AFR operator, complying with these rules, is afforded a great deal of operational flexibility. AFR aircraft can fly at economical altitudes and speeds along optimal routes, bypass miles-in-trail restrictions and many ground delay programs, find better routes through weather, and fly with a freedom similar to VFR operations, but in instrument meteorological conditions. The benefits from reduced restrictions and increased operational flexibility are expected to significantly exceed the cost of equipage and the burden of taking on separation responsibility to fly AFR [4].

AFR as developed thus far is primarily an en-route concept of operations. An aircraft may depart under a conventional IFR clearance, transition to an AFR clearance for climb, cruise, and initial descent, and regain an IFR clearance upon entrance to the terminal area. In future work, it is envisioned that AFR will be integrated with flight deck interval management and spacing concepts for arrivals [5].

The word *autonomous* with respect to the AFR concept refers to the operator's autonomy to fly with great independence from outside supervision or control. It does not refer to the sense of "autonomous" machine control that is sometimes associated with robotic systems such as unmanned aerial vehicles. While there are parallels between aerial robotics and the automation tools that are enabling of manned autonomous flight, the pilot is firmly in command and in the loop with respect to the AFR concept.

AFR has been developed and evaluated in a number of studies over the past decade. Computer simulations have addressed the effects of traffic density and encounter geometry [6], pilot delay [7], ADS-B range and interference [8], wind uncertainty [9], and dynamic weather avoidance [10]. AFR-IFR mixed operations have also been studied through human-in-the-loop experiments from the perspective of both the aircrew and ground controllers [11].

ADS-B State and Intent Broadcasts

AFR is an application of broadcast information sharing. Direct, two-way, explicit communication is not required between AFR and traffic aircraft to perform this operation. Each AFR aircraft is equipped with a compatible ASAS that operates on the ADS-B OUT messages from surrounding traffic; in this fashion, the functions for separation assurance are distributed and coordinated implicitly by implementing the same procedures and compatible computer algorithms on each AFR aircraft.

The FAA 2020 ADS-B OUT mandate [12] requires all aircraft operating in Mode C transponder airspace to broadcast a state vector report via ADS-B OUT. This vector includes the latitude, longitude, altitude, airspeed, ground track, and other information necessary to define the state of the aircraft. State information is essential for self-separation, as it enables the basic awareness of traffic location and direction of flight.

The current state of an aircraft can be projected forward into the future along a steady-state trajectory for conflict detection. However, the state alone does not provide any information as to whether the aircraft is planning an upcoming turn or altitude change. Thus, in addition to the current state of a nearby aircraft, it is useful (but not necessarily required) to know what the aircraft is planning to do in the future; this is referred to as intent information.

The Minimum Aviation Performance Standards (MASPS) for ADS-B provides for two kinds of intent broadcast: short-term target state reports [13, §3.4.7] and long-term trajectory change reports [13, §3.4.8]. The target state (TS) report is depicted in Figure 1. When the aircraft is currently performing a maneuver, the TS report includes information about the intended state of the aircraft once the maneuver has completed. In the case of a turn, the TS report includes the target track angle, and in the case of a climb or descent, it includes the level-out altitude. The TS report is only useful during such a maneuver; when the aircraft is flying straight and level, the TS report is identical to the aircraft's current state vector.

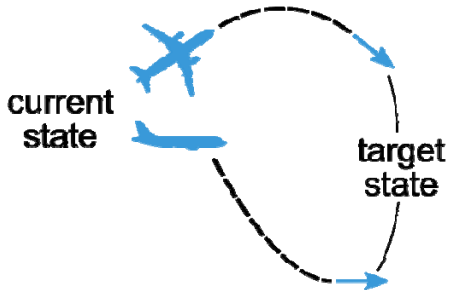

Figure 1. ADS-B state and target state reports.

Trajectory change (TC) reports are shown in Figure 2. Aircraft broadcast a trajectory change point (TCP) for each upcoming turn and altitude change, up to a fixed number of points. Each TCP includes information to define a turn or altitude change point in time and space. Due to bandwidth limitations in the 1090 MHz Mode S channel, these points are broadcast one-at-a-time, sequentially. After the last point in the sequence is broadcast, the cycle begins again with the first point. The time interval between the broadcast of these points is constrained by the channel bandwidth.

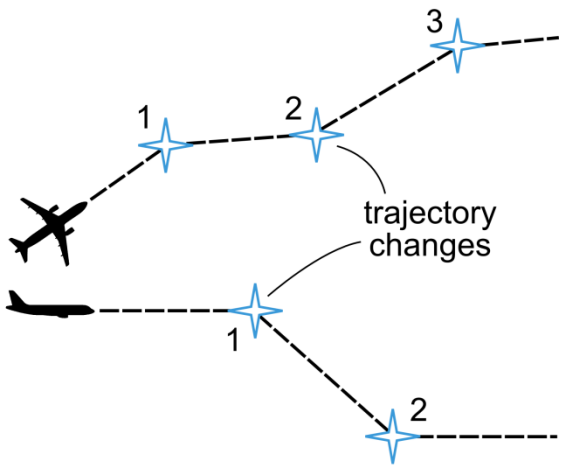

Figure 2. ADS-B trajectory change reports.

Each lateral TCP includes both the track-in angle and the track-out angle, so that the last waypoint in the sequence (such as point 3 in Figure 2) includes some information with which to project a path beyond itself. However, in the case of vertical waypoints, the current ADS-B TC report specification [13, §3.4.8] does not include a final altitude, vertical speed, or other indication that would help another aircraft to predict its vertical path given one point alone. Thus, while broadcasting only one TCP is sufficient to characterize one turn, it is not enough to define an upcoming altitude change. This limitation in the ADS-B specification has implications for the analysis of the results of the experiment described later in this paper.

Equipage for ADS-B Intent Broadcast

The ADS-B mandate requires aircraft to broadcast a state vector only; the TS intent broadcast is optional, and the TC message is not included in the mandate. While the equipage of AFR aircraft for intent broadcast is specified in the AFR concept and can be expected as part of the "price of admission" to achieve AFR benefits, it is more difficult to motivate IFR operators to equip. There are additional expenses to equip for intent broadcast over the mandated state-only broadcast. For example, intent information generally resides in separate equipment from the ADS-B transponder, such as in the FMS or a panel-mounted GPS navigation system, and data connections between these systems are not trivial.

These expenses compound the problem that IFR operators do not receive a clear and tangible benefit to equipping for intent broadcast: the intent information is to the benefit of the AFR operator, and it is not required for IFR operations. While it can be argued that there are numerous indirect benefits for IFR operators if AFR is successful, this is likely not to be a sufficient motivator.

There are other potential sources of IFR intent information. One NextGen development is the System Wide Information Management (SWIM) network, which is being designed to facilitate the exchange of a variety of information between National Airspace System users [14]. One SWIM component is the "flight object," which may eventually contain trajectory intent information for IFR flights as maintained by controllers and computers on the ground. It is conceivable that this intent information could be made available to an AFR consumer on the flight deck through a rebroadcast or on-demand uplink. IFR intent available through SWIM could take the place of ADS-B TC reports, without additional equipage by IFR operators.

Autonomous Operations Planner

Early work in this area suggested that a simple traffic display is not sufficient for human pilots to perform self-separation effectively. Support from computer automation on the flight deck is essential for solving complex traffic problems. Since 1998, NASA Langley has been developing a research prototype decision support tool for AFR operations known as the Autonomous Operations Planner (AOP).

AOP receives ADS-B state and intent reports from aircraft within surveillance range and continually scans for conflicts with traffic out to a specified look-ahead horizon (e.g., 10 minutes). In the event a conflict is detected, AOP alerts the aircrew and provides guidance for resolving the conflict. Depending on the situation, AOP may provide a strategic resolution via a modified FMS route, or it may give tactical commands to be flown using the MCP.

AOP also provides a conflict prevention capability for pilot-initiated maneuvers, probing FMS and MCP plans to ensure they are conflict-free before they are executed. In addition to protecting against conflicts, routes generated by AOP are optimized to maximize flight performance and to meet FMS trajectory constraints. In addition to traffic, AOP has

also been developed to take polygonal special-use airspace and convective weather areas into account.

A more complete description of AOP can be found in [15].

Simulation Environment

To investigate the impact of IFR intent availability on AFR conflict detection and resolution, a simulation was conducted in the NASA Langley Air Traffic Operations Laboratory (ATOL). The ATOL has facilities for conducting large-scale human-in-the-loop and unpiloted computer simulations of advanced air traffic management concepts.

In this simulation, each aircraft was implemented as one instance of the NASA Aircraft Simulation for Traffic Operations Research (ASTOR) platform. The ASTOR is a multi-fidelity flight simulator supporting realistic displays and controls representative of a Boeing 777, twin-engine transport aircraft performance model, fully-function autoflight system and FMS, and a digital avionics data bus emulating realistic internal and external data communications, including ADS-B IN and OUT.

Each ASTOR executes on a separate rack-mounted Windows PC blade server. ASTORs communicate with one another during the simulation using an implementation of the High-Level Architecture (HLA) standard over a local network. The running network of ASTOR simulators is referred to as the Airspace and Traffic Operations Simulation (ATOS).

Each ASTOR contains a separate instance of AOP performing conflict detection and resolution. The aircraft is flown by a rule-based "pilot model" software agent, which is responsible for following AOP's guidance in the event a conflict resolution is required.

The simulated experiment airspace is shown in Figure 3, and the progression of a typical flight is shown in Figure 4. Aircraft are instantiated at random points on the 320 nmi diameter outer circle (A). Each aircraft is randomly assigned a route that crosses the inner 160 nmi experiment area circle, and onward to a metering fix and destination airport beyond.

AOP is permitted a "grace period" to ignore conflicts between A and B because those conflicts would have been detected and resolved in the unsimulated flight prior to arriving at A. Approximately 10 minutes of flying time is allowed before reaching B. Flying time across the inner circle (between B and C) is approximately 20 minutes. Upon reaching the inner circle boundary again (C), the flight is terminated.

The size of the experiment area was chosen to give approximately thirty minutes of flying time to each ASTOR (from A to C). This gives each aircraft a chance to perform typically one complete conflict detection and resolution cycle during the simulation; sometimes several.

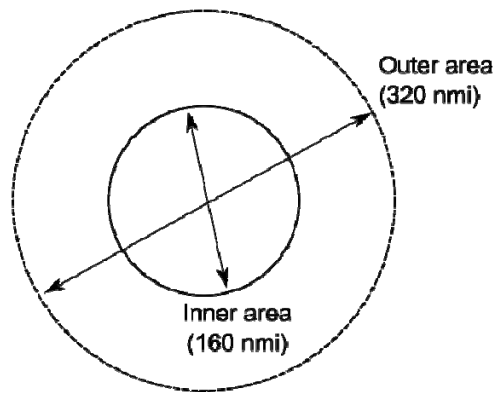

Figure 3. Simulation area.

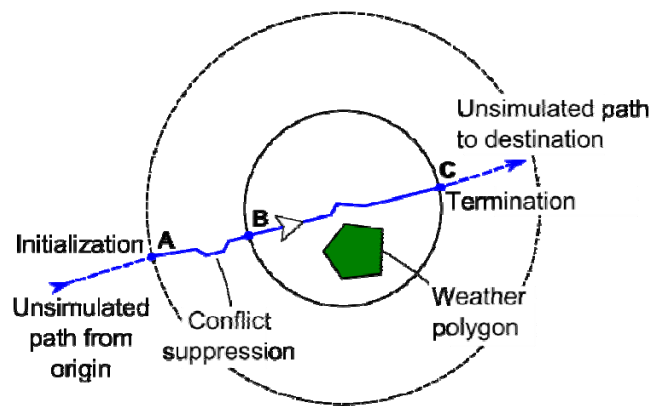

Figure 4. Sample flight.

Aircraft are initialized in cruise at altitudes uniformly distributed between FL 290 and 390. Aircraft generally remain in cruise throughout the simulation, although they may occasionally execute step climbs and descents for conflict resolution.

Aircraft are generated at a rate to create a density of 60–70 aircraft in the inner circle. Based on the size of the experiment area, this density is roughly 2X the level of traffic of high-altitude sectors in today's NAS. The purpose of the simulation environ-

ment is not to match the pattern of traffic or traffic density in a particular airspace; rather, it is to create enough traffic so as to exercise AOP's CD&R algorithms under random conflict geometries and surrounding traffic patterns.

Each simulation included a moving weather cell represented by a translating and morphing avoidance area polygon. Although this study was not focused on weather avoidance, the weather polygon provided a disturbance to encourage aircraft trajectory changes, which are important for this study. The pilot model used AOP to avoid the weather polygon as it does for traffic.

Each simulation case executed for 2 hours and 20 minutes. In this time, roughly 550 aircraft were simulated. As aircraft flew through the experiment and terminated, the ASTOR was recycled for a new flight. In this way, the simulation was accomplished with a set of approximately 120 machines. The size of the simulation was restrained primarily by the number of machines and the duration of all of the simulation cases, which ran in real time.

To study mixed operations, it is necessary to simulate both AFR and IFR aircraft. AFR operations are well represented by the pilot model and AOP. However, no ground-side for IFR operations was simulated here. In lieu of a true IFR simulation, a pseudo-IFR simulation capability was employed, depicted in Figure 5.

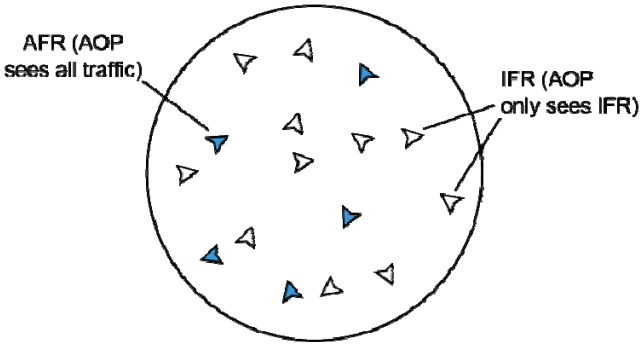

Figure 5. Pseudo-IFR simulation approach.

In this approach, "IFR" aircraft in the simulation used AOP and the pilot model in the same way as the AFR aircraft. However, "IFR" aircraft received only the ADS-B messages from other "IFR" aircraft; they were essentially ignorant of AFR traffic. In this way, the "IFR" aircraft remain separated from each other, as would aircraft managed by the ground, while AFR-IFR and AFR-AFR separation is performed only by AFR aircraft.

This environment represents a very rough approximation of mixed operations, but it is considered acceptable for studying the mixed-operations performance of just the AFR aircraft. While the pseudo-IFR aircraft in the simulation are not performing true IFR operations with a ground controller in the loop, their behavior is suitable to serve a specific purpose in this study: to provide IFR targets with some element of trajectory unpredictability with which to evaluate AFR conflict detection and resolution.

For the remainder of this paper, the pseudo-IFR aircraft will be referred to simply as IFR.

Experiment Design

Research Question

This experiment was designed to answer the following general research question: What are the effects of varying the availability of IFR aircraft intent information on the safety and efficiency performance of AOP conflict detection and resolution for AFR aircraft? In particular, availability of intent information is defined as the number of TCPs shared and the time interval between each TCP broadcast.

Independent Variables

Number of TCPs broadcast by IFR aircraft. The most important factor in defining the amount of intent information available is the maximum number of TCPs broadcast via ADS-B in each cycle. This independent variable was tested with values of 0, 2, 3, and 4 TCPs. In the zero TCP case, The IFR aircraft broadcast only the state vector report. The maximum setting of four TCPs was chosen here based on experience gained in early testing that found diminishing returns on the benefit of broadcasting more than four TCPs. Only the number of IFR TCPs was varied here; the number of TCPs broadcast by each AFR aircraft was fixed at four.

Time interval between TCP broadcasts. The other parameter defining availability of intent information is the time interval between TC message broadcasts and thus the total time to receive a complete intent update. Interval settings of 3, 10, and 20 seconds were tested.

The shortest interval of three seconds was selected based on an analysis of the maximum number of Mode S extended squitters (messages) per second available for ADS-B TCPs in the Minimum Operational Performance Standards (MOPS) for ADS-B [16]. The MOPS places a cap of 6.2 extended squitters per second (ES/s) on the rate of ADS-B message transmissions [16, §2.2.3.3.1.3]. 5.4 ES/S of this allowance is used for other messages (such as the state vector and aircraft status messages), leaving 0.8 ES/s available for TC messages [16, §O.4]. This budget is halved since a complete set of TC information requires both a "basic" and "supplemental" message, giving a TCP budget of 0.4 ES/s. The inverse of the message rate gives the interval of (1 / 0.4) = 2.5 s/ES. This value was rounded up to 3 s in this study, giving the minimum setting for the TCP broadcast interval.

The interval of 10 s is a maximum interval based on a requirement on the allowable amount of time to receive TCP changes in the MASPS [13]. For a worst-case head-on encounter between two aircraft, the MASPS makes a safety-based suggestion that the first TCP be received within 41 s of entering an assumed 90 nmi ADS-B reception range [13, §N.11.1]. Since the aircraft are broadcasting a maximum number of four TCPs in this simulation, the maximum allowable interval between TCPs in order to meet this recommendation is about 10 s. This was chosen as the intermediate setting for the interval in this study.

Finally, 20 s was chosen as a large value for the purpose of probing the limits of AOP's CD&R performance with long delays between TCP broadcasts. Only the interval for IFR broadcasts is varied in this experiment; the AFR TCP broadcasts are fixed with a 3 s interval.

Dependent Variables

While data was recorded for all aircraft in the simulation, the results presented in the following section focus only on the AFR aircraft and their interactions with IFR aircraft. Dependent variables related to conflicts were only analyzed for those encounters between an AFR and an IFR aircraft; AFR-AFR and IFR-IFR conflicts are not considered here. For the other variables related to the overall flight, only the AFR aircraft were considered.

Number of traffic conflicts between AFR and IFR aircraft. During each simulation, all AOP conflict detections are recorded, and these events may be counted to give an indication of the amount of activity required by the AFR aircraft in performing self-separation.

Having more conflicts leads to more pilot alerts, more required AOP actions, and thus a greater workload on the aircrew. More conflict alerts may reduce the number of "missed detections" by providing more opportunities for the aircrew to employ AOP in ensuring separation. By contrast, more conflict alerts may also increase the number of "false detections," wherein the crew is prompted to maneuver unnecessarily for a conflict that has not been accurately predicted.

Number of losses of separation between AFR and IFR aircraft. In contrast to the number of conflicts which only measures *predicted* LOSs, the simulation also counts *actual* LOSs that occur. This number is directly related to the level of safety of the overall AFR operation, and it is the goal of the AFR concept, procedures, and the AOP tool to drive the number of losses of separation to as close to zero as possible, while still enabling an efficient and beneficial mode of operations to the user.

Time remaining to first loss of separation at initial conflict detection. At the first detection of a conflict, there is a certain time remaining until the aircraft are predicted to lose separation. This essentially represents the amount of warning time that the pilot has to act to resolve the conflict. Having less time-to-LOS gives the pilot less time to act to resolve the conflict, possibly leading to more tactical maneuvering.

Percentage of traffic conflicts resolved tactically. Most conflicts are resolved through some pilot action. Depending on the situation, the conflict may be resolved strategically through a modified FMS route, or tactically by MCP direction (e.g., a 15° heading change). The type of resolution is recorded for each AFR-IFR conflict. Tactical conflict resolution is more workload intensive and results in a temporary loss of the aircraft's predictability as perceived by neighboring traffic, since the it requires the aircraft to depart the FMS route.

Percent time in tactical flight mode. The flight mode of each AFR aircraft over time is recorded, and it can be determined how much time the autoflight system was coupled vs. uncoupled with the FMS. Again, strategic flight is preferred over tactical flight,

for reasons of pilot workload, trajectory predictability, as well as the fact that flying coupled with the FMS allows the pilot to take advantage of the flight optimization features of the FMS and AOP.

Lateral flight-path deviation. For each AFR aircraft as-flown path length from start to finish can be compared to the aircraft's original route at initialization. The deviation represents the amount of additional distance flown on the ground due to traffic and weather avoidance during the simulation. Increased flight-path deviation is an indication of inefficiency induced on the performance of the flight due to the presence of other aircraft and the required execution of conflict resolution maneuvers.

Full Factorial Design

The experiment was designed around two independent variables: the number of IFR TCPs and the IFR TCP broadcast interval. A *test power* analysis was performed to evaluate several experimental design options. The *test power* is the ability for the experiment to detect the magnitude of the "signal" of each dependent variable amidst sources of random noise; i.e., the power to detect the signal-to-noise (SNR) ratio. A higher power is preferred; it is an indication of how well the design is able to detect the SNR correctly. For example, a power of 80% for an SNR of 2:1 means that there is an 80% chance that the design will correctly detect that the signal is larger than the random noise by a factor of 2. In many fields of research, an SNR of 2:1 is common, and a power of at least 80% is preferred.

A two-factor factorial design with four settings for the number of IFR TCPs (1, 2, 3, and 4) and three TCP interval settings (3, 10, and 20 s) was evaluated (the interval is irrelevant for the zero TCP case where no TCPs are broadcast). Figure 6 shows the test power for several design options, given an SNR of 2:1.

Each combination of independent variables requires a certain number of simulation repetitions in order to provide a given experiment power. Figure 6 shows that to achieve at least 80% power, five repetitions are required; this number of repetitions was chosen for this experiment. The zero TCP case with five repetitions was then added to this design, resulting in a total of 65 simulation runs. Another power analysis was performed to ensure that the test power for the zero and non-zero TCPs was at least 80%. This second power analysis on the one-factor design of the number of TCPs confirmed a test power of 83% with an SNR of 3:1. This process ensured that the selected design balanced available computing time and the quality of the data to meet statistical testing criteria.

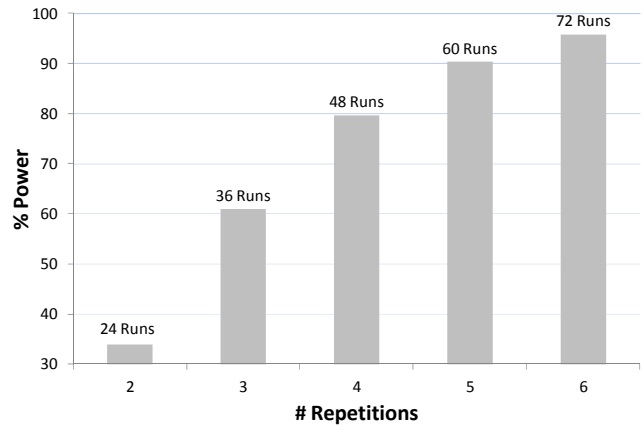

Figure 6. Experiment design test power analysis.

The one TCP case presented a problem for the analysis of the results. As discussed previously, broadcasting only a single TCP is insufficient to characterize climbs and descents, making this setting unsuitable for consideration here. For these reasons, the one TCP cases are omitted from further presentation in this paper, even though they were part of the original experiment design and data collection. The test power of the experiment remains unimpacted by this omission.

The final experiment design, consisting of ten independent variable combinations at five repetitions each, is shown in Table 1.

Table 1. Experiment runs.

Scenario	# IFR TCPs	IFR TCP Interval [s]	Reps
1	0	n/a	5
2	2	3	5
3	3	3	5
4	4	3	5
5	2	10	5
6	3	10	5
7	4	10	5
8	2	20	5
9	3	20	5
10	4	20	5

Results

Number of Conflicts

Figure 7 shows a box plot of the number of conflicts per flight hour across all experiment repetitions. The box plot gives a visual presentation of the center, spread, and outliers for each data set. The 25th and 75th percentiles of the data are shown by the ends of the rectangular box. A line is drawn through the middle the box at the 50th percentile (median). Dashed lines extend from the box to the extreme values.

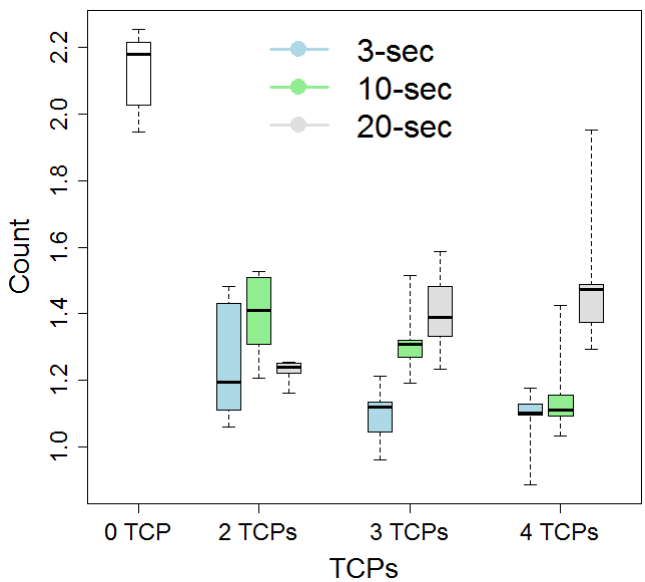

Figure 7. No. of traffic conflicts per flight hour.

One-way ANOVA revealed a significant effect of the presence vs. absence of IFR TCPs on the number of traffic conflicts ($p < 0.0001$). There are fewer AFR-IFR conflicts detected by AFR aircraft when IFR aircraft trajectory intent is available (TCP > 0) than when it is not (TCP = 0): roughly 2 conflicts per hour without TCPs, and between about 1 and 1.5 conflicts per hour with TCPs. With IFR TCPs available, an AFR aircraft has more information about IFR trajectory changes, and AOP can suggest improved routes to better avoid IFR traffic during conflict resolutions and weather avoidance route changes. Fewer conflicts means fewer required AFR pilot actions, reduced workload, and fewer chances for error. Fewer conflict resolutions leads to fewer trajectory changes and therefore a more stable airspace.

For the cases with more than one IFR TCP, two-way ANOVA indicated no significant main effect for the number of TCPs ($p = 0.213$), indicating that the effect of this variable must be considered in its interaction with the broadcast interval.

ANOVA showed a significant main effect of the broadcast interval by itself ($p < 0.0001$); in general, shorter broadcast intervals yield fewer conflicts. This is intuitive; with less delay in the intent information shared by IFR aircraft, the AFR AOP is better able to suggest conflict-free routes.

Finally, ANOVA revealed a significant interaction between the number of IFR TCPs and the broadcast interval ($p = 0.002$). The interaction is roughly as follows: When the interval is lower (3 or 10 s), adding more TCPs to the IFR broadcast is beneficial; it reduces the conflict count. However, the trend is reversed when the interval is 20 s. With this interval setting, adding TCPs noticeably increases the number of conflicts. This is counter to the intuition that adding TCPs and thus providing more information about IFR intent to the AFR aircraft should improve the operation.

This effect can be explained as follows. With a sequence of four TCPs, each TCP is broadcast in turn with the given delay. With a 20 s interval, it follows that it can take up to $20 \times 4 = 80$ s for the sequence to restart and re-broadcast the first TCP. Compare this to a 3 s interval, where it takes only 12 s to complete and restart the sequence. The effect is, for more TCPs in the sequence, the AFR aircraft must wait longer to receive a complete set of TCPs. This is compounded with the fact that the AFR aircraft will not always receive each TCP due to ADS-B signal interference and other disturbances [17]; it may occasionally have to wait for one or more subsequent broadcast cycles to receive a complete set of TCPs.

The additional time to receive a complete TCP set is related to the observation that the first and second TCPs in the sequence are more important to the AFR aircraft than the third and beyond. The first and second TCPs represent the nearest upcoming trajectory changes for the IFR aircraft, and they are more relevant in the time horizon of conflict detection and resolution. The ultimate conclusion to be drawn from this interaction between the number of TCPs and the broadcast interval is that it is better to have fewer TCPs at a shorter interval than more TCPs at a longer interval.

Losses of Separation

Table 2 shows the number of losses of separation across the runs as compared with the total flight time of AFR aircraft and the number of AFR-IFR conflicts. The LOSs are rare events in the simulation. While there are not enough of these events to draw statistically significant conclusions about the effects of the independent variables, the numbers generally appear to be consistent with the conclusions about the conflict counts in the previous section.

Table 2. Losses of separation.

IFR TCPs	IFR TCP Interval [s]	AFR Flight Hours	AFR/ IFR Conflicts	LOS
0	n/a	528.7	1122	21
2	3	521.0	655	3
	10	546.5	758	9
	20	550.5	674	8
3	3	554.2	605	3
	10	538.7	711	4
	20	550.9	722	13
4	3	561.9	605	0
	10	563.9	654	2
	20	545.5	826	29

LOSs are caused by a number of factors. As only AFR-IFR LOSs are shown here, chief among the contributors is the fact that IFR aircraft make trajectory changes without regard to AFR aircraft. Additionally, encounters between climbing and descending aircraft tend to be troublesome due to the increased difficulty in predicting vertical trajectories. Many LOSs can also be explained where the pilot model's rule-based logic fails in a situation that could be managed by a human pilot. All LOSs observed in simulation, however rare, feed into the improvement of the AFR concept and the AOP software to drive this number as close to zero as possible.

Time to First Loss of Separation

Figure 8 shows the box plot for the median time to first loss of separation at initial conflict detection, across all experiment runs for AFR–IFR conflicts. Median analysis was used to reduce the effect of outliers on the results. As before with the number of conflicts, there is a significant, beneficial effect of the presence of IFR TCPs on the time to LOS ($p < 0.0001$). When broadcasting IFR TCPs, the median time to LOS is between 20 and 50 s greater than without TCPs (7.7 min vs. 8 to 8.5 min). This indicates that the broadcast of IFR TCPs provides an overall increase in the amount of warning time LOS when a conflict is initially detected. Additional warning time provides more time to resolve the conflict and the potential for more efficient resolution maneuvers.

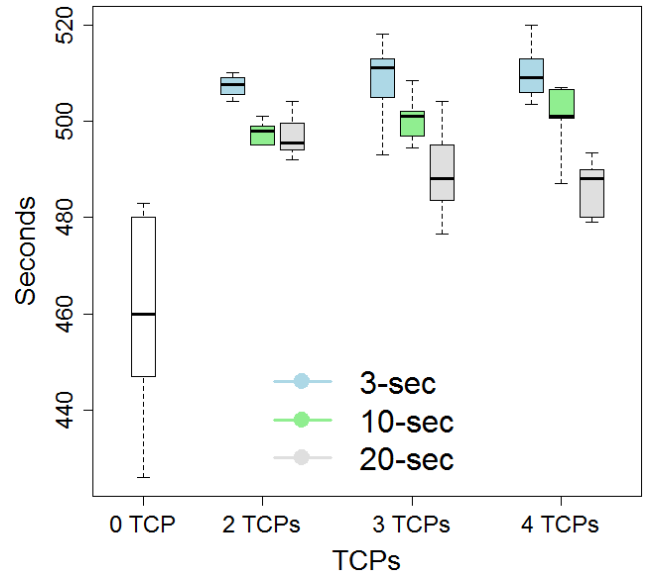

Figure 8. Median time to first loss of separation at initial conflict detection.

The main effect of the broadcast interval is also significant; a shorter interval leads to a greater time to first LOS ($p < 0.0001$), following intuition. Loosely, the amount of time saved by not waiting longer to receive critical TCPs translates into extra warning time on average.

The main effect of the number of TCPs, along with the interaction between TCPs and interval, are not significant ($p = 0.783$ and 0.139 respectively), indicating that although having TCPs is better than not, no significant difference was found between 2, 3, or 4 TCPs on the warning time to the predicted LOS. In other words, the warning time enhancement benefit from the availability of intent information was achieved with just two TCPs.

In addition to the median, the distribution of the time to LOS can also show the benefit to intent availability. For example, with zero TCPs, there are 24% of conflicts that have a time to LOS less than five minutes, compared to 10% to 19% when intent is

available. Intent availability significantly reduces the number of conflicts that appear with short warning times.

Percentage of Conflicts Resolved Tactically

Another indicator of CD&R performance is whether a conflict is resolved strategically (via FMS routes) or tactically (via MCP track and altitude changes). In this simulation, tactical resolutions were triggered for unresolved conflicts with less than five minutes remaining to LOS. This box plot is shown in Figure 9.

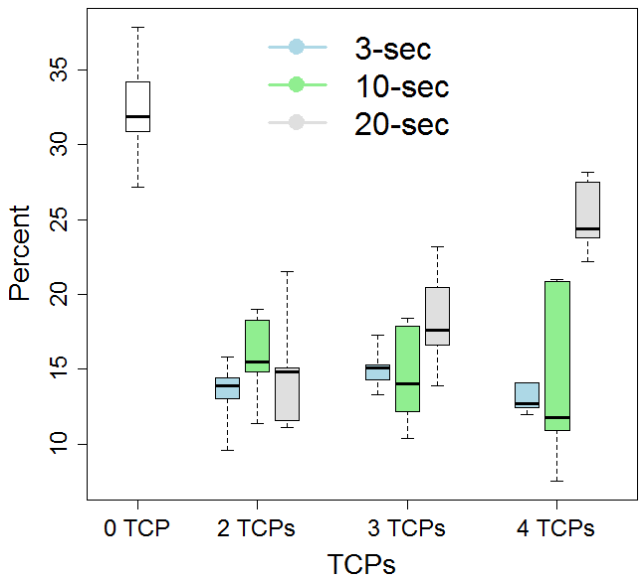

Figure 9. Percentage of conflicts resolved tactically.

The effects here are similar to those for the number of conflicts discussed previously. The benefit of having vs. not having IFR TCPs is significant ($p < 0.0001$). Roughly 30% of conflicts are resolved tactically when no IFR TCPs are broadcast vs. approximately 10% to 20% of conflicts when IFR TCPs are available. Fewer tactical resolutions generate less work for the pilot; in many situations it is preferred to use strategic (FMS-based) route changes. The strategic flying mode is also indicative of less "urgency" in pilot interactions with AOP. Tactical maneuvering also makes the AFR aircraft less predictable to other AFR aircraft in the vicinity because future tactical turns are not known and thus cannot be broadcasted.

The main effect of the TCP broadcast interval is also significant ($p = 0.005$), indicating fewer tactical resolutions when the interval is shorter. The interaction between the number of TCPs and the interval is also significant ($p = 0.026$), and reveal that it is better to have fewer TCPs at a shorter interval rather than more TCPs at a longer interval, as discussed before.

Percent Time in Tactical Flight Mode

In addition to conflict resolutions, the amount of tactical vs. strategic operation can be investigated via the amount of time spent flying in either mode, as shown in the box plot in Figure 10. As with the previous figures, the presence of IFR TCPs has a significant effect on this variable ($p = 0.003$). With IFR TCPs, AFR aircraft spend on average more than 6% of their time flying tactically, vs. 2% to 4% when flying strategically and coupled with the FMS. No other effects were statistically significant here.

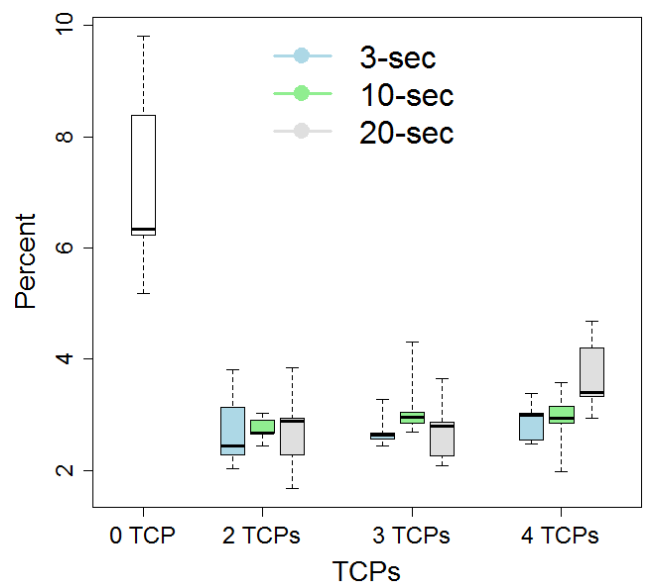

Figure 10. Mean percent time flying tactically.

The effects observed for this variable are likely due to the same causes as for the percentage of conflicts resolved tactically; more tactical resolutions by definition leads to more time spent flying tactically.

Lateral Flight-Path Deviation

Finally, we analyzed a metric related to flight efficiency: the median AFR lateral flight-path deviation, shown in Figure 11. This variable represents the amount of excess path length of the aircraft from initialization to destination when the recorded flight is compared with its initial trajectory. Due to traffic and weather avoidance maneuvers during the simulation, aircraft will experience this path length deviation.

Excessive path length deviation is an indicator of flight inefficiency.

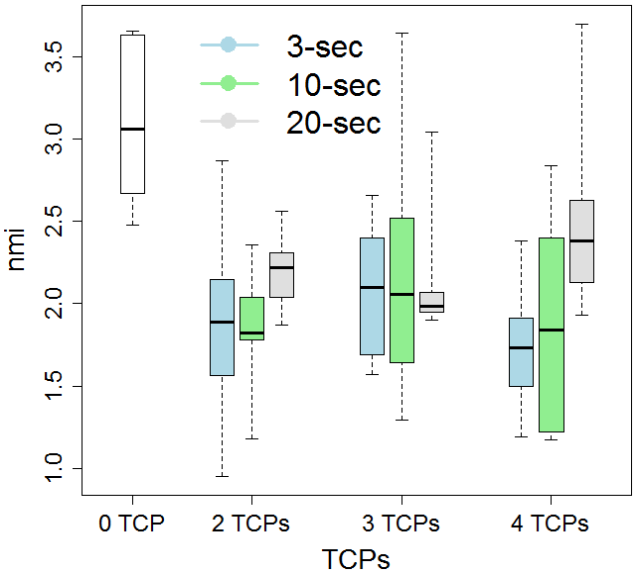

Figure 11. Median lateral flight-path deviation.

The presence of TCPs is significant as in all of the other results presented ($p < 0.033$). Without IFR TCPs, the median flight-path deviation is roughly 3 nmi vs. 1 to 2.5 nmi with IFR TCPs during the roughly 240 nmi flight across the simulation area. The potential cost savings to be had from the availability of intent information are appreciable, although the results here should be regarded in the context of the high traffic density and non-realistic traffic patterns in this simulation.

No other effect of the number of TCPs and the interval was found to be significant on this variable.

Discussion and Conclusions

The results presented show the benefit of IFR intent information in the form of ADS-B trajectory change reports to the conflict detection and resolution performance of an AFR aircraft using the Autonomous Operations Planner. In this mixed-operations simulation, the availability of IFR TCPs reduced the number of AFR-IFR conflicts per hour, increased the amount of warning time before loss of separation when those conflicts were first detected, and reduced the amount of tactical maneuvering. On the flight efficiency side, in cases where IFR TCPs were available, the availability of IFR intent information reduced the AFR aircraft flight-path deviation due to conflict resolutions.

It was found that two or more IFR TCPs is better than none at all, little benefit was found to having additional TCPs beyond two. In fact, if the interval between TCP broadcasts is long enough, increasing the number of TCPs in the sequence can be detrimental to CD&R performance.

While the results demonstrate some benefit of IFR intent availability for AFR operation, they do not show that a solution based solely on ADS-B state reports is infeasible. The experiment showed that safe operations are achievable with no IFR intent information, as evidenced by the low number of losses of separation in those cases. Improvements to separation assurance procedures and computer algorithms may be needed to handle unannounced IFR trajectory changes that could result in very-short-notice conflicts.

The value of making IFR intent information available to airborne AFR consumers as shown in this experiment motivates the inclusion of this information as part of the ADS-B message set, or by an uplink from a ground-based service such as the System Wide Information Management network. As only a small amount of information appears to be required (in the form of the first two trajectory change points), the bandwidth impact of this recommendation should be minimal.

References

[1] Cotton, William B., Aug. 1965, “New Directions in Air Traffic Control at Kennedy Airport”, Master's thesis, Course XVI, M.I.T.

[2] Wing, David J. and William B. Cotton, Nov. 2011, “Autonomous Flight Rules: A Concept for Self-Separation in U.S. Domestic Airspace,” NASA/TP–2011-217174.

[3] Wing, David J. and William B. Cotton, Sep. 2011, “For Spacious Skies: Self-Separation with ‘Autonomous Flight Rules’ in US Domestic Airspace,” AIAA Paper 2011-6865, 11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference.

[4] Stouffer, Virginia, Shahab Hasan, and Daniel Kozarsky, Sep. 2004, “Initial Life-Cycle Cost/Benefit Assessments of Distributed Air/Ground Traffic Management Concept Elements,” AIAA Paper 2004-6452, 4th AIAA Aviation Technology, Integration, and Operations (ATIO) Forum.

[5] Swieringa, Kurt, Jennifer Murdoch, Brian Baxley, and Clay Hubbs, Sep. 2011, "Evaluation of an Airborne Spacing Concept, On-board Spacing Tool, and Pilot Interface," AIAA Paper 2011-6902, 11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference.

[6] Consiglio, Maria C., Sherwood T. Hoadley, David J. Wing, and Brian T. Baxley, Sep. 2007, "Safety Performance of Airborne Separation: Preliminary Baseline Testing," AIAA Paper 2007-7739, 7th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference.

[7] Consiglio, Maria, Sherwood Hoadley, David Wing, Brian Baxley, and Danette Allen, Sep. 2008, "Impact of Pilot Delay and Non-Responsiveness on the Safety Performance of Airborne Separation," AIAA Paper 2008-8882, 26th Congress of International Council of the Aeronautical Sciences (ICAS).

[8] Idris, Husni, Maria C. Consiglio, and David J. Wing, Nov. 2011, "Surveillance Range and Interference Impacts on Self-Separation Performance," IEEE/AIAA 30th Digital Avionics Systems Conference (DASC).

[9] Consiglio, Maria, Sherwood Hoadley, and B. Danette Allen, Jun. 2009, "Estimation of Separation Buffers for Wind-Prediction Error in an Airborne Separation Assistance System," 8th USA/Europe Air Traffic Management Research and Development Seminar (ATM2009).

[10] Consiglio, Maria C., James P. Chamberlain, and Sara R. Wilson, Oct 2011, "Integration of weather avoidance and traffic separation," IEEE/AIAA 30th Digital Avionics Systems Conference (DASC).

[11] Barhydt, Richard, Parimal Kopardekar, Vernol Battiste, Nathan Doble, Walter Johnson, Paul Lee, Thomas Prevot, and Nancy Smith, Jun. 2005, "Joint NASA Ames/Langley Experimental Evaluation of Integrated Air/Ground Operations for En Route Free Maneuvering," 6th USA/Europe Air Traffic Management Research and Development Seminar (ATM2005).

[12] Department of Transportation, Federal Aviation Administration, May 2010, "Automatic Dependent Surveillance–Broadcast (ADS-B) Out Performance Requirements To Support Air Traffic Control (ATC) Service; Final Rule," 14 CFR Part 91.

[13] RTCA SC-186, Jun. 2002, "Minimum Aviation System Performance Standards For Automatic Dependent Surveillance–Broadcast (ADSB)," RTCA/DO-242A.

[14] Federal Aviation Administration, May 2007, "System Wide Information Management (SWIM) Final Program Requirements, Segment 1."

[15] Karr, David, Robert Vivona, David Roscoe, Stephen DePascale, and Maria Consiglio, Aug. 2009, "Experimental Performance of a Genetic Algorithm for Airborne Strategic Conflict Resolution," AIAA Paper 2009-5748, AIAA Guidance, Navigation, and Control Conference.

[16] RTCA SC-186, Dec. 2009, "Minimum Operational Performance Standards for 1090 MHz Extended Squitter Automatic Dependent Surveillance–Broadcast (ADS-B) and Traffic Information Services (TIS-B)," RTCA/DO-260B.

[17] Chung, William W. and Ronald Staab, Aug. 2006, "A 1090 Extended Squitter Automatic Dependent Surveillance – Broadcast (ADS-B) Reception Model for Air-Traffic-Management Simulations," AIAA Paper 2006-6614, AIAA Modeling and Simulation Technologies Conference and Exhibit.

Email Addresses

timothy.a.lewis@nasa.gov
nipa.phojanamongkolkij@nasa.gov
david.wing@nasa.gov

2012 Integrated Communications Navigation and Surveillance (ICNS) Conference

April 24-26, 2012