HANDLING TRAJECTORY UNCERTAINTIES FOR AIRBORNE CONFLICT MANAGEMENT

Richard Barhydt, NASA Langley Research Center, Hampton, VA
Nathan A. Doble and David Karr, Titan Corporation, Hampton, VA
Michael T. Palmer, NASA Langley Research Center, Hampton, VA

Abstract

Airborne conflict management is an enabling capability for NASA's Distributed Air-Ground Traffic Management (DAG-TM) concept. DAG-TM has the goal of significantly increasing capacity within the National Airspace System, while maintaining or improving safety. Under DAG-TM, "autonomous" aircraft maintain separation from each other and from "managed" aircraft unequipped for autonomous flight.

NASA Langley Research Center has developed the Autonomous Operations Planner (AOP), an onboard decision support system that provides airborne conflict management (ACM) and strategic flight planning support for autonomous aircraft pilots. The AOP performs conflict detection, prevention, and resolution from nearby traffic aircraft and area hazards. Traffic trajectory information is assumed to be provided by Automatic Dependent Surveillance Broadcast (ADS-B).

Reliable trajectory prediction is a key capability for providing effective ACM functions. Trajectory uncertainties due to environmental effects, differences in aircraft systems and performance, and unknown intent information lead to prediction errors that can adversely affect AOP performance.

To accommodate these uncertainties, the AOP has been enhanced to create cross-track, vertical, and along-track buffers along the predicted trajectories of both ownship and traffic aircraft. These buffers will be structured based on prediction errors noted from previous simulations such as a recent Joint Experiment between NASA Ames and Langley Research Centers and from other outside studies. Currently defined ADS-B parameters related to navigation capability, trajectory type, and path conformance will be used to support the algorithms that generate the buffers.

Introduction

The National Aeronautics and Space Administration (NASA) is studying a new concept of operations for the National Airspace System known as Distributed Air-Ground Traffic Management (DAG-TM) [1]. This concept is designed to significantly improve capacity, while maintaining or improving safety.

One component of DAG-TM, *En Route Free Maneuvering*, represents a paradigm shift between a centralized ground-based system to a distributed system. Under this concept, flight crews of appropriately equipped "autonomous" aircraft fly under Autonomous Flight Rules (AFR). These aircraft are able to choose their own route and altitude, subject to maintaining separation from all other aircraft. Controllers continue to provide separation between "managed" aircraft unequipped for autonomous flight and traffic flow management services for all aircraft. Managed aircraft fly under conventional Instrument Flight Rules (IFR). Flow management is normally applied in the form of waypoint speed, altitude, and time constraints at a terminal area meter fix. Inside the terminal area, AFR flights revert to IFR.

Airborne Conflict Management (ACM) functions including conflict detection, prevention, and resolution are key enabling capabilities for DAG-TM. A conflict is defined as a predicted loss of separation (LOS) between the ownship and another aircraft. Conflict detection alerts pilots to the presence of a conflict, prevention alerts to a potential ownship maneuver that would cause a conflict if initiated, and resolution gives a recommended avoidance maneuver to resolve the conflict.

All ACM functions compare trajectories from the ownship and nearby traffic aircraft. Supporting information including state vector and intent are

presumed to be broadcast over Automatic Dependent Surveillance Broadcast [2]. Various uncertainties are associated with all trajectories and must be properly managed by onboard systems. The ACM system should accurately predict the trajectories within the limits of the information provided and properly cope with regions of uncertainty.

Because trajectory errors cannot be eliminated, an ACM system must be aware of their potential impact on any alerts or guidance provided to the pilot. NASA Langley Research Center has developed the Autonomous Operations Planner (AOP) as a prototype avionics system for AFR pilots [3]. The AOP interacts with the Flight Management System (FMS) and other aircraft systems to provide necessary conflict management and strategic flight planning capabilities. As part of these functions, a new AOP capability assesses features of both the ownship and traffic trajectories that influence the magnitude of along-path, lateral, and vertical uncertainties. It applies an appropriate safety buffer that varies along each trajectory and is based on the local three-dimensional uncertainty. All ACM functions then work with this modified trajectory.

AOP Conflict Management

To accommodate different types of flight operations and a range of conflict scenarios, the AOP provides both tactical and strategic conflict management capabilities. The tactical system works with the own aircraft's state vector and any available traffic aircraft intent sent over ADS-B. The strategic system considers the ownship's active four-dimensional trajectory and also uses available traffic intent. These paths usually consist of the FMS flight plan and include multiple target states. Activation of the tactical or strategic systems are dependent upon the time to predicted separation loss and the ownship's current operating flight mode [3].

Tactical Conflict Management

The tactical system is geared toward flight involving inner loop heading, speed, vertical speed, and altitude commands given by the pilot. When flying in these modes, pilots have normally left the programmed flight plan. Tactical control is often used for shorter term goals such as convective weather avoidance or simple maneuvers around traffic conflicts. During flight in a tactical mode such as "Heading Select" or "Altitude Hold", the AOP conflict management functions consider only the position and velocity of the ownship. They consider the four dimensional intent-based trajectories of nearby traffic, if available [4].

In addition to the current flight mode, the time to predicted separation loss also affects AOP choice of a tactical or strategic resolution. Tactical conflict resolutions in the form of heading, vertical speed, and altitude commands are offered when this time is less than approximately three minutes. Inside this point, there is little time for the pilot to consider and execute a flight plan amendment.

Trajectory uncertainties for tactical control can include position and velocity vector inaccuracies [2]. Due to the shorter time horizon of tactical operations, however (currently defined to be within 5 minutes), the AOP emphasizes trajectory uncertainty handling for its strategic system.

Strategic Conflict Management

A major anticipated benefit of DAG-TM En Route Free Maneuvering is the pilot's ability to optimize his route and altitude without coordinating with air traffic control. Outside the terminal area (where these operations occur), pilots are likely to use the FMS flight plan to control the aircraft. If an airspace or traffic conflict occurs in this case, the AOP proposes a flight plan change that resolves all conflicts, continues to meet flow management constraints assigned by ATC, and considers pilot efficiency preferences.

For all strategic conflict management functions, the AOP compares the four-dimensional *command trajectory* of both the ownship and traffic aircraft [4]. The *command trajectory* refers to the path the aircraft will fly if the pilot doesn't change any automation modes or settings actively supporting aircraft guidance. This path may include multiple flight mode transitions. For example, an aircraft may be flying an FMS descent with a flight control panel altitude between the aircraft's current altitude and the end of descent. Its descent will be limited by the selected altitude [5]. In this case, the

vertical command trajectory consists of the FMS descent until the aircraft reaches the flight control panel altitude and then level flight at that altitude until the end of the prediction.

Changes to the command trajectory normally result from a pilot input. However, a non-programmed mode transition may also occur that affects the command trajectory, such as reversion to speed priority on descent if the intended vertical path results in an over-speed condition. These intent changes can affect conflict predictability. In the former case, the intent may change as the pilot resets the flight control panel altitude during an FMS climb or descent. In the latter, a change to speed priority will take the aircraft off the planned descent profile. Despite these potential uncertainties, use of the command trajectory to represent aircraft intent was proposed by the FAA and Eurocontrol in a 2000 Technical Interchange Meeting [6] and is also supported by RTCA [2].

In addition to pilot control actions or autoflight mode reversions that may affect the intent-based trajectory, several factors can lead to differences between an aircraft's predicted trajectory and that actually flown. In order to study the DAG-TM concept, as well as the performance of contributing prototype avionics systems, NASA Langley has developed the Airspace and Traffic Operations Simulation (ATOS) hosted by the Air Traffic Operations Lab [7]. This simulation consists of 12 medium fidelity pilot workstations that incorporate the AOP. Information is shared over a High Level Architecture between subject pilot, pseudo-pilot, and ground stations (the latter located outside the ATOL). Within this architecture, AOP performance can be evaluated in the presence of several forms of trajectory uncertainty.

Sources of Trajectory Uncertainty

A recent study by Mondoloni and Bayraktutar identified several types of trajectory uncertainty [8]. Categories include environmental effects, variations in aircraft performance, and incomplete intent information. Those uncertainty types that are modeled within ATOS and therefore impact AOP conflict management performance are discussed below. Performance is evaluated based on the comparison of the predicted trajectory (as derived from command trajectory elements broadcast over ADS-B) with the actual trajectory.

Environmental

ATOS allows entry of a horizontal wind field at three dimensional grid points. This grid can also vary with time during the simulation.

When an aircraft is flying on an FMS flight plan, open loop trajectory change points (TCPs) such as top of climb and end of descent are affected by the wind field. If the pilot fails to enter the wind conditions into the FMS prior to a climb or descent or if the wind is different from that entered, the along track position or arrival time of these points will change from the original prediction. Since pilots fly at constant Mach or calibrated airspeed (CAS) during cruise, a change in wind will affect the arrival time at subsequent waypoints.

Aircraft Performance

The AOP uses a single aircraft performance model to re-construct the trajectories of traffic aircraft. During experiments conducted with the ATOS, there are typically several different aircraft models flying within the simulation. The Langley subject pilot and pseudo pilot aircraft models (the latter being developed by the National Aerospace Laboratory of the Netherlands [9]) have different characteristics. A recent Joint Experiment with NASA Ames resulted in a third aircraft model when Ames aircraft were added to the simulation [10].

Significant variability in path construction and guidance features between aircraft types [11] can lead to trajectory prediction errors. These differences can include turn anticipation at a waypoint, descent initiation near top of descent, and path construction between two waypoints containing speed and altitude constraints.

Incomplete Intent Information

Aircraft flying FMS descents within the ATOS typically broadcast one vertical TCP at the top of descent and another at the end of descent. Lacking any information about the middle of descent (such as the planned descent speed or deceleration points), AOP makes a straight line assumption.

This path assumes constant groundspeed and constant flight path angle between the two points.

Figure 1 shows a nominal FMS descent profile with a speed and altitude restriction at the end of descent (assumed to occur above the mandatory 250 knot speed restriction below 10,000 ft). A two point prediction is superimposed. Assuming the descent and cruise Mach are the same, the aircraft begins its descent at the top of descent. When the descent Mach equals the descent CAS, the aircraft tracks a constant CAS until reaching a deceleration point where the descent shallows prior to the end of descent.
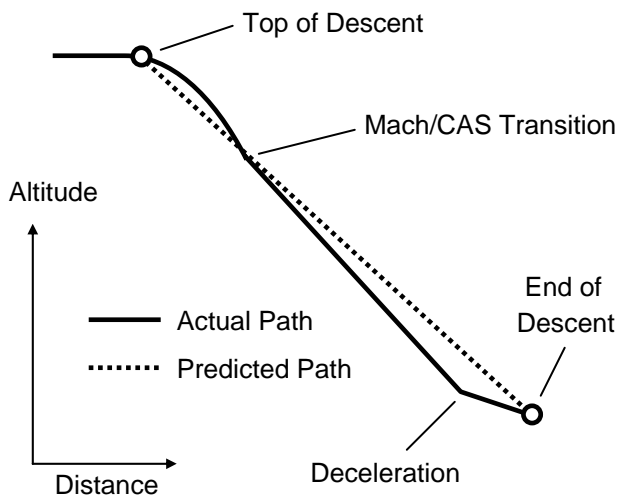

Figure 1. Nominal FMS Descent Profile Compared with Two Point Prediction

Path features such as a difference between cruise and descent Mach or a speed reduction to meet the 250 knot restriction may cause additional variation. These differences can be on the order of hundreds of feet.

AOP conflict prediction performance was analyzed for distinct cases that occurred during the Joint Experiment between Ames and Langley Research Centers. The experiment provided a good platform for comparing differences between predicted and actual trajectories and the operational impact of those differences for aircraft operating in a DAG-TM environment.

Joint Experiment

Design

The Joint Experiment between NASA Langley and Ames Research Centers evaluated air/ground coordination issues during DAG-TM en route and arrival operations. It also served as a test platform for the AOP conflict management and flight planning functions. Corresponding to the DAG-TM goal of substantially increasing traffic levels in a mixed equipage environment (autonomous and managed aircraft flying in the same airspace), the experiment included two primary factors:

- Mixed Operations (comparison of mixture of autonomous and managed aircraft in same airspace vs. managed aircraft alone, at the same traffic level).

- Scalability (addition of increasing levels of autonomous aircraft to level of managed aircraft occurring under current operations).

Traffic scenarios were run at three different traffic levels ("L1," "L2," and "L3"), with the L1 level repeated for the "all managed" and "mixed" autonomous/managed traffic conditions. This design led to a total of four experimental conditions. The L2 and L3 levels included the same number of managed aircraft as the L1 mixed condition, while adding autonomous aircraft. The L1 level was set to a challenging level based on all-managed operations. En route autonomous aircraft were added at L2 and L3, with L3 representing about twice current day capacity in the sectors modeled. The number of aircraft descending to the terminal area remained the same for all conditions.

The Langley ATOL and Ames Airspace Operations Lab were connected for the experiment and operations at both labs were conducted simultaneously as part of an overall traffic environment. Twelve subject pilots participated at Langley (each flying a pilot workstation equipped with AOP). Five subject controllers and nine subject pilots participated at Ames. Pseudo pilots monitored pre-programmed background aircraft added to achieve the desired traffic level. Subject pilots flew four different traffic scenarios for each of the four conditions, using a within-subjects design.

The experimental airspace modeled the Dallas Ft. Worth area and is shown in Figure 2. Subject controllers staffed the Amarillo, Ardmore, and Wichita Falls High Sectors and the Bowie Low sector. A pseudo controller handled each "ghost" sector and performed handoff duties to the subject controllers. Aircraft flew in the ghost sectors, but no data were collected there. Each subject pilot flew two overflights and two arrivals for each experimental condition. Overflights crossed either the Amarillo or Ardmore High sectors. For arriving aircraft, two streams began in level flight in Amarillo or Ardmore and both included a descent to cross the BAMBE meter fix at 250 knots and 11,000 ft. The subject controller used a scheduler to assign required times of arrival at BAMBE to each arriving aircraft. The scenario ended for each subject pilot when he or she crossed the high altitude sector for overflights or the BAMBE fix for arrivals. Controllers continued to work traffic until all subject pilots from Ames and Langley had finished.
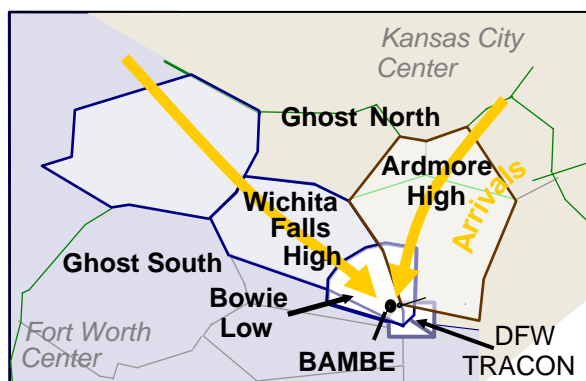

Figure 2. Joint Experiment Airspace

Pilot Responsibilities and Use of AOP

When flying as a managed aircraft (25% of all runs), pilots followed trajectory instructions from controllers using conventional IFR procedures. As in today's operations, the controller was responsible for separation assurance and traffic flow management under those conditions.

Consistent with AFR, autonomous aircraft pilots were required to maintain separation from all other aircraft and meet the assigned speed, altitude, and time constraints at the BAMBE meter fix. In addition, AFR pilots were not allowed to maneuver

in a way that would cause a near-term conflict for another aircraft (defined to be within 4 minutes to predicted LOS). To meet these responsibilities, pilots used the AOP conflict detection and prevention functions to recognize and prevent conflicts, respectively. For the descent scenarios, the AOP integrated with the FMS to provide conflict resolutions that met the BAMBE meter fix constraints. During en route scenarios, the AOP looked for minimum path changes to the flight plan.

Because AFR pilots were solely responsible for resolving all traffic conflicts, reliable AOP conflict management performance was required to ensure separation. A missed conflict detection would likely result in a separation violation because controllers were not notified of AFR/AFR conflicts. As a safety buffer, they were alerted to a near-term AFR/IFR conflict, but were not required to take action to resolve it.

At the time the experiment took place, conflict prediction was based solely on a comparison between the predicted trajectories of the ownship and traffic aircraft. Pilots were only notified of a conflict if the two aircraft were predicted to be within the minimum required separation between the current position and the look-ahead horizon (10 min). No capabilities existed to make a more conservative conflict assessment in regions of high uncertainty.

Performance Metrics

The primary overall experimental performance metrics were the occurrence of separation violations (subject pilot aircraft within 1000 ft and 5 NM of traffic) and the pilot's ability to meet the required speed, altitude, and time restrictions at the terminal area meter fix. The latter were only applicable to descending aircraft. These results as well as other metrics related to DAG-TM concept feasibility are annotated separately [12].

Relative to AOP conflict management performance, comparisons were made between the predicted and actual trajectories. In a few cases where a separation violation was attributed to a conflict prediction error, a detailed analysis was performed to determine the error category and likely cause. Two such cases are described below.

AOP Conflict Prediction Case Studies

In both cases described below, the ownship and intruder aircraft were predicted to maintain adequate separation. Therefore, the pilot was not alerted to an impending conflict. The actual trajectories differed from the predictions sufficiently to cause an LOS.

Case 1

In the first case, the ownship (a Langley subject pilot) was flying straight and level at Flight Level (FL) 330 after having previously initiated a lateral strategic resolution maneuver. It was approaching its top of descent point. Traffic (an Ames pseudo-pilot) was flying straight and level at FL 320, and would cross underneath slightly before ownship's top of descent. Figure 3 superimposes the actual and predicted trajectories for both the ownship and traffic aircraft. Two minute tick marks allow direct comparison between the two aircraft and between the lateral and vertical trajectories (top and bottom pane of Figure 3, respectively).

As shown in Figure 3, the ownship "rounded the corner" at the top of descent even more than AOP anticipated and initiated its descent slightly before AOP predicted it would. This difference resulted in an LOS (highlighted) with ownship descending to within 1000 ft of the traffic before the required 5 nm of lateral separation was regained.
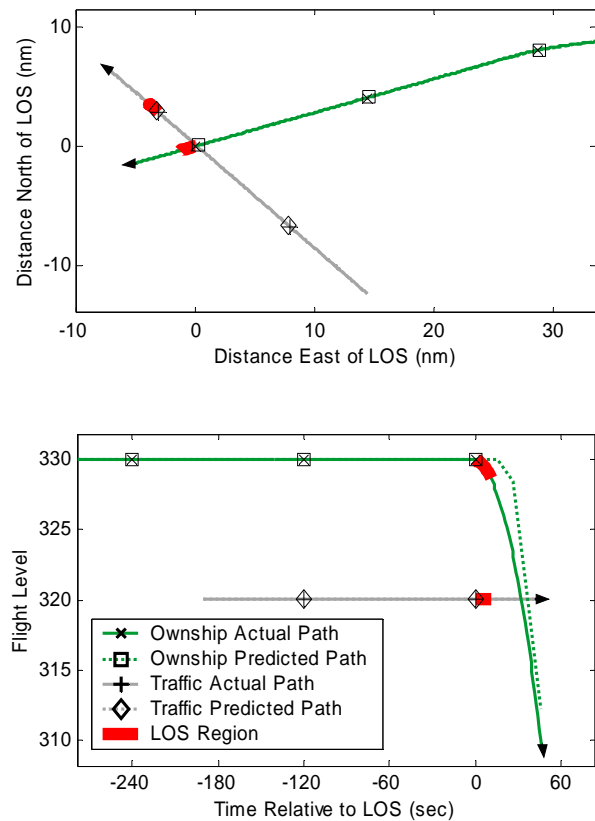

Figure 3. Case 1: Comparison of Actual and Predicted Trajectories

Figure 4 shows the actual and predicted lateral and vertical separations between the two aircraft. If the ownship had followed the AOP trajectory as predicted, the two aircraft would have had over 5 nm of lateral separation before ownship initiated its descent and no LOS would have occurred.

In this case, the trajectory prediction error can be traced to a discrepancy in the assumed behavior near the ownship's top of descent. After further study, it was determined that the AOP assumed the descent would commence precisely at the top of descent. The ownship guidance algorithm smoothed this transition and called for the aircraft to start its descent prior to the top of descent point provided to the AOP. This top of descent 'round-off' in the guidance algorithm had previously been designed to be consistent with a particular FMS type. Even if the AOP trajectory predictor were tweaked to accommodate this nuance, it would likely not capture the transition behavior of another FMS aboard a different aircraft.
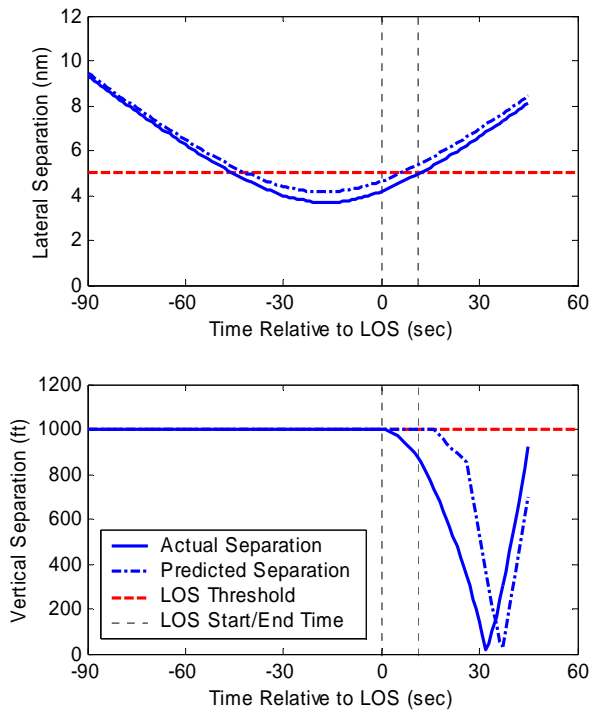
Figure 4. Case 1: Comparison of Actual and Predicted Separation

Figure 5. Case 2: Comparison of Actual and Predicted Trajectories

Case 2

In the second case, the ownship (a Langley subject pilot) had previously initiated a lateral strategic resolution maneuver for a different conflict. At the time of the separation loss, it was flying straight and level at FL 330. Traffic (an Ames pseudo pilot) descended out of FL 380 and was predicted to cross ownship's altitude with greater than the required lateral separation. However, the traffic flight path was closer to ownship than predicted. These effects resulted in an LOS, with traffic passing within 5 nm of ownship while descending through its altitude. Figure 5 superimposes the actual and predicted trajectories for both aircraft.

In this case, the LOS was attributed to a speed discrepancy and a difference in the way the top of descent was defined by the two aircraft. The source of the speed discrepancy is still under investigation.
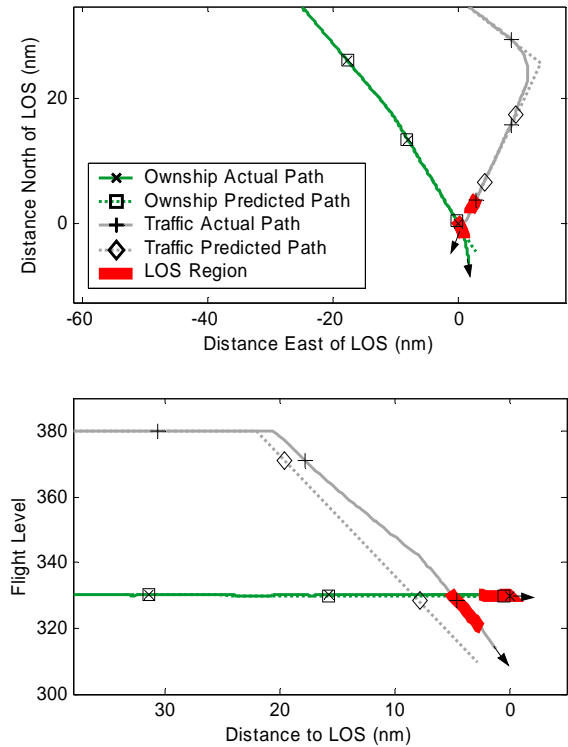
Concerning the top of descent, it was later determined that the guidance algorithm for the traffic aircraft called for it to fly slightly over the top of descent, then gradually intercept the profile from above. This behavior is different, but no less correct, than that of the descending aircraft in Case 1. In the former case, the top of descent resembled a fly-by waypoint, whereas it was similar to a fly-over waypoint in the latter. Two different FMS systems and guidance algorithms yielded these consequential trajectory prediction errors.

Figure 6 shows the actual and predicted lateral and vertical separations between the two aircraft. If the aircraft speed discrepancy had not occurred, the difference in top of descent behavior would only have resulted in a higher descent relative to the prediction. If this had been the only prediction error, an LOS would not have occurred, because the lateral separation would never have gone below the minimum required distance. Because the traffic aircraft was faster than expected, it was closer to ownship when the former passed behind and underneath the latter.
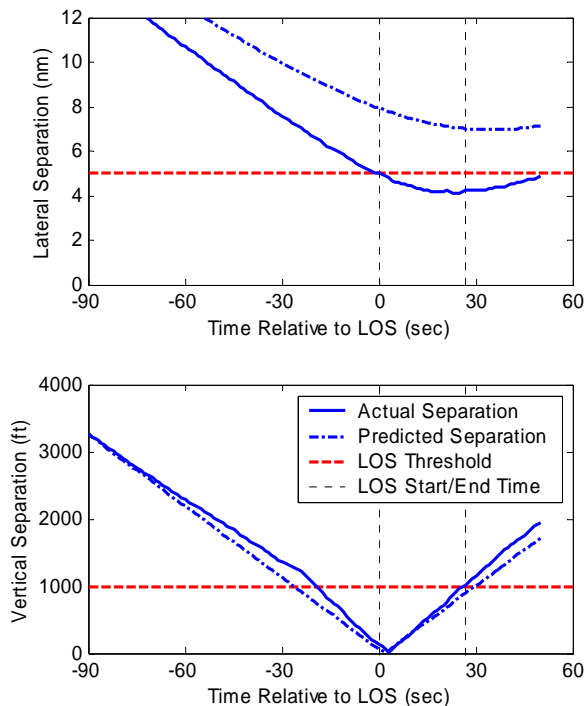
Figure 6. Case 2: Comparison of Actual and Predicted Separation

Although the LOS in this case can be attributed to the along-path prediction error, the vertical prediction is worth noting because such differences can lead to LOS events, as occurred in Case 1. In addition, the Mondoloni and Bayraktutar study determined that top of descent placement has a significant impact on the accuracy of the predicted vertical trajectory [8].

These two cases represent only a small subset of potential variation between predicted and actual trajectories. It would be impractical to model the precise flight path techniques of all flight management systems currently in operation [11]. Even for a given aircraft, the tracking performance is not guaranteed to be consistent [13].

Although some improvements to the fidelity of the AOP trajectory generator will be made as a result of this experiment, the development of trajectory uncertainty-based buffers are seen as the primary factor towards mitigating prediction error.

Trajectory Buffers

AOP capability has recently been expanded to add dynamic buffers to the predicted trajectories of both ownship and traffic aircraft. These buffers account for the varying uncertainty that occurs in three-dimensional space along a predicted path. The buffers allow AOP to represent predicted trajectories as three-dimensional regions of airspace instead of simple lines and curved segments as occurred in previous versions.

Buffers along a predicted path can be assigned locally using a trajectory-relative coordinate system. The axes include cross-track (normal to the path in the horizontal plane), vertical, and along-path (used to represent time uncertainty).

Buffers are assigned at progressive points along a trajectory, with linear growth or shrink rates describing the rates of change between points. When assembling the traffic aircraft trajectory, AOP may generate intermediate points in between the received TCPs. These points often capture details of a trajectory change such as the start and end of turn or the transition region near top of descent. For the ownship, AOP uses internal performance models and communications with its FMS. It then uses the same criteria for assigning buffers as those used for traffic aircraft.

After adding the buffers, each trajectory can be represented as a three-dimensional tube through space, the center of which is the aircraft's predicted path. The trajectories have an associated time range at each position along the path. AOP's conflict detection routines compare the buffer-enhanced trajectories as four-dimensional objects to determine regions of potential conflict. When making this comparison, the entire trajectory region is given equal weight. This approach differs from other conflict probes that determine conflict probability [14].

AOP first compares the two trajectories to see whether their vertical separation is ever less than the minimum required. If so, these regions are evaluated to determine if there is a time overlap. Any regions of time overlap are further compared to see whether they lie within the minimum required lateral separation. Remaining segments are marked as potential conflicts and the pilots are notified accordingly.

The buffers have been designed as a way to represent the types of discrepancies observed in simulation between the actual and predicted

trajectories. They focus on accounting for differences due to environmental conditions and aircraft dynamics, performance, and guidance techniques. Trajectory errors associated with potential intent changes after the predictions are made are not currently accommodated.

Cross-track Buffers

While on an FMS-path, cross-track buffers can be used to accommodate uncertainty near a waypoint involving a course change. Potential errors can include turn anticipation differences or modeling the radius of turn. The latter varies for fly-by and fly-over turns. These buffers can also represent tracking error between waypoints.

Figure 7 shows an example of how these buffers can be applied, in this case to account for errors in modeling the radius of a turn at a waypoint. The predicted path is shown as a heavy line, with a thinner line outlining the region swept by the cross-track buffer along each segment. In this example, the straight-flight segments before and after the turn have relatively narrow cross-track buffers; the buffer grows during the first segment of the turn, remains wide during the segment in the middle of the turn, and shrinks again in the last segment of the turn.
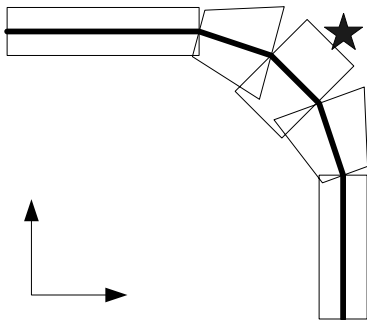

Figure 7. Cross-track Buffers at a Waypoint Turn

When extended to cover non-FMS paths, cross-track buffers will be able to account for lateral uncertainty while an airplane is flying in an air-referenced mode (such as Heading Hold). In this situation, the aircraft's ground track is subject to changing wind fields.

Vertical Buffers

Vertical buffers are used to represent altitude uncertainty along a path. They are particularly helpful in addressing the complex and variable trajectories associated with FMS climbs and descents. For instance, they can be used to provide vertical tolerance for the multiple inflection points shown in Figure 1. If an unexpected tailwind is encountered during a descent and the pilot does not add sufficient drag, the aircraft may leave the path in order to avoid an overspeed condition.

Figure 8 shows how the altitude buffer might account for errors in modeling the vertical profile of an FMS descent without intermediate altitude constraints. The vertical buffer is relatively small during level flight. Approaching the top of descent, the buffer below the path grows in order to protect against an early descent or the "rounding off" effect prior to the top of descent. As the aircraft descends, the buffer above the path gradually grows until they are large both above and below. At the end of descent, the buffer below the path shrinks to reflect the assumption that the aircraft will level off at the target altitude.
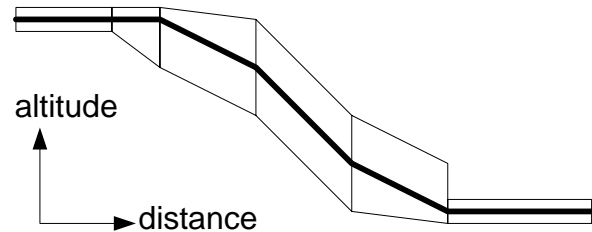

Figure 8. Vertical Buffers during an FMS Descent without Intermediate Altitude Constraints

If the descent has an intermediate waypoint altitude constraint, the buffers can be locally reduced to reflect higher vertical path certainty when crossing the waypoint. A comparable process can be applied for subsequent descent segments.

Along-Path Buffers

Along-path buffers are used to represent time uncertainty along the path. This component assumes the aircraft follows its predicted path

exactly but arrives at each point earlier or later than expected.

This buffer need not grow uniformly on all segments. For example, the error may grow much faster during a descent segment than while in level flight. The along-path buffer may shrink as the aircraft approaches a meter fix with a required time of arrival on the assumption that the aircraft will comply with the restriction.

Application to Previous Cases

The differences between the actual and predicted trajectories in the two cases described previously can be related to this coordinate system. In Case 1 (Figures 3-4), there is good lateral and along-track agreement. The top of descent "round-off" can be classified as a pure vertical error.

In Case 2, there is generally good cross-track agreement with the predicted trajectory. One exception is the area near the traffic aircraft turn that wasn't modeled by the existing version of AOP. (Turn modeling has since been added). There are both vertical and along-path errors for this trajectory. The aircraft was faster than expected (along-path error), but also began its descent later than expected (vertical error). In addition, some inflection points followed by this aircraft were not modeled by AOP. When converted back to a common time reference, the along-path error impacts both the aircraft's lateral and vertical predicted position. The differences between the predicted and actual separation are the cumulative effect of the cross-track and vertical components of the along-path error (at a given time) with the independent cross-track and vertical errors. The lateral separation error results from the aircraft being further along the path (and therefore closer to ownship) than expected. The top of descent occurs later than expected, but the aircraft also traveled faster than expected. These two effects negate each other and the resulting vertical error is minimal during the LOS region.

Application of Trajectory Uncertainty Elements Available over ADS-B

Although the architecture to generate buffers has been implemented within AOP, the algorithms to create them have not yet been developed. These algorithms will be designed to use trajectory elements available for broadcast over ADS-B [2].

The RTCA Minimum Aviation System Performance Standards (MASPS) for ADS-B provides industry standards for intent broadcast [2]. The latest version of this document (RTCA/DO-242A) has defined several additional parameters in the mode status, state vector, target state, and trajectory change reports that can be used to assess trajectory uncertainty.

Mode Status and State Vector Reports

The mode status and state vector reports provide information on an aircraft's operational capabilities and dynamic state, respectively. Parameters describing the transmitting aircraft's navigation accuracy and integrity are contained within these reports. These parameters indicate the capability of an aircraft's navigation sensors to accurately determine its position and velocity (accuracy) and to know when its position exceeds a containment radius (integrity).

AOP could use this information to apply cross-track and vertical buffers to the current flight segment. A sudden drop in integrity may require the buffers to expand significantly. Buffer adjustments due to accuracy changes would likely be smaller.

Target State Report

The target state report contains the target altitude and target heading or target track angle of the transmitting aircraft. During straight and level flight, these elements represent the current commanded states. For a climb, descent, or turn, they are the level-off altitude or roll-out heading or track angle commanded by the guidance system.

Horizontal and vertical target source indicators state whether the respective targets come from the flight control panel, FMS, or are the current state values commanded by the autopilot. This information could be used by AOP to assess the potential for future intent changes. In general, flight control panel settings represent tactical operations that may be more subject to change than the programmed flight plan.

Two flags are reserved for the transmitting aircraft's assessment of its conformance to the broadcast horizontal and vertical paths. A non-conforming aircraft would require larger buffers than a conforming one under similar circumstances.

Trajectory Change Report

When an aircraft is flying using strategic control (having multiple future target states), it may broadcast a series of messages that provide information on its predicted position, altitude, and time at upcoming TCPs. As part of this message set, it sends the horizontal and vertical type of each broadcast trajectory change point. These parameters give information on the waypoint type and the preceding flight segment.

The ADS-B MASPS currently defines five horizontal and five vertical change types. Horizontal types include direct to fix and course or track to fix legs (terminated at the fix or concluding with a fly-by turn) and radius to fix turns. Vertical types distinguish between open loop points such as top of climb and descent and those with constraints that the aircraft is trying to meet (such as a target altitude or waypoint altitude constraint).

Trajectory change types give AOP valuable information needed to assess the uncertainty at various locations along the path. Trajectory changes occurring at less certain locations will necessarily have larger buffers. Knowledge of a top of descent point and the occurrence of any successive altitude constraints will enable AOP to develop descent buffers comparable to Figure 8.

This report has also reserved space for several conformance assessments made by the transmitting aircraft. An "Able/Unable Altitude Constraint" indicates whether the aircraft predicts it will be able to make its next waypoint altitude crossing restriction. This parameter helps mitigate the effect of sending misleading information about an aircraft's altitude at a waypoint. With this information, the AOP can structure the vertical buffer around this point accordingly. Two additional flags provide notification of estimated horizontal and vertical conformance. These values are expected to assess the validity of other predicted trajectory elements broadcasted by that aircraft.

Use of these parameters that are already defined industry standards should give AOP significant insight into defining the buffers for each nearby aircraft.

Conclusions

Feasibility assessments of advanced operational concepts such as DAG-TM will need to incorporate reliable airborne conflict management capabilities. Trajectory prediction is a key element of all of these capabilities.

Environmental effects, variations in aircraft performance and avionics system design, and lack of necessary intent information make it difficult to precisely predict each aircraft's trajectory. Until the National Airspace System becomes entirely reliant on Required Navigation Performance [15], these uncertainties will continue to exist.

A recent experiment has validated some of these uncertainty types and led to the creation of a trajectory uncertainty-based buffer system that becomes a characteristic of both the ownship and traffic aircraft predicted trajectories. Algorithms for creating these buffers will rely heavily on industry-specified ADS-B parameters that give information on navigation capabilities, trajectory types, and expected path conformance.

Current AOP development work will concentrate on developing and refining these algorithms. Guidance from industry documents and identification of additional trajectory error cases from previous simulations will help drive the design process. One focused area of testing will assess the acceptability of false alarms created by expanding the predicted trajectory space for each aircraft. This area can be investigated through both follow-on human-in-the-loop experiments and with batch studies using the ATOS.

References

[1] NASA, 1999, Concept Definition for Distributed Air/Ground Traffic Management (DAG-TM), NASA, http://www.asc.nasa.gov/aatt/dagconop.pdf

[2] RTCA, 2002, Minimum Aviation System Performance Standards for Automatic Dependent

Surveillance Broadcast (ADS-B), RTCA/DO-242A, Washington, DC, RTCA.

[3] Barhydt R. and K. Krishnamurthy, 2004, Design of a Multi-mode Flight Deck Decision Support System for Airborne Conflict Management, International Conference on Human-Computer Interaction in Aeronautics, Toulouse, France, Eurisco.

[4] Barhydt R. and A. Warren, 2002, Newly Enacted Intent Changes to ADS-B MASPS: Emphasis on Operations, Compatibility, and Integrity, AIAA-2002-4932, AIAA.

[5] ARINC, 2000, Advanced Flight Management Computer System, ARINC Characteristic 702A-1, Annapolis, MD, ARINC.

[6] Casaux, F., 2000, Report of the Focus Area 3, FAA/Eurocontrol Technical Interchange Meeting, Shared Flight Intent Information and Aircraft Intent Data, Atlantic City, NJ, FAA.

[7] Barhydt, R., M. Palmer, and T. Eischeid, 2004, Development and Evaluation of an Airborne Separation Assurance System for Autonomous Aircraft Operations, Proceedings of the 24th International Congress of the Aeronautical Sciences, Yokohama, Japan, ICAS.

[8] Mondoloni, S. and I. Bayraktutar, 2005, Impact of Factors, Conditions and Metrics on Trajectory Prediction Accuracy, 6th USA/Europe Air Traffic Management R&D Seminar, Baltimore, MD, FAA and Eurocontrol.

[9] Hoekstra J., 2001, Designing for Safety the Free Flight Air Traffic Management Concept, NLR-TP-2001-313, National Aerospace Laboratory of the Netherlands (NLR).

[10] Prevot, T. et al., 2002, A Multi-Fidelity Simulation Environment for Human-in-the-Loop Studies of Distributed Air-Ground Traffic Management, AIAA-2002-4679, AIAA.

[11] Ottobre, R., J. O'Neill, and A. Herndon, 2005, Analysis of Advanced Flight Management Systems (FMSs), FAA Contract #DTFA01-01-C-00001, MITRE, http://www.airlines.org/restricted/fmstf/d.aspX?nid=3582, User ID = RNAV and Password = taskforce

[12] Barhydt, R. and P. Kopardekar, 2005, Joint NASA Ames/Langley Experimental Evaluation of Integrated Air/Ground Operations for En Route Free Maneuvering, 6th USA/Europe Air Traffic Management R&D Seminar, Baltimore, MD, FAA and Eurocontrol.

[13] Reynolds, T. and R.J. Hansman, 2003, Investigating Conformance Monitoring Issues in Air Traffic Control Using Fault Detection Approaches, Report No. ICAT-2003-5, Cambridge, MA, MIT.

[14] Erzberger, H., R. Paielli, D. Isaacson, and M. Eshowl, 1997, Conflict Detection and Resolution in the Presence of Prediction Error, 1st USA/Europe Air Traffic Management R&D Seminar, Saclay, France, FAA and Eurocontrol.

[15] FAA, 2003, Roadmap for Performance-Based Navigation, Version 1.0, FAA, http://www.faa.gov/avr/afs/afs400/RNProadmap.pdf

Email Addresses

Richard Barhydt: richard.barhydt@nasa.gov

Nathan Doble: n.a.doble@larc.nasa.gov

David Karr: David.Karr@titan.com

Michael Palmer: michael.t.palmer@nasa.gov