Autonomous Coordinated Airspace Services for Terminal and Enroute Operations with Wind Errors

Todd A. Lauderdale*
NASA Ames Research Center, Moffett Field, CA, 94035, USA

Christabelle S. Bosson†, Yung-Cheng Chu‡, and Heinz Erzberger§
USRA/NAMS, NASA Ames Research Center, Moffett Field, CA, 94035, USA

As current operations expand and novel uses of the airspace continue to multiply, increasingly autonomous systems will be required to ensure safety and efficiency. This paper presents the current status of an autonomous airspace service designed to ensure safe and efficient trajectories for aircraft in both enroute and terminal airspaces, the Advanced Airspace Concept. Previous papers have demonstrated the efficacy of this algorithm for enroute operations and for handling commercial operations in a complex metroplex when there is no uncertainty present. This study extends that work to demonstrate coordinated enroute and terminal operations and to demonstrate the performance of the algorithm under high levels of wind-magnitude uncertainty. Important algorithm improvements to enable coordinated enroute and terminal operations are discussed along with methods for handling uncertainty for constrained arrival operations. Results of simulations of operations into and out of the Dallas-Fort Worth area are presented that show the performance of the coordinated system both with and without uncertainty. The level of uncertainty included in these simulations resulted in a more than 100% increase in delay over baseline operations. Arrival schedule deviations were also significantly increased, but a new arrival conformance monitoring system mitigated these deviations to some extent.

I. Introduction

Airspace operations are continuing to evolve towards higher densities and higher levels of complexity, especially in the lower altitude airspace over urban areas [1, 2]. A shift will be required towards higher levels of autonomy for both aircraft and airspace systems if these levels of complexity and density are to be safely and efficiently handled. One set of concepts of how to enable the shift towards airspace autonomy is known as the Advanced Airspace Concept (AAC) [3] and the related Terminal Advanced Airspace Concept (TAAC) [4]. Both of these concepts include a strategic separation assurance, sequencing, and scheduling algorithm known as the Autoresolver. The TAAC Autoresolver has been demonstrated in simulation to effectively separate, sequence, and schedule aircraft arriving into the D10 TRACON over the Dallas-Fort Worth metropolitan area [5]. Operations in D10 include complex interactions between multiple major airports (Dallas-Fort Worth International and Dallas Love Field), efficient scheduling of aircraft to constrained arrival runways, and efficient departure management to ensure safe and continuous climb trajectories (see Fig. 1). To date, the development of these terminal operation management systems has mainly focused on large, transport aircraft, but the concepts are applicable to many different operations and aircraft types [6]. Also, development of the Terminal and Enroute algorithms have proceeded independently, and the interactions and coordination of the two algorithms have not been studied.

One common property of most automation systems for separation and efficiency is their reliance on trajectory predictions (predictions of the future states of the aircraft) to determine future safety issues and to generate efficient schedules for constrained resources. In the terminal area, there are many constraints and interactions among aircraft, and these interactions are very difficult to handle, even if perfect predictions are used. In real-world operations though, there will always be uncertainties leading to trajectory prediction errors. For an autonomous system to be considered

---

*Aerospace Engineer, Aviation Systems Division, NASA Ames Research Center
†Aerospace Engineer, Unversities Space Research Association, NASA Ames Research Center, AIAA Member
‡Senior Software Engineer, Unversities Space Research Association, NASA Ames Research Center
§Senior Scientist, Unversities Space Research Association, NASA Ames Research Center, AIAA Fellow
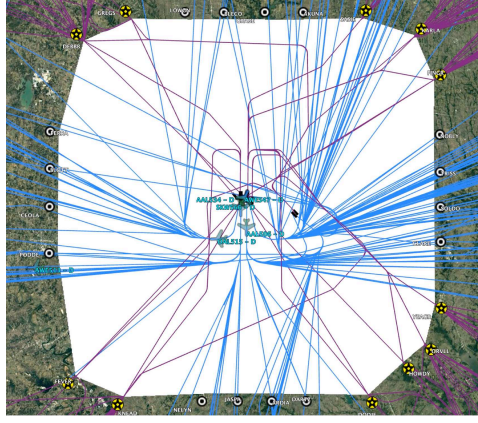
Fig. 1   An image of simulated trajectories into and out of the D10 TRACON including traffic to both DFW and DAL. Magenta trajectories are for arrivals and blue ones are for departures.

for operational use, it must be shown that the algorithms are robust and efficient in the presence of these real-world issues.

In this paper, we present the current status of development of a coordinated AAC/TAAC system. We describe the coordination of the two Autoresolver algorithms including some improvements to ensure smooth transition between the two control paradigms. We also look at the effects of trajectory prediction errors on conflict detection performance in terminal airspace and on the TAAC resolution process. This is an extension to the terminal airspace of previous analyses of enroute conflict detection and resolution in the presence of trajectory prediction error for the enroute AAC Autoresolver algorithm [7, 8]. While this study focuses on resolutions provided by a specific algorithm, the intention is to understand more generally the difficulties surrounding controlling aircraft operations in dense, complex airspace in the presence of trajectory prediction errors due to uncertainty.

II. Coordinated Enroute and Terminal Operations

Until recently, development of AAC and TAAC proceeded separately with each type of operation handled independently. For this paper, these operations were coordinated with 20 enroute AAC Autoresolver agents responsible for separation in each of the 20 mainland enroute control centers and one TAAC Autoresolver agent responsible for aircraft in the D10 TRACON. Methods for coordination between enroute centers have been demonstrated previously [9], so the main question for this work is how to coordinate enroute and terminal operations. As in the previous research, for this paper each agent can only maneuver aircraft currently contained in its airspace.


Fig. 2   The 20 mainland air traffic control centers and the D10 TRACON.

A. Coordination Rules

Terminal and enroute operations are subject to significantly different separation requirements [10], including 5-nmi horizontal separation in enroute airspace and 3-nmi separation in the terminal, making coordination more difficult. Since the terminal agent cannot control enroute aircraft and vice versa, it is important to ensure that aircraft transitioning from one region of control to the other are afforded a reasonable amount of time where the new controlling agent will not need to act on them.

As an initial attempt at smooth coordination, three relatively simple rules were implemented and followed by all agents. The first is that two flights departing the terminal from the same departure fix are subject to the enroute separation requirement inside the terminal area (Fig. 3(a)). The second rule is that enroute flights arriving at the terminal must be free of conflict (using terminal standards) with terminal flights for at least 65 seconds beyond the terminal area boundary crossing (Fig. 3(b)). The third rule is that aircraft inside the terminal must be conflict free (using terminal requirements) with aircraft arriving into the terminal for at least 135 seconds after the time at which the arriving aircraft enters into the terminal (Fig. 3(c)). The time cutoffs discussed above are just initial guesses as to values that allow enough time for conflicts to be resolved between entering and exiting flights. Further research will be necessary to identify the best values for these parameters.


Fig. 3   Illustration of the three Terminal/Enroute coordination rules. The green aircraft indicates the aircraft being checked for conflicts. In (a), two aircraft departing the terminal at the same fix are subject to enroute separation requirements through the terminal area. In (b), the enroute agent must ensure that a flight arriving to a terminal is conflict-free for at least 65 seconds beyond the terminal crossing. In (c), the terminal agent must not create a conflict with an aircraft entering within the first 135 seconds of arrival in the terminal.

To enable these rules, each agent must be aware of the aircraft in the neighboring airspace and near the boundary. They also must both agree on the predicted trajectories of the aircraft in the other airspace. Differences in predicted trajectories may lead to inconsistencies in conflict predictions including pop-up, near-term predicted losses of separation that can be very difficult for separation systems to resolve.

B. Effects of Coordination Rules

To illustrate the efficacy of these coordination rules, two different simulations were performed using the fast-time Airspace Concept Evaluation System (ACES) [11]. For both of these simulations, perfect trajectory knowledge is assumed so there are no trajectory prediction errors. A flight dataset with more than 350 operations into and out of D10, at approximately current arrival and departure rates, was used with both enroute and terminal agents active. In the first simulation, there were no coordination rules implemented, so terminal and enroute agents did not consider aircraft beyond their boundaries. In the second simulation, the rules as outlined above were followed.

The results of the two runs are shown in Fig. 4. When coordination was implemented, the average alert lead time before first loss of separation increased by 7.2%, from less than 9.8 minutes to about 10.5 minutes. This is an average across all conflicts, not just conflicts across boundaries, so this a relatively large increase and generally simplifies the resolution process as more time is available to solve each conflict. More importantly, the number of losses detected with a lead time of less than 1 minute decreased from 10 to 2 when was coordination is implemented. The remaining two near-term losses indicate that some additional work is required to ensure smooth coordination between enroute and terminal operations.

Fig. 4 Comparison of results when coordination rules were turned on and off. (a) shows the average time to loss when conflicts are first detected and (b) shows the number of conflicts detected with less than one minute to loss of separation.

III. Automation and Trajectory Errors

Both the enroute Autoresolver and the terminal Autoresolver were developed in ACES to provide realistic conflict scenarios that are based on actual, flown routes. ACES can be configured with or without trajectory-prediction error modeling, and the initial development of both algorithms began using perfect knowledge of the future state of aircraft and, therefore, no trajectory prediction errors. Arrivals and departures to/from both airports, as flown by ACES, are illustrated in Fig. 1.

As development of the enroute Autoresolver progressed, trajectory-error modeling was introduced into the simulation to mature the algorithm [8]. These trajectory errors are included in the development of the terminal Autoresolver for the first time in this paper. Only wind-speed errors were modeled for this study involving enroute and terminal Autoresolver agents (other possible sources of errors include: aircraft weight uncertainty, descent speed uncertainty, and pilot timing uncertainty) . To simplify the analysis, a constant, uniform wind field of 25 knots from the south was applied as shown in Fig. 5(a). The magnitude of the wind was perturbed by 50%, but the wind direction was not perturbed.

Fig. 5 Depictions of the constant, uniform wind fields used in the simulation. (a) shows the actual wind field modeled in the simulation and (b) shows predicted winds (dashed lines) used for conflict detection and resolution. To simplify the analysis, the predicted winds in all simulations in this paper differ from the actual simulated winds only in the wind magnitude.

The effects of these 50% wind-speed errors on trajectory predictions for a representative flight are shown in Fig. 6. The aircraft is arriving at Dallas Love Field from the northwest (Fig. 6(a)). During the first portion of the trajectory the aircraft is flying largely perpendicular to the wind resulting in little difference between the predicted and actual ground-speed due to wind errors, and the along-track error shown in Fig. 6(b) is negligible. When the aircraft turns

4

south, the along-track errors start to accrue as the predicted ground speed is different than the actual ground speed. The effects on the vertical profile are shown in Fig. 6(c). Since the wind is a headwind and the speed is predicted to be greater than it actually is, the time to the descent point in the predicted trajectory is longer than the actual time to the descent point. This also shows that this wind error results in a predicted arrival time that is 45 seconds too late.

Fig. 6 **Plots comparing the actual, flown trajectory (solid lines) and the predicted trajectory (dashed lines) for a flight arriving at Dallas Love Field. (a) shows the arrival path from the northwest, (b) shows a comparison of the flown, along-track distance for the two trajectories including the arrival time error, and (c) shows the differences in altitude profile and the error in the predicted arrival time.**

IV. Managing Arrivals with Trajectory Errors

For general interactions between cruising aircraft, trajectory prediction errors can largely be handled by adding detection buffers for predicting conflicts [8, 12]. For aircraft merging across a fix (the arrival fix for enroute operations and the arrival runway threshold for terminal operations), though, the problem is more difficult. The complexity of the arrival problem, even with no prediction errors, was the initial motivation for the Arrival Manager component of the Autoresolver [3]. For arrival management, the Autoresolver creates and maintains a schedule of aircraft arrival times, based on the predicted fix- or threshold-crossing time obtained from trajectory predictions, and then creates conflict-free trajectories that satisfy that arrival schedule. With no prediction errors, once the schedule and satisfying trajectories are created no other maneuvers are required to ensure continuous, conflict-free arrivals across the fix.

Trajectory prediction errors cause this original schedule to be incorrect, as the originally predicted crossing times are the product of erroneous wind data (see Fig. 6(c)). As the aircraft continues to approach the fix, the effects of these errors become more apparent, and conflicts between leading and trailing aircraft become likely. As in normal operations, a conflict prediction buffer can be used to predict these conflicts in the presence of errors, but large buffers are inefficient for arrival management, and resolutions late in the arrival process can often lead to lost slots across metering points and correspondingly higher delay.

With this in mind, an arrival time conformance monitoring function was added to the Arrival Manager to try to maintain the schedule in the presence of errors while reducing the need for additional separation buffers. Fig. 7(a) shows an arrival timeline where a newly scheduled aircraft (A2) is scheduled behind A1 with an additional time

spacing added to the required separation. As time progresses, the newly predicted arrival times are compared against the originally scheduled times, and checked to see if the predicted time is beyond a conformance allowance. In Fig. 7(b) the predicted time for A2 is outside of this allowance, so A2 is deemed to be out of conformance.

Fig. 7   Arrival timelines at two different epochs. (a) At time t0, aircraft A2 is scheduled with a spacing buffer added to the required spacing. (b) At time t1, the predicted times of arrival have drifted due to prediction errors.

During every resolution cycle of the Autoresolver, for any aircraft that is out of conformance, the Arrival Manager attempts to create a resolution to move this aircraft back into conformance before any other resolution process is attempted. In most cases, this means of keeping aircraft in conformance precludes predicted losses of separation, but if a loss is not resolved in this way, it might require the aircraft to be rescheduled. Two new resolution maneuvers were developed to aid in the process (described in Ref. [4]). One is a small-speed-change maneuver. In this maneuver, smaller increments of speed change are applied than are used for general speed-change maneuvers (typically down to 1 knot increments). The other new maneuver is a small-delay path stretch. This maneuver combines a new route with a speed increase or decrease to try to achieve the desired arrival time. Overall, this concept is similar to an autonomous version of the human-centric Terminal Sequencing and Spacing System developed by NASA and the FAA [13].

V. Results of Fast-Time Simulation with Trajectory Prediction Errors

To understand the effects of wind-speed errors on the performance of conflict detection and resolution, sequencing, and scheduling in the terminal, the analysis was partitioned into two studies: (1) detection-only and (2) detection and resolution. In the resolution study, the results are compared with and without arrival conformance monitoring. For all studies, the D10 dataset with over 350 operations into and out of D10, discussed previously, was used. Also, enroute and terminal coordination rules were used in all simulations.

A. Detection Only

In the first study, only conflict detection is evaluated with no conflict resolutions performed. Four different detection-only simulations were performed with two different wind errors: 0% (baseline) and 50% wind magnitude error and two different detection criteria: 3 nmi (no buffer) and 4 nmi (a 1-nmi buffer). A geometric conflict detection system was used comparing aircraft trajectories point-by-point to determine times when the aircraft are within a proscribed distance of one-another. In the terminal, the separation standard is generally 3 nmi and 1000 ft, but this standard is complicated by other terminal-specific rules [10].

In the following discussion, imagine creating two trajectories for every aircraft at any point in time, one trajectory containing errors and one perfect trajectory. Detection between the perfect trajectories never involves a detection buffer, but the buffer is added to detections between the trajectories with errors. A missed alert is counted any time an alert is detected between the perfect trajectories but not between the trajectories containing errors. A false alert is counted when a conflict is detected between the non-perfect trajectories (and applying the buffer) but not between the perfect trajectories (applying the standard separation criteria).

For the detection study, the effects of a horizontal buffer of 1 nmi were determined. Fig. 8(a) shows the percentage of true conflicts that were not detected as functions of time to first loss for the two different wind error levels (different

6

colors) and for the no-buffer (solid lines) and the 1-nmi buffer (dashed lines) cases. As expected, as the aircraft approach the loss of separation time there are fewer missed alerts. Also, there are significantly more missed alerts with the 50% wind error than with no errors. Adding the 1-nmi detection buffer results in a little more than half as many missed alerts than with no buffer. Note that with no error (with and without buffers) the conflict detection system does not miss any alerts, and these curves both lie along the horizontal axis.

Fig. 8   Detection-only results where the solid lines are for no-buffer results and the dashed lines include a 1-nmi buffer. (a) The missed alerts as functions of time to first loss for the two wind error levels. In this figure, black solid and dashed lines coincide. (b) False alerts as functions of predicted time to first loss for two wind error levels.

Interestingly, not all conflicts are detected as you approach the loss time. This stems from cases where the wind errors in the terminal make it impossible to satisfy all of the constraints in the terminal (e.g. arriving at the altitude of the runway at the location of the runway threshold). The trajectory predictor used in this simulation would not return a trajectory in this case, resulting in these missed alerts. An enhancement to the system would be to rely on dead-reckoning trajectories in such instances.

Another important metric to understand conflict detection performance is the rate of false alerts. For the simulated scenarios, the percentage of alerts that are false as a function of predicted time to first loss are shown in Fig. 8(b). The horizontal buffer has a large effect on false alerts - basically creating a constant increase over the no-buffer cases (since more airspace is being protected). The level of wind prediction error also has a significant effect on the false alerts, with higher errors leading to more incorrect alerts.

B. Effects of Errors on Resolutions

For the next study, the resolution process of the TAAC algorithm was used to attempt to keep aircraft separated and to space the aircraft appropriately at the arrival fixes and the runway threshold. As demonstrated previously [5], the terminal algorithm is capable of separating complex metroplex flows when there are no trajectory errors modeled. This study builds on these previous studies by evaluating the performance of the algorithm when trajectory errors due to wind prediction errors are included. For the resolution runs, no detection buffers were applied.

The baseline case is composed of coordinated enroute and terminal operations with no wind errors. The next simulation included wind errors of 50%, as discussed previously, with no arrival conformance monitoring, and the final simulation included the errors and the conformance monitoring capability.

Fig. 9(a) shows how the wind errors affect the total number of resolutions required to separate aircraft. The 50% wind error leads to approximately a 50% increase in the number of resolutions over the baseline run. The resulting effects on delay are shown in Fig. 9(b). Because of the large number of constraints in the terminal area, the 50% wind error results in an increase in delay of over 100%.

The effects of including the arrival conformance monitoring capabilities described in Sec. IV are also shown in Fig. 9(a) and Fig. 9(b). For these simulations, a scheduling buffer of 15 seconds was used and a conformance bound of ±5 seconds was used. Since the conformance bound is so tight, over 6 times as many resolutions had to be performed to maintain the schedule. These additional resolutions reduced total delay by approximately 16%.

Fig. 9   For the three different cases simulated, (a) shows the number of resolutions, and (b) shows the total delay, comparing the baseline values to the increases present in the other simulations.

To get more detail on the performance of the conformance monitoring system, the stability of aircraft order and arrival times in both terminal and enroute operations are shown in Fig. 10. In this figure, the number of times an aircraft has been scheduled and then needs to be rescheduled due to interactions with other arrival aircraft is counted. Fig. 10(a) shows that for the baseline case there are only two times when an aircraft needs to be rescheduled. It also shows that when errors are present but no monitoring there are 75 schedule changes, and that the monitoring system reduces these schedule changes by nearly half. A breakdown for the two error cases between enroute and terminal is shown in Fig. 10(b). It can be seen that the monitoring system is more effective at reducing schedule changes in the enroute system than in the terminal. This is due to the complexities and interactions of having arrivals to relatively closely spaced parallel runways in the terminal while the enroute arrival fixes are spaced reasonably far apart.

Fig. 10   Results focusing on the stability of arrival schedules for (a) the three different simulation cases and (b) the two error cases partitioned by terminal and enroute airspace.

It may seem contradictory that the monitoring system maintains the schedule but still has significant delay, but the nature of the errors (and the definition of delay) in this simulation explain this contradiction as follows. Imagine that we have a single aircraft arriving at a runway. This aircraft is scheduled some distance from the runway based on a predicted (estimated) time of arrival (ETA) derived from a trajectory prediction at that point. Assume that in this case the trajectory prediction errors cause the ETA to be later than it would be with no errors. With no conformance monitoring the aircraft will just fly its actual trajectory and arrive earlier than predicted. This would result in no delay. Now, if conformance monitoring is used, the system will try to maintain the original ETA, and it will have to delay the aircraft from its default trajectory to hit this time.

If the ETA errors are random with some positive and some negative, then the net effect on delay of the schedule conformance monitoring should be relatively low, with some aircraft being sped up and some slowed down. In this simulation, though, as this is a south-flow arrival stream with a positive northward error in wind magnitude, nearly all aircraft have a later predicted ETA than their no-error ETA (as in the example trajectory in Fig. 6(c)). The conformance

monitoring in this situation results in significant delays as a result.

The conformance monitoring system was successful at reducing both total delay and the number of schedule changes in the presence of wind errors, but at the cost of significantly more resolutions. Further work is necessary to balance the increased number of resolutions against the scheduling benefits through the use of appropriately sized buffers and bounds. Also, the complex interactions in the terminal need to be studied in more detail to ensure that schedules can be maintained.

VI. Conclusions

Autonomous systems will be required in order to meet the projected demand of novel airspace uses enabled by unmanned vehicles and new power and propulsion systems - especially in terminal airspace over larger metropolitan areas. These autonomous systems will rely on predictions of the future state of all neighboring aircraft, and these predictions will necessarily be imperfect. This paper presented the current status of the Advanced Airspace Concept (AAC) algorithms for automating separation, sequencing, and scheduling services in both the enroute and terminal environments. Similar automation systems could be used to manage these new operations and may be necessary for legacy operations in this new, complex environment.

Two contributions are discussed in this paper. The first is the coordination of terminal and enroute algorithms to provide trajectory-based separation capabilities from takeoff to landing. Rules designed to enable coordination between terminal and enroute operations are discussed and shown in simulation to drastically reduce the pop-up conflicts present in the absence of these rules.

The second contribution is an initial demonstration of the terminal and enroute algorithms working together in the presence of trajectory prediction errors. Enroute Autoresolver was evaluated previously with trajectory-modeling errors, but this is the first time that the terminal algorithm has been analyzed with errors and the first time that the two together have been evaluated in that scenario.

Looking just at conflict detection in the presence of 50% wind-speed errors, initial results show a significant adverse effect on performance in terms of both missed and false alerts, with approximately 5% of alerts missed even with a 1-nmi detection buffer and about 20% of alerts being unnecessary. Conflict resolution was also significantly affected by wind-speed errors with the total delay increased by over 100%. A schedule conformance monitoring system was discussed and evaluated in this paper. Conformance monitoring provided a moderate reduction in total delay in the presence of trajectory-prediction errors, and it reduced schedule disruptions but at a cost of significantly more resolutions implemented.

The results in this paper demonstrate that trajectory prediction errors can have a profound impact on trajectory-based operations in terminal airspace because of the complex constraints and interactions there. Development of the AAC Autoresolver algorithms will continue to attempt to handle these errors more effectively in both terminal and enroute operations.

References

[1] Duffy, M. J., Wakayama, S., Hupp, R., Lacy, R., and Stauffer, M., "A Study in Reducing the Cost of Vertical Flight with Electric Propulsion," *17th AIAA Aviation Technology, Integration, and Operations Conference*, 2017, p. 3442.

[2] Holden, J., and Goel, N., "Uber Elevate: Fast-Forwarding to a future of On-Demand Urban Air Transportation," Tech. rep., 2016.

[3] Erzberger, H., Lauderdale, T., and Chu, Y., "Automated conflict resolution, arrival management, and weather avoidance for air traffic management," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of aerospace engineering*, Vol. 226, No. 8, 2012, pp. 930–949.

[4] Erzberger, H., Nikoleris, T., Paielli, R. A., and Chu, Y.-C., "Algorithms for control of arrival and departure traffic in terminal airspace," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 230, No. 9, 2016, pp. 1762–1779.

[5] Nikoleris, T., Erzberger, H., Paielli, R. A., and Chu, Y.-C., "Performance of an Automated System for Control of Traffic in Terminal Airspace," *16th AIAA Aviation Technology, Integration, and Operations Conference*, Washington, D.C., 2016.

[6] Bosson, C. S., and Lauderdale, T. A., "Simulation Evaluations of an Autonomous Urban Air Mobility Network Management and Separation Service," *AIAA Aviation Technology, Integration, and Operations Conference*, Atlanta, Georgia, 2018.

[7] Lauderdale, T. A., "The Effects of Speed Uncertainty on a Separation Assurance Algorithm," *AIAA Aviation Technology, Integration, and Operations Conference*, Fort Worth, Texas, 2010.

[8] Lauderdale, T. A., Cone, A. C., and Bowe, A. R., "Relative Significance of Trajectory Prediction Error on an Automated Separation Assurance Algorithm," *9th USA/Europe Air Traffic Management R&D Seminar*, Berlin, Germany, 2011.

[9] Lauderdale, T. A., and Wang, T., "Coordination between multiple ground-based separation assurance agents," *2013 Aviation Technology, Integration, and Operations Conference*, 2013, p. 4368.

[10] Tang, H., Robinson, J. E., and Denery, D. G., "Tactical conflict detection in terminal airspace," *Journal of Guidance Control and Dynamics*, Vol. 34, No. 2, 2011, p. 403.

[11] George, S., Satapathy, G., Manikonda, V., Palopo, K., Meyn, L., Lauderdale, T. A., Downs, M., Refai, M., and Dupee, R., "Build 8 of the Airspace Concept Evaluation System," *AIAA Modeling and Simulation Technologies Conference*, Portland, Oregon, 2011.

[12] Cone, A. C., Bowe, A. R., and Lauderdale, T. A., "Robust Conflict Detection and Resolution around Top of Descent," *AIAA Aviation Technology, Integration, and Operations Conference*, Indianapolis, Indiana, 2012.

[13] Robinson, J. E., Thipphavong, J., and Johnson, W. C., "Enabling Performance-Based Navigation Arrivals: Development and Simulation Testing of the Terminal Sequencing and Spacing System," *Air Traffic Control Quarterly*, Vol. 23, No. 1, 2015, pp. 5–27.