Validation of an Automated System for Arrival Traffic Management

Robert Windhorst[1], Todd Lauderdale[2], Richard Coppenbarger[3], and Heinz Erzberger[4]
NASA Ames Research Center, Moffett Field, CA, 94035

The fuel-efficiencies of arrival flights that were managed by an automated system were compared to the fuel-efficiencies of arrival flights that were managed by air traffic controllers. It was infeasible to have the automated system control arrivals in real operations, so the comparison was accomplished by setting up a fast-time simulation where the automated system could manage arrivals with the same initial conditions and flight plans as those that operated in real operations during a selected comparison period and in the same background traffic. For this study, Newark Liberty International Airport was selected as the arrival airport because its high traffic load and constrained arrival procedures were expected to highlight fuel-efficiency benefits of an automated system. In the simulation, the automated system managed Newark arrivals, and the other flights (arrivals to other airports, departures, and overflights) composed the background traffic. To match the simulation and the real operations background traffic, the other flights flew in simulation the same trajectory that they flew in real operations during the comparison period. Fuel-efficiency was measured by calculating fuel burns of the arrival trajectories. The fuel-efficiencies of arrival trajectories produced in the simulation were compared with the estimated fuel-efficiencies of arrival trajectories recorded from real operations during the comparison period. Results showed that automation managed arrivals burned 346 lbs less fuel per flight on average than controller managed arrivals.

I. Introduction

Improving air traffic fuel efficiency lowers airline costs and reduces greenhouse gas emissions. Flight trajectories in high traffic conditions in current arrival operations often have excessive fuel burn because traffic is organized in a way that assists human controllers with safely and efficiently sequencing and spacing traffic. An automated system for arrival traffic management, due to its ability to anticipate downstream conflicts and to concurrently manage horizontal and vertical spacing of large numbers of arrivals, would not need the same organization and has the potential to allow more fuel-efficient trajectories.

This paper presents an initial validation and benefits study of AutoResolver[1], which is the part of an automated arrival traffic management system that generates and manages trajectories. The algorithms within AutoResolver sequence and space arrival traffic by periodically predicting flight trajectories, detecting conflicts (future violations of separation and schedule criteria) along those trajectories, and generating resolutions (flight maneuvers that eliminate conflicts). In this study, the fuel-efficiency of AutoResolver in managing arrival traffic at a major hub airport was validated by comparing it against the fuel-efficiency of human air traffic controllers in managing the same traffic. Since it was not feasible for safety reasons to validate AutoResolver as part of the automated traffic management system in an operational environment with live traffic, the approach used fast-time simulation and recorded data produced by traffic in the National Airspace System (NAS) as a surrogate for testing with live traffic.

Several studies have showed that managing arrival traffic with automated versus human control reduced fuel-burn, noise, and emissions, including the greenhouse gas carbon-dioxide. A flight test of arrivals into San Francisco International Airport [2] demonstrated savings of 227, 358, and 3,219 lbs of fuel burn per flight on average for light,

[1] Aerospace Engineer, Aerospace High Density Operations Branch, AIAA Senior Member.
[2] Aerospace Engineer, Flight Trajectory Dynamics and Controls Branch.
[3] Aerospace Engineer, Flight Trajectory Dynamics and Controls Branch, AIAA Associate Fellow Member.
[4] Ames Associate, Fellow, AIAA.

medium, and heavy traffic, respectively. In addition, a human-in-the-loop study of arrivals into Denver International Airport [3] produced savings of 65 and 170 lbs of fuel-burn per flight on average in medium and heavy traffic, respectively. Savings in [2] were larger than savings in [3] in part because [2] focused on heavier weight class aircraft. Analysis of arrival trajectories from both studies showed that those produced by automated control began descent later, had steeper descents, had fewer costly level-offs during descent, and absorbed delays at higher altitudes, where aircraft are more fuel efficient.

In [2] and [3], specialized and initial version algorithms generated the trajectory solutions to the arrival management problem, whereas in this study, the algorithms in AutoResolver were used. AutoResolver algorithms generalize and extend the conflict resolution and scheduling capabilities of those used in [2] and [3] and have been studied in simulation and flight test. Flight tests were recently conducted that evaluated key functions of AutoResolver in managing arrival traffic. Specifically, its functions for generating flight plan changes that produced descent trajectories with specified arrival times at given fixes were validated during the ecoDemonstrator campaign [4]. In addition, the flight test showed that the AutoResolver trajectory generator could predict fix crossing times with an error of less than plus or minus 30 seconds, which was as good as, and in some cases better, than the error of the onboard FMS trajectory generator.

Simulation studies of AutoResolver have demonstrated its ability to successfully maintain separation of traffic in various parts of the NAS and for different types of aircraft. AutoResolver was shown to manage conventional jet and prop traffic in the NAS [5,6,7,8,9,10], Cleveland Air Route Traffic Control Center (ZOB) airspace [11,12,13], Fort Worth Air Route Traffic Center (ZFW) airspace [14,15,16,17], and terminal airspace surrounding Dallas Fort Worth International Airport (DFW) [18,19,20] and Newark Liberty International Airport (EWR) [21,22]. In addition, it was shown to manage electric Vertical TakeOff and Landing (eVTOL) traffic in class B, D, and G airspace surrounding DFW [23,24] (an Urban Air Mobility [25] scenario). Furthermore, it was shown to manage handoffs across center/center [6,9] and center/terminal [19] airspace boundaries and to be robust to trajectory prediction errors [7,8,14,19,24,26].

Differences between AutoResolver and controller maneuvers for the same conflict were investigated in [16,17]. The comparison was accomplished by having controllers and AutoResolver alongside each other manage flights in a human-in-the-loop simulation.

In these studies, several approaches were used to make the simulations realistic in the sense that they modeled real operations in the NAS. Flights followed FAA-approved departure and arrival procedures. Standard navigation fixes and airways and actual airport locations, as well as center and sector airspace boundaries, were used. Simulations incorporated realistic wind and weather models, aircraft performance models, and an atmospheric model.

Similarly in these studies, historical and operational data were incorporated into simulations to make their traffic realistic. The scenarios, input files that define the set of flights to be simulated and their plans, were created by copying flight plans from historical data recorded from real operations in the NAS. These data included routes, takeoff times, types of aircraft, origin and destination airports, and cruise altitudes and speeds.

Furthermore in [14,15], flight tracks from historical data were played back in simulation until a specified transition boundary, where control of flights was handed over to the simulation and AutoResolver. This process matched the initial position of an AutoResolver controlled flight in simulation with its position recorded in real operations. The transition boundary was the lower boundary of a high-altitude airspace, specifically ZFW airspace. All flights that crossed the boundary were transitioned so that all flights in ZFW airspace were under AutoResolver control. Hence, AutoResolver always detected and resolved conflicts between two flights controlled by the AutoResolver.

However, in these simulation studies, the performance of AutoResolver, which is measured in this paper by the fuel-efficiency of trajectories and runway arrival rates, was not compared directly to that of air traffic controllers managing the same traffic. One reason for this was that the focus of the previous studies was AutoResolver algorithm development, and the goal was to identify losses of separation and enhance AutoResolver algorithms to prevent them.

Another reason for not comparing performance was that AutoResolver experienced different traffic conditions in simulation than those experienced by controllers in real operations. Although traffic conditions at the start of simulation were the same as those in real operations, due to the use of historical data as discussed above, after the start, traffic conditions experienced by AutoResolver in simulation diverged from those experienced by air traffic controllers in real operations. This was because small differences at the start of simulation propagated into larger differences as the simulation evolved. This was especially true for flights that AutoResolver maneuvered. Maneuvers issued by AutoResolver were generally not the same, in terms of timing, type, and magnitude of turn, speed, or altitude change, as those that were issued by human controllers.

In this study, a simulation environment was created that minimized differences between simulation and real operations traffic environments so that AutoResolver fuel-efficiency in simulation could be compared with air traffic controller fuel-efficiency in real operations. Fuel-efficiency was measured by calculating the fuel burn of trajectories, which was done post simulation for AutoResolver and post operations for controllers. This approach to validation was used as the best available alternative to having AutoResolver control actual flights in the NAS and extends the algorithm development and maneuver comparison achieved in previous papers. This paper describes the simulation and its results in the following sections: Simulation Setup, Results, Next Steps, and Conclusions.

II. Simulation Setup

A simulation was setup to compare the fuel-efficiency of arrival trajectories produced by AutoResolver and the fuel-efficiency of arrival trajectories produced by air traffic controllers. AutoResolver managed arrivals in the simulation, while air traffic controllers managed arrivals in real operations during the comparison period. To make the comparison valid, the arrivals managed by AutoResolver in simulation had to have the same initial conditions and flight plans as the arrivals managed by controllers in real operations and the simulation and real operations background traffics had to be the same. Different flight modeling approaches and a strategy for transitioning between those approaches were used in simulation to make the comparison valid. Sections II.A and II.B explain the modeling and transition approaches.

Section II.C and Section II.D present the tools and the sources of historical data, respectively. Section II.E presents the terminal airspace and arrival procedures used at Newark Liberty International Airport (EWR). EWR was selected as the arrival airport for the study because its constrained arrival procedures and high traffic was expected to highlight the fuel-efficiency of traffic controlled by AutoResolver. Finally, Section II.F explains the hierarchical scheduling process.

A. Approaches for Modeling Flights

To make the simulation a valid model of arrival operations at EWR during the comparison period, different approaches were used to model flights. The approaches were named automation-managed and controller-managed to assist the reader with keeping track of how they apply to the comparison. The following describes the approaches.

1. Controller-managed

Controller-managed flights composed the background flights. They moved along trajectories that were gathered from historical data. As the simulation stepped forward in time, it updated each controller-managed flight's position according to the one recorded for that time in the historical data, a process known as playback. These flights moved in simulation exactly as they did in real operations during the comparison period.

AutoResolver detected controller-managed flight positions and predicted their trajectories perfectly. It was able to predict flight trajectories perfectly because it was set to reproduce the prediction from historical data. In real operations, perfect predictions are not possible. Perfect predictions were used to gain an understanding of how AutoResolver performed in ideal conditions.

Although AutoResolver detected controller-managed flights, it was configured to not maneuver them because their trajectories were predetermined by those in real operations. If AutoResolver had detected a conflict between two controller-managed flights, it would have ignored the conflict because it could not maneuver them. However, these types of conflicts did not occur because controller-managed flight trajectories matched those in real operations, where controllers were keeping flights separated.

2. Automation-managed

Automation-managed flights were controlled by AutoResolver and used in the fuel-efficiency comparison. They were propagated using the simulation trajectory generator. Generally, they were not expected to move as they did in real operations because AutoResolver could maneuver them. However, prior to or in absence of AutoResolver maneuvering, automation-managed flights were expected to move as they did in real operations, and they did so within the accuracy of the trajectory generator.

AutoResolver detected automation-managed flight positions and predicted their trajectories. Predicted trajectories did not accumulate prediction errors because AutoResolver was configured to use the same trajectory generator for prediction that was used by the simulation for propagation. This was done to understand ideal AutoResolver performance.

AutoResolver maneuvered automation-managed flights to meet arrival time constraints and to resolve conflicts, both against other automation-managed flights and controller-managed flights. It achieved that by changing their speeds, altitudes, or routes. When AutoResolver detected a conflict between an automation-managed flight and a controller-managed flight, it maneuvered only the automation-managed flight.

The simulation capability that detected conflicts between automation-managed and controller-managed flights was novel to this study ([14,15] only had automation-managed vs automation-managed conflicts) and helped make the traffic environment in simulation closer to the traffic environment in real operations during the comparison period because background flight trajectories were the same in simulation and real operations and were mixed with automation-managed flights.

B. Transition from Controller-Managed to Automation-Managed

The assignment of a flight to be controller-managed or automation-managed was also designed to make the simulation a valid model of real operations during the comparison period. The flights selected to be automation-managed did not start the simulation with that modeling approach. They started as controller-managed so that when they made the transition to automation-managed their initial conditions in simulation were the same as those in real operations like in [14,15]. The transition point was on the boundary of the region where AutoResolver assumed control of the arrivals.

EWR arrivals were selected to transition from controller-managed to automation-managed and were the flights used in the fuel-efficiency comparison. All other flights were considered part of the background traffic and were controller-managed throughout the simulation. EWR arrivals were selected to be controlled by AutoResolver because arrival flows are known to be challenging to control. Arrival flows are also where [2,3] showed improved fuel-efficiency by automated arrival management.

The region used to define when AutoResolver began managing a flight was the airspace within a 250 nm circle centered on EWR, as illustrated in Fig. 1. The blue lines in Fig. 1 are the approach routes into EWR. EWR arrivals transitioned from controller-managed to automation-managed when crossing into the circle. A 250 nm radius was selected because at that distance from the airport most long-distance flights are at cruise altitude and have not begun descent. For short-distance flights, it was observed that some flights departing from airports close to EWR were still in climb when they crossed into the circle. Shortly after crossing the circle, long distance flights began descent. By restricting the AutoResolver control region to only the last approximately 250 nm miles until landing and in descent, a balance was struck between capturing benefits and restricting the size of the region to minimize propagation of differences between simulation and real operations.

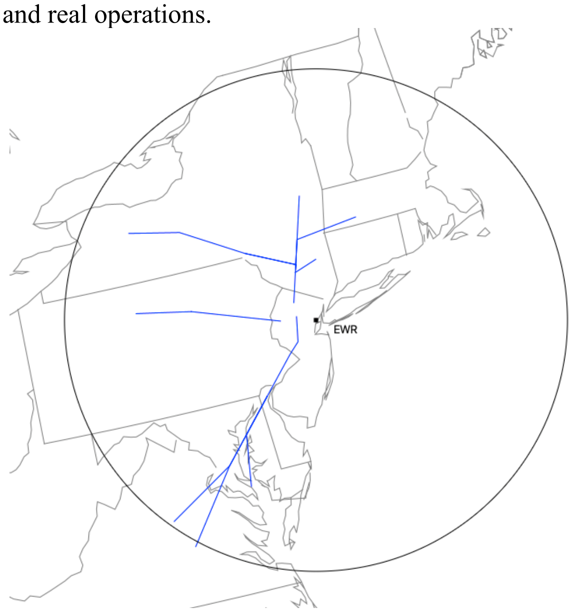

Figure 1. 250 Nautical Mile Circle

Figure 2 notionally depicts transition for a single arrival. It shows the arrival as a controller-managed target before it crosses the circle. After the arrival crosses the circle, it becomes an automation-managed target. This is analogous to handoffs between controllers at sector boundaries.

After the crossing, another target named controller-managed-reference appears in Fig. 2. The dashed green circles and connecting line show that the controller-managed-reference target corresponds to the automation-managed target. Controller-managed-reference targets and trajectories were used for visual and post analysis comparison. They represent the position and trajectory of the arrival in real operations under controller management. They were the part of the controller-managed trajectory that formed the basis for comparison with automation-managed trajectories and thus may be regarded as “truth” trajectories. They constituted the essential validation data for this study.

Although important to the fuel-efficiency comparison, the controller-managed-reference target was not visible to AutoResolver's conflict detection and resolution function. After crossing, the arrival in Fig. 2 was modeled in simulation by the automation-managed target, not the controller-managed-reference target.

The controller-managed-reference trajectory began at the transition point and time and terminated at the final approach fix (IDACE). The final approach fix was used in place of the runway threshold because track data for some trajectories were not available all the way to the runway threshold. Similarly, the automation-managed trajectory began and ended at the same points, respectively. Controller-managed-reference and automation-managed trajectories did not generally end at the same time due to differences in speeds and distances. In the fuel-efficiency comparison, fuel burn calculated for controller-managed-reference trajectories measured fuel-efficiency of controllers in real operations, and fuel burn calculated for automation-managed trajectories measured fuel-efficiency of AutoResolver in simulation.

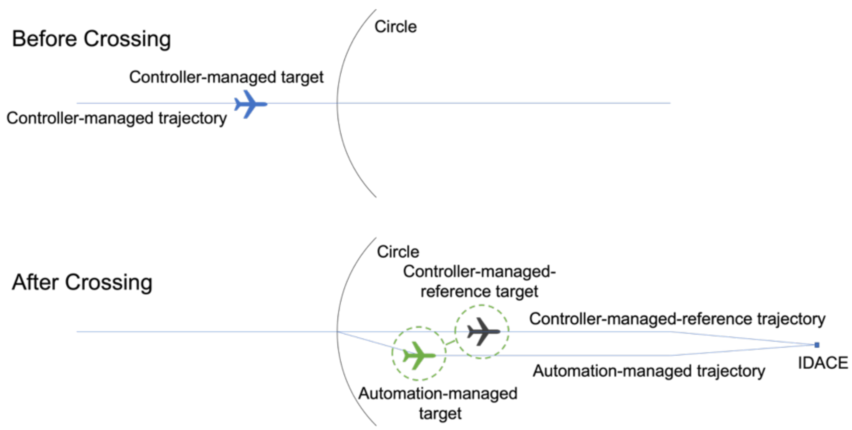

Figure 2. Transition From Controller-Managed to Automation-Managed

The simulation started at the beginning of the comparison period and ended at the end of the comparison period. At start, background flights and EWR arrivals outside the circle were controller-managed. Historical data included flight plans and tracks for all IMC flights in the NAS so the simulation included flights located throughout the NAS.

For EWR arrivals already inside the circle at the start of the comparison period, an initial transition was executed in the same way as if they had crossed the circle. The automation-managed flights were initialized with the states and flight plans extracted from historical data at beginning of the comparison period (start of simulation time). The beginning of the comparison period was selected at a time when arrival traffic at EWR was low. This was done to minimize unrealistic traffic conflicts arising from the untimely transitions taking place inside the automation-controlled airspace.

C. Simulation Tools

The simulation environment used in this study was the NASA Testbed simulation platform [27]. It was augmented with a suite of embedded simulation components known as the Autonomy Development toolKit (ADK). Both the Testbed platform and the suite of ADK components were developed at NASA Ames Research Center. Testbed and ADK included models of airspace, airports, aircraft performance, wind, weather, and atmospheric conditions. ADK contained trajectory generators [28,29], one each for enroute and terminal airspace, that propagated flights according to flight plans, aircraft performance, and atmospheric models. The enroute trajectory generator was the same used in [3].

Testbed and ADK read flight plan and trajectory data collected from recorded historical data. Routes, cruise speeds, and cruise altitudes were extracted from the flight plans and used to propagate automation-managed flights and to predict trajectories. The flight plan and trajectory data used in this study and the approach for using it is described in Section II.D.

The Testbed also included a visual display of the simulation named the Traffic Viewer [30]. It was used to monitor the simulation as it progressed and to observe differences between controller-managed-reference targets and automation-managed targets.

AutoResolver was included as a core component of ADK and was responsible for detecting and resolving schedule and separation conflicts. Testbed and ADK provided the air traffic environment for testing AutoResolver via its abilities to incorporate historical data into simulations. AutoResolver performed the functions of air traffic controllers by keeping flights separated by at least 5 nm, laterally, or 1000 ft, vertically, in enroute airspace and 3 nm, laterally, or 1000 ft, vertically, in terminal airspace. In addition, AutoResolver contained a built-in single-point scheduler. The built-in scheduler generated scheduled times of arrival at the runway threshold for flights in terminal airspace.

ADK also included the Collaborative Seamless Manager of Airspace Resources and Traffic (CSMART) [21,22]. CSMART was a multi-point scheduler. It generated scheduled times of arrival at the runway threshold and the arrival meter fixes, which are presented in Section II.E. The CSMART scheduled times of arrival were passed to AutoResolver, which used them to keep flights sequenced and spaced in the enroute airspace. The overall approach for combining the AutoResolver built-in and CSMART schedulers and using them to sequence and space arrivals is described in Section II.F.

D. Historical Data

Historical data was used in several ways in simulation. The first way was to use historical data to set the initial conditions and the flight plans of automation-managed flights at the transition points. The second way was to play back the entire trajectories of controller-managed flights.

Historical data for this paper was recorded on September 16, 2019. This was a clear-weather, general (before the pandemic) traffic day. The comparison period started at 11:00 UTC and ended at 13:30 UTC. During this time, arrivals into EWR landed towards the south. The start and end times were chosen because EWR was in south flow in between those times and that was the arrival flow direction that ADK was set up to simulate. No analysis of traffic and weather on different days was done to characterize the day that was selected.

The source of the historical data was the Traffic Flow Management Data (TFMData) Service provided by the FAA's Traffic Flow Management System (TFMS). The TFMData used in this paper was downloaded from NASA's Sherlock data warehouse [31]. TFMData contained flight plan and track reports for flights following Instrument Meteorological Conditions (IMC) rules that operated in the NAS during the comparison period. Flight plan data included routes, cruise speeds, and cruise altitudes. These were used by ADK to propagate the automation-managed flights and by the AutoResolver conflict detection function to predict their trajectories.

Track data was used to compose and replay trajectories for controller-managed flights. Track reports were reported at one-minute intervals. ADK read in track data and positioned controller-managed flights in the simulation at the track positions at the track times. This is known as simulation playback. In addition, track data was used by the AutoResolver trajectory prediction function. Instead of using a trajectory generator to build trajectory predictions for background flights, the function gathered future track hits and passed them as trajectory predictions.

The source of the wind data was Rapid Refresh (RAP). RAP contained wind vectors at points along a 13 nautical mile grid in the NAS. Wind vectors were updated every hour. The simulation read the data from file. Trajectory generators interpolated the wind data and used it to in part propagate the automation-controlled flights and predict their trajectories.

E. Newark Liberty International Airport Surrounding Airspace and Arrival Procedures

Figure 3 shows the EWR surrounding airspace and arrival procedures for south flow arrival conditions. There were three Standard Terminal ARrivals (STARs): SHAFF7, FQM3, and PHLBO3. STARs defined the route and altitude and speed constraints that each EWR arrival must follow. Flights approaching from the north used SHAFF7. Those from the west used FQM3, and those from the south used PHLBO3. The STARs ended at an arrival meter fix. SHAFF7 ended in SAX. FQM3 ended in SWEET, and PHLBO3 ended in METRO. The arrival meter fixes were located on the boundary between the enroute and terminal airspaces.

The airspace outlined in green is the class B airspace surrounding Newark, LaGuardia, and Kennedy airports. The merging of the three traffic flows from arrival meter fixes SAX, SWEET, and METRO at the initial approach fix GIBTE presented a challenging spacing problem.

Once flights exited a STAR, a controller guided them to the initial and final approach fixes, GIBTE and IDACE, via vectors and speed advisories. The orange curves in the figure depict nominal paths from the arrival meter fixes to the approach fixes. In real operations during low traffic conditions, when no scheduling conflicts occurred, traffic controllers allowed arrivals to closely follow the nominal paths. As traffic increased, path stretches were used to delay flights and achieve required spacing at the merge point.

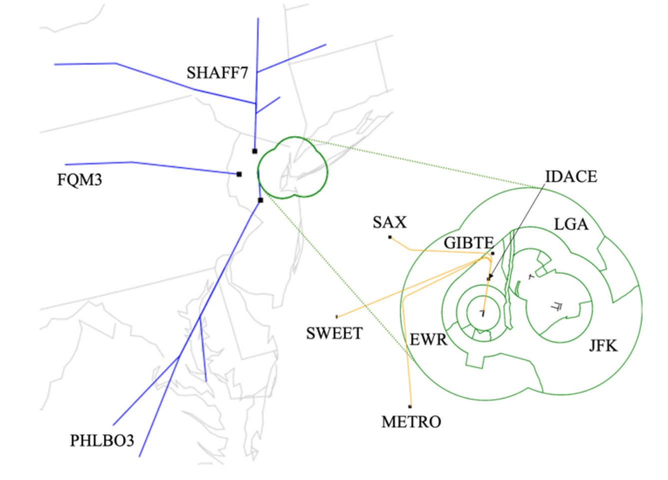

Figure 3. Newark Airspace and Arrival Procedures

During the comparison period, the airport was in south flow condition. Arrivals landed on runway 22 left, while departures took off from 22 right. The percent of arrivals through SWEET, METRO, and SAX were 35%, 30%, and 35%, respectively.

F. Scheduling

Using two schedulers to schedule arrivals reduced the impacts of meter fix crossing time errors, differences between actual and scheduled crossing times. If a flight missed its scheduled meter fix crossing time, it received an updated scheduled time at the runway threshold as it crossed the meter fix.

CSMART computed scheduled times of arrival at the runway and arrival meter for each flight as it crossed the 250 nm circle and transitioned from controller-managed to automation-managed. The arrival meter fix scheduled time of arrival was consistent with meter fix and runway spacing requirements and the time required for the flight to travel from the meter fix to the runway threshold. CSMART did not update the scheduled times of arrival, and it passed the scheduled times of arrival to AutoResolver. AutoResolver monitored flights in enroute space for meter fix scheduled time of arrival conformance and resolved, using speed changes or path stretches, any conflicts.

The AutoResolver built-in scheduler computed a scheduled time of arrival at the runway threshold for each flight as in crossed the meter fix. The runway threshold scheduled time of arrival met the runway spacing constraint, and it replaced the one previously computed by CSMART. AutoResolver monitored flights in the terminal airspace for runway scheduled time of arrival conformance and resolved non-conformances.

The AutoResolver scheduler was set to space arrivals at the runway threshold by 75 seconds, and the CSMART scheduler was set to space arrivals at the meter fixes and runway threshold by 80 seconds. The 5 seconds of extra spacing that the CSMART scheduler provided slightly throttled the flow of aircraft into the terminal. This helped AutoResolver with spacing and sequencing flights in the terminal.

III.Results

This section presents a fuel-efficiency comparison of arrival trajectories produced by AutoResolver in simulation and those produced by air traffic controllers in real operations. The following sections present metrics comparison, traffic viewer screen shots, plots for an example arrival, sequence switches, and crossing fix arrival rates.

A. Metrics Comparison

Time duration, distance, and fuel burn were computed for automation-managed and controller-managed-reference trajectories. Fuel flow rates came from Base of Aircraft DAta (BADA) version 3 [32]. The three metrics were computed for each arrival and averaged across all arrivals. There was a total of 54 arrivals. The averages are shown in Table 1.

Table 1. Time, Distance, and Fuel Burn Metrics for Controller-managed-reference and Automation-managed Trajectories

	Time (min:sec)			Distance (nm)			Fuel burn (lbs)		
	Enroute	Terminal	Total	Enroute	Terminal	Total	Enroute	Terminal	Total
Controller-managed-reference	27:50	11:40	**39:30**	190	51	**241**	2,255	705	**2960**
Automation-managed	31:24	10:21	**41:45**	209	46	**255**	2,066	549	**2615**
Difference	-3:34	1:19	**-2:15**	-19	5	**-14**	190	156	**346**

According to Table 1, automation-managed versus controller-managed-reference trajectories saved on average 190 lbs and 156 lbs of fuel burn in enroute and terminal airspaces, respectively. Savings in enroute occurred because automation-managed trajectories delayed start of descent and stayed at cruise altitude, where aircraft were more fuel efficient, longer. Enroute savings occurred despite automation-managed trajectories having longer time durations and distances on average. Automation-managed arrivals that required a delay and received a speed reduction gained additional savings because the speed reduction was achieved by reducing thrust and, hence, fuel burn rate.

In terminal airspace, both automation-managed and controller-managed-reference trajectories had similar altitudes and speeds, which made their fuel burn rates similar. Terminal airspace fuel burn savings were not due to fuel efficiency. Rather, they were due to shorter distances and time durations.

Figure 4 shows a histogram of the automation-managed vs. controller-managed total fuel savings. The arrival with fuel savings in the 2,160-2,400 lbs bin was in the heavy weight class. Its actual fuel savings was 2,398 lbs.

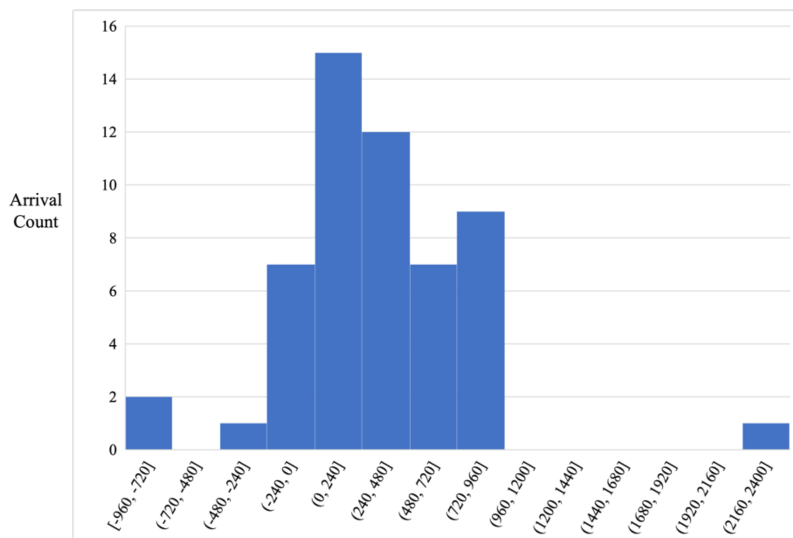

Figure 4. Automation-managed vs Controller-managed Fuel Burn Savings Histogram

B. Traffic Viewer Screen Shot

Figure 5 illustrates a Traffic Viewer screen shot zoomed out from Newark Liberty International Airport. The green curves are automation-managed trajectories, and the white curves are controller-managed-reference trajectories. The blue targets depict controller-managed flights. These flights were the background traffic within which AutoResolver and air traffic controllers managed arrivals.

Automation-managed trajectories had path stretches occurring just after entering the 250 nm circle. This happened because as the arrival crossed the circle a scheduled time of arrival at the arrival meter fix was generated by CSMART for it, and AutoResolver immediately gave the arrival a maneuver to meet the scheduled time of arrival. The automation-managed flights flew the path stretches at cruise altitude where they were fuel-efficient.

As opposed to automation-managed trajectory path stretches, Controller-managed-reference trajectories closely followed the STAR routes after they enter the circle and before crossing the meter fixes. The trajectories do not show evidence of maneuvering for spacing at the flow merge point before the final approach fix.

Figure 6 shows a Traffic Viewer screen shot zoomed in to Newark Liberty International Airport. Controller-managed-reference trajectories had wide turns in terminal airspace. In real operations, air traffic controllers used wide turns in the terminal airspace to properly space arrivals as they merged with other arrival flows just before beginning final approach to the runway. Air traffic controllers did not begin delaying flights for merge spacing until after they had crossed the meter fix and entered the terminal airspace. Automation-managed trajectories did not use wide turns into the final approach fix because they had already been spaced by the path stretches that AutoResolver gave them upstream in enroute airspace.

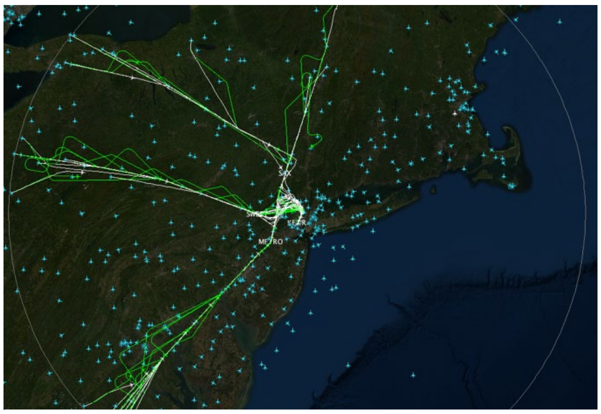

Figure 5. Traffic Viewer Screen Shot Zoomed Out from Newark Liberty International Airport

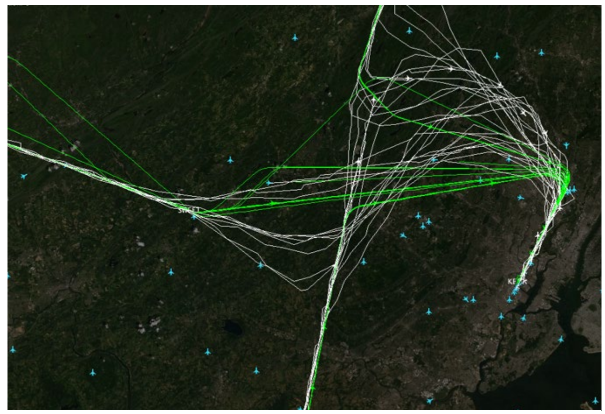

Figure 6. Traffic Viewer Screen Shot Zoomed in to Newark Liberty International Airport

C. Plots for Example Arrival

One of the arrivals that had 712 lbs fuel-burn savings was selected for illustration. Figure 7 shows the horizontal profile of the arrival. The blue curve depicts the automation-managed trajectory, and the green curve is the controller-managed-reference trajectory. The grey curves are outlines of states, and the orange curves show the class B and terminal airspace surrounding KEWR. The automation-managed trajectory had a large path stretch occurring before crossing SWEET, with a direct path to GIBTE after passing SWEET. On the other hand, the controller-managed-reference trajectory closely followed the FQM3 until crossing SWEET. After crossing, it took a wide S shaped path to get to GIBTE. GIBTE is approximately where traffic flows coming from the north and south (not shown in the figure) merge before final approach.

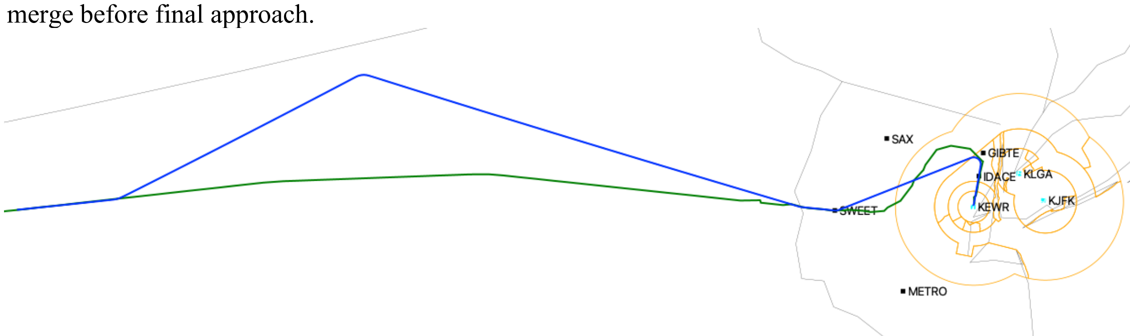

Figure 7. Horizontal Plot of Automation-managed (blue) and Controller-managed-reference (green) Arrival Trajectories

Figure 8 shows the vertical profile of the trajectories. The automation-managed trajectory began descent later than the controller-managed-reference trajectory and in general was steeper.

ADK was configured to level off the automation-managed arrivals descent at 9000 ft, which is where the arrival crossed SWEET. The controller-managed-reference trajectory had two small level offs around 7,000 ft and 6,500 ft. These level offs were part of the arrival procedure and were designed to keep the arrival flow segregated from other nearby traffic flows. ADK level off altitude was different from the operational level off altitudes. More investigation is needed to better configure the ADK level off altitude.

The controller-managed-reference trajectory has an additional level off at 30,000 ft. It is not known why the controller asked for this level off.

Figure 9 shows the indicated airspeed profiles of the trajectories. The deceleration from 290 kts to 240 kts in the automation-managed trajectory was part of the maneuver given to the arrival by AutoResolver when the arrival entered the 250 nm circle. This maneuver was for keeping the arrival's meter fix crossing time on schedule and increased the arrival's fuel burn savings, see that its savings was above average. The speed change was in addition to the path stretch, see Fig. 7.

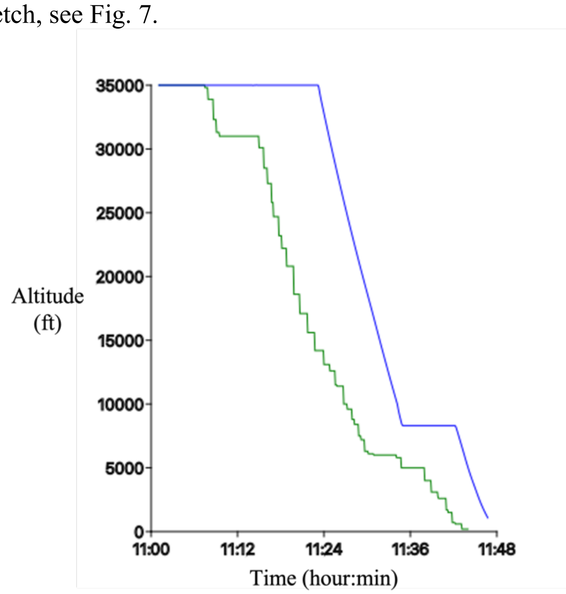

Figure 8. Altitude Verses Time for Automation-managed (blue) and Controller-managed-reference (green) Trajectories

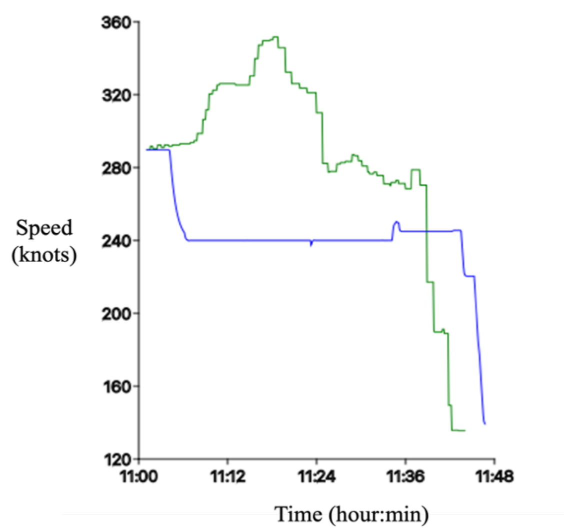

Figure 9. Indicated Airspeed Versus Time for Automation-managed (blue) and Controller-managed-reference (green) Trajectories

D. Arrival Sequence Switches

Figure 10 compares the automation-managed arrival sequence at IDACE to the controller-managed-reference arrival sequence. The markers in the figure represent IDACE crossing times relative to the start of simulation. The two markers and flat line on the left represents the automation-managed arrival time and the two markers and flat line on the right represents the controller-managed-reference arrival time for an arrival. A line connects the automation-managed and controller-managed-reference arrival times for each arrival. Instances of lines crossing in the middle of the figure represent switches of sequence between the automation-managed and controller-managed-reference sequences.

Although there are sequence switches, the first and last flights of both automation-managed and controller-managed-reference arrive at approximately the same time. This means that the all the arrivals required about the same time-duration to land.

Lines that slope up toward the left represent automation-managed arrivals that crossed IDACE later than their corresponding controller-managed-reference arrivals. In Fig. 10, there are two automation-managed arrivals that cross IDACE around 12 min and 24 min. AutoResolver delayed these arrivals too much, and they delayed the later parts of the arrival sequence. More investigation is needed to understand why these flights were delayed too much.

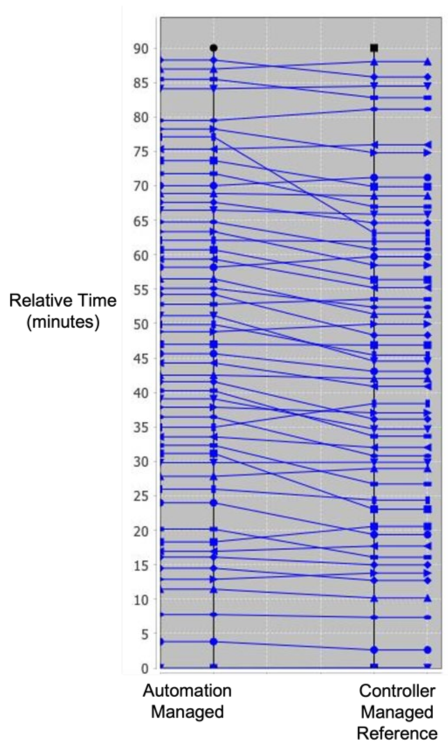

Figure 10 Relative Crossing Times at the Approach Fix (IDACE) for Automation-managed and Controller-managed-reference Trajectories

E. Crossing Fix Arrival Rates

Figure 11 shows the count of final approach fix crossings per 5-minute time bin for automation-managed and controller-managed-reference trajectories. The bins centered on 11:37:30 and 12:07:30 had 5 controller-managed crossings, whereas the greatest number of automation-managed crossing per bin was only 4. This indicates that the automation-managed crossing rate did not peak as high as the controller managed crossing rate. This may have been caused because the 80 and 75 second spacing of the automation schedulers was too large. More investigation is required.

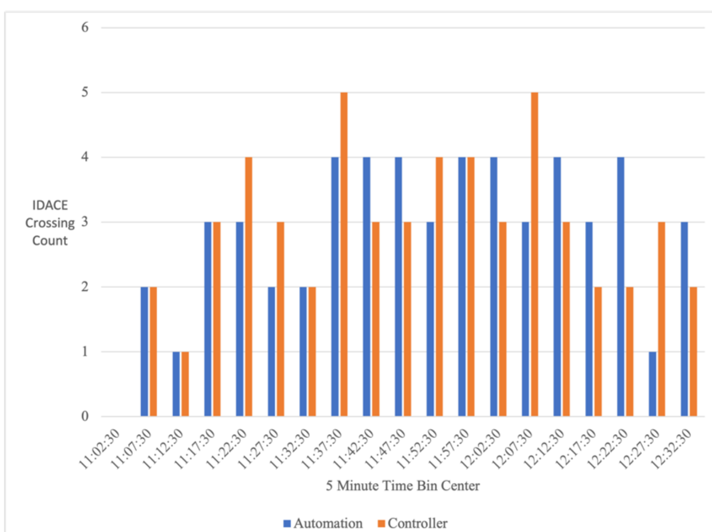

Figure 11. Final Approach Fix Crossings per 5 Minute Time Bin

IV.Next Steps

This study presents the first time automated-system fuel-efficiency was compared to air traffic controller fuel-efficiency using fast-time simulation. While many steps were taken to make the simulation traffic environment as close as possible to the traffic environment in real operations during the comparison period, future work can take the following additional steps.

The trajectory predictions in the simulation were perfect. Future simulation can model uncertainties in trajectory predictions. When trajectory predictions have errors, AutoResolver may need to issue additional maneuvers to arrivals to correct spacing errors.

The crossing runway at Newark was not used. The comparison period was selected such that there were no arrivals observed on the crossing runway at Newark. Future studies can investigate diverting arrivals to the crossing runway. This could increase the arrival rate and possibly save additional fuel by requiring fewer arrival delays.

Similarly in this study, Newark airport was landing flights toward the south during the comparison period. Future studies can perform the comparison for operations when Newark is landing flights toward the north. In addition, there was no convective weather in this study. Future studies can include convective weather.

It was observed that automation-managed trajectories had one less operation per 5-minute bin at times than the controller-managed-reference trajectories. Future studies can investigate changing the spacing parameters in the AutoResolver and CSMART schedulers to further reduce delays as well as refining the AutoResolver delay maneuver function so that it does excessively delay flights.

This study investigated a one and half hour period on a single day at a single airport. A future study could determine fuel savings at more airports and days of operations and characterize the traffic and weather on those days to produce an annualized benefits assessment for the United States.

A detailed assessment of losses of separation was not done for this study. Any losses of separation that may have occurred in the simulation would require resolution measures, which could decrease the savings result.

V. Conclusion

The fuel-efficiency of an automated system controlling Newark arrivals was compared to the fuel-efficiency of air traffic controllers managing the same arrivals in real operations. Automation-managed trajectories had 346 lbs less fuel burn on average. This savings was consistent with previous flight test [2] and human-in-the-loop [3] simulation results. Fuel burn savings in enroute airspace were due to trajectories that stayed higher in altitude where airplanes are more fuel efficient. Fuel burn savings in terminal airspace were due to shorter more direct paths from the arrival meter fix to the final approach fix. A detailed loss of separation analysis needs to be conducted to further support the savings results.

In addition to saving fuel, automation-managed trajectories produced similar arrival rates to controller-managed-reference trajectories at the runway. Future work can improve automation-managed arrival rates by adjusting the flight spacing parameters used in the automated-system scheduling algorithms and improving the AutoResolver delay maneuver function.

The fuel-efficiency savings were measured for a morning arrival rush at Newark airport. Although future studies are needed to determine a more accurate National Airspace System benefits assessment, there were no limitations in the study that indicate this fuel-efficiency savings is not achievable during other high traffic periods at Newark and at other large airports in the United States.

Not only could the automated arrival management system produce fuel-efficiency benefits, but recently deployed technologies have also made it increasingly feasible. The system requires datalink, FANS-1A, FMS, and the Time-Based Flow Management (TBFM) system (which would replace the enroute scheduler used in this paper). These technologies have already been deployed, and [3] demonstrated the integration of the automated arrival management system, FANS-1A, and datalink. Accommodations for aircraft not equipped with FANS-1A would need to be developed.

Moreover, the automated arrival management system could extend the usability and benefits of another technology the FAA is currently planning to deploy named Terminal Sequencing and Spacing (TSAS). TSAS manages arrivals in the terminal area, as opposed to the automated system studied in this paper which managed arrivals in both the enroute and terminal areas. In an operational setting, the automated system studied in this paper could be configured to handoff arrivals to TSAS as they passed from enroute to the terminal area.

References

[1] Erzberger, H., "Automated Conflict Resolution for Air Traffic Control," 25th International Congress of the Aeronautical Sciences, 2006.

[2] Coppenbarger, R., Mead, R., and Sweet, D., "Field Evaluation of the Tailored Arrivals Concept of Datalink-Enabled Continuous Descent Approach," Journal of Aircraft, Vol. 46, No. 4, 2009, pp 1200-1209.

[3] Coppenbarger, R., Hayashi, M., Nagle, G., Sweet, D., and Salcido, R., "The Efficient Descent Advisor: Technology Validation and Transition," Proceedings of the 12th AIAA Aviation Technology, Integration, and Operations Conference, September 2012, Indianapolis, IN.

[4] Coppenbarger, R., Aweiss, A., et al, "Flight Demonstration of the Tailored Arrival Manager," AIAA Aviation Forum, AIAA 2021-2373, 2021.

[5] Windhorst, R., et al, "Fast-time Simulation of An Automated Conflict Detection and Resolution Concept," 6th AIAA Aviation Technology, Integration, and Operations Conference, AIAA 2006-7825, September 2006, Wichita, Kansas.

[6] Lauderdale, T., "The Effects of Passive Coordination on Distributed Separation Assurance," 9th AIAA Aviation Technology, Integration, and Operations Conference, AIAA 2009-6908, Sept. 2009, Hilton Head, South Carolina.

[7] Lauderdale, T., Cone, A., and Bowe, A., "Relative Significance of Trajectory Prediction Errors on an Automated Separation Assurance Algorithm," Ninth UAS/Europe Air Traffic Management Research and Development Seminar, ATM 2011.

[8] Cone, A., Bowe, A., and Lauderdale, L., "Robust Conflict Detection and Resolution Around Top of Descent," 12th AIAA Aviation Technology, Integration, and Operations Conference, AIAA 2012-5644, Indianapolis, Indiana, 2012.

[9] Lauderdale, T., and Wang, T., "Coordination Between Multiple Ground-Based Separation Assurance Agents," Aviation Technology, Integration, and Operations Conference, Los Angeles, CA, 2013.

[10] Holladay, J., Lewis, T., and Lauderdale, T., "Modeling and Simulation of Function Allocation Concepts for Separation Assurance," Proceedings of the International Conference on Human-Computer Interaction in Aerospace, Article No. 2, Pages 1-9, July 2014.

[11] Erzberger, H., Lauderdale, T., and Chu, Y-C., "Automated Conflict Resolution, Arrival Management, and Weather Avoidance for Air Traffic Management," Proc. IMechE Part G: Journal of Aerospace Engineering, Vol. 226, 2011.

[12] Farley, T., Kupfer, M., et al, "Automated Conflict Resolution: A Simulation Evaluation Under High Demand Including Merging Arrivals," 7th AIAA Aviation Technology, Integration, and Operations Conference, Belfast, Northern Ireland, 2007.

[13] Farley, T., et al, "Fast-time Simulation Evaluation of a Conflict Resolution Algorithm Under High Air Traffic Demand," 7th USA/Europe Air Traffic Management R&D Seminar, Barcelona, Spain, July 2-5, 2007.

[14] McNally D., and Thipphavong, D., "Automated Separation Assurance in the Presence of Uncertainty," 26th International Congress of the Aeronautical Sciences, 2008.

[15] McNally, D., Mueller, E., Thipphavong, D., Paielli, R., Cheng, J., Lee, C., Sahlman, S., and Walton, J., "A Near-term Concept for Trajectory-based Operations with Air/Ground Data Link Communication," 27th International Congress of the Aeronautical Sciences, 2010.

[16] Gong, C., Santiago, C., and Bach, R., "Simulation Evaluation of Conflict Resolution and Weather Avoidance in Near-Term Mixed Equipage Datalink Operations," 12th AIAA Aviation Technology, Integration, and Operations Conference, AIAA2012-5618, Indianapolis, Indiana, 2012.

[17] Mueller, E., McNally, D., Rentas, T., Aweiss, A., Thipphavong, D., Gong, C., Cheng, J., Walton, J., Walker J., Lee, C., Sahlman, S., and Carpenter, D., "Controller and Pilot Evaluation of a Datalink-Enabled Trajectory-Based Operations Concept," 9th UAS/Eurpoe Air Traffic Management Research and Development Seminar, ATM2011.

[18] Aweiss, A., and Lauderdale, T., "Evaluation of an Integrated Arrival Scheduling and Automated Conflict Detection and Resolution System," 12th AIAA Aviation Technology, Integration, and Operations Conference, Indianapolis, Indiana, 2012.

[19] Lauderdale, T., Bosson, C., Chu, Y-C, et al, "Autonomous Coordinated Airspace Services for Terminal and Enroute Operations with Wind Errors," AIAA Aviation Forum, Atlanta, GA, 2018.

[20] Erzberger, H., Nikoleris, T., Paielli, R., and Chu, Y-C., "Algorithms for Control of Arrival and Departure Traffic in Terminal Airspace," Proc. IMechE Part G: Journal of Aerospace Engineering, Vol. 230, pp. 1762-1779, 2016.

[21] Windhorst, R., Zelinski, S., Lauderdale, T., Sadovsky, A., Chu, Y., Phillips, J., Zheng, Y., Nguyen, T., and D'Amore, J., "Initial Validation of a Simulation System for Studying Interoperability in Future Air Traffic Management Systems," AIAA Aviation Forum, AIAA 2021-2352, 2021.

[22] Windhorst, R., Lauderdale, T., Sadovsky, A., Phillips, J., and Chu, Y., "Strategic and Tactical Functions in an Autonomous Air Traffic Management System," AIAA Aviation Forum, AIAA 2021-2355, 2021

[23] Bosson, C., and Lauderdale, T., "Simulation Evaluations of an Autonomous Urban Air Mobility Network Management and Separation Service," AIAA Aviation Forum, 2018-3365, Atlanta, Georgia, 2018.

[24] Lauderdale, T., Predeep, P., Edholm, K., and Bosson, C., "Separation at Crossing Waypoints Under Wind Uncertainty in Urban Air Mobility," AIAA Aviation Forum, AIAA 2021-2351, 2021.

[25] Muellar, E., Kopardekar, P., and Goodrich, K., "Enabling Airspace Integration for High-Density On-Demand Mobility Operations, 17th AIAA Aviation Technology, Integration, and Operations Conference, Denver, CO, 2017.

[26] Lauderdale, T., "Probabilistic Conflict Detection for Robust Detection and Resolution," 12th AIAA Aviation Technology, Integration, and Operations Conference, AIAA 2012-5643, Indianapolis, Indiana, 2012.

[27] Robinson, J.E., Lee, A., and Lai, C.F., "Development of a High-Fidelity Simulation Environment for Shadow-Mode Assessments of Air Traffic Concepts," Royal Aeronautical Society: Modeling and Simulation in Air Traffic Management Conference, London, UK, 14-15 November 2017.

[28] Meyn, L., Windhorst, R., Roth, K., Van Drei, D., Kubat, G., Manikonda, V., Roney, S., Hunter, G., Huang, A., Couluris, G., "Build 4 of the Airspace Concept Evaluation System," AIAA Modeling and Simulation Technologies Conference and Exhibit, AIAA 2006-6110, August 2006, Keystone CO.
[29] Zhang, Y., Satapathy, G., Manikonda, V., Nigam, N. "KTG: A Fast-time Kinematic Trajectory Generator for Modeling and Simulation of ATM Automation Concepts and NAS-wide System level Analysis," AIAA Modeling and Simulation Technologies Conference, August 2110, Toronto, Ontario, Canada.
[30] Lai, C.F., "Air Traffic Management TestBed Traffic Viewer: Developer's Guide," NASA/TM-2020-220511, March 2020.
[31] Eshow, M, Lui, M., and Ranjan, S., "Architecture and Capabilities of a Data Warehouse for ATM Research," IEEE 33rd Digital Avionics Systems Conference, DASC.2014.6979418, October 2014.
[32] Base of Aircraft Data, https://www.eurocontrol.int/model/bada