

MODELING OF AIR TRAFFIC ARRIVAL OPERATIONS THROUGH AGENT-BASED SIMULATION

Sanjiv Shresta
Ralf H. Mayer

The MITRE Corporation
7515 Colshire Drive
McLean, VA 22102, USA

ABSTRACT

This paper reports on the development and validation of an agent based simulation model of air traffic control arrival operations. The simulation model includes modeling of both the structure and procedures of air traffic operations. It is thus suitable for evaluating the impacts of shifts in those structures and procedures. Three key operational metrics are introduced which are sensitive to the internal workings of air traffic arrival operations. The simulation model is validated by demonstrating agreement in those key metrics between the simulation and a set of baseline arrival operations radar data. After the simulation model has been shown to reproduce actual operations, select details of the simulation can be altered to incorporate proposed operational changes. The impact of the changes on the computer simulation will offer a prediction of how the operational changes will affect actual operations.

1 INTRODUCTION

The volume of air traffic and the demand for Air Traffic Control (ATC) services is expected to grow steadily over the next several years (Federal Aviation Administration 2006). The expected increase in demand will be felt most acutely in the already congested and highly constrained terminal airspace surrounding busy airports. It has been widely recognized that current operational structures and procedures will not be able to handle the expected increased demand in air traffic. In response, many concepts for operational shifts in ATC, aimed at increasing the efficiency of ATC operations, have been proposed.

Evaluating the effects of these proposed operational shifts is the subject of several ongoing studies (Human-In-The-Loop, AviationSimNet, etc), and presents a fruitful realm for simulation. That is especially true since ATC systems involve complex interactions between multiple stakeholders, each of which operates under a separate set of constraints and priorities, and the effects of changes in the operational structures and procedures may not be apparent prior to actual implementation of proposed opera-

tional shifts. However, given the dire possible consequences of flawed implementations and the high cost of extensive field trials, prediction of the effects of operational changes through simulation is an attractive first step to judging their suitability.

This paper describes the creation of a Monte-Carlo simulation model of terminal arrival operations which seeks to include enough detailed modeling of the structures and procedures of ATC to evaluate a wide range of proposed operational changes. The simulation model builds on the Integrated Terminal Research, Analysis and Evaluation Capability (iTRAEC) platform, which was developed at The MITRE Corporation (Mayer 2006) in the Simulation Language with Extensibility (SLX) discrete event programming language (Henriksen 1998, Wolverine 2006).

Several authors have already used simulation to evaluate operational shifts in terminal ATC arrival operations. The simulation algorithms and operational shifts evaluated include: Trajectory-Centered Simulator (TCSim) used to evaluate time-based metering (Callantine et al. 2003), Terminal Area Route Generation Evaluation and Traffic Simulation (TARGETS) used to evaluate Area Navigation (RNAV) benefits (Becher et al. 2004), Traffic Manager (TMX) used to evaluate airborne spacing (Krishnamurthy et al. 2005), the Enhanced Airfield Capacity Model (EACM) to estimate airport capacity (FAA 2004), link-node network models to evaluate airspace redesign (Boesel et al. 2004), and agent-based queuing models to evaluate reduced separation standards (Xie, Shortle, and Donohue 2004). These evaluations have demonstrated the power and utility of simulation based evaluation of novel terminal operations concepts. The simulation model of the present work seeks to add to these by creating a simulation model which includes detailed mimicking of the stakeholders in ATC and broad validation of the model by comparison with baseline operations.

The organization of this paper is as follows. Section 2 describes the development methodology followed in constructing the simulation model including some general

background description of a typical terminal ATC arrival operation and three key operational metrics. Section 3 describes the details of the simulation model, which is then validated in section 4. Finally, section 5 discusses some on-going evaluations that have been accomplished using the simulation model.

2 MODEL DEVELOPMENT

In order to enable evaluation of changes in the structure and procedures of arrival ATC, the agent-based simulation described here was designed to mimic the tasks of air traffic controllers with sufficient fidelity to recreate details relevant to operational shifts. Thus, it was necessary to model the details of how arriving aircraft are guided and fly through the terminal airspace, including the different stakeholders involved in arrival operations and their behavior and interactions. Modeling of the agents was accomplished by comparison with a set of baseline operational data and consultation with subject matter experts. This section reviews the methodology used to construct the model agents which reproduce the functions and decisions of stakeholders in terminal ATC arrival operations.

2.1 Case Study: Hartsfield-Jackson Atlanta International Airport

The simulation model described in this paper is applicable to a range of airspace configurations. However, since a concrete set of arrival operational data was required for comparison and validation purposes, the development has been based on arrival operations at Hartsfield-Jackson Atlanta International airport (ATL). The choice of ATL was made for two reasons. First, ATL arrival operations has a commonly used runway layout and airspace configuration, which allows comparison with several other airports. Second, ATL arrival rates are often near the maximum sustainable rate. This fact ensures that a simulation validated against ATL arrival operations will have been stress tested in high density traffic situations.

The structure of arrival operations at ATL has a four corner post terminal airspace design and a dual parallel runway landing layout. Arriving aircraft are funneled into the terminal airspace over one of four waypoints located at the corners of the terminal airspace. As shown in Figure 1, the arriving aircraft are directed from the corner posts down to one of two landing runways.

A simplified description of ATC procedures in a typical aircraft arrival begins with the choice of runway on which to land the inbound aircraft, which is made by a traffic flow manager before the aircraft reaches a corner post, and is based on pilot request and the need to balance the throughput on the available runways. The arriving aircraft is then issued directional vectors and speed adviso-

ries by a terminal air traffic controller, who guides the aircraft down to the final approach course of the assigned runway while maintaining safe separations from all other concurrently arriving aircraft. The dispersion of aircraft arrival tracks seen in Figure 1 illustrates the range of vectoring experienced by a typical arriving aircraft in the baseline operations.




Figure 1: A sample of arriving aircraft radar tracks from the ATL operations used as a baseline to create the simulation model is shown. All the tracks over the indicated 3 hour period are superimposed on top of each other.

2.2 Key Operational Metrics

From the track data, metrics can be computed which characterize the internal workings of how terminal arrival operations are conducted. Such metrics are of particular interest in creating a computer simulation, since they can be used to guide development of the model. The metrics used in this work are discussed in the following subsections.

2.2.1 Separation

The separations imposed between aircraft by terminal controllers can be characterized by the **inter-arrival separation distribution** of aircraft as they fly along the final approach course. This distribution is a good measure of the applied separations since it is sensitive to not only the average separations applied, but also to the variability of the applied separations around the average. The variability of aircraft separations can have many sources, including the different required minimums between different aircraft types and the deviation of aircraft from their assigned routes and speeds. All of these sources are captured by the inter-arrival separation distribution along the final approach. The final approach is particularly important because it is where aircraft are moving most slowly, are most closely spaced, and have their routes most constrained.

2.2.2 Vectoring

In current operations, one of the methods used by terminal controllers to establish separation between arriving aircraft is to issue radar vectors for the aircraft to fly. Vectoring allows the controllers to delay aircraft by increasing their track lengths. A good measure of how much vectoring is applied is thus the **distribution of track lengths** of arriving aircraft. As for the inter-arrival separations, the distribution of track lengths is sensitive to not only the average track lengths, but also to the variation around the average. It contains information about how often straight-in arrivals were conducted, how often longer arrivals were flown, and how often all track lengths in between were observed.

2.2.3 Speed Control

In addition to vectoring, terminal controllers can maintain separations between arriving aircraft by assigning speeds. As for vectoring, speed assignments allow the terminal controllers to control the delay arriving aircraft experience in the terminal airspace. However, speed assignment can not be separated from vectoring, since assigning a speed imposes delay only after some distance has been flown. Rather than characterizing speed assignment on its own, the speed assignment descriptions should be made conditionally for each track length. A good way to accomplish that measurement is to construct a **scatter plot**, placing track length, as a measure of vectoring, along the horizontal axis and average speed, as a measure of speed assignment, along the vertical axis. The scatter plot will have a point for each arriving aircraft. How those points are distributed on the scatter plot can then be compared between the simulation and actual operations.

3 MODEL DESCRIPTION

3.1 Input Data

The simulation of this work allows input traffic which defines the fleet mix, sequence and timing of flight arrivals over the corner posts. The fleet mix used in the validation of the simulation was the daily fleet mix from baseline operations. Using the baseline input traffic guaranteed that the simulation was compared to actual operations for the same arrival demand.

In order to match delivery altitudes, the altitudes of arriving aircraft as they cross over a corner post were set to the most commonly observed altitude in baseline operations over the same corner post. In the case of delivery speeds, the model air traffic controller was allowed to assign speed over the corner post, but was constrained to stay within the range of speeds observed in the baseline data. This insured that flights entered the terminal air-

space with a range of speeds that matched the range of speeds in baseline operations.

3.2 Model Flow Manager

The flow manager is responsible for several tasks which are completed before an arriving aircraft is handed off to the model terminal controller. The tasks include assigning a runway, sequencing flights, and assigning the default route and speed. The following sections describe how the model terminal air traffic controller completes the runway balancing and sequence optimization tasks.

3.2.1 Runway Balancing

The model flow manager chooses the landing runway for each inbound aircraft based on the need to balance throughput on the available runways. Balancing the throughput on all the runways of an airport is important because overwhelming one runway while another remains underutilized will lower the overall throughput of the airport. In order to choose which runway would be best, based on balancing considerations, the model flow manager first calculates an estimated demand on each runway and then decides whether re-assigning the aircraft to other than the nearest runway would be desirable. The estimated demand is calculated as a weighted moving average constructed so that nearer flights contribute more to the estimated demand than distant flights, as illustrated in Figure 2. The aircraft currently being assigned is the one on the far right. The other aircraft are those that are scheduled to land before it on the runway being considered. The number of aircraft included in the weighted average is called the demand number and is set to 15 flights in the simulation. The flow manager uses the expected inter-arrival times (denoted by T_j) over the runway threshold to calculate the weighted demand estimate (denoted by D) as shown.




Figure 2: Diagram of the expected demand calculation.

Once estimated demands have been calculated for each available runway, the flow manager will reassign the incoming flight to the runway with the lowest estimated demand if the demand estimate on that runway is lower than the demand estimate on the nearest runway by at least a specified amount called the demand barrier. The demand barrier used in the simulations is 2 flights per

hour. The purpose of the demand barrier is to ensure that flights are not reassigned to other than their closest runway unless there is a clear advantage in doing so.

3.2.2 Local Sequence Optimization

By default, the model flow manager sequences flights in first-come-first-served order based on expected runway arrival time. However, the flow manager may increase overall throughput by locally optimizing the sequence of incoming flights. For example, if there are two large and two small aircraft to be landed, the cumulative inter-arrival separations will be lower and the throughput higher if the two small aircraft are landed first, followed by the two large aircraft. The decreased total separation comes from the fact that a large aircraft following a small aircraft has a lower minimum required separation than a small aircraft following a large aircraft. The flow manager can be allowed to institute a limited version of sequence optimization by transposing immediately sequential flights when a throughput gain can be attained.

3.3 Model Air Traffic Controller

The model terminal air traffic controller is responsible for several tasks which are completed after the arriving aircraft is handed off by the model flow manager but before it reaches a corner post. The tasks include assigning a route and speed to ensure separation and re-sequencing flights arriving in different streams. The following sections describe how the model terminal air traffic controller completes each task.

3.3.1 Speed Assignment

Speed assignment in the computer simulation of arrival operations is implemented in a simplified manner. Unlike in actual operations, the model terminal air traffic controller may assign a speed at only a single point along the arriving aircraft's route, namely the speed when the aircraft crosses over a corner post. The speed of the arriving aircraft along the rest of the route follows from the flight profile of the arriving aircraft as it descends and decelerates toward the runway. For each aircraft type, the range of assignable corner post speeds is bounded by the range of speeds seen in actual operations and according to the particular aircraft performance specifications.

As observed in actual operations, the speeds assigned are loosely dependent upon the route, such that the quickest arrivals will be the fast and short ones, while the most prolonged arrivals will be the slow and long ones. The advantage of using route dependent speed assignment method is an extended range of delay available in the terminal airspace. Delay in this context is defined as the extra time required for an aircraft to fly from the corner post to

the runway beyond the time required on the default route and speed. This principle is illustrated in Figure 3. On the horizontal axis is the track length and on the vertical axis is the average speed. Marked on the track length axis are the shortest (d_l) and longest (d_h) possible track lengths. Marked on the average speed axis are the fastest (v_h) and slowest (v_l) possible average speeds. The interior of the rectangle created by the intersection of the limits defines the space of assignable speed and route length pairs. The radial lines extending out from the origin are lines of constant time flown. The two red lines demarcate the lowest (t_l) and highest (t_h) time flown possible. The range between the red lines is the maximum range of delay that can be imposed on a flight in the terminal airspace. In order to extend the utilized delay as widely as possible, the average speed can be tied to the length of the route so that the pair fall on the diagonal line inside the rectangle. Doing so allows the terminal controller to use the available airspace efficiently.

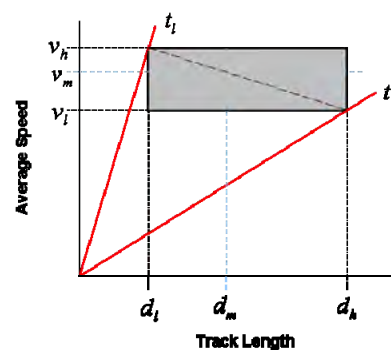


Figure 3: Route length dependent speed assignment.

3.3.2 Route Assignment

The mimicking of vectoring is accomplished in the computer simulation by offering the model terminal air traffic controller a range of flight paths which aim to reproduce the flight tracks found in actual operations. The range of flight tracks in actual operations and the range of routes matched to them are shown in Figure 4. The paths in the simulation include mimicking of base leg and upwind leg vectoring on the short-side arrivals, as well as crosswind leg and downwind leg vectoring on the long-side arrivals.




Figure 4: Example tracks in the baseline operations and the routes which mimic them in the simulation.

Route assignment is then accomplished by searching for the shortest route which avoids loss of separation conflicts. Once such a route is found, it is assigned to the arriving aircraft, and the speed is set in accordance with the route length concurrently. If no suitable route is found, the aircraft is assigned to the route which has the least severe loss of separation conflict. Severity of conflict used in the simulation is similar to Federal Aviation Administration (FAA) Order 7210.56C, in which greater spacing violations give greater severities. The simulation defines severity as the square of the normalized amount by which the minimum required spacing is violated.

3.3.3 Streamwise Re-sequencing

The model terminal controller has a limited ability to re-sequence the landing order of arriving flights. The limitation is that the terminal controller may only re-sequence aircraft in different streams. The reason for that restriction is that having flights overtake each other in the terminal airspace is undesirable. Interchanging the order of flights on different streams can be advantageous if the airspace available for aircraft arriving from one direction is more constrained than from another. In ATL operations that situation exists. Aircraft arriving on the short-side, which in west flow configuration includes eastern arrivals, have less maneuvering space than those arriving on the long-side from the west (see Figure 4).

This principle is illustrated in Figure 5. In this example, interchanging the order of landing will allow the terminal controller to transfer delay from a constrained stream to a less constrained stream. The aircraft at the far left is the lead aircraft. The next two aircraft must be assigned routes and speeds so they are properly spaced behind the lead aircraft and behind each other. The green boxed aircraft is a long-side arrival and D_L is the range of delay that is available on the long-side. The red boxed aircraft is on the short-side and D_S is the range available on the short-side. As they are ordered in the figure, the green flight is scheduled to land before the red flight, at the times labeled ETA_L and ETA_S , respectively. If the terminal controller leaves the order as it is, the green flight will be delayed to the blue line labeled ATA_L and the red flight will lack sufficient delay range to be properly spaced behind the green flight. The solution is to interchange the landing order of green and red flights. All flights can then be properly spaced. This method of re-sequencing landing order will give no advantage for flights in the same stream, since they have the same amount of delay available.




Figure 5: Stream-wise re-sequencing diagram.

3.4 Model Aircraft

After the model flow manager and model terminal air traffic controller have completed all the needed assignments, the model aircraft will fly the assigned route and speed. The trajectory model used in the simulation is the trajectory model of the iTRAEC platform (Mayer 2006). During arrival, the trajectory model adjusts the deceleration and descent of the arriving aircraft according to the performance specifications of the aircraft type and the phases of flight through which an aircraft must progress as it approaches and lands at a runway. Those phases begin with a descent, in which the aircraft slows and descends to the glide slope intercept speed and altitude, followed by glide slope intercept, then final approach, in which the aircraft slows to landing speed, and finally touchdown.

The variability of aircraft conformance to the assigned routes and speeds is included in the computer simulation as an additional external parameter. The variability is simplified, and is implemented as an in-trail position uncertainty, which does not include lateral and vertical uncertainty. The in-trail variability combines the cumulative effect of all variations in aircraft conformance, including the effect of wind. That the lateral and vertical position uncertainties can be neglected is an assumption of the simulation model, and is justified by the agreement between baseline operations and the simulation of baseline operations. A diagram illustrating variability in aircraft position is shown in Figure 6. In that figure, each aircraft is spaced behind the one in front of it by a required separation plus a buffer separation. Each aircraft also has a variability of their actual position around their expected positions. The actual position of each aircraft can be anywhere within the envelope of variability. The exact location of the aircraft is unknown by both the model flow manager and the model terminal controller.




Figure 6: Buffer and aircraft performance variability.

3.5 Chronology

As an incoming flight approaches the destination airport it is controlled by the model flow manager and by the model terminal air traffic controller. Each performs tasks on the arriving aircraft related to guiding it to a safe landing. The progression in which the aircraft is controlled is determined by the planning horizons of the model flow manager and model terminal controller. The chronology is as follows.

First, the model flow manager takes control of arriving aircraft when they are some specified time interval away from reaching a corner post. The time defines a flow manager planning horizon, which is illustrated schematically in Figure 7. Before an incoming flight reaches the flow manager planning horizon the flow manager has no knowledge of it and does not consider it in any decisions. After the incoming flight has reached the flow manager planning horizon, but before it reaches the terminal controller planning horizon, the flow manager is able to control the flight and begins the assignment tasks it must complete for the flight. All assignment tasks are finished before the incoming flight has reached the terminal controller planning horizon.

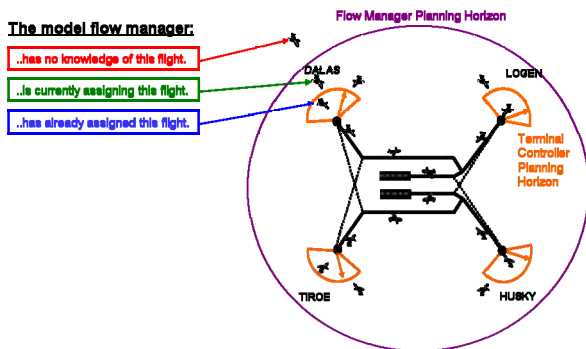


Figure 7: Planning and freeze horizon of the model flow manager.

The hand-off from the model flow manager to the model terminal controller takes place at the terminal controller planning horizon, which is illustrated schematically in Figure 8. The terminal controller planning horizon is

also defined by a specified time interval before a flight reaches a corner post. As for the flow manager, the time interval for the terminal controller is a tunable parameter and is called the terminal horizon interval. Prior to a flight crossing the terminal controller planning horizon, the terminal controller has no knowledge of it and does not consider it in any decisions. After the arriving flight has reached the terminal controller planning horizon, but before it reaches the corner post, the terminal controller begins the assignment tasks for the flight. All assignment tasks are finished before the arriving flight has reached the corner post.

After a flight passes the corner post the model aircraft flies the assigned route and speed according to the specific aircraft type characteristics. Variability in the flight performance due to influences such as wind is included as an in-trail position uncertainty, and was previously discussed (see Section 3.4).

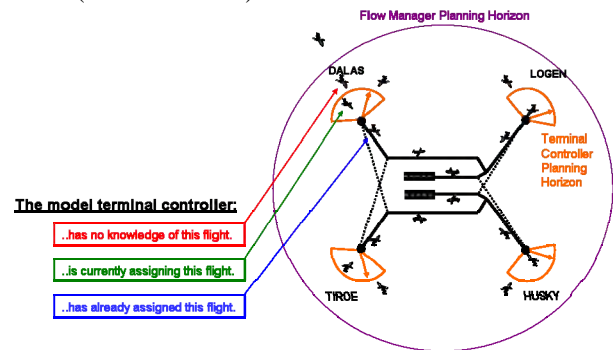


Figure 8: Planning and freeze horizon of the model terminal air traffic controller.

4 MODEL VALIDATION

Validation of the computer simulation is a critical step since the reliability of predictions derived from it depends on the fidelity of the simulation to actual operations. Validation is achieved by demonstrating agreement of the simulation with actual operations in four key operational aspects.

4.1 Separation

Validation of separation enforcement in the computer simulation of arrival operations is achieved by comparison with separation enforcement in baseline operations, as shown in Figure 9. The red curve is the distribution of inter-arrival times for baseline operations along the final approach course. The blue curve is the distribution of runway threshold inter-arrival times generated by the simulation of baseline operations. As can be seen, they agree well. Of particular interest are the left edges of the two distributions which describe the separation between the most closely spaced aircraft. The two distributions

match almost exactly, indicating that the simulation emulates the separation enforced between aircraft well.




Figure 9: Distribution of arriving aircraft inter-arrival times along the final approach course.

4.2 Vectoring

Validation of vectoring in the computer simulation is achieved by comparison of the track length distributions in baseline operations with those generated through simulation of baseline operations. The track length distributions are a sensitive measure of vectoring since greater vectoring leads to longer track lengths and less vectoring leads to shorter track lengths. Thus the distributions of track lengths measure the distributions of amounts of vectoring. The comparisons are shown for short-side streams in Figure 10, and for long-side streams in Figure 11 and Figure 12. Alongside each distribution the streams included in the distribution are illustrated. In each plot, the red curve is the distribution in baseline operations, and the blue curves are the distributions generated through the simulation of baseline operations.

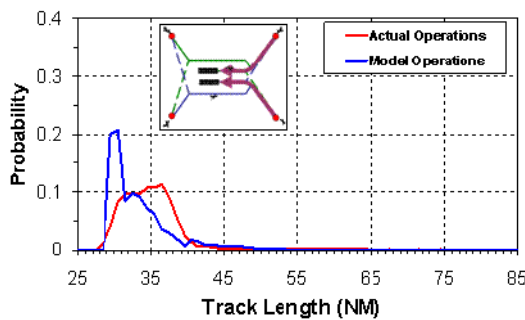


Figure 10: Distribution of track lengths flown along the flow streams indicated in the inset diagram.




Figure 11: Distribution of track lengths flown along the flow streams indicated in the inset diagram.

As can be seen, the range over which the distributions extend match well in all the sets of streams. The distributions themselves match very well on the long-side. However, on the short-side, the track length distribution tends toward shorter lengths in the simulation than in baseline operations. The reason for that is the preference of the simulation to always assign the shortest possible approach in all cases. Despite that mismatch, the short-side distribution of the simulation is in reasonably good agreement with baseline operations in the range of track lengths assigned. Combined with the near perfect agreement on the long-side, this comparison shows that the simulation emulates vectoring in baseline operations well.




Figure 12: Distribution of track lengths flown along the flow streams indicated in the inset diagram.

4.3 Speed Control

Validation of the speed control in the computer simulation of arrival operations is achieved by comparison of the speed versus track length scatter plots in baseline operations with those generated through simulation of baseline operations. The scatter plot comparisons are shown for short-side streams in Figure 13, and for long-side streams in Figure 14 and Figure 15. In each figure, the streams which are included in the scatter plot are illustrated inset to the scatter plots. Each red circle is an average speed and track length pair from baseline operations. Each blue

circle is a speed and length pair from the simulation of baseline operations. Inspection of the scatter plots shows good agreement between the simulation and baseline operations. In each case the blue circles of the simulation fall on top of the red circles of actual operations. This shows that speed control, as implemented in the simulation, emulates speed control in baseline operations well.




Figure 13: Scatter plot of speed assignment along the flow streams indicated in the inset diagram.

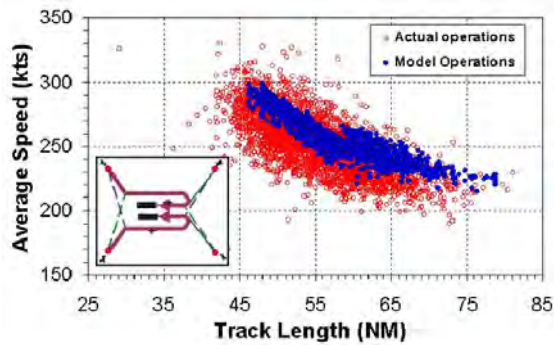


Figure 14: Scatter plot of speed assignment along the flow streams indicated in the inset diagram.

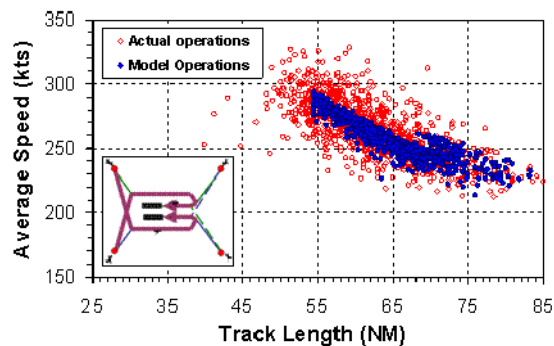


Figure 15: Scatter plot of speed assignment along the flow streams indicated in the inset diagram.

4.4 Capacity

As a final validation metric, the airspace capacity predicted in the simulation can be compared against what is observed in actual operations. In the simulation, capacity is estimated by the maximum sustainable throughput at infinite demand (Swedish 1981). In the case of actual operations, each facility reports what it has observed to be the maximum throughput in saturated demand conditions. Those numbers are listed for ATL in Table 1. The facility reported capacity and the capacity as estimated by the Enhanced Airfield Capacity Model (EACM) were obtained from the Airport Capacity Benchmark Report of 2004 (FAA 2004). All three numbers are similar, indicating that the capacity estimate of the simulation developed in this study is in agreement with those estimates.

Table 1: Comparison of capacity predictions.

Source	Capacity (flights per hour)
Actual Operations	92
Model Operations	89
EACM	88

5 CONCLUSION

The computer simulation of arrival operations developed in this study emulates actual arrival operations in both form and function. In form, the simulation models the structure of terminal arrival routes and the schedules and fleet mix of incoming traffic. In function, the simulation mimics the stakeholders involved in arrival operations, including flow managers, terminal air traffic controllers, and aircraft. The stakeholders are modeled by agents who complete their tasks and interact with each other in a way that follows actual operations. The result is a computer simulation which is realistic and flexible enough to evaluate the impacts of structural and procedural changes, as demonstrated through extensive validation.

The simulation model of this study is currently being used to quantify some of the challenges that will be faced in high density terminal operations. As a first evaluation, the operational consequences associated with metered delivery of arriving flights has been analyzed, answering the question of what level of metering error can be tolerated in metered delivery operations and how much vectoring flexibility will be required to absorb errors in metering. Subjects of ongoing work include investigation of offset vectoring in confined terminal corridors, the effect of Required Time of Arrival (RTA) on ATC operations, and the impact of aircraft performance on Continuous Descent Arrival (CDA) operations. In general, the terminal arrival operations simulation model of this study provides a framework for evaluating a range of terminal domain operational concepts.

ACKNOWLEDGEMENTS

The authors are grateful for the support and direction of Dr. Tom Becher and the numerous others at The MITRE Corporation's Center for Advanced Aviation System Development (CAASD) and the FAA whose expert advice was invaluable.

The contents of this material reflect the views of the author and/or the Director of the Center for Advanced Aviation System Development. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of the views expressed herein.

REFERENCES

- Becher, T. A., B. M. Hatli, K. R. Sprong, J. S. DeArmon. 2004. Simulating terminal RNAV: Assessing time and distance benefits. In *Proceedings of the 2004 Summer Computer Simulation Conference*, ed. Agostino G. Bruzzone and Edward Williams, 405-411. Society for Computer Simulation.
- Boesel, J., and D. Bodoh. 2004. Simulating Airspace Re-design for Arrivals to Detroit-Wayne County Airport. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R. G. Ingalls, M. D. Rossetti, J. S. Smith and B. A. Peters, 1318-1325. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers.
- Callantine, T. J., and E. A. Palmer. 2003. Fast-time simulation studies of terminal-area spacing and merging concepts, In *Proceedings of the 22nd Digital Avionics Systems Conference*, 3-10.
- Federal Aviation Administration. 2004. *Airport Capacity Benchmark Report 2004* United States Department of Transportation, Washington, DC.
- Federal Aviation Administration. 2006. *FAA Aerospace Forecasts 2006-2017* United States Department of Transportation, Washington, DC.
- Henriksen, J. O. 1998. Stretching the boundaries of simulation software. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D. J. Medeiros, E. Watson, M. S. Manivannan, and J. Carson, 227-234. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers.
- Krishnamurthy, K., B. Barmore, and F. Bussink. 2005. Airborne Precision Spacing In Merging Terminal Arrival Routes: A Fast-Time Simulation Study. In *Proceedings of the 6th USA/Europe Air Traffic Management R&D Seminar*, 1-10. Baltimore, MD.
- Mayer, R. H. 2006. Estimating operational benefits of aircraft navigation and air traffic control procedures using an integrated aviation modeling and evaluation platform. In *Proceedings of the 2006 Winter Simulation Conference*, ed. L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 1569-1577. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers.
- Swedish, W. J. 1981. *Upgraded FAA Airfield Capacity Model - Volume I, Supplemental User's Guide; Volume II, Technical Description of Revisions*, MTR-81W16, The MITRE Corporation, McLean, VA.
- Wolverine. 2006. Wolverine Software Corporation, Alexandria, Virginia.
- Xie, Y., J. Shortle, and G. Donohue. 2004. Airport terminal-approach and capacity analysis using an agent-based model. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R. G. Ingalls, M. D. Rossetti, J. S. Smith and B. A. Peters, 1349-1357. Piscataway, New Jersey: Institute of Electrical and Electronic Engineers.

AUTHOR BIOGRAPHIES

SANJIV SHRESTA is a Senior Simulation Modeling Engineer at The MITRE Corporation's Center for Advanced Aviation System Development (CAASD). He holds a B.S. and M.S. in Mechanical Engineering from Virginia Tech and Texas A&M University, respectively, and a Ph.D. in Physics from the University of Maryland at College Park. At CAASD he works on agent-based simulation modeling, fuel flow and trajectory modeling, and analyses of aircraft separation. He can be contacted by e-mail at sshresta@mitre.org.

RALF H. MAYER leads the RNAV/RNP Operational Assessment and Simulation Group at The MITRE Corporation's Center for Advanced Aviation System Development (CAASD). He received a Ph.D. degree in experimental physics from Purdue University, has over 10 years of research experience in physics and aviation, and is an author of over 30 scientific journal publications. Dr. Mayer holds commercial pilot and flight instructor certificates and served on the staff of Purdue University's Aviation Technology Department before joining CAASD where he has worked on capacity/delay/fuel burn benefit evaluations of airspace and procedure designs. He can be contacted by e-mail at rmayer@mitre.org.