Proceedings of the 2004 Winter Simulation Conference
R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

AIRPORT TERMINAL-APPROACH SAFETY AND CAPACITY ANALYSIS USING AN AGENT-BASED MODEL

Yue Xie
John Shortle
George Donohue

Department of Systems Engineering and Operations Research
George Mason University
Fairfax, VA 22030, U.S.A.

ABSTRACT

The consistent growth of air traffic demand is causing the operational volumes at hub airports to approach their maximum capacities. With this growth, delays are increasing, and safety is becoming a more crucial problem. The terminal approaching and landing phases are especially important since the airspace is more crowded and operational procedures are more complicated compared with the en route phase. We have developed an agent-based stochastic simulation model which is useful to analyze the relationship among airport arrival capacity, delay, and safety. We first present a simplified queue model to demonstrate key ideas. Then, we give a detailed agent-based model that is calibrated to Hartsfield Atlanta International Airport. We use the model to evaluate several operational scenarios and examine the trade-offs between system capacity and safety.

1 INTRODUCTION

With the consistent growth of air traffic demand, the operational volumes in hub airports are approaching their maximum capacities. Evidences of the saturation include the decline of airline on-time (DoT 2004), increased air traffic controllers' workload, and more separation violations (Haynie 2002). Compared with the en route phase, the terminal area has a very complex operating environment. This is due to shorter separation between aircraft, more complicated airplane routes with ascending, descending, turning, and accelerating aircraft, and multiple hand-off procedures between controllers. Thus, terminal delays contribute substantially to overall capacity problems in the National Airspace System.

Airport capacity and delay have been analyzed extensively. For example, the NASA ASAC (Aviation System Analysis Capability) Airport Capacity and Delay Models support the analysis and evaluation of new technologies that are designed to address airport capacity constraints due to runway occupancy times and minimum airborne separations (Lee, Nelson, and Shapiro 1998). The ASAC capacity model uses a family of departure-arrival capacity curves that apply under various meteorological conditions or with different levels of technology. If more detailed operations and situations need to be considered, another model, TAAM (Total Airspace and Airport Modeler), provides facilities to customize the modeled airports and airspace. TAAM can be configured to assess the performance of an airport under various operating scenarios. To meet the specific requirement of analysis on airport terminal TRACON airspace, MITRE CAASD has been developing a model to evaluate the capacity and delay benefits from innovative operating procedures, airspace redesign, or navigational improvements (Boesel 2003).

However, when safety becomes an issue to consider, very few models analyze safety metrics in combination with capacity and delay. With regard to safety, we are concerned with events such as wake-vortex separation violations, simultaneous runway occupancies (SRO's), and collisions on the runway. Furthermore, the aviation system is a complicated stochastic system.

The purpose of this paper is to understand the trade-off between safety and capacity. That is, we seek to understand what happens to safety when throughput at an airport is increased. To do this, we first give a simplified queuing model to describe basic system characteristics and explain qualitatively observed inter-arrival times at the runway threshold. Then, we give a detailed agent-based model that is calibrated to Hartsfield Atlanta International Airport. We use the model to evaluate several operational scenarios and examine the trade-offs between system capacity and safety. We compare these trade-offs under different agent models for the local controller.

2 DATA DESCRIPTION

Flight Explorer is a tool that allows the user to display information about commercial, passenger, and private flights on IFR flight plans. The Explorer screen refreshes every 10 seconds with any position changes that have been received in the previous 10-second window. We collected data using Flight Explorer in three randomly chosen workdays in the year 2003: 5/7/2003, 5/24/2003, and 6/28/2003. The recording period in each day was 7:00 through 24:00 (local time). The data accuracy is subject to the refresh frequency of the Flight Explorer mentioned above and an observation error around ± 5 seconds.

Figure 1 shows a snapshot of ATL terminal airspace. We identify the time when an aircraft passes the TRACON boundary as TRACON_time. The time between two successive aircraft's TRACON_time is TRACON_IAT. We identify the time when an aircraft finishes the downwind leg and turns to the base leg as BASE_time. The time difference between two successive aircraft's BASE_time is BASE_IAT. The time between an aircraft's TRACON_time and BASE_time is identified as FLIGHT_length.

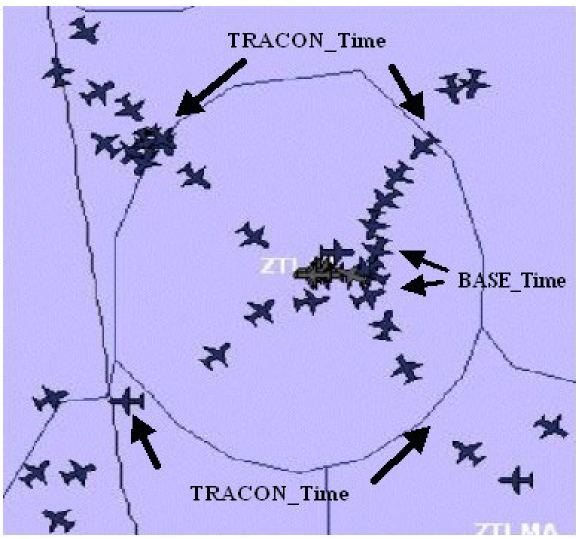

Figure 1: A Snapshot of ATL Terminal Airspace (Courtesy of Flight Explorer)

We recorded TRACON_time and BASE_time for each aircraft, and then calculated TRACON_IAT, BASE_IAT and FLIGHT_length. Tables 1 and 2 show the results of this collection. We fitted TRACON_IAT of each stream with an exponential distribution, and FLIGHT_length with a Gaussian distribution. The average arrival rate for the two runways combined is about *58* approaches per hour.

Table 1: TRACON_IAT

Stream	Mean (sec.)	std.dev	# of data points
Northeast	199	195	908
Northwest	232	269	818
Southwest	354	405	541
Southeast	252	256	721

Table 2: FLIGHT_length

Stream	Mean (sec.)	std.dev	# of data points
Northeast	563	79	911
Northwest	780	133	821
Southwest	673	95	544
Southeast	548	94	724

Figure 2 shows the distribution of landing time intervals on a runway at ATL. This data is from a separate study (Haynie 2002) in which 364 landings were observed on location at the airport. One simultaneous runway occupancy (SRO) was observed from this set. A rough estimate for the probability of a SRO is 1/364 = 0.0027, which is consistent with a 90% confidence interval 0.0035 ± 0.0007 calculated from a landing simulation model (Xie, Shortle, and Donohue 2003). This study also observed runway occupancy times which were fitted using a Gaussian distribution $N(48, 8^2)$ in seconds.

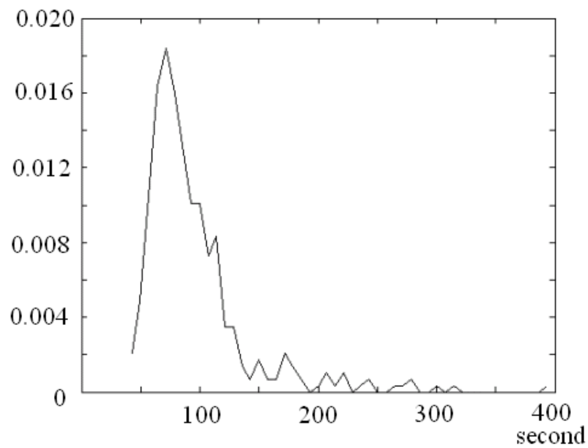

Figure 2: Histogram of Landing Time Interval from Haynie's Observation

3 A SIMPLIFIED MODEL

In this section, we give a simplified model that predicts the distribution of inter-arrival times at the runway (e.g., Figure 2) based on inter-arrival times at the TRACON (e.g., Table 1). Figure 3 summarizes the model. We assume the two arrival runways of ATL are independent and equivalent, and that each aircraft chooses either runway with equal probability, so we only need to consider one of the runways. The arrival stream has an exponential inter-arrival time, and the service time includes two parts. One is the separation S between aircraft at the runway threshold, which is modeled as an M/G/1 queue, where the service time represents the separation between two aircraft applied by the air traffic controller (ATCo). The other is the runway occupancy time, modeled as a random delay. Each service time is modeled as a Gaussian distribution.

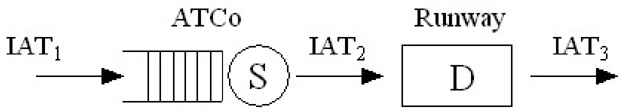

Figure 3: A Simplified Queuing Model

IAT_1, IAT_2 and IAT_3 are the inter-arrival times for each phase in the model. IAT_1 is the time interval between two aircraft as they are ready to land. It is like the TRACON_IAT and the model ignores the flight length. IAT_1 is assumed to be an exponential distribution with mean 124, where 124 = 2/(1/199 + 1/232 + 1/354 + 1/255), similar to the observed arrival process at the TRACON (Table 1). (In this model, we ignore the movement of aircraft from the TRACON boundary to the final approach and runway threshold. Instead we assume that aircraft arrive directly at a "queue" where they immediately land once they leave the queue.) We model the separation time S as a Gaussian distribution $N(80, 10^2)$ in seconds, where 80 is assumed from an average separation distance of 4 nm and an average speed 180 knots. The 10-second standard deviation is chosen to be close to the measured arrival time error 9.85 seconds at final approach fix by NASA Langley (Oseguera-Lohr and Williams 2003). IAT_2 is the time interval of two aircraft at the runway threshold. The runway occupancy time D is assumed to be $N(48, 8^2)$ (Haynie 2002). IAT_3 is the time interval between aircraft exiting the runway.

Figure 4 shows the histogram of IAT_2 based on simulation of the model. It is similar qualitatively to Figure 2. Also, since the model allows for the possibility that two aircraft are on the runway at the same time, we can estimate the probability of a simultaneous runway occupancy. The model gives a 90% confidence interval for the SRO probability of 0.0033 ± 0.0002.

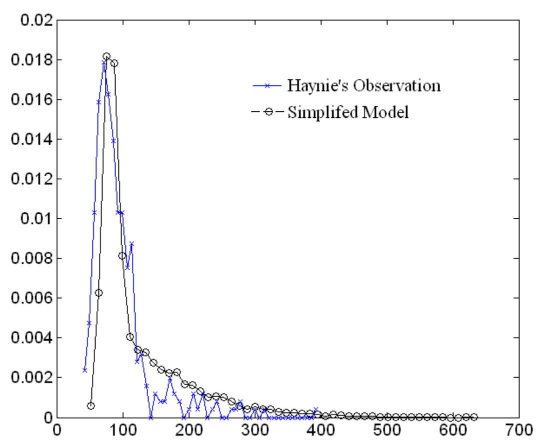

Figure 4: Histogram of Landing Time Interval from the Simplified Model

Although the simplified model leaves many important factors out of consideration, the similarity of the landing time interval histogram with Haynie's observation and the consistency of the SRO probability indicate that the model captures some fundamental characteristics of the terminal approaching and landing phase.

However, the model is insufficient to investigate the system sensitivity to arrival traffic volume or controllers' operating procedures. For example, the simplified model does not account for the variance in applied separation based on arriving aircraft types (heavy, B757, large, small), nor kinematic drift of velocities and positions, nor human operating errors. In order to provide a more accurate estimation of system performance, we propose an agent-based simulation model.

4 AN AGENT-BASED MODEL

A generally acceptable definition of a software agent is a computer program that performs tasks on behalf of another entity without direct supervision or control (Sichman 1998). Here, we assume an agent has a goal, is able to make decisions and perform actions to achieve its goal under various environments, and has the ability to communicate with other agents. In other words, an agent is *autonomous, pro-active, adaptive, and social*. In this paper, we model the local controller as an agent. Its goal is to vector as many aircraft as possible to land without compromising safety. It has the ability to judge and adjust the separation between aircraft, and may reduce separation under high traffic pressure. It is also able to communicate with other controllers and pilots.

4.1 Model Description

Figure 5 shows the available trajectories for arriving aircraft. ATL has four arrival traffic streams from the northeast, northwest, southwest and southeast. Historical data indicate that the northeast stream has a larger volume than the other three streams. In our model, the input inter-arrival times have exponential distributions with means given in Table 1. Once entering the TRACON area, the aircraft

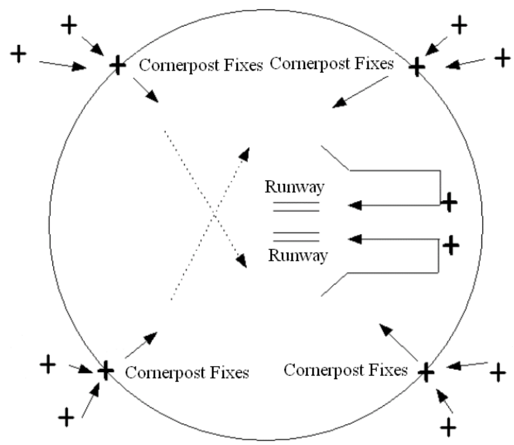

Figure 5: Arrival Aircraft Trajectory

follow a standard approach pattern to land on one of the two runways with the instruction of controllers. For more detailed information on standard approach pattern, please see (Nolan 1999).

Figure 6 shows the basic logic flow of the modeled approach process from the corner-post fix to the start of the final approach.

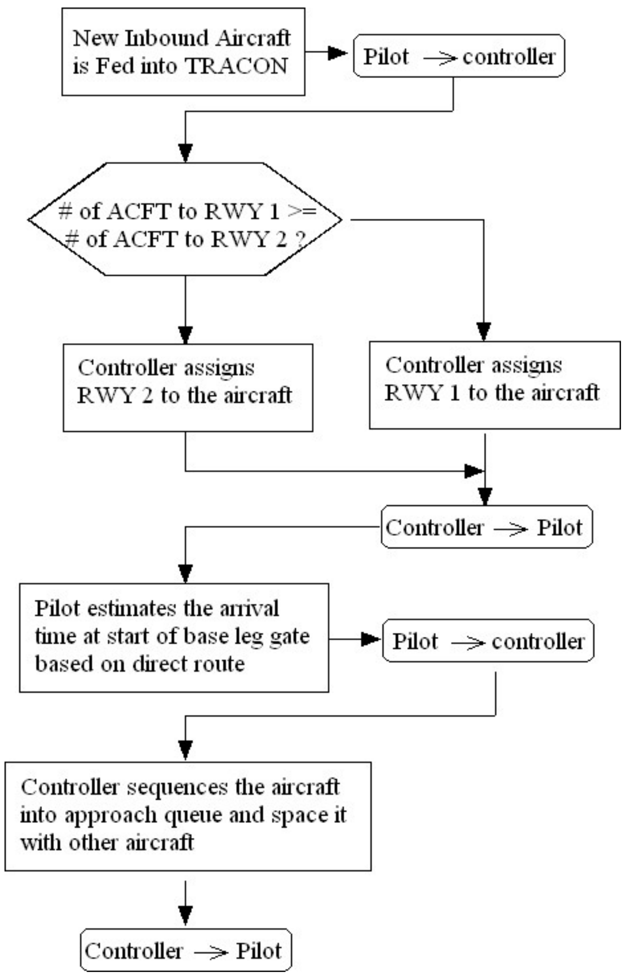

Figure 6: Logic Flow of Aircraft Approach Process

The controllers choose a desired or target separation between two successive aircraft beginning their final approaches. In the model, this separation is drawn from a Gaussian distribution $\tau \sim (60, 8^2)$ in seconds.

The *actual* time when an aircraft begins its final approach is calculated based on the following algorithm. We suppose that an aircraft is assigned to runway A. Let *M* be the flight number of this aircraft, and let [A, *j*] be the flight number of the aircraft which is the *j*th to land on runway A. Let $T_{est}(k)$ be the estimated arrival time at the final approach fix for flight *k*. The actual approach time $T_{act}(k)$ is given by the following algorithm which ensures that the target separation τ is kept between successive aircraft at the final approach fix. In addition, the algorithm sequences aircraft on a first-come-first-served basis, based on estimated time to arrive at the final approach fix.

$$\begin{aligned}
&\text{Sequencing:}\\
&[\mathrm{A}, \mathrm{k}+1] = [\mathrm{A}, \mathrm{k}], \mathrm{k} >= \mathrm{j},\\
&(\mathrm{j} \mid \mathrm{T}_{\mathrm{act}}([\mathrm{A}, \mathrm{j}]) > \mathrm{T}_{\mathrm{est}}(\mathrm{M}) > \mathrm{T}_{\mathrm{act}}([\mathrm{A}, \mathrm{j}-1])),\\
&[\mathrm{A}, \mathrm{j}] = \mathrm{M};\\
&\text{Spacing:}\\
&\mathrm{T}_{\mathrm{act}}([\mathrm{A}, \mathrm{j}]) = \max(\mathrm{T}_{\mathrm{est}}([1, \mathrm{j}]), \mathrm{T}_{\mathrm{act}}([\mathrm{A}, \mathrm{j}-1]) + \tau),\\
&\mathrm{T}_{act}([\mathrm{A}, \kappa]) = \max(\mathrm{T}_{act}([1, \kappa]), \mathrm{T}_{act}([\mathrm{A}, \kappa-1]) + \tau), \kappa > j.
\end{aligned}$$

The delay then is the difference between the actual final approach time and the estimated time.

The model establishes landing separation at the runway threshold as follows: All aircraft are categorized into four classes of aircraft, which are heavy, large, B757 and small. The modeled speed of each type of aircraft at the final approach fix is listed in Table 3. Each speed variable follows a Gaussian distribution. The variance models different kinetic characteristics of different aircraft types in the same category and random deviations from desired speed. Once an aircraft begins its final approach, the controller advises the pilot to maintain a certain separation from the previous aircraft. The separation should be kept along the entire approach.

To avoid the situation where a faster aircraft catches up with a slower aircraft, an extra separation is required for the following faster aircraft at the beginning of its final approach. FAA has an official separation requirement for each aircraft mix, and the separations are chosen to avoid a wake vortex encounter. The *target* separation at the threshold is drawn from a Gaussian distribution ~ N(μ, σ^2), where the mean μ is given in Table 4 based on the leader and trailer aircraft types, and the standard deviation σ is assumed to be 10 seconds. The *actual* separation at the threshold depends on several factors including aircraft speed and time of arrival at the final approach fix.

The process of air traffic controller's sequencing and spacing aircraft is portrayed in Figure 6.

Figure 7 shows a histogram of landing time intervals for a runway, as output from the model. The histogram has

Table 3: Modeled Aircraft Speed Matrix (knots)

Speed(knots)\Category	Heavy		Large		B757		Small	
	Mean	Std.Dev	Mean	Std.Dev	Mean	Std.Dev	Mean	Std.Dev
Final Approach Gate	175	8	155.5	8	169	6	152	4
Runway Threshold	145	6	140	6	140	4	130	4

Table 4: Separation Standard Matrix (Seconds)

	Trailer			
Leader	Heavy	B757	Large	Small
Heavy	99 (4nm)	129 (5nm)	129 (5nm)	166 (6nm)
B757	99 (4nm)	103 (4nm)	103 (4nm)	138 (5nm)
Large	62 (2.5nm)	64 (2.5nm)	64 (2.5nm)	111 (4nm)
Small	62 (2.5nm)	64 (2.5nm)	64 (2.5nm)	69 (2.5nm)

multiple modes, which indicate the distinct separation differences between different aircraft mixes. For instance, the separation of a small aircraft after a heavy one, 166 seconds, is much larger than the separation of a heavy after a small, 62 seconds. The tail results from gaps in the arrival process.

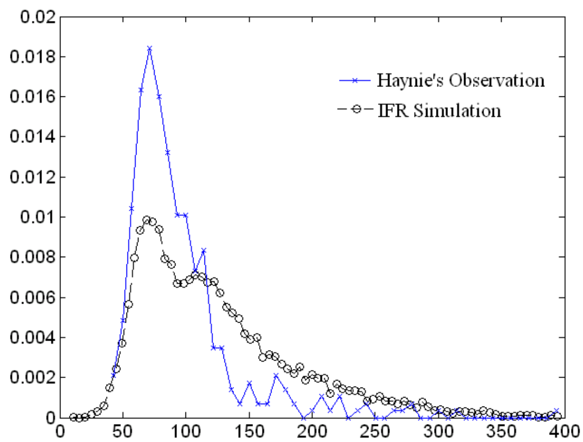

Figure 7: Histogram of Landing Time Interval when Using Separation Standard Matrix

However, the simulation result is not consistent with field observation, which was conducted by Haynie at ATL in 2002. The histogram of the landing time interval is plotted in Figure 2. In Haynie's histogram, there is only one distinct mode around 80 seconds. Specifically, the mean of separation for near separated aircraft mixes is around *12 seconds more* than the separation standard, while the mean for far separated aircraft mixes is around *17 less* than the separation standard (Xie, Shortle, and Donohue 2003). Therefore, in real VFR operation, controllers may allow pilots to blur the separation difference between different aircraft mixes, and it will reduce workload especially at busy traffic situation. Acknowledging such a phenomenon, we adjusted the working logic model of the local controller agent, and its separation allocation mechanism is listed in Table 5.

From the histogram of the simulation result, Figure 8, one can see that the model has been calibrated to be very similar to the real situation. However, it must be kept on mind that the system is hybrid and stochastic, and both the observation and a simulation are only a realization of the stochastic system.

Table 5: Hypothetical Reduced Separation (Seconds)

	Trailer			
Leader	Heavy	B757	Large	Small
Heavy	82 (3.3nm)	85 (3.3nm)	85 (3.3nm)	91 (3.3nm)
B757	82 (3.3nm)	85 (3.3nm)	85 (3.3nm)	91 (3.3nm)
Large	67 (2.7nm)	69 (2.7nm)	69 (3.3nm)	91 (3.3nm)
Small	67 (2.7nm)	69 (2.7nm)	69 (2.7nm)	74 (2.7nm)

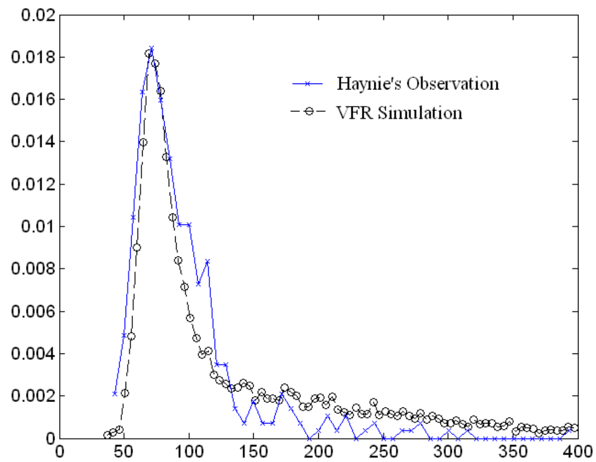

Figure 8: Histogram of Landing Time Intervals from Simulation

Since current model configuration has produced reasonably consistent result with the real observation, we will take the current model as the baseline situation. We are interested in several system performance metrics that reflect the capacity and safety of the system, shown in Table 6.

Table 6: System Performance Metrics

Capacity	Airborne delay, Runway landing rate, maximum landing rate
Safety	Simultaneous runway occupancy probability per landing

An aircraft has to wait in holding pattern if the landing glide slope is occupied and no more aircraft can enter the queue of final approach, thereby the induced waiting time is airborne delay. Runway landing rate is the number of aircraft landed on a runway per hour. When an aircraft touches down on a runway before the pre-landing aircraft exits the same runway, it is a case of simultaneous runway occupancy.

In the baseline case, 20 trials were run, and in each trial, 20,000 aircraft were landed. The average landing rate is 27 arrivals per hour, and maximum rate is 44.

The average airborne delay is 18 seconds. The delay is so small and it means that the system has enough capacity to accommodate the current traffic demand.

The probability of simultaneous runway occupancy is 0.0031 per landing. It is worthy to mention that the estimate of SRO probability here is consistent with that from the Haynie's observation.

4.2 Experimental Analysis

To evaluate the system performance under different air traffic volumes, we carried out three experiments for sensitivity analysis on safety and capacity with respect to the change of aircraft arrival rates.

The current traffic volume is defined as the baseline, and the volume change is set in terms of a multiple of the baseline. The scenarios considered include *lighter-than-baseline* situations, where the volumes are 0.1, 0.25, 0.5, and 0.75 times the baseline, and *heavier-than-baseline* situations, where the volumes are 1.25, 1.35, 1.45, 1.55, 1.75, 1.85, and 2 times the baseline. We scale all four-arrival streams uniformly.

4.2.1 Experiment 1

With the first experiment, air traffic controllers do not respond to the traffic volume change. They perform a constant maneuver strategy regardless of traffic pressure. The desired separation matrix is the one used in Table 5.

Figure 9 shows the results of this simulation. The runway landing rate exhibits a linear relationship between landing rates and traffic volume growth ratio until the growth ratio reaches 1.75. When traffic becomes more than 1.75 of the current volume, the system is almost operated at its maximum capacity, and the landing rate rests on a plateau at 47 landings per hour on average. The maximum landing rate, which achieves 54 landings per hour, is determined by the separation strategy deployed by controllers. Since the service rate cannot be increased any more, while the arrival rate is higher than the service rate, the system losses the stationarity. In this case, when traffic doubles, each runway has arrival rate of 58 per hour on average, while a runway only can land 54 aircraft per hour at the most. Therefore, although the average delay among the simulated 20,000 landings seems to be endurable, the accumulated delay might be huge since the system is not steady anymore. The safety measure (simultaneous runway occupancy) changes almost linearly with the arrival rate with the consistent separation strategy. The correlation coefficient of them is 0.9876. As traffic volume grows, more and more aircraft have to be separated close to the minimum acceptable distance, and more cases of simultaneous runway occupancy are likely to happen.

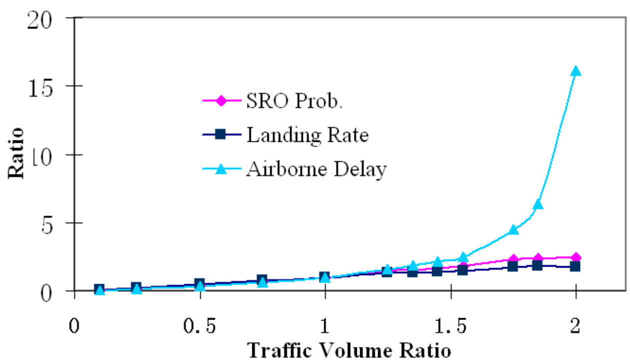

Figure 9: System Performance vs. Traffic Growth

The distributions of landing time intervals are different in light, medium, and high traffic volume situations. When the traffic is very light, the distribution has a long tail due to long periods of time when no airplanes arrive (Figure 10). When the traffic is heavy, the distribution displays more symmetry, see Figure 11, and variance mainly comes from system disturbance, such as airplane kinetic deviation, human performance limitation, weather noise, and so forth.

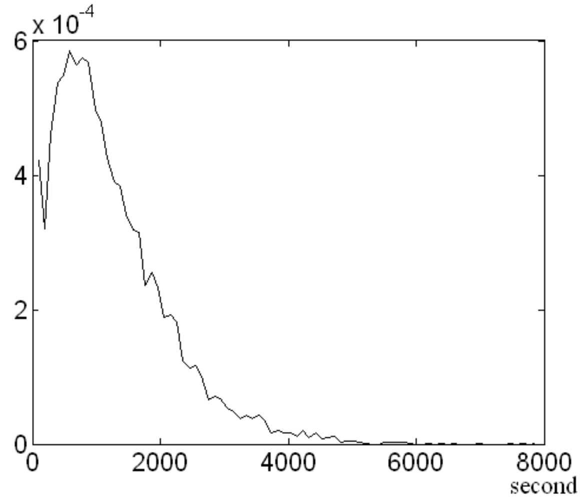

Figure 10: Histogram of Landing Time Interval when Traffic is Very Light

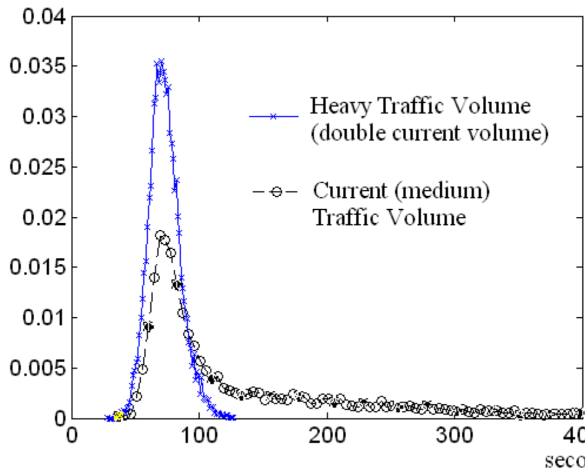

Figure 11: Histogram of Landing Time Interval when Traffic is Medium and Heavy

4.2.2 Experiment 2

The second experiment allows air traffic controllers to be more adaptive to the traffic pressure. Figure 12 shows a hypothetical continuous adaptive model, in which the VMC separation allowed by controllers changes as a function of the traffic volume. That is, when traffic is light, controllers employ the separation standards given in Table 5. When traffic is heavy, the controllers decrease the standards in an effort to accommodate the traffic. The formulas are

$$\text{Far-Separated:} \quad y = 0.6 * \frac{e^{-5*(x-1.5)} - e^{5*(x-1.5)}}{e^{-5*(x-1.5)} + e^{5*(x-1.5)}} + 2.7$$

$$\text{Near-Separated:} \quad y = 0.6 * \frac{e^{-5*(x-1.5)} - e^{5*(x-1.5)}}{e^{-5*(x-1.5)} + e^{5*(x-1.5)}} + 2.1\,.$$

The near-separated mix includes aircraft pairs whose separation standard is less than 4 nautical miles. Others are far-separated mixes. The separation variance is converted from the variance in time domain using expected average speed (Oseguera-Lohr and Williams 2003).

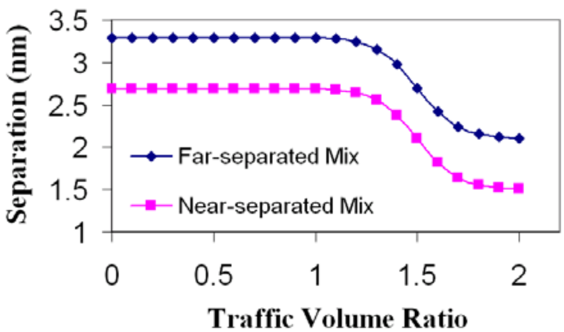

Figure 12: A Hypothetical Controller Adaptive Model to Traffic

Figure 13 shows the output of this experiment. The airborne delay is not as serious as in the previous example (compare Figures 9 and 13 for large traffic volumes). However, the safety risk (probability of an SRO) increases much more dramatically. In this experiment, doubling the traffic volume results in a 6-fold increase in P(SRO), with only a 2-fold increase in capacity. In experiment 1, where controllers maintain a consistent separation distance, SRO probability only doubles when the arrival rate doubles.

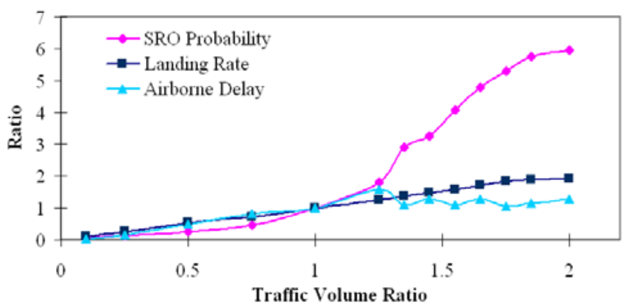

Figure 13: System Performance vs. Traffic Growth

Figure 14 shows a histogram of the landing separation times for the two experiments with a traffic volume at 1.75 times the current level. The histogram of the scenario with the adaptive separation shifts to the left of the histogram of that with consistent separation. The long right tails of the two curves dilute the differences on the left tails and make the global averages very close. However, it is the left tail of the curve that influences the runway safety significantly. Therefore, when we are estimating runway safety, the information about landing rate is not enough, and the distribution of separation time is very important.

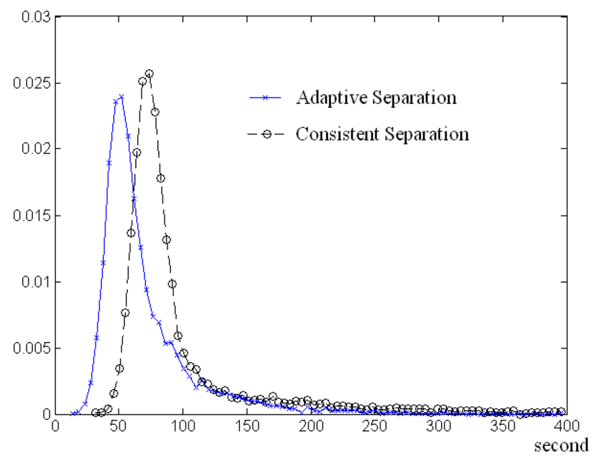

Figure 14: Histograms of Landing Time Interval from Experiment 1, 2

4.2.3 Experiment 3

The third experiment is to demonstrate the advantages of a more synchronized schedule, where traffic is metered uniformly in the TRACON area. In particular, we use a Gaussian distribution with a small variance (instead of an exponential distribution) to model the inter-arrival times at each of the four arrival streams. We keep the mean arrival times the same. Although the four streams do not synchronize with each other, the arrival schedules will not be as unsynchronized as those will so far. With the controller's strategy used in the Experiment 1, we will compare the safety, capacity and delay using synchronized arrival times with those using exponential arrival times. Without losing generality, we only choose a high traffic volume situation, which doubles the current volume, to analyze the influence. The results are shown in Table 7. The means and standard deviations are calculated from the 20 trials of the simulations.

Synchronized schedules are more efficient than the current random schedule because of the higher landing rate and less delay. Although average simultaneous runway occupancy probability is a little higher than random schedules, the standard deviation is much lower. Higher SRO probability comes from the fact that fewer aircraft must fly a further distance.

Table 7: System Performance in Experiment 3

Double Traffic Volume	Random Schedules		Synchronized Schedules	
	Mean	Std.Dev	Mean	Std.Dev
SRO Prob.	0.0075	0.0009	0.0081	0.0001
Landing Rate /hour	47	5	50	0.1
Delay	289	125	89	1

The runway efficiency is also improved, which can be shown in the histogram of runway landing rates. See Figure 15.

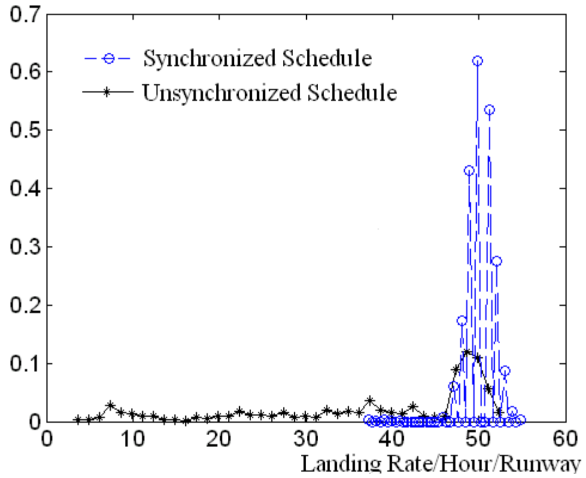

Figure 15: Histogram of Landing Rate when Traffic is Very Heavy

A synchronized schedule can lead to a landing rate with smaller variance. That is the reason that 3 more aircraft are able to land per hour on average.

It is worthy to point out that the simultaneous runway occupancy probability *0.0081* is possibly very close to the true risk level of current operations. It is higher than the field-measured SRO risk *0.0031*, and it warns us that the current operations are not as safe as what we are feeling right now. The reasons for the higher risk involve two points. The first is that current traffic volume is not large enough; another is that random schedule leads to a longer right tail of landing time interval distribution.

5 SUMMARY

This paper presents an agent-based simulation to model airport terminal approach and landing operations. The paper first gives a simplified queuing model to demonstrate the essence of terminal approach and landing process. To give a more detailed and accurate analysis of the system performance, we build an agent-based model. We then calibrate the model to the current operational situation at ATL airport by inputting real arrival statistics and outputting landing time intervals and estimates of simultaneous runway occupancy probability that are consistent with field observation data.

After that, we conduct a sensitivity analysis to examine the performance of the terminal approaching and landing system under different traffic loads. If controllers maintain the current separation standards that they are deploying, delay increases non-linearly with the traffic load. This is consistent with the theoretical delay of a queue that is near its steady-state operating capacity.

A reasonable reaction that controllers might take under high traffic pressure is to relax the separation distance between aircraft to gain more capacity. With a simplified adaptive model, the simulation illustrates that by relaxing separation distances, it is possible to eliminate the non-linear growth in delay, but this comes at the cost of a non-linear increase in simultaneous runway occupancies and substantially reduced safety.

A potential way to increase capacity without compromising safety is to synchronize the flight schedules. Even the strategy that only independently synchronizes each arrival stream can demonstrate the advantage. Our experiments illustrated a substantial reduction in delay with little change in safety for a system with synchronized arrival streams.

ACKNOWLEDGMENTS

This work has been partially supported by the WakeVAS program at NASA Langley Research Center, under Wayne Bryant, and the National Institute for Aerospace, contract NAS1-02117.

REFERENCES

Boesel, J. R. 2003. An Air Traffic Simulation Model That Predicts And Prevents Excess Demand. In Proceedings of the AAIA Modeling, Simulation and Technology Conference, Austin, Texas.

Department of Transportation. 2004. Available Aviation Delay. 2004. Available online via <www.dot.gov/PerfPlan 2004/mobility_delay.html> [accessed August 18, 2004].

Haynie, C. R. 2002. An Investigation of Capacity and Safety in Near-Term Airspace for Guiding Information Technology Adoption, doctoral dissertation, George Mason University, Fairfax, Virginia.

Lee, D. A., C. Nelson, and G. Shapiro. 1998. The Aviation System Analysis Capability Airport Capacity and Delay Models. Logistics Management Institute, McLean, Virginia, NASA/CR-1998-207659.

Nolan, M.S. 1999. Fundamentals of Air Traffic Control, Pacific Grove, California.

Oseguera-Lohr, R. M., and D. H. Williams. 2003. Evaluation of Trajectory Errors in an Automated Terminal-Area Environment. NASA Langley Research Center. 2003.

Sichman, J. S. 1998. Multi-Agent Systems and Agent-Based Simulation. Springer.

Xie, Y., J. Shortle, and G. Donohue. 2003. Runway Landing Safety Analysis: A Case Study of Atlanta Hartsfield Airport. In Proceedings of the 22nd Digital Avionics Systems Conference, Indianapolis, Indiana.

AUTHOR BIOGRAPHIES

YUE XIE is currently a Ph.D candidate in Department of Systems Engineering and Operations Research, George Mason University. His research interests include modeling and simulation of aviation system and the application on aviation safety analysis. He received his M.S. in 2001 and B.S. in 1998 both from HuaZhong University of Science and Technology, China. His e-mail address is <yxie@gmu.edu>.

JOHN SHORTLE is currently an assistant professor of systems engineering at George Mason University. His research interests include simulation and queueing applications in telecommunications and air transportation. Previously, he worked at US WEST Advanced Technologies. He received a B.S. in mathematics from Harvey Mudd College in 1992 and a Ph.D. and M.S. in operations research at UC Berkeley in 1996. His e-mail address is <jshortle@gmu.edu>.

GEORGE DONOHUE, formerly Associate Administrator of Research and Acquisition in the Federal Aviation Administration, who has broad experience in managing major research and technology projects in both the public and private sector, was named the FAA visiting Professor for Air Transportation Technology and Policy in July, 1998. He assumed his current position as Professor of Systems Engineering and Operations Research in February of 2000. His e-mail address is <gdonohue@gmu.edu>.