Accuracy of Reinforcement Learning Algorithms for Predicting Aircraft Taxi-out Times

(A Case-study of Tampa Bay Departures)

Poornima Balakrishna, Rajesh Ganesan, Lance Sherry
Center for Air Transportation Systems Research
Department of Systems Engineering and Operations Research
George Mason University
Fairfax, U.S.A
pbalakri@gmu.edu, rganesan@gmu.edu, lsherry@gmu.edu

Abstract—**Taxi-out delay is a significant portion of the block time of a flight. Uncertainty in taxi-out times reduces predictability of arrival times at the destination. This in turn results in inefficient use of airline resources such as aircraft, crew, and ground personnel. Taxi-out time prediction is also a first step in enabling schedule modifications that would help mitigate congestion and reduce emissions. The dynamically changing operation at the airport makes it difficult to accurately predict taxi-out time. In this paper we investigate the accuracy of taxi out time prediction using a nonparametric reinforcement learning (RL) based method, set in the probabilistic framework of stochastic dynamic programming. A case-study of Tampa International Airport (TPA) shows that on an average, with 93.7% probability, on any given day, our predicted mean taxi-out time for any given quarter, matches the actual mean taxi-out time for the same quarter with a standard error of 1.5 minutes. Also, for individual flights, the taxi-out time of 81% of them were predicted accurately within a standard error of 2 minutes. The predictions were done 15 minutes before gate departure. OOOI data available in the ASPM database maintained by the FAA was used to model and analyze the problem. The prediction accuracy is high even without the use of detailed track data.**

Keywords-taxi-out delay; prediction; reinforcement learning.

I. INTRODUCTION

Flight delays have a significant impact on the nation's economy. The United States National Airspace System (NAS) is a complex system consisting of several components including the administration, control centers, airlines, aircraft, and passengers. Flight delays propagate over the NAS and increases with time over the length of the day due to the cascading effect. Stakeholders, particularly the ground and tower controllers, are overwhelmed during peak hours when the number of departures and arrivals increase; at times beyond capacity. Taxi-out delay is a major component of flight delays. Predictability of taxi-out time would help ease congestion and mitigate delays via better gate departure planning. Taxi-out time of a flight is defined as the time between gate pushback and time of takeoff. Increased predictability of taxi-time at the departure airport will also increase planning efficiency at arrival airports.

Delays are caused by several factors. Some of these include increased demand, weather, near-capacity operation of major hub airports, and air traffic management programs such as Ground Delay Programs (GDPs) and Ground Stops (GS). The delay phenomenon is continuously evolving and is both stochastic and elastic in nature. The stochastic nature is due to the uncertainties that lie at the local level (such as the local control tower, arrival/departures movements on ground, and human causes), system level (such as GDP), and in the environment (weather). The elastic behavior is due to the fact that delay could be adjusted (positively or negatively) by flying speed, taking alternate routes, turnaround time on the ground, and position in the departure clearance queue especially during busy hours of the airport. In order to minimize the taxi-out delay component of the total delay, it is necessary to accurately predict taxi-out under dynamic airport conditions. This information in turn will allow the airlines to better schedule and dynamically adjust departures, which minimizes congestions, and the control towers will benefit from smoother airport operations by avoiding situations when demand (departure rates) nears or exceeds airport capacity.

II. LITERATURE REVIEW

ISBN: 978-0-615-20720-9

Several research attempts have been documented to understand the departure process at airports. These include both simulation models and analytical formulations. The Departure Enhanced Planning And Runway/Taxiway Assignment System (DEPARTS) [1] developed at MITRE Corporation attempts to reduce taxi times by generating optimal runway assignments, departure sequencing and departure fix loading. Results of their analysis also indicate that pushback predictability could influence all phases of flight and traffic flow management.

A simulation based study of queueing dynamics and "traffic rules" is reported in [2]. They conclude that flow-rate restrictions significantly impact departure traffic. The impact of downstream restrictions is measured by considering aggregate metrics such as airport throughput, departure congestion, and average taxi-out delay. Other research that has focused on departure processes and departure runway balancing are available in [3,4]. Many statistical models that consider the probability distribution of departure delays and aircraft takeoff time in order to predict taxi-time have evolved in recent years [5,6].

In [7] a queueing model for taxi-time prediction is developed. They identify takeoff queue size to be an important factor affecting taxi-out time. An estimate of the takeoff queue size experienced by an aircraft is obtained by predicting the amount of passing that it may experience on the airport surface during its taxi-out, and by considering the number of takeoffs between its pushback time and its takeoff time. However, this requires prior knowledge of actual takeoff times of flights and hence may be unsuitable for planning purposes. The model is valid for a specific runway configuration since the runway configuration at the future time of taxi-time prediction is unknown. Suggested extensions to the model include a runway configuration predictor. A queuing model based on simulation to test different emissions scenarios related to duration of taxi-out was developed in [8]. Some of the scenarios that are considered are redistribution of flights evenly across the day, and variation in number of departures under current capacity. The study showed that lower taxi-out times (and thus lower emissions) are experienced by airlines that use less congested airports and don't rely on hub-and-spoke systems. Other research that develops a departure planning tool for departure time prediction is available in [9-13].

Direct predictions attempting to minimize taxi-out delays using accurate surface surveillance data have been presented to literature [14,15]. Recent work using surface surveillance data presented in [16] develops a bivariate quadratic polynomial regression equation to predict taxi time. In this work data from Aircraft Situation Display to Industry (ASDI) and that provided for Northwest Airlines for DTW (Flight Event Data Store, FEDS) were compared with surface surveillance data to extract gate OUT, wheels OFF, wheels ON, and gate IN (OOOI) data for prediction purposes.

A Bayesian networks approach to predict different segments of flight delay including taxi-out delay has been presented in [17]. An algorithm to reduce departure time estimation error (up to 15%) is available in [18], which calculates the ground time error and adds it to the estimated ground time at a given departure time. A genetic algorithm based approach to estimating flight departure delay is presented in [19].

It is useful to keep in mind that a lot of the data is proprietary and the different attempts in the literature use different data sources depending on accessibility.

III. RL METHODOLOGY

In this research a machine learning approach is used for the task of taxi-out time $a \in A$ prediction, where *A* denotes the action space. The evolution of system state $x \in X$ is modeled as a Markov chain, where *X* denotes the system state space. The decision to predict the taxi-out time based on the system state is modeled as a Markov decision process (MDP). For the purpose of solving the MDP, it is necessary to discretize *X* and *A*. Due to the large number of state and action combinations *(x,a)*, the Markov decision model is solved using a machine learning (reinforcement learning (RL), in particular) approach.

The purpose of the RL estimator is to predict taxi-out time given the dynamic system state. The input to RL is the system state and the output of the learning process is a reward function *R(x,a)* where $a \in A$ is the predicted taxi-out values. The utility function (reward) *R(x,a)* is updated based on the difference between the actual and predicted taxi-out values *r(x,a,j)*. We define reward *r(x,a,j)* for taking action *a* in state *x* at any time *t* that results in a transition to state j, as the absolute value of error *r(x,a,j)* = |*Actual Taxi-out – predicted Taxi-out*| resulting from the action. The transition probability in a MDP can be represented as *p(x,a,j)*, for transition from state *x* to state *j* under action *a*. Then the prediction system can be stated as follows. For any given $x \in X$ at time *t* there is a prediction *a* such that the expected value of error *(Actual – predicted Taxi-out)* is zero. Theoretically, the action space for the predicted taxi-out could have a wide range of numbers. However, in practice, for a non-diverging process, the action space is quite small, which can be discretized to a finite number of actions.

For practical implementation since transition probabilities *p(x,a,j)* are not known, we use the reinforcement learning

 ISBN: 978-0-615-20720-9

version of the Bellman's optimality equation [21] to update *R(x,a)* as follows

$$R^{t+1}(x,a) = (1-\alpha)R^t(x,a) + \alpha[r(x,a,j) + \beta \min_{b \in A} R^t(j,b)] \quad x, j \in X \; a \in A$$

where α is a learning parameter that is decayed over time, and β is the discount parameter.

The state variables $x=\{x_1,x_2,x_3,x_4,x_5\}$ for the taxi-time prediction problem were determined by analyzing the available data. Analysis of the data suggests that for a specific aircraft that is scheduled to pushback, the number in queue at the runway (x_1), the number of departure aircraft co-taxiing (x_2), and the number of arrival aircraft co-taxiing (x_3) are the major factors that influence taxi-out time. In addition, taxi-out time changes gradually over the day. The taxi-out time during a given quarter was found to depend on the taxi-out times of the previous two quarters. So the average taxi-out time of the previous two quarters was considered as a factor influencing taxi-out time (x_4). Along these lines, the time of day (x_5) was also included as a factor. Thus, there are 5 variables that comprised the state vector.

Several measures of performance such as discounted reward, average reward, and total reward can be used to solve a MDP. At the beginning of the learning process, the R-values are initialized to zeros. When the process enters a state for the first time, the action is chosen randomly since the R-values for all actions are zero initially. In order to allow for effective learning in the early learning stages, instead of the greedy action (action with lowest R-value) the decision maker, with probability P_t, chooses from other actions. The choice among the other actions is made by generating a random number from a uniform distribution. The above procedure is commonly referred to in RL literature as exploration. The RL based functional block diagram is shown in Fig. 1. Theoretical details of the RL algorithm can be obtained from [20-24].

A. Obtaining Predicted Taxi-Out Time

Once learning is completed, the R-values (reward) provide the optimal action choice for each state. At any time *t* as the process enters a state, the action *a* corresponding to the lowest non-zero R-value indicates the predicted taxi-out time *a*. In what follows we present the steps of the RL algorithm in the implementation phase. The RL estimator was coded in MATLAB®.

B. Steps in RL

- ***Step 1:*** Once the states, actions, and the reward scheme are set up, the next step is to simulate the *t+45* look-ahead window. Assume 15 minute decision (prediction) epochs *i.e.* prediction was done for flights in a moving window of length *t* to *t+15* minutes. This means that for each departing flight in the 15 minute interval from current time, the airport dynamics was simulated for 30 minutes from its scheduled departure time.
- ***Step 2:*** Simulate the first 15 minute window. For each flight in the window obtain the system state *x*. To calculate average taxi-out times before current time *t*, actual flight data between *t* and *t-30* are used. Initialize *R(x,a)* to zeros.
- ***Step 3:*** If exploration has decayed go to step 4, else choose arbitrary actions (predictions from set *A*). The window is then moved in 1 minute increments and all flights in the window are predicted again. This means that every flight, unless it leaves before scheduled time, has its taxi-out time predicted at least 15 times. Simulate the new window of 15 minutes. Find the next state *j* for each flight. Compute *r(x,a,j)*. Update reward *R(x,a)* using the fundamental Robbin-Monro's stochastic approximation scheme [25] that is used to solve Bellman's optimality equation [21] provided earlier.
- ***Step 4:*** If learning phase is in progress, choose greedy action *a* from set *A* (action corresponding to the lowest R-value). The window is then moved in 1 minute increment and all flights in the window are predicted again. Simulate the new window of 15 minutes. Find the next state *j*. Compute *r(x,a,j)*. Update R(x,a).
- ***Step 5:*** Continue learning by simulating every 15 minute interval, until 45 minutes have been completed. Next, move the window of width 60 minutes by a fixed time increment (say 15 minutes) and repeat learning by going to Step 2.
- ***Step 6:*** Continue learning with several months of ASPM data until a stopping, or a near-optimal criterion is reached such as $| R^{t+1}(x,a) - R^t(x,a) | \le \varepsilon$ where ε is a very small number.
- ***Step 7:*** Once learning is complete, the optimal prediction for a given state is the one that corresponds to the minimum R-value for that state.

Figure 1. Reinforcement Learning Based Functional Block Diagram for Taxi-Out Time Prediction

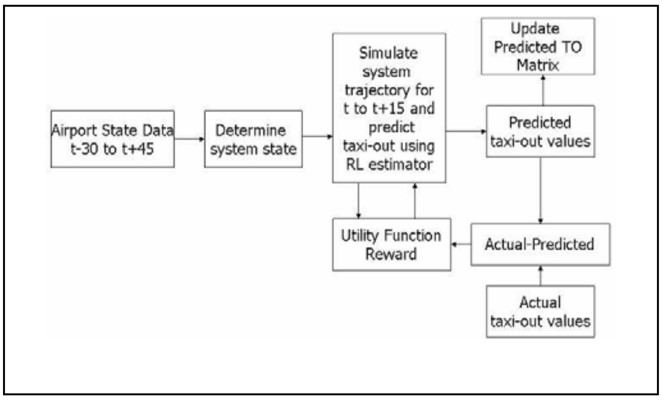

IV. RESULTS

ISBN: 978-0-615-20720-9

A. Data Source

OOOI (Out,Off,On,In) data was obtained for Tampa International Airport (TPA) from the ASPM (Aviation System Performance Metric) database maintained by the FAA (Federal Aviation Administration). Data from June 1st 2007 up to August 25th 2007 was used to train the RL based taxi-time estimator, and data for August 26th to August 31st was used for testing the accuracy of prediction. OOOI data provides the following information for each recorded flight – Scheduled pushback time from the gate, Actual pushback time from the gate, Actual Wheels Off time, Actual Wheels On time at the arrival airport, and Actual In time which is recorded when the aircraft reaches the gate after the taxi-in process. In addition, the ASPM database also provides an airline (not individual flight) specific seasonal average for the nominal or unimpeded taxi-out time and taxi-in time. In this research we assume that if an aircraft completes the corresponding nominal taxi-out time, it joins a runway queue.

It is possible that an aircraft pushes back from the gate and for varying reasons may have to return to the gate and pushback again. It is unclear as to whether the actual pushback time reported by the airlines indicates the first pushback or the second pushback. The Bureau of Transportation Statistics (BTS) recently issued a directive [26] requiring all airlines reporting data to ensure that the first pushback time be recorded as the actual gate-out time. This clearly influences the measured taxi-out time.

B. Observations

We adopt two methods to analyze the results. First we compute the taxi-out time prediction accuracy for individual flights on a given day. Second, we evaluate the prediction accuracy of average taxi-out times in 15 minute intervals of the day. Table 1 below summarizes the prediction accuracy for individual flights for six days in August 2007.

TABLE 1. A comparison of prediction accuracy for individual flights across days of August 2007 for TPA.

Day (August 2007)	**26th**	**27th**	**28th**	**29th**	**30th**	**31st**
Mean Actual Taxi-Out time	11.53	11.73	11.17	11.21	11.06	16.06
Mean Predicted Taxi Time (min)	10.82	11.20	10.09	10.38	10.74	11.96
Std. Dev. Actual Taxi Time (min)	3.80	4.76	3.30	4.02	6.29	8.93
Std. Dev. Predicted Taxi Time (min)	1.86	4.23	1.29	1.76	1.93	3.26
Median Actual Taxi Time (min)	10.8	10.2	10.2	10.2	10.2	12
Median Predicted Taxi Time (min)	10	10	10	10	10	12
% of Flights with RMSE of 2 min	87.5	81.62	83.15	79.18	84.48	69.83

The results in Table 1 indicate that the mean of predictions for a given day are comparable to the mean of actual taxi-out times. We note that the standard deviations of predicted taxi-out time values are not very closely matched with the standard deviations of actual taxi-out times. A possible reason for this is that we consider a flight to enter the runway queue if it has not taken off by the end of its nominal or unimpeded taxi-out time. The nominal taxi-out time data available for this research is however a seasonal average specified for each airline, and not for each individual flight. This is also our only measure in some sense of gate to runway distance. This will undoubtedly introduce further uncertainty in our predictions since factors such as runway configuration are not captured in this average. Also, predictions of taxi-out times are made 15 minutes prior to scheduled pushback of flights. In what follows, we analyze an alternative method to compare the prediction results by considering the following four cases.

Case 1: Consider all flights *scheduled* to pushback in a specific quarter (15 minute interval) of the day. Plot their corresponding mean predicted and mean actual taxi-out time with respect to the same quarter. Note that all flights that are scheduled to pushback in a certain quarter may not take off together (around the same time). This is because taxi-out time is defined as the time elapsed between pushback from the gate and take-off time; and thus depends on several other factors influencing individual flights such as distance of gate from runway, enforcement of downstream restrictions such as Ground Delay Programs (GDPs), Ground Stops (GS) and Miles-In-Trail (MIT).

Case 2: Consider all flights that *actually* took off in a specific quarter of the day. Plot their corresponding mean predicted and mean actual taxi-out time with respect to the same quarter. In this case the flights being considered would have pushed back from the gate at different times spread over different quarters.

Case 3: Consider all flights that were *predicted* (by the algorithm) to take off in a certain quarter (predicted off time can be computed by adding the predicted taxi-out time to scheduled gate-out time). Plot their corresponding mean predicted and mean actual taxi-out times with respect to that quarter.

Case 4: Consider all flights that were predicted to take off in a certain quarter and plot their corresponding mean *predicted* taxi-out times. Now, consider all flights that actually took off in that same quarter. Plot their corresponding mean *actual* taxi-out times. Here we note that the flights that *actually* took off in the quarter being analyzed may not exactly match the set of flights that were *predicted* to take off in that same quarter – this is an inherent limitation of the data available in the ASPM database. Information regarding downstream

 ISBN: 978-0-615-20720-9

restrictions affecting individual flights is not available. Hence we cannot account for passing of aircrafts on the taxiway. Plots of the four cases are provided in Fig. 2-5 for 26th August 2007.

It is easy to see that cases 1-3 represents the average accuracy of prediction for a specified group of flights per quarter, while case 4 discusses average accuracy of prediction for a specified time interval of day which indicates behavior of the airport. In this study, Case 4 is extremely useful in predicting average airport taxi-out time trends approximately 30 minutes in advance of the given time of day (specifying the take off quarter).

For each of the days and for each of the four cases, the accuracy of prediction was measured as the percentage of the time for which the mean predicted taxi-out time per quarter matched the mean actual observed taxi-out time for the same quarter with a standard error of 1.5 minutes. The results are tabulated in Table 2.

TABLE 2. A comparison of prediction accuracy of averages across days of August 2007

Day	**Case 1**	**Case 2**	**Case 3**	**Case 4**
(August 2007)	*% accuracy Standard Error of 1.5 min*	*% accuracy Standard Error of 1.5 min*	*% accuracy Standard Error of 1.5 min*	*% accuracy Standard Error of 1.5 min*
26th	97.1014	97.1014	100.0000	95.6522
27th	88.4058	84.0580	88.4058	95.6522
28th	88.4058	92.7536	78.2609	92.7536
29th	89.8551	92.7536	92.7536	92.7536
30th	98.5507	98.5507	98.5507	95.6522
31st	89.8551	86.9565	94.2029	89.8551

V. CONCLUSIONS

The analysis using the artificial intelligence methodology (reinforcement learning) indicates that on an average, for a given day, the accuracy of mean predicted taxi-out time per quarter in comparison to the actual taxi-out time for the same quarter is approximately 93.7% (case 4) with a standard error of 1.5 minutes. It is to be noted that the prediction was done 15 minutes before scheduled departure for individual flights which were then averaged in quarter time intervals. Prediction accuracy for individual flights was also tested, and on average, 81% of flights were predicted within a root mean square error value of 2 minutes.

It is expected that control tower operations, surface management systems, and airline scheduling can benefit from this prediction by adjusting schedules to minimize congestion, delays, and emissions, and also by better utilization of ground personnel and resources. Especially, with airport dynamics changing throughout the day in the face of uncertainties such as weather, prediction of airport taxi-out time averages combined with individual flight predictions, could help airlines manage decisions such as incurring delays at the gate as opposed to increasing emissions due to longer taxi times. Air Traffic Control would also benefit from this knowledge when making decisions regarding holding flights at the gate or ramp area due to increased congestion. This could improve the performance of air traffic flow management both on ground and in air across the entire NAS in the US and worldwide. It can be integrated to support the futuristic Total Airport Management concepts beyond Collaborative Decision Making [27] that envisions automation of several airport operations.

As part of future work, accuracy of predications will be improved by incorporating runway direction. This is because runway configurations change during the day which could alter the gate to runway distance. Also a sensitivity analysis of the learning parameters will be conducted. Further analysis to capture seasonal trends and incorporation of runway and gate assignments could improve prediction accuracy. Also study of other majors hubs are part of this ongoing research.

ACKNOWLEDGMENTS

The authors would like to thank FAA for their permission to obtain data from the ASPM database.

ISBN: 978-0-615-20720-9

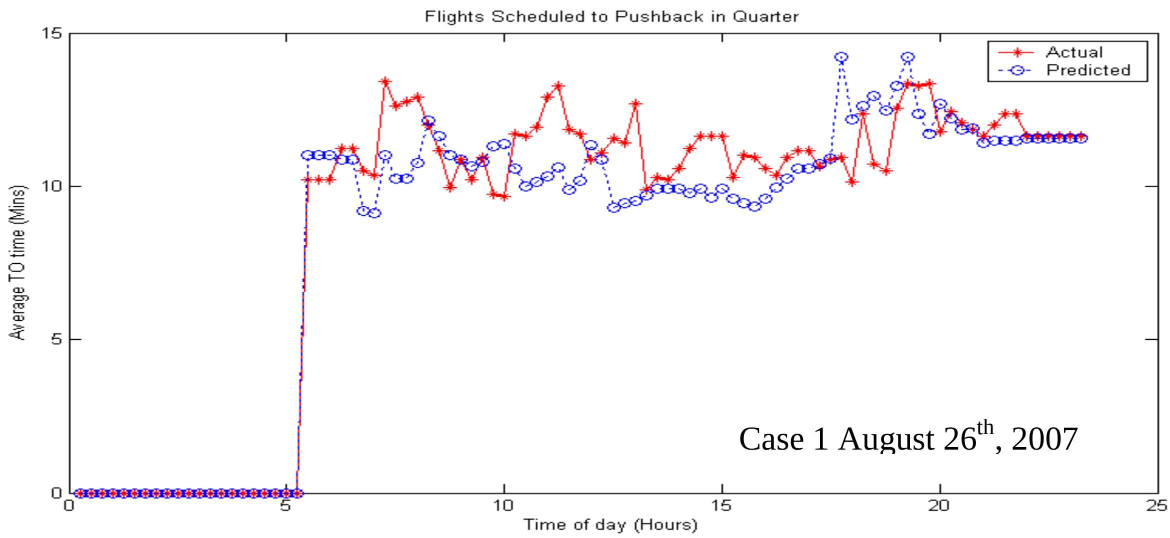

Figure 2. Case 1: Plot of Actual Taxi-Out Time vs. Predicted Taxi-Out Time

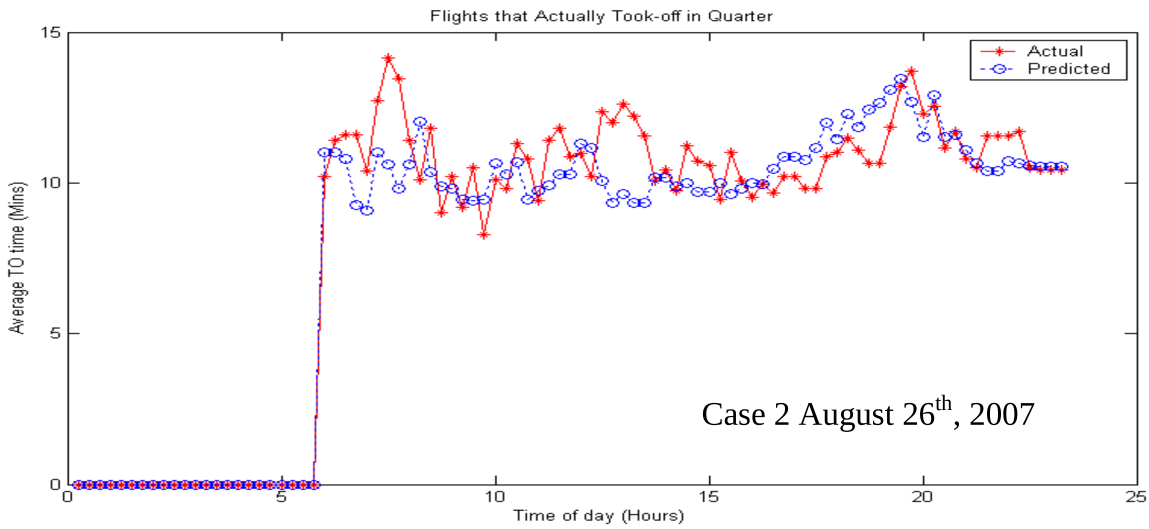

Figure 3. Case 2: Plot of Actual Taxi-Out Time vs. Predicted Taxi-Out Time

ISBN: 978-0-615-20720-9

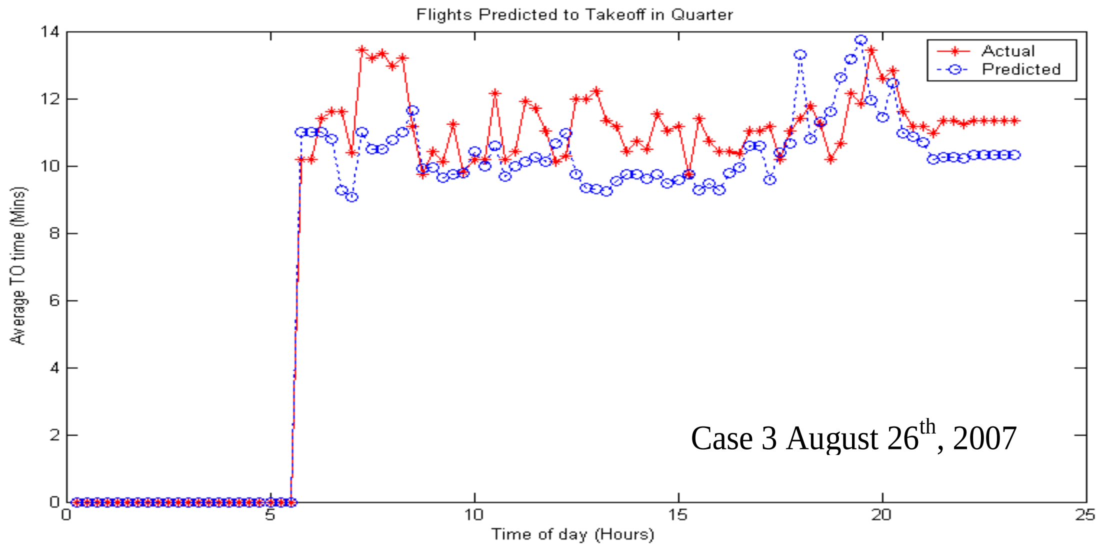

Figure 4. Case 3: Plot of Actual Taxi-Out Time vs. Predicted Taxi-Out Time

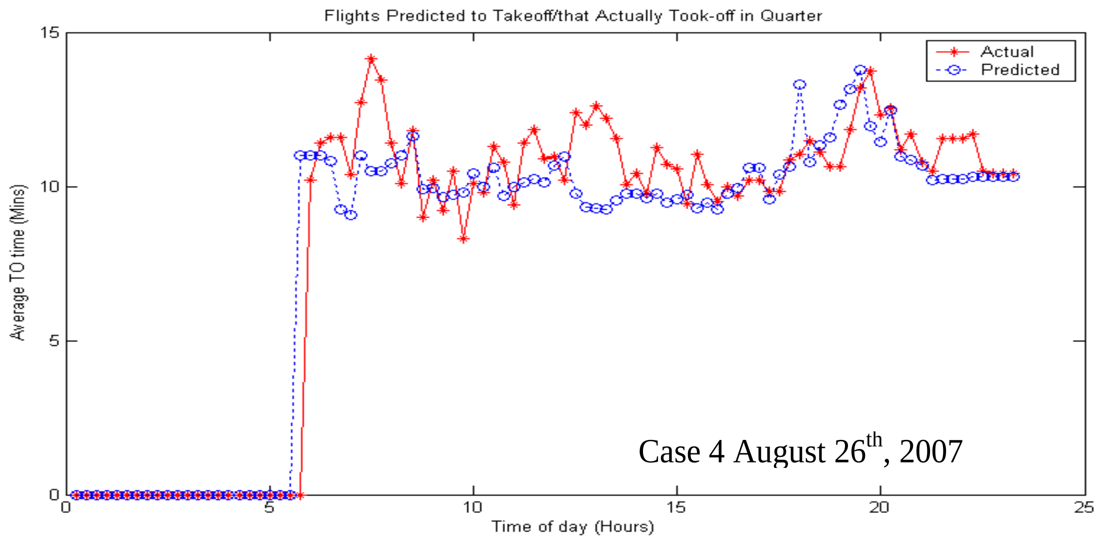

Figure 5. Case 4: Plot of Actual Taxi-Out Time vs. Predicted Taxi-Out Time

ISBN: 978-0-615-20720-9

REFERENCES

[1] Cooper, W.W. Jr., E.A. Cherniavsky, J.S. DeArmon, J.M. Glenn, M.J. Foster, S.C. Mohleji, and F.Z. Zhu. *Determination of Minimum Push-Back Time Predictability Needed for Near-Term Departure Scheduling using DEPARTS*. The MITRE Corporation,2001,url: http://www.mitre.org/work/tech_papers/tech_papers_01/cooper_determination/cooper_determination.pdf

[2] Carr, F., A. Evans, J-P. Clarke, E. Feron. Modeling and Control of Airport Queueing Dynamics under Severe Flow Restrictions. *Proceedings of the American Control Conference*, AACC, Anchorage, AK, 2002.

[3] Atkins, S., and D. Walton. Prediction and Control of Departure Runway Balancing at Dallas Fort Worth Airport. *Proceedings of the American Control Conference*, Anchorage, AK, 2002.

[4] Idris, H., I. Anagnostakis, B. Delcaire, J.P. Clarke, R.J. Hansman, E. Feron, and A. Odoni. Observations of Departure Processes at Logan Airport to Support the Development of Departure Planning Tools. *Air Traffic Control Quarterly*, Vol. 7(4), pp. 229-257, 1999.

[5] Tu, Y., M O. Ball, and J. Wolfgang. *Estimating Flight Departure Delay Distributions - A Statistical Approach with Long-Term Trend and Short-Term Pattern*. Robert H. Smith School Research Paper, RHS 06-034, 2005, url: http://ssrn.com/abstract=923628

[6]

[7] Shumsky, R.A. *Dynamic Statistical Models for the Prediction of Aircraft Take-off Times*. Ph.D. Thesis, Operations Research Center, MIT, Cambridge, MA., 1995.

[8] Idris, H., J.P. Clarke, R. Bhuva, and L. Kang. Queuing Model for Taxi-out Time Estimation. *Air Traffic Control Quarterly*, 2002.

[9] Levine, B.S. and O.H. Gao. Aircraft Taxi-Out Emissions at Congested Hub Airports and the Implications for Aviation Emissions Reduction in the United States. CD-ROM. Submitted to the TRB 2007 Annual Meeting, 2007.

[10] Barrer, J.N., G.F. Swetnam, W.E. Weiss. *The Feasibility Study of using Computer Optimization for Airport Surface Traffic Management*. The MITRE Corporation, MTR89W00010, 1989.

[11] Idris, H.R., B. Delcaire, I. Anagnostakis, W.D. Hall, N. Pujet, E. Feron, R.J. Hansman, J.P. Clarke, and A.R. Odoni. Identification of Flow Constraints and Control Points in Departure Operations at Airport System. *Proceedings AIAA Guidance, Navigation and Control Conference*, AIAA 98-4291, Boston, MA, 1998.

[12] Anagnostakis, I., H.R. Idris, J.P. Clarke, E. Feron, R J. Hansman, A R. Odoni, and W.D. Hall. A Conceptual Design of a Departure Planner Decision Aid. Presented at the *3rd USA/Europe Air Traffic Management R&D Seminar*, Naples, Italy, 2000.

[13] Shumsky, R.A. Real Time Forecasts of Aircraft Departure Queues. *Air Traffic Control Quarterly*, Vol. 5(4), 1997.

[14] Lindsay, K., D. Greenbaum, C. Wanke. Pre-departure Uncertainty and Prediction Performance in Collaborative Routing Coordination Tools. *Journal of Guidance, Control, and Dynamics*, Vol. 28(6), 2005.

[15] Clow, M., K. Howard, B. Midwood, and R. Oiesen. *Analysis of the Benefits of Surface Data for ETMS*. Volpe National Transportation Systems Center, VNTSC-ATMS-04-01, 2004.

[16] Welch, J., S. Bussolari, S. Atkins. Using Surface Surveillance to Help Reduce Taxi Delays. *AIAA Guidance, Navigation & Control Conference*, AIAA-2001-4360, Montreal, Quebec, 2001.

[17] Signor, D.B., and B.S. Levy. *Accurate OOOI Data: Implications for Efficient Resource Utilization*. Sensis Corporation, 25th Digital Avionics Systems Conference, 2006.

[18] Laskey, K.B., N. Xu, C-.H. Chen. *Propagation of Delays in the National Airspace System*. Technical Report, 2006, url:http://ite.gmu.edu/~klaskey/papers/UAI2006_Delay.pdf

[19] Futer, A. Improving Etms' Ground Time Predictions. 25th Digital Avionics Systems Conference, 2006 IEEE/AIAA, pp.1-12, 2006.

[20] Tu, Y., M.O. Ball and W. Jank. Estimating Flight Departure Delay Distributions -- a Statistical Approach with Long-Term Trend and Short-Term Pattern. Technical Report, U. of Maryland, 2005.

[21] Gosavi, A. *Simulation Based Optimization: Parametric Optimization Techniques and Reinforcement Learning. Norwell*. MA: Kluwer Academic, 2003.

[22] Bellman, R. *The Theory of Dynamic Programming*. Bull. Amer. Math. Soc., vol. 60, 1954, pp. 503-516.

[23] Howard, R. *In Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA, 1960.

[24] Bertsekas, D., and J. Tsitsiklis. *In Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1995.

[25] Puterman, M.L. *Markov Decision Processes*. Wiley Interscience, New York, 1994.

[26] Robbins, H., and S. Monro. A Stochastic Approximation Method. *Ann. Math. Statistics*, vol. 22, 1951 pp. 400-407.

[27] http://www.bts.gov/programs/airline_information/accounting_and_reporting_directives/technical_directive_15.html

[28] Meier, C., P. Eriksen. *Total Airport Management: A Step Beyond Airport Collaborative Decision Making*. 2006,url: http://www.eurocontrol.int/eec/public/standard_page/EEC_News_2006_3_TAM.html .

 ISBN: 978-0-615-20720-9