Application of reinforcement learning algorithms for predicting taxi-out times

Poornima Balakrishna, Ph.D. Candidate

George Mason University, Fairfax, VA, U.S.A
Sensis Corporation, Reston, VA, U.S.A

Dr. Rajesh Ganesan, Ph.D.
Dr. Lance Sherry, Ph.D.

Center for Air Transportation Systems Research
George Mason University, Fairfax, VA, U.S.A

ATM Seminar 2009
Napa, California, U.S.A

Structure of the presentation

1) Introduction and problem definition

2) Modeling the taxi-out time estimation problem using approximate dynamic programming (reinforcement learning)

3) Prediction accuracy results
   a. Individual flights
   b. Average in 15 minute intervals of the day

4) Overview of the research methodology

Taxi-out time predictions: Java interface

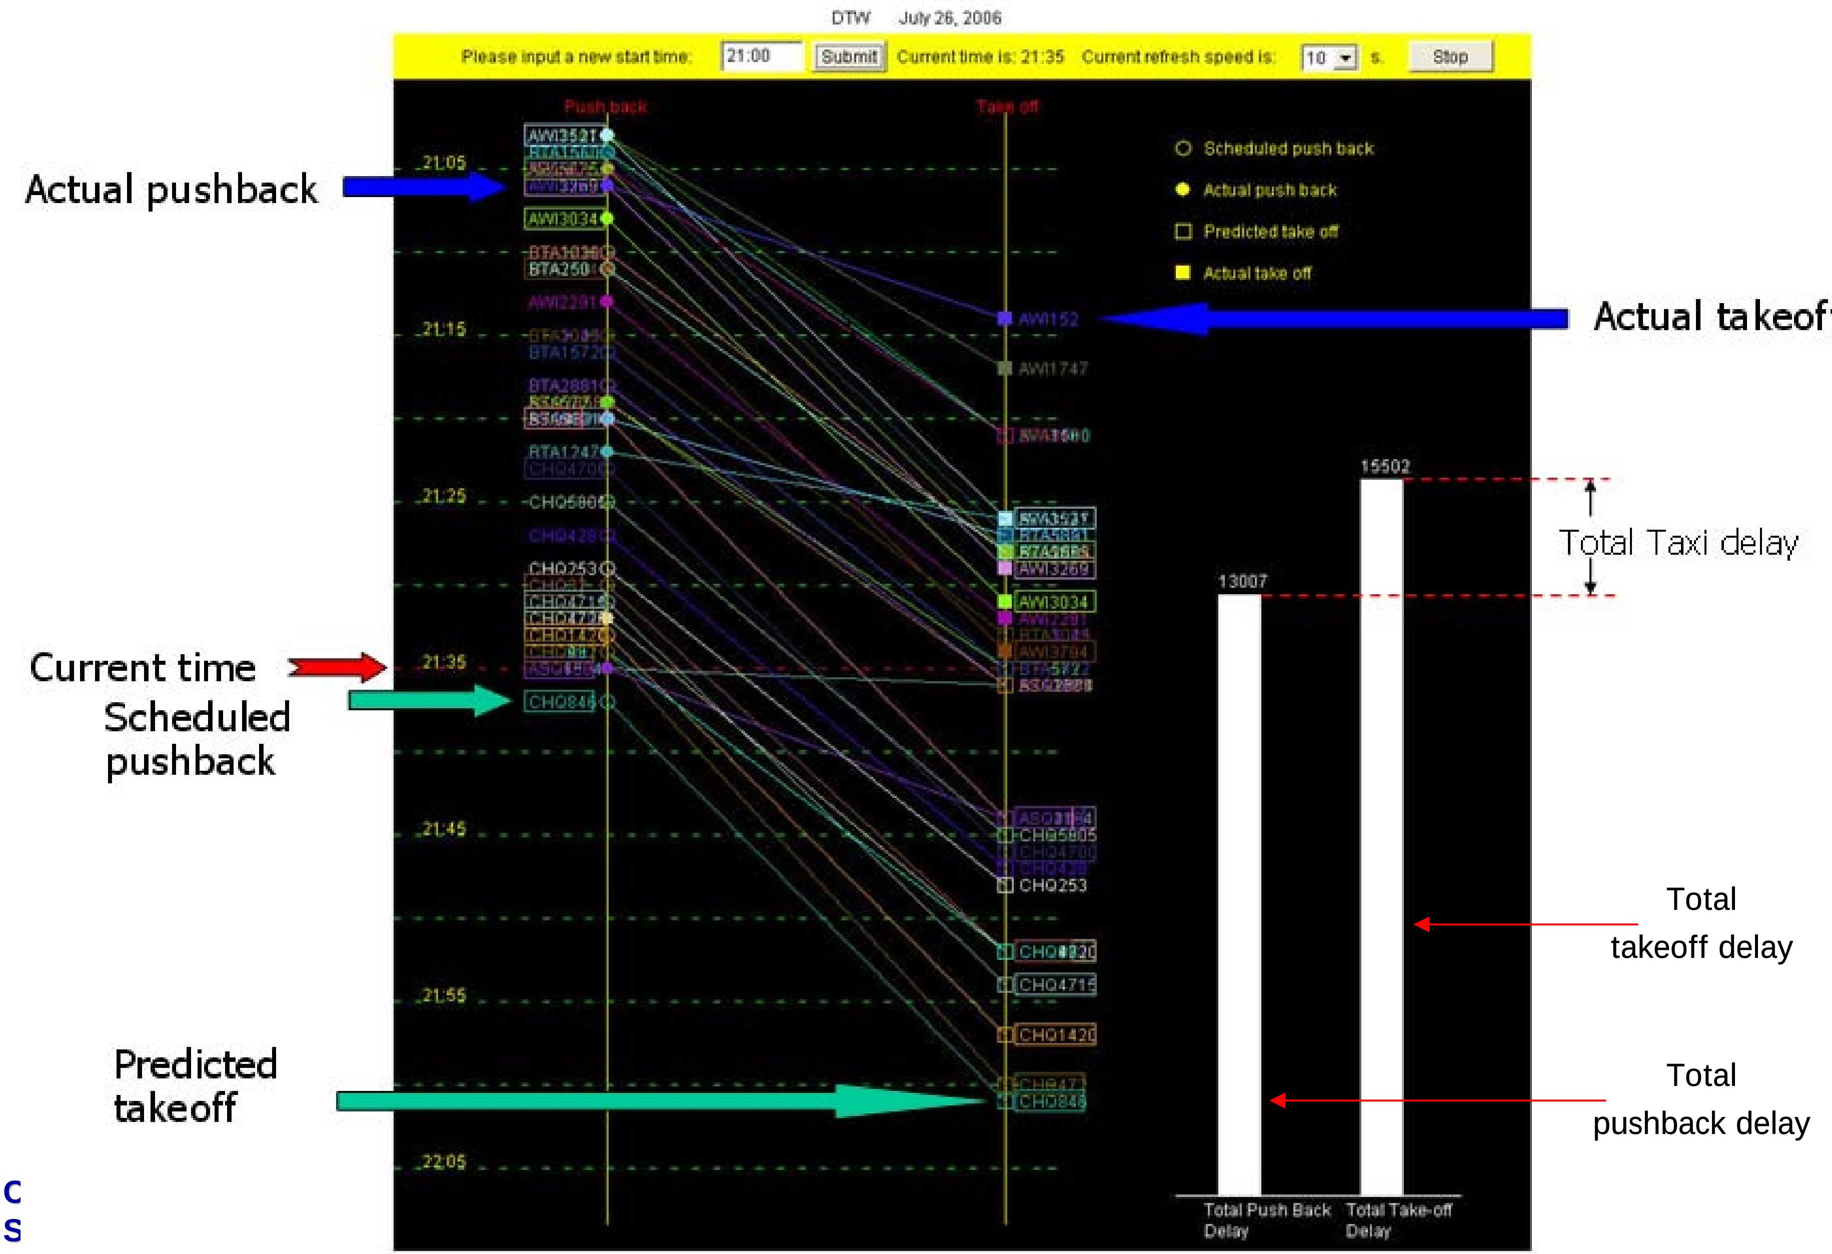

Taxi-out time prediction: motivation

- Taxi-out time (duration between gate-pushback and takeoff) contributes to 20 % of total flight delay (FAA report).
- In order to minimize taxi-out delay it is necessary to first predict taxi time under dynamic airport conditions.
- Major benefits of accurately predicting taxi-out times:
  - Congestion mitigation by avoiding near-capacity operations (DOT)
  - Emissions compliance through optimal adjustment of departure schedules (NAAQP)
  - Reduction in returns to the gate for refueling (Dalton personnel communication)
  - Efficient resource utilization (ground personnel, gates) (NextGen)
  - Estimated departure clearance time compliance (Sherry and Belle, 2008)

Need for a stochastic dynamic approach

The FAA is looking towards a modernized ATC system with increased automation (NextGen, JPDO):

- Predictions in real-time, as the system evolves.
- Method that is adaptive to changing airport dynamics
- Due to uncertainties involved, and the complex nature of airport operations, it is often difficult to obtain mathematical models to completely describe airport dynamics.

The above can be addressed using reinforcement learning (a strand of stochastic dynamic programming):

- The problem of sequential prediction is well-suited to the stochastic dynamic programming formulation
- RL is a model free approach that is adaptive to changing airport dynamics.
- RL learns by interacting with the environment (alleviating need for good training data in neural networks).
- Suitable for large-scale optimization due to its simple recursive formulation.

Comparison of stochastic dynamic programming method with literature

Year, author	Approach	Data range/airport	Data used	Additional information
Idris et al. (2002)	Queueing model	BOS, August 1998	ASPM, downstream restrictions, PRAS	Predictions are not a-priori of pushback
Signor and Levy (2006)	Bi-variate quadratic regression	DTW	Surface surveillance	Taxi-out time defined from ramp area/spot to takeoff
Futer (2006)	Running average	Used by the FAA at several airports	ETMS (based on flight plan)	
Balakrishna, Ganesan, Sherry (2009)	Stochastic dynamic programming	DTW: Sep2005-Aug2008 JFK: May-Oct 2007 May-Oct 2008	ASPM	Predictions at least 15 min in advance of gate pushback, sequential predictions as time evolves

BOS: Boston Logan International airport

DTW: Detroit International airport

JFK: John F. Kennedy International airport

Acronyms:

PRAS: Preferential Runway Advisory System

ASPM: Aviation System Performance Metrics

ETMS: Enhanced Traffic Management System

FAA: Federal Aviation Administration

Data Source

- OOOI (Out,Off,On,In) data from the Aviation System Performance Metrics (ASPM) database maintained by the Federal Aviation Administration (FAA) was used.

Snapshot of the ASPM data (Departing and Arriving Flights)

FLTNO	SCHOUTTM	ACTOUTTM	NOMTO	ACTTO	SCHOFFTM	ACTOFFTM
Flight Number	Scheduled Out time	Actual Out time	Nominal TO time	Actual TO time	Scheduled Off time	Actual Off time
275	07:19	07:30	14.4	13	07:33	07:43

ACTONTM	NOMTI	ACTTI	ACTINTM
Actual On time	Nominal TI time	Actual TI Time	Actual In Time
10:11	5.5	6	10:17

This research does not use surface surveillance (track) data

Stochastic Dynamic Programming Formulation

$P(s_t, x, s'_{t+1})$ → transition probability

$C(s_t, x)$ → error

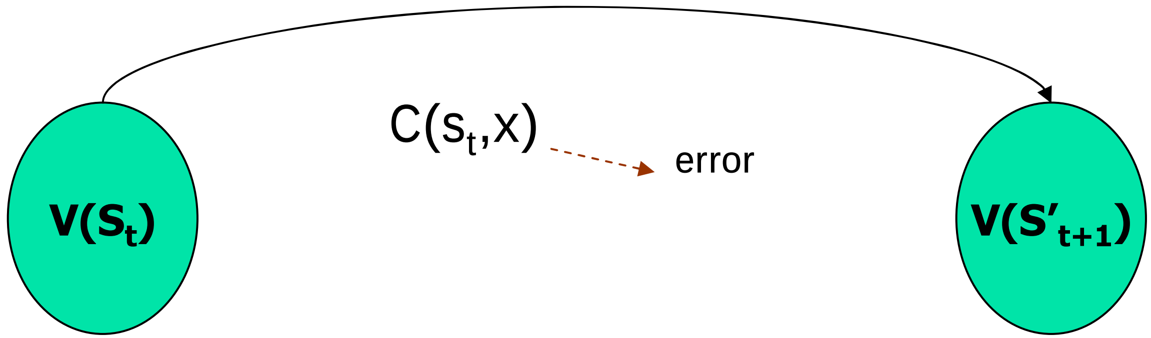

(The idea) Value Function

$$V(S_t) = \min_x [C(S_t, x) + \gamma E[V(S'_{t+1})]]$$

γ is a fixed discount parameter

C=Error= |actual TO-predicted TO|

Learning version based on Q-factors

$Q(S,x) = (1-\alpha)Q(S,x) + \alpha[C(S,x) + \gamma\{\min_b Q(S,b)\}]$

C(S,x)= |actual TO-predicted TO|

α is the learning parameter

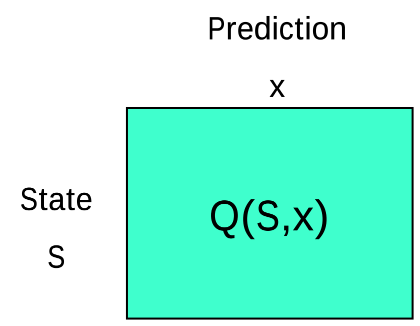

$$V(S) = \min_x Q(S, x)$$

Implementation– RL Estimating Taxi-out

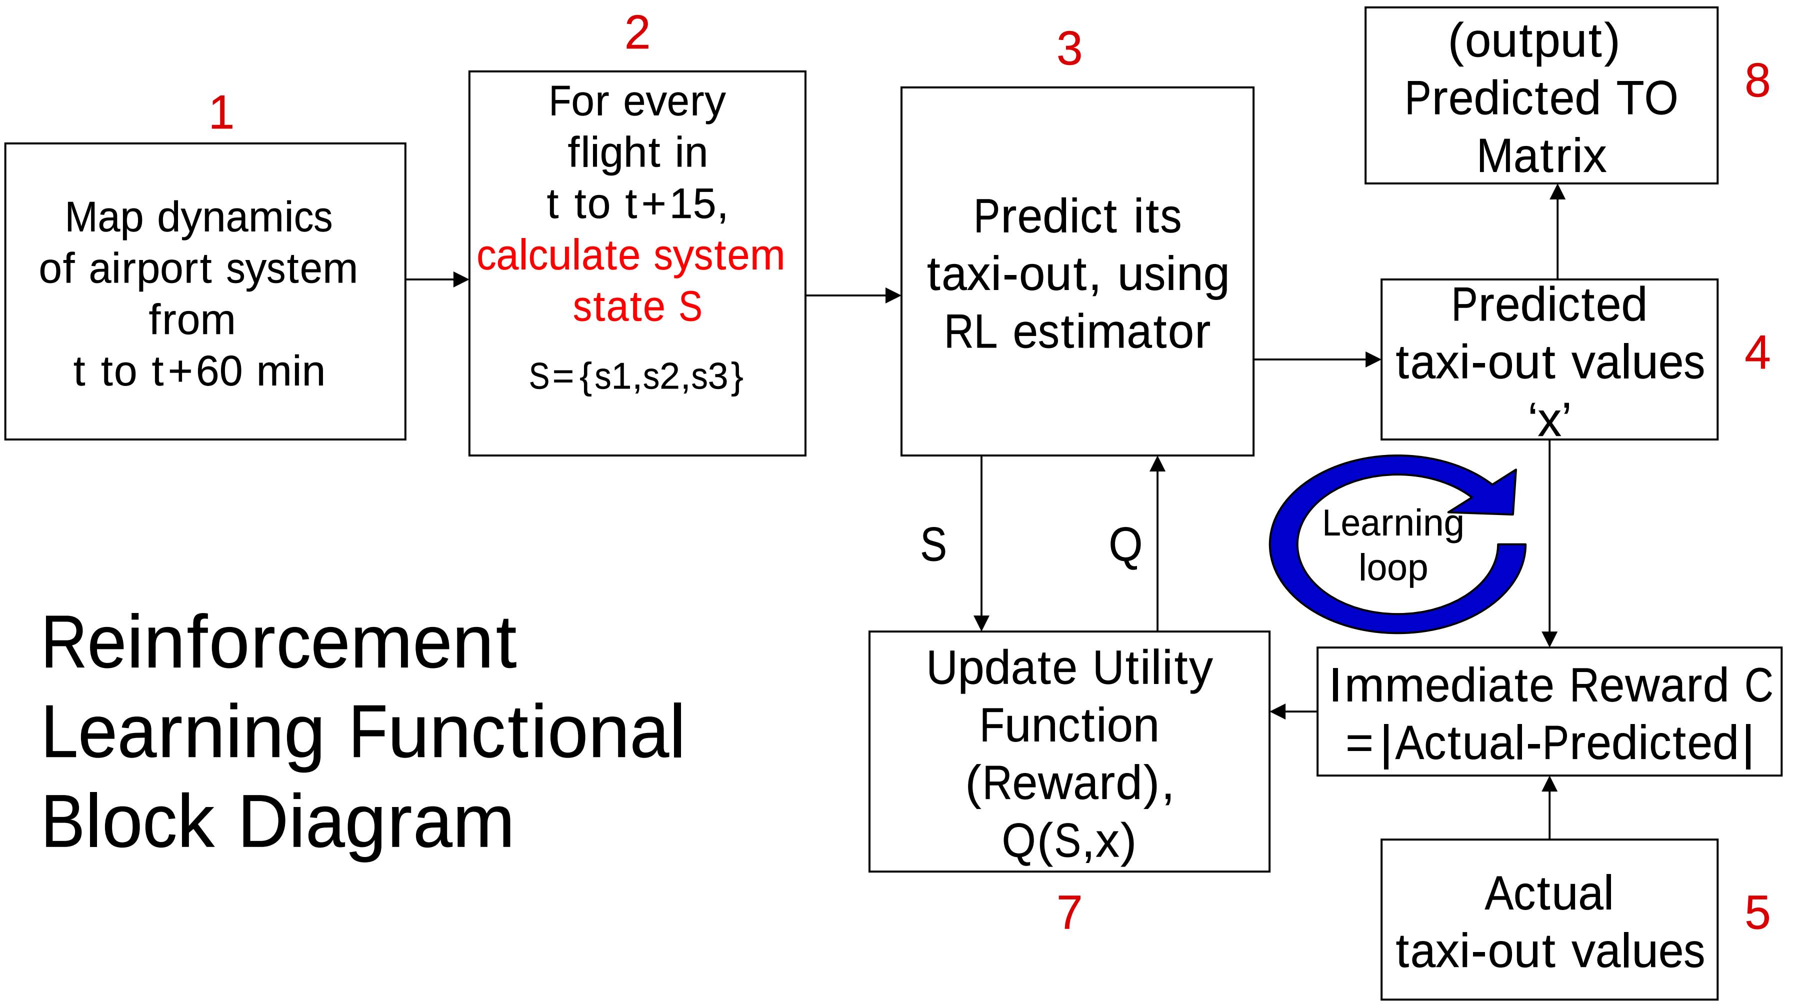

Reinforcement Learning Functional Block Diagram

Learnt phase of the Q-Learning Approach

- To obtain taxi-out time prediction for a flight,
  - Identify its system state
  - Look for smallest non-zero Q value in corresponding row
  - The corresponding prediction x^* is the taxi-out time estimate

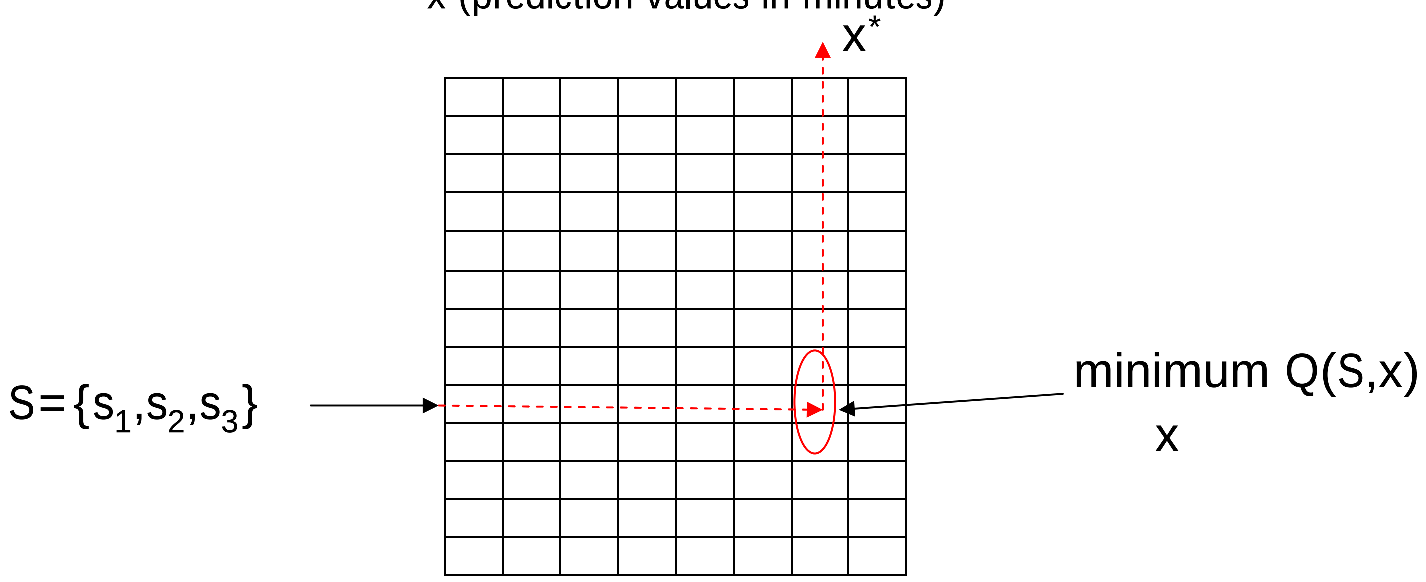

Q matrix

Airports analyzed

- DTW (Detroit International)
- JFK (John F. Kennedy International)

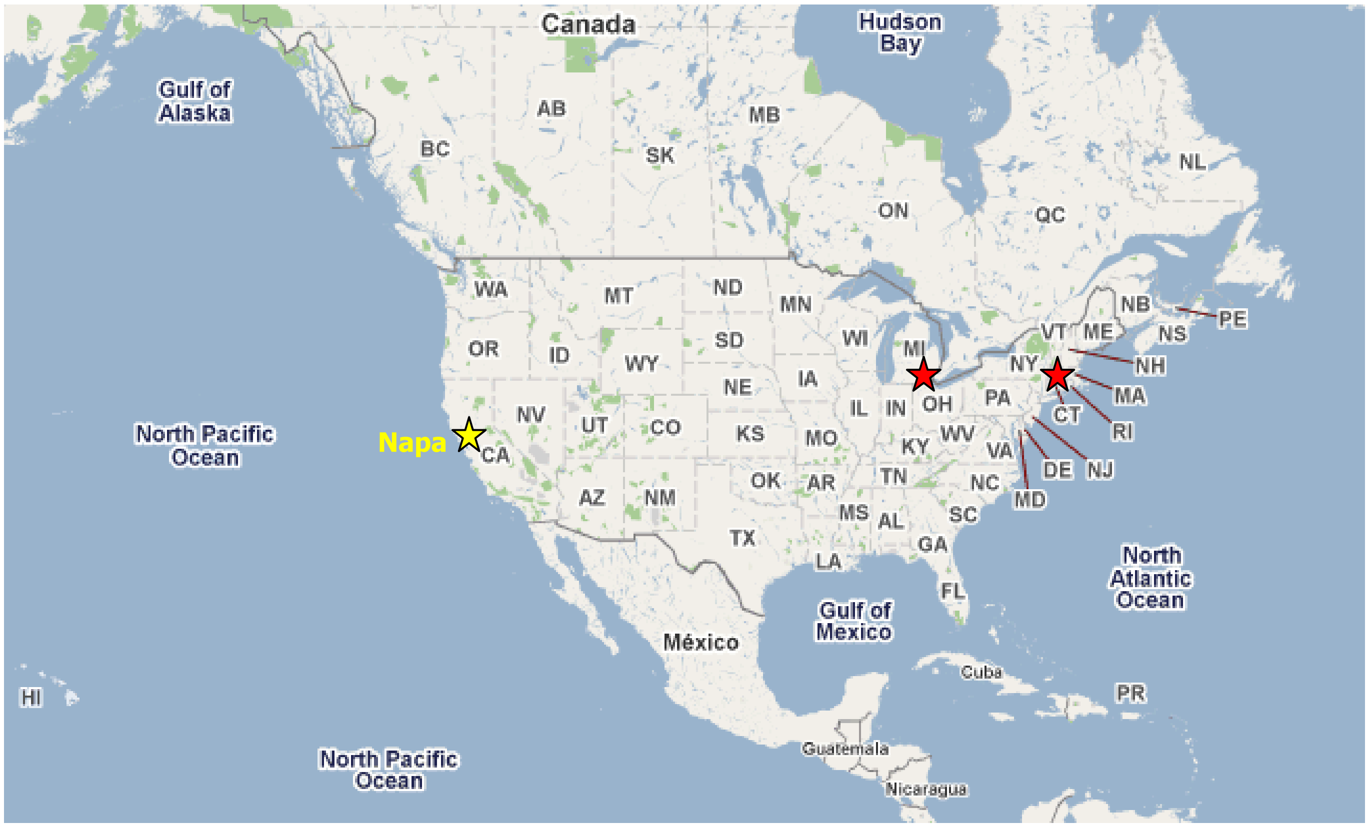

maps.google.com

DTW: Detroit International airport

- Days used for learning –
  - Scenario 1: March-August (Spring-Summer)
  - Scenario 2: September – February (Fall-Winter)
- In each scenario, 42 days were selected at random for testing
- Years studied:
  - September 2005 – August 2008

Daily airport analysis: percentage prediction accuracy within ± 4 min, DTW airport

Individual flight prediction accuracy across 42 days of testing

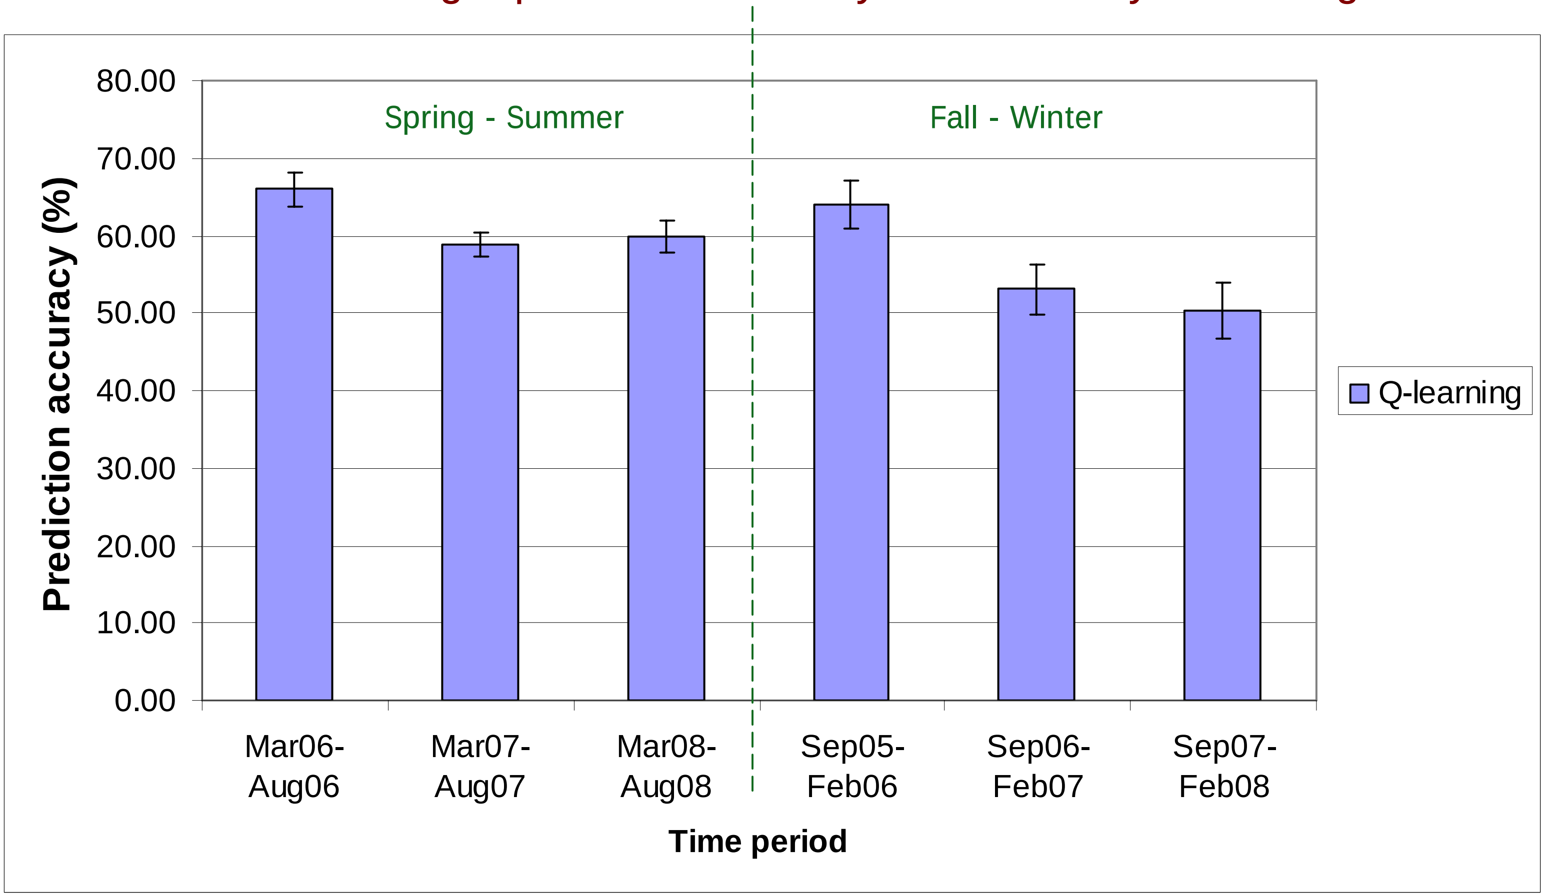

Daily airport analysis: percentage prediction accuracy, DTW airport

Average prediction accuracy in 15 min intervals of the day

Year	Fall –Winter (2007)			Spring – Summer (2008)		
Date	Dec 15th	Nov 8th	Dec 4th	May 2nd	June 13th	May 27th
Prediction accuracy within ± 4 min (%)	53.62	68.12	85.51	63.77	79.71	98.55

DTW airport: Daily taxi-out time standard deviation (Fall-Winter)

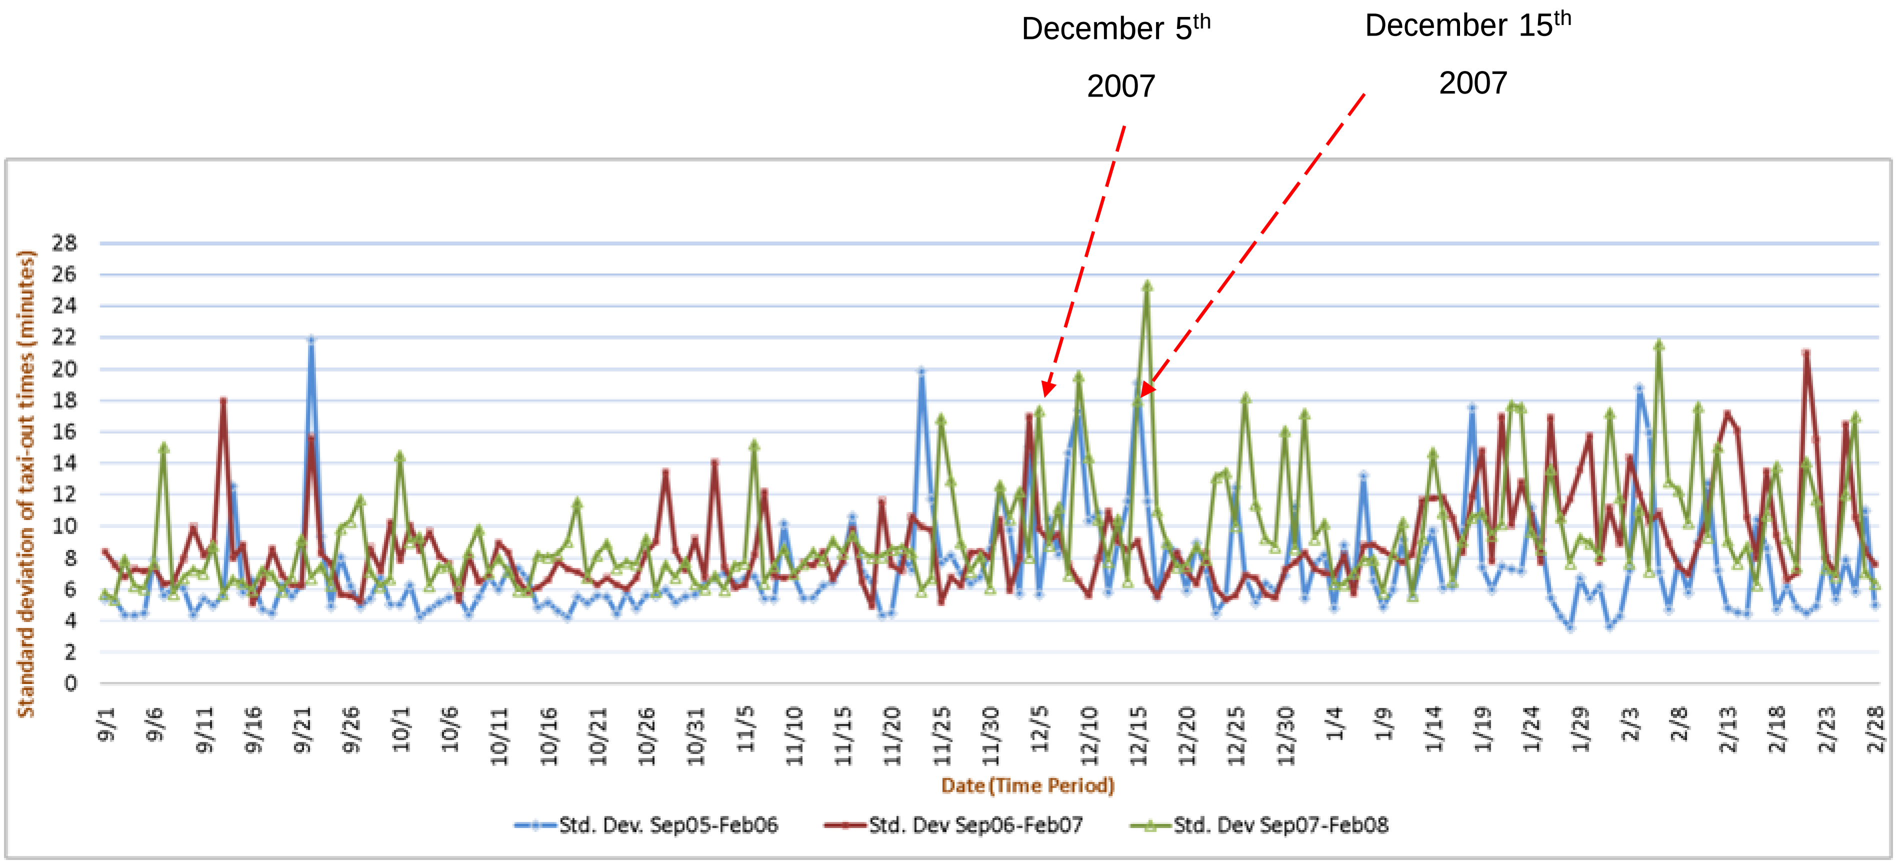

Source: ASPM database, FAA

DTW airport behavior, December 5, 2007

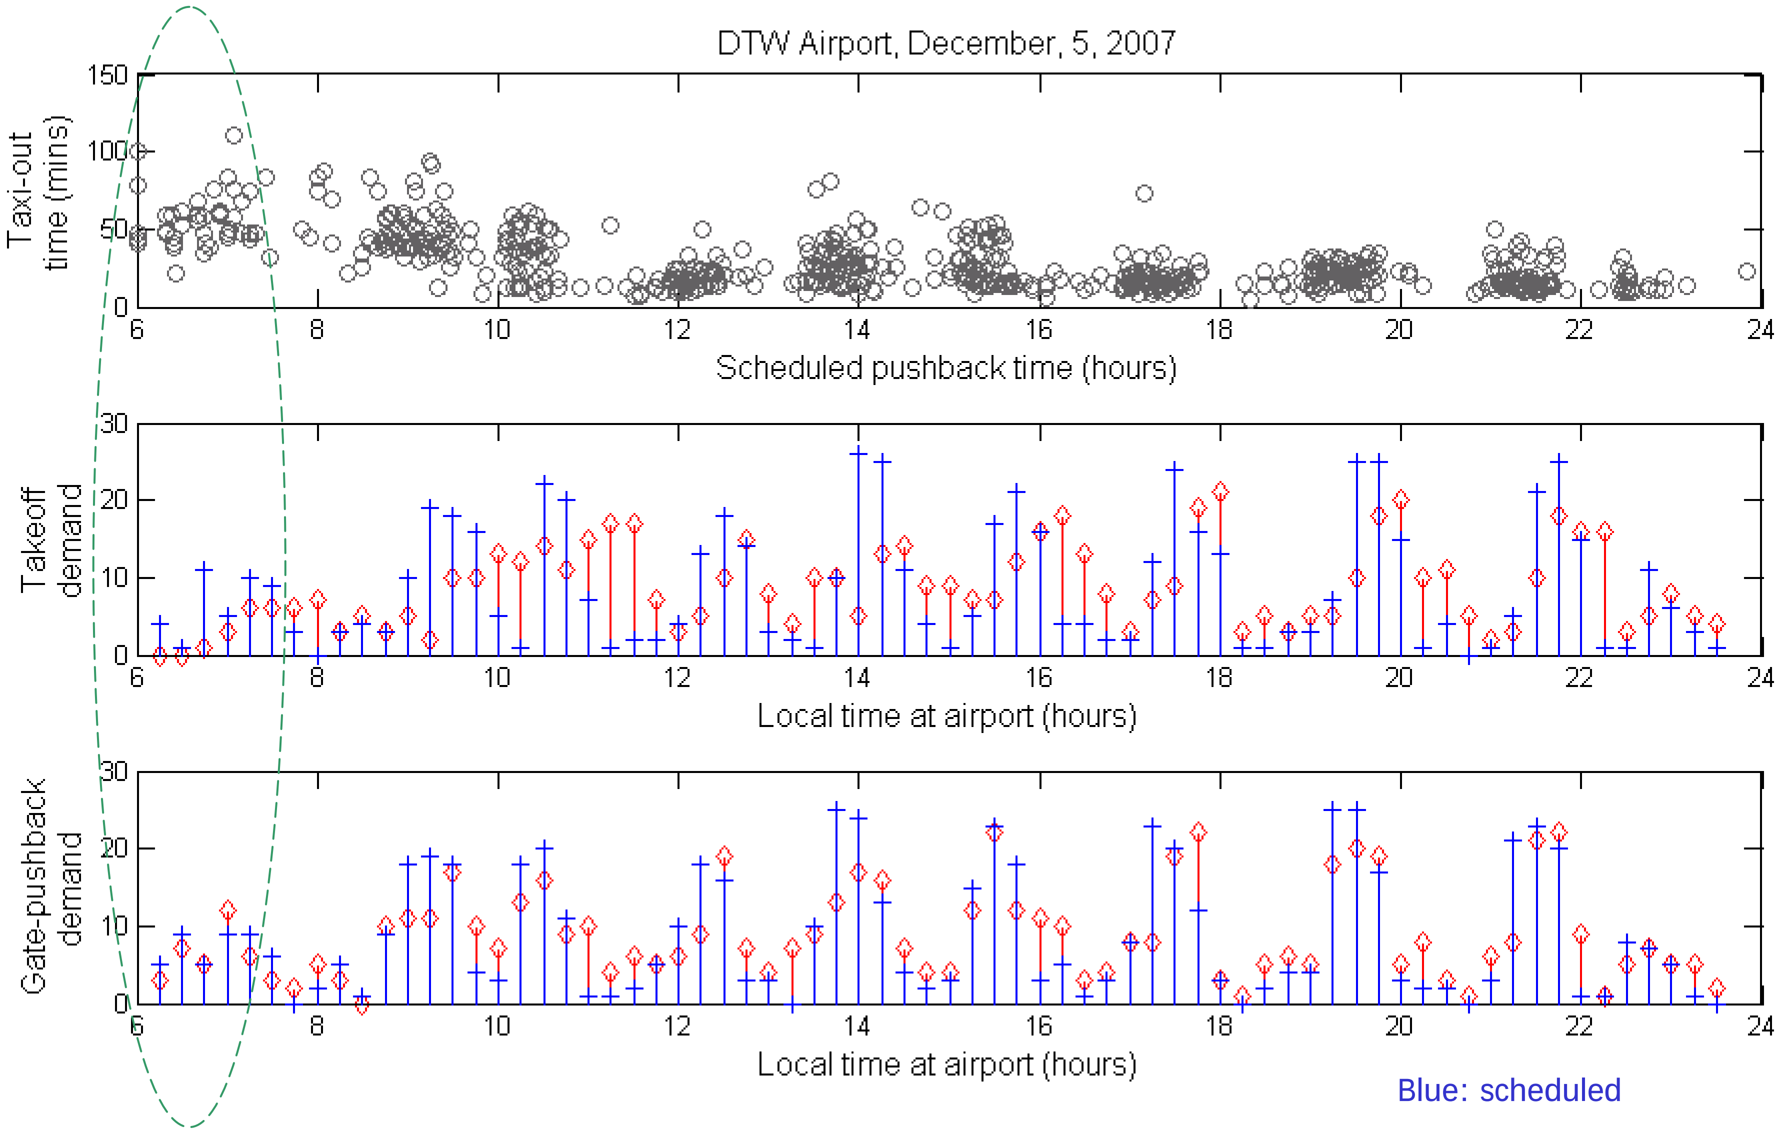

Blue: scheduled

Red: actual

Data source: ASPM

Results: Detroit International Airport (DTW)

JFK: John F. Kennedy International airport (challenges)

- Literature (Andrew Compart, Aviation Week) suggests that
  - Percentage of on-time departures at JFK is about 65 % (both in July 2007 and July 2008)
  - On time departure performance worsened significantly in the late-afternoons and evenings.
- Wide variations in taxi-out times across a single day (about 20 min to as high as 130 min)

Airport	Date	Mean of actual TO times (min)	Standard deviation of actual TO times (min)
DTW	29th Jan 2006	17.9	8.9
DTW	27th Jul 2006	17.1	8.2
JFK	4th Jul 2007	31.0	14.2
JFK	4th Dec 2007	37.3	19.5

JFK airport: actual operations

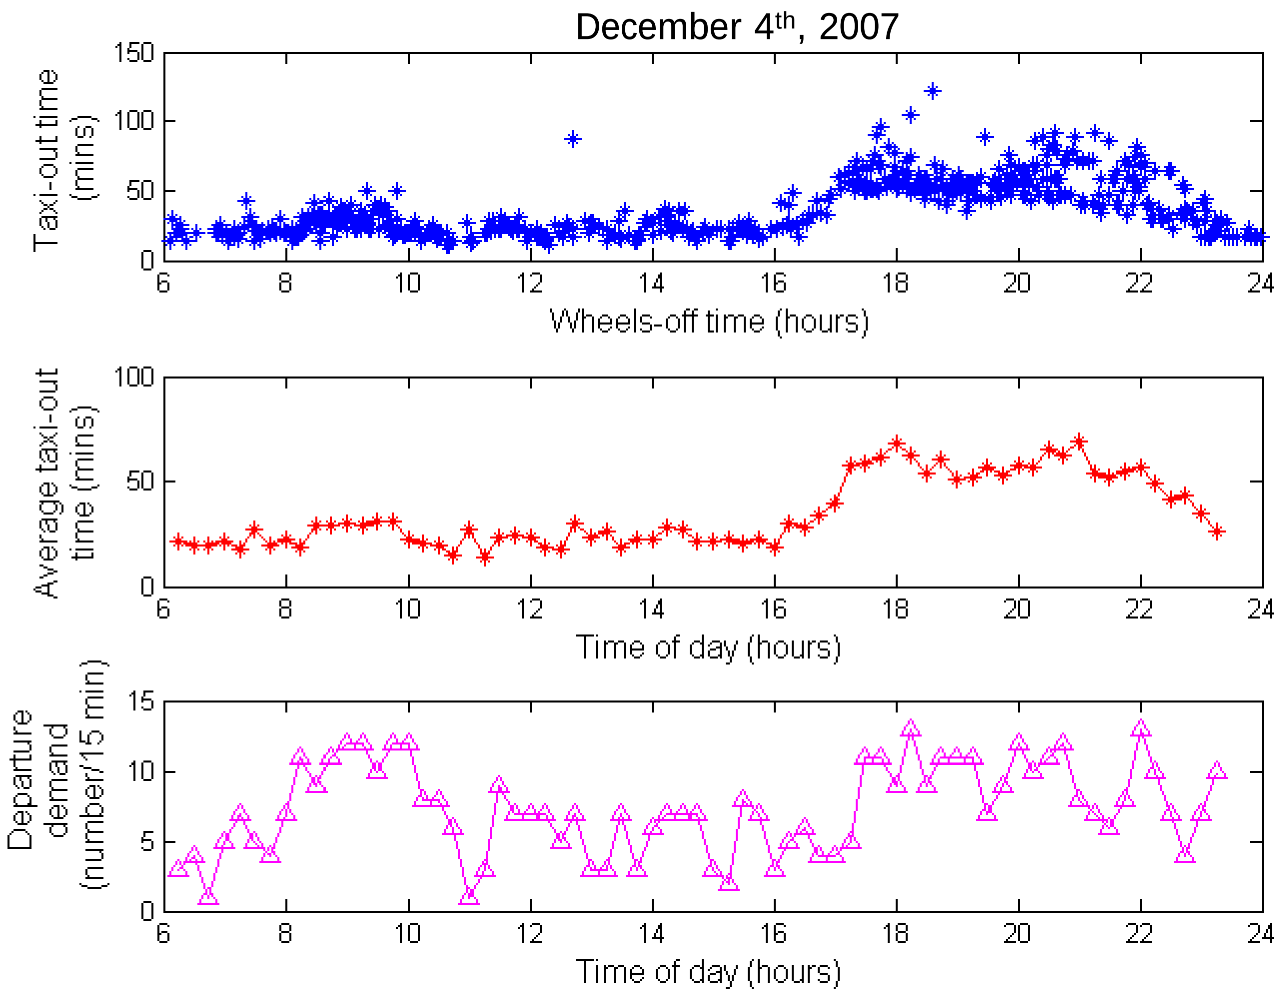

JFK: John F. Kennedy International airport

- Days used for learning –
  - Scenario 1: May-October
  - 42 days were selected at random for testing
  - Years studied: 2007, 2008

  - Scenario 2: April-November
  - 10 days were selected at the end of the period for testing
  - Years studied: 2007

Daily airport analysis: percentage prediction accuracy within accuracy of ± 5 min and ± 8 min, JFK airport

Individual flight prediction accuracy across 42 days of testing

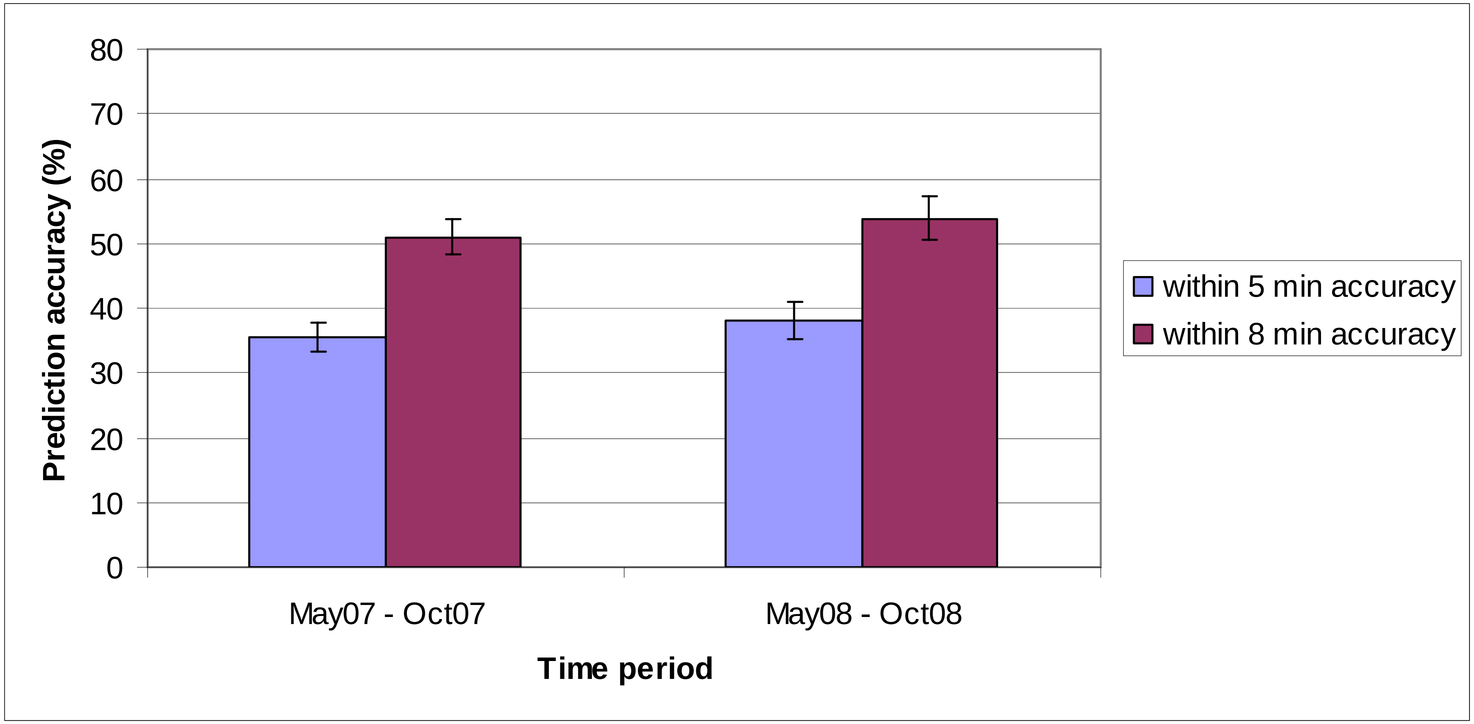

Results: John F. Kennedy International Airport (JFK)

Daily airport analysis: percentage prediction accuracy, JFK airport , November - December 2007

Average prediction accuracy in 15 min intervals of the day

Year 2007 Date	Time Period of Day	Dec 5th	Dec 7th	Nov 29th
Prediction accuracy within ± 5 min (%)	Before 4:00 P.M	67.50	77.50	100.00
	After 4:00 P.M	41.38	58.62	55.17
	Across Whole Day	56.52	69.57	81.16

JFK airport: Effect of ADR, Change in runway configuration

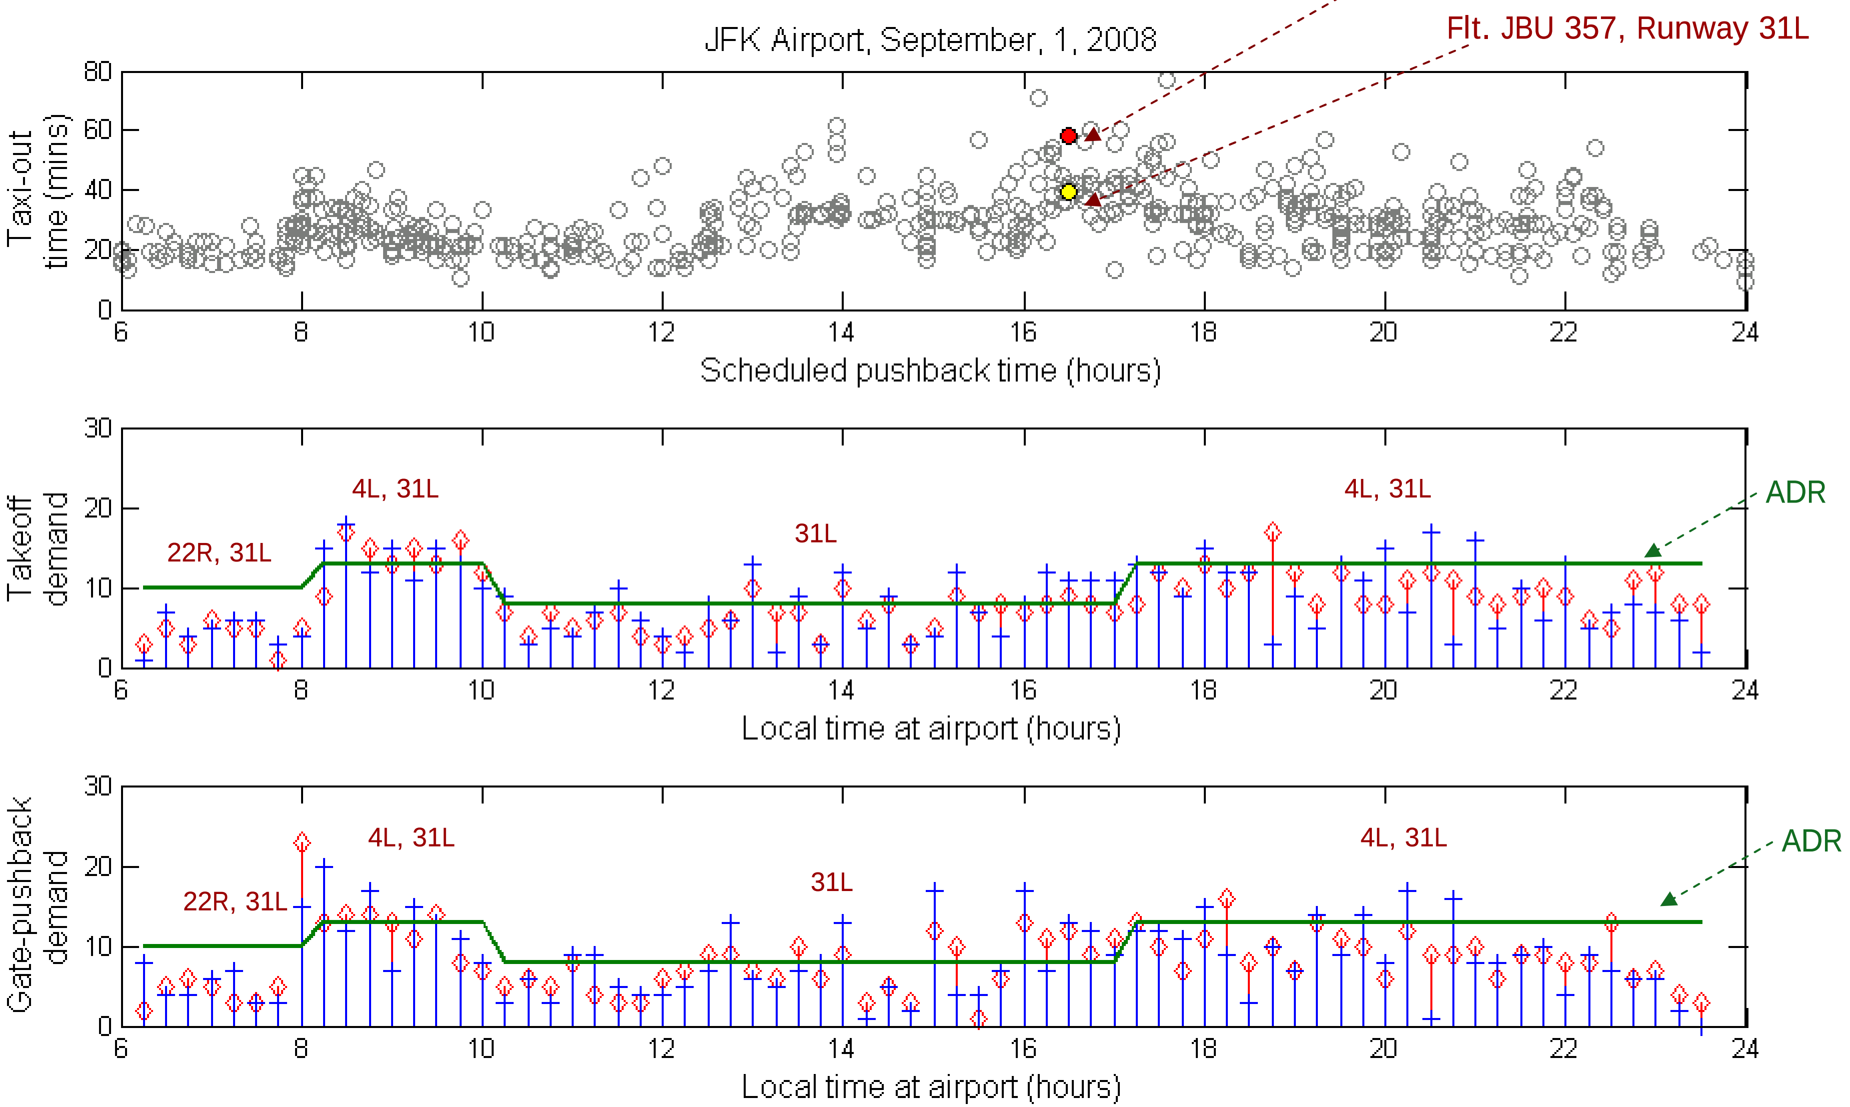

Blue: scheduled

Red: actual

Data source: ASPM

(Observations based on viewing AEROBAHN®, Sensis Corporation)

Macro-view of the research

Sequential decision making

Mathematical Programming

Dynamic Programming

(Markov decision process)

Classical DP

Curse of modeling arises due to need for transition probabilities

Suitable for deterministic problems and small scale problems

Approximate DP

Alleviates Curse of modeling

Curse of dimensionality persists

Q-learning

Explicit Storage

Value function (Reward learning)

$V = \underset{x}{\text{minimum}}(Q)$

Value function Approximation

(Approach for large scale problems)

Aggregation

Interpolation

Function fitting

Linear, nonlinear, piecewise linear, neural networks

Diffusion Wavelets

Thank you!

Results - Prediction Accuracy of **average TO** in 15 min intervals

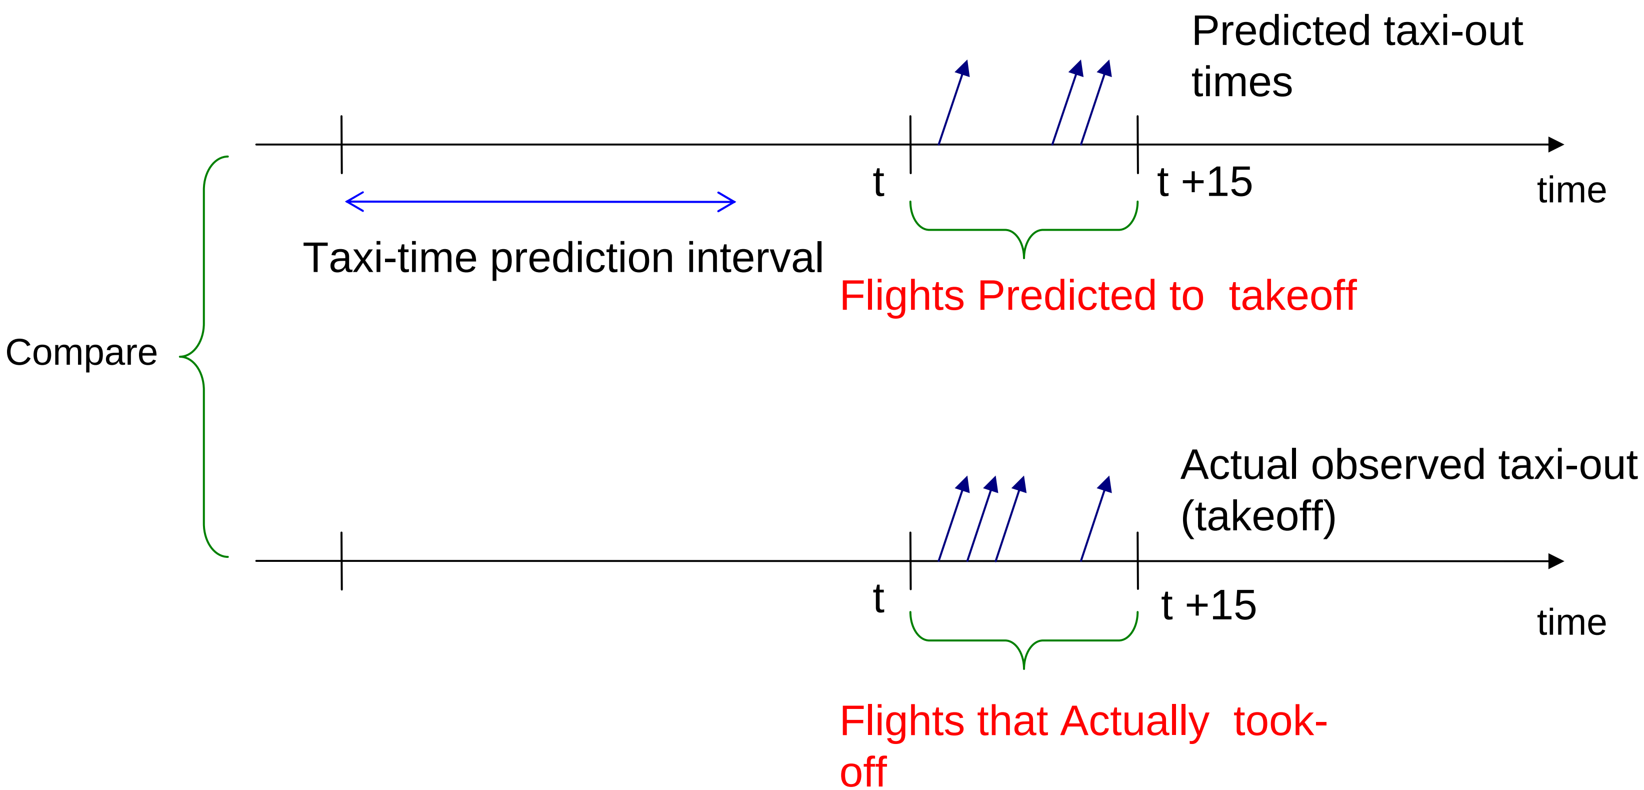

This is indicative of how well airport behavior is predicted in advance

Adapting RL to Taxi-out Prediction

- Markov Chain: The evolving airport system dynamics is perceived as a Markov chain (i.e. Future system state depends on the present system state and not on the past)
- Markov Decision Process: A taxi-out time prediction x is made 15 minutes before scheduled gate departure of a flight based on a system state S (explained later)
- The objective is to find the best prediction in every state S, which minimizes the error in prediction
- The above is modeled as a stochastic dynamic programming problem

Macro-view of analysis

Taxi-out time estimation (TOTE)

Regression model for TOTE

Approximate Dynamic Programming model

Q-learning

(explicit storage of Q-matrix)

Value function (Reward learning)

V = minimum(Q)

Value function Approximation (VFA)

(for large state spaces)

Diffusion wavelets

(for large and multi-dimensional state spaces)

DTW airport: individual flight prediction accuracy within $\pm$ 4 min

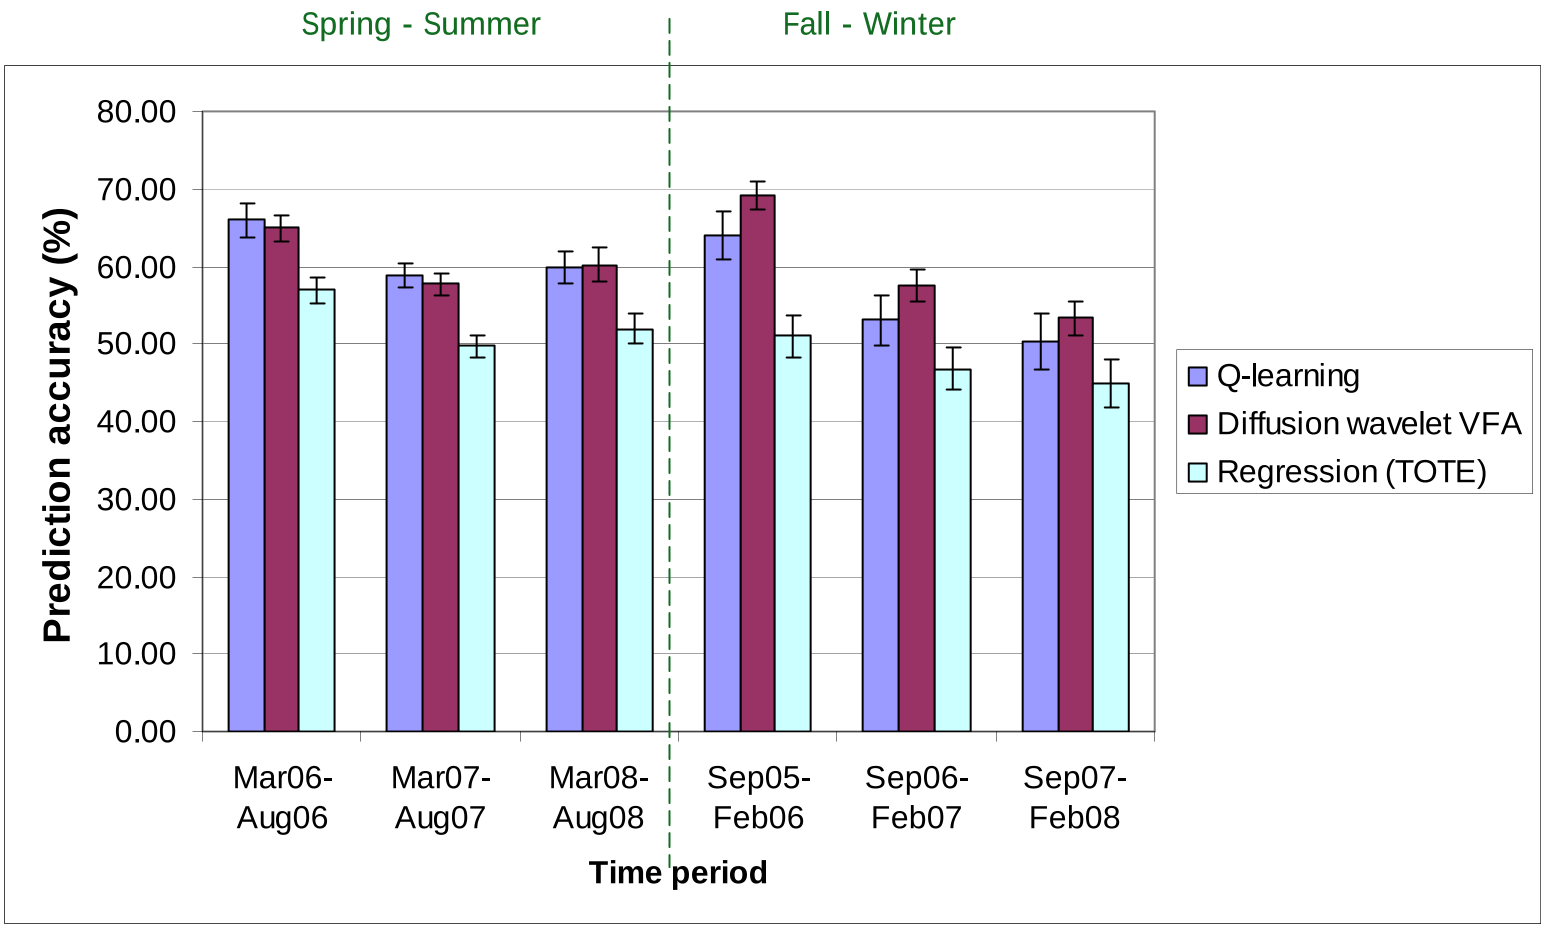

JFK airport: individual flight prediction accuracy within ± 8 min

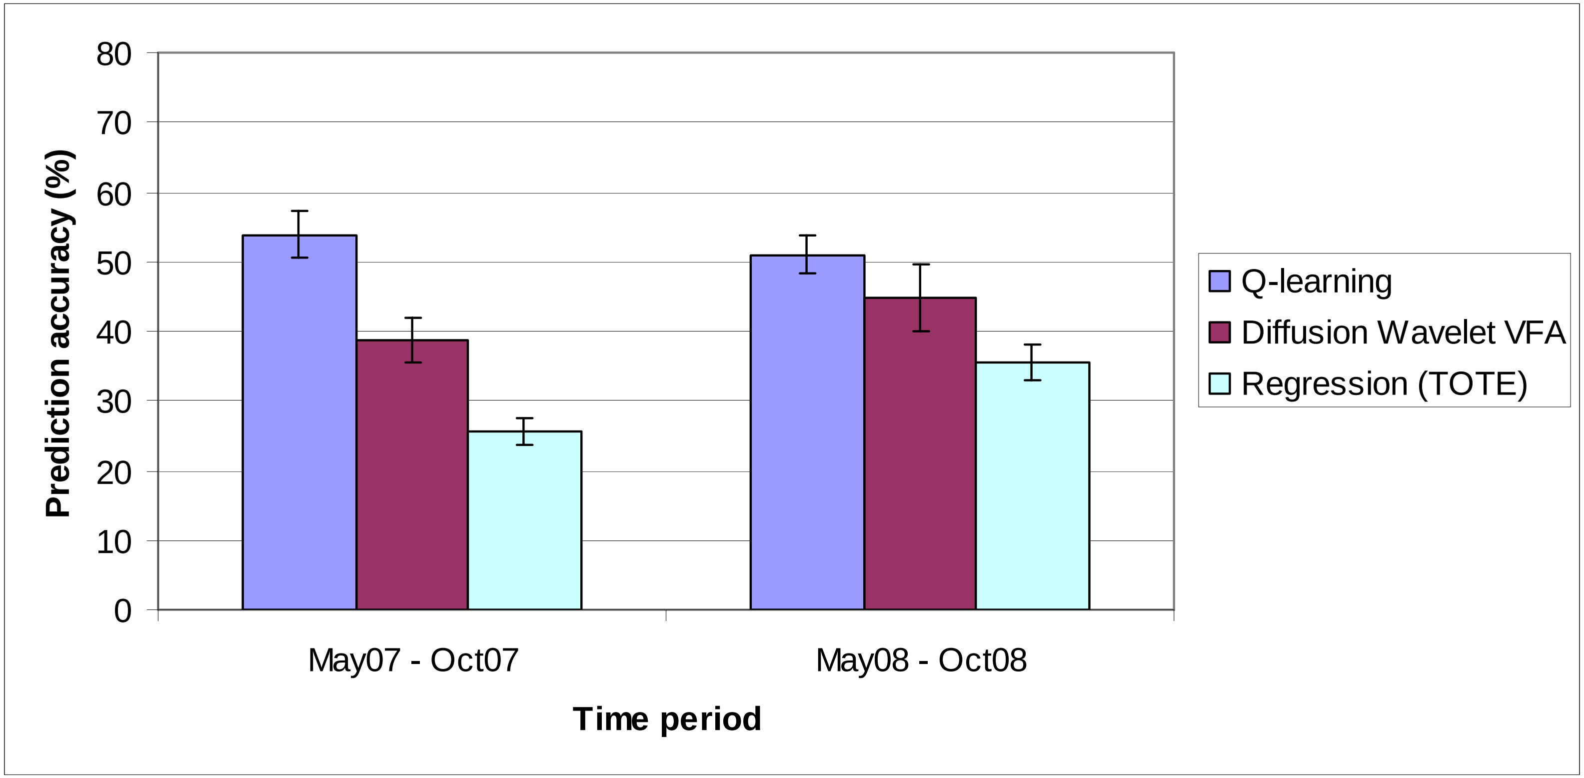

Future work

- Expansion of the state space for the taxi-out time estimation model, to include features such as day of the week and forecasted weather information.

- Application of wavelet analysis for detection of anomalies in the airport departure process data, to gain a better understanding of taxi-out time behavior.

Future work

- Use of taxi-out predictions as inputs to a decision making problem for departure sequencing and downstream adjustments.

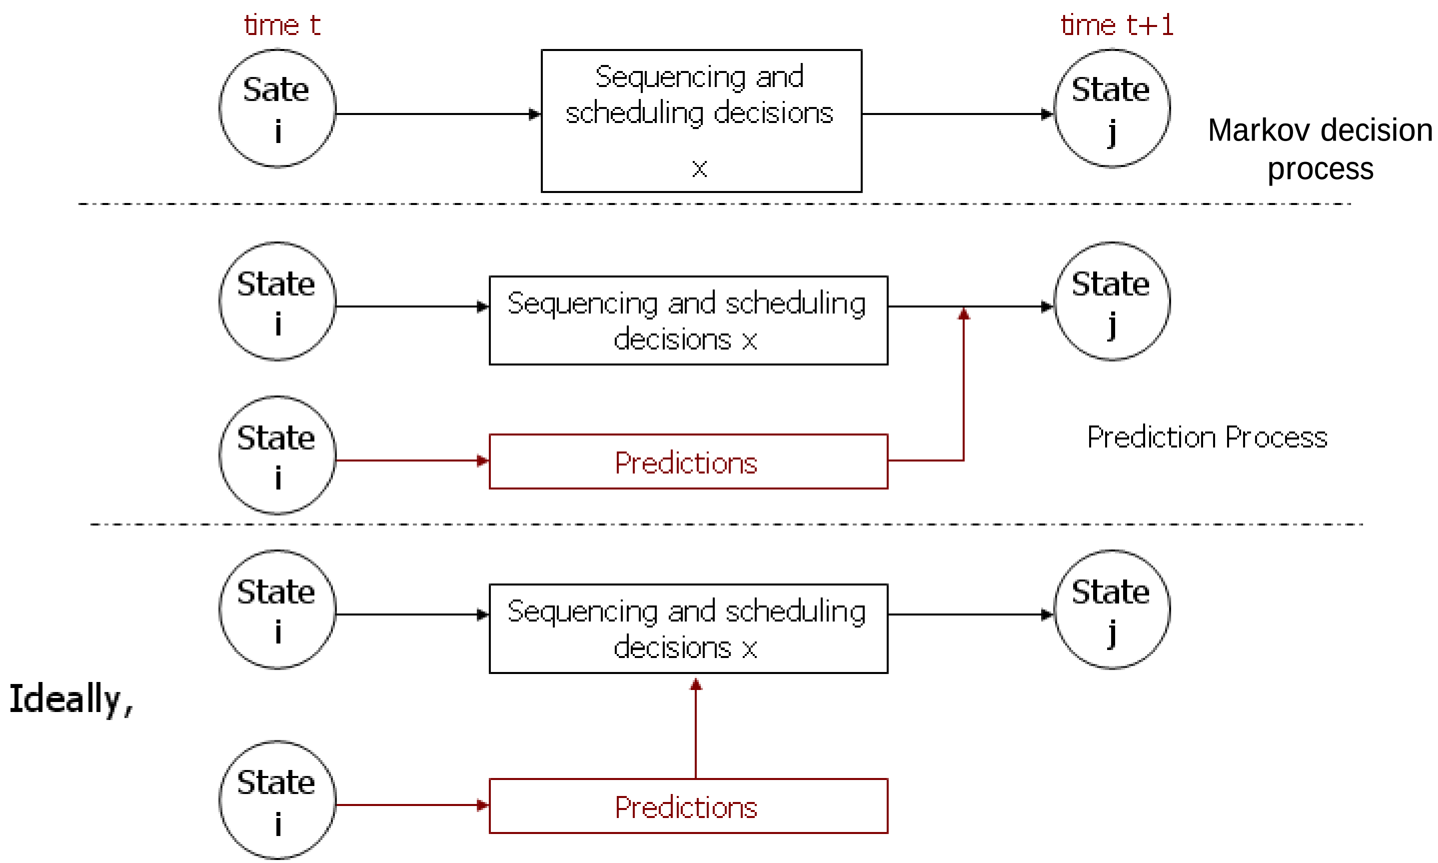

Future work: Surface Surveillance Data

- Surveillance data supplements ASPM data by providing detailed information on nominal taxi times between specific gate/ramp area and runway pairs.

  Implications:

  - Improved accuracy of the RL algorithm in determining when a flight enters the runway queue
  - More accurate predictions on an individual flight basis
  - More precise records of OOOI event times which are more readily accessible than proprietary data (from airlines for example)