Performance Evaluation of SARDA: An Individual Aircraft-based Advisory Concept for Surface Management

Yoon Jung, Waqar Malik, Leonard Tobias, Gautam Gupta, Ty Hoang, and Miwa Hayashi

Surface operations at airports in the US are based on tactical operations, where departure aircraft primarily queue up and wait at the departure runways. NASA's Spot And Runway Departure Advisor (SARDA) tool was developed to address these inefficiencies through air traffic control tower advisories. The SARDA system is being updated to include collaborative gate-hold, either tactically or strategically. This paper presents the results of the human-in-the-loop evaluation of the tactical gate-hold version of SARDA in a 360° simulated tower setting. The simulations were conducted for the east side of the Dallas/Fort Worth airport. The new system provides gate-hold, ground controller and local controller advisories based on a single scheduler. Simulations were conducted with SARDA on and off, the off case reflecting current day operations with no gate-hold. Scenarios based on medium (1.2x current levels) and heavy (1.5x current levels) traffic were explored. Data collected from the simulation were analyzed for runway usage, runway queue size, delay for departures and arrivals, and fuel consumption. Further, traffic management initiatives were introduced for a subset of the aircraft. Results indicated that runway usage did not change with the use of SARDA, i.e., there was no loss in runway throughput

Yoon Jung, Ty Hoang and Miwa Hayashi are with NASA Ames Research Center, Moffett Field, CA; Gautam Gupta, Waqar Malik, and Leonard Tobias are with the University of California, Santa Cruz, NASA Ames Research Center, Moffett Field, CA.

as compared to baseline. At the same time, runway queue size did not exceed seven aircraft under SARDA, compared to a maximum of sixteen in the baseline. Taxiing delay was significantly reduced with the use of advisory by 45% in medium scenarios and 60% in heavy. Observed gate-holds were less than 15 minutes in all but one scenario, and even in this scenario 95% of the aircraft had a gate-hold of less than 15 minutes. Arrival delay was unaffected by the use of advisory. Total fuel consumption was also reduced by 23% in medium traffic and 33% in heavy traffic. TMI compliance appeared unaffected by the advisory.

INTRODUCTION

Surface operations at airports in the US National Airspace System (NAS) are often based on tactical operations, with aircraft often being controlled reactively. Although there are variations in procedures at different airports, the basic procedures remain the same: departure aircraft are moved from gate to runway whereas arrival aircraft are moved from touchdown to the gate. Numerous intermediate steps are involved in these processes with some degree of connectivity in the steps. Because of the mostly reactive nature of air traffic control (ATC) at the airport and the lack of adequate decision support tools for surface management, however, surface operations can be inefficient. One cause of inefficiency is peak traffic: multiple aircraft pushback at around the same time and contest for the limited resource (the runway). This situation leads to many aircraft taxiing to the runway simultaneously causing long runway queues as well as adverse congestion effects on taxiways. Departure metering (holding the aircraft in a holding area, probably in the ramp) as a method to address this inefficiency has recently received attention in the literature. A recent Federal Aviation Administration (FAA) sponsored study estimated that at eight major US airports, departure metering could yield cumulative fuel savings of 2.3 billion USD from 2012 to 2030 [A. Nakahara and T.G. Reynolds, 2012]. Another study estimates potential fuel savings of 42 to 300 million USD in 2011 at 43 top US airports [H. Daniel and M. Tim, 2012].

Given the potential benefits of departure metering, methods to implement it at US airports are being explored by federal organizations (the Federal Aviation Administration [FAA] and the National Aeronautics and Space Administration [NASA]), academia, and the aviation industry [C. Brinton et al., 2011; Jung et al., 2011; A. Nakahara et al.; I. Simaiakis et al., 2011]. These concepts can be divided into two broad categories: air traffic control tower (ATCT) based and airline based. ATCT-based concepts provide "advisories" to ATCT controllers at the airport, specifically to the ground controller (broadly responsible for taxiway movements) and the local controller (broadly responsible for departure and arrival aircraft on

runways). One such concept is based on controlling the rate of aircraft being released in to the taxiways by the ATCT controllers [I. Simaiakis et al., 2011], as well as the number of aircraft in runway queues. In this rate control concept, rate advisories are provided to the ground controller only. Another concept, Spot And Runway Departure Advisor (SARDA), gives aircraft specific sequence and time advisories to ATCT to reduce the number of aircraft on the taxiways and runway queues [Y. Jung et al., 2011]. In the SARDA concept, advisories are provided to both ground and local controllers. Both the rate control and SARDA concepts are based on moving the delay from the taxiways and runway queues to the ramp area by providing guidance to the ATCT, and do not cause a reduction in overall delay. Further, both are "tactical" tools that respond to the current traffic scenarios with little strategic planning.

Another set of concepts reduces the number of aircraft in taxiways and runway queues by altering the pushback times for departure aircraft; these alterations are done in collaboration with the airlines. One such method, called Collaborative Departure Queue Management (CDQM), manages the length of the runway departure queues by giving the flight operator an allocation of slots to enter the taxiways [C. Brinton et al.]. The assignment of the slots is done through a "ration by schedule" approach, which effectively is a first-scheduled-first-serve approach. Another method with some similarities to CDQM was developed and implemented at the John F. Kennedy Airport in New York, and is currently under use [A. Nakahara et al., 2011]. The slot allocation is conducted using ration by schedule, and slots are assigned two hours in advance. Any swaps or changes in this allocation are managed by the "slot allocation manager," a neutral third party. The idea behind both these approaches was to hold aircraft at the gate or preassigned holding pads with engines off as much as possible, reducing the delays on the taxiway as well as the fuel consumption and emissions. Compared to the above ATCT advisory-based tools, these airline collaborative decision making (CDM) tools are strategic in nature.

Tools for departure management have been studied in European airports before the studies began in the US. The EUROCONTROL and DLR (German Aerospace Center) DMAN (Departure MANager) tool [D. Böhme, 2005] does tactical departure management through a heuristics-based optimization scheme. The tool provides a "managed time" for pushback, and initial results showed no change in runway throughput while reducing taxi-out times. In Böhme [D. Böhme, 2005], however, it is emphasized that DMAN does not use surveillance information and, hence, could provide target runway sequences that are not achievable because of interactions between aircraft on the taxiways. These interactions have been the focus of recent research [R. Dean et al. 2009, I. Gerdes and A. Temme, 2012] where the DMAN tool is

being integrated with SMAN (Surface MANager); the position of each aircraft is taken into account (through surveillance data) and the runway schedule is updated. This frequent updating becomes even more important at US airports, where "with the exception of taxi-in times, variability in times for all flight phases is higher" as compared to airports in Europe [J. Gulding et al., 2009].

Recently, a new version of the SARDA system was introduced [G. Gupta et al., 2012] that combined the tactical controller advisories with collaborative gate-hold to develop an integrated system for airport surface management. As before, this system is *aircraft specific*. The SARDA system is capable of collaborative gate-holding either strategically (where pushback times are provided an hour before scheduled pushback time) or tactically (where pushback hold is recommended, if necessary, only when the aircraft is ready to pushback). In Gupta's paper [G. Gupta et al., 2012], the tactical gate-hold version of the new SARDA system was tested in an automated simulation setup with varying levels of uncertainty in gate pushback.

In this paper, the results from a human-in-the-loop (HITL) evaluation of the tactical gate-hold version of the new SARDA system are presented. The HITL study was conducted at the NASA Ames Research Center in May 2012 with recently retired controllers and pseudo pilots in a 360° simulated tower environment. Numerous improvements were made in both the SARDA concept and experiment setup as compared to the previous tests [Y. Jung et al., 2011]:

- Holding at gate was added as compared to the holding at spot. This means that most taxi delays from the runway queue were shifted all the way to the gates. In this experiment, ramp area complexities caused by uncertainties associated with flight ready time, gate conflicts, aircraft being blocked by the aircraft that had pushed back from next gate were not modeled.
- The SARDA scheduler was updated to a unified single scheduler providing gate-hold, ground controller, and local controller advisories.
- Taxi movements were more realistic because of taxi speed uncertainty.
- An electronic flight strip (EFS) based interface was designed for the ATC controllers, and SARDA advisories were incorporated in EFS.
- The effect of traffic management initiatives (TMI) on certain aircraft was evaluated. TMI is used here as a generic term for extra timing constraints on certain aircraft because of weather or other phenomenon, and is described further later in this report.
- A complete out-of-the-window view was provided to ATC controllers in the simulated tower environment.

A variety of data were collected in this study, and this paper presents some of the system performance data from the experiments. First the tactical gate-hold SARDA concept is described, followed by a description of the experimental setup, system performance results,

and a short summary of effect on controller workload. The paper concludes with a discussion and directions for future work.

SARDA CONCEPT DESCRIPTION

Previous work on SARDA focused on developing ground and local controller advisories for improved surface traffic management [Y. Jung et al., 2011]. The new version of SARDA under development is the augmentation of ground and local controller advisories through sharing flight movement and related operations information among airport operators, flight operators and ATC. The goal is to improve airport surface operations by maximizing available airport and airspace capacity while minimizing adverse effects on stakeholders (e.g., excessive fuel costs for airlines), passengers, and the environment. In previous work, the taxi and runway queue delay for departure aircraft was moved to the transition point where ATC tower takes control of the aircraft from the ramp controller. Here, the delays are moved to the gate to provide benefits to the airlines in fuel savings and potentially better connection time for passengers. The concept includes sharing data between ATCT and airline operators (including updated gate pushback readiness and ATC constraints due to weather), which enables these benefits.

Some of the key assumptions for the SARDA concept include:

- Ramp area operations are managed by airlines or airport authorities and, therefore, ATCT does not have direct control of gate pushback for departure aircraft.
- Voice is still the main means of communication between ATCT controllers and pilots. Assuming that a data exchange mechanism exists between ATC and ramp control personnel, traffic control decisions, such as pushback sequence, can be made collaboratively.
- Aircraft position data may not be available in the ramp area through a surface surveillance system, such as Airport Surveillance Detection Equipment–Model X (ASDE-X). Actual gate pushback time for departure aircraft is known, however. Aircraft position data at the spots and the movement area is available through ASDE-X, with the assumption that the current accuracy levels in ASDE-X is sufficient.
- Prediction of arrival times of departure aircraft at runway queue entrance is available. In the current implementation, predictions based on nominal unimpeded travel from scheduled gate pushback to the runway are used. This could be changed in the future, and better prediction models can be used.
- For the cases where arrival aircraft cross the departure runway, prediction of earliest runway crossing time is available. These data

can be inputs from a system like Traffic Management Advisor (TMA) [H. Swensen et al., 1997]. Swensen et al. [H. Swensen et al., 2011] discuss the extension of TMA in the terminal area and performance of the controller decision support tool that provides precision spacing and scheduling of arrival flights to the runway with enhanced accuracy.

The version of SARDA being discussed here includes tactical gate-hold. For strategic gate-hold, a different version of SARDA has been described by Gupta [G. Gupta et al., 2012].

SARDA Scheduler

The core computational engine behind SARDA, called the SARDA scheduler, is based on the Spot Release Planner (SRP) [W. Malik et al., 2012], a method to provide metering advisories. SRP is a two-stage algorithm: The first stage is a runway scheduler [G. Gupta et al., 2010; J. Montoya et al., 2011], which gives the best sequence and times for runway usage by a set of departure aircraft ready for take-off and arrival aircraft waiting to cross the same departure runway. The second stage of the SRP determines times to release aircraft from gates or assigned spots to meet the optimal departure schedules [W. Malik et al., 2012]. The implementation of the SARDA scheduler is airport specific. Some airports may have multiple runways with some degree of interdependence in operations; some airports might have multiple runways that can operate independently; others might have separate arrival and departure runways. The runway scheduler in the first stage of SRP is tailored to the runway layout as well as the runway configuration. The second stage of the SRP is tailored to the airports' taxiway layout. Where applicable, the second stage of SRP can also provide taxi route advisories. A nominal taxi time from gate and spot was used to calculate gate pushback time.

The scheduler takes as input the current snapshot of the airport, aircraft-specific parameters, separation constraints, scheduled pushback times, and scheduled arrival times for the aircraft in the next 15 minutes. The choice of 15 minutes for planning horizon was made based on computational demand (i.e., because of the size of optimization problem to solve) and effectiveness of solution because of uncertainties in future traffic situation. The most efficient choice for this parameter value, however, is an open question. Uncertainties in aircraft movement pose a challenge to generating advisories. To mitigate the effect of uncertainties, the scheduler gets an updated airport condition snapshot every 10 seconds, which is then used to recalculate the schedule. Thus, the scheduler follows an update cycle of 10 seconds. The choice of 10 seconds for schedule update frequency

was made aligned with the model update cycle of the decision support tool's software platform. To reduce the frequent changes in advisories for the controller because of the frequent updates, a "freeze horizon" is implemented: the schedule for certain aircraft does not change based on the location of aircraft and the current time. For example: in the local controller advisory, the sequence for the first three aircraft scheduled to use the runway is fixed between successive scheduler calls. The scheduler does not provide intermediate advisories for those aircraft while they are taxiing. The choice of freeze horizon parameter was made mainly based on controllers' feedback.

As previously mentioned, the SRP's first stage is a runway scheduler that provides the best runway usage schedule for both arrivals (crossings) and departures. The runway scheduling problem has numerous constraints (wake vortex separation, miles-in-trail and others) and can be solved for multiple objectives including throughput (runway usage time for last aircraft) and system delay (total delay for all aircraft). G. Gupta et al. [G. Gupta et al., 2010] show that optimizing for system delay results in small deviations from optimal throughput, whereas optimizing for throughput results in large deviations in system delay. For this reason, system delay was chosen as the objective for the scheduler. Further, taxi time estimates were required to predict the earliest take-off times for the first stage of SRP, and for calculating spot and gate releases in the second stage. These predictions were based on the historical aircraft movement on the surface: numerous unimpeded trajectories on surface were analyzed to estimate the distribution of unimpeded speeds. The speed corresponding to 30th percentile was chosen for the estimates in the first stage of SRP, and for the second stage as well. This particular percentile value was derived through analysis of previous simulations and was chosen to provide sufficient buffer in the runway queue while ensuring that the runway is not underutilized.

SARDA Tactical Gate-Hold Walkthrough

The SARDA tactical gate-hold concept is illustrated using an example of flight XYZ101 with scheduled pushback time of 10:00am.

10:00am   Scheduled pushback time for the aircraft. Because of ongoing boarding or baggage loading, the pilot is not ready to pushback.

10:03am   Pilot communicates pushback readiness. Ramp controller inputs readiness in the ramp controller interface. Through collaboration with ATC, this interface reflects any TMI constraints that aircraft might have. Based on current traffic, SARDA modifies actual pushback to 10:12am. Ramp controller accepts this and communicates this to the pilot.

10:12am	Ramp controller interface prompts for XYZ101 pushback. Ramp controller communicates this to the pilot. Pilot accepts and commences pushback.
10:14am	Pilot reaches the spot to transition to the active movement area. Ground controller is aware of XYZ101, but SARDA advises spot release close to 10:15am because of taxiing traffic next to the spot.
10:15am	Ground controller initiates communication with XYZ101 and clears it to taxi using one of the standard taxi routes. Pilot acknowledges, and commences taxiing.
10:18am	Pilot reaches departure queue. Local controller is aware of XYZ101. SARDA advises aircraft is 3^{rd} in the departure queue. Local Controller can communicate this to the pilot at his discretion.
10:20am	Local controller clears XYZ101 to proceed to the runway, line up and wait.
10:21am	Local controller clears XYZ101 for take-off. Pilot acknowledges and begins take-off procedure.

EXPERIMENTAL SETUP

This section describes the setup for the HITL experiments to evaluate the tactical gate-hold version of SARDA. The experiments were conducted at the NASA Ames Research Center in the Future Flight Central (FFC) facility, which offers a 360-degree, full-scale, real-time simulation of an airport control tower. Multiple software systems were integrated at FFC to enable pseudo-pilot control of the aircraft, ATCT controller displays and computer generated imagery (CGI) based out-of-the-window view from the ATC tower. In the following subsections, the system architecture, controller user interface, and experiment scenarios are discussed.

An underlying assumption for these experiments was 100% airline compliance with SARDA-generated gate pushback times. We assumed that pilots would call in pushback readiness at scheduled pushback times; SARDA would then generate pushback times to which pilots would comply. The effects of various levels of airline compliance with SARDA pushback times has been studied in an automated setup [G. Gupta et al., 2012], but because of certain system limitations, as well as ongoing development of the ramp controller interface, uncertainty in meeting SARDA-generated pushback times was not tested.

System Architecture

The tactical gate-hold version of SARDA was implemented within the Surface Management System (SMS) [Raytheon, 2004]. SMS was

originally developed as a decision support tool to assist ATCT controllers and managers as well as airline operators in managing and controlling airport surface operations [S. Atkins et al., 2004]. For this simulation, SMS exchanged flight information and scheduling solutions with the optimization algorithms over the network. Existing SMS user interfaces were modified to provide advisories to the ground and local controller positions.

The Airspace Traffic Generator (ATG) system was used to generate motions of aircraft either on the surface or in the airspace near the airport and to send position data to SMS for display [SAIC, 2006]. ATG is a high-fidelity, real-time aircraft simulation tool that can move the aircraft on the airport surface and generate and display targets of the aircraft. The ground pilot stations (GPS), components of ATG, were used by the pseudo-pilots to taxi aircraft manually when clearances were issued by the controllers over voice.

Experiment Scenarios

The experiments were conducted using the east side of the Dallas/Fort Worth International Airport (DFW), with FFC providing the view from the east DFW ATC tower. Figure 1 shows the map view of the east side of DFW with various aspects marked.

Four traffic scenarios were generated based on January 2012 DFW traffic; two of these were called medium traffic (1.2 x current traffic), and two were called heavy traffic (1.5 x current traffic). These scenarios were labeled Medium 1, Medium 2, Heavy 3 and Heavy 4. The medium scenarios had 40 departure aircraft in 50 minutes, and heavy scenarios had 50 departure aircraft in


Figure 1. East side of DFW showing departure runway (17R), arrival runway (17C), arrival exits, and departure taxi routes. (Only a limited number of spots are shown.)

Figure 2. Cumulative departure aircraft counts for the four experiment scenarios.

50 minutes. Figure 2 shows the cumulative departure counts for the four scenarios at 2-minute intervals. These were plotted assuming unimpeded travel from scheduled pushback. It shows that Heavy 3 has 2 peaks between the 10th and 30th minute, whereas Heavy 4 has only one large peak starting at the 25th minute. The traffic scenarios were generated by augmenting the current day traffic density of DFW because of the surface traffic in DFW that occurred in clear weather days did not experience significant congestion. This situation has existed since American Airlines, the dominant carrier operating in DFW, decided to depeak at the airport in 2003 [T. Reed, 2009].

Scenarios were generated such that departure traffic begins at the gates upon activation in the simulation and arrival aircraft appear about 10 nmi from the runway threshold. After a departure aircraft pushes back from the gate, it maneuvers in an automated mode towards its assigned spot and stops before it, unless the ground controller issues the pilot a taxi clearance while the aircraft is still moving. Scenario data for a departure aircraft contain callsign, aircraft type, flight plan route, departure fix, activation time (pushback or first track hit), initial position, gate, spot, and runway. Uncertainty in taxiway movement was included, with each aircraft having a nominal taxi speed from 12 to 17 knots. Of course, the aircraft slowed down when instructed by the pseudo pilot e.g., when an aircraft in front is taxiing more slowly; the above nominal speed is the default taxi speed for the aircraft unless instructed otherwise by the pseudo pilot. For repeatability, the same aircraft was given the same nominal speed in multiple runs of the same scenario.

To test the effect of the SARDA advisories, a non-SARDA or baseline case was also tested. In this case, the SARDA scheduler was not running, and advisories were absent with no gate-holding. This situation essentially represents current day operations. Since the baseline scenarios were generated by augmenting the current day traffic (by 20-50%), direct validation with real operations was not feasible. Instead, the real-time HITL simulation was validated by human subjects (i.e., retired controllers from DFW tower) in a subjective manner by assessing the realism of simulated traffic and operational procedures. To determine the performance envelope of the concept, controllers were asked to follow the advisories, when present, as long as it did not compromise aircraft safety.

Six recently retired controllers participated in the experiments over a three-week period, with each pair running scenarios for an entire week. Controllers rotated between the ground controller and local controller positions, with a controller running each scenario twice at the same position, once in baseline and once in SARDA or advisory case. This results in a total of 48 runs, with six runs of each scenario in either baseline or advisory mode.

Controller Interface

Displays for both ground and local controllers were provided for the HITL simulation. The basic display for the ground controller is composed of an existing SMS map display that shows spots and taxiways under the controller's responsibility. The basic display for the local controller consists of a surface map of the responsible area (i.e., runway queue and crossing queues) and a map of terminal airspace that covers portions of final approach and initial climb paths. Instead of using paper strips, however, an electronic flight strip (EFS) system was developed for both the advisory and baseline case. The EFS system is touch screen enabled, and mimicked certain functionalities of the paper strips. In the advisory case, the controller advisories were integrated in the EFS system and displayed to the controller. The ground controller was shown the spot release sequence and window; the local controller was shown the sequence of runway 17R operations including departure takeoffs and arrival-crossings. Further, advisories on departure taxi route were also generated and presented to the controller within the EFS system. Taxiing and airborne aircraft are shown on the map displays with a data tag attached to the aircraft icon. Figure 3 shows a sample picture of the ground controller display with both the map and EFS. A detailed description of the controller

Figure 3. Map (left) and EFS (right) displays for the ground controller.

interface can be found in work done by Hayashi, et al. [M. Hayashi et al., 2013].

Both the ground and local controllers working with SARDA displays were responsible for safe separation of traffic in the movement area and were interacting with pseudo-pilots via a voice communication. Traffic in the ramp area, from pushback to arrival at designated spots, however, was managed by automation.

RESULTS

This section details the results observed during the HITL experiments. The results are split into runway usage, delays, fuel, and the effect on TMI aircraft. The section concludes with a brief note on controller workload.

Runway Usage

Runway usage is defined as the number of runway operations in a given time period. Operations would include departure take-off or arrival-crossing on runway 17R. With the SARDA concept enabling departure metering at the gates, it is important to compare the runway usage to detect cases of excessive metering. Figure 4 shows the cumulative runway usage for the four scenarios at five-minute intervals. Each bar represents the mean over the six runs of the scenario in either baseline or advisory case, and the whiskers represent the max and min over the six data points.

The results show almost no difference in throughput between the advisory and baseline runs of the same scenario. Hence, no

Figure 4. Throughput comparison in advisory and baseline runs for all scenarios (bar represents the mean; whiskers represent the min and max).

loss in runway usage was observed for the SARDA runs with departure metering.

Runway Queue Size

Since SARDA is enabling departure metering by holding the aircraft at the gate, it is expected that the number of aircraft in the runway queue would be reduced. The east side runway queue area for DFW airport is shown in Figure 1. The number of aircraft in the runway queue (at one-minute intervals) for both baseline and SARDA runs are compared in Figure 5. Both the average across six runs and the min and max values are plotted. Results show that although the baseline queue size varies widely and depends on the scenario, the use of an advisory makes the queue size insensitive to the scenario or traffic load, and the queue size never exceeds seven aircraft.

Delays

Delay is defined as the difference between the observed travel time and the unimpeded travel time for the same route or travel section. Since the nominal speed of the aircraft could vary from 12 to 17 knots (see section titled Experiment Scenarios), an unimpeded speed of 17 knots is used for calculating unimpeded travel times to avoid negative delay values. First, the delay in departure aircraft is considered and then the arrival delay.
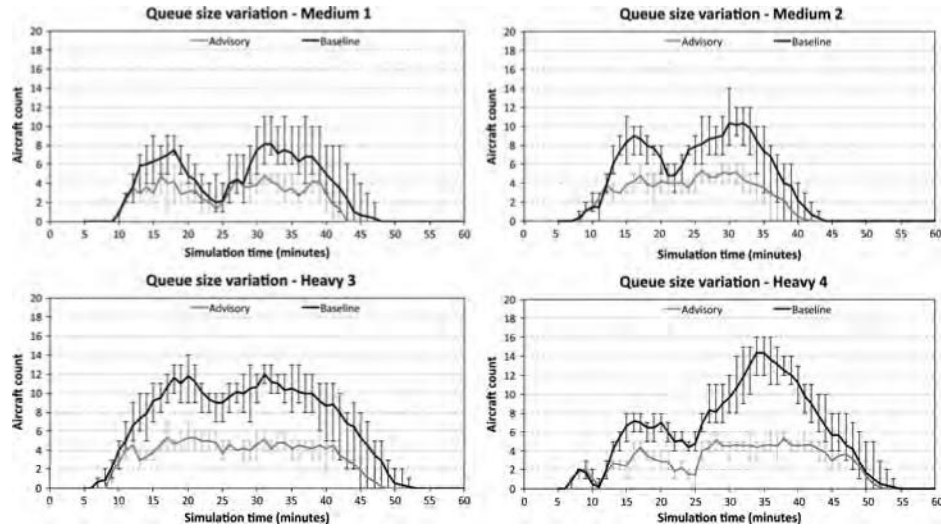
Figure 5. Runway queue-size comparison between advisory and baseline for all scenarios (mean, min, and max values are plotted across all runs at one-minute intervals).

Departure Aircraft. For departure aircraft, first the scheduled delay is considered. Scheduled delay is defined as the observed take-off time minus the unimpeded take-off time calculated from the scheduled pushback time. Scheduled delay is the overall delay experienced by the aircraft at gate, ramp, taxiways, and runway queue. Figure 6 shows the box and whisker plots for both advisory and baseline runs for the four scenarios. As expected, scheduled delay increases with increasing traffic volume. It appears that scheduled


Figure 6. Scheduled departure delay. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25th and 75th percentile; whiskers show 10th and 90th percentile; dots show min and max; diamonds show mean.)

Departure taxiing delay

Figure 7. Taxiing delay for departures. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25^{th} and 75^{th} percentile; whiskers show 10^{th} and 90^{th} percentile; dots show min and max; diamonds show mean.)

delay is slightly less with the use of advisory; this effect is further discussed in subsection 4.7.

One of the prime motivations for departure metering is to move the delay from the taxiways and runway queues to the gate, to enable less congestion in the active movement area and fuel savings. To gauge this effect, taxiing delay is evaluated and is defined as the difference in observed and unimpeded take-off times, when the unimpeded take-off has been calculated from the actual pushback time. Therefore, taxiing delay calculation represents the delay in the ramp, taxiways, and runway queues. For baseline runs, this is equivalent to the scheduled delay, but is less for the advisory runs. Figure 7 presents the taxiing delay results. As expected, taxiing delay is noticeably reduced with the use of advisory: an average 45% (3 min per aircraft) reduction in medium traffic and 60% (5.5 min per aircraft) reduction in heavy traffic were observed. Further, the variation in taxiing delay increases with increasing traffic in the baseline cases, but is insensitive in the advisory case.

The current FAA Department of Transportation (DOT) on-time performance metrics[1] would preclude gate-holds of more than 15 minutes beyond the scheduled pushback time, since these would be reported as airlines delay. Although a policy change can be implemented in the future after discussions with various stakeholders, an immediate solution would be to limit gate-hold. The SARDA system is capable of

---

[1] Code of Federal Regulations (CFR) Title 14-Aeronautics and Space, Chapter II, Part 234 – Airline Service Quality Performance Reports

Figure 8. Gate-hold for departures. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25^{th} and 75^{th} percentile; whiskers show 10^{th} and 90^{th} percentile; dots show min and max; diamonds show mean.)

limiting gate-hold, but this functionality was not implemented in the experiments. Thus, it becomes important to investigate the gate-hold on each aircraft. Figure 8 shows the gate-hold values across all the experimental runs. It should be noted that even without limiting the amount of gate-hold, in all scenarios except Heavy 4, gate-hold was never more than 15 minutes. Even in Heavy 4, 95% of the aircraft had a gate-hold of less than 15 minutes.

Arrival Aircraft. The delay in arrival aircraft is investigated to check if the use of an advisory has an effect on arrival aircraft. Since arrival-crossing advisories are being provided to the local controller, one might reasonably be concerned as to whether the improvements in departure taxiing delays are being achieved by increasing arrival delays. Figure 9 shows the delay in arrival aircraft across all runs, including delays in waiting to cross as well as taxiing to the spot after crossing. Arrival taxiing in the ramp is not considered here, since this was handled automatically by the ATG software. Results show that the delay in arrival aircraft were insensitive to the use of advisory or traffic level.

Fuel Consumption and Emissions by Departure Aircraft

The amount of fuel consumed by departure aircraft was evaluated using the method described by T. Nikoleris et al. [Nikoleris et al., 2011]. This method is an augmentation of the current ICAO emissions databank, and includes the effects of stops and acceleration events. For departure aircraft, fuel consumption was evaluated

Arrival delay (crossing and taxiing)

Figure 9. Arrival delay. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25^{th} and 75^{th} percentile; whiskers show 10^{th} and 90^{th} percentile; dots show min and max; diamonds show mean.)

assuming all-engine taxiing from actual pushback. It was assumed that if the aircraft is being held at the gate, the engines are off. Figure 10 shows the total fuel used by departure aircraft over all the experimental runs. Fuel consumption is smaller in the advisory case as compared to the baseline: 23% average reduction in medium traffic and 33% average reduction in heavy traffic was observed. Moreover, fuel consumption in the baseline case seems more sensitive to traffic level as compared to advisory.

Total fuel used by departure aircraft

Figure 10. Total fuel used by departure aircraft. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25^{th} and 75^{th} percentile; whiskers show 10^{th} and 90^{th} percentile; dots show min and max; diamonds show mean.)

Total CO emissions by departure aircraft

Figure 11. Total CO emissions by departure aircraft. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25^{th} and 75^{th} percentile; whiskers show 10^{th} and 90^{th} percentile; dots show min and max; diamonds show mean.)

Along with fuel consumption, emissions from departure aircraft were also evaluated. Emission estimates are generated using emission indices, which are the amount of emissions generated per kilogram of fuel at certain engine thrust levels [T. Nikoleris et al., 2011]. Hence, the trends in emissions are similar to those in fuel consumption. Because of restricted space, only carbon monoxide (CO) emissions are provided here; hydrocarbon (HC) and nitrogen oxides (NOx) show similar trends. Figure 11 shows CO emissions calculated for all departure aircraft in the experimental runs. As in fuel consumption, CO emissions are lower for the advisory case with a 24% average reduction in medium traffic and 34% average reduction in heavy traffic.

Effect on TMI Aircraft

As discussed earlier, certain aircraft in each scenario had TMIs associated with them; these were the desired take-off times given to a subset of departure aircraft. Five aircraft had TMIs in the medium traffic scenarios, and seven aircraft had TMIs in the heavy scenarios. The window specified to the controllers for TMI adherence was ± 60 seconds. If the controllers were somehow unable to meet the TMI time, they were asked to get the take-off close to the TMI time, and no new TMI was generated.

In this section, the controllers' adherence to the prescribed TMI times is explored. For each aircraft, the observed take-off time was compared to the prescribed time. The observed take-off time is defined

Figure 12. Difference in take-off and TMI time. (Horizontal axis "a" and "b" represent advisory and baseline. Bars show median, 25th and 75th percentile; whiskers show 10th and 90th percentile; dots show min and max; diamonds show mean.)

as the time the departure aircraft crosses a threshold on the runway, close to taxiway Z. There is a small delay between the controller command for take-off and the time the aircraft crosses this threshold; this could be because of variations in pilot response. Thus, when considering adherence to TMI time, a strict ± 60 second window would be incorrect. For this purpose, the actual difference in take-off and prescribed time is considered. Figure 12 shows the plot of this difference for the various runs.

In almost all cases, 90% of the TMI aircraft appear to meet the prescribed window. Besides the few outliers in the heavy scenario, there does not seem to be any evidence that TMI compliance was affected by the use of the advisory. It was expected, however, that the primary effect of the advisory on TMI aircraft would be a reduction in delay. Possibly, to reduce the chance of missing a TMI time, controllers might get the TMI aircraft to the runway sooner and then make it wait in the departure queue. Using the advisory, it was expected that such cases would decrease. This is explored further in the next section.

Statistical Significance of Taxiing Delay Reduction

Figure 7 shows that the use of advisory decreases the taxiing delay, and the effect is more prominent for the heavy traffic cases. To parse

and understand such effects, tests of statistical significance must be done.

The first step in testing the statistical significance of the benefits of advisory and the effect of traffic level on taxiing delay would be a t-test. A requirement for the use of t-test, however, is the independence within the sample, i.e., there should be no dependence within the dataset. For taxiing delay it is highly probable that part of the delay from the leading aircraft in the runway take-off sequence might propagate to the trailing aircraft. Thus, the delay of the n^{th} aircraft in the runway take-off sequence might depend on the delay of the $(n\text{-}1)^{th}$ aircraft. To address this effect while conducting tests for significance, linear regression is used with the delay of the previous departure aircraft as a variable within regression.

Using the delay of the previous aircraft as an independent variable is a case of the lagged variable model where disturbances are serially correlated [S.P. Washington et al., 2009]. Ordinary least squares regression would not give unbiased estimates of the coefficients in this case. To address this, a first order autoregressive process was assumed and generalized least squares were used, with estimates evaluated using maximum likelihood estimation.

Besides the use of advisory and the traffic level, the effect of aircraft weight class and an aircraft having a TMI was also tested. Further, it is possible that a longer taxi distance will lead to larger delays, since it increases the possibility of encountering other aircraft while taxiing. For this reason, total taxi distance was also included as an explanatory variable. Below is a list of all the variables used in the regression model, and the description of the variables. Along with these variables, the interaction terms with the indicator variables are also explored.

A   Advisory or baseline, 1 if advisory, 0 otherwise
H   Traffic scenario, 1 if heavy, 0 otherwise
W   Aircraft weight class, 0 if large, 1 otherwise
TMI   1 if aircraft had a TMI
L   Distance traveled in ramp, taxiways, and queues
$D_{n\text{-}1}$ Taxiing delay for the immediate predecessor in runway sequence (departure aircraft only)

Table 1 gives the results from the linear regression. The estimate, standard error, t-value, and significance level are presented; rows in white are significant to the 5% level, rows in grey are not. Following are some of the observations from the regression:

• The effect of the advisory is significant, and the use of SARDA advisory decreases taxiing delay in all cases.

- Delays are higher in heavy traffic; however, the use of advisory reduces delay in heavy traffic also. The magnitude of the interaction term A×H is close to that of H; this suggests that the use of advisory substantially reduces the increased delay effect of heavy traffic.
- TMI aircraft typically have a higher delay than non-TMI aircraft, and this effect is significant. The use of the advisory, however, seems to reduce this effect (A×TMI), and the reduction is statistically significant. Further, heavy traffic scenarios further increase the delay in TMI aircraft. A potential reason for this is the presence of three runway queues at DFW. Controllers use one of the queues ("full length") to stage TMI aircraft, a procedure that allows for a TMI departure at anytime but increases the possibility of more TMI delays in runway queue.
- Aircraft in the heavyweight class seem to get larger delays in both advisory and baseline case, although the effect is of borderline significance.
- Delay does not depend on the taxi distance in the baseline case, but there is some correlation in the advisory case (A×L).
- The delay experienced by the previous aircraft in runway sequence significantly affects the delay of the successor. This is true for both baseline and advisory, but magnitude of the effect is reduced in the advisory case (A×D_{n-1}). The effect seems to reduce in the heavy scenarios, but this is not statistically significant.

Table 1. Results from Regression on Taxiing Delay

Variable	Coefficient Estimate	Standard Error	t-value	Pr(>|t|)
Intercept	154.70	30.44	5.08	0.00
A	−159.82	31.49	−5.08	0.00
H	104.58	32.88	3.18	0.00
W	86.69	46.39	1.87	0.06
TMI	309.43	53.19	5.82	0.00
L	−0.01	0.01	−0.51	0.61
D_{n-1}	0.65	0.04	16.58	0.00
A×H	−74.93	16.07	−4.66	0.00
A×W	31.72	24.62	1.29	0.20
A×TMI	−160.09	24.50	−6.53	0.00
A×D_{n-1}	−0.23	0.07	−3.27	0.00
A×L	0.06	0.01	5.42	0.00
H×W	−5.99	22.73	−0.26	0.79
H×TMI	84.59	21.88	3.87	0.00
H×L	−0.01	0.01	−0.81	0.42
H×D_{n-1}	−0.07	0.04	−1.51	0.13
W×L	−0.03	0.01	−1.76	0.08
W×D_{n-1}	−0.23	0.05	−4.45	0.00
TMI×L	−0.06	0.02	−3.55	0.00
TMI×D_{n-1}	−0.30	0.06	−5.25	0.00

Statistical Significance of Scheduled Delay Reduction

Figure 6 hints at some reduction in scheduled delay. To explore this further, a similar linear regression treatment of scheduled delay was conducted, as done in the previous section. All the same indicator and continuous variables are used, with the exception that instead of D_{n-1} (taxiing delay of previous departure), D'_{n-1} was used, which represents the scheduled delay of the last departure aircraft based on runway sequence.

Table 2 presents the results of the regression. Results show that advisory indeed reduces scheduled delay, and the effect is significant. This can be attributed to the use of system delay as the objective in the SARDA scheduler, as discussed in the section, SARDA Scheduler. The effect of advisory in heavy traffic cases, however, is not significant. Another interesting observation is that TMI aircraft have higher scheduled delays; this can be explained from the facts that 1) the aircraft have to wait for the desired take-off time and 2) the scenarios themselves were designed such that TMI times could be achieved, artificially inducing delay.

Controller Workload

While investigating decision support tools for air traffic control, it is necessary to examine the effect of the tool on the controller in terms of usability, workload, and acceptance. Data were collected from the

Table 2. Results from Regression on Scheduled Delay

Variable	Coefficient Estimate	Standard Error	t-value	Pr($>$|t|)
Intercept	155.60	38.06	4.09	0.00
A	−85.79	37.84	−2.27	0.02
H	72.90	40.90	1.78	0.07
W	74.52	58.22	1.28	0.20
TMI	384.65	67.73	5.68	0.00
L	0.01	0.01	0.45	0.65
D'_{n-1}	0.58	0.04	13.05	0.00
A×H	8.12	18.66	0.43	0.66
A×W	184.54	27.81	6.63	0.00
A×TMI	56.16	27.64	2.03	0.04
A×D'_{n-1}	−0.11	0.04	−2.79	0.01
A×L	0.02	0.01	1.68	0.09
H×W	−9.89	30.03	−0.33	0.74
H×TMI	183.87	28.97	6.35	0.00
H×L	−0.02	0.01	−1.06	0.29
H×D'_{n-1}	0.02	0.05	0.39	0.70
W×L	−0.03	0.02	−1.38	0.17
W×D'_{n-1}	−0.20	0.06	−3.65	0.00
TMI×L	−0.11	0.02	−5.22	0.00
TMI×D'_{n-1}	−0.27	0.06	−4.44	0.00

participating controllers for this purpose. For the efficiency aspect, real-time workload ratings were taken every five minutes; NASA task load index (TLX) workload ratings [Hart and Staveland, 1988] were collected at the end of each run, and other post-run questionnaire responses were analyzed to assess how much internal resources (e.g., spare capacity for workload, spare attention) the controllers felt they had during each run. For the acceptance aspect, the controllers' responses to the post-run and post-study questionnaire regarding their subjective judgment on the helpfulness of the SARDA advisories and ease of use of the user interface were examined.

Detailed results from this analysis can be found in [M. Hayashi et al., 2013]. In summary, the results do not show any increase in workload with the use of advisory. In fact, the NASA TLX workload rating results exhibited clear reductions of workload levels in terms of temporal demand (time pressure), effort (how hard controllers had to work physically and mentally), physical demand (e.g., using EFS, communicating on the radio), and mental d (e.g., thinking, deciding, calculating, remembering, looking). In all four ratings, the magnitude of the mean-score reductions from the baseline runs to the advisory runs was approximately 2 points, which may have been large enough to be sensed by the controllers.

CONCLUSIONS

The 360° simulated tower HITL experiments of tactical gate-hold using SARDA showed promising results. Taxiing delay reduction of 60% was observed in heavy traffic scenarios. Estimated fuel reduction of 33% was observed in the heavy scenarios, with emissions showing a similar trend. Reductions in delay and fuel were also observed for the medium traffic level. Even though the amount of gate-hold for every aircraft was not limited, less than 5% aircraft in one scenario had a gate-hold of more than 15 minutes. Arrival aircraft delay remained unaffected, and the runway throughput with the use of advisory seems the same as the baseline case. TMI compliance was not much different between the advisory and the baseline case. The taxiing delay reductions because of the use of advisory are statistically significant. Some evidence suggests that the use of advisory decreases the overall scheduled delay in the system, although further tests are needed to evaluate this finding.

Although the results are promising, there are limitations in this study. The effect of pushback uncertainty was not studied. Although this has been studied in offline simulations, studying this in a HITL environment would be invaluable in transitioning the concept to any airport. Connected to this is the need for a ramp controller interface, which is currently under development at NASA Ames Research

Center. As of submission of this article, the HITL simulation has been adapted to the ramp operations of a major US airport and a series of simulations have been conducted. A variety of uncertainties, including flight ready time and pilot behavior in pushback operations, as well as uncertainties in aircraft maneuver in the ramp area due pilot-controller interactions were added. The scheduler has also been modified to address these uncertainties. Results from these high fidelity simulations and description of scheduler function will be reported in future publications. Of course, there is a need for evaluating this concept in a field trial, where many more sources of ramp movement uncertainty are present, including baggage and fuel carts, uncertain pushback paths, and others. Lastly, with increasing acceptance of departure metering as a method for reducing surface fuel consumption and emissions, policy debates and changes might be required to allow gate-holding without adversely affecting airline performance metrics.

ACKNOWLEDGMENTS

ACRONYMS

ASDE-X	Airport Surveillance Detection Equipment – Model X
ATC	air traffic control
ATCT	air traffic control tower
ATG	Airspace Traffic Generator
CDM	collaborative decision making
CDQM	Collaborative Departure Queue Management
CGI	computer generated imagery
CO	carbon monoxide
DFW	Dallas/Fort Worth International Airport
DLR	German Aerospace Center
DMAN	Departure MANager
DOT	Department of Transportation
EFS	electronic flight strip
FAA	Federal Aviation Administration
FFC	Future Flight Central

GPS   ground pilot stations
HC   hydrocarbon
HITL   human-in-the-loop
NAS   National Airspace System
NASA   National Aeronautics and Space Administration
NOx   nitrogen oxide
SARDA   Spot And Runway Departure Advisor
SMAN   Surface MANager
SMS   Surface Management System
SRP   Spot Release Planner
TLX   task load index
TMA   Traffic Management Advisor
TMI   traffic management initiatives

REFERENCES

Atkins, S. et al. (2004), "Surface management system field trial results," AIAA 4th Aviation Technology, Integration, and Operations Forum. Chicago, IL: American Institute of Aeronautics and Astronautics Inc.

Böhme, D. (2005), "Tactical departure management with the EUROCONTROL/DLR DMAN," 6th USA/Europe Air Traffic Management Research and Development Seminar, Baltimore, MD.

Brinton, C., Provan, C., Lent, S., Prevost, T., and Passmore, S. (2011), "Collaborative Departure Queue Management," 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), Berlin, Germany.

Daniel, H. and Tim, M. (2012), "Estimating the Achievable Benefits of Airport Surface Metering," 12th AIAA ATIO and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Indianapolis, IN.

Dean, R., Gotteland, J.B., and Durand, N. (2009), "Airport surface management and runways scheduling," 8th USA/Europe Air Traffic Management R&D Seminar, Napa, CA.

Gupta, G., Malik, W., and Jung, Y. (2012), "An Integrated Collaborative Decision Making and Tactical Advisory Concept for Airport Surface Operations Management," 12th AIAA ATIO and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Indianapolis, IN.

Gerdes, I. and Temme, A. (2012), "Taxi routing for aircraft: Creation and Controlling," The Second SESAR Innovation Days, Braunschweig, Germany.

Gulding, J. et al. (2009), "US/Europe comparison of ATM-related operational performance–An initial harmonized assessment by phase of flight," 8th USA/Europe Air Traffic Management Research and Development (ATM2009), Napa, CA.

Gupta, G., Malik, W., and Jung, Y. (2010), "Incorporating Active Runway Crossings in Airport Departure Scheduling," AIAA GNC Conference, Toronto, Canada.

Hart, S.G. and Staveland, L.E. (1988), "Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research," Human mental workload, **1**: p. 139-183.

Hayashi, M. et al. (2013), "Usability Evaluation of Spot and Runway Departure Advisor (SARDA) Concept in Dallas/Fort Worth Airport Tower Simulation," 10th USA/Europe Air Traffic Management Research and Development Seminar (ATM2013), Chicago, IL.

Jung, Y. et al. (2011), "Performance Evaluation of a Surface Traffic Management Tool for Dallas/Fort Worth International Airport," 9th USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany.

Malik, W., Gupta, G., and Jung, Y. (2012), "Spot Release Planner: Efficient Solution for Detailed Airport Surface Traffic Optimization," 12th AIAA ATIO and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Indianapolis, IN.

Montoya, J., Wood, Z., and Rathinam, S. (2011), "Runway Scheduling Using Generalized Dynamic Programming," AIAA GNC Conference, Portland, OR.

Nakahara, A. and Reynolds, T. G. (2012), "An Approach for Estimating Current and Future Benefits of Airport Surface Congestion Management Techniques," 12th AIAA ATIO and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2012. Indianapolis, IN.

Nakahara, A., Reynolds, T.G., White, T., Maccarone, C., and Dunsky, R. (2011), "Analysis of a Surface Congestion Management Technique at New York JFK Airport," 11th AIAA ATIO Conference, Virginia Beach, VA.

Nikoleris, T., Gupta, G., and Kistler, M. (2011), "Detailed estimation of fuel consumption and emissions during aircraft taxi operations at Dallas/Fort Worth International Airport," Transportation Research Part D: Transport and Environment.

Raytheon (2004), CTO-05 Surface Management System, CTOD 24 Final Report, R.A. Team, Editor, May.

Reed, T. (2009), American Strategy Flies Against the Flock. The Street.

SAIC (2006), Airspace Traffic Generator: User's Manual Supplement, Revision 2.6.

Simaiakis, I. et al. (2011), "Demonstration of Reduced Airport Congestion Through Pushback Rate Control," 9th USA/Europe Air Traffic Management Reserch and Development Seminar, Berlin, Germany.

Swensen, H., Thipphavong, J., Sadovsky, A., Chen, L., Sullivan, C., Martin, L. (2011), "Design and Evaluation of the Terminal Area Precision Scheduling and Spacing System," 9th USA/Europe Air Traffic Management Research and Development, 2011. Berlin, Germany.

Swensen, H., Hoang, T., Engelland, S., Vincent, D., Sanders, T., Sanford, B., and Heere, K. (1997), "Design and Evaluation of the Traffic Management Advisor at the Fort Worth Air Route Traffic Control Center," 1st USA/Europe Air Traffic Management Research and Development, Saclay, France.

Washington, S.P., Karlaftis, M.G., and Mannering, F.L. (2009), Statistical and economic methods for transportation data analysis, Chapman & Hall/CRC.

BIOGRAPHIES

Dr. Yoon Jung is an aerospace engineer at NASA Ames Research Center, Moffett Field, CA, since 2000, specializing in research on airport surface operation. Prior to joining NASA, he worked at Advanced Rotorcraft Technology and Raytheon. He received his B.S. and M.S. degrees from Seoul National University, Seoul, Korea, and his Ph.D. in Mechanical Engineering from the University of California, Davis. Dr. Jung is a Senior Member of the American Institute of Aeronautics and Astronautics.

Dr. Waqar Malik received his B.Tech. degree in mechanical engineering from the Indian Institute of Technology, Kharagpur, India, in 2001 and his M.S. and Ph.D. degrees in mechanical engineering from the Texas A&M University, College Station, Texas, in 2003 and 2007, respectively. He is currently a Research Scientist at University of California, Santa Cruz, where he is involved in developing a Surface Management System aimed at enabling significant increase in efficiency and throughput of airport surface operations while maintaining safety. He has previously worked on the synthesis of fixed structure controllers, combinatorial motion planning, collision avoidance, and path planning for unmanned vehicles.

Dr. Leonard Tobias received his PhD. from the Polytechnic Institute of Brooklyn in 1969. He has been engaged in ATM research at NASA Ames Research Center for more than 40 years.

Dr. Gautam Gupta received a B.Tech. in Civil Engineering from the Indian Institute of Technology Bombay, India, and an M.S. and Ph.D. in Transportation Engineering from University of California, Berkeley. He is currently a Research Scientist at UARC-NASA Ames Research Center, researching methods of improving airport efficiency. His research interests include airline operations and planning, airline economics, and large-scale optimization problems in transportation.

Ty Hoang received his Bachelor (1991) and Masters (1994) of Science degrees in aeronautical engineering at the California Polytechnic State University at San Luis Obispo. He started his career working on a pseudo- aircraft piloting system then participated in the field-testing of the Traffic Management Advisor (TMA) system at DFW in the 1990s. Thereafter, he assembled the Multi-Center TMA (McTMA) HITL simulator leading to successful field tests. He now leads the SARDA simulator work.

Dr. Miwa Hayashi is an aerospace engineer at NASA Ames Research Center with a background in aerospace human factors. She received a B.Eng. from Tokyo Metropolitan Institute of Technology in 1994, an M.S. in Aerospace Engineering from the University of Texas at Austin in 1997, and a Ph.D. in Humans and Automation from Massachusetts Institute of Technology in 2004. At Ames, she works to design and evaluate user interfaces for spacecraft cockpit and air traffic controllers' workstations.