Terminal Flight Data Manager (TFDM) Runway Balancing Capability Assessment

A.E. Vela
M.J. Sandberg
T.G. Reynolds

6 February 2015

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS

1. Report No. ATC-421	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Terminal Flight Data Manager (TFDM)		5. Report Date 6 February 2015
		6. Performing Organization Code
7. Author(s) A. E. Vela, M.J. Sandberg, and T.G. Reynolds, MIT Lincoln Laboratory		8. Performing Organization Report No. ATC-421
9. Performing Organization Name and Address MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108		10. Work Unit No. (TRAIS)
		11. Contract or Grant No. FA8721-05-C-0002
12. Sponsoring Agency Name and Address Department of Transportation Federal Aviation Administration 800 Independence Ave., S.W. Washington, DC 20591		13. Type of Report and Period Covered Project Report
		14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

Under the Terminal Flight Data Manager program, new operational improvements are envisioned at a number of large airports. One operational improvement manifests through the Airport Resource Management tool, which seeks to balance departure demand at runways. Another related operational improvement is runway balancing, which is expected to provide greater flexibility in tactical runway assignments. Both improvements are expected to reduce surface delays for departing aircraft. This report provides a study into the potential delay-reduction benefits of both capabilities at three case-study airports (DFW, LAX, and MCO). Through a series of simulation studies, it is found that the benefits associated with each operational improvement are closely linked to departure demand and imbalances in demand across filed aircraft departure procedures. So, while large delay-reduction benefits are expected at LAX—which exhibits both large demand and departure imbalances—smaller benefits are expected at DFW where departure operations are already well-balanced. Meanwhile at MCO, the operational improvements are not expected to reduce delays due to limited departure demand at the airport.

17. Key Words		18. Distribution Statement

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 87	22. Price

FORM DOT F 1700.7 (8-72)   Reproduction of completed page authorized

This page intentionally left blank.

REVISION HISTORY

Version	Description	Release Date
0.5	Internal development versions	N/A
0.6	Initial draft release version for comment	08/22/14
1.0	Revised draft release	09/30/14

This page intentionally left blank.

EXECUTIVE SUMMARY

Two of the capabilities being considered for the Terminal Flight Data Manager (TFDM) automation system are an airport resource information platform for departure-demand allocation and a runway balancing tool. Both the information platform for departure-demand allocation and the runway balancing tool are expected to aid in balancing demand at airports with multiple departure runways. This document reports the potential delay-reduction benefits of both capabilities at three case-study airports that were considered representative of the range of airports considered for TFDM deployment: Dallas-Fort Worth International Airport (DFW), Los Angeles International Airport (LAX), and Orlando International Airport (MCO).

The Airport Resource Management Tool (ARMT) is a TFDM core capability that serves as an information and decision platform; ARMT aids air traffic control in managing airport resources. In one particular use-case, ARMT users can observe projected demand at each runway and over each departure fix. Based on the demand information, air traffic control can then adjust aircraft runway assignment rules to dynamically allocate departure demand to manage delays at each runway. To better understand the potential benefit of ARMT and departure-demand allocation, three runway assignment algorithms are considered: a static runway mapping that only considers departure information (and does not update in response to departure demand); a traditional runway mapping that considers both airline information and aircraft departure information; and a dynamic mapping that adjusts runway assignments rules (based on aircraft departure procedures) at 30-minutes intervals to balance runway demand. The third runway assignment algorithm, which coincides with dynamic departure-demand allocation, best represents one approximation of the decision-making process enabled by ARMT, i.e., dynamic adjustment of the runway assignment rules in response to expected demand.

The goal of runway balancing is to reduce take-off delays by reassigning select aircraft to alternative runways relative to nominal runway assignments. This goal is achieved by utilizing open take-off slots when runways are unused or by reassigning aircraft to alternative runways for the purpose of advancing expected take-off times. Two candidate runway-balancing algorithms are evaluated: OpenSlot and GreedySlot. OpenSlot reassigns aircraft to runways in order to improve departure times, while ensuring no impact or delays to other aircraft. GreedySlot is a relaxed implementation of OpenSlot; it also seeks to improve the departure times of aircraft, but may potentially impact other individual aircraft in the process. Both runway balancing algorithms reduce overall surface delays.

The underlying procedure for the benefits analysis is to simulate multiple days of traffic at an airport using different runway assignment algorithms, with and without runway balancing; evaluate the resulting aircraft delays; and scale the results to calculate the annual delay savings of dynamic departure-demand allocation and runway balancing. Such a simulation-based analysis approach requires that airport-specific models for the case-study sites are constructed. These airport-specific models are generated by tailoring a generic airport queuing model using historical operational data.

After applying the simulation model and aggregating the benefits between 2017 and 2030, dynamic departure-demand allocation is projected to reduce surface delays by at least 30,000 hours at DFW and 515,000 hours at LAX. In 2030, the expected delay savings at DFW and LAX translate to an average delay reduction of 2.3 minutes and 20 minutes per aircraft. At MCO, simulations indicate that traditional runway assignments outperform departure-demand allocation by 35,000 hours cumulatively over the years 2017 through 2030. The lack of a benefit at MCO does not invalidate the positive impact of ARMT, but rather indicates that usage of ARMT will be different at MCO then at DFW or LAX. Additionally, there may be other benefits that are not considered here, like workload benefits.

The benefits of runway balancing are not as significant as dynamic departure-demand allocation. Cumulatively, from 2017 until 2030, the greatest expected delay saving from runway balancing is 52,000 hours at LAX. At DFW, the delays saving are less significant; through 2030, runway balancing will cumulatively reduce delays by 5,600 hours. At MCO, the benefits of runway balancing are less clear as traditional runway assignments are still able to out-perform dynamic departure-demand allocation even when combined with runway balancing. The greatest positive impact of runway balancing is observed at LAX, where operations across departure procedures are unevenly distributed leading to demand imbalances at departure runways, which runway balancing is able to take advantage of. Less potential benefit is observed at DFW due to a pre-existing departure demand that is well-balanced. At MCO, while asymmetric demands sometimes lead to runways imbalances, the excess departure capacity removes the expected benefit of formalized runway balancing.

Based on the evaluation of DFW, LAX, and MCO, this report recommends that further analysis should focus on a select set of additional airports with multiple departure runways – especially at airports projected to near their overall departure capacity. By expanding the benefit analysis, key airport and departure demand features associated with departure-demand allocation and runway-balancing benefits may be identified and generalized. Using the generalized features, a NAS-wide benefits assessment is then possible to help inform the TFDM investment analysis activities.

TABLE OF CONTENTS

Page

This page intentionally left blank.

LIST OF ILLUSTRATIONS

LIST OF ILLUSTRATIONS
(Continued)

LIST OF ILLUSTRATIONS
(Continued)

This page intentionally left blank.

LIST OF TABLES

This page intentionally left blank.

1. INTRODUCTION

1.1 TFDM PROGRAM BACKGROUND

The Terminal Flight Data Manager (TFDM) is a Federal Aviation Administration (FAA) program to develop an advanced automation platform for air traffic controllers to safely and efficiently manage aircraft operations on the airport surface and in the terminal area. Potential TFDM capabilities include [1]

- Surface Situation Display
- Electronic Flight Data
- Decision Support Tools, including:
  - Surface metering
  - Departure scheduling & sequencing
  - Runway assignment/balancing
  - Taxi routing
  - Airport configuration planning

In order to assess the suitability of TFDM to become an acquisition system for deployment in the field, it is undergoing an investment analysis process consistent with the FAA Lifecycle Management Process shown in Figure 1.
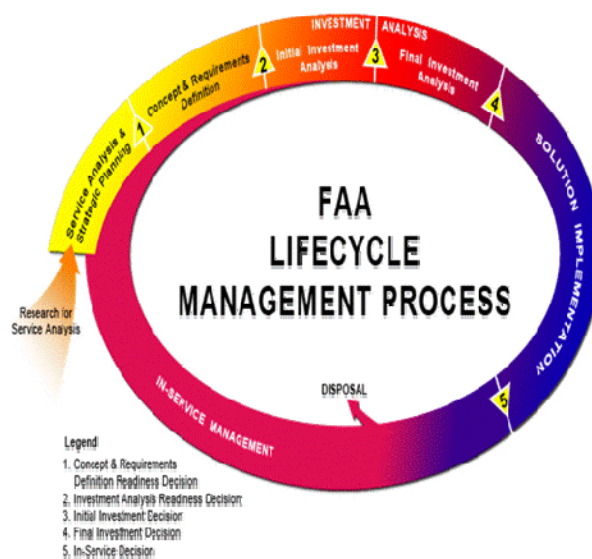
Figure 1. FAA Lifecycle Management Process [source: FAA].

MIT Lincoln Laboratory has been supporting the TFDM Program Office in the early stages of this process, including the design, implementation, and test of a prototype TFDM system deployed to Dallas Fort Worth airport in 2010–11 to assist in concept and requirements definition, and computer modeling and analysis to support the investment analysis processes. This included shortfall analysis studies to inform the Investment Analysis Readiness Decision (IARD) in 2010 [2], and a more detailed benefits assessment study to support Initial Investment Decision (IID) activities, with a report released to the FAA in 2012 [3]. This included analyses of the potential benefits of TFDM decision support tools in the areas of departure metering, sequence optimization and airport configuration, as well as a TFDM safety assessment. Support during FY14 is in the areas of benefits assessment of runway balancing, environmental impact mitigation, and electronic flight data, with the runway balancing assessment being the focus of this report.

1.2   DEPARTURE-DEMAND ALLOCATION AND RUNWAY BALANCING

Under TFDM, departure runway assignments for aircraft follow airport-specific rule-sets that consider factors such as airline preference, departure procedure, and aircraft type. At the same time, the runway assignments must be consistent with an airport's runway configuration at departure time, standard operating procedures, noise abatement procedures, and any aircraft restrictions. Should any considered factors change (e.g., airport configuration), TFDM will automatically adjust runway assignments where operationally feasible and practical [4].

The Airport Resource Management Tool (ARMT) is a TFDM core capability aligned with airport-specific runway assignment rule-sets. ARMT serves as an information and decision platform for managing airport resources like arrival and departure runways. In particular, air traffic control can use ARMT to observe projected demand at each runway and over each departure fix. Based on available demand information, air traffic control can adjust aircraft runway assignment rules to better balance departure demand. One of the more common means to allocate departure demand across runways is to re-map departure fixes to departure runways; typically at large airports the primary driver of an aircraft's runway assignment is its departure fix.

In addition to basic runway assignments, the TFDM concept of operations (CONOPS) allows for aircraft to be reassigned to alternative runways in an effort to reduce delays across active runways. By identifying delay imbalances, candidate aircraft are considered for new runway assignments prior to entering the active movement area so they can be taxied by ground control to the runway that results in the most efficient operation. A condition of the runway assignments is that they must not engender airborne crossovers (i.e., crossing streams of departure traffic from other runways to get to a given departure fix) or any other unsafe conditions.

The goal of runway balancing, as a direct result of assigning aircraft to alternative runways, is to reduce overall departure delay. This goal is achieved by utilizing open take-off slots when runways are unused (i.e., runway starvation is present) or by reassigning aircraft to alternative runways with shorter

taxi times and departure queues. Depending on how runway balancing is accomplished, the delay reduction may be distributed over all aircraft or can be targeted to a specific subset of aircraft.

Runway balancing is distinct from allocating departure demand at runways; runway balancing operates as an ad-hoc tool and is not a pre-planning activity. Prior to any runway assignment occurring, the rule-sets discussed above are generated to allocate scheduled runway demand in accordance to expected push-back times. If the push-back schedule is realized and the runway assignment rule-sets were generated evenly, then runway demand and delays are well-balanced. However, in the case of inaccurate push-back times or inaccurate taxi time predictions, the planned runway demands may not be realized, and significant delay imbalances may result. It is here that runway balancing performs a separate and distinct function from demand allocation; runway balancing provides fine-tuned, real-time corrections to runway assignment in order to balance demand and provide opportunistic substitutions for filling open departure slots. Departure-demand allocation is a form of operational decision-making, while runway balancing is ad-hoc tactical decision making.

1.3   SUMMARY OF ANALYSIS OBJECTIVES AND METHOD

The purpose of this document is to report the potential delay-savings as a direct result of departure-demand allocation and runway balancing. In support of the benefit analysis results presented in this report, a structured framework and method for performing the analysis is presented. So while the benefit analysis in this report is performed at a select set of airports, the structured framework can be applied to other airports as required.

In the process of executing the benefits analysis, a generic airport model is created. The generic airport model is then tailored to specific high-value case-study airports. Using operational data and airport-specific simulation models, appropriate airport performance metrics are calculated for current and future operations. Using the case-study airports as a foundation, necessary characteristics to derive benefits from runway balancing are identified. Additionally, operational constraints that limit the potential benefits of departure-demand allocation and runway balancing are documented.

1.4   AIRPORTS OF INTEREST

The TFDM program office has highlighted a number of airports where TFDM will be deployed (see Table 1 for a list of prioritized sites). Assuming no new runway construction, a subset of these TFDM airports commonly run multiple departure runways allowing them to be candidates for runway balancing. Figure 2 indicates the percentage of departures that occur in configurations with two or more departure runways for the complete set of TFDM airports. For the three major airports in the New York area: JFK, EWR, and LGA, the majority of departures occur in a single departure runway configuration. Meanwhile, airports like ATL, ORD, DEN, DFW and LAX are potential candidates for departure-demand allocation and runway balancing, as all of the airports make use of multiple departure runways, often times completely independent from any arrival operations. Other airports with fewer operations or interacting

arrival and departure runways are also of interest; example airports include CLT, MCO, and MIA. (See Glossary of Airports for airport codes.)

For the benefit analysis in this report, three airport sites were selected as case-study airports: DFW, LAX, and MCO. These case-study airports are highlighted in the prioritized list of TFDM airport in Table 1. The selected airports represent a range of total annual departures. It is worth noting that the case-study airports exclude airports with crossing arrival and departure runways or airports with dependent departure runways (e.g., SFO). Previously, CLT was selected as part of the case study; however, recent and coming changes to the airport (closure of runway 8/23, planned construction of a fourth parallel runway, and a new international terminal) limit any meaningful analysis. Any simulation-based analysis might be erroneous if based on current operations or speculative if considering future operations.

TABLE 1

Prioritized List of TFDM Airports

Sites	ATL (highest priority), JFK, EWR, IAH, ORD, SFO, LGA, LAX, CLT, PHX, PHL, MCO, DEN, BOS, LAS, DFW, DCA, MIA, IAD, BWI, MSP, SLC, SAN, DTW, FLL, SEA, MDW (lowest priority)
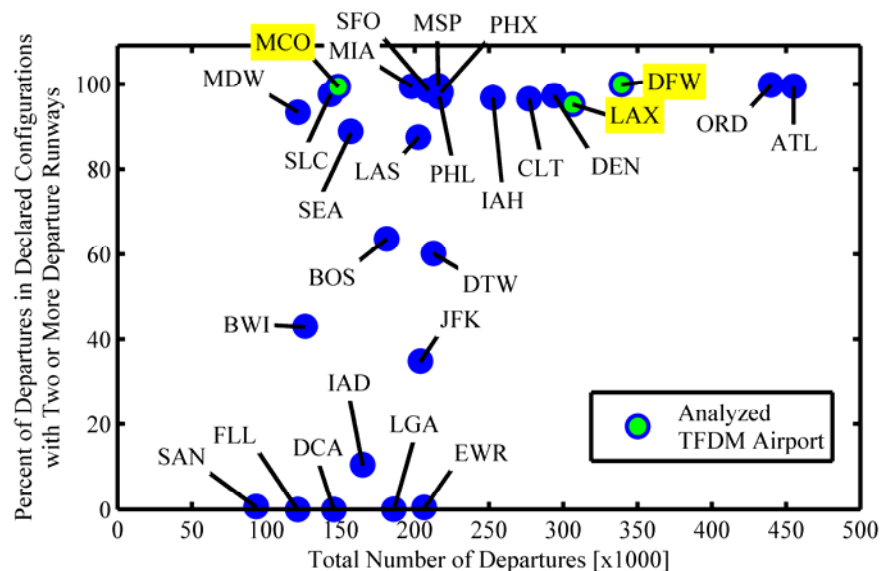
Figure 2. Percent of departures in configurations with two or more departure runways (according to ASPM, 2013).

1.5 REPORT OUTLINE

The remainder of the report is structured as follows:

Section 2 provides a description of the overall process used to perform the benefits analysis. As part of the process, an airport queuing model is created via a data-driven approach. The input and output of the queuing model, as well as the behavior of individual model elements, are described in the section. The section ends with a description of the calculations used to estimate the annual benefits of runway balancing.

Section 3 follows with benefit-analysis results for the airports DFW, LAX, and MCO. A comparison of two different runway balancing methods is provided. Additionally, characteristics and limiting factors associated with runway balancing are identified.

Section 4 discusses additional opportunities for further analysis.

Appendix A provides additional details describing how components of the airport queuing model are generated.

This page intentionally left blank.

2. TFDM BENEFITS ANALYSIS METHOD

The general method for the TFDM Benefits Analysis is to simulate multiple days of traffic at an airport consistent with forecast demand in various future years, with different departure-demand allocation schemes, and with and without runway balancing; evaluate the resulting aircraft delays; and scale the results to yield an annual benefit. An overview of the process is illustrated in Figure 3. Such a simulation-based analysis approach requires that airport-specific models for each site are constructed, and that the process for establishing annual costs is clearly defined.
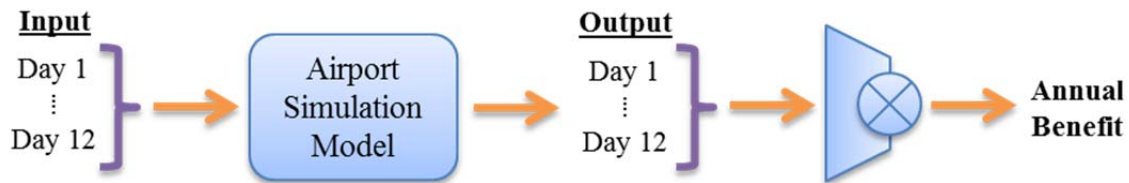
Figure 3. Benefit analysis process.

Each airport-specific simulation model (i.e., the process step shown in Figure 3) is based on the generic airport queuing model illustrated in Figure 4. Components of the airport model depicted in Figure 4 germane to the analysis work performed at MIT Lincoln Laboratory are encircled by the dashed line. The model components of particular interest to this report are the departure-to-runway mapping and the runway assignment modules; it is here that departure-demand allocation and runway balancing functionality resides. The inputs to the airport model, including the arrival and departure schedules for each day are exogenous inputs to the model provided by outside data sources; these inputs are discussed in Section 2.1.1. The outputs of the model are performance statistics and metrics for one given day, details of which are provided in Section 2.1.2. Using the outputs for a set of 12 days, the total benefits are proportioned to extract a total annual benefit of departure-demand allocation and runway balancing for each airport. Ideally, to calculate a total National Airspace System benefit, the same analysis should be performed at all relevant airports, or a sufficient subset to enable high-confidence extrapolation.
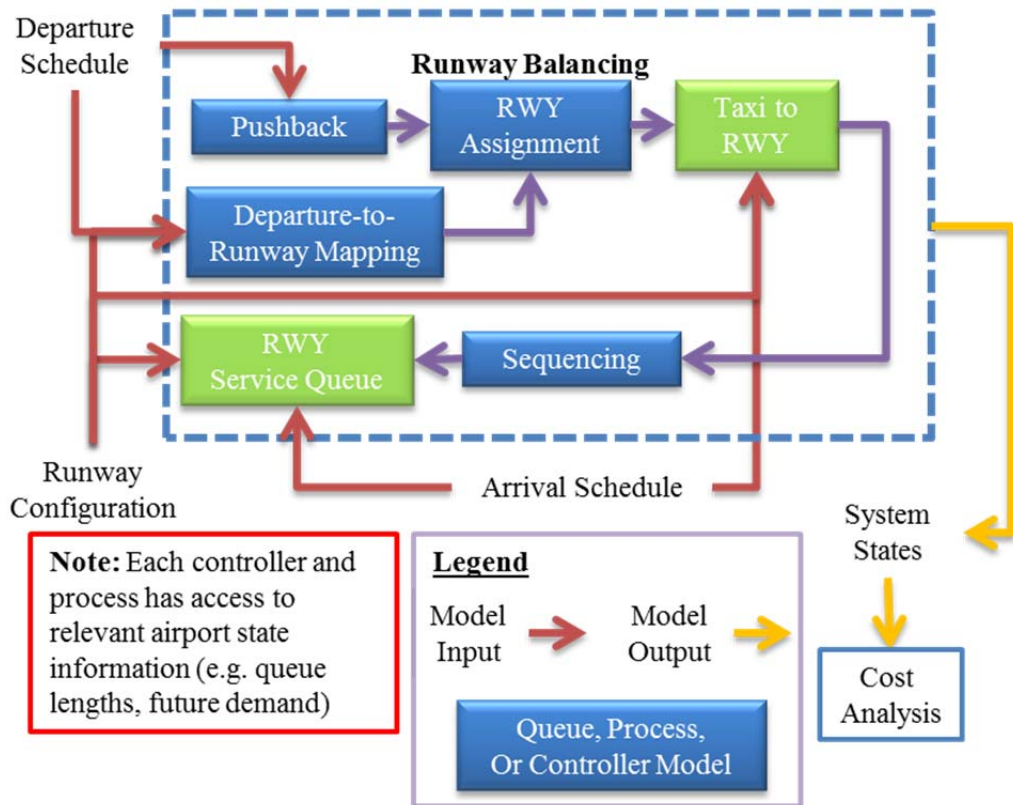
Figure 4. Generic airport queuing model.

The next portions of this section detail the airport simulation model. Later, relevant airport model settings and assumptions are provided. The section then proceeds with a description of how the baseline and the benefit pool are established, and concludes with the procedure for calculating the annual predicted benefits at an airport.

2.1   AIRPORT SIMULATION MODEL

For the generic airport queuing model represented in Figure 4, aircraft are processed through the model via a number of processes and control actions. Example processes include taxiing between spot areas and the runway, and waiting in a runway departure queue for take-off. Control actions include runway assignments and take-off clearances. Control actions can move aircraft from one process to another (e.g., from spot to runway queue, or from the front of the runway queue to take-off). Control actions can also indirectly move aircraft within a process (e.g., advancing aircraft within a departure queue after each take-off). Some of the controls and processes are airport-specific and require that they be

generated for each site; these airport-specific models include, but are not limited to taxi time models, runway assignment control, and airport configurations models.

As an aircraft moves through the queuing model, information about the aircraft's state is stored. Relevant information includes the aircraft's runway assignment, the aircraft's entrance time to the movement area and runway queue, and the aircraft's exit time from the runway queue to its subsequent take-off. Likewise, pertinent airport state information like queue lengths and average wait times are updated with each aircraft movement. Using both aircraft and airport information, performance metrics related to runway balancing are extracted.

The next two portions of this subsection detail the necessary model inputs and expected model outputs. Afterwards, the behaviors of each component of the airport queuing model are described.

2.1.1 Model Inputs

The primary inputs to the simulation model are arrival and departure schedules for the appropriate analysis year and an airport runway configuration description. The schedules are provided by the System-Wide Analysis Capability (SWAC) model. Airport runway configurations are taken from the Aviation System Performance Metrics (ASPM) database.

Airport Arrival and Departure Demand Schedule

The SWAC model is a fast-time system-wide model that forecasts airport demand and performance over the NAS [5]. Select days from the SWAC model covering 2010 until 2030 in 5 year increments was used in this analysis. The SWAC model takes into account a number of evolving factors such as fleet mix, aircraft equipage, and most importantly demand schedules in creating the output for each day. The resulting arrival and departure schedules factor in resource constraints at both airports and en route airspace (which may be the result of weather conditions). The airport simulation model described in this report takes a portion of the output from the SWAC to formulate a schedule.

While the SWAC model produces a large number of outputs (approximately 90), only a select few are required to simulate operations at an airport of interest. The needed SWAC outputs for the airport-specific models are

- Destination airport for departure aircraft
- Aircraft type
- Airline
- Scheduled push-back time for departure aircraft
- Actual push-back time for departure aircraft
- Wheels-on time for arrival aircraft

- Nominal taxi time

Additional required inputs to the airport simulation model include departure procedures for all aircraft, given the departure fix is one of the key drivers of runway assignment. While aircraft departure procedures are provided by the SWAC model, they sometimes involve unpublished procedures. The cases of unpublished procedures limit modeling baseline simulations because they represent unknowns. Additionally, the unpublished procedures are not consistent with historical data used to build components of the airport simulation model. As such, a departure procedure pre-processing step is required to regenerate valid departure procedures for each flight. A description of the departure procedure pre-processing step is given in Appendix A.1.

Airport Configuration

An airport's runway configuration declares which runways are active and the associated operation type (arrivals, departures, and shared). Over a given day, an airport's runway configuration may change in response to traffic demands and weather conditions (especially wind). Table 2 demonstrates an example of the variation in MCO's runway configuration over a sample day.

TABLE 2

Runway Configuration for MCO on July 3, 2010

Time	Runway Configuration Arr. RWYs | Dep. RWYs
12:00 AM	17L,18R|17L,18L
12:15 AM	17L,18R|17R,18L
3:00 PM	35R,36R|35L,36L
9:45 PM	17L,18R|17L,18L

At some airports, in the case of multiple arrival and departure runways, the day may begin with only one active arrival and departure runway. However, as demand increases, so does the number of active runways. At nighttime, the number of runways scales down while still fulfilling arrival and departure demand. Operationally, the number of active runways is reduced during off-peak hours to satisfy staffing constraints.

Because the benefit analysis in this report covers demand allocation and runway balancing, and not staff requirements or airport configurations, the benefits analysis is performed under the assumption that airports operate in a runway configuration that maximizes throughput during the day. In the case of MCO,

a reasonable runway configuration schedule would be 12:00 AM–3:00 PM: **17L,18R|17R,18L,** followed by 3:00 PM–12:00 PM: **35R,36R|35L,36L**. With this configuration schedule, the maximum number of operationally acceptable runways is active throughout the day so that available runway resources are not an artificial constraint on the arrival and departure demand. Otherwise, if the runway configuration schedule in Table 2 were used to simulate increased traffic demand in 2030, there would be the possibility that increased demand after 9:45 PM could lead to large runway queue delays due to shared runway operations on 17L.

2.1.2   Model Outputs

The primary outputs of the airport simulation model are the surface-delay times for all aircraft on a given simulation day. The aircraft surface-delay times are extracted from event times as each aircraft is processed through the airport queuing model. The surface-delay time takes into account both taxi times between spot locations and the runways, and the aircraft wait-time at runway departure queues. Airport state information is also maintained to better understand when runway balancing serves as a useful tool for managing aircraft delays and when runway balancing proves to be ineffective.

Using multiple simulation days, the expected benefit is extrapolated over a year; Section 2.4 documents this calculation process.

2.1.3   Model Component: Runway Assignment Mapping

Airports typically have standard operating procedures that map an aircraft's filed departure procedure to the runway from which the aircraft will depart. For example, according to standard operating procedures at DFW, if an aircraft files a flight plan using the Area Navigation (RNAV) departure procedure AKUNA, the aircraft will most likely depart on the eastern departure runway 17R/35L. The mapping between departure procedures and departure runways is established to provide clear separation between all arrival and departure flows, but also to help manage controller workload and improve operational efficiency. An example baseline departure fix-to-runway mapping for a generic airport is illustrated on the left in Figure 5. Groupings of departure procedures are represented by numbered colored-coded nodes, each mapping to a similar colored runway.

Departure fix-to-runway mappings are not necessarily static. Air traffic control often adjusts mappings to balance workload and to manage the demands and delays at each runway. For example, consider the case of the generic airport in Figure 5 in which both runways have the same available capacity. Supposing departures to the north constitute the majority of the demand (60%), long runway queues at 9L will develop. Air traffic control can adjust the mapping such that flights in departure group (5) take off from the southern runway 9R without causing complications from crossing traffic once airborne. By remapping those departure groups on the boundaries, the demand at each runway can be approximately balanced (from 60/40 to 53/47). Dynamic mapping essentially performs the role of departure-demand allocation.

11
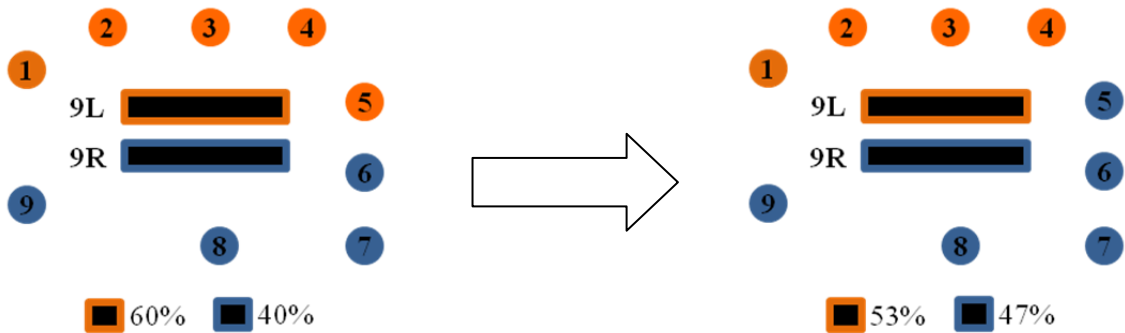
Figure 5. Departure fix-to-runway mapping for a generic airport.

The inclusion of a departure fix-to-runway mapping, i.e., departure-demand allocation, within the airport simulation model takes a dynamic approach that adjusts the mapping according to traffic demand. The dynamic mapping attempts to replicate the actions of air traffic control to manage runway demand and delays. At a regular interval of 30 minutes, the dynamic mapping function observes the current runway queues and forecasts the departure demand up to 30 minutes in advance to find an optimal departure fix-to-runway mapping. A time horizon of 30 minutes is selected because the TFDM ARMT decision-support tool for resource balancing also uses 30 minutes for the planning horizon [4]. While the function of departure fix-to-runway mapping and construction of runway assignment rules are much more nuanced and complicated in practice, the mapping procedure described here provides an approximate proxy for real-life dynamic mapping.

The dynamic mapping process begins from a static departure fix-to-runway mapping as a baseline representing the default current operation. Using this static mapping, a select set of departure groups near the boundaries is considered to be changeable – in that departure groups may be re-associated with another runway over the 30-minute time horizon. In order to ensure operationally suitable mappings, the dynamic departure fix-to-runway mappings should prohibit crossing departure routes. For the example in Figure 5, one option for changeable departures is shown in Figure 6; changeable groups are indicated by the dashed nodes.
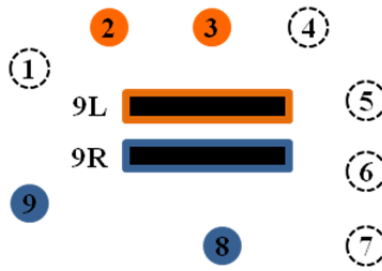
Figure 6. Departure fix-to-runway mapping with changeable departure groups.

For the simulations, the static departure fix-to-runway mappings are extracted from historical data based on the airport runway configuration. Details of extracting the mapping are provided in Appendix A.2; the static mappings for DFW, LAX, and MCO are provided later in Section 3. The departure groups near the split points are selected as changeable.

At each 30-minute decision point, the dynamic departure fix-to-runway mapping is obtained by selecting the mapping that proportions the departure demand closest to the departure capacity at each runway. Additionally, a penalty is incurred for any changes in the departure fix-to-runway mapping. Penalizing changes helps maintain continuity and consistency of operations over time. For any departure group change, the penalty is valued at a 5-aircraft reduction in runway imbalances. That is to say, a remapping that adjusts the 30-minute runway demand from a 10-to-14 aircraft split at the runways to a 12-to-12 split is not optimal as the runway imbalance is only reduced by four aircraft. Figure 7 illustrates a solution for the same generic airport when the departure capacities at the runways are asymmetric: 20 aircraft per 30 minutes for 9L, and 10 aircraft per 30 minutes for 9R. In the figure, the current runway queues and 30-minute scheduled demand at the departure groups are noted, as both are considered in the dynamic mapping. The future demand and current runway demands should be as close to 2-to-1 as possible, corresponding to the 20-to-10 departure capacity ratio. The solution indicated in Figure 7 maps 26 aircraft to 9L, and 12 aircraft to 9R, resulting in a 2.166-to-1 ratio.
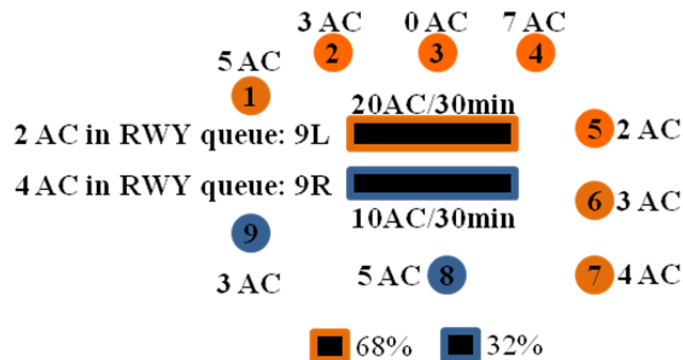
Figure 7. Example solution dynamic mapping problem for runways with asymmetric departure rates.

13

Proportioning traffic according to the overall departure capacity on each runway is only one possible method for performing dynamic departure fix-to-runway mappings. Other methods could consider delay sensitivity and uncertainty in the departure capacity, or a model could take detailed temporal approach to balance traffic on the runways over each minute, instead of balancing over the aggregate of the 30-minute time horizon. While departure fix-to-runway mappings provide guidelines for runway assignments, there can be exceptions to the assignments. Reasons for the exceptions may include aircraft types, runway length, airline preferences, noise abatement, or advancing take-off times. The next section will address when departure fix-to-runway mappings are overridden, particularly for the expressed purpose of runway balancing.

2.1.4 Model Component: Runway Balancing

In practice, aircraft typically follow the departure fix-to-runway mappings described in Section 2.1.3. However, when operations permit there are cases when the mapping rules may be overridden and additional flexibility in runway assignment is allowed. Aircraft are sometimes assigned alternative runways according to airline preferences (e.g., to minimize departure and taxi time when an airlines gates are much closer to one departure runway than another), or to help balance operations. Furthermore, aircraft may be assigned an alternative runway in order to segregate large and small aircraft, or because weather conditions and aircraft take-off weights require a longer runway.

As part of the runway balancing benefits analysis, the runway assignment procedure will be adjusted and analyzed. Three different options are considered: a runway assignment mapping, and a runway assignment mapping with one of two augmented approaches, OpenSlot and GreedySlot. The baseline mapping assigns aircraft to runways strictly based on the dynamic departure fix-to-runway mappings. OpenSlot and GreedySlot represent two possible options for implementing runway balancing on top of the dynamic departure fix-to-runway mapping. Both assignment procedures seek to reduce total departure delay. OpenSlot and GreedySlot are associated with particular rules that influence when each method is applied and which aircraft derive benefits. The two runway assignment procedures are detailed below, but first, the notion of a flexible departure group must be introduced.

For the runway balancing procedures OpenSlot and GreedySlot to override default runway assignment mapping, a set of candidate aircraft must first be identified as available to be reassigned runways. In the simulation model, the set of candidate aircraft available for runway balancing is determined by the aircraft's departure procedure. For a specific departure fix-to-runway mapping, select departure groups are deemed flexible. Any aircraft filing a departure procedure within a flexible departure group has the potential to be reassigned to an alternative runway according to the OpenSlot and GreedySlot procedures.

Figure 8 provides an illustrative example for the previously described generic airport with asymmetric departure capacity; the flexible departure groups are indicated by the split-color nodes. The process of selecting flexible departure groups occurs after setting the mapping. The selection criteria for flexible departure groups require that the group be at a boundary, and that the departure groups be

associated with the imbalanced runway. In this case, 9L has proportionally more traffic than its departure capacity allows, as such departure groups (1) and (7) are deemed flexible. Flights associated with the flexible departure groups could be reassigned runways by OpenSlot or GreedySlot. All other flights that leave along inflexible departure groups strictly follow the default departure fix-to-runway mapping.
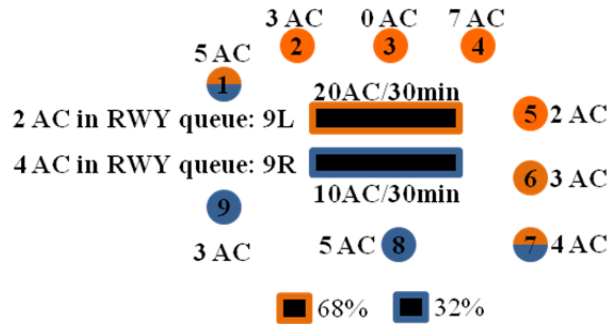
Figure 8. Flexible departure groups.

OpenSlot Runway Assignment

The OpenSlot procedure replicates the runway balancing scheme described in the TFDM System Specification Document [6]. First, for an aircraft to be considered a potential candidate for runway balancing, the aircraft must belong to a flexible departure group. If an aircraft is a potential candidate for runway balancing, then two other conditions must hold for an alternative runway assignment:

1. By changing runway assignments, the aircraft must have an earlier expected departure time compared to if they stayed on their default runway.

2. The aircraft shall not delay any other aircraft that would normally take off from the alternative runway.

The second condition requires a large gap in the expected take-off times to exist at the alternative runway. Otherwise, uncertainties in taxi times for the candidate aircraft may cause the aircraft to violate the second condition by delaying others. When applying the two conditions, OpenSlot approximately ensures that no individual aircraft is worse-off than if OpenSlot was not applied (there are potential second-order interactions with the dynamic load-balancing).

An example application of OpenSlot is depicted in Figure 9 for two generic runways 9L and 9R. In the figure, runway 9R receives a steady demand of aircraft, while 9L has a large open take-off slot. One of the flexible aircraft that would typically take-off from 9R, according to the runway assignment rules, is

reassigned to 9L. The reassigned aircraft is able to advance its expected take-off time, and does not impede any other aircraft that would normally take-off from 9L. Meanwhile, the two other aircraft that would have followed on 9R are advanced one slot as a result of the opening made by reassigning the candidate aircraft.
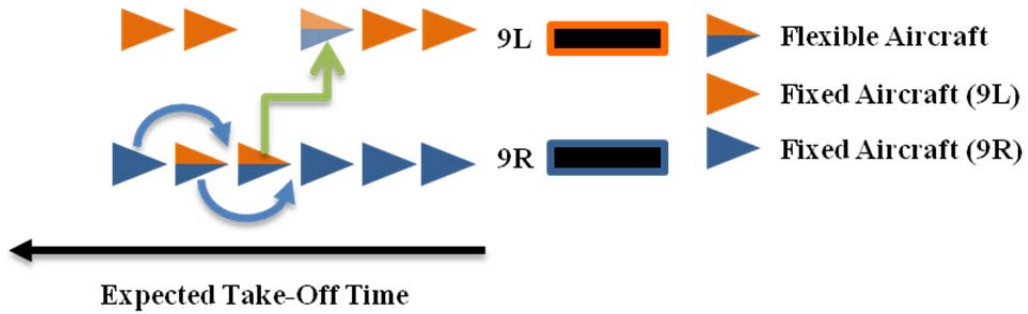
Figure 9. Example runway balancing with OpenSlot.

GreedySlot Runway Assignment

The GreedySlot procedure seeks to improve the overall system performance at an airport by taking a greedy approach to runway balancing. A relaxed version of OpenSlot, condition (2) is removed, and condition (1) is applied. An additional constraint is added that only allows aircraft from impacted runways to be reassigned to less impacted runways (according to the predicted allocated 30-minute schedule). Following the conditions, each aircraft that is a potential candidate for runway balancing seeks out a runway assignment that reduces its expected take-off time. Such a greedy approach is expected to distribute delays more evenly across runways and improve overall system performance. However, with GreedySlot there is possibility that some aircraft will have later take-off times when compared to the case when no runway balancing is applied, but the overall delay is reduced. Figure 10 depicts one outcome of GreedySlot. In the scenario, two aircraft switch runways from 9R to 9L, advancing their expected take-off times. However, two aircraft set to depart from 9L are delayed slightly to accommodate the reassigned aircraft. In total, there is still a net reduction in delays.
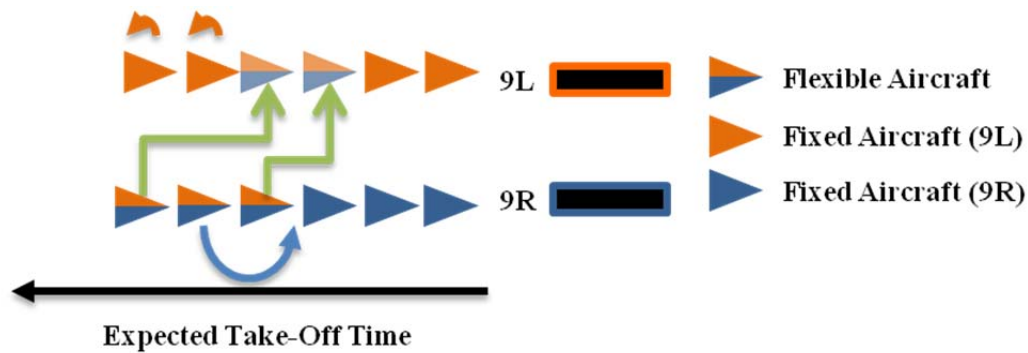
Figure 10. Example runway balancing with GreedySlot.

For both runway-balancing procedures, OpenSlot and GreedySlot, a measure of the *expected departure time* is required. For the purposes of simulation, the expected departure time is calculated by summing the taxi time from the spot to the runway queue, and the expected wait-time spent within the departure queue. The wait-time within the departure queue is calculated by considering aircraft already in the queue, taxiing aircraft, and any aircraft that will be pushing back within the next 10 minutes. Predicted taxi times and estimated arrival times at the runway queue for the previously queued aircraft are estimated using the taxi time model described in Section 2.1.5. In calculating the expected wait-times, aircraft within the departure queue are assumed to depart at a constant rate according to the 30-minute departure capacity of the runway.

2.1.5   Model Component: Taxi Times

The taxi time component generates times for an aircraft to taxi from a spot location to the runway queue. The process for generating taxi times requires two steps: generating spot locations for each aircraft, and calculating the taxi time from the spot location to the runway.

In practice, spot locations are determined by the aircraft's departure gate, assigned runway, surface traffic, as well as spot availability. Because historical gate information is not available for all airports, a simpler model is used. The model generates random spot locations using only airline information and flight type (international or domestic). Next, accounting for the spot location, the assigned departure runway, a regression model assigns the aircraft a taxi time. The spot location model is built using historical spot location distributions for each airline, while the taxi time regression model uses a random forest regression model. Both models are constructed from 100 days of historical ASDE-X data (see Appendix A.3 for details). By constructing a regression model that only considers spot location and departure runway, the effects of surface congestion are omitted, which is preferred in order to isolate the direct benefits of departure-demand allocation and runway balancing. Otherwise the inclusion of surface congestion begins to draw in other operational improvements, such as the TFDM core capability of

17

departure metering, which is expected to reduce surface congestion and taxi times by regulating the number of taxiing aircraft.

To assist in understanding the sensitivity of benefits to taxi times, the surface-delay time is divided up between the taxi time and the departure queue wait-time. Also, in Section 3.5, a sensitivity analysis is performed to demonstrate that many results are invariant to the taxi time model.

2.1.6 Model Component: Sequencing

Sequencing of aircraft is a function commonly performed at major airports to improve runway throughput. Otherwise, without proper aircraft ordering there is a risk of incurring additional system delay. The additional delay is derived from spacing requirements to prevent wake turbulence hazards at take-off, and to ensure aircraft separation for flights traveling along the same departure route.

The sequencing component in the model checks for opportunities to reduce runway delays via two-aircraft swaps. If two consecutive aircraft can be reordered such that the take-off time of subsequent aircraft are advanced, then the swap occurs. This sequencing process is depicted in Figure 11. To prevent aircraft from being swapped backwards or forwards multiple times, aircraft are restricted to only one swap.
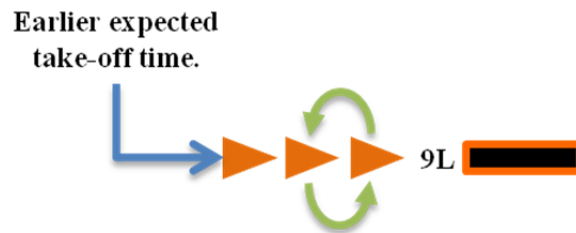
Figure 11. Sequencing is applied when expected take-off times are reduced.

In the airport simulation model, sequencing is assumed to occur at the runway threshold. In reality, sequencing can occur at spot locations or even along taxiways. From a modeling perspective there is no need to distinguish where swapping occurs because throughput and wait-times remain the same, as long as the aircraft are sequenced by the time they reach the runway threshold. Additionally, while advanced and aggressive forms of sequencing could constitute an operational improvement, the sequencing implementation presented here is intended to be simple and limited to single aircraft swaps. This type of swapping procedure is commonly used at major airports when provided with sufficient taxi space. Accordingly, inclusion of the simple sequencing algorithm does not represent an operational improvement; the simulation results presented are independent of future operational improvements with regards to sequencing.

2.1.7 Model Component: Runway Service Queue

One of the more critical features of the airport queuing model is the runway service queue, which simulates how aircraft are serviced at the runway for take-off. Once a sequence is established, aircraft within each runway departure queue take off one at a time. For dedicated departure runways, the time between successive aircraft take-off times is determined by the spacing needed to prevent wake turbulence hazards and spacing need to maintain separation along the same departure route. Using historical data, average spacing times between all aircraft operations are extracted. These times are then utilized to model service times as part of the departure process.

For aircraft departing in different departure groups, the look-up table, Table 3, is used to determine runway service times. The service times are extracted from an analysis of PASSUR data and Enhanced Traffic Management System (ETMS) flight plan data. Appendix A.4 contains additional details into how the values in the table were generated. Note that the values are not expressed as the minimum distance separation standards often found in aviation literature, but rather the table represents separation times used in practice.

TABLE 3

Spacing Time for Aircraft on Different Departure Routes (seconds)

L\F	Small	Large	757	Heavy
Small	60	60	60	60
Large	60	60	60	60
757	105	105	90	90
Heavy	120	115	115	95

When aircraft are departing along the same departure route and/or over the same fix, additional time may be required to ensure separation once airborne. Again, using historical PASSUR operations data and flight plans an additional spacing table is generated; see Table 4. The values shown in bold are larger than corresponding spacing times in Table 3, while values in red indicate some uncertainty due to lack of data or missing relevant data. The spacing times in Tables 3 and 4 are used by the sequencing process discussed in Section 2.1.6.

TABLE 4

Spacing Times for Aircraft on the Same Departure Route (seconds)

L\F	Small	Large	757	Heavy
Small	*60*	*60*	*60*	*60*
Large	90	90	90	*90*
757	105	105	90	90
Heavy	*120*	115	115	95

While not relevant to the airports operations modeled in this report, a similar analysis was performed to understand the departure process in shared runway operations. For this case the required time-spacing between arrival and departure aircraft is estimated.

2.2 STANDARD SIMULATION SETTINGS AND ASSUMPTIONS

Outside of the modeling parameters and behaviors listed previously, the airport simulation model makes use of additional modeling assumptions. In addition to assumptions embedded in the SWAC model, some key assumptions are stated here so that future simulation results can be understood within the proper context.

- The demand schedule within the simulation follows the SWAC schedule without adjustment
  - Gate-out times are not adjusted in response to the simulation
  - Wheel-on times are not adjusted in response to the simulation
- All operations are assumed to occur under ideal conditions
  - Airports are using an airport configuration consistent with high-demand operations, using multiple departure runways
  - The airport has at least one dedicated departure runway
  - Airports are operating under visual instrument conditions, regardless of simulation inputs found in the SWAC model
- The actual departure demand within the next 30 minutes is always known
- The actual gate-out times for aircraft are known within 10-minutes of its occurrence
- The exact taxi times for aircraft are considered known; however, an uncertainty window is used in all calculations for identifying open departure slots
- Spot location distributions remain the same for all airlines in the future

2.3 BASELINE AND BEST-CASE MODELING

To understand the scale of departure-demand allocation potential benefits of runway balancing, simulation results for OpenSlot and GreedySlot must be placed in relation to a baseline scenario and to a best-case scenario. One baseline scenario represents what would occur if current-day operations continued without runway balancing. The best-case scenario represents what would occur when airport throughput is maximized and delays are evenly distributed across runways.

Moving forward, one baseline will be partially defined by the dynamic departure fix-to-runway mapping described in Section 2.1.3. Every 30 minutes, the mapping is adjusted according to the arrival and departure demand. Aircraft will strictly follow the departure fix-to-runway mapping without exception, as no runway balancing is applied. Traditional runway assignments based on airline and aircraft departure procedure is also considered.

For the best-case scenario, all departure fixes will be considered flexible and runway assignments will be generated according to GreedySlot. In this way, aircraft have the option of departing from any runway, regardless of departure procedure; accordingly this scenario is referred to as *fully flexible*. A representation of a fully flexible DFW mapping is depicted in Figure 12. Such a balancing framework is operationally infeasible, as it would engender extensive crossing departure procedures. That said, the best-case fully flexible scenario establishes an upper bound performance reference against which OpenSlot or GreedySlot can be judged.
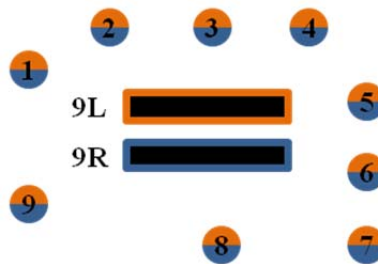
Figure 12. A fully flexible runway assignment scheme.

2.4 ANNUAL AIRPORT ASSESMENT CALCULATION

There are 12 available SWAC simulation days for the years 2010, 2015, 2020, 2025, and 2030 representing different typical demand days in each quarter – these days were provided courtesy of MCR and the FAA. The annual benefit of runway balancing over each year is calculated by averaging metrics over the 12 simulation days, adjusting according to the baseline, and scaling the aggregate benefits across these 12 days to estimate annualized results. The primary metric considered in this study is the total surface-delay time. As depicted in Figure 13, the surface-delay time begins when the aircraft leaves the

spot location and ends at take-off. As such, surface delay includes both the taxi time to the runway, in addition to any time spent waiting within the departure queue. As desired, the taxi time or departure-queue wait time could be the metric considered in any calculations or benefit assessment.
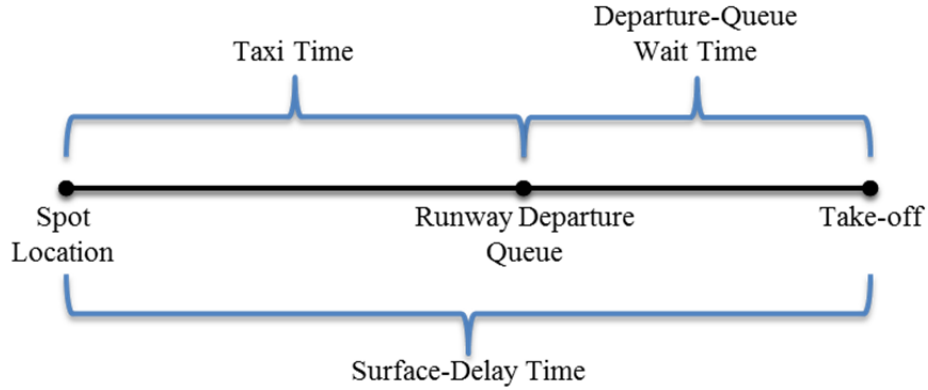
Figure 13. Simulation delay metrics.

For any given metric (e.g., surface delay), the daily average for year y is given by

$$\overline{M}_y = \frac{1}{12} \sum_{d=1}^{12} M_{d,y} \, ,$$

where $M_{d,y}$ is the calculated metric for a single simulation day with runway balancing. Similarly a baseline measure $\overline{M}_y^{base}$ is calculated according to the simulation settings described in Section 2.3. Annually, the total measure of the benefit is

$$\bar{B}_Y = \frac{\#VMC}{365} \left(\overline{M}_y - \overline{M}_y^{base} \right) .$$

The annual benefit $\bar{B}_Y$ is a scaling of the average daily benefit by the fraction of days the airport operated under visual meteorological conditions (VMC); instrument meteorological conditions (IMC) are mostly excluded. The VMC scaling factor is required because during inclement weather, operations at an airport are often adjusted. These adjustments were not explicitly considered in the simulation model. For example, possible complicating factors include a reduction in the number available runways due to missing Instrument Landing Systems, and increased separation requirements for instrument procedures.

As such, the benefits calculation includes only benefits realized on days with VMC, assuming the average number of VMC days per year. Because instrument meteorological conditions days are not considered, the estimated annual benefits are likely conservative.

2.5 SUMMARY

This section describes the overall method and underlying airport simulation model used for the TFDM Runway Balancing Benefits Analysis. The airport simulation model is a queuing model framework which can be adapted to specific airports. The model is based on knowledge of airport operations, and where possible, supported by and reflective of historically observed and measured data.

Moving forward, the airport simulation model is tailored for DFW, LAX, and MCO and simulated. Following simulation, potential benefits of runway balancing are extracted using the process described in Section 2.4.

This page intentionally left blank.

3. TFDM RUNWAY BALANCING ANALYSIS RESULTS

SWAC traffic forecasts project a significant increase in departure demand at DFW, LAX, and MCO over the next twenty-five years. As shown in Figure 14, from 2010 to 2030 the number of departures at MCO is expected to double to just over 800 operations per day, while at DFW, and LAX the average number of daily departure operations will exceed 1200 flights by 2030. As demonstrated in this section, increases in departure demand will in certain cases necessitate changes in operations, otherwise departure delays will become unsustainable.

In this section a review of simulation results and analysis are presented for the case-study airports, with special focus on DFW. For each airport a brief description of operations is provided, along with key characteristics relating to the potential implementation of and benefits of departure-demand allocation and runway balancing. The analysis focuses on departure delays with and without runway balancing.
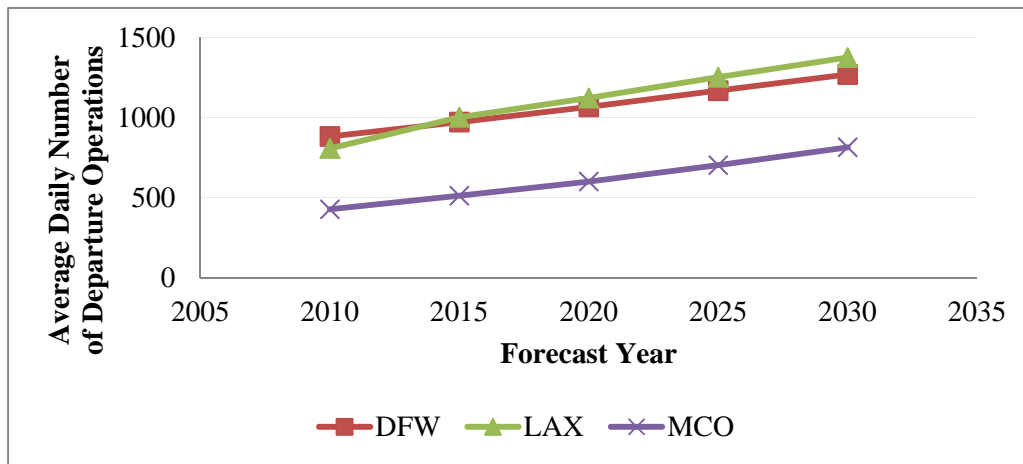
Figure 14. SWAC departure demand forecast.

3.1   DFW

Airport Background

The operations at DFW are the most structured and predictable of the three airports analyzed in this report. The airport's most common runway configurations and corresponding annual departure operation counts are listed in Table 5. In the two most common departure configurations (by annual departure count), historical aircraft runway assignments can be predicted with just over 92% and 86% certainty using only knowledge of the aircraft's filed departure procedure. If both the departure procedure and

airline identifier are considered, the historical runway prediction rate increases marginally. Such a result suggests that runway assignments are primarily dictated by departure fix-to-runway mappings. In contrast, at airports like LAX, the airline and departure runways are highly correlated via the departure terminal; the desire for shorter taxi times from their terminal gate and earlier departure times guide aircraft preferences towards utilizing the nearest available runway, assuming equal departure queues. Moving forward, runway predictions based on departure procedure and airline will be referred to as *traditional runway assignment*.

TABLE 5

Predicting Airport Operations at DFW

Departure Configuration	Total Number of Departures	RWY Prediction Rate [Percent] (Departure Group)	RWY Prediction Rate [Percent] (Departure Group + Airline)
17R-18L	208498	92.8	93.0
31L-35L-36R	82090	86.6	86.6

In the departure runway configurations 17R-18L and 31L-35L-36R, departure operations are performed independently, with large aircraft utilizing the inner runways, and when allowed, small aircraft segregated to 31L for non-RNAV departures. These two departure configurations often map to the arrival runway configurations 17C-18R or 36L-35C. A diagram of the airport is provided in Figure 15 for visualization. For departing aircraft, the difference in taxi distance between the two departure runways (18L/36R and 17R/35L) is approximately three-quarters of a mile. As such, even if an aircraft's departure procedure typically maps to a further runway assignment, the additional taxi time incurred is not always significant, especially when persistent departure queues exists (i.e., more than 2–3 aircraft).
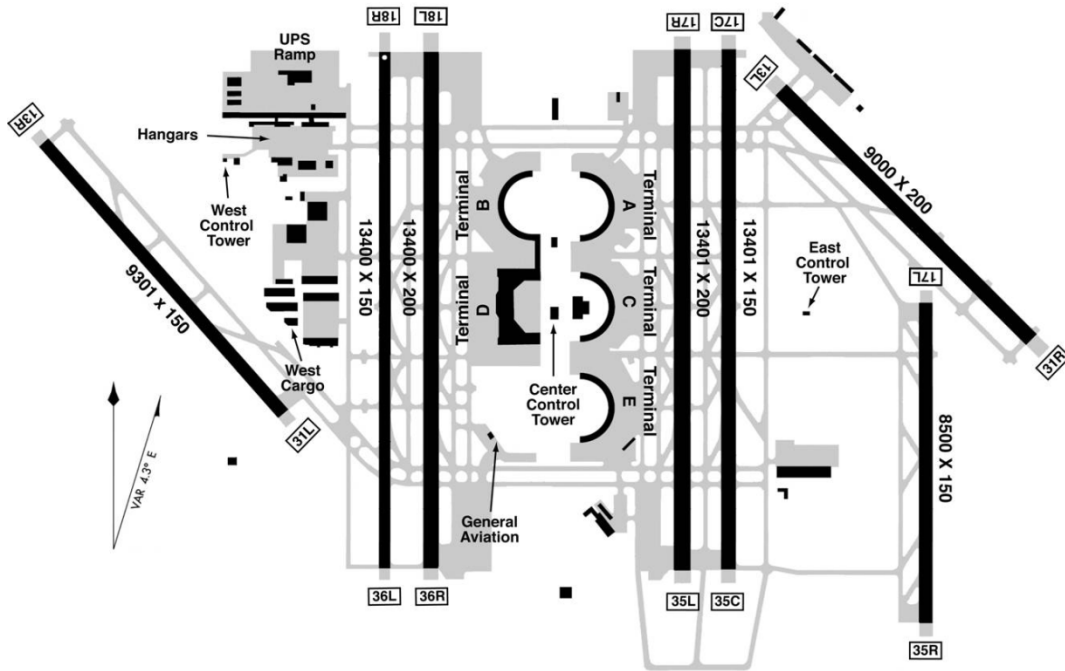
Figure 15. DFW airport diagram.

The static baseline departure fix-to-runway mapping for DFW in the 17R-18L departure configuration is depicted in Figure 16; western departure fixes map to 18L, while eastern departure fixes map to 17R. The diagram for the departure fix-to-runway mapping represents 16 departure groups, which correspond to more than 70 departure procedures with transition points (e.g., TEX.ROLLS and BLECO.DECKK). Many departure groups include both RNAV procedures and non-RNAV equivalents that map to the same departure fix.

When applying the static departure fix-to-runway mapping to the SWAC flights (which are inputs to the airport simulation model), the aggregate runway demand over the 30 years is a 44%–56% split between 18L and 17R. It is worth noting that the demand over the northern and southern departure groups are almost evenly distributed across (1)–(3) and (9)–(12), and in fact, over all departure groups, no departure group accounts for more than 11% percent of the total departure demand.
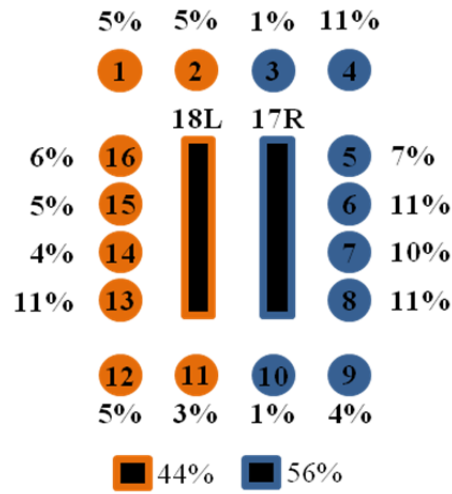
Figure 16. DFW static baseline departure fix-to-runway mapping.

Airport Simulations and Analysis

The DFW simulations are run in a 17R-18L departure configuration. Four different candidate options for runway assignments are first considered, none with runway balancing. The first two are static options: the static departure fix-to-runway mapping illustrated in Figure 16, and the traditional runway mapping based on aircraft and departure procedure (corresponding to Table 5). The other two runway assignment candidates make use of dynamic departure fix-to-runway mappings to allocate departure demand on the runways. The dynamic mapping options use the static baseline mapping as a foundation, but allows for changeable departure groups to be re-associated with the runways. The dynamic mapping options are illustrated in Figure 17. For Dynamic Mapping 1, only two northern and two southern departure groups are changeable, while for Dynamic Mapping 2 all northern and southern departure groups are changeable. As stated previously, the dynamic mapping process occurs every 30 minutes and serves as an approximation to air traffic control responding to the ARMT information platform.

Figure 17. Dynamic mapping options for DFW.

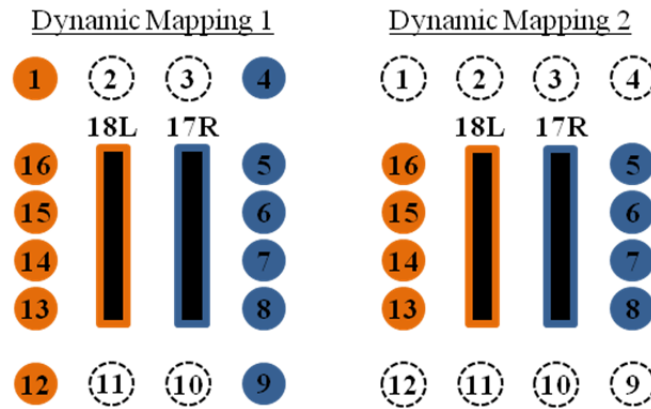Simulation results for DFW under the four runway assignment algorithms are shown in Figures 18 through 20. With exception to the fully flexible case, the delay measures in the figures indicate that at present-day demands, none of the runway assignment methods result in a substantial surface-delay reduction compared to traditional runway assignments. For the 2010 simulation year, extracted taxi times indicate that aircraft spend 6.15 minutes in the movement-area from spot to take-off (see Figure 18), of which 40 seconds are waiting in a runway departure queue. The 2010 per-aircraft surface-delay time translates to an average daily total surface-delay time of about 90 hours. By 2030, simulation results indicate an increase of the per-aircraft surface-delay time to about 9 minutes for the traditional runway assignment, static mapping, and Dynamic Mapping 1, signifying an average runway queue-time of approximately 3.75 minutes. Dynamic Mapping 2 is able to better manage per-aircraft surface-delay times to 8 minutes in 2030. The difference in the average daily total surface-delay time between Dynamic Mapping 2 and all others is more than 22 hours. Also included in the figures are simulation results for the fully flexible case in which aircraft can depart from any runway regardless of departure procedure. While the fully flexible case is operationally infeasible, it sets a best-case performance bound. Under fully flexible operations there would be a reduction in total surface-delay time of 50 hours, doubling the benefit of Dynamic Mapping 2. The results depicted in Figures 18 through 20 provide an indication of scale for later when runway balancing is introduced to Dynamic Mapping 1 and Dynamic Mapping 2.

In the context ARMT, simulation results indicate that dynamic-demand allocation has limited benefits in 2010 and 2015. However, by 2030, the benefit of demand allocation will be substantial, accounting for a 1-minute reduction in surface delays per aircraft for Dynamic Mapping 2. In contrast to Dynamic Mapping 2, when compared to traditional runway assignments, Dynamic Mapping 1 results in longer surface delays (see Figure 20). The poor performance of Dynamic Mapping 1 is due to only allowing for four flexible departure groups. On average, these four changeable departure groups only account for 10% of all departure demand, of which at most 8% can be reallocated from 18L to 17R

(departure groups 2 and 11), and 2% from 17R to 18L (departure groups 3 and 10). And because of the persistent 44%–56% demand imbalance, moving 2% of demand does not provide sufficient leeway to balance operations. During periods of demand imbalance, greater latitude is required for reassigning departure fixes to runways to allocate departure demand evenly. In the case of Dynamic Mapping 2, where all northern and southern departures are changeable, up to 25% of the flights can be reassigned to runways.

Moving forward, two options are considered for runway balancing at DFW: Dynamic Mapping 1 and Dynamic Mapping 2. Dynamic Mapping 1 is selected because it approximates current-day operations with regards to load balancing. Additionally, because the simulation results for Dynamic Mapping 1 are within range of the static baseline mapping and the traditional mapping, the dynamic variant can serve as a stand-in for these options. Dynamic Mapping 2 provides a comparison reference to Dynamic Mapping 1 to better understand how more aggressive departure-demand allocation influences the benefits of ad hoc runway balancing methods like OpenSlot and GreedySlot.
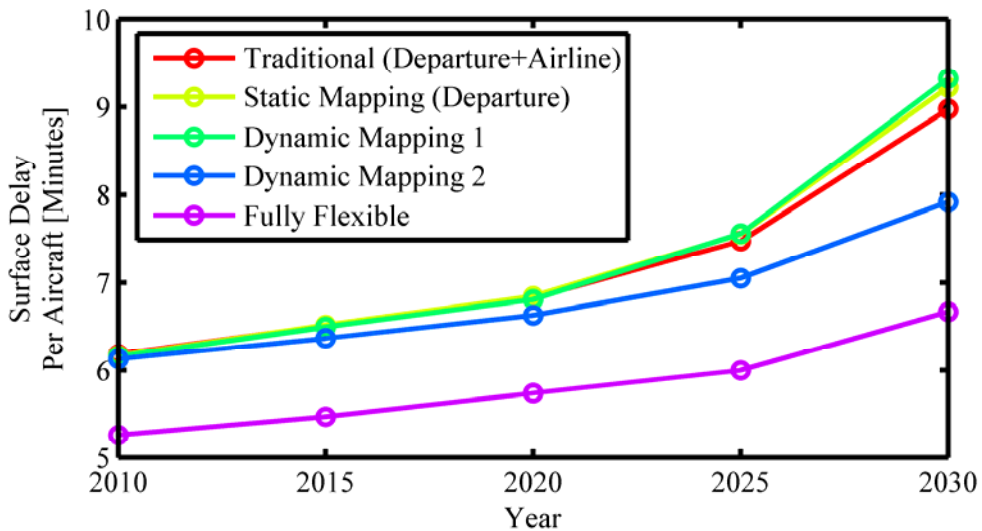

Figure 18 . Average surface delay per aircraft under different runway assignment rules at DFW.
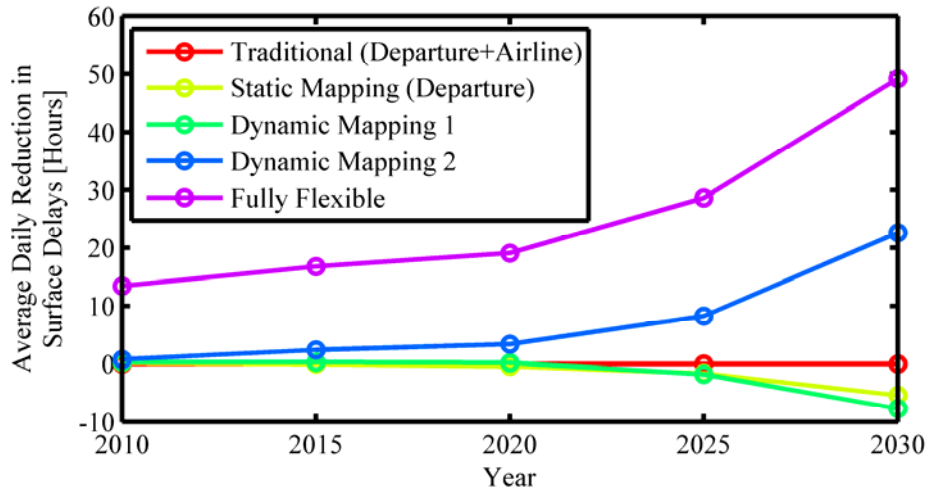
30

Figure 19. Average daily reduction in surface delays relative to traditional runway assignments at DFW.
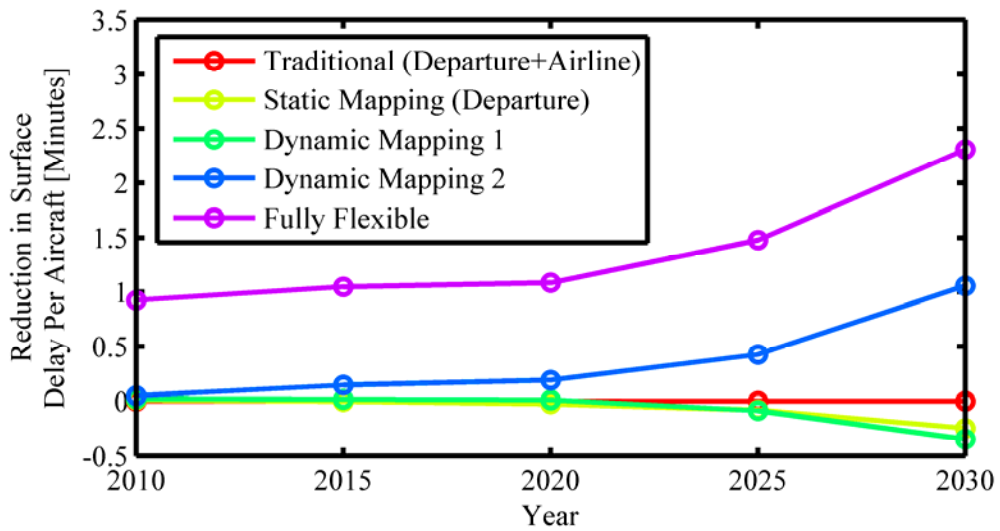
Figure 20. Average reduction in surface delay per aircraft relative to traditional runway assignments at DFW.

31
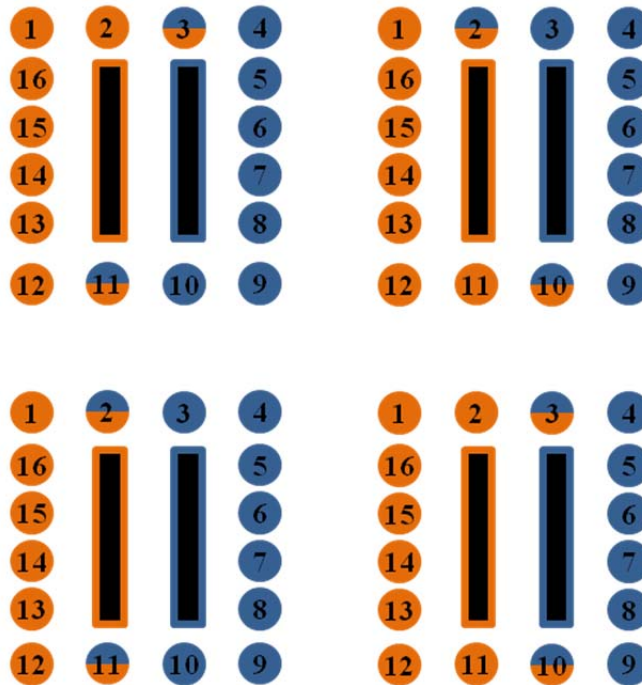
Figure 21. All flexible options for DFW Dynamic Mapping 1.

To introduce runway balancing into operations, potential flexible departure groups must first be selected. In the case of Dynamic Mapping 1 and Dynamic Mapping 2, all departure groups that are changeable are deemed to be potentially flexible. To be consistent with the TFDM CONOPS and allow for safe operations, not all potential flexible departure groups can in fact be flexible. Otherwise, if the case of adjacent flexible departures were permitted (e.g., departure groups (2) and (3) are flexible), then it becomes possible for flight paths to cross, which would constitute a safety hazard. As a reference example, the four flexible options for Dynamic Mapping 1 are illustrated in Figure 21. To maximize the likely usage of runway balancing, the flexible mapping with the potential to off-load the busiest runway is selected. For example, if there is a 30-minute demand imbalance with the western runway 18L in excess, then the optimal mapping selects departure groups (2) and (11) to be flexible, potentially shifting excess demand to 17R, the eastern runway. With each 30-minute update of the departure fix-to-runway mapping, new flexible departure groups are selected.

Figure 22. Reduction in total aircraft surface-delay time with runway balancing for Dynamic Mapping 1 at DFW.

Figure 23. Reduction in total aircraft surface-delay time with runway balancing for Dynamic Mapping 2 at DFW.

Figures 22 and 23 provide the additional benefit of runway balancing when applied to Dynamic Mapping 1 and Dynamic Mapping 2. As such, the total benefit of runway balancing and Dynamic
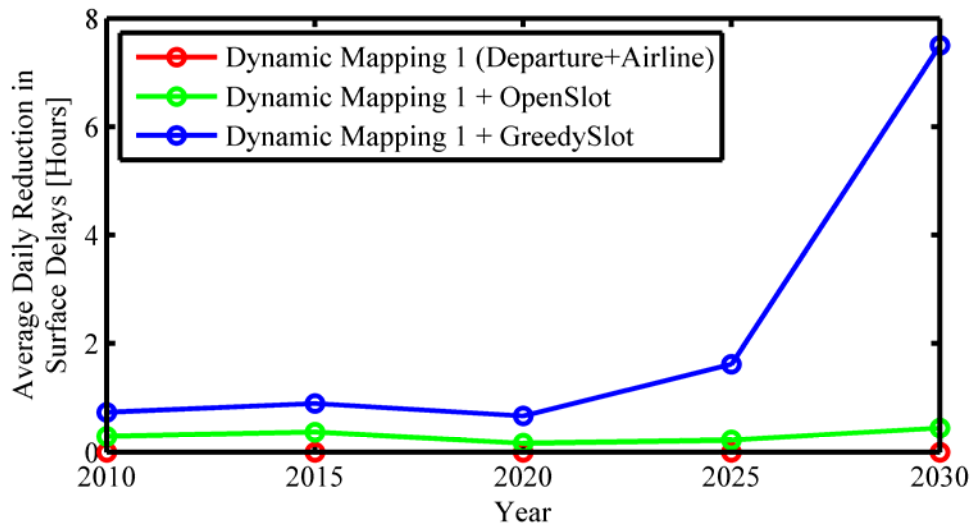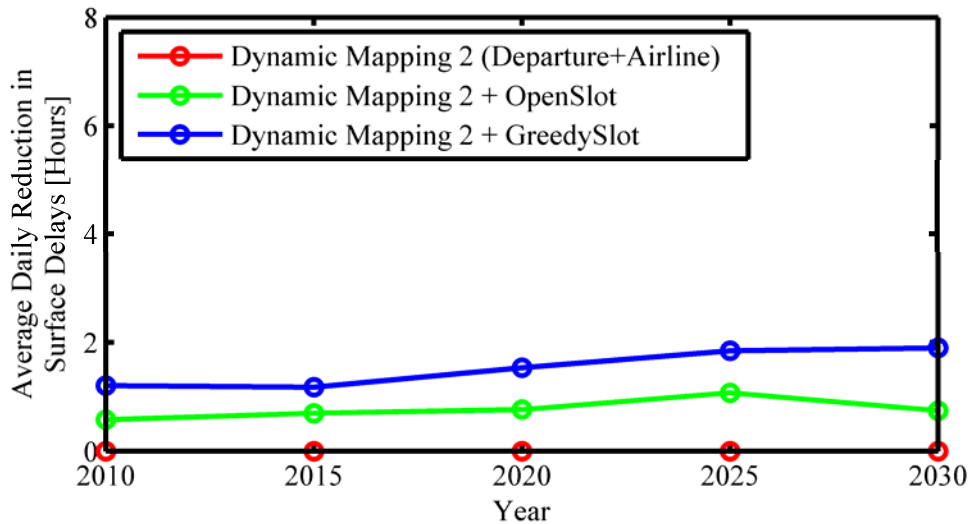
Mapping 1 (Dynamic Mapping 2) is the sum of the benefits presented in Figure 19 for the Dynamic Mapping option with the added benefit of runway balancing in Figure 22 (Figure 23). As indicated in Figure 22, when allocating demand with Dynamic Mapping 1, the introduction of runway balancing reduces the total aircraft surface-delay time by less than 2 hours for the average day until 2030. In 2030, the benefit of runway balancing increases to 7.5 hours per day for GreedySlot, while remaining steady for OpenSlot. When departure-demand allocation is executed via Dynamic Mapping 2, the introduction of runway balancing provides an average daily benefit of about 1 hour for OpenSlot and 2 hours for GreedySlot. Focusing on GreedySlot, the corresponding daily average per-aircraft saving in 2030 is limited to 5.4 seconds per aircraft. Compared to the total daily surface-movement time for Dynamic Mapping 1, a reduction of 7.5 hours with GreedySlot represents an overall reduction of 3.87%; for Dynamic Mapping 2, the surface-delay time savings is less than 1.13%. The relatively low benefit of runway balancing with Dynamic Mapping 2 is a result of Dynamic Mapping 2 evenly distributing demand to both runways, thereby limiting any benefit from balancing.

A comparison of the derived benefits of runway balancing for Dynamic Mapping 1 and Dynamic Mapping 2 supports the conjecture that more aggressive load-balancing through operational decision-making (i.e., dynamic departure-demand allocation) decreases the need for tactical actions via runway balancing.

Figure 24. Average daily number of flexible aircraft for Dynamic Mapping 1 at DFW.
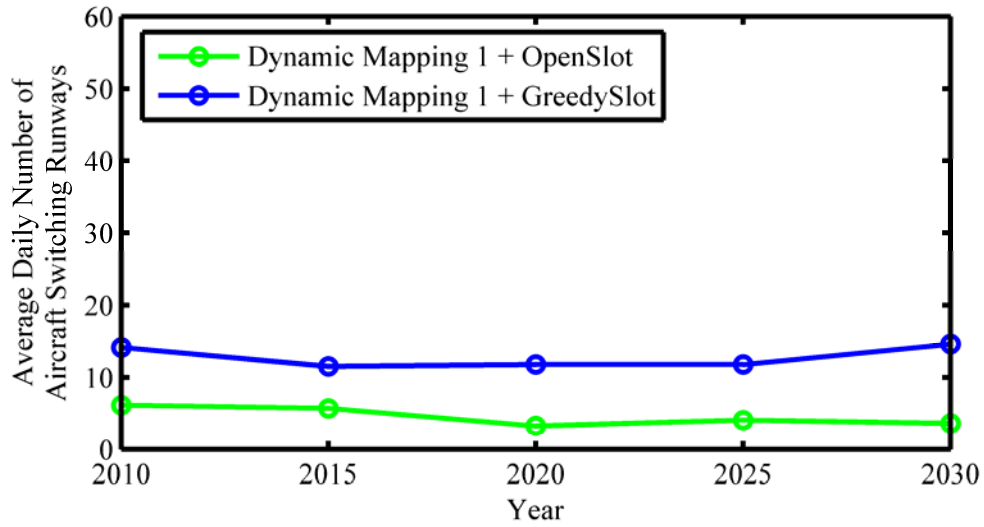
Figure 25. Average daily number of runway reassignments for Dynamic Mapping 1 at DFW.

The reduction in surface-delay time for both OpenSlot and GreedySlot derives from alternative runway assignments for a limited set of aircraft. Over the 2010–2030 simulation years, under Dynamic Mapping 1, the average number of aircraft departing on flexible departure groups increases from 43 aircraft to 53 aircraft for OpenSlot and GreedySlot see Figure 24). For both GreedySlot and OpenSlot, few of the opportunities for reassigned aircraft are taken. As shown in Figure 25, OpenSlot averages fewer than 10 reassignments on a given day, and GreedySlot averages less than 15. Over a range of simulation years, even with a near-constant number of aircraft reassignments, the trickle-down benefit of each GreedySlot movement increases significantly; the 7.5 hours of saved surface-delay time is the result of 15 aircraft being assigned to different runways.
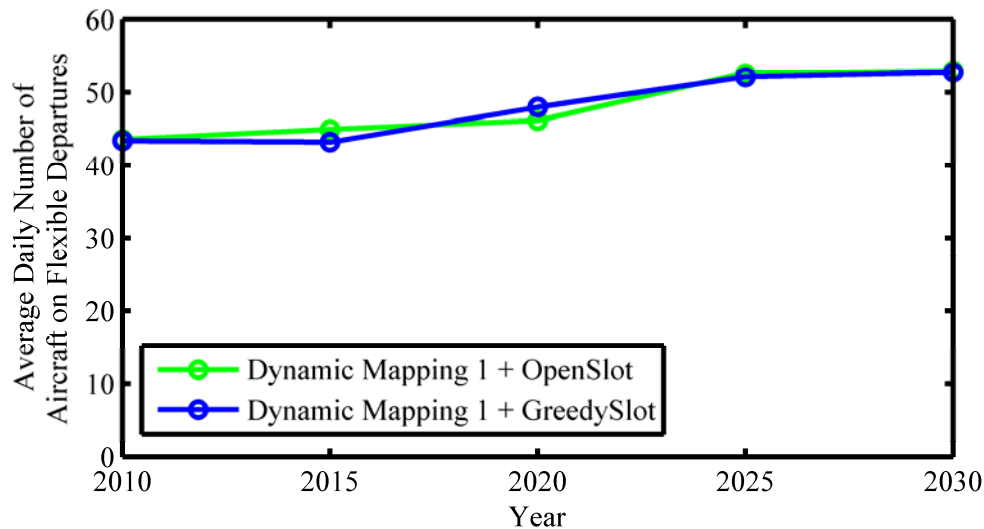
Figure 26. Average daily number of flexible aircraft for Dynamic Mapping 2 at DFW.

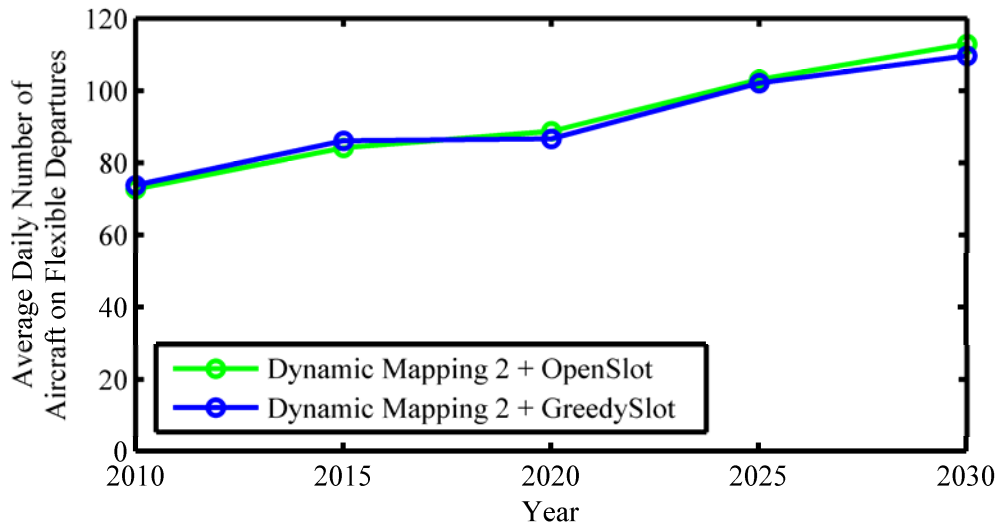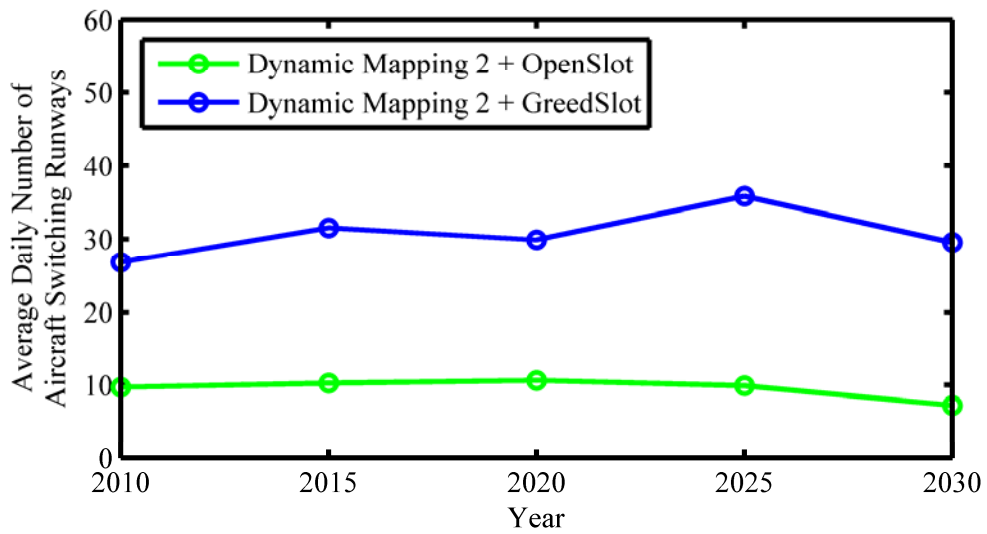*Figure 27. Average daily number of runway reassignments for Dynamic Mapping 2 at DFW.*

With Dynamic Mapping 2, the number of aircraft that can potentially be reassigned runways is substantially larger when compared with Dynamic Mapping 1 (compare Figure 24 and Figure 26). Because GreedySlot has relaxed conditions relative to OpenSlot, it is more often able to take advantage of the flexibility. As shown in Figure 27, with Dynamic Mapping 2 GreedySlot reassigns between 27 and 36 aircraft daily from 2010 to 2030, while OpenSlot is limited to about 10. It is also worth noting that as traffic demand increases, the applicability of OpenSlot decreases because of fewer available departure slots. From 2010 to 2030, for both Dynamic Mapping 1 and Dynamic Mapping 2, the number of aircraft switching with OpenSlot decreases annually.

Comparatively, none of the dynamic mappings or runway balancing combinations is able to significantly reduce surface-delay times to values close to the fully flexible option. However, with more options in load balancing and greater flexibility in aircraft reassignments, departure-demand allocation and runway balancing enables a maximum decrease in surface delays of 12.9% at DFW by 2030.

3.2 MCO

Like DFW, MCO has four parallel runways, two of which are often used for independent departure operations. However, unlike DFW, operations at MCO are not nearly as structured or balanced. The airport's two most common runway configurations along with annual departure operation counts are listed in Table 6. For the most common configuration, historical aircraft runway assignments can only be predicted with 59% certainty when only using the aircraft's filed departure procedure. However, when both the departure procedure and airline identifier are considered, the historical runway prediction rate increases significantly to 89%. The importance of airline identifier is closely linked to the terminal gate. Observation of the airport diagram in Figure 28 shows that aircraft departing from Airside 2 and Airside 4 probably have a preference to depart from 17R, instead of taxiing to 18L. Departing on 17R reduces the taxi distance by approximately 1.5 miles. And since departure demand at MCO is significantly less that at DFW, air traffic control can probably accommodate such preferences.

TABLE 6

Predicting Airport Operations at MCO

Departure Configuration	Total Number of Departures	RWY Prediction Rate [Percent] (Departure Group)	RWY Prediction Rate [Percent] (Departure Group + Airline)
17R, 18L	69603	59.6	88.9
35L, 36L	32243	87.1	93.4

The static baseline departure fix-to-runway mapping for MCO in the 17R-18L configuration is depicted in Figure 29; departure groups (1)–(6) map to 17R, while all others map to 18L. The dominant departure group is (3), which accounts for 40% of all demand. The skewed demand distribution complicates balancing, as sequential departures on 17R departing to (3) can significantly decrease the overall departure rate due to extended aircraft spacing along the same routes. Overall, there is a runway imbalance (34% to 66%) that could be difficult to correct even with dynamic departure-demand allocation.
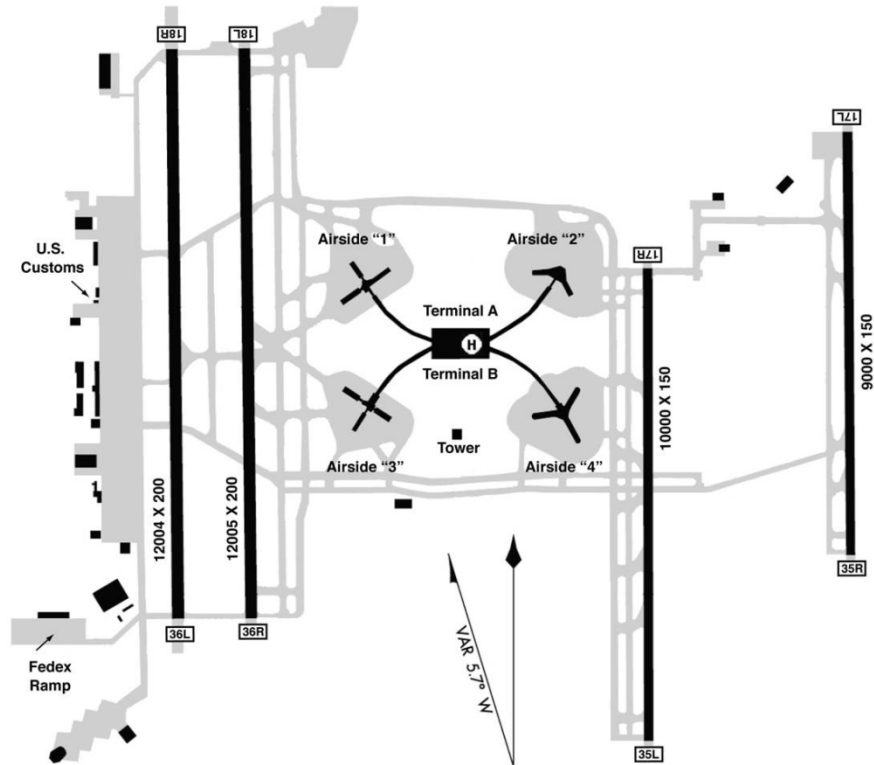
Figure 28. Airport diagram for MCO.
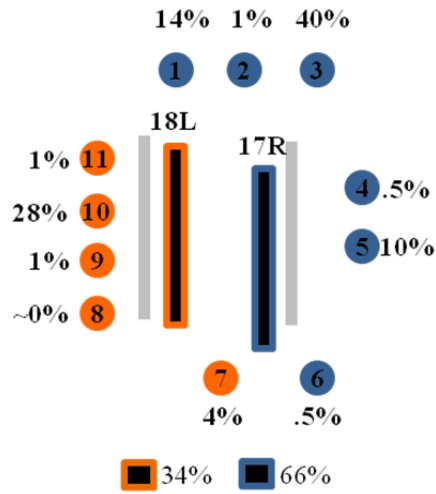
Figure 29. Static baseline mapping for MCO.

Starting from a 17R-18L configuration, four different candidate options for baseline runway mapping assignments at MCO are considered: a traditional mapping (using the departure procedure and airline), a static baseline mapping depicted in Figure 29, and two dynamic mappings illustrated in Figure 30. Dynamic Mapping 1 has changeable groups (1), (5), and (6), while for Dynamic Mapping 2, groups (1), (2), (3), (7), (8), and (9) are changeable. For Dynamic Mapping 1 and Dynamic Mapping 2, the changeable departure groups are also allowed to be flexible when runway balancing is introduced.
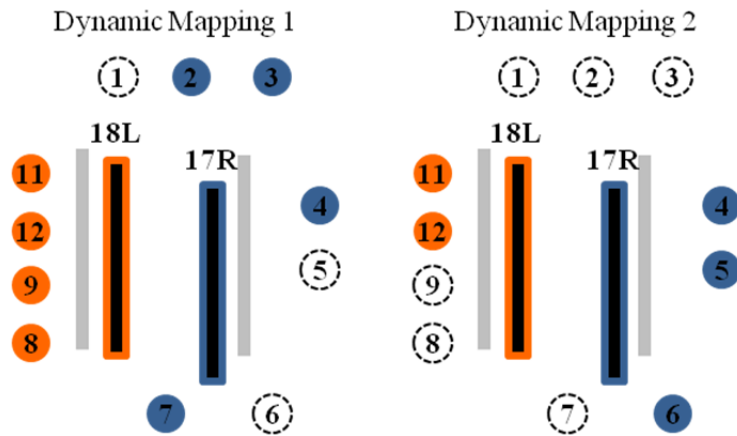
Figure 30. Dynamic mappings for MCO.
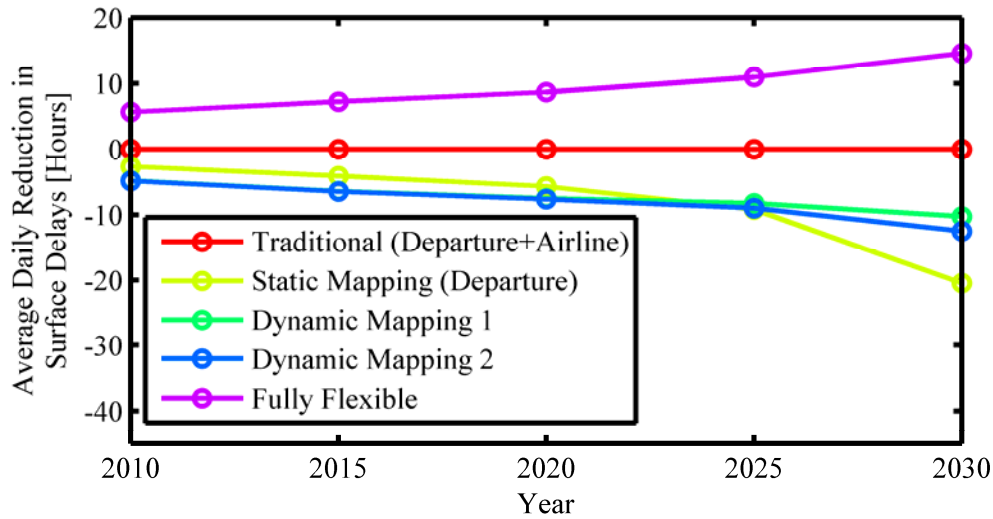
Figure 31. Average daily reduction in surface delays relative to traditional runway assignments at MCO.

Comparing the runway assignment options in Figure 31, the traditional runway mapping actually performs the best (excluding the fully flexible mapping), even surpassing the dynamic mapping options. It is likely that airline preferences already lead to relatively well-balanced demand at the runways. The static baseline mapping yields the largest delays by 2030, which is expected since by default the mapping is imbalanced with a 34%–66% split across 18L and 17R. Moving forward, Dynamic Mapping 2 and traditional mapping are maintained as part of the analysis, with the Dynamic Mapping 2 reserved as the baseline when introducing runway balancing. Dynamic Mapping 2 is selected over Dynamic Mapping 1, despite being outperformed in some cases, because it allows for a greater number of flexible aircraft departing from group (3), which is associated with 40% of all flights.
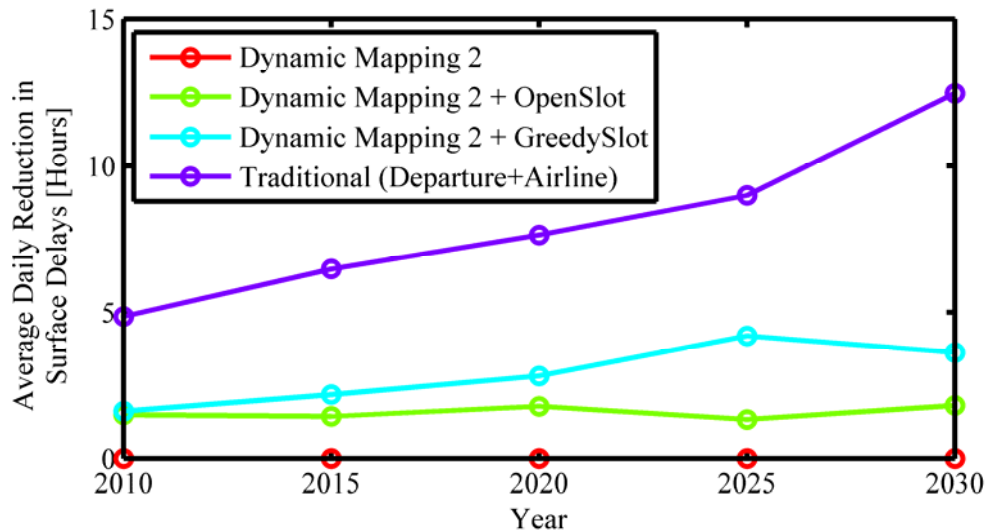
Figure 32. Relative benefit of runway balancing at MCO.

Figure 32 plots the average daily reduction in surface-delay time with the inclusion of OpenSlot and GreedySlot relative to Dynamic Mapping 2 (the total benefit of runway balancing with Dynamic Mapping 2 would require summing the results depicted in Figure 31 and Figure 32). Also included in Figure 32 is the relative benefit of the traditional mapping assignment compared to Dynamic Mapping 2. Even when applying runway balancing in conjunction with Dynamic Mapping 2, surface-delay times are unable to match surface delays observed with traditional runway assignments. Thus, the simulations indicate that operations at MCO are likely already well-balanced through the preferences and actions of airlines and air traffic control; the inclusion of greater structure through departure-demand allocation would only serve to inhibit efficient operations.

3.3 LAX

As shown in the airport diagram in Figure 33, LAX has four parallel runways divided by a central terminal. The most common configuration for the airport is to run all four runways for both arrivals and departures. That said, even when in a 24L-24R-25L-25R configuration not all runways are equally used for both arrival and departure operations. Typically, departures will strongly favor 24L and 25R. And because of the relatively long taxi times between the north and south sides the airport, and the addition of a new international terminal slowing traffic near the crossing taxiway, there is a strong preference for flights to depart on the same side of the airport as their terminal gate. So for United Airlines, who is located in the south-east terminals 7 and 8, their flights typically depart from runway 25R. Meanwhile, Southwest Airlines and American Airlines located in terminals 1 and 2 depart on 24L. The preference to

depart on the closest runway is expressed in the prediction error for runway assignments shown in Table 7; with only knowledge of the departure procedure, at best the correct departure runway is selected 68% of the time. Yet, when information about the airline is included, the correct runway can be predicted with up to a 93% success rate.
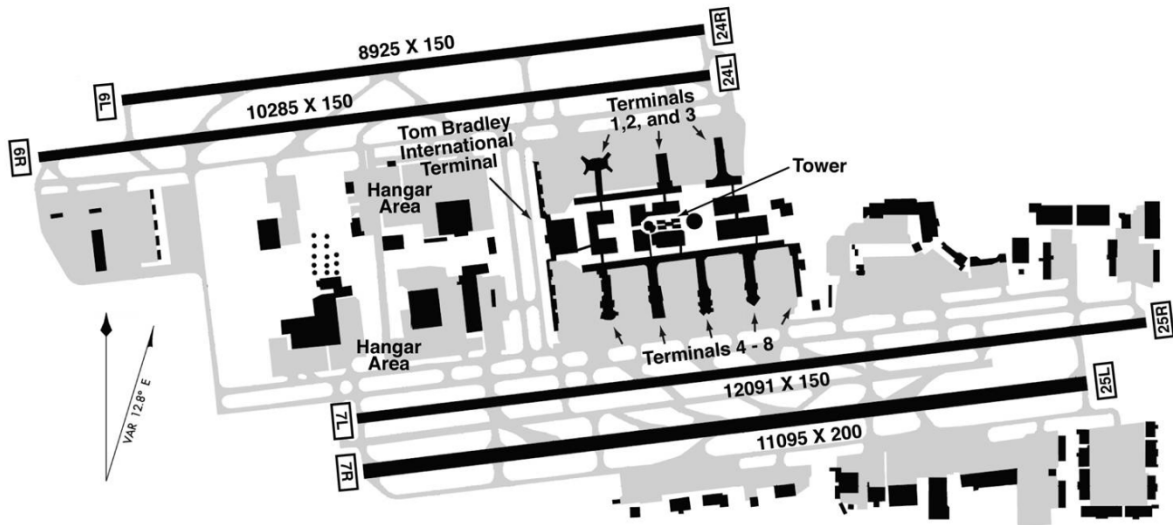
Figure 33. Airport diagram for LAX.

TABLE 7

Predicting Airport Operations at LAX

Departure Configuration	Total Number of Departures	RWY Prediction Rate [Percent] (Departure Group)	RWY Prediction Rate [Percent] (Departure Group + Airline)
24L, 24R, 25L, 25R	146523	60.9	83.5
24L, 25R	113318	64.3	93.4
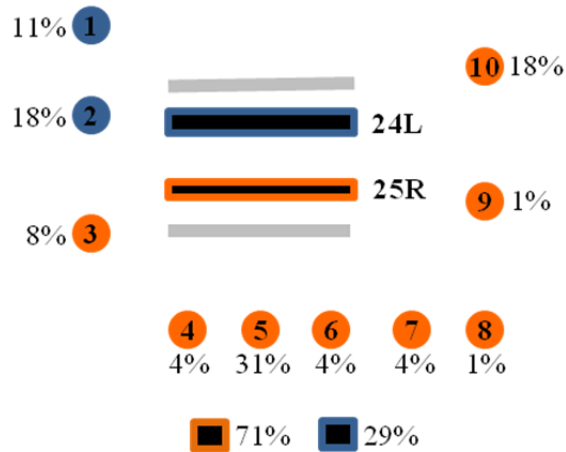
Figure 34. Static baseline mapping for LAX.

When in a 24L-25R departure configuration, the static baseline departure fix-to-runway mapping for LAX is illustrated in Figure 34; however, the error rate for this mapping is large (as indicated in Table 7). Geographically, departure groups (1) and (2) are associated with flight paths heading north along the California coast or west for a short time period; group (3) departs out to the ocean; and the remaining departure groups curl around south to east.

It is worth noting that under current traffic conditions air traffic control is able to depart aircraft on either runway. However, it is unknown if this type of flexible operation is sustainable at future traffic levels. Currently the number of aircraft departing on alternative runways can be roughly estimated using the average number of daily departure operations from Figure 14 and the error rates in Table 7. Approximately 30% of the 800 daily flights depart on an alternate runway, which accounts for about 240 departures.
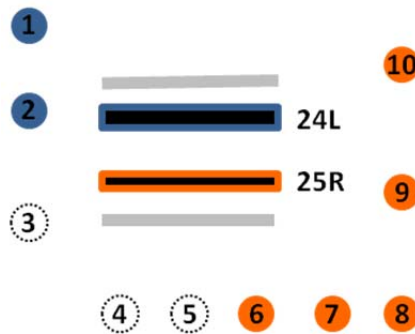
Figure 35. Dynamic mapping for LAX.

43

For a 24L-25R departure configuration, three different candidate options for runway mapping assignments at LAX are considered: a traditional mapping (using the departure procedure and airline), a static baseline mapping depicted in Figure 34, and a dynamic mapping with changeable groups (3), (4), and (5) as illustrated in Figure 35. The same changeable departure groups are also allowed to be flexible when runway balancing is introduced. The goal with this mapping is to allow flights in departure group (5) to take-off from either 24L or 25R.
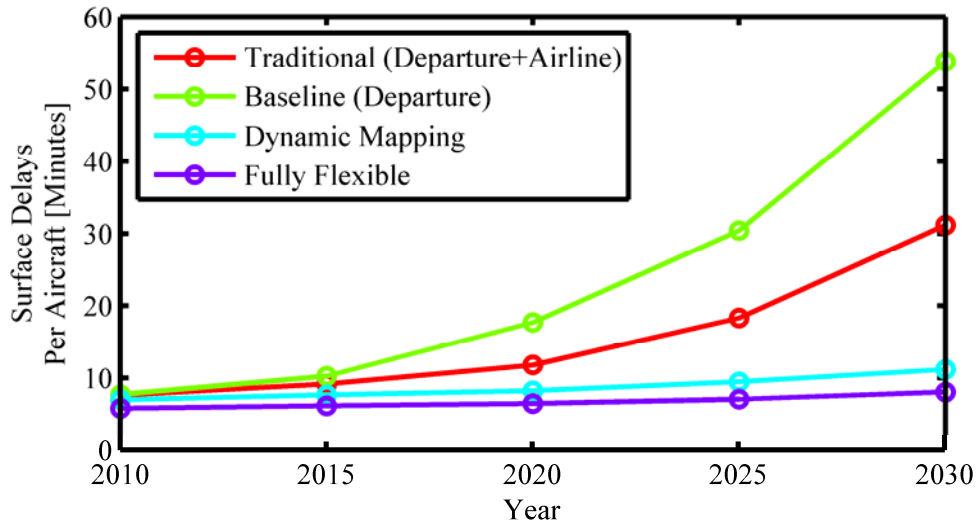
Figure 36. Average surface delay per aircraft at LAX.

The surface delay per aircraft at LAX under the different runway assignment and departure allocation schemes are shown in Figure 36. The surface-delay time measurements indicate that at present-day demands, excluding the fully flexible case, all departure-allocation and runway assignment approaches result in approximately the same average surface-delay per aircraft, between 7 and 8 minutes, which corresponds to 86 and 96 total hours of delay each day. If today's traditional approach to assigning runways continues to 2030, as traffic demand increases so are delays expected to increase, up to 300%. Delays will be up to 31 minutes per aircraft, translating to 810 total hours each day. When dynamic demand-allocation is applied to departure operations, surface delays decrease substantially, and actually approach the best-case fully flexible scenario. Furthermore, as indicated in Figure 36, the per-aircraft delay only peaks at just above 10 minutes in 2030.
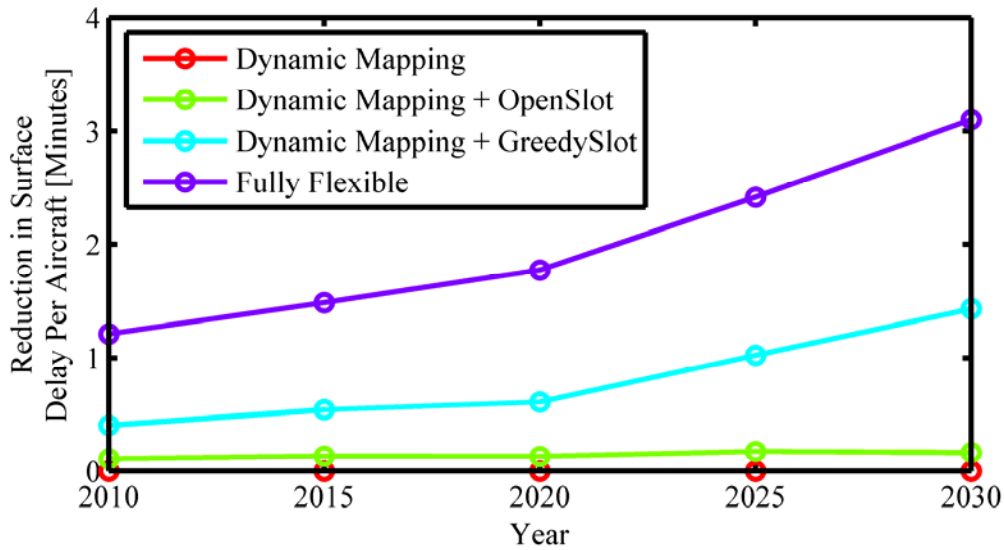
Figure 37. Reduction in total surface-delay time at LAX with runway balancing.

As shown in Figure 37, the introduction of OpenSlot or GreedySlot is able to reduce the surface delays for the average day. In the case of GreedySlot for the year 2030, runway balancing reduces surface delays by an additional 1.5 minutes per aircraft, which accounts for an additional 12.9% decrease in surface delays. When compared to the traditional case at LAX, which is similar to present day operations, the combined benefit of the Dynamic Mapping (~20 minutes per aircraft) and GreedySlot (1.5 minutes per aircraft) approaches 70% by 2030. However, as previously noted, traditional runway assignment is unsustainable, and will likely be forced to change under traffic demand and the resulting delays; ARMT can help support that change. That said, the scale of the 12.9% reduction in surface-delay time stands in contrast to the benefits observed at DFW and MCO, which in the best cases did not exceed 8 hours or about 3.87% reductions in delays.

3.4 CUMULATIVE BENEFITS OF DEPARTURE-DEMAND ALLOCATION AND RUNWAY BALANCING

The simulation results for DFW, MCO, and LAX presented in Sections 3.1–3.3 provide average daily reductions in surface-delay time for the twelve SWAC simulation days assuming VMC conditions. The cumulative benefit of runway balancing at each airport can be calculated by scaling the daily results by the number of VMC days in a year (provided in Table 8), then integrating from 2017 to 2030.

TABLE 8

Number of Days of VMC per Year

Airport Site	# Days of VMC between Hours of 8 AM and 11 PM for 2013 [ASPM]
DFW	225
LAX	183
MCO	233

Figure 38. Cumulative reduction in surface delays from departure-demand allocation at DFW.
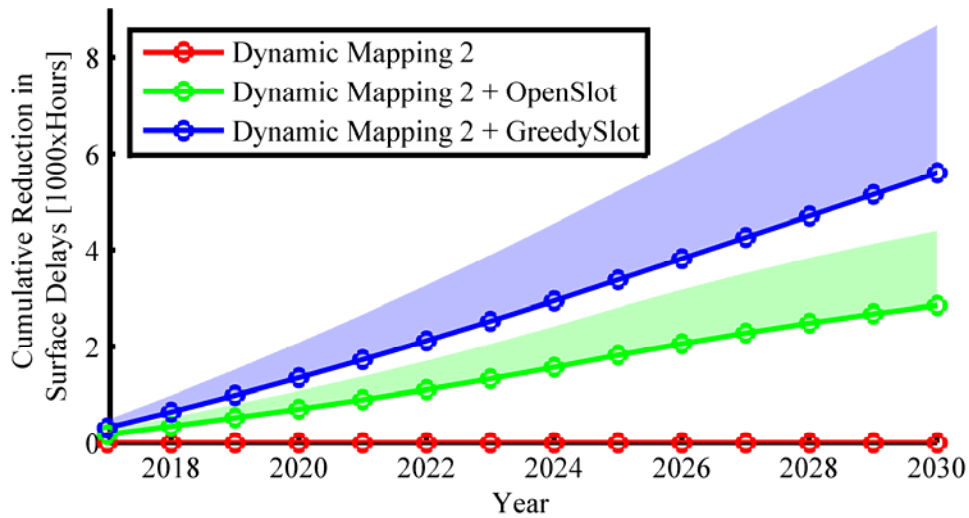
Figure 39. Cumulative reduction in surface delays from runway balancing at DFW.

Using the scaling in Table 8, the cumulative benefits of the departure-demand allocation at DFW are provided in Figure 38. Results indicate that by 2030, Dynamic Mapping 2 is expected to save over 30,000 hours when compared to traditional runway assignments. The shaded portion associated with each line indicates the potential benefit when the delay savings are scaled over the complete year, not just over the VMC days – the upper boundary serves as a proxy for consideration of IMC days (assuming equal benefits on these days). As such, it is possible that Dynamic Mapping 2 in fact saves 47,000 hours when integrated from 2017 until 2030. The cumulative benefits of runway balancing are not as significant as departure-demand allocation at DFW. As shown in Figure 39, by 2030 OpenSlot and GreedySlot would save approximately 2,800 and 5,600 hours, respectively.

Figure 40. Cumulative reduction in surface delays from demand allocation at LAX.
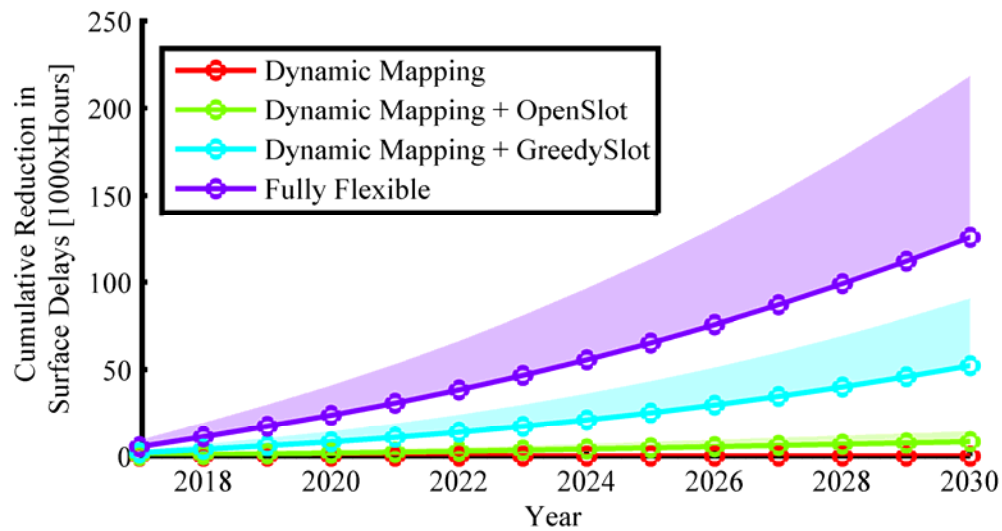
Figure 41. Cumulative reduction in surface delays from runway balancing at LAX.

At LAX, the cumulative benefits of the departure-demand allocation from 2017–2030 reach over 515,000 hours (possibly up to 890,000 hours), see Figure 40, which is far more than that observed at DFW. Likewise, the cumulative benefits of runway balancing are significant for GreedySlot. As shown in

Figure 41, GreedySlot could save an additional 52,200 hours from 2017 to 2030. For OpenSlot, the benefits are significantly less, yet in line with the benefits observed at DFW. Simulations indicate that OpenSlot will save 8,200 hours over the same 13-year period.

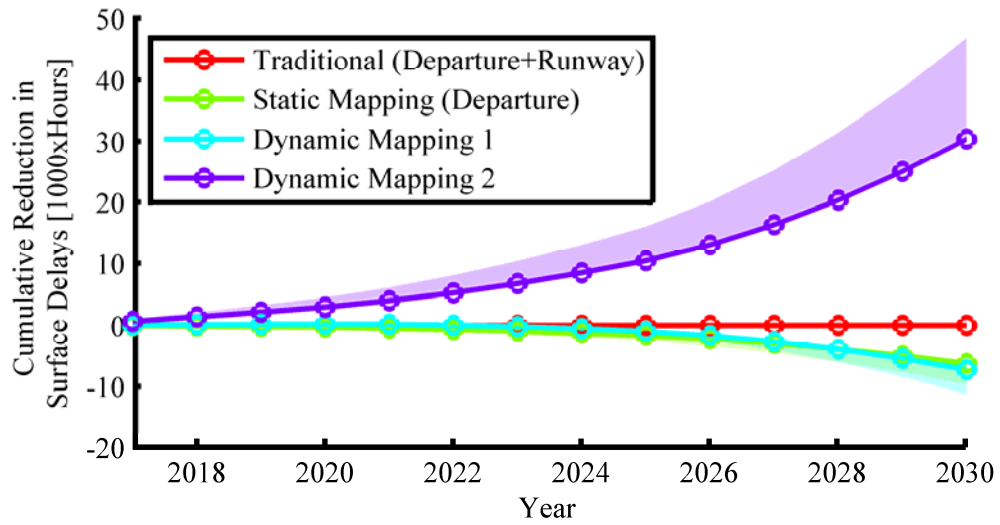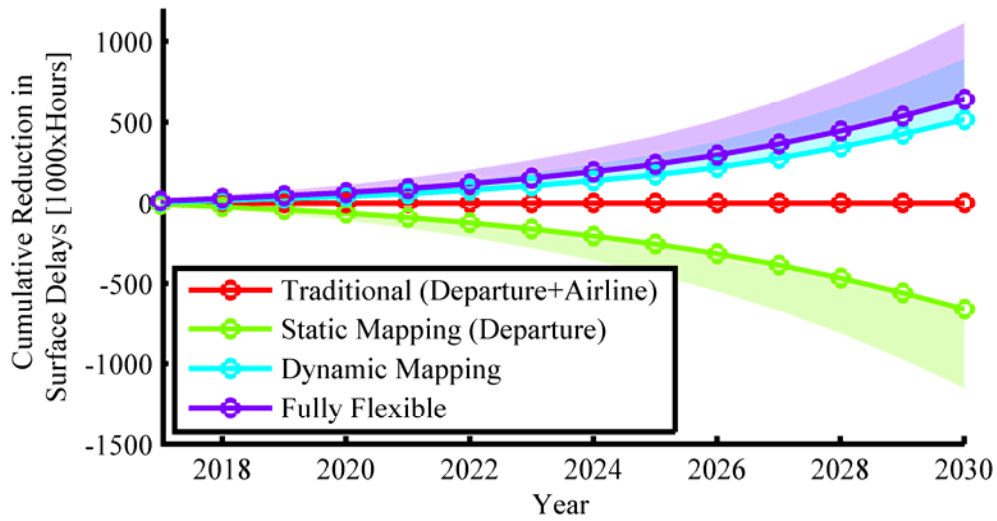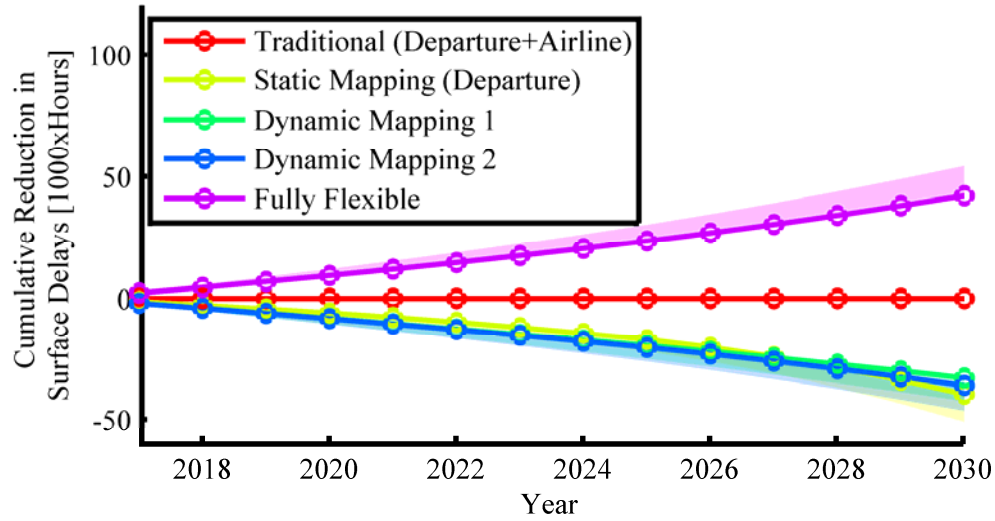*Figure 42. Cumulative reduction in surface delays from demand allocation at MCO.*
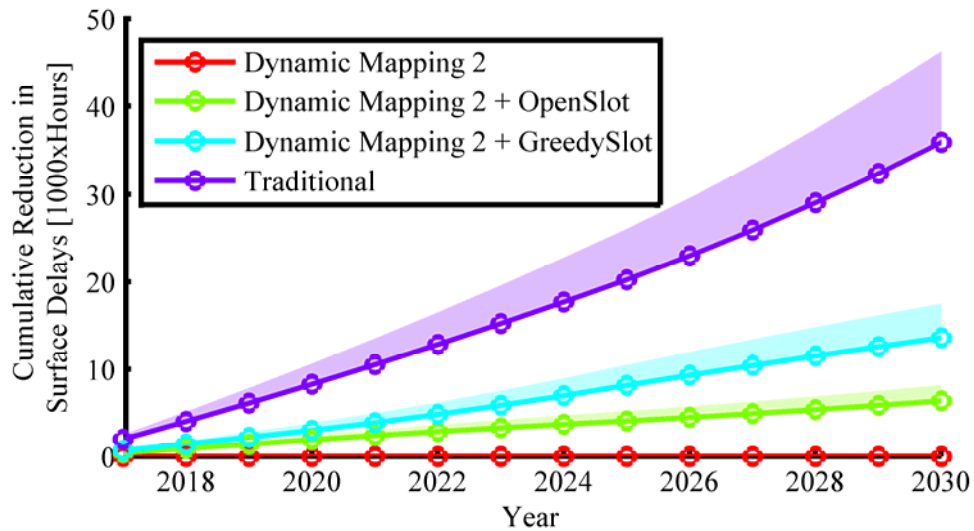
Figure 43. Cumulative reduction in surface delays from runway balancing at MCO.

The dynamic departure-demand benefits at MCO are nonexistent when compared to traditional runway assignment. As indicated in Figure 42, the only runway assignment scheme that out-performs traditional runway assignment is the operationally infeasible fully flexible mapping. That said, there is no guarantee that at future traffic levels the traditional runway assignment scheme is still sustainable with regards to controller workload and safety. At some point, a more structured approach to runway assignments may be required, instead of allowing aircraft to depart on any runway, regardless of departure procedure. When runway balancing is added to departure-demand allocation, there is an additional benefit, as indicated in Figure 43; however, the added benefit does not surpass traditional runway assignments.
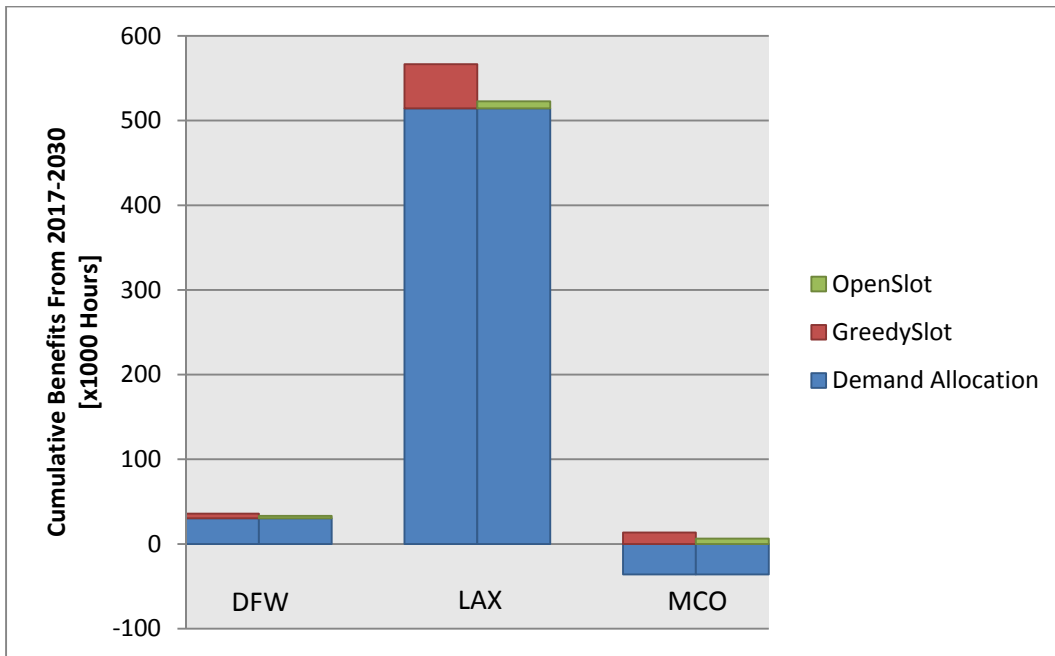
Figure 44. Summary of cumulative benefits for dynamic demand allocation and cumulative benefits of runway balancing at DFW, LAX, and MCO.

Figure 44 summarizes Figures 38 through 43 by depicting the cumulative benefits from 2017 to 2030 of both demand allocation and runway balancing at DFW, LAX, and MCO. In total, the magnitude of influence of demand allocation surpasses that of runway balancing (for both GreedySlot and OpenSlot). As shown in Figure 44, at DFW and LAX the positive effect of implementing demand allocation (30,000 hours and 515,000 hours) is significantly larger than the greatest observed benefit of GreedySlot at DFW and LAX (5,600 hours and 52,000 hours). Even in the case at MCO where demand allocation yields a negative benefit of 35,000 hours, the effect of runway balancing is more than a factor of two smaller (13,600 hours).

Furthermore, for the three simulated airports, GreedySlot consistently out-performs OpenSlot, yielding greater reductions in surface-delay time at each airport. As shown in Figure 44, the difference between GreedySlot and OpenSlot can be quite significant. At DFW and MCO, the benefit of GreedySlot is approximately double that of OpenSlot. And at LAX, the delay savings from GreedySlot is five times greater than OpenSlot. GreedySlot is not necessarily the only option or best option for runway assignment for all airports. While GreedySlot provides a positive impact, in the case of MCO, when compared to current-day operations, traditional runway assignments provide substantially greater benefits than dynamic departure-demand allocation with runway balancing (see Figure 43). Thus, for MCO, the advantage of GreedySlot over OpenSlot may be irrelevant.

Comparatively, the derived benefits of departure-demand allocation and GreedySlot runway balancing at LAX far exceed those observed at DFW and MCO; three times and greater. There are two key differences that may account for such a distinction. First, demand across the different departure groups at LAX are imbalanced and skewed when compared to DFW (see Figure 16 for DFW and Figure 34 for LAX demand distributions). One departure group accounts for 31% of all traffic at LAX, while at DFW demand is evenly spread out, with no one departure group carrying more than 11% of the demand. So while departure-demand allocation with dynamic mapping is able to yield 50%–50% demand across departure runways at DFW, achieving equal runway demand at LAX is more challenging. It is the resulting runway imbalance that GreedySlot and OpenSlot are able to take advantage of and help correct. On the other hand, while MCO has a skew demand across its departure groups, overall departure demand at the runways is relatively low. So while there are opportunities for runway balancing, the aggregate benefit of each reassigned aircraft is limited. As such, the second key distinction LAX has is significant departure demand.

3.5   SENSITIVITY TO TAXI TIMES

For the DFW, MCO, and LAX simulations, taxi times are calculated using a random forest regression based on historical data. In order to help validate the results of the benefit analysis, it is necessary to examine if the taxi time models skew the results or obfuscate any additional benefit of departure-demand allocation or runway balancing. To test the hypothesis that the demand-departure allocation and runway balancing results are insensitive to taxi times, a set of experimental simulations are executed. In the experiment, instead of using the taxi time regression model, all taxi times are assumed to be 5 minutes (with ±1 minute uncertainty padding used for OpenSlot), regardless of departure gate, runway assignment, and surface congestion. Forcing uniform taxi times also tests how sensitive GreedySlot and OpenSlot are to future airport changes (e.g., new taxiways). If the simulation results are insensitive to the change in taxi times, then there is evidence that the benefits of departure-demand allocation and runway balancing are not heavily influenced by taxi times. As a corollary then, complex taxi time models may not be necessary for future runway balancing benefits analysis at other airports.

TABLE 9

Taxi Time Sensitivity on the Cumulative 2017–2030 Relative Benefits at DFW

	Random Forest Model [1000×Hours]	Uniform Model [1000×Hours]	Absolute Difference [1000×Hours]	PCT. Difference [%]	Avg. Daily Difference [Hours]
Dynamic Mapping 2	30.3	34.43	4.13	4.1	0.81
Dynamic Mapping 2 + GreedySlot	5.6	5.81	0.21	0.2	0.04
Dynamic Mapping 2 + OpenSlot	2.85	1	1.85	64.5	0.36

TABLE 10

Taxi Time Sensitivity on the Relative Benefits at LAX

	Random Forest Model [1000×Hours]	Uniform Model [1000×Hours]	Absolute Difference [1000×Hours]	PCT. Difference [%]	Avg. Daily Difference [Hours]
Dynamic Mapping 1	514.4	489.7	24.7	4.8	4.8
Dynamic Mapping 1 + GreedySlot	52.2	56.5	4.3	−8.3	0.8
Dynamic Mapping 1 + OpenSlot	8.3	12.0	3.7	−44.3	0.7

TABLE 11

Taxi Time Sensitivity on the Relative Benefits at MCO

	Random Forest Model [1000×Hours]	Uniform Model [1000×Hours]	Absolute Difference [1000×Hours]	PCT. Difference [%]	Avg. Daily Difference [Hours]
Dynamic Mapping 2	−35.9	−4.76	31.1	86.7	6.1
Dynamic Mapping 2 + GreedySlot	13.6	7.3	6.3	46.2	1.2
Dynamic Mapping 2 + OpenSlot	6.4	2.5	3.9	60.9	0.8

The experiment is executed at all three of the case study airports using GreedySlot and OpenSlot; the results are shown in Table 9, Table 10, and Table 11 for DFW, LAX, and MCO, respectively. Each table contains a comparison of the cumulative relative benefits of departure-demand allocation and runway balancing for each taxi time model (Random Forest and Uniform); the random forest results are the same as those in Figure 20, Figure 23, Figure 31, Figure 32, and Figure 37. As shown in Table 9, both the random forest taxi time model and the uniform taxi time model result in similar benefits for both departure-demand allocation and runway balancing at DFW. The most significant difference only results in an average daily difference in surface delay estimates for Dynamic Mapping 2 of .81 hours per day. For runway balancing at DFW, the difference in taxi time models only results in a .04 and .36 hour difference in benefits.

In the case of LAX (see Table 10), the different taxi models do not have a significant effect on the qualitative conclusions. The benefit difference for runway balancing between the two taxi time models is less than an hour per day for both GreedySlot and OpenSlot. As such, the observed benefit of runway balancing remains the same. For the case of dynamic departure-demand allocation, the difference in magnitude of estimated benefits is significant: 514,000 hours versus 489,000 hours. However, the estimated benefits are within 10% of each other, and furthermore, the benefit of departure-demand allocation is apparent regardless of taxi time model.

The most significant difference between the benefits for the two taxi models occurs at MCO (see Table 11). In the case of dynamic departure-demand allocation, when a uniform taxi model is applied, there is an average daily difference in surface delays of 6.1 hours compared to the random forest taxi time model. That said, the less accurate model appears to infer a greater benefit, which implies that even if the taxi models were improved it is unlikely that a previously hidden benefit to departure-demand allocation would appear at MCO. Traditional runway mapping still out-performs dynamic departure-allocation.

Compared to DFW, LAX and MCO have larger absolute differences in the observed benefits between the two taxi time models. The larger differences are a result of highly asymmetrical taxi times present at each airport; there exist terminals directly adjacent to one runway, but far from another. However, in the case of DFW, regardless of which active runway an aircraft departs from, taxi times are approximately equal, since the difference in taxi distance is less than three-quarters of a mile (see Figure 15).

Ultimately, the experimental simulation results suggest that taxi times are not critical to establishing the scale and qualitative nature of the potential benefits of dynamic departure-demand allocation and runway balancing at the airports considered. That said, for future simulations, airports with significant distinctions in taxi times between active runways should make use of a more accurate taxi time model when possible. When it is not possible to construct more complex taxi time models (due to a lack of data), it should be understood the accuracy of the potential benefits is reduced when taxi times are substantially different between available runways.

3.6   SUMMARY

The underlying procedure for the departure-demand allocation and runway-balancing benefits analysis is to simulate multiple days of traffic at an airport, with and without runway balancing; evaluate the resulting aircraft delays; and scale the results to yield an annual and cumulative delay savings.

According to the simulations, the cumulative benefit of departure-demand allocation could potentially reduce delays at LAX by over 514,000 hours, while runway balancing could further save 52,200 hours of delay. Expected delay savings at DFW are 30,300 hours and 5,600 hours for departure-departure demand allocation and runway balancing. In contrast to LAX and DFW, there is no observed benefit for departure-demand allocation and runway balancing at MCO when compared to traditional runway assignment.

The greatest positive impact of runway balancing is observed at LAX, where operations across departure procedures are unevenly distributed leading to demand imbalances at departure runways. Less potential benefit is observed at DFW due to a pre-existing departure demand that is almost balanced. At MCO, while asymmetric demands sometimes lead to runway imbalances, the excess departure capacity removes the expected benefit of formalized runway balancing. More so, there is the possibility that departure-demand allocation and runway balancing do not reduce departure delays over current-day operations at the airport.

4. OPPORTUNITIES FOR ADDITIONAL ANALYSIS

The benefits analysis in this report focused on DFW, MCO, and LAX. Of the three airports, LAX showed the greatest potential for improvement using departure-demand allocation and runway balancing. As noted in Section 3.3, the delay reductions observed at LAX derive from a combination of heavy departure demand and a skewed demand distribution over the departure groups. Unlike LAX, at DFW the large departure demand is easily balanced over the runways, resulting in fewer opportunities to apply runway balancing. And while MCO has a skewed demand over the departure groups, the magnitude of the total departure demand on each runway does not stress airport operations. The lack of departure pressure means that when departure-demand allocation and runway balancing are applied, there is limited to no benefit over current-day operations at MCO. Because of the large spread in observed benefits at each airport, there is a need to expand the study to other airports to better understand what features are associated with reduced surface-delay times. Only then can the NAS-wide benefit of runway balancing be understood and calculated. Therefore, this report concludes by advising that an additional analysis is performed on at least 5 targeted airports.
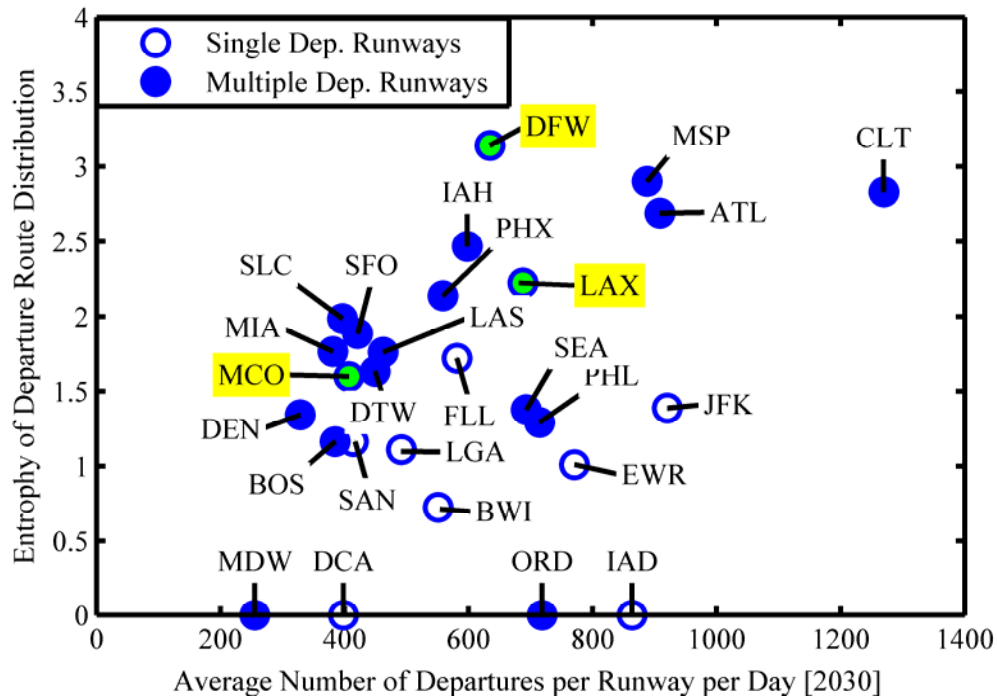

Figure 45. Departure demand and entropy of departure distributions for TFDM airports.

Based on the three case-study airports, two possible features have been identified as associated with runway-balancing benefits: departure demand at the runways, and an unbalanced departure group distribution. These two features can be used to identify additional airports to be considered in a future study.

One measure of departure demand is the number of daily departures per runway (where shared runways are counted as half a runway). If this measure is large, then the airport may be operating at or near capacity. However, if the daily departure demand per runway is low, then it is unlikely that there are departure queues or imbalances that would benefit from runway balancing. Without queues or reason to apply runway balancing, there is little total delay reduction; MCO is the exemplar. In Figure 45, the average number of daily departures per runway for the year 2030 is plotted for each TDFM airport. DTW, which is clustered near MCO and also has four parallel runways, will likely not profit from runway balancing because it lacks the necessary departure pressure. However, there is the potential for airports like SLC, MIA, and LAS to yield runway-balancing benefits; these airports run shared arrival/departure runway operations. During arrival or departure pushes at the airport, temporal imbalances may make runway balancing advantageous. At least one of these airports should be considered for future study – the delay reductions observed at one will likely apply to the others.

The second axis on Figure 45 is the *entropy* of the departure group distribution. Here entropy is an indication of how evenly traffic over the departures groups is distributed. The larger the entropy value the more evenly traffic is distributed, while a lower value indicates uneven and heavily trafficked departure groups. For example, at LAX, 30% of the traffic uses one departure procedure, and at MCO, the most traveled departure procedure accounted for 40% of the departure traffic – these airports have a lower entropy than DFW where traffic is almost evenly distributed over the departure fixes. For airports along the zero-axis (MDW, DCA, ORD, and IAD), there is no available information on the departure distributions. Airports like DFW and CLT spread demand evenly across all departure groups, and as a result, traffic loads can be evenly split across the runways. When traffic is split 50%–50% on the departure runways, there are fewer opportunities for runway balancing. So even though MSP, ATL, and CLT have significant departure demand, the demand can probably be balanced at the runways. Therefore, it is unlikely that the potential benefits of runway balancing at these airports are similar to LAX, but rather are more likely aligned to the delay reductions observed at DFW, but scaled according to the total demand.

The airports most likely to benefit the greatest from runway balancing are those on the lower-right of the figure relative to DFW and MCO. Considering departure demand per runway and entropy of the departure distribution, potential airports that may exhibit runway balancing benefits include IAH, SEA, PHL, and PHX. And while IAD is traditionally a single departure runway airport, it is likely that operations will expand to an additional runway to accommodate the increase in departure demand.

Because there still remain questions surrounding the NAS-wide benefit of departure-demand allocation and runway balancing, this report suggests additional case studies. According to the analysis in this section it is recommend that further studies focus on at least five other airports, which could possibly include SEA, PHL, IAH, MIA, MSP, and LAS.

This page intentionally left blank.

APPENDIX A
GENERATING MODEL COMPONENT BEHAVIOR

A.1   DEPARTURE PROCEDURE PREPROCESSING

The departure procedures provided by the SWAC model sometimes involve procedures inconsistent with the airport simulation model. These inconsistent departure procedures entail unpublished procedures or procedures that are not observed in historical flight plans. Example inconsistent departure procedures for MCO airport include WORMS, EPCOT, FSHUN, GAUSP, and LEWRD. None of these departure procedures have been observed in historical ETMS flight plans for 2013. Instead of using the departure information provided by the SWAC, an historical generative model is used to associate flights with a new departure procedure and transition fix. The historical generative model establishes preferred routings by the popularity of departure procedures found in historical ETMS flight plans. The historical generative model is based on a Random Forrest implementation.

A historical generative model for each airport is created by first gathering the following input features for all departing flights:

- Airline

- Three character shortened aircraft type (e.g., B72 for B727)

- Weight class (Small, Large, Full, Heavy)

- Engine type (Jet, Propeller, Turboprop)

- Distance to destination (great circle)

- Heading angle to destination (great circle)

- Planned cruise altitude

- Hour of day (local time)

The relevant output features for the historical generative model are a departure procedure and transition fix.

The airline and aircraft input features are mapped to one of 31 categories. The top 30 categories for each feature represent the top 30 airlines and top 30 aircraft types operating out of the airport of interest. All other airlines and aircraft that do not fall into the 30 most common categories are grouped together in an 'other' category.
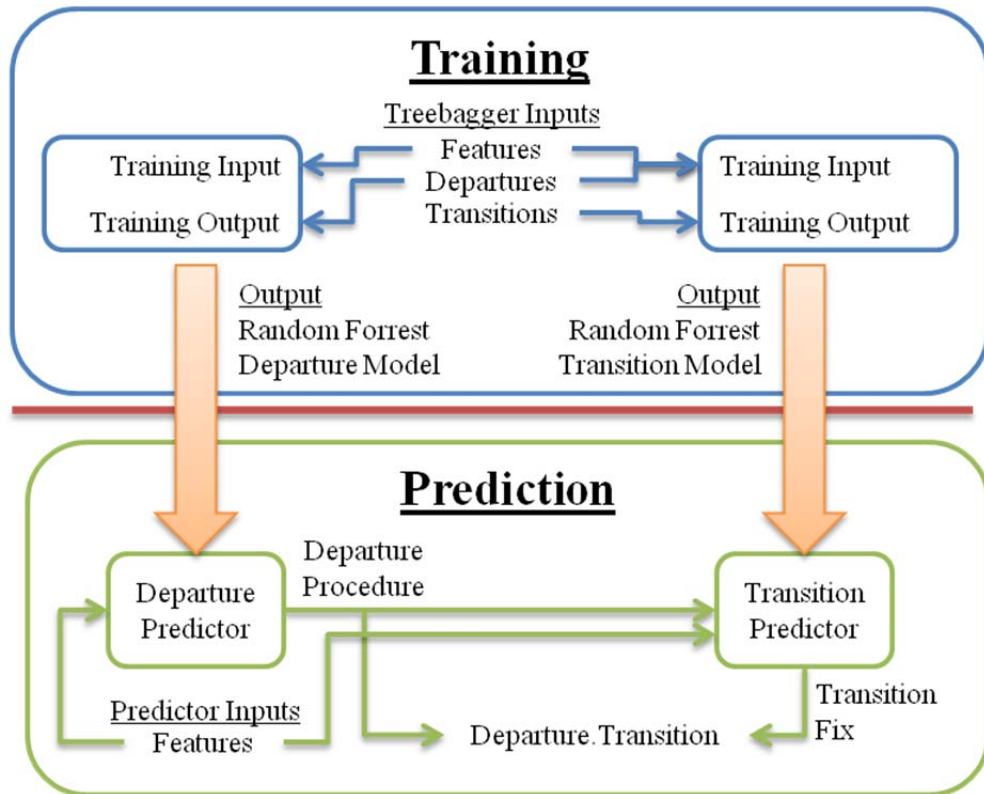
Figure 46. Process flow for departure and transition modeling.

All the collected and categorized data is processed through MATLAB's TreeBagger function using the default settings (with 50 random trees). The output of the TreeBagger function is a mapping function that takes as its input the above listed features and returns a departure procedure. After generating departure procedures for each flight, a transition fix is also generated using a similar process. A pictorial representation of the process is illustrated in Figure 46.

TABLE 12

Prediction Errors for Generative Departure Model

Airport	Training Set Size	Departure Prediction Error	Transition Prediction Error (given the departure)
DFW	149172	17.54%	2.78%
LAX	97446	13.86%	1.99%
MCO	49921	7.09%	.55%

One metric for the accuracy of the generative model is the out-of-bag prediction error; the error value is a measure of the percentage of training samples incorrectly labeled according to the generative model. The out-of-bag prediction errors for DFW, LAX, and MCO are provided in Table 12. Overall, the method provides accurate prediction of the departure procedure and transition fix for all the airports considered.

Using the historical generative model to assign consistent departure procedures and transition fixes to SWAC flights assumes that aircraft maintain the same mappings characteristics as they do today; that is to say that the features correlated to departure procedures and transition fixes today remain the same in the future. Furthermore, the modeling process assumes that the departure procedures remain static, and that new procedures are not introduced. Should the assumption that departure procedures remain static not hold, then for the analysis it is likely sufficient to assume that departure procedures are at least replaced with procedures that are spatially similar true (which is likely as non-RNAV procedures are replaced with RNAV procedures).

A.2   GENERATING BASELINE DEPARTURE FIX-TO-RUNWAY MAPPINGS

Two common runway assignment mappings are used in the simulation model for all case-study airports. The first runway assignment mapping is the static departure fix-to-runway mapping, which links departure procedures to assigned runways, given a specific airport runway configuration. The second runway mapping, referred to as a tradition runway mapping, considers both the departure procedure and the airline of each flight. Both mappings are derived from historical flight data using the procedure depicted in Figure 47.

To derive both the static departure fix-to-runway mapping and the tradition runway mapping, three data sources are required: PASSUR fight data, which includes departure times and departure runways; ETMS flight plans containing each aircraft departure procedure; and reported airport configurations contained within ASPM. All three sources of data are cross-referenced so that each flight is associated with a runway, airport configuration, and departure procedure. Any flights departing on runways that are not considered active by the reported airport configuration are excluded in the data collection.
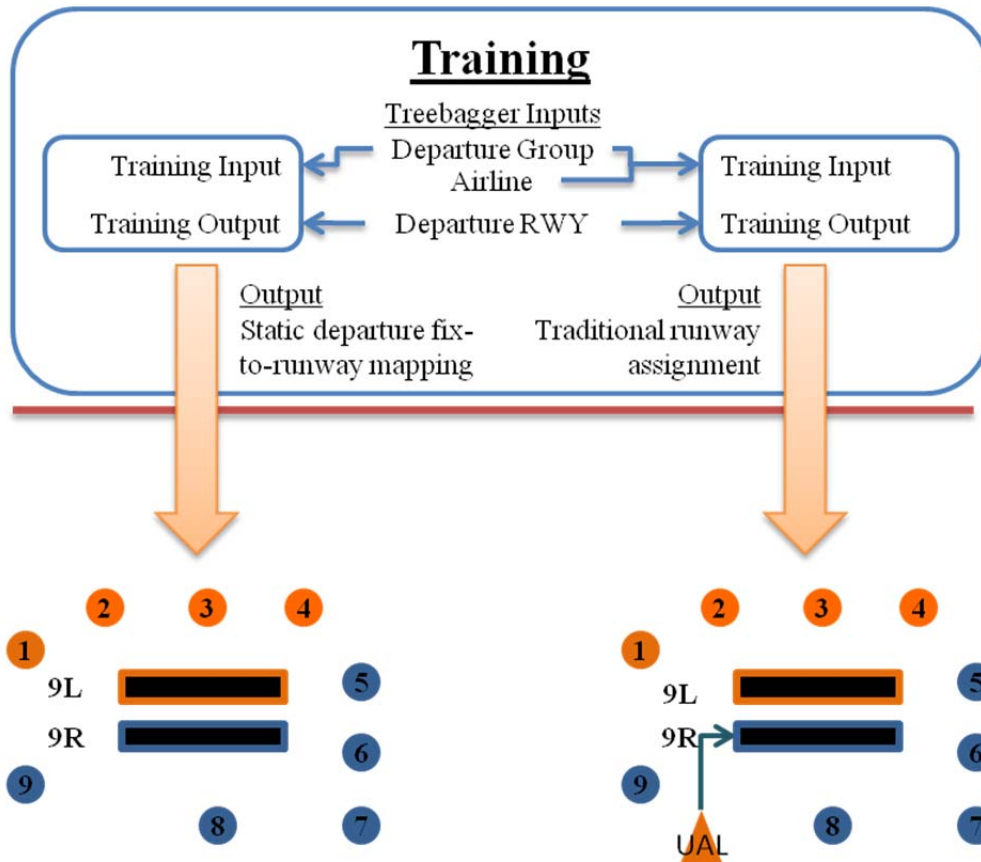
Figure 47. Process for generating runway assignment mappings.

For each airport and runway configuration, the static departure fix-to-runway mapping and tradition mapping are generated using the necessary corresponding data. The overall process is illustrated in Figure 46. The aggregated departure procedure and departure runway information for each airport and runway configuration are passed to the TreeBaggger function in MATLAB. The output of the TreeBagger function is a predictive mapping. The mapping approximately represents the most likely departure runway based on the available departure information. The traditional runway assignment mapping is generated in a similar fashion, except that airline information is included in forming any runway prediction. Inclusion of the airline information takes into account any airline preferences that might exist due to geographic location of an airline's terminal relative to the runway.

A.3   TAXI TIME MODELING

This appendix describes the use of ASDE-X flight track data for both taxi time validation and for predicting airline spot assignments.

In order to compare predicted taxi times with actual taxi times, three airports were studied: DFW, LAX, and MCO. Surface track data coming from ASDE-X feeds was analyzed for these airports over a period of 70 days from May 2014 until December 2014. The main purposes were to determine which ramp area aircraft depart from, which runway they are assigned to, and what the taxi times for those spot-runway pairs are. This was also further analyzed to determine how the flights for a given airline are divided amongst the spots.

The first step was defining the critical locations of the airport. This includes the runways, ramp areas, and queue boxes. The latitude/longitude coordinates for these were found using mapping software. These were manually defined by drawing boxes around the runways and ramp areas. Any ramp areas that were not separated by a movement area (such as a taxiway) were grouped together into one large ramp area. The queue boxes were also manually defined with the area chosen being the primary taxiway leading up to the runway threshold.

Total departure taxi times are defined as the difference in time from when the aircraft exits the ramp area box and when it enters the runway box. The departure runway is assigned by finding when the aircraft is inside a runway box and traveling above a certain speed threshold.

Other metrics being generated from the ASDE-X data analysis are

- Time entered in queue box
- Time stopped in queue box
- Spot-to-runway distance
- Spot-to-runway time

The spot locations were not manually determined from airport diagrams, rather automatically generated based on the flight tracks. For each departure flight track, the latitude/longitude coordinates are recorded for the last time it is within the ramp area. These locations are collected for all departures at a given airport and used as input for generating spot locations. The mean shift algorithm was used to translate all these ramp exit points to a set of spot locations.

Once the spots were determined for each airport, each flight was mapped to a spot. Using this information, it was possible to determine the distribution of spot locations for each airline. The flights were grouped by airline and then further divided into international and domestic flights. Then percentages of flights per spot were calculated; for example, the percentage of domestic American Airlines flights going to spot 3 at DFW. This can be used going forward to predict the number of flights per airline departing from various spots on the airport.

Figures 48 through 50 show the key areas for LAX, DFW, and MCO. The ramp areas are shown by the red boxes, the runways are shown by the blue boxes, and the spots are shown as the blue circles.
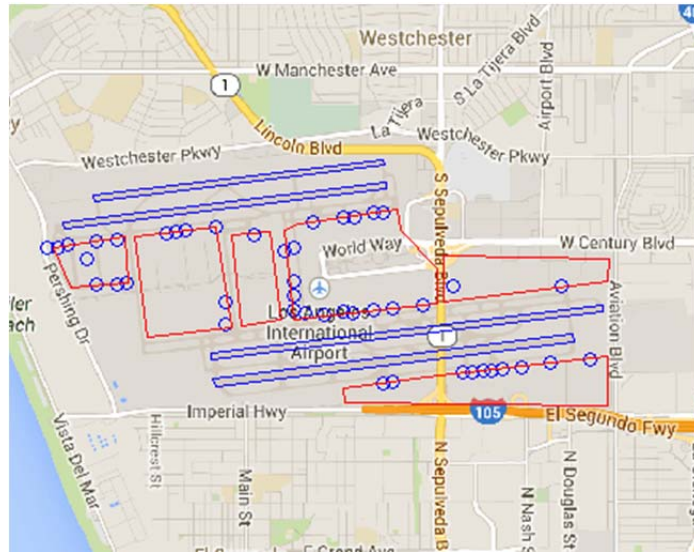

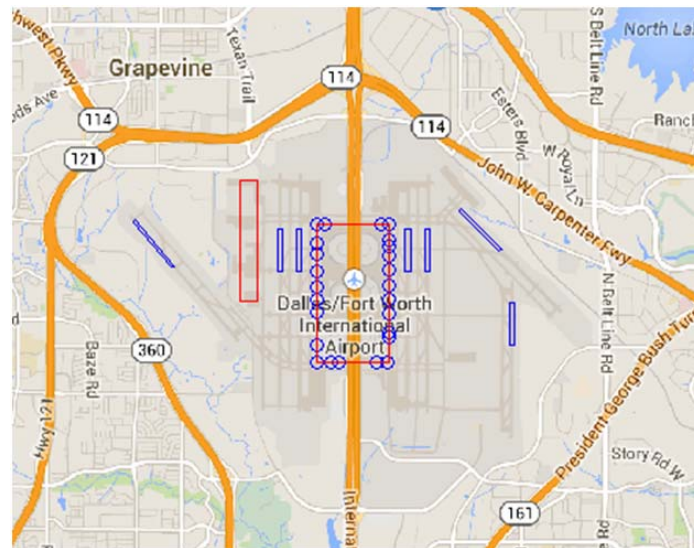*Figure 48. Ramp areas, identified spot locations, and runways at LAX.*

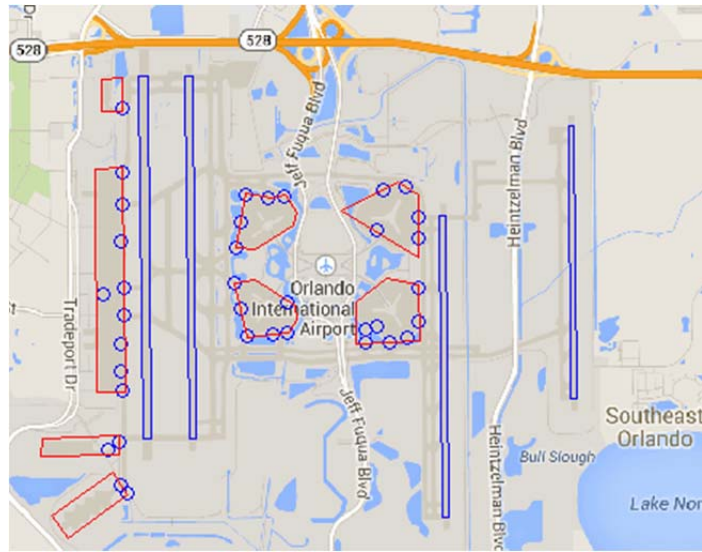*Figure 49. Ramp areas, identified spot locations, and runways at DFW.*

Figure 50. Ramp areas, identified spot locations, and runways at MCO.

A.4 RUNWAY SERVICE TIMES

Runway service times dictate the temporal process by which aircraft departures are modeled at runways. Traditionally, air traffic controllers release aircraft for take-off according to separation requirements. The standard separation requirements, defined in units of distance, are provided in Table 13 for aircraft flying different departure procedures. The separation distances are established to prevent wake turbulence hazards and to ensure the 3 NM separation within the terminal airspace. When aircraft depart along the same departure procedure additional spacing may be required to ensure 5 NM separation as aircraft enter en route airspace. The added spacing can be achieved by delaying aircraft departures so that the aircraft are separated by 5 NM at the time of take-off or by vertically separating aircraft.

TABLE 13

Spacing for Aircraft on Different Departure Procedures (NM)

L\F	Small	Large	757	Heavy
Small	3	3	3	3
Large	3	3	3	3
757	5	4	4	4
Heavy	5	5	4	4

65

The required departure spacing in the table above cannot be used in the time-based airport queuing model. Instead, time-based separations must be generated for the same aircraft classes.

The first step in generating equivalent departure time-spacing tables is to cross-reference and mine historical departure times and departure procedures at a number of airports (CLT, DFW, LAX, MCO, and PHX). The relevant data sources are PASSUR flight-ops data, which covers aircraft type, departure runway, and take-off time; and ETMS, which includes aircraft flight plans. The wait time between consecutive take-offs on the same runway are recorded, along with each aircrafts wake category. Additionally, each consecutive departure pair is labeled as having the same departure procedure or different departure procedures. Once all the wait times are organized according to departure status and leader-follower aircraft types, probability distributions are generated. The wait-time distributions for aircraft flying along different departure procedures are shown in Figure 51, and the wait-time distributions for aircraft flying along the same departure procedure are shown in Figure 52.
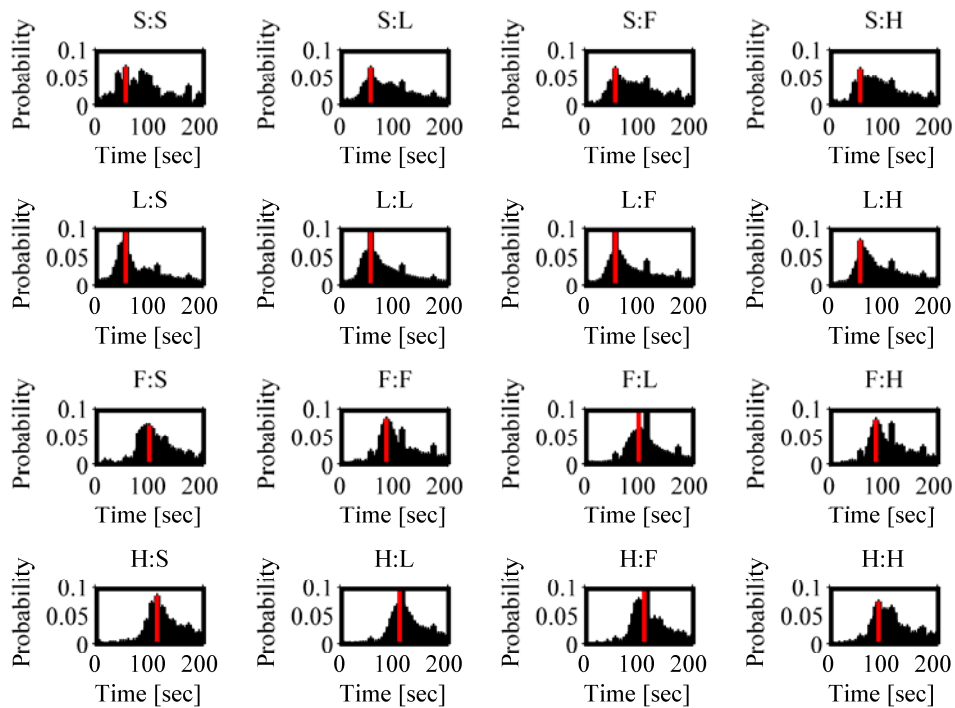

Figure 51. Wait-time distributions for aircraft on different departure routes.
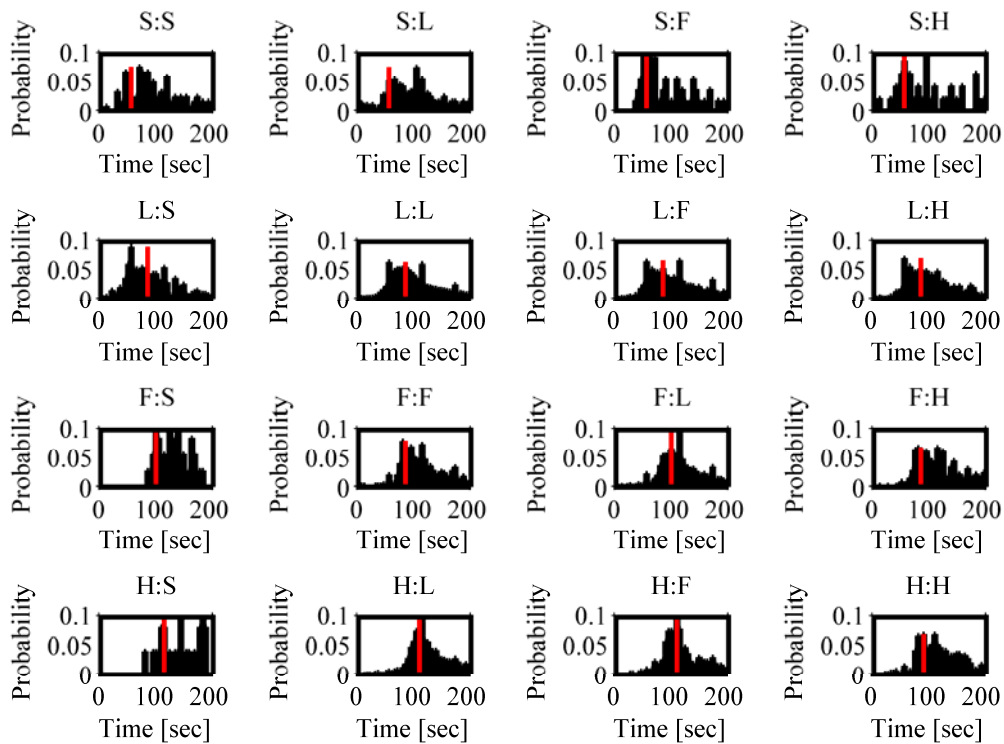
Figure 52. Wait-time distributions for aircraft on the same departure route.

One-by-one each wait-time distribution is assessed with the goal of establishing an average wait time, or more specifically, a central tendency. The average wait time is manually established by seeking out an approximate mode, while considering operational constraints. In cases of insufficient data (e.g., heavy followed by a small along the same departure), a best-estimate is made. The average wait time for each distribution is indicated by the red lines. A summary of the manually established average wait times are provided in Tables 14 and 15.

TABLE 14

Time Spacing for Aircraft on Different Departure Routes (seconds)

L\F	Small	Large	757	Heavy
Small	*60*	*60*	*60*	*60*
Large	90	90	90	*90*
757	105	105	90	90
Heavy	*120*	115	115	95

TABLE 15

Spacing Times for Aircraft on the Same Departure Route (seconds)

L\F	Small	Large	757	Heavy
Small	*60*	*60*	*60*	*60*
Large	90	90	90	*90*
757	105	105	90	90
Heavy	*120*	115	115	95

The procedure described above makes use of a number of assumptions. First, there is no distinction made between IMC and VMC conditions at airports. The implicit assumption is that the majority of departures at the airports occur during VMC conditions, and it follows that the resulting average wait times are more closely associated with VMC condition than with IMC conditions. Additionally, the process assumes that departure wait times at each airport are reflective of a common set of rules and free of any dependencies (e.g., dependent-departure operations at SFO).

It is worth noting that the wait-time distributions in Figure 51 and Figure 52 are based on imperfect take-off time data; the distributions show some inter-departure times of less than 15 seconds, which is operationally not acceptable and unlikely to have actually happened.

In the simulation model for each airport, the average wait times are used to define how long a departure aircraft claims a runway resource, and how much time is needed until the next aircraft must wait before it is able to depart. While aircraft speeds vary within each aircraft wake category, the wait times are intended to represent average operations. And while the average wait times do not correspond to actual distances, they do have the advantage of taking into account actualized pilot delays and air traffic controller delays in releasing aircraft for take-off.

GLOSSARY

ARMT	Airport Resource Management Tool
ASDE-X	Airport Surface Detection Equipment, Model X
ASPM	Aviation System Performance Metrics
CONOPS	Concept of Operations
ETMS	Enhanced Traffic Management System
FAA	Federal Aviation Administration
FID	Final Investment Decision
IARD	Investment Analysis Readiness Decision
IID	Initial Investment Decision
IMC	Instrument Meteorological Conditions
NAS	National Airspace System
SWAC	System-Wide Analysis Capability model
RNAV	Area Navigation
TFDM	Terminal Flight Data Manager
VMC	Visual Meteorological Conditions

This page intentionally left blank.

GLOSSARY FOR AIRPORT CODES

ATL	Hartsfield-Jackson Atlanta International Airport
BOS	Boston Logan International Airport
BWI	Baltimore Washington International Airport
CLT	Charlotte Douglas International Airport
DCA	Washington National Airport
DEN	Denver International Airport
DFW	Dallas-Fort Worth International Airport
DTW	Detroit Metro Airport
EWR	Newark Liberty International Airport
FLL	Fort Lauderdale-Hollywood International Airport
IAD	Washington Dulles International Airport
IAH	George Bush Intercontinental Airport
JFK	John F. Kennedy International Airport
LAS	McCarran International Airport
LAX	Los Angeles International Airport
LGA	LaGuardia Airport
MDW	Chicago Midway International Airport
MIA	Miami International Airport
MCO	Orlando International Airport
MSP	Minneapolis- St. Paul International Airport
ORD	Chicago O'Hare International Airport
PHL	Philadelphia International Airport
PHX	Phoenix Sky Harbor International Airport
RNAV	Area Navigation
SAN	San Diego International Airport
SEA	Seattle-Tacoma International Airport
SFO	San Francisco International Airport
SLC	Salt Lake City International Airport

This page intentionally left blank.

REFERENCES

[1]   Federal Aviation Administration, "Terminal Flight Data Manager Request for Information," RFI-13-00013, May 2013, https://faaco.faa.gov/index.cfm/announcement/view/14141.

[2]   Federal Aviation Administration, "Initial Investment Decision Benefits Basis of Estimate for Terminal Flight Data Manager (TFDM)," Version 0.99, March 10, 2014.

[3]   Federal Aviation Administration, "Terminal Area Forecast 2013-2040," February 2014, http://aspm.faa.gov.

[4]   Federal Aviation Administration, "Terminal Flight Data Manager Core for Air Traffic Control Towers Concept of Operations," Revision 2.1, November 4, 2013.

[5]   Noonan, K., "FAA's System-Wide Analysis Capability," Integrated Communications, Navigation and Surveillance Conference (ICNS), 10-12 May 2011, doi:10.1109/ICNSURV.2011.5935365, http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5935365&isnumber=5935117.

[6]   Federal Aviation Administration, "Terminal Flight Data Manager: Draft – System Specification Document," March 3, 2014.