ATD-2 Phase 3 Scheduling in a Metroplex Environment Incorporating Trajectory Option Sets

William J. Coupe, Yoon Jung
NASA Ames Research Center
Moffett Field, CA
william.j.coupe@nasa.gov,
yoon.c.jung@nasa.gov

Liang Chen
Moffett Technologies Inc.
Moffett Field, CA
liang.chen@nasa.gov

Isaac Robeson
Mosaic ATM
Leesburg, VA
irobeson@mosaicatm.com

Abstract—The NASA Airspace Technology Demonstration 2 Phase 1 and 2 Field Evaluations have successfully demonstrated new technologies developed to manage the Integrated Arrival, Departure, and Surface traffic flows at a single airport. The Phase 3 Field Evaluation extends the capabilities to a Metroplex environment where multiple airports are interacting and sharing resources along the terminal boundary. This paper describes the scheduling algorithm enabling the coordinated scheduling and describes the interaction between airports within the Metroplex and the terminal boundary. We describe the metrics developed to inform flight operators about reroute opportunities and discuss the potential benefits to the rerouted flight and the system-wide aggregate benefits of a single reroute. We believe that the capabilities developed and the lessons learned during the Phase 3 Field Evaluation will set up the National Airspace System for future success.

Index Terms—Airspace Technology Demonstration 2, metroplex scheduling, trajectory option set

I. Introduction

Concepts and technologies to manage arrival, departure, and surface operations have been under development by NASA [1]–[5], the Federal Aviation Administration (FAA) [6], and industry to improve the flow of traffic into and out of the nation's busiest airports. NASA is conducting the Airspace Technology Demonstration-2 (ATD-2) to evaluate an Integrated Arrival, Departure, and Surface (IADS) traffic management system that demonstrates these technologies [7], [8].

The IADS system was deployed to Charlotte Douglas International Airport (CLT) for a three-year field evaluation that is divided into three distinct Phases lasting one year each. During Phase 1 the IADS system successfully demonstrated three key capabilities [9]: 1) data exchange and integration, 2) departure scheduling and electronic negotiation of release times of controlled flights for overhead stream insertion, and 3) tactical departure surface metering [10]. In Phase 2 of the field evaluation the scheduling capabilities were improved [11] to enable strategic Surface Metering Programs (SMP). The IADS system at CLT with Phase 1 and 2 capabilities is a precursor of the FAA Terminal Flight Data Manager (TFDM), which is slated to be rolled out across the United States' busiest airports starting in 2020, after the conclusion of NASA's field evaluation.

The Phase 3 field evaluation extends the coordinated scheduling of arrivals, departures, and surface traffic from a single airport at CLT [11] to a Metroplex environment in North Texas containing Dallas Love International Airport (DAL), Dallas Fort Worth International Airport (DFW), and other small satellite airports in the D10 Terminal Radar Approach CONtrol (TRACON). The challenges in the D10 Metroplex are fundamentally different than the challenges addressed by the Phase 1 and Phase 2 IADS capabilities in CLT. In CLT surface congestion and constraints from controlled flights are the main challenges, whereas in the D10 Metroplex the main constraint is the departure fix capacity as multiple major airports compete for the same limited resources. This problem can be magnified when inclement weather impacts D10 and reduces the capacity at the terminal fixes which can propagate delay to the surface of each airport in the D10 Metroplex.

When inclement weather constrains the capacity at a given fix in the D10 Metroplex, there are often alternative fixes that are not impacted by the weather that have available capacity. When this situation occurs, a flight that is originally routed through the constrained fix can reroute through the alternative fix with little to no delay. The tradeoff for the flight to reroute to the alternative fix is often a longer route in terms of air miles and requires additional fuel. By comparing the additional mileage of the alternative route to the delay savings, the airline operators can make an informed decision about when it is advantageous to fly the alternative route.

The IADS Phase 3 system in the D10 Metroplex aids the decision to reroute aircraft over an alternative fix by assessing the delay savings on each alternative route defined by each flight operator's Trajectory Option Set (TOS). The TOS is a set of alternative routes the flight is willing to fly and each route has an associated Relative Trajectory Cost (RTC). The delay savings for each route in the TOS is compared to its RTC to determine when the delay savings on an alternative route rises above the RTC threshold value. In addition to computing the delay savings for individual flights, the IADS Phase 3 system also calculates the overall savings at the system level resulting from a reroute of a single flight. The savings at the system level is important for the flight operators as they are able to see how rerouting a single flight can benefit their fleet.
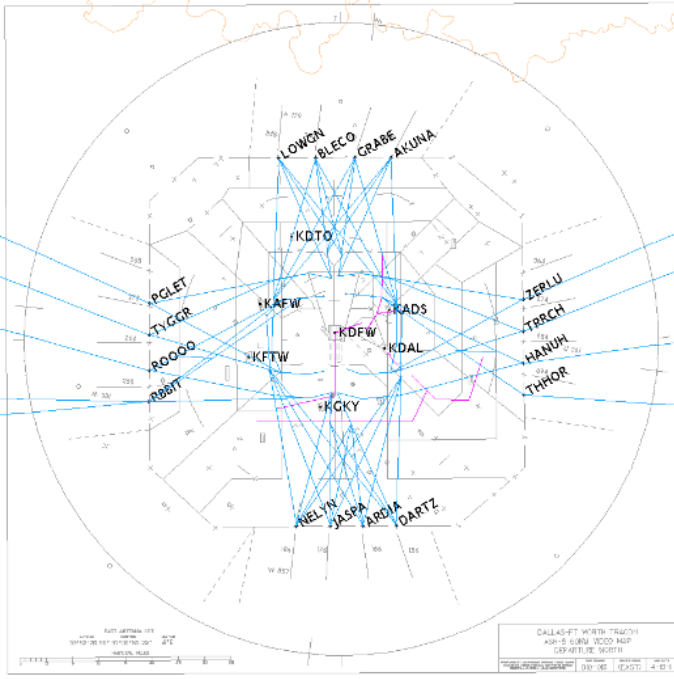
In this paper, we describe the IADS Phase 3 scheduling

Fig. 1. D10 airspace showing how airports share the 16 departure fixes along the terminal boundary.

Fig. 2. Top: Demand for departure gate. Bottom: Demand for departure fix

algorithm which provides the coordinated schedules between each airport surface and the terminal boundary in the D10 Metroplex. This is a challenging scheduling problem because constraints imposed by fix closures and Miles-In-Trail need to be properly accounted for at both the terminal boundary and each airport surface. In addition to describing the scheduling algorithm, we define the metrics used to inform the flight operators about opportunities to reroute aircraft and metrics used to assess the performance of the system including potential benefits. Using the operational data, we illustrate the benefits of the Phase 3 IADS scheduling concept in the D10 Metroplex.

II. BACKGROUND INFORMATION ON D10 AND TOS

A. D10 Airspace and Demand at the Terminal Boundary

The D10 TRACON is centered on Dallas/Fort Worth International airport (DFW) and extends outward approximately forty miles. It contains two major airports, DFW and Dallas Love Field (DAL), which are separated by approximately ten miles, see KDFW and KDAL in Fig. 1. Several busy general aviation airports, a regional cargo hub, and a Naval Air Station Joint Reserve Base are also located within the D10 TRACON, contributing to operational complexity [12].

Along the boundary of the D10 terminal boundary there are 16 departure fixes that are shared among all the airports. Although each airport contributes to demand at the departure fixes, the majority of flights originate from DFW and DAL. The departure fixes are arranged in groups of four along the North, East, South, and West departure gates, see Fig. 1. Traffic through the departure fixes and gates is not evenly dis-
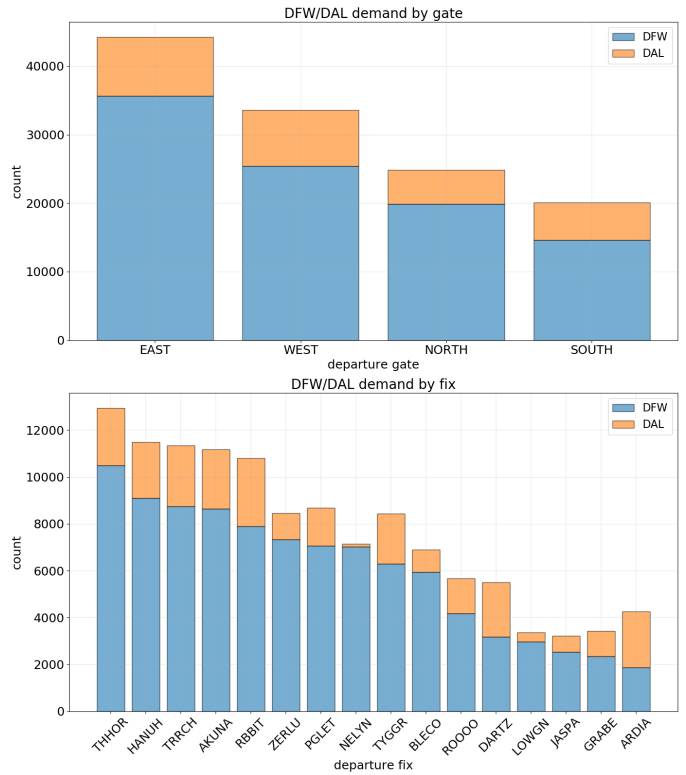
tributed and is dictated by the demand to different geographic locations of destination airports, see Fig. 2.

Fig. 2 shows a summary of the demand by airport and contains all departures from DFW and DAL between 2019-07-01 through 2019-09-30. During this time period there were 95,791 and 27,657 departure operations from DFW and DAL, respectively. The East gate was utilized most, with 44,231 operations (36% of overall demand). Each departure with a route filed through the East gate competes for the fix with all other airports. This competition for a shared resource at the terminal boundary can increase delays.

Due to the imbalanced demand among the departure gates and fixes there is often unused capacity at the terminal boundary. For example, the South gate was utilized less than half as often as the East gate. During periods where more departures are competing for the East gate than can be accommodated, flight operators can often route flights through the South or North gate where they will not be subject to the demand/capacity imbalance and will experience less delay. This reroute comes at a cost, however, as the route through the South or North gate typically requires the flight to fly farther along a less direct route.

B. Capacity and Restrictions at the Terminal Boundary

Capacity at the terminal boundary is defined by minimum separation constraints and Traffic Management Initiative (TMI) restrictions that are enforced at the departure fix. TMIs at the terminal boundary are typically triggered by weather
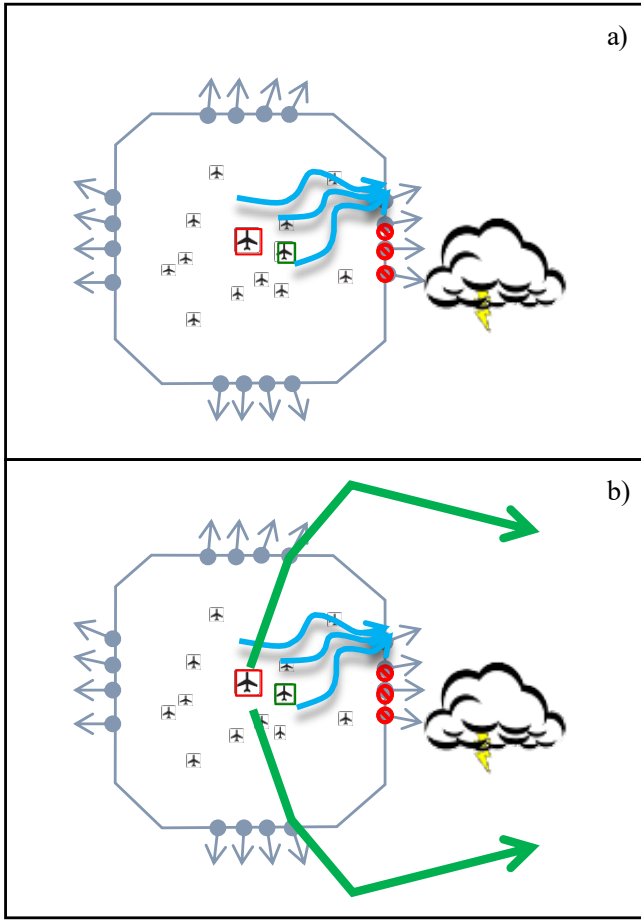
Fig. 3. a) D10 airspace with weather impacting the East gate. b) Available TOS routes not impacted by weather constraints.

Fig. 4. a) Daily count of restrictions. b) Percentage of flights subject to MIT restriction. c) Percentage of flights subject to fix closure.

events or downstream flow constraints that propagate back to the TRACON environment [12], and ultimately the departure airports.

In response to weather events around or near the terminal boundary the Traffic Management Unit within Air Traffic Control (ATC) will close departure fixes which result in the departure gate being partially or completely blocked, see Fig. 3(a). The Fig. illustrates a situation where three of the four East departure fixes have been closed and traffic through these fixes is rerouted to the single remaining fix along the East gate. This compression of the departure fixes reduces the capacity at the terminal boundary and delays can be amplified when additional MIT restrictions are enforced at the departure fix.

In the D10 TRACON, weather events often lead to multiple, dynamic TMI restrictions being issued by ATC. Fig. 4(a) shows the count of fix closure and MIT restrictions during the time period 2019-07-01 through 2019-09-30. Some days during this time period had 30 or more unique fix closures throughout the day. For a given day, we see fewer MIT restrictions than fix closures because ATC often responds to weather events by closing multiple fixes and putting a single MIT restriction on the compressed flow, see Fig. 3(a).
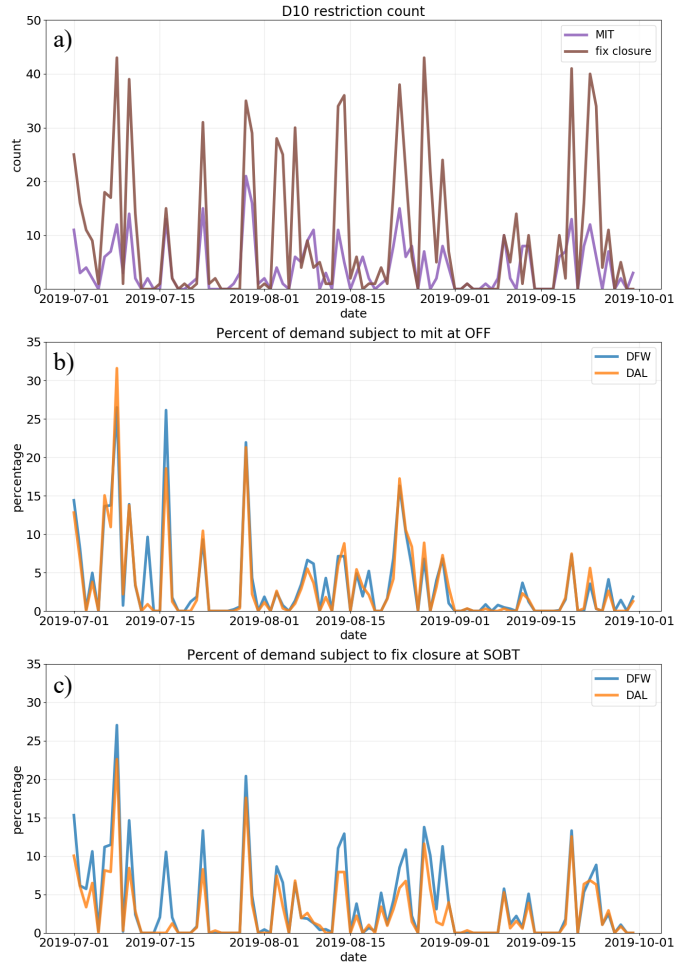
Fig. 4(b) shows the percentage of departure flights from DFW and DAL that are subject to a MIT restriction at the OFF event. Similarly, Fig. 4(c) shows the percentage of departure flights from DFW and DAL that are subject to a fix closure at their Scheduled Off Block Time (SOBT). Since fix closures and MIT restrictions are often enforced together the percentage of flights that are subject to the different types of restrictions is similar.

C. Trajectory Option Set (TOS)

When TMI restrictions reduce the capacity at the terminal boundary there are often opportunities to route around the restrictions and reduce the delay. Consider Fig. 3(b) which shows the situation where the East gate is limited to a single fix with a MIT restriction, while the North gate and South gate have all four fixes available. Since the traffic volumes through the North and South gate are relatively light, see Fig. 2, and green routes are not impacted by a TMI restriction, a flight could reroute through the North or South gate with little to no delay.

A flight operator defines the Trajectory Option Set (TOS) which is the set of feasible routes for a given flight. The filed route is typically the most direct route and is preferred by the flight operators under nominal operations. The cost of each route, often a function of the additional mileage needed to fly the route, is provided by the flight operators in the form of a Relative Trajectory Cost (RTC). The RTC is a way for the flight operators to express their willingness to fly a more costly route when the delay savings exceeds the RTC threshold.

To inform flight operators about reroute opportunities the predicted delay on the filed route and each TOS alternative route is needed. This difference in delay between the filed route and the alternative route is compared to the RTC threshold to determine if and when a reroute is warranted. The delay savings on the alternative route minus the RTC represents the net savings to the rerouted flight.

D. Aggregate System-Wide Benefits from TOS Reroute

There are often additional system-wide benefits from a single flight reroute in addition to the net savings to the individual flight that is rerouted. When the terminal boundary is operating as a significant constraint to the flow of traffic, in contrast to situations where the runway is acting as the main constraint, there exists benefits to the set of flights that are not rerouted and subject to the terminal constraints.

The system-wide aggregate benefits materialize if MIT restricted flights are able to move one slot earlier owing to the rerouted flight giving up its slot. If the MIT restrictions at the terminal boundary are operating as the the main constraint on the system, then there is often available capacity at the runway to accommodate the rerouted flight without delaying other flights. The system-wide aggregate delay savings will be discussed in more detail in Section IV-C. The remainder of this paper describes our approach to estimate delay among the different routes and the metrics created to inform flight operators about TOS reroute opportunities.

III. SCHEDULER

The Terminal Scheduler generates the delay estimate on the filed route and the TOS alternative routes. The Terminal Scheduler is composed of an Orchestrator, Trajectory Modeler, Departure Fix Scheduler, and Airport Scheduler for each airport in D10, see Fig. 5. The Fig. illustrates Loop k, which is a single iteration of the Terminal Scheduler that is run every 10 seconds and generates an Estimated Take Off Time (ETOT) for every flight on all filed and TOS alternative routes.

The ETOT is generated by an iterative process between the Departure Fix Scheduler and the Airport Schedulers and accounts for all terminal and surface constraints. Within each Loop k, the Orchestrator is initialized with the set of known flights and then executes three sub-routines, referred to as Loop $k1$, Loop $k2$, and Loop $k3$.

A. Loop $k1$

After initialization we execute Loop $k1$ which loops through each route including filed and TOS routes and calls the Trajectory Modeler. For a departure aircraft, the Trajectory Modeler generates an Unimpeded Off-Block Time (UOBT), Unimpeded Taxi Time (UTT), Unimpeded Take Off Time (UTOT), Unimpeded Flight Time (UFT), and Unimpeded Fix Crossing Time (UFXT) estimate. The off-block time refers to the time the aircraft initiates the pushback from the gate. The model is provided with an Earliest Off-Block Time (EOBT) prediction from the airline operators. The UOBT represents the best estimate of the time the aircraft will initiate the pushback process. It is defined as the maximum of the EOBT and the current time because if the EOBT estimate is in the past, then the current time is the earliest the flight would be available to initiate the pushback process. The UTT and UFT are derived from nominal taxi / flight speeds and the expected taxi / flight route and are used to generate the UTOT defined as the UOBT + UTT and the UFXT defined as UTOT + UFT.

One of the core functions of the Trajectory Modeler is computing the four-dimensional (4D) (x,y,z,t) surface + terminal trajectory from the gate to the runway to the departure fix based on the expected airport/runway configuration, gate/runway assignment, and fix/runway mapping. The Trajectory Modeler uses surface surveillance data, when available, to detect the actual surface trajectory and update the trajectory prediction all the way to the fix. In the Phase 3 system DFW is the only airport using surface surveillance in D10. The Trajectory Modeler uses coded taxi routes defined by the adaptation using the airport resource information to select the available routes or default to the shortest path when the coded taxi routes are not available in the adaptation.

B. Loop $k2$

The purpose of Loop $k2$ is to build a schedule at the terminal boundary and then try to enforce the terminal constraints within the local Airport Schedulers. If the airport surface constraints violate the terminal constraints we communicate this back to the Departure Fix Scheduler by redefining the UFXT as the surface constrained Target Take Off Time (TTOT) plus transit time to the terminal boundary. The Departure Fix Scheduler then builds a new schedule knowing about the surface constraints and we again check the feasibility at the airport surface. Through this iterative process the scheduler converges to a schedule that satisfies all known terminal and surface constraints.

After the trajectories are generated in Loop $k1$, we execute Loop $k2$ which is the core scheduling loop. The output of Loop $k2$ is an ETOT on the filed route for each departure flight. We begin by sorting flights by UFXT and send this sorted list to the Departure Fix Scheduler. The Departure Fix Scheduler applies a simple First Come First Served (FCFS) heuristic to schedule flights according to minimum separation constraints and MIT restrictions. In the absence of MIT restrictions, the Departure Fix Scheduler enforces a minimum separation of 5NM for all flights through the East gate, but does not apply minimum separation constraints for flights through the North, South, and West gates as we assume ATC can accommodate the flow.
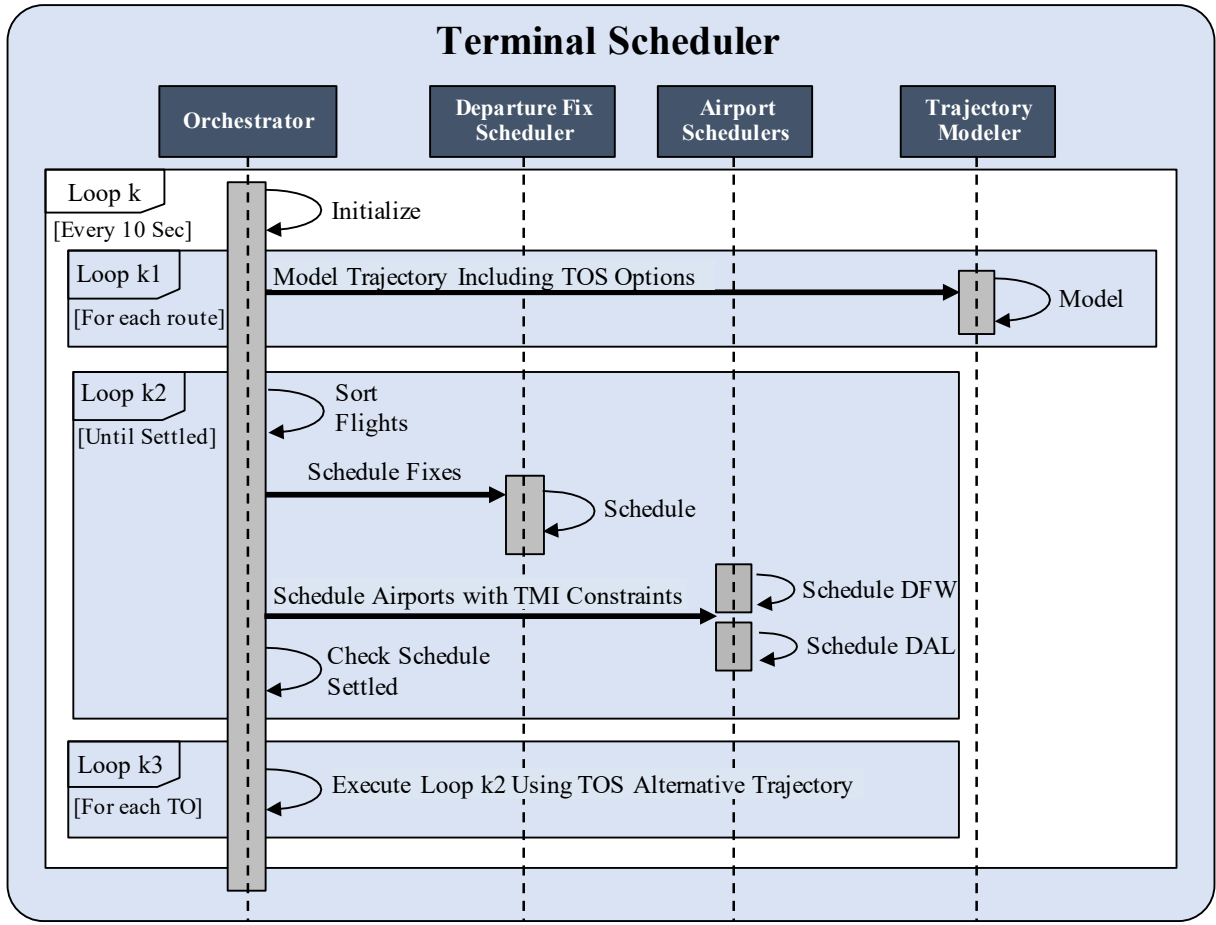
Fig. 5. Terminal Scheduler diagram showing how Orchestrator, Trajectory Modeler, Departure Fix Scheduler, and Airport Scheduler work together.

Flights that are subject to MIT restrictions or that travel through the East gate and thus subject to minimum separation constraints must be able to depart their origin airport at the required time to comply to the scheduled Target Fix Crossing Time (TFXT). These flights are assigned a Terminal Controlled Off Time (TCOT) defined as TCOT = TFXT - UFT. The output of the Departure Fix Scheduler is a TCOT for the set of flights restricted at the terminal boundary. These TCOTs will be passed to the Airport Schedulers as input.

After scheduling flights at the terminal boundary the scheduler checks to see if the assigned TCOTs violate any known surface constraints. To check if the terminal schedule is feasible, the Orchestrator passes the TCOT constraints to the Airport Schedulers which each build a surface schedule that tries to honor the TCOT constraints. The methodology of how we apply TCOT constraints will be described in Section III-D. The output of the Airport Scheduler is a TTOT for each flight.

For the Airport Scheduler, arrivals are inserted into the schedule and assigned a Target LanDing Time (TLDT) before departures. The departures are then assigned TTOTs in order based on a selection criteria defined by the UTOT and Scheduling Group. The scheduler is modular to allow for different selection criteria to be implemented. Once a departure is selected to be inserted into the schedule, the departure is assigned a feasible TTOT such that the TTOT satisfies constraints including aircraft type (i.e., taxi speed, wake vortex separation), dual-use runways, converging runway operations, any TMIs, and conflicts at the runway thresholds [11].

After building the surface schedule at both airports, the scheduler checks to see if the schedule has settled. The schedule is defined to settle when the output of the Departure Fix Scheduler (TCOT) matches the output of the Airport Schedulers (TTOT). If the schedule did not settle, the UFXT is redefined as UFXT = TTOT + UFT and Loop $k2$ continues. If the schedule has settled, the scheduler breaks out of Loop $k2$ and returns an ETOT on the filed route for each flight.

C. Loop $k3$

After the ETOTs on the filed routes are generated, we execute Loop $k3$ which computes the ETOT on the TOS alternative routes. For each TOS alternative route, Loop $k3$ takes the underlying trajectories from the Trajectory Modeler used in Loop $k2$ and changes a single trajectory. For each TOS alternative trajectory we substitute the alternative trajectory for the original filed trajectory and this is passed as input to the core scheduling Loop $k2$.
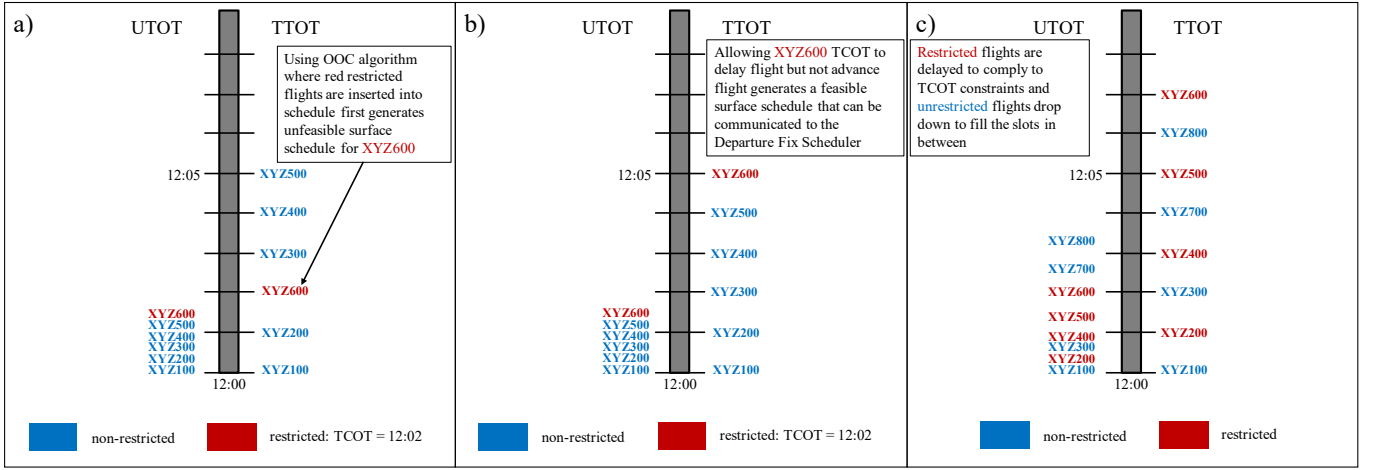
Fig. 6. (a): Schedule generated with OOC heuristic. (b): Schedule generated using heuristic that TCOT can delay flight but not advance flight in UTOT sequence. (c): Schedule showing that non-restricted flights can fill the gap between restricted flights.

We execute the scheduling loop and return an ETOT on the TOS alternative route for the flight in question. A set of ETOTs for all other flights given the reroute are also returned. As explained in Section II-D, a single reroute can impact the ETOT of other flights in the schedule. The set of ETOTs given the reroute will be used in Section IV-C to measure the system-wide aggregate benefits for a given TOS alternative option.

D. Applying TCOT Constraints in the Airport Scheduler

The Terminal Scheduler requires the Airport Scheduler to enforce TCOT constraints when possible and communicate surface constraints when not possible. If we used a simple Order Of Consideration (OOC) algorithm where the restricted TCOT flights are inserted into the scheduler before any non-restricted flight then we could fail to communicate the surface constraints to the Departure Fix Scheduler, see Fig. 6.

In Fig. 6(a) we show a vertical timeline with 12:00 at the bottom and times later than 12:00 above. We show a situation with 5 blue non-restricted flights that are in a physical queue and a red restricted flight that arrived at the back of queue. We imagine that the red restricted flight has a TCOT from the terminal scheduler at 12:02. If the Airport Scheduler uses a simple order of consideration and first assigns the red flight to the 12:02 slot and then schedules the blue flights around we could have the TTOT timeline shown.

Fig. 6(a) shows that even though the red restricted flight is physically in the back of the queue and can not take off at 12:02, the Airport Scheduler has assigned the 12:02 slot to the red flight and failed to communicate the surface constraints. Instead, the Airport Scheduler implements a heuristic that the TCOT constraint can delay a flight but can not advance the flight in the UTOT surface sequence. The result is the schedule shown in Fig. 6(b). The TTOT from the schedule on the right can then be used to defined UFXT = TTOT+UFT and thus the surface constraints are communicated back to the Departure Fix Scheduler.

If the TCOT constraints result in a restricted flight being delayed the Airport Scheduler will delay the restricted flight while allowing non-restricted flights to drop down the timeline and fill the gap between the restricted flights, see Fig. 6 c). The restricted flights are colored red and non-restricted flights colored blue. In the UTOT sequence there are three restricted flights in a row: XYZ400, XYZ500, and XYZ600. To comply to the MIT restrictions these flights will be separated leaving one slot between any two restricted flights. This allows the non-restricted flights XYZ700 and XYZ800 to fill the gap.

Delaying the restricted flights while allowing the non-restricted flights to drop in between we create a mismatch between the UTOT sequence on the left hand side of the timeline and the TTOT sequence on the right hand side of the timeline. We find this mismatch acceptable as we make the assumption that ATC ground or local controllers will delay the restricted flights to comply to the MIT restrictions while simultaneously filling the slots between the restricted flights to ensure there is no wasted capacity at the runway.

As flights get closer to taking off we expect the UTOT sequence to align with the TTOT sequence. For example, XYZ100 through XYZ400 are the first four flights on the timeline and ATC ground controllers have delivered a feasible sequence that complies to the MIT restrictions. The flights higher up on the timeline XYZ500 through XYZ800 are the flights where the UTOT sequence does not match the TTOT sequence. Our expectation is that through the actions of the ATC ground or local controller, as these flights get closer to taking off the UTOT sequence would change to align with the TTOT sequence.

IV. METRICS

A. Estimated Take Off Time Accuracy

The system relies on accurate ETOT predictions to evaluate the delay savings between the filed route and TOS alternative routes. To provide flight operators and ATC a tool to judge the accuracy of the system we measure and display the ETOT
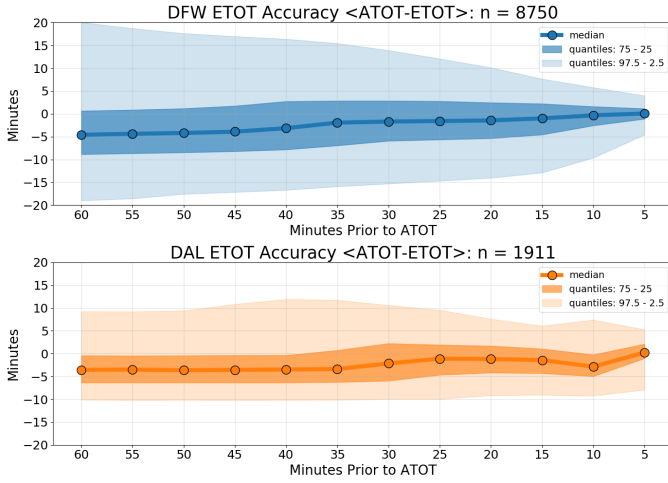
Fig. 7. ETOT accuracy as a function of lookahead time prior to ATOT

Fig. 8. Top: Scenario 1 $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$ represent when we think the flight will take off on the filed and TOS alternative route, respectively. Bottom: Scenario 1 delay savings $\mathcal{DS}$ and associated CDF

accuracy as a function of lookahead time prior to the Actual Take Off Time (ATOT). As an example, Fig. 7 contains all departure flights from DFW and DAL between 2020-05-01 through 2020-05-30. A filter is applied to only include predictions for flights scheduled on the runway the flight actually used. We also apply a filter to eliminate any observations that are beyond the median value $\pm 3.5 \times$ IQR where IQR is the difference between the 25^{th} and 75^{th} quantile.

The horizontal axis of Fig. 7 represents the lookahead time from 60 minutes prior to ATOT to 5 minutes prior to ATOT. The vertical axis is the accuracy of the ETOT prediction measured as ATOT − ETOT in minutes. A positive / negative value means the aircraft took off later / earlier than the prediction. The median is illustrated with a solid line and the IQR is shown as a dark shade around the line. A light shade around the median and IQR illustrates the 2.5 and 97.5 quantiles.

The difference in ETOT accuracy at DFW and DAL is likely due to many factors including the complexity and number of operations and also the availability of surface surveillance at DFW but not DAL. The surface surveillance enables the airport model within the Terminal Scheduler to consistently update UTOT predictions which help improve the ETOT accuracy as the flight gets closer to take off.

The ETOT accuracy provides useful information to the flight operators when considering a TOS reroute. The predicted delay savings should be considered in the context of accuracy of the ETOT predictions. If the delay savings is much larger / smaller than the underlying uncertainty in the ETOT prediction then we should have more / less conviction that the delay savings will materialize. To help facilitate this comparison we developed a new metric that measures the probability that the delay savings will exceed the RTC. The following Section IV-B describes our approach to measure this value.
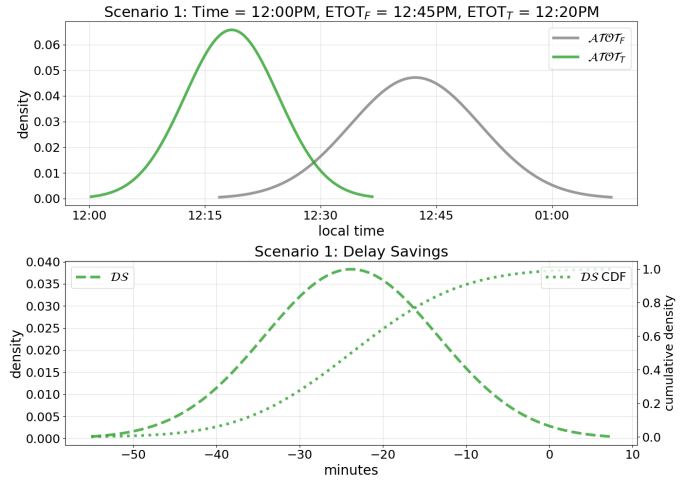
B. Probability that Delay Savings Exceeds RTC

The ETOT predictions on the filed and TOS alternative route are used in combination with the accuracy measurements shown in Fig. 7 to construct the probability that the delay savings will exceed the RTC. In Scenario 1, at 12:00 local time a DFW flight has ETOT_F = 12:45 and ETOT_T = 12:20 where ETOT_F and ETOT_T are the ETOT on the filed and TOS alternative route, respectively. This implies a lookahead value of 45 minutes on the filed route and 20 minutes on the TOS alternative route.

To approximate the uncertainty in the prediction for the filed and TOS alternative route we fit a Normal distribution to the residuals defined at the 45 and 20 minute lookahead, respectively. Using the Normal distributions we construct two new variables:

$$\mathcal{ATOT}_F = \text{ETOT}_F + \mathcal{N}(\mu_{45}, \sigma_{45}) \tag{1}$$
$$\mathcal{ATOT}_T = \text{ETOT}_T + \mathcal{N}(\mu_{20}, \sigma_{20}) \tag{2}$$

where the F subscript denotes the filed route, the T subscript denotes the TOS alternative, and μ_X and σ_X represent the mean and standard deviation of the ATOT − ETOT residuals at X minute lookahead (derived from the data shown in Fig. 7). Variables $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$ are defined as a constant plus a Normal, and are thus Normal themselves.

For Scenario 1 if we plug in the values ETOT_F = 12:45 and ETOT_T = 12:20 to (1) and (2) we can visualize $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$, see Fig. 8. In the top subgraph we plot $\mathcal{ATOT}_F$ with a grey line and $\mathcal{ATOT}_T$ is plotted with a green line. The distributions represent where we think the flight would take off on the filed and TOS alternative route given the ETOT predictions and the underlying accuracy.

We define the delay savings distribution $\mathcal{DS}$ from $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$ as follows:

$$\mathcal{DS} = \mathcal{ATOT}_T - \mathcal{ATOT}_F \tag{3}$$

where a negative / positive value represents the TOS alternative route experiencing less / more delay compared to the filed route. The $\mathcal{DS}$ is simple to construct as it is the difference between two Normal distributions. For Scenario 1 we plot $\mathcal{DS}$ with a dashed green line in the bottom subgraph of Fig. 8. $\mathcal{DS}$ is centered near -25, the difference between $\text{ETOT}_F = 12{:}45$ and $\text{ETOT}_T = 12{:}20$. The $\mathcal{DS}$ distribution mean is not exactly at -25, however, as the mean values of $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$ account for the bias in the residuals.

The $\mathcal{DS}$ distribution from (3) is converted into a Cumulative Distribution Function (CDF) to calculate the probability that the delay savings is greater than the RTC value. The CDF for $\mathcal{DS}$ is shown with a green dotted line in the bottom subgraph of Fig. 8. Once the CDF is constructed we can estimate the probability that the delay savings $\mathcal{DS}$ exceeds any given RTC value as follows:

$$\text{pr}(\mathcal{DS} \leq \text{RTC}) = \text{CDF}(\text{RTC}) \tag{4}$$

where $\text{CDF}(\text{RTC})$ is the CDF evaluated at the RTC value. For example, if the TOS alternative route in Scenario 1 was set as $\text{RTC} = -10$ the $\text{pr}(\mathcal{DS} \leq -10) = 0.908$, i.e. there is a 90.8% chance that the delay savings on the TOS alternative route would exceed the RTC.

We can compare the results of Scenario 1 shown in Fig. 8 to Scenario 2 where at 12:00 local time a DFW flight has $\text{ETOT}_F = 12{:}45$ and $\text{ETOT}_T = 12{:}40$. Scenario 2 differs from Scenario 1 only in the value of the ETOT_T on the alternative route. Fig. 9 contains the same information as Fig. 8 and can be compared to understand how the $\mathcal{ATOT}_T$ and $\mathcal{DS}$ distributions are impacted by the change in ETOT_T.

Fig. 9 shows that if the difference between ETOT_F and ETOT_T is small compared to the underlying uncertainty in the ETOT predictions, then the $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$ distributions overlap and the $\mathcal{DS}$ distribution is centered near zero with density for positive values. Evaluating (4) for Scenario 2, we calculate $\text{pr}(\mathcal{DS} \leq -10) = 0.314$, i.e. there is a 31.4% chance that the delay savings on the TOS alternative route would exceed the RTC.

C. System-Wide Aggregate Delay Savings

For each TOS alternative trajectory, in Loop $k3$ we calculate an ETOT_T for the rerouted flight and ETOT_R for the entire schedule given the reroute. We define the system-wide aggregate delay savings ADS as:

$$ADS = \Big(\text{ETOT}_T - \text{ETOT}_F\Big) + \sum_{\mathbb{F}} \Big(\text{ETOT}_R - \text{ETOT}_F\Big) \tag{5}$$

which is the delay savings to the rerouted flight plus a sum over the set of flights $\mathbb{F}$ of the ETOT difference between the ETOT_R given the reroute and the schedule on the filed route ETOT_F. When a single flight is rerouted and the reroute results
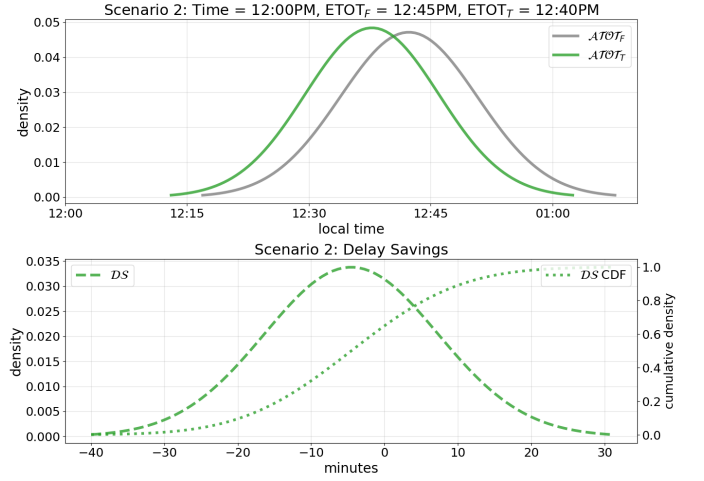

Fig. 9. Top: Scenario 2 $\mathcal{ATOT}_F$ and $\mathcal{ATOT}_T$ represent where we think the flight will take off on the filed and TOS alternative route, respectively. Bottom: Scenario 2 delay savings $\mathcal{DS}$ and associated CDF

in $\text{ETOT}_T \neq \text{ETOT}_F$, once the change propagates through the schedule, the result can be that flights that are not rerouted have $\text{ETOT}_R \neq \text{ETOT}_F$, thus the system-wide ADS measure differs.

The set of flights $\mathbb{F}$ that we include in the ADS summation can be defined to provide different flavors of the metric. We can include all flights in D10, flights only from DFW, flights only from DAL, or any other constraints. The different versions of ADS could be valuable to different decision makers. ATC might be interested in looking at ADS summed over the D10 TRACON to understand the impact of a single reroute to the flow through the terminal whereas flight operators might be more interested in the set of flights $\mathbb{F}$ from a specific airport or even a specific flight operator to understand the impact of the reroute decision on their fleet.

An additional constraint enforced on the set of flights $\mathbb{F}$ is that the flight must provide an EOBT and the UTOT must be within 60 minutes of current time. This constrains the calculation to only include flights with high quality trajectories driven by the EOBT predictions and within a reasonable lookahead time.

An illustration of the ADS metric across D10, DAL, and DFW is shown in the top subgraph of Fig. 10. We show the distribution of ETOT_F in blue and the distribution of ETOT_R in orange. The distribution of ETOT_R is a result of a single reroute from a DFW flight. The distributions are shown for the set of flights $\mathbb{F}$ within D10, DAL, and DFW as you go from left to right across the Fig.. As can be seen, rerouting a single flight impacts the ETOT_R for many flights in the schedule and the orange distribution is shifted to the left of the blue distribution. The result is the ADS across D10, DAL, and DFW is 56.4, 15.5, and 40.9 minutes.

The results for the ADS shown in the top subgraph of Fig. 10 were measured during a time period where the terminal boundary was the main constraint to the system. This occurs

	D10	DAL	DFW	Rerouted Flight
Total Savings	56.4	15.5	40.9	27.7
Average Savings	1.4	1.3	1.4	

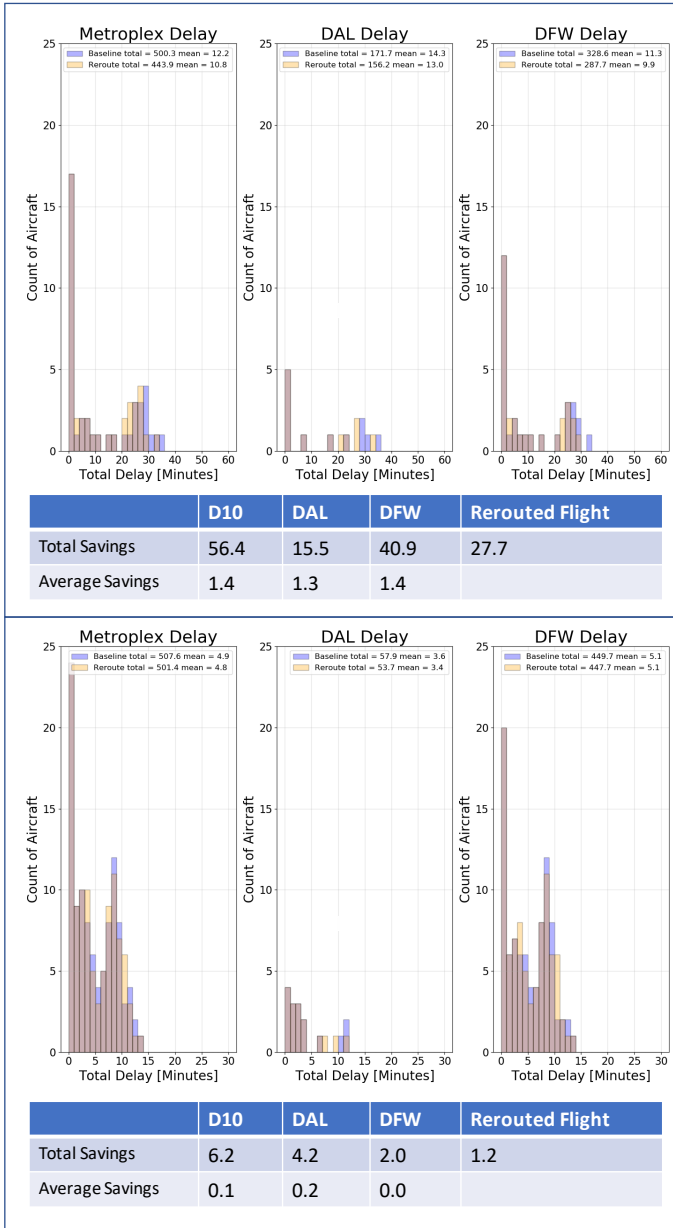|	D10	DAL	DFW	Rerouted Flight
Total Savings	6.2	4.2	2.0	1.2
Average Savings	0.1	0.2	0.0	

Fig. 10. Top: Example of system-wide aggregate delay savings when the terminal boundary is the main constraint. Bottom: Example of system-wide aggregate delay savings when the runway is the main constraint.

when the majority of the flow through the system is subject to MIT restrictions. For example, consider a situation where 100% of the flights are subject to MIT restrictions. The runway schedule would then show that between any two flights there is an empty slot at the runway. If a restricted flight is rerouted, the rerouted flight could occupy one of these empty slots at the runway without impacting any other flights. This would result in the rerouted flight receiving a benefit and in addition all flights on the restricted route that were originally behind the rerouted flight would move up one slot. The reroute of the single flight results in an ADS above and beyond the benefit to the individual flight.

If the runway is the main constraint on the system, then an ADS above and beyond the savings to the individual flight might not materialize because all the available slots at the runway would already be assigned. This typically happens when there is a mix of restricted and non-restricted flights where a non-restricted flight is able to occupy every available slot between any two restricted flights. In this situation, if a restricted flight is rerouted we will observe a benefit to the other restricted flights, as previously described, where each restricted flight moves up one slot. The rerouted flight, however, will no longer be able to occupy a vacant slot at the runway because all slots are occupied. The rerouted flight will jump ahead of some non-restricted flights and occupy a slot that was previously assigned to a different flight. The result is that at the system level the set of restricted flights will see a delay reduction, however, the set of non-restricted flights will see an increase in delay as the rerouted flight will end up delaying some non-restricted flights.

The bottom subgraph of Fig. 10 shows an example of the ADS when the runway is the main constraint to the system. If we compare the bottom subgraph of Fig. 10 to the top subgraph we see that the bottom subgraph contains more flights and more delay than the top subgraph. If the ADS was simply a function of demand we would expect the ADS to be much higher for the bottom subgraph. Instead, we measure an ADS across D10, DAL, and DFW as 6.2, 4.2, and 2.0 minutes, respectively.

The low value of ADS is driven by a situation where the runway is the main constraint. This can be observed in the histogram of delay values showing a bi-modal structure where the left mode is formed from the non-restricted flights and the right mode formed by the restricted flights. There is enough non-restricted demand at the runway that all slots are assigned the non-restricted flights are experiencing delay. This is in contrast to the top subgraph where the non-restricted flights have zero delay and the restricted flights are the only flights experiencing delay in the system.

V. CONCLUSION

In this paper, we described NASA ATD-2 Phase 3 scheduling in a Metroplex environment incorporating TOS options provided by flight operators. We began by describing the D10 TRACON and summarizing the demand and the restrictions that are the main driver of delay during inclement weather events. We described the TOS concept and explained the strategy of using TOS alternative trajectories to route around fix closures and MIT restrictions and explained the benefits to the rerouted flight and the aggregate system-wide benefits.

Next, we described at a high level the Phase 3 Terminal Scheduler which is composed of an Orchestrator, Departure Fix Scheduler, Trajectory Modeler, and Airport Scheduler for each airport in D10. We showed how the different components of the Terminal Scheduler are working within three subroutines and explained how the Departure Fix Scheduler and Airport Scheduler are working together within the core scheduling loop. We then described how the core scheduling loop is

leveraged for each TOS alternative route to calculate the individual and aggregated delay savings that would result from the reroute. We provided high level details describing the heuristics we follow when applying TCOT constraints.

Lastly, we defined the metrics that were developed to inform flight operators and ATC about the performance of the system and reroute opportunities. We showed the ETOT accuracy as a function of lookahead time prior to ATOT and recommended that the delay savings estimate should be considered in context with the accuracy of the system. We described our methodology to build a distribution of delay savings that incorporates both the delay savings prediction and the accuracy of the system and allows us to calculate the probability that the delay savings will exceed the RTC. We defined the system-wide aggregate delay savings metric and showed how the system-wide savings can be different when the terminal boundary or runway is the main constraint to the system.

Future work will focus on analyzing the results from the Phase 3 Field Evaluation. The Phase 3 Field Evaluation was intended for the Summer of 2020 but is now extended through the Summer of 2021 due to the impact of COVID-19. In addition to analyzing the results from the Field Evaluation, future work will also focus on extending the Phase 3 concept to other Metroplexes across the NAS. Of particular interest is the New York TRACON (N90) which has three major airports sharing capacity at the terminal boundary and interacting with each other. We believe that the lessons learned during the Phase 3 Field Evaluation in D10 will put us in a good position to tackle future challenges in the NAS.

REFERENCES

[1] Thipphavong, J., Jung, J., Swenson, H. N., Witzberger, K. E., Lin, M. I., Nguyen, J., Martin, L., Downs, M. B., and Smith, T. A., "Evaluation of the controller-managed spacing tools, flight-deck interval management, and terminal area metering capabilities for the ATM Technology Demonstration 1," *11th USA/Europe Air Traffic Management Research and Development Seminar, 2015*.

[2] Engelland, S. A., Capps, R., Day, K. B., Kistler, M. S., Gaither, F., and Juro, G., "Precision Departure Release Capability (PDRC) Final Report," 2013.

[3] Jung, Y., Malik, W., Tobias, L., Gupta, G., Hoang, T., and Hayashi, M., "Performance evaluation of SARDA: an individual aircraft-based advisory concept for surface management," *Air Traffic Control Quarterly*, Vol. 22, No. 3, 2014, pp. 195–221.

[4] Hayashi, M., Hoang, T., Jung, Y. C., Malik, W., Lee, H., and Dulchinos, V. L., "Evaluation of pushback decision-support tool concept for Charlotte Douglas International Airport ramp operations," *11th USA/Europe Air Traffic Management Research and Development Seminar, 2015*.

[5] Capps, A., Kistler, M. S., and Engelland, S. A., "Design characteristics of a terminal departure scheduler," *14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014, p. 2020.

[6] FAA Air Traffic Organization Surface Operations Office, "U.S. airport Surface Collaborative Decision Making (CDM) Concept of Operations (ConOps) in the near-term: application of the surface concept at United States airports," 2014.

[7] Jung, Y., Engelland, S., Capps, A., Coppenbarger, R., Hooey, B., Sharma, S., Stevens, L., and Verma, S., "Airspace Technology Demonstration 2 (ATD-2) phase 1 Concept of Use (ConUse)," 2018.

[8] Ging, A., Engelland, S., Capps, A., Eshow, M., Jung, Y., Sharma, S., Talebi, E., Downs, M., Freedman, C., Ngo, T., Sielski, H., Wang, E., Burke, J., Gorman, S., Phipps, B., and Morgan Ruszkowski, L., "Airspace Technology Demonstration 2 (ATD-2) Technology Description Document (TDD)," 2018.

[9] Jung, Y., Coupe, W., Capps, A., Engelland, S., and Sharma, S., "Field evaluation of the baseline integrated arrival, departure, surface capabilities at Charlotte Douglas Interntional Airport," *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)*, 2019.

[10] Coupe, W. J., Bagasol, L., Chen, L., Lee, H., and Jung, Y., "A data driven analysis of a tactical surface scheduler," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, 2018.

[11] Coupe, W. J., Jung, Y., Lee, H., Chen, L., and Robeson, I., "Scheduling improvements following the Phase 1 field evaluation of the ATD-2 integrated arrival, departure, and surface concept," *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)*.

[12] Kistler, M. S., Capps, A., and Engelland, S. A., "Characterization of Nationwide TRACON Departure Operations," *14th AIAA Aviation Technology, Integration, and Operations Conference*, 2014, p. 2019.