TOWARDS SUSTAINABLE AVIATION WITH EFFICIENT AIRSPACE OPERATIONS

William J. Coupe[1] & Swati Saxena[1]

[1]NASA Ames Research Center, Moffett Field, CA 94043

Abstract

In November 2021, the Federal Aviation Administration published the United States Aviation Climate Action Plan to accelerate innovation across the U.S. aviation ecosystem. In response, NASA has established the Sustainable Flight National Partnership to engage with industry, academia, and other agencies to accomplish net-zero carbon emissions by 2050. As part of the SFNP Mission, NASA is conducting a series of real world operational demonstrations in the current National Airspace System with a focus on delivering real world sustainability benefits. This paper presents the current work and future plan for the SFNP Ops Demo Series and describes the cloud based infrastructure used for virtual deployment of decision support tools to flight operators and Air Traffic Controllers to help improve operational efficiency of the National Airspace System. Validation results are shared along with the key sustainability benefits such as jet fuel savings and reduction in CO_2 emissions. The operational efficiency metrics such as delay savings are also reported.

Keywords: NASA Sustainable Flight National Partnership, Sustainable Aviation Benefits, Digital Transformation, Machine Learning

1. Introduction

In November 2021, the Federal Aviation Administration (FAA) published the United States Aviation Climate Action Plan [1] to accelerate innovation across the U.S. aviation ecosystem and contribute towards the goal of net-zero greenhouse gas emissions from the U.S. aviation sector by 2050. Key stakeholders responsible for implementing the plan include airlines, manufacturers/suppliers, airports, energy companies, and government agencies.

In response to the United States Aviation Climate Action Plan, NASA established the Sustainable Flight National Partnership (SFNP) to engage with industry, academia, and other agencies to accomplish net-zero aviation carbon emissions by 2050 [2]. The SFNP Mission draws expertise across the portfolio of NASA Aeronautics Research Mission Directorate to develop new technologies in advanced vehicle technologies, Sustainable Aviation Fuels (SAF), and efficient air traffic operations, see Figure 1.

For aircraft infrastructure, NASA plays a primary role in collaboration with Boeing to develop the X-66, the first X-plane specifically focused on delivering to net-zero aviation emissions. The Transonic Truss-Braced Wing concept aims to achieve 30% less fuel consumption and reduced emissions when compared with today's best-in-class aircraft [3].

To validate and advance SAF research, NASA performs a combination of computational modeling, emissions and combustion laboratory testing, and flight testing. NASA is supporting the research to better understand the sustainability benefits of Sustainable Aviation Fuels (SAF) through different ongoing efforts such as by participating in flight campaigns to characterize emissions for different aviation fuels. As an example, NASA in collaboration with Boeing and other partners (GE Aerospace, the German Aerospace Center, the FAA, United Airlines) conducted a flight test in 2023 that collected crucial sensor data about the sustainable aviation fuel and its effects on condensation trail formation.

Figure 1 – Sustainable Flight National Partnership (SFNP) Mission.

For efficient air traffic operations, NASA plays a primary role through the Airspace Operations and Safety Program (AOSP) with work being executed by the Air Traffic Management - eXploration (ATM-X) project. Within the ATM-X project, the Digital Information Platform (DIP) sub-project has established partnerships with the FAA, the National Air Traffic Controllers Association (NATCA), five major US airlines and one regional airline (American Airlines, Delta Air Lines, jetBlue Airways, Southwest Airlines, United Airlines, and Envoy Airlines) to conduct a series of SFNP Operational Demonstrations (SFNP Ops). SFNP Ops brings together a cloud based ecosystem of digital services, test vehicles, and partnerships with major US flight operators to develop and demonstrate, in an operational environment, concepts and technologies delivering to the goal of net-zero greenhouse gas emissions.
This paper presents the work done so far for the first SFNP Ops Demo and the plan for the future demos with a focus on delivering real world sustainability benefits. Section 2. introduces the series of Ops demos that DIP will be conducting to contribute to NASA's SFNP Mission. Section 3. introduces NASA's Digital Information Platform which is cloud infrastructure enabling virtual deployment of Decision Support Tools (DSTs) to both flight operators and Air Traffic Control (ATC) to improve efficiency of the airspace operations. Section 4. provides additional details of the first demo executed between January 2022 and May 2024. Additionally, validation results are provided and benefits in both delay savings and fuel and emissions reduction are reported. Section 5. provides a summary discussion of activities and results.

2. Sustainable Flight National Partnership Ops Demo Series

SFNP Ops is a series of four operational demos NASA is conducting between 2022 and 2030, see Figure 2. A core concept of the demo series is that each demo does not live in isolation, rather the demos are designed to build on and extend capabilities developed in previous demos. This enables the introduction of incremental capabilities that can be tested and evaluated in a crawl, walk, run approach.
The first demo focuses on single flight pre-departure Trajectory Option Set (TOS) rerouting. To enable this, NASA has developed the Collaborative Digital Departure Reroute (CDDR) service that alerts flight operators to pre-departure reroute opportunities and enables electronic coordination between flight operators and ATC. The pre-departure TOS reroutes deliver benefits in delay savings and both fuel and emissions reduction. The CDDR service was originally deployed to North Texas (Dallas-Fort Worth Terminal Radar Approach CONtrol (TRACON)) in 2022 and will later be deployed to Houston

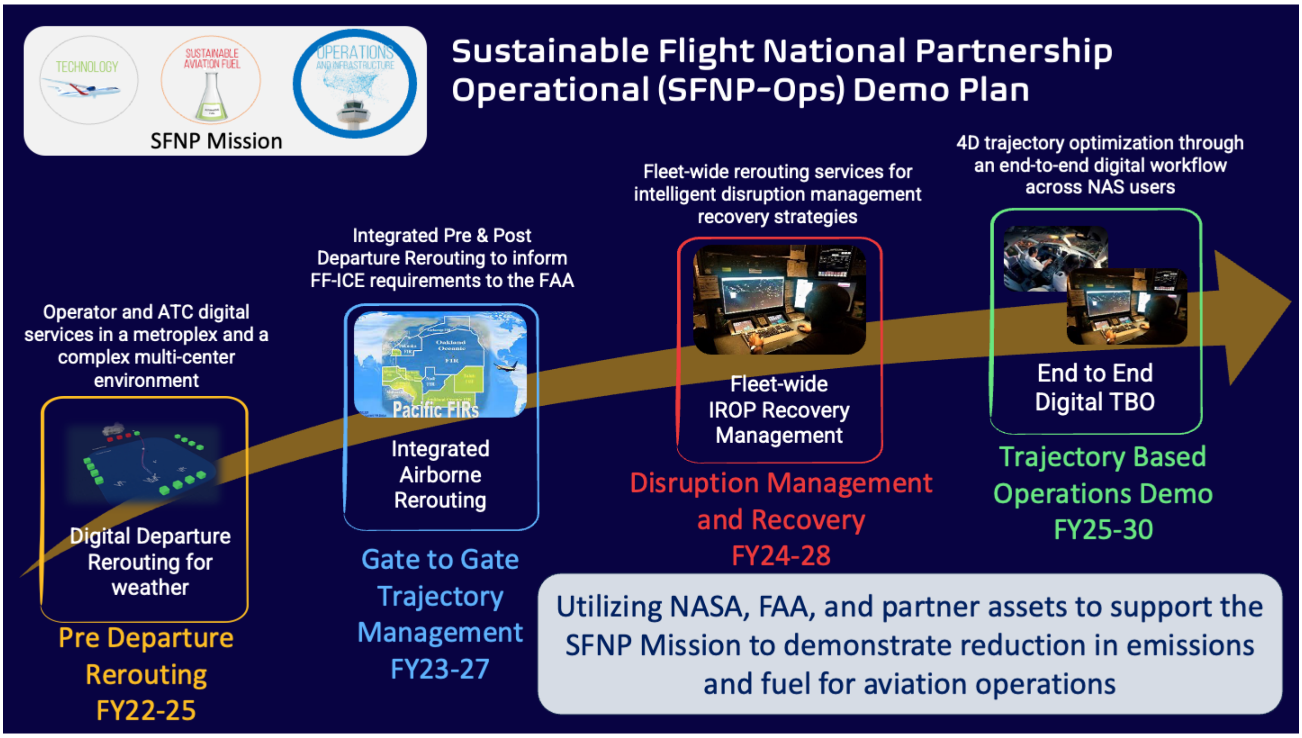

Figure 2 – DIP SFNP Ops Demo Series.

TRACON in 2025.
The second demo builds upon the single flight pre-departure capability to introduce airborne trajectory management and reroutes. This enables gate-to-gate trajectory management for sustainability use cases including but not limited to convective weather avoidance, turbulence avoidance, and contrail avoidance. To enable the reroutes, the demo will prototype and evaluate future National Airspace System (NAS) concepts including Flight and Flow - Information for Collaborative Environment [4] (FF-ICE) to help inform future trajectory management requirements.
The third demo extends the single flight use case to develop a fleet-wide Decision Support Tool (DST) to help flight operators and ATC manage large scale disruptions and Irregular Operations (IROP) across the network. This demo will be focused on both helping stakeholders anticipate future IROP disruptions and recover from ongoing disruptions. Large scale disruptions of interest include, but are not limited to, weather events that impact large portions of the NAS and space launch events.
The fourth demo is a capstone demo that pulls together capabilities from the previous demos while introducing increasing levels of automation. A special focus of the fourth demo will be to define the end-to-end digital workflow between flight operators and ATC within the emerging NAS architecture which enables Trajectory Based Operations (TBO). The digital workflow will leverage automation tools to help reduce workload of key decision makers while maintaining situational awareness of the operation.

2.1 Previous NASA Operational Demos

NASA has a long history of developing and field testing new technologies in the NAS to help manage operations. This work started in the mid 1990's with the Center TRACON Automation System (CTAS) [5] and the Traffic Management Advisor (TMA) [6]. The CTAS/TMA tools developed by NASA were evaluated at the Ft. Worth Air Route Traffic Control Center (ARTCC) [7] and later tech transferred to the FAA and became Time Based Flow Management (TBFM) [8]. TBFM is a core DST for time-based management in the en route and terminal environments.
Building on work done for arrivals, departures and surface operations were incorporated using TBO concepts by NASA, the FAA, and industry to improve the flow of traffic into and out of the nation's busiest airports. NASA technologies for specific phases of flight were integrated[9] across surface[10, 11, 12] and airspace domains[13] and deployed as the Integrated Arrival Departure and Surface

(IADS) traffic management system[14, 15] in 2017 to Charlotte Douglas International Airport. The IADS system was developed in alignment with FAA's Surface Collaborative Decision Making (S-CDM) Concept of Operations[16] and refined over time[17]. This technology was tech transferred to the FAA and became the surface management solution known as Terminal Flight Data Manager (TFDM) [18]. Although these demos extended across three decades and addressed different use cases and capabilities, all the previous NASA work shared a common challenge. Deployment and evaluation of new capabilities required physical deployment of NASA hardware and network components to enable delivering of DST to users of the system. While this approach enabled field evaluations, it set a high bar for deployment and evaluation. To address this challenge, NASA has developed the Digital Information Platform to enable virtual deployment of DSTs into both flight operator and ATC facilities. This reduces complexity of the deployment process and enables accelerated evaluation of new technologies across the NAS.

3. Digital Information Platform

Future needs of the NAS require DSTs to adopt an architecture as described by the FAA's vision for an Info-Centric NAS[19]. To align with the vision, DSTs need to migrate toward learning, adaptable, and lightweight interacting systems. To achieve this, many existing systems will need to undergo a digital transformation from a monolithic decision support tool to a service-oriented architecture where individual services are exposed through a well defined Application Programming Interface (API). This enables industry services to be deployed alongside FAA infrastructure and services to supplement existing investments in TBO capabilities.

NASA has developed the Digital Information Platform (DIP)[20] to enable the transformation towards an Info-Centric NAS. The overall goal of DIP is to support transformation of the NAS through the development of a cloud-based platform for advanced, data-driven, digital services for both traditional and emergent operations. DIP provides access to real-time and historical data upon which data-driven services for NAS users are developed and operated. The DIP ecosystem helps accelerate the development of Artificial Intelligence (AI) and Machine Learning (ML) techniques applied to aviation and provides an opportunity for AI/ML experts from outside traditional aviation to make an immediate impact through the development of novel services.

3.1 Predictive Engine: Machine Learning Airport Surface Model

A core capability enabling SFNP Ops is an airport surface model which generates demand predictions as a function of constraints defined by ATC input. Much of this capability was originally developed as the IADS system deployed as part of the Airspace Technology Demonstration 2 (ATD-2) Sub-project[14].

The IADS systems generated predictions including but not limited to airport configuration, runway assignment, unimpeded taxi times, and arrival ON times[17]. The airport surface model predictions were used as input by the IADS Terminal Scheduler which applied all known constraints across each airport surface and the terminal boundary[21] to generate predictions for the Estimated Take Off Time (ETOT) for each departure flight. To generate the airport surface model predictions the IADS system relied upon detailed adaptation developed for each individual airport and the terminal airspace.

Adaptation for each airport requires creating a detailed structure of a link-node network defining the gate locations, ramp and taxiway structure, and runway locations. The adaptation goes beyond defining the physical structure of the airport and also requires detailed knowledge from ATC encoded in decision trees including departure fix to runway mappings and other local knowledge that might be applied to the airport or airspace. Creation of the adaptation is often a manual process that requires significant time and effort to both build and maintain.

The system as developed under ATD-2 was designed as a monolithic DST leveraging physics based models with adaptation which created a bottleneck to scalability [22]. To align with the FAA's vision for an Info-Centric NAS and to address scalability challenges, the DIP Sub-project applied digital transformation of the IADS system. The digital transformation resulted in the Machine Learning Airport Surface Model which was developed with a service-oriented architecture and deployed in a cloud environment. Modern Machine Learning techniques were used in place of legacy physics based models relying on adaptation.

3.1.1 Digital Transformation

Figure 3 illustrates the digital transformation of the ATD-2 IADS system to the DIP Machine Learning Airport Surface Model. Starting with the monolithic decision support tool, key capabilities were broken out as individual services. For predictive services, ML was applied to replace the physics based services that relied on adaptation. Outputs of each individual service are made available through well defined APIs deployed on NASA's Digital Information Platform (DIP). Each individual service can be consumed and used as building blocks for downstream service developers. By making the services available through API, others can benefit from the ML Airport Surface Model accelerating the development cycle for new capabilities that require these foundational data elements.

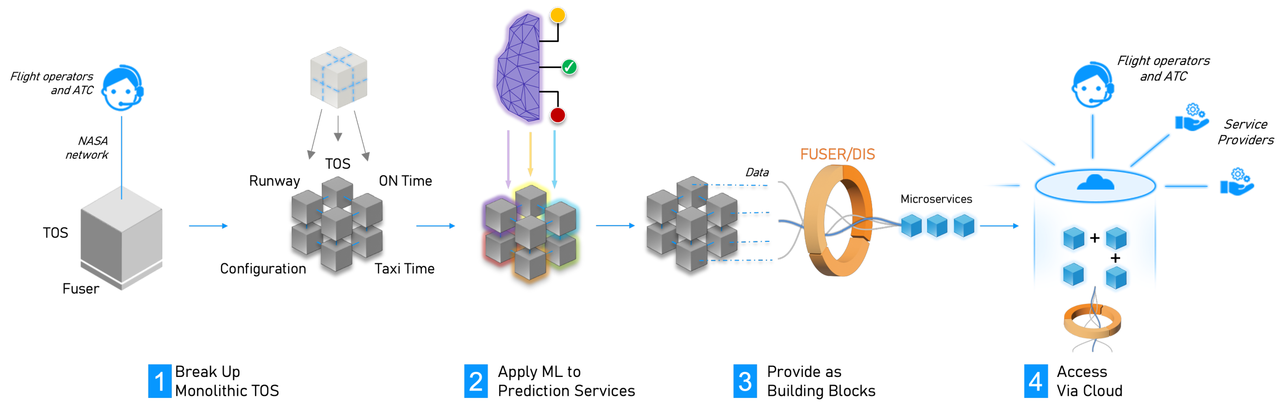

Figure 3 – Digital transformation of NASA's predictive engine towards FAA's vision for an Info-Centric NAS: Learning, adaptable, and lightweight interacting systems.

Figure 4 shows a detailed view of the ML Airport Surface Model architecture. The ML Airport Surface Model is deployed in a service-oriented-architecture in which each logical service is distinct with well-defined inputs and outputs. The ML Airport Surface Model starts at the bottom of the figure from a foundation of raw data feeds including FAA System Wide Information Management (SWIM) data feeds and other available airline or airport data feeds. The raw feeds contain valuable data, but can provide inconsistent information on the same flight that is difficult to reconcile in a real-time environment.

To address this challenge, NASA developed logic that could resolve data processing and mediation complexities. Much of this work is embodied in the Fuser service [23]. Both the Fuser and the ML Airport Surface Model are intended to supplement existing and planned FAA capabilities such as the SWIM data feeds. The Fuser framework mediates between the disparate sources of data, pulling in the right data, at the right time. The Fuser leverages heuristics and analysis on which data source is best to use for a specific need and provides access to the information in a well-defined, common data model.

The Fuser data is used as input to the airport surface model orchestrator. The orchestrator also pulls in weather data, current airport configuration data in the form of Digital Automatics Terminal Information Service (D-ATIS), and restriction data from NASA's Traffic Management Initiative (TMI) service. The TMI service combines restriction data from FAA SWIM data feeds in addition to local restrictions only available on the Operational Information System (OIS) page. The restriction data is parsed to identify individual restrictions correlated and assigned at the flight level by the TMI service prior to being passed as input to the orchestrator.

The orchestrator is responsible for collecting the inputs required by each ML prediction service and also responsible for calling the ML services in the proper order. Even though each service is distinct with well defined inputs and outputs there are dependencies between the different ML prediction services that need to be accounted for. Figure 4 shows the dependencies between the services as the output of the airport configuration service is used as input to the runway service. Similarly, the output of the runway service is used as input by the taxi time service and the arrival ON time service. Outputs of the runway, taxi-time, and arrival ON time service are used as input by the NAS Terminal Scheduler.

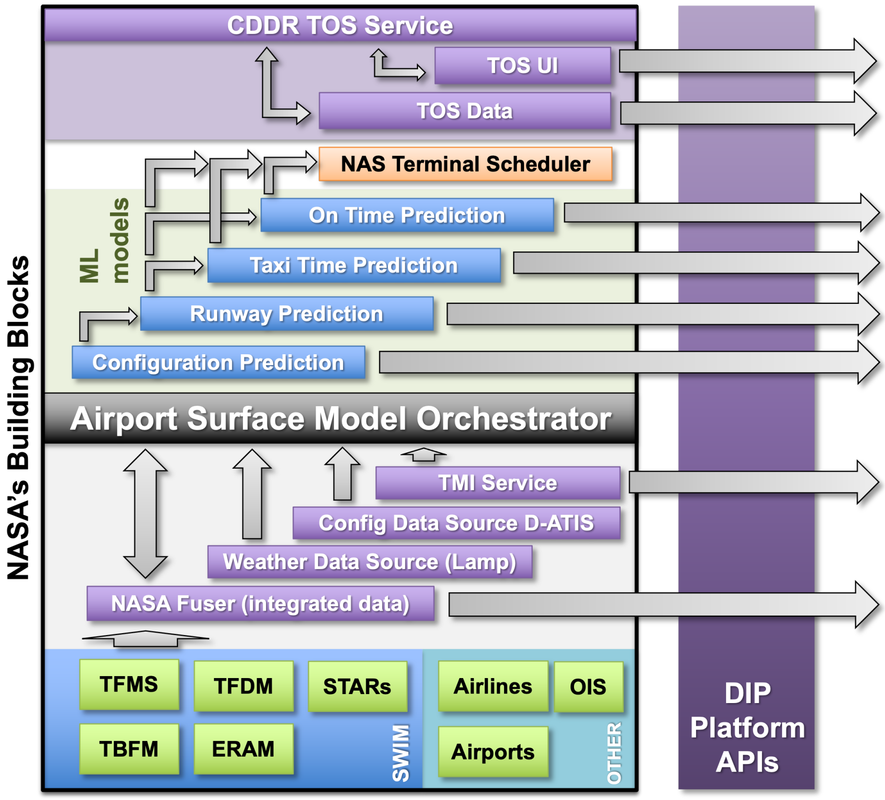

Figure 4 – Predictive engine service-oriented architecture. TFMS: Traffic Flow Management System, TFDM: Terminal Flight Data Manager, STARs: Standard Terminal Automation Replacement System, TBFM: Time-Based Flow Management, ERAM: En Route automation modernization, SWIM: System Wide Information Management, OIS: Operational Information System, D-ATIS: Digital-Automatic Terminal Information Service

3.1.2 Machine Learning Operations (MLOps)

ML techniques have been applied to aviation problems for many years[24] to develop prediction models. However, the real challenge isn't building an ML model; the challenge is building an integrated ML system and to continuously operate it in production[25]. Without the proper approach, it is easy to incur massive ongoing maintenance costs at the system level when applying machine learning[26]. To address this challenge, in recent years there has been focused work on Machine Learning Operations (MLOps) to develop infrastructures and platforms for end-to-end life-cycle management of ML[25, 27, 28].

For deployment of the ML Airport Surface Model, we adopt MLOps best practices across the real-time system and the off-line training infrastructure. Figure 5 illustrates the real-time system and the off-line training infrastructure and the relationship between the two. The adoption of MLOps best practices described in this Section helps reduce risk in deployment to ensure both the models and the pipelines feeding the models are consistent between the off-line training infrastructure and the real-time deployment. MLOps best practices also allow for automation, reproducibility, monitoring, and continuous integration of ML into production software.

The off-line training infrastructure is shown in the top of Figure 5 and begins with a historical archive of the data used as input to the Orchestrator shown in Figure 4. These data are ingested by the ML model training pipelines that are composed of three phases: data query, data engineering, and model building and training. The trained models are then validated against out of sample test sets using k-fold cross validation and other validation analysis. The trained models are then stored in MLFlow and version control of the model training pipelines is maintained via Bitbucket.

The real-time system is shown in the bottom of Figure 5 and begins with streaming data feeds that are fused and mediated in real-time by the Fuser prior to being used as input to the Orchestrator shown in Figure 4. It is important to ensure that any data engineering techniques applied to the historical

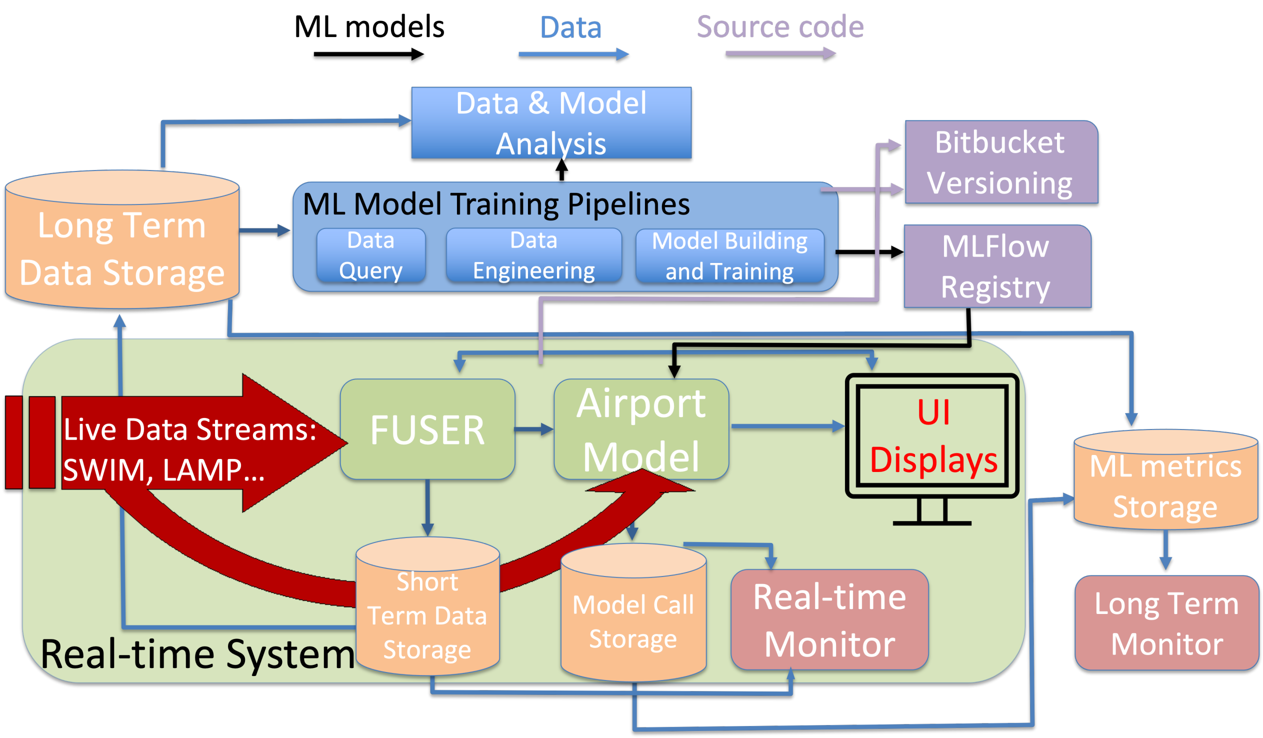

Figure 5 – Machine Learning Ops (MLOps) infrastructure.

data are replicated in the real-time feeds or else model performance will degrade as the distribution of features in the training data will not match the features provided to the model in real-time. The real-time fused data feeds and the features delivered to the on-line model are stored in databases and used as reference in a real-time performance monitor. Nightly queries are performed to summarize model performance and stored in a ML metric storage database allowing for long-term monitoring of model performance.

The real-time performance monitor has been a valuable tool during deployment of the system to provide a tight feedback loop between the system and engineers. Often times it is easier to identify and debug issues while tracking the real-time performance as opposed to aggregate performance. To track real-time performance, we adopted a strategy to measure performance of models conditional on the input received. As shown in Figure 4, the ML models have dependencies with upstream models and it is important to understand these dependencies if the performance of a particular model is degrading or if the inputs fed to the model are degrading.

4. SFNP Ops Field Evaluation

The SFNP Ops demos as described in Section 2. will be executed in a series of real world field evaluations. The remainder of Section 4. provides additional details on the concept and field evaluation results for the Collaborative Digital Departure Reroute demonstration.

4.1 SFNP Ops 1: Collaborative Digital Departure Reroute Concept

In response to weather events around or near the terminal boundary, the TRACON Traffic Management Unit (TMU) will close departure fixes affected by inclement weather which results in the departure gate being partially or completely blocked. The departure gate is the collection of departure fixes along each side of the terminal boundary. Figure 6a illustrates a situation where three of the four departure fixes along the top departure gate have been closed and traffic through these fixes is rerouted to the single remaining open fix within the departure gate. This compression of the departure fixes reduces the capacity at the terminal boundary and delays can be amplified when ATC enforces additional departure fix restrictions such as Miles-In-Trail (MIT).

When Traffic Management Initiative (TMI) restrictions, e.g. departure fix closures, reduce the capacity at the terminal boundary, there are often opportunities to route around the restricted area and reduce

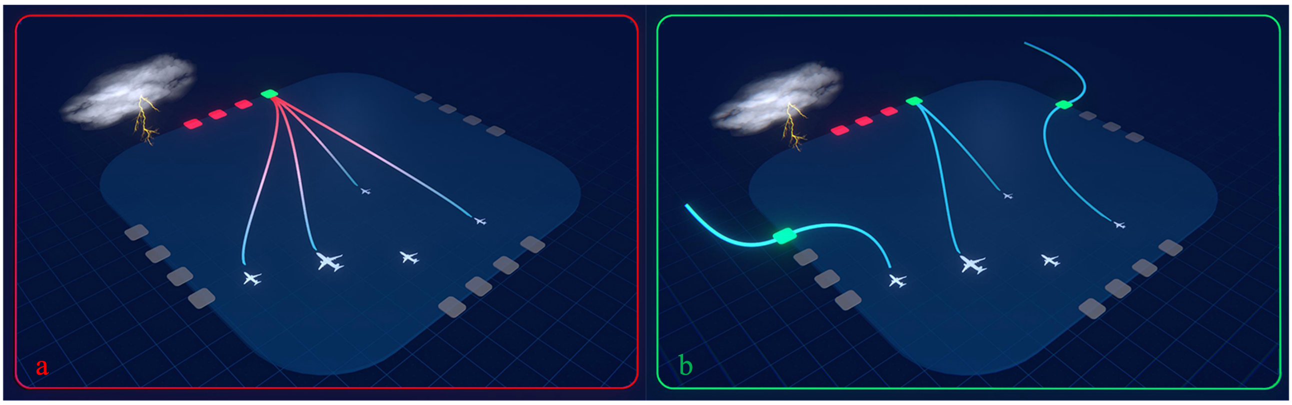

Figure 6 – SFNP Ops Collaborative Digital Departure Reroute a) problem and b) solution.

the surface delay. Figure 6b shows the situation where the top departure gate is limited to a single fix with a potential MIT restriction, while the left or right departure gate have all four fixes available. When the traffic volumes through the left and right departure gate are relatively light and the routes are not impacted by a TMI restriction, a flight could reroute through the left or right departure gate with little to no delay on the airport surface.

A flight operator defines the TOS which is the set of feasible routes for a given flight. During earlier NASA field evaluations, flight operators and ATC agreed to use vetted Coded Departure Routes (CDR) as the available TOS routes for departures. The filed route is typically the most direct route and is preferred by the flight operators under nominal operations. The cost of each TOS alternative route option, 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 longer route in exchange to reduce surface delay. When the delay savings on the surface exceeds the RTC threshold, the alternate route provides enough delay savings to be worth the cost of the longer route.

The predictions are delivered to users through a custom User Interface (UI)[29] developed by NASA. The UI is delivered to facilities including: Dallas/Fort Worth International Airport (KDFW) and Dallas Love Field Airport (KDAL) Air Traffic Control Towers, Dallas-Fort Worth TRACON (D10), Fort Worth Air Route Traffic Control Center (ZFW), American Airlines Integrated Operations Center, Southwest Airlines Network Operations Center, and Envoy Airlines Headquarters. The UI provides the predicted delay savings on the filed route and each TOS alternative route, which forms the basis for the reroute recommendations provided by the tool. In addition to the delay savings predictions, the UI enables digital coordination between flight operators and ATC[29, 15].

4.2 CDDR Field Evaluation Prediction Accuracy

The CDDR service was deployed to North Texas and used by American Airlines, Southwest Airlines, Envoy Airlines, and ATC in an operational field evaluation. For evaluation of the ML Airport Surface Model (pure-ml) performance, we compare against performance of the legacy IADS adaptation based Airport Surface Model (legacy) as baseline. We use data collected during the Stormy 2022 Field Evaluation between April 29 2022 through September 16 2022. During this time period at KDFW, there were 127,898 departures and 126,721 arrivals used for validation of the prediction services.

Figure 7 shows the performance of the predictive services deployed as part of the ML Airport Surface Model. The left and right sides of Figure 7 show arrival flight predictions sampled at crossing of the arrival fix and departure flight predictions sampled at the push back event, respectively. The top row illustrates runway prediction accuracy, the middle row illustrates the distribution of the arrival ON time error and departure ETOT error, and the bottom row shows the arrival ON time Standard Deviation (STD) of error and departure ETOT STD of error. Predictions from the legacy IADS system and the Machine Learning Airport Surface Model are shown in blue and orange, respectively.

For runway prediction accuracy shown in the top row of Figure 7, on each day we measure the percentage of flights with correct runway predictions and plot that with a small circle. The 14 day

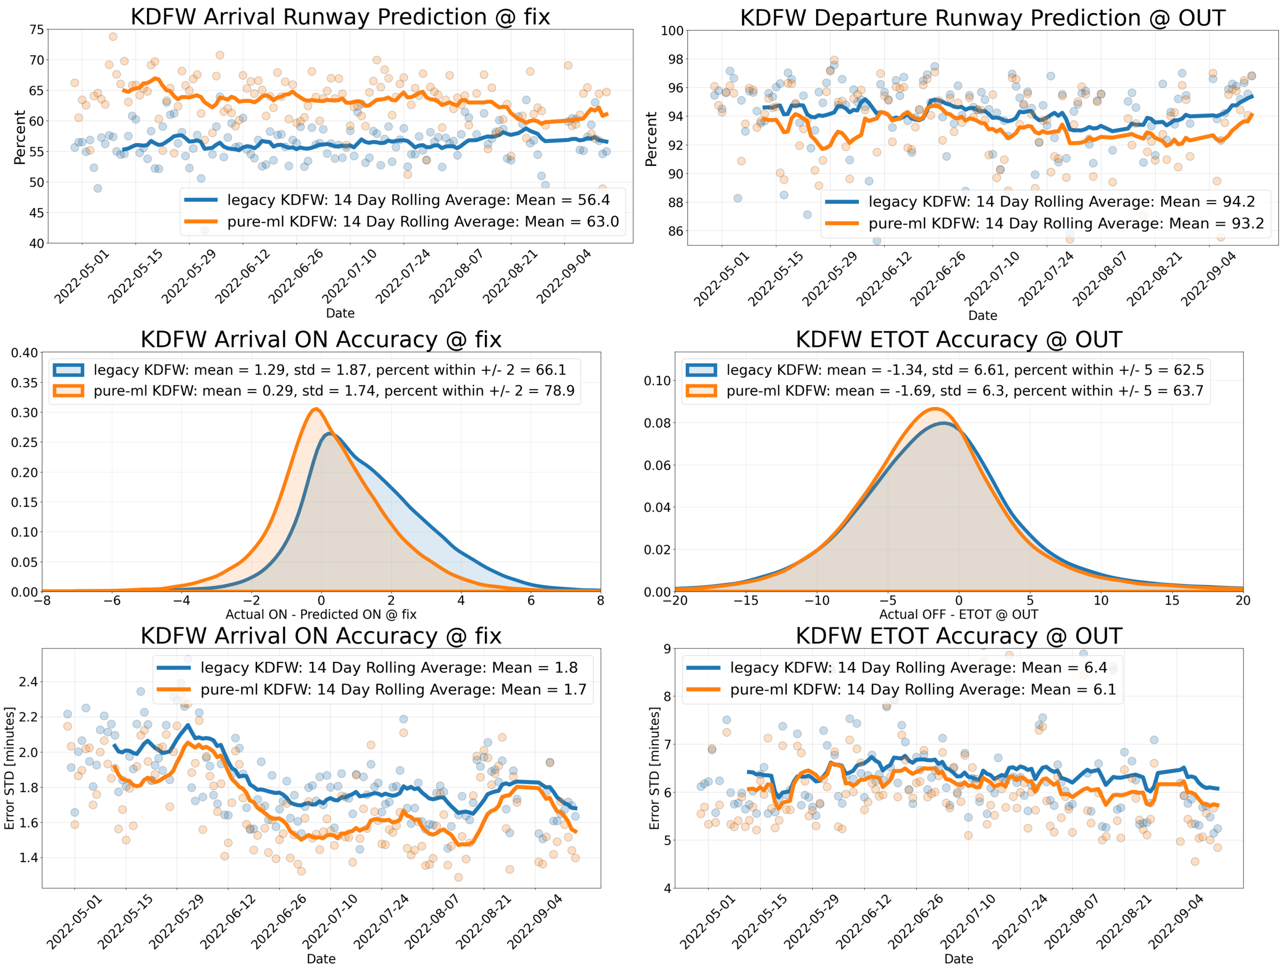

Figure 7 – Machine learning vs. physics based models.

rolling average of the prediction percentage is plotted with a solid line and provides an indication of performance averaged across multiple days and situations.

The overall accuracy of the ML arrival runway prediction model (63.0%) outperforms accuracy of the legacy arrival runway model (56.4%). The legacy arrival runway model uses TBFM predictions to generate the arrival runway. Not only is the overall percentage performance of the pure-ml system above the legacy system, but the cloud of orange dots shows separation and sits above the cloud of blue dots, indicating that the pure-ml system is outperforming the legacy system for arrival runway on a consistent day-by-day basis.

The overall accuracy of the ML departure runway prediction model (93.2%) is very close to the accuracy of the legacy departure runway model (94.2%). Not only is the overall accuracy quite close, but the rolling average shows similar trends between the pure-ml model performance compared to legacy model performance, which indicates both systems are impacted in similar ways.

The performance of the departure runway model is quite encouraging because the legacy system at KDFW is not truly predicting the departure runway. Instead, the legacy system allows for ATC to set a taxi-plan that defines a departure fix to runway mapping which *assigns* departures to runways according to the way in which ATC wants to load balance the demand across runways. The ML departure runway model is only one percentage point below the accuracy of the legacy ATC assignment via the taxi-plan. This is encouraging and gives us confidence that the ML system could be deployed with or without receiving ATC input regarding load balancing strategies for departures and maintain comparable performance.

For arrival ON time predictions shown in the middle row left subplot of Figure 7, the legacy arrival ON model uses a decision tree which selects TBFM Scheduled Time of Arrival (STA) when available, else uses TFMS Earliest Time of Arrival (ETA). The arrival ON model is evaluated using the error in prediction measured as the difference between the actual ON time minus the predicted ON time sampled at the arrival fix crossing. The horizontal axis is the error measured in minutes where a positive value indicates the actual ON time was later than the predicted ON time. As can be seen comparing the legacy distribution to the ML distribution, the ML arrival ON time prediction is outperforming the legacy arrival ON time prediction with a standard deviation of 1.74 minutes for the pure-ml system compared to 1.87 minutes for the legacy system.

The main improvement for the arrival ON time prediction model using the ML system is the mean error. The legacy system has mean error of 1.29 minutes compared to the ML system with mean error 0.29 minutes. Due to the reduction in mean error the percentage of arrival flights with prediction within a $+/-$ 2 minute window increased from 66.1% for the legacy system up to 78.9% for the ML system.

For analysis of departure ETOT prediction accuracy, we restrict our attention to American Airlines major flights at KDFW, as these flights were participants in the field evaluation and provided Earliest Off-Block Time (EOBT) predictions. For the middle row right subplot of Figure 7, the horizontal axis represents the difference in minutes between the actual OFF time and ETOT where a positive value represents the flight took off later than the prediction.

The departure ETOT prediction accuracy is quite similar between the ML system and the legacy system. The ML system has a slightly lower standard deviation of 6.3 minutes compared to the legacy system of 6.61 minutes. The ML system has mean error -1.69 minutes compared to -1.34 minutes. Even with the slightly larger bias, the tighter standard deviation results in the ML system having a higher percentage of flights fall within $+/-5$ minutes (63.7%) compared to the legacy system (62.5%).

The bottom row of Figure 7 shows the date on the horizontal axis and the STD of error for the arrival ON time prediction and departure ETOT prediction on the vertical axis. The STD for each day is plotted with a small circle and the 14 day rolling average is plotted with a solid line. For both arrival and departure predictions, the ML system shows an improved STD of error for the ML system throughout the time range during a variety of operating conditions.

4.3 CDDR Field Evaluation Delay Savings Benefits

During the field evaluation, the ML Airport Surface Model was used to reroute departure flights according to the concept described in Section 4.1. A core capability of the ML Airport Surface Model is predicting an individual flight's OFF Delay Savings (ODS) for a TOS route. The ML Airport Surface Model is responsible for generating ETOT predictions on the filed route and each TOS alternative route which can be used to calculate delay savings predictions. The OFF event is the time the departure takes off from the origin airport. The predicted ODS on a given TOS route is sampled at the pushback event and defined as:

$$ODS_T = TT_F - TT_T \tag{1}$$

where TT_F and TT_T represent the predicted Taxi Time (TT) on the original filed route and the TOS alternative route, respectively. A positive value represents the TOS route is beneficial as the predicted taxi time on the TOS alternative route is less than the predicted taxi time on the original filed route.

The system also provides an IN Delay Savings (IDS) prediction with respect to the benefit of getting to the destination airport on the TOS alternative route. The IN event is the time that the flight arrives at the arrival gate (or stand). The IDS is defined as:

$$IDS_T = ODS_T - AFT_T \tag{2}$$

where AFT_T represents the Additional Flight Time on the TOS alternative route. The AFT_T is calculated based on the difference between the ground miles of the TOS route and the original filed route divided by the filed flight speed. A positive value for the IDS_T represents the TOS route would arrive at the destination earlier than the original filed route.

In addition to the benefit to the individual rerouted flight, the system provides a prediction of the benefit at the system level. For each TOS alternative trajectory we calculate an Estimated Take Off Time on the TOS route $ETOT_T$ for the rerouted flight and $ETOT_R$ for the rest of the flights in the schedule under the assumption of the TOS reroute. We define the system-wide Aggregate Delay Savings (ADS) associated with a given TOS route as:

$$ADS_T = ODS_T + \sum_{\mathbb{F}} \left(TT_F^* - TT_T^* \right) \tag{3}$$

which is the OFF Delay Savings to the rerouted flight plus a sum over the set of flights $\mathbb{F}$ of the difference in taxi time $TT_F^* - TT_T^*$ for other flights under the assumption of the TOS reroute. When a single flight is rerouted and the reroute results in $ETOT_T$ on the TOS route not equal to $ETOT_F$ on the filed route, the change propagates through the schedule and other flights' ETOTs can be updated. The result can be that flights that are not rerouted have taxi time TT_T^* (assuming the TOS reroute) not equal to TT_F^* (assuming the original filed route), thus the system-wide ADS_T measure changes.
When a flight is rerouted, the system calculates the *actual* delay savings which could differ from the predicted delay savings. The actual delay savings is calculated by comparing the predictions of taxi time on the filed route to the actual taxi time on the TOS alternative route. The Actual OFF Delay Savings ($\mathbb{ODS}$) is defined as:

$$\mathbb{ODS}_T = TT_F - \mathbb{TT}_T \tag{4}$$

where TT_F represent the predicted Taxi Time on the filed route, sampled at the OUT event, and $\mathbb{TT}_T$ represents the actual taxi time on the TOS alternative route.
Similarly, the Actual IN Delay Savings ($\mathbb{IDS}$) is defined as:

$$\mathbb{IDS}_T = \mathbb{ODS}_T - AFT_T \tag{5}$$

which is the Actual OFF Delay Savings $\mathbb{ODS}_T$ minus the Additional Flight Time associated with the TOS route AFT_T.
Between January 1st, 2022 and May 4th, 2024 there were 108 flights rerouted and used for analysis. The list of flights included for benefits analysis was agreed upon between NASA, flight operators, and FAA participants. Reasons that flights were removed from benefits analysis include but are not limited to pilots refusing the reroute, challenges with non-Controller Pilot Data Link Communications (CPDLC) equipped flights changing the route in the Flight Management System (FMS), ATC not amending the flight plan after approval, airport flow change after the approval, and general outliers.
Figure 8 shows the distribution of predicted and actual delay savings for rerouted flights. The left column shows the distribution of OFF Delay Savings, IN Delay Savings, and Aggregate Delay Savings in the top, middle, and bottom row, respectively. The average predicted ODS_T was 12.8 minutes and the average actual $\mathbb{ODS}_T$ was 8.7 minutes. The IDS shows similar results, with predicted IDS_T benefit of 13.2 minutes and actual $\mathbb{IDS}_T$ benefit of 9.1 minutes. The average system-level ADS_T benefit was 29.8 minutes which is more than double the benefit to the individual flight.
This difference in the system-wide benefits indicates that when there are restrictions on the system a single reroute can have large benefits in reducing system level delay. When the individual flight is rerouted off the restricted route, all subsequent flights move up one slot and the system makes better use of the available capacity.
There is a slight reduction between the average predicted and actual delay savings, but overall the majority of predicted benefit was captured by the rerouted flight. It is also encouraging to see the overall shape of the green (predicted) and purple (actual) benefits distributions look so similar. This is an indication that the system predictions are relatively accurate and the benefits can be realized.
It is interesting to see that the Actual IN Delay Savings $\mathbb{IDS}_T$ of 9.1 minutes is greater than the Actual OFF Delay Savings $\mathbb{ODS}_T$ of 8.7 minutes. Given that the filed route is typically the most direct route we expect the Additional Flight Time AFT_T on the TOS route to be greater than zero, thus Equation (5) should be less than Equation (4). For flights rerouted at the end of Severe Weather Avoidance Programs (SWAP) events, however, the TOS route is typically much shorter than the SWAP route and the benefit of the shorter route results in a larger amount of $\mathbb{IDS}_T$ at the destination.

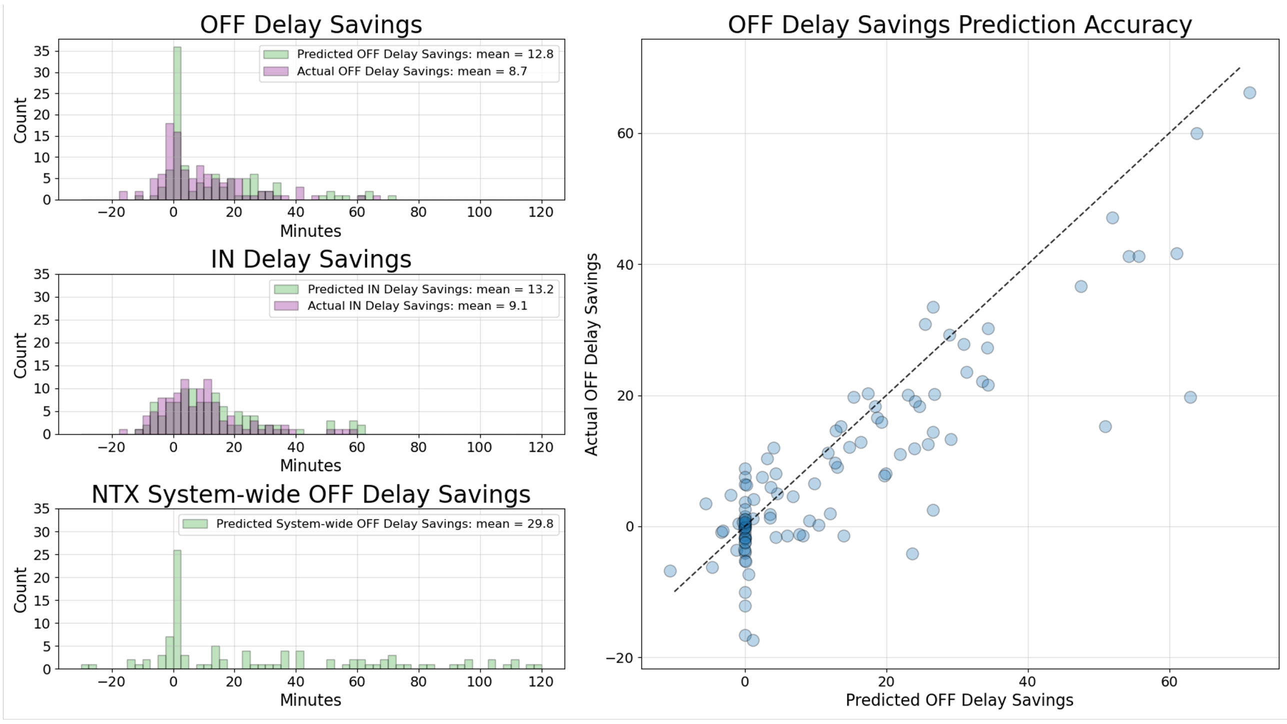

Figure 8 – Predicted vs actual delay savings benefits.

The right column of Figure 8 further illustrates the relationship between the predicted ODS_T and the actual $\mathbb{ODS}_T$ on a flight by flight basis. The horizontal and vertical axis represent the predicted and actual OFF Delay Savings, respectively. Each blue dot represents a rerouted flight and it is encouraging to see the data follow the pattern along the diagonal dashed line which represents the values where the actual values are equal to the predicted values.

The benefits of reroutes at the tail end of SWAP events can also be seen in the right column of Figure 8. There is a large chunk of flights which were rerouted with zero predicted ODS_T on the horizontal axis. These flights appear to have no benefit when only considering the ODS_T, however, they were rerouted to *shorter* routes. The average flight with zero predicted ODS_T was rerouted to a TOS alternative route with 4.7 minutes shorter flight time resulting in both fuel and emissions benefits. In total, all flights rerouted during the field evaluation saved 15.7 hours of actual $\mathbb{ODS}_T$ and 16.4 hours of $\mathbb{IDS}_T$. The actual $\mathbb{IDS}_T$ for rerouted flights is converted to passenger value of time savings $109,009 and $24,063 in flight crew cost savings [30].

4.4 CDDR Field Evaluation and Estimated Environmental Benefits

The delay savings benefits reported in Section 4.3 are converted into estimated environmental benefits for each rerouted flight. The environmental benefits account for both the surface fuel flow (OFF Delay Savings) and the airborne fuel flow (Additional Flight Time). The $\mathbb{ODS}_T$ is used with a NASA developed surface fuel flow model to calculate the surface fuel savings [30]. The AFT_T associated with the TOS route is combined with flight operator provided airborne fuel flow models to calculate the airborne fuel cost. The surface fuel savings plus the airborne fuel cost is combined to calculate the rerouted flight fuel savings associated with the TOS reroute.

The NASA developed surface fuel flow model was used previously by ATD-2 to estimate fuel and emissions benefits for departure surface metering and overhead stream insertion. The model begins by using the tail number of a flight to identify the specific engines on the aircraft. NASA collaborated with flight operators to identify what percentage of flights use single engine vs. double engine taxi during the taxi out phase and encoded this information in decision trees. Knowing the engine type and the single vs. double engine taxi details, we calculate a fuel flow rate on the surface. This fuel flow rate is then multiplied by the $\mathbb{ODS}_T$ to calculate surface fuel savings.

For the airborne fuel cost NASA collaborated with flight operators which provided detailed airborne fuel burn tables. The airborne fuel burn tables allow for calculations that incorporate a variety of inputs including the flight range, flight time, total payload, load factor, and fuel weight. Each month the flight operators provide estimates of the load factor which we use to estimate the total payload for a specific flight. Given the flight range and the total payload, we use the flight operator provided lookup tables to obtain the airborne fuel burn rate. Multiplying the airborne fuel burn rate by the AFT_T provides the airborne fuel cost.

For each flight that is rerouted, we calculate the total fuel savings for the individual rerouted flight and then add the fuel savings at the system level. Since flights that are not rerouted will fly the same route, the fuel savings at the system level is calculated using the summation in Equation (3) which represents the OFF Delay Savings summed over all flights in the system. At the system level, we assume each flight is a Boeing 737-800 (the most frequent aircraft type in NTX) and apply the NASA surface fuel flow model to obtain the system level surface fuel savings. The total fuel savings accounts for the fuel savings of both the individual rerouted flight and the system level surface fuel savings.

The total fuel savings is converted to CO_2 emission savings using the conversion that 3.08 pounds of CO_2 is generated for each pound of jet fuel burned. The total CO_2 emissions savings is converted into the equivalent number of urban trees using the conversion 134.48 pounds of CO_2 is equivalent to 1 urban tree grown for ten years.

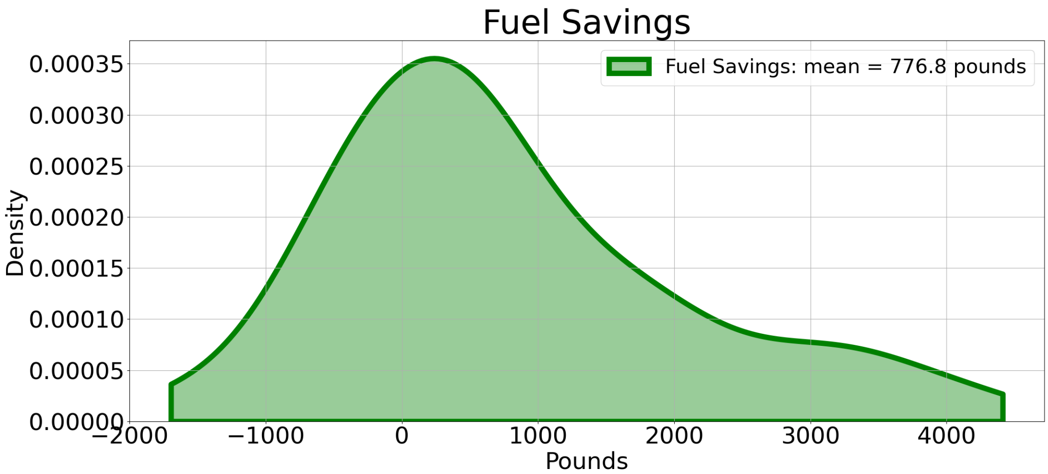

Figure 9 – Estimated environmental benefits of rerouted flights. The horizontal axes represents the pounds of fuel savings as a result of the reroute decision.

Figure 9 shows the distribution of fuel savings for the rerouted flight. The average flight saved 776.8 pounds of fuel and reduced CO_2 emissions by 2,392 pounds, which is equivalent to planting 17.8 urban trees. In total, all flights rerouted during the field evaluation saved 83,893 pounds of fuel, reduced CO_2 emissions by 258,392 pounds, which is equivalent to planting 1921 urban trees.

5. Discussion

This paper introduced the Sustainable Flight National Partnership Ops Demo Series as part of the SFNP Mission delivering to the United States Aviation Climate Action Plan and net-zero carbon emissions by 2050. SFNP Ops is a series of four operational demos NASA is conducting between 2022 and 2030 with a focus on delivering real world sustainability benefits.

NASA has a long history of building and deploying operational systems as part of field evaluations, including core Decision Support Tools such as Time Based Flow Management and Terminal Flight Data Manager. To evaluate new concepts in the field required deployment of physical hardware and networks which set a high bar for demonstration. To address this, NASA developed the Digital Information Platform which enables virtual deployment of Decision Support Tools for evaluation in the National Airspace System.

A core capability enabling the SFNP Ops Demo Series is the Machine Learning Airport Surface Model. The Machine Learning Airport Surface Model is the result of digital transformation from

NASA's IADS system developed as a monolithic Decision Support Tool to a service-oriented architecture deployed in a cloud environment. The Machine Learning Airport Surface Model also replaces legacy physics based prediction algorithms that rely on hand crafted adaptation with machine learning algorithms that learn directly from the data. This architecture aligns with the FAA's vision for an Info-Centric NAS and addresses scalability challenges introduced by legacy adaptation.

The first demo in the SFNP Ops series, Collaborative Digital Departure Reroute, focuses on pre-departure Trajectory Option Set rerouting. The Machine Learning Airport Surface Model was deployed in January 2022 as part of CDDR and predictions were validated in a side-by-side comparison to the legacy IADS system. Performance of the ML Airport Surface model compared to legacy IADS system was very encouraging.

For arrival predictions, the ML outperformed the legacy approach with the arrival runway prediction increasing from 56.4% to 63.0% accuracy and the arrival on time prediction improving where the legacy approach had 66.1% flights and ML approach had 78.9% flights with prediction within $+/-$ 2 minutes of the actual value. For departure predictions, ML departure runway predictions were within one percent of ATC assignments of departure runways and the ETOT predictions at both KDFW and KDAL showed improvement with the ML approach when measuring the standard deviation of prediction error and the percent of flights with prediction within $+/-$ 5 minutes of the actual value.

The predictions from the Machine Learning Airport Surface Model were used to recommend pre-departure reroutes that resulted in benefits in both time and environmental savings. Between January 2022 and May 2024 the system rerouted 108 flights resulting in average actual OFF Delay savings of 8.7 minutes, actual IN Delay Savings of 9.1 minutes, and system-level Aggregate Delay Savings of 29.8 minutes. This delay savings were calculated to save, on average, 776.8 pounds of fuel and reduced CO_2 emissions by 2,392 pounds, which is equivalent to planting 17.8 urban trees. In total, all flights rerouted during the field evaluation saved 83,893 pounds of fuel, reduced CO_2 emissions by 258,392 pounds, which is equivalent to planting urban 1921 urban trees.

Future work for SFNP Ops will include a benefits extrapolation of the CDDR results across the NAS to all major TRACONS. Future SFNP Ops demos are planned through 2030, and each demo will iteratively build upon capabilities developed in previous demos. The second demo builds upon the single flight pre-departure capability to introduce airborne trajectory management and reroutes. The third demo extends the single flight use case to develop fleet-wide DSTs to help flight operators and ATC manage large scale disruptions across the network. The fourth demo is a capstone demo that pulls together capabilities from the previous demos while introducing increasing levels of automation and end-to-end digital workflows.

Contact Author Email Address

Corresponding author: william.j.coupe@nasa.gov

References

[1] United states 2021 aviation climate action plan. `https://www.faa.gov/sites/faa.gov/files/2021-11/Aviation_Climate_Action_Plan.pdf`. Accessed: 2024-05-15.

[2] Sustainable flight national partnership. `https://www.nasa.gov/directorates/armd/sfnp/`. Accessed: 2024-05-15.

[3] New look at nasa, boeing sustainable experimental airliner. `https://www.nasa.gov/image-article/new-look-at-nasa-boeing-sustainable-experimental-airliner/`. Accessed: 2024-05-15.

[4] Ff-ice: Flight and flow information for a collaborative environment. `https://www.faa.gov/sites/faa.gov/files/FactSheet-FF-ICE.pdf`. Accessed: 2024-05-15.

[5] Heinz Erzberger. Ctas: Computer intelligence for air traffic control in the terminal area. Technical report, 1992.

[6] William Nedell, Heinz Erzberger, and Frank Neuman. The traffic management advisor. In *1990 American Control Conference*, pages 514–520. IEEE, 1990.

[7] Harry N Swenson, Danny Vincent, and Leonard Tobias. Design and operational evaluation of the traffic management advisor at the ft. worth air route traffic control center. In *United States/Europe Air Traffic Management Research and Development Seminar*, 1997.

[8] Time based flow management. https://www.faa.gov/air_traffic/publications/atpubs/foa_html/chap18_section_25.html. Accessed: 2024-05-15.

[9] Rich Coppenbarger, Yoon Jung, Tom Kozon, Amir Farrahi, Waqar Malik, Hanbong Lee, Eric Chevalley, and Matt Kistler. Benefit opportunities for integrated surface and airspace departure scheduling: a study of operations at charlotte-douglas international airport. In *Digital Avionics Systems Conference (DASC)*, 2016.

[10] Yoon Jung, Waqar Malik, Leonard Tobias, Gautam Gupta, Ty Hoang, and Miwa Hayashi. Performance evaluation of sarda: an individual aircraft-based advisory concept for surface management. *Air Traffic Control Quarterly*, 22(3):195–221, 2014.

[11] Miwa Hayashi, Ty Hoang, Yoon C Jung, Waqar Malik, Hanbong Lee, and Victoria L Dulchinos. Evaluation of pushback decision-support tool concept for charlotte douglas international airport ramp operations. In *11th USA/Europe Air Traffic Management Research and Development Seminar*.

[12] Susan Lockwood, S Atkins, and Nancy Dorighi. Surface management systems simulations in nasa's future flight central. In *AIAA Modeling and Simulation Technologies Conference and Exhibit*, page 4680, 2002.

[13] Shawn A Engelland, Richard Capps, Kevin Brian Day, Matthew Stephen Kistler, Frank Gaither, and Greg Juro. Precision departure release capability (pdrc) final report. 2013.

[14] Yoon Jung, William Coupe, Al Capps, Shawn Engelland, and Shivanjli Sharma. Field evaluation of the baseline integrated arrival, departure, surface capabilities at charlotte douglas interntional airport. In *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)*, 2019.

[15] William J Coupe, Divya Bhadoria, Yoon Jung, Eric Chevalley, and Greg Juro. Atd-2 field evaluation of pre-departure trajectory option set reroutes in the north texas metroplex. In *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, pages 1–10. IEEE, 2022.

[16] 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.

[17] William J Coupe, Yoon Jung, Hanbong Lee, Liang Chen, and Isaac J Robeson. Scheduling improvements following the phase 1 field evaluation of the atd-2 integrated arrival, departure, and surface concept. In *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)*, 2019.

[18] Terminal flight data manager. https://www.faa.gov/air_traffic/technology/tfdm. Accessed: 2024-05-15.

[19] Charting aviation's future: Operations in an info-centric national airspace system. https://www.faa.gov/sites/faa.gov/files/Charting-Aviations-Future-Operations-in-ICN_0.pdf. Accessed: 2023-01-22.

[20] Mohana M Gurram, Pallavi Hegde, and Swati Saxena. Nasa's digital information platform (dip) to accelerate nas transformation. In *AIAA AVIATION 2023 Forum*, page 3400, 2023.

[21] William J Coupe, Yoon Jung, Liang Chen, and Isaac Robeson. Atd-2 phase 3 scheduling in a metroplex environment incorporating trajectory option sets. In *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, pages 1–10. IEEE, 2020.

[22] William J Coupe, Alexandre Amblard, Sarah Youlton, and Mathew Kistler. Machine learning airport surface model. In *2023 IEEE/AIAA 42nd Digital Avionics Systems Conference (DASC)*, pages 1–10. IEEE, 2023.

[23] Fuser high level software description. https://aviationsystems.arc.nasa.gov/atd2-industry-workshop/fuser/High-Level-Software-Description_84377871.html. Accessed: 2024-05-15.

[24] Banavar Sridhar. Applications of machine learning techniques to aviation operations: Promises and challenges. In *2020 International Conference on Artificial Intelligence and Data Analytics for Air Transportation (AIDA-AT)*, pages 1–12. IEEE, 2020.

[25] Mlops: Continuous delivery and automation pipelines in machine learning. https://cloud.google.com/architecture/mlops-continuous-delivery-and-automation-pipelines-in-machine-learnin Accessed: 2023-01-22.

[26] David Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, and Michael Young. Machine learning: The high interest credit card of technical debt. 2014.

[27] Yue Zhou, Yue Yu, and Bo Ding. Towards mlops: A case study of ml pipeline platform. In *2020 International conference on artificial intelligence and computer engineering (ICAICE)*, pages 494–500. IEEE, 2020.

[28] Meenu Mary John, Helena Holmström Olsson, and Jan Bosch. Towards mlops: A framework and maturity

model. In *2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*, pages 1–8. IEEE, 2021.

[29] Eric Chevalley, Gregory L Juro, Deborah Bakowski, Isaac Robeson, Liang X Chen, William J Coupe, Yoon C Jung, and Richard A Capps. Nasa atd-2 trajectory option set prototype capability for rerouting departures in metroplex airspace. In *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, pages 1–10. IEEE, 2020.

[30] Divya Bhadoria. Atd-2 phase 3 benefits mechanism. In *ATD-2 Tech Transfer*, 2021.