Fifteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2023)

Risk-Adjusted Traffic Management Strategies for Convective Weather Conditions*

James C, Jones, Zachary Ellenbogen
MIT Lincoln Laboratory
244 Wood St, Lexington, MA 02421, United States

Abstract— **Weather is a significant source of disruption and uncertainty for air traffic. This unpredictability can present significant challenges to traffic managers when managing airport and airspace resources. The lack of data-driven decision support tools to advise stakeholders on how to best deal with weather impacts also represents a critical shortfall in enabling improved decision-making within air traffic management. In this paper, we present an epsilon-greedy approach that incorporates risk-adjusted objectives into recommendations of Traffic Management Initiative (TMI) parameters during uncertain weather conditions. The method attempts to achieve the best performance within the context of some of the worst-case weather outcomes. The method is compared to a standard epsilon greedy approach that attempts to maximize the expected value of an objective. The two approaches are evaluated using a parallel fast-time simulation framework over various weather scenarios. A set of TMIs at airports and airspace resources is applied and tested against seven case days in which the airspace capacity in the Northeast United States was affected by convective weather. The risk-adjusted method is generally able to achieve a higher number of operations with lower amounts of airborne holding in adverse weather conditions. The results suggest that the approach could potentially aid more risk-averse air traffic stakeholders by supporting their operational planning.**

Keywords-Reinforcement learning, risk, epsilon greedy, weather, Simulation, airspace capacity

I. INTRODUCTION

Weather poses a considerable source of both disruption for Air Navigation Service Providers (ANSPs) and airlines when planning air traffic management operations. As the weather can limit the capacity of the airport and airspace resources, its presence may induce mismatches between the flight demand for specific resources and the available capacity. In these situations, ANSPs will preemptively assign delays to flights. These delays are known as Traffic Management Initiatives (TMIs) such as Ground Delay Programs (GDPs), Airspace Flow Programs (AFPs), Ground Stops (GSs) and Collaborative Trajectory Options Programs (CTOPs) pushing the excess traffic demand back to later times during the day with the goal of balancing the air traffic operations with the available resource capacity.

This problem of weather-related disruptions is compounded by the lack of predictability associated with the information available. Weather forecasts can often be quite uncertain, especially over long time-horizons. This uncertainty manifests in a number of different ways. The uncertainty associated with the weather forecasts issued may imply significant variation in the possible temporal and spatial locations of the affecting weather patterns as well as the intensity of the weather patterns when they arrive at airspace and airport resources. Even when the uncertainty of these forecasts is low, the forecasts themselves do not directly describe the impact that the weather will have on the resources that need to be managed. This information must be obtained by other models that translate the weather into an estimate of the capacity of the affected resource. As this process is imprecise, the estimation of weather impacts adds an additional dimension of uncertainty to the problem. When facing this uncertainty, stakeholders must make decisions on how to adjust flight schedules and air traffic flow without a clear indication of how these decisions will impact their operational and business objectives.

The scope of the weather impact can be substantial at times affecting many different resources. In these situations, air traffic managers often apply TMIs to attempt to manage demand in a coordinated fashion across multiple resources. These TMIs can often have considerable complexity. Decision-makers need to determine what resources to control, when to control them, the duration of the proposed intervention and the levels of adjustment at each hour over the period of intervention. As the range of options available for each of these decisions is often substantial, traffic

* DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.
This material is based upon work supported by the Department of the Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of the Air Force.

managers are faced with a problem of considerable complexity.

In the United States, traffic managers have a limited number of decision support tools available to them to aid them in developing a strategy. The Traffic Flow Management System (TFMS) provides the users with an estimate of the demand at each airport and the capability to visualize forecast traffic, while the Network Manager serves a similar function in Europe [1]. Weather forecasts such the Corridor Integrated Weather System [2] are also available to visualize the trajectory of the anticipated convective weather. These tools can be combined with experiential knowledge to provide some understanding of the effect of applying potential TMIs at a very coarse level but do not offer any direct prediction of TMI performance.

In the presence of weather impacts, traffic managers generally act on this imperfect information by making an initial decision on the TMI parameters and revising the TMI as the day progresses. This challenge can be viewed as a stochastic sequential staged decision-making problem, where the decision-maker needs to adjust the flight demand given the potential set of capacity scenario evolutions in a manner that optimizes a stated objective function. This topic has been studied widely using both descriptive and prescriptive approaches to address the problem. The descriptive studies attempt to predict the capacity of the airport and airspace resources when weather impacts are present [3]–[7]. While this information can enhance the level of situational awareness for traffic managers, it does not directly advise them on what to do with the additional information provided. The prescriptive studies seek to determine the optimal action to take in the context of potential capacity disruptions and uncertainty [8]–[15]. They often employ integer programming models to search for a solution that best advances the stated objective. While such approaches could be useful if they accurately depict the operational states of the airspace systems they try to model, the approaches often assume some theoretical distribution of the capacity that does not map to the actual impact of weather.

Other approaches have tried to apply data-driven descriptions of the weather forecast into models that optimize for performance. Integer programming models that incorporate weather-based capacity constraints have been used to optimize decision-making for ground delay programs [16]–[18]. Recent work has also adapted reinforcement learning algorithms to optimize the performance of air traffic across a set of airport and airspace resources across a geographic region. Our own work leveraged the Traffic Flow Impact (TFI) product that combines several forecast models to predict the impact on the capacity of airspace resources over time (see discussion in the next section). The information was used to construct a set of weather scenarios that informed various search methods with the goal of selecting GDP and AFP parameters that optimize operational throughput across a set of airports while limiting airborne holding[19]–[21]. In another line of research, the Short-Range Ensemble Forecast (SREF) has been used to probabilistically model weather scenarios that were used to optimize the parameters on a ground delay program with metering using a cost function that weighted the impact of air, ground and air traffic induced delay[22], [23].

In both of these cases the models used considered the impact of various weather scenarios, while attempting to optimize for the expected cost in terms of the metrics considered. Neither of these approaches, however, considered the problem by optimizing the recommended decision to the level of risk-tolerance of the decision-maker. As the weather represents an uncertain prediction into the future, with varying levels of impact, the adoption of any specific recommendation from a model may not manifest. When this happens, traffic managers may be left to question the utility of the model even though the model provided the best recommendation given the information available at the time it was issued. In this context, one might argue that providing decision-makers with the relative risk associated with adopting a particular decision may help to quell an adverse reaction if a worst-case outcome were to occur.

In this paper, we propose a methodology for recommending TMIs that allows the user to optimize the TMI performance over a set of worst-case outcomes. In our case, we define the worst-case as the TMI that yields the lowest throughput in the affecting weather scenario. We then compare the proposed method with one that issues recommendations based on an expected value and assess the relative utility of using such a risk-averse objective relative to expected value. Both methods use variants of an epsilon greedy (ε-greedy) algorithm and a parallel fast-time simulation to search for TMI parameters and seek to optimize performance over a set of weather forecast scenarios derived from the TFI weather impact model. In Section II, we describe our modeling framework and our methodology for incorporating risk levels into TMI recommendation. Section III, describes the computational experiments used to test the approach. We evaluate our method against a set of case days in the Northeastern United States when traffic was compromised by convective weather and discuss utility of our method to operational decision-making to support various levels of risk-aversion.

II. METHODOLOGY

In this section we describe our methodology for incorporating aspects of risk into the recommendation of Traffic Management Initiatives. We begin by describing our approach for translating the weather forecast into a prediction of the weather impact of airspace capacity. We discuss two ways to use this forecast information to develop TMIs that directly address the uncertainty in the forecast. We also present our modeling framework that we use to evaluate the two approaches.

A. Traffic Flow Impact

The inability to reliably provide users with greater situational awareness during convective weather events has historically served as a significant barrier to the advancement for strategic decision-making in air traffic management. Earlier work to predict convective weather impact focused on the propensity of flights to avoid specific storm cells [24], [25]. This information has been used to reroute traffic around convective weather [26], [27] and assess sector capacity [28], [29]. While such information can be useful in making tactical decisions, at the strategic level it is somewhat unreliable due to the lack of predictability in the movement of convective weather over long time horizons. The Traffic Flow Impact (TFI) tool was developed to address this shortfall by translating the weather to an impact level over larger regions of airspace. TFI produces predictions of the percentage of weather-free airspace known as permeability as forecast quantiles that are issued every hour. The permeability forecast is applicable to en route flights at an altitude of 35,000 ft over 25 geographic regions across the eastern United States. These regions, known as flow constrained areas (FCAs) are located at chokepoints where the weather can greatly affect the throughput of traffic within the National Airspace System (NAS).

The permeability represents the degree that traffic flows in a given airspace region will be constrained by convective weather based on the vertically integrated liquid values and echo top levels of the FCA regions at 35,000 ft. The TFI forecast predicts this permeability value using a supervised machine learning approach [30]. The feature set for the models consists of inputs from deterministic weather models such as the Consolidated Storm Prediction for Aviation (CoSPA) and the High-Resolution Rapid Refresh (HRRR) model. Other features are received from the Localized Aviation Model Output Statistics Program (LAMP) and the Short-range Ensemble Forecast model (SREF). The model is trained with weather data from prior convective seasons using two steps. In the first step the actual permeability is used to fit the features using ridge regression. In the second step the features are fit using quantile regression.

The TFI forecast represents a distribution of permeability over time that can be sampled to generate the potential evolution of the weather pattern and its impact on the FCA airspace region. The samples are conditioned to have a consistent time-correlation to match the behavior of the actual permeability. Thus, the value of the sampled permeability is dependent on the state in the previous forecast lead. By sampling such draws at the available FCA regions, we can characterize the impact on the permeability over the eastern U.S. airspace.

The permeability provides an estimate of a flight's ability to traverse the region but does not directly inform us of how the weather is expected to impact the airspace capacity. We can convert the estimates of permeability to a flow rate describing the number of flights that can cross the airspace over an hour using a set of look-up tables that were derived through studies on the pilot avoidance of convective weather [31]. The avoidance was observed to depend on the permeability of the airspace and the length of time that the permeability level was present in the affected FCA region. The presence of a permeability level for longer periods of time is associated with lower flow rates for the same level of permeability. The flow rate is also driven by other airspace characteristics (e.g. size, traffic density) that differ between regions.

B. Searching for TMI Parameters

The translation of weather forecasts to the impact they will have on specific regions of airspace provides us with an estimate of the airspace capacity. As this forecast and its impact is uncertain, there are a number of potential evolutions of weather that could manifest within the airspace that will be suited to different types of traffic management strategies. Severe weather days may require strategies that impose heavy ground delays on flights at multiple airport and airspace resources in order to better match demand with capacity preemptively. On the other hand, light weather impacts may require little intervention and can often be managed tactically.

We would like to use the information available in the TFI forecast of airspace capacity to design TMIs that account for the range of potential outcomes that may affect a set of managed resources. One way to accomplish that is to sample from the TFI distribution and search for the TMI that will maximize an expected value of some performance metric(s) when averaged over all the capacity scenarios that are selected. We can view these samples in the context of a scenario tree where each node represents a set of feasible capacities during some hour at a set of airspace resources across the NAS. Letting W_{ij} be the forecast from scenarios *i* in hour *j* we can envision a scenario constructed by sampling from a set of correlated draws at FCA resources as a single branch along the tree. An illustration of the concept is shown in Figure 1. We can incorporate these samples into our model that characterizes that state of the NAS over time. By running instantiations of the model, we can gain a sense of how the NAS will perform given a weather forecast and the assumed demand.

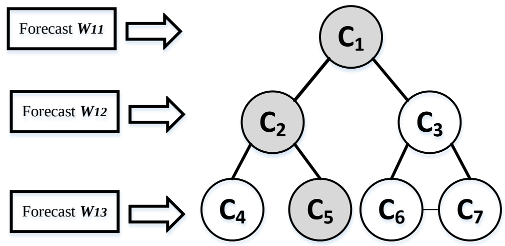

Figure 1. A set of sampled capacities on a scenario tree where C_i represents the capacity for the resource.

If we had the ability to sample from all of the scenario tree branches, we could compute the value of each state so that we can maximize the expected value. Due to the computational complexity of the problem of managing a large set of resources in the presence of uncertianty, is it only possible to consider a small number of potential scenarios. While we can determine the optimal action given what we know about the state of the world, we do so at the cost of not learning about the value of the unsampled states and actions. The tradoff between sampling the unknown states and sampling states that will likely perform well given current knowledge, is known as the exploration-exploitation problem.

ε-greedy policies:

There are a number of algorithms that have been proposed to deal with the exploration-exploitation problem. One of the more widely used approaches is to adopt an ε-greedy policy. The approach attempts to strike a balance between exploration and exploitation by favoring each aspect at different points in the process. The algorithm begins by sampling from a random distribution. When the value of the sample is greater than some predefined level ε, the policy selects a random action, when the sample is less than ε, the algorithm chooses the optimal decision based on an estimate of the value of each choice that was fit with information collected prior to the decision. Initially the threshold value is usually set to a high level to force the algorithm to favor random selection. When the actions are largely random, the algorithm can learn more about the value of being in various states. As the algorithm continues, the value of the threshold is lower, making the chance of selecting exploitative action higher. This change is often advantageous as the value of new information tends to drop as the more iterations are completed. An adaptation of the algorithm for TMI decision-making problem is shown in Table I with the ε threshold occurring in step 5. When new information is collected, we need to fit the value of the information using some model. In our implementation, we train non-parametric supervised learning method known as Gradient Boosting Regression by fitting the TMI parameters to operational metrics of interest (e.g. number of operations, number of flights experiencing excessive airborne holding, hours of ground delay).

The variable definitions used in Table I are provided as follows:

Ω≡The set of all scenarios

$S_t^{n,m}$≡The system state at time *t* in scenario *m* during trial *n*

W_t≡The airspace capacity described in forecast in period *t*

J ≡The set of all simulation instances

ω^m≡ The scenario forecast *m*

$X_j^{\pi,n}$≡ The policy function X^π that maps air traffic states to TMIs for instance *j* of the air traffic simulation

$x^{n,m}$≡ The TMI decision made in scenario *m* for instance *j* on trial *n*

$C(S^{n,m}, x_j^{n,m})$ ≡ The cost of the decision made

$\hat{v}_j^{n,m}$ ≡ The value of the decision in scenario *m* for instance *j* on trial *n*

TABLE I. AN ε-GREEDY APPROACH FOR ASSIGNING TMIS

Step 0. Initialization
  Step 0a. Initialize $\bar{V}_j^0, j \in J$.
  Step 0b. Initialize $S_j^1, j \in J$.
  Step 0c. Choose an initial policy $X^{\pi,0}$.
  Step 0d. Set $n = 1$.
Step 1. Repeat for $m = 1, 2, \dots, M$.
  Step 1a. Choose a sample path ω^m.
Step 2. Do for j =0, $1, \dots, J$.
  Step 2a. Find $x_j^{n,m} = X^{\pi,n-1}(S_j^{n,m})$.
  Step 2b. Update the state variable by simulating the air traffic
  $S_j^{n,m} = S^M(S_j^{n,m}, x_j^{n,m}, W(\omega^m))$.
  Step 2c. Set $\hat{v}_j^{n,m} = 0$
  $\hat{v}_j^{n,m} = C(S_j^{n,m}, x_j^{n,m})$
Step 3. Compute the average value from starting in state S_j^1:

$$\bar{v}_j^n = \frac{1}{M}\sum_{m=1}^{M} \hat{v}_j^{n,m}.$$

Step 4. Update the value function approximation by using the average values by fitting an estimate.
  $\bar{V} \leftarrow U^V(\bar{V}_j^{n-1}, S_j^{x,n}, \bar{v}_j^n)$.
Step 5. With probability ϵ, choose J decisions x^n at random from X. With probability $1 - \epsilon$, choose J decisions x^n using the following procedure. Let j =0
  Step 5a. For j find

$$X_j^{\pi,n}(S) = arg_{x\in X}\max(C(S_j^n, x) + \bar{V}_j^n(S^{M,a}(S_j^n, x)))$$

  Step 5b. Remove x such at $x \notin X$ and let $x \in D$, where *D* is the set of TMI decisions to be taken.
  Step 5c. Update the value of ϵ=α/(α+n+1), where α is the learning rate.
  Step 5d. Increment j. If $j \leq J$ go to step 6, if not go to step 5a.
Step 6. Increment n. If $n \leq N$ go to step 1.
Step 7. Return the value functions $\bar{V}_j^n$.

Sample-ranked constrained ε-greedy policies:

Another approach is to assume the capacity of the airspace resources are assigned to quantiles of the distribution. This approach is commonly adopted in chance-constrained programming. In this framework, FCA regions would be assigned a certain quantile based on the risk tolerance of the decision-maker. Once the capacities are defined, they will stay fixed and the TMI parameters can be adjusted to optimize the desired objective. The solution would then describe the best possible performance assuming the airspace capacity did not exceed the assumed constraints. One drawback to this approach is that it is difficult to construct scenarios that limit the capacity bounds of 25 airspace resources without over-constraining certain resources to levels that represent a realistic outcome.

A potentially less stringent alternative to the chance-constrained approach is to draw from the distribution and grade the performance of each scenario based on some performance metric. The goal is then to achieve the best possible performance when averaging over the subset of worst-case scenarios that were selected. In this approach, we evaluate the performance of the TMIs under each scenario and rank the scenarios from best to worst based on a given performance metric (e.g. number of operations, delay cost). Once the scenarios are ranked, we establish a threshold β that equates to an assigned percentile. We then eliminate the scenarios that demonstrated better performance than the one in the chosen quantile and average over the remaining scenarios. We will refer to this method as the Ranked-sample constrained ε-greedy approach. A description of the method is shown in Table II.

TABLE II. A RANKED-SAMPLE CONSTRAINTED ε-GREEDY APPROACH FOR ASSIGNING TMIS

Step 0. Execute Steps 0, 1 and 2 in Table I.
Step 1. Order samples of $\hat{v}_j^{n,m}$ for all sample paths m.
  Step 1a. Let $V(\beta)$ be the value of the sample realization in the β-percentile.
  Step 1b. If $\hat{v}_j^{n,m} \leq V(\beta)$ then $m \in Q$ where $Q$is the set of scenarios achieving performance below quantile β.
Step 2. Compute the average value from starting in state S_j^1:

$$\bar{v}_j^n = \frac{1}{[Q]} \sum_{m \in Q} v_j^{n,m}.$$

Step 3. Update the value function approximation by using the average values by fitting an estimate.

$$\bar{V} \leftarrow U^V(\bar{V}_j^{n-1}, S_j^{x,n}, \bar{v}_j^n).$$

Step 4. With probability ϵ, choose J decisions x^n at random from X. With probability $1 - \epsilon$, choose J decisions x^n using the following procedure. Let j =0
  Step 4a. For j find

$$X_j^{\pi,n}(S) = arg_{x \in X} \max(C(S_j^n, x) + \bar{V}_j^n(S^{M,a}(S_j^n, x)))$$

  Step 4b. Remove x such at $x \notin X$ and let $x \in D$, where D is the set of TMI decisions to be taken.
  Step 4c. Update the value of $\epsilon = \alpha/(\alpha+n+1)$, where α is the learning rate.
  Step 4d. Increment j. If $j \leq J$ go to step 5, if not go to step 4a.
Step 5. Increment n. If $n \leq N$ go to step 0.
Step 6. Return the value functions $\bar{V}_j^n$.

Sample Generation:

In both variants of the ε-greedy algorithm we need to generate a sample path from the TFI distribution. We also need to estimate the value of various TMI decisions. We consider two types of TMI decisions: GDPs and AFPs. GDPs assign ground delays to flights that are scheduled to fly to specified airports, while AFPs assign ground delay to flights that plan to fly through various airspace regions. Each of these programs can be defined by a start time and end time and a rate. In order to estimate a TMI value, we need to generate a set of candidate TMIs. We begin with a baseline vector describing the characteristics of the TMI program. The baseline vector element includes the hourly controlled rates at the set of resources under consideration, in this case a set of airports and/or flow constrained areas in an AFP. Another vector element is the size of the exemption radii for GDPs. When an exemption radius is applied to a GDP, all flights at origins that are more than the specified distance away from the GDP airport exempted from the program, while flights at origin airports inside of the radius receive controlled ground delays. The TMI baseline vector sets the size of the exemption radius to one of four distances for each GDP. Another element included in the baseline vector is the start time of the GDP or AFP for each managed resource. The end time of the program is defined by adding the number of elements that define the rate of the start time divided by 24 and taking the remainder. Once all the elements of a program have been included in the baseline vector, the vector elements continue to define the parameters for the other managed resources until all the TMI parameters have been listed. The TMI sample generation process is depicted in Figure 2.

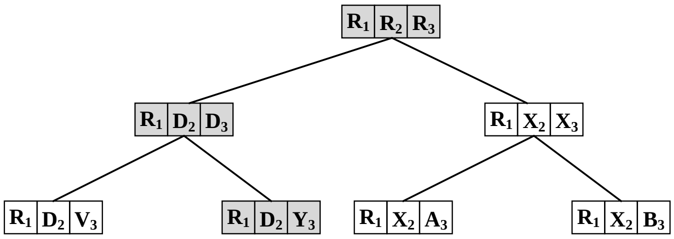

Figure 2. An illustration of a sample perturbation of the TMI baseline vector.

C. Search Strategy

A fast-time simulation framework was used to evaluate the effects of weather impacts on airspace against various traffic demand scenarios. The simulation can import either the flow rate that is based on a sampled TFI forecast draw or the actual weather. The simulation also ingests wind forecast models such as the High-Resolution Rapid Refresh or the Global Forecast System to account for the impact of wind on the four-dimensional flight trajectories. Controller workload behavior based on a set of analytical models described in [28] and [29] are used to set the sector capacities across the NAS. Flight demand is defined by historical flight plans from the Traffic Flow Management System. The Base of Aircraft Data (BADA) 3.6 model is used to derive the speed and fuel burn profiles for each aircraft. Coded instrument flight navigation procedures from the Federal Aviation Administration (FAA) provide the basis for the four-dimensional trajectories used in the simulation. The TMI responses are implemented using Collaborative Decision Making (CDM) procedures such as ration-by-schedule, cancellation and substitution and compression [32] Tactical TMIs such as Approval Requests (APREQ) that are used in the US impose ground-holding to adjust the demand at nearby sectors were also implemented within the simulation.

When the simulation initializes, a TMI can be applied to the scenario by adjusting the scheduled arrival times prior to take-off. A set of concurrent instantiations of the air traffic simulation are run in parallel on a computing cluster to provide broader scenario exploration. A set of TFI-correlated FCA draws based on the forecast weather is sampled by each simulation. After our initialization, the algorithmic approach under evaluation generates a new set of TMI parameters to control a set of airspace and airport resources in the scenario. The performance of the evaluated TMI is scored using a set of operationally-relevant metrics (e.g., number of operations, holding events, delay) and logged in a database. As the process iterates, the framework continues the series of air traffic management simulations by using our algorithmic selection method to choose different sets of TMI parameters until we have reached our last simulation run. A diagram of the architecture is shown in Figure 3.

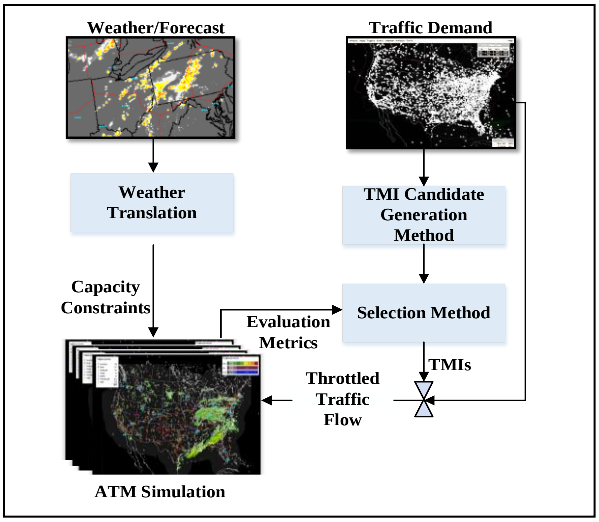

Figure 3. A simulation framework to facilitate exploration of Traffic Management Initiative parameters.

III. RESULTS AND DISCUSSION

In this section we describe our computational experiments used to evaluate the approaches for selecting TMI parameters. We discuss the relative performance of two methods against forecast weather scenarios over a set of case days. We also report the performance of the Ranked-Sample Constrained e-greedy method relative to the TMIs that were actually implemented using the actual weather.

A. Experimental Description

A set of computational experiments were performed using the methodology described in Section II. Seven case days in the summer of 2019 were chosen as test scenarios. The weather impacts on those days were either moderate or severe and all days had sufficient convective blockage to motivate the traffic managers on duty to initiate a set of TMIs at airports and airspace resources within the Northeastern United States. These weather impacts provide the type of airspace capacity disruptions that our proposed methodology was designed to mitigate. The TFI capacity data was injected into a set of simulations configured to evaluate the traffic on the days examined over a period lasting between 04:00GMT to 3:59GMT the next day. Historical TFMS flight plans for an entire day of traffic in the U.S. were ingested and simulated along the route in the flight plans to deterministically emulate the demand levels for the day. The TMIs identified with our search methods were imposed at New York John F. Kennedy (JFK), LaGuardia (LGA), Newark (EWR), Philadelphia (PHL), Boston Logan (BOS) airports and FCA regions called A01 and A08. A diagram showing the geographic locations of the airports and FCAs used is shown in Figure 4. In many instances, the actual case days often included other GDPs. We wanted to include them for comparison purposes when evaluating the tested TMIs relative to the implemented TMIs, so we also included the GDPs that were put in place at these resources on the day being evaluated. Since we did not intend to test the GDPs at these airports, the parameters on the GDP, resources remained constant and consistent with the implemented rates in all trials.

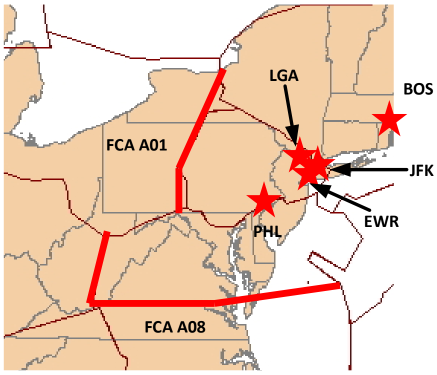

Figure 4. The location of the airports and FCA regions that were managed by the TMI programs under test.

An instantiation of our simulation framework was run on a high-performance computing cluster [33]. A set of TFI FCA forecast distributions were sampled 30 times to emulate 30 potential spatial-temporal evolutions of the weather. A group of 15 TMI parameter configurations was also selected to evaluate the performance of 15 different TMIs with TFI forecast samples. Each TMI was tested with 30 TFI forecast samples. The combination of TFI forecast samples and TMI configurations composed 450 concurrent simulation instances, each of which was run on an Intel Xeon-phi 7210 processor. The outcomes of each instance describe the operational performance of a single sampled TFI forecast when one of the 15 TMIs is implemented.

A nominal baseline TMI was constructed and perturbed with a 1st-order Gaussian auto-regressive random process whose parameters can be described by a mean, standard deviation and autocorrelation. The baseline settings were chosen to generate samples that provided moderate restriction on the resources, while future iterations were modified to evaluate more benign or restrictive control of flight demand. The baseline TMI had a mean value of 80% of the visual flight rules airport capacity for all 5 airport resources under test, a rate of 115 flights/hour for FCAA01 and 120 flights/hour in the case of FCAA08. The distribution that perturbed the baseline had a standard deviation of 3 flights for all GDPs and 5 flights for all AFPs since the FCA resources generally operate at higher flow rates than the airports under test. The exemption radii of the GDPs were discretized to distances of 1000, 1500, 2000 and 2500 NM. The baseline TMIs had a duration of 12 hours. During each strategy evaluated, the simulation iterated over 40 batch runs. Since the weather impacts for the days evaluated were later in the day, the initial start time of the TMIs were set to 16:00 GMT. The TMI baseline could also be adjusted by cancelling the program at any given resource prior to the scheduled end time.

We sought to identify TMI decisions that best yield high operational throughput. This goal can be obtained by over-delivering the traffic to the airports and airspace resources. However, when the amount of holding is excessive it can affect controller workload and may also lead to significant airborne holding. Since these effects are operationally undesirable, we would also like our selected TMIs to limit demand in a manner that results in a lower amount of holding. Thus, our ideal operating point would be achieved when we can replicate the throughput when the airspace is relatively saturated relative to its capacity, while delivering a low number of holding events. A depiction of the concept is shown in Figure 5.

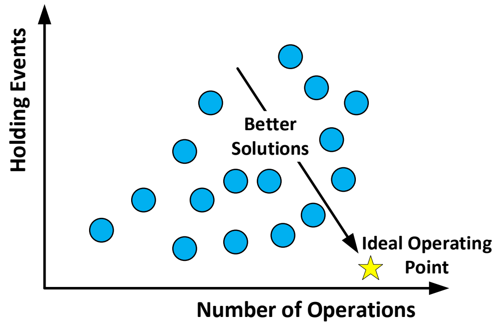

Figure 5. An illustration of an operational performance objective under evaluation.

The value of the TMI decisions for the ε-greedy method was estimated with Gradient Boosting Regression using a least-squares loss function. The model was trained at each exploitation step based on the performance data from the prior TMI implementations. At each iteration, we sampled from a uniform distribution from zero to one and compared the result to a threshold to determine whether an exploration or exploitation step was selected. The threshold level to choose an exploitation step started at one and dropped with each subsequent iteration using a learning rate of α=5. While other rates could be evaluated, due to the large dimensionality of the problem we did not investigate the sensitivity of this parameter. A gradient tree boosting model was created to predict the number of operations at the five airports evaluated. These TMIs were filtered based on the criterion that the TMIs implemented ranked in the bottom 25% of flights with more than 15 minutes of airborne holding. Each model was trained with an initial sample set of 5 random samples. These trained models were then used to predict the performance of an additional 100,000 samples.

In the case of the Ranked-Sample Constrained ε-greedy method, the performance of each TMI sample was sorted over the 30 TFI samples based on the number of operations that occurred in the simulations over the day at the five controlled airports after the simulation executed. Forecasts that yielded higher number of operations were ranked more highly and mapped to higher percentiles than forecasts that induced lower number of operations when using the same TMI. Samples with values below a specified threshold were removed from the sample pool for the purposes of calculating the objective. This threshold was set to the 90th, 80th and 70th percentile worst case outcomes with respect to operations. These thresholds filtered the pool of simulations down to the worst 3 samples (90th percentile) and the worst 6 samples (80th percentile) and worst 9 samples (70th percentile). The expected value of each of the filtered sample pools was used to fit the value function of the Gradient-Boosting Regression algorithm. We then sought to maximize the expected value of the filtered samples using the same process as the ε-greedy algorithm. A summary of the various test instances is shown in Table III.

TABLE III. SIMULATION FRAMEWORK TEST PARAMETER CONFIGURATIONS

Selection Method	*Simulated TFI, TMI, Trials*	*Exemption Radii NM*	*Objective Calculation*	*Std Rate Perturbation*
Ranked-sample constrained ε-greedy	30, 15, 40	1000, 1500, 2000, 2500	Samples with TMI performance in the worst70th, 80th and 90th percentiles	3 GDP, 5 AFP
ε-greedy	30, 15, 40	1000, 1500, 2000, 2500	All TFI forecast samples	3 GDP, 5 AFP

B. Results and Discussion

The results of the experiments described in the previous sub-section were aggregated to identify the best performing TMI. For each TMI tested in our simulation framework, we calculated our performance metrics (e.g. number of operations, number of holding events) by taking the expected value over the appropriate number of TFI samples. Once the individual statistics were computed for the individual TMI programs, the resulting performance data from each TMI was used to identify the best performing parameter settings based on our metrics of interest. We also calculated the expected value that we would have obtained if we used a TMI selected with the ε-greedy method and evaluated the performance with another TFI sample test configuration (e.g. 70th percentile). Although this calculation was not used in the ε-greedy objective function calculation it did enable us to project how we would have done with an ε-greedy approach when measured against more severe weather scenarios. We can see the relative improvement in the outcome by subtracting the performance we would have scored in a reduced set of more pessimistic outcomes from the performance obtained using all of the TFI samples. Figures 6 and 7 illustrate the relative change in performance obtained using the TMI solutions derived from the ε-greedy approach in Table I when considering each of the percentile outcomes. Negative values indicate that there are fewer operations or holds when evaluating the percentile scenarios against the mean, while positive values correspond to more operations and holds. We would like to see small negative values for the number of operations in Figure 6 and small positive values for holding in Figure 7 as such outcomes would suggest that the ε-greedy approach is relatively robust against the worst weather scenarios. Thus, the large adverse magnitude of both operations and holding indicate that the approach performs poorly against these more extreme weather impacts. In the seven days tested, our best mean outcome yields more than 350 operations while achieving significantly lower levels of holding.

The plots in Figures 6 and 7 suggest that there is a considerable gap in performance when the selected TMIs are evaluated with more pessimistic scenarios. Given this difference, more risk-averse parties might ask whether anything can be done to perform better in these situations. We attempt to address this issue by comparing the performance of the Ranked-Sample Constrained (RSC) ε-greedy algorithm relative to the ε-greedy approach. Figure 8 shows the relative number of operations obtained with the RSC ε-greedy method. Figures 9 and 10 list the number of instances when airborne holding exceeds 15 minutes (holding events) and 45 minutes (extreme holding events) respectively. The relative performance of the two algorithms is also listed in Table IV. This table lists the mean number of operations and holding events (holding greater than 15 minutes) for the best performing TMI selected with each configuration.

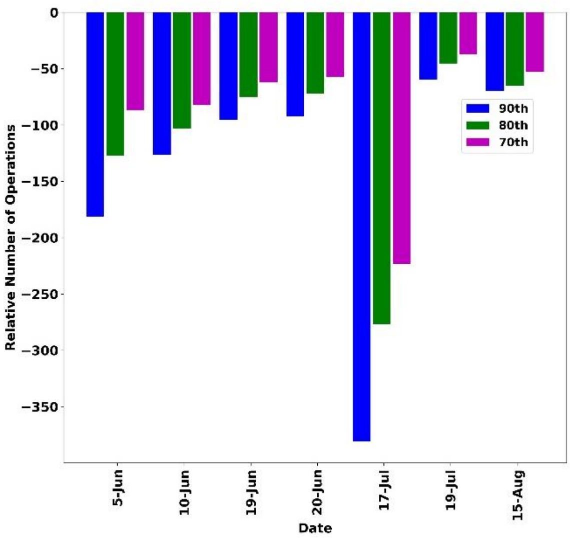

Figure 6. Number of operations obtained in each percentile scenario set relative to the 30 sample mean scenario set when using the ε-greedy method.

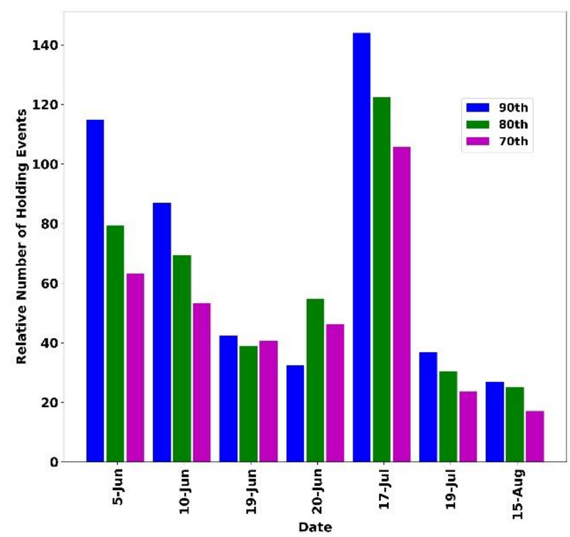

Figure 7. Number of additional holding events obtained in each percentile scenario set relative to the 30 sample mean scenario set when using the ε-greedy method.

The data in Figure 8 suggests that the RSC ε-greedy method yields a higher number of operations than the ε-greedy method in all but one instance. In the one instance where the performance does not exceed that of the e-greedy approach it is only marginally lower. This suggests the algorithm is able to successfully find better solutions from the standpoint of operational throughput in cases where the weather impacts are generally more limiting. The performance improvement is better with the 90th percentile on most days. This may relate to the fact that the ε-greedy algorithm is valuing a metric that is less similar to the 90th percntile than 70th percentile. However there are two notable exceptions on June 19th and July 17th. The increase exceeds 140 more operations on July 17th. On this day the weather impact was large, wide-spread and lasted for a long time. Under these circumstances, the convective weather blockage may have limited the potential for improvement in operational throughput. On June 19th the improvement was lower at 38 more operations. On this day the blockage in certain areas lasted a long time. In these circumstances, it is also more likely for different scenarios to yield worst-case outcomes for different TMIs making it harder to establish an optimal that performs noticably better for the higher percentiles. The number of holding events where the holding exceeded 15 minutes and the extreme holding events where the holding exceeded 45 are plotted in Figures 9 and 10. We can see from Figure 9 that in most cases the RSC ε-greedy method produces lower numbers of holding events, suggesting that the improved throughput of the method is generally not gained at the expense of additional operational holding. There is no clear trend to the relative performance of the percentile, however, this is not surprising as we are not explicitly trying to minimize holding in our value function. Figure 10 suggests that the method is even more successful at limiting extreme holding. In nearly every instance the number dropped when using the RSC ε-greedy method. This improvement may indicate that the RSC ε-greedy method could also be better at limiting diverted flights in addition to providing higher throughput in these worst case situations.

The performance suggests that the approach could be used to provide decision-makers with alternative choices to potentially mitigate more extreme weather outcomes. When applying the concept to the operational environment, the risk-tolerance of traffic managers could be assessed through interviews, human-in-the-loop studies, training and user-feedback to inform the settings of a decision-support tool. Decision-makers could then review the projections of the potential outcomes, which may help improve expressions of risk-tolerance within the context of the weather forecast and the expected traffic demand. Such discussions may motivate choices to hedge more aggressively and make decisions that are more consistent with their actual preferences and beliefs. In these instances the RSC ε-greedy algorithm might provide a more targeted alternative than one that seeks to address more benign scenarios that may be of less concern to the user.

It should be noted that while the results of this study focused on applying a set of GDPs and AFPs to address convective weather impacts in the Northeast United States, the methodology could be applied to a range of air traffic management situations and control measures. The impacts from non-convective weather (e.g., strong winds, low ceiling/visibility) can be mitigated through a combination of airport capacity uncertainty forecasts and reinforcement learning. Similarly, scenarios involving demand uncertainty could be addressed using the proposed concept. As ANSPs adopt more trajectory-based operational concepts to manage traffic, the technique could also be applied to address more targeted flight pools by using a combination of ground delay and speed control. In this framework the associated reinforcement learning-based recommendations could issue controlled-times-of-arrival rather than controlled-times-of-departure and be paired with optimization models [34]–[36] that apply speed control to realize more tactical adjustments based on airline preferences.

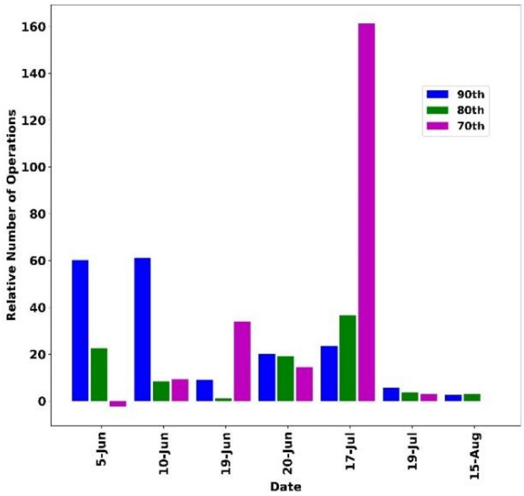

Figure 8. Number of operations obtained when using the ranked-selection constrained ε-greedy method relative to the ε-greedy method.

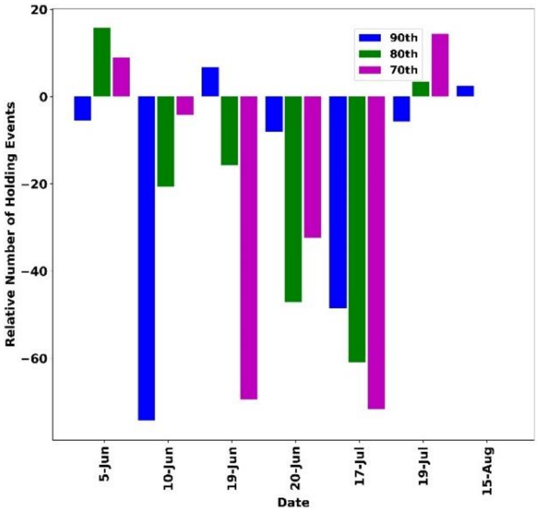

Figure 9. Number of holding events obtained when using the ranked-selection constrained ε-greedy method relative to the ε-greedy method.

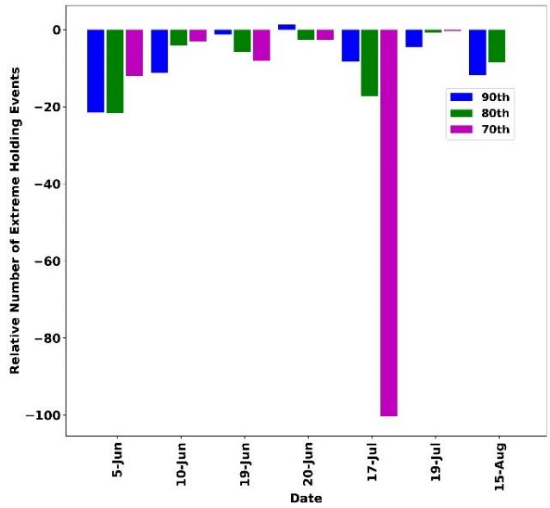

Figure 10. Number of extreme holding events obtained when using the ranked-selection constrained ε-greedy method relative to the ε-greedy method.

TABLE IV. OPERATIONAL PERFORMANCE OF THE RANKED-SELECTION CONSTRAINED ε-GREEDY METHOD RELATIVE TO THE ε-GREEDY METHOD

Day	***90^{th} Percentile (Operations/ Holding Events)***	***80^{th} Percentile (Operations/ Holding Events)***	***70 Percentile (Operations/ Holding Events)***
June 5	60.25/-5.50	22.61/15.71	-2.40/8.97
June 10	61.25/-74.25	8.43/-20.71	9.20/-4.20
June 19	9.00/6.75	1.14/-15.71	34.03/-69.47
June 20	20.00/-8.00	19.14/-47.14	14.55/-8.00
July 17	23.50/-48.60	36.62/-60.90	161.17/-71.67
July 19	5.75/-5.75	3.57/3.43	2.90/14.40
August 15	2.50/2.50	2.86/0.0	0.0//0.0

IV. SUMMARY AND FUTURE WORK

In this paper, we presented a Rank-Sample Constrained ε-greedy method for searching for the TMI parameters that attempts to optimize an objective under a set of impactful convective weather scenarios. The operational performance of the approach was evaluated against samples from a TFI weather impact forecast distribution. The method was compared to a prior implementation of an ε-greedy approach that was adapted to selected TMI parameters. The resulting performance indicated that the proposed method can outperform a traditional ε-greedy approach in these situations by facilitating higher levels of throughput relative to the ε-greedy approach while often reducing the level of holding.

There are a number of areas of potential exploration that could be pursued based on our proposed methodology. Given the challenges of developing an algorithm that can be practically used on a time-scale needed to support operational decision-making, one might explore alternative ways to improve the computational performance of the simulation framework in order to train the algorithm more efficiently. Alternatively, pre-training the current approach offline with additional historical weather and traffic scenarios may also provide another way to improve the overall performance and efficiency of the methodology. Another area of potential inquiry is to examine the distributions in the TFI forecast model with the goal of developing better ways to sample and account for weather impacts from the tails of the distribution. Other input from stakeholders such as airlines could be incorporated to understand and quantify their own risk-preferences and their tolerance for and response to particularly impactful weather scenarios. By examining the question of how stakeholders perceive risk we can develop decision-support technologies that more acutely map to the concerns and potential responses of the air traffic management community in dealing with potentially unfavorable weather events.

ACKNOWLEDGMENT

The authors would like to thank Yan Glina for his help in generating some of the flight plan data used in this study. The authors would also like to thank Mike Matthews and Joe Venuti for their assistance in classifying of weather impact days used in this paper.

AUTHOR BIOGRAPHIES

James Jones is Technical Staff in the Air Traffic Control Systems Group at MIT Lincoln Laboratory. He has a PhD in Civil and Environmental Engineering from the University of Maryland.

Zachary Ellenbogen is Associate Staff in the Cyber System Assessments Group at MIT Lincoln Laboratory. Zachary received his BA in Math and Computer Science and his MA in Computer Science from Columbia University.

REFERENCES

[1] S. Ruiz, H. Kadour, and P. Choroba, "A novel air traffic flow management model to optimise network delay," in *The 13th USA/Europe ATM R&D Seminar*, Vienna, Austria, 2019, p. 10.

[2] J. E. Evans, K. Carusone, M. Wolfson, B. Crowe, D. Meyer, and D. Klingle-Wilson, "THE CORRIDOR INTEGRATED WEATHER SYSTEM (CIWS)," in *10th Conference on Aviation, Range, and Aerospace Meteorology*, 2001, p. 6.

[3] C. A. Provan, L. Cook, and J. Cunningham, "A probabilistic airport capacity model for improved ground delay program planning," in *2011 IEEE/AIAA 30th Digital Avionics Systems Conference*, Oct. 2011, pp. 2B6-1-2B6-12. doi: 10.1109/DASC.2011.6095990.

[4] R. Kicinger, J.-T. Chen, M. Steiner, and J. Pinto, "Airport Capacity Prediction with Explicit Consideration of Weather Forecast Uncertainty," *Journal of Air Transportation*, vol. 24, no. 1, pp. 18–28, Jan. 2016, doi: 10.2514/1.D0017.

[5] J. Cox and M. J. Kochenderfer, "Optimization Approaches to the Single Airport Ground-Holding Problem," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 12, pp. 2399–2406, Dec. 2015, doi: 10.2514/1.G001081.

[6] S.-L. (Alex) Tien, C. Taylor, E. Vargo, and C. Wanke, "Using Ensemble Weather Forecasts for Predicting Airport Arrival Capacity," *Journal of Air Transportation*, vol. 26, no. 3, pp. 123–132, Jul. 2018, doi: 10.2514/1.D0105.

[7] M. P. Matthews, M. S. Veillette, J. C. Venuti, R. A. DeLaura, and J. K. Kuchar, "Heterogeneous Convective Weather Forecast Translation into Airspace Permeability with Prediction Intervals," *Journal of Air Transportation*, vol. 24, no. 2, pp. 41–54, Apr. 2016, doi: 10.2514/1.D0025.

[8] M. O. Ball, R. Hoffman, A. R. Odoni, and R. Rifkin, "A Stochastic Integer Program with Dual Network Structure and Its Application to the Ground-Holding Problem," *Operations Research*, vol. 51, no. 1, pp. 167–171, Feb. 2003, doi: 10.1287/opre.51.1.167.12795.

[9] B. Kotnyek and O. Richetta, "Equitable Models for the Stochastic Ground-Holding Problem Under Collaborative Decision Making," *Transportation Science*, vol. 40, no. 2, pp. 133–146, May 2006, doi: 10.1287/trsc.1050.0129.

[10] O. Richetta and A. R. Odoni, "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control," *Transportation Science*, vol. 27, no. 3, p. 228, Aug. 1993, doi: 10.1287/trsc.27.3.228.

[11] A. Mukherjee and M. Hansen, "A Dynamic Stochastic Model for the Single Airport Ground Holding Problem," *Transportation Science*, vol. 41, no. 4, pp. 444–456, Nov. 2007, doi: 10.1287/trsc.1070.0210.

[12] D. Bertsimas, G. Lulli, and A. Odoni, "An Integer Optimization Approach to Large-Scale Air Traffic Flow Management," *Operations Research*, vol. 59, no. 1, pp. 211–227, Feb. 2011, doi: 10.1287/opre.1100.0899.

[13] A. S. Estes and M. O. Ball, "Equity and Strength in Stochastic Integer Programming Models for the Dynamic Single Airport Ground-Holding Problem," *Transportation Science*, vol. 54, no. 4, pp. 944–955, Jul. 2020, doi: 10.1287/trsc.2020.0975.

[14] G. Lulli and A. Odoni, "The European Air Traffic Flow Management Problem," *Transportation Science*, vol. 41, no. 4, pp. 431–443, Nov. 2007, doi: 10.1287/trsc.1070.0214.

[15] A. Agustı´n, A. Alonso-Ayuso, L. F. Escudero, and C. Pizarro, "On air traffic flow management with rerouting. Part II: Stochastic case," *European Journal of Operational Research*, vol. 219, no. 1, pp. 167–177, May 2012, doi: 10.1016/j.ejor.2011.12.032.

[16] G. Buxi and M. Hansen, "Generating Probabilistic Capacity Profiles from weather forecast: A design-of-experiment approach," in *The 9th USA/Europe Air Traffic Management Research and Development Seminar*, Berlin, Germany, 2011, p. 10.

[17] J. C. Jones and R. DeLaura, "Predicting Airport Capacity in the Presence of Winds," in *17th AIAA Aviation Technology, Integration, and Operations Conference*, Denver, Colorado: American Institute of Aeronautics and Astronautics, Jun. 2017. doi: 10.2514/6.2017-3595.

[18] J. C. Jones, R. DeLaura, M. Pawlak, S. Troxel, and N. Underhill, "Predicting & Quantifying Risk in Airport Capacity Profile Selection for Air Traffic Management," in *The 12th USA/Europe Air Traffic Management Research and Development Seminar*, Seattle, WA, Jun. 2017, p. 10.

[19] J. C. Jones, R. DeLaura, Y. Glina, and E. Hassey, "Learning Airspace Flow Rates through Fast-time Simulation," in *2018 Aviation Technology, Integration, and Operations Conference*, Atlanta, Georgia: American Institute of Aeronautics and Astronautics, Jun. 2018. doi: 10.2514/6.2018-2881.

[20] J. C. Jones and Y. Glina, "Estimating Flow Rates in Convective Weather: A Simulation-Based Approach," in *13th USA/Europe ATM R&D Semina*, Vienna, Austria, 2019, p. 10.

[21] J. C. Jones, Ellenbogen, Zachary, and Y. Glina, Yan, "Recommending Strategic Air Traffic Management Initiatives in Convective Weather," in *The 14th USA/Europe Air Traffic Management Research and Development Seminar*, VIRTUAL EVENT, Sep. 2021. Accessed: Jan. 24, 2023. [Online]. Available: https://drive.google.com/file/d/1esb9B4My2DflVSjHdCMjDDNNrRRBlVO3/view?usp=sharing&usp=embed_facebook

[22] C. Taylor, E. Vargo, Bromberg, Emily, and Carson, Everett, "Reinforcement Learning for Traffic Flow Management Decision Support," in *The 14th USA/Europe Air Traffic Management Research and Development Seminar*, VIRTUAL EVENT, Sep. 2021. Accessed: Jan. 24, 2023. [Online]. Available: https://drive.google.com/file/d/1xD26ZTN00FtOP_rsrEMGdNTSHI2lm7o2/view?usp=sharing&usp=embed_facebook

[23] C. Taylor, T. Masek, C. Wanke, and S. Roy, "Designing Traffic Flow Management Strategies Under Uncertainty," in *The 11th USA/Europe Air Traffic Management Research and Development Seminar*, Chicago, 2015, p. 10.

[24] M. P. Matthews and R. DeLaura, "Evaluation of enroute convective weather avoidance models based on planned and observed flights," *14th American Meteorological Society Conference on Aviation, Range, and Aerospace Meteorology*, p. 17, 2010.

[25] M. Rubnich, M. Matthews, and R. DeLaura, "Use of the Convective Weather Avoidance Polygon (CWAP) to Identify Temporally Coherent Convective Storm Boundaries," in *2013 Aviation Technology, Integration, and Operations Conference*, Los Angeles, CA: American Institute of Aeronautics and Astronautics, Aug. 2013. doi: 10.2514/6.2013-4215.

[26] M. P. Matthews and R. DeLaura, "Decision Risk in the Use of Convective Weather Forecasts for Trajectory-Based Operations," in *14th AIAA Aviation Technology, Integration, and Operations Conference*, Atlanta, GA: American Institute of Aeronautics and Astronautics, Jun. 2014. doi: 10.2514/6.2014-2717.

[27] C. Gong and D. Mcnally, "Dynamic Arrival Routes: A Trajectory-Based Weather Avoidance System for Merging Arrivals and Metering," in *15th AIAA Aviation Technology, Integration, and Operations Conference*, Dallas, TX: American Institute of Aeronautics and Astronautics, Jun. 2015. doi: 10.2514/6.2015-3394.

[28] J. Y. N. Cho, J. D. Welch, and N. K. Underhill, "Analytical Workload Model for Estimating En Route Sector Capacity in Convective Weather," in *The 9th USA/Europe Air Traffic Management Research and Development Seminar*, Berlin, Germany, 2011, p. 10.

[29] J. D. Welch, J. Y. N. Cho, N. K. Underhill, and R. A. DeLaura, "Sector Workload Model for Benefits Analysis and Convective Weather Capacity Prediction," in *The 10th USA/Europe Air Traffic Management Research and Development Seminar*, 2013, p. 10.

[30] M. P. Matthews, M. S. Veillette, J. C. Venuti, R. A. DeLaura, and J. K. Kuchar, "Heterogeneous Convective Weather Forecast Translation into Airspace Permeability with Prediction Intervals," *Journal of Air Transportation*, vol. 24, no. 2, pp. 41–54, Apr. 2016, doi: 10.2514/1.D0025.

[31] M. Matthews, R. DeLaura, M. Veillette, J. Venuti, and J. Kuchar, "Airspace Flow Rate Forecast Algorithms, Validation, and Implementation," MIT Lincoln Laboratory, Lexington, MA, Project Report ATC-428, 2015.

[32] T. Vossen and M. Ball, "Optimization and mediated bartering models for ground delay programs," *Naval Research Logistics*, vol. 53, no. 1, pp. 75–90, Feb. 2006, doi: 10.1002/nav.20123.

[33] A. Reuther *et al.*, "Interactive Supercomputing on 40,000 Cores for Machine Learning and Data Analysis," in *2018 IEEE High Performance extreme Computing Conference (HPEC)*, Waltham, MA: IEEE, Sep. 2018, pp. 1–6. doi: 10.1109/HPEC.2018.8547629.

[34] J. C. Jones, D. J. Lovell, and M. O. Ball, "Combining Control by CTA and Dynamic Enroute Speed Adjustment to Improve Ground Delay Program Performance," in *The 11th USA/Europe Air Traffic Management Research and Development Seminar*, Lisbon, Portugal, Jun. 2015, p. 10.

[35] J. C. Jones, D. J. Lovell, and M. O. Ball, "En Route Speed Control Methods for Transferring Terminal Delay," in *The 10th USA/Europe Air Traffic Management Research and Development Seminar*, Chicago, IL, Jun. 2013, p. 10.

[36] J. C. Jones, D. J. Lovell, and M. O. Ball, "Stochastic Optimization Models for Transferring Delay Along Flight Trajectories to Reduce Fuel Usage," *Transportation Science*, vol. 52, no. 1, pp. 134–149, Jan. 2018, doi: 10.1287/trsc.2016.0689.