Fourteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2021)

# Recommending Strategic Air Traffic Management Initiatives in Convective Weather*

James Jones, Zachary Ellenbogen, Yan Glina
MIT Lincoln Laboratory
244 Wood St, Lexington, MA 02421, USA

*Abstract*—**The presence of uncertainty in weather forecasts poses significant challenges for air traffic managers. These challenges can have major repercussions on stakeholders in terms of their impact on the delay within the system. In this paper, we discuss an approach for recommending Traffic Management Initiative (TMI) parameters during uncertain weather conditions. We propose four methods for TMI selection. The first two favor random exploration of TMI decisions. An epsilon-greedy approach and a softmax algorithm are also evaluated against the two random exploration approaches. A parallel fast-time simulation framework is presented for evaluating the proposed methods over a range of weather forecast scenarios. A set of regional TMIs is applied and tested against a case day in which the airspace capacity in the Northeast United States was compromised by convective weather. Both the softmax and epsilon-greedy approaches demonstrate strong performance relative to the other methods. The results suggest that the approach could potentially aid air traffic stakeholders in understanding how to best deal with weather forecast uncertainty.**

**Keywords—reinforcement learning, epsilon-greedy, softmax, traffic management initiatives, weather, simulation.**

## I. INTRODUCTION

Weather presents significant challenges to managing airspace and airport resources. When present, it limits the capacity at airspace and airport resources, which often causes airlines and Air Navigation Service Providers (ANSPs) to delay or cancel flights imposing significant costs to passengers and carriers annually. As weather forecasts are uncertain, particularly at longer time horizons, weather-related disruptions force stakeholders to alter flight schedules and the flow of air traffic around weather without clear knowledge of how these decisions will impact the flights and resources that they manage. In the presence of these challenges, traffic managers often adjust the flight demand for weather-compromised resources by imposing Traffic Management Initiatives (TMIs). In the US, these include initiatives such as Ground Delay Programs (GDPs), Airspace Flow Programs (AFPs), Ground Stops (GSs) and the new Collaborative Trajectory Options Programs (CTOPs), while in Europe the Network Manager applies demand-capacity balancing to flights [1] in order to better match the air traffic with the available capacity. These programs impose delays on flights thereby pushing the excess traffic demand back to later times during the day.

In the United States, traffic managers have traditionally managed traffic at the strategic level by consulting decision support systems such as the Traffic Flow Management System (TFMS) and weather forecasts such as the Corridor Integrated Weather System [2] to enhance their situational awareness of the traffic demand and weather. Traffic Managers use the information they receive from these tools along with their mental models and experiential knowledge of how the weather blocks the traffic flows along corridors within the airspace to manage traffic through the affected resources. Through consultation with other stakeholders, they attempt to make collaborative decisions about what actions to take in order to best mitigate the impact of the weather.

Making these traffic management decisions is an extremely complex endeavor. In this context, decision-makers are charged with translating uncertain weather forecast information into a time-varying estimate of the airspace blockage across all weather-affected resources. They must then translate this blockage information into an estimate of the future capacity of all of the airspace resources that they need to manage. Once they translate this weather to a capacity, they must then balance the demand and capacity at each resource across time. As weather forecasts are uncertain and mapping even perfect weather information to a capacity is an imprecise process, the inherent uncertainty in the problem grows with each passing stage of translation. While some of this demand mismatch can be managed tactically through mechanisms like path stretch, holding and time-based metering [3]–[9], when the demand/capacity mismatches are significant, most of the impact is not easily recoverable.

Beyond the challenges with using weather information to make actionable decisions, traffic managers are often forced to make not one but many decisions on how to control the traffic. They must decide what resources to control, when to impose control, what flights to control, where to control these flights, when to release control and at what rates the affected resources should be controlled to in order to match the projected demand. When paired together, this group of decisions represents an extremely computationally complex problem. This problem is further complicated by the fact that the decision-makers do not act as one body, but many with sometimes competing undisclosed objectives. For example, although there is some level of coordination, traffic managers may not know how many flights airlines will cancel. They also may not know how people in other facilities will respond as they are impacted by the initial traffic management plans.

As a result of all of these challenges, decision-makers are rarely able to perfectly match capacity to demand at the strategic level on the initial set of decisions. Accordingly, they often revise TMIs throughout the day as new information becomes available, circumstances change, the forecast horizon of the decision shortens and the reliability of the information that they are using increases. When cast in this context, the problem can be viewed as a set of sequential decisions over a series of stages where each subsequent decision depends on the outcome of the prior decision and the new information that is presented to the decision-maker.

Over the past three decades, a number of studies have focused on the subject of airspace and airport resource capacity management. This body of research includes both descriptive studies that aim to predict the capacity of airport and airspace resources with the goal of providing better translation of weather forecasts into actionable information, and prescriptive methods that issue specific recommendations for more optimally managing flight demand given the capacity uncertainty. The prescriptive approaches often employ integer programming to assign arrival times to flights at airports and/or airspace resources over a fixed time horizon while constraining the flight demand so that it does not exceed the capacity level of the managed resource(s) [10]–[21]. In these problems, the decision-maker either treats the problem deterministically, assuming all resource capacities are known, or they adopt a scenario-based approach in which they represent the capacity as a set of profiles that represent one potential manifestation of the evolution of resource capacities over time. In many cases the proposed approaches are capable of generating optimal solutions given the assumptions of the problem while demonstrating considerable computational tractability. The utility of the methods to support operational decision-making, however, has been limited to some extent due to the inability of weather forecasts and translational models to generate accurate estimates of resource capacities.

More recently there have been a number of advances in forecasting airport acceptance rates. These approaches issue predictions based on weather forecast products coupled with historical data on the airport acceptance rates, using machine learning to predict the airport acceptance rates given the forecast conditions [22]–[26]. The success of such developments have allowed researchers to leverage this information in data-driven integer programming models that support GDP planning [26]–[29].

Other efforts to predict resource capacity have focused on the airspace. At the tactical level, translational products such as the Convective Weather Avoidance Model (CWAM) [30], [31] provide additional visibility into the degree of airspace blockage imposed by the weather. At the strategic level, the Traffic Flow Impact (TFI) tool maps convective weather forecasts of up to 12 hours to predictions of en route airspace flow rates. The tool also provides a set of uncertainty bounds associated with the estimate [32]. The methodology has also been extended to forecast the airspace capacity in the terminal area [33]. While these estimates are a significant step towards providing capacity information to decision-makers, the tool does not provide any direct recommendation of how to better coordinate the flow of traffic in the form of a traffic management initiative given the forecast capacity impact.

One means of evaluating the impact of various interventions on air traffic is to use fast-time simulations [34] [26]. Given the appropriate weather and traffic information, these models can be used to statistically characterize the effect of weather on operational metrics such as delays and cancellations when different control strategies are implemented. One example of this type of model is NASPlay, an agent-based simulation with strategic en route weather translational models from TFI [35]. This tool allows users to simulate various traffic management initiatives and gauge the impact of these programs on the flight resources of interest. This capability was paired with automation to select TMI parameters to optimize flow rates over a set of flow constrained areas (FCAs) in the terminal area, [36], [37] however, the approach treated the forecasts as deterministic rather than stochastic profiles that evolve over time. Another approach injected weather scenarios from the Short-Range Ensemble Forecast (SREF) model into a fast-time simulation and used genetic algorithms to select a combination of TMIs over the Northeastern United States [38]. The approach demonstrated the use of automated design of TMIs and discretization of TMI parameters and examined a limited range of discretized TMI parameter values. While probabilistic forecasts such as the SREF can be used to model weather uncertainty, the forecast skill of the model has shown limitations due to its 6-hour resolution.

In this paper, we build on the needs identified above by developing methods for selecting TMI parameters to manage regional convective weather impacts on a set of airspace and airport resources. The proposed methods combine reinforcement learning with significant parallel processing to achieve strong TMIs performance over a range of stochastic weather scenarios. The first uses an epsilon-greedy (ε-greedy) algorithm, while the second uses a softmax approach to identify potential TMI candidate parameters. A fast-time simulation inspired by (and improving on) NASPlay, which leverages the TFI model, is used to study the impact of weather forecast uncertainty on operational performance.

In Section II, we describe our modeling framework and the methodology behind the approaches that we explored. Section III describes our computational experiments and presents an evaluation of our proposed methods using a case day in the Northeastern United States with significant convective weather. Within this scenario, we study the ability of our TMI parameter selection methods to mitigate the operational impacts of the affecting weather. Section IV provides a summary and proposed future expansions of this work.

## II. Methodology

### A. Characterizing the Weather Uncertainty

The shortcomings in weather translational decision support tools represent a critical gap in the capability of air traffic management operations. Although there has been some effort to develop weather-aided decision support tools, much of the work done to translate weather forecasts into airspace flow rate estimates has been limited to the sector level [39]–[41]. TFI generates hourly forecast probability quantiles of the percentage of weather-free airspace, or *permeability*, along corridors of en route traffic. The permeability forecast is provided, for an altitude of 35,000 ft., in twenty-five idealized rectangular regions overlapping major traffic management boundaries in the Eastern portion of the National Airspace System (NAS). These discrete regions, called *flow constrained areas* (FCAs), represent choke points upon which convective weather has a large impact on NAS throughput.

The actual permeability within an FCA at any given time is the fraction of airspace outside of clustered convective storm cells detected in a scalar Weather Avoidance Field translated from radar Vertically Integrated Liquid (VIL) and Echo Tops for 35,000 ft. TFI provides a forecast permeability via a Machine Learning approach [32]. A feature set is derived from the imagery of deterministic models such as NOAA's High Resolution Rapid Refresh (HRRR) product, as well as from probabilistic (lightning) models such as the Localized Aviation Model Output Statistics Program (LAMP). The permeability forecast model is trained with data from one or more convective seasons in a two-step process. In the first step, the input features are fit to the actual permeability for each model type by means of a Ridge Regression technique. In the second step, Quantile Regression among the model types is used to establish distribution boundaries for combinations of input models as a function of time and FCA.

Figure 1 shows an example TFI forecast distribution for an FCA near the transition to New York airspace from the Midwest. The origin of the forecast lead axis is relative to the issue time at 18 July 2019 at 13:00GMT. The solid black curve gives the median and the dashed lines give the 20$^{th}$ and 80$^{th}$ percentiles of the permeability distribution predicted for each hour after issue time. The presentation of a distribution provides a measure of *forecast confidence*, or alternatively, the uncertainty in the prediction. The light colors in the plot background represent three categories of permeability (low, medium, high) generally indicating the severity of the weather impact.

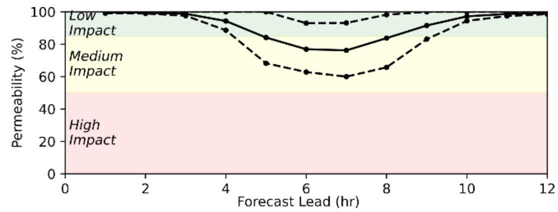


Figure 1. An example TFI forecast for New York transitional airspace.

In order to consolidate the forecast information into a form that is more consumable for fast-time simulations, the TFI distribution forecast is capable of generating time-series permeability values using Monte Carlo sampling over random draws from each quantile of the distribution. To obtain realistic variation vs. forecast lead hour, each time-series draw is conditioned to have the same time correlation behavior as the actual permeability observed in the training data for that FCA. In these calculations, the prediction interval at each forecast lead hour is assumed to be Gaussian, which is a good approximation in the body of the distribution but somewhat under-predicts the highest impact tails.

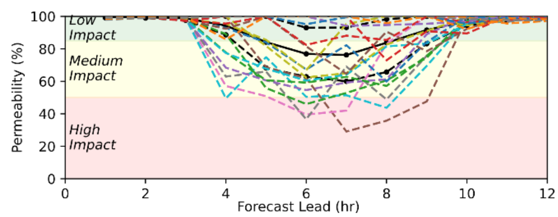


Figure 2. Correlated random samples from the TFI prediction distribution.

An example of twenty such draws is shown in Figure 2. Note that although the draws represent a large variation in possible airspace capacity, the general indication of the ensemble is that a medium impact event will begin about four hours from the issue time and last for five hours. These draws are currently uncorrelated across FCAs. An active area of research is to apply spatial correlation.

Draws of permeability are converted to flow rate (i.e., the number of aircraft estimated to be able to cross the FCA per hour) using tables established in studies of pilot avoidance of convective weather fields [42]. The major dependencies of that avoidance were observed to be the permeability in the airspace and the length of time the permeability existed at a given level. Table I shows the conversion table for an FCA region in the New York transitional airspace. The result is that the relationship between permeability and flow rate is largely linear, with a dampening of flow rate recovery depending on the amount of time the airspace capacity has been degraded. The flow rate is also driven by the airspace characteristics (e.g. size, traffic density) and will differ between regions.

TABLE I. EXAMPLE AIRCRAFT FLOW RATE (PER HR) AS A FUNCTION OF PERMEABILITY (%) FOR A REGION IN THE NEW YORK AIRSPACE

| Length of Impact (minutes) | Permeability Percentage | | | | | |
|---|---|---|---|---|---|---|
| | 0-20 | 21-40 | 41-60 | 61-80 | 81-99 | 100 |
| 1-15 | 62 | 73 | 97 | 116 | 120 | 120 |
| 15-30 | 39 | 55 | 73 | 86 | 94 | 120 |
| 30-45 | 30 | 44 | 57 | 70 | 94 | 120 |
| 45-90 | 18 | 34 | 45 | 70 | 94 | 120 |
| 90+ | 11 | 21 | 45 | 70 | 94 | 120 |

### B. *Searching for TMI Parameters*

The translation of weather forecasts to airspace capacities provides us with one representation of uncertainty. There is however, another dimension of uncertainty that can be captured by the decision-making process itself. Traffic managers will often make decisions that are continually revised throughout the day. The ground delay programs that are put in place can be cancelled or morph into ground stops. Likewise the rates on AFPs can become more severe or the programs can be relaxed or cancelled before they were initially scheduled to end. While some of these changes are caused by unexpected changes in the weather, other changes could be completely independent. We would like to capture these two dimensions of uncertainty by considering a set of staged decisions that evolve over time. At the first stage we make an initial decision and in the subsequent stages we receive new information about the weather, traffic demand and stakeholder objectives and can revise our previous decisions to improve the performance relative to our objectives. If we let $W_{ij}$≡The forecast from scenarios *i* in period *j* and T≡The set of all time stages, we can describe the process using a tree where the branches represent choices and the nodes represent states. In this construct, the decision-maker receives forecast information at each stage and progressively steps through a set of options. A notional depiction is shown in Figure 3.

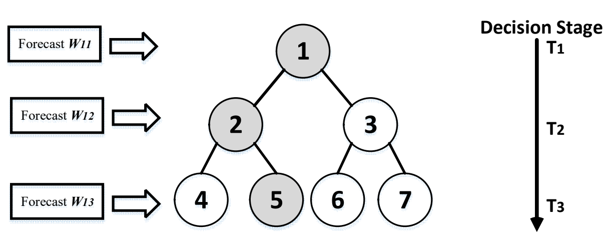


Figure 3. A decision tree for a single scenario with forecast inputs at each decision stage.

Under ideal circumstances we would like to be able to sample from all of the branches of the tree and compute the expected value of reaching each state so that we can maximize our expected reward. However, there is simply not enough time to consider even a large subset of the possible scenarios in an operational setting due to the computational complexity of multi-resource TMIs, the uncertainty associated with the weather forecasts and translation, the uncertainty associated with the later-stage decisions and the relatively short horizon of the decisions being made. We can still select an optimal action given the information that we have about the system and the likelihood that future states will occur. However, we ignore learning more about the value of other states because we do not choose to visit them and may miss out on better solutions. This is commonly referred to as the exploration-exploitation problem.

We propose four potential search strategies to identify the most promising TMIs. The first two favor heavy exploration of potential decisions, while the second two provide a more balanced mixture. Although the two exploration strategies run slightly faster as they do not have to learn from the data, all the approaches are able to generate solutions in less than 1 minute. While we believe the latter two methods will perform better since TMI performance has some dependency on the TMI parameters used, the degree to which this relationship between TMI parameters and TMI performce can be learned from the data is not clear given the computational complexity of the problem.

Random Local Exploration (RLE):

The Random Local Exploration approach begins with a baseline vector that describes the program charateristics of the TMI (e.g. resource, rate, scope, start time, end time). This baseline TMI vector is perturbed with a randomly sampled autoregressive distribution for time period $t=\{n...T\}$ while the initial period is kept constant. Initially $n$ is set to 1 but the value is updated as the decision-maker makes new decisions and receives new information about the state of the air traffic management system. This fixed initial vector allows us to explore and learn the potential recourse decisions more thoroughly at the expense of exploring other dissimilar TMI options. The value of the decision is then obtained by averaging over the resulting samples. A depiction of the process is shown in Figure 4.

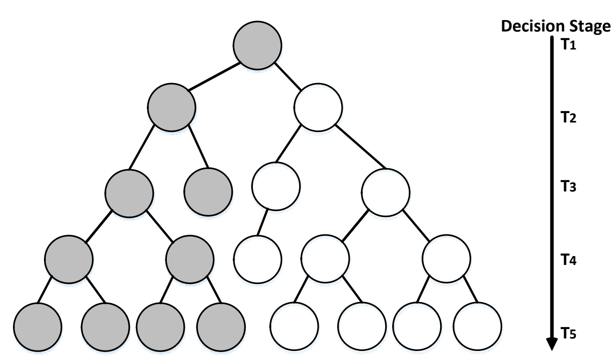


Figure 4. A set of samples take using random local exploration (RLE) over a set of staged decisions.

Random Global Exploration (RGE):

This approach begins by perturbing a baseline vector that describes the program charateristics of the TMI (e.g. resource, rate, scope, start time, end time) with an autoregressive distribution. Unlike the Random Local Exploration (RLE) method, the TMI vector is randomly sampled over all time periods t={0...T}. This less restrictive sampling range permits a more universal exploration of states and different TMI options. Another difference between the prior approach and this one is that the value of the decisions is obtained by taking the maximum value of all sampled forecast scenarios. An example of the search process is shown in Figure 5.

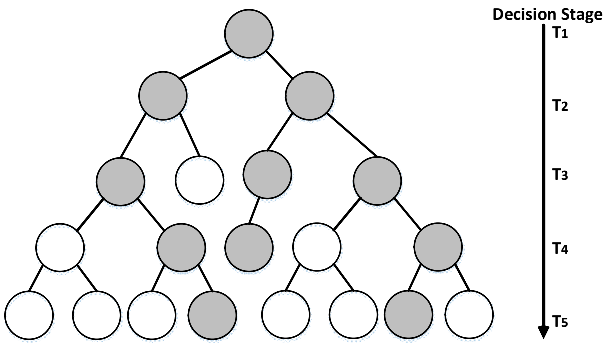


Figure 5. A set of samples take using random global exploration (RGE) over a set of staged decisions.

ε-greedy policies:

The prior two methods dealt with the issues of exploration vs. exploitation by heavily favoring exploration. One drawback of these approaches is that they do not make use of the data that they have collected as they go through the search process. As a result they may explore many poor solutions that could have been foreseen based on previous history. On the other hand, it may be undesirable to learn from the data when we do not have enough information as we may train our search to look for the wrong things. ε-greedy policies are commonly used to deal with these conflicting concerns. Under this approach, the decision-maker samples from a random distribution. If the sample value exceeds a level ε, then the policy selects a decision based on random selection, otherwise the algorithm selects the optimal action given the information currently available at the time of the decision based on an estimate of the value of each choice. Initially the algorithm typically favors a selection of random actions. When this happens, the algorithm can learn the value of the action and update the estimates accordingly. As more decisions are made and more samples are collected, the value of acquiring new information often diminishes. As a result, the value of epsilon is often set to decrease after each action to reflect an increase in the value of exploiting the existing information. A description of the algorithm is shown in Table II with the ε threshold occurring in step 5. The value of each decision can be fit using a variety of methods. In our implementation we train our model using non-parametric supervised learning methods (e.g., Random Forest Regression, Gradient Boosting Regression, Support Vector Regression etc.).

If we let:

$\Omega \equiv$The set of all scenarios

$T \equiv$The set of all time stages

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

$W_t \equiv$The airspace capacity described in forecast in period $t$

$J \equiv$The set of all simulation instances

$\omega^m \equiv$ The scenario forecast $m$

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

$x_t^{n,m} \equiv$The TMI decision made at time $t$ in scenario $m$ for instance $j$ on trial $n$

$C(S_t^{n,m}, x_{tj}^{n,m}) \equiv$ The cost to the decision made at time $t$

$\hat{v}_{t,j}^{n,m} \equiv$ The value of the decision at time $t$ in scenario $m$ for instance $j$ on trial $n$

TABLE II. AN ε-GREEDY APPROACH FOR ASSIGNING TMIs

| |
|---|
| Step 0. Initialization |
| Step 0a. Initialize $\bar{V}_{tj}^0, t \in T, j \in J$ |
| Step 0b. Initialize $S_{0,j}^1, j \in J$ |
| 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 $\omega^m$ |
| Step 2. Do for $t = 0, 1, \dots, T$ and j =$0, 1, \dots, J$ |
| Step 2a. Find $x_{tj}^{n,m} = X^{\pi,n-1}(S_{t,j}^{n,m})$ |
| Step 2b. Update the state variable by simulating the air traffic $S_{t+1,j}^{n,m} = S^M(S_{t,j}^{n,m}, x_{tj}^{n,m}, W_{t+1}(\omega^m))$ |
| Step 2c. Set $\hat{v}_{0j}^{n,m} = 0$, and $t = 0, 1, \dots, T$ and j =$0, 1, \dots, J$ $\hat{v}_{t,j}^{n,m} = C(S_{t,j}^{n,m}, x_{tj}^{n,m}) + \hat{v}_{t-1,j}^{n,m}$ |
| Step 3. Compute the average value of starting in state $S_{0,j}^1$ $\bar{v}_{0,j}^n = \frac{1}{M}\sum_{m=1}^M \hat{v}_{0,j}^{n,m}$ |
| Step 4. Update the value function approximation by using the average values by fitting an estimate $\bar{V} \leftarrow U^V(\bar{V}_{0,j}^{n-1}, S_{0,j}^{x,n}, \bar{v}_{0,j}^n)$ |
| Step 5. With probability $\epsilon$, 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 $X_j^{\pi,n}(S) = arg_{x \in X}\max(C(S_{0,j}^n, x) + \bar{V}_{0j}^n(S^{M,a}(S_{0,j}^n, x)))$ |
| 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 $\epsilon = \alpha/(\alpha+n+1)$, where $\alpha$ 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 $(\bar{V}_{t,j}^n)_{t=1}^T$ |

Softmax Exploration:

One issue with the ε-greedy approach is that when it selects an exploration step, it is equally likely to sample from all potential state-actions. As noted with the prior exploration-based strategies, this uniform sampling can result in the selection of some very poor choices where we may learn little from the decision. The softmax approach gets around this issue by weighting the probability of selecting a given decision by the estimated value of that decision. An expression for this probability is shown in below:

$$Q(s,x) = C(S,x) + \bar{V}^n(S,x) \quad (1)$$

$$P^n(s,a) = \frac{e^{\frac{Q(s,x)}{T}}}{\sum_{x'=X} e^{\frac{Q(s,x')}{T}}} \quad (2)$$

where $Q(s,x)$ is the value of choosing decision $x$ and $T$ is called the temperature. Similar to the ε-greedy approach, we would like to reduce the scaling parameter $T$ value as we iterate. Decrementing $T$ increases the odds of selecting the best choices as sampling continues. We can apply the softmax approach using the same algorithmic steps shown in Table II and substituting the ε–greedy selection with this probabilistic selection of TMI decisions in Step 5. In this context we would select from $R$ possible TMI decisions $J$ times.

### C. *Simulating the Air Traffic*

A custom fast-time simulation framework was used to evaluate the effects of our algorithmic selection of TMI parameters. The simulation ingests airspace flow rates from TFI based on either forecast or actual weather. The simulation can also use wind forecast models such as the High Resolution Rapid Refresh model or the Global Forecast System model to adjust the four-dimensional flight trajectories. The sector workload constraints are enforced using the analytical models in [40] and [41]. The simulation also consumes flight plans based on historical data from the Traffic Flow Management System. Aircraft speed profiles derived from the BADA 3.6 model. FAA coded instrument flight navigation procedures are used to generate the initial four dimensional trajectories required for the traffic management step of the simulation. The traffic management initiatives are modeled based on CDM procedures such as ration-by-schedule, cancellation and substitution and compression [43].

The simulation initializes the algorithm by generating a TMI description that is applied to the scenario. Once the simulation is initialized, a set of concurrent instantiations of the air traffic simulation are run in parallel. Each simulation samples a set of TFI-correlated FCA draws that are based on the forecast weather. After our initialization, we then use one of the selection methods to choose a set of TMI parameters based on the airspace and airport resources evaluated in the scenario. When a TMI is applied, the traffic flight demand for the affected resources is throttled and the temporal dimension of the traffic flow changes, introducing a number of follow-on effects. We observe the performance of this strategy in the simulation and score the resulting metrics (e.g., number of operations, holding events, delay) that characterize the TMI performance. As this process iterates, the framework continues the series of air traffic management simulations by using our selection method to choose different sets of TMI parameters until we have reached our last simulation run. A diagram of the process is shown in Figure 6.

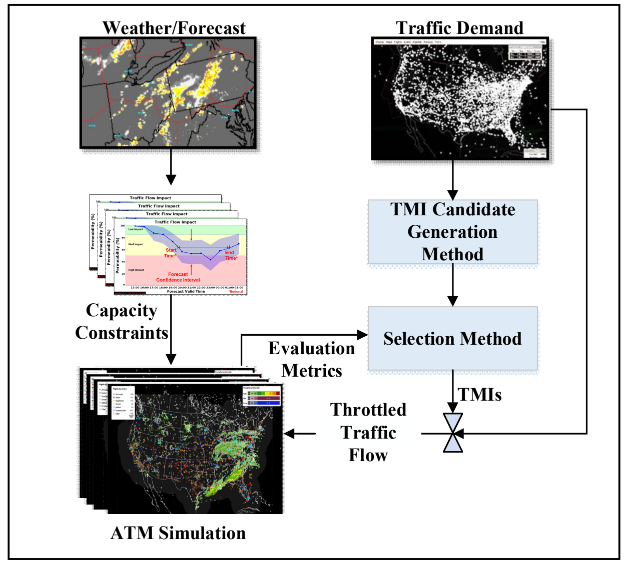


Figure 6. A simulation framework to facilitate exploration of Traffic Flow Management Initiative parameters.

## III. Results and Discussion

A set of computational experiments were conducted to evaluate the performance of each of the proposed methods using the fast-time air traffic simulation discussed in the previous section. In this section, we study the relative performance of each method at a number of TMI decision stages. We also discuss the implications of the results for TMI decision support.

### A. *Experimental Description*

A computational experiment was conducted using the modeling framework described in the previous section. The selected scenario used the weather that occurred on June 10, 2019. This day was chosen because the weather impacts were moderate but the convective blockage prompted traffic managers to impose a number of rolling TMIs throughout the Northeast. As such it was a good case day for the type of progressive air traffic management decision-making that our approach could be exercised on. A snapshot of the weather is shown in Figure 7.

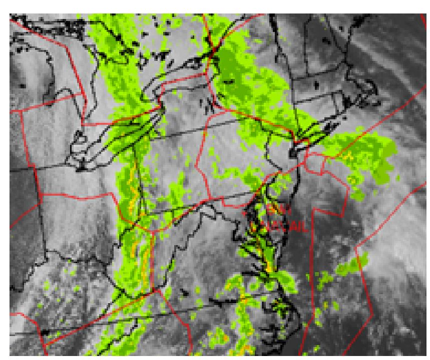

Figure 7. Weather image at 20:00 GMT on June 10, 2019.

A set of simulations was configured to evaluate the traffic over a period lasting between 04:00GMT on June 10 to 3:59GMT on June 11. Flight plans from the entire set of traffic in the U.S. National Airspace were injected and flown in the simulation. A set of TMIs were imposed at 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 8.

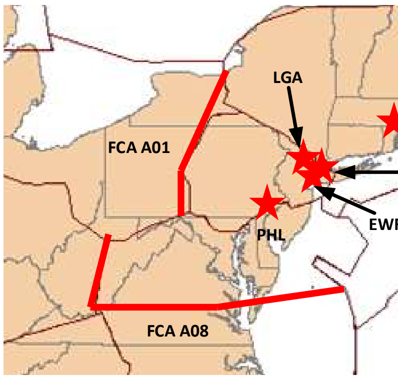


Figure 8. The location of the set of airports and FCA regions in the TMI programs during the case day under examination.

The modeling framework was developed on a high-performance computing cluster [44]. During each run, the model sampled 30 randomly drawn TFI FCA forecasts representing 30 different spatial-temporal evolutions of the weather with 15 different TMI parameter configurations representing 15 different TMIs. Each TMI was tested on all 30 TFI sample draws for a total of 450 concurrent instances, each on a single Intel xeon-e5 processor.

The TMIs were configured on a nominal baseline using a 1$^{st}$-order Gaussian auto-regressive random process that produces time difference values that are parameterized with the mean, standard deviation and auto-correlation. A nominal baseline TMI was configured to a mean value that was set to 80% of the airport capacity in the case of the airport resources, a rate of 115 flights/hour in the case of FCAA01 and 120 flights/hour in the case of FCAA08. The rates were then perturbed using a distribution with a standard deviation of 3 in the case of the airports and 5 in the case of the AFPs as the FCA resources typically have higher flow rates associated with them. The exemption radii of the GDPs were limited to distances of 1000, 1500, 2000 and 2500 NM. The initial TMIs lasted for 12 hours. During each strategy evaluated, the simulation cycled through a set of 40 runs. All the TMIs implemented in this experiment were initiated at a start time of 18:00GMT. In order to study the effect of revising the TMIs at various decision stages, the simulation was configured to issue revisions to the programs at 21:00 GMT and 00:00GMT, which we shall refer to as the 0-hour, 3-hour and 6-hour decision stages. During these revisions, the rate and exemption radius decisions made from 18:00GMT up until the revision time are kept fixed at whatever value they were set to by the prior decision. Looking forward beyond that time, the simulation has the option of revising the rate or cancelling the program at any given resource. Note that, while the planned rates were implemented at hourly intervals, we only revise our decisions every 3 hours. A summary of the various test instances is shown in Table III.

The Gradient Boosting Regression method was used to fit the value of various TMI decisions in our implementations of the ε-greedy and softmax algorithms as it demonstrated strong performance for related work [37]. A least-squares loss function was selected to fit the model estimates. As mentioned previously, the ε-greedy and softmax algorithms both vary the threshold for selection using a learning rate of $\alpha$=5, in the case of the ε-greedy, and T=4(1-n/40)+1, in the case of the softmax approach. While there are many potential parameter values that could be explored, since the dimensionality of the problem space is already very large, we will leave the subject of sensitivity analysis as an area for future study. Two gradient tree boosting models were created, one to predict the number of operations and the other to predict the number of flights that had 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.

| **Selection Method** | TABLE III.: SIMULATION TEST PARAMETER CONFIGURATIONS | | | |
|---|---|---|---|---|
| | ***TFI,TMI, Trials*** | ***Exemption Radii NM*** | ***Decision Stages*** | ***Std Rate Perturbation*** |
| **RLE** | 30,15,40 | 1000,1500, 2000, 2500 | 0 hour, 3hour, 6 hour | 3 GDP, 5 AFP |
| **RGE** | 30,15,40 | 1000,1500, 2000, 2500 | 0 hour, 3hour, 6 hour | 3 GDP, 5 AFP |
| **ε-greedy** | 30,15,40 | 1000,1500, 2000, 2500 | 0 hour, 3hour, 6 hour | 3 GDP, 5 AFP |
| **Softmax** | 30,15,40 | 1000,1500, 2000, 2500 | 0 hour, 3hour, 6 hour | 3 GDP, 5 AFP |

## *B. Results*

A set of metrics were used to quantify the relative performance of the different approaches. As we decided to prioritize throughput at the airports, we tallied the number of operations (arrivals and departures) at the five airports where TMIs were applied. Although the weather impact was limited during the earlier part of the day, operations numbers were calculated over a 24-hour period. As it is possible to maintain a high level of throughput at an airport and still have operationally undesirable conditions if there is significant holding present, we also computed the number of flights that needed to hold for longer than 15 minutes. As a general objective, we would like to keep the number of operations as high as possible while limiting the amount of holding whenever possible. An initial simulation was run with additional demand to understand the system throughput limits. From this limit we take the ideal operating point to be the throughput under saturation with no holding events. While it is debatable whether the system could actually operate under those conditions, for the purposes of this study we would like our TMIs to generate metrics that are as close as possible to this level. The results of the 40 trials with each method are shown in Figures 9-11 for the 0 hour, 3 hour and 6 hour cases respectively, while a set of summary statistics are listed in Table IV.

Since we aim to maximize operations while limiting holding events, we would like our TMIs to produce values in the lower right corner of the graphs. By that measure, the RLE method generally underperforms. This is not surprising as it is the least sophisticated of the approaches implemented and can be viewed as the relative baseline. The solutions identified by the RGE, however, perform significantly better. This is likely due to the broader nature of the search. Since we have a large range of TMI parameters to consider due to the dimensionality of the problem, it is generally better to consider more solutions. There is a significant increase in the number of operations and drop in the amount of holding from stage to stage suggesting that the forecast may have over-predicted the weather impact. Yet, while the performance improves in all cases, the ε-greedy and softmax approaches outperform the two more exploratory methods. In both cases the mean values of operations are consistently higher in all instances and the relative standard deviation of the number of operations are lower suggesting that the selection performance is more consistent. Although in the case of the 3-hour decision stage the RGE exhibits marginally lower holding numbers than the ε-greedy approach, it also does so with an average of over 20 fewer operations. More importantly, the maximum values of the resulting method for the two learning-based approaches are consistently higher, suggesting that they are more likely to find the best solutions. It is somewhat unclear which learning-based method performs better. In the 0-hour decision stage, the softmax method finds TMIs that lead to a higher number of operations. However, in the two later stages the ε-greedy outperforms the others, exhibiting noticeably higher numbers of operations.

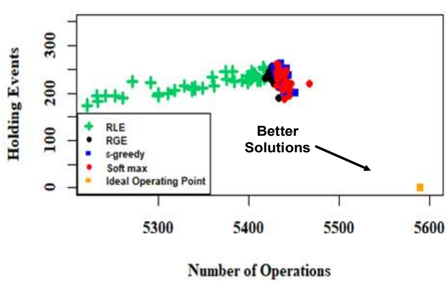


Figure 9. TMI program preformance after revisions in the 0 hour stage.

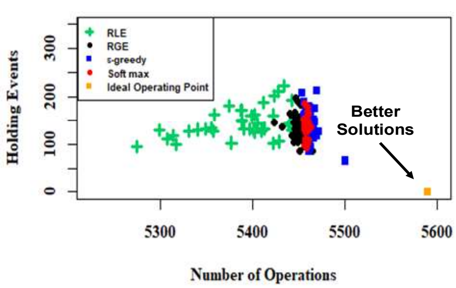


Figure 10. TMI program preformance after revisions in the 3 hour stage.

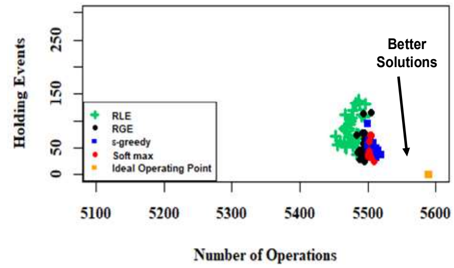


Figure 11. TMI program preformance after revisions in the 3 and 6 hour stages.

TABLE IV. AGGREGATED OPERATIONS PERFORMANCE OF EACH SELECTION METHOD

| Hour | Selection Method | *Number of Operations Mean (STD)* | *Number of Holds Mean (STD)* | *Maximum Ops. Case Ops./Holds* |
|---|---|---|---|---|
| 0 | **RLE** | 5345 (66) | 219 (21) | 5420/236 |
| | **RGE** | 5428 (18) | 239 (14) | 5440/218 |
| | **ε-greedy** | 5436 (6) | 233 (19) | 5452/199 |
| | **Softmax** | 5439 (7) | 229 (18) | 5467/220 |
| 3 | **RLE** | 5389 (44) | 144 (30) | 5443/131 |
| | **RGE** | 5448 (7) | 138 (27) | 5465/86 |
| | **ε-greedy** | 5464 (7) | 139 (32) | 5501/66 |
| | **Softmax** | 5459 (2) | 137 (20) | 5463/135 |
| 6 | **RLE** | 5477 (11) | 80 (27) | 5497/131 |
| | **RGE** | 5495 (4) | 49 (20) | 5505/39 |
| | **ε-greedy** | 5511 (5) | 46 (10) | 5519/37 |
| | **Softmax** | 5502 (2) | 43 (15) | 5509/21 |

The performance of the TMIs in terms of the number of operations should be taken in context. While we can infer relative improvement with respect to the objectives, the intervention may still be undesireable if the improvement in the objectives results in significantly worse performance in other metrics.

One might question what type of delay impact these initiatives have on the airports. To answer this question, we computed the mean and standard deviation of the total airport delay in hours at each decision stage. The resulting performance is shown in Table V. The data suggests that the two learning-based methods (of ε-greedy and softmax) yield lower amounts of delay with significantly less sample variance than the two more exploratory methods. The softmax function exhibits a particularly low level of sample variance. One interesting feature of the data is that, although the number of operations increases as the weather plays out and the forecast and TMI decisions are revised, the amount of delay increases.

This is particularly true of the two exploratory methods. The one exception to this trend is the softmax method which remains fairly stable despite yielding considerable improvement in terms of operations and holding.

The strong performance in all three planning stages suggest that the proposed parameter identification approach could potentially be used at various points throughout the day to aid traffic managers in supporting strategic traffic planning, as there is a clear mapping of weather forecasts and their associated uncertainty to the range of operational performance when a set of operational plans is enacted. Instead of prescribing flow rates that vary across time and a broad set of resources, the approach provides specific recommendations on actions that the decision-maker can take to deal with the implied weather impact. By speaking to stakeholders in this context, we can more readily communicate the implications of weather forecast uncertainty in terms which they may more easily relate and integrate into their operational decision-making processes.

TABLE V. AGGREGATED DELAY HOURS OF EACH SELECTION METHOD

| **Selection method** | ***0 hour Mean (STD)*** | ***3 hour Mean (STD)*** | ***6 hour Mean (STD)*** |
|---|---|---|---|
| **RLE** | 6924 (190) | 7026 (135) | 7009 (48) |
| **RGE** | 6859 (210) | 6868 (193) | 6979 (191) |
| **ε-greedy** | 6772 (90) | 6795 (51) | 6829 (26) |
| **Softmax** | 6777 (73) | 6809 (31) | 6754 (10) |

## IV. SUMMARY AND FUTURE WORK

In this paper we presented four methods for selecting TMI program parameters in the presence of weather forecast uncertainty. Two of the methods used exploration-based search strategies in order to identify the appropriate rates and program exemption radii. The other two used a more balanced mixture of exploration and exploitation to search through the parameter space using reinforcement learning approaches. A highly-parallelized air traffic simulation was used to evaluate the proposed approaches with a range of weather translational forecasts. A set of TMIs was applied to this framework over the range of regional airport and airspace resources. The resulting performance suggests that the two learning-based approaches may provide some enhanced capability to select TMI parameters in the presence of uncertainty.

There are a number of potential areas of research that could be explored to further the objectives proposed in this paper. As the selection of TMIs is a problem plagued by dimensionality, there is a need to develop methods to accelerate the pace of parameter exploration and the evaluation of potential search options. These advancements could include improvements in the computing architecture and simulation runtime, and more effective search policies. Additionally, future studies could examine a framework that considers a broad range of objectives such as cancellations, predictability and more active incorporation of the relative preferences and risk tolerance of various stakeholders. The proposed search strategies could then be modified to include these relative weightings in their prioritization of TMIs. As the effectiveness of the approach is only as good as the credibility of the simulation that we are using, future studies will need to validate the model through a comparison of metrics obtained over a set of case days. Finally, it should be noted that more user feedback is needed from stakeholders as this methodology evolves to ensure the operational realism of the approach and reflect the concerns of the system decision-makers.

## ACKNOWLEDGMENTS

The authors would like to thank Mark Worris for his help in generating the TFI forecast samples used in this analysis. The authors would also like to thank Amy Alexander, Mike Matthews and Joe Venuti for their participation in many thoughtful discussions.

## 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 Air Traffic Management Research and Development Seminar*, Vienna, Austria, Jun. 2015, 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]L. Delgado, X. Prats, and B. Sridhar, “Cruise speed reduction for ground delay programs: A case study for San Francisco International Airport arrivals | Elsevier Enhanced Reader,” *Transp. Res. Part C*, vol. 36, pp. 83–96, 2013.

[4]L. Delgado and X. Prats, “Effect of Radii of Exemption on Ground Delay Programs with Operating Cost Based Cruise Speed Reduction,” *Transp. Res. Part C*, p. 10.

[5]J. C. Jones and D. J. Lovell, “Methods for Curbing Exemption Bias in Ground Delay Programs through Speed Control,” *Transp. Res. Rec.*, p. 8.

[6]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.

[7]Y. Xu and X. Prats, “Effects of linear holding for reducing additional flight delays without extra fuel consumption | Elsevier Enhanced Reader,” *Transp. Res. Part D*, vol. 53, pp. 388–397, 2017.

[8]N. M. Smith, C. Brasil, P. U. Lee, N. Buckley, C. Gabriel, C. P. Mohlenbrink, F. Omar, B. Parke, C. Speridakos, and H.-S. Yoo, “Integrated Demand Management: Coordinating Strategic and Tactical Flow Scheduling Operations,” presented at the 16th AIAA Aviation Technology, Integration, and Operations Conference, Washington, D.C., Jun. 2016.

[9]G. Hodell, H. Yoo, C. L. Brasil, N. Buckley, C. V. Gabriel, S. Kalush, P. U. Lee, and N. M. Smith, “Evaluation of Multiple Flow Constrained Area Capacity Setting Methods for Collaborative Trajectory Options Program,” in *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*, Sep. 2018, pp. 1–8.

[10]O. Richetta and A. R. Odoni, “Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control,” *Transp. Sci.*, Aug. 1993.

[11]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,” *Oper. Res.*, vol. 51, no. 1, pp. 167–171, Feb. 2003.

[12]B. Kotnyek and O. Richetta, “Equitable Models for the Stochastic Ground-Holding Problem Under Collaborative Decision Making,” *Transp. Sci.*, vol. 40, no. 2, pp. 133–146, May 2006.

[13]A. Mukherjee and M. Hansen, “A Dynamic Stochastic Model for the Single Airport Ground Holding Problem,” *Transp. Sci.*, vol. 41, no. 4, pp. 444–456, Nov. 2007.

[14]D. Bertsimas, G. Lulli, and A. Odoni, “An Integer Optimization Approach to Large-Scale Air Traffic Flow Management,” *Oper. Res.*, vol. 59, no. 1, pp. 211–227, Feb. 2011.

[15]G. Lulli and A. Odoni, "The European Air Traffic Flow Management Problem," *Transp. Sci.*, vol. 41, no. 4, pp. 431–443, Nov. 2007.
[16]A. Agustı´n, A. Alonso-Ayuso, L. F. Escudero, and C. Pizarro, "On air traffic flow management with rerouting. Part I: Deterministic case," *Eur. J. Oper. Res.*, vol. 219, no. 1, pp. 156–166, May 2012.
[17]A. Agustı´n, A. Alonso-Ayuso, L. F. Escudero, and C. Pizarro, "On air traffic flow management with rerouting. Part II: Stochastic case," *Eur. J. Oper. Res.*, vol. 219, no. 1, pp. 167–177, May 2012.
[18]R. Hoffman, B. Hackney, P. Wei, and G. Zhu, "Enhanced Stochastic Optimization Model (ESOM) for Setting Flow Rates in Collaborative Trajectory Options Programs (CTOP)," presented at the 2018 Aviation Technology, Integration, and Operations Conference, Atlanta, Georgia, Jun. 2018.
[19]G. Zhu, P. Wei, R. Hoffman, and B. Hackney, "Centralized Disaggregate Stochastic Allocation Models for Collaborative Trajectory Options Program (CTOP)," in *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*, Sep. 2018, pp. 1–10.
[20]G. Zhu, P. Wei, R. Hoffman, and B. Hackney, "Risk-hedged Multistage Stochastic Programming Model for Setting Flow Rates in Collaborative Trajectory Options Programs (CTOP)," presented at the AIAA SciTech 2019 Forum, San Diego, California, Jan. 2019..
[21]A. S. Estes and M. O. Ball, "Equity and Strength in Stochastic Integer Programming Models for the Dynamic Single Airport Ground-Holding Problem," *Transp. Sci.*, vol. 54, no. 4, pp. 944–955, Jul. 2020.
[22]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.
[23]R. Kicinger, J.-T. Chen, M. Steiner, and J. Pinto, "Airport Capacity Prediction with Explicit Consideration of Weather Forecast Uncertainty," *J. Air Transp.*, vol. 24, no. 1, pp. 18–28, Jan. 2016.
[24]J. Cox and M. J. Kochenderfer, "Optimization Approaches to the Single Airport Ground-Holding Problem," *J. Guid. Control Dyn.*, vol. 38, no. 12, pp. 2399–2406, Dec. 2015.
[25]S.-L. (Alex) Tien, C. Taylor, E. Vargo, and C. Wanke, "Using Ensemble Weather Forecasts for Predicting Airport Arrival Capacity," *J. Air Transp.*, vol. 26, no. 3, pp. 123–132, Jul. 2018.
[26]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.
[27]L. Shisler, C. Provan, W. N. Chan, S. Grabbe, D. A. Clark, K. Venzke, C. Riley, D. Gilani, and E. Corcoran, "An Operational Evaluation of the Ground Delay Program Parameters Selection Model (GPSM)," *Air Traffic Control Q.*, vol. 21, no. 3, pp. 299–324, Jul. 2013.
[28]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.
[29]J. C. Jones and R. DeLaura, "Predicting Airport Capacity in the Presence of Winds," presented at the 17th AIAA Aviation Technology, Integration, and Operations Conference, Denver, Colorado, Jun. 2017.
[30]M. Matthews and R. DeLaura, "Assessment and Interpretation of En Route Weather Avoidance Fields from the Convective Weather Avoidance Model," presented at the 10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Fort Worth, Texas, Sep. 2010.
[31]M. P. Matthews and R. DeLaura, "Evaluation of enroute convective weather avoidance models based on planned and observed flights," *14th Am. Meteorol. Soc. Conf. Aviat. Range Aerosp. Meteorol.*, p. 17, 2010.
[32]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," *J. Air Transp.*, vol. 24, no. 2, pp. 41–54, Apr. 2016.
[33]M. P. Matthews, R. DeLaura, and J. Venuti, "Strategic Forecasts of TRACON Airspace Capacity during Convective Weather Impacts," presented at the 17th AIAA Aviation Technology, Integration, and Operations Conference, Denver, Colorado, Jun. 2017.
[34]K. D. Bilimoria, B. Sridhar, S. R. Grabbe, G. B. Chatterji, and K. S. Sheth, "FACET: Future ATM Concepts Evaluation Tool," *Air Traffic Control Q.*, vol. 9, no. 1, pp. 1–20, Jan. 2001.
[35]H. J. Davison Reynolds, R. DeLaura, and B. Soulliard, "'How Good Is Good Enough?' Exploring Validation for Serious Gaming in Air Traffic Control," in *Advances in Human Factors, Business Management, Training and Education*, Cham, 2017, pp. 687–697.
[36]J. C. Jones, R. DeLaura, Y. Glina, and E. Hassey, "Learning Airspace Flow Rates through Fast-time Simulation," presented at the 2018 Aviation Technology, Integration, and Operations Conference, Atlanta, Georgia, Jun. 2018.
[37]J. C. Jones and Y. Glina, "Estimating Flow Rates in Convective Weather: A Simulation-Based Approach," in *13th USA/Europe ATM R&D Seminar*, Vienna, Austria, 2019, p. 10.
[38]C. Taylor, T. Masek, C. Wanke, and S. Roy, "Designing Traffic Flow Management Strategies Under Uncertainty," in *The 10th USA/Europe Air Traffic Management Research and Development Seminar*, Chicago, 2013, p. 10.
[39]L. Song, C. Wanke, and D. Greenbaum, "Predicting Sector Capacity for TFM Decision Support," presented at the 6th AIAA Aviation Technology, Integration and Operations Conference (ATIO), Wichita, Kansas, Sep. 2006.
[40]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.
[41]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.
[42]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.
[43]T. Vossen and M. Ball, "Optimization and mediated bartering models for ground delay programs," *Nav. Res. Logist.*, vol. 53, no. 1, pp. 75–90, Feb. 2006.
[44]A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein, L. Milechin, J. Mullen, A. Prout, A. Rosa, C. Yee, and P. Michaleas, "Interactive Supercomputing on 40,000 Cores for Machine Learning and Data Analysis," in *2018 IEEE High Performance extreme Computing Conference (HPEC)*, Waltham, MA, Sep. 2018, pp. 1–6.

## 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 a B.A. in Math and Computer Science and a M.A. in Computer Science from Columbia University.

**Yan Glina** is Associate Staff in the Air Traffic Control Systems Group at MIT Lincoln Laboratory. He received his M.S. degree in Electrical and Computer Engineering in 2000 from Northeastern University.