

# An Enhanced Autonomy Approach to Automated Trajectory Negotiation

Alicia Fernandes<sup>1</sup>, Brian Capozzi<sup>2</sup> and George Hunter<sup>3</sup>  
*Mosaic ATM, Leesburg, VA*

Terique Barney<sup>4</sup>  
*NASA Langley Research Center, Hampton, VA*

**Improved real-time decision-making capabilities in both civil transport avionics and airline dispatcher workstations are helping to improve commercial aviation operations in many ways. One emerging capability is the automated derivation of alternate flight plans that yield both flight cost savings and smoother workflow. This emerging capability may be complemented with automated trajectory negotiation capabilities. Here we explore such automated trajectory negotiation capabilities, potential benefit mechanisms, and the required tools, procedures, and architectures.**

## I. Introduction

Completing the Federal Aviation Administration's (FAA's) transition to Trajectory Based Operations (TBO) requires frequent, fast negotiations to adjust flight trajectories to account for the many uncertainties in National Airspace System (NAS) operations. These negotiations will be required not only pre-departure, but also while aircraft are en route. Furthermore, an increasing number of flight deck applications and improved air/ground connectivity will make it easier than ever for airspace users (AUs) to request trajectory amendments. But the current FAA automation infrastructure for managing flight plans is not designed to handle that scale of exchanges. This paper describes an effort to develop capabilities to support automated trajectory negotiation, leveraging new data exchange architectures and increasingly autonomous capabilities. Our approach supports trajectory negotiation that incorporates airspace user and FAA constraints and preferences, yet takes place outside the FAA's ecosystem of legacy automated capabilities that are difficult to upgrade, allowing rapid deployment and scaling to achieve automated trajectory negotiation for an increasing number and variety of AUs.

Currently, AUs (flight dispatchers or third-party flight plan providers) file flight plans with the En Route Automation Modernization (ERAM) system pre-departure, using proprietary flight planning and filing systems that exchange data with FAA systems via legacy point-to-point interfaces. Reroutes are sometimes required due to changes in constraints, and current capabilities require dispatchers to submit flight plan amendments more than 30 minutes before departure. Emerging technologies may allow dispatchers to submit amendments after this time, but anything that affects preparation of the aircraft for departure may in turn affect the time at which the flight is ready to depart. Aside from controller-initiated reroutes, after that time airspace user requests for flight plan amendments currently are either made verbally by the flight crew to the controller (the most frequent approach), or must be coordinated separately—via telephone—between the dispatcher (or air traffic control [ATC] coordinator) and FAA traffic management.

There are several problems with the current reroute process. First, the process is manual and time-consuming, which is especially problematic during periods of significant uncertainty in the NAS when flexibility is most needed. This causes airspace users to limit the number of requests they make to avoid rejection, and they may not get more

---

<sup>1</sup> Principal Analyst and AIAA Member.

<sup>2</sup> Principal Analyst and AIAA Member.

<sup>3</sup> Principal Analyst and AIAA Member.

<sup>4</sup> Aerospace Research Engineer, Crew Systems and Aviation Operations Branch.

detailed reasons for rejection than a controller responding, “Unable traffic.” Second, flight crews and controllers have the narrowest area of responsibility relative to NAS operations—a single aircraft and a single sector, respectively. Requiring flight crews and controllers to negotiate reroutes makes it difficult to ensure that the resulting route amendments are desirable from a wider system perspective. Third, the manual process is not scalable to a TBO environment or to a NAS that successfully accommodates the full range of emerging airspace users. Updates to conform with Flight and Flow-Information for a Collaborative Environment (FF-ICE) [1] will better support airspace user-initiated reroutes (pre- and post-departure), but may not be scalable to the anticipated volume of requests once en route trajectory optimization capabilities are widely adopted by AUs and the initial vision for TBO [2] is achieved.

This volume of requests will only be exacerbated by emerging users such as unmanned aircraft systems (UAS) and urban air mobility (UAM) operations.

This paper presents a cloud-based Automated Trajectory Negotiation Service (ATNS) that operates outside the legacy NAS infrastructure, supporting rapid deployment of the capability and speeding the transition to TBO. As an automated capability, it relieves the human participants from being actively involved in every trajectory proposal and increases the likelihood that AUs will be able to adapt aircraft trajectories as needed to respond to uncertainty. In the near term, the AU will request the reroute from the air navigation service provider (ANSP) via methods such as submitting a Trajectory Option Set (TOS) to account for the current FAA automation infrastructure [3]. However, in the long term, the ATNS is envisioned to be able to negotiate directly with the ANSP on behalf of the AU such as via the Unified Flight Plan Filing (UFPF) capability [1]. The following sections examine scalability and trajectory negotiation processes, and present our prototype reroute evaluation and modification algorithms, ATNS architecture, and a demonstration system.

## II. Scalability – A Motivation for Automated Trajectory Negotiation

Significant research has focused on the problem of deriving efficient post-departure reroutes [4-8]. These decision support tools (DSTs) can provide users with substantial flight cost savings. The majority of the total savings obtained in a flight are provided by numerous reroutes that each offer a small amount of savings, such as 1-3 minutes. Therefore, whereas each proposed trajectory may be only slightly more efficient than a flight’s current flight plan trajectory, several such reroutes can accumulate to provide substantial AU savings. Furthermore, the alternate routes proposed by these tools may compete with other AUs’ reroute requests and/or interact with ANSP plans for managing traffic, and therefore may require AU-ANSP negotiations to come to agreement on a trajectory that achieves the AU’s goals while also meeting ANSP traffic management constraints. Therefore, the emerging reroute DSTs may pose workload and scalability challenges that the ATNS can help overcome. We conducted an initial analysis, and found that approximately 1,000 reroutes of greater than 2 minutes may be requested on a moderately heavy day. This complements NASA’s finding that 65% of sampled flights in an operational evaluation requested, and received, reroutes averaging approximately one minute of flight time savings [9]. This trajectory negotiation process is likely to become even more pressing with the entry of non-traditional, emergent users and operations into the national airspace system. In the early stages of integration, these new entrants may impose relatively low densities, but this likely will be followed by much higher traffic loading, leading to high density traffic scenarios, particularly at lower altitudes and in urban environments by the time a full TBO transition has been realized.

Therefore, it will be important for the trajectory negotiation process to be scalable to levels beyond the capabilities of human controllers to handle. Whereas the process will remain human-driven (i.e., AU and ANSP personnel specify high-level constraints, preferences, costs, etc.), the process cannot be human-centric (i.e., humans directly evaluating every proposed time-saving reroute relative to AU and/or ANSP objectives, constraints, and other requested trajectories). An automated negotiation process can support the rapid exchange of detailed data and resolution of complex cost functions that will allow consideration of multiple, potentially simultaneous, reroute requests.

This leads to some scalability requirements for the ATNS. The ATNS must be robust to multiple AUs simultaneously involved in their own negotiations while the underlying NAS conditions are time-varying. Therefore, the ATM-imposed constraints may vary during the negotiation process. Anything that slows the negotiation process increases the likelihood that the NAS will change before the clearance is issued and the reroute is initiated. This could be addressed by using technologies that promote speed of data exchanges (high bandwidth, high speed, and high reliability), beyond traditional data communications channels. Clearly, the system integrity and stability must be safeguarded. For example, negotiated changes to a flight should be recorded as flight plan modifications which should be filed. Also, the iterative trajectory negotiation process likely will need to avoid unbounded iterations in favor of a fixed, deterministic, process. A practical implementation approach would be to implement a timeout (e.g., 2 minutes) for trajectory proposals whose negotiation has not been completed.

Furthermore, a sequential approach could be used in which flight plan modifications, for different flights, cannot be simultaneously filed. Such a first-come, first-serve process could ensure that the system has not evolved during the negotiation process. Such an approach, however, would require an extremely fast negotiation process, and may preclude the desired scalability. Alternatively, a more sophisticated approach could allow simultaneous filings which are decoupled. That is, simultaneous filings could be allowed for negotiated trajectories which are found to be independent in the sense that they do not impose constraints on each other. Such decoupling could be based on geography and/or time. Geographic independence could be defined as no shared sectors or Air Route Traffic Control Center (ARTCC) airspace in the current or requested routes (avoid conflicting sector congestion estimates). Time independence could be achieved when traversal of shared sectors or ARTCCs in the current or requested routes do not overlap in time. If decoupling is impossible, the ATNS may need to revert to a first come first served (FCFS) protocol. Future research could consider more advanced approaches to arbitrating dependencies, e.g., ranked solutions relative to individual AU preferences or sequential scheduling [10].

The requirement for scalability has implications for ATNS architecture and design. Architecture and design requirements also are derived from use cases associated with trajectory negotiation and data exchange, as discussed in the next section.

### III. Trajectory Negotiation Processes and Use Cases

This section summarizes the use cases explored and implemented in the ATNS prototype. This helps to motivate later discussion of functional, algorithmic, and user interaction requirements. Fig. 1 illustrates the following use case sequence:

- i. Aircraft progresses along its original flight plan
- ii. Airspace user proposes a reroute for the ATNS to evaluate
- iii. ATNS determines that the reroute is not operationally acceptable due to reported turbulence (or some other airspace constraint) along the AU-proposed route
- iv. ATNS identifies several common routes near the AU-requested route that avoid known constraints (the common routes approach supports operational acceptability and is discussed in Section IV.B.)
- v. ATNS selects a common route that is nearly as time-efficient as the AU-proposed route and provides sufficient separation from known constraints
- vi. ATNS provides the recommended alternate route to the AU, along with the reason the AU-proposed route was rejected
- vii. AU concurs with the ATNS-proposed route
- viii. AU requests the reroute from the ANSP and the route is amended

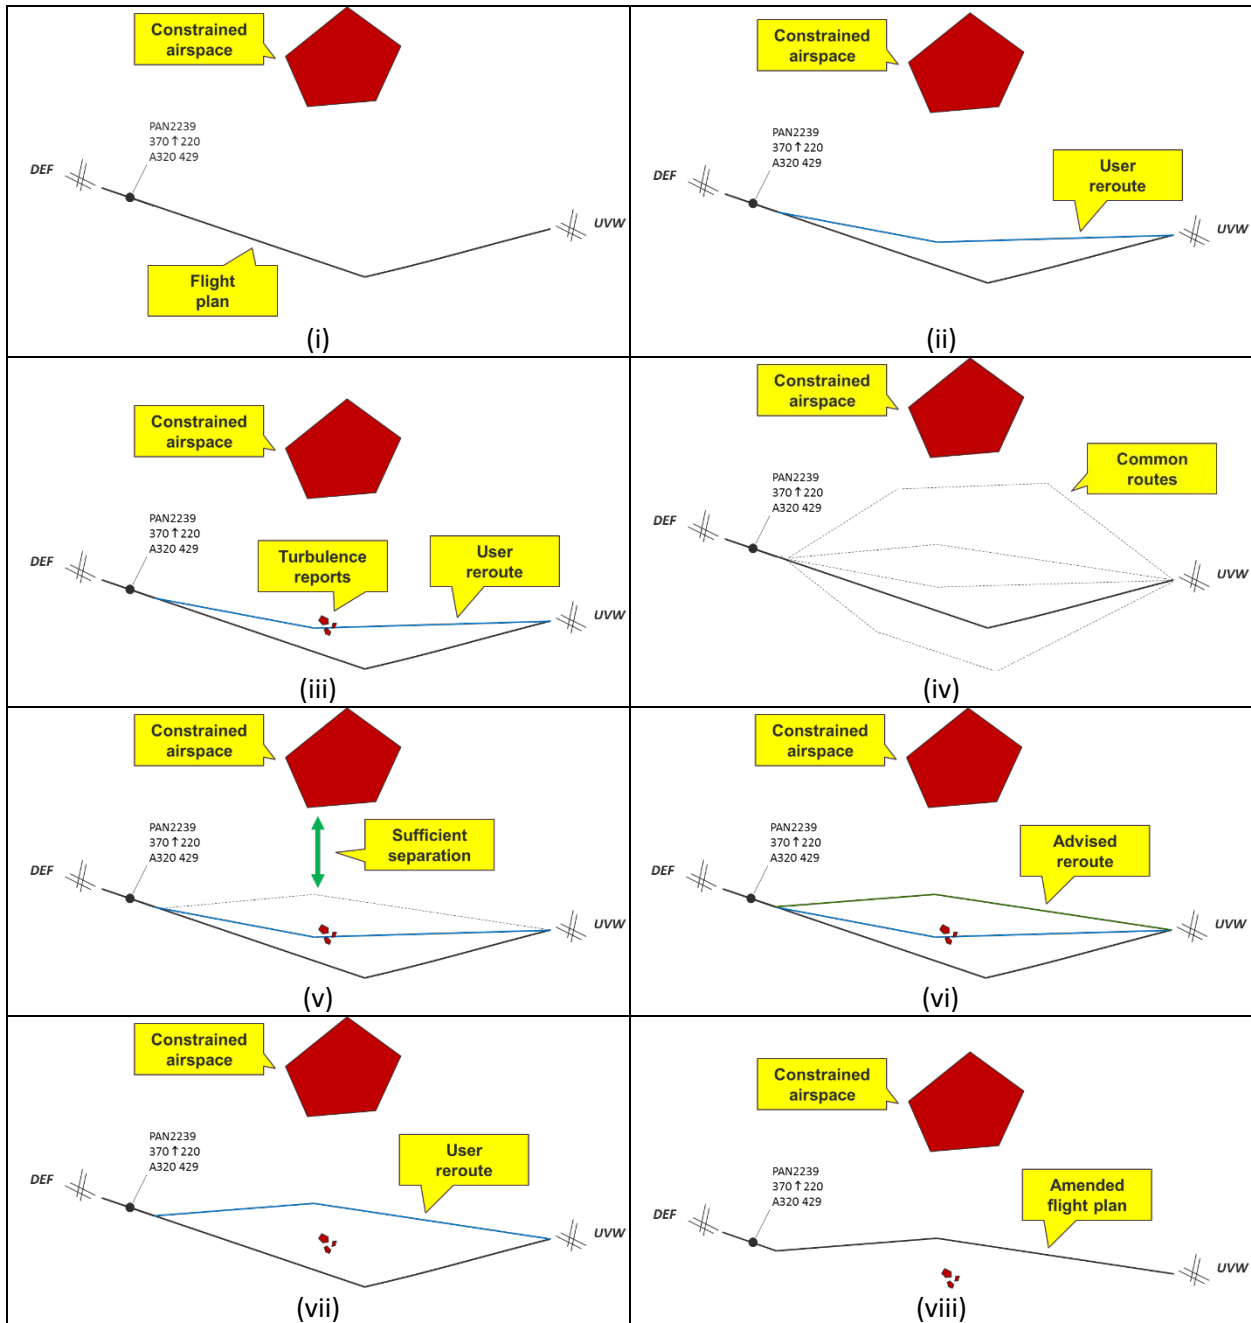
This simple sequence supports evaluation of questions such as the following:

- What are the required data flows between the AU, ATNS, and the ANSP?
- What is the full list of reasons the ANSP and/or the AU would reject a proposed trajectory?
- What are appropriate algorithms for evaluating a proposed trajectory and proposing an alternate when the initially proposed trajectory is rejected?
- What is an appropriate language for communicating the reasons for rejection as well as AU and ANSP constraints and preferences?

The required data flows are discussed in the following subsection, and the remaining questions are considered key areas of ongoing research.

#### A. ATNS Data Flows

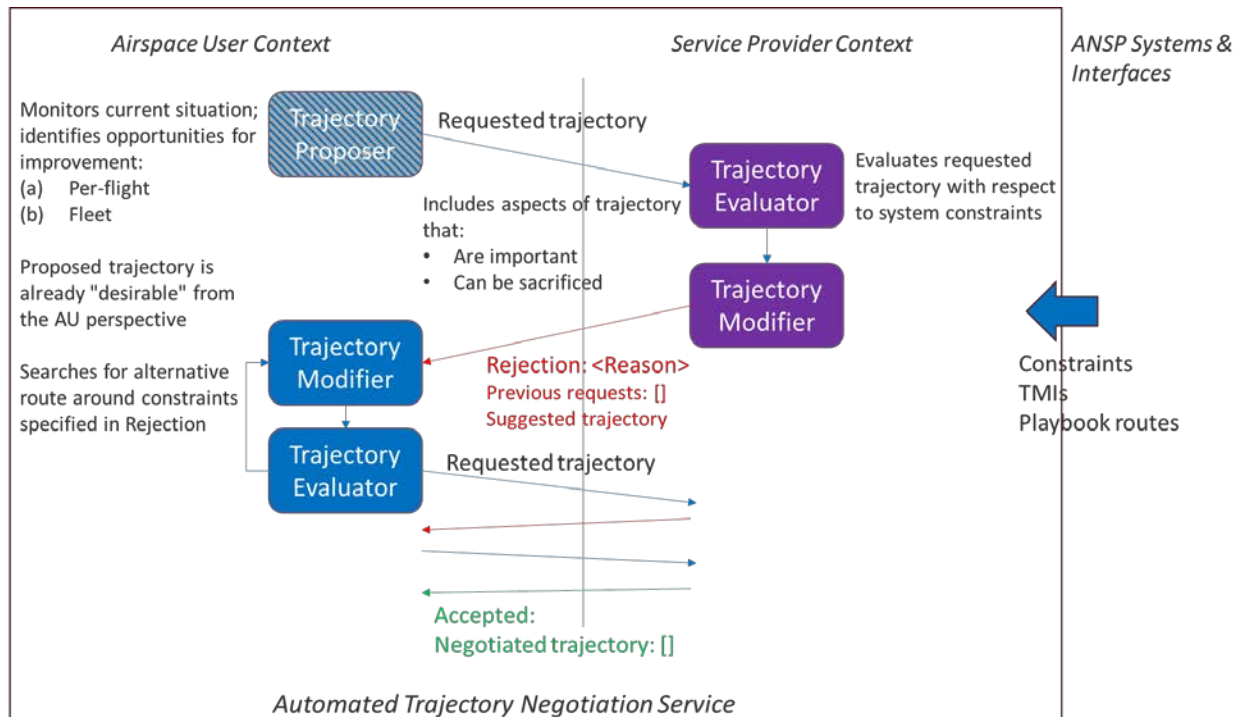
High-level data flows for the use case illustrated in Fig. 1 are shown in Fig. 2, starting at the upper left of the diagram and flowing downward. The airspace user provides a trajectory to the ATNS for evaluation, likely proposed by a trajectory optimization decision support tool (DST) [3,5-9], represented in Fig. 2 by the Trajectory Proposer function. Note that such a function could also be incorporated into the ATNS. Its position as either within or external to the ATNS is denoted by the striping pattern in the figure. The trajectory request includes information about aspects of the trajectory that are important (hard constraints) and aspects that the AU is willing to sacrifice if necessary for the trajectory to be operationally acceptable (soft constraints).



**Fig. 1 Use Case Sequence Implemented to Demonstrate Technical Feasibility**

The Trajectory Evaluator function, shown in purple as a function within the Service Provider instance of the ATNS, evaluates the requested trajectory with respect to known system constraints and other factors affecting ANSP acceptability of the proposed route. Currently, the ATNS receives these constraints from the ANSP, shown in the blue arrow on the left of the diagram, via System Wide Information Management (SWIM) but in the future could be more tightly integrated with ANSP automation systems if the ANSP chooses to participate. If the Trajectory Evaluator function determines that the proposed trajectory is not likely to be operationally acceptable to the ANSP, the Trajectory Modifier function identifies an alternative that satisfies all known constraints and best meets the AU's preferences. In the current demonstration environment, discussed in Section VI, the modified route is that which most closely achieves (or exceeds) the time savings associated with the AU's initially proposed route while also meeting known ANSP constraints. The Trajectory Modifier function then notifies the AU that the AU-proposed route is likely to be rejected if requested from the ANSP, provides a reason for the likely rejection, and also recommends the alternate route. The

AU instance of the Trajectory Modifier and Trajectory Evaluator functions within the ATNS, shown in blue to distinguish it from the Service Provider instance, then evaluates this proposed alternative to determine whether it meets all constraints held internally to that AU's instance of the ATNS. The functions also evaluate whether there is a more preferable trajectory relative to those constraints and preferences held internally. If not, the AU instance of the ATNS requests a new alternative trajectory for the Service Provider context of the ATNS to evaluate, continuing to iterate until either a trajectory is agreed upon by both the AU and Service Provider contexts or the negotiation is terminated. Once a trajectory has been found that is acceptable to both instances, the AU would request it from the ANSP (not shown in Fig. 2 because in the near term this must occur outside the ATNS). In the end state vision, the ATNS could, when appropriate, request the reroute from the ANSP on behalf of the AU. The ANSP retains the final authority in whether the proposed trajectory is acceptable; the ATNS attempts to replicate as many of the operational considerations as possible to maximize ANSP acceptability.



**Fig. 2 High level data flows for the ATNS**

This high level data flow and functional architecture support design of algorithms to support the Trajectory Evaluator and Trajectory Modifier functions. In particular, the algorithms need to consider reasons why a given trajectory proposal may be rejected by the AU or the ANSP. Some reasons are discussed in the next section.

#### IV. Algorithms for Trajectory Evaluation and Modification

The heart of the ATNS is the capability to evaluate proposed trajectories against known constraints (user-provided constraints as well as those available via SWIM data), other aircraft trajectories, and participant preferences. Acceptable trajectories satisfy all hard constraints and as many of the soft constraints as possible, and several different approaches to this problem have been proposed [5,11-16].


Based on the functional requirements, we identified initial algorithmic approaches. For example, we determined that the ATNS is most appropriately applied to proposed trajectory amendments with a maneuver start point beyond a tactical time horizon (10-15 minutes). This is consistent with current thinking on the FF-ICE project [1] as well as results from previous FAA trajectory negotiation concept development [17]. It ensures that the negotiation automation is not approving trajectory modifications that may affect the current sector controller's planning horizon and preserves the sector controller's authority over the aircraft and airspace for which he/she is responsible.

A side benefit of this approach is that it eliminates the need for trajectory-trajectory conflict detection within the ATNS. Beyond 10-15 minutes, it is difficult to predict trajectories with enough certainty to detect trajectory conflicts

with reasonable reliability. Thus, rather than conflict detection, our ATNS prototype uses a forecast of sector congestion as the key determinant of operational acceptability of a trajectory. The sector congestion forecast is based on the likely exceedance of expected sector capacity, with both demand and capacity represented as distributions rather than point values [15]. A sector permeability estimate (0.0 – 1.0, for completely impassable to not at all affected by convective weather, respectively) is used as a means to adjust the sector Monitor Alert Parameter (MAP) value to account for weather [16]. The next section describes how the Sector Congestion Cost (SCC) is calculated.

### A. Sector Congestion Cost Calculation

Fig. 3 shows a notional calculation by the SCC calculator as implemented in the ATNS. It combines probabilistic, time-dependent predictions of sector demand and capacity to produce predicted sector congestion over time.



**Fig. 3 Overview of Sector Congestion Cost Calculation**

Our SCC formulation may be thought of as the Sector Overloading Risk, a time-varying quantity we refer to as  $O_i(t)$ . We calculate it for each sector and each time interval. It is based on the formulation in [15-16], and provides a measure of the overlap of demand probability mass exceeding capacity probability mass:

$$O_i(t) = \sum_{d=0}^{|D_i^\omega(t)|} \sum_{c=0}^d (d - c) P_{C_i^\omega(t)}(c) P_{D_i^\omega(t)}(d) \quad (1)$$

In Equation 1,  $i \in I$  represents the sector in question,  $t \in \tau$  represents the time interval, and the notation  $P_A(a)$  is used to indicate the probability of random variable  $A$  taking on value  $a$ .  $O_i(t)$  thus represents the expectation of the amount by which demand for sector  $i$  in time interval  $t$  will exceed capacity, ignoring negative values. High values of  $O_i(t)$  indicate a high risk of sector overload, although they do not necessarily imply expected demand exceeding expected capacity, since demand and capacity may be highly uncertain. The greater the probability and degree of capacity-demand imbalance, the higher will be the SCC produced for the sector in the relevant time interval.

In the ATNS, a proposed trajectory that causes sector congestion to exceed a threshold during the time the aircraft will traverse the sector is rejected. In reality, Traffic Management Coordinators would not automatically reject a proposed trajectory that caused estimated sector demand to exceed the MAP value by one aircraft for a short time period. Therefore, we plan to explore algorithmic approaches to account more carefully for such operational considerations.

When the Trajectory Evaluator deems a proposed trajectory unacceptable due to its sector congestion cost, the Trajectory Proposer function proposes an alternative trajectory that is operationally acceptable to the ANSP while also meeting the AU's constraints and preferences. Our initial approach is to use a search algorithm constrained by the Clearable Routes Network (CRN) database [12] (discussed next), to identify paths around constraining sectors.

### B. Trajectory Proposal Function – The Clearable Routes Network

The CRN contains a network representation of actual flight clearances issued over a given period of time in the NAS. Each link in the network has a set of properties, including aircraft type, overall frequency of use (including daily and seasonal), and climb/descent/cruise information. Each property contains a count which is accumulated over all aircraft with that property using the same link. In this way we construct a network representation of where and how aircraft are allowed to fly in the NAS.

Most scheduled operators have developed over time a list of historically operationally acceptable routings for each origin-destination pair. They typically evaluate each route from their pre-designed list and select one to file.

Unscheduled operators, who collectively fly many hundreds of thousands of origin-destination pairs each year, are often guided by lists of previously filed routes maintained by their flight planning provider. Neither approach provides an adequate number of alternate routes that can be used to circumnavigate the day’s bad weather or congestion efficiently.

In order to maintain route acceptability and route flexibility and choice, a larger number of acceptable routes are needed. Our CRN is an empirically-derived database of commonly-used, familiar routes. The CRN provides a large set of feasible reroute options for rerouting tools to search over [12].

We construct the CRN using a record of historically accepted and flown routes. From historical track data, we process the route clearances to determine the routes which were cleared and flown. We also capture the frequency with which each segment of each such route was assigned in a historical route network. Rather than being based on the origin-destination pair as are current flight operator solutions, our approach mines historical clearance data to determine the acceptable routings in each part of the airspace.


A challenge in implementing this approach is that for many clearances, we do not want to capture the entire clearance because many clearances made by upstream controllers are unacceptable to, and subsequently amended by, downstream controllers. Thus, we parse and piece together only the parts of clearances that were flown, and record those in the frequency information in the historical route network.

During this process, we ignore clearances that serve only to shortcut the ATC-preferred route with direct clearances past the next few fixes. Although these direct clearances were cleared and flown, they are typically issued on a time, workload and traffic-permitting basis and cannot be counted upon to be acceptable.

The CRN does include both historically filed and amended routes. Step one is to read historically filed, and amended, routes. Step two is to add the finally cleared route segments to the network. By following these steps, the final cleared route for each flight is stored in the database. Importantly, this eliminates cleared routes that were never actually flown, or cleared routes that were accepted and immediately amended.

Where segments of a flight’s route do not exist in the network, additional vertices and links are created as necessary. Where segments do exist, the usage count of those segments is incremented. The result is a route network representing the frequency with which clearances are granted and flown throughout the airspace for different types of flights at different altitudes. The method tracks counts of each aircraft category (e.g., jet, prop, or turboprop) using each link in the route network and the altitudes—climbs, descents, or cruise—at which the link was flown. During route generation, the tool can exclude links with inappropriate aircraft types, links too seldom used, etc.

Our software loads airway definitions from the National Flight Data Center (NFDC) and interprets a large historical sample of filings, clearances, and amendments. The algorithms merge route clearances and amendments together to construct the historical route network. Fig. 4 illustrates how the CRN is used for automated trajectory negotiation.



**Fig. 4 Use of the CRN to improve acceptability of a computed reroute.**

In Fig. 4, the AU seeks a flight plan amendment, as a computed route has been found to be more efficient than the flight plan route. But the computed route may not be acceptable. Instead, a search of the CRN finds a nearby common route—almost as efficient as the computed route, but more likely to be accepted. The CRN may also be used to find commonly used flight segments, which can improve the acceptability of a reroute.

Fig. 5 illustrates this use case. In Fig. 5, the flight plan route from airport DEF to UVW will transit a constrained airspace. Therefore, a reroute is desired. The GHIJK-LMNOP leg is commonly used, and can be used as part of the reroute.

The SCC and CRN algorithms support the Trajectory Evaluator and Trajectory Modifier functions of the ATNS shown in Fig. 2. In the next section, we discuss how these functions are situated in a cloud-based architecture.

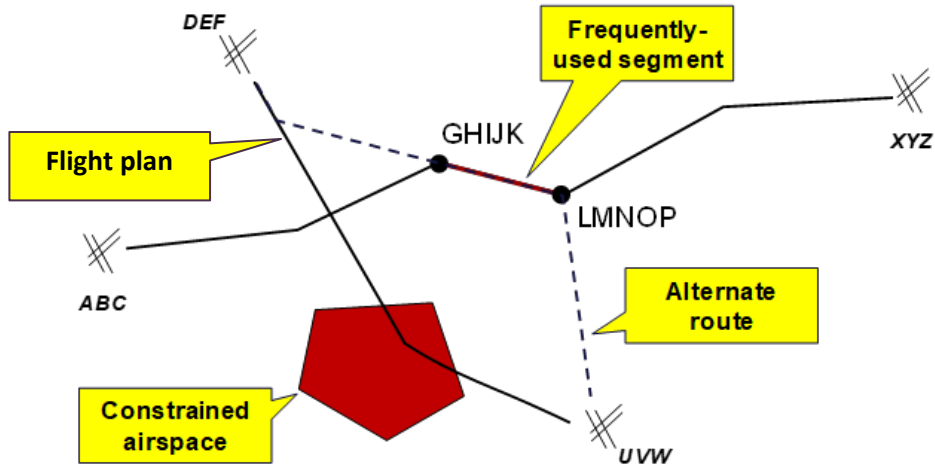


Fig. 5 Notional Use of the CRN to identify a frequently-used segment.

## V. System Architecture for Automated Trajectory Negotiation Service

A key aspect of implementing a trajectory negotiation service is defining its architecture. Fig. 6 shows a high-level version of our cloud-based architecture for deploying the ATNS prototype. The AU plugin implements trajectory evaluation and modification services that are focused on concerns specific to AU preferences and constraints. The ATNS includes at least one AU plugin for each AU accessing the Service. Similarly, the ANSP plugin implements trajectory evaluation and modification services focused on concerns and constraints specific to the ANSP. Implementing the ATNS as a collection of plugins specific to each user ensures that proprietary AU data is protected. It is expected that each user will have different constraints and preferences for determining acceptability and desirability of a given reroute that will require different parameters governing the Trajectory Evaluator and Trajectory Proposer functions, if not different algorithms altogether.

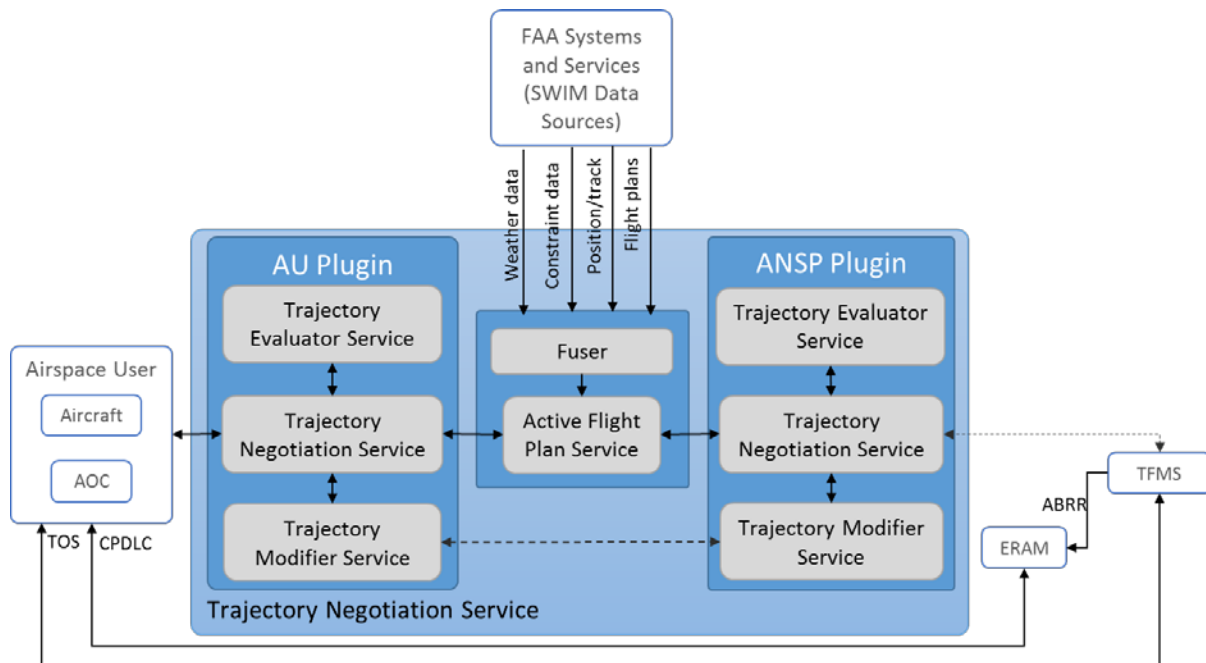


Fig. 6 High-level ATNS architecture.

The Fig. 6 architecture suggests the need for an application programming interface (API) that allows the ATNS to exchange proposed trajectories and AU constraints and preferences with an external source. On the ANSP side, constraints are expected to come from several sources, such as those listed in Table 1.

**Table 1. Summary of ANSP constraints that the ATNS should consider.**

Data source	Purpose
Airspace adaptation	Sector-level loading and capacity modeling, routes.
Convective weather	Nowcast and forecast of heavy storms.
Turbulence	Nowcast and forecast of CAT and CIT.
Weather translation server	Nowcast and forecast of weather-impacted capacity.
Flight plans and traffic	Nowcast and forecast of loading; tactical conflicts.
SAAAs and real-time usage	Nowcast and forecast of airspace closures.
Clearable routes network	Set of candidate route legs to search over.
Traffic Management Initiatives	Forecast of routing delays due to ATC.

On the AU side, preferences may be defined in terms of flight time, flight distance, and fuel burn. These are in addition to constraints and considerations such as landing weight restrictions, arrival time preferences, gate availability, and intra-fleet sequencing. To the extent the AU chooses to share such data with the ATNS, the ATNS will use them to evaluate proposed trajectories and propose alternate trajectories when the Trajectory Evaluator determines that the requested trajectory is operationally unacceptable.

In the near term, the AU will need to request the agreed upon route from the ANSP using voice (pilot to controller) or by submitting a TOS (dispatcher/AOC to ATCSCC/TMU). However, the end state vision includes the ability for the ATNS to propose the negotiated reroute to the ANSP.

## VI. Demonstration System

To evaluate the feasibility of the ATNS concept, we implemented an initial prototype that leverages the SCC and CRN algorithms described above. Initially our focus is on iteratively creating a prototype that implements the functional requirements and the selected algorithms, demonstrates the functional flows of the system, and supports the feasibility evaluation discussed below.

The ATNS prototype is implemented as a representational state transfer (REST) service that provides a REST endpoint for evaluating a user reroute request and generating suggested alternatives when necessary to allow the negotiated route to avoid modeled constraints. To support evaluation and demonstration, the ATNS prototype was implemented in a demonstration environment that included apps that provide supporting services, shown in Fig. 7, including a Reroute Candidate Finder App, a Scenario Driver App, and a Demo Web App. The Reroute Candidate Finder App serves as a proxy for a real-time reroute DST. The Reroute Candidate Finder application processes a specified historical route snapshot (either contained in a local file or via a connection to an external application) to identify potential reroute opportunities using the Trajectory Proposal Heuristic. Upon completion, this application outputs a scenario file to a user-specified location, comprised of a unique identifier, a time span, and the list of user reroute requests found. This scenario file serves as the input for the Scenario Driver App.

The Scenario Driver application parses previously identified user reroute requests from the local scenario file and sends the simulated user requests at random intervals over a web interface to the Demo Web App to emulate a real-time user and a “live feed.”

The Demo Web App is a simple user interface that supports demonstration of basic ATNS functionality. Given user reroute requests (obtained from local file or from the Scenario Driver App), the Demo Web App allows a user to visualize the current flight plan and proposed reroute, submit reroutes to the ATNS for evaluation, and visualize the response of the ATNS—either accepted or modified.

The Negotiation Service backend can be executed in a live mode or playback mode. The negotiated route, negotiation status, and other relevant data are provided to the User System.

In addition to demonstrating the feasibility of the automated trajectory negotiation concept, implementing an initial set of algorithms to carry out negotiation, and providing an environment for exploring functional and data exchange requirements, the ATNS prototype and demonstration environment were used to gain an initial insight into the ability for the proposed architecture to scale to the speed necessary to support automated trajectory negotiation. Chiefly, we performed an initial analysis of the response times associated with individual reroute requests, discussed next.

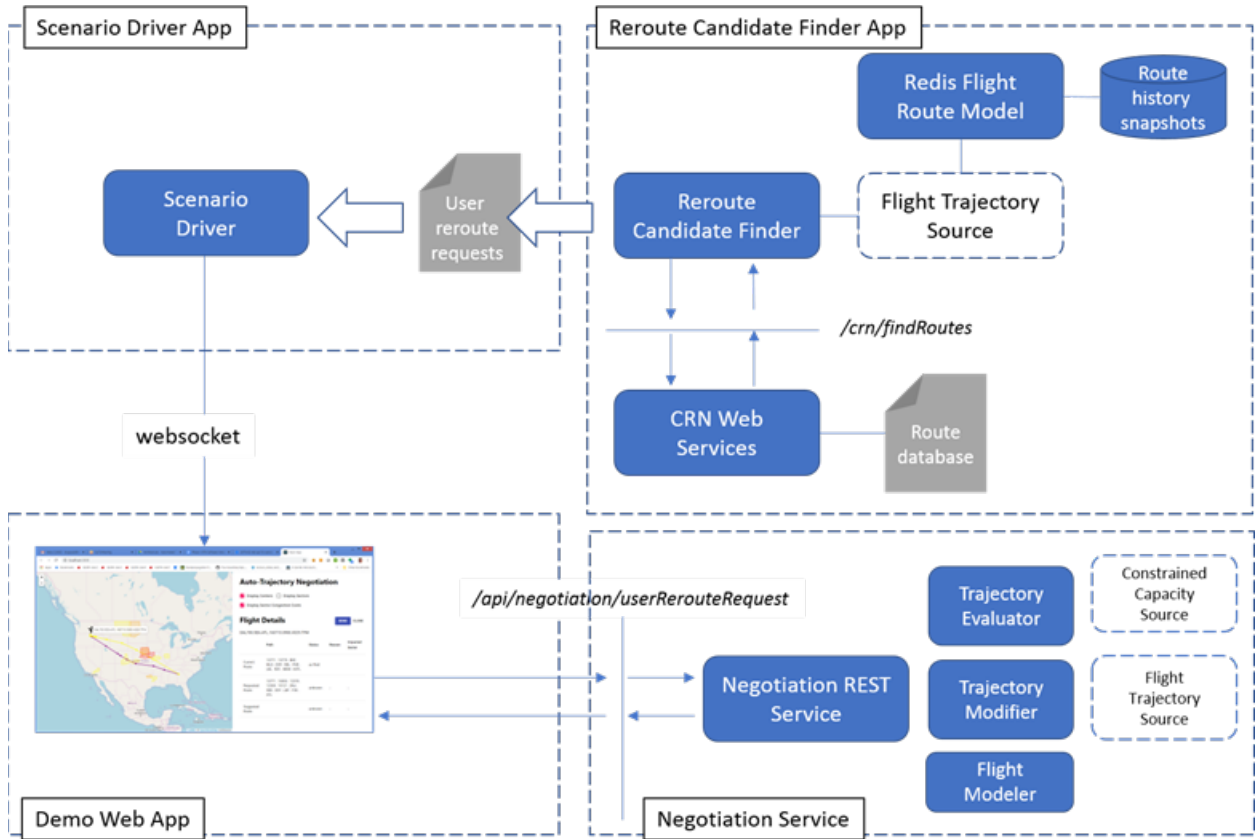


Fig. 7 Applications that are deployed as part of the ATNS and how they relate.


### A. Response Time Analysis

In order to gain an initial understanding of the factors influencing the response time of our prototype ATNS, we sent a series of reroute requests for a specific scenario to the ATNS and measured the CPU time associated with key functional areas that comprise the processing of a given reroute request. In this experiment, we are not attempting to measure the ability to scale across multiple, simultaneous requests. Rather, we focus on the end-to-end processing of a single request at a time.


#### 1. Response Time Analysis Results

Fig. 8 shows the processing time associated with each of the 96 requests in the scenario. Each horizontal line corresponds to a given request, and the different colors within a given line reflect different processing phases. Note that not all requests require all the processing steps. Specifically, only a small fraction of requests requires a search for an alternate trajectory, since this is only required when the original request is deemed to be unacceptable by the ATNS. Also, such a small amount of the processing is associated with loading the set of flights currently operating in the NAS (blue) that no blue lines are visible in the graph.

The response time analysis was run on a laptop with an Intel i7 2.7GHz processor with 4 cores and 16GB of RAM. There is considerable variation in the end-to-end processing time across the requests. The response time is largely dominated by the black portion of the line, corresponding to the computations associated with evaluating sector congestion cost. This computation scales linearly with the number of sectors crossed by the requested reroute. It requires the system to identify and estimate the entry and exit times for all flights to the sectors expected to be crossed by the requested trajectory. Fig. 9 shows the distribution and cumulative distribution for SCC computation time, with 90% of the requests requiring two seconds or less of time for the calculation.



**Fig. 8 Measuring CPU time for processing a series of reroute requests.**



**Fig. 9 Empirical distribution and cumulative distribution for CPU time associated with estimating sector congestion costs for the 96 requests in the scenario.**

This provides us a path for further improving performance and scalability as the project progresses. In addition to technical requirements, we also have considered requirements for supporting human users of the ATNS, discussed next.

## VII. Supporting Human Users of the ATNS

This section uses the example reroute request procedures shown in Fig. 1 and Fig. 2 to examine the requirements for human participation in the negotiation process. This allows us to start to lay out a path to requirements for supporting human autonomy teaming [18-21] in automated trajectory negotiation. We use this process to derive factors affecting reroute acceptability to each party and consider the role of automation support in the automated process.

Automated processes, however and especially in the near term, are hardly likely to be fully automated. A collaborative, human-autonomy team is necessary for trajectory negotiation to be truly fast, efficient, and reliable.

Thus, in addition to respecting the workload problems outlined above and following best practices to support human autonomy teaming (e.g., transparency [21]; effective allocation of authority and autonomy [19]; and engendering appropriate calibration of trust [20]), the following requirements are proposed that are specific to this context.

- Distribution of responsibilities between humans and automated systems must be changeable based on context to support varying workload.
- Constraints and preferences should be expressed in a manner that is straightforward for the human users and fits into their workflows.
- Support human-automation and human-human coordination and communication of constraints and preferences.
- The human operator should be able to direct the Service to explore the solution space for new trajectory options.
- Provide sufficient information when human review/approval of a trajectory request is required.
- Allow humans to maintain an understanding of flights' routes, negotiation status, performance of existing trajectories relative to the solution space and user preferences, even when human evaluation of trajectory requests is not required.

Workload is expected to be a significant issue for personnel in determining whether to consider a proposed reroute. To be operationally acceptable, proposed trajectories must fit into the plans, objectives, and workflows of dispatchers, pilots, controllers, and traffic managers [22]. Plans and objectives may or may not be made explicit through any published data. For example, dispatchers typically have dozens of applications open and running on their monitors, and receive hundreds of messages that require attention throughout the day [23]. They need to make a decision and reach closure on messages and alerts, and typically within 15-30 seconds. Key information elements include current and proposed routing, the estimated time and/or fuel savings of the reroute, the route string, and relevant routing constraints such as unavailable airspace, convective weather, turbulence, nearby traffic, and upcoming congested sectors. In addition, dispatchers must consider the proposed trajectory in the context of the airline fleet.

Workload is also an important factor in whether the air traffic controller accepts the reroute. Other factors here include the volume and complexity of neighboring traffic flows, whether the route would take the aircraft outside of sector boundaries, whether the reroute is consistent with the current routing plan, and airspace closures and hazardous weather. Although hazardous weather is, technically, a problem for the flight crew to reckon with, controllers also pay attention. This is because they are aware of, and seek to avoid, a scenario in which a flight crew requests an aggressive reroute sending them into a region with hazardous weather, and after receiving approval and approaching the weather, requests high-priority assistance in deviating due to the safety threat. Such a scenario is disruptive to work flow and traffic flow. Therefore, controllers are more likely to grant the request if they have assurances that the flight crew is aware of the weather, and really will fly the reroute. A key goal of the ATNS concept is to minimize the burden on these personnel associated with evaluating proposed trajectories and division potential route amendments to improve operations relative to their goals. This naturally has implications for user interaction design to support human-autonomy teaming in this context. These requirements will be used to update the demonstration system in a future phase of work.

## VIII. Conclusions and Next Steps

This paper introduces a concept and demonstration system for automated strategic trajectory negotiation (beyond 15 minutes look-ahead). Enhanced autonomy is needed to support the expected volume of negotiations at the speed necessary to complete the FAA's transition to TBO. We have developed and evaluated initial algorithms and a prototype framework to support automated trajectory negotiation. We also proposed an architecture for an Automated Trajectory Negotiation Service to be further developed in future efforts. Several aspects of the concept and prototype need further evaluation, including the specific parameters and associated algorithms for determining operational acceptability of reroutes and requirements for human evaluation of proposed reroutes.

For example, the CRN described here assumes that the key to ANSP acceptance of a proposed reroute is a series of route segments previously used with some frequency. However, the CRN currently does not guarantee that the route segments are easily recognizable to the ANSP personnel responsible for evaluating whether the proposed reroute is, in fact, acceptable and/or desirable. Additionally, while the SCC is a good measure of whether a proposed reroute will create demand that exceeds capacity, sector demand relative to capacity is not the only measure of operational acceptability and should be fortified with other considerations. Future research will explore additional metrics associated with operational acceptability and associated algorithm formulations.

Future research also will explore appropriate interaction designs for human participants with the ATNS to support:

- 1) expressing constraints and preferences in a manner that is straightforward for the human users and that fits into

their workflows; 2) directing the ATNS to explore the solution space for new trajectory options; and 3) communication from the ATNS of negotiation status and “performance” of existing trajectories relative to the solution space and user preferences.

These efforts will support the FAA in realizing the full vision of TBO, allowing airspace users to negotiate frequent, small trajectory modifications to continually tune flight trajectories to achieve trajectories that are as close as possible to their preferred trajectories.

## Acknowledgments

This work was funded by NASA via contract number 80NSSC19C0370. The authors wish to thank the team at NASA Langley Research Center, particularly Matthew Underwood, who provided significant insights.

## References

- [1] Ngo, T., Zuniga, A., Biel-Goebel, J. J., Miller, M. E., “Technical Evaluation of FF-ICE/R2,” *2019 Integrated Communications Navigation and Surveillance (ICNS) Conference*, Reston, VA, 2019.
- [2] JPDO, “JPDO Trajectory-Based Operations (TBO) Study Team Report,” Next Generation Air Transportation System Joint Planning and Development Office, Washington, DC, 2011.
- [3] Talebi, E., Ging, A., Engelland, S., Capps, A., Eshow, M., Jung, Y., Sharma, S., Downs, M., Freedman, C., Ngo, T., et al., “Airspace Technology Demonstration 2 (ATD-2) Technology Description Document (TDD),” NASA/TM-2019-220355.
- [4] Wing, D., “The TASAR Project: Launching Aviation on an Optimized Route Toward Aircraft Autonomy,” NASA/TP-2019-220432, 2019.
- [5] Ballin, M., Wing, D., “Traffic Aware Strategic Aircrew Requests (TASAR),” *12th AIAA Aviation Technology, Integration, and Operations Conference*, 2012.
- [6] Wing, D., Burke, K. A., Henderson, J., Vivona, R. A., Woodward, J., “Initial Implementation and Operational Use of TASAR in Alaska Airlines Flight Operations,” *Aviation Technology, Integration, and Operations Conference*, AIAA, Atlanta, 2018.
- [7] Sheth, K., McNally, D., Somersall, P., Morando, A., Clymer, A., Shih, F. T., “Assessment of a National Airspace System Airborne Rerouting Tool,” *Eleventh USA/Europe Air Traffic Management Research and Development Seminar (ATM2015)*, 2015.
- [8] McNally, D., Sheth, K., Gong, C., Sahlman, S., Hinton, S., Lee, C., Sterenchuk, M., Shih, F. T., “Dynamic Weather Routes: Two Years of Operational Testing at American Airlines,” *Eleventh USA/Europe Air Traffic Management Research and Development Seminar (ATM2015)*, 2015.
- [9] Henderson, J., Wing, D. J., Ballard, K. M., “Alaska Airlines TASAR Operational Evaluation: Achieved Benefits,” NASA/TM-2019-220400, 2019.
- [10] Sadovsky, A., Windhorst, R., “A Scheduling Algorithm Compatible with a Distributed Management of Arrivals in the National Airspace System,” *38th Digital Avionics Systems Conference (DASC)*, AIAA, San Diego, CA, 2019.
- [11] Evans, A., Lee, P., “Using Machine-Learning to Dynamically Generate Operationally Acceptable Strategic Reroute Options,” *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)*, 2019.
- [12] Hall, W., Hunter, G., “Trajectory Optimization and the Clearable Route Network,” *18th AIAA Aviation Technology, Integration, and Operations Conference*, June, 2018.
- [13] Evans, A., Sridhar, B., McNally, D., “Improving Operational Acceptability of Dynamic Weather Routes Through Analysis of Commonly Used Routings,” *16th AIAA Aviation Technology, Integration, and Operations Conference*, 2016.
- [14] Taylor, C., Wanke, C., “Dynamically Generating Operationally Acceptable Route Alternatives Using Simulated Annealing,” *Air Traffic Control Quarterly*, Vol. 20, No. 1, 2012, pp. 97-121.
- [15] Hunter, G., “Interactions Between Single-Center Dynamic Resectorization and NextGen TFM,” *28th Digital Avionics Systems Conference (DASC)*, October, 2009.
- [16] Kaler, C., Hall, W., Brinton, C., Capozzi, B., Hunter, G., “Collaborative Trajectory Options Program Within the NAS Flow Advisory Manager,” *2018 Aviation Technology, Integration, and Operations Conference*, June, 2018.
- [17] FAA, “NextGen Trajectory Negotiation Concept of Operations,” Federal Aviation Administration, 2014.
- [18] Mosier, K., L., Fischer, U., Burian, B. K., Kochan, J. A., “Autonomous, Context-Sensitive, Task Management Systems and Decision Support Tools I: Human-Autonomy Teaming Fundamentals and State of the Art,” NASA/TM-2017-219565, 2017.
- [19] Smith, P. J., “Making Brittle Technologies Useful,” *Cognitive Systems Engineering: The Future for a Changing World*, edited by P. J. Smith and R. R. Hoffman, CRC Press, Boca Raton, FL, 2018, pp. 181-208.
- [20] Lyons, J. B., Koltai, K. S., Ho, N. T., Johnson, W. B., Smith, D. E., Shively, R. J., “Engineering Trust in Complex Automated Systems,” *Ergonomics in Design*, Jan. 2016, pp. 13-17.

- [21] Springer, A., Whittaker, S., “Progressive Disclosure: Empirically Motivated Approaches to Designing Effective Transparency,” *Proceedings of the 24th International Conference on Intelligent User Interfaces (IUI 2019)*, Marina Del Rey, CA, 2019.
- [22] Grudin, J., “Why CSCW applications fail: Problems in the design and evaluation of organizational interfaces,” *Proceedings of the 1988 ACM Conference on Computer-Supported Cooperative Work*, ACM, 1988.
- [23] Munro, P., “A Field Study of Airline Dispatcher Work Practices and Workload Issues,” NASA/TM-2018-219924, 2018.