

Improved Dynamic Generation of Operationally Acceptable Reroutes Using Network Optimization

Christine Taylor* and Craig Wanke†
MITRE Corporation, McLean, Virginia 22102

DOI: 10.2514/1.52957

This paper develops a methodology for generating operationally acceptable reroutes in the presence of convective weather. A decision support system that defines flight-specific reroutes would aid traffic managers in efficiently maneuvering flights, but only if the reroutes provided were acceptable alternatives. As such, this research describes a methodology that captures, through the network definition and metric evaluations, properties of reroutes that are both flexible and operationally acceptable. By employing network optimization techniques, reroutes are developed that can estimate and avoid dynamic weather constraints and do so by optimizing a set of operational acceptability metrics. This paper first illustrates how prioritizing weather avoidance as compared with other cost factors in network generation alters the reroutes generated, which can aid in parameter setting for automation. The reroutes generated through network optimization are then compared with traditional reroute alternatives by evaluating each set of reroutes against a set of operationally acceptable metrics. The overall approach of the proposed methodology is reviewed and details are provided on the modeling formulation and solution algorithm employed by the network optimization method. Multiple flights with weather-impacted routes are examined to show the effectiveness of the proposed methodology in producing reroutes with better operational acceptability.

I. Introduction

MANAGING congestion, especially during weather events, requires improved methods for assisting decision makers and increasing operational efficiency. Hazardous weather requires traffic managers to reroute flights that plan to pass through the weather, while balancing demand through sectors with reduced capacity or increased traffic volumes (resulting from other flights deviated from their original routes). Today's methods for rerouting traffic are mostly manual: air traffic managers employ their expertise to handle a single flight, or entire flows of traffic are rerouted using National Playbook routes [1]. These methods have been historically employed due to the complexity of defining an operationally acceptable route in real time; however, the reroute alternatives provided are limited. As the need to maximize all available airspace capacity is imperative, it is necessary to widen the set of operationally acceptable reroutes provided to decision makers or decision support systems.

Research in decision support systems for both the strategic [2] and tactical [3] time frames models the impact on congestion due to weather and investigates how and when to act to maintain safe airspace throughput. When reroute options are limited to traditional route alternatives such as coded departure routes (CDRs), National Playbook plays, and historically flown reroutes, and dynamic rerouting is not employed, the solutions derived provide a limited set of resolution options as compared with real tactical operations.

Including the capability to dynamically generate reroutes provides a larger solution space, but the additional computation expense can be significant. Even in a nonreal-time decision support system, such as NASA's Airspace Concept Evaluation System [4] the dynamic rerouting performed must be simplified for computation considerations. The Constrained Airspace Rerouting Planner discussed in [4] considers trajectories and weather constraints evolving in time

and generates the shortest deviation reroute around the weather. However, considering only distance during the evaluation may produce a reroute that does not satisfy other traffic management concerns such as sector or traffic flow coordination.

Alternatively, [5] develops a fast-time simulation that provides reroutes derived from a grid network overlaying the original flow. A reroute is defined as the shortest path through the network that avoids areas restricted by weather; however again the operational acceptability of these proposed reroutes remains in question. Using a fully populated grid network is a common modeling technique for rerouting as it is fast to implement and solve. In [6], a similar modeling technique is used to generate a reroute for a general aviation (GA) aircraft in free flight. The assumption of free flight works well for flights operating visual flight rules, but does not necessarily extend for flights operating instrument flight rules and under air traffic control (ATC).

Reference [7] considers both a grid network and a waypoint-based network when defining reroutes for aircraft flows approaching the arrival terminal. The nodes in a waypoint-based network consist of existing waypoints, or slight permutations created from original waypoints to enable additional options. The reroutes are constructed using the flow-based route planner [8] with constraints governing interactions with other flows and therefore accurately simulate the restrictions of the transition approach environment. A similar approach was used in [9] where the development of weather-specific CDRs for predeparture flights was investigated. Using the flow-based route planner, new CDRs were defined that better matched current weather predictions, but were still limited by the restrictive constraints placed on the network reroutes to ensure operational acceptability.

As is seen from previous research into this area, one of the major difficulties in dynamic rerouting is capturing conventional guidelines for quality route development while recognizing that rerouting around a weather event may require maneuvers that deviate from these guidelines. Specifying end-to-end route options provides acceptable, but limited reroute alternatives for a given flight. Alternatively, defining a grid network produces a larger number of potentially operationally unacceptable reroutes. Any decision support system developed for dynamic rerouting must bridge this divide by defining operationally acceptable reroutes in real time to effectively aid decision makers in managing air traffic.

Terrestrial rerouting research [10] inspires a compromise approach, as the network is derived from existing roadways and therefore a reroute that uses these roadways is operationally feasible.

Presented as Paper 2010-8077 at the AIAA Guidance, Navigation, and Control Conference, Toronto, Ontario, 2–5 August 2010; received 2 November 2010; revision received 1 March 2011; accepted for publication 1 March 2011. Copyright © 2011 by The MITRE Corporation. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/11 and \$10.00 in correspondence with the CCC.

*Senior Simulation Modeling Engineer, M/S 450. Member AIAA.

†Senior Principal Simulation Modeling Engineer, M/S 450. Senior Member AIAA.

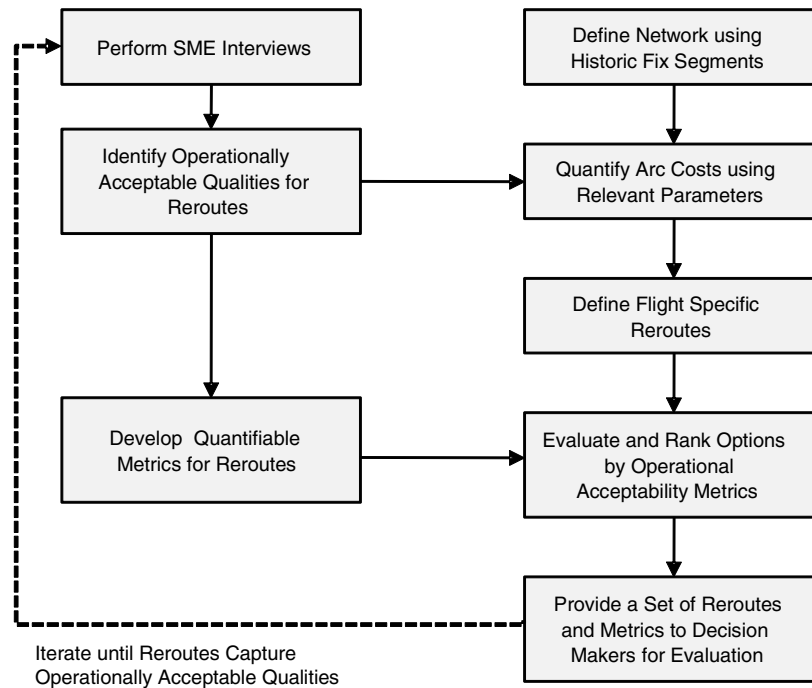


Fig. 1 Process diagram for dynamic generation of operationally acceptable reroutes.

Extending this methodology to air traffic rerouting would imply that a previously flown segment can be considered as a component of a feasible reroute, even if it is not normally used for traffic between the origin and destination pair of the flight.

The research presented in this paper is part of a greater research effort to bridge this divide by defining operationally acceptable route alternatives in real time to effectively aid decision makers in managing air traffic. The proposed approach for generating reroutes is based on the definition of route segments derived from historically flown connections between existing fixes [11] and evaluated using metrics of operational acceptability, derived from analyses based upon subject matter expert advice [12].

The purpose of this work is to define a methodology for optimizing flight-specific reroutes, based on metrics of operational acceptability, and to demonstrate the improvements from a global optimization approach over traditional reroute options. Section II provides an overview of the greater research initiative that this work is a part of and discusses the specific definitions of operational acceptability considered. Section III provides a description of the problem formulation and models developed. The algorithm and metrics used to generate and evaluate reroutes are presented in Sec. IV. Section V first presents an example rerouting problem and analyzes the impact of prioritization of weather avoidance on solution quality. Given this analysis, Sec. V compares the reroute alternatives generated using the network optimization approach presented in this work to traditional routing options using the metrics of operational acceptability. Section VI presents the conclusions drawn from this research and the ongoing work in this area.

II. Defining Operationally Acceptable Reroutes

The goal of this research is to define dynamic, flight-specific reroutes for predeparture or enroute flights that provide operationally acceptable alternatives to decision makers. Defining operational acceptability, however, can be difficult as it requires extracting the essence of quality route design, as understood by subject matter experts[‡] (SMEs) into quantifiable and generic evaluation metrics or constraints. Issues of route geometry, the interaction of the route with

current sector geometries and with other traffic makes this problem difficult to tackle.

This research implements a dual-track approach for defining operationally acceptable reroutes, as shown in Fig. 1. The extraction of SME knowledge into metrics defining operational acceptability is described on the left side of Fig. 1. The right side of Fig. 1 shows the process of constructing reroutes, beginning with the network model that is developed using only historic fix-pair segments to produce reroutes. The set of reroutes generated are further evaluated by the operational acceptability metrics developed, and the best of these reroute alternatives are provided to decision makers. A set of reroute options, as opposed to a single best reroute, are returned as inevitably there will be air traffic requirements specific to a given situation that cannot be captured by the general operational acceptability metrics. Finally, iterations with SMEs will be performed during the prototyping stage of the research to ensure that the operational acceptability metrics defined produce useful reroute alternatives to decision makers.

The collaboration with SMEs which formed the basis of the operational acceptability metrics considered in this research was conducted as part of the larger research initiative and is described in greater detail in [12]. This paper focuses on the methodology associated with generating the reroute alternatives that capture the operational acceptability metrics defined. Through the iterative interaction with SMEs shown in Fig. 1, modifications to the set of operational acceptability metrics are envisioned; however, the methodology for generating the reroutes is defined in a flexible manner to incorporate such changes. The remainder of this section provides an overview of the current operational acceptability metrics considered in this analysis.

A. Route Distance

The most frequently considered metric of reroute quality is the distance of the reroute as compared with the original route. As both the original route and the reroute alternatives are defined by a list of fixes along the route, the distance of a route is defined as the sum of the distances between each consecutive pair of fixes. The distance between a fix-pair can be interpreted as the air distance, assuming both the wind and aircraft velocity profiles are provided. However, for the purposes of this research, we assume a simple great circle path between consecutive fix-pairs in the route, and therefore define the

[‡]Five subject matter experts were queried for this analysis and consisted of former traffic managers, air traffic controllers and airline dispatchers.

reroute distance as the sum of the great circle fix-pair distances between consecutive fixes.

B. Origin-Destination Flow Factor

The flow conformance of a route is a measure of how consistent the route is to historical routing. Although all fix-pairs defined in the network were historically flown, these fix-pairs may not be historically used by flights traveling between a particular pair of regions. To compute this conformance measure, each fix-pair in the network is assigned an origin-destination (O-D) flow factor. The flow factor [12] is computed as follows.

Approximately 4000 airports were grouped into 35 geographically distinct regions, and for each of these region pairs, the historical usage of each fix-pair segment was analyzed. The O-D flow factor assigned to each fix-pair segment is not a count of usage, but a relative comparison of usage ranging between high usage (O-D flow factor approaching zero) and almost no usage (O-D flow factor approaching one). All fix-pair segments not used between a region pair are assigned an O-D flow factor of 1. The region pair for a given flight is determined by the flight's departure and arrival airports, which in turn, determines the O-D flow factors of the fix-pair segments.

C. Route Blockage

Route blockage defines the forecasted usability of each fix-pair in the reroute to quantify how likely a given reroute is to be weather impacted. As weather is a dynamic event, fix-pair segments may be intermittently available during the period of interest. By capturing the dynamic nature of route blockage, efficient reroutes around weather-impacted areas can be generated.

The blockage of a fix-pair segment is determined using the methodology defined in [13]. An overview of this methodology is as follows. A grid of 1 km by 1 km cells is overlaid on a fix-pair segment and the blockage of each grid cell is forecasted using the convective weather avoidance model (CWAM) [14]. CWAM determines the flight altitude below which pilots would likely deviate around the cell, based on precipitation and echo top forecasts from the Corridor Integrated Weather System. The "minimum traversable altitude" is computed for each grid cell, for 15-min time bins over a look-ahead time (LAT) of two hours.

For each LAT, the minimum clear altitude for each cross section of the grid is computed, where it is assumed that 10 km of contiguous cells are needed to denote a passable cross section. Figure 2 illustrates this concept. The clear altitude for the fix-pair is simply the maximum clear altitude for each cross section. A fix-pair is considered blocked if the flight altitude is less than the clear altitude for the segment during the LAT the flight will traverse the fix-pair; otherwise the fix-pair is considered open and weather-free.

The authors note that while the use of a two-hour forecast may limit the time frames for which this methodology can be used to reroute flights around weather events, the weather prediction uncertainties present at longer LATs does not enable accurate predictions of fix-pair availability. If better forecasts were available, the methodology described above would enable predictions at longer LATs. Until such forecasts are available, flights travelling greater distances could benefit from the proposed reroute generation method when closer to the weather event. Alternatively, more aggregated forecasts,




Fig. 2 Depiction of route blockage computation.

where larger areas of airspace are deemed weather-blocked, could also be used to evaluate potential reroutes. This approach would be similar to defining the surrounding area of a predicted weather event as impassable to all flights, as was previously developed in [15].

D. Lateral Deviation

The lateral deviation of a reroute is a measure of how differently the reroute will use the airspace, as compared with the original route, and thus is a surrogate for coordination complexity. To define lateral deviation L^p mathematically, we begin by computing the maximum lateral distance d_L , or cross-track, between the two routes. This distance is then scaled using the linear regression provided in Eq. (1)

$$L^p = \min(1, 1.26 - 0.033d_L + 0.027(d_L - 25)(d_L > 25) + 0.048(d_L - 70)(d_L > 70) + 0.00095(d_L - 180)(d_L > 180)) \quad (1)$$

Equation [1], which was derived using a correlation analysis described in [12], defines a zero-one scale metric where a score approaching zero is assigned to reroutes with small lateral deviations and a score approaching one is assigned to reroutes with large lateral deviations.

E. Global Flow Conformance

When large severe weather systems occur, greater deviation from the original route may be required to define a good reroute that avoids the weather. In these cases, the O-D flow conformance may not provide a useful measure of the route's flow conformance, as many of the possibly useful fix-pairs may never have been flown for the O-D pair of interest. In these cases, we still wish to differentiate between reroutes, preferring to choose fix-pairs that are more commonly used than others.

As every fix-pair in the network has been historically flown in some context, there is an associated usage for the fix-pair. It is desirable to consider how predominant the usage is on a particular fix-pair, relative to all other fix-pairs. As such, we define the global flow factor of a fix-pair to be the percentage of usage on that fix-pair as compared with all usage on the fix-pairs in the network.

III. Problem Formulation and Network Model Development

The dynamic rerouting problem considered in this research begins with the identification of a flight either predeparture or enroute to the destination airport that must deviate from its route because of a weather event. Dynamic reroutes are generated as paths through a network that is constructed for the specific flight under consideration. To promote the generation of operationally acceptable reroutes, the network is derived from a database of historically flown fix-pair segments [11]. As such, the nodes of the network comprise the set of fixes identified in this historical fix-pair segment database and the connections between these nodes, or directed arcs, are taken from the connections that have been previously flown. This section explains in greater detail the mathematical equations and modeling assumptions required to formulate this problem.

A. Rerouting Problem

The flight requiring a reroute is identified in the model by its route string,[§] which can be translated to a set of fixes from the departure airport to the destination airport. As the flight can be anywhere between the departure airport and the arrival airport when a reroute is initiated, the deviation point and the rejoin point of the original route must be specified. The deviation point is the fix along the original route where the reroute can begin, which is any fix including the departure airport (when the flight is predeparture), that occurs before the weather intersects the original route. Similarly the rejoin point is

[§]The route string of a flight refers to the origin airport, a series of airways and fixes, and the destination airport.

the final fix along the original route where the reroute reconnects to the original route and can be any fix along the original route after the weather event, including the destination airport.

B. Defining the Flight-Specific Network

The fix segment database is defined for the entire National Airspace System. However, for a given flight, only a subset of these segments is useful in defining the network. To scope the search area appropriately without unnecessarily constraining it, an ellipse was chosen as the search space. This definition provides a meaningful area for deviation from the original route while limiting the search area to reduce the computational expense associated with generating and solving the network. The elliptical search area is defined as follows.

Given the route string of a flight, the deviation point, and the rejoin point, a unique network is defined as an ellipse containing the allowable set of nodes and arcs. The parameters defining the ellipse, namely the origin, semimajor axis a and semiminor axis b are defined as follows.

Given the deviation point of the route R^d and the rejoin point of the route R^r , additional search buffers are appended to the great circle connecting these points to ensure that all reasonable connections are identified. The buffer length at the deviation point b_d is defined to be 100 nm if the deviation point is the first point on the original route, corresponding to the departure airport; otherwise the buffer is 25 nm. This distinction recognizes that if the deviation point is the departure airport, a greater search area that encompasses multiple standard instrument departures is necessary to evaluate alternative routes; however if the deviation point is a fix enroute, the search area for feasible connections, especially opposite the direction of travel, is more limited. The buffer length at the rejoin point b_r is 100 nm to provide a large area encompassing possible reconnections to the original route or destination airport.

Figure 3 depicts the construction of the ellipse. The origin of the ellipse is defined as the midpoint of the line connecting the great circle formed between the buffered deviation point and the buffered rejoin point. Buffered points simply refer to the extension of the great circle between the deviation point and rejoin point by the buffer distance b_d or b_r , respectively. The semimajor axis of the ellipse a is the distance between the origin of the ellipse and either buffer point. The semiminor axis b is defined as the maximum value of either half the semimajor axis or 100 nm greater than the maximum lateral

distance, or maximum cross-track distance, of the original route from the great circle connecting the origin and destination airports.

C. Defining Arc Costs

For each arc defined in the network, the arc cost represents the value of selecting that particular arc over another when constructing the set of best paths through the network. As the cost of the arc must be defined as an inherent property of the arc itself and independent of the path, it is desirable to identify factors that provide a measure of the quality of selection. Thus, for the arc costs, we choose three of the five metrics of operational acceptability presented in Sec. II, namely distance, O-D flow factor, and route blockage. These three metrics were chosen as they represent important features of an individual arc (as opposed to the entire route) and therefore are desirable to use when generating the set of k -shortest paths.

The distance of an arc is defined as the great circle distance from the starting fix or node to the end fix or node. As stated previously, the distance computation is simplified to be independent of wind, but this is a modeling simplification and not a limitation of the methodology. As such, we define the distance of an arc between node i and node j as

$$d_{i,j} = 2 * \sin^{-1} \sqrt{\left[\sin \left[\frac{(\theta_i - \theta_j)}{2} \right] \right]^2 + \cos(\theta_i) * \cos(\theta_j) * \left[\sin \left[\frac{(\varphi_i - \varphi_j)}{2} \right] \right]^2} \quad (2)$$

where θ_i and φ_i are the latitude and longitude of node i , respectively, and θ_j and φ_j are the latitude and longitude of node j , respectively. The distance cost of an arc from node i to node j is normalized by the total distance of the original route between the deviation and rejoin points d^R , as shown in Eq. (3)

$$c_{i,j}^d = \frac{d_{i,j}}{d^R} \quad (3)$$

The flow factor of an arc represents how often that arc is used when traveling from the cluster containing the departure airport to the cluster containing the arrival airport, as discussed in Sec. II. The flow cost of the arc $c_{i,j}^f$ is defined as the flow factor of the arc from node i to node j between the departure airport cluster and the arrival airport cluster $f_{i,j}^{DA}$ multiplied by the normalized distance of the arc $c_{i,j}^d$, as shown in Eq. (4)




Fig. 3 Description of network search area ellipse.

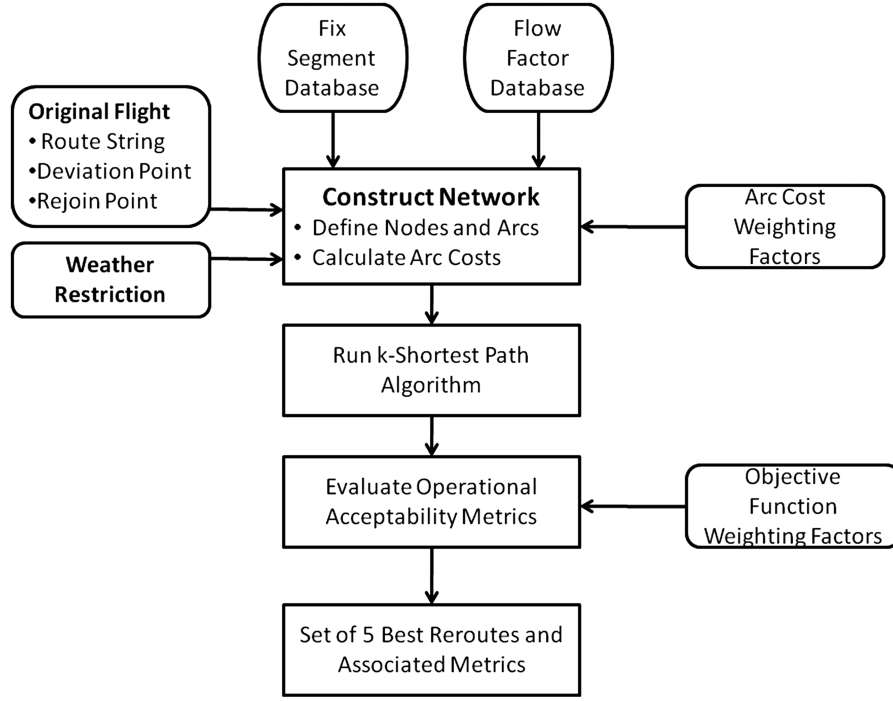


Fig. 4 Flow diagram for computation of operationally acceptable reroutes.

$$c_{i,j}^f = \frac{f_{i,j}^{DA} * d_{i,j}}{d^R} \quad (4)$$

The flow factor is scaled in this manner to emphasize that the longer the arc, the more important it is for the arc to have a low flow factor.

Defining the route blockage of an arc is similar to the methodology described in Sec. II with the exception that, for an individual arc, the current flight altitude and time of arrival to the arc is unknown. As such, we must estimate the blockage of each arc in the network. The route blockage estimate is a consequence of using a static network to define the reroutes, as opposed to a time expanded network [16]; however, given the computational expense associated with generating and solving such networks, and that when the reroute is evaluated for selection, the route blockage is known, we have chosen to employ the following estimation technique. Specifically, we estimate the flight altitude by assuming Eq. (1) the cruise altitude, if the fix-pair is outside the transition radius of either airport, or Eq. (2) zero ft, if within the transition radius. The transition radius is defined as the distance it takes a flight to reach cruise altitude, which is assumed as 3 miles for every 1000 ft of cruise altitude. Alternatively, the transition radius could represent top of climb and top of descent, at the origin and destination airports, respectively, if this information is provided for the flight under consideration. If the flight altitude is above the minimum clear altitude, the segment is unblocked during this time period; otherwise the segment is considered blocked.

Mathematically, we define the minimum cleared altitude of a segment at time t as $h_{i,j}^t$. For each time t , we determine if the segment is considered blocked by comparing the cleared altitude to the flight altitude, where

$$b_{i,j}^t = \begin{cases} 1, & h_{i,j}^t \leq h^f \\ 0, & h_{i,j}^t > h^f \end{cases} \quad (5)$$

where h^f is the flight altitude assumed for the segment.

The blockage cost of the arc, as expressed in Eq. (6), is defined using a weighted average of blockage values from the time bin containing the earliest possible arrival time t^e to a two-hour LAT T

$$c_{i,j}^b = \frac{\sum_{t \geq t^e}^{T-t+1} f_b^t * b_{i,j}^t}{\sum_{t \geq t^e}^{T-t+1} f_b^t} \quad (6)$$

where the blockage weighting factors are defined as $f_b^t = [1, 1, 1, 1, 0.5, 0.5, 0.5, 0.5]$ to provide a greater emphasis on blockage values closer to the earliest arrival time to the arc. The earliest arrival time is simply estimated using the cruise velocity and the distance from the original deviation point of the route to the fix-pair. If the earliest possible arrival time is later than two hours, no blockage information is available and the segment is considered unblocked.

The total cost of an arc from node i to node j is then represented as

$$c_{i,j} = k_d * c_{i,j}^d + k_f * c_{i,j}^f + k_b * c_{i,j}^b \quad (7)$$

where k_d , k_f , and k_b are the weighting factors for arc distance cost, arc flow factor cost, and arc blockage cost, respectively.

IV. Generating Operationally Acceptable Solutions

This research defines a set of reroute alternatives, as opposed to a single option for each flight. Specifically, a set of alternatives that best meet the operational acceptability metrics, defined in Sec. II, are returned. This allows decision makers to evaluate different route alternatives in the context of the larger traffic management situation and determine the best reroute for a flight.

The overall flow of the algorithm is shown in Fig. 4. Specifically, the user supplies the original flight information as well as the weather information required to determine route blockage. The databases are accessed when constructing the network to determine which arcs exist and what is the cost of these arcs, as defined by the arc cost weighting factors. The reroute alternatives are then produced by a k -shortest path algorithm and evaluated, based on the objective function weighting factors, on operational acceptability. The five best reroute options are returned.

In this section, we first describe the generation of the reroute alternatives using a k -shortest path approach. We then discuss how the aspects of reroute operational acceptability presented in Sec. II provide quantitative metrics to measure the overall quality of the route alternatives and inform decision makers.

A. Generating K-Shortest Paths

The generation of a single or multiple paths through a network can be performed by implementing a variety of solution methods. Dijkstra's algorithm [17] is the classic means of defining a path

through the network, as shown in [5]. Dynamic programming, which is a related technique, is used by Ng et al. [18] for path construction. An A* search method is implemented to construct the paths defined by the flow-based route planner, described in [8], and used in [7,9]. Modifications to the A* search that include a heuristic estimation function to speed the shortest path search are employed by [6,19]. Reference [10] uses a heuristic method known as multi-agent systems to search and update path performance in dynamic networks.

In this research we construct our set of reroute options using a k -shortest paths algorithm,[¶] which employs a Bellman–Ford algorithm. We note that although for a single shortest path, Dijkstra's algorithm would be more appropriate as there are no negative costs in the network, the more general k -shortest path problem can take advantage of the process employed by the Bellman–Ford approach. The algorithm approach can be summarized as follows.

Given the source node and sink node, which are defined as the route deviation point R^d and route rejoin point R^r , respectively, we seek the k -shortest paths connecting these nodes. To compute the set of paths, the Bellman–Ford algorithm updates the distance of each node in the network to the destination node, while tracking the predecessor relationship to the source node. These updates are performed at most $|N| - 1$ times, where $|N|$ is the number of nodes in the network. This methodology can be used to define all shortest paths in the network, where the k -shortest paths are returned.

B. Constraining Path Feasibility

Before evaluating the operational acceptability of a path, we need to assure that the path is operationally feasible. The first condition on feasibility is that there are no cycles, or repeated fixes in the path. For a shortest path problem with positive arc costs, this constraint is automatically satisfied; however, as we seek multiple paths, it may be beneficial from a cost perspective to add a subcycle to a path, instead of constructing a different path. As this is not operationally feasible, and simply a consequence of the algorithm employed, we remove all paths that contain a cyclic subpath.

In previous research [15], the analysis showed that without an explicit constraint, paths with large turn angles could be generated as reroute options. As these paths are often not operationally feasible, and are highly undesirable, we added a turn-angle constraint, where paths that do not meet the constraint are removed from the set of paths to be examined and evaluated as reroute options. Specifically, for every arc in the path, we evaluate the change in heading between two subsequent arcs and ensure that the absolute angle change is less than the maximum angle change permitted.

C. Evaluating Operational Acceptability

Once the k -shortest feasible paths are defined, the operational acceptability metrics discussed in Sec. II can evaluate each reroute to determine its quality. To compare the different aspects of reroute quality, a multimetric path objective function is defined that incorporates the five operational acceptability metrics, namely distance, O-D flow factor, route blockage, lateral deviation, and global flow factor.

The distance metric D^p is defined as the accumulated scaled arc distance from the deviation point to the rejoin point, as shown in Eq. (8)

$$D^p = \sum_{(i,j) \in p} c_{i,j}^d \quad (8)$$

Here, the arc distance cost $c_{i,j}^d$ is as defined in Sec. III and p is the reroute defined by the set of arcs. The O-D flow factor metric F^p is the accumulated distance-weighted flow factor $c_{i,j}^f$ defined in Sec. III and is provided in Eq. (9)

$$F^p = \sum_{(i,j) \in p} c_{i,j}^f \quad (9)$$

Unlike the route blockage calculation for the arc costs, we know precisely which arcs are blocked in the reroute, as these can be determined by analyzing the resulting trajectory of the reroute. Thus, route blockage metric B^p is the number of blocked arcs in the reroute, as defined in Eq. (10)

$$B^p = \sum_{(i,j) \in p} b_{i,j} \quad (10)$$

The lateral deviation L^p of the reroute is the fourth metric considered. The computation of the reroute lateral deviation and translation into the zero–one scale is as described in Sec. II and [12]. The global flow conformance G^p is the distance-weighted usage fraction of the arc within the network. If we define the usage of the arc as $u_{i,j}$, where usage is computed as frequency of use for the historical time period evaluated in [11], then the global flow conformance can be defined, as shown in Eq. (11)

$$G^p = \sum_{(i,j) \in p} \left(1 - \frac{u_{i,j}}{u_{\max}} \right) * c_{i,j}^d \quad (11)$$

Here, $c_{i,j}^d$ is as defined in Sec. III and u_{\max} is the maximum usage factor in the fix-pair database.

Combining these five performance metrics into an overall objective function for the reroute operational acceptability yields the expression in Eq. (12)

$$C^p = w_d * D^p + w_f * F^p + w_b * B^p + w_l * L^p + w_g * G^p \quad (12)$$

where w_d , w_f , w_b , w_l , and w_g are the relative weighting factors for reroute distance, weighted average O-D flow factor, route blockage, scaled lateral deviation, and global flow factor, respectively.

V. Results

The methodology developed in this paper aims to capture aspects of operational acceptability in the development of dynamically generated reroutes. To this end, we will first examine how altering the relative priorities of the arc cost components provide different routing options for a flight impacted by weather. By varying the relative weighting of route blockage we will show how different user priorities can be accommodated in reroute selection. Drawing on these results, we then generate reroutes for multiple flights using the methodology described in this paper and compare the performance of the network-optimized reroutes with traditional reroutes using the metrics of operational acceptability. Finally, an initial computation effort assessment is provided to illustrate the ability of this methodology to provide an alternative approach for dynamically generating quality reroutes to aid in decision making.

All examples presented in this paper are derived from reroutes generated for scheduled flights on 20 April 2009 at 20:00 GMT. All network-optimized reroutes are defined using the k -shortest path approach, where 200 paths are generated using the arc cost definition in Eq. (6) and the arc weightings as described in the following section as well as the time-based weights for route blockage provided in Sec. III. From the 200 paths generated, all infeasible paths are culled and the remaining paths are evaluated against the operational acceptability metrics. To provide a strict ranking of reroute options, the operational acceptability metrics are weighted equally and combined into an overall objective function, as expressed in Eq. (12). The five reroutes with the lowest objective values are returned.

Note that the choice of employing a single objective function weighting, where all metrics are weighted equally is to provide a strict comparison between reroutes. If the goal is to provide a variety of routes for use by human decision makers, it may be desirable to make multiple runs with different weights on the objective function factors. This will provide a more diverse set of routes to choose from, in the event that there are other operational factors, unknown to the automation, that constrain the ultimate route choice.

[¶]Additional data available at www.jgraphpt.org [retrieved 2/28/2011].




Fig. 5 Original route for flight A with weather blockage.

A. Weighting Factors for Reroute Generation

The arc cost weighting factors defined in Eq. (6) represent the relative value of distance, O-D flow factor, and route blockage when computing the cost of an arc for reroute generation. Changing the relative value has a direct impact on the set of reroutes generated since this would change the relative costs between the arcs and therefore which arcs define the shortest path (and k -shortest paths). Based on previous research [15], we have found that a useful arc cost weighting ratio corresponds to a 3:1 distance to O-D flow factor prioritization.

In this research, we consider the inclusion of route blockage as an arc cost and thus investigate how different weightings of this metric in the overall arc cost impact the quality of the reroutes. To illustrate this, we examine the rerouting of “Flight A,” which is scheduled to depart from Dulles International Airport (IAD) to Greensboro (GSO) at 20:16 GMT. Figure 5 displays the planned route for the flight (shown as a solid black line), along with the weather at that time. The eastward-moving weather will partially obstruct the original route (shown as the thick black segment between GVE and LYH), necessitating a reroute. As the flight is scheduled for departure shortly,

HAFNR is chosen as the deviation point of the original route and we define the rejoin point as the destination airport (GSO). Table 1 provides the associated ellipse parameters as well as the details of the resulting fix-pair network shown in Fig. 6.

Table 2 provides the three sets of arc cost weighting factors considered, where the only change between the three cases is the increasing importance of route blockage in the arc costs. Using the

Table 1 Network detail for Fig. 6

Network parameters for flight A	Parameter values
Semimajor axis	169.7 nm
Seminor axis	129.8 nm
Origin	Latitude = 36.7, Longitude = -79.7
Rotation angle	44.7 deg
Number of nodes	875
Number of arcs	7580

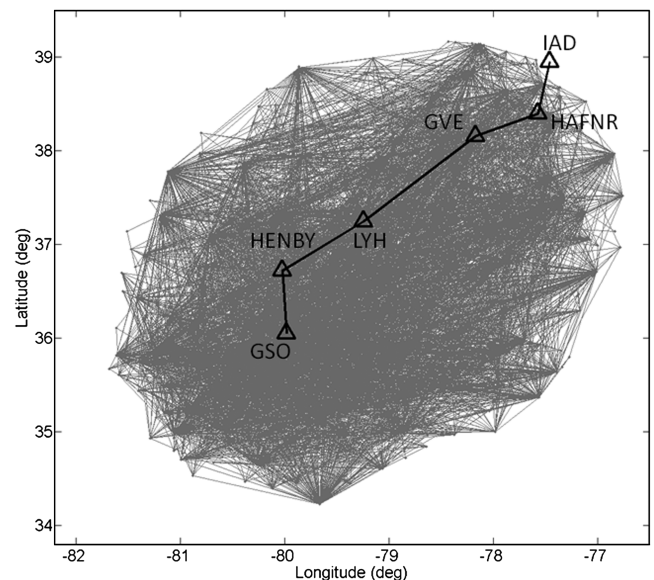


Fig. 6 Network search areas with original route for flight A.

Table 2 Relative weightings on arc cost for three cases

Arc cost parameter	Case 1	Case 2	Case 3
Arc cost distance	3	3	3
Arc cost O-D flow factor	1	1	1
Arc cost route blockage	0.5	1	3

three cases defined in Table 2, we compare the five best reroutes generated when all objective function components are weighted equally and Figs. 7–9 show the reroutes generated for cases 1 through 3, respectively.

Examining Fig. 7, we see that the first network reroute avoids the weather by directly heading from GVE to the destination airport while the remaining four reroutes head further east to DAN before travelling via different pathways to GSO. Specifically, the second and third reroutes differ only slightly in that the second reroute travels from DAN to UFFIN to LEAKS while the third reroute defines a direct connection from DAN to LEAKS. The fourth reroute suggests a direct route from DAN to GSO. The fifth reroute returns to GSO via MAYOS, following the last part of the original route.

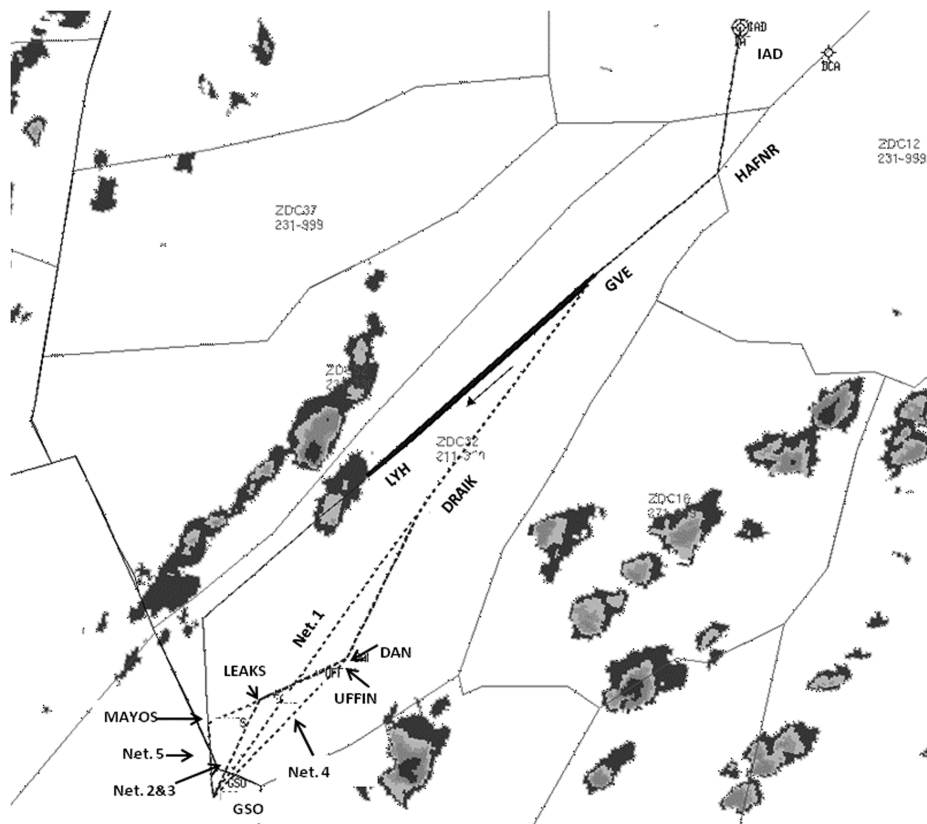
Examining Fig. 8, which corresponds to case 2, we see that the first network reroute from case 1 is also defined as the first network reroute in case 2. The remaining four reroutes in case 2 travel further east than the reroutes in case 1, proceeding from GVE to SBV before heading west to LEAKS then GSO. Here we note that of these four reroutes, only two can be distinguished (they differ by the inclusion of UFFIN in the route string) because of the existence of collinear fix-pairs in the database. Specifically, ARVON, COLAB, and SHEPS all reside on the same path, but multiple combinations of these fix-pairs exist in the database, thus producing separate but identical options. Removing these collinear fixes is desirable as an implementation issue To produce more diverse reroute options; however this does not impact the underlying methodology. Figure 9 shows the reroutes generated for case 3, where again all five reroutes provide effectively the same option, using different combinations of connections of the

collinear fix-pairs ARVON, COLAB, and SHEPS, as well as varying the connection between SBV and LEAKS by including either UFFIN, DAN, or both. The first through fourth of the reroutes from case 3 are identical to the second through fifth reroutes in case 2.

Comparing Figs. 7–9 shows that as the arc weighting factor on route blockage is increased, route options tend to move further east of the predicted weather, reducing the risk of potential route blockage. Although none of the reroutes suggested are blocked by weather, the choice of this weighting parameter is important given the uncertainty in the weather progression. By weighting route blockage higher in the arc cost, we are effectively trading between having more certainty in a nonweather-blocked reroute in the future with a more efficient reroute in terms of distance or conformity. Figure 10 compares the metric values for the reroutes of the three cases and depicts exactly this tradeoff. Examining Fig. 10 shows that the best reroute, in terms of overall objective function value, corresponds to the first network reroute provided in case 1 and case 2. The reroute options unique to case 1 provide lower overall objective function values than the remaining reroutes in case 2 and case 3, however, if the weather should develop differently than expected, these reroutes are at a greater risk of being weather-impacted as they are closer to the weather event. Case 2 produces both the best reroute option that is provided in case 1 as well as the more conservative reroute options provided in case 3. Although these conservative options (network options 2–5 of case 2) have higher objective function values, this arc cost weighting ratio provides a compromise in options between improving operational acceptability and averting risk and as such, the remainder of this paper will use the arc weightings in case 2. We note that, if the uncertainty in the weather forecast were known, a more informed tradeoff between route blockage and the other factors could be determined.

B. Comparison of Reroute Generation Methods

Having evaluated the impact of the route blockage weighting in arc costs, we now compare the results obtained using the network optimization approach to traditional rerouting approaches. In today's

**Fig. 7** Five best network-optimized reroute alternatives for case 1.

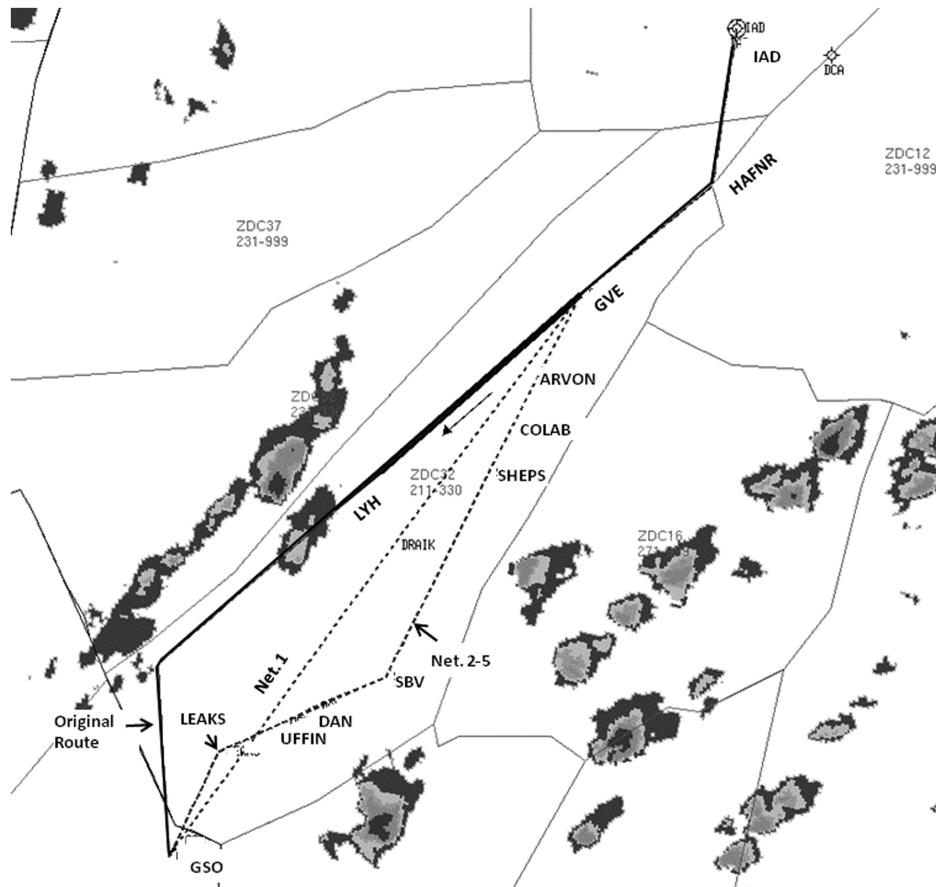


Fig. 8 Five best reroute network-optimized reroute alternatives for case 2.

operations, CDRs and historic reroutes are often used to reroute a flight around weather. Using these readily available reroutes enables traffic managers to reduce the workload associated with defining flight-specific operationally acceptable reroutes as CDRs are precoordinated. However, CDRs do not exist for many origin-destination pairs, and can only be used for predeparture flights. As such, historically used reroutes are often employed to supplement these traditional options as the operational acceptability of these reroutes have been tested in so much as they have been previously used. Reference [11] discusses the parameters used to generate the database of historically flown reroutes and it is this database we will use to compare the network-optimized reroutes to the traditional reroutes.

Using the same weather shown in Fig. 5, two additional flights were identified as requiring a reroute for the same segment blockage. Note that the previous example of Flight A cannot be used for comparison, as no historic reroutes were found for this flight; there is simply too much weather blockage. This in itself provides support for the robustness of the network optimization approach in finding solutions for difficult reroute problems.

The second example is of "Flight B," enroute from Logan International Airport (BOS) to Charlotte Douglas International (CLT). As the original flight path, shown in Figs. 11 and 12 as the solid black line, is blocked by weather, reroutes are generated by deviating at JERSY and rejoining at the destination airport. Figure 11 shows the historic reroute alternative found for this flight and Fig. 12 shows the five best reroute alternatives defined by the network optimization. The reroute alternatives are shown as dashed lines in both figures.

Examining the historic reroute in Fig. 11, we see that the one historic reroute alternative suggestion travels east of the original route and heads directly to CLT, without rejoining the STAR. In comparison, we see that the all network-optimized reroutes deviate east from GVE to SHEPS returning to LYH. The first four network

reroutes proceed along the original route, connecting to the Standard Terminal Arrival Route (STAR)** at MAYOS while the fifth option travels further west to HENBY before reconnecting at MAYOS. The "zig-zag" maneuvers used by the network-optimized reroutes effectively provide delay, permitting weather to pass before retuning on the original route. Furthermore, the fifth reroute provides a more conservative estimate of the weather movement, by performing the second "zig-zag" maneuver to avoid potential additional weather on the original route. Examining the network-optimized reroutes in Fig. 12 we see no distinction between the five reroutes between GVE and LYH as they are composed of the same collinear fix-pairs identified in the Flight A example. Furthermore, although the network-optimized reroutes rejoin the STAR, the final specification of the STAR is not followed as it violates the turn angle constraint. To resolve this issue, the turn-angle constraint can be relaxed or removed in areas surrounding the terminal airport; however, this may produce undesirable reroutes. A more robust approach would be to create and use a database of STARs for each destination airport and permit connections only to these complete reroute segments. However, these alterations remain as continuing work for implementation and do not impact the overall validation of the methodology.

Figure 13 provides the metric value comparison for the two sets of reroutes. Examining Fig. 13 reveals that all reroutes perform approximately the same, including the historic reroute. Of the five network-optimized reroutes, only slight differences arise between the specific fix-pair segments incorporated. The network-optimized reroutes perform between 15–18% better in O-D flow conformance than the historic reroute; however they perform about 10% worse in global flow conformity. This suggests that the network-optimized reroutes

**A Standard Terminal Arrival Route provides a detailed set of waypoints to the airport To precoordinate arrival approaches and have a specific set of entry points or transition fixes from which a flight enters this defined final flight segment.




Fig. 9 Five best network-optimized reroute alternatives for case 3.

travel segments typically used by flights between BOS and CLT while the historic reroute employs segments more frequently used in general.

The final example considered is of “Flight C,” scheduled for departure at 20:08 GMT from IAD and heading to CLT. Given the imminent departure, GVE is chosen as the deviation point of the original route and CLT is defined as the rejoin point. Figures 14 and 15 show the original flight path as a solid black line as well as the

historic reroutes and network-optimized reroutes (as dashed lines), respectively.

Comparing the reroutes shown in Figs. 14 and 15, we see that the network-optimized reroutes rejoin the original STAR at GIZMO, but again exhibit the same performance as before in that they directly connect to the arrival airport to avoid the turn angle constraint. The historic reroutes show a progression of alternatives moving east of the weather and connecting directly from an alternate fix to the airport


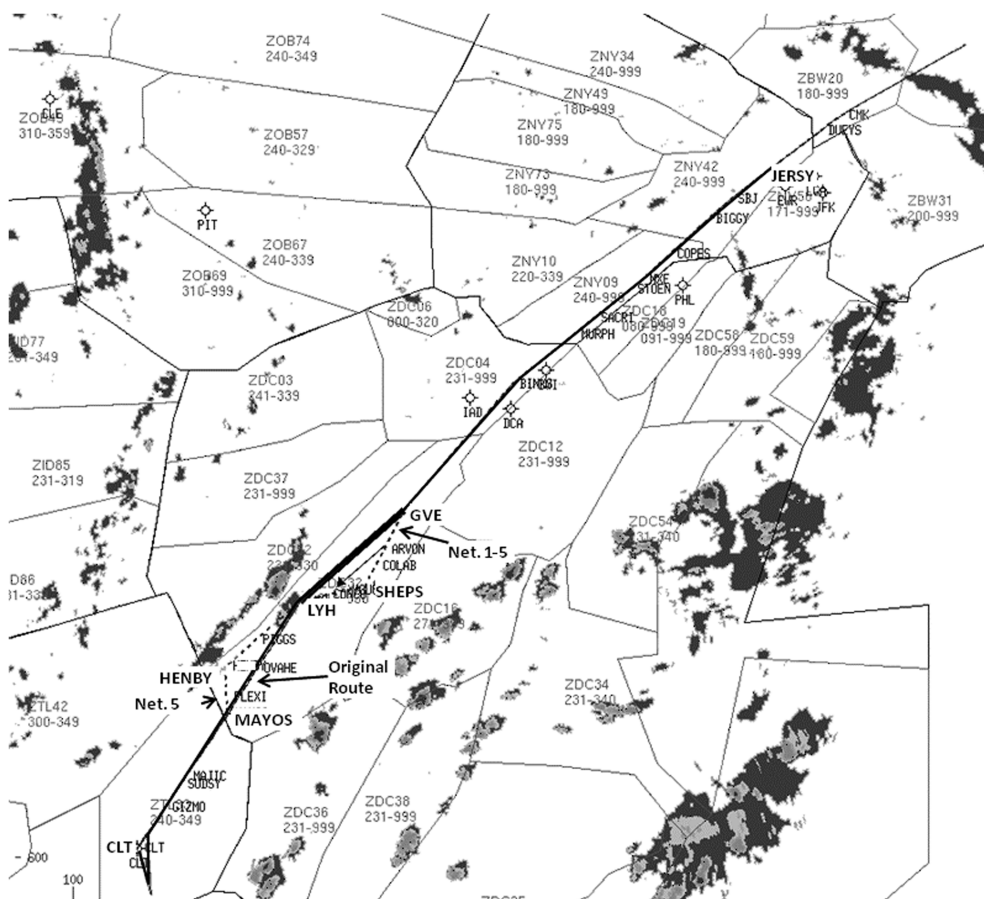
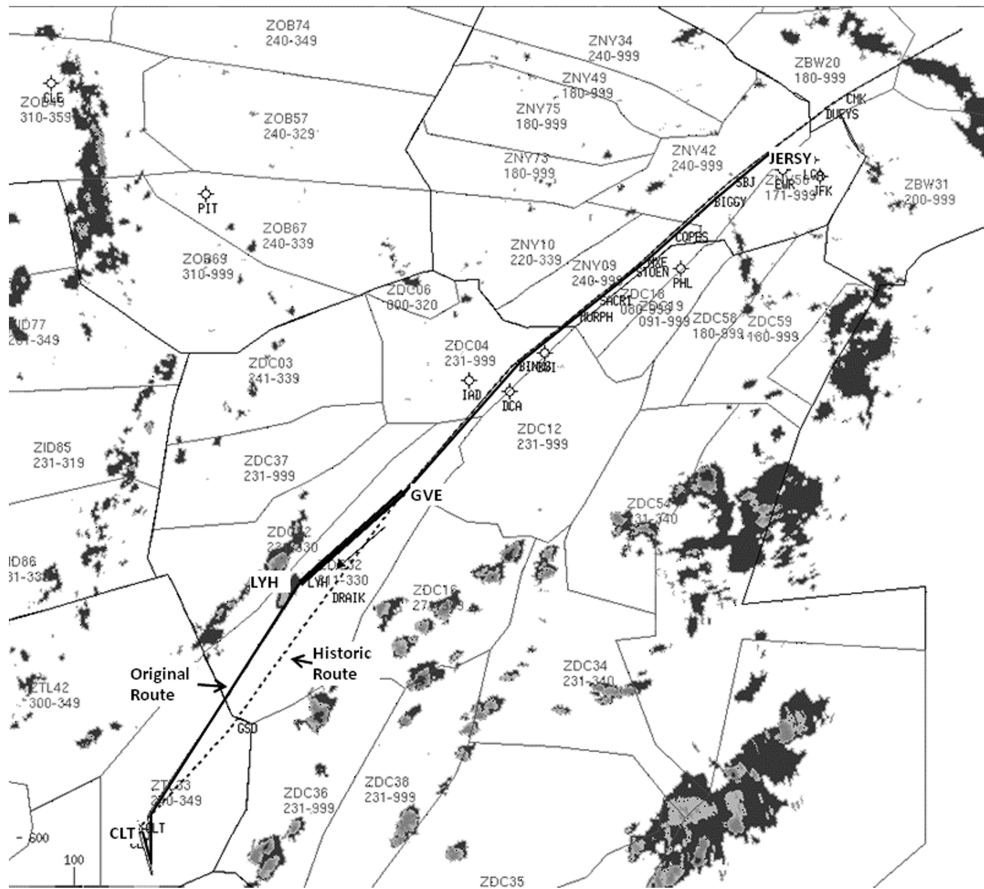


Fig. 10 Comparison of metric values for cases 1-3 reroute options.






Fig. 13 Comparison of network-optimized and historic reroutes for flight B.

without rejoining a STAR. Figure 16 compares the operational acceptability of these reroutes. By examining Fig. 16 we see that although all reroute options avoid the weather blockage, the network optimization reroutes all provide better (lower) objective function values than the historic reroutes. The five network-optimized reroutes all have similar

distance, deviation, O-D and global flow conformity, as can be seen in Fig. 15. Comparing the network-optimized reroutes to the historic reroutes, the historic reroutes perform worse in most categories. Specifically, only the first historic reroute provides a lower metric value on distance than the network-optimized reroutes as it defines a direct connection from the deviation point to the destination airport;

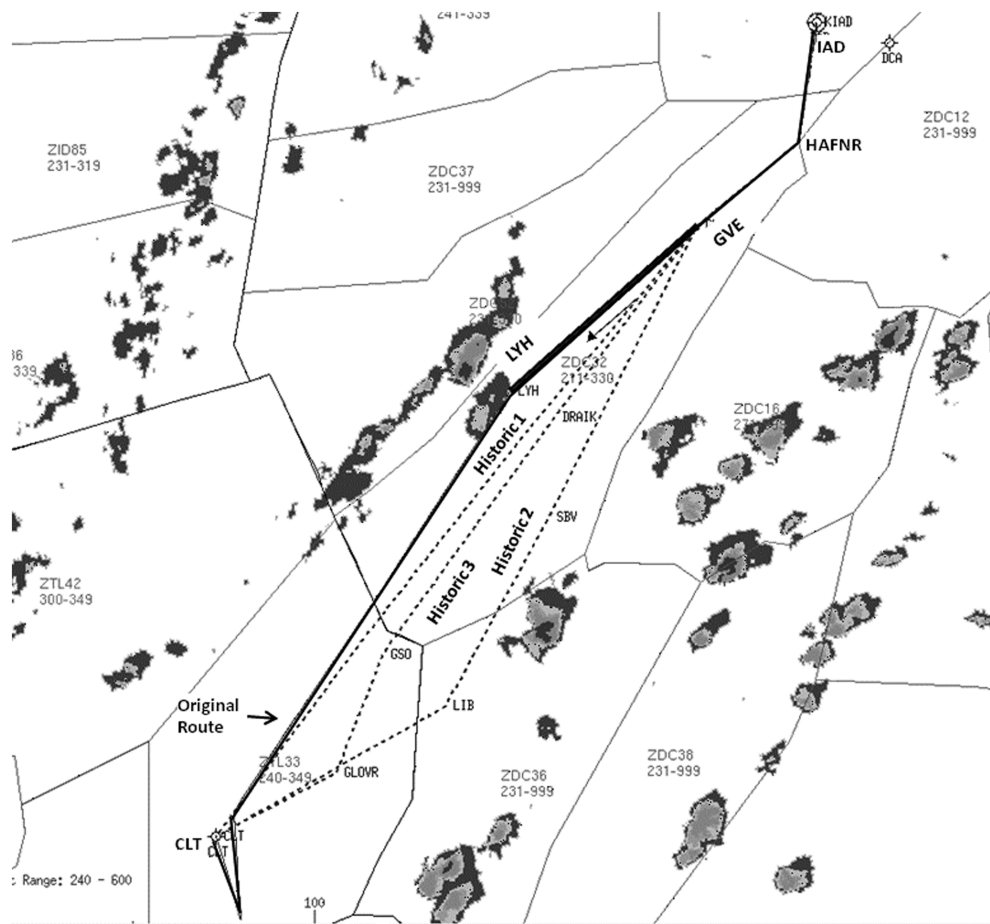


Fig. 14 Historic reroute options for flight C.

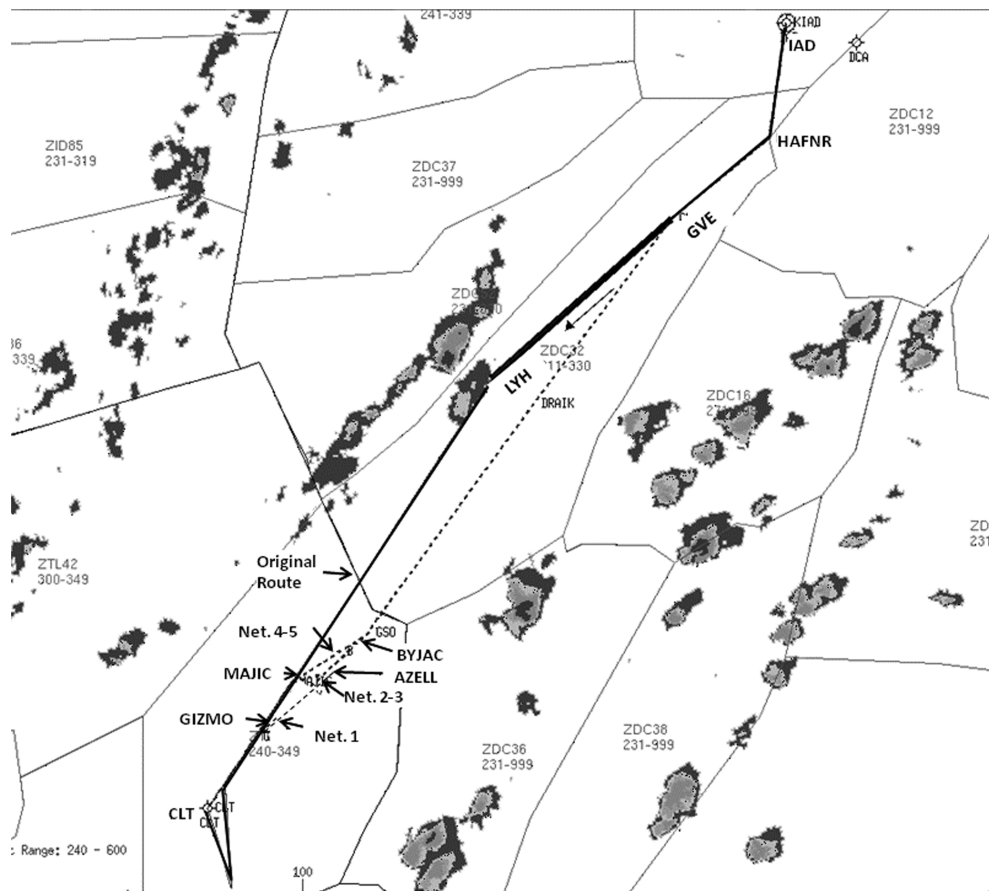


Fig. 15 Network-Optimized reroute options for flight C.

however, this minor improvement is offset by poorer performance in both O-D and global flow conformity. Comparing the first network-optimized reroute to the first historic option, we see that although the network-optimized reroute performs approximately 1.2% worse in the distance metric, the O-D and global flow conformity metrics show approximately a 68 and 19% improvement, respectively. The remaining historic options perform no better than any network-optimized reroutes, and often perform worse.

C. Analysis of Computation Requirements

The methodology proposed in this paper aims to address the need of decision makers for defining flight-specific reroutes in a real-time environment. As the previous analysis showed, the reroutes generated by the proposed methodology can provide options when no historic options are available, as well as produce equivalent or better operationally acceptable alternatives than the historic reroutes. By defining operationally acceptable reroutes using automation, the




Fig. 16 Comparison of network-optimized and historic reroutes for flight C.

Table 3 Analysis of computation effort

Examples	Number of nodes	Number of arcs	Network creation time, sec	Reroute generation time, sec	Number of reroutes generated	Reroute evaluation and selection time, sec	Total time, sec
Flight A	875	7580	10.86	2.36	109	0.82	14.05
Flight B	2235	28695	51.93	14.24	89	1.9	68.08
Flight C	797	7129	10.2	2.69	182	1.9	14.84

significant workload associated with manual rerouting can be reduced; however the methodology must be computationally-efficient.

Table 3 provides the details on the network size and computational effort required to define the network-optimized reroutes for the three flights analyzed in this paper. All computation times represent an average of ten runs for each flight. The methodology was implemented in the Java programming language and executed on a nondedicated server. As the code was not optimized for performance, this analysis is intended for illustrative purposes only and not as a measure of the absolute performance of the methodology.

Examining Table 3 we see that overall the overall computation time is related to the network size. Specifically Flight A and Flight C have similar network sizes (in terms of number of nodes and number of arcs defined), and similar computation times. Flight B, on the other hand, requires over 1 min of computation time as the network is much larger. Recalling that Flight B is enroute from LGA to CLT, the scope of the network is much greater than the other two flights.

VI. Conclusions

The research presented in this paper extends previous work on developing a methodology for the computer-based design of operationally acceptable flight-specific reroutes. Dynamically generating reroutes is challenging as it requires capturing aspects of quality route design that are generalizable to any flight. As such, the methodology proposed in this research begins with consultations with SMEs to define metrics of reroute operational acceptability. In addition, the network model is developed using only operationally acceptable segments of historic routes. Finally, acknowledging that there will be aspects of a given traffic scenario that cannot be captured by the metrics developed, the methodology is proposed as a component to a decision support tool for traffic managers or decision makers, where multiple reroute options for each flight will be generated for further evaluation and selection.

The methodology presented here aims to improve the reroute alternatives provided for a flight by employing a network optimization approach, where metrics of operational acceptability are used to evaluate and select alternatives. To enable a more robust reroute definition, a dynamic definition of weather blockage was first defined, and how changing the prioritization of weather avoidance in reroute generation could change the alternatives provided was analyzed. For the example considered, it was shown that as the arc weighting parameter or route blockage increased, more conservative reroutes were generated. As such, the decision was made to implement the moderate weighting, which produced reroutes that spanned the more and less risk-adverse weightings, providing both options to decision makers.

The purpose of this paper was to analyze how the reroute options generated by this method perform compared with traditional reroute alternatives. From the analysis presented here, it can be seen that there is significant potential for enhanced reroute development using the network optimization approach, as these reroutes provide alternatives that are event specific, as opposed to the historic reroutes. Furthermore, the reroutes generated by the network optimization possess similar or better operational acceptability characteristics than the historic reroutes. This is expected as the traditional reroutes do not consider these metrics when constructed. Although refining the definition of the operational acceptability metrics is an area of ongoing research, capturing these metrics in the design of reroutes provides a quantitative evaluation of proposed alternatives, and the methodology presented here is flexible in so much that any metric

that can be measured using the information available can be implemented and used to generate network-optimized reroutes. Finally, flights with difficult reroutes will still be provided with reroute options since the network optimization approach is guaranteed to return a feasible reroute, if there is a connected network.

In this paper, individualized reroutes for multiple flights under the same weather constraint were developed. However, an area of ongoing research is to extend these results for solving multiple flights simultaneously. Moving toward a generalization of the network model for multiple flights will enhance the decision support capabilities provided to decision makers when considering alternatives for groups of flights impacted by weather.

Acknowledgments

The authors would like to thank Jonathan Anderson, James DeArmon, Dan Greenbaum, Joe Nardelli, Laurel Rhodes, and Stephen Zobell of MITRE Corporation for their contributions to the research described here.

The contents of this material reflect the views of the authors and MITRE Corporation and do not necessarily reflect the views of the FAA or the DOT. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of these views.

References

- [1] *Severe Weather and Route Management 2006*, Federal Aviation Administration: Air Traffic Control System Command Center, http://www.fly.faa.gov/Operations/Strategic_Planning/svrwx_handbook.html [retrieved May 2011].
- [2] Wanke, C., and Greenbaum, D., "Sequential Congestion Management with Weather Forecast Uncertainty," *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2008-6327, Aug. 2008.
- [3] Wanke, C., "Continual Probabilistic Airspace Congestion Management," *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2009-5863, Aug. 2009.
- [4] Mayhew, D., and Manikonda, V., "Constrained Airspace Rerouting Planner for ACES," *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2008-6346, Aug. 2008.
- [5] Grabbe, S., Sridhar, B., and Mukerjee, Av., "Sequential Traffic Flow Optimization with Tactical Flight Control Heuristics," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 3, May-June 2009, pp. 811-820.
- [6] Bokadia, S., and Valasek, J., "Severe Weather Avoidance Using Informed Heuristic Search," *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2001-4232, Aug. 2001.
- [7] Krozel, J., Penny, S., Prete, J., and Mitchell, J., "Automated Route Generation for Avoiding Deterministic Weather in Transition Airspace," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 1, Jan.-Feb. 2007, pp. 144-173. doi:10.2514/1.22970
- [8] Prete, J., and Mitchell, J., "Safe Rerouting of Multiple Aircraft Flows in the Presence of Time-Varying Weather Data," *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2004-4791, Aug. 2004.
- [9] Krozel, J., Prete, J., Mitchell, J., Smith, P., and Andre, A., "Designing On-Demand Coded Departure Reroutes," *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2006-6767, Aug. 2006.
- [10] Weyns, D., Holvoet, T., and Hellebooghe, A., "Anticipatory Vehicle Routing Using Delegate Multi-Agent System," *Proceedings of the 2007 IEEE Intelligent Transportation Systems Conference*, IEEE Publications, Piscataway, NJ, Sept.-Oct. 2007, pp. 87-93.
- [11] Taber, N., Jackson, C., Masalonis, A., Song, L., Taylor, C., Wanke, C., and Zobell, S., "Advanced TFM Congestion Management Performance

- Analysis and Research Results,” The MITRE Corporation MTR090220, Sept. 2009.
- [12] DeArmon, J., Wanke, C., and Nardelli, J., “Quantification of Desirable Air Route Attributes in a Reroute Generation Capability,” *AIAA Aviation Technology, Integration, and Operations Conference*, AIAA 2009-7092, Sept. 2009.
 - [13] Wanke, C., Taber, N., Greenbaum, D., Jackson, C., Jacobs, G., Masalonis, A., Rhodes, L., Song, L., Stewart, T., Taylor, C., and Zobell, S., “Research Results for Advanced Congestion Prediction and Resolution Capabilities,” The MITRE Corporation MTR080087, Aug. 2008.
 - [14] Delaura, R., and Evans, J., “An Exploratory Study of Modeling Enroute Pilot Convective Storm Flight Deviation Behavior,” NASA A-6, 2006.
 - [15] Taylor, C., and Wanke, C., “Dynamic Generation of Operationally Acceptable Reroutes,” *AIAA Aviation Technology, Integration, and Operations Conference*, AIAA 2009-7092, Sept. 2009.
 - [16] Jarvis, J., and Ratliff, D., “Some Equivalent Objectives for Dynamic Flow Problems,” *Management Science*, Vol. 28, No. 1, 1982, pp. 106–109.
doi:10.1287/mnsc.28.1.106
 - [17] Ahuja, R., Magnanti, T., and Orlin, J., *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall Inc., Englewood Cliffs, NJ, 1993.
 - [18] Ng, H., Grabbe, S., and Mukherjee, A., “Design and Evaluation of a Dynamic Programming Flight Routing Algorithm Using the Convective Weather Avoidance Model,” *AIAA Guidance Navigation and Control Conference and Exhibit*, AIAA 2009-5862, Aug. 2009.
 - [19] Liu, G., and Ramakrishnan, K., “A*Prune: An Algorithm for Finding K Shortest Paths Subject to Multiple Constraints,” *Proceedings of the 2001 IEEE INFOCOM*, Vol. 2, IEEE Publications, Piscataway, NJ, pp. 743–749.