

Automated Route Generation for Avoiding Deterministic Weather in Transition Airspace

Jimmy Krozel* and Steve Penny†

Metron Aviation, Inc., Herndon, Virginia 20170

and

Joseph Prete‡ and Joseph S. B. Mitchell§

Stony Brook University, Stony Brook, New York 11794

DOI: 10.2514/1.22970

We investigate the problem of synthesizing weather avoidance routes in the transition airspace, from the 200 n mile range to the metering fixes of an airport, given a deterministic weather forecast. The investigation of the problem is motivated by the desire to maximize airspace capacity while ensuring safe separation between aircraft and between aircraft and hazardous weather. Three solution methods are compared with current day operations. Emphasis is placed on a comparison of the arrival traffic weather avoidance routing method and metrics associated with such routes. Historical weather avoidance paths (baseline) are compared with three alternatives: variations of the standard arrival routes, a geometric optimization solution synthesizing multiple nonintersecting routes, and two free-flight approaches in which aircraft fly weather avoidance routes using a greedy prioritization method. Results indicate that increases in capacity over today's system are achievable, while maintaining safety. However, such increases are limited by the method, the required supporting infrastructure, and weather severity and location.

I. Introduction

AIR traffic management (ATM) in the airport transition airspace works well under normal operating conditions. However, during inclement weather, the reduction in available nonconvective airspace limits capacity, adversely affecting throughput. The Aviation Capacity Enhancement Plan [1] lists weather as the leading cause of delays greater than 15 minutes, with terminal volume as the second leading cause. Furthermore, weather-related delays are likely to get worse as the volume of air traffic increases in the future. The future national airspace system (NAS) requires that airports must accommodate greater capacity; hazardous weather is a challenging limiting factor.

Whereas weather phenomena have not fundamentally changed over the years, delays have steadily increased due to greater demands placed on limited resources such as runways, navigational fixes, and jet routes. Those resources experience significant capacity reductions during severe weather events. Observations of weather avoidance maneuvering typically reveal reactive (tactical) deviations around hazardous weather cells. Safety constraints dictate that aircraft must remain separated from one another as well as from hazardous weather. As aircraft deviate around weather cells (it is the pilot's responsibility to avoid hazardous weather) they stray from well-defined jet routes and require neighboring aircraft to "keep out of their way." Air traffic control (ATC) provides the function of safely separating aircraft. Thus, neighboring aircraft are typically spaced out at greater distances behind a leading aircraft along a jet route flow as ATM personnel increase the mandated miles-in-trail (MIT). This effectively reduces the flow's throughput and adds to the delays experienced during severe weather events. Thus, not only are

aircraft delayed due to a path stretching effect of weather avoidance maneuvering, they are also delayed due to MIT restrictions imposed to make the situation more manageable during weather events.

This paper compares three solution approaches to synthesizing (by computer automation) weather avoidance maneuvers for aircraft in the transition airspace: within a roughly 200 n mile range to the metering fixes (MFs) of an airport. Using current NAS operations as a baseline, new ATM solutions are investigated which may change the way routes are defined or the roles and responsibilities between pilot and controller:

1) Variations on standard terminal arrival routes (STARs): The navigation aids (NavAids) that constitute the STARs of an airport are varied by a set amount of lateral separation to design a new STAR as a function of time of day that safely avoids hazardous weather. Aircraft are metered at the 200 n mile range ring so that only one aircraft arrives at a MF at a time, regardless of the weather avoidance route chosen.

2) Flow-based route planner (FBRP): Nonintersecting routes are designed to lead from a 200 n mile range to airport MFs. Multiple routes are synthesized for each MF, based on maximizing the total number of routes that both avoid hazardous weather and are nonintersecting (to reduce airspace complexity). Given a set of routes leading to the MFs, the route that minimizes the distance traveled, subject to the constraints, is chosen for each aircraft. Aircraft are metered at the 200 n mile range so that only one aircraft arrives at a MF at a time, regardless of the weather avoidance route chosen.

3) Free flight: Free-flight [2] routes are synthesized to avoid hazardous weather and to separate aircraft from those already routed in the transition airspace. The algorithm is "greedy" in the sense that each aircraft acts independently, selecting a best safe path between the 200 n mile range ring and the MF, subject to separation constraints imposed by the routes chosen by earlier aircraft arriving into the airspace. Separation requirements implicitly create the metering of aircraft at the MF. Two forms of free flight are studied: a) required time of arrival (RTA), where aircraft are required to arrive at a MF at the RTA allocated by ATC to maximize throughput while avoiding conflicts at the fix. They are given an appropriate RTA at the 200 n mile range to facilitate the RTA at the MF; and b) first come first served (FCFS), where aircraft are routed to the MF on a FCFS basis as they arrive at the 200 n mile range ring.

All approaches assume a common, shared model of weather constraints based on a perfect weather forecast. In the future as

Presented as Paper 4790 at the AIAA Guidance, Navigation, and Control Conference, Providence, RI, 16–19 August 2004; received 3 February 2006; revision received 11 April 2006; accepted for publication 24 April 2006. Copyright © 2006 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. 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 \$10.00 in correspondence with the CCC.

*Senior Engineer, Research and Analysis Department, 131 Elden Street, Suite 200, Associate Fellow AIAA.

†Analyst, Research and Analysis Department, 131 Elden Street, Suite 200.

‡Research Assistant, Department of Applied Mathematics and Statistics.

§Professor, Department of Applied Mathematics and Statistics.

net-centric operations become the standard for the NAS, this might be the case; however, currently, the flight deck, airlines, and controllers do not share a common view of the weather. We assume a perfect weather forecast and a common shared model of weather constraints to form a fair comparison between the route planning algorithms. The analysis of the robustness of these routing algorithms to weather forecasting uncertainties remains a future research effort.

II. Related Literature

A number of approaches to synthesizing weather avoidance routes are found in the literature. Several approaches apply an optimal path algorithm based on grid search methods [3–7]. The problem can be formalized as a weighted regions problem in which routes obey Snell's law of refraction [5,8,9,11] as a local optimality criterion. Algorithms can exploit the fact that optimal routes bend at boundaries between regions of varying weather severity in analogy with light rays that refract as they pass through regions of varying refractive index [5–11]. Another approach [3] searches for paths having at most k turns (waypoints), while avoiding hazardous weather, thereby bounding the workload of the pilot and controller required to track the solution. In general, these approaches mostly address the routing of a single aircraft rather than designing routes for flows of aircraft or for system optimization, as addressed in this paper.

III. Theoretical Problem Statement

The transition airspace problem for ATM is a four-dimensional (4-D) problem that includes arriving and departing traffic heading to/from an airport. However, due to current ATM practices, the dimensionality and complexity of this problem are reduced and simplified in many ways. For example, aircraft are routed only on well-defined jet routes that turn at Navaids and intersections. Analysis of traffic leading to and from Atlanta's Hartsfield International Airport (ATL), for instance, indicates that the arrival and departure traffic have standard climb and descent profiles that separate the arrivals and departures vertically. Horizontally, within roughly 50 n mile of ATL [as shown in Fig. 1 based on enhanced traffic management system (ETMS) data], the arrival and departure traffic are also separated by ATC low altitude sectors designated as either arrival or departure sectors. Some of these same practices that reduce the dimensionality of the problem are retained in our problem statement.

Optimizing traffic flow in the transition airspace must take into consideration safety, efficiency, workload (complexity), operational feasibility (ease in verbally coordinating maneuvers between controllers and pilots), and roles and responsibilities (separation between aircraft by pilots vs controllers and separation between aircraft and hazardous weather by pilots vs controllers). These factors are taken into consideration as mathematical problem

statements are developed for new methods of routing traffic around hazardous weather in transition airspace.

A. Problem Geometry and Constraints

The geometry and constraints are modeled first, as illustrated in Fig. 2. Whereas the geometry is generally described by 3-D space, we model the problem in a 2-D horizontal plane primarily due to the fact that we plan routes using standard ascent and descent rate profiles. Our model extends to 3-D airspace domains and those that take into account the curvature of the Earth; however, this is not needed in the transition problem statement described in this paper. Let A be the airspace domain, a subset of the real plane. Domain A consists of the union of one or more ATC sectors describing the transition airspace in the proximity of an airport centered at point $a_0 \in A$ with a range ring R (radius 200 n mile). At time t there exists a set of constraints, $C(t) = \{C_1(t), C_2(t), \dots, C_N(t)\}$, with each connected component $C_i(t) \subset A$, a region of airspace through which aircraft cannot safely traverse. Airspace constraints are of various types, including 1) no-fly zones and special use airspace (SUA) (typically stationary, not varying with t); 2) departure-only constraints or arrival-only constraints, which specify airspace regions that apply only to departing or to arriving aircraft (constraining departures to avoid airspace designated primarily for arrivals and vice versa); and 3) weather constraints, typically growing, decaying, or moving weather cells (varying with t).

Weather constraints arise from convective weather severe enough (above an intensity threshold) to pose a safety hazard for aircraft. Whereas the criteria for weather avoidance depend on pilot preferences and airline guidelines, research [12] shows that pilots generally avoid National Weather Service (NWS) Level 3 and higher weather cells (greater than 13.3 mm/hr rainfall or reflectivity greater than 41 dBZ). Whereas the altitude of cloud tops in severe storms is also an important factor [13] that pilots consider in determining which storm cells to avoid, it is not modeled in our problem statement.

Because of weather forecasting errors, weather constraints are not usually known with certainty. The uncertainty is smaller for short-range forecasts (up to 1 h in time), but the uncertainty increases and becomes substantial for long-range forecasts (beyond 1 h in time). In this paper, we model weather constraints as deterministic constraints, varying with time according to a piecewise-constant function that is based on a weather forecast model: the most recent short-range weather forecast data are considered to be valid (and static) from the moment it is made until the moment the next updated short-range weather forecast is made, nominally at a 5 min update rate. Ideally, a fully stochastic and dynamic model of weather would be available and usable for route synthesis and periodic reroute planning. However, we leave it as an important (and challenging) future research item to extend our investigation to establish remaneuvering requirements [14] associated with weather forecast uncertainties. In a manner similar to the work of Love et al. [14], we account for a small

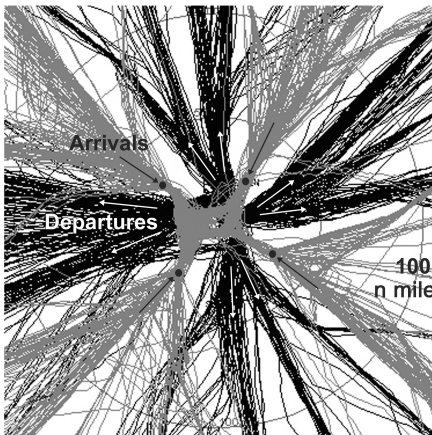


Fig. 1 Arrivals and departures are typically separated horizontally near an airport.

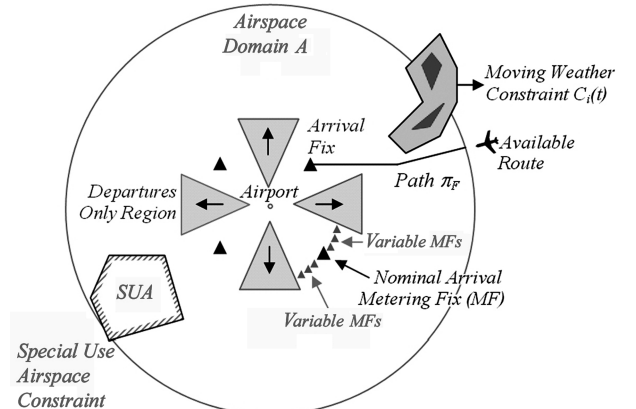


Fig. 2 Geometry and constraints for the arrival transition airspace problem.

amount of uncertainty for weather constraints by imposing a hazardous weather safety margin $\varepsilon_i > 0$ for each constraint $C_i(t)$. Each planned route must avoid approaching within distance ε_i of constraint $C_i(t)$. In our experiments, we use a single weather safety margin, $\varepsilon = \varepsilon_i$, for all constraints $C_i(t)$, $i = 1, 2, \dots, N$. Mathematically, this corresponds to replacing each $C_i(t)$ with the constraint $C_i(t) \oplus B(\varepsilon_i)$, where \oplus denotes the Minkowski sum, and $B(\varepsilon_i)$ is the ball of radius ε_i centered at the origin.

Because this study focuses on synthesizing routes from 200 to 40 n mile from an airport, a horizontal separation requirement δ (5 n mile) requires two aircraft to be horizontally separated by a distance δ at each instant of time t .

B. Aircraft Dynamics and Aircraft Flows

Aircraft are modeled as points in motion within A . Aircraft dynamics are specified in terms of bounds on the speed and the magnitude of acceleration. Whereas acceleration bounds give rise to bounds on the radius of curvature of flight, we assume that the scale of our solution space is large enough that we can approximate aircraft dynamics with a simple representation of piecewise-linear flight legs connected at waypoints, as is the standard notation for flight plans today. We assume that there is a constant speed profile $v_j: A \rightarrow [v_{\min}, v_{\max}]$ in each of the four quadrants Q_j surrounding the airport, on each historical day studied. This reflected large-scale wind fluctuations that occurred geographically during a single experiment day, or within each quadrant Q_j over multiple experiment days. Note that our model does not explicitly take into account the fact that different classes of aircraft may travel at different speeds. We assume all aircraft are expected to travel according to a single ascent or descent speed profile along any given route.

A set of flow demands D represents flow requirements for aircraft passing through A . Each demand $D = (o_D, d_D, [t_D^s, t_D^e]) \in D$ consists of an origin of the flow $o_D \subset A$, a destination of the flow $d_D \subset A$, and a time interval $[t_D^s, t_D^e]$ over which the demand occurs, from start time t_D^s to end time t_D^e . During the time interval $[t_D^s, t_D^e]$ a stream of aircraft are arriving at o_D heading to the destination d_D . We define o_D and d_D to be sets, in general, rather than singleton points, to allow flexibility in the exact location a flow is to originate or terminate. In transition airspace, for example, we may define the origin o_D for arriving flights to lie within an arc on the range ring R , and the destination to be either a fixed MF location, a destination point d_D fixed in space for all time t , or a variable MF location (Fig. 2), and a destination d_D variable in location and time t , chosen to maximize throughput. When varying the MF location, one of either three or nine possible discrete MF locations along an arc centered at the original MF location are considered.

Define a flow F to be a path $\pi_F \subset A$, from a point o_F to a point d_F , and an associated time interval $[t_F^s, t_F^e]$ during which the flow is active. A flow F is feasible if, for any $t_0 \in [t_F^s, t_F^e]$, an aircraft that arrives at o_F at time t_0 and flies exactly along the path π_F according to the speed profile function, is at all times safe with respect to the set of constraints $C(t)$. More precisely, flow F is feasible if for any $t_0 \in [t_F^s, t_F^e]$ and any $t \in [0, t_\pi]$, we have

$$[x(t_0 + t), y(t_0 + t)] \in AC \in C(t_0 + t) \quad (1)$$

locating at time $t_0 + t$ a point along π_F that arrives at point o_F at time t_0 and moves along π_F according to the speed profile function (thus, this location may be obtained by integrating the speed profile function along the path π_F). Here, t_π denotes the time it takes to fly along π_F , from o_F to d_F , with speed that matches the speed profile function.

Flows are designed so that they do not interfere with each other in the transition airspace. A set of flows F is feasible if each flow $F \in F$ is feasible and for any pair of flows, $F_1, F_2 \in F$ and any time $t \in [t_{F_1}^s, t_{F_1}^e + t_{\pi_1}] \cap [t_{F_2}^s, t_{F_2}^e + t_{\pi_2}]$, the position of an aircraft at time t along the path π_1 for F_1 is separated from the position of an aircraft at time t along the path π_2 for F_2 by at least the horizontal separation distance δ .

A feasible flow F is said to satisfy demand D for time interval Δt if it starts and ends according to the demand ($o_F \in o_D$ and $d_F \in d_D$),

and its time interval has an overlap of length Δt with the time interval of D , i.e., $\|[t_F^s, t_F^e] \cap [t_D^s, t_D^e]\| = \Delta t$.

The throughput associated with a feasible set of flows is defined to be the total sum of the lengths of the demand time intervals that correspond to feasible flows that satisfy the demand. More precisely, the throughput, $\mu_F(D)$, of F provided for demand D is defined to be the total length (measure) of the subintervals of $[t_D^s, t_D^e]$ that are satisfied by at least one flow in a given feasible set F of flows. The total throughput of the set F is defined to be $\sum_{D \in D} \mu_F(D)$.

A flow can be thought of as an opportunity for routing aircraft to satisfy a demand: a “highway in the sky” open for a specified interval of time. In space-time, a flow corresponds to a continuous set of trajectories having the same geometry in space (following the same waypoints), but having starting times that vary within a certain window of time; thus, a flow can be visualized as a swept trajectory in space-time, for which there is an opportunity for multiple aircraft to follow safely the same trajectory, spaced according to MIT restrictions. A feasible set of flows represents an opportunity for multiple trajectories having distinct, safe geometry, both in terms of weather constraints and in terms of horizontal separation standards.

To simplify the presentation of results, we only compare results for arrival routes, and do not analyze the coupling of arrival routes and departure routes; nonetheless, our algorithms and solution approaches are quite general and are easily applied to the problem of synthesizing arrival and departure routes simultaneously. For instance, the same algorithm that is used to design variable STARs can be used to design variable standard instrument departure routes (SIDs). Furthermore, we do not investigate the synthesis of weather avoidance routes from the MFs to the runways. These routes can be synthesized using similar approaches to those that are applied here; see Krozel et al. [3] for an algorithmic solution. Finally, we do not explicitly address vertical separation issues, because real-world data suggest that the arrival and departure routes can be separated based on standard arrival descent profiles and standard departure ascent profiles that are designed to avoid vertical conflicts.

IV. Algorithmic Solutions

Algorithms were developed for the three solution methods for routing either flows of aircraft or individual aircraft (for free-flight methods). A set W of waypoints defines the points at which a flow/route is allowed to turn. Each of the methods involves computing one or more routes in a search graph, $G = (W, E)$, with node set W and edge set E . Historically, these locations have been specified by the physical position of Nav aids, physical navigation aid devices located on the surface of the Earth. With the latest navigation equipment, though, points of W can now be located at essentially any point of airspace A . Depending on the model described next, the waypoints will be defined as either fixed locations or they will vary as a function of the location of the constraints. Further, we consider cases in which the MF is either fixed in location, as dictated by current operational standards, or variable in location, as may be the case for the future, allowing the MF to move in location within a small tolerance.

A. Variable STARs

As shown in Fig. 3, the variable STARs method creates alternate routes around hazardous weather using the STARs as a baseline and additional waypoints that are varied within a fixed lateral offset from the STAR. The underlying search graph $G = (W, E)$ is defined as follows: the set W of nodes of G consists of STAR Nav aids, together with waypoints defined as shifted copies of each Nav aid, plus or minus a lateral offset. The lateral offset varies from 8 n mile at the MF to 16 n mile at the 200 n mile range, defined by a linear function $\omega[d(a_0, w)]$, where $d(a_0, w)$ is the distance from the airport a_0 to the waypoint $w \in W$. Let w_i be a waypoint that is part of a STAR, and let $d(a_0, w_i)$ be the distance from the airport a_0 to the waypoint w_i . If c_1 is a circle centered at the airport with radius $d(a_0, w_i)$, and c_2 is a circle centered at the waypoint w_i with radius $\omega[d(a_0, w_i)]$, then the two points of intersection of circles c_1 and c_2 define alternate (shifted) waypoints for w_i .




Fig. 3 The ORD weather avoidance search graph based on SIDs and STARs.

Define the edges E within graph G as follows. For each STAR, denote the STAR's original $(n+1)$ base waypoints by w_0, w_1, \dots, w_n , with $w_0 = a_0$, and the $2n$ new alternate waypoints by $w_{1,1}, w_{1,2}, w_{2,1}, w_{2,2}, \dots, w_{n,1}, w_{n,2}$. Edges E in graph G connect w_0 to each of the waypoints $w_1, w_{1,1}$, and $w_{1,2}$, and connect each vertex at stage j (i.e., $w_j, w_{j,1}$ and $w_{j,2}$), to each vertex at stage $(j+1)$ (i.e., $w_{j+1}, w_{j+1,1}, w_{j+1,2}$). For each time interval $[t_F^s, t_F^e]$, we search the graph, using an A^* algorithm [15], for an optimal path from w_0 to one of three goal points, $w_n, w_{n,1}$, and $w_{n,2}$. The resulting path π_F defines a flow F over $[t_F^s, t_F^e]$ available to satisfy the demand over this time interval.

In ATM operations, the lateral offsets may be communicated easily between controllers and pilots (verbally or with automation), and familiarity with SID and STAR Navaids is not lost. The method maintains close similarities with current ATM operations while providing some flexibility and automation for hazardous weather avoidance. Such similarities are clearly a benefit during the deployment stage of a route-synthesizing algorithm.

A limitation of this method is that synthesized routes may have no choice but to penetrate hazardous weather. In such a case, the flow is not feasible, and aircraft are restricted from using the route. Subsequently, aircraft are not permitted to enter the airspace until the algorithm considers the next time interval and can synthesize a feasible route that avoids hazardous weather. In operation, this would lead to route closures with consequent delays assigned upstream. To compute a maximum throughput, MIT restrictions of 5 n mile (consistent with current day minimums) are maintained along planned routes.

The variable STAR method may use either fixed MF or variable MFs. When varying the MF, the variable STAR method uses one of three possible discrete MF locations 8 n mile apart, centered at the original MF location.

B. Flow-Based Route Planner (FBRP)

The FBRP algorithm [16] determines if there are one, two, or three nonintersecting weather avoidance routes that lead from the range ring R to either one fixed MF or one of k_{MF} ($k_{MF} = 9$ in our experiments) evenly spaced variable MFs along an arc through the




Fig. 4 The ORD weather avoidance search identifies nonintersecting routes forming independent flows.

nominal MF (Fig. 2). Technically, we can design up to a parameter limit of k routes, which we limit in our application to $k = 3$. Each single route is partitioned into a set of flows having short time gaps between their respective arrival time windows, to account for any MIT restrictions. The creation of multiple nonintersecting routes in multiple flows adds another layer of complexity to the algorithm, because horizontal separation standards must be enforced. The algorithm computes multiple routes incrementally, enforcing separation standards with respect to each route already in place.

Because conflicts can arise between already routed flows and flows yet to be routed, to increase the ability to discover a set of routes, the algorithm searches over several possible orderings of routes, in hopes of finding solutions to all routes (if such hazard-free routes exist). Note that whereas all three routes may potentially be in use, by the time aircraft arrive at the MF, they will have been controlled by speed to be separated longitudinally by 5 n mile. Thus, at the MF a continuous flow of aircraft is achieved even though the aircraft may be using different routes to arrive at the MF. Figure 4 illustrates an example solution from this algorithm.

Routes are planned in two stages. First, we approximate a shortest path between the source and the destination by searching for a route that uses the points of a grid. Second, we refine this route by searching for another, simpler route, nearby the original route.

Stage 1. An approximation is established for a shortest path $\pi_F \subset A$, from a start point o_F to a destination MF point d_F , by searching a graph using the A^* algorithm [15]. The search is guaranteed to terminate with an optimal path (just as breadth-first methods, such as Dijkstra's algorithm [17]), but the search is goal-directed, exploring the most promising extensions to the path first. The path must satisfy all the specified constraints: avoiding hazardous weather by at least the hazardous weather safety margin $\varepsilon_i > 0$, maintaining the horizontal separation requirement δ with respect to already established flow routes, and obeying turn/orientation constraints.

The search graph $G = (V, E)$ is based on an enhanced-connectivity rectilinear grid. The grid can be any size, but the fineness of the grid must be balanced against both the quality of the solution and the running time. Grid points were spaced roughly 6.25 n mile apart; for a 400×400 n mile airspace A around an airport, roughly $64 \times 64 = 4096$ grid points. Points of the grid become the vertices of a graph and are interconnected in specific ways, with the connections forming the edges of the graph, which represent candidate flight legs. Interconnectivity is defined with interconnectivity constant K as illustrated in Fig. 5. For each grid point P , the set $N(P)$ includes all points within K units of P along each axis; this results in a square group of grid points, $2K + 1$ units on a side, centered on P . We connect P to each of the grid points of the neighborhood $N(P)$, but connect P to only one of the grid points in any particular direction from P , namely, the closest one in that direction. For example, if $P = (i, j)$ and $K = 3$, connect P to $(i + 1, j + 1)$, but not to $(i + 2, j + 2)$ or $(i + 3, j + 3)$, because these points are collinear with the connection to $(i + 1, j + 1)$, in the direction at 45 deg with respect to the x -axis. Higher values of K produce a larger number of edges linking a point to its neighbors at a wider variety of orientations.

As the search proceeds through the graph, each node that is reached by a path of valid edges (connections in the graph that do not violate any of the constraints) is labeled with a time (or cost) value. This value indicates the minimum time (or cost) required by a path from the start U to this node, along a path of valid edges. The same




Fig. 5 Connections from a central point to its neighbors for connectivity $K = 1, 2, 3, 4$.


position with different time values is the same node; thus, at no time are we ever searching a 3-D space of nodes in space-time. When a shortest-path has been found from start point o_F to a destination MF point d_F , every node on that path has been labeled with the time at which an aircraft can first reach it.

Hazardous weather is retained in a time series of weather forecasts. The weather grid is independent of the search space grid and is usually at a different resolution. Each weather grid point represents the forecasted weather severity at one particular point in time. Thus, time varying weather is represented as a series of such grids sequenced evenly at times T_i across the time of interest. We have taken the conservative position that the weather sample from time T_i is valid from the earliest possible time to the latest; thus, it is valid, with above-threshold hazardous weather being impenetrable, from T_{i-1} to T_{i+1} , and this is the valid time window of the weather. If the weather data are the first or the last over the time horizon of the experiment, then their time window has no limit in the appropriate direction.


Determining hazardous weather conflicts requires locating the part of a path leg relevant to the time window of a weather sample, and then checking the path against the weather sample to determine if the path abides by the constraints. Each candidate path in the grid search extends over both space and time. In 3-D space-time, then, a path occupies a region that is a union of connected quadrilaterals that cross from the start waypoint to the goal waypoint while moving upwards (i.e., forward in time). To maintain the horizontal separation standard δ , each path can instead be treated as if it were an extruded cylinder in space-time.

It is computationally difficult to perform exact calculations to determine whether two paths conflict in space-time. Therefore, the extruded-cylinder shape is approximated by a bounding polytope, as shown in Fig. 6. These extruded cylinders become the input to a polygon collision detection algorithm based on bounding-volume hierarchies of k -dops [18]. Figure 6 illustrates only one path leg; the entire flow is represented by a chain of such polyhedra.

Turn angle constraints can be handled by augmenting the search graph to take into account heading turn limits [3]. However, the results we report here are based instead on using another form of angle constraint: each flight leg is required to stay within a specified angle of the vector from the start point o_F to a destination MF point d_F . The parameter Θ indicates the maximum angle that any flight leg (edge along the path π_F) can make with respect to the line from point



a) Continuous representation



b) Bounding Discrete Polytope

Fig. 6 The space-time constraint region corresponding to a flight leg for aircraft in a flow.




Fig. 7 Refinement of candidate turn points.

o_F to point d_F . This constraint captures several important good qualities of paths that are difficult to quantify. It also greatly restricts the area that will be searched, which is useful for reducing running time.

Stage 1 outputs a path π_F which may have too many waypoints to be useful or practical for ATM applications. So, stage 2 of the algorithm refines the path π_F to reduce its complexity, searching for a new path with fewer turns. Reducing the number of turns is intended to make the maneuver easier to clear by air traffic controllers.

Stage 2. The refinement stage searches for a path over a graph of a specifically defined set of points S (candidate waypoints). For each point p_i in π_F , we define a set S of points as follows (Fig. 7): 1) each point p_i is in S , and 2) M evenly spaced points on each side of p_i , within a user-specified path tolerance distance τ , along the angular bisector of the path τ_F at point p_i .

Whereas the stage 1 search sought paths of low distance, the stage 2 search minimizes the number of links (flight legs, or waypoints) in a path π_F . This may increase the length of the output path by a small amount, because the output path may find a simpler but slightly longer path by using one of the additional vertices in the search space.

In addition to the search criteria from the initial search, we add a length criterion to the refinement phase; a link that is too short is ignored. Admissible links in the graph, and hence legal paths, have the following properties: 1) avoids hazardous weather, with intensity above a user-specified threshold; 2) obeys the separation standards for aircraft; 3) travels within a particular arc determined by the direction to the goal point; and 4) has at least a specified minimum length, L .

C. Free Flight

The free-flight algorithm (Fig. 8) is designed to synthesize routes that enable individual flights to be routed conflict-free and clear of dynamic, hazardous weather constraints. As each flight arrives at the 200 n mile boundary, it is individually routed to the arrival MF corresponding to its quadrant of entry. Each aircraft is optimally routed without regard for how its route might adversely affect upstream aircraft (e.g., a nonoptimal sequence). The solution is designed to mimic a distributed system in which each aircraft noncooperatively plans its own route.




Fig. 8 A search identifies the optimal free-flight route avoiding hazardous weather.

The free-flight algorithm searches space-time (using the same principles as in the space-time flow method [19]), treating already routed flights downstream of an aircraft as constraints to be avoided (according to minimum separation standard between aircraft), along with severe weather constraints, no-fly zones, and quadrant constraints (confining an aircraft to the quadrant initially entered).

Free-flight routing allows maximum flexibility in terms of individual aircraft routing (user preferences); however, by allowing each flight to be greedily routed, the free-flight solution can result in some routes being particularly efficient, at the expense of blocking off a passage for upstream aircraft. Furthermore, there is no attempt in the free-flight algorithm to sequence aircraft optimally.

Two variations of free flight are studied: RTA and FCFS. The RTA method routes individual aircraft based on an optimized schedule at the MF that maximizes throughput given the minimum separation requirement. Operationally, this method has two drawbacks: 1) routes must be generated in advance of the aircraft arriving at the 200 n mile range, and 2) flights are given MF RTAs at the 200 n mile range to which they must adhere. The FCFS method is less restrictive, generating routes for the aircraft as they arrive at the 200 n mile range. Both free-flight solutions implement horizontal maneuvers only for resolving aircraft conflicts (no vertical maneuvers were used in our study). Whereas free flight allows more freedom for the flight operators, there is a potential for loss in overall system capacity.

Algorithmically, FCFS free-flight route synthesis is treated as a special case of the FBRP algorithm. A free-flight route is essentially a flow-based route for a flow of zero duration that starts at the same time and position as the arriving aircraft; the zero length (duration) of the flow is based on the (approximate) model of the aircraft as being a point, having no physical size. The same search concepts of the FBRP algorithm are applied: a search space of interconnected grid points is searched in stage 1, using weather and angular constraints, and then the resulting path is simplified in stage 2. Each of these steps uses the same algorithms and code as the FBRP.

RTA flights are treated slightly differently. Instead of being routed from an arrival point o_D on the range ring to a MF, they are routed in reverse from the MF to a point near o_D . By computing routes by searching backwards in time, from the MF backwards towards o_D , we guarantee that each individual route arrives at the MF at the RTA. This does, however, mean that the time of arrival at the range ring is not known in advance but is, in essence, back-propagated from the RTA at the MF. Aside from being routed in reverse, the core algorithm is unchanged from FCFS and FBRP.

The free-flight method may use either fixed or variable MFs. When varying the MF, the free-flight method uses one of nine possible discrete MF locations along an arc centered at the original MF location.

V. Comparison

In this section, we first describe the weather data used in our comparison, then the metrics, and finally we describe the results of comparing the three solution approaches.

A. Weather Data

National convective weather diagnostic (NCWD) data on a 4 × 4 km grid with a 5 min update cycle were used in this study. Weather data were selected from a variety of real weather days (22 May 2002, 26 June 2002, and 27 June 2002) for a total of 36 h of data. Additionally, synthetic weather data were generated based on processing the real weather data with a simple expansion algorithm to amplify severe weather data by expanding these regions by a specified amount $\varepsilon = 1, 2, \dots, 10$ n mile in eight neighboring pixel directions. By creating synthetic weather data, we were able to study severe weather coverage values up to 100% coverage to test the limits of each route planning algorithm. Experiments spanned time periods from 12:00 Zulu to 0:00 Zulu (8 a.m. to 8 p.m. local time) spanning periods of no weather to periods that complete a full cycle of growth and decay of weather cells as storms pass over the ATL airport and surrounding airspace. In general, the convective organization of the

real weather data studied ranged from isolated single and multicellular storms to some storms with linear organization: this is often the situation in the southeastern portion of the United States.

B. Metrics of Comparison

Two metrics of comparison are used. We recorded throughput as a function of weather severity and airspace complexity as a function of weather severity.

Throughput. Throughput is measured at MFs only, computed as the number of aircraft crossing the MF (calculated at the closest point of approach to the MF) over a fixed period of time (15 min). The maximum throughput is $\tau_{\max} = vt_s/60\delta$, where v is the velocity (n mile/h) at the MF (assumed constant), δ is the minimum separation (5 n mile), and t_s is the sampling period (min).

Airspace Complexity. Airspace complexity is often measured in terms of dynamic density. As stated by the RTCA [2], dynamic density is described as the “essential factor affecting conflict rate in both the en route and terminal airspace.” These factors are traffic density, complexity of flow, and separation standards. In several investigations on dynamic density, the relative importance of factors affecting dynamic density were determined. These investigations [20–22] typically determine dynamic density as a linear combination of multiple factors. However, reviews of the literature [23,24] do not report any single agreed-upon model for dynamic density.

Most studies consider only a few of the top ranking factors in a dynamic density measurement for airspace complexity:

1) Density $\rho = N/A_{\text{ref}}$ of aircraft, which is the number N of aircraft per reference area A_{ref} . The reference area is typically defined to be a sector, center, or a circular region.

2) Average proximity of neighboring aircraft δ_{NN} (the average nearest neighbor distance): $\delta_{\text{NN}} = \frac{1}{N} \sum_{i=1}^N \min_{j \neq i} \{d_{ij}\}$, where d_{ij} is the distance between aircraft i and aircraft j at time t .

3) Average points of closest approach (PCA) distribution: $\delta_{\text{PCA}} = \frac{1}{N} \sum_{i=1}^N \text{PCA}_i$, where PCA_i is the closest approach distance between aircraft i and j , over a look-ahead window of time.

Studies often use one of these factors or some linear combination of these factors, with the relative weighting determined by various issues relevant to the study.

In our study, all of the aircraft in a quadrant (with respect to an airport as the origin) are heading to the same arrival gate. The algorithms are designed to maximize MF throughput, so we expect the density of aircraft to approach a maximum as the ATM methods are pushed to their limits. Further, we are comparing methods of flow management for arrivals only or departures only, so there are not mixed arrival and departure traffic in our experiments. Because of these issues, metrics from the literature that focus on density or proximity to neighboring traffic will not add discriminative value to our research. Density is designed to be maximized in our work, and proximity is forced to be at the separation limit as aircraft pass over the MF.

The complexity metric we use for our experiments is defined on a grid in the airspace A , with the complexity associated with a grid point p defined in terms of a weighted average of the variance of the velocity vectors of aircraft in the neighborhood of p , scaled according to distance from p . More specifically, for each aircraft a_i , a scaling factor $s_i(p) = \{\max[R - d(a_i, p), 0]\}/R$ is defined according to its distance from each grid point $p \in A$ for some specified radius R centered at p , where $d(a_i, p)$ is the distance from aircraft a_i to point p . This scaling is applied so that the contribution of an aircraft in the complexity metric falls off linearly, from 1 to 0, with distance from p , up to the radius R . Technically, any function monotonically decreasing to zero (such as a Gaussian function) would achieve a similar result, although not all would work equally well (e.g., a metric that gives equal weights to all flights within a particular radius will tend to produce complex discontinuities in the complexity metric). Flights for which $d(a_i, p) > R$ are not considered in the calculation of the complexity metric for point p . The average contribution of all aircraft is computed, with aircraft a_i 's contribution scaled according to $s_i(p)$. For an instant in time t , define the average local velocity and variance of velocity as follows. Let

$V_{\text{avg}}(p, t)$ be the local average velocity vector in the neighborhood of grid point p , at time t , scaled according to the factors $s_i(p)$:

$$V_{\text{avg}}(p, t) = \frac{\sum s_i(p) v_i(t)}{\sum s_i(p)} \quad (2)$$

where $v_i(t)$ is the velocity (vector) of aircraft a_i at time t . The (scalar) quantity $\|v_i(t) - V_{\text{avg}}(p, t)\|^2$ gives the squared deviation of the velocity of aircraft a_i from the local average velocity vector in the neighborhood of grid point p ($\|u\|$ is the Euclidean length of vector u). Intuitively, the larger this quantity, the more variation there is in the velocity vectors, as contributed by aircraft a_i , in the neighborhood of point p . Summing over all aircraft, and scaling by $s_i(p)$ to account for the distance from point p , we obtain the expression for the scaled-contribution velocity variance at point p , at time t :

$$\text{Var}(p, t) = \sum s_i(p) \|v_i(t) - V_{\text{avg}}(p, t)\|^2 \quad (3)$$

Our complexity metric is based on a linear combination of this variance term and a density term, $N(p, t)$, where $N(p, t)$ is the number of aircraft at time t within a specified radius R of grid point p . The overall composite complexity metric takes into account velocity variation and scaled density:

$$C(p, t) = \lambda_1 \sum s_i(p) \|v_i(t) - V_{\text{avg}}(p, t)\|^2 + \lambda_2 N(p, t) \quad (4)$$

In our experiments, we use $\lambda_1 = 0.36$, $\lambda_2 = 2$, and $R = 35$ n mile. The overall complexity at time t in a given region (e.g., quadrant) of airspace is obtained by summing $C(p, t)$, over all grid points p within the region. The time t is discretized into 1-min time intervals to get a better time-averaged view of complexity. We sum the complexity values over 15-min intervals to get the complexity values reported next.

Weather Severity. Weather severity is defined for each of the four quadrants Q_j ($j = 1, 2, 3, 4$) around an airport. If the weather severity level for grid point p , at time t is given by $L(p, t)$, then the weather severity, $\varepsilon_j(t)$, at time t , for quadrant Q_j in the vicinity of the airport is given by:

$$\varepsilon_j(t) = \frac{\sum_{p \in \bar{Q}_j} \max_{\tau \in [t-\bar{t}, t]} \{L(p, \tau)\}}{\text{Area}(\bar{Q}_j)} \quad (5)$$

where the summation is over all grid points $p \in \bar{Q}_j = Q_j \cap \alpha(r_1, r_2)$ in quadrant Q_j that lie in the annulus $\alpha(r_1, r_2)$ centered at the airport of inner radius r_1 and outer radius r_2 . The radii used in the experiments are $r_1 = 40$ n mile and $r_2 = 60$ n mile; this choice allows the weather severity to depend on hazardous weather within a reasonable neighborhood of the MF location. The time aggregation of the weather (indicated by the “max” over the time interval $[t - \bar{t}, t]$) allows an instantaneous event (e.g., an aircraft crossing the MF) to be mapped to the severity of weather that an aircraft has encountered over its flight history. The parameter $\bar{t} = 30$ min has been chosen to approximate the average flight time from the 200 n mile range entry point to the MF.

C. Comparison of Algorithms

Each approach for weather avoidance routing [historical, STAR, FBRP, and free flight (FCFS, RTA)] is compared based on 1) throughput as a function of weather severity and 2) airspace complexity as a function of weather severity. In general, the computation of any route or flow takes between 0.1 and 4 s,

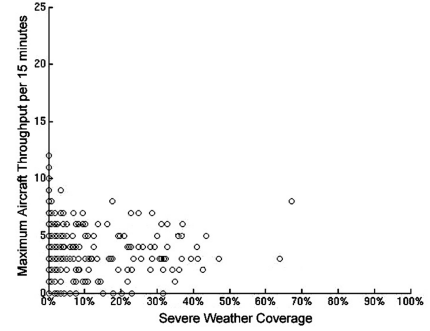


Fig. 9 Historical throughput rates vs severe weather coverage (ATL airport, ETMS data).

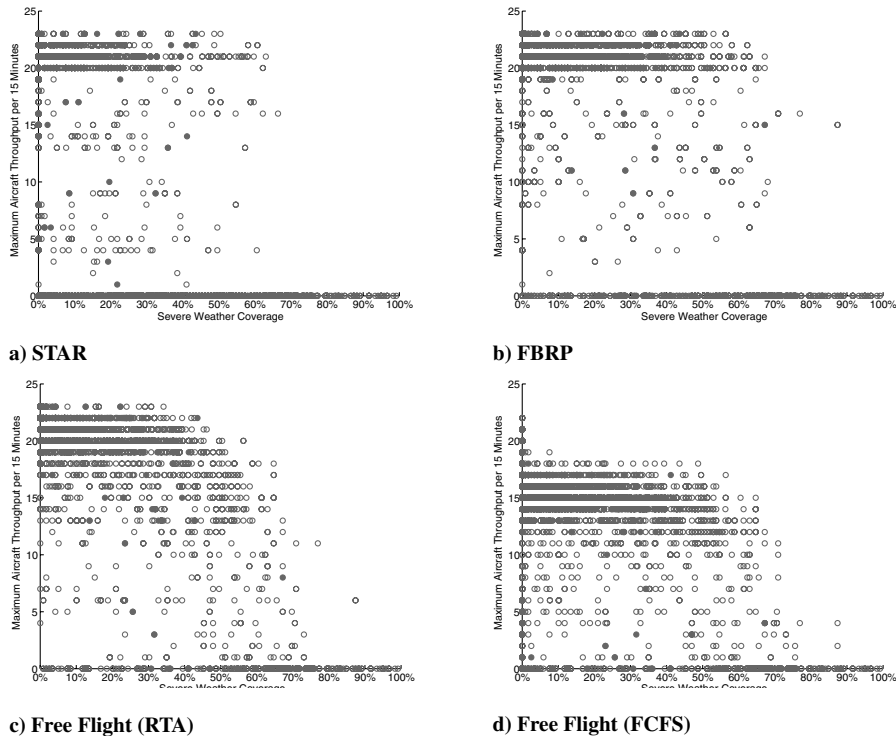


Fig. 10 Maximum realizable throughput rates at fixed MFs ($\varepsilon = 1, 15$ min updates).

averaging approximately 1 s (on a 2-GHz Athlon with 1.5 gigabytes of RAM). Under any of these models, the transition airspace for a given airport could be managed in real-time on a single modern personal computer.

Figure 9 shows historical throughput data for ATL, and how the average throughput decreases with increasing severe weather coverage. Each point represents a 15-min time interval during the experiment period in one of four quadrants centered at ATL. The actual historical throughput rates are dependent on demand. In clear weather, periods with lower demand will clearly exhibit lower throughput. We observe periods with higher weather coverage experiencing lower average throughput and conclude there is some combination of lowered demand and lowered capacity in the transition airspace during these periods. Using our algorithms, we intend to restore capacity to the transition airspace, thereby increasing demand during periods of severe weather coverage.

Each algorithm was applied to identical weather scenarios corresponding to the historical weather data set. The maximum throughput was computed for each time period in the experiment scenario, as shown in Fig. 10. Each data point represents a 15-min time interval during the experiment period in one of four quadrants centered at the airport, exactly correlated with the measurements of historical throughput rates (Fig. 9). The maximum throughput computed with algorithmically generated routing is shown in Fig. 10 for real and synthetic weather data spanning zero to 100% severe weather coverage. Both fixed and variable MF locations were used in experiments, although the experiments with variable MF locations did not generate much of a difference from those shown in Fig. 10.

Because the speed profiles varied for each of the four arrival MFs, there are four distinct throughput bands representing maximum capacity (with 5 n mile MIT) for the flow-based techniques. These bands are less pronounced for the free-flight (RTA) method. They are not discernible for the free-flight (FCFS) method because the throughput does not reach levels near maximum capacity.

Some maximum throughput rates are zero. There are two causes of this phenomenon. First, the MF may be impacted by hazardous weather, thereby blocking any routes in that quadrant. The actual historical data show cases where aircraft may either penetrate, or violate the separation from weather that we have imposed on the algorithmic solutions, thus we see historical aircraft registering

higher throughput in some of these cases. Alternately, this phenomenon is a result of the nature of flow-based routing, and is an indication that a single fixed route cannot be created that will be feasible for the full 40-min time period (approximately) required to maintain a flow of traffic through the transition area.

Given its greater flexibility, the FBRP method has far fewer time periods with zero throughput when compared with the variable STAR approach. Again, whereas feasible routes could not be generated in these cases, a few historical flights managed to traverse the transition airspace. This is either because they penetrate weather that the algorithms are forced to avoid, or because a single aircraft may safely pass through the weather which an entire flow could not. Essentially, the explanation for the cases of reduced throughput compared with historical data is that the historical flights are less constrained than the algorithmic approaches.

With greater flexibility comes greater throughput; this is evident by the flexibility of the routing method as well as by the selection of the hazardous weather safety margin ε , as illustrated in Fig. 11. Among the flow-based methods, starting with the variable STAR approach (fixed MF), to variable STAR (variable MF), to FBRP (fixed MF), to FBRP (variable MF), the number of periods with improvement over historical throughput rates steadily increases as the hazardous weather safety margin ε is reduced to zero.

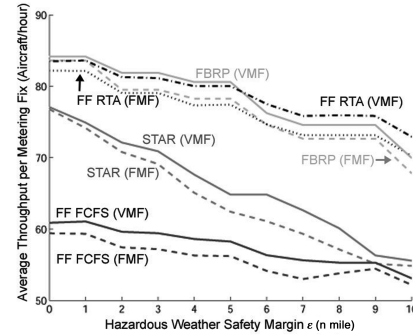


Fig. 11 Average throughput with fixed MFs (FMF) and variable MFs (VMF) (15 min updates).

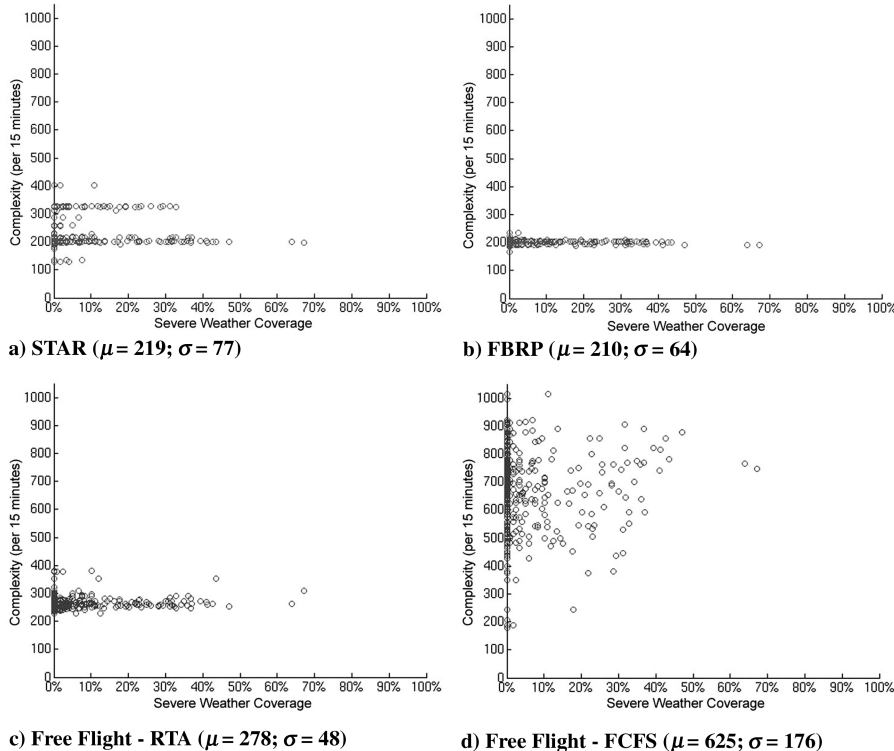


Fig. 12 Complexity metric for maximum throughput rates (real weather data, FMF method, $\varepsilon = 1$, 15 min updates).

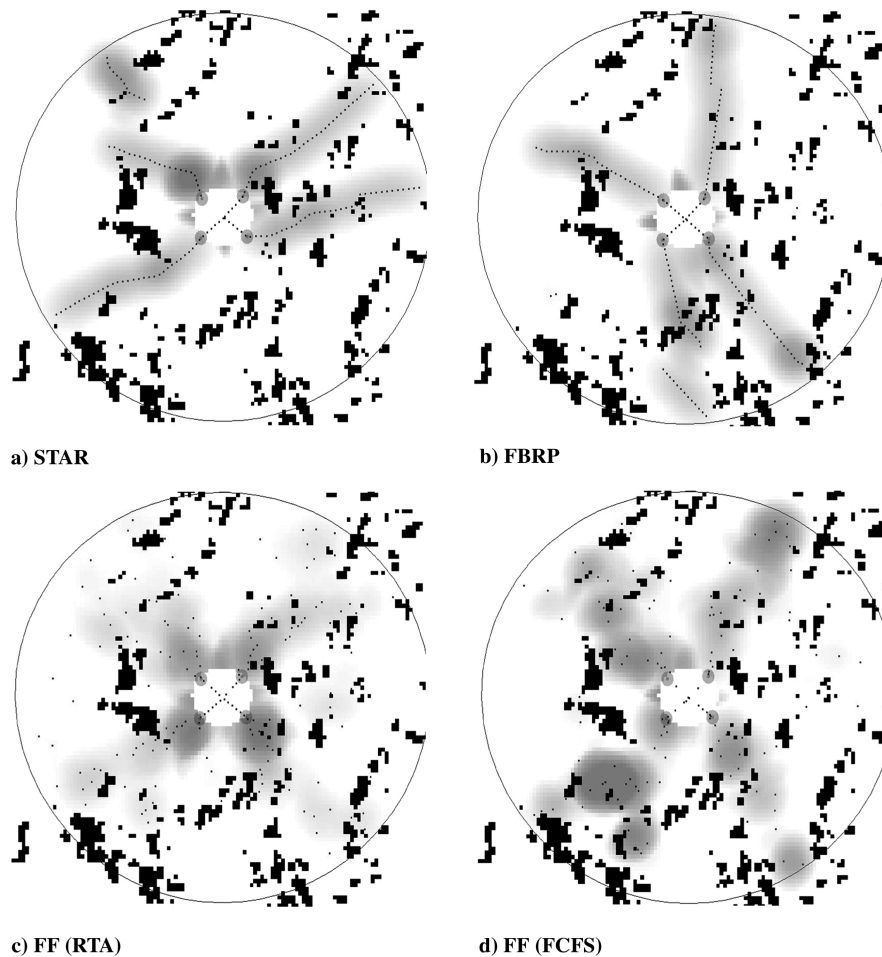


Fig. 13 Example complexity metric (gray) for maximum throughput rates ($\epsilon = 1, 15$ min updates; black indicates hazardous weather).

Free-flight methods show varying performance. Free flight (RTA) makes maximum use of arrival slots at the arrival MF. The limiting factor for this method compared with flow-based techniques was the potential for conflicts with other aircraft routed through the airspace.

Free flight (FCFS) experienced frequent underuse of capacity at the MF. Because the flights were routed as they arrived at the 200 n mile range, with no organization of how or when to arrive, minimum separation at the MF could not typically be achieved. Because we required speed to be constant for these experiments, sequencing could only be achieved through vectoring aircraft. Increased vectoring would result in greater constraints for all following aircraft attempting to generate routes to the arrival MF.

Whereas greater flexibility in routing results in an increase in maximum throughput, it also increases airspace complexity. As shown in Fig. 12, the free-flight methods show the greatest degree of complexity, as the flow of aircraft was designed to be maximized, not organized. The flow-based techniques have by design a greater degree of organization, and therefore lower complexity. The complexity values represent a composite complexity, combining a density term with a variance in velocity vector term, as described earlier. The density terms are comparable across methods (in rough proportion to the throughputs), resulting in a “base” complexity value around 185–190. Complexity values above this are due to the velocity vector term, which measures a level of disorganization between routes. Free flight (FCFS) shows the highest complexity, due to the high level of disorganization in routing. The bimodal effect in the STAR method arises because one of the quadrants has a STAR route that has a relatively sharp turn, causing a consistent addition to the variance in velocity term for all flights flying in that quadrant.

In Fig. 13, plan view displays of aircraft and complexity are shown for a given instant in time. With flow-based techniques, complexity is uniformly distributed around the defined flows, and greater

complexity may be observed at waypoints that require heading changes. Complexity is lower for the free-flight (RTA) approach far from the airport, where aircraft are scattered with sufficient distance from each other. However, complexity increases as aircraft funnel into the arrival MF. Free flight (FCFS) experiences the greatest complexity, due to the vectoring that must be performed to achieve orderly flow at the MF that also conforms to separation standards.

VII. Conclusions

Three algorithmic weather avoidance routing strategies have been compared. Compared with today’s routing practices, these methods demonstrate improved throughput with increased safety during hazardous weather events in the transition airspace. The close proximity of severe weather to critical resources (e.g., airport, arrival and departure fixes, and jet routes) limits throughput, and therefore overall capacity. Making national airspace system resources such as routes, arrival fixes and departure fixes flexible, predictable, and consistently available facilitates maintained capacity and increased throughput during severe weather events.

A tradeoff exists between centralized flow-based routing and the distributed free-flight solutions. Flow-based techniques are less computationally intensive, as a single synthesized route may apply to many aircraft. A controlled version of free-flight (using required times of arrival for aircraft) can perform on par with the highest performing flow-based techniques. However, free flight requires routes to be generated for each aircraft that simultaneously avoid hazardous weather and aircraft conflicts, therefore requiring greater computational time. Further, free-flight techniques exhibit greater complexity than flow-based techniques, and therefore imply greater workload to monitor.

Acknowledgments

This research was funded by NASA Ames Research Center under contract NAS2-02075 for the virtual airspace modeling and simulation (VAMS) project. The guidance of M. Jardin at NASA Ames Research Center is greatly appreciated. Additionally, J. S. B. Mitchell acknowledges support from the National Science Foundation (CCR-0098172, CCF-0431030, CCF-0528209) and NASA Ames Research Center (NAG2-1620).

References

- [1] Federal Aviation Administration, *2003 Aviation Capacity Enhancement Plan*, Federal Aviation Administration Office of System Capacity, Washington, DC, Dec. 2003.
- [2] RTCA, *Report of the RTCA Board of Directors' Select Committee on Free Flight*, RTCA, Inc., Washington, DC, Jan. 1995.
- [3] Krozel, J., Lee, C., and Mitchell, J. S. B., "Estimating Time of Arrival in Heavy Weather Conditions," AIAA Paper 1999-4232, Aug. 1999.
- [4] Dixon, M., and Weiner, G., "Automated Aircraft Routing Through Weather-Impacted Airspace," *5th International Conference on Aviation Weather Systems*, American Meteorological Society, Boston, 1993, pp. 295–298.
- [5] Krozel, J., Weidner, T., and Hunter, G., "Terminal Area Guidance Incorporating Heavy Weather," AIAA Paper 97-3541, Aug. 1997.
- [6] Lanthier, M., Maheshwari, A., and Sack, J.-R., "Approximating Weighted Shortest Paths on Polyhedral Surfaces," *Proceedings of the 13th Annual ACM Symposium on Computational Geometry*, Association for Computing Machinery, New York, 1997, pp. 274–283.
- [7] Mata, C., and Mitchell, J. S. B., "A New Algorithm for Computing Shortest Paths in Weighted Planar Subdivisions," *Proceedings of the 13th Annual ACM Symposium on Computational Geometry*, Association for Computing Machinery, New York, 1997, pp. 264–273.
- [8] Mitchell, J. S. B., and Papadimitriou, C. H., "The Weighted Region Problem: Finding Shortest Paths Through a Weighted Planar Subdivision," *Journal of the Association for Computing Machinery*, Vol. 38, No. 1, Jan. 1991, pp. 18–73.
- [9] Richbourg, R. F., Rowe, N. C., Zyda, M. J., and McGhee, R. B., "Solving Global, Two-Dimensional Routing Problems Using Snell's Law and A* Search," *IEEE International Conference on Robotics and Automation*, Institute of Electrical and Electronics Engineers, Piscataway, NJ, 1987, pp. 1631–1636.
- [10] Rowe, N. C., and Richbourg, R. F., "An Efficient Snell's Law Method for Optimal-Path Planning Across Multiple Two-Dimensional, Irregular, Homogeneous-Cost Regions," *International Journal of Robotics Research*, Vol. 9, No. 6, Dec. 1990, pp. 48–66.
- [11] Mitchell, J. S. B., "An Algorithmic Approach to Some Problems in Terrain Navigation," *Artificial Intelligence*, Vol. 37, Nos. 1–3, 1988, pp. 171–201.
- [12] Rhoda, D. A., and Pawlak, M. L., "The Thunderstorm Penetration/Deviation Decision in the Terminal Area," *American Meteorological Society's 8th Conference on Aviation, Range, and Aerospace Meteorology*, American Meteorological Society, Boston, 1999, pp. 308–312.
- [13] DeLaura, R., and Evans, J., "An Exploratory Study of Modeling En Route Pilot Convective Storm Flight Deviation Behavior," *American Meteorological Society's 12th Conference on Aviation, Range, and Aerospace Meteorology*, American Meteorological Society, Boston, 2006.
- [14] Love, W. D., Arthur, W. C., Heagy, W. S., and Kirk, D. B., "Assessment of Prediction Error Impact on Resolutions for Aircraft and Severe Weather Avoidance," AIAA Paper 2004-6265, Sept. 2004.
- [15] Nilsson, N. J., *Principles of Artificial Intelligence*, Tioga Publishing Company, Palo Alto, CA, 1980, Chap. 2.
- [16] Prete, J., and Mitchell, J. S. B., "Safe Routing of Multiple Aircraft Flows in the Presence of Time-Varying Weather Data," AIAA Paper 2004-4791, Aug. 2004.
- [17] Dijkstra, E. W., "A Note on Two Problems in Connexion with Graphs," *Numerische Mathematik*, Vol. 1, 1959, pp. 269–271.
- [18] Klosowski, J., Held, M., Mitchell, J. S. B., Sowizral, H., and Zikan, K., "Efficient Collision Detection Using Bounding Volume Hierarchies of *k*-DOPs," *IEEE Transactions on Visualization and Computer Graphics*, Vol. 4, No. 1, 1998, pp. 21–36.
- [19] Chiang, Y.-J., Klosowski, J. T., Lee, C., and Mitchell, J. S. B., "Geometric Algorithms for Conflict Detection/Resolution in Air Traffic Management," *36th IEEE Conference on Decision and Control*, IEEE Control Systems Society, Piscataway, NJ, 1997, pp. 1835–1840.
- [20] "An Evaluation of Air Traffic Control Complexity," Final Report, NASA Contract No. NAS2-14284, Wyndemere Corp., Boulder, CO, Oct. 1996.
- [21] Laudeman, I., Shelden, S., Branstrom, R., and Brasil, C., "Dynamic Density: An Air Traffic Management Metric," NASA TM-1998-112226, April 1998.
- [22] Sridhar, B., Kapil, S., and Grabbe, S., "Airspace Complexity and its Application in Air Traffic Management," *2nd USA/Europe Air Traffic Management R&D Seminar*, 1998, http://atm-seminar-98.eurocontrol.fr/paper_013 [cited 13 July 2006].
- [23] Mogford, R., Guttman, J., Morrow, S., and Kopardekar, P., "The Complexity Construct in Air Traffic Control: A Review and Synthesis of the Literature," U.S. Dept. of Transportation, Report DOT/FAA/CT-TN95/22, Federal Aviation Administration, Atlantic City, NJ, July 1995.
- [24] Rodgers, M. D., Mogford, R. H., and Mogford, L. S., "The Relationship of Sector Characteristics to Operational Errors," *Air Traffic Control Quarterly*, Vol. 5, No. 4, 1997, pp. 241–263.