

Robust Dynamic Routing of Aircraft under Uncertainty¹

Arnab Nilim, Department of EECS², University of California, Berkeley

Laurent El Ghaoui, Department of EECS², University of California, Berkeley

Vu Duong, Eurocontrol Experimental Centre, Bretigny-sur-Orge, France

Summary

Much of the delay in the US National Airspace System (NAS) arises from convective weather. One major objective of our research is to take a less conservative route, where we take a risk of higher delay to attain a better expected delay, instead of avoiding the bad weather zone completely. We address the single aircraft problem using a Markov decision process model and a stochastic dynamic programming algorithm, where the evolution of the weather is modeled as a stationary Markov chain. Our solution provides a dynamic routing strategy for an aircraft that minimizes the expected delay. A significant improvement in delay is obtained by using our methods over the traditional methods. In addition, we propose an algorithm for dynamic routing where the solution is robust with respect to the estimation errors of the storm probabilities. To the Bellman equations, which are derived in solving the dynamic routing strategy of an aircraft, we add a further requirements: we assume that the transition probabilities are unknown, but bounded within a convex set. The uncertainty described in our approach is based on likelihood functions. This makes the robust Dynamic Programming (DP) "tractable" (that is, not much more complicated than the original DP) and yet not conservative. Our algorithm optimizes the performance of the system, given there are errors in the estimation of the probabilities of the storms.

Overview

Air traffic delay due to convective weather has grown rapidly over the last few years. According to the FAA, flight delays have increased by more than 58 percent since 1995, cancellations by 68 percent. The delay distribution can be described in the figure below:

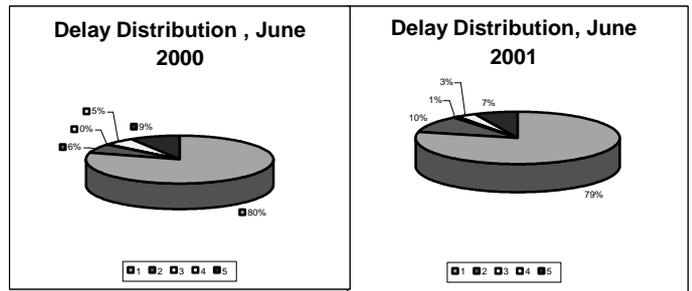


Figure 1: Air Traffic delay distribution in US in the month of June, 2000 and 2001. 1:Weather 2:Volume 3: Equipment 4:Runway 5.Other.

It is obvious from the above graph that weather (which is around 80% of the total delay) is the main contributor for the delay in the US Air Traffic Control system. The cost of delays to airlines and passengers are billions of US dollars per year. The air traffic flow management (TFM) problem under deterministic environment is a well addressed problem [2], [3], [4], [13], [14], [15]. As the major portion of this delay is due to bad weather, which is not deterministic in nature. As a consequence, TFM problems cannot be addressed in a deterministic setting.

¹Research supported by Eurocontrol -014692 and DARPA-F33615-01 C-3150

²Electrical Engineering and Computer Sciences

In the National Airspace System, Traffic flow management (TFM) decisions for release and Federal Aviation Administration controllers and airline dispatcher make routing of aircraft fleets 3-4 hours in advance of the actual operation. Knowledge of the location and the intensity of the hazardous convective weather 3-4 hours ahead is key to select air routes. Since, weather is stochastic in nature, there is an urgent need for automation tools, which explicitly deals with the random dynamics of the storms and provides solutions that reduce the expected delay in the air traffic control system. A dynamic routing strategy is provided in [1], where the expected delay is minimized for an aircraft. It is shown in [1] that a significant improvement in delay is be obtained by using Dynamic Programming algorithms over the traditional methods.

However, the algorithm in [1] is based on a very strong assumption: the storm probabilities are exactly known and we can have the perfect point estimates of transition probabilities. In this paper, we will see that the optimal routing of aircraft is very sensitive to these storm probabilities. A slight mistake in estimating the storm transition probabilities will change the optimal solution significantly. So the algorithm can work well if the exactly accurate probabilities of storms are obtained.

In the recent years, FAA, the national weather service (NWS), and airline operations centers have collaborated to produce the best convective forecast available under the current forecasting state of art [16]. These forecast, named as Collaborative Convective Forecast Product (CCFP), are improvements over other conventional products but are far from perfect forecast. According to the real time verification statistical techniques being employed by National Oceanic and Atmospheric Administration Forecast Systems Laboratory, the current forecast typically verifies at 0.75 False Alarm Rate and 0.28 probability of detection. For any kind of strategic or tactical route planning, FAA believes a maximum False Alarm Rate of 0.20 and a minimum probability of detection of 0.80 is highly desirable. For example, some of the days when the predicted and actual weather varied significantly are given below:

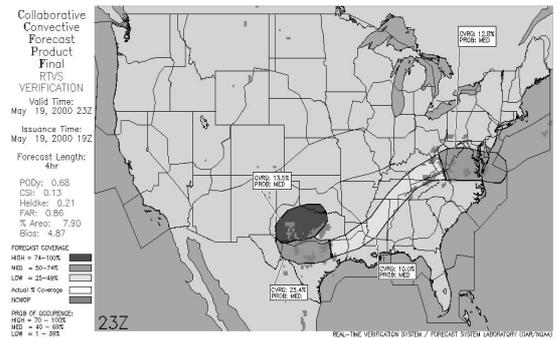


Figure 2: CCFP vs. Actual weather on May 19, 2001. Yellow polygons predict the probability of weather occurring at forecast time 20-50%, with a predicted coverage of 25-49%. A huge airspace was unused due the prediction, which actually turned out to completely good.

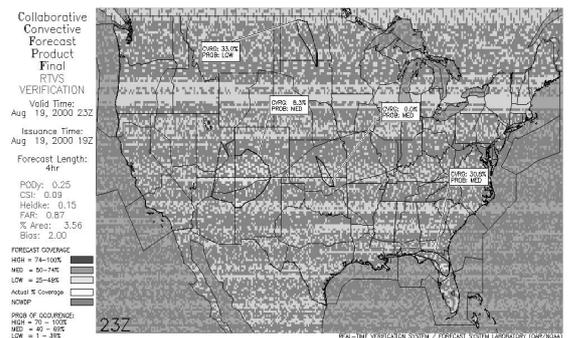


Figure 3: Example where actual weather was much more widespread than the forecast

The optimal solution derived in [1] is potentially sensitive with the estimation error of the transition probabilities and a very accurate estimation of storm probabilities is required to obtain a result that can be used in practice. In this paper, we present a robust Dynamic Programming approach, where the algorithm optimizes the aircraft routing of the system, despite there are errors in the estimation of the probabilities of the storms.

Problem Set-up

We consider a two-dimensional flight plan of an aircraft. We are interested in finding the optimal path of a single aircraft in the en-route portion. There are inherent uncertainties in the En-route portion of the flight. We address the problem as in a decentralized fashion, as a large-scale stochastic dynamic programming problem. Consider the following scenario: an aircraft is flying from the origin to the destination. There are many obstacles that might interfere with the shortest possible route. These obstacles can be both stochastic and deterministic in nature. Examples of the stochastic "no zones" include storm zones, intersections of the flight plan of other aircraft, strong wind zones and the deterministic "no zones" include military and national security airspace. Given those constraints, the optimal route of the aircraft is the one that provides the minimal cost (time, fuel) while avoiding all of the obstacles.

In the optimization problem, it is easier to deal with the deterministic "no zones", as we can assign them an infinite cost if aircraft were to penetrate through the zones. But crux of the weather problem is to deal with the stochastic "no zones". In the current practice, those stochastic zones are assumed to be unusable, and solution proceeds as if they are deterministic constraints. As those zones were just predicted to be of unusable with a certain probability, it often turns out that the zones were perfectly usable. The rerouting strategies do not use these resources and as a result, the airspace resources are under-utilized, leading to congestion in the remaining airspace.

Various weather teams (CCFP, ITWS etc) produce predictions that some zones in the airspace will be unusable in certain time and their predictions are dynamically updated with time. However, as described in the previous section, those predictions are often incorrect which makes it even harder to select an optimal routing.

In our proposed model, we will not exclude the zones, which are predicted to be unusable (with some probability) at a certain time. Instead, we will assign some cost to those zones, a cost that will depend upon the state of the system. Our solution will take into consideration the fact that there will be more updates with the course of flight and

recourse will be applied accordingly [10 11 12]. Moreover, we will analyze the previous data and count the number of times it goes from each state to different states. From this counting, we form a likelihood constraint, which describes the uncertainty set within which the transition matrix is bounded. We take a route that will provide us the best-expected delay, where the time varying transition probability matrix bounded within the uncertainty set is acting as an opponent.

For this class of problems, we look for the "best policy", not the "best path". Determining the "best policy" is deciding where to go next given the currently available information. We consider the set of decisions facing an aircraft that starts moving towards the destination along a certain path, with the recourse option of choosing a new path whenever new information is obtained.

Weather Model

Various weather teams provide the probability of a storm at a particular place at a particular time. The weather information is updated in about every 15 minutes. The further away the prediction is from the event, the more unreliable it becomes and vice versa. We can discretize the time in a number of 15 minutes time intervals. From the weather science, we can assign a probability " P_r " or " Q_r " for a particular region such that,

P_r = probability (there will be a storm in that region in the next 15 minutes time interval/there is a storm in the current time in the region), or

Q_r = probability (there will be no storm in that region in the next 15 minutes time interval/there is no storm in the current time in the region)

It is also realistic to assume that the aircraft has a perfect knowledge about the regions that are 15 minutes (15 times the velocity of the aircraft provides the distance) away from it. Whenever an aircraft is in 15 minutes away from the storm region, the pilot knows for sure whether there is a storm in that region or not. The probability of storm at a particular region is time varying and takes a value 0 or 1 when it reaches 15 minutes away from the region. If P is defined as the probability of the storm in the next interval, the dynamics of P will follow a path described below,

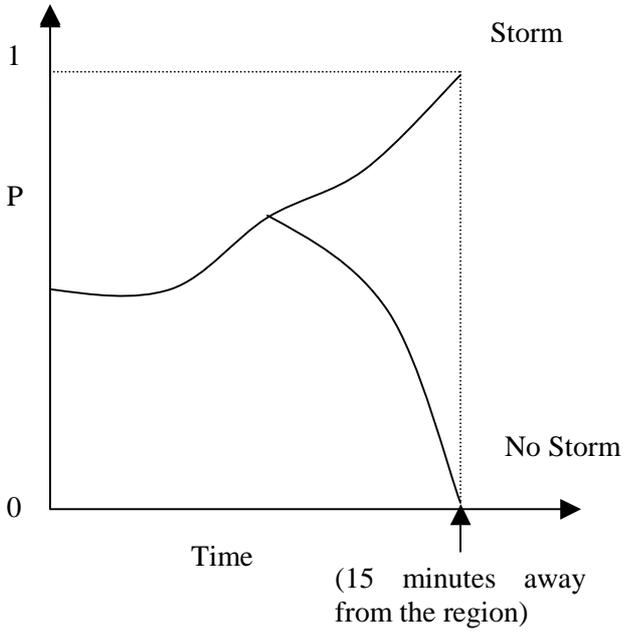


Figure 4: Variations of the true probability of storm with time

We can discretize time as $1, 2, \dots, n$ stages, where "1" corresponds to the time 0-15 minutes from the current time, "2" corresponds to the time 15-30 minutes from the current time and so on. Let, T is the time required to go from the origin to destination in the worst possible routes (where the worst possible route is defined as the route where nominal path is followed assuming no storm and the recourse is applied just before the storm zone to avoid the storm). No policy will take more time than T . Total number of stages can be calculated by the following formula,

$$n_{\max} = \frac{T - \text{mod}(\frac{T}{15})}{15} + 1 \quad (1)$$

Suppose there are $1, 2, \dots, m$ storms that are predicted to happen that might force the aircraft to deviate from its nominal path. We can define state '1' corresponding to the state as having storm in a region at a particular stage and state '0' corresponding to state as having no storm at a particular stage in that region. As we know the status of any storm in the time interval of 0-15 minutes (stage 1), we can assign 0(1) to every storm at the stage 1. Again we also know the conditional probability of having (not having) a storm in a

region in the next 15 minutes time interval (stage 2), given there is a storm (no storm) in the current time (stage 1) in the region. If we assume that the probability of storm in a particular zone varies in a Markovian fashion, we can represent the transition model of each storm in the following manner.

We can define, p_1 is the probability of no storm in the next stage if there is no storm in the current stage in a particular region and q_1 is the probability of having a storm in the next stage if there is a storm in the current stage in that particular region.

In the same way, we can define the probabilities $p_2, q_2, p_3, q_3, \dots, p_m, q_m$. If there are m storms, it will be a 2^m state Markov chain. A 2^m tuple vector can describe the storm situation completely, i.e., $[1 \ 0 \ 0 \ 0 \dots 0]$ denotes that there is a storm in zone '1' and the rest of the predicted bad weather zones are storm free.

We can define $S = \{ \mathbf{1}, \mathbf{2}, \dots, \mathbf{M} \}$ as the set of all the states where, $\mathbf{1} = [0 \ 0 \ 0 \ 0 \dots 0]$, $\mathbf{2} = [1 \ 0 \ 0 \ 0 \dots 0]$, $\mathbf{3} = [0 \ 1 \ 0 \ 0 \dots 0]$, \dots , $\mathbf{M} = [1 \ 1 \ 1 \ 1 \dots 1]$ (2^m th state).

The transition matrix P for the Markov chain is $2^m \times 2^m$ matrix whose (i, j) th entry p_{ij} denotes the probability of ij th state in the next stage, given the current state is i . The transition matrix can be defined as follows,

$$P = \begin{bmatrix} p_{11} \dots p_{12^m} \\ \dots \dots \dots \\ p_{2^m 1} \dots p_{2^m 2^m} \end{bmatrix} \quad (2)$$

If the current state $X_1 = i$ is given, the probability that the k th state $X_k = j$ can be determined as follows,

$$p(X_k = j / X_1 = i) = P_{ij}^k \quad (3)$$

Which is the (i, j) th component of P_k matrix.

The P matrix is an input in our algorithm. We can form a P matrix in a storm situation where the storms are moving and are not probabilistically independent. But for the demonstration in this

paper, we assume that there is no movement of storms, i.e., the zones where no storms are predicted to happen are assumed to remain as perfectly usable by the aircraft over all time and storms don't expand or contract over time (relaxation of this assumption will increase the size of the probability matrix).

Uncertainty model of the transition probability

$P = (p_{ij})$ is the $|S| \times |S|$ dimensional transition probability matrix. We assume that we can observe the initial state, which is the current status of the storms in the airspace.

By exploring the past data, we obtain a matrix, which comprise the number of times each state (storm condition) makes a transition from one state to other (number of time from state i to state j is N_{ij}).

The likelihood function can then be written as

$$L(P) = \prod_{i,j} p_{ij}^{N_{ij}} \quad (4)$$

The log-likelihood function is,

$$l(P) = \sum_{i,j} N_{ij} \log(p_{ij}) \quad (5)$$

In addition, we require that the estimated probabilities are positive

$$p_{ij} \geq 0, \forall i = 1, 2, \dots, |S|, j = 1, 2, \dots, |S| \quad (6)$$

and satisfy the constraints of a transition matrix:

$$\sum_j p_{ij} = 1, \forall i = 1, 2, \dots, |S| \quad (7)$$

The maximum likelihood estimate can now be obtained by solving the convex optimization problem

$$\max_p l(P) : s.t. (6)(7) \quad (8)$$

This is a separable convex optimization problem that can be solved analytically. The optimal solution to the optimization problem is,

$$\hat{p}_{ij} = \frac{N_{ij}}{\sum_j N_{ij}} \quad (9)$$

This is the Maximum Likelihood Estimate (MLE) and the optimal value of the log likelihood function corresponding to these probabilities is defined as β_{\max} .

A classical description of uncertainty in a maximum likelihood setting is via the likelihood region is as follows,

$$\{P \in R^{(|S| \times |S|)} : P \geq 0, Pe = e, \sum_{i,j} N_{ij} \log(p_{ij}) \geq \beta\} \quad (10)$$

Where, β is chosen as $\beta < \beta_{\max}$ and e is the vector of all ones.

In our problem,, we only need to work with the un certainty on the each row p_i , that is, with projections of the set above. Due to the separable nature of the maximum likelihood problem, the projection of the above set onto the variable p_i of the matrix P can be given explicitly, as

$$P_i(\beta) \equiv \{p \in R^{|S|} : p \geq 0, p^T e = e, \sum_j N_i(j) \log p_i(j) \geq \beta_i\}, \text{ where,} \quad (11)$$

$$\beta_i \equiv \beta + \sum_{k \neq i} \sum_j N_{kj} \log N_{kj}$$

We can calculate β that will correspond to a desired level of confidence in the estimate. In order to do so, let's define a vector θ obtained by stacking the first $|S| - 1$ elements of the each row of the probability matrix. Provided some regularity conditions hold [18], it is possible to make Laplace approximation of the Likelihood function and we can make the following asymptotic statement about

the distribution of θ , that is $\theta \sim N(\hat{\theta}, I(\theta))$. That

is θ is normally distributed with the mean $\hat{\theta}$ given by (9) and $I(\theta)$ is the Fisher Information matrix given by ,

$$I(\theta) = E_{\theta} \left(-\frac{\partial^2}{\partial \theta_j \partial \theta_k} l(\theta) \right) \quad (12)$$

We can approximate $I(\theta)$ with the observed information matrix, which is meaningful in the neighborhood of $\hat{\theta}$. The equation of the observed information matrix is given by,

$$I_o(\theta)_{j,k} = -\frac{\partial^2}{\partial \theta_j \partial \theta_k} l(\theta) \quad (13)$$

This is true for large number of sample [18]. We further define, $H^{-1} = I_o(\theta)$. Then the parameter β is chosen to be the smallest such that, under the probability distribution $\mathfrak{N}(\theta, \hat{H}^{-1})$, the set,

$$\xi_\beta = \{\theta; \bar{l}(\theta) \geq \beta\}$$

where $\bar{l}(\theta)$ is the quadratic approximation to $l(\theta)$ around $\theta = \hat{\theta}$, that is,

$$\bar{l}(\theta) = \beta_{\max} - \frac{1}{2}(\theta - \hat{\theta})^T H(\theta - \hat{\theta})$$

has the probability larger than a threshold $(1 - \alpha)$, where (say) $\alpha = 5\%$ in order to obtain the 95% confidence level.

It turns out that, we can solve for such a β explicitly,

$$(1 - \alpha) = F_{\chi^2_{|S||S-1|}} \left(\frac{2(\beta_{\max} - \beta)}{\sum_{i,j} N_{ij}} \right) \quad (14)$$

where $F_{\chi^2_{|S||S-1|}}$ is the cumulative χ^2 distribution with the degrees of freedom $|S||S-1|$.

Graphical representation of the airspace

Our problem is to find out the direction of the aircraft in a way that will take into account the fact that more information will be received in the

course of the flight. The aircraft will stick to the direction till it receives another weather update. In the previous section, we have already discretized the time and divided the decision horizon in a number of stages. An aircraft will not receive any information until it goes to another stage. We simulate the variation of the probability of storm by using our Markovian model. The direction to be taken in each stage is a continuous problem, which is hard to solve. As it is a common practice in ATC process that aircraft follow some fixed waypoints, it is an acceptable formulation to discretize the airspace. We represent the airspace by a rectangular grid where the arc length is 8 n.mi. When the aircraft is at the origin, the solution of the routing problem lies within the set of grid points that can be reached by the aircraft in 15 minutes, before the aircraft receives the next update. Once, the point is decided, the aircraft is recommended to fly straight to the point and the vector direction of the aircraft can be determined. The point to be noted that the arc length is a design parameter and can be chosen by the planner. Reducing the distance between two grid points result in better solution, but it incurs more expensive calculation. Also, reducing the distance between two grid points to zero gives the continuous solution.

Markov Decision Process

For a system in which the uncertainty can be described in a Markovian manner, the stochastic control problem can be formulated in the well known Markov Decision Problem(MDP). Markov Decision Processes capture several attractive features that are important in decision-making under uncertainty: they handle risk in sequential decision-making via a state transition probability matrix, while taking into account the possibility of information gathering and recourse corresponding to this information during the multi-stage decision process [5 6 7 8 9].

If there are N states in the system and n stages to go, a policy that delivers the maximal(minimal) value is called an optimal policy. $v(i, n)$ is the value function where i is the current state of the system and n is the stages to go. $v(i, n)$ is usually unique but there can be more than one policy that can provide this value.

Dynamic programming is an inductive approach. An optimal policy for a process whose current state is i and n stages to go, must make use of optimal policies for the system with $n - 1$ step remaining. If after the initial decision and transition, the system is in state j , the original optimal policy now constitute an optimal policy for the system with the initial state j and $n - 2$ stages remaining. There are A_i alternatives out of state i on the first transition. If alternative k is pre-described, the expected gain on the initial transitions would be q_i^k and the probability of moving to state j from state i would be p_{ij}^k . If the system does in fact move to state j , from the principle of optimality, the total expected cost over the optimal policy over the last $n - 1$ stages would be $v(j, n - 1)$. Hence the total expected cost is,

$$q_i^k + \sum_{j=1}^N p_{ij}^k v(j, n - 1).$$

It follows that $n - 1$ satisfies the recursive equation,

$$v(i, n) = \min_{1 \leq k \leq A_i} \left\{ q_i^k + \sum_{j=1}^N p_{ij}^k v(j, n - 1) \right\} \quad (15)$$

Markov Decision Process Algorithm for dynamic routing of aircraft (with no uncertainty in the transition matrices)

In this section, we will present the dynamic routing of aircraft under convective weather, where the probabilities of the storms are exactly known. In the next section, we will present the robust version of the problem, i.e., dynamic routing of aircraft under convective weather, where the probabilities of the storms are not known but bounded within a set described by the likelihood function.

We are mainly concerned with the enroute part of the aircraft flight where the velocity remains almost constant (say V). So we can consider the velocity to be constant. In this way, minimizing expected delay is same as minimizing expected

distance to be traveled to go to the destination from the origin. Our objective function is the expected distance to be traveled. Decision variables are the nodes to go at the end of each stage time till it reaches the destination. We define $\delta \approx 15V$ as the distance from the storm from where the pilot knows for sure whether there is a storm or not. As previously described, the aircraft is at origin and the destination is fixed. There are m regions, labeled k_1, \dots, k_m , which are predicted to have storm so that it might obstruct the nominal path of the aircraft. We get weather update in every 15 minutes. The transition matrix P is given. Our algorithm provides the routing strategy of the aircraft in order to obtain the minimum expected distance. The algorithm is described below.

Step 1: Calculating the total number of stages

$$n_{\max}, \text{ where } n_{\max} = \frac{T - \text{mod} \frac{T}{15}}{15} + 1 \text{ and } T \text{ is the}$$

time required in the worst possible route which is defined as the route where nominal path is followed assuming no storm and the recourse is applied just before the storm zone to avoid the storm).

Step 2: Discretizing the airspace with a rectangular grid (of spacing $\lll 15$ min).

Step 3: Pruning the search space which is obtained by the convex combination of all the shortest path routes corresponding to different states.

Step 4: Determining the points that can be reached in the next 15 minutes. This can be approximately calculated if we draw an annular region with $15(V \pm \epsilon)$ as radii, with a predefined angle θ and checking which grid points fall in the region. For the first stage, the angle can be obtained from the orientation of the storms. For the next stages, the angle θ could be the maximum permissible turning angle of the aircraft. Let, (x, y) is the coordinates of the origin and $(z_1, w_1), (z_2, w_2), \dots, (z_r, w_r)$ are r such points that can be reached at the end of first stage.

Step 5: Assigning appropriate costs. Costs in our algorithm should be such that they enforce the solution will take a path through the predicted storm zone if there is no storm and avoid that if there is a

storm. Cost from going to a point from a point is a function of the state of storm and the Euclidean distance between the two points.

Define $c(i, x, y, z_j, w_j)$ as cost to go from (x, y) to (z_j, w_j) if the storm state is i .

Starting from origin (x, y) ,

Check:

For $j = 1, 2, \dots, r$

If $\{(z_j, w_j) \in (k_1, k_2, \dots, k_m)\}$ or {the straight line $\lambda(x, y) + (1 - \lambda)(z_j, w_j)$ and $(0 \leq \lambda \leq 1)$

connecting (x, y) and the point (z_j, w_j) cut any of the predicted storm zone} and {state of the intersected zone corresponds to a storm at that particular zone}

$c(i, x, y, z_j, w_j) = \{Very_High_Cost\}$

else

$c(i, x, y, z_j, w_j) = \|(x, y) - (z_j, w_j)\| \approx 15V$

endif

endfor

Proceeding this cost assignment till it reaches the destination point.

Step 6: Defining the value function for our dynamic program, $v(i, x, y, n) =$ Expected minimum distance to go if the aircraft is at the initial point (x, y) with the initial state i and it n stages to go to reach the destination point.

Step 7: Assigning the boundary value to the value functions that guarantee the desired solution.

Case 1: We have to make sure that we always get a complete path (path containing the origin and the destination) as a solution. If (p, q) are the coordinates of the destination point, the conditions below guarantee that the solution will provide a complete path.

For any state i ,

if $\{u = p\} \cap \{v = q\}$

$v(i, u, v, 0) = 0$

else

$v(i, u, v, 0) = \infty$

endif

Case 2: In the iteration process we need to put values of the value function for the points which are less than $15V$ apart from the destination point.

For any n , for any points (l, m) such that

$\|(l, m) - (p, q)\| \leq 15V$

if { There is no storm zone in the straight line $(\lambda(l, m) + (1 - \lambda)(p, q) \& 0 \leq \lambda \leq 1)$ }

$v(i, l, m, n) = \|(l, k) - (p, q)\|$ for any i

elseif { for all i corresponds to the state with no storm at that zone }

$v(i, l, m, n) = \|(l, k) - (p, q)\|$ for any i

else

$v(i, l, m, n) = \infty$

endif

Step 8: Implementing the recursive equations,

$v(i, x, y, n) = \min_{(z_1, w_1), \dots, (z_r, w_r)} V$, where $0 \leq n \leq n_{\max}$

and

$$V = \begin{pmatrix} c(i, x, y, z_1, w_1) + \sum_{j=1}^{2^m} p_{ij} v(j, z_1, w_1, n-1) \\ c(i, x, y, z_2, w_2) + \sum_{j=1}^{2^m} p_{ij} v(j, z_2, w_2, n-1) \\ \dots \\ c(i, x, y, z_r, w_r) + \sum_{j=1}^{2^m} p_{ij} v(j, z_a, w_a, n-1) \end{pmatrix} \quad (16)$$

However, we don't know the values of $v(j, z_1, w_1, n-1), \dots, v(j, z_r, w_r, n-1)$. To obtain those values, we need other recursive relations. Our calculation moves forward till we reach the boundary conditions and then we back track and calculate all of the value functions. From

this value functions and we can trace the minimum expected distance path. The aircraft will keep continue proceeding according to the solution till a new update is obtained. After receiving a new weather update, the aircraft can run this model again to obtain a new updated route with the new input (position and vector direction of the aircraft at that time, the new updated weather). In this way, an aircraft will follow a trajectory which is updated in every 15 minutes. The aircraft avoids the bad weather zone if there is actually a storm, but takes a less circuitous route if there is no storm. As a result, the expected distance traveled is minimized.

Robust Markov Decision Processes

In obtaining the optimal solution of a Markov Decision Problem with the uncertainty in the transition matrices, we propose an uncertainty model which results in an algorithm that is both statistically accurate and numerically tractable. Precisely, we develop a robustness that can be handled at very moderate additional computational cost. That means the method can be readily applied to the routing problem with negligible extra computational cost.

If we define a new variable $I = \{i, x, y\}$ as the state of the system, $a \in A$ denotes the points to go in the later stages to reach to the destination and $V_n(\cdot)$ as the vector obtained by stacking all the $v(\cdot, \dots, n)$ in all possible I in the system (there are $N_s = |S||X||Y|$ number of states), the equation 15 can be written as,

$$V_n(I) = \min_{a \in A} (c(I, a) + \sum_{J=1}^{N_s} P_{IJ}^a V_{n-1}(J)) \quad (17)$$

The complexity of the Bellman recursion is $O(|A||S|)$.

Now consider the case when the collection of transition matrices $P = (P^a)_{a \in A}$ lie in the set described by the equation (10). For a given action a , the state I , we denote by p_I^a the next state distribution drawn from P^a corresponding to the state I , thus p_I^a is the I th row of matrix P^a .

Hence, the robust counter part of the nominal Bellman Equation (16) is given by,

$$V_n(I) = \max_{P \in \mathcal{P}} \min_{a \in A} (c(I, a) + \sum_{J=1}^{N_s} P_{IJ}^a V_{n-1}(J)) \quad (18)$$

Using the standard duality arguments [17], we can show that the robust Bellman recursion can be written as,

$$\begin{aligned} V_n(I) &= \min_{a \in A} \max_{P \in \mathcal{P}} (c(I, a) + \sum_J P_{IJ}^a V_{n-1}(J)) \\ &= \min_{a \in A} (c(I, a) + \phi_{P_I^a}(V_{n-1})) \end{aligned} \quad (19)$$

Where $\phi_{P_I^a}$ denotes the support function of the convex set P_I^a .

The challenge of solving the problem reduces to computing value of the support function of the set \mathcal{P} :

$$\phi_P(v) = \max_{p \in P} v^T p \quad (20)$$

where the variable p corresponds to a particular row of a specific transition matrix, \mathcal{P} is the set describes the uncertainty on this row given by the equation (10) and v is an appropriately defined vector with non-negative components, containing the elements of the value function. We refer to the above problem as the inner problem.

Solution to the Inner Problem

The inner problem of the recursion:

$$\{\max_p p^T v : p \geq 0, p^T e = e, \sum_j N(j) \log p(j) \geq \beta\}$$

Here, we have dropped the subscript i of the empirical frequency vector N_i and in the lower bound β_i .

Since the above problem is convex and has a feasible set with non-empty interior, there is no Lagrange duality gap.

$$\begin{aligned} \min_{\lambda, \mu, \gamma} & \mu - (1 + \beta)\lambda + \lambda \sum_j N(j) \log \frac{\lambda N(j)}{\mu - v(j) - \gamma(j)} \\ \text{s.t.} & \lambda \geq 0, v(j) + \gamma(j) \leq \mu, j = 1, 2, \dots, n \end{aligned} \quad (21)$$

By monotonicity argument, we obtain that the optimal dual variable $\gamma = 0$, which reduces the number of variable to two:

$$\phi = \min_{\lambda, \mu} f(\lambda, \mu)$$

where,

$$f(\lambda, \mu) := \begin{cases} \mu - (1 + \beta) + \lambda \sum_j N(j) \log \frac{\lambda N(j)}{v(j) - \mu} & \text{if } \lambda \geq 0, \mu > v_{\max} := \max_j v(j) \\ \infty, & \text{otherwise} \end{cases} \quad (22)$$

For further reference, we note that f is twice differentiable on its domain and that its gradient is given by,

$$\nabla f(\lambda, \mu) = \begin{bmatrix} \sum_j N(j) \log \frac{\lambda N(j)}{\mu - v(j)} - \beta \\ 1 - \lambda \sum_j \frac{N(j)}{\mu - v(j)} \end{bmatrix} \quad (23)$$

From the expression of the gradient above, we obtain that the optimal value of λ for a fixed μ , $\lambda(\mu)$, is given analytically by,

$$\lambda(\mu) = \frac{1}{\sum_j \frac{N(j)}{\mu - v(j)}} \quad (24)$$

which further reduces the problem to a one-dimensional problem:

$$\phi = \min_{\mu \geq v_{\max}} \phi(\mu) \quad (25)$$

The function ϕ is convex in its domain $(v_{\max} + \infty)$. Hence, we can use the bisection to minimize ϕ . The gradient of ϕ ,

$$\nabla \phi(\mu) = \frac{\partial f}{\partial \mu}(\lambda(\mu), \mu) + \frac{\partial f}{\partial \lambda}(\lambda(\mu), \mu) \frac{d\lambda(\mu)}{d\mu} \quad (26)$$

To initialize the bisection algorithm, we need finite upper and lower bound. A lower bound

is clearly $\mu_- = v_{\max}$. If $v^- := \frac{\sum_j N(j)v(j)}{\sum_i N(i)}$, then

upper bound can be calculated as ,

$$\mu_+ = \frac{v_{\max} - e^{\beta - \beta_{\max}} v^-}{1 - e^{\beta - \beta_{\max}}} \quad (27)$$

The bisection algorithm goes as follows,

1. Set $\mu_- = v_{\max}$ and μ_+ as in (26). Let ϵ be a small convergence parameter.
2. While $\mu_+ - \mu_- > \epsilon$, repeat
 - (a) Set $\mu = (\mu_- + \mu_+) / 2$.
 - (b) Compute the gradient as in (25)
 - (c) If $\nabla \phi(\mu) > 0$, set $\mu_+ = \mu$, otherwise, set $\mu_- = \mu$
 - (d) Go to 2a

The above algorithm converges in at most $O(\log(|S|))$ iterations. Each iteration requires the computation of the function value and its gradient. This can be done at most $O(|S|)$. Hence, the inner problem under the likelihood uncertainty model can be solved at a worst-case cost of $O(|S| \log(|S|))$, which is a very moderate increase from the nominal problem $O(|S|)$.

Robust Markov Decision Process Algorithm for dynamic routing of

aircraft (with likelihood uncertainty in the transition matrices)

We will present our algorithm of the robust dynamic routing of aircraft under convective weather, where the probabilities of the storms are not known but bounded within a set described by the likelihood function. As previously described in the nominal problem, we are mainly concerned with the enroute part of the aircraft flight where the velocity remains almost constant (say V). So we are trying to minimize the expected delay. Decision variables are the nodes to go at the end of each stage time till it reaches the destination. $\delta \approx 15V$ is the distance from the storm from where the pilot knows for sure whether there is a storm or not. There are m bad zones and we get weather update in every 15 minutes. The nominal transition matrix P and the past data matrix N is given.. The algorithm is described below.

Step1-Step7: Follow the same procedure as in the nominal case.

Step 8: Instead of solving the (15), solve the equation (17) by using the algorithm described in the previous section in order to obtain the optimal nodes.

The aircraft will keep continue proceeding according to the solution till a new update is obtained. After receiving a new weather update and the corresponding past data, the aircraft can run this model again to obtain a new updated route with the new input (position and vector direction of the aircraft at that time, the new updated weather). The complexity of the algorithm is $O(|A||S|\log(|S|))$. In comparison, the complexity of the non robust(nominal) algorithm is $O(|A||S|)$. Hence, the robustness incurs a very moderate additional computational cost in the path planning algorithm. As the algorithm provides a very fast solution (for moderate number of storms), this algorithm can be implemented in real time routing of aircraft under uncertainty.

Simulation

We have implemented both of the algorithms for dynamic routing of aircraft under

uncertainty in MATLAB. One assumes a perfect prediction, which is called the nominal MDP algorithm, and the other assumes uncertainty in the prediction, which is called the robust MDP algorithm. We ran our program in a simplified scenario where the nominal route of an aircraft is obstructed by a predicted convective zone. In the scenario, an aircraft's current position is at $[0,0]^T$ and the destination is at $[360,0]^T$. All the units are in n.mi. The velocity of the aircraft is 480 n.mi/hour. There is a prediction that a storm might obstruct its nominal flight path. The storm zone is a rectangular space with the corner points at $[160,192]^T$, $[160,-192]^T$, $[168,192]^T$ and $[168,-192]^T$. We are assuming that the weather information of the portion of the airspace that can be reached in 15 minutes is deterministic. The probability of storm propagates in a Markovian fashion with time.

In the algorithm where the probabilities are assumed to be perfectly known, the probabilities are given in the following manner: if there is a storm currently, the probability that the storm will stay there in the next 15 minutes is .80 (and consequently, there will not be any storm with a probability .20). Moreover, if there is no storm currently, the probability that there will be a storm in the next 15 minutes is .25 (and consequently, there will not be any storm with a probability .75).

In the algorithm where the probabilities are assumed uncertain but bounded within a set described by the likelihood function, we also obtain $N = (N_{ij})$ by exploring previous data. N_{ij} is the number of times the system went from state i to state j . The matrix is as follows

$$N = \begin{pmatrix} 80 & 20 \\ 25 & 75 \end{pmatrix}.$$

$N_{11} = 80$ means 80 out of 100 times the storm stayed in the next period, given the current situation; i.e., There is a storm and the analysis of the meteorological situation provides the probability of storm in the next stage as 0.80. Similarly, $N_{22} = 75$ means 75 out of 100 times the storm did not pop up in the next period, given the current situation; i.e, There is no storm and the

analysis of the meteorological situation provides the probability of storm not occurring in the next stage is .75. The values of N_{12} & N_{21} provide a similar meaning.

In the current practice the storm zone is avoided completely as if it is a deterministic obstruction. This traditional strategy incurs 51.5% delay (as a percentage of the nominal flight path). If we take $\beta = \beta_{\max}$, the uncertainty set becomes a singleton, and hence we obtain the maximum likelihood solution in the robust MDP algorithm, which provides the same solution as obtained in the nominal MDP algorithm. As β deviates from β_{\max} , the uncertainty set gets bigger. The robust solution performs better than the nominal solution if the prediction error increases. We can see in Figure 5, that the optimal value is very sensitive in the range of values of β close to β_{\max} . In real life, it is almost impossible to pinpoint the probability of a storm. Hence, we can never obtain 8.02% delay corresponding to the perfect prediction; instead we will obtain a delay of at least 25.1 % if we do not use the robust algorithm. The robust solution provides almost 10% less delay than the nominal solution when uncertainty is present in the probabilities. However, both of the strategies produce much less delay than the traditional strategy. If there is no error in the prediction, both of the algorithms provide a strategy that produces 43.3% less delay than the traditional strategy. Even with the presence of a significant estimation error, the robust strategy performs much better than the traditional strategy. Moreover, the running time for the nominal MDP algorithm was 3.83 sec, while the running time for the robust MDP algorithm was 8.32 sec. The extra computation cost for robustness was just 4.49 seconds, which is very moderate.

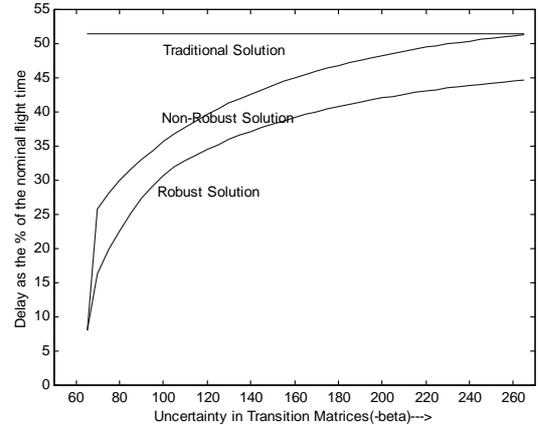


Figure 5: A comparison in delays between robust, non-robust and traditional solutions with varying uncertainty

Conclusion

We provide a tool that can be used by Air Traffic controller or Airline Dispatcher to dynamically route a single aircraft under uncertain weather. Our solution provides a less circuitous route for any aircraft which is subjected to bad weather and hence, restricts the overloading of aircraft in the neighboring sectors of the predicted storm zones.

The complexity of the computation depends on the origin-destination pair, size and location of the storms, level of discretization, and the stages of information updates. The complexity of the algorithm is exponential with the number of storms. So our algorithm can be applied in a scenario where the number of storms is moderate (around 10), which is sufficient in most of realistic situations.

In practice, we know that the probabilities that are generated by various weather prediction agencies are often incorrect. In the simulation, we showed that the optimal delay is very sensitive to the prediction error. Hence, we proposed a robust dynamic routing strategy where the best solution can be obtained even when there are errors in the estimation of storm probabilities.

Moreover, our algorithm adds very moderate additional computational complexity for adding robustness in the routing problem. We have proposed a Joint Set estimation/Robust optimization approach where

- The robust routing problem is as tractable as original routing problem.
- The likelihood functions allow a nice tradeoff between robustness and accuracy.
- We can capture asymmetric/correlated errors in estimates of storm probabilities.

In future work, we propose an extension of this model that will provide a dynamic routing strategy of multi-aircraft under weather uncertainty

References

- [1] A. Nilim, L. El Ghaoui, M. Hansen, and V. Duong, *Trajectory-based Air Traffic Management (TB-ATM) under Weather Uncertainty*, USA/EUROPE ATM R&D Seminar, Santa Fe, NM, December 2001.
- [2] Dimitris Bertsimas, Sarah Patterson. *The air Traffic flow management problem with enroute capacity*, Operations Research, 46:406-422,1998.
- [3] J. Goodhart. Ph.D thesis, Department of Industrial Engineering and Operations Research, University of California, Berkeley, 1999
- [4] M. Bielli, G. Calicchio, B. Nicoletti and S. Riccardelli. *The air traffic flow control problem as an application of network theory*. Computation and Operations Research, 9:265-278,1982.
- [5] Stuart Dreyfus *The art and theory of dynamic programming*, Academic Press, New York, 1977
- [6] M. Vidyasagar. *A theory of learning and generalization: with applications to neural networks and control systems*. Springer- Verlag, London, 1997.
- [7] D. P. Bertsekas. *Dynamic Programming and Stochastic Control*, Academic Press, New York, 1976.
- [8] D. P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, 1987.
- [9] Martin Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley Inter-Science, New York, 1994.
- [10] G. H. Polychronopoulos and J. N. Tsitsiklis. *Stochastic Shortest Path Problems with Recourse*, Networks, 27:133-143, 1996.
- [11] H.N. Psaraftis and J.N. Tsitsiklis, *Dynamic Shortest Paths in Acyclic Networks with Markovian Arc Costs*, Operations Research, 41:91-101 1,1993
- [12] D.P. Bertsekas and J.N. Tsitsiklis, *An Analysis of Stochastic Shortest Path Problems*, Mathematics of Operations Research, 16:580-595, 1991.
- [13] C. Burlingame, A. Boyd, K. Lindsey. *Traffic flow management modeling with the Time Assignment Model*, MP 93W0000036, The MITRE CORPORATION, Virginia, January, 1994
- [14] B. Lucio. *Multilevel Approach to ATC problems: On-line Strategy Control of Flights*. International Journal of Systems Science, 21:1515-1527, 1990
- [15] H. Marcia. *The Potential of Network Flow Models for Managing Air Traffic*, MP 94W0000058, The MITRE CORPORATION, Virginia, June, 1995
- [16] J. Evans, *Tactical Weather Decision Support to Complement "Strategic Traffic Flow Management for Convective Weather"*, USA/EUROPE ATM R&D Seminar, Santa Fe, NM, December 2001.
- [17] S. Boyd, *Convex Optimization*, Course Reader, EECS 290N, Department of Electrical Engineering and Computer Sciences, Fall 2001
- [18] E. Lehmann and G. Casella " *Theory of point estimation*", Springer-Verlag, NY, 1998