Airspace Throughput Analysis Considering Stochastic Weather

Joseph S. B. Mitchell, Ph.D.* and Valentin Polishchuk†
Applied Mathematics & Statistics, Stony Brook University, Stony Brook, NY 11794
Jimmy Krozel, Ph.D.‡
Metron Aviation, Inc., Herndon, VA, 20170

The estimation of the capacity of an airspace region during weather events is an important part of air traffic management. This problem must be solved ahead of time with predicted traffic demands and weather forecasts. In order to capture the uncertainty of the weather, a stochastic weather model is used. We investigate the problem of estimating the maximum capacity of an airspace region by analyzing the sector airspace geometry and a stochastic weather model. Using algorithms for computing geometric flow capacity in 2-dimensional regions, we compute the maximum capacity for an airspace having a given (deterministic) set of weather constraints. Then, we extend our results to a stochastic weather model, obtaining analytical results for weather constraints that form constraints along a line segment (e.g., placed along the flow bottleneck or along a squall line) and obtaining simulation results for a more general two-dimensional stochastic weather model. Our results allow us to determine the probability distribution of the throughput capacity of an airspace, given a probabilistic weather forecast.

I. Introduction

In Air Traffic Management (ATM), the National Airspace System (NAS) is decomposed into centers and sectors. An Air Route Traffic Control Center (ARTCC), often called simply a *center*, is typically partitioned into several *sectors*. Air traffic controllers are assigned to control the airspace within sectors, and a Traffic Management Unit (TMU) controls traffic between (across) sectors. Convective weather systems determine constraints ("obstacles") through which aircraft are not allowed to fly, for safety considerations, thereby reducing the throughput of a sector.

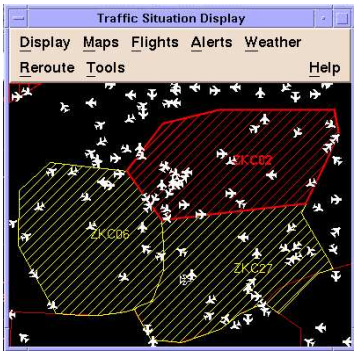

Figure 1. Traffic Situation Display (TSD) monitors the demand of sectors relative to MAP values.

In order to manage the throughput across a sector during weather events, controllers look upstream at aircraft on jet routes leading towards the sector, to estimate the number of aircraft planning to use the airspace, and consult a weather forecast to see what the predicted weather constraint is going to be. Based on this (and other relevant information about the traffic flows and constraints), controllers determine what strategy to employ to safely route traffic through the center and sectors. In order to set up conditions that do not overload the sector controllers, the number of aircraft expected within a sector (the *demand*) is monitored. The *capacity* of the sector is determined by what number of aircraft can safely be handled by the air traffic controllers for the sector, given the weather and other constraints. Each sector is assigned a Monitor Alert Parameter (MAP), a constant, which identifies the peak number of aircraft that can be safely handled for a given 15 minute time interval. Traffic Management Coordinators (TMCs in the TMU) monitor the demand on each sector using a Center Monitor (CM) tool (Fig. 1). The CM tool

*Professor, Stony Brook University.
†Ph.D. student, Stony Brook University.
‡Senior Engineer, Research and Analysis Department, 131 Elden Street, Suite 200, AIAA Associate Fellow.

shows 15-minute periods as being green, yellow, or red. Green means the sector volume is within acceptable levels (as determined by the MAP value), yellow means that some of flights that are causing the sector's predicted volume to exceed acceptable levels have not taken off yet, and red means that all of the flights that are causing the sector's predicted volume to exceed acceptable levels are already airborne. TMCs typically look 2.5 hours out on the CM, but the TMU does not get concerned with yellow and red alerts until 30 min – 1 hr out. The goal is plan ahead of time to reduce the demand to be below the MAP. This problem gets difficult when time varying weather constraints with varying degrees of geometric complexity (size and shape of the weather cells) add uncertainty to the situation.

Another ATM application that requires the controller to estimate the capacity of an airspace is the use of Flow Evaluation Areas (FEAs). These regions of airspace are typically polygonal regions – usually larger than the size of a single sector – where controllers draw the FEA into a map and analyze the traffic expected to enter into the area. Often, an FEA is drawn in order to evaluate an area of airspace that is impacted by weather, with the goal to estimate the capacity within the FEA as a function of the weather forecast. If the demand for the FEA exceeds the capacity of the FEA, a Flow Constrained Area (FCA) will possibly be set up by the controller, to regulate the demand to be no greater than the expected capacity of the airspace. Excess demand can be handled by routing aircraft around the FCA or by holding aircraft on the ground until a later take-off time. The capacity of a FEA (or FCA) can be established with the methods described in this paper.

In both the sector capacity and FEA capacity estimation problems, there is a polygonal region for which the airspace capacity is estimated given a forecast of the weather. If the look-ahead time for the weather forecast is short, e.g., 30 minutes or less, it is possible to reason with a deterministic weather forecast. However, if the look-ahead time for the weather forecast is long, e.g., 1-6 hours, then a stochastic weather forecast is needed. In this paper, we present algorithmic approaches to computing capacity of an airspace region in both deterministic and stochastic models. The error in the weather forecast introduces uncertainty both in the position of each weather cell and in its size. The further into the future the forecast is extended, the larger the variance of the forecast. Thus, in the stochastic problem statement, we study the effect of the variance of the weather forecast on the theoretical maximum capacity of an airspace.

A. Related Work

1. Related Work in Aviation

Wanke et al.[1] investigate the problem of probabilistic congestion management, which includes the prediction of traffic levels and airspace capacity. Their work warns that MAP values should not be considered a measure of airspace capacity. The statistics of sector peak count prediction uncertainty was also studied in previous work of Wanke et al.[2] Krozel et al.[3] analyzed the source of sector demand prediction errors. Additionally, Krozel et al.[4] investigate the problem of estimating the time of arrival of an aircraft to an airport given weather constraints, which could be applied to estimating the time of arrival to a sector as well. Moreau and Roy[5] present a stochastic model of a controlled stream of traffic leading to a fix or a sector. Additionally, Meyn[6] presents probabilistic methods for air traffic demand forecasting.

The work of Rhoda and Pawlak[7] studied how pilots decide to deviate or penetrate weather cells as they near the metering fixes of airports. In general, they find that pilots tend to avoid National Weather Service (NWS) Level 3 or higher weather cells; this has been adapted as the assumption of our paper for modeling weather constraints. While the altitude of cloud tops in severe storms is also an important factor[8] that pilots consider in determining which storm cells to avoid en route, it is not modeled explicitly in our problem statement.

2. Related Work on Weather Forecasting

Deterministic weather forecasts are the norm for ATM, available at almost all ATM facilities in the NAS. These forecasts are made available by the Weather and Radar Processor (WARP), the Integrated Terminal Weather System (ITWS), the Corridor Integrated Weather System (CIWS), and other tools used by air traffic controllers and meteorologists at ATM facilities.

Probabilistic weather forecasts are in place today and are also an active research topic. The National Convective Weather Forecast (NCWF-2), for instance, currently provides experimental 0-2 hr probabilistic forecasts for the NAS. A test version, NCWF-6, extends the forecast period to 6 hrs based on a combination of observations and Rapid Update Cycle (RUC) numerical weather prediction. NCWF-6 updates every 15 min

and provides a probabilistic map of forecast storm locations. Additional NCWF-6 forecast fields will include echo tops and storm organization. An ensemble of forecasts allows each forecast to be used individually or combined into probabilities. For example, ensembles from Numerical Weather Prediction (NWP) forecasts are created by varying the NWP model, the model's boundary conditions, and/or parameterizations. Each model run provides a different realization that is used to measure uncertainty in the forecast. An ensemble is constructed that blends observation based extrapolation forecasts (an extension of NCWF-2) and the RUC Convective Probability Forecast (RCPF).[9] Both NCWF-2 and the RCPF forecast products are given as probability maps. An approach to compiling ensembles based on observations is to vary the sources of uncertainty using stochastic models to generate an ensemble of perturbations.[10] The probability that convection will occur at a given point is determined by a calculation using the predicted field in the region near the point. This allows a straightforward means of calculating probabilities based on coverage within a distance of a point or a spatial filter. For both observation-based and RCPF systems, a threshold determines the areas affected by convection strong enough to impact aviation. For instance, in the NCWF-2, a threshold corresponding to NWS Level 3 is applied to the extrapolated convective interest field to obtain convective probability forecasts. The probability that convection will be present at a given point is determined by dividing the number of points where the threshold is exceeded within the filter area by the total number of points encompassed by the filter. The filter area is increased with forecast look-ahead time due to increased uncertainty in the forecast. Two forecasts are combined using linear weights – the weighting changes based on forecast period and time of day.

3. Related Work in Computational Geometry

The solution approach of this paper relies on maxflow/mincut theory.[11,12,13,14] A *capacitated network* $N = (V, E)$ is a graph on a set of nodes V and a set of edges E, such that each edge $e \in E$ has an associated positive number c_e, called the *capacity* of the edge (Fig. 2). An *s-t cut* in N is the partition of V into disjoint sets S and T, $V = S \cup T$, $S \cap T = \emptyset$, such that $s \in S$, $t \in T$. An edge $e \in E$ connecting the vertices i and j of V is said to *cross* the cut if $i \in S$, $j \in T$ or $i \in T$, $j \in S$. The *capacity* of the cut is the sum of the capacities of all edges that cross it. The *minimum cut*, or *mincut* through the network is a cut of minimum capacity. The mincut/maxflow theorem states that the value of the maximum *s-t* flow equals to the capacity of the mincut through the network. The mincut and the maxflow can be computed efficiently.[11]

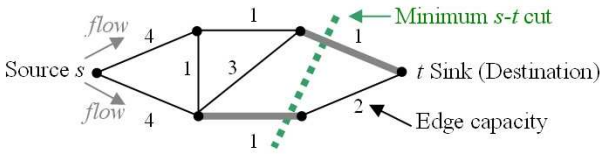

Figure 2. A network in which the minimum *s-t* cut (crossing over the bold edges) determines a cut value of 2 and thus a maximum *s-t* flow of 2.

Iri,[12] Strang[13] and Mitchell[14] studied the generalization of the network mincut/maxflow. They considered the continuous version of the problem. In this version, instead of a discrete network, a polygon P (possibly, with holes) is given. Two edges of P are marked as a *source* s and a *sink* t (Fig. 3). A *flow* in P is a divergence-free field σ with the support in P. The *value* of σ is the integral of σ across s (or t). An *s-t cut* of P is the partition of P into two sets S and T, $P = S \cup T$, $S \cap T = \emptyset$, such that $s \in S$, $t \in T$. The *capacity* of the cut is the length of the boundary between S and T, measured according to the 0/1 metric,[15] which assigns cost 1 for traveling through P and cost 0 for traveling through the holes. The *minimum cut*, or *mincut* through P is a cut of minimum capacity. Iri[12] and Strang[13] proved the mincut/maxflow theorem for the continuous flows. Mitchell[14] gave a different proof of the mincut/maxflow theorem and provided algorithms for finding the mincut and maxflow.

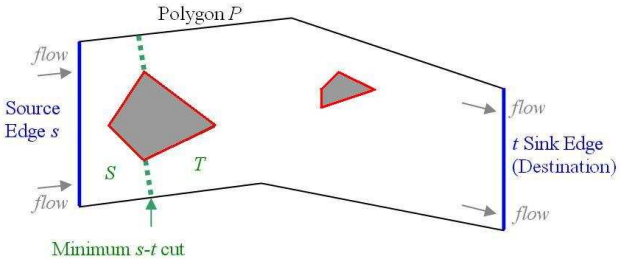

Figure 3. A continuous (polygonal) domain where the mincut is shown in green.

II. Modeling and Algorithms

A. Modeling the Capacity Estimation Problem

We model the aircraft as a point object traveling on a flight path π, represented with a set of line segments. Let w be the Required Navigation Performance (RNP) – the minimum separation required to be maintained between lanes of aircraft. Then the *air lane* Π associated with the flight π is the locus of points within distance $\frac{w}{2}$ from π: $\Pi = \cup_{x \in \pi} \mathcal{C}(x, \frac{w}{2})$, where $\mathcal{C}(c, r)$ is the disk of radius r centered at a point c.

We consider here a simplified version of the problem. In order to estimate the maximum capacity, we assign an air lane to be swept out "permanently" (over the time horizon of interest) so that it becomes a constraint for routing any subsequent air lanes. An air lane does not represent a jet route, nor is it collocated with any jet route, but it is a mathematical construct used to determine the capacity. Maximum throughput is achieved if each air lane is occupied by a sequence of aircraft "head to tail".

In the example illustrated in Figure 4, we consider a simple case of a rectangular sector in which we route air lanes West to East or East to West. Since the width of the rectangle-sector is approximately 30 nmi, and the normal RNP is about 5 nmi, we estimate the capacity of each air lane to be about 6 aircraft if the aircraft are routed "head to tail". While this spacing is not practical nor likely in the real world, it is used to obtain a theoretical upper bound on capacity; safety margins and/or Miles In Trail (MIT) restrictions may be added to make the problem statement more practical to observed average aircraft separations. A new MAP value of a sector in a real-world situation is influenced by a number of parameters pertinent to the traffic pattern, e.g., the entry and exit points, the number of aircraft in each air lane, etc. In this work our goal is not to determine a new MAP value; rather, we investigate how the *maximum possible* sector capacity changes with the weather coverage and forecast uncertainty.

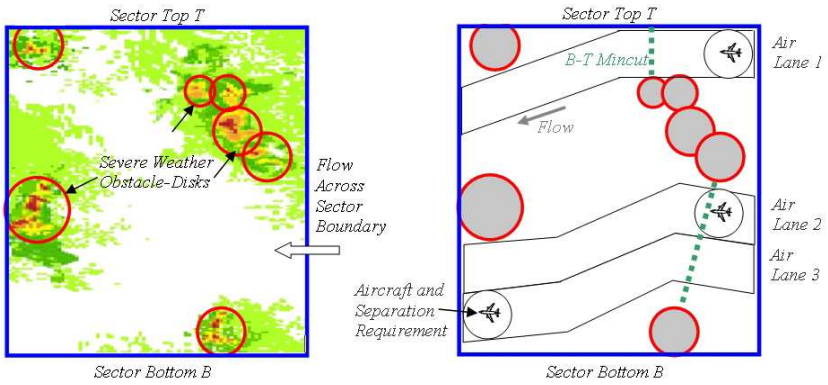

Figure 4. A simplified rectangular model for a sector is used to set up our analysis transforming weather data into a simple disk-constraint model in which air lanes are synthesized between weather constraints.

We model the sector as a simple polygon in the plane and we model the weather constraints as unions of (possibly overlapping) disks (or squares); see Fig. 4. In the experiments we conducted, the sectors are simply rectangles, but our algorithms apply to any sector shapes. Considering rectangular sectors simplifies the exposition. In particular, following the notation in,[14] the bottom (South) and top (North) edges of the rectangle are called B and T respectively and are treated as constraints such that the flow may only pass across the sector East to West or West to East. The left and right (West and East) sides of the rectangle will serve as the source s and the sink t.

The weather constraints may take on any arbitrary shape and size. We model the set of weather constraints as a union of disks that overlap the hazardous weather cells (NWS Level 3 or above). These disks serve as constraints ("disk-constraints"), since they specify points where aircraft are not safe to fly. The safety margin around each disk-constraint is included in the radius of the disk. If the knowledge about the weather were perfect, the radius of each disk would be a (known) constant. To model the uncertainty in the weather prediction, we make the radius of a disk a random variable. The variance of this random variable represents the uncertainty in the weather forecast. We do not focus on the exact nature of the uncertainty increase, which is dependent on the weather prediction look-ahead time, weather model grid size, and the nature of the weather system (e.g., isolated cells vs. squall line, etc.).

As a limitation of our approach, we note that in our experiments the weather cells positions and safety

margins around them are assumed to be independent identically distributed random variables, which is not true in reality. Our experiments could be extended, however, to generate all of the random variables from any predefined joint distribution.

B. Algorithms

We are interested in determining the *maximum throughput* across a rectangular airspace while avoiding (possibly overlapping) disk-constraints, denoted by $\mathcal{C}$. By the continuous maxflow/mincut theorem,[12,13,14] in order to find the maxflow value it is enough to compute the mincut length. This in turn is evaluated by computing the length of a shortest path from B to T in the 0/1-metric in which the constraints are assigned the weight 0 (path length is "free" inside a constraint) and the remainder of the rectangle (the "free space") is assigned weight 1.[15,14] The weights indicate the cost per unit distance of a path that determines the mincut. Throughout the paper, the mincut is used to define the capacity of an airspace.

Computing the mincut is equivalent to finding a shortest path from B to T in the *critical graph*, $G = (V, E)$, of the domain.[15,14] The graph G has a set $V = \{B, T\} \cup \mathcal{C}$ of nodes, one corresponding to each disk-constraint in the set $\mathcal{C}$ of constraints, and a node corresponding to B and a node corresponding to T. Edges E join each pair of nodes (the graph is complete).

In the *continuous mincut model*, the cost (length) of edge $e = (i, j)$ corresponding to constraints i and j is defined to be the (Euclidean) distance,

$$l(i,j) = \begin{cases} ||c_i - c_j|| - r_i - r_j & \text{if } ||c_i - c_j|| > r_i + r_j \\ 0 & \text{otherwise (i.e., } c_i \cap c_j \neq \emptyset) \end{cases}, \tag{1}$$

between disk-constraint i and disk-constraint j, with respective centers c_i and c_j and radii r_i and r_j. (Here, $||\cdot||$ denotes Euclidean length.) The cost of an edge between B or T and a constraint is the distance between the disk and the segment B or T. An edge of the graph G corresponds to a line segment, joining the closest points of one constraint to another constraint (or to B or T). If this line segment intersects some other constraint, then it is easy to see that the corresponding edge will not be an edge in any shortest path from B to T (since it can be replaced by two shorter edges with a lower sum of costs).

In the *discrete mincut model*, the cost, $l'(i, j)$, of edge $e = (i, j)$ corresponding to constraints i and j is defined to be the number of air lanes that fit side by side within a gap of length $l(i, j)$. Specifically, $l'(i, j) = \lfloor l(i, j)/w \rfloor$, where $\lfloor x \rfloor$ denotes the floor function (the largest integer less than or equal to x).

The continuous and discrete mincut models are similar but different. In the continuous mincut model, we do not discretize the throughput metric, i.e., we accept an answer like "5.78 air lanes may be routed through the sector", this flow may arise as a sum of multiple fractional edge costs (e.g., $5.78 = 1.2 + 3.08 + 1.5$). Using a continuous flow model, the cost of an edge (i, j) between constraints i and j is equal to the length, $l(i, j)$, of the edge, and the length of a mincut is exactly equal to the amount of continuous flow that can pass through the domain. In the discrete mincut model, we take into consideration the aircraft RNP, w, and express the capacity of an edge in terms of an integer number of air lanes. The number of the air lanes that can possibly pass between the gap in the weather defined by disk-constraint i and disk-constraint j is $l'(i, j) = \lfloor l(i, j)/w \rfloor$. The maxflow/mincut theory applies equally to the discrete case, permitting us to compute the number of air lanes by computing a shortest path in the critical graph. Currently, the width of an air lane is assumed to be $w = 5$ nmi, based on the horizontal separation standard between aircraft. To study the effect of the minimum required aircraft separation on the capacity, we repeated our experiments with $w = 3, 5$, or 10, and we refer to these cases as the reduced, normal, and expanded lateral separation RNP requirements.

Throughout the paper, the mincut is used to denote the capacity of an airspace. This may also be referred to at times as the maximum throughput. It should be noted that the location of the mincut may change over time or for different safety margins, as illustrated in Fig. 5. Additionally, the mincut line may be interpreted as the line that locates the theoretical "bottleneck" caused by the weather constraints.

1. Computing Capacity in a Deterministic Weather Model

We consider first the case in which there is a single set $\mathcal{C}$ of constraints, assumed to be known exactly. Our algorithm constructs the critical graph G from $\mathcal{C}$, and then searches it for a shortest path across the domain.

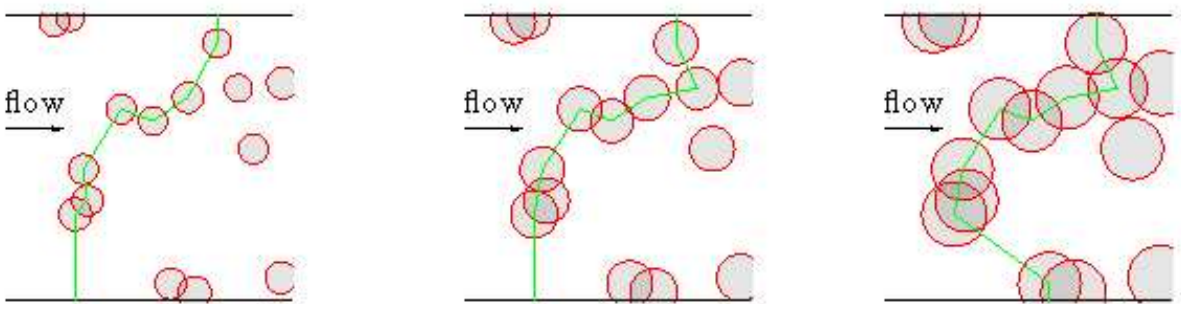

(a) low variance forecast (b) average variance forecast (c) high variance forecast

Figure 5. The mincut location changes as the safety margin around the weather systems changes (due to a more severe weather, forecast inaccuracy, etc.).

The edge costs in G may be either the lengths $l(i,j)$ or the number of air lanes, $l'(i,j)$. Our algorithm to compute the mincut is summarized as follows:

Algorithm Mincut

1. Assign costs ($l(i,j)$ or $l'(i,j)$) to the edges of G.
2. Compute a minimum-cost path from B to T in G using Dijkstra's shortest path algorithm. Output the length of the path (the mincut, which equals the maxflow).

Figure 6 shows a domain with its corresponding critical graph and the result of searching it with Dijkstra's algorithm to obtain a tree of shortest paths rooted at B, which includes a mincut path to T.

2. Computing the Distribution of Capacity in a Stochastic Weather Model

We are also able to compute the probability distribution of throughput for a given set of weather maps, each having an associated probability. This *scenario-based* stochastic weather model allows us to use a priori beliefs associated with multiple weather forecasts to obtain both expected capacity and a probability distribution of capacity. This is done by simply applying our capacity algorithm to each of the weather scenarios, thereby obtaining a discrete probability distribution of capacity, based on the belief probabilities.

In more detail, we extend our algorithm to a discrete stochastic weather model, as follows. Consider a model in which there is a discrete probability distribution on a set of forecasts (scenarios), each of which has an associated constraint set, $\mathcal{C}_k$. Suppose that forecast k has an a priori probability p_k associated with it, with $\sum_k p_k = 1$; i.e., p_k represents the probability (belief) that forecast k is accurate. Then, the capacity of the domain is a random variable, X, and we are interested in its expected value, $E(X)$, its variance, $var(X)$, and its probability distribution function, $F_X(x) = P(X \leq x)$. From $F_X(x)$ we can compute quantities like the probability that at least 5 air lanes can be accommodated (namely, $P(X \geq 5) = 1 - F_X(4)$), etc.

Algorithm Mincut Probability Distribution

1. For each forecast scenario k:
   (a) Assign costs ($l(i,j)$ or $l'(i,j)$) to the edges of the critical graph, G_k, associated with the set $\mathcal{C}_k$ of constraints.
   (b) Compute a minimum-cost path from B to T in G_k using Dijkstra's shortest path algorithm. Let x_k be the length of the path.
2. Output x_k and compute: $E(X) = \sum_k x_k p_k$, $var(X) = \sum_k x_k^2 p_k - [E(X)]^2$, and $F_X(x) = \sum_{k:x_k \leq x} p_k$.

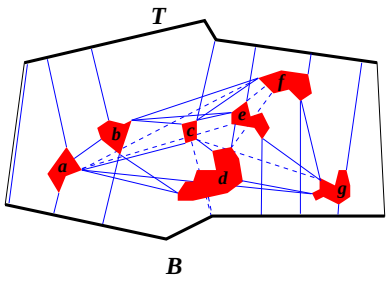

(a) Blue edges (defining the critical graph) connect points on hazardous weather constraints (red polygons) that are closest; dashed lines are edges that intersect constraints and are excluded from the search graph.

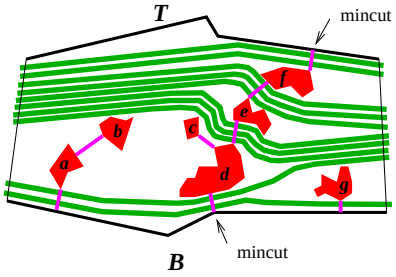

(b) The edges (thick magenta) of the shortest path tree obtained from Dijkstra's algorithm applied to find the maximum capacity of 10 air lanes (green).

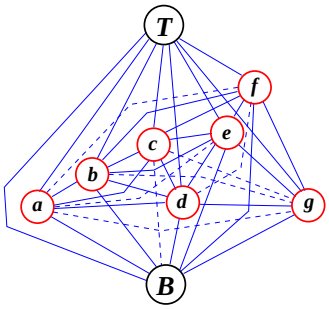

(c) The critical graph where nodes are defined by the 7 weather constraints and the boundaries B and T.

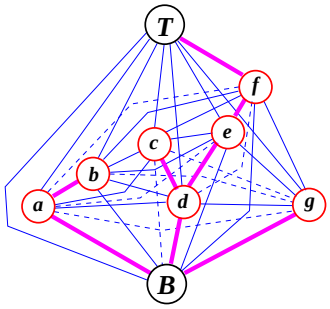

(d) The shortest path tree (magenta) within the critical graph.

Figure 6. Given hazardous weather constraints within a polygon airspace boundary, the objective is to compute the capacity of the airspace for traffic from left to right (right to left) within the airspace boundaries.

A simple example illustrating the above algorithm is presented in Fig. 7. The algorithm computes the capacities for each $k = 1, 2, 3$, obtaining 4, 4, and 7 air lanes, respectively. Thus, the expected capacity is $E(X) = 70\% \cdot 4 + 10\% \cdot 4 + 20\% \cdot 7 = 4.60$, the variance of capacity is $var(X) = 70\% \cdot 16 + 10\% \cdot 16 + 20\% \cdot 49 - 4.6^2 = 1.44$, and the probability distribution of capacity X is given by $P(X = 4) = 0.8$ and $P(X = 7) = 0.2$.

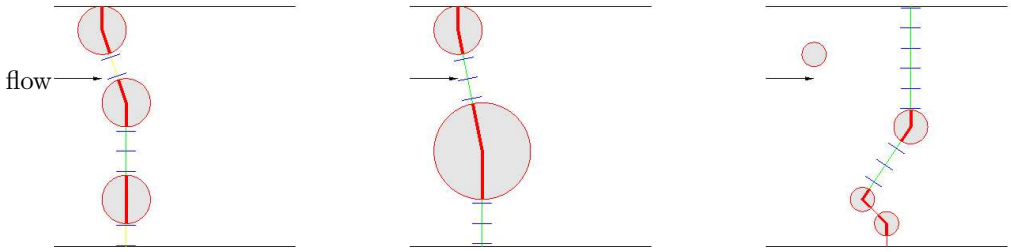

Figure 7. Three weather forecast (scenarios) with a priori probabilities (beliefs) 70%, 10%, and 20%, and computed capacities of 4, 4, and 7 air lanes.

The weakness of this simplified stochastic weather model is that it assumes that exactly one of the scenarios happens exactly. In the experiments we conducted, we addressed this concern partially by synthesizing a very large set of scenarios from a single forecast dataset, using randomization to generate a set of disk-constraints that have a distribution given by a real weather dataset. Each randomly generated set of disk-constraints represents a scenario that is similar to, but not exactly equal to, the real weather data. The scenarios are equally weighted (all p_k values are equal). In our experiments, we then use the above algorithm to compute the expected capacity $E(X)$, the variance $var(X)$, and the probability distribution of X.

Our methods allow us also to use a given (small) set of different weather forecasts (scenarios), each with a prior belief probability, and then use random disk-constraint generation within each forecast to generate a set of synthesized weather data that is statistically similar to the set of forecasts. From the algorithm, then, we obtain the distribution of the capacity X.

In future work, we expect to extend our model to work directly with a probabilistic weather map (or set of maps, each weighted by a priori probabilities of belief) as input and to compute the distribution of capacity directly from an algorithm, rather than rely, as we do here, on Monte Carlo simulation to obtain the empirical distribution.

III. 1D Analytical Solution

We begin with a study of a special case of our problem, in which the weather constraints are modeled as collinear line segments within a sector. We refer to this special case as the one-dimensional (1D) version of the problem. Studying the 1D problem not only allows us to obtain some analytic understanding of the problem, but also can be used to model certain weather cell configurations in two dimensions.

Consider the following scenario. Within a sector, there are $n+1$ weather cells, indexed $i = 0 \ldots n$, and each weather cell i is a line segment lying on a common line, ℓ. The i^{th} segment has length $2r_i$ and is centered at a point T_i on ℓ. (The segments may overlap along ℓ.) We consider the positions T_i to be random variables that define a renewal process (Fig. 8):

$$T_0 = 0 \tag{2}$$

$$T_i - T_{i-1} = X_i \qquad i = 1 \ldots n \tag{3}$$

$$X_i'\text{s} - \text{independent identically distributed} \qquad X_i \sim F_X \tag{4}$$

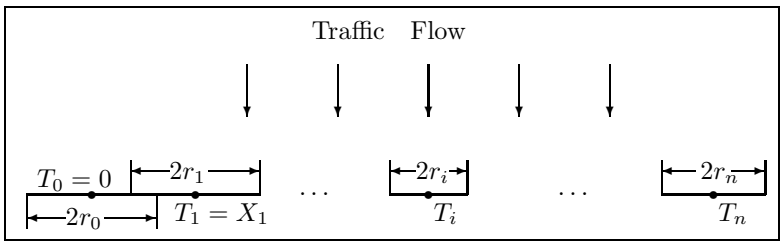

Figure 8. 1D problem drawn with the segments shown horizontally.

Our goal is to determine the maximum number of air lanes that can be safely routed in the presence of the 1D weather constraints, avoiding conflicts (i.e., intersections) between air lanes and weather cells. The 1D *mincut*, or *capacity* value is then the total length, $\sum_i^n \Delta_i$, of the interval $[0, T_n]$ not occupied by the segments, where $\Delta_i = [X_i - (r_i + r_{i-1})]^+$, with $[x]^+ = \frac{x+|x|}{2}$ denoting the positive part of x (i.e., $[x]^+ = x$ if $x \geq 0$ and $[x]^+ = 0$ otherwise). Letting $f_X(x) = F_X'(x)$ denote the probability density function for X, we can compute the expected capacity and the variance of capacity as follows:

$$E[\Delta_i] = \int_{r_i + r_{i-1}}^{\infty} f_X(x)(x - (r_i + r_{i+1}))dx \tag{5}$$

$$E[capacity] = \sum_{i=1}^{n} \int_{r_i + r_{i-1}}^{\infty} f_X(x)(x - (r_i + r_{i+1}))dx \tag{6}$$

$$Var[\Delta_i] = \int_{r_i + r_{i-1}}^{\infty} f_X(x)(x - (r_i + r_{i+1}))^2 dx \; - \; E[\Delta_i]^2 \tag{7}$$

$$Var[capacity] = \sum_{i=1}^{n} Var[\Delta_i] \tag{8}$$

In the appendix, we do some explicit calculations of the quantities $E[capacity]$ and $Var[capacity]$, as functions of the parameters of the renewal process (e.g., for X_i exponential, as in a Poisson process). Further, we compute the explicit dependence for the model in which the segment length parameters r_i are Gamma distributed and for the case in which there are a random number, n, of segments. In each case, the expected capacity is seen to decrease with increasing weather uncertainty.

The 1D model arises as a special case of the 2D model on which our experiments are based. Specifically, consider air traffic that moves across a rectangular sector of height L; the constraints in the sector are modeled as disks of different radii with the centers aligned along a vertical line within the sector (Fig. 9). The weather coverage $Wx\_coverage$ is defined as the portion of the line that is covered with the disks. The capacity of the sector is the length of the mincut, which is exactly the length of the portion of the line that is not covered (the green dotted line):

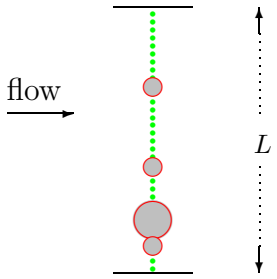

Figure 9. 1D problem.

$$capacity \quad = \quad L - Wx\_coverage \tag{9}$$

i.e., the capacity decreases linearly with the weather coverage.

IV. 2D Experiments with Synthesized Weather

In this section, we report on MATLAB simulation results for 2D experiments of the effects of uncertainty in the weather prediction on a sector capacity. There are two parameters in our experiments. The first one, the weather coverage, *Wx_coverage*, is the portion of the airspace that is occupied by the hazardous weather. We ran the experiments for the values of *Wx_coverage* equal to $0.1, 0.2, \ldots, 0.7$, and 0.8. If *Wx_coverage* is equal to, say, 0.7, then we draw the disks, representing the weather cells, until the union of the disks occupies at least 70% of the area of the airspace. The second parameter is the *variance* of a disk radius. The *variance* represents the uncertainty in the weather prediction. Below we describe the details and the results of the experiments.

A. Weather Models

We run our 2D experiments on two types of data: synthesized and real. In both cases, the airspace is represented by a 60×60 square. In the experiments with synthesized data, the weather cells centers are drawn from some probability distribution (we give the details below). In the experiments with real data, the weather cells centers are distributed according to a typical distribution of the weather intensity; we took the weather intensity distribution from a sample of National Convective Weather Diagnostic (NCWD) data from NWS and scaled it to the 60×60 square.

Our experiments on synthesized data model two different types of thunderstorm distribution. The first one is called *Popcorn Convection*. In a Popcorn Convection thunderstorms form on a scattered basis with little or no apparent organization (usually during the afternoon in response to diurnal heating). The individual weather systems within the airspace are modeled with disks of random (identically distributed) radii. The centers of the disks are distributed as follows: the abscissa and the ordinate of a center are uniformly distributed between 0 and 60; thus center of a disk is uniformly distributed over the whole square.

The second type of thunderstorm distribution is the *Squall Line* – a solid or nearly solid line or band of active thunderstorms. We model a sector with a Squall Line by a rectangle. The individual weather systems within the sector are still modeled with disks of random (identically distributed) radii. The centers of the disks are distributed differently: the abscissa of a center is uniformly distributed between 21 and 39, the ordinate of a center is uniformly distributed between 0 and 60; thus a center of a disk is uniformly distributed over the (vertical) rectangle centered in the square.

B. Experimental Setup

The experimental setup was as follows. For values of $Wx\_coverage = 0.1, 0.2, \ldots, 0.7, 0.8$, and for $variance = 0.0001, 0.001, 0.1, 1, 10, 100, 1000$ we repeat the following experiment (100 times). A 60x60 square airspace

is populated with disks to represent weather constraints. Disk weather constraints are added to the airspace until the union of disks occupies *Wx_coverage* of the area of the square. If the Popcorn Convection is modeled, the centers of the disks are uniformly distributed within the square. If the Squall Line is modeled, the centers are distributed uniformly within the 18×60 rectangle centered in the square. The radius of a disk is a random variable r, with $r = r_U + r_G$, where $r_U \sim U[1.5, 3]$ is uniformly distributed between 1.5 and 3, and $r_G \sim 0.75 \cdot Gamma(1/variance, variance)$ is proportional to a Gamma random variable. After generating just enough disks to occupy fraction *Wx_coverage* of the area, a mincut from B to T is found and its value is recorded. The number of disks necessary to reach the desired coverage is recorded separately.

Experiments were run under the following conditions. The capacity was found by computing the B-T shortest path in the 0/1 metric. To compute the shortest path we used the critical graph as explained in Section II: the graph has a node for each disk-constraint, a node for B and a node for T, and the length of the edge between constraints i, $\mathcal{C}(c_i, r_i)$ and j, $\mathcal{C}(c_j, r_j)$ is $[||c_i - c_j|| - r_i - r_j]^+$.

The Gamma random variable is parameterized so that $E[Gamma(a,b)] = ab$, $Var[Gamma(a,b)] = ab^2$. This way we had $E[r] = 3$, $Var[r] = Var[r_U] + Var[r_G] = \text{const} + 0.75 \cdot variance$.

Since running time was not a primary concern of this study, no particular care was made to optimize implementations of algorithms. The actual time spent to collect the experimental data was around 1 hour. One iteration took ~ 0.3 sec for low weather coverage and ~ 1 sec for high coverage; the increase in the running time was due to the larger number of constraints.

C. Results

Altogether, we ran 44800 experiments based on synthesized weather data. The results are presented in Fig. 10. As expected, as the weather coverage increases, the average capacity decreases, reaching zero at

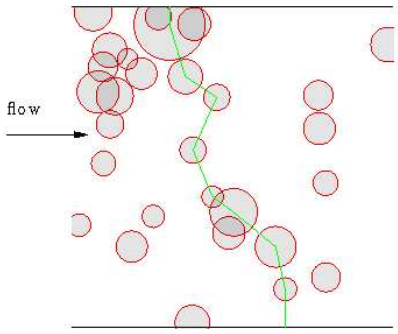

(a) Typical instance of Popcorn Convection model

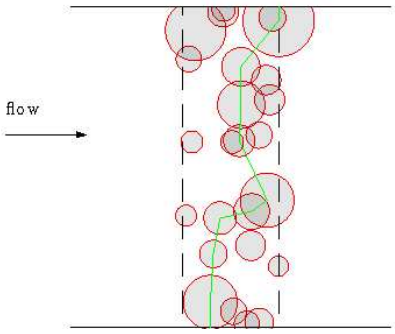

(b) Typical instance of Squall Line model

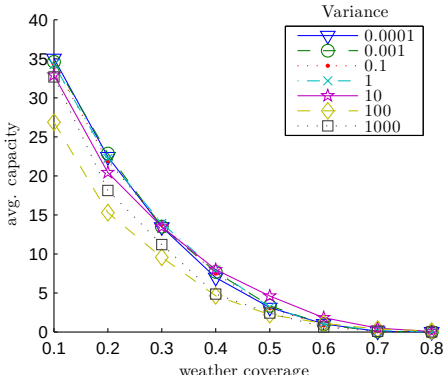

(c) Ave. capacity for Popcorn Convection model

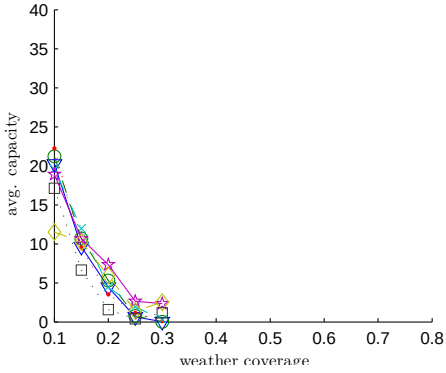

(d) Ave. capacity for Squall Line model

Figure 10. Experiments with synthetic weather.

some point prior to 100% weather coverage. As the weather coverage increases, the variance of the capacity goes down. The capacity is also a function of the type of weather system and its orientation. An isolated single weather cell or a multi-cell storm (e.g., popcorn convection) generally has greater capacity than a dense storm with linear organization (e.g., a squall line). With a severe weather constraint that has a linear organization, the orientation of the flow relative to the linear organization is an important factor: a flow perpendicular to the linear organization is more constrained than a flow parallel to the linear organization.

For a given weather coverage the capacity does not change much as the Gamma *variance* changes. For low weather coverage (10% -40%) the capacity goes slightly down at Gamma *variance* equal to 100 and then goes slightly up. For high coverage the capacity goes slightly up at Gamma *variance* equal to 100–1000 and then goes slightly down.

1. *Comparing Popcorn Convection vs. Squall Line*

The results of the experiments are different depending on whether the weather represents Popcorn Convection or Squall Line (Fig. 11). We observe the following:

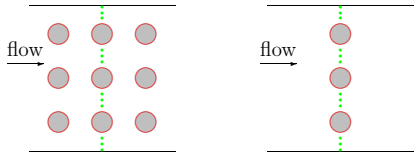

Figure 11. Abstractions of the weather distributions in the long run.

(1). When modeling the Squall Line the disks' centers only lay within a rectangle with the area .3 of the area of the airspace. Thus, the weather coverage never reached levels higher than .3 (except for pathological cases when there was a very large disk).

(2). For the Popcorn Convection the decrease of the capacity with the increase of the weather coverage is slower than linear. Indeed, one could imagine that in the long run the disk centers are distributed on a uniform grid within the square (Fig. 11, left). If there are k (unit) disks within an $L \times L$ square, then $Wx\_coverage \approx k\pi/L^2$ and

$$capacity \quad \approx \quad L - \sqrt{k} \quad \sim \quad 1 - \sqrt{Wx\_coverage} \tag{10}$$

Although (10) provides only a crude estimation of the dependence of the capacity on the coverage level, it is exactly this dependence that we observe experimentally (Fig. 10).

(3). The rate of the capacity decrease is almost the same for any Gamma *variance* value.

(4). For a given weather coverage the capacity value does not vary much with the Gamma *variance*.

2. *Taking into Account Aircraft RNP*

In Fig. 12 the results are presented for the case when the aircraft RNP was taken into account. The conclusions we made earlier are generally true also when the capacity is thresholded. The main difference is

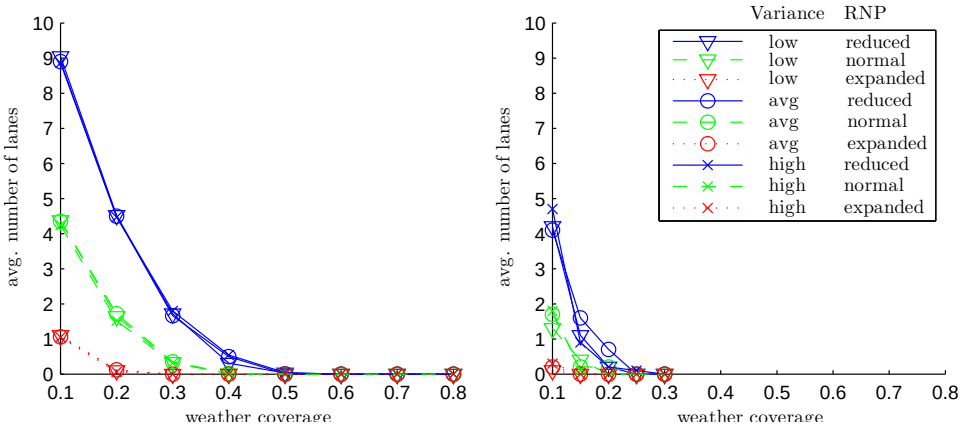

(a) Popcorn Convection (b) Squall Line

Figure 12. The average number of air lanes that fit into the airspace, experiments with the synthetic weather.

that the picture has become less sharp, especially for the case of modeling Squall Line. This is due to the discrete discontinuous nature of the floor function $\lfloor x \rfloor$.

V. 2D Experiments with Real Weather Data

In the experiments described above the centers of the disk-constraints were drawn from a uniform probability distribution. In real data, weather cells are scattered in highly non-uniform distributions. Thus, we conducted additional experiments using samples of real weather data (June 27, 2002, in the vicinity of the Atlanta airport) as the underlying probability distribution for generating the centers of disk-constraints. Specifically, the probability density function for centers of disks is taken to be proportional to the weather intensity map: Regions of high weather intensity are much more likely to have centers of disk-constraints than regions of low weather intensity. See Fig. 13. Thus, our technique of random generation of disks results in multiple instances of constraints that are statistically similar to real weather instances.

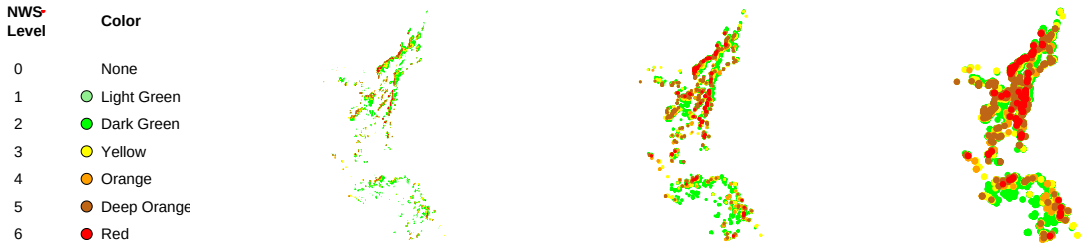

(a) low variance forecast (b) average variance forecast (c) high variance forecast

Figure 13. Example in which disk-constraints are generated according to a real weather data distribution.

- We chose the centers of the disks only from the points where the measurements of the storm intensity indicated that the weather intensity is at or above NWS=3. We call the set of these points $\mathcal{S}$; the weather intensity at the points in $\mathcal{S}$ ranges from 3 to 6.
- To model the influence of the weather intensity on the presence of constraints we chose the center of a disk to lie at a point with intensity $I \in \{3 \dots 6\}$ with probability $I/18 = I/(3+4+5+6)$.
- All the points with equal weather intensity were equally likely to serve as a weather cell center.
- The coverage of the airspace with the weather was defined as the proportion of the points in $\mathcal{S}$ covered with the disks.

A. Results

Altogether, we ran 22,400 experiments based on real weather data. The average values of the capacity as a function of the weather coverage are shown in Fig. 14.

The experiments based on the real data lead to the same conclusions as the experiments with the disks uniformly distributed within the square. Below we describe the differences that we observed.

1. As expected, the capacity decreases as the weather coverage increases. It does not decrease to 0 though, since we did not allow the disk-constraints to reside too close to the boundary. Indeed, the regions of the highest weather intensity (red on Fig. 13) are quite far from the boundary of the airspace; thus, almost always there was space to "fit in" some throughput close to the bottom and the top boundaries of the airspace.

2. The decrease of the capacity with the increase of the weather coverage is almost linear. This is a unique feature of the real Squall Line – the problem is essentially one-dimensional (see Fig. 9)! Indeed, one could imagine that in the long run the disks' centers are distributed uniformly on a line across the rectangle (Fig. 11, right). If there are k (unit) disks within a $W \times L$ rectangle, then $Wx\_coverage \approx k\pi/WL$ and

$$capacity \quad \approx \quad L - k \quad \sim \quad 1 - Wx\_coverage \tag{11}$$

Although (11) provides only a crude estimation of the dependence of the capacity on the coverage level, it is exactly this dependence that we observe experimentally (Fig. 14).

3. As before, the rate of the capacity decrease is almost the same for any Gamma *variance* value.

4. As before, for a given weather coverage the capacity does not change a lot as the Gamma *variance* changes.

Note that the given sample of real weather data behaves statistically like a Squall Line.

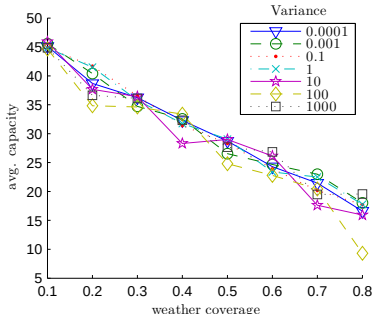

Figure 14. Experiments with real weather data determine the average capacity.

B. Taking into Account Aircraft RNP

Figure 15 shows the average number of air lanes that fit into the region as a function of weather coverage. The conclusions we made earlier are generally true also for the case of the thresholded capacity. The main difference is that the picture has become less sharp. This is due to the discrete discontinuous nature of the floor function $\lfloor x \rfloor$, which is used to model the integer number of air lanes. As expected, for larger RNP the capacity reaches zero when the weather coverage is high enough.

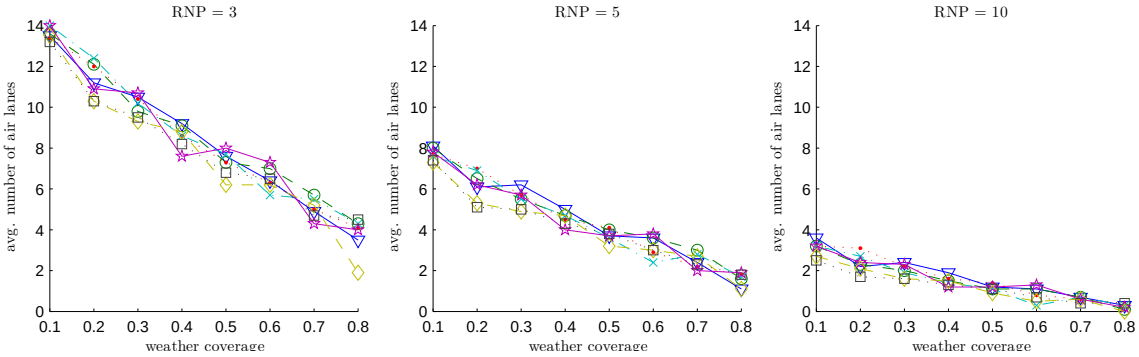

Figure 15. Experiments with real weather data determine the average number of air lanes for different values of RNP.

We define the *lost capacity* parameter as the Continuous Capacity defined by the mincut minus the number of air lanes ($n_{airlanes}$) times the RNP width:

$$Lost\_Capacity = MincutCapacity - RNP \cdot n_{airlanes} \quad (12)$$

This represents the tradeoff of the RNP requirement with the capacity resource. Figure 16 shows the lost capacity as a function of the weather coverage with level curves for RNP=3, 5, 10. The capacity loss appears to be almost independent of the weather coverage.

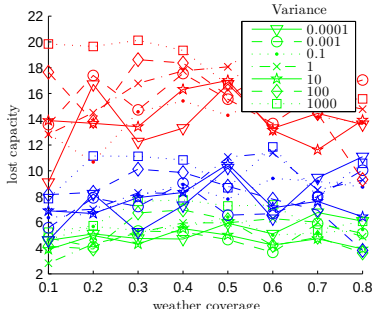

Figure 16. Lost capacity: green – reduced RNP, blue – normal RNP, red – extended RNP.

In Fig. 17 the probability that it will be possible to route at least a certain number of air lanes through the airspace is presented as a function of the weather coverage. These probabilities were determined by empirical measurements from experiments in which we measured the frequency with which it was possible to route a given number of air lanes. The probability decreases with the weather coverage.

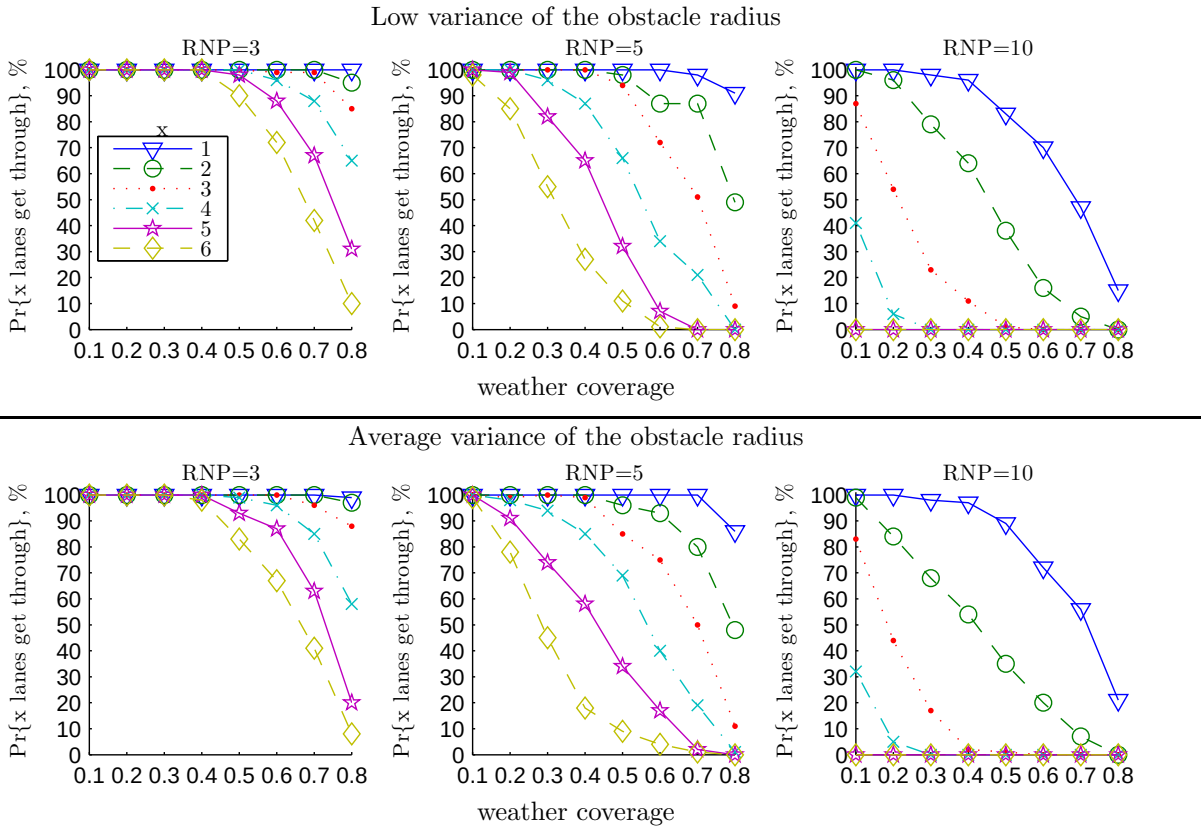

Figure 17. Probability (empirically determined) that a certain number of air lanes, x, fits into the airspace vs. weather coverage. Disk centers are randomly selected using the real weather data distribution. The disk radius is taken to be the sum of a uniform and a normal random variable.

VI. Application to Traffic Flow Management

In this section, we describe the application of our capacity estimation algorithm to a Traffic Flow Management (TFM) problem in the NAS. An Airspace Flow Program (AFP)[16] is a TFM strategy in which the capacity of an FCA needs to be estimated to address airspace constrained by severe weather (e.g., as shown in Figure 18). The AFP strategy determines which aircraft must fly around vs. through the FCA, and when aircraft are allowed to pass through the FCA based on the estimated capacity of the FCA. The AFP needs

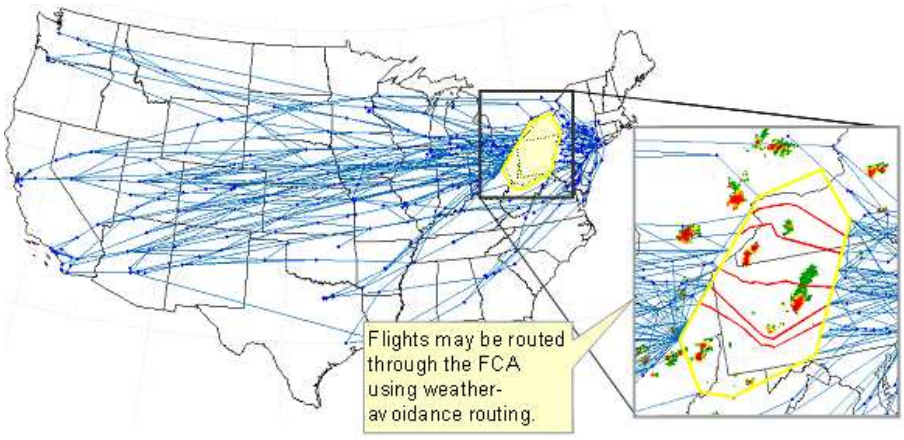

Figure 18. The AFP requires an estimate of the capacity of an FCA so that it can maximize the number of routes through the FCA.

(1) an estimation of the capacity for strategic 2+ hour forecasts so that it can plan for how many aircraft to route into the FCA, and (2) a refined estimate of the capacity of the FCA within 2 hours to establish the exact routes that aircraft should use to fly through the FCA. In this application, a probabilistic estimate of the capacity of the FCA is useful for strategic planning ($\geq$ 2 hours look-ahead), while a deterministic analysis is required for tactically defining the routes across the FCA. Our algorithmic approach to estimating capacity of an airspace is ideal for application in the AFP problem statement.

For this example, we do not illustrate the strategic (probabilistic) estimation of the FCA capacity, rather, we present the second stage of planning where the precise routes across the FCA are planned within the last hour of flight prior to entering the FCA. (Note: refer to Figure 17 for an example of how to establish the probability of x air lanes being available for the FCA.) We consider only the traffic that is passing from West to East (a similar process would be used to analyze the traffic from East to West, flying according to an alternate altitude rule).

The AFP plans routes through the FCA ahead of time, and repeats the planning process as new weather forecasts become available and as the desired flight plans of aircraft change throughout the day (or time period of the AFP). As the weather forecast indicates that the weather is moving across the FCA (Figure 19), routes must be established ahead of time to be given to the airlines and pilots for safe passage through the FCA. The AFP requires that an air lane be established to be "open" for 30 minutes, although the weather forecast may provide "snapshots" of what the weather constraint is expected to look like every 5 minutes. Given the 5 weather forecast snapshots, based on the modeling described in this paper, we determine the lowest airspace capacity that the FCA may allow during the course of the 30 minute time horizon. The minimum capacities for the 5 forecast snapshots are combined into one problem statement of planning routes around a 2D weather constraint within a given flight level with time taken into account. A viable solution approach to this problem is the Flow-Based Route Planner,[17] which is capable of routing one set of air lanes that are valid (conflict free) over the entire 30 minutes time horizon. A new set of air lanes is established in a similar way for the next 30 minute time horizon. This planning process is repeated every 30 minutes as long as the AFP is in place.

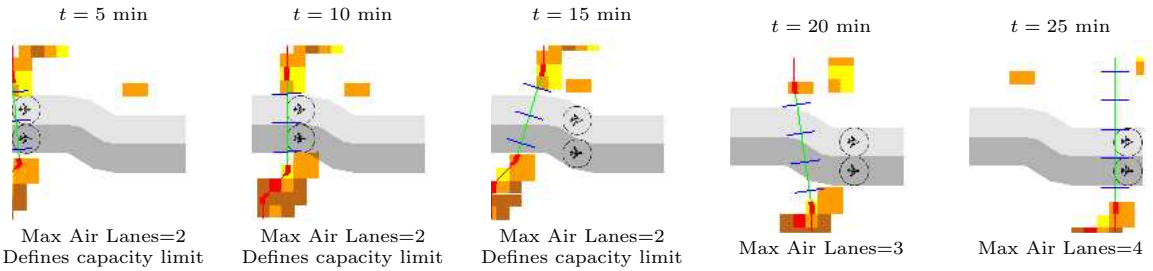

Figure 19. Weather forecasts for 5-minute intervals (only severe weather forecast data are shown) and the air lanes that sustain for the time period.

VII. Conclusions

We have shown how the throughput capacity of an airspace can be computed exactly algorithmically for a given domain and set of hazardous weather constraints. We have also examined how the stochastic nature of weather affects capacity. We have analyzed a model of stochastic weather in which weather constraints are synthesized either according to a user-specified probability distribution or according to actual weather data samples. We were able to analyze explicitly the expectation and variance of throughput in a simplified 1D model, which can approximate the effect of a squall line weather system. In more general 2D settings, we analyzed the throughput as a function of weather coverage and as a function of aircraft Required Navigation Performance (RNP). As expected, as the weather coverage goes up, the throughput goes down.

The throughput appeared to be almost independent of the *variance* of the constraint size. As expected, as the *variance* increases, the results become less "sharp". Taking into account the aircraft RNP has the same effect. The rate of the throughput decrease with the weather coverage depends on the type of weather organization. For squall line convection the decrease is linear, but for popcorn convection it is slower than linear. We propose to treat the rate of the decrease as the indicator of how close a weather organization is to the squall line model.

VIII. Future Research

Several directions remain for current and future research, including:

(1). Our current stochastic model is a simplified static scenario-based model. In future work, we expect to be able to compute the probability distribution of capacity as a function of either a single probabilistic weather map or a set of probabilistic weather maps, each weighted by an a priori probability (belief). Even more challenging is to study stochastic models of *dynamic* weather systems, explicitly modeling the uncertainty in movement when calculating the probability distribution of capacity.

(2). Our current results on 2D capacity are based on applying our algorithm repeatedly in Monte Carlo simulations. While we have obtained some analytic results for 1D, can anything be said *analytically* about the capacity in 2D for arbitrary polygon shaped airspace geometries as they occur in typical sectors in the NAS? Some upper/lower bounds may be possible.

(3). How should upstream and downstream weather constraints be accounted for in the 2D analysis to determine the maximum flow through an airspace? It is possible that the bottleneck for a sector might occur in the adjacent sector or sectors further away.

(4). For applications in the NAS, what is an appropriate set of values to use for the RNP requirements between air lanes (including any buffers that are natural in operations) and a practical number to use for the MIT requirements between aircraft in an air lane flow? Such parameters are easily included in the analysis described in this paper and would allow for practical estimates of throughput to be generated as opposed to the maximum throughput (capacity) that is described in this paper.

(5). Future research should explore the capacity problem in a full 3D space, including the altitude dimension of the problem. Echo tops data should be included into the definition of the weather hazard constraint, and flight levels should be modeled such that the flow across the airspace should consider flows of aircraft (or individual aircraft) that wish to pass over a weather constraint by climbing to new flight levels as opposed to avoiding weather constraints solely through horizontal maneuvering. Such a model is important to understand the increased capacity associated with a Reduced Vertical Separation Minimum (RVSM) that have recently been implemented in the NAS.

(6). The algorithm to compute capacity is based on searching a critical graph for a shortest path. The maxflow/mincut theory provides a method (described theoretically in[14]) to compute an associated set of actual air lanes to achieve the mincut capacity. However, the air lanes produced by such an algorithm may not conform to operational constraints and may not be desirable. Thus, it is important to implement and test new methods for computing a near-maximal capacity set of air lanes that obey operational constraints. Recent algorithmic results on thick paths[18] and minimum-cost flows[19] may be of help here. The Flow-Based Route Planner (FBRP) of[17,20] (employed also in[21]) is a practical method to search for many air lanes among *deterministic* constraints (even moving constraints). We expect to extend the FBRP to stochastic weather models and to determine how close to maximum theoretical capacity its sets of air lanes are in practice.

Acknowledgments

This research was funded by NASA Ames Research Center under contract NAS2-02075. We are greatly appreciative to our NASA Technical Monitor, Matt Jardin, Ph.D., for his guidance, which helped focus the effort. Finally, we appreciate the financial support of the sponsor of the research, NASA Ames Research Center and the Virtual Airspace Modeling and Simulation (VAMS) Project Manager, Mr. Harry Swenson. J. Mitchell also acknowledges support from the National Science Foundation (ACI-0328930, CCF-0431030, CCF-0528209) and NASA (NAG2-1620).

References

[1]Wanke, C., Song., L., Zobell, S., Greenbaum, D., and Mulgund, S., "Probabilistic Congestion Management," *6th USA/Europe Seminar on ATM R&D*, Baltimore, MD, June 2005.

[2]Wanke, C., Callaham, M., Greenbaum, D., and Masalonis, A., "Measuring Uncertainty in Airspace Demand Predictions for Traffic Flow Management Applications," *AIAA Guidance, Navigation, and Control Conf.*, Austin, TX, Aug. 2003.

[3]Krozel, J., Rosman, D., and Grabbe, S., "Analysis of En Route Sector Demand Error Sources," *AIAA 2000-5016, AIAA Guidance, Navigation, and Control Conf.*, Monterey, CA, Aug. 2002.

[4]Krozel, J., Lee, C., and Mitchell, J., "Estimating Time of Arrival in Heavy Weather Conditions," *AIAA Guidance, Navigation, and Control Conf.*, Portland, OR, Aug. 1999.

[5]Moreau, D. and Roy, S., "A Stochastic Characterization of En Route Traffic Flow Management Strategies," *AIAA Guidance, Navigation, and Control Conf.*, San Francisco, CA, Aug. 2005.

[6]Meyn, L., "Probabilistic Methods for Air Traffc Demand Forecasting," *AIAA Guidance, Navigation, and Control Conf.*, Monterey, CA, Aug. 2002.

[7]Rhoda, D. and Pawlak, M., "The Thunderstorm Penetration / Deviation Decision in the Terminal Area," *American Meteorological Society, 8th Conf. on Aviation, Range, and Aerospace Meteorology*, Dallas, TX, Jan. 1999, pp. 308–312.

[8]DeLaura, R. and Evans, J., "An Exploratory Study of Modeling En route Pilot Convective Storm Flight Deviation Behavior," *American Meteorological Society, 12th Conf. on Aviation, Range, and Aerospace Meteorology*, Atlanta, GA, Jan./Feb. 2006, Paper P12.6.

[9]Weygandt, S. and Benjamin, S., "RUC model-based convective probability forecasts," *11th Conf. on Aviation, Range and Aerosopace Meteorology*, Hyannis, MA, 2004.

[10]Seed, A., "A dynamic and spatial scaling approach to advection forecasting," *J. Appl. Meteor.*, Vol. 42, 2003, pp. 381–388.

[11]Ahuja, R. K., Magnanti, T. L., and Orlin, J. B., *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, Englewood Cliffs, NJ, 1993.

[12]Iri, M., *Survey of mathematical programming*, North-Holland, Amsterdam, Netherlands, 1979.

[13]Strang, G., "Maximal flow through a domain," *Math. Program.*, Vol. 26, 1983, pp. 123–143.

[14]Mitchell, J. S. B., "On maximum flows in polyhedral domains," *J. Comput. Syst. Sci.*, Vol. 40, 1990, pp. 88–123.

[15]Gewali, L., Meng, A., Mitchell, J. S. B., and Ntafos, S., "Path planning in $0/1/\infty$ weighted regions with applications," *ORSA J. Comput.*, Vol. 2, No. 3, 1990, pp. 253–272.

[16]Krozel, J., Jakobovits, R., and Penny, S., "An Algorithmic Approach for Airspace Flow Programs," *Air Traffic Control Quarterly*, Vol. 14, No. 3, 2006.

[17]Prete, J. and Mitchell, J. S. B., "Safe Routing of Multiple Aircraft Flows in the Presence of Time-Varying Weather Data," *AIAA Guidance, Navigation, and Control Conf.*, Providence, RI, 2004.

[18]Mitchell, J. S. B. and Polishchuk, V., "Thick Non-Crossing Paths," Submitted to *Association for Computing Machinery Society for Industrial and Applied Mathematics Symposium on Discrete Algorithms*, New Orleans, LA, Jan., 2007.

[19]Mitchell, J. S. B. and Polishchuk, V., "The Continuous Flow Decomposition Theorem," Submitted to *Association for Computing Machinery Society for Industrial and Applied Mathematics Symposium on Discrete Algorithms*, New Orleans, LA, Jan., 2007.

[20]Krozel, J., Penny, S., Prete, J., and Mitchell, J. S. B., "Automated Route Generation for Avoiding Deterministic Weather in Transition Airspace," *Journal of Guidance, Control, and Dynamics*, Fall, 2006.

[21]Krozel, J., Prete, J., Mitchell, J. S. B., Smith, P., and Andre, A. D., "Designing On-Demand Coded Departure Routes," *AIAA Guidance, Navigation, and Control Conf.*, Keystone, CO, Aug. 2006.

Appendix

Here we derive, using equations (6) and (8), explicit expressions for the expected length and variance of the minimum cut. We begin with the basic model in which the segment length parameters, r_i, are given (not random), but the locations of the segments (their center points) are determined by a renewal process with interevent random variables X_i.

Example 1 *$X_i \sim exp(\lambda)$, i.e., T_i form a Poisson process with the rate λ. Then*

$$E[\Delta_i] = \int_{r_i+r_{i-1}}^{\infty} \lambda e^{-\lambda x}(x-(r_i+r_{i-1}))dx = e^{-\lambda(r_i+r_{i-1})}\int_0^{\infty} \lambda e^{-\lambda u}u du = \frac{1}{\lambda}e^{-\lambda(r_i+r_{i-1})} \tag{13}$$

$$E[capacity] = \frac{1}{\lambda}\sum_{i=1}^{n} e^{-\lambda(r_i+r_{i-1})} \tag{14}$$

$$Var[\Delta_i] = \int_{r_i+r_{i-1}}^{\infty} \lambda e^{-\lambda x}(x-(r_i+r_{i-1}))^2 dx - \left(\frac{1}{\lambda}e^{\lambda(r_i+r_{i-1})}\right)^2 = \tag{15}$$

$$= \frac{1}{\lambda^2}\left(2e^{-\lambda(r_i+r_{i-1})} - e^{-2\lambda(r_i+r_{i-1})}\right)$$

$$Var[capacity] = \frac{1}{\lambda^2}\sum_{i=1}^{n}\left(2e^{-\lambda(r_i+r_{i-1})} - e^{-2\lambda(r_i+r_{i-1})}\right) \tag{16}$$

Suppose, for simplicity, that all weather cells have the same safety margin (radius) R. Note that the standard deviation of the capacity, normalized by its expectation,

$$\frac{STD[capacity]}{E[capacity]} = \frac{1}{\sqrt{n}}\sqrt{2e^{2\lambda R} - e^{-2\lambda R}} \tag{17}$$

increases with R. Indeed, the larger R is, the less we know about the weather and hence have to assume the larger portion of the space around a weather cell to be "no-fly" zone.

A. Variable Constraint Size

Let now each of the r_i be a random variable. We assume r_i's to be continuous independent identically distributed random variables; let f_r be their common density function. Then formulae (5) and (6) give the expected value of the length of one "window" and the capacity *conditioned* on $r_0, \ldots, r_n$:

$$E[\,\Delta_i \,|\, r_i + r_{i-1}] = \int\limits_{r_i + r_{i-1}}^{\infty} f_X(x)(x - (r_i + r_{i-1}))dx \tag{18}$$

$$E[\,capacity \,|\, r_0, \ldots, r_n] = \sum_{i=1}^{n} \int\limits_{r_i + r_{i-1}}^{\infty} f_X(x)(x - (r_i + r_{i-1}))dx \tag{19}$$

By the conditional expectation formula

$$E[\Delta_i] = E\left[E[\,\Delta_i \,|\, r_i + r_{i-1}]\right] = \int\limits_0^{\infty} f_{r_i + r_{i-1}}(r)dr \int\limits_r^{\infty} f_X(x)(x - r)dx \tag{20}$$

and

$$E[capacity] = \sum_{i=1}^{n} E[\Delta_i] \; = n \int\limits_0^{\infty} f_{r_i + r_{i-1}}(r)dr \int\limits_r^{\infty} f_X(x)(x - r)dx \tag{21}$$

Similarly,

$$Var[\Delta_i] = E[\Delta_i^2] - E[\Delta_i]^2 = E\left[E[\,\Delta_i^2 \,|\, r_i + r_{i-1}]\right] - E[\Delta_i]^2 = \tag{22}$$

$$= \int\limits_0^{\infty} f_{r_i + r_{i-1}}(r)dr \int\limits_r^{\infty} f_X(x)(x - r)^2 dx \quad - \quad E[\Delta_i]^2$$

and

$$Var[capacity] = nVar[\Delta_i] \tag{23}$$

What would be a good choice for the model distribution of r_i's? Of course, r_i should be non-negative. Also recall, that we are interested in the impact, which the *variance* (and not the *expected value*) of r_i has on $Var[capacity]$. Thus, we need a family of non-negative distributions with the same expected value and different variance. A natural choice is to model the r_i's with a Gamma distribution.

Example 2 *$X_i \sim exp(\lambda)$, $r_i \sim Gamma(1/v, v)$, $v \in \mathbf{R}$; i.e., $f_{r_i}(r) = \frac{1}{\Gamma(1/v)}\frac{1}{v}e^{-\frac{1}{v}r}\left(\frac{1}{v}r\right)^{\frac{1}{v}-1}$, $E[r_i] = \frac{1}{v}v = 1$, $Var[r_i] = \frac{1}{v}v^2 = v$. Then $r_i + r_{i-1} \sim Gamma(2/v, v)$ and*

$$E[\Delta_i] = \int\limits_0^{\infty} \frac{1}{\Gamma(2/v)}\frac{1}{v}e^{-\frac{1}{v}r}\left(\frac{1}{v}r\right)^{\frac{2}{v}-1} dr \int\limits_r^{\infty} \lambda e^{-\lambda x}(x - r)dx = \tag{24}$$

$$= \frac{1}{\lambda\Gamma(2/v)v^{2/v}}\frac{\Gamma(2/v)}{\left(\lambda + \frac{1}{v}\right)^{\frac{2}{v}}} \quad = \quad \frac{1}{\lambda(\lambda v + 1)^{2/v}}$$

$$Var[\Delta_i] = \int\limits_0^{\infty} \frac{1}{\Gamma(2/v)}\frac{1}{v}e^{-\frac{1}{v}r}\left(\frac{1}{v}r\right)^{\frac{2}{v}-1} dr \int\limits_r^{\infty} \lambda e^{-\lambda x}(x - r)^2 dx - E[\Delta_i]^2 = \tag{25}$$

$$= \frac{2}{\lambda^2(\lambda v+1)^{2/v}} - \frac{1}{\lambda^2(\lambda v+1)^{4/v}}$$

$$E[capacity] = \frac{n}{\lambda(\lambda v+1)^{2/v}} \tag{26}$$

$$Var[capacity] = nVar[\Delta_i] \tag{27}$$

Observe that $f(v) = (\lambda v+1)^{2/v}$ *is monotonically decreasing from* $e^{2\lambda}$ *to* 1 *on* $(0,+\infty)$*. Thus,*

$$Var[\Delta_i] = \frac{1}{\lambda^2}\frac{1}{f(v)}\left(2 - \frac{1}{f(v)}\right)$$

is monotonically increasing from

$$\lim_{v\to 0} Var[\Delta_i] = \frac{1}{\lambda^2}(2e^{-2\lambda} - e^{-4\lambda}) \tag{28}$$

to

$$\lim_{v\to\infty} Var[\Delta_i] = \frac{1}{\lambda^2} \tag{29}$$

on $(0,+\infty)$*.*

B. Random Number of Constraints

Finally, we extend our analysis to the case in which the number, N, of segments that model hazardous weather cells is also a random variable. In this case, equations (26) and (27) give the expected value and variance of the capacity *conditioned* on the event that the number of constraints is n (plus one constraint at zero). By the conditional expectation formula

$$E[capacity] = \frac{E[N]}{\lambda(\lambda v+1)^{2/v}}, \qquad \text{and} \qquad Var[capacity] = E[N]\,Var[\Delta_i]. \tag{30}$$

Example 3 *$X_i \sim exp(\lambda)$, but this time we are watching the throughput between constraints in the interval $[0,1]$. To get rid of "boundary effects", suppose that the last (i.e., N-th) constraint extends to the right up to 1. Then $N \sim Poisson(\lambda)$, and we compute*

$$E[capacity] = \frac{1}{(\lambda v+1)^{2/v}}, \qquad \text{and} \qquad Var[capacity] = \frac{2}{\lambda(\lambda v+1)^{2/v}} - \frac{1}{\lambda(\lambda v+1)^{4/v}}. \tag{31}$$

C. Expected Number of Air Lanes

Let the radii R (the lengths) of the constraints be fixed. Let the number, n, of constraints be fixed. Then the expected number of air lanes that can be routed through the region is

$$E\left[\sum_{i=1}^{n}\left\lfloor\frac{[X_i-2R]^+}{w}\right\rfloor\right] = nE\left[\left\lfloor\frac{[X_i-2R]^+}{w}\right\rfloor\right] \tag{32}$$

$$E\left[\left\lfloor\frac{[X_i-2R]^+}{w}\right\rfloor\right] = \int_0^\infty f_X(x)\left\lfloor\frac{[x-2R]^+}{w}\right\rfloor dx = \tag{33}$$

$$= \sum_{j=1}^{\infty}\int_{2R+jw}^{2R+(j+1)w} f_X(x) j\,dx = \sum_{j=1}^{\infty} j\,[F_X(2R+(j+1)w) - F_X(2R+jw)] = \tag{34}$$

$$= \frac{e^{-2\lambda R}}{e^{\lambda w}-1} \tag{35}$$

Thus,

$$E\left[\sum_{i=1}^{n}\left\lfloor\frac{[X_i-2R]^+}{w}\right\rfloor\right] = n\frac{e^{-2\lambda R}}{e^{\lambda w}-1} \tag{36}$$