Quantifying Desirable Air Route Attributes for a Reroute Generation Capability

James S. DeArmon,[*] Craig R. Wanke,[†] and Joseph A. Nardelli[‡]
The MITRE Corporation, McLean, VA, 22102

When airspace is partially or completely blocked by severe weather, a major challenge for airspace managers and users is the creation and selection of viable alternate routing for air traffic. Research is underway for an automated capability to help to meet this challenge. An important component of such a capability is the quantification of the attributes of operationally-acceptable routes. In this paper we describe two such quantifications, for route lateral deviation, and for flow preference.

I. Introduction

The automatic generation of operationally-acceptable route options would be a useful decision support capability for air traffic flow managers. This capability could be employed, for example, when severe weather impacts en route airspace, blocking or partially blocking normal routing. Such a capability would be useful both for mid-term Next Generation Air Transportation System (NextGen) capabilities already in the design stage and for the advanced 4-D Trajectory-Based Operations (TBO) envisioned for the 2025 timeframe.[1] In this paper, we propose an initial approach to evaluating the operational acceptability of route options, applicable to the planned mid-term NextGen traffic management toolset.

This work is part of a larger project to address the overall need for flight option generation in both near-term and far-term decision support systems. Many NextGen concepts call for automatic generation and/or negotiation of alternate flight routings, to respond to weather hazards, metering and volume constraints, and air-to-air conflicts. These options must be rapidly generated, negotiated, and revised in the presence of uncertain predictions and changing conditions. This project proposes to develop a flexible, multi-constraint algorithm for developing flight options, to include not only changes to lateral routes but also ground delays, altitude maneuvers, or 4D time constraints. The resulting concept and prototype will be applicable to a wide range of NextGen decision support research, especially where 4-D trajectory negotiation is involved.

In the near-term, given current en route air traffic control (ATC) automation, “good” route choices must satisfy several qualitative operational constraints. Route choices must not induce excess airspace complexity (e.g., departures cannot cross arrival flows), they must obey procedural restrictions defined within and between the Federal Aviation Administration (FAA) facilities, they should not involve excessive inter-facility coordination, and they must be acceptable to the pilot and dispatcher. The latter includes consideration of fuel, schedule, and weather avoidance constraints.

For a decision support tool to generate options which adhere to these constraints, we need quantitative measures of how well particular options fare with respect to the constraints. An operational-acceptability score for a route option can be computed from these measures, and used by automation and/or human decision-makers to develop operationally-acceptable solutions to traffic management problems. Here, we describe quantification approaches for two important route attributes.

II. Background

Much research has been done on the problem of flight rerouting during adverse en route weather[2,3,4]. The authors know of no research regarding the automatic assessment of operational acceptability of candidate routes.

To begin our research in this area, we interviewed several subject matter experts, and found a fair amount of consensus on the important operational considerations for alternate route construction by traffic managers. In this

---

[*] Multi-Discipline Sys Eng, Principal, Terminal/TFM System Eng, MS N450
[†] Simulation Modeling Eng, Sr Principal, Terminal/TFM System Eng., MS N450, Senior Member AIAA
[‡] Simulation Modeling Eng, Lead, Perf and Eng Modeling Analysis, MS N590

paper we describe the analysis undertaken for two factors. The final result is the quantification into continuous values from 0 to 1. The 0/1 range allows for consistent assignment of weighting coefficients in a final evaluation scheme. (The description of the scheme is beyond the scope of this paper.)

III. Lateral Deviation

The goal is to map lateral deviations during severe weather into a 0/1 score. To develop this relationship, good weather flight tracks were compared to bad weather flight tracks. For a known bad weather day (7/27/06) and a known good weather day (8/17/2006), both Thursdays, flights were paired based on flight ID, and origin and destination airport. For a 3-hour period of the bad weather day, severe weather polygons were constructed, plus a margin of 200 nautical miles (NM) around them. All flights that horizontally intersected these expanded polygons during the 3-hour period were considered as candidates for the computation. Lateral deviation was evaluated as the *maximum* lateral distance between the good and bad weather flight tracks. It would be expected that bad weather tracks, in avoiding weather, had non-trivial lateral deviations from an associated good weather track.

But how do we know that the lateral deviations were due to bad weather? As a test of this notion, another comparison was made, this time between the same candidate tracks (again identified by flight ID, origin and destination airport) for two good weather days. A computation was again made deriving maximum lateral deviation per track pair. Next, the two distributions were compared, bad/good and good/good, and it was found that the mean values were statistically significantly different at $p<0.001$. So, as a result of the bad/good comparison, we have a collection of values varying between 0 and 400 NM lateral deviation, for flights that went somewhere near the many severe weather regions on the 7/27/06 date. To generate the mapping into 0..1, 1 minus the cumulative distribution function (CDF) is calculated. In Fig. 1, this is plotted as red Xs.

A mapping function $y = f(x)$ was desired to make easy use of this data. As it happened, the data did not fit any classical probability distribution, nor did they fit a non-linear regression well. Therefore, a piecewise-linear regression was modeled. These are shown as blue squares in Fig. 1. The R-squared of that model was > 0.99.

The final functional form can be expressed as:

$$y = 1.26 - 0.033x + 0.027(x-25)(x>25) + 0.0048(x-70)(x>70) + 0.00095(x-180)(x>180);\quad y = \min(1,y)$$

where:

- x is the maximum lateral deviation in NM, a positive value either right or left of the undeviated good-weather course
- y is the 0..1 response, and
- true/false comparison terms in the equation, of the form "(x>25)," evaluate into 1/0 values, respectively

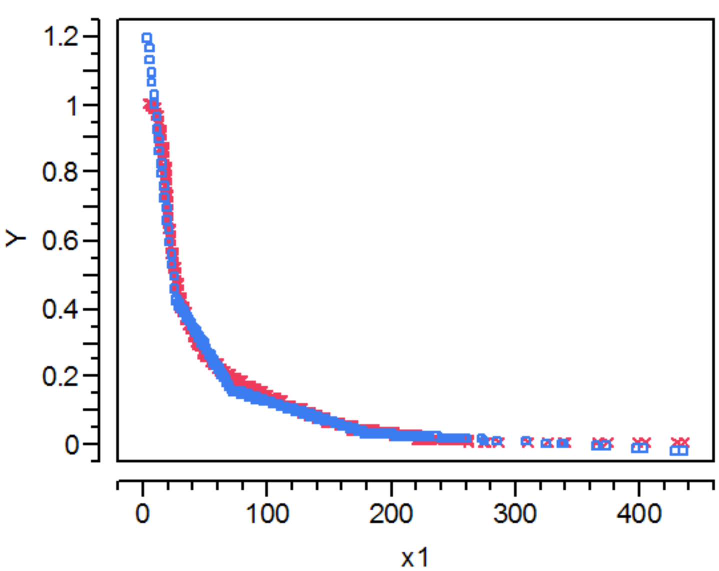

Figure 1. One minus CDF for lateral deviation (red X's), plus piece-wise linear regression model (blue squares)

The other part of the analysis was to evaluate altitude deviations for reroutes. In an analysis just like the one for lateral deviations, we found no relationship – there did not exist any difference in altitude for flights whose lateral deviation was small (<10 NM), i.e., the set of flights that stayed on route, but used altitude (only) to get past the weather region. Unlike lateral deviation, altitude difference is signed, and the resultant altitude deviation distribution was symmetric, with a near-zero mean. Discussions with a subject matter expert confirmed this as an expected result – altitude change does not offer much relief in the face of convective en route weather. Convective weather can rise to very high altitudes, and pilots like to clear the weather tops by several thousand feet because of turbulence above the storms.

IV. Flow Preference

A second important operational-acceptability factor opined by subject matter experts was that a suggested route should "go with the flow." That is, there exist in the National Airspace System (NAS) frequently-used paths, and these paths are well-known to air traffic management and readily acceptable if proposed as (part of) a flight route. So we sought to quantify "going with the flow," a 0/1 scoring function that would return a high value near 1 if a flow is used frequently, and a low value near 0 if a flow is not used frequently. The details on converting a frequency count to a 0/1 score are presented below. We decided on the granularity of a simple, single route segment, as defined by a fix pair ordered as: (FromFix,ToFix). This element is referred to hereafter in this paper as a Fix Pair Route Segment (FPRS).

As an aside, note that in other research detecting/defining air traffic flows in the NAS, a longer portion of a route (or an entire route) has been considered;[5] the rationale for our finer treatment here, using a route segment spanning a fix pair, is as follows. Most flights in the NAS that will require weather-induced rerouting will likely keep major portions of their original routing; the part before the flights approach the weather region, and, on the other side of the weather, rejoining the original routing to destination. So some flexibility is needed in creating routes with respect to "going with the flow." Some parts of the route must use high-frequency segments (e.g., standard terminal arrival routes) but other parts of the routes may need to use less-traveled paths, accessing low-frequency segments to try to get back onto a dominant flow as quickly as possible.

Regarding computation of count frequencies, one must consider the total population of interest – this will be the denominator in calculating the frequencies. It was suspected that the total population of interest for the problem at hand needed to be more focused, than, say, using the entire NAS. It was decided that a way to focus the definition of total population was to consider subsets of flights, by using groupings or clusters of proximate airports. The idea here is that flights from one region to another region of the NAS will likely be using similar route segments and routes, making for a fairly homogenous population. See Fig. 2 wherein a K-means clustering was applied to 3000+ airports in the conterminous U.S. using latitude and longitude as attributes. Cluster numbers are arbitrarily assigned.

An important question arising when performing cluster analysis is the selection of number of clusters. See the appendix, which shows a rationale for 35 clusters for the problem at hand.

Frequency tabulations were computed for all 35 x 35 = 1225 combinations – this includes directionality, e.g., 8 to 5 and 5 to 8, as well as within a single cluster, e.g., 9 to 9. The data for tabulating frequencies was 100 days of route data from flight plans for April—July 2008. This time period would capture some number of days that are considered to be in the severe weather season with respect to aviation. The route data were expanded, and fix pairs were considered in turn. For each cluster pair, the set of applicable route data was expanded and fix pairs were tallied and rank-ordered from most to least frequent. In order to map into a 0/1 variable with 1 corresponding to the highest-frequency fix pair, the cumulative distribution function (CDF)[§] was calculated, and then 1 minus the CDF was substituted, yielding, as an example, that shown in Fig. 3. Labels along the x-axis are illegible in the figure, but each of the tic marks corresponds to a fix pair for the 23-to-9 cluster pair. A single fix pair, KLUBB_DWINE is called-out in Fig. 3.

As an informal validation of the efficacy of the procedure, Figs. 4 through 6 are shown. Plotted on top of the NAS center boundaries are colored lines, wherein both line width and "hotness" of color (yellow to red) indicate higher frequencies as ten classes. Figure 4 shows flows mainly from the New York City area to Houston. Note the dominance of the single, rather direct path from New York City area to Houston, but also note the other lines that sometimes go rather far afield in getting to the destination cluster – these may be viable options if the major flow is blocked by weather. Figure 5 shows the reverse, and demonstrates the importance of direction in denoting the cluster pair – flows are from Houston to New York, and show the highly structured routes feeding the busy New York-area airports. Finally, consider Fig. 6 which displays route segment frequencies within cluster 9, in the Pacific Northwest. A reasonable-looking network can be seen, with major flows between metropolitan areas of Seattle, Portland, Bellingham, etc., and many minor flows that make secondary connections. (As in the other geographic figures, west longitude and north latitude are the x and y axes, respectively.) Although the within-cluster sets have no implied directionality, the route segments themselves indicate this. For example for cluster pair 9-9, both of these fix pairs appear: BEEZR_ELN and ELN_BEEZR, indicating direction; each fix pair has its own frequency score.

[§] To compute the CDF, the histogram of counts for fix pairs is ordered from high to low count along the x-axis. Next, these counts are substituted with a cumulative value starting from the origin; the percentage of total observations is on the y-axis.

The resulting product of this analysis is 1225 files, specific to the 35 x 35 cluster pairs, containing data like that shown in Fig. 3, and 1 file containing a lookup-table that relates airport IDs to cluster number 1-35. The information is ready for use in a search algorithm aimed at constructing operationally viable alternate routing.

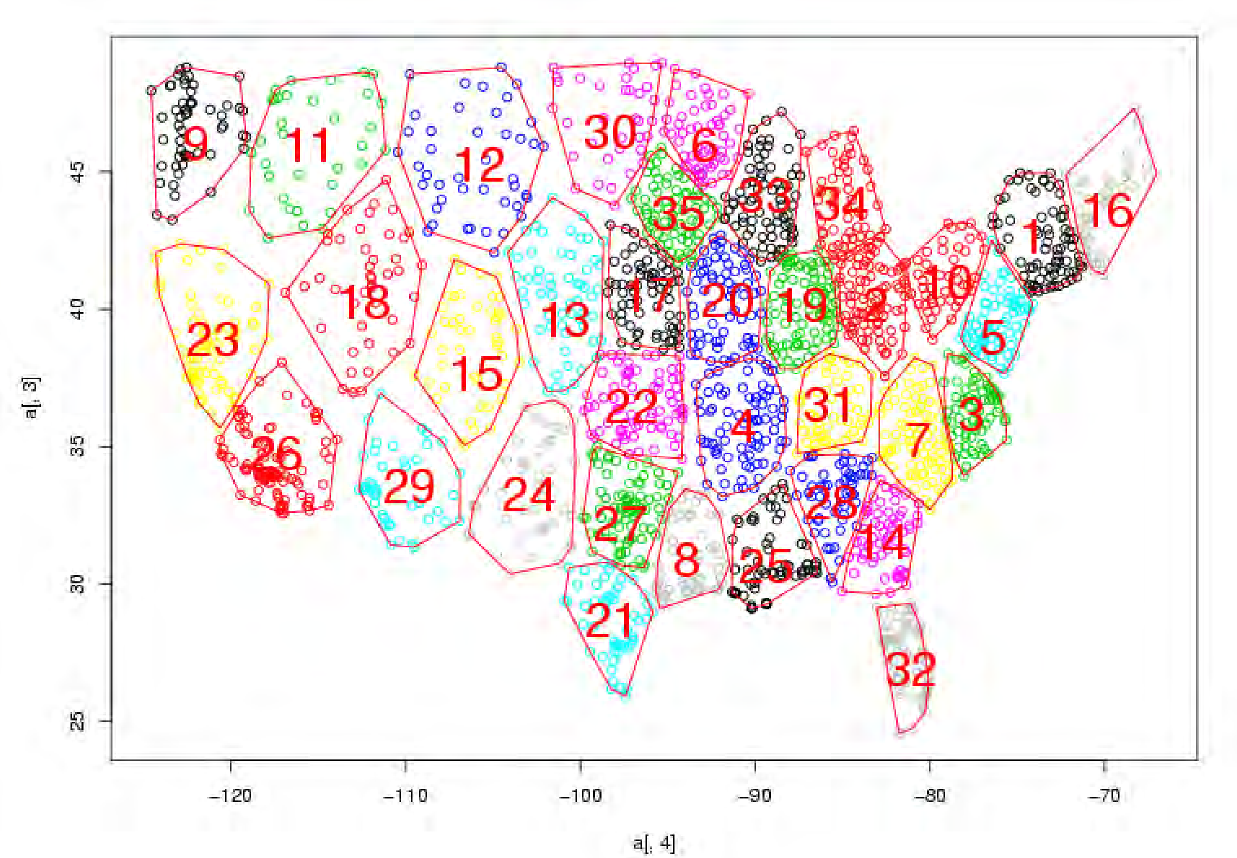

Figure 2. Three-thousand+ airports of the CONUS, grouped into 35 clusters geographically

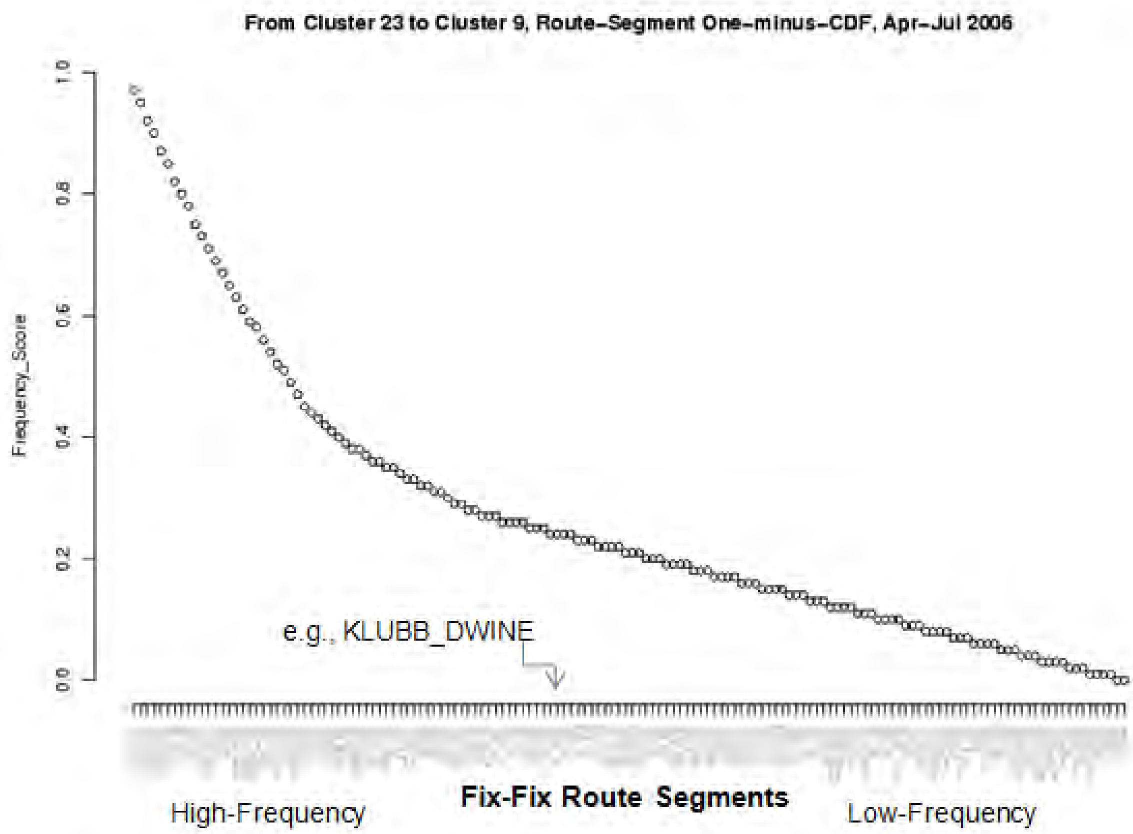

Figure 3. Mapping route segments into a "frequency score" for the 23-9 cluster pair

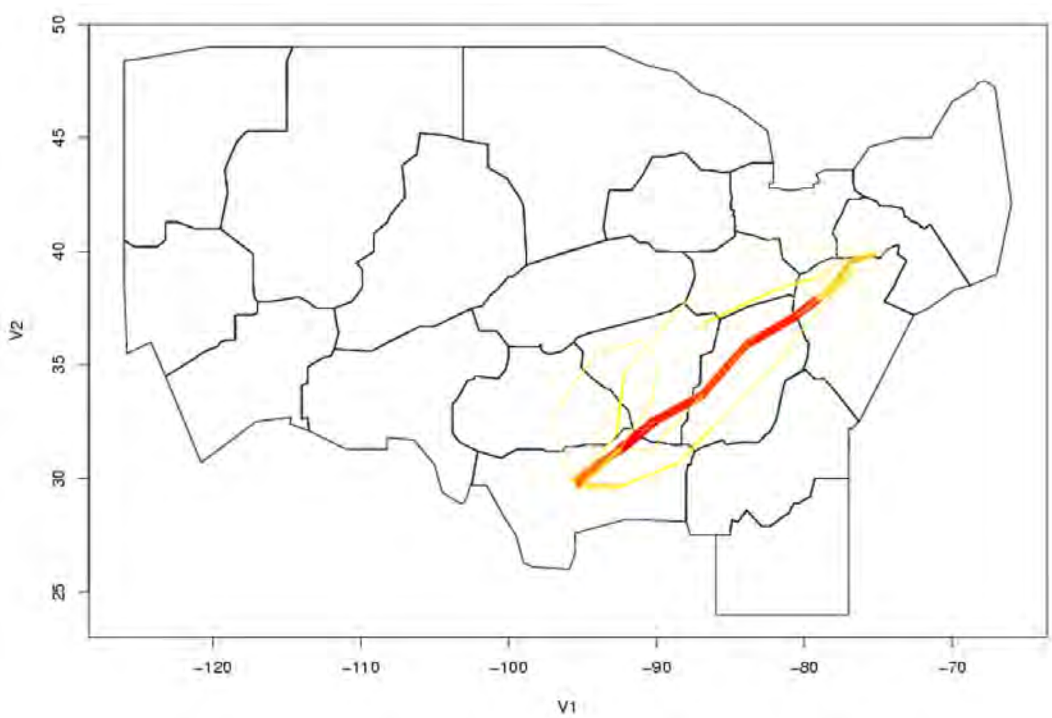

Figure 4. Route segment frequencies, cluster pair 5-8 (major flow is NYC area to Houston)

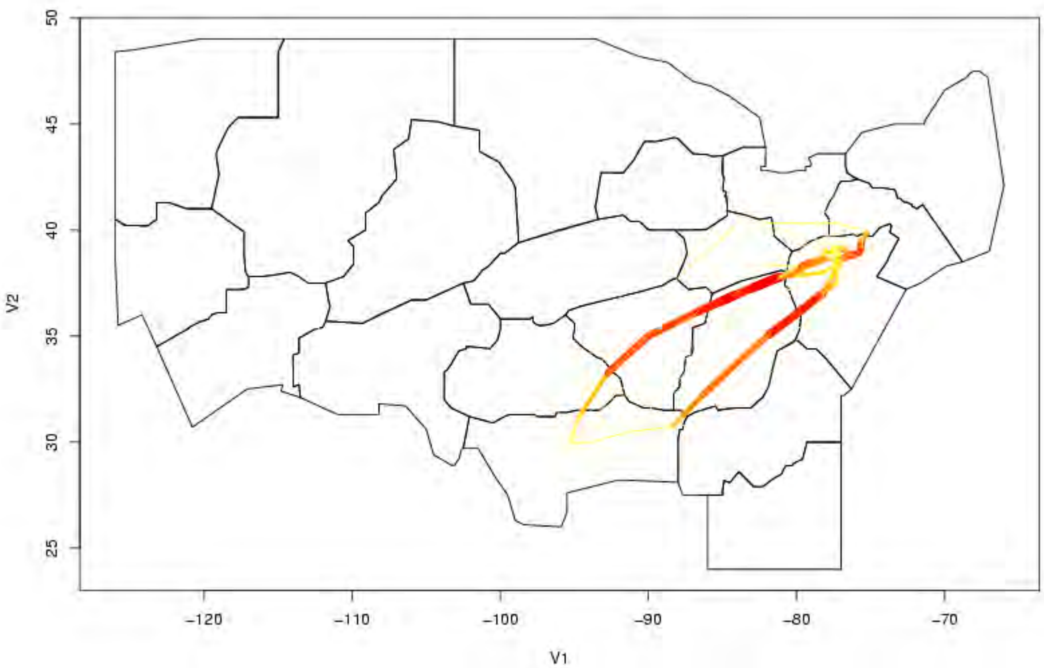

Figure 5. Route segment frequencies, cluster pair 8-5 (major flow is Houston to NYC area)

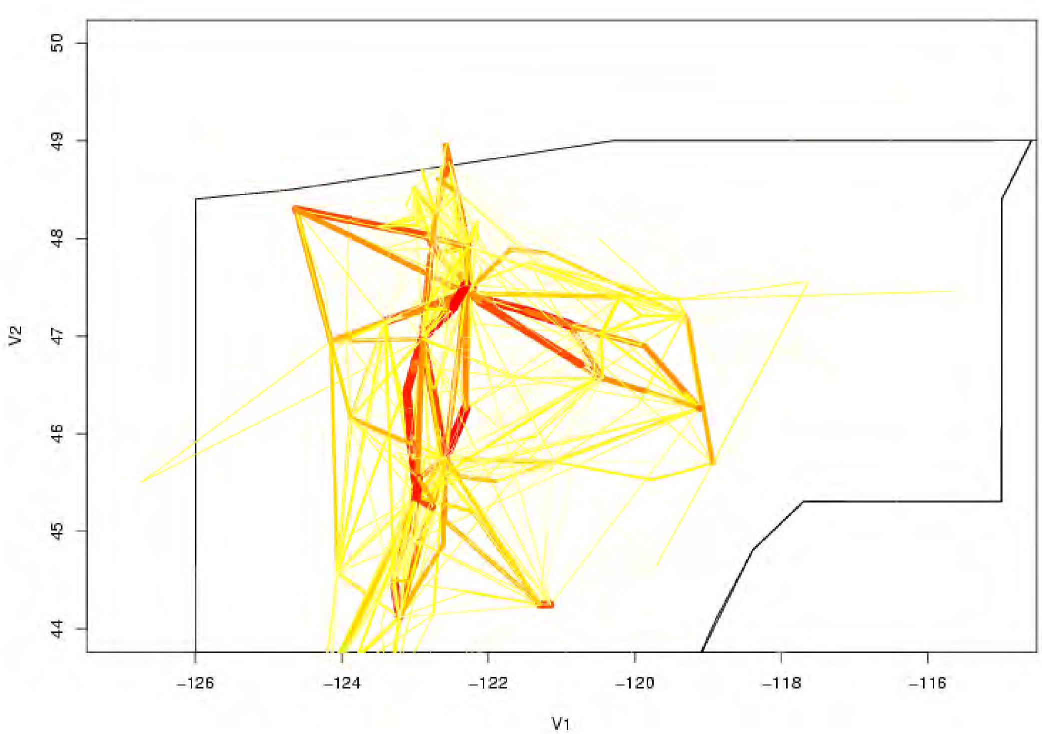

Figure 6. Route segment frequencies, cluster pair 9-9 (flows within the extreme north-west corner of the conterminous U.S.)

V. Summary

The exposition here has highlighted two of several important factors in automatic scoring of the operational acceptability of candidate alternate routing. The research is continuing to quantify more factors, and to use these in automation-driven decision support tools for the mid-term NextGen traffic management toolset. In fact, these factors have been used as input to a new dynamic route generation algorithm.[6]

Appendix: Selecting the Number of Clusters of Airports

In another study requiring geographic grouping of airports for cluster-to-cluster examination of flights,[7] it was found that 25 clusters was a reasonable number. It was suspected that a number of clusters somewhere in the neighborhood of 25 might be a good number for the problem at hand, so an experiment was undertaken considering the range of scores for each FPRS across cluster pairings.

As an example of a large range, consider Table A-1, which shows the range of scores across cluster pairs, for the 25-cluster grouping of airports, for the FPRS JARRD-to-TOOOK (near Cincinnati). The first column of the table shows the score, between 0 and 1, for the FPRS in the various cluster pairs shown in column 2. For example, at the extremes, the FPRS scores 0.02 (flights not often "with the flow") in the cluster pair 19-12, but scores 0.82 (flights often "with the flow") in the cluster pair 4-12. This difference in scores indicates a wide range of desirability of the route segment, varying as a function of origin/destination airport cluster. The large range likely offers a breadth and richness in the solution set of flight reroute possibilities.

Table A-1. Variety of Usage Scores for Fix Pair Route Segment JARRD_TOOOK, for 25 Clusters. The range is 0.82-0.02 = 0.8.

Score	Cluster Pair
0.02	6-12
0.02	19-12
0.05	10-12
0.05	18-12
0.08	17-12
0.11	20-12
0.11	21-12
0.14	23-12
0.17	12-12
0.20	2-12
0.21	1-12
0.32	22-12
0.33	16-12
0.46	11-12
0.56	13-12
0.57	3-12
0.61	8-12
0.62	7-12
0.77	14-12
0.82	4-12

The range of scores across the set of cluster pairs can be computed for each FPRS, and a trade-off exists between number of clusters and the average range of scores per FPRS for a given number of clusters. It was decided to determine the appropriate number of clusters using two measures of average score range:

- Average range for a FPRS across cluster pairs, for ranges of values where $n >= 1$
- Average range for a FPRS across cluster pairs, for ranges of values where $n > 4$

These two measures are denoted AvgRngNge1 and AvgRngNgt4 in Fig. A-1; note that they track similarly.

It can be seen from the figure that there is a large response in average score range as the number of clusters increases from 1 to 50, and little response as number of clusters increases from 50 to 200. At about 35 clusters, there was a dampening in this response and so 35 was selected as a reasonable number of clusters to be used in the study.

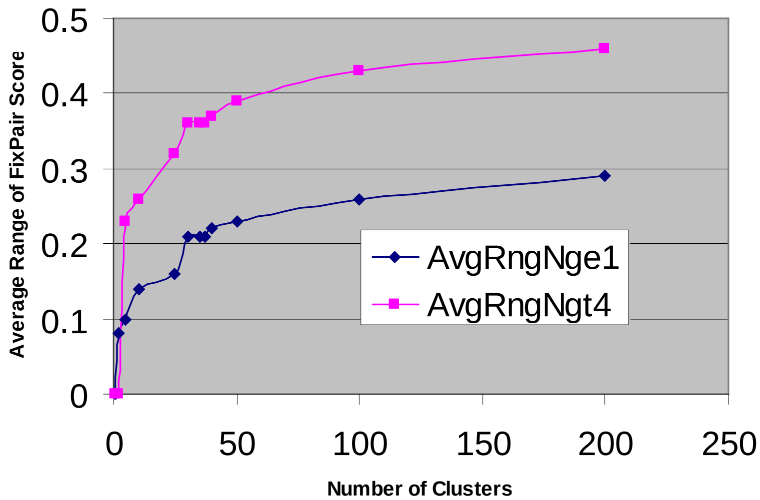

Figure A-1. Number of Clusters vs. Average Range of Fix Pair Route Segment (FPRS) Score

Acknowledgments

The authors appreciate helpful discussions with Mike Klinker, Jon Mintzer, Laurel Rhodes, and Steve Wyloge.

References

[1]Joint Program Development Office (JPDO), "Concept of Operations for the Next Generation Air Transportation System," Version 2.0, 13 June 2007.

[2]Prete, J., Krozel, J., Mitchell, J., and Kim, J., "Flexible, Performance-based Route Planning for Super-Dense Operations," *AIAA Guidance, Navigation and Control Conference and Exhibit,* Honolulu, HI, 18-21 August, 2008.

[3]Kuo, Vincent H., "Flight Plan Optimization in the Presence of Inadmissible Airspace," *AIAA Guidance, Navigation and Control Conference and Exhibit,* Honolulu, HI, 18-21 August, 2008.

[4]Berge, M. et al., "Benefits of Collaborative Flow Management During Convective Weather Disruption," ATM Conference, 2007, Barcelona, Spain.

[5]Rhodes, L. et al., "Probabilistic Automation-Assisted Congestion Management for En Route (PACER): Phase 1 Updated Operational Concept," MP06W000144, The MITRE Corporation, McLean, Virginia, 2006.

[6]Taylor, C. and Wanke, C. "Dynamic Generation of Operationally Acceptable Reroutes" *AIAA Aviation Technology, Integration, and Operations Conference*, Hilton Head, South Carolina, Sept. 2009.

[7]DeArmon, J. et al., "Excess Flying Time in the National Airspace System," ATC Quarterly, 2006.

Disclaimer

This work was produced for the U.S. Government under Contract DTFA01-01-C-00001 and is subject to Federal Aviation Administration Acquisition Management System Clause 3.5-13, Rights In Data-General, Alt. III and Alt. IV (Oct. 1996).

The contents of this material reflect the views of the authors. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty, guarantee, or promise, expressed or implied, concerning the accuracy of the views expressed herein.