Solving Probabilistic Airspace Congestion: Preliminary Benefits Analysis

Jim DeArmon, [*] Craig Wanke,[†] Dan Greenbaum, Jr.,[‡] Lixia Song,[§]
Sandeep Mulgund,[**] Steve Zobell,[††] and Neera Sood[‡‡]
The MITRE Corporation, Center for Advanced Aviation System Development (CAASD)
7515 Colshire Drive, McLean, Virginia, 22102-7508

In the U.S. National Airspace System (NAS) a function called traffic flow management (TFM) seeks a balance between resource capacities and the demands placed upon them by air traffic. In general, capacity cannot be manipulated, and it is necessary for demand to be altered to meet a reduced capacity. Typically, demand can be altered in time (via delay, i.e., slowing flights so that the number per unit time is reduced) or space (via rerouting, when specific airspace sector capacity is reduced, e.g., during severe en route weather). This paper discusses the use of probability modeling for assessing airspace capacity, and discusses comparison of three techniques for generating solutions to the problem of demand allocation during reduced airspace capacity caused by severe en route weather.

Nomenclature

ATC	=	Air traffic control
ATM	=	Air traffic management
GA	=	Genetic algorithm
LAT	=	Look-ahead time
MAP	=	Monitor/alert parameter
NAS	=	National airspace system
TFM	=	Traffic flow management
TMU	=	Traffic management unit

I. Introduction

IN the U.S. National Airspace System (NAS) a function called Air Traffic Management (ATM) consists of air traffic control (ATC) and traffic flow management (TFM). ATC provides separation services, keeping a minimum distance or altitude between proximate aircraft. By contrast, the purview of TFM is more strategic, seeking a balance between resource capacities and the demands placed upon them by air traffic. In general, capacity cannot be manipulated, and it is necessary for demand to be altered to meet a reduced capacity. Typically, demand can be altered in time (via delay, i.e., slowing flights so that the number per unit time is reduced) or space (via rerouting, when specific airspace sector capacity is reduced, e.g., during severe en route weather). This paper discusses the use of probability modeling for assessing airspace capacity, and discusses comparison of three techniques for generating solutions to the problem of demand allocation during reduced airspace capacity caused by severe en route weather.

A probabilistic approach to addressing air traffic flow problems is recognition of inherent uncertainty. Until now, traffic flow management (TFM) decisions have relied on a simple verdict: "At future time t, NAS resource A (e.g., a fix, sector, route segment, runway) is forecast to be over demand." However, modeling of the situation can be improved by explicit consideration of multiple future "states of nature," i.e., forecast demand and capacity of

---

[*] Multi-Discipline Sys Eng, Prin, Terminal/TFM System Eng, N450.
[†] Simulation Modeling Eng, Sr Prin, Terminal/TFM System Eng, N450, Senior Member AIAA.
[‡] Software Systems Eng, Lead, Terminal/TFM System Eng, N450.
[§] Simulation Modeling Eng, Sr, Terminal/TFM System Eng, N450.
[**] Info Systems Eng, Lead, HSI, Visualization & Dec Supt, M315.
[††] Software Systems Eng, Lead, Terminal/TFM System Eng, N450.
[‡‡] Netwk Sys&Dist Sys Eng, Sr, Modeling, Simulation & Econ, N590.

NAS resources.[1,2] Through improved modeling, more efficient TFM solutions for airspace users may be expected. Alternatively, some situations would be better left alone, with TFM *not* intervening—these cases should also be more easily recognized with improved modeling. Continuing improvements in weather forecast accuracy and flight trajectory modeling can also be used directly to improve decision-making.

II. Managing Airspace Capacity/Demand

In today's system, sector capacities are characterized by a Monitor/Alert Parameter (MAP), an aircraft count threshold. Though many factors determine true sector capacity, e.g., number of potential conflicts, number of flights in altitude transition, and traffic complexity, this scalar MAP value is used. And the use is fairly simplistic: a visual display (see Fig. 1) shows cells containing future quarter-hourly forecast aircraft counts per sector, and color-codes these cells, as follows:

1) Green: no alert (predicted count below or equal to the MAP value)
2) Yellow: total forecast count exceeds MAP value, but among them, active flights alone do not exceed the MAP value
3) Red: alert—the forecast count contains active flights which, even without including inactive flights, exceed the MAP threshold

Note that this scheme supports simple decision rules: green—do nothing; red—examine the situation: intervention may be indicated; yellow—monitor the situation. The probabilistic approach seeks a more complete characterization of forecast information using statistical distributions.

As a notional example, consider Figs. 2 and 3.[1] In Fig. 2, using weather forecasts, contours have been constructed to delineate regions of likelihood of severe en route weather, at levels > 50% and > 75% likelihood. Considering these impacts on airspace sector capacity and staffing, as well as expected air traffic demand, a probabilistic score can be associated with affected sectors, here > 50% and > 75% probability of congestion. Figure 3 shows the demand vs. capacity situation for Sector 2 over time. The blue boxes show capacity, expressed as an expected or mean value in the middle of the box, plus 50% error bounds at the top and bottom of the box. Green, yellow, and red boxes, corresponding to increasing probability of congestion, show expected demand, using the same format of mean and confidence interval bounds. (Boxes have mean lines not evenly splitting the vertical extent of the box, since the probability distributions are typically non-normal and even non-symmetric). A traffic flow manager could manage resources using an automation tool with a display as in Fig. 3, by reducing demand (either in time via delay or in space via alternate routing) until the red boxes become lower on the display (the color would thence become yellow), matching better the blue underlying capacity distributions.

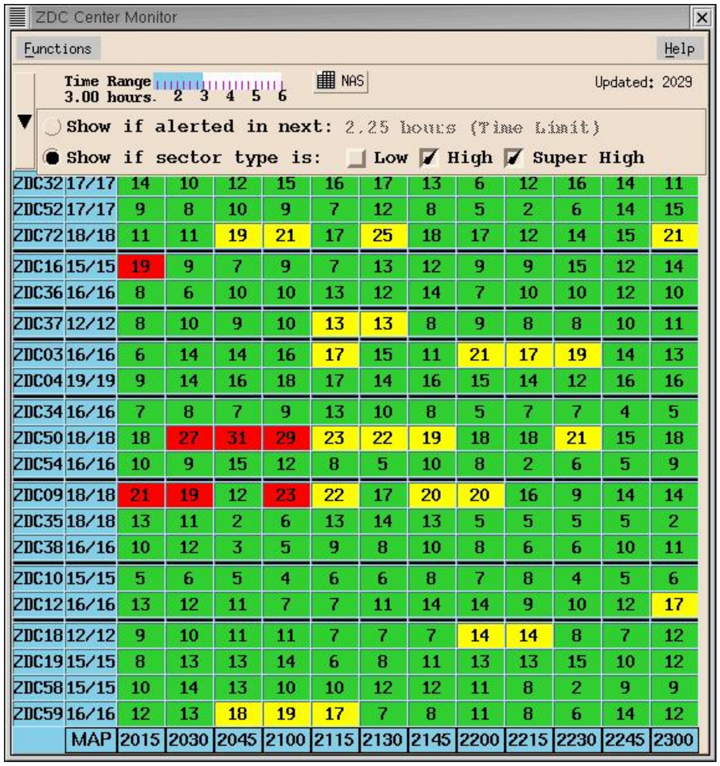

ZDC32 17/17	14	10	12	15	16	17	13	6	12	16	14	11
ZDC52 17/17	9	8	10	9	7	12	8	5	2	6	14	15
ZDC72 18/18	11	11	19	21	17	25	18	17	12	14	15	21
ZDC16 15/15	19	9	7	9	7	13	12	9	9	15	12	14
ZDC36 16/16	8	6	10	10	13	12	14	7	10	10	12	10
ZDC37 12/12	8	10	9	10	13	13	8	9	8	8	10	11
ZDC03 16/16	6	14	14	16	17	15	11	21	17	19	14	13
ZDC04 19/19	9	14	16	18	17	14	16	15	14	12	16	16
ZDC34 16/16	7	8	7	9	13	10	8	5	7	7	4	5
ZDC50 18/18	18	27	31	29	23	22	19	18	18	21	15	18
ZDC54 16/16	10	9	15	12	8	5	10	8	2	6	5	9
ZDC09 18/18	21	19	12	23	22	17	20	20	16	9	14	14
ZDC35 18/18	13	11	2	6	13	14	13	5	5	5	5	2
ZDC38 16/16	10	12	3	5	9	8	10	8	6	6	10	11
ZDC10 15/15	5	6	5	4	6	6	8	7	8	4	5	6
ZDC12 16/16	13	12	11	7	7	11	14	14	9	10	12	17
ZDC18 12/12	9	10	11	11	7	7	7	14	14	8	7	12
ZDC19 15/15	8	13	13	14	6	8	11	13	13	15	10	12
ZDC58 15/15	10	14	13	10	10	12	12	11	8	2	9	9
ZDC59 16/16	12	13	18	19	17	7	8	11	8	6	14	12
MAP	2015	2030	2045	2100	2115	2130	2145	2200	2215	2230	2245	2300

Figure 1. Sector monitor highlights sectors that may need intervention.

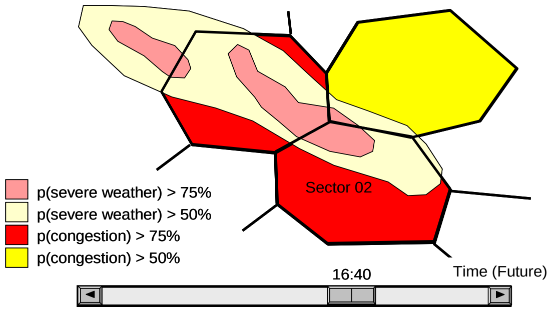

Figure 2. A probabilistic airspace congestion forecast for three ATC sectors

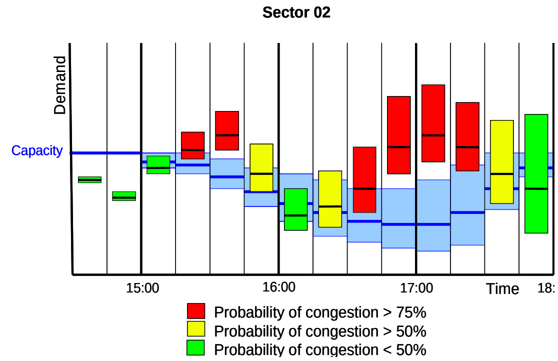

Figure 3. The congestion forecast over time shows transient and longer term conditions

Given there is an improved characterization of capacity, the question remains how best to manipulate the demand to meet capacity. For the problem at hand, airspace capacity loss caused by severe en route weather, it would be some combination of ground delay, air delay, rerouting, and "do nothing" on a per-flight basis. In today's system, flow managers assess the weather forecast, collaborate with air carriers, and then implement pre-stored routings from what is called the National Playbook. These routings are pre-coordinated among the various centers, and typically qualified to include some set of origin and destination airports. This solution approach is performed as a single action several hours before the weather event, with most affected flights still on the ground (implying that their pre-departure flight plan route is amended to comport with the TFM solution). As time moves on and the weather event becomes more predictable, adjustments may be performed on the original initiative. In addition, local actions by the Traffic Management Units (TMUs) further fine-tune the plan. There have, in recent years, been improved visualization products and more automated communication with airlines' flight planning systems, but the procedure is still mostly manual.

III. Solution Approaches

But looking ahead to the next wave of automation support for TFM, one can imagine greater harnessing of significant computing power and algorithms. This paper examines the potential benefits, by comparing today's solution (hereafter called Manual Approach) with two variants using advanced computation systems: a Heuristic Approach and a Genetic Algorithm (GA) Approach.

At the disposal of this analysis is a simulation platform for exploring prototype tools and ideas. This platform can represent airspace and its constraints, as well as 4-D trajectories of flights, and can evaluate future sector capacities probabilistically. To perform a comparison for benefits calculation, the three approaches—Manual, Heuristic, and GA—are executed, starting with an actual NAS traffic day and an actual severe en route weather event. (The input traffic data are actually taken from a different day—a good weather day when air traffic flew undeviated, since we want the competing approaches to begin with "pristine" intent data—what the airlines wish they could do, i.e., fly straight through the weather and expedite to the destination. The solution approaches will deviate the traffic to meet reduced sector capacities associated with the weather event.) Note that all three approaches implement a single action, a set of deviations for all affected flights. More realistic modeling is certainly plausible—to solve part of the problem, wait a while for events to unfold, resolve the new problem situation, and repeat. This approach being explored further by the authors.

To compare approaches, ground delay and airborne delay (caused by circuitous routing) are converted to dollar costs and ranked to evaluate the relative merit of each approach. In addition, metrics regarding the success in achieving target sector loading are compared.

A. Manual Approach

To represent the Manual Approach in our simulation system, daily logs of TFM actions were extracted for the bad weather day, and the Playbook routings for that day were used. For flights not put onto a Playbook routing, the Heuristic Approach was used, as described next.

B. Heuristic Approach

The Heuristic Approach may be an improvement over the Manual approach. Its logic is characterized as a "greedy" algorithm,[3] in that the first feasible option available for a flight is the one selected, without either look-ahead or back-tracking to improve the solution. The logic is as follows (see Fig. 4):

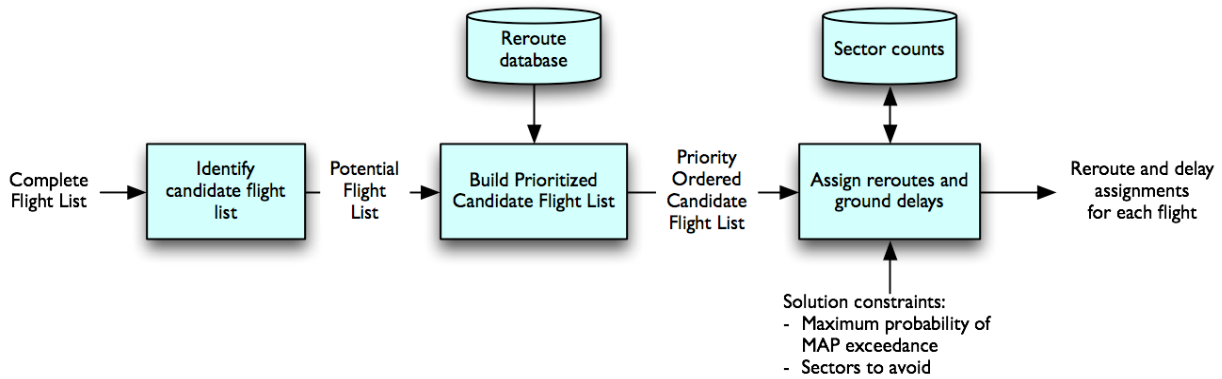

Figure 4. Overview of Heuristic Congestion Management Algorithm.

The first step is to identify all of the flights that enter the weather region in the time of interest. All such flights that are eligible for reroutes and delays are sorted into a processing order that is based on increasing time of arrival to the weather region; each flight is then removed, in order, from counts in the sector tally, and then added back one at a time. If the original flight plan does not violate any congestion constraints (e.g., the maximum acceptable probability of congestion), then the original flight plan is accepted. If it *does* violate the constraints, reroute and delay options are examined to find the option that provides the earliest arrival time to the flight's destination that satisfies the problem constraints. If no workable option is available, the flight is left on its original route. Alternative routings are selected from a database of historical routes between the appropriate origin and destination.

As the algorithm proceeds, ever fewer "degrees of freedom" remain, and the inclusion of ground delay is increasingly necessary. Hence the classification of the approach as "greedy"—without any consideration of flights not yet assigned a route and a (possibly delayed) take-off time, the current flight set is locked into the current solution.

C. Genetic Algorithm (GA) Approach

A GA Approach may further improve overall results, because the myopia of the Heuristic Approach is overcome. In brief, a GA Approach mimics, in certain ways, genetic mechanisms in the real world. A solution to the problem is represented as a string of values (in the natural world, a chromosome comprised of genes). For the problem at hand, the values are possible route/delay solutions for each flight. For example (see Fig. 5), the string: 0 6 4 2… means that flight #1 gets its solutions #0, flight #2 gets its solution #6, flight #3 gets its solution #4, flight #4 gets its solution #2… So ultimately, some string will be selected as having the lowest cost (ground + air delay) among all strings examined, and will be the final result of the algorithm. This string is not limited to a single priority ordering, but instead is chosen to meet an overall performance metric.

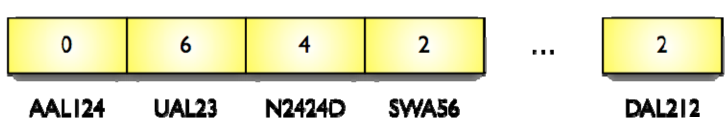

Figure 5. GA representation of potential solution.

The selection and manipulation of strings mimics the natural world: a population of strings is generated using random assignment other appropriate method**,** so if there were 100 flights that required intervention, and each flight had a solution set of 15 possible delay/routing values, then a single member of the population would be a vector of length 100, containing integers valued 0 to 14. As the algorithm proceeds, new strings are created by cutting and splicing between pairs of strings (analogous to genetic recombination in the natural world). Offspring from the current generation can become parents in the next generation, subject to a fitness test. Fitness is evaluated using the cost function—low cost strings are selected as having the potential to generate ever-better solutions. The analogue in the natural world is the "survival of the fittest" notion. And, as in the natural world, a final step of mutation or random perturbation of some string values has the potential to further improve fitness. (See Fig. 6 for a pictorial of the generational transition—recombination and mutation.) The GA Approach has been shown in the literature[4] to be a robust way to search complex spaces, and that is what is called for when the airspace congestion problem is sufficiently complex.

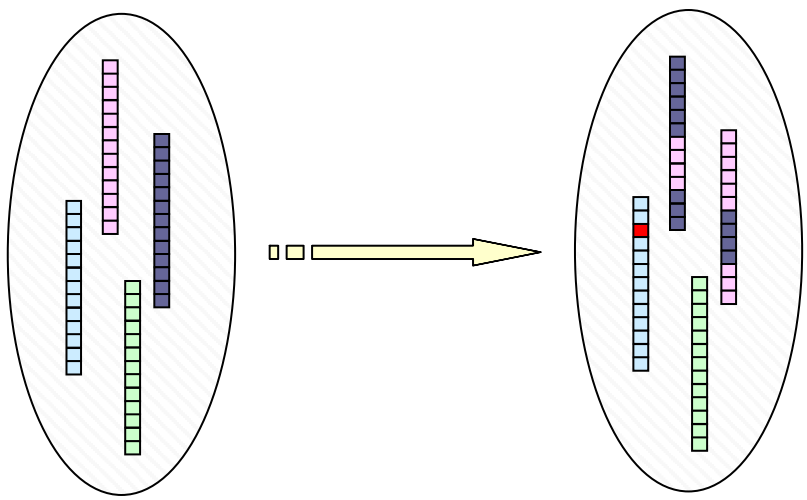

Figure 6. Recombination, mutation to make a new generation.

IV. Earlier Results

Several earlier analyses have been undertaken. Initial comparisons of Manual vs. Heuristic show promise—Heuristic bests Manual for two different bad-weather days and three different randomly selected flight sets, using a simple cost function of dollar value of ground delay and airborne delay.[5]

A comparison of Heuristic vs. GA[6] on three test problems (5 to 35 manipulated flights) examined two solutions per problem—one solution was constrained to be delay-only, the other solution allowed both delay and re-routing. In summary, the study showed the following:

1) Both Heuristic and GA were effective at solving the congestion problems, although Heuristic failed to find a feasible solution for one of the six (3 test problems × 2 solutions each).
2) Heuristic was most effective with the delay-only solutions; GA generally performed better when re-routes allowed.

The observation that GA sometimes did not best the Heuristic was initially surprising—it was initially assumed that GA should always best Heuristic, since it was a more elaborate approach, and was well suited to our problems, as described above. It was concluded that the GA was not converging, and perhaps not even getting close to an optimal solution. A new approach was considered— "seeding" the GA with the Heuristic solution. This was thought to be reasonable, because it would force the GA to start searching in a fruitful neighborhood of the search space. It would also guarantee that the GA would do no worse than the Heuristic**.** Initial experiments show promise.

These results suggest that the benefits of an optimizing algorithm may be largest for congestion scenarios with many degrees of freedom, and mutual objectives such as cost and equity that must be balanced in arriving at a solution. This is the type of problem most likely encountered in actual practice. In the following sections, a large, complex real-world problem is described and solutions are attempted by the three approaches.

V. A Complex Real-World Problem

A real-world problem was selected for comparing the three approaches described above. The day of May 31, 2004 was selected as being a challenging bad en route weather scenario. A line of thunderstorms hundreds of miles long, from Texas north to Ohio, formed and slowly moved east (see Fig. 7). Because it was a well organized system, and was moving east at a steady pace, it was well suited to TFM solutions that involved ground delay and/or pre-departure re-route formulation. The look-ahead time, per the activity logs of the TMU of the en route ATC facilities showed TFM initiatives being implemented a couple of hours before the weather event. Moreover, the TMU logs also reflected the complexity and demanding nature of the event—many local TFM initiatives were necessary to fine-tune the earlier strategic solutions as the day wore on, and the weather moved and morphed. The problem was thus deemed appropriate for comparing our three approaches, especially for exercising the GA Approach, because we had conjectured that that approach would best show its mettle in solution of a complex problem.

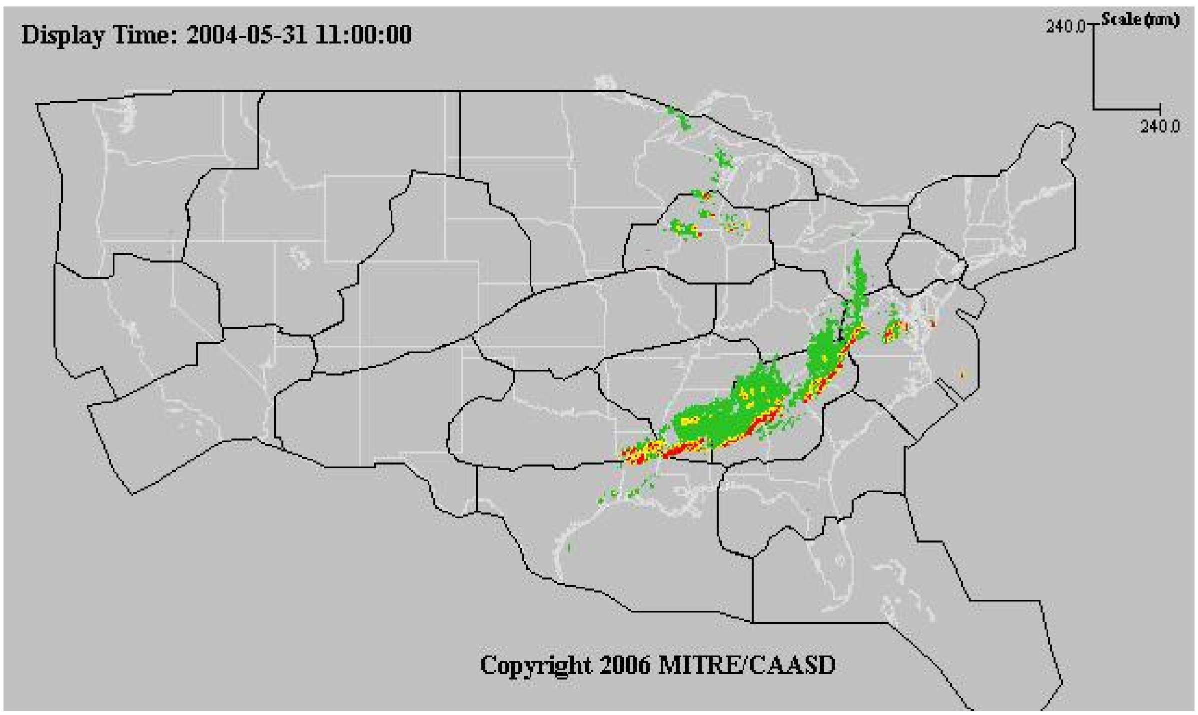

Figure 7. En route weather at 11:00 GMT on May 31, 2004.

To represent this problem for our computer-implemented algorithms, two geographic regions were identified. See Fig. 8 for a layout of the airspace sectors and overlays of the two regions. The inner region, circumscribed by a dotted line is the nominal estimate of the location of the severe en route weather (en route meaning that the weather is high-altitude, and interferes with flights at their cruise altitude.) For simplicity, our scenario definition had the weather stationary and fixed in form, even though it moved and morphed in the real-world. Within the inner region, sectors are subject to reduced capacity because of the weather—a rough estimate of capacity loss was made, using simply the area of coverage. If a sector was 50% covered by the inner region of Fig. 8, then it was assumed that the sector had a 50% capacity loss resulting from severe weather. The outer region, inscribed by the bolder, dashed line is the "area of interest"—any flights whose good-weather flight plan path penetrates this region are subject to manipulation. All of the sectors covered by the outer region are monitored for load, and flight assignment for sector traversal is compared against the already assigned demand in the sector. Eight hundred twenty-eight flights penetrated the outer region, making the problem quite complex.

Next, the three approaches to problem solution are described.

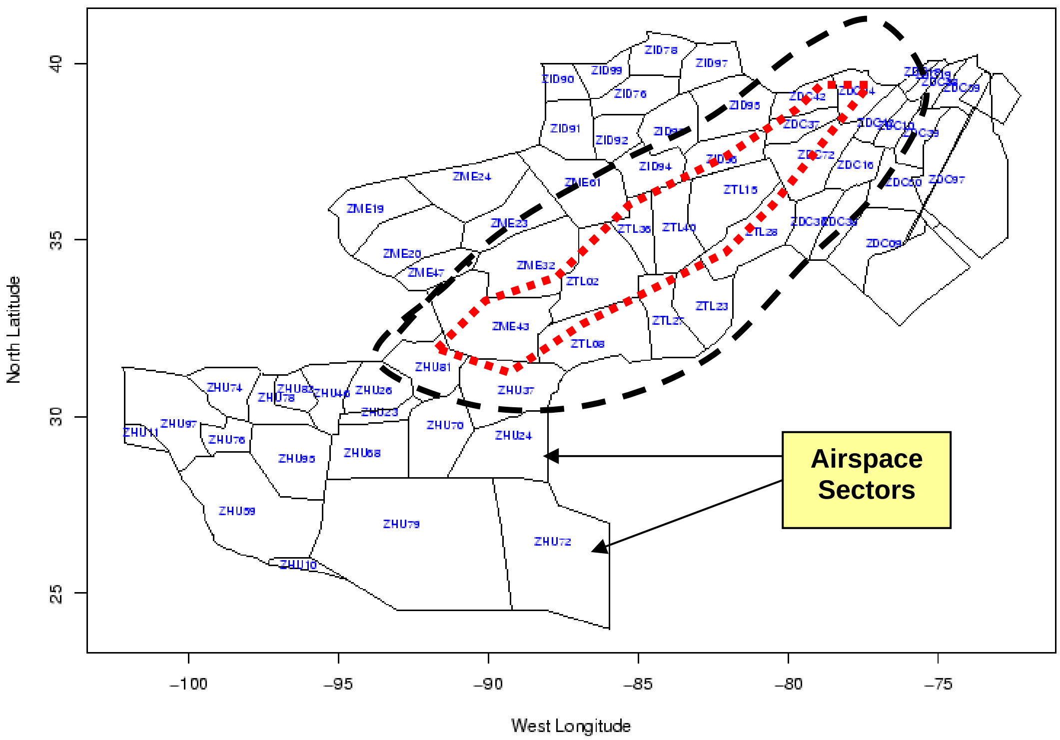

Figure 8. Estimating effect of weather on sector capacity.

A. Manual Approach

The Manual Approach was supposed to mimic the current-day solution of:

1) Selection of strategic Playbook routes to offload major demand components onto pre-stored and pre-coordinated routes and
2) Local fine-tuning of the remaining demand via ground delay and tactical (shorter look-ahead) re-routing.

To effect this, the TMU activity logs were examined, and the major Playbook routings of that day were considered. As shown in Fig. 9, Playbook routings can go well wide of the actual weather location. Note also in that figure the re-routes from Florida to the Nation's Capital. Because traffic diverted east of the weather would be interfering with the major flows north from Atlanta and Florida, it was necessary to move some Florida traffic even further east, even involving some tracks over the ocean.

Playbook routings are along major existing airways, and are *pre-coordinated*, meaning that those facilities affected by the route have all agreed to the operational feasibility of the routing beforehand. Even though there are scores of Playbook routes available, they cannot possibly route around all weather situations efficiently. This phenomenon is one of the hypothesized benefits mechanisms at play—the Playbook routes are quick and easy to implement, but they can often involve an excursion away from the weather. By utilizing a database of heterogeneous historical routes, the Heuristic and GA Approaches had the opportunity to gain efficiency over Playbooks two ways:

1) By just skirting the weather region, avoiding costly excursions and
2) By spreading traffic out, avoiding the bottlenecks resulting from many deviated flights taking the same re-route.

An open research question is the operational acceptability of a heterogeneity of flights in light of the uncertainties of the weather day.

Using our off-line TFM planning facility, over 100 flights were put onto these routes. The remainder of flights were assumed to have received more tactical actions—ground hold/delay and local re-routes with shorter look-ahead, as time passed and the weather event neared. To mimic the tactical TFM initiatives, the Heuristic Approach was applied to the flights not assigned to a Playbook route.

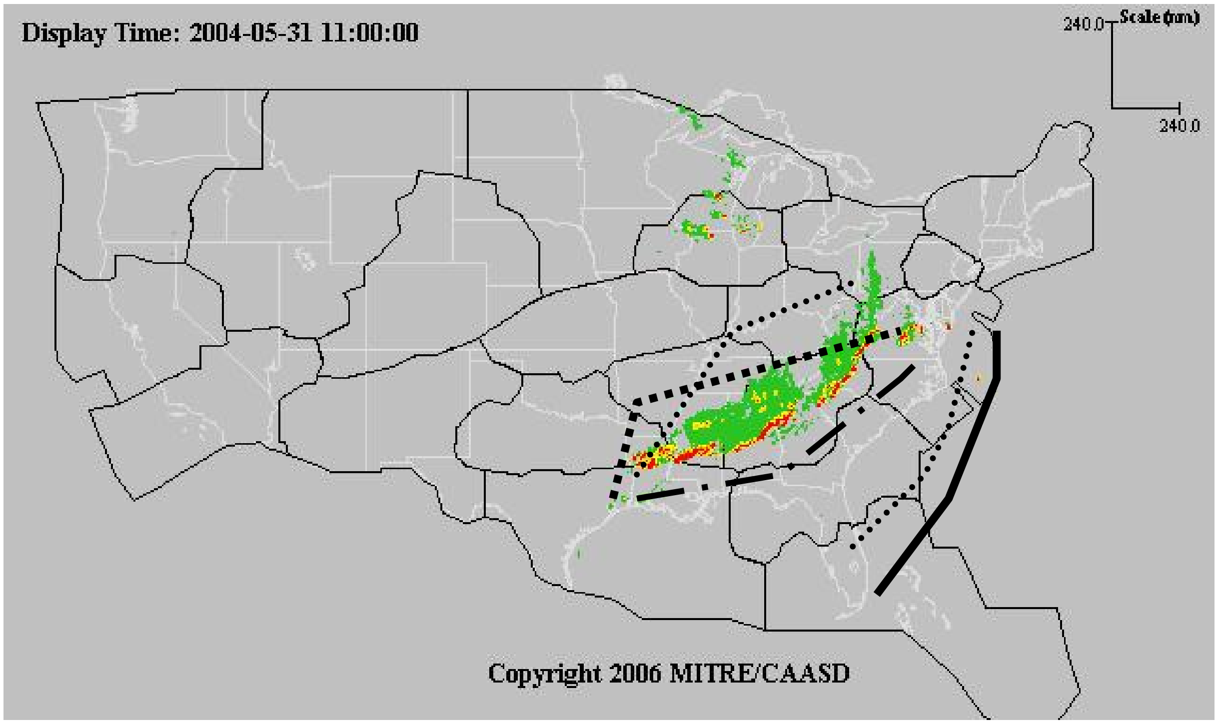

Figure 9. Playbook routing around severe weather.

B. Heuristic Approach

The Heuristic Approach was applied to the selected scenario. Flights were sorted on their nominal unperturbed time of first encountering the weather system. Then, one-by-one, a re-route/delayed take-off was assigned to each flight, on a "first-fit" basis. That is, the first feasible re-route/delayed take-off combination was selected for a flight. As successive flights on the sorted list of eight hundred flights were handled, fewer degrees of freedom were available. Flights near the end of the list might receive delay and routing that was extreme and likely not very acceptable to the air carrier, or, alternately, it happened that some flights could generate no feasible solution. That is, there remained no option, in terms of ground delay and re-routing, that could still satisfy the goal of a fixed probability of sector congestion for all sectors in the outer region. This situation was allowed, and was accounted for in the "probability of remaining congestion"-style metrics used to evaluate and compare the solution approaches.

C. GA Approach

It was supposed that the GA Approach would do well on this large, complex scenario. Given 828 flights needing manipulation, and 8 or 10 delay/route options for each, the sample space was vast and considerably non-smooth. A GA is well suited to this situation—it efficiently searches large, ill-behaved sample spaces. However, either because of insufficient run-time, or the challenge of setting the tuning parameters correctly, the GA Approach did not best the Heuristic Approach, as discussed below. The idea of "seeding" the GA with results from the Heuristic approach are being pursued, but no results are yet available.

VI. Results

Four metrics were computed for the three solution approaches. Figure 10 shows the delay costs in thousands of dollars, associated with the ground delay and re-routes (re-routes typically add time to a flight, because the undeviated path is generally very direct.) Of the three approaches, Heuristic has the lowest cost, with the Best GA second lowest. (For these metrics, several variant runs of the GA were undertaken, with varying parameter values. The best of these variant runs was used to represent the GA result. Because no single GA solution generated the four metrics values described in this section, it may be optimistic to assume that, with proper parameter settings, a single GA solution will be found that can exhibit the performance shown here.)

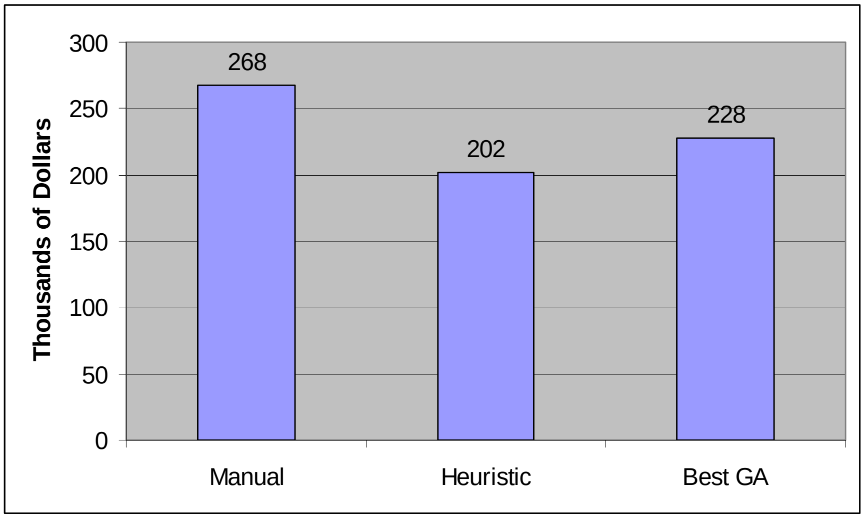

Figure 10. Delay costs, in thousands of dollars (smaller is better).

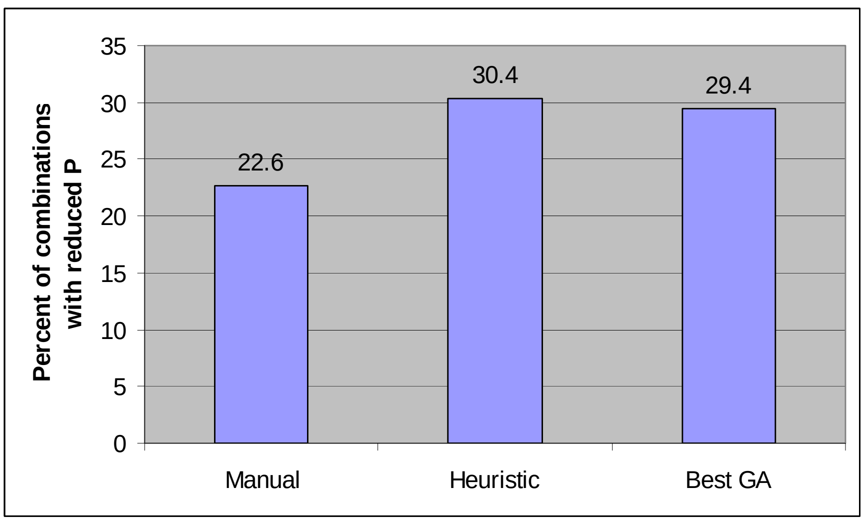

Figure 11. Percent of sector/look-ahead time combinations wherein P(overage) reduced to < 50%, given that unperturbed P was > 50% (larger is better).

Figure 11 examines one aspect of congestion. It shows the percent of sector/look-ahead-time combinations in which the probability of sector congestion was reduced to below 50%, given that the unperturbed probability was above 50%. (Each sector of interest is examined at ten different look-ahead times: from 60 minutes to 3 hrs 15 minutes. Although there could easily be correlation in the results—if a sector is congested 60 minutes from now, it might similarly be congested 75 minutes from now—all of the sector/look-ahead-time combinations are evaluated on an equal basis.) This metric measures the success in lowering the probability of congestion, for the sectors that need it, and a larger value is better. In this case, Heuristic slightly bests the best GA solution.

Figure 12 shows a metric that is somewhat the reverse of the prior one. It shows the percent of sector/look-ahead-time combinations in which the probability of sector congestion was *increased to above* 50%, given that the unperturbed probability was *below* 50%. So, this is the amount of "trouble caused" by the approach, in its attempt to solve the congestion problem. In this case, a smaller value is better, and the best GA solution is the winner.

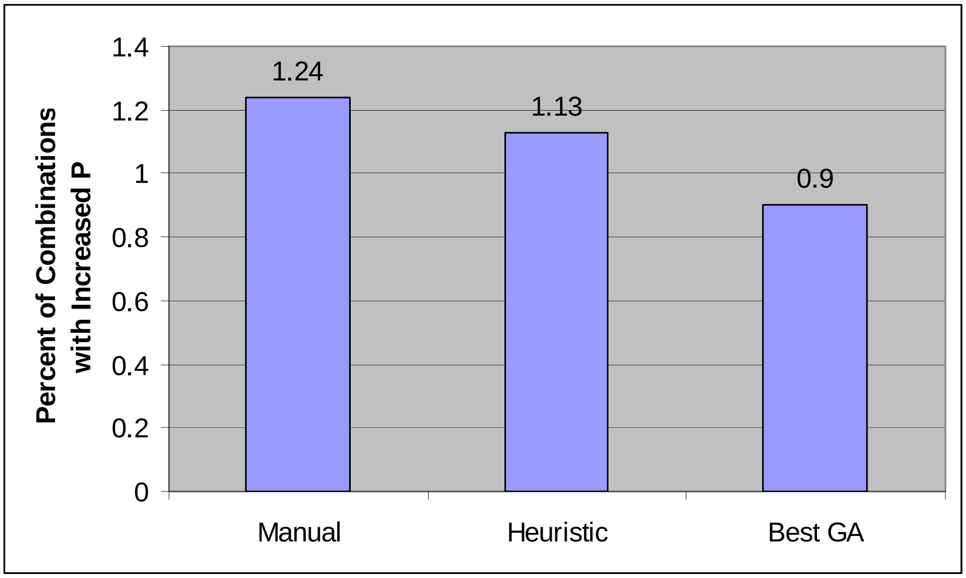

Figure 12. Percent of sector/look-ahead time combinations wherein P(overage) increased to > 50%, given that unperturbed P was < 50% (smaller is better).

Finally, Fig. 13 shows the average change (over applicable sector/LAT combinations) in P(congestion), given that the unperturbed P was > 50%. The most favorable result is a large negative value. (Note the reversed sense of the Y-axis in the figure.) Heuristic clearly bests the Best GA here.

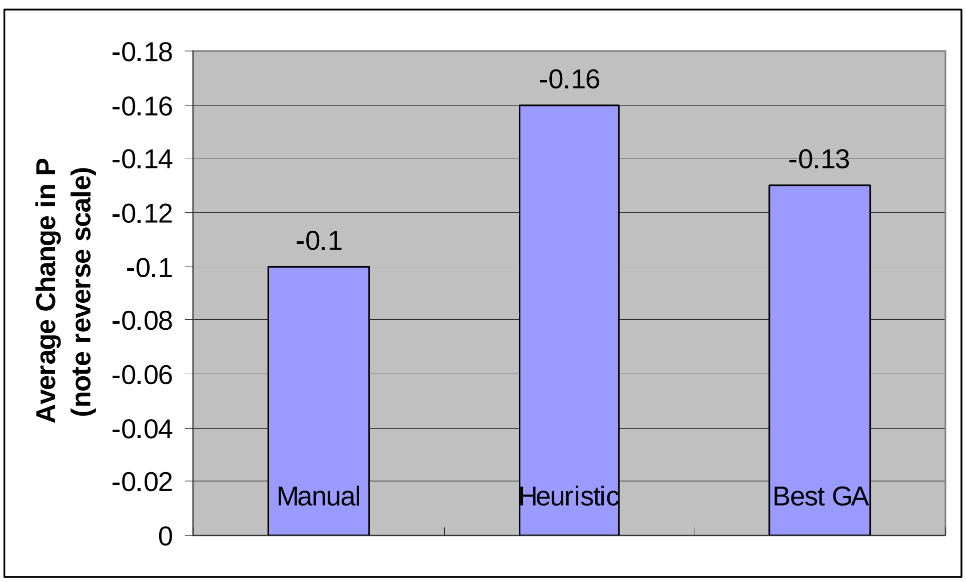

Figure 13. Average change (treatment minus unperturbed) in P(overage), given that unperturbed P was > 50% (large negative is better).

VII. Summary and Next Steps

A large, complex real-world problem has been attempted using three algorithms. The Manual Approach is a modeled representation of how today's TFM system might handle a severe en route weather situation. Two alternative, probability-based machine solutions, called Heuristic Approach and Genetic Algorithm Approach were also modeled and compared. The problem was quite challenging – a large severe weather mass from Texas to Ohio, affecting hundreds of commercial flights.

Four metrics were captured and compared to assess the three approaches, one metric on the dollar cost (to the air carriers) of flight deviations, and three metrics characterizing the "probability of sector congestion" criterion active in the model mechanisms. As a general ranking among the three approaches, Heuristic was best, GA was second, and Manual was last. It is somewhat surprising for the Heuristic Approach to best the GA Approach – the GA is a more sophisticated approach, and should do well on the large complex problem. It is assumed that a combination of insufficient computer run time and ill-tuning of internal parameters are the causes of the GA Approach's poor performance.

A idea being pursued is the "seeding" of the GA initial population with final results of the Heuristic Approach. This may lead the GA to neighborhoods in the search space which are promising in terms of objective function. This stratagem has been successfully applied to smaller problems, and is being adapted to this one.

Finally, the application of these solutions needs to be explored. Certainly they could be used as a "prescription" to alleviate a congestion problem, but they may also serve as the starting point for a collaborative solution process among FAA facilities and aircraft operators. Appropriate mechanisms for this process are being explored.

References

[1]Wanke, C., Zobell, S., and Song, L., "Probabilistic Airspace Congestion Management," 5th Aviation Technology, Integration, and Operations Conference, Arlington, Virginia, 26-28 September 2005.

[2]Krozel,J., D. Rosman, and S.R. Grabbe, "Analysis of En Route Sector Demand Error Sources", AIAA 2002-5016, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, California, USA 5-8 August 2002.

[3]National Institute of Standards and Technology, *Dictionary of Algorithms and Data Structures*, URL: http://www.nist.gov/dads/,

[4]Goldberg, D., *Genetic Algorithms in Search, Optimization, and Machine Learning,* Addison-Wesley Publishing Co, New York, 1989.

[5]Federal Aviation Administration, "Benefits Basis of Estimate (BOE) for Traffic Flow Management Modernization (TFM-M)," Washington, D.C., March 2005.

[6]Mulgund, S., Wanke, C., Greenbaum, D., and Sood, N., "A Genetic Algorithm Approach to Probabilistic Airspace Congestion Management," AIAA Guidance, Navigation and Control Conference, Keystone, Colorado, August 2006.

Notice