MIT Open Access Articles

*Dynamic reconfiguration of terminal
airspace during convective weather*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Michalek, Diana, and Hamsa Balakrishnan. "Dynamic Reconfiguration of Terminal Airspace During Convective Weather." 49th IEEE Conference on Decision and Control (CDC). Atlanta, GA, USA, 2010. 4875-4881. © Copyright 2010 IEEE

As Published: <http://dx.doi.org/10.1109/CDC.2010.5718007>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/66132>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

Dynamic Reconfiguration of Terminal Airspace During Convective Weather

Diana Michalek and Hamsa Balakrishnan

Abstract—Dynamic airspace configuration (DAC) algorithms strive to restructure the U.S. National Airspace System (NAS) in ways that allow air traffic control to better manage aircraft flows. Although past research has largely focused on enroute airspace in clear weather conditions, the principle of better matching airspace structure to ambient conditions has potential to benefit airport terminal areas, which are often impacted by congestion due to convective weather, especially during summer months when travel demand is high. This paper studies the problem of dynamic airspace configuration in the terminal area given a stochastic model of route availability during convective weather conditions. An integer programming model is proposed for the dynamic reconfiguration of the terminal area. This model recommends small changes to airspace structure that alleviate airspace congestion, while limiting disruptions to air traffic control procedures. The model is tested against actual weather scenarios, and shows promising benefits to operations.

I. INTRODUCTION

The growth in demand for air traffic operations in the United States has made the system particularly susceptible to weather-related disruptions. In 2009, 44% of the total minutes of flight delays in the U.S. were due to weather [1]. Convective weather, in particular, is responsible for large delays and widespread disruptions in the National Airspace System (NAS), especially during summer months when travel demand is high. Efficiently operating the airspace system in the presence of weather requires the integration of weather forecast products into air traffic management decision-making. One strategy for managing aircraft during periods of decreased airspace capacity due to the presence of storms is to relax the rigid structure of airspace and reconfigure airspace more effectively given the specific demand and weather conditions.

Currently, aircraft flying under instrument flight rules follow filed flight plans which are represented by standard way-points connected by airways. In the terminal-area, a flight follows a Standard Instrument Departure (SID) when departing an airport, and a Standard Terminal Arrival Route (STAR) into its destination airport. The corresponding arrival fixes, departure fixes, and terminal airspace sectorization, are fixed, even when the presence of hazardous weather renders them unusable. There is clear potential to recover

lost capacity by dynamically altering the terminal airspace structure in the presence of adverse weather. This concept of relaxing the rigid structure of airspace and the resultant potential increase in airspace capacity has been identified by the NextGen ATM-Weather integration plan [2].

The goal of this paper is to identify and evaluate gentle strategies for reconfiguring airspace, without drastically rearranging airspace structure. We explore potential benefits of making relatively small displacements in existing sector boundaries and fixes (displacements that are performed in practice on an ad hoc basis in order to temporarily increase arrival or departure throughput), thereby limiting disruption to existing air traffic control procedures. To this end, we develop an integer programming approach to optimally choose terminal area arrival and departure fixes as well as sector boundaries, for a given weather forecast, subject to constraints on displacement from today's fixed airspace structure. In prior work, the authors have developed a route availability forecast that uses the Lincoln Lab CIWS weather product [3], and predicts the probability p_r that a given route r will be open (subject to certain assumptions and constraints to be described later) for a fixed horizon [4]. In this paper, the route availability forecast will guide the selection of fixes that are likely to be open when weather materializes, although this selection will be traded off against the deviation from the default terminal area configuration.

The structure of this paper is as follows: Section II discusses background research, Section III introduces the terminal area model as well as the weather forecast used as input to the problem, Section IV describes the integer programming model, Section V tests the model on real weather data for Atlanta Hartsfield-Jackson Airport (ATL), and Section VI gives conclusions and future directions.

II. BACKGROUND

Airspace sectorization and dynamic airspace configuration has been a growing area of study, as researchers have sought to find methods to partition and repartition airspace in a way that allows for the safe and efficient management of aircraft flow by air traffic controllers [5].

Past research has focused on enroute airspace, and has typically modeled the problem as one of partitioning a geometric space subject to convexity, connectivity, and minimum-time-in-sector constraints. The objectives used have served as proxies of overall controller complexity and workload, and involve balancing sector workload and minimizing inter-sector crossings.

This work was supported by NASA under the NGATS-ATM Airspace Program (NNA06CN24A) and by the NSF (ECCS-0745237)

Diana Michalek is a Ph.D. candidate at the Operations Research Center, at Massachusetts Institute of Technology, Cambridge, MA, USA dianam@mit.edu

Hamsa Balakrishnan is an assistant professor in the department of Aeronautics and Astronautics at the Massachusetts Institute of Technology. hamsa@mit.edu

Researchers have used many different solution techniques to solve the resulting NP-hard problem, including genetic algorithms that partition airspace using Voronoi tessellations which are found by successively moving 2D coordinates [6], and mathematical programming formulations that partition 2D airspace into hexagons and then assign the hexagons to a set of sectors [7]. In [8], the authors develop a method that recursively partitions a geometric space to build sectors, and mention that the pie-cut has potential for sectorization in the terminal area. However, to the best of our knowledge, there has not been a focus on either the unique challenges and characteristics of resectorizing the terminal area, or on the effect of weather on resectorization.

III. MODELING APPROACH

A. Sectorization of terminal airspace

In the context of the terminal area, the constraints and objectives of a good sectorization differ somewhat from that of enroute airspace studied in previous research.

First, aircraft trajectories can be approximated by line segments (without turns) from the terminal fix to the airport (or vice versa in the case of departures), with some maneuvering within 20 km of the airport. This simplifies the convexity and connectivity constraints of the sectorization problem, and allows airspace sectors to be constrained to pie slices.

Second, while flows are allowed to merge, aircraft crossings are rarely allowed in the terminal area, and arrivals and departures are kept in separate airspace to minimize complexity and maintain safety.

Finally, in any given time interval (say, 30 minutes), there is an inherent imbalance of arriving and departing traffic. This means that spreading of controller workload among all sectors is not an objective (though it may be desirable to spread the arrival demand across all *arrival* sectors). However, controller workload is still an important factor, and is incorporated in this research by limiting the deviations from existing airspace structure.

A more appropriate objective for the terminal area is that of meeting demand by (for example) expanding sectors when arrival demand is larger than departure demand, or by moving sectors or fixes during periods of weather activity. In the face of weather, a predicted storm cell may render an entire sector (or more) impenetrable by pilots.

As the terminal area is the bottleneck for airspace operations, the potential benefits gained from resectorization are especially relevant there, and are the focus of this paper.

B. Terminal Area Setup

This section introduces the terminal area model used in this paper. Consider the terminal depicted in Figure 1, which models the terminal airspace T using two concentric circles: an outer circle C_O of radius R , and an inner circle C_I of radius r . C_O represents the points at which arriving (departing) aircraft enter (exit) the terminal airspace, hence R represents the distance between the airport and the arrival (departure) fix of each sector. The inner circle C_I represents the points at which aircraft start their final approach into the airport.

Fig. 1. Model of terminal area, partitioned into a set of default arrival and departure sectors (dashed lines) and fixes (gray arrows).

Fig. 2. Arrival and departure sectors and fixes for ATL. The outer circle denotes the boundary of the terminal-area (100 km from the airport). © Google, Map Data US Navy. Image© 2009 DigitalGlobe. Image USDA Farm Service Agency.

The dashed gray lines represent the division of T into a set S of m sectors, where each $s \in S$ contains a fix, whose position and direction (either arrival or departure) are indicated by the placement of the gray arrow. The solid line in the southeast arrival sector indicates the route aircraft take from the fix to the airport. Note that these routes as well as all sector boundaries lie along radii of the circle C_O , and are of length $R - r$.

This abstract model of terminal airspace is motivated and modeled on the ATL terminal, which is the central case study of this paper. Figure 2 contains a diagram of this airspace, showing the (typical) four corner-post structure with 4 arrival sectors alternating with 4 departure sectors. Each sector contains a fix (depicted by a green triangle) approximately $R = 100$ km away from the airport, and typically all merges and landing patterns occur within $r = 20$ km of the airport.

With this model, we can now ask the following questions: Given a weather forecast, how can we restructure the terminal airspace to minimize disruptions to scheduled airspace usage? Can we make small changes to airspace structure (for instance, by moving a sector boundary and/or fix) that avoid or mitigate the effects of blocked airspace?

In order to answer these questions, we begin with a description of the weather forecast model used as input to this problem.

C. Weather model

As input to this problem we use a model of route robustness developed by the authors in previous work [4]. For a route r through terminal airspace, the model predicts a probability p_r that the route will be open (with some wiggle room) when the actual weather materializes. This data-driven

Fig. 3. Illustration of relocating fixes and associated routes to increase probability of remaining open.

model is built, using techniques from machine learning, on top of the state of the art Lincoln Lab CIWS forecast product. CIWS provides a frequently-updated deterministic forecast of convective weather for the NAS, at a resolution of 1 km x 1km and a time horizon of 0-2 hours [3].

This forecast p_r has been shown to be a good predictor of level 3+ weather along r , which is typically avoided by pilots in enroute airspace. We note that pilot deviation in the terminal area is not as well understood by researchers as in enroute airspace, and level 3+ weather may just be one predictor for deviation, among other factors such as demand.

Figure 3 shows an example of how this model can be used to optimize fix placement in a given sector, where the default fix (SID fix) may be replaced by a fix that is more likely to be open (chosen from a large set of potential fixes).

IV. FORMULATION

This section describes an integer programming formulation for the problem of assigning airspace sector boundaries in order to pick terminal fixes likely to be open during convective weather activity, while limiting deviation from existing airspace layout.

The integer program (IP) essentially partitions the terminal airspace into pie slices and assigns sector boundaries and fixes to these slices, subject to constraints on the maximum displacement from the original boundaries and the minimum distance between adjacent fixes. The IP strives to select fixes with high probability of being open (at the least, any fix selected must be forecast to be open) while at the same time limiting deviation from the original airspace layout.

A. Input

We start with terminal airspace T partitioned into n pie slices, each corresponding to a potential fix and associated undirected route (a route is a line segment between C_I and C_O , along a radius of C_O). The set of these potential fixes (and routes) is $F := \{1, \dots, n\}$. n should be chosen to be large enough to provide many options (say, 360 for a terminal area), but small enough so that each route is at least 1 km wide (the granularity of the weather forecast). We are also given a set $S := \{1, \dots, m\}$ of sectors.

Because the terminal area is circular, fix 1 is adjacent to fix n geometrically (distance is 1). To deal with this “wrap around” effect in the formulation, we introduce an augmented set of n^+ fixes, $F^+ := \{1, \dots, n^+\} \supset F$, referred to as wedges, and an augmented set of sectors $S^+ := \{1, \dots, m+1\} \supset S$.

The final input is a weather forecast for T , for a specific date and time horizon.

B. Parameters

The IP uses the following parameters:

For each $s \in S$ and $i \in F^+$, the weather forecast (the probability that i will be open in s) is p_{si} . Note that arrival and departure fixes may have different values of p .

For each $s \in S$ and $i \in F^+$, $d_{s,i}^{\text{fix}}$ is the distance of wedge i to the original fix for sector s , while $d_{s,i}^{\text{sect}}$ is the distance of wedge i to the original boundary for sector s . The boundary of sector s always refers to the clockwise-first boundary (or, the boundary with minimum wedge number, mod n), so that sector s extends from its boundary to the boundary of sector $s+1$. All distances in this formulation are in number of wedges.

K is the maximum displacement of a sector boundary, L is the minimum distance between any two fixes, and M is a large constant. Finally, the parameters α , β , γ , and λ are used to control the weight given to the various objectives (discussed in Section IV-D).

C. Variables

The binary variables are

$$x_{si} = 1 \text{ iff sector } s \text{ is assigned fix } i$$

$$y_{si} = 1 \text{ iff the boundary of sector } s \text{ is at wedge } i$$

$$z_s = 1 \text{ iff sector } s \text{ is open}$$

where $s \in S^+$ and $i \in F^+$.

D. Objectives

Four main objectives are desirable in the sectorization of terminal airspace:

- 1) Maximizing the probability that selected terminal fixes are open
- 2) Limiting the distance between the new sector boundaries and their default locations
- 3) Limiting the distance between the new fixes and their default locations
- 4) Keeping each sector open if feasible

These objectives can conflict with each other, since a high-probability fix in sector s may only be possible (depending on the weather forecast) if the boundary of sector s is moved clockwise, rendering sector $s+1$ blocked (say, if all remaining potential fixes are blocked: $p_{(s+1)i} \leq 0.5 \forall i \in (s+1)$). Thus, we create a linear combination of these objectives, and explore trade-offs in Section V. The overall objective function is therefore to minimize:

$$\sum_{s \in S} \left[\sum_{i \in F^+} \alpha x_{si} p_{si} - \beta y_{si} d_{s,i}^{\text{sect}} - \lambda x_{si} d_{s,i}^{\text{fix}} \right] - \gamma (1 - z_s) \quad (1)$$

In (1), the parameter α is the weight given to the first objective, namely, maximizing the probability that the fixes are open, β and λ are the penalties for deviation from the default sector and fix locations respectively, and γ is the penalty for closing down a sector, along with its associated arrival or departure fix.

E. Constraints

The first two constraints ensure that each sector has exactly one boundary, and that a fix is selected for each open sector, respectively.

$$\sum_{i \in F^+} y_{si} = 1 \quad \forall s \in S \quad (2)$$

$$\sum_{i \in F^+} x_{si} = z_s \quad \forall s \in S \quad (3)$$

The next set of constraints ensure that an optimal fix f^* for sector s is feasible for s : f^* is contained within s , and f^* is forecast open ($p_{sf^*} \geq 0.5$).

$$\sum_{i \in F^+} i y_{si} \leq \sum_{i \in F^+} i x_{si} + M(1 - z_s) \quad \forall s \in S \quad (4)$$

$$\sum_{i \in F^+} i x_{si} \leq \sum_{i \in F^+} i y_{(s+1)i} \quad \forall s \in S \quad (5)$$

$$0.5z_s \leq \sum_{i \in F^+} x_{si} p_{si} \quad \forall s \in S \quad (6)$$

The next two constraints ensure that fixes are at least L wedges apart, and that sector boundaries are moved no more than K wedges from their default locations, respectively.

$$L \leq \sum_{i \in F^+} i x_{(s+1)i} - \sum_{i \in F^+} i x_{si} + M(1 - z_{s+1}) \quad \forall s \in S \quad (7)$$

$$\sum_{i \in F^+} d_{si}^{\text{sect}} y_{si} \leq K \quad \forall s \in S \quad (8)$$

Finally, we have constraints to take care of the wrap around effect due to the circular airspace structure by essentially setting sectors 1 and $m+1$ to be equal. We assume without loss of generality that sector 1 is always defined to have its boundary at wedge $K+1$, so that its boundary will be kept between 1 and $2K+1$ by the IP.

$$x_{(m+1)i} = 0 \quad \forall i \in \{1, \dots, n\} \subseteq F^+ \quad (9)$$

$$x_{(m+1)i} = x_{1(i-n)} \quad \forall i \in \{(n+1), \dots, n^+\} \subseteq F^+ \quad (10)$$

$$y_{(m+1)i} = 0 \quad \forall i \in \{1, \dots, n\} \subseteq F^+ \quad (11)$$

$$y_{(m+1)i} = y_{1(i-n)} \quad \forall i \in \{(n+1), \dots, n^+\} \subseteq F^+ \quad (12)$$

$$z_{m+1} = z_1 \quad (13)$$

Equations (9) and (10) assign sector $(s+1)$ the same fix as sector 1 (mod n). Equations (11) and (12) assign sector $(s+1)$ the same sector boundary as sector 1 (mod n). Finally, (13) ensures that sector $(s+1)$ is open iff sector 1 is open.

V. RESULTS

Although the model presented is an integer program, which is NP-hard to solve in general, the problem size can be kept small in practice, thus eliminating computational issues. In particular, $n = 360$ and $m \leq 8$ is a realistic terminal airspace problem size, which we were able to solve in under one second using CPLEX. This section contains computational results.

The optimization algorithm was run on 28 weather scenarios over ATL, taken from the 8 most weather-impacted days during the months of June through August 2008 (see Section III-C). The stochastic route robustness model was

calibrated for a terminal area with $r = 20$ km, and $R = 100$ km, on an independent set of 2007 weather data.

Due to the nature of the multi-objective optimization, we begin with a detailed description of results for one setting of the objective function parameters. Afterwards, we look at how the weightings of various objectives affect the sectorization results.

A. Results for fixed parameter settings

This section describes results of the model when $(\alpha, \beta, \lambda, \gamma) = (100, 1, 1, 1)$. This parameter setting emphasizes the selection of a robust fix (one with high probability of being open), with small penalties for the displacement of fixes and sectors from their default positions, and a small penalty for a blocked sector.

Figure 4 shows diagrams of the sectorization found by the algorithm for two weather scenarios. The top shows results for August 26 2008 1100 Zulu, at the 60-minute time horizon. The algorithm makes a small sector boundary adjustment in the northwest sector, illustrating its potential to open up fixes that would otherwise be blocked. We see that the sector containing fix ROME is predicted to be blocked. The algorithm recommends a move of the sector boundary which results in a fix predicted to be open with probability 0.875. Fix CADIT is then moved to the far side of its sector so as to maintain separation. In the observed weather (on the right), this turns out to have been a good decision, as both the new ROME and CADIT fixes are open, while the original ROME fix is blocked by a large weather cell.

In the bottom of Figure 4, we see a scenario where the algorithm does not move any sectors, but does move fix ROME further away from weather activity, giving it a higher probability of being open. Moreover, two sectors in the southwest are (correctly) declared blocked.

Table I shows the overall performance of the algorithm for varying time horizons, and for arrival and departure sectors. Each row corresponds to one time horizon and direction (arrival/departure) combination, and represents 224 data points (28 weather scenarios, each with 8 sectors per flight direction). The computed metrics reflect the effectiveness and trade-offs of the optimization model.

The first metric reported, fix movements, refers to the percentage of fixes moved, and gives a measure of how often the algorithm recommends an alternate fix. This number tends to be larger for departures than for arrivals, but shows little variance within each direction. The second metric, sector movements, reports the percentage of instances that a sector boundary was moved. This value is small (under 7%) for this setting of parameters, reflecting the fact that sector movements only occur when there is a large gain in fix probability.

The next metric, forecast fix blockages, refers to the percentage of (original) fixes which were predicted as blocked. This number increases with increasing time horizon. Of course, a predicted fix blockage does not necessarily mean the fix will actually be blocked, and this situation is captured in the next metric, the percentage of actual blocked fixes

Fig. 4. Sectorization results for two illustrative weather scenarios, one with a sector boundary shift (top), and one without (bottom). The magenta routes correspond to those passing through the original fixes, the blue ones are those passing through the new fixes, and the purple routes represent their overlap.

TABLE I
OVERALL RESULTS FOR FIX OPTIMIZATION AND SECTOR CHANGES.

	Horizon (min)	Fix movements (%)	Sector movements (%)	Fix forecast blocked (%)	Fix blocked given forecast blocked (%)	Potential avoidable blockage (%)	Avoided blockage (%)
Arr	10	26	1	19	48	62	48
	30	28	4	18	55	45	32
	60	25	5	21	29	54	46
	90	25	3	24	30	54	39
Dep	10	24	7	21	42	74	70
	30	33	7	26	31	65	57
	60	37	4	32	22	71	61
	90	34	4	31	23	55	47

given that the fix is forecast to be blocked. Here we find that the longer time horizons are accompanied by lower values, reflecting the fact of lower forecast accuracy at longer time horizons.

Potential avoidable blockage shows the percentage of predicted-blocked fixes for which the algorithm recommends an optimal fix (which is predicted to be open). We find that at shorter time horizons, the potential to avoid blockages is predicted to be greatest. The percentage of avoidable blockages is above 50% in all cases, except for an outlier at arrivals with 30-min time horizon, meaning that the algorithm gives an

alternate routing possibility more than half the time. Finally, avoided blockages refers to the percentage of predicted-blocked fixes for which the optimal fix recommended by the algorithm is open in actual weather. This statistic tends to decrease with time horizon, and once again tends to be higher for departures than arrivals. The gap between the last two columns gives a measure of accuracy on predicted-blocked routes, though it does not distinguish between fixes assigned a 90% probability of being open and those with 60% probability. The accuracy should clearly depend on these probabilities, and this correlation is explored later.

TABLE II
ANALYSIS OF FIX MOVEMENTS AND SECTOR CHANGES.

	Horizon (min)	Number of Movements	Orig. & Opt. Open (OO)	Orig. Blocked, Opt. Open (BO)	Orig. Open, Opt. Blocked (OB)	Orig. & Opt. Blocked (BB)	OO + BO (%)
Arr	10	29	18	6	0	5	83
	30	31	23	4	0	4	87
	60	28	22	1	2	3	82
	90	28	22	1	2	3	82
Dep	10	27	17	9	1	0	96
	30	37	29	5	2	1	92
	60	41	34	2	2	3	88
	90	38	34	1	1	2	92

TABLE III

VALIDATION RESULTS FOR FIX OPTIMIZATION. EACH COLUMN CORRESPONDS TO A RANGE OF VALUES FOR THE PREDICTED PROBABILITY OF A FIX BEING OPEN, WHILE THE ENTRY IS THE EMPIRICAL PROBABILITY. THE STANDARD ERROR IS ALSO REPORTED IN PARENTHESES.

	Horizon (min)	% open	$p \in (0.95, 1.00]$	% open	$p \in (0.75, 0.95]$	% open	$p \in (0.50, 0.75]$
Arr	10	99.03	(0.01)	-	-	-	-
	30	95.96	(0.02)	-	-	-	-
	60	91.67	(0.03)	-	-	-	-
	90	100.00	(0.00)	89.47	(0.04)	84.62	(0.10)
Dep	10	99.02	(0.01)	-	-	-	-
	30	95.83	(0.02)	90.00	(0.09)	-	-
	60	95.29	(0.02)	93.75	(0.06)	-	-
	90	95.24	(0.03)	92.59	(0.05)	-	-

Table II provides a closer look at the fixes that are moved to some optimal fix by the algorithm. When a fix is moved, there are four possible outcomes: the original fix and the optimal fix are both open in the observed weather (OO), both are blocked (BB), the original fix is blocked while the optimal is open (BO), or the original fix is open while the optimal is blocked (OB). Ideally, we would want that the cases where the algorithm makes a mistake in moving a fix, BB and OB, be few in number, while BO (especially) and OO be many.

The table indicates several trends. First, OO accounts for more than 62% of fix movements across all categories, while OO and BO together account for more than 82% of fix movements, indicating that the optimal fix is usually likely be at least as good as the original. A movement of a fix that turns out to be open may seem undesirable, but the confidence in the optimal fix and associated route is greater than the original fix, making it the more conservative and robust choice. There are very few data points in the other three categories, indicating possibly large sampling error, so we only perform modest analysis of these cases. Nevertheless, at the 10-min time horizon for both arrivals and departures, and at the 30-min horizon for departures, it is a good decision to move the fix. This is consistent with the findings of prior research, which validated the short-term accuracy of 1 km $\times$ 1 km, pixel-based, forecasts [9]. Thus, tactical decisions to move fixes can be relied on, although more care and validation must be employed at longer and more strategic time horizons.

Table III shows algorithm performance as a function of prediction probability. The empirical percentage of open optimal fixes is listed, based on the probability with which they were predicated to be open by the classifier. Blank

entries correspond to cells with fewer than 10 data points, which were removed to eliminate cells with standard error greater than 0.10.

The uneven spread of data points among the three probability levels is an artifact of the the behavior of the underlying weather model, which is less likely to assign high-probability predictions as the time horizon increases. The table shows that the percentage of open routes tends to stay within the predicted percentage when there are enough data points. The table also shows that the validation is less accurate with increased time horizon and with decreased probability interval, as expected based on the behavior of the weather forecast model. Thus, the predicted probabilities correlate well with actual rates of route availability, and can be used to inform fix movement decisions in marginal weather conditions.

B. Weighting the objective

This section discusses results of the optimization as the parameters in the objective function vary.

Figure 5 shows how four key result statistics are effected as each of the four objective function parameters is varied while the others are fixed. When fixed, parameters are set to values studied in Section V-A: $\alpha = 100$, and $\beta = \lambda = \gamma = 1$. The results focus on departures at a 60-min horizon, and all other model parameters are unchanged from Section V-A.

The figure shows that parameters α (preferring increases in fix probability) and λ (penalty on distance of fix movement) have the largest effect on the percentage of fix movements (top left), the percentage of sector boundary movements (bottom left), and percentage of potential avoided blockages (bottom right). These parameters directly oppose each other, and a clear trade-off in results is evident in the plots.

Fig. 5. Sectorization results as a function of objective function parameters

In contrast, as γ increases, there is a very modest effect of under 3% on these first three result statistics. This can be explained by the relatively small number of total sector boundary movements available – when a fix is moved, there is often a high-probability fix within the original sector boundaries, making a boundary movement unnecessary. β , the penalty on the distance of a sector boundary movement, also shows modest effects on these three statistics, for similar reasons.

The top-right plot shows average increase in probability for a fix movement, which is surprisingly invariable to the four parameter values. Low values of α show the largest effect, but in general this statistic stays around 0.45. This indicates that when a fix is moved, on average there is a high gain in robustness. We note that the lack of variability in this plot could be due to the choice of the fixed parameters.

Overall, we see that the parameters have expected results in key result statistics, and decision-makers can tune these to reflect their preferences.

VI. CONCLUSION

The focus of this paper has been the development and evaluation of an integer programming approach to terminal airspace sectorization in the presence of convective weather, subject to constraints on the deviation of sector boundaries and fixes from today's fixed airspace structure.

In future work, we plan to relax these constraints and investigate the benefits of sectorization "from scratch", similar in spirit to research performed for enroute airspace, where the current sectors are typically not considered as constraints. Furthermore, we plan to incorporate arrival and departure demand into the formulation, and consider pilot

deviation behavior and air traffic controller workload as demand changes.

VII. ACKNOWLEDGMENTS

The authors gratefully acknowledge Marilyn Wolfson, Rich De Laura, Mike Matthews, and Mike Robinson at MIT Lincoln Lab for help with CIWS data and many fruitful discussions.

REFERENCES

- [1] Bureau of Transportation Statistics, "Understanding the reporting of causes of flight delays and cancellations," <http://www.bts.gov/help/aviation/html/understanding.html>, 2010.
- [2] Joint Planning and Development Office, *Joint Planning and Development Office ATM-Weather Integration Plan, Version 0.8*, July 2009.
- [3] M. Wolfson, B. Forman, K. Calden, W. Dupree, R. Johnson, R. Boldi, C. Wilson, P. Bieringer, E. Mann, and J. Morgan, "Tactical 0-2 hour convective weather forecasts for FAA," in *11th Conference on Aviation, Range and Aerospace Meteorology*, Hyannis, MA, 2004.
- [4] D. Michalek and H. Balakrishnan, "Identification of Robust Routes using Convective Weather Forecasts," in *USA/Europe Air Traffic Management R&D Seminar*, Napa, CA, June 2009.
- [5] K. Leiden, J. Kamienski, and P. Kopardekar, "Initial implications of automation on dynamic airspace configuration," in *7th AIAA Aviation Technology, Integration and Operations Conference (ATIO)*, 2007.
- [6] D. Delahaye and S. Puechmorel, "3D airspace sectoring by evolutionary computation: real-world applications," in *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, 2006.
- [7] A. Yousefi, "Optimum airspace design with air traffic controller workload-based partitioning," Ph.D. dissertation, George Mason University, 2005.
- [8] A. Basu, J. S. B. Mitchell, and G. Sabhnani, "Geometric algorithms for optimal airspace design and air traffic controller workload balancing," in *ALENEX*, 2008.
- [9] D. Michalek and H. Balakrishnan, "Building a stochastic terminal airspace capacity forecast from convective weather forecasts," in *Aviation, Range and Aerospace Meteorology Special Symposium on Weather-Air Traffic Management Integration*, Phoenix, AZ, January 2009.