Background

A primary goal of air-traffic management (ATM) is to balance airspace capacity and demand. Traditionally ATM has approached this problem assuming constant airspace configuration. A new ATM paradigm known as dynamic airspace configuration (DAC) analyzes how to reconfigure airspace dynamically in response to demand fluctuations, adverse weather, and the limitations of human cognition. DAC changes the airspace configuration to shift capacity where it is most needed.

DAC does not supersede risk-related ATM decision making. Rather, when there is a risk of adverse weather or route overloading, airspace-*independent* traffic flow management (TFM) decisions precede and constrain DAC decisions to manage risks appropriately. Regardless, airspace-*dependent* TFM decisions follow DAC decisions, to support equitable, optimal distribution of airspace capacity across air carriers, given the current airspace configuration.

Approach

Decision-support architecture. This case study describes development of a decision support system (DSS) that recommends DAC and TFM decisions in the presence of uncertain weather and demand patterns. The DSS' architecture differs from other DAC models by acknowledging that decision makers (such as air traffic controllers) have a limited ability to evaluate alternatives in real time. The DSS therefore pre-computes a relatively small set of DAC decision **templates**. During ATM operations the decisions makers choose the template that best matches weather and demand forecasts. As time progresses the decision makers can make modest, incremental changes to a template, in response to actual weather and demand conditions. The DSS thus incorporates a three-step decision process:

> **Prepare** optimal DAC templates.
> **Strategically select** a template for the upcoming time period.
> **Tactically adjust** template parameters as conditions warrant.

The preparation model. The DSS incorporates several subsystems, each requiring its own data-science tools and techniques. For brevity we focus on the development of the subsystem that defines optimal airspace-configuration templates. This **sectorization** subsystem partitions weather and demand forecasts into equivalence classes. Each equivalence class has its own airspace configuration (routing and re-routing) template.

The project team employed several interesting data-science concepts and methods to cluster historical weather and demand scenarios. These included ensembles, sampling, stochastic optimization, and k-medoids clustering using the **earth-mover distance (EMD) metric**. The EMD metric defines a distance between two discrete probability distributions in terms of the

amount of mass that must be shifted to transform one distribution into the other. Computing the metric requires solving a linear-programming problem, but the metric's technical advantages make this worthwhile. Figure 1 illustrates the clustering process for EMD-based clustering of traffic patterns.
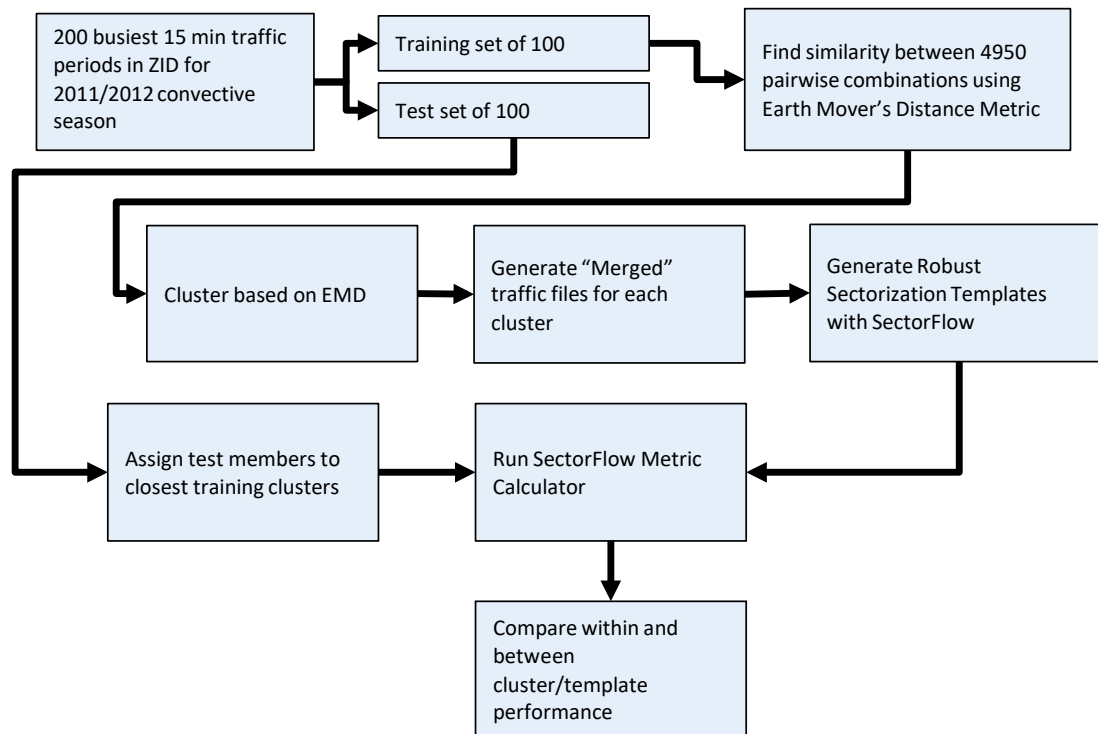
```
┌────────────────────┐     ┌──────────────────┐     ┌────────────────────┐
│ 200 busiest 15 min │     │ Training set of  │     │ Find similarity    │
│ traffic periods in │ ──▶ │ 100              │ ──▶ │ between 4950        │
│ ZID for 2011/2012  │     ├──────────────────┤     │ pairwise           │
│ convective season  │     │ Test set of 100  │     │ combinations using │
│                    │     │                  │     │ Earth Mover's      │
└────────────────────┘     └──────────────────┘     │ Distance Metric    │
                                                     └────────────────────┘

┌────────────────────┐     ┌──────────────────┐     ┌────────────────────┐
│ Cluster based on   │ ──▶ │ Generate "Merged"│ ──▶ │ Generate Robust    │
│ EMD                │     │ traffic files for│     │ Sectorization      │
│                    │     │ each cluster     │     │ Templates with     │
│                    │     │                  │     │ SectorFlow         │
└────────────────────┘     └──────────────────┘     └────────────────────┘

┌────────────────────┐     ┌──────────────────┐
│ Assign test members│ ──▶ │ Run SectorFlow   │
│ to closest training│     │ Metric Calculator│
│ clusters           │     │                  │
└────────────────────┘     └──────────────────┘

                           ┌──────────────────┐
                           │ Compare within   │
                           │ and between      │
                           │ cluster/template │
                           │ performance      │
                           └──────────────────┘
```

Figure 1: EMD Clustering of Traffic Patterns

The amount of EMD computation for 4,950 scenario comparisons would take 10 days of execution time on available hardware. The team used a much simpler center-of-mass comparison to avoid computing the EMD distance between obviously dissimilar traffic patterns.

The clustering analysis led to selecting 10 clusters. Figure 2 presents two of the clusters' mediods (central exemplars), by way of example:
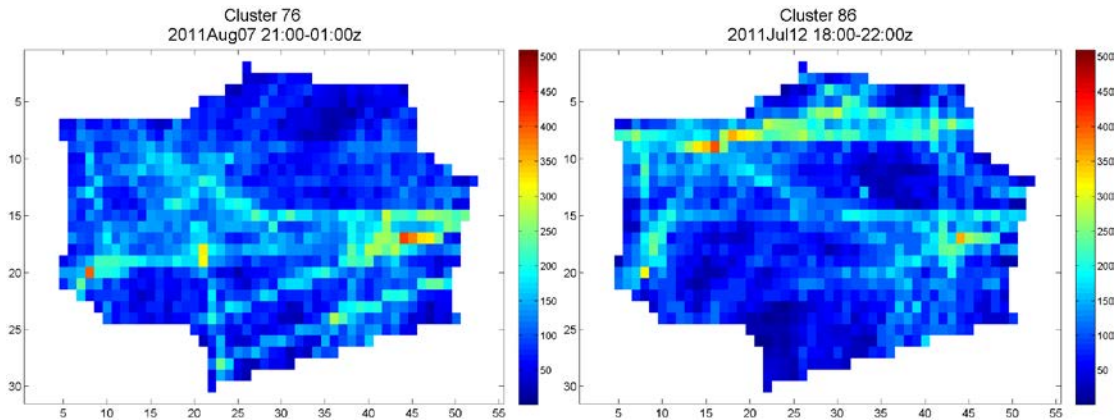
Figure 2: Sample Medoids

Next, the team used a variant of the SectorFlow algorithm to optimize each traffic cluster's sectorization. SectorFlow was developed under NASA contract from 2006 to 2010. Its objective function aggregates dynamic density and penalizes over-capacity conditions; see Figure 3:
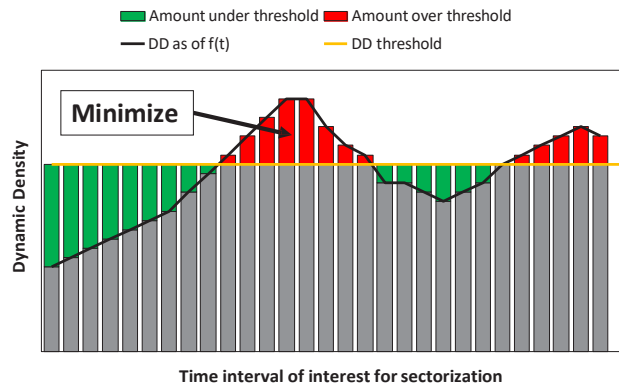

Figure 3: SectorFlow Objective Function

SectorFlow uses NASA's standard Welch capacity algorithm to compute sector capacity. Here are SectorFlow's steps:

1. Generate airspace demand profile based on user-preferred trajectories.
2. Cluster sampled aircraft positions.
3. Use computational-geometry techniques to form the initial boundary of the airspace that most efficiently encloses each cluster's aircraft positions.
4. Adjust sector boundaries heuristically to attempt to reduce the objective function.
5. If no reduction is achieved, exit; otherwise, repeat step 4.

Figure 4 depicts this process:
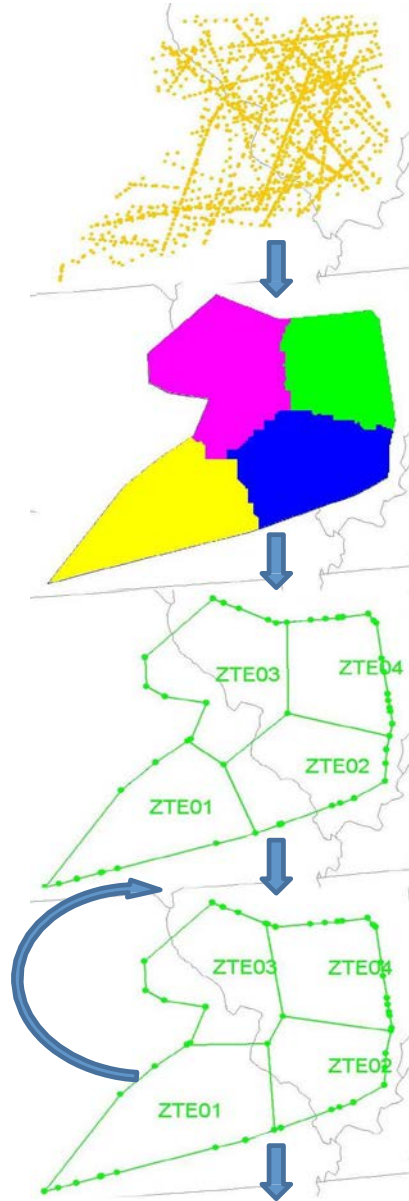
Figure 4: SectorFlow

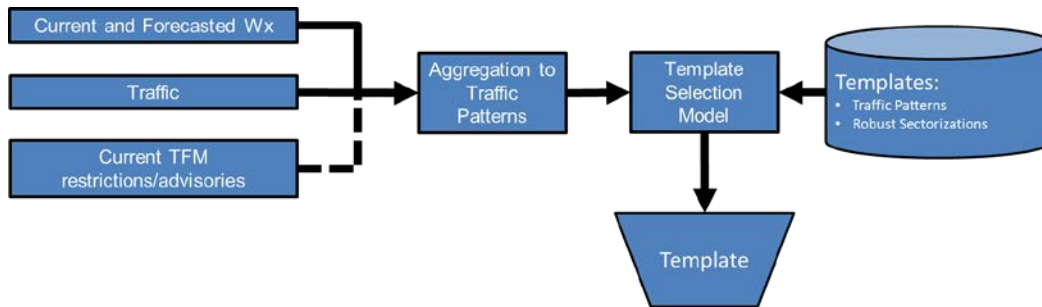**The strategic model.** Figure 5 depicts the DSS' template-selection workflow.

Figure 5: Template-Selection Workflow

One interesting challenge the team faced in implementing this workflow was standardizing convective weather forecast data from different kinds of forecasting products, including

- gridded probabilistic products,
- gridded deterministic products,
- area-based probabilistic products, and
- the National Convective Weather Detection (NCWD) product.

The DSS converts all of these into a quasi-NCWD representation that accounts for spatial and temporal variability, and that scales convective forecasts into maximum possible hourly NCWD scores for hexagonal cells. Figure 6 illustrates the result.
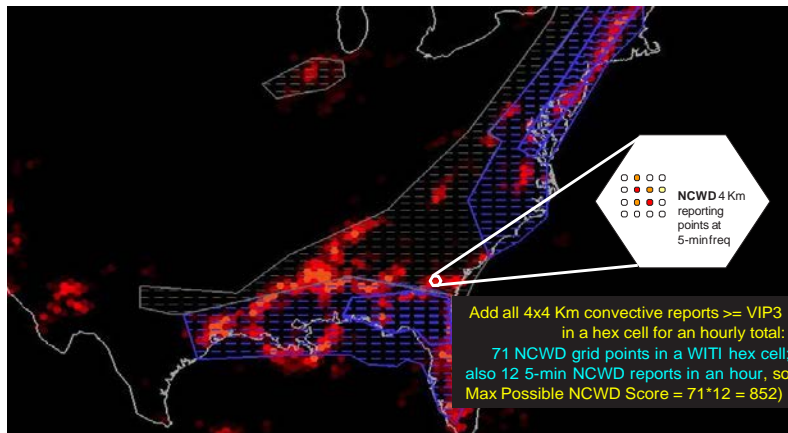
Figure 6: Hourly NCWD Score for a Hex Cell

The tactical model. The DSS limits tactical airspace-configuration adjustments to modest boundary changes made in pre-defined increments called **dynamic-airspace unit (DAU) slices**. The hypothesis was that operational decision makers can manage DUA adjustments well, while

they might struggle to respond correctly to dramatic operational changes to airspace configuration. We omit further details.

Results

To test the DSS, the project team relied on the superfast-time DART traffic and weather simulation model developed for NASA and the FAA by Air Traffic Analysis Inc. in 2010. DART is remarkably efficient; it can run a highly realistic simulation of a full day of U.S. air traffic and weather on a desktop computer in two minutes. Figure 7 is a sample DART display. (Yellow flights are on normal routes; blue flights have been re-routed; adverse weather cells are red.)
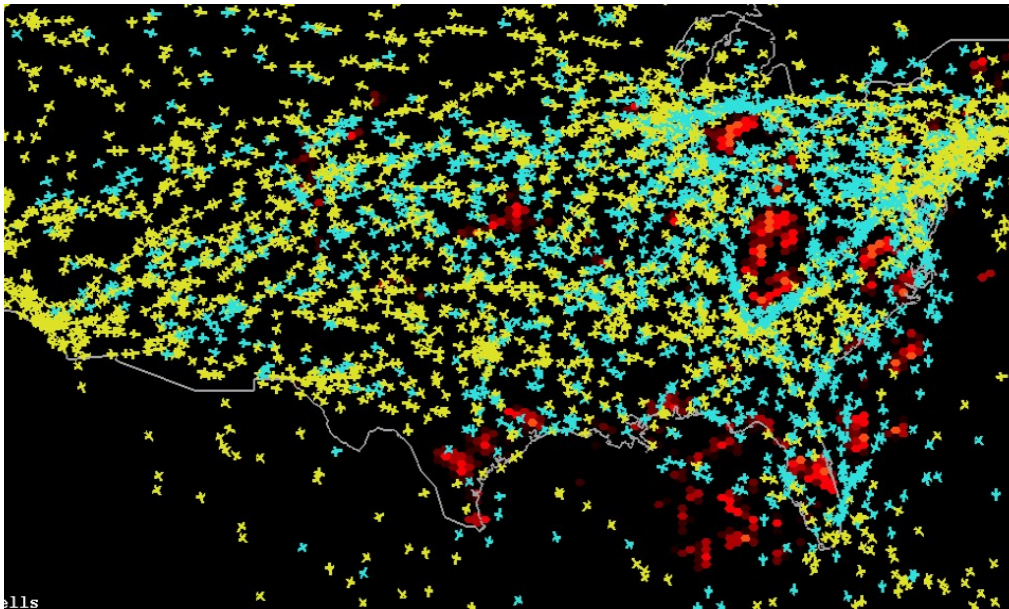
Figure 7: Sample DART Simulation of ~ 4,200 Flights in FAA Airspace

The project produced many simulation results; here are some highlights. First, the decision templates reduced delay compared to baseline by an average of 58%, and by as much as 78%. See Figure 8.
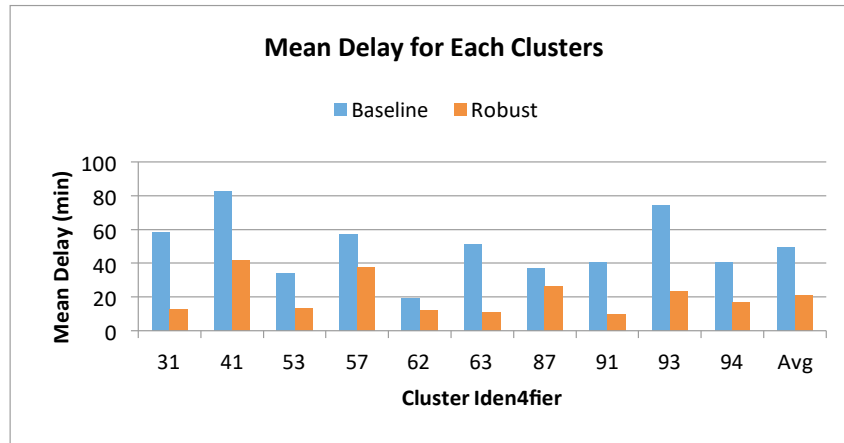
Figure 8: Mean Delay by Cluster

Depending on the computational technique used for a certain step in the sectorization process, traffic performance improved as often as 86% of the traffic days studied. See Figure 9.
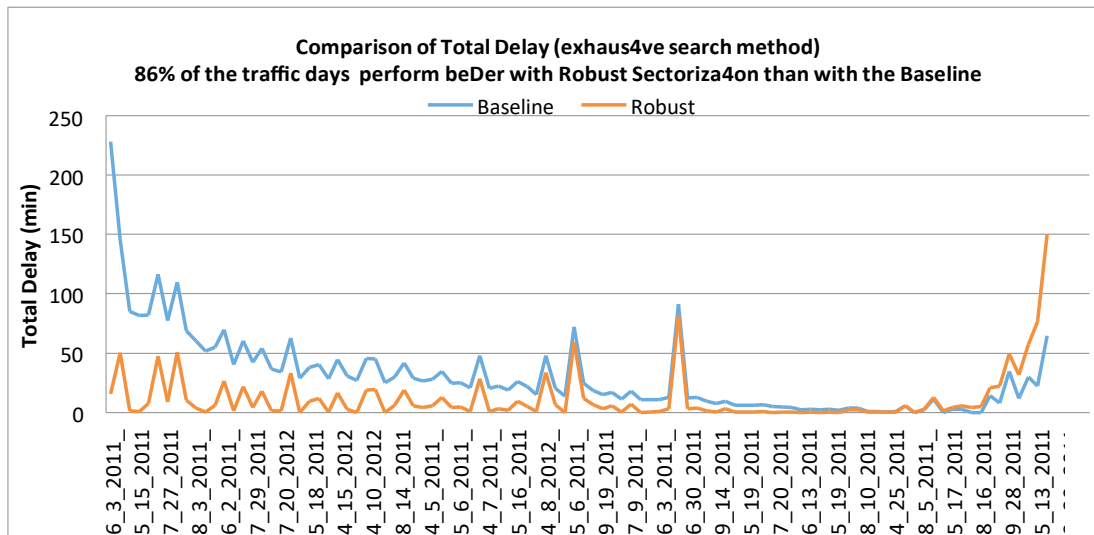

Figure 9: Total Delay by Traffic Day (Exhaustive Search)

Most important, the project computed overall annual operational cost reductions over the entire national airspace for Airline Service Quality Performance (ASQP) carriers, using several methods. Figure 10 depicts savings by cluster, for each of 10 buckets of dates having similar traffic patterns.
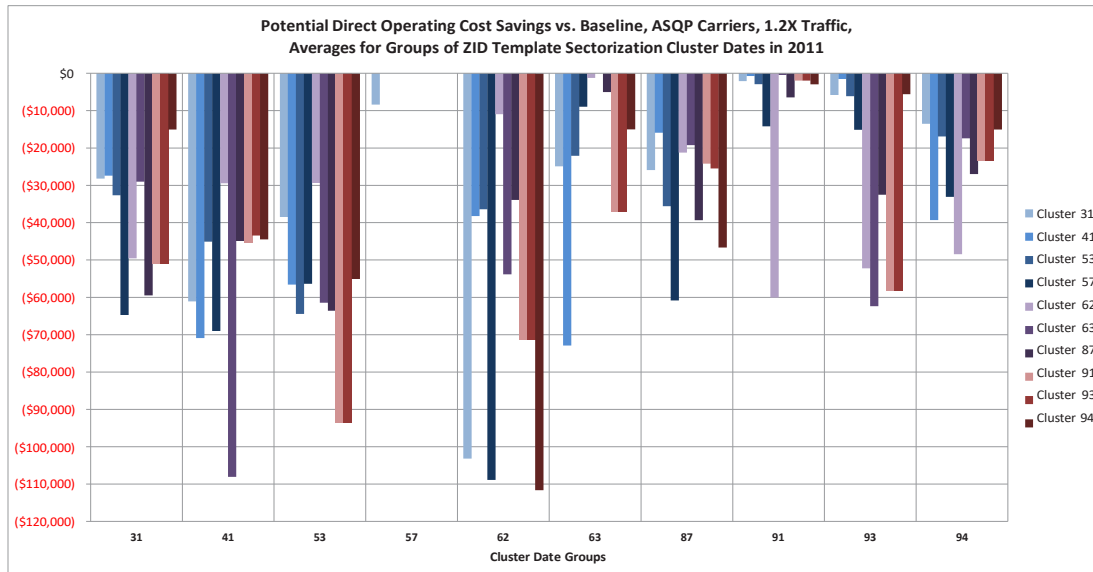
Figure 10: Daily Savings by Cluster and Date Bucket

This analysis estimated **annual nationwide benefits ranging between 51 and 268 million dollars**, depending on certain assumptions. These savings will increase if the analysis accounts for all national-airspace traffic, not just ASQP carriers.

For More Information

The foregoing glosses numerous technical details. If you'd like to learn more about this case, please contact Mosaic and request the following document:

> Mara Asencio et/al, "A Template Based Approach to Dynamic Airspace Configuration in the Presence of Uncertainties: Allocation Response to Weather Events (Final Report)," NASA Contract NNA11AC18C, 2013.

Contact Mosaic today to learn how we can help your organization achieve equally dramatic returns on your data science investment!