

ANALYSIS OF DIFFERENT COST FUNCTIONS IN THE GEOSCT AIRSPACE PARTITIONING TOOL

Gregory L. Wong, NASA Ames Research Center, Moffett Field, California

Abstract

A new cost function representing air traffic controller workload is implemented in the Geosct airspace partitioning tool. Geosct currently uses a combination of aircraft count and dwell time to select optimal airspace partitions that balance controller workload. This is referred to as the aircraft count/dwell time hybrid cost function. The new cost function is based on Simplified Dynamic Density, a measure of different aspects of air traffic controller workload.

Three sectorizations are compared. These are the current sectorization, Geosct's sectorization based on the aircraft count/dwell time hybrid cost function, and Geosct's sectorization based on the Simplified Dynamic Density cost function. Each sectorization is evaluated for maximum and average workload along with workload balance using the Simplified Dynamic Density as the workload measure. In addition, the Airspace Concept Evaluation System, a nationwide air traffic simulator, is used to determine the capacity and delay incurred by each sectorization.

The sectorization resulting from the Simplified Dynamic Density cost function had a lower maximum workload measure than the other sectorizations, and the sectorization based on the combination of aircraft count and dwell time did a better job of balancing workload and balancing capacity. However, the current sectorization had the lowest average workload, highest sector capacity, and the least system delay.

Introduction

Airspace is divided into sectors to distribute air traffic controller workload. The goal is to balance the workload amongst all controllers, while not overwhelming any single controller. Overloading an air traffic controller is unsafe, and aircraft rerouting and delays are used to prevent overloading. Optimally partitioning the airspace, hereafter known as sectorization, requires a metric that quantifies controller workload. Previous research identified

such a quantification, known as Dynamic Density, that includes a wide variety of air traffic and airspace metrics [1-6]. This has been narrowed to a subset of key metrics known as the Simplified Dynamic Density [7, 8].

Geosct is a tool that partitions the airspace into sectors, while optimizing a cost function representing the magnitude of controller workload for each sector [9-12]. Geosct accomplishes this using computational geometric techniques and constraints to produce sector shapes that are acceptable to air traffic controllers. In its current implementation, Geosct uses a cost function that is a combination of the number of aircraft predicted to be in a given sector (aircraft count) and the average amount of time each aircraft is predicted to spend in each sector (dwell time). While aircraft count and dwell time are both components of the Simplified Dynamic Density metric, they do not account for all factors that contribute to air traffic controller workload. For example, having ten aircraft evenly spaced and flying in the same direction in a sector causes less workload than having five aircraft all converging on the same point from different directions.

The focus of this study is to evaluate the benefits of using Simplified Dynamic Density as a cost function in Geosct to partition airspace. Benefits are measured in terms of sector capacity and average system delay. Benefits are also measured using the Simplified Dynamic Density metric as a gauge of controller workload. This sort of benefit analysis is similar to that of Zelinski [13]. At the time the Zelinski study was conducted, Geosct had not reached the necessary level of maturity to be included in the study.

This paper gives a brief overview of the Geosct Airspace Partitioning Tool, followed by an explanation of the cost functions implemented in Geosct. The method by which the cost functions were tested is explained, followed by an analysis of the results. This is then followed by a discussion of future work. Finally, conclusions are drawn from the analysis.

Overview of Geosect

Geosect partitions or sectorizes airspace using computational geometric techniques. The particular version of Geosect examined in this study uses a top down approach with binary partitions. Given a region of airspace, Geosect partitions that airspace into two sectors. Each of these sectors is then evaluated and, if necessary, partitioned into two smaller sectors. Geosect continues partitioning until the desired number of sectors (an input parameter) has been created. Various geometric constraints ensure that the resulting sector shapes are acceptable to controllers in the presence of the anticipated traffic flow. Geosect's current implementation balances controller workload through the use of a cost function based on aircraft count and dwell times.

Geosect begins with a single airspace region (usually the boundaries of an Air Route Traffic Control Center or "Center"), predicted aircraft track hits within the region, polylines representing the dominant air traffic flows through the region (as determined by the user), and the location of major airports and special use airspaces (SUAs) in the region (see Figure 1). In order to partition the region into sectors that are geometrically acceptable to air traffic controllers, Geosect sets up exclusion areas and search nodes.




Figure 1. Geosect Input Components

Exclusion areas (blue circles in Figure 2) are created around major airports, SUAs, and dominant flow intersections. Geosect does not create partitions that encroach on these exclusion areas. This ensures that critical points such as airports, SUA corners, and dominant flow intersections are not too close to a sector boundary. Moreover, as partitions are tested and accepted by Geosect, exclusion zones are set up

around the intersections of the partitions. This prevents the creation of points where four or more sectors meet. Such a point is undesirable because it creates hand-off ambiguities for air traffic controllers.




Figure 2. Exclusion Areas in Cleveland Center

Geosect constructs a series of external and internal search nodes (black squares in Figure 3). The external search nodes are equally spaced on the perimeter of the region to be partitioned. The internal search nodes are placed midway between the dominant flows. The segments joining these search nodes make up the search space for segments of each candidate partition. This ensures two properties of the partition. First, the external search nodes ensure that the partition extends from a point on the perimeter to another point on the perimeter. Second, the internal search nodes constrain bends in the partition to occur as far away from the dominant flows as possible.

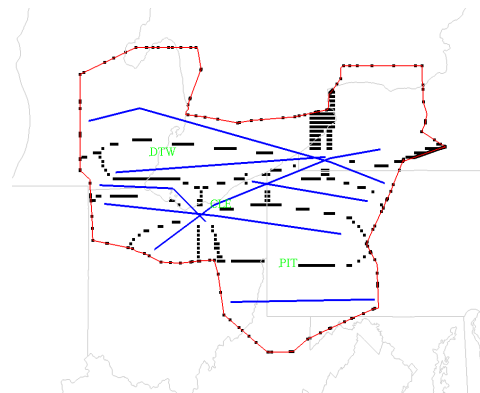


Figure 3. Initial Search Nodes in Cleveland Center

Geosect further constrains candidate partitions to those whose orientation is orthogonal to the long axis of the region. Thus, if the region to be partitioned is long and skinny, Geosect will not partition the region

into two skinnier sectors. In addition, Geosect excludes candidate partitions that cross dominant flows at small angles (i.e. nearly parallel). Allowing small angles can create hand-off ambiguities for air traffic controllers.

Given these geometric constraints, Geosect examines a series of partitions that divide the region into two smaller sectors. The cost function is applied to each of the sub-regions, and the highest cost of the two is associated with the partition. From the set of possible partitions for a given region, Geosect selects the partition with the lowest cost associated with it. The resulting smaller sectors are then placed on a priority queue, which is ordered by a priority function, for further partitioning. The process of partitioning the highest priority sector in the queue and then adding the resulting smaller sectors back on the queue is repeated until the user-defined number of sectors has been created.

Cost Functions

This section describes the cost functions implemented in Geosect for this study. The inputs to the cost functions are aircraft radar track hits. Each track hit includes information on the aircraft's location, altitude, ground speed and heading. The track hits are one minute apart.

Aircraft Count/Dwell Time Hybrid

The first cost function examined by this study is a hybrid of aircraft dwell time and average aircraft count. The aircraft dwell time is the total time all aircraft spend in a given sector. It is desirable to maximize dwell time to reduce the number of hand-offs required between sectors and thus the amount of controller coordination required. Therefore, the dwell time function is implemented as an inverse cost function where lower dwell times translate to higher cost. The aircraft dwell time function tends to create sectors that conform to the dominant flow of traffic. The average aircraft count function is the number of aircraft present in a sector averaged over the period the sectorization will take place. As described in the previous section, each candidate partition is evaluated based on the largest cost of the resulting two sectors. If the number of sectors created is less than half the desired number of sectors, then the aircraft dwell time function is used. After half the desired number

of sectors has been created, then the average aircraft count cost function is used to evaluate the candidate partitions. The goal of this hybrid cost function is to create large super-sectors that align with the dominant flows. These super-sectors are then divided into smaller sectors to reduce and balance each controller's workload.

The sectorization based on a hybrid of aircraft count and dwell time uses maximum aircraft count as its priority function. Maximum aircraft count is the greatest number of aircraft predicted to be within a sector at any given moment. The sector in the priority queue with the greatest maximum aircraft count is selected as the next sector to partition.

Occupancy Count Component

The Simplified Dynamic Density (SDD) metric is a weighted combination of seven components. The first SDD component, occupancy count, is the number of aircraft track hits within a sector averaged over a 15-minute period. More aircraft present in a sector at the same time implies higher controller workload. Occupancy count, $x1_{s,k}$, for sector s and 15-minute period k is given by

$$x1_{s,k} = \frac{n_{s,k}}{15}, \quad (1)$$

where $n_{s,k}$ is the number of aircraft track hits in s during k . $x1_{s,k}$ is a component of the SDD cost function. It, along with the other components, appears as a term in Eq. (4), shown later in this paper.

Composite Proximity Level Component

Proximity level is a quantification of how close the track hits of two aircraft are to each other in space and time. Aircraft that are predicted to be near each other in time and space could potentially come into conflict and increase the air traffic controller's workload. Even if they do not conflict, aircraft pairs with high proximity levels require monitoring by the air traffic controller.

The track hits of each pair of aircraft are compared and assigned a proximity level. Table 1 lists the proximity levels from more severe to less severe.

Table 1. Proximity Level Criteria

Prox. Level	Vertical Separation (ft.)	Horizontal Separation (nm)	Time Separation (seconds)
1	< 1000	< 5	< 10
2	< 1000	5 to 7.5	< 10
		< 5	< 20
		7.5 to 10	< 10
3	< 1000	5 to 7.5	< 20
		< 5	< 30
		< 5	< 10
4	any	< 5	< 10

For example, if two aircraft are separated by between 5 and 7.5 nautical miles horizontally and less than 1000 feet vertically, and their track hits' timestamps were less than 10 seconds apart, then the aircraft pair are assigned a proximity level of 2.

The composite proximity level, $x2_{s,k}$, for sector s and 15-minute period k is a composite of the identified Proximity Levels according to

$$x2_{s,k} = \frac{4p_{1,s,k} + 2p_{2,s,k} + p_{3,s,k} + p_{4,s,k}}{4}, \quad (2)$$

where $p_{1,x,k}$, $p_{2,x,k}$, $p_{3,x,k}$, and $p_{4,x,k}$ represent the number of level 1, 2, 3, and 4 proximities, respectively.

Altitude Transition Count Component

Altitude transition count, $x3_{s,k}$, is the number of aircraft track hits where the aircraft's absolute altitude change rate is greater than 500 feet per minute. Aircraft that are changing flight levels can increase an air traffic controller's workload as they merge or cross through streams of other aircraft.

Sector Transfer Count Component

Sector transfer count, $x4_{s,k}$, is the number of aircraft that enter and exit sector s during 15-minute period k . Note that an aircraft leaving sector r and entering sector s is counted in the sector transfer count for both s and r . Sector transfers require communication and coordination between sector controllers and pilots, and this adds to controller workload.

Sector Density Component

Sector density, $x5_{s,k}$, is the number of aircraft track hits in sector s during 15-minute period k (same as occupancy count $x1_{s,k}$) divided by the volume of the sector (in km^3). Higher densities imply higher workload because the air traffic controllers have less airspace available to resolve conflicts.

Heading Variance Component

Heading variance, $x6_{s,k}$, is a quantification of the variation in aircraft headings within a sector. The heading variance is given by

$$x6_{s,k} = \frac{1}{N_{s,k}} \sum_{i=1}^{N_{s,k}} (h_{i,s,k} - m_{s,k})^2, \quad (3)$$

where $m_{s,k}$ is the mean of all $N_{s,k}$ headings, $h_{i,s,k}$, in sector s during 15-minute period k . The motivation behind this metric is that it is easier to control streams of aircraft that are flying in the same direction than aircraft with a variety of headings and potentially crossing trajectories.

Speed Variance Component

Speed variance, $x7_{s,k}$, is based on the aircraft ground speeds in a manner similar to heading variance.

Simplified Dynamic Density

The seven components of the SDD are combined in a weighted sum, $\bar{I}_{s,k}$, for sector s and 15-minute period k according to

$$\begin{aligned} \chi_{s,k} = & 2.2(x1_{s,k}) + 1.2(x2_{s,k}) + 0.2(x3_{s,k}) + \\ & 0.4(x4_{s,k}) + 3000(x5_{s,k}) + 0.0005(x6_{s,k}) + \\ & 0.0005(x7_{s,k}) \end{aligned} \quad (4)$$

The SDD cost function, X_s , is computed as the mean of $\bar{I}_{s,k}$ over all of the 15-minute periods.

Method for Evaluating Cost Functions

Geosect generated sectors in two altitude strata for Cleveland Center. Cleveland Center was selected for this study because its en route airspace includes a variety of crossing, climbing, and descending air traffic. High altitude sectors covered the first strata from 24,000 to 34,900 feet. Super high sectors

covered the second strata from 35,000 to 60,000 feet. This is the same stratification used by most of Cleveland's present day sectorization. The locations of the airports in Detroit, Cleveland, and Pittsburgh were input to Geosect to provide partitioning exclusion zones. Aircraft track hit data were from a day in 2005 that was not impacted by weather and spanned approximately 32 hours.

The current sectorization includes 29 total sectors, covering the high and super high altitude strata. Since this averages to 15 sectors per stratum, Geosect was configured to generate 15 sectors for each altitude stratum using the aircraft count/dwell time cost function and the SDD cost function described above. Each Geosect generated sector was evaluated using SDD as the workload measure. Each sector of the current sectorization was similarly evaluated using SDD as the workload measure. The maximum, mean, and standard deviation over all of the sectors were computed for each sectorization.

In addition, for each sectorization, the high and super high sectors were input into a simulation using the Airspace Concept Evaluation System (ACES) [14]. First, ACES simulated current (2005) levels of air traffic unconstrained by sector capacities. This unconstrained traffic data was then used to compute the Monitor Alert Parameter (MAP) values for each sector. The Federal Aviation Administration (FAA) uses MAP values as an indication of sector capacity [15]. MAP values range from 5 (lowest capacity) to 18 (highest capacity). Next, ACES simulated the same traffic but imposed delays so that each sector's capacity, as given by its MAP value, was not exceeded. The resulting average delay over the entire system was recorded.

Results and Analysis

The SDD function is used as the workload measure in the analysis of each Geosect generated sectorization as well as the current day Cleveland sectorization. The maximum is a sectorization's worst-case workload, and the mean is a sectorization's average workload over the entire period under consideration. The standard deviation is interpreted as a sectorization's workload balance among its individual sectors. The smaller the standard deviation, the better the balance. In addition, the average MAP value and average system delay for

each sectorization are derived from ACES simulations and compared.

Workload

Figure 4 shows the maximum SDD measure for each of the sectorizations. (Note that the sectorization based on the aircraft count/dwell time hybrid cost function is abbreviated as AC/DTH in the following figures.) Using SDD as a measure of workload, both Geosect sectorizations had smaller worst-cases than the current sectorization. The sectorization generated by Geosect using SDD as its cost function had the least worst-case.




Figure 4. Maximum SDD

Figure 5 shows the mean SDD measure for each of the sectorizations. The current sectorization has a significantly smaller average workload than either of the Geosect sectorizations. Geosect's average workload was slightly lower using the SDD cost function versus the sectorization based on the aircraft count/dwell time hybrid.




Figure 5. Mean SDD

Figure 6 shows the SDD measure's standard deviation, an indication of workload balance. Both Geosect generated sectorizations had better workload balance (lower standard deviations) than the current sectorization. The workload balance was better when Geosect used the aircraft count/dwell time hybrid cost function than the SDD cost function.




Figure 6. SDD Standard Deviation

Capacity and Delay

Mean MAP values are shown in Figure 7. Interpreting MAP value as an indication of capacity, neither of the Geosect sectorizations had an average sector capacity as high as the current sectorization. However, as shown in Figure 8, the capacity is better balanced between the sectors created by Geosect using the aircraft count/dwell time hybrid than the current day sectors. Given that dwell time is closely related to the computation of MAP values, it makes sense that the aircraft count/dwell time hybrid sectorization scored better than the SDD sectorization.




Figure 7. Mean MAP Value




Figure 8. MAP Value Standard Deviation

ACES also computed the average system delay for each sectorization, and the results are shown in Figure 9. The current day sectorization has less average delay than either Geosect sectorizations. Among the Geosect generated sectorizations, the aircraft count/dwell time hybrid sectorization had less delay than the SDD sectorization. Again, this may be due to the fact that dwell time is closely related to the MAP value calculations used by ACES.




Figure 9. Average System Delay

Summary of Results

Table 2 summarizes which sectorizations did the best under the various evaluations used in this study. Using the aircraft count/dwell time cost function, Geosect does a good job of balancing workload and capacity. However, current day sectorization outperforms Geosect in the areas of average workload, sector capacity and system delay.

Table 2. Summary of the Best Sectorizations

Comparison	Best Sectorization
maximum workload	SDD
average workload	current
workload balance	aircraft count/ dwell time hybrid
Average Sector Capacity	current
Sector Capacity Balance	aircraft count/ Dwell time hybrid
Average System Delay	current

Design of the current sectorization is an art form requiring the efforts of an experienced airspace designer who designs the sectors in conjunction with the design of the air traffic routes. As a result, the current sectorization scored well in this study since the traffic used to evaluate the sectorizations followed these routes. Given a different set of traffic that do not follow the routes, either due to weather or some sort of wind-optimized routing, the evaluations may have yielded different results. The current sectorization has also benefited from years of evolution while the design of automated sectorization tools is fairly new.

The ACES simulations show that the current sectorization design strives to optimize airspace capacity while reducing system delay. While Geosect does a good job of balancing workload and capacity, further development is required to bring it up to the level of the current sectorization, and this is covered in the next section.

Future Work

Within each altitude stratum (high and super high), the current day sectorization included sectors that occupied only portions of the stratum. In those cases, sectors overlapped each other with different sectors covering different portions of the altitude stratum. The version of Geosect used in this study was limited by its two-dimensional nature. Each sector created by Geosect was expected to cover the full altitude stratum, and the stratum is specified by the user. A future study should look at using Geosect to generate three-dimensional sectors, where Geosect would pick its own altitude strata. The resulting sectors would overlap would not necessarily all have the same floors and ceilings.

At the time of this study, the dominant flows used by Geosect were created by hand. In future, Geosect will use automatically generated dominant flows. This will make dominant flows more consistent between different Centers, and will make Geosect easier to use with different altitude strata.

Finally, at the time of this study, the user specified how many sectors for Geosect to generate. A future enhancement would allow Geosect to create an undetermined number of sectors until some criterion was met. This would alleviate the problem

of small sectors being generated merely because Geosect had not reached its user-specified limit.

Conclusion

Two sets of sectorizations of Cleveland Center were generated by Geosect using different cost functions. These sectorizations and the current day sectorization were compared by applying the Simplified Dynamic Density measure (representing workload) to each sector. In addition, the Airspace Concept Evaluation System was used to simulate current air traffic. From these simulations, sector capacity and average system delay were computed for each sectorization.

Geosect generated sectorizations using the aircraft count/dwell time hybrid cost function gave a better workload balance and sector capacity balance than the current sectorization. The sectorization resulting from the Simplified Dynamic Density cost function had a lower maximum workload measure than the other sectorizations. However, Geosect's sectorizations incurred greater delay and did not generate as much sector capacity as the current sectorization.

References

- [1] Laudeman, I. V., S. G. Shelden, R. Branstrom, C. L. Brasil, April 1998, Dynamic Density: An Air Traffic Management Metric, NASA-TM-1998-112226, Moffett Field, California.
- [2] Sridhar, B., K. S. Sheth, S. Grabbe, 1-4 December 1998, Airspace Complexity and its Application in Air Traffic Management, 2nd USA/Europe Air Traffic Management R&D Seminar, Orlando, Florida.
- [3] Chatterji, G. B., B. Sridhar, 16-18 October 2001, Measures for Air Traffic Controller Workload Prediction, AIAA-2001-5242, Proceedings of the First AIAA Aircraft Technology, Integration, and Operations Forum, Los Angeles, California.
- [4] Kopardekar, P., S. Magyarits, June 2003, Measurement and Prediction of Dynamic Density, 5th USA/Europe Air Traffic Management R&D Seminar, Budapest, Hungary.
- [5] Masalonis, A., M. Callahan, C. Wanke, June 2003, Dynamic Density and Complexity Metrics for Realtime Traffic Flow Management, 5th USA/Europe Air Traffic Management R&D Seminar, Budapest, Hungary.
- [6] Kopardekar, P., A. Schwartz, S. Magyarits, J. Rhodes, 2-5 July 2007, Airspace Complexity Measurement: an Air Traffic Control Simulation Analysis, 7th USA/Europe Air Traffic Management R&D Seminar, Barcelona, Spain.
- [7] Klein, S., M. Rodgers, H. Kaing, P. Lucic, K. Leiden, 3 December 2008, DAG CE-6 Golden Nuggets Fast-Time Modeling and Simulation Studies, Final Report Part 3: Dynamic Airspace Configuration Analysis, Deliverables for Subtasks 13 & 14 under Contract GS-35F-0308K; Order No. A65616D, NASA Ames Research Center, Moffett Field, California.
- [8] Klein, A., M. Rodgers, K. Leiden, Simplified Dynamic Density: a Metric for Dynamic Airspace Configuration and NEXTGEN Analysis, 28th Digital Avionics Systems Conference, Orlando, Florida.
- [9] Mitchell, J. S. B., G. Sabhnani, J. Krozel, R. Hoffman, A. Yousefi, 18-21 August 2008, Dynamic Airspace Configuration Management based on Computational Geometry Techniques, AIAA 2008-7225, AIAA Guidance, Navigation, and Control Conference, Honolulu, Hawaii.
- [10] Basu, A., J. S. B. Mitchell, G. Sabhnani, 2008, Geometric Algorithms for Optimal Airspace Design and Air Traffic Controller Workload Balancing, Proc. Tenth Workshop on Algorithm Engineering and Experiments, ALENEX'08, SIAM Proceedings in Applied Mathematics, (Journal version submitted to ACM Journal of Experimental Algorithmics, 2008).
- [11] Basu, A., J. S. B. Mitchell, G. Sabhnani, 10-11 November 2006, Geometric Algorithms for Optimal Airspace Design and Air Traffic Controller Workload Balancing, Proc. 16th Fall Workshop on Computational and Combinatorial Geometry, Smith College, Northampton, Massachusetts.
- [12] Sabhnani, G. R., 2009, Geometric Algorithms for Dynamic Airspace Sectorization. PhD Dissertation, Stony Brook University, Stony Brook, New York.
- [13] Zelinski, S., 29 June – 2 July, 2009, A Comparison of Algorithm Generated Sectorizations, Eighth USA/Europe Air Traffic Management

Research and Development Seminar (ATM2009), Napa, California.

[14] Meyn, L., R. Windhorst, K. Roth, D. Van Drei, G. Kubat, V. Manikonda, S. Roney, G. Hunter, A. Huang, G. Couluris, 21-24 August 2006, Build 4 of the Airspace Concept Evaluation System, AIAA 2006-6110, AIAA Modeling and Simulation Technologies Conference, Keystone, Colorado.

[15] Federal Aviation Administration, Order JO 7210.3V, 2008, Section 8. Monitor Alert Parameter, [online publication]
URL: http://www.faa.gov/air_traffic/publications/atpubs/FAC/Ch17/s1708.html, effective date: February 14, 2008, includes change effective July 31, 2008 and change 2 effective March 12, 2009.

Acknowledgements

The author thanks Girishkumar Sabhnani at State University of New York (SUNY) at Stony Brook (now at Metron Aviation) for his assistance and expertise with Geosect, and Charlene Cayabyab for her ACES support at NASA Ames Research Center.

Email Addresses

Gregory L. Wong: Gregory.L.Wong@nasa.gov

*28th Digital Avionics Systems Conference
October 25-29, 2009*