Infotech@Aerospace 2012
19 - 21 June 2012, Garden Grove, California

AIAA 2012-2540

Visual Analysis of Air Traffic Data Using Aircraft Density and Conflict Probability

Georg H. Albrecht* and Hak Tae Lee†
University of California, Santa Cruz, Moffett Field, CA, 94035
Alex Pang‡
University of California, Santa Cruz, Santa Cruz, CA, 95064

This paper presents visual analysis tools intended to aid researchers and policy makers by producing a visual representation of air traffic density and, by taking into account the required aircraft separation minima, the conflict probability. The tools support visualization of time-varying air traffic density over an area of interest using different time granularity. This visual analysis platform is used to investigate how potential changes to the aircraft separation policy can reduce congestion while maintaining key safety requirements. The same platform can also be used as a decision aid when assessing potential flight plans for unmanned aerial vehicles.

I. Introduction

Aviation and air travel has established itself as a key economic and social resource in modern times. As the world population increases and becomes ever more interconnected, the demand for air travel will only increase. On average there are over 87,000 flights within the United States' national airspace system (NAS) every day. Of these flights, approximately 28,500 are commercial, 27,000 are private, 24,000 are air taxis, and 2000 are freight,[1] moving some 1.7 million passengers and 54 million pounds of freight.[2] The number of passengers flying within the U.S. is predicted to grow an average of 4.5% annually, and the number of freight aircraft is expected to grow some 2% anually, while general aviation is expected to grow 1% annually.[3] In addition, there is increasing interest, from both government and commercial sectors, in integrating Unmanned Aircraft Systems (UAS) into the NAS.[4,5] Though full UAS integration poses its own unique set of complications, it is only a matter of time before they contribute to the air traffic within the NAS. This expected overall increase in air traffic will require new standards and procedures to efficiently accommodate the increasing air traffic. To address this the US Congress approved plans for the development of the Next-Generation Air Transportation System (NextGen),[6] over the next 15 years. Among other things, the NextGen system is intended to provide the infrastructure to allow aircraft to safely fly closer together, thereby making more efficient use of limited airspace. With pieces of the NextGen infrastructure coming into place there is an opportunity to further their benefits by developing software tools that provide added value.

This paper will discuss methods and tools used to calculate and render air traffic density and, by taking into account aircraft separation minima, conflict probability over an area of interest. This paper also discuss methods for aggregating such traffic densities over different time scales to extract fluctuations and periodic cycles in traffic patterns. In general these tools can be used to aid evaluation processes to determine, and increase, potentially underutilized airspace. The use of these tools is then demonstrated with two possible applications. The first application involves examining the effects of possible modifications to the current en-route aircraft separation requirements. These modifications, which are based on the characteristics of large fixed wing aircraft, have the potential to increase the amount of available airspace, allowing for future

*Software Engineer, ARC-AF, Ames Research Center M/S 210-8 Moffett Field, CA 94035, Member
†Associate Scientist, ARC-AFM, Ames Research Center M/S 210-8 Moffett Field, CA 94035, Member
‡Professor, Baskin School of Engineering, 1156 High St, Santa Cruz, CA 95064.

increases in overall air traffic numbers. The second application shows how these tools can be used to provide a quick visual inspection of a potential UAS operation during different times of the day. This is done to help minimize the impact of a UAS on the air traffic during peak hours.

II. Background

The main thrust of this paper is on visual analysis of air traffic data. Hence, this section focuses on work related to visualizing air traffic data. One of the most popular techniques for visualizing air traffic data is to represent the trajectory of each aircraft as an animated particle. Many such visualizations are available on the web via sites such as youtube. A version that was designed by Aaron Koblin,[7] and accessible through http://www.youtube.com/watch?v=H2qTwvaQ_F4, demonstrates several techniques and embellishments for presenting the flight trajectories. More recently, the discrete nature of the flight tracks were smoothed out to obtain a continuous estimate of air traffic density using a view-dependent kernel density estimator.[8] Representing air traffic data as a density plot is not new. Kellner[9] also used density plots of the arrival and departure rates of aircraft at different airports to assess their capacity. Another source of air traffic visualizations is from air traffic simulators, such as FACET[10] and ACES.[11] These simulators are mainly used by government agencies to help develop and evaluate advanced air transportation concepts. While they do provide visual representations of air traffic, they are often particle based, showing the path of an aircraft over time. In addition, while the aircraft separation minima is taken into account during the simulation, it is not shown in the visualizations. This paper will discuss a method which explicitly factors in the separation minima, allowing for a visualization which depicts the surrounding airspace occupied by an aircraft. This, represented as a conflict probability volume, will allow the user to quickly identify how much airspace is actually being utilized.

There are many factors affecting air traffic congestion and airport capacity. One of those that is controllable and affects policy decisions is the specification of minimum separation between aircraft. Currently, this is set to 5 nautical miles horizontally, and 1,000 feet vertically when the aircraft is en-route.[12] This limit is adjusted as the aircraft approaches an airport and can drop to 3 miles horizontally on landing approaches to airports. The relative weight class of the leading and following aircraft are also taken into consideration in such situations in order to reduce risks due to wake turbulence.[13] The more stringent en-route limit accounts for higher aircraft speeds (typical passenger jets fly at average speed of 500 miles per hour or just over 8 miles per minute), weather impact on visibility, and wake turbulence from leading aircraft, among other factors. With the expected capabilities the NextGen surveillance and weather system, and integrated information systems, one can theoretically safely reduce the minimum separation requirements between aircraft. This paper provides visual analysis tools to examine the effects of different shapes and parameters describing the minimum separation criteria for aircraft.

Over the past few years, interest in UAS has rapidly increased. This is because of the possibilities they offer to both government and commercial interests. They would enable a broad range of satellite-like abilities, but at a much lower cost. Aerial photography, communications, environmental monitoring, and security are some of the abilities that UAS deployment could make possible on a large scale. Currently, UAS are predominantly used by the Department of Defense and the Department of Homeland Security, and often outside of the NAS. A handful of UAS are allowed to operate inside the NAS, though almost exclusively for national security or research purposes. However, each UAS operation must be pre-approved by the FAA on a case by case basis. This process is very tedious and does not scale well to large numbers of flights. There are a few studies on risk management of operating UAS. A recent study uses site-specific non-uniform probabilistic background air traffic to study these risks.[14] Using the visual analysis tools presented in this paper, checking whether the flight plan for a UAS will allow for a safe operation within the NAS can be accomplished expeditiously.

III. Air Traffic Density Volumes

Air traffic data usually consists of flight telemerty as well as specific flight plan information for each aircraft and aircraft type. The data set used to test and demonstrate the visual analysis tool contained 349,992 unique flight path records. Each flight path was comprised of latitude, longitude, and altitude information at 10 second intervals for the duration of the flight within the area of interest. The time of day and month in which the flights took place are specified. However, the specific date, aircraft type, and

departure and arrival airport of a flight was not included. The area spanned by the data is approximately 190 miles by 165 miles, between the altitudes of 0 and 60,000 feet. This area is centered around the cities of Devils Lake and Grand Forks in north eastern North Dakota. The data was collected from July 2008 to June 2009.

As previously mentioned, an intuitive way to visualize flight paths is by representing each aircraft as a particle and animating the trajectory of each aircraft. An alternative way using only a single static visualization is to trace and overlay individual flight trajectories. Figure 1 shows an example of such a visualization. Note that because there are so many flight trajectories in this data set, displaying each path at full intensity will produce an extremely cluttered display. Instead, each path is rendered at only 10% opacity. It is only then that clusters of similar flight paths can be observed. These clusters can in turn be associated with local and regional airports and show the prevailing takeoff and landing patterns. The flight trajectory data set in its current form is not directly suitable for such analysis. What is needed is a conflict probability volume that characterizes the probability that an aircraft can be found in any location in space (and time) and would therefore be in conflict with another aircraft (or UAS). In order to create such a volume, the discrete flight tracks need to be converted into continuous volumetric data.

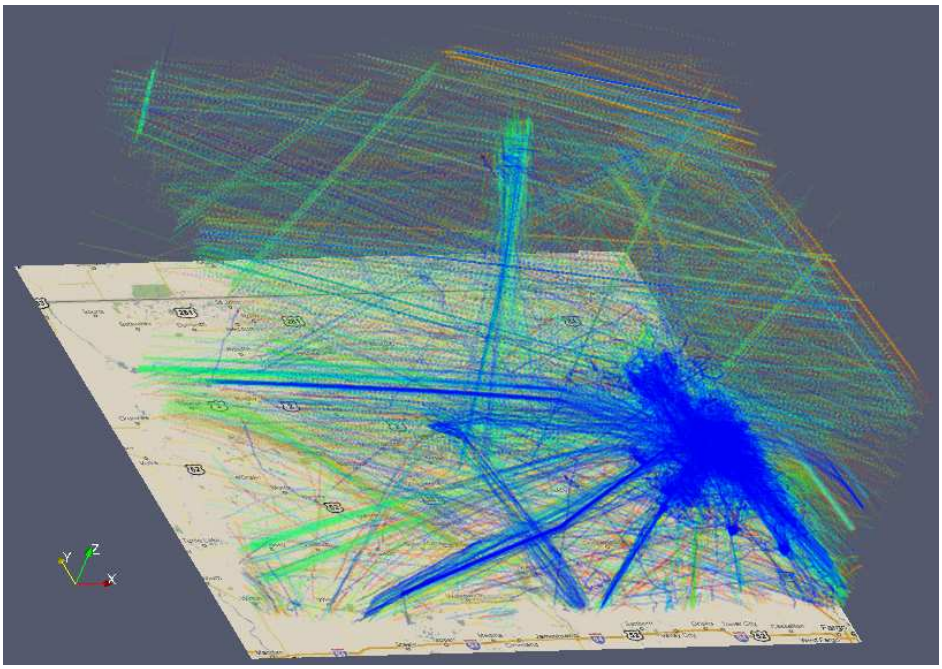

Figure 1. Close to 350,000 flight tracks during the span of 1 year are overlayed into a single image. The tracks are colored depending on which month the flight occurred. The data set covers an area of approximately 190 miles by 165 miles in the north eastern portion of North Dakota. This figure represents altitudes between 0 and 60,000 feet, with a vertical exaggeration of 10.

The first issue in creating such a volume is to determine an appropriate discretization of the volume in question. Another issue is how the discrete flight tracks are used to produce a smooth continuous volume. There were two main factors that were considered in order to determine the spatial discretization. The first factor is the FAA's minimum aircraft separation criterion. This is the required minimum separation between two in-flight aircraft, and is referred to as the conflict boundary, or more generally the separation minima. For en-route flights, this is specified to be 5 nautical miles (nmi) horizontally (30,380 feet), and 1,000 feet vertically, above and below each aircraft. Within the terminal area, horizontal separation is reduced to 3 nautical miles. This criterion would have discretized the volume of interest (190 mi x 165 mi x 60,000 feet) into a 38 x 32 x 60 grid. The second factor is determining a good resolution to represent the approximate distance traveled by an aircraft in 10 seconds, the time resolution of the data. Given the presence of both commercial aircraft and many relatively slower general aviation and training aircraft used by the University of North Dakota Aerospace aviation school, there is no single, all-purpose resolution. Speeds can vary between 140 mph (2,050 feet per 10 seconds) for the slower personal and training aircraft, to 600 mph (8,800 feet per 10 seconds) for the faster passenger airliners. Taking these factors into consideration, it is clear that the speed of the aircraft requires a much finer discretization than the separation criteria by itself. Therefore each voxel was made to represent a volume of 5,000 feet by 5,000 feet by 500 feet of airspace, which discretized the volume into a 208 x 174 x 121 grid. It should be noted that the volume of interest covers a much larger lateral area compared to its vertical extent. Thus in all the shown visualizations, there is a

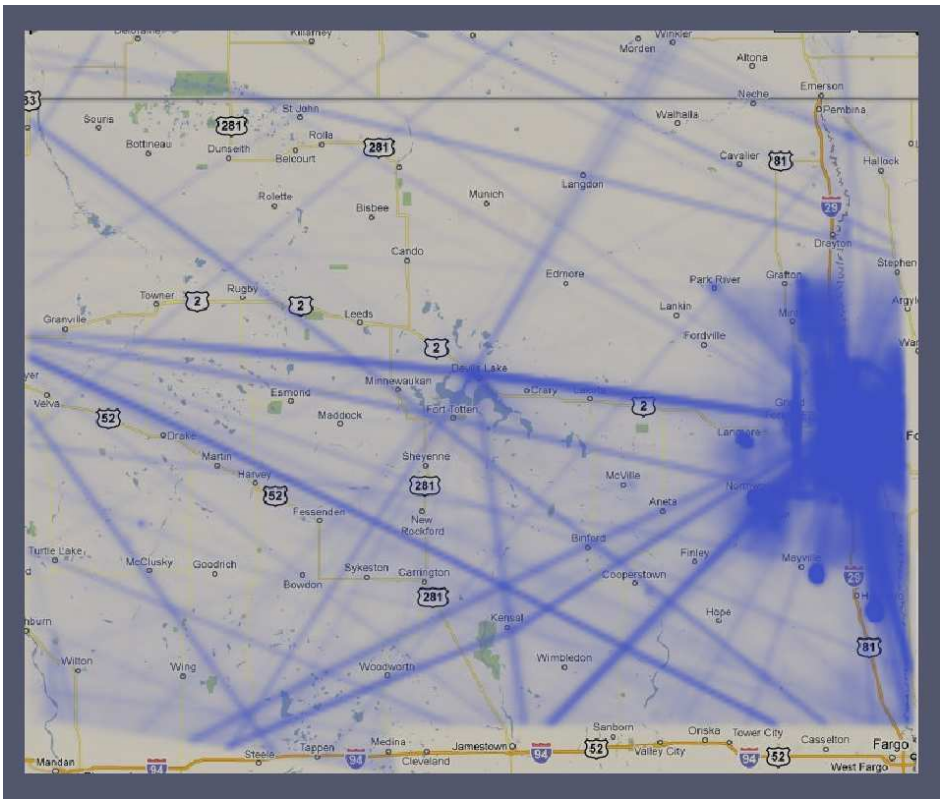

Figure 2. Close to 350,000 flight tracks during the span of 1 year are overlayed into a single image. The tracks are colored depending on which month the flight occurred. The data set covers an area of approximately 190 miles by 165 miles in the north eastern portion of North Dakota.

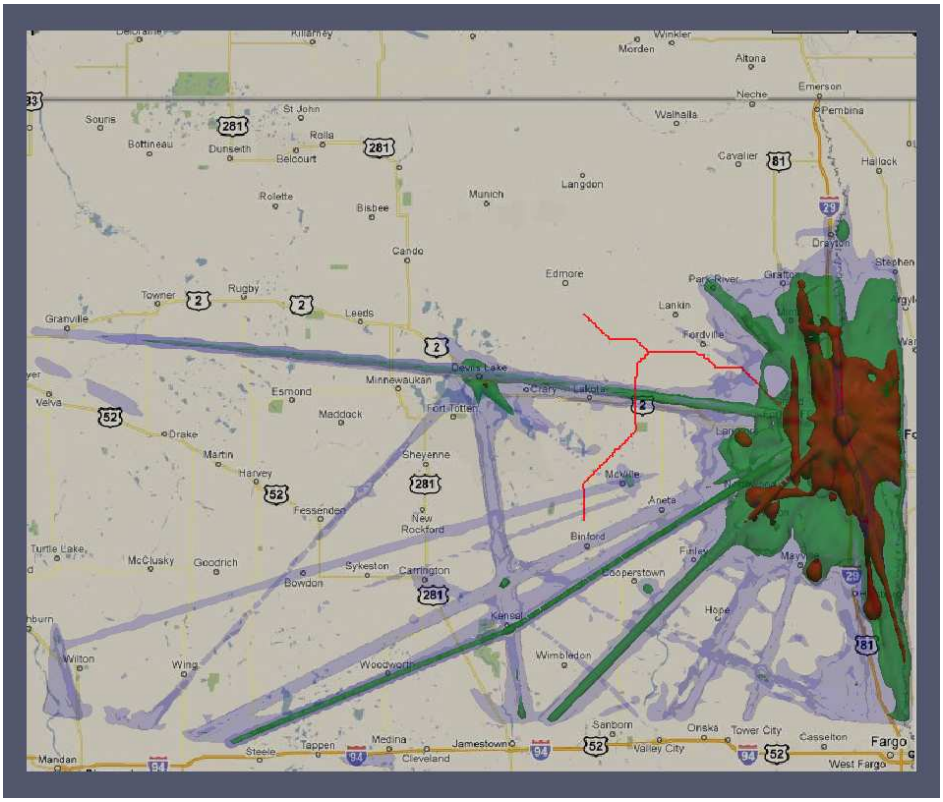

Figure 3. Similar to Figure 2 but with applied isosurfaces. The isosurfaces colored blue, green, and red, represent densities of 25, 125, and 625 aircraft per voxel, respectively. Also visible are two proposed UAS tracks, shown in bright red. Both UAS tracks are below 10,000 feet so only the flight tracks below 15,000 feet elevation are depicted.

vertical exaggeration by a factor of ten. This exaggeration allows for easily discernible variations in altitude when viewing the volume from lower angles.

The next step in calculating the conflict probability volume is to record the flight path information onto the grid. This process is discussed in the next section. Here, an alternative is discussed that also provides information about air traffic density, as well as explain some visualization methods which will be used to show the conflict probability. In essence what is required is a density volume that records the number of aircraft found at each grid point over the duration of interest (e.g. every 10 second interval, hourly, monthly, or over the entire year). For this purpose, cloud-in-cell (CIC) interpolation is used.[15] The CIC method works by first determining the voxel that contains the flight data point. It then distributes the contributions of the flight data point among the eight grid points of the voxel.The weighting function contains the product of three terms, each of which is linear in a dimension, with the total weighting normalized to unity. For example, if a data point is respectively 25%, 40%, and 70% of the way across the cell from the lower front left corner in the three dimensions, one attributes 0.75 x 0.6 x 0.3 of the particle's mass to that corner.[16] This produces a smoother density volume, and does not incur the cost of more expensive methods such as kernel density estimates.[17] These contributions are weighted according to proximity of the data point to each grid point.

Once processed, the air traffic density volume (and later on, the conflict probability volume) needs to be visualized. To do this, ParaView[18] and the open-source visualization toolkit Vtk[19] were used to help with the analysis. Two main visualization techniques were used. The first is direct volume rendering (for example, see Figure 2). While this technique does make relative high and low traffic areas discernible, the internal 3D structures of the volume, particularly at busy regions such as near airports, are very difficult to discern without the aid of additional tools such as cutting planes. This can be alleviated to some extent by adjusting the transfer function used in the rendering. However, because the distribution of air traffic density is skewed (towards very low aircraft counts, unless right over the airport), some form of contrast adjustment, such as histogram equalization, is needed. This can make it easier to distinguish slight variations in aircraft densities. An alternative is to use isosurfaces. An isosurface is a three dimensional contour surface formed over the points of a certain threshold value. With it one can specify different threshold values corresponding to different densities of interest. The different isosurfaces corresponding to these different thresholds are then rendered as surfaces with different colors and opacity (see Figure 3). As can be seen in Figure 2, the main thoroughfares are clearly visible as blue lines across the map. Areas with denser air traffic are also visible as blue shaded areas. However, it is hard to see distinct variations between low and high traffic areas other than as subtle variation of opacity in the shades of blue. Once isosurfaces are applied, the resulting renderings are much more informative. Figure 3 shows a rendering with three distinct isosurfaces. High and low traffic areas have become easily discernible, while still showing the main thoroughfares.

The aircraft density volumes previously described and depicted in Figures 2 and 3 provide a sense of the air traffic density over the region of interest. However, these density volumes only account for the location of the aircraft itself. They do not take into account the separation boundary surrounding an aircraft. As previously stated, this separation boundary is 5 nmi (30,380 feet) horizontally, which is much larger than the horizontal length of a single voxel (5,000 feet). Because of this an aircraft effectivly occupies a much larger area than the physical aircraft itself. When applied to a density volume, this means that the aircraft influnces a much larger area than just the voxel within which it is contained. The next section describes an alternate formulation that explicitly factors in the separation minima surrounding each aircraft.

IV. Calculating Conflict Probability

First, the conflict probability as used by Lee and Meyn[14] is described. The separation minima essentially describes a boundary in the shape of a circular cylinder, with a radius of 5 nmi and a height of 2000 ft (1000 ft above and below the aircraft), which is centered around an aircraft. The conflict probability, C_i, refers to the probability that another aircraft can be found within the separation boundary of an aircraft i. It can be computed by integrating the probability distribution functions over the separation boundary.

When calculating the probability distribution functions, horizontal uncertainty is considered to be independent of the vertical uncertainty since the horizontal coordinates come from radar data, while the vertical altitudes come from transponder data. The radar data came from the Advanced Synthetic Aperture Radar with a conservative uncertainty estimate of 0.125 nmi, while the transponder altitude uncertainty is at most 100 feet.

For the following set of equations, assume an aircraft i is at position (x_i, y_i, z_i) and one is looking to calculate the probability of a conflict with another aircraft at point (x, y, z). The horizontal probability distribution p_{h_i} is modeled as a 2D Gaussian with standard deviation σ_h, and the vertical probability distribution p_{v_i} is modeled as a 1D Gaussian with standard deviation σ_v.

$$p_{h_i} = \frac{1}{2\pi\sigma_h^2} exp\left(-\frac{(x-x_i)^2+(y-y_i)^2}{2\sigma_h^2}\right) \quad (1)$$

$$p_{v_i} = \frac{1}{\sqrt{2\pi}\sigma_v} exp\left(-\frac{(z-z_i)^2}{2\sigma_v^2}\right) \quad (2)$$

To calculate the conflict probability one first needs to determine the horizontal and vertical components. This is done by integrating the probability distributions, p_{h_i} and p_{v_i}, over the regions they span. In this case, the region is the circular cylinder representing the separation minima. The horizontal component is integrated over the circular area R_h, and the vertical component is integrated over the vertical distance R_v. Then, multiply the horizontal and vertical probabilities to produce the probability of conflict, C_i, of aircraft i, with another aircraft at point (x, y, z).

$$C_i = \int\int_{R_h} p_{h_i}(x,y)dxdy \int_{R_v} p_{v_i}(z)dz \quad (3)$$

$$R_h = (x-x_i)^2+(y-y_i)^2 \leq s_h^2 \quad (4)$$

$$R_v = |z-z_i| \leq s_v \quad (5)$$

When more than one aircraft is involved, the total conflict probability at a point is the sum of the contributions from each aircraft divided by the number of days spanned by the data used. To calculate the conflict probability volume, calculate the conflict probability contributions from each aircraft onto each grid point, normalized by the number of days (365 in this case). If there is interest in a particular time span within a day, then only consider aircraft data from that time span.

This volume has some similarities and differences when compared to the air traffic density volume obtained using CIC. They are similar in that both produce a smooth volume based on the air traffic data and provide information about local traffic density or probability. They are different in that CIC uses a linear dropoff function of contribution within the voxel that contains the aircraft, while the conflict probability uses a Gaussian dropoff function spread out over several voxels, up to 5 nmi from the aircraft position and 1000 ft above and below the aircraft. Also, the CIC volume provides average number of aircraft, while the conflict probability volume provides a probability of being in conflict with another aircraft.

V. Modifying Separation Minima

Next this visualization method will be used to investigate the effects that changes to the separation minima have on the airspace. In particular, the case when the shape of the conflict boundary volume is changed from a circular cylinder to an elliptical cylinder is examined. Note that this paper does not propose that the separation criteria be changed, it is merely showing how the visualization techniques can be modified to accommodate such changes.

Minor Axis \ Conflict Prob.	>0	≥0.005	≥0.01	≥0.05
5 nmi	14%	4.36%	1.68%	0.34%
4 nmi	12%	3.45%	1.33%	0.27%
3 nmi	9%	2.51%	0.98%	0.19%
2.5 nmi	8%	2.04%	0.81%	0.14%
2 nmi	7%	1.58%	0.64%	0.09%
1 nmi	3%	0.72%	0.29%	0.02%

Table 1. Effects of changing the separation minimum along the minor axis (orthogonal to aircraft heading) on the percentage of conflict probability volume using different safety margins.

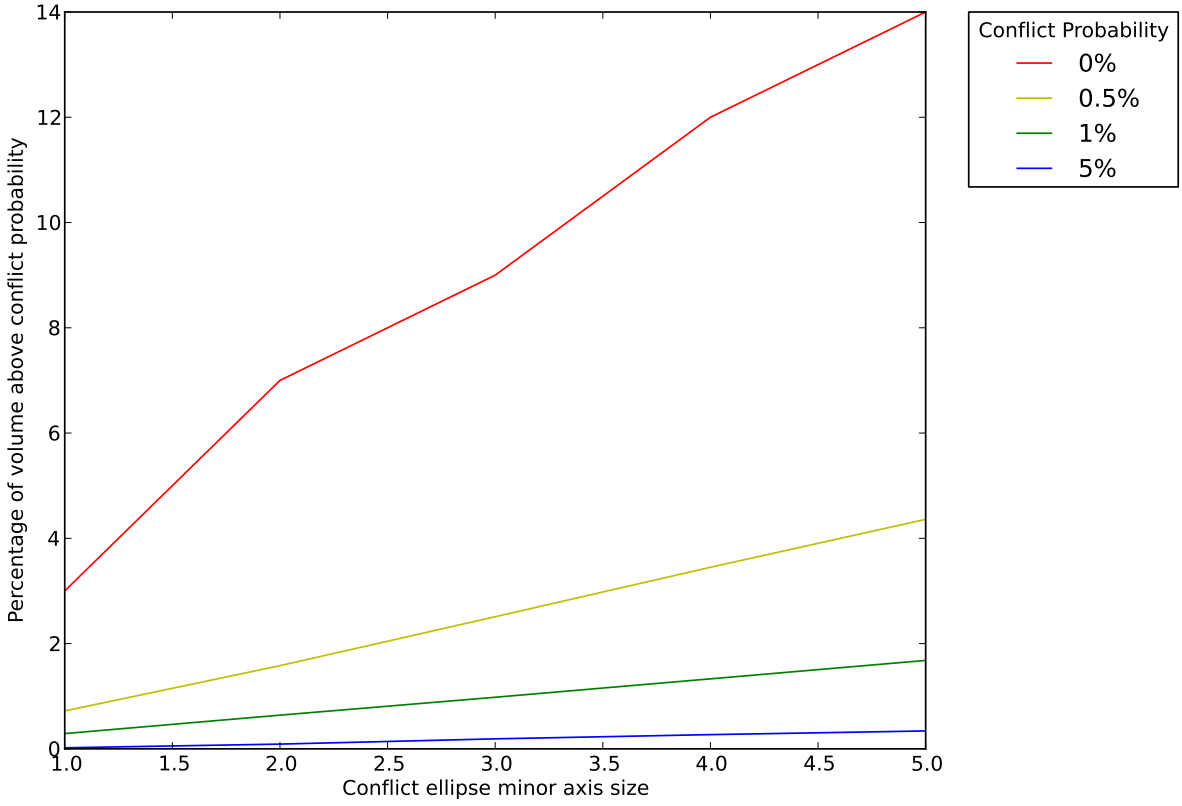

Figure 4. The same information as in Table 1 showing the linear relationship between the reduction of the percentage of airspace volume conflict as the minor axis is reduced. The different curves are for different thresholds of the conflict probability.

Aircraft, particularly large passenger or commercial planes, tend to maintain their current heading, altitude, and speed when en-route and make much smaller changes to their bearing. In addition, uncertainties in the direction of flight tend to be greater than laterally. Thus, the 5 nmi separation along the plane's heading is kept and align this with the major axis of an ellipse. The minor axis is then aligned orthogonal to major axis and can have a smaller separation minimum. The plane is still assumed to be centered in the ellipse to account for potential danger due to wake turbulence. The vertical components remain unchanged.

The formulation for the contribution of individual aircraft to the conflict probability volume is modified to be the product of the integral of three one-dimensional Gaussian distributions. These represent the Gaussian distributions along the major axis, the minor axis, and the vertical axis of each aircraft.

Various values of the minor axis separation minimum for were used and their impact on the reduction of conflict probability (which is proportional to air traffic congestion) quantified. Table 1 shows the effects of varying the minor axis from the original 5 nmi down to 1 nmi. The column headings indicate different thresholds on the conflict probability volume. Entries in the table indicate the percentage of grid points that are at or above the safety margin. For example, if the minimum separation along the minor axis is set to 4 nmi, 1.33% of the volume of airspace in the study region has a conflict probability of .01 or more Not surprisingly, as the separation minimum along the minor axis is reduced, a reduction in the conflict probability is expected. These effects are better seen in Figure 4 which shows the mostly linear relationship between separation minimum of the minor axis and the reduction in the conflict probability (expressed as a ratio of the conflict probability using a different minor axis against the conflict probability where the minor axis is 5 nmi). These results are bolstered by the visualizations showing how the conflict probability volume appears for two different separation minima of the minor axis, illustrated in Figure 5.

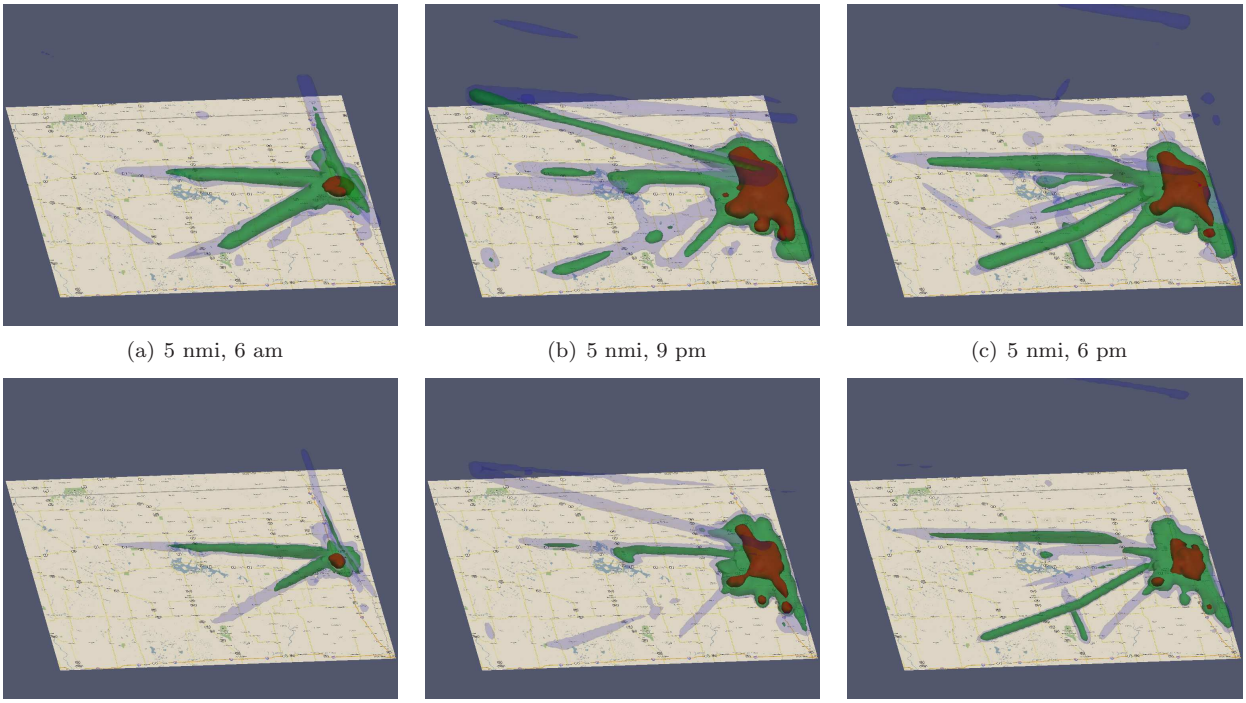

(a) 5 nmi, 6 am (b) 5 nmi, 9 pm (c) 5 nmi, 6 pm

(d) 2.5 nmi, 6 am (e) 2.5 nmi, 9 am (f) 2.5 nmi, 6 pm

Figure 5. The first row shows the conflict probability volumes using 5 nmi for the minor axis separation, while the second row uses 2.5 nmi for the minor axis separation. The first column shows the conflict probability volumes for 6am CMT when one can see the first surge of morning traffic. The next two columns show the conflict probability volumes for 9am and 6pm CMT respectively. These images show the spatial context for the reduction in conflict probabilities (and the air traffic density) as a result of reducing the separation minimum of the minor axis from 5 nmi to 2.5 nmi. Isosurface colors correspond to thresholds of .5% in blue, 1% in green and 5% in red. The quantitative changes in conflict volumes are shown in Table 1.

VI. Analysis of UAS Tracks

The conflict probability volume constructed based on the separation minima can also be used in the analysis of UAS trajectories. Current UAS operations are granted approval based on the impact on their safe operation within the NAS. This can be achieved by minimizing the cumulative conflict probability along their path, and ensuring that the conflict probabilities everywhere along their trajectories are below some threshold.

Two theoretical UAS tracks originating from the Grand Forks Air Force Base (RDR) were obtained for this study. One track goes towards the Tiger military operation area (MOA) along the northern corridor while the other goes towards the Devils Lake MOA along the southern corridor (See Figures 6 and 7). Both tracks share a common initial track from the RDR. These theoretical tracks are meant to show the least conflicting routes from RDR to the north and south MOAs. These tracks can be seen in Figures 6. Note that unlike the aircraft trajectory data, the UAS tracks were derived based on the resolution of the conflict probability grid. Hence, one can observe the jagged nature of their trajectories.

The two UAS paths do not go above 10,000 feet. To get a less cluttered visualization, air traffic above 15,000 feet is not shown. Once the initial trajectory clears the local RDR terminal airspace, one can observe that the northbound UAS path is relatively clear of any high density air traffic areas. On the other hand, the southbound UAS path not only leads to an area surrounded by high density air traffic, but this path even passes through a heavily traveled thoroughfare, depicted by the green isosurface in Figure 6. After showing this potential problem with the southern track to the track developers, two reasons were given: (i) The UAS headed to the Devil Lake MOA along the southbound path initially headed north in order to reduce the amount of time spent in the high traffic region. This is because the amount of time in the high traffic region is directly proportional to the conflict probability of the track. (ii) The grid that was used to develop the tracks essentially had horizontal information only. The vertical component of conflict probability was modeled as a simple linearly decreasing function up till 6,000 feet. Using the visual analysis tools, one can quickly examine and see potential conflicts of planned UAS trajectories against the prevailing air traffic patterns for a given area of operation.

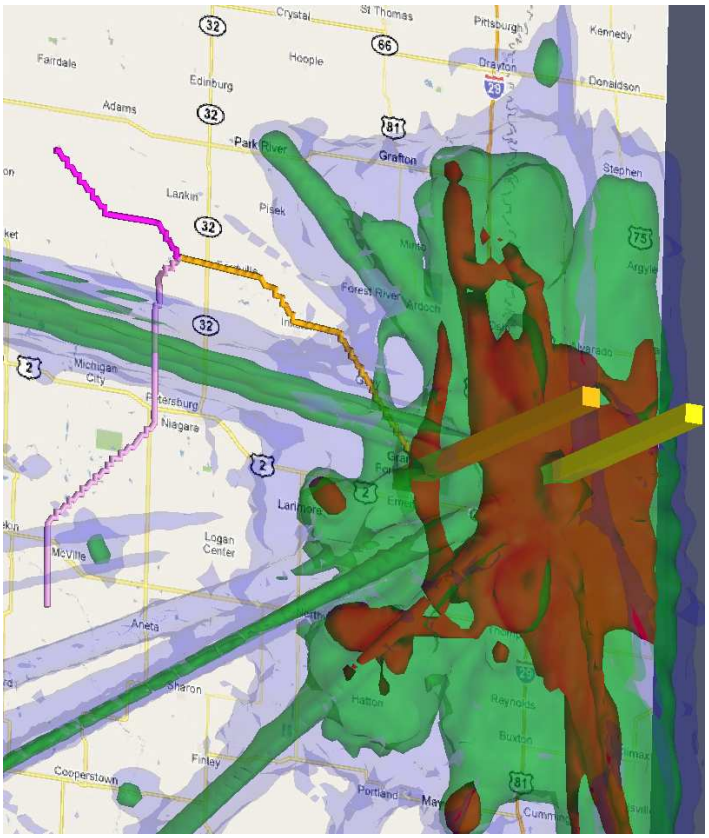

Figure 6. The UAS tracks are now colored differently to indicate the overlapping segment as well as the northern and southern forks. The two tall rectangular volumes indicate the locations of two airports : Grand Forks Intl. and Grand Forks AFB. The origination airport of the two UAS tracks is the Grand Forks Air Force Base on the left. The UAS headed along the south bound path clearly intersects the green isosurface, which represents a density of at least 125 aircraft per day on average.

Since Figure 6 is based on daily averages, one ask whether there are particular times during the day when it is safer relative to others. Obviously, flying during non-peak hours e.g. before 6am or late at night would reduce potential conflicts. In fact, from 10 pm to 5 am, the southbound UAS track does not intersect with

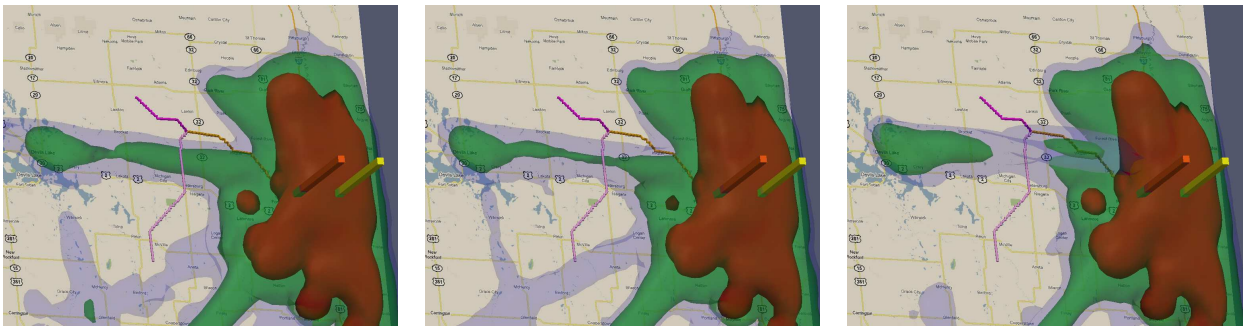

(a) 10 AM (b) 11 AM (c) 12 PM

Figure 7. Conflict probability with isosurface colors correspond to thresholds of 0.5% in blue, 1% in green and 5% in red in conflict probabilities. Here, one can see that the southbound UAS track has conflict probability of less than 1% at 10 AM and 11 AM, and less than 0.5 at 12 PM.

the green traffic region. However, these may not be the times required for the UAS operations. Between 6am and 9pm, the southbound UAS track has significantly reduced conflict probability between 10am and 5pm CMT (see Figure 7). Figure 7 shows that the southbound UAS penetrated the green zone at 10 am and 11 am, but it did not penetrate the green zone at 12 pm.

VII. Conclusion

This paper presented visual analysis tools for studying air traffic data. These tools can be used to view airspace not only in terms of aircraft density, but also for potentially underutilized airspace and the probability of conflict. This is because the analysis takes into account the FAA-required separation minima, not just the position of an aircraft. The utility of the tools has been demonstrated using two different applications. The first application explored the impact of changing the separation minima for aircraft on the conflict probability volume. This is one of the first considerations if aircraft are to fly closer to each other using the NextGen infrastructure, and can be used to study future proposed changes to the separation minima. The second application demonstrated the use of the visual analysis tool to quickly analyze a potential UAS trajectory in relation to the air traffic density (and hence conflict probability) in its theater of operation. This will help to quickly identify potential conflicts of a specific UAS path, as well as make known general areas of high risk that UAS operators should try to avoid.

The tools developed so far offer a glimpse of what visual analysis has to offer. There are some obvious rooms for enhancements. The conflict probability volume is currently based on air traffic trajectories only, but does not take into account the flight phase of the aircraft, possible variations in aircraft equipage, or environmental constraints such as prevailing winds, mountain ranges, etc. Taking these other relevant factors into account can provide a more accurate conflict probability volume. Also, while the visual analysis of UAS requests allows one to quickly identify potential conflicts, it operates under the assumption of an advance request. Future enhancements can incorporate dynamic modifications to the UAS flight plan that would take into account varying weather conditions as well as unforeseen deviations from the normal traffic patterns on flight day.

References

[1]National Air Traffic Controllers Association, "Air Traffic Control: By the Numbers," http://www.natca.org.

[2]Research and Administration, I. T., "TranStats," http://www.transtats.bts.gov.

[3]Price, H. J., "Fact Sheet FAA Forccast Fact Sheet Fiscal Years 2011-31," 2011, http://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=12440.

[4]Atkins, E., "Risk identification and management for safe UAS operation," *Systems and Control in Aeronautics and Astronautics (ISSCAA)*, 2010, http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5632396.

[5]Jr., L. D. and Duquette, A., "Fact Sheet – Unmanned Aircrafts Systems (UAS)," http://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=6287.

[6]"NextGen Implementation Plan 2011," March 2011, http://www.faa.gov/nextgen.

[7]Koblin, A., "Flight Patterns," http://users.design.ucla.edu/~akoblin/work/faa.

[8]Lampe, O. and Hauser, H., "Interactive visualization of streaming data with kernel density estimation," *Pacific Visualization Symposium (PacificVis)*, 2011, http://www.youtube.com/watch?v=1N5kg_tcG1s.

[9]Kellner, S., "Airport Capacity Benchmarking by Density Plots," *Deutscher Luft- und Raumfahrtkongress (German Aerospace Congress)*, 2009.

[10]Karl Bilimoria, Banavar Sridhar, G. C., "FACET: Future ATM Concepts Evaluation Tool," Vol. 9, No. 1, 2001, pp. 1–20.

[11]Kee Palopo, Larry Meyn, T. L., "Build 8 of the Airspace Concept Evaluation System," AIAA, AIAA Modeling and Simulation Technologies Conference, August 2011.

[12]FAA, "Order JO 7110.65U Air Traffic Control," 2012, http://www.faa.gov/documentlibrary/media/order/atc.pdf.

[13]FAA, "Pilot and Air Traffic Controller Guide to Wake Turbulence," http://www.faa.gov/training_testing/training/media/wake/04SEC2.PDF.

[14]Hak Tae Lee, Larry Meyn, S. K., "Safety Assessment of Unmanned Aerial System Operations over the Grand Forks Air Force Base," 2012, accepted as an Engineering Note to Journal of Guidance Control, and Dynamics.

[15]Hockney, R. W. and Eastwood, J. W., *Computer simulation using particles*, Taylor & Francis, Inc., Bristol, PA, USA, 1988.

[16]Faber, J., James C. Lombardi, J., and Rasio, F., "StarCrash: A parallel smoothed particle hydrodynamics (SPH) code for calculating stellar interactions," http://ciera.northwestern.edu/StarCrash/manual/usersmanual.pdf.

[17]Silverman, B. W., *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall, 1986.

[18]*ParaView Guide*, Kitware, http://www.paraview.org.

[19]Schroeder, W., Martin, K., and Lorensen, B., *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, Prentice Hall, New Jersey, 1996.