

A Comparison of Wake-Vortex Models for Use in Probabilistic Aviation Safety Analysis

J. F. Shortle, Ph.D.; Systems Engineering & Operations Research, George Mason University; Fairfax, Virginia, USA

Keywords: wake vortices, aviation, wake encounters, simulation

Abstract

The goal of this paper is to address part of a larger safety issue in aviation: What is the probability that a landing aircraft flies through a wake vortex generated by the aircraft in front of it and subsequently crashes? An important element needed to address this overall research question is a model to predict the evolution of wake vortices in time. Several such models have been developed and published in the literature. However, accurate prediction of wakes is inherently difficult. The objective of this paper is to assess the performance of different wake models *when imbedded within a larger probabilistic safety analysis*. This paper does not attempt to assess the accuracy of wake models by comparison to actual data. The purpose here is to compare the models with each other, and in particular, evaluate differences in performance when used within a safety assessment.

Introduction

Wake vortices are circulating flows of air produced by airplanes as they fly. An aircraft that flies through a wake may experience a roll which, if severe enough, can lead to a fatal crash. In order to mitigate this hazard, minimum separation distances are required between aircraft to ensure that wakes decay to a low strength. These separation requirements are a key bottleneck to improving capacity at an airport. This is particularly true for airports with closely spaced parallel runways, where adverse weather can reduce two runways to an effective capacity of one runway. The Federal Aviation Administration (FAA) and the National Aeronautics and Space Administration (NASA) are currently investigating changes to the wake separation rules for closely spaced parallel runways (ref. 1). Of course, before changes can be made, they must be demonstrated to be safe.

In principle, separation standards between successive aircraft guarantee that wake encounters do not occur during landings. In practice, there is some variability both in the locations of the airplanes and in the evolution of the wake vortices, so encounters may occur with some small probability. To estimate the probability of a wake encounter, two critical pieces of information are needed: (1) the probabilistic distribution of the locations and separations of aircraft, and (2) the probabilistic distribution of wake locations and strengths. In a previous paper (ref. 2), Shortle and Jeddi presented a methodology for estimating wake encounter probabilities using this information. They used actual flight-track data to model the locations and separations of aircraft, and they used a wake model estimate the locations of strengths of the generated wakes. Combining these sets of data in an overall simulation gave estimates for the probabilities of wake encounters.

A key limitation of this work was that only a single wake model was used in the analysis (specifically, the P2P model discussed in the next section). Wakes are very difficult to predict, and there is no clear consensus of a “best” model. The contribution of this paper is to compare several models of wake evolution against each other and to show how safety analysis can be sensitive to the choice of the wake model. We do not attempt to identify a “best” model, but merely seek to understand qualitatively how the choice of the wake model affects the resulting safety analysis.

Wake Modeling

The measurement and modeling of wakes is a key element in building the safety case for separation standards. However, these are difficult tasks. Direct measurement of wakes can be challenging because one cannot directly “see” a wake. A wake is a mass of air that is rotating in a coherent fashion. Such patterns of air also occur naturally in the atmosphere. Thus, to identify and track a wake, one must first identify a circulating pattern of air within the velocity field of the atmosphere. Then, one must determine that this circulation pattern is due to an airplane and not some other cause. There are two main technologies for measuring the velocity field of a region of air for the purpose of wake detection. One is a LIDAR (Light Detection And Ranging). A LIDAR is similar to a radar but using light

In *Proceedings of the 25th International System Safety Conference*, eds. A. G. Boyer and N. J. Gauthier, Baltimore, MD, 2007, System Safety Society, 581-589.

waves instead of radio waves. A LIDAR emits a pulse of light and then measures properties of the reflected light, including the time to return. Wakes can also be measured with sound using a SODAR (Sonic Detection And Ranging). For examples of wake measurements using these technologies, see (refs. 3-5).

A substantial amount of research has also focused on developing *models* of wakes and their behavior. Reference 6 gives a history of some of the early wake models for use in setting aircraft-separation standards. The basic challenge is that wake behavior is complex and depends on a variety of atmospheric conditions such as wind, temperature, temperature gradient, pressure, and so forth. These conditions vary in both space and time. Thus, it is difficult to *precisely* predict the path and circulation strength of a wake throughout its lifetime of several minutes. Naturally, a critical question for the use of wake models in safety analysis is: How accurate are the wake models compared with actual data?

It is not the objective of this paper to directly assess the accuracy of the models compared with data. The objective of this paper is to compare the use of different wake models *when imbedded within a larger probabilistic safety analysis*. We assume that while wake models may not be able to precisely track observed wakes on a case-by-case basis, they may be able to generate accurate trends over a large ensemble of modeled wakes. Again, the purpose here is not to directly evaluate the accuracy of individual models, but to understand differences between the models. Validation is still a significant issue, and these limitations must be kept in mind when evaluating the results presented in this paper.

This paper compares three different wake-evolution models:

- TDAWP (TASS-Driven Algorithms for Wake Prediction, ref. 7),
- D2P and P2P (Deterministic and Probabilistic Two-Phase Model, ref. 8).
- APA (AVOSS Prediction Algorithm, ref. 9), and

TDAWP and APA have been developed by NASA. D2P / P2P was developed by DLR, German Aerospace Research Center.

TDAWP was developed around physical principles using guidance from a more computationally expensive model, TASS (Terminal Area Simulation System, ref. 10). The TASS model is a full-scale fluid-dynamics model that numerically solves the Navier-Stokes equations for wake vortices. A major drawback of TASS is that it is computationally expensive. This is a significant problem for safety analysis where a large number of runs may be required to generate the rare events in question. TDAWP was developed with the idea of creating a faster model that generates results similar to TASS. TDAWP consists of a small set of ordinary differential equations. These equations can be solved numerically much faster than the original TASS model. And the solutions provide a good approximation to the TASS output over a range of input conditions.

A similar approach was used to develop D2P. D2P consists of a small set of differential equations. As with TDAWP, these equations were obtained from physical principles using the output of a large eddy simulation model (ref. 11) to guide the process. The differential equations can be solved quickly and they provide a good approximation to the large-scale model.

Compared with TDAWP and D2P, the APA model was developed more by using empirical observations rather than model output as guidance. APA integrates several sub-models that have appeared in the literature (e.g., refs. 12 and 13). The model contains four phases of wake-vortex evolution: away from the ground (phase 1), near the ground (phase 2), and close to the ground (phases 3 and 4). Each phase is modeled by a set of differential equations that can be evaluated numerically. A key feature of the APA model is the inclusion of the *ground effect*. In general, wakes descend vertically after being generated. If the wakes get near the ground, they tend to spread apart horizontally. This is significant for landing aircraft, and particularly for arrivals on closely-spaced parallel runways. In such configurations, it is possible for a wake to descend, hit the ground, and then travel across the ground to the other runway.

For this research, we have implemented the TDAWP and D2P models in MATLAB, based on the model descriptions in publicly available papers (refs. 7 and 8). The APA model (version 3.2) was obtained directly from NASA in the form of C++ source code. All models were imbedded within an overall MATLAB implementation (where the APA C++ code was compiled and then called from within MATLAB).

The models are relatively fast to simulate. Using the TDAWP model on a 2 GHz PC, it takes about 38 seconds to simulate 100,000 wakes (where each wake is simulated for 3 minutes of simulation time). Using the D2P model, it takes about 94 seconds to simulate 100,000 wakes. Using the APA model, it takes about 89 seconds to simulate 100,000 wakes. Because of the relatively fast simulation times, these wake models can be imbedded within larger safety simulations. However, *very* rare events can still be difficult to predict. For example, the approximate time to simulate 10^9 wakes (which is the approximate number of simulations needed to observe one event with probability 10^{-9}) is about 4.5 days using TDAWP, 11 days using D2P, and 10.5 days using APA.

For the purposes of this paper, we are concerned with the following outputs generated by all three models:

- Circulation strength as a function of time (measured in units of m^2 / sec),
- Vertical position as a function of time (measured in m / sec).

Figures 1 and 2 show sample output from the three models. We make several observations from Figure 1. First, the models all give the same *initial* circulation (about $320 \text{ m}^2 / \text{sec}$). The reason is that all three models use the same formulas for the initial vortex strength and speed (this will be discussed in a moment). However, after this initial agreement, there is a wide disparity in the circulation curves. Even the qualitative shapes of the curves differ. For example, the APA model shows what looks like exponential decay. The TDAWP and D2P models have a period of slow decay followed by a period of rapid decay (followed by a leveling-off). The D2P model in particular has a distinct “kink” in the curve at about 40 seconds. These two distinct phases are artifacts of directly approximating the output of TASS, which itself predicts a period of slow decay followed by a distinct transition to a period of rapid decay.

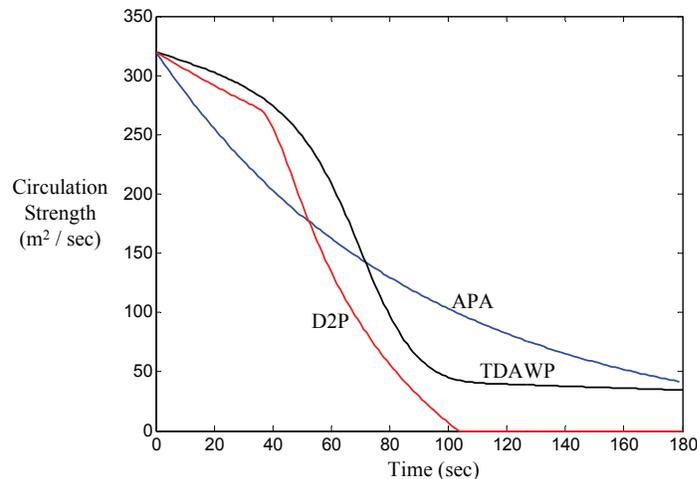


Figure 1 – Sample Output from Models: Circulation Strength

Figure 2 shows the vertical position, relative to the initial height, as a function of time (that is, zero denotes the altitude of the airplane at the point where the wake is initially generated). Here, there is much better agreement between the three models. The main difference is that TDAWP and D2P have two distinct phases – a period of rapid descent, followed by a period of slow descent – whereas the descent predicted by APA is more uniform.

Figures 1 and 2 illustrate two competing trends. For wakes that are over 100 seconds old, the APA model predicts the highest circulation strength of the three models (in this particular example). It predicts the most dangerous situation in this sense. On the other hand, the APA model also predicts the lowest descent height. It predicts the safest situation in this sense, because a wake that descends farther is more likely to be out of the approach path of the trailing aircraft. The reason for these competing trends is that the descent rate of a wake vortex is proportional to its circulation strength. Thus, a wake with a higher circulation is more dangerous once it is hit (higher severity), but it is less likely to be hit (lower likelihood).

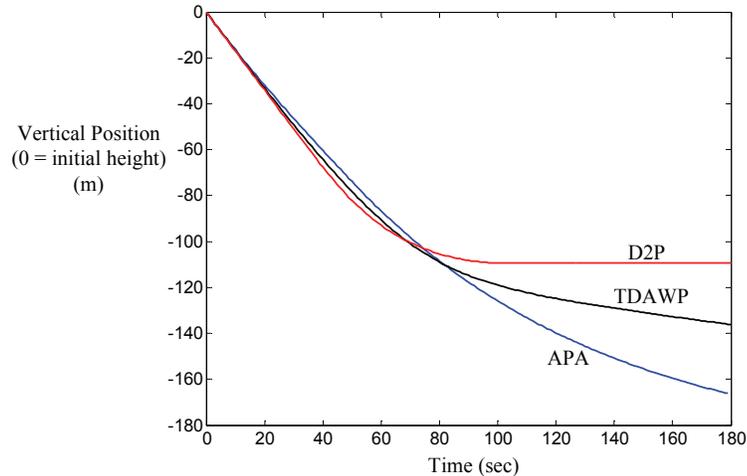


Figure 2 – Sample Output from Models: Vertical Position

The three wake models are driven by two main types of inputs: aircraft-related inputs and atmospheric-related inputs. The relevant airplane inputs are the velocity (V), mass (M), and wingspan (B) of the airplane that is generating the wake. The initial conditions of the wake are determined entirely by these parameters and the air density ρ (kg / m^3). These relationships are the same for all three models:

- The initial wake circulation is $\Gamma_0 = Mg / (\rho s B V)$, with units of m^2/s , where $s = \pi / 4$ is a constant, and g is the acceleration of gravity,
- The initial lateral spacing of the vortices is $b_0 = sB$,
- The initial descent velocity is $w_0 = \Gamma_0 / (2\pi b_0)$.

The relevant atmospheric inputs to the wake models are turbulence, stratification, wind profile, and air density. In this paper, we ignore the wind profile, which primarily affects the lateral transport of the wakes. Turbulence is quantified by the *eddy dissipation rate* ε , where higher values correspond to greater turbulence. Stratification is quantified by the Brunt-Vaisala frequency N . It is a measure of the vertical gradient of potential temperature in the atmosphere. The basic idea is that as a wake sinks, it encounters increasingly dense surrounding air. This causes an upward buoyancy force on the wake, leading to a vertical oscillation. When $N = 0$ (per second), there is no buoyancy force and the air is said to be neutrally stratified. Higher values of N correspond to a greater buoyancy force. In particular, sinking wakes may “bounce” in such conditions. The basic effect is that higher values of turbulence ε and stratification N cause wakes to decay more quickly. This also causes them to sink more slowly.

The specific parameters used to generate Figures 1 and 2 are: airplane velocity $V = 65$ m/s (about 130 knots), airplane mass $M = 76,000$ kg, and airplane wingspan $B = 38$ m. The atmospheric parameters used are air density $\rho = 1.2$ kg/m^3 , eddy dissipation rate $\varepsilon = .001$ m^2/s^3 , and Brunt-Vaisala frequency $N = 0$ / sec. All atmospheric parameters are assumed constant as a function of altitude.

Now, Figures 1 and 2 have shown differences between the models *for a single set of inputs*. We now give a comparison across a range of atmospheric parameters. Figure 3 shows a comparison of the three models by varying the turbulence parameter ε and the stratification parameter N . The values shown in the contour plots are the times (in seconds) for the wake circulations to decay to 40% of their original values. For example, using the TDAWP model with $N = 0.02$ / sec and $\varepsilon = 0.004$ m^2/s^3 , it takes about 90 seconds for the circulation strength to decay to 40% of its original value. All figures were generated using the aircraft parameters: airplane velocity $V = 80$ m/s (about 160 knots), airplane mass $M = 273,000$ kg, and airplane wingspan $B = 64.4$ m, corresponding to a landing Boeing 747-400.

We make the following observations from this figure. The APA model yields the largest *range* of decay times. For small values of ε and N , the decay times are longer than the other models (160 seconds or more). For large values of ε and N , the decay times are shorter than the other models (40 seconds or less). TDAWP and D2P yield decay times

in roughly the same range. This is somewhat expected since both of these models are derived as a fit to the model TASS. The decay times for TDAWP are slightly longer than for D2P.

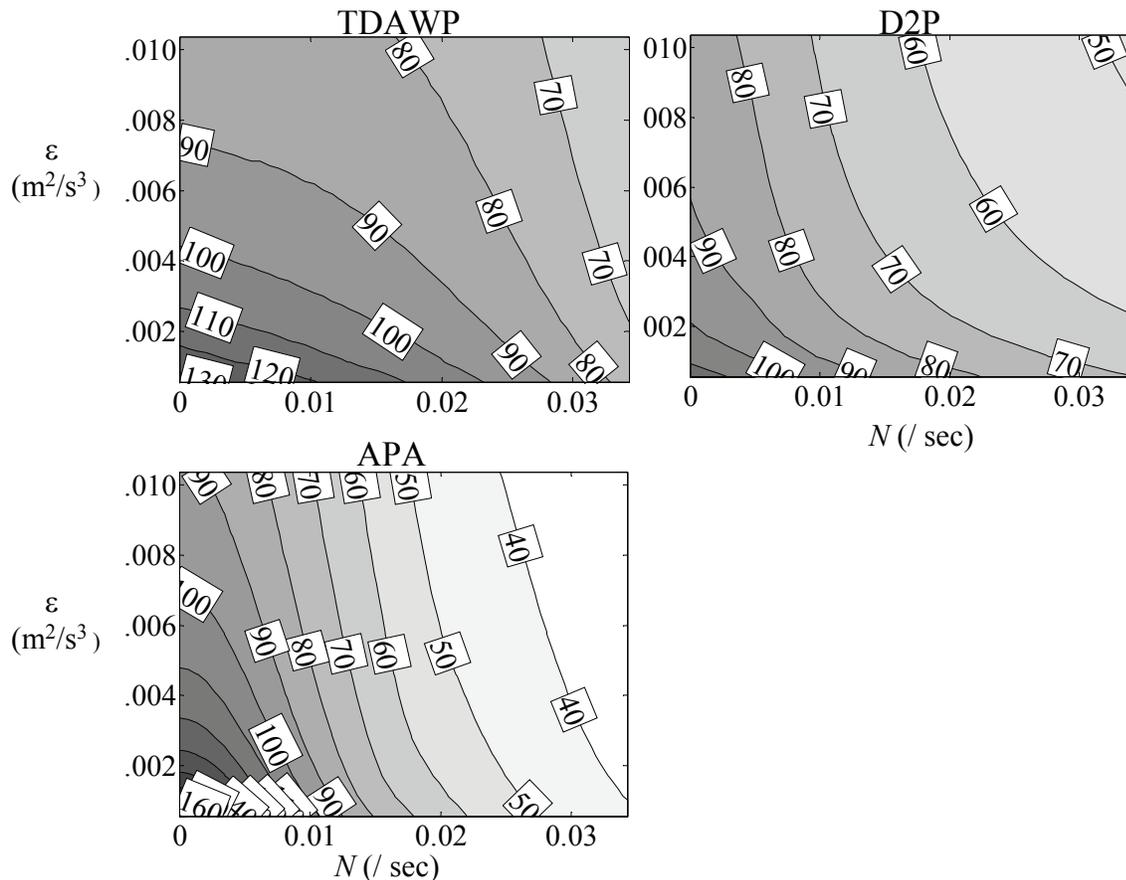


Figure 3 – Sample Output from Models: Time to Decay to 40% of Initial Circulation (sec)

Analysis Approach

The specific analysis approach used in this paper is discussed in more detail in a previous paper (ref. 2). Here, we simply highlight the key concepts. There are two key pieces of information needed to predict a wake encounter – the position of the trailing aircraft and the position and strength of the wake from the leading aircraft. (In this analysis, we ignore complicating factors such as wake-encounter geometry, encounter duration, size of the trailing aircraft, control authority of the trailing aircraft, pilot response, and so forth. These pieces of information are needed in the overall safety analysis – that is, to estimate the likelihood of a fatal accident due to a wake encounter. Here, we are simply concerned with predicting the probability of an encounter with a wake whose strength is greater than some threshold.)

The positions of aircraft are obtained from actual operations through multilateration flight-track data. These data give the positions of aircraft in the vicinity of the airport at one-second intervals. This particular analysis makes use of about 1,200 landings on a single runway at Detroit airport (DTW) over the span of one week. This gives a measure of the probabilistic variation of the tracks. From a safety perspective, the most dangerous situation is when (a) two successive aircraft are very close, possibly violating the separation standard, and (b) the trailing aircraft is low, while the leading aircraft is high. Because wakes generally sink, this leads to the possibility that the wake from the leading aircraft descends into the flight path of the trailing aircraft. The positions of wakes, on the other hand, are not obtained from direct observation, but rather through computer simulation of wake models as discussed in the previous section. In summary, the overall analysis is partly data-based and partly model-based.

In *Proceedings of the 25th International System Safety Conference*, eds. A. G. Boyer and N. J. Gauthier, Baltimore, MD, 2007, System Safety Society, 581-589.

The model simulates each airplane by moving it along its measured trajectory as specified by the multilateration data. Simultaneously, the model simulates the wakes generated from each airplane using one of the three wake models. This gives the wake strength and location as a function of time. By comparing the locations of the airplanes and the wakes, it is possible to count the number of wake encounters predicted by the dataset and models. To simplify the analysis, we only count wakes observed at a specific distance from the runway threshold (we have chosen 3 nautical miles, though other distances could be used equally well).

Since we only use a one-week dataset for this analysis, the number of wake encounters predicted is small. In order to generate more wake events, we define an event which is less severe than a wake encounter. A *wake alert* occurs when the following events occur simultaneously:

- The vertical position of the trailing aircraft is *below* the predicted vertical position of the wake, *and*
- The predicted circulation of the wake has not decayed below a specified threshold.

That is, a wake alert is defined to be an event where the trailing aircraft enters a *region* of space where the wake is *likely* to be. This is an event corresponding to increased danger, but it does not necessarily imply any realized harm or even a wake encounter. We expect that the frequency of wake alerts is correlated to the frequency of wake encounters.

We further extend this definition in a conservative sense by using *upper bounds* to specify the vertical positions and circulations of the wakes, rather than their original predicted values. That is, we revise the previous definition to say that a wake alert occurs when:

- The vertical position of the trailing aircraft is below the predicted *upper bound* of vertical position of the wake, *and*
- The predicted *upper bound* of circulation of the wake has not decayed below a specified threshold.

This conservatism yields a higher number of wake alerts, making it useful to see trends in the analysis. To generate these upper bounds, we take an approach that has been used to extend the *deterministic* D2P model to the *probabilistic* P2P model. Specifically, we generate an upper bound on circulation for each model by multiplying 1.2 times the predicted value given by the original model. We generate an upper bound on vertical position for each model by adding one initial vortex spacing b_0 to the original model.

Finally, a limitation of the analysis is that aircraft are assumed to have the same wingspan and landing mass, representative of an aircraft on the high end of the “large” weight category. More accurate estimates could be obtained by linking the Mode-S identifier in the multilateration data to the specific type of aircraft, but this has not been done yet.

Safety Analysis

As we have mentioned previously, the probabilistic variation of the aircraft positions are completely determined by the input flight-track data. The wake evolution, on the other hand, is driven by a variety of parameters as input into the wake models. Thus, a valid safety analysis must take into account the possible range of these values. Because measured turbulence and stratification parameters are not available for the observed period of the multilateration data, we vary these parameters over a range to see the effect on wake alerts. Figure 4 shows an example using the TDAWP model as the underlying wake model. The x -axis is the stratification parameter N . The y -axis is the turbulence parameter ϵ . Each point on the figure is the fraction of aircraft landings which resulted in a wake alert, based on the modeling approach described in the previous section. (Specifically, a wake alert occurs when the trailing aircraft flies below the wake and where the wake has a circulation strength of at least $70 \text{ m}^2/\text{s}$ at the time of passing, where $70 \text{ m}^2/\text{s}$ is the approximate background circulation occurring naturally in the atmosphere.) The highest values of ϵ and N correspond to the lowest probabilities of wake alerts. This is because higher turbulence and stratification cause the wakes to decay more quickly. Somewhat surprisingly, the lowest values of ϵ and N do not correspond to the highest probabilities of wake alerts. This is because even though the wakes are stronger, they are descending more quickly, and thus are more likely to be clear of the path of the trailing aircraft. The most dangerous atmospheric conditions appear to be low turbulence, with moderate values of stratification.

Figures 5 and 6 show similar analyses, but using the P2P and APA models. From a qualitative perspective, figures 4 and 5 (TDAWP and P2P) are much more similar than figure 6 (APA). This is somewhat expected, since the

TDAWP and P2P models are derived from the same large-scale model TASS. Both figures show a similar trend: the highest values of turbulence ϵ and stratification N correspond to the lowest alert rates. The highest alert rates are generated by either low ϵ and moderate N or low N and moderate ϵ . The APA model (figure 6), on the other hand, looks much different qualitatively and predicts much lower alert rates (also note that the scale is also slightly different in this figure). This is because the APA model tends to predict lower vertical descents of the wakes (see figure 2), particularly for the lower values of ϵ and N where the wake alert probabilities are higher. Even though the wakes are predicted to be stronger, the fact that the wakes are more likely to be clear of the approach path dominates the overall alert-probability calculation.

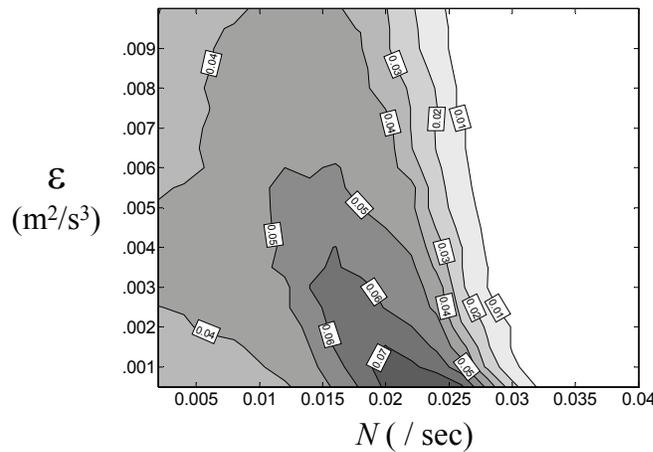


Figure 4 – Wake Alert Probability Using TDAWP Model

Figure 7 illustrates this more concretely. Recall that a wake alert occurs when two conditions are satisfied: (1) The vertical position of the trailing aircraft is below the wake, *and* (2) The wake strength is above a specified threshold. Figure 7 shows how many times each condition is violated for each wake model, varying the computation over the turbulence parameter ϵ . We have fixed $N = 0.01$ / sec, so the figure represents a vertical slice from figures 4-6, breaking down the overall wake alert probability into individual violations of the two conditions. As ϵ increases, the number of circulation violations decreases (because the wakes are decaying faster), but the number of vertical violations increases (because the wakes are not sinking as far). The vertical condition is violated less frequently. Since both conditions must occur, the vertical condition dominates the overall probability. Specifically, for lower values of ϵ , the APA model has a smaller number of vertical violations which explains why the APA model yields less overall wake alerts.

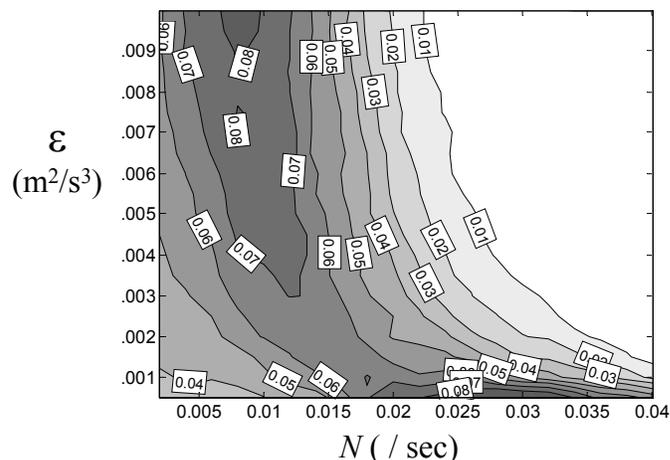


Figure 5 – Wake Alert Probability Using P2P Model

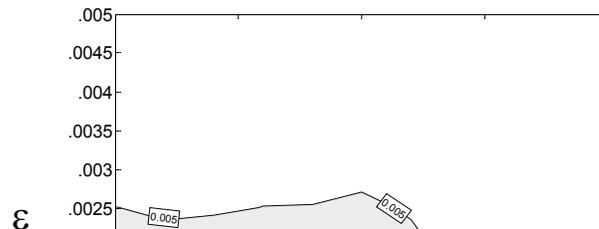


Figure 6 – Wake Alert Probability Using APA Model

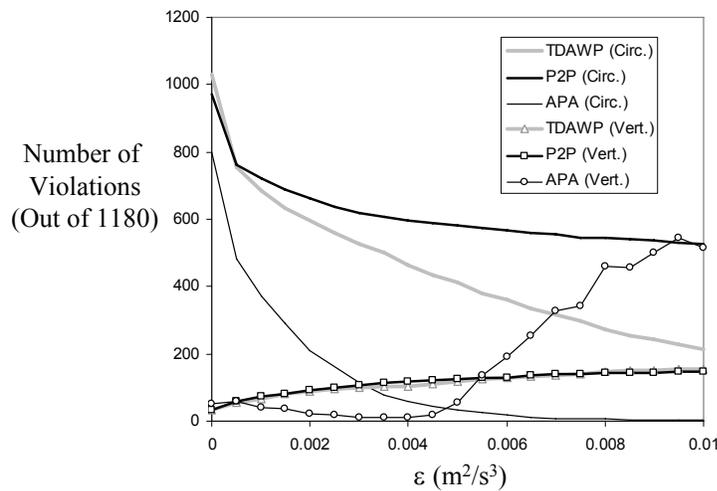


Figure 7 –Wake Alert Conditions

Conclusions

This paper has compared the use of three different wake models within a safety analysis for wake encounters of landing aircraft. The main conclusions are:

1. There are non-trivial qualitative differences among the models (figures 1-3). This leads to different estimates of safety levels with respect to estimating wake-vortex alerts (figures 4-6). Thus, any wake-safety analysis that uses a wake model should include a sensitivity analysis with respect to the models.
2. In the analysis conducted, the APA model predicts a smaller number of mild encounters than TDAWP and P2P. This is because the model tends to predict lower vertical descents at typical separations of about 90 seconds. The vertical separation appears to be the dominating condition in driving the overall probability of a wake alert. We conjecture, however, that if the circulation threshold that defines a wake alert is increased from $70 \text{ m}^2/\text{s}$ to say $200 \text{ m}^2/\text{s}$, then the APA model would predict a higher probability of wake alerts. In this case, the circulation criterion would play a stronger role in the overall probability. Because the APA model tends to predict longer decay times when the turbulence parameter ϵ and the stratification parameter N are low, we conjecture it would predict a higher wake-alert probability than the other models.

Finally, this paper highlights the continued need for further research in wake models. Wake prediction is inherently difficult, but wake models play an important role in safety assessments associated with capacity improvements.

Acknowledgments

This paper solely represents the opinions of the author and does not necessarily reflect the opinions of the United States government or NASA. The author thanks Wayne Bryant and Ed Johnson at NASA Langley for support of this research. The author also thanks Babak Jeddi for processing the flight-track data used in this analysis.

References

1. Lang, S., A. Mundra, W. Cooper, B. Levy, C. Lunsford, A. Smith, and J. Tittsworth. A Phased Approach to Increase Airport Capacity through Safe Reduction of Existing Wake Turbulence Constraints. *Air Traffic Control Quarterly*, Vol. 11, No. 4, 2003, pp. 331-356.
2. Shortle, J., and B. Jeddi. Probabilistic Analysis of Wake Vortex Hazards for Landing Aircraft Using Multilateration Data. To appear in *Transportation Research Record*, 2007.
3. Dasey, T., S. Campbell, R. Heinrichs, M. Matthews, R. Freehart, G. Perras, and P. Salamitou. A Comprehensive System for Measuring Wake Vortex Behavior and Related Atmospheric Conditions at Memphis, Tennessee. *Air Traffic Control Quarterly*, Vol. 5, No. 1, 1997, pp. 49-68.
4. Chalson, M., J. Hallock, S. Mackey, M. Soares, F. Wang, D. Burnham, and S. Hannon. Data Supporting STL Proposed Change, WakeNet USA 2005, Boca Raton, FL, 2005.
5. Wang, F., H. Wassaf, and A. Gulsrud. Acoustic Imaging of Aircraft Wake Vortex Dynamics, 23rd AIAA Applied Aerodynamics Conference, Toronto, Ontario Canada, AIAA 2005-4849, 2005.
6. Burnham, D., and J. Hallock. Wake Vortex Separation Standards: Analysis Methods. Publication DOT/FAA/ND-97-4. FAA and U.S. Dept. of Transportation, 1997.
7. Proctor, F., D. Hamilton, and G. Switzer. TASS Driven Algorithms for Wake Prediction. 44th AIAA Aerospace Sciences Meeting and Exhibit, AIAA 2006-1073, Reno, NV, 2006.
8. Holzapfel, F. Probabilistic Two-phase Wake Vortex Decay and Transport Model. *Journal of Aircraft*, Vol. 40, No. 2, 2003, pp. 323-331.
9. Robins, R., and D. Delisi. NWRA AVOSS Wake Vortex Prediction Algorithm Version 3.1.1. Publication NASA/CR-2002-211746, NASA, 2002.
10. Switzer, G., and F. Proctor. Numerical Study of Wake Vortex Behavior in Turbulent Domains with Ambient Stratification. 38th Aerospace Sciences Meeting and Exhibit, Reno, NV, AIAA-2000-0755, 2000.
11. Holzapfel, F., T. Gerz, and R. Baumann. The Turbulent Decay of Trailing Vortex Pairs in Stably Stratified Environments. *Aerospace Science and Technology*, Vol. 5, No. 2, 2001, pp. 95-108.
12. Greene, G. An Approximate Model of Vortex Decay in the Atmosphere. *Journal of Aircraft*, Vol. 23, No. 7, 1986, pp. 566-573.
13. Sarpkaya, T. New Model for Vortex Decay in the Atmosphere. *Journal of Aircraft*, Vol. 37, No. 1, 2000, pp. 53-61.

Biography

J. F. Shortle, Ph.D., Associate Professor, Systems Engineering & Operations Research, George Mason University, 4400 University Dr., MS 4A6, Fairfax, VA 22030, USA, telephone – (703) 993-3571, facsimile – (703) 993-1521, e-mail – jshortle@gmu.edu.

Dr. Shortle is currently an associate professor of systems engineering at George Mason University. His research interests include simulation and queueing applications in telecommunications and air transportation. Previously, he worked at US WEST Advanced Technologies. He received a B.S. in mathematics from Harvey Mudd College in 1992 and a Ph.D. and M.S. in operations research at UC Berkeley in 1996.