Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

RARE EVENT SIMULATION FOR POTENTIAL WAKE ENCOUNTERS

Azin Zare-Noghabi
John F. Shortle

Department of Systems Engineering and Operations Research
George Mason University
Fairfax, VA 22030, USA

ABSTRACT

A flying aircraft produces two coherent rotating vortices of air in its trail. If another aircraft flies into one of these vortices, it can experience an un-commanded roll. Because of the potential risk, wake separation standards exist to significantly reduce the probability of such events. Thus, wake encounters are inherently rare. As new procedures and technologies are proposed to increase the capacity of the airspace, rare-event simulation is necessary to evaluate to the safety of proposed changes. This paper explores the performance of a rare-event splitting technique in the context of estimating probabilities of potential wake encounters. The goal is to identify good strategies for the splitting method while using a relatively simple model for the wakes. Suggestions for the choice of the level function and the locations of levels are given.

1 INTRODUCTION

When an aircraft flies, it produces two coherent rotating vortices of air in its trail. These vortices can persist for several minutes at high altitudes. If another aircraft flies into a wake vortex, it can experience an un-commanded roll. Because of the potential risk, wake separation standards – which specify a minimum distant or minimum time between aircraft – exist to significantly reduce the probability of such events. Thus, wake encounters are inherently rare events.

While separation standards ensure safety, they also impose a limiting constraint in terms of increasing capacity of the airspace system. As new procedures, concepts, and technologies are proposed to improve airspace capacity, such changes must be demonstrated to be safe prior to implementation. Thus, common questions arise: Will a given procedure maintain the same level of safety that is observed today with respect to wake vortices? What is the minimum separation that can be achieved while still maintaining a given level of safety? Addressing such questions often involves Monte-Carlo simulation of rare events.

The main problem in estimating very small probabilities by standard Monte Carlo simulation is efficiency – to get a reasonable estimate of the probability, a huge number of simulation replications is needed, which might be too time consuming depending on the complexity of the system. For example, a simulation of several thousand flights might generate zero simulated encounters. In the literature, two common methods for rare-event simulation are splitting and importance sampling (e.g., Rubino and Tuffin 2009).

The objective of this paper is to explore the performance of rare-event splitting techniques to estimate probabilities associated with wake encounters models. The goal is to identify good strategies for the splitting method while using a relatively simple model for the wakes. Future work will adapt the methodology to more complex wake models. The splitting method has been applied to other real world problems such as collision risk in aviation (Blom et al. 2006, Blom, Bakker, and Krystal 2009, de Oliveira et al. 2010), cascading blackouts (Shortle 2013), queueing networks (Garvels 2011) and network reliability (Murray, Cancela, and Rubino 2013).

The risk of wake vortex encounters has been studied extensively in the terminal area during take-off and landing. One example a simulation-based risk assessment approach for landing aircraft is given in Kos et al. (2001). Fewer studies have focused on wakes at cruise altitudes. One approach is to combine track data, which can either be historical surveillance data or synthesized tracks generated from a computer simulation, and a model of wake vortex behavior, such as the model from Robins and Delisi (2002). Hoogstraten et al. (2015) developed a simulation framework using recorded track data and a wake simulation model to analyze the risk of wake turbulence in upper airspace. Results include estimating the probability of encounter, determining the main factors contributing to risk and recommending mitigating measures. It was found that encounter geometry is an important contributing factor after weight, and most encounters happen when one or both aircraft are climbing or descending. Nelson (2006) reviewed a number of studies related to en-route wakes and argued that en-route encounters are likely to increase over time with increasing disparity in aircraft sizes (e.g., super heavy aircraft which generate powerful wakes sharing airspace with very light jets which are more susceptible to wakes).

2 AIRCRAFT DYNAMICS AND WAKE VORTEX MODEL

2.1 Aircraft Model

In this paper, aircraft are modeled as point masses, where aircraft states include the along-track position, along-track airspeed, across-track position, and altitude. We consider the acceleration of each aircraft as the control parameter in the along-track dimension. Movements of the aircraft in all three dimensions are assumed to be independent from each other. An example of a more complex model with six coupled state variables, derived from basic aerodynamic principles, is given in Glover and Lygeros (2004).

2.1.1 Along-Track Movements

Our model simulates two aircraft – a leading aircraft and a trailing aircraft. Different control laws govern the along-track movements of the two aircraft. The leading aircraft tries to maintain a preferred target speed without considering the position or the speed of the trailing aircraft. Conversely, the trailing aircraft adjusts its velocity to maintain a target separation from the leading aircraft using a proportional-derivative (PD) controller. To model the along-track movement of the leading aircraft, we adopt the Ornstein-Uhlenbeck (OU) process to control the velocity, which is given by the following stochastic differential equation:

$$dv_l = -\rho(v_l - \mu)dt + \sigma dW_t .$$

Here, v_l is the velocity (airspeed) of the leading aircraft, μ is the target velocity, ρ is the reversion rate, σ is the volatility parameter and W_t is the Wiener process (standard Brownian motion). Under this process, random perturbations occur, but the process is mean reverting, so the velocity tends to return to the preferred target average. In steady state, it can be shown that the velocity has a normal distribution with mean μ and variance $\sigma^2/2\rho$. The reversion rate and volatility parameter can be chosen so that the velocity stays in a desired interval around the target velocity with some specified probability.

For the following aircraft, a PD controller is used to make the trailing aircraft track the leading aircraft as the reference point. The trailing aircraft tries to maintain a desired along-track separation D and to fly at the same speed as the leader. This yields the following stochastic differential equation

$$dv_f = k_p(x_l - D - x_f)dt + k_d(v_l - v_f)dt + \sigma dW_t ,$$

where x_f and v_f (x_l and v_l) are the along-track position and velocity of the following (leading) aircraft, and D is the desired along track separation. In steady state, it can be shown that the mean and variance of the along track separation between the leading and trailing aircraft are

$$E[x_l - x_f] \qquad \text{and} \qquad \text{var}[x_l - x_f] = \frac{\sigma^2}{2k_p k_d},$$

and the variance of the velocity of the following aircraft is

$$\text{var}[v_f] = \frac{\sigma^2}{2k_d}.$$

Using these results, the parameters k_d, k_p, and σ can be chosen so that the aircraft separation and velocity have some prescribed variability.

2.1.2 Vertical Movements

For the vertical dimension, each aircraft tries to maintain a target altitude. A similar mean reverting OU process, with parameters μ_z, ρ_z and σ_z, is used to model vertical position. The parameters are chosen so that each aircraft remains within h feet of the target altitude 95% of the time. The stationary distribution is a normal distribution where the half-width of the 95% bound is $1.96\sigma_z/\sqrt{2\rho_z} = h$. Figure 1 shows a sample path of vertical movements of an aircraft. The axes are not to scale, and the actual motion would be represented as a smoother curve.

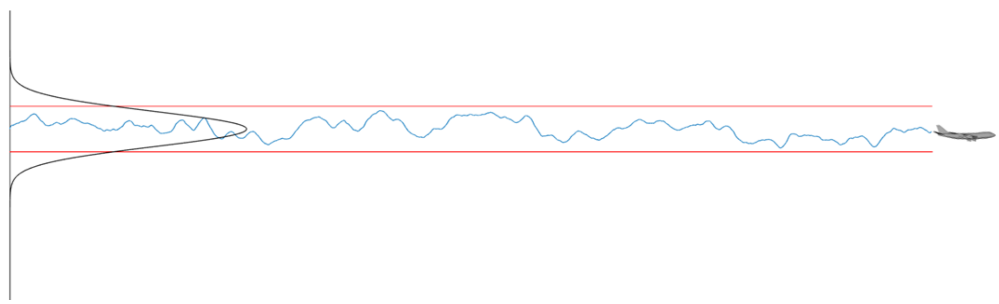

Figure 1: Sample path for vertical trajectory of an aircraft (axes not to scale).

For two aircraft in trail, both z_l and z_f (the altitudes of the leading and trailing aircraft) have normal distributions with mean μ_z and variance $\sigma_z^2/2\rho_z$ so the relative vertical position of the trailing aircraft to the leading aircraft $z_l - z_f$ has a normal distribution with mean zero and variance σ_z^2/ρ_z, since the vertical movements of the aircraft are assumed to be independent. Figure 2 shows a snapshot of the steady-state relative position of two aircraft in the along-track and vertical dimensions. The figure is obtained by simulating trajectories of 50,000 pairs of aircraft, according to the previously described control laws for 30 simulated minutes, and recording their relative position at the end of the simulation. Figure 3 shows normalized histograms for the along track and vertical distances of the aircraft in steady state. The fitted probability density function matches the theoretical probability density function.

2.1.3 Lateral Movements

Similar to the vertical dimension, we use a mean reverting OU process for modeling the lateral motion of aircraft. This dimension is treated independently of the other two dimensions. The parameters of the process are again chosen in a way such that specified navigation performance targets are met (e.g., the aircraft remains within a certain distance of the centerline 95% of the time).

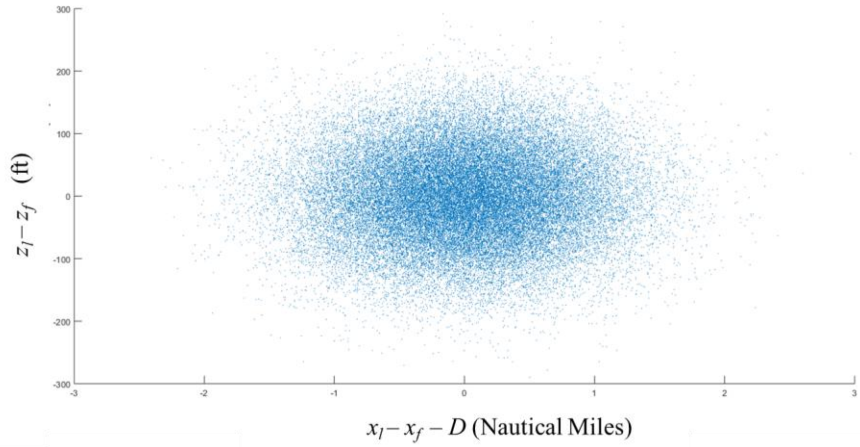

Figure 2: Snapshot of relative location of aircrafts in steady state.

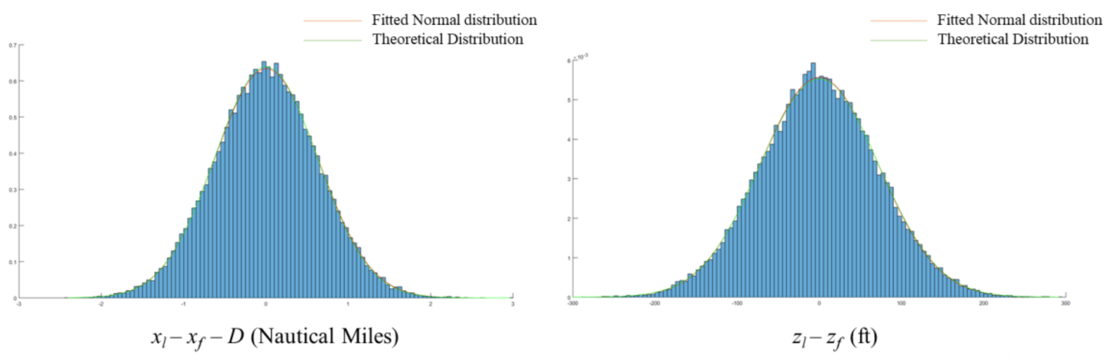

Figure 3: Histograms for relative positions (along track and vertical) in steady state.

2.2 Modelling the Wake Vortex Region

Every aircraft in flight generates wake vortices as a by-product of generating lift. The pressure difference below and above the wings causes the high pressure air from below the wing to go around the wingtip into the region where there is lower pressure above the wing. This rollup generates counter-rotating vortex trails from the wingtips. Wake vortices generally descend in altitude and can persist for several minutes.

Wake vortex behavior can be quite complex. Factors affecting the wake behavior include the weight, air speed, wing shape and wing span of the generating aircraft as well as meteorological conditions such as air density, wind speed, wind direction, turbulence, and temperature stratification. A variety of models have been developed in the literature to capture this behavior (e.g., Robins and Delisi 2002). Many of these models involve numerically solving a set of differential equations. One challenge is that embedding these models into a simulation requires resolving the system of differential equations at every time step for every active aircraft in the simulation. This is because each new time step results in a new set of initial conditions and a new “starting point” for the wake of each aircraft. An aircraft may change altitude, heading, or airspeed at each time step resulting in new initial conditions for the wake.

The purpose of this paper is to provide a proof-of-concept and some preliminary results on using a rare-event splitting method in estimating the probability of potential wake vortex encounters. Since the focus is on implementation of the rare-event simulation methodology, the wake model is simplified as much as possible. This characterization allows us to focus tests on basic properties of the splitting methodology. Future work will embed the more complicated models into this framework.

We first consider a two dimensional model in which only along-track and vertical movements are considered for the aircraft (lateral movement is ignored). Wang and Shortle (2012) have shown via sensitivity analysis that the variability parameters in these two dimensions are key parameters in terms of reducing encounters. Figure 4 shows a high level geometric model of the wake region, characterized by a triangle in two dimensions. The dimensions of the region – namely the length of the wake region, the minimum predicted descent, and the maximum predicted descent – may be different for different aircraft with different velocities at varying meteorological conditions. The region does not predict the exact position of the wake, but rather defines an area that is likely to contain the wake. So if another aircraft enters the triangle, it may or may not hit the wake. We refer to these events as *potential wake encounters*.

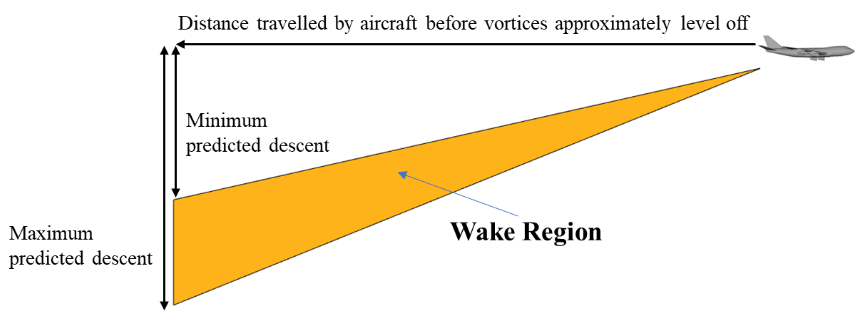

Figure 4: Wake region behind the leading aircraft in two dimensional model.

In the three-dimensional model we allow for lateral movement. In higher altitudes, the vortices remain spaced slightly less than a wingspan apart, drifting with the crosswind. To take into account the possible lateral drift of the wake vortices, the 3D model is a polyhedron in shape of a wedge which is shown from different views in Figure 5. The maximum lateral drift is an input for the model. The symmetry of the model assumes that the crosswind is unknown with equal uncertainty in either direction.

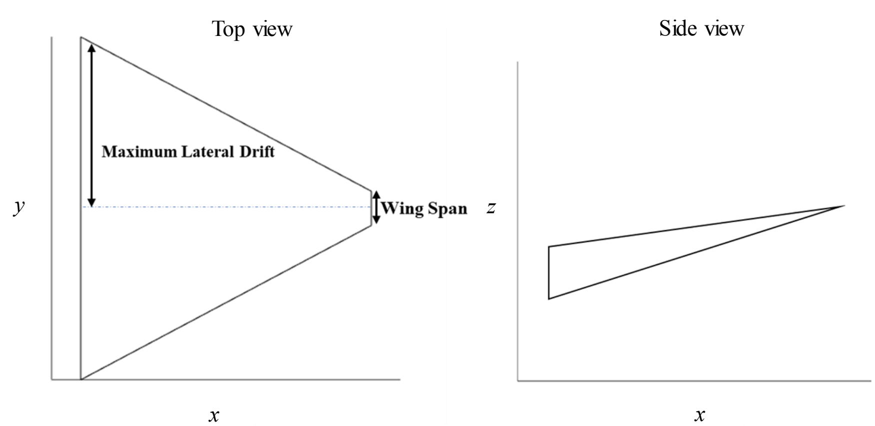

Figure 5: Wake region behind the leading aircraft in 3D model.

3 SPLITTING METHODOLOGY

The main challenge in using standard Monte Carlo simulation to estimate rare event probabilities is computation time. A typical approach is to simulate n i.i.d. replications and let $X_i = 1$ if the event occurs, and 0 otherwise. Then $\hat{\gamma} = \sum_{i=1}^{n} X_i / n$ is an unbiased estimator of the event probability γ, with $\text{var}[\hat{\gamma}] = \gamma(1 - \gamma) / n$. For a rare event, the relative error, which is the standard deviation of the estimator divided by its mean, is approximately $1/\sqrt{\gamma n}$. For example, if $\gamma = 10^{-9}$, then to achieve a relative error of 10% requires at least 10^{11} replications, which may be intractable, particularly if the simulation time for each replication is non-trivial. Thus, other methods are needed.

The main idea of the splitting technique (e.g., L'Ecuyer et al. 2009) is to split the simulation into separate independent runs when trajectories get "near" the rare event. This tends to focus computation effort on runs that are more likely to hit the rare event. One advantage is that the probability laws of the system remain unchanged and so the stochastic model capturing the system evolution can be developed independent of the splitting method.

We define our problem similar to the basic setting of splitting problems in the literature. We have a stochastic process $X \equiv \{X_t, t \geq 0\}$ with state space χ, where X_t is a strong Markov process. In our case, X_t consists of the positions and airspeeds of each aircraft (see state variables defined in the previous section). We define two disjoint subsets of χ, S and W, where W is the rare event set and S is a set of "safe" states considered as initial conditions to the simulation. In our case, the rare-event set W is the wake region. The set S is defined as a region that is behind and above the leading aircraft (Figure 6), where the lower right corner of the region is the target location of the trailing aircraft.

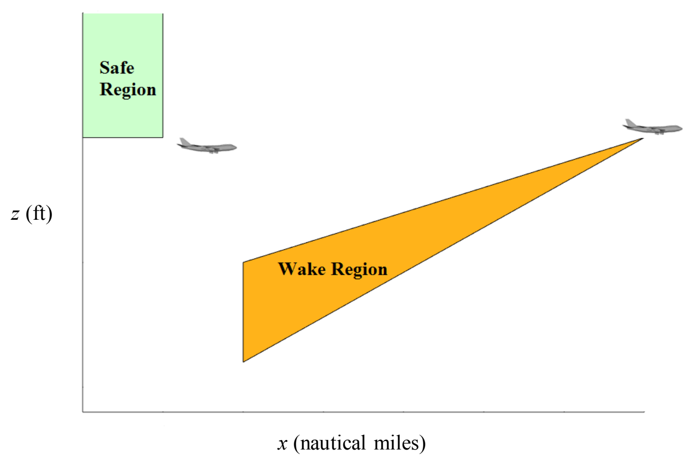

Figure 6: Side view of the wake region behind the leading aircraft.

The process starts at some state on the boundary of the safe set. It then leaves S and either eventually returns to S or reaches W. The goal is to estimate the probability γ that a trailing aircraft enters W prior to returning to S, starting from a point on the boundary of S. Since the process tends to revert to the lower right-hand corner of S, the termination time of the simulation is not an issue. The rare-event probability is

$$\gamma = \Pr\{X(t) \text{ reaches wake region before returning to safe region} \mid \text{trailing aircraft has just left safe region}\}.$$

To implement splitting, a set of intermediate level sets must be defined (as an example, see Figure 8). The simulation splits whenever it gets to a new level. We define a sequence of m levels $\{l_0, l_1, l_2, \ldots, l_m\}$ such that l_0 is the boundary of S and l_m is the boundary of W. Let D_k be the event that the trailing aircraft crosses level l_k before returning to the safe state. Let $p_1 = \Pr\{D_1\}$ and $p_k = \Pr\{D_k \mid D_{k-1}\}$ for $k = 2, \ldots, m$. That is, p_k is the probability of crossing level k given that level k-1 has been crossed. Since $D_M \subset D_{M-1} \subset \cdots \subset D_1$, the probability of reaching the rare event before returning to the safe state is $\gamma = \Pr\{D_m\} = \Pr\{D_1\}\Pr\{D_2|D_1\} \ldots \Pr\{D_m|D_{m-1}\} = p_1p_2\ldots p_m$.

There are several variations of splitting described in the literature. A brief review of these variations can be found in L'Ecuyer et al. (2009). We use a fixed-effort splitting method, where the number n of simulation runs from each stage are fixed and pre-selected. In the first stage, we start n independent runs from the initial states (i.e., the boundary of the safe set), where the initial states are randomly generated based on the steady state probability distribution for the relative position of the trailing aircraft. At time zero, the leading aircraft is assumed to be at a fixed point. At each time step we update the position and velocity of both aircraft and continue until the relative position reaches level 1 or returns to the safe set. If the run reaches level 1, the simulation is stopped and the end state is saved in a set called L_1. Let R_1 be the number of runs reaching the first level. The probability of reaching level 1 is estimated as $\widehat{p_1} = R_1/n.$ In stage $k \geq 2$, the initial states are generated by randomly sampling n states with replacement from the set L_{k-1}. Each run is simulated independently until it reaches level k or goes back to level k-2, in which case the run is truncated. This is done to reduce time simulating trajectories back to the safe set.

In doing this truncation, we have to correct the bias. This is done as follows. For each stage k, we randomly sample r_k chains from all the killed chains in the stage, where r_k is a pre-selected integer. For these selected chains, we continue simulating them until they return to the safe set or reach level k. Every chain that reaches level k gets a weight $W_k = M_k / r_k$, where M_k is the number of killed chains in that stage. For example, if we have $M_k = 100$ killed chains and we select $r_k = 5$ of these to simulate to completion, then each of these is representative of 100/5 = 20 killed chains. Therefore if a sampled chain reaches level k, it is cloned to $\lfloor W_k \rfloor$ or $\lfloor W_k \rfloor + 1$ copies (since the number of copies must an integer), with probabilities $\delta = W_k - \lfloor W_k \rfloor$ and $1 - \delta$ respectively. For $k > 2$, the probability of reaching level k form level k-1 is estimated as $(R_k + W_k S_k) / n$, where R_k is the number of chains that reach level k without down crossing level k-2 and S_k is the number of chains that reach level k after down-crossing level k-2.

4 TEST CASES AND RESULTS

4.1 Test Case

To test our methodology we consider a pair of aircraft in cruise phase. Using the NWRA AVOSS wake vortex prediction algorithm (Robins and Delisi 2002) we calculate approximate boundaries of our wake region. We choose a B737-800 aircraft as the leading aircraft, with a wing span of 34.32 m and a mass of 66.36 tons (maximum landing weight). The altitude is set at 35,000 feet with air density equal to 0.38 kg/m^3. The atmospheric parameters Eddy Dissipation Rate (EDR) and Brunt-Vaisala Frequency (BVF) are set to 10^{-7} m^2/s^3 and 0 per second, respectively. Figure 7 shows the predicted altitude and circulation of wake vortices using the AVOSS model with the selected parameters. The results are used as a rough reference point for setting the dimensions of the wake zone in Figure 4.

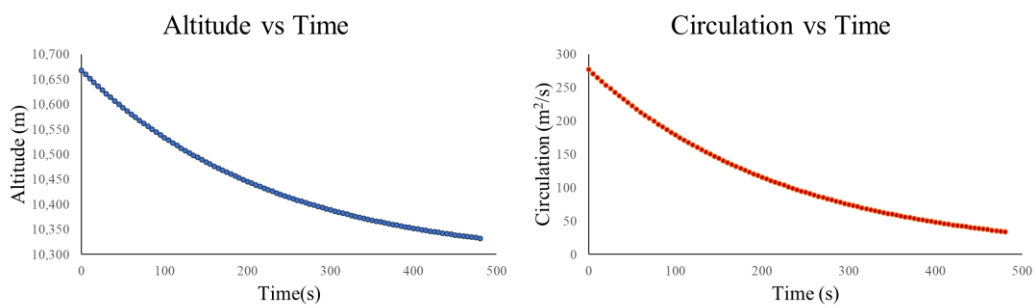

Figure 7 : Decay and transport of wake vortices in time.

The leading aircraft flies at a target speed of 450 knots. The following aircraft flies on the same path maintaining a two minute separation (which equates to 15 nm) with a standard deviation of 5 seconds (.625 nm). In the vertical dimension, the aircraft remain within 100 feet of the desired altitude at least 95% of the time.

Assuming a critical circulation threshold of 180 m^2/s for the back end of the wake zone equates to about 100 seconds in Figure 7. We set the length of the wake zone as 12.5 nm, which is the distance equivalent of 100 seconds, assuming a speed of 450 knots. Since the trailing aircraft is trying to maintain two-minute separation, this corresponds to a 20-second buffer beyond the back end of the wake region. The altitude chart in Figure 7 indicates that the vortices descend about 135 m (or about 450 feet) in the first 100 seconds and about 325 m in the first 500 seconds (which equates to a rate of about 130 feet in 100 seconds). Since the shape of the wake region is assumed to be triangular (Figure 4), we use the initial descent rate (450 feet per 100 seconds) multiplied by 100 seconds to set the maximum predicted descent and the average descent rate (130 feet per 100 seconds) to set the minimum predicted descent.

In simulating the motion of the two aircraft, we assume that the wake region is rigidly attached to the lead aircraft. That is, if the lead aircraft moves up, the *entire* wake region instantly shifts with it. In reality, only the portion of the zone very near the aircraft where the wake is generated would shift. Defining the wake region in this way is a simplification, but it eliminates the need for keeping a time history of the wake at each point. In future work, this assumption will be relaxed using more complex wake models, though we do not expect the main results to change much.

4.2 Splitting Schemes and Results

The first splitting scheme is inspired by the triangular shape of the wake region and uses nested triangles around the wake region as intermediate levels (Figure 8). Two variations are implemented. In the first variation, the intermediate levels are evenly spaced *in distance* from each other. In the second variation, the intermediate levels are spaced *in probability* – meaning that the probability of reaching level j from level j-1 is approximately the same for all j. It should be noted that the probability of reaching level j depends on the starting point in level j-1 and not all points on the contour of a given level have the same probability of reaching the next level.

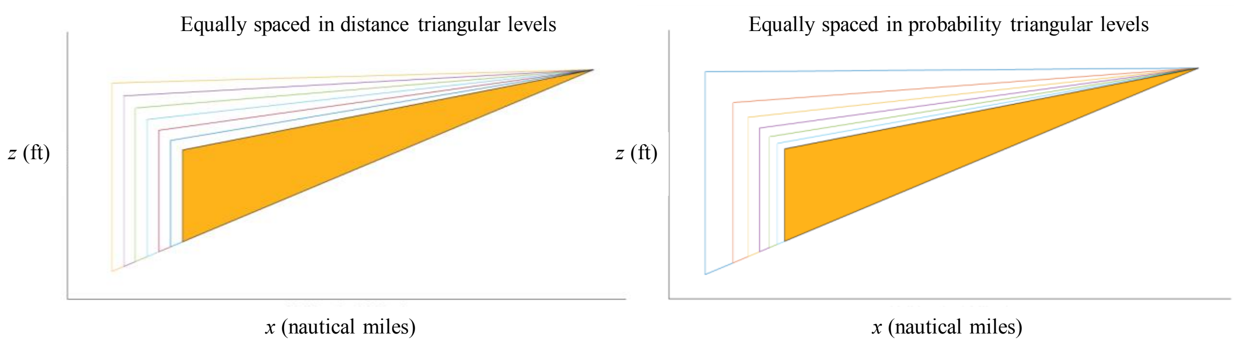

Figure 8: Geometry of intermediate levels in first splitting scheme.

Figure 9 shows the variance of the simulation estimator for different simulation experiments. The variance is proportional to the computation time required to achieve a given relative error, so an n-fold reduction in variance corresponds to an n-fold increase in simulation efficiency. Three different numbers of levels and two different level sets are used: 5, 7 and 10 levels and equal-distance and equal-probability levels. In each experiment, 100 replications are simulated with n = 2,000 runs simulated in each level for each replication. (For example, with 10 levels, each replication consists of 2,000 runs at each level for a total of 20,000 runs. With 5 levels, there are half as many runs. However, the levels are further apart, so it takes longer to simulate between one level and the next. So the total run time is similar in each case.) The levels that are evenly spaced in probability result in a smaller variance in comparison with levels that are evenly spaced in distance.

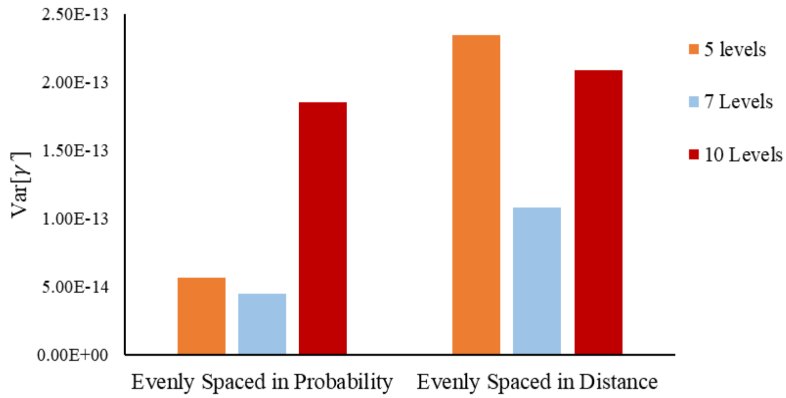

Figure 9 : Sample variance using splitting with triangular intermediate levels.

Figure 10 shows the level probabilities from the previous experiments – that is, the probability of reaching level *j* starting from level *j*-1 prior to returning to the safe set. The probability of reaching the first level is significantly smaller than the other level probabilities, even when we try to use levels that are evenly spaced in probability. This is a result of the difference in geometry between the boundary of the safe region where the simulation starts and the boundaries of the levels as shown in Figure 6.

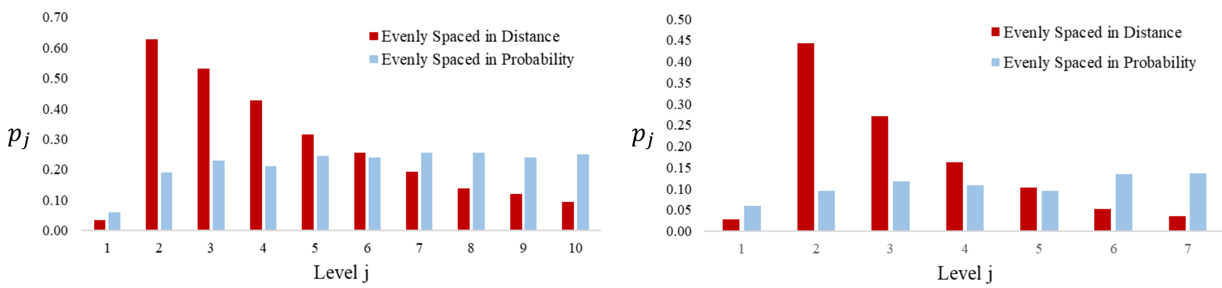

Figure 10: Level probabilities in 10-level and 7-level sets.

To solve this problem, we introduce a different scheme for intermediate levels as shown in Figure 11. The idea is to have levels with similar geometry near the safe region boundary and then transition to levels with similar geometry to the wake region closer to the rare event. Figure 12 compares the variances achieved by this new level set with the variances achieved by the previous level set (nested triangles). This new scheme offers smaller variance and in this experiment is also about 2 times faster than the previous schemes due to simpler calculations that are needed to check if the trailing aircraft reached the first three intermediate levels.

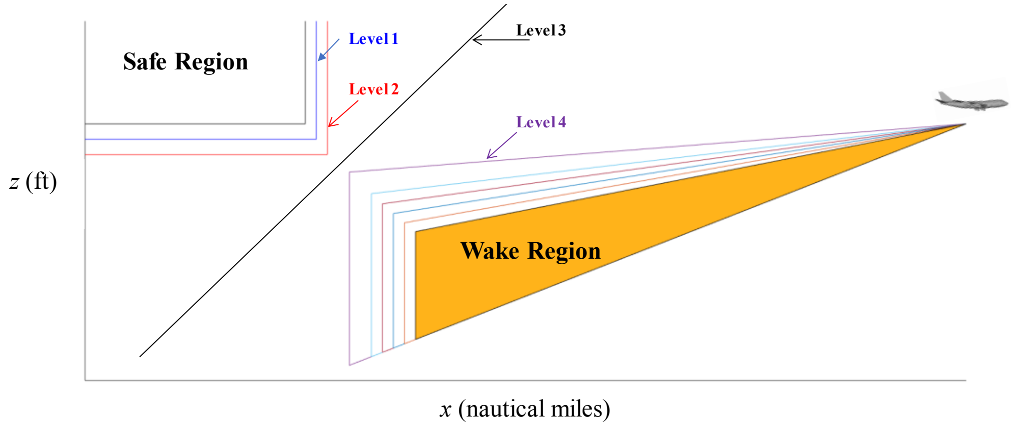

Figure11: New level sets for splitting method.

In addition, a number of standard Monte Carlo simulation experiments are performed with a similar computing budget. The standard simulation approach is significantly less efficient than splitting, since *no* potential wake encounters are observed in any of these experiments (so the estimated encounter probability is zero).

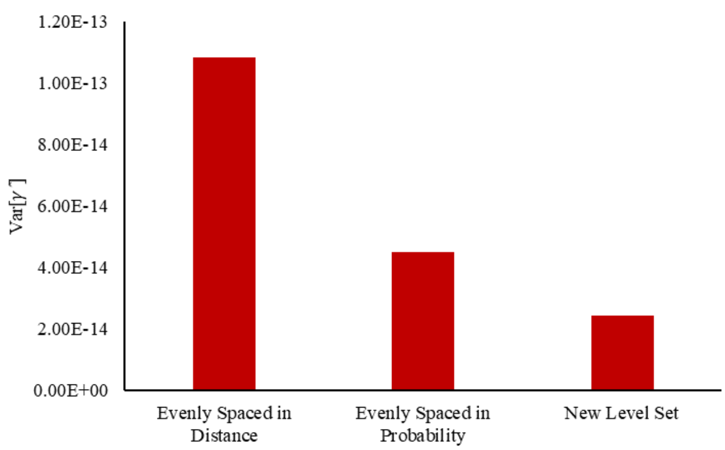

Figure 12 : Comparison of variances for 7 intermediate levels with different locations and geometries.

Experiments with a three-dimensional model are done in a similar manner to the first splitting scheme adding a lateral dimension to the trajectories of the aircrafts and the boundaries of the wake region. Since aircraft are trying to remain at the centerline of the path and perturbations are not extreme, the lateral (across-track) deviation of the aircraft is not large enough to help the trailing aircraft avoid the wake region (which is conservative in size). The lateral dimension is unique in that the aircraft is *trying* to flying through the middle of the zone in this dimension. In the along-track and vertical dimensions, the aircraft is trying to fly behind and above the zone. The probability of a potential encounter is similar to the two dimensional model.

5 CONCLUSIONS

This paper demonstrated a proof-of-concept for using a rare-event splitting technique in simulating potential wake encounters. This work can be used in evaluating risk in high density environments where encounters are more likely to occur. While standard Monte-Carlo simulation did not generate any hits of the rare event set (for the problems and computational budgets considered in this paper), the splitting method was able to generate results in reasonable time. Key decisions in implementing splitting are the choice of the level function and the locations of the levels. It was found that levels that are equally spaced in probability provide lower variance than levels that are equally spaced in distance. This is consistent with the standard theory of splitting (e.g., L'Ecuyer et al. 2009). A level function that attempted to mirror the shape of both the rare-event set and the safe set was also found to reduce the variance.

The wake model used in this approach was very simple – a geometric triangle. The rationale for this choice was to focus testing on the splitting method using a "first-order" approximation of the wake. Future work will integrate more complex wake models into this framework. For example, one of the authors has developed a stochastic wireframe implementation of the NASA APA models, called WakeWISE. Future work will also test dynamic wake separation concepts in which the separation requirement is not a static distance, but a distance that changes as a function of atmospheric parameters (e.g., crosswinds) and/or the state of the leading aircraft (speed and weight). More advanced variations of splitting dealing with hybrid stochastic systems can also be used to estimate encounter probabilities in which discrete system failures are considered.

REFERENCES

Blom, H. A. P., J. Krystul, G. J. Bakker, M. B. Klompstra, and B. Klein Obbink. 2006. "Free Flight Collision Risk Estimation by Sequential Monte Carlo Simulation". In *Stochastic Hybrid Systems; Recent Developments and Research Trends*, edited by C. G. Cassandras and J. Lygeros, 247–279. New York: Taylor & Francis.

Blom, H. A. P., G. J. Bakker, and J. Krystul. 2009. "Rare Event Estimation for a Large-Scale Stochastic Hybrid System with Air Traffic Application". In *Rare Event Simulation Using Monte Carlo Methods,* edited by G. Rubino and B. Tuffin, 193–214. West Sussex, United Kingdom: Wiley.

De Oliveira, I. R., L. F. Vismari, P. S. Cugnasca, J. B. Camargo, G. J. Bakker, and H. A. P. Blom. 2010. "A Case Study of Advanced Airborne Technology Impacting Air Traffic Management". In *Computational Models, Software Engineering and Advanced Technologies in Air Transportation, Next Generation Applications*, edited by L. Weigang et al., 177–214.

Garvels, M. J. J. 2011. "A Combined Splitting-Cross Entropy Method for Rare-Event Probability Estimation of Queueing Networks". *Annals of Operations Research* 189 (1):167–85.

Glover, W., and J. Lygeros. 2004. "A Stochastic Hybrid Model for Air Traffic Control Simulation". In *Hybrid Systems: Computation and Control*, edited by R. Alur and G. J. Pappas, 372–386. Berlin: Springer-Verlag.

Hoogstraten, M., H. G. Visser, D. Hart, V. Treve, and F. Rooseleer. 2015. "Improved Understanding of En Route Wake-Vortex Encounters". *Journal of Aircraft* 52:981–989.

Kos, J., H. A. P. Blom, L. J. P. Speijker, M. B. Klompstra, and G. J. Bakker. 2001. "Probabilistic Wake Vortex Induced Accident Risk Assessment". *Air Transportation Systems Engineering*, edited by G. L. Donohue and A. G. Zellweger, 513–531. Reston, VA: American Institute of Aeronautics and Astronautics.

L'Ecuyer, P., F. Le Gland, P. Lezaud, and B. Tuffin. 2009. "Splitting Techniques". In *Rare Event Simulation Using Monte Carlo Methods,* edited by G. Rubino and B. Tuffin, 39–61. West Sussex, United Kingdom: Wiley.

Murray, L., H. Cancela, and G. Rubino. 2013. "A Splitting Algorithm for Network Reliability Estimation". *IIE Transactions* 45(2):177–89.

Nelson, R. 2006. "Trailing Vortex Wake Encounters at Altitude - A Potential Flight Safety Issue?". AIAA Atmospheric Flight Mechanics Conference and Exhibit.

Robins, R., and D. Delisi. 2002. "NWRA AVOSS Wake Vortex Prediction Algorithm Version 3.1.1". National Aeronautics and Space Administration, Langley Research Center.

Rubino, G., and B. Tuffin. 2009. "Introduction to Rare Event Simulation". In *Rare Event Simulation Using Monte Carlo Methods,* edited by G. Rubino and B. Tuffin, 1–13. West Sussex, United Kingdom: Wiley.

Shortle, J. 2013. "Efficient Simulation of Blackout Probabilities Using Splitting". *Electrical Power and Energy Systems* 44:743–751.

Wang, Z., and J. Shortle. 2012. "Sensitivity Analysis of Potential Wake Encounters to Stochastic Flight-Track Parameters". In *Proceedings of the International Conference on Research in Air Transportation*, Berkeley, CA, 1–8.

AUTHOR BIOGRAPHIES

AZIN ZARE NOGHABI is a Ph.D. student in the department of Systems Engineering and Operations Research and works as a graduate research assistant in the Center for Air Transportation Systems Research at George Mason University. Her research interests are in air transportation safety and applications of rare event simulation in the aviation domain. Her email address is azarenog@gmu.edu.

JOHN SHORTLE is a professor of Systems Engineering and Operations Research at George Mason University and a member of the Center for Air Transportation Systems Research. His research interests include simulation, queueing, and analysis of rare events with applications in air transportation, energy, and telecommunications. He is the current president of the INFORMS Simulation Society. He received his Ph.D. in industrial engineering and operations research from U.C. Berkeley. His e-mail address is jshortle@gmu.edu.