AIAA Guidance, Navigation and Control Conference and Exhibit
18 - 21 August 2008, Honolulu, Hawaii

AIAA 2008-6630

Hazard Alerting using Line-of-Sight Rate*

Mykel J. Kochenderfer†, J. Daniel Griffith‡,
and James K. Kuchar§

Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA 02420

This paper presents an analysis of an electro-optical hazard alerting system based on intruder line-of-sight rate. We use a recently-developed airspace encounter model to analyze intruder line-of-sight rate behavior prior to near miss. We look at a simple hazard alerting system that alerts whenever the line-of-sight rate drops below some set threshold. Simulations demonstrate that such an approach, regardless of the chosen threshold, leads to frequent false alerts. We explain how the problem of hazard alerting can also be formulated as a partially observable Markov decision process (POMDP) and show how such an approach significantly decreases the false alert rate.

I. Introduction

NORTHROP Grumman and the Air Force are considering a variety of sensors onboard Global Hawk for a collision avoidance system (CAS). These include the Traffic alert and Collision Avoidance System (TCAS), automatic dependent surveillance-broadcast (ADS-B), electro-optical (EO) and infrared (IR) systems, and radar. TCAS and ADS-B provide a satisfactory means of sensing transponder-equipped aircraft but lack the ability to detect aircraft that are not equipped with a transponder. EO, IR and radar sensors are attractive solutions for detecting traffic because they do not require that intruders have special equipage. This paper evaluates the performance of a system that uses an EO sensor for hazard alerting, i.e. discriminating intruders on a collision course, using Monte Carlo simulation of thousands of encounters.

In this study, an intruder aircraft is defined as a hazard if a near-mid air collision (NMAC) will occur unless the unmanned aircraft performs an avoidance maneuver. An NMAC occurs when two aircraft come within 100 ft vertically and 500 ft horizontally of each other. One approach to hazard alerting using an EO system is to declare an intruder aircraft as a hazard if the inertial line-of-sight (LOS) angular rate is near zero.[1] For an autonomous system, once an intruder is declared a hazard, the CAS initiates some resolution action to mitigate the hazard. The current configuration proposed for Global Hawk performs a small self-maneuver to obtain a more accurate estimate of the intruder's range and state vector. The CAS then determines whether or not a collision avoidance maneuver is necessary based on the LOS rate and range to the intruder. In its simplest form, however, an EO-based hazard alerting approach requires a minimum LOS angular rate threshold—if the intruder LOS rate is below the threshold then then CAS logic declares the intruder as a hazard and initiates some resolution action. If the minimum LOS rate in the system's logic is too low, then the CAS may fail to identify and mitigate a real hazard. However, setting the minimum rate too high may cause the CAS to declare many non-hazardous aircraft as potential NMACs and the Global Hawk vehicle will unnecessarily deviate from its planned course or air traffic control assigned corridor, potentially endangering other surrounding aircraft.

This paper shows that building a robust hazard alerting system is not simply a matter of choosing a threshold LOS rate and declaring an intruder as a hazard if the rate is below that threshold. Simply setting a LOS rate threshold and alerting whenever the rate drops below this threshold results in many unnecessary alerts, even for very low thresholds. We propose formulating the problem of hazard alerting as a partially

*This work is sponsored by the Air Force under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

†Technical Staff, Surveillance Systems, 244 Wood Street.
‡Associate Staff, Surveillance Systems, 244 Wood Street.
§Assistant Group Leader, Surveillance Systems, 244 Wood Street. Senior Member AIAA.

observable Markov decision process (POMDP).[2] The advantage of such an approach is that it takes into account a reward function, an observation model, and an underlying state model. POMDPs take into account uncertainty in the state of the system, e.g. whether or not an intruder will result in an NMAC if no avoidance maneuver is taken.

We find that the POMDP approach significantly outperforms the simple threshold approach. If we are able to estimate range and range rate then we can further improve performance. Unfortunately, obtaining satisfactory estimates of range and range rate with an EO sensor is very difficult. Radar sensors can provide very accurate estimates of range and range rate but are typically more expensive than EO sensors and are too heavy for small unmanned aircraft.

The outline of this document is as follows: Section II briefly describes the encounter model that we use to simulate intruder behavior and calculate LOS rates for intruder aircraft. Section III analyzes the LOS rate distributions of intruder aircraft prior to near miss. This analysis includes a comparison with the LOS rate distribution of all non-hazardous intruder aircraft. We also present the system operating characteristic (SOC) curve[3,4] of a hazard alerting system that alerts based on a threshold LOS rate. Finally, Section IV describes a POMDP formulation for hazard alerting and presents simulation results.

II. Encounter Generation

The analysis provided in this document utilizes Lincoln Laboratory's Collision Avoidance System Safety AssessmenT Tool (CASSATT) to simulate encounters that are statistically representative of intruder aircraft behavior based on a visual-flight-rules (VFR) encounter model developed at Lincoln Laboratory.[5]

An encounter occurs when an intruder penetrates an encounter cylinder centered on the aircraft. The appropriate size for the encounter cylinder is determined by the aircraft dynamics and CAS. If the cylinder is too small, then the CAS does not have sufficient time to detect and track an intruder; however, if the cylinder is too large, then computation time is wasted.

We use two collections of 1 million encounters below 12,500 ft each as the basis for our analysis. One collection consists of encounters between pairs of VFR aircraft generated by the encounter model. Figure 1 shows histograms of the trajectory characteristics at the start of the encounter—during the course of an encounter, the turn rate, vertical rate, and airspeed acceleration is permitted to change according to the model. The other collection consists of encounters between Global Hawk and VFR aircraft. The Global Hawk trajectories are a mixture of four profiles, as shown in table 1. The first two profiles were extracted from radar data of an actual Global Hawk flight. The other two profiles were based on Global Hawk performance specifications. We used a larger encounter cylinder for encounters involving Global Hawk because the airspeed and vertical rates of Global Hawk are typically greater than VFR aircraft. Table 2 summarizes the characteristics of the two encounter sets. Of particular note is that the time of minimal cylindrical distance (TMCD)[a] for the encounter sets occurs at least 25 s after the start of the encounter, which is sufficient time for a CAS to sense and avoid an intruder.

Table 1. Profiles used to generate Global Hawk trajectories. The first two profiles in the table have average vertical rate, turn rate, and airspeed listed since they vary during the course of the trajectory. All other values in the table are constant.

Mixture	Vertical rate	Turn rate	Airspeed
25%	3392 ft/min	1.5 deg/s	191 kt
25%	-1279 ft/min	0.2 deg/s	145 kt
25%	3100 ft/min	0 deg/s	170 kt
25%	-1300 ft/min	0 deg/s	150 kt

III. Intruder Line-of-Sight Rate Analysis

One approach to intruder hazard alerting is to measure the LOS rate of the intruder aircraft. If the LOS rate is near zero, then the two aircraft are assumed to be on a collision course.[1] While it is commonly agreed that the LOS rate can be used to discriminate intruder aircraft based on whether or not they pose threat, even aircraft that pose no threat can produce low LOS rates briefly over the course of an encounter.

[a] Cylindrical distance is $\max(r_h/5, r_v)$, where r_h is horizontal range and r_v is vertical range.

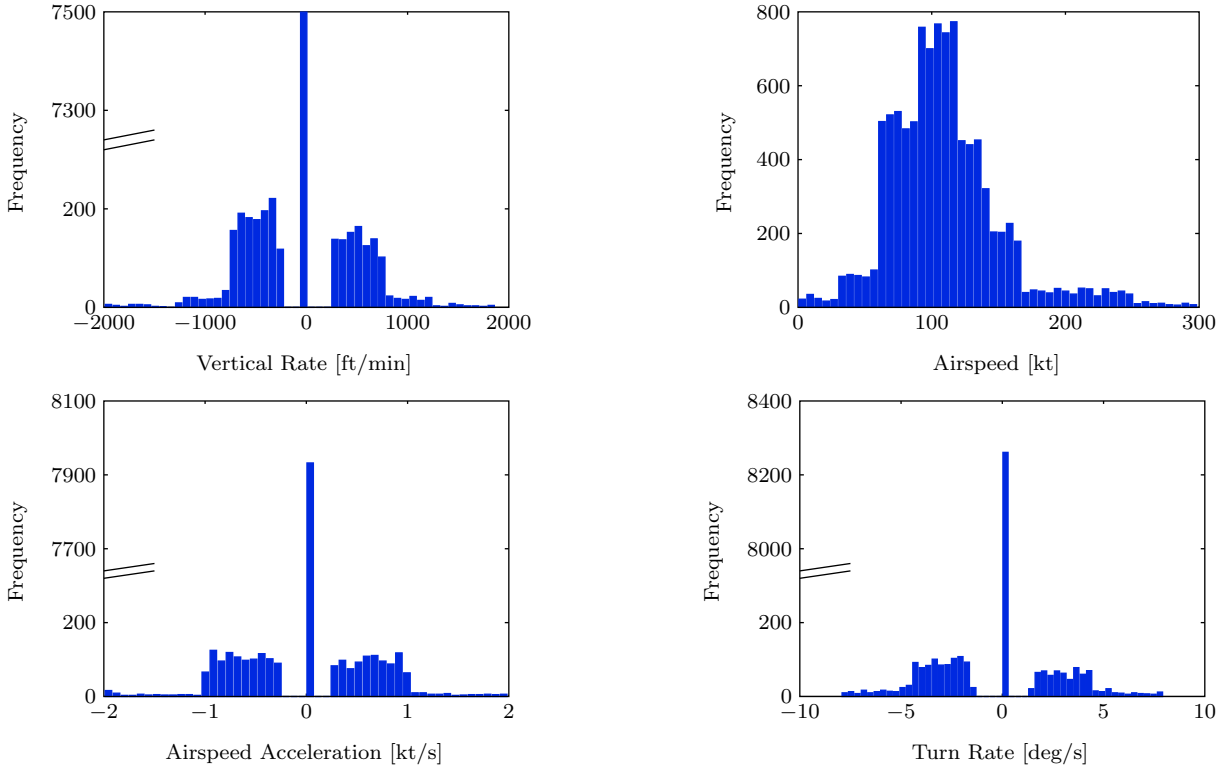

Figure 1. Histograms of VFR trajectory characteristics at the start of the encounter.

Table 2. Characteristics of the two collections of 1 million encounters used for our analysis.

	VFR/VFR	Global Hawk/VFR
Encounter cylinder radius	5 NM	5 NM
Encounter cylinder height	± 1500 ft	± 3300 ft
Minimum TMCD at start	27.9 s	30.4 s
Mean TMCD at start	166.4 s	89.4 s
NMACs	541	364

False alerts are unavoidable. Similarly, maneuvering intruder aircraft that result in an NMAC can have high LOS rates. In this section, we analyze the utility of LOS rate measurements for discriminating hazardous intruders. First, we compare the distribution of the LOS rate for NMAC-inducing intruders with non-NMAC intruders. We also evaluate the system operating characteristic (SOC) curve of a CAS system that declares an intruder a hazard if the LOS rate falls below some threshold rate. Finally, we compare this SOC with a hazard alerting system based on range and range rate measurements to assess the relative utility of LOS rate measurements versus range and range rate measurements for hazard alerting.

III.A. Distribution of Intruder Line-of-Sight Rate

The inertial LOS rate, $\omega(t)$, of an intruder aircraft is

$$\omega(t) = \frac{|(r_i(t) - r_o(t)) \times (v_i(t) - v_o(t))|}{|r_i(t) - r_o(t)|^2}, \tag{1}$$

where r_o and v_o are the three-dimensional position and velocity vectors for the own aircraft. Similarly, r_i and v_i are the vectors for the intruder aircraft.

We considered several families of distributions valid on the semi-infinite interval $[0, \infty)$ to characterize the LOS rates of intruder aircraft: Gamma, Weibel, and log normal. We compared the distributions using the negative log-likelihood and found that the log normal family provided the best match with simulation data.

The log normal is a family of distributions for a random variable whose logarithm is normally distributed. A log normal distribution is completely defined by the mean μ and standard distribution σ of the random variable's logarithm. The probability density function (pdf) and cumulative distribution function (cdf) for a log normal distribution are

$$f(x;\mu,\sigma) = \frac{e^{-(\ln x-\mu)^2/(2\sigma^2)}}{x\sigma\sqrt{2\pi}} \tag{2}$$

$$F(x;\mu,\sigma) = \frac{1}{2}+\frac{1}{2}\mathrm{erf}\left[\frac{\ln x-\mu}{\sigma\sqrt{2}}\right]. \tag{3}$$

Maximum likelihood estimation tells us which μ and σ best explains the data. We found that $\mu = -1.55$ and $\sigma = 1.14$ fits the distribution over LOS rate at 20 s prior to TMCD (figure 2).

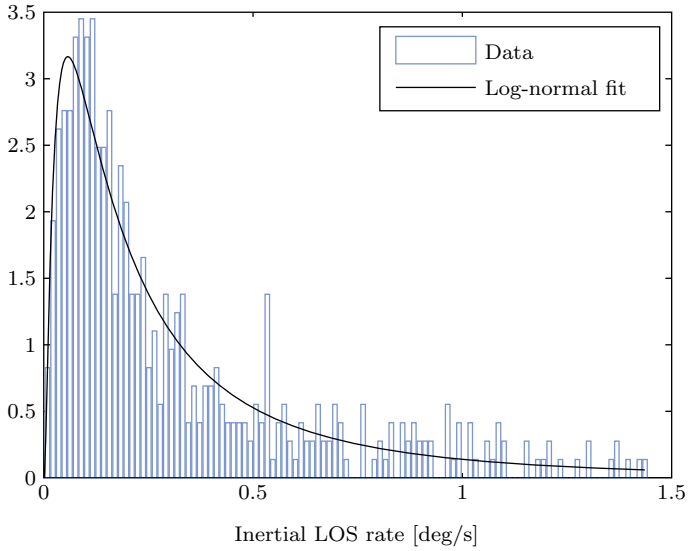

Figure 2. Probability density function of NMAC intruder LOS rates 20 s prior to TMCD. Shown is a log-normal distribution ($\mu = -1.55, \sigma = 1.14$) that best fits the data.

Figure 3 presents the LOS rate pdfs of NMAC-inducing intruders at 30 and 10 s prior to TMCD. These pdfs are compared to the LOS rate distribution of non-NMAC intruders at 30 and 10 s prior to TMCD in the VFR/VFR encounter set. The LOS rate distributions for NMAC-inducing and non-NMAC intruders are similar suggesting that given a single LOS rate measurement one cannot confidently determine whether or not the intruder will cause an NMAC. We observed similar results from the Global Hawk/VFR encounter set, which are not shown. Table 3 presents all relevant log normal distribution parameters for NMAC-inducing intruders.

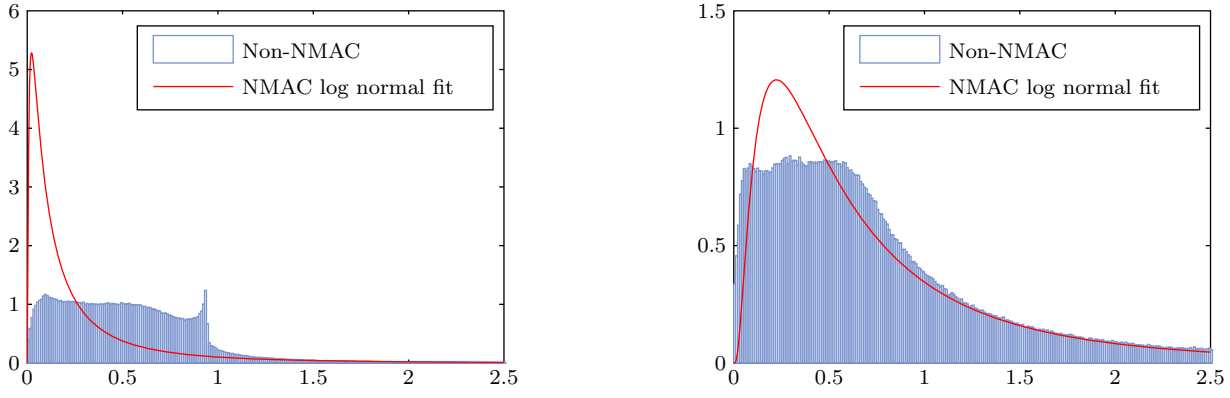

(a) Thirty seconds prior to TMCD.

(b) Ten seconds prior TMCD.

Figure 3. Probability density function of the LOS rate of NMAC-inducing intruders compared to non-NMAC intruders.

Table 3. LOS rate lognormal distribution parameters for NMAC-inducing intruders.

	VFR		Global Hawk	
Time until TMCD (s)	μ	σ	μ	σ
40	-2.193	1.445	-1.984	1.429
30	-1.977	1.343	-1.742	1.344
20	-1.552	1.142	-1.320	1.126
10	-0.603	0.945	-0.551	0.869

There are several factors that account for the lack of discernibility between NMAC and non-NMAC intruder LOS rates. First, non-NMAC intruders are generally farther away. The LOS rate of an intruder is smaller when the intruder is farther away since the LOS rate is inversely proportional to range. In addition, the hypothesis that a near-zero LOS rate can be used to determine that two aircraft are on a collision course is based on the assumption that neither aircraft is accelerating or turning at the beginning of an encounter. Only 69% of the VFR trajectories in our model are neither accelerating or turning. Thus, there is approximately a 48% chance that both aircraft in a single encounter are not maneuvering.

For comparison, we consider range and range rate measurements for discriminating NMAC-inducing intruder from non-NMAC intruders. If a CAS can measure both range and range rate, then the system may estimate time of closest approach (TCA) by dividing range by range rate. TCAS, for example, utilizes this metric to determine whether or not a resolution advisory is necessary.[6] In this analysis we consider TCA estimates

$$p(t) = \frac{r(t) \cdot \Delta t}{r(t - \Delta t) - r(t)}, \tag{4}$$

where Δt is the time-step between consecutive measurements. If $p(t)$ is positive, then the aircraft are closing together. If $p(t)$ is negative, then the range between the two aircraft is increasing. We observed that TCA measurements obtained by dividing range by range rate provide a good estimate of TMCD for NMAC-inducing intruders. Estimates are usually within a couple of seconds of the true TMCD. The distribution of TCA estimates for non-NMAC intruders, however, is approximately uniform up to several thousand seconds regardless of TMCD. The estimated TCA distributions between NMAC-inducing intruders and non-NMAC intruders are largely discernible from each other suggesting that range and range rate measurements provide more suitable information about the intruder aircraft for hazard alerting than LOS rate measurements.

III.B. Threshold System Operating Characteristics

Figure 4 shows the SOC curves for different hazard alerting systems. A point on the SOC plane is a plot between the probability of a successful alert (an alert at least 10 s prior to NMAC with an intruder) against the probability of a false alert (an alert for an intruder where no NMAC would occur). Each point corresponds to a hazard alert system with different alerting criteria. The ideal operating point is in the upper left hand corner. However, a system will rarely operate at this point; a hazard alerting designer can only hope to be as close as possible to this point.[3,4]

The solid blue curve in figure 4 corresponds to a system that uses a simple LOS rate threshold, while the dashed red curve corresponds to a system that estimates TCA. A correct alert occurs if the hazard alerting system alerts for an NMAC-inducing intruder at least 10 s prior to TMCD. We assume that if an alert is issued at least 10 s prior to TMCD, then the CAS selects and performs a successful collision avoidance maneuver. A false alert occurs if the hazard alerting system alerts for an encounter with an intruder that would not have resulted in an NMAC. These curves are generated by varying the threshold values and calculating the corresponding correct alert and false alert rates. We assume that the measurements are errorless and that there are no field of view limitations.

Because brief dips in LOS rate measurements might be due to an intruder turning in the field of view, we experimented with alerting systems that alert only if the LOS rate drops below some threshold for longer than some set amount of time. Figure 4 shows the SOC curves for systems that alert when the LOS rate drops below some varying thresholds for longer than 3 seconds and 6 seconds. Requiring the LOS rate to drop below some threshold for a minimum amount of time will help eliminate some false alerts but will

decrease the correct alert rate as well. The SOC curves show that alerting after 3 and 6 seconds typically impairs performance.

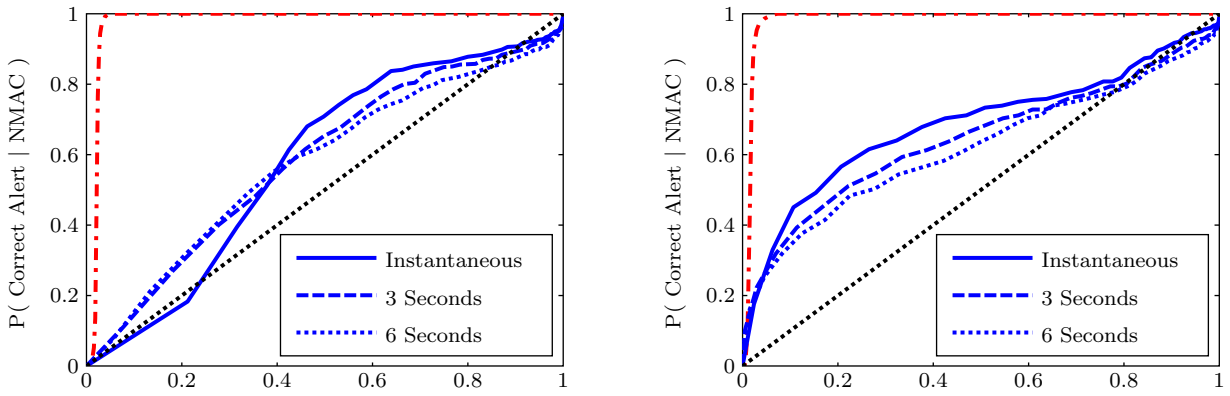

(a) VFR pair encounter set SOC.

(b) Global Hawk/VFR encounter set SOC.

Figure 4. SOC for hazard alerting systems based on LOS rate measurements (in blue) and TCA estimates (in dashed red). The hazard alerting systems based on LOS rate measurements were set to alert instantaneously when dropping below a threshold (sold), after 3 seconds of staying below a threshold (long dash), and after 6 seconds of staying below a threshold (short dash). Points in the top left indicate better performance.

These curves suggest that a hazard alerting system that alerts when the intruder LOS rate falls below some threshold is not effective at discriminating NMAC intruders from non-NMAC intruders. A conservative alerting system might require that the correct alert rate is at least 0.95, if not higher. This operating point results in a similar false alert rate. In contrast, the false alert rate for a hazard alert system based on TCA estimates is 0.02. The false alert rate is reduced by more than an order of magnitude compared to a hazard alerting system based on LOS rate measurement. It is important to note that some false alerts may be acceptable, if not desirable. If a hazard alerting system alerts for an intruder that is only a couple feet away from being an NMAC, then this is considered a false alert in figure 4. In practice, however, alerting for this example situation may be considered a successful alert.

These findings suggest that a hazard alerting system that only uses a minimum LOS rate threshold for hazard alerting is not effective. In order to achieve an acceptable missed alert rate, such a system system would need to alert for nearly every intruder. This observation suggests that a better hazard alerting system is necessary—one that allows for the LOS to briefly fall to near zero and not result in an alert. A system that tracks the intruder and estimates some underlying state to decide whether or not to alert for the intruder will likely improve performance. It has been demonstrated recently that a gentle maneuver can improve the range estimate to an intruder when only LOS measurements are available.[1] An even better hazard alerting system may fuse together LOS rate, range, and range rate measurements to improve hazard alerting performance.

IV. Hazard Alerting Using POMDPs

POMDPs have been used to model a wide variety of problems.[7] Earlier work at MIT[8] investigated the use of POMDPs for hazard alerting when the horizontal and vertical distances to the intruder are known exactly at each time step. Formulating the problem of hazard alerting based on LOS rate as a POMDP has several advantages. One advantage is that we are able to explicitly balance the cost of unnecessary alerts with NMACs by defining a reward function. Another advantage is that we can dynamically update our uncertainty about the underlying state (i.e. whether an NMAC is imminent) based on observations (i.e. LOS rate) using Bayes' rule. POMDP solution methods tell us exactly how to behave in order to maximize our expected reward. This section explains how we formulate the problem of hazard alerting based on LOS rate as a POMDP and presents results.

IV.A. POMDP Formulation

A variety of POMDP formulations exist for a hazard alerting system. Our POMDP formulation is intentionally simple in order to demonstrate how one may implement such a hazard alerting system and demonstrate the benefits of using a POMDP approach. Our formulation uses only one state variable, but more sophisticated formulations may include additional variables.

Our POMDP formulation is as follows:

State Space: The state space consists of the closed interval $[0, \tau]$. For convenience, we discretize time at 1 Hz making $S = \{0, \ldots, \tau\}$. Being in state τ means that the aircraft is at least time τ away from a collision. Upon reaching state 0, the process terminates.

State Transition Model: The state transition model $P(s_{i+1} \mid s_i)$ is represented in figure 5 and follows these two rules:

- $P(\tau \mid \tau) + P(40 \mid \tau) = 1$
- $P(s - 1 \mid s) = 1$ for all $s < \tau$

Observation Space: The observation space is the space of possible LOS rates. We discretize LOS rate into 35 bins. Between 0 and 2.04 deg/s are 34 bins of equal width, and the last bin consists of all LOS rates above 2.04 deg/s.

Observation Model: The observation model $P(o \mid s)$ is estimated from simulation data.

Action Space: The action space consists of $\{\text{alert}, \neg\text{alert}\}$. Once an alert is issued, the process terminates.

Reward: Reward is only issued at termination. Alerting in state τ is given reward $-c_{\text{false}}$. Reaching state 9 results in reward $-c_{\text{collision}}$. The values for c_{false} and c_{collide} are chosen to balance false alerts with risk of collision. We experiment with different values for these two variables in Section IV.B.

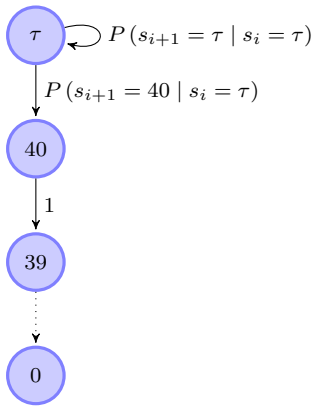

Figure 5. Transition state model for POMDP. The state s_i corresponds to the time to TCA.

The state transition and observation models are estimated from simulated data between two VFR aircraft. We gathered data at 1 Hz from each aircraft in the 1 million encounters described in table 2. After each encounter we first determine the state at each data point depending on if and when an NMAC occurs. We estimated $P(40 \mid \tau)$ by dividing the number of times an encounter was in state 40 by the number of times an encounter was in state τ. Next we calculated $P(o \mid \tau)$ from both tracks in any encounter that does not results in an NMAC and time greater than 40 s until TMCD for encounters that do result in an NMAC. Each track from the 541 NMAC encounters is used to estimate $P(o \mid s)$ where $s < \tau$.

Since we cannot directly observe the state, we are uncertain about the underlying state with the available observations. However, by using Bayes' rule, we can update our beliefs about the state with each observation. Our beliefs are represented by the function b, where $b(s)$ is the degree of belief that we are in state s. Our

beliefs are updated by the following rule

$$\begin{aligned} b(s_{i+1}) &\leftarrow P\left(s_{i+1} \mid o, b\right) \\ &\leftarrow P\left(o \mid s_{i+1}, b\right) P\left(s_{i+1} \mid b\right) / P\left(o \mid b\right) \\ &\leftarrow \kappa P\left(o \mid s_{i+1}\right) \sum_{s} P\left(s_{i+1} \mid b, s_i\right) P\left(s_i \mid b\right) \\ &\leftarrow \kappa P\left(o \mid s_{i+1}\right) \sum_{s} P\left(s_{i+1} \mid s_i\right) b(s) \end{aligned}$$

where κ is a normalized constant equal to $1/\sum_s b(s)$. With each time step, we update our distribution over states. Initially $b(\tau) = 1$, corresponding to the belief that we begin in a state where there is no threat of collision within time τ.

We solved the POMDPs using Heuristic Search Value Iteration (HSVI)[9] using the ZMDP software package available at http://www.cs.cmu.edu/~trey/zmdp. We used a discount factor of 0.99. The regret bound for the approximately optimal policies found by ZMDP were lower than 0.03. The policies tell us which action to take, i.e. whether to alert or not, given our distribution over states b.

IV.B. POMDP Hazard Alerting Results

Solving and simulating the POMDP hazard alerting system with various reward functions allows us to estimate the SOC for various POMDP formulations. We use all combinations of $c_{\text{false}} = 0.01, 0.1, 1, 10, 100, 1000$ and $c_{\text{collide}} = 10, 100, 1000$ that provide a wide range of operating points to generate the SOC. In order to prevent overfitting, we use one set of encounters to estimate the observation and state transition models of the POMDP and a completely separate set of encounters to estimate the SOC curve.

Figure 6 presents the SOC of the POMDP hazard alerting system. The POMDP hazard alerting system outperforms a system that simply alerts when the LOS rate falls below a particular threshold value. The reason POMDP hazard alerting performs better is that it takes the entire history of LOS rates into account—implicitly through the belief function b—so that brief dips in the LOS rate do not necessarily result in alerts. The shaded region in figure 6 represents the improved performance that we observe with the range of operating points we tested. Further tuning of the reward function may increase the area of the region between the two operating curves. There are significant gains when the hazard alerting problem based on

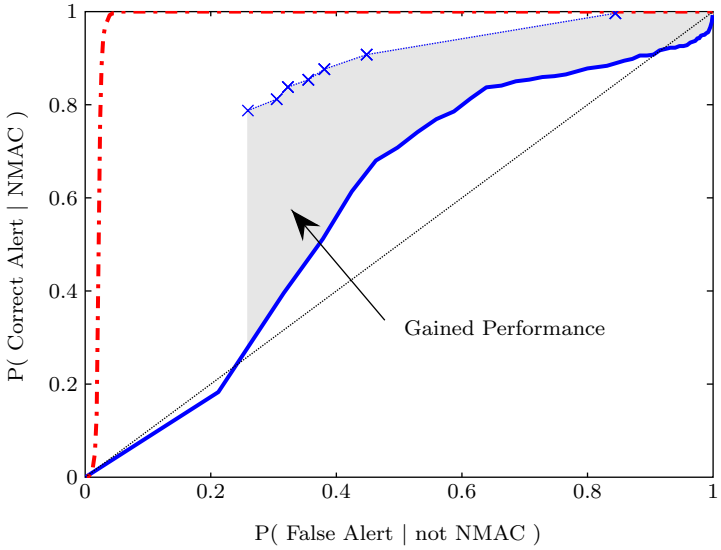

Figure 6. SOC for the POMDP hazard alerting system. Each blue "x" denotes a single operating point. The grey region represents the improved performance between the POMDP hazard alerting system relative to a simple system that alerts if an intruder's LOS rate falls below a threshold value.

LOS measurements is posed as a POMDP solution. For example, when $P(\text{Correct Alert} \mid \text{NMAC}) = 0.91$ for each system, the false alert rate for the basic alerting system is 0.90 while the false alert rate for the POMDP hazard alerting system is 0.45—the false alert rate is cut in half.

The POMDP hazard alerting system with only LOS rate measurements, however, still does not perform as well as a system based on an estimate of TCA. This observation reinforces the importance of accurately measuring range and range rate for hazard alerting. Hazard alerting based on intruder LOS rate is difficult because non-threatening intruders are farther away and many intruders are either accelerating or turning. In addition, the LOS rates for hazardous intruders are often not near zero. Similar difficulties were noted by Shakernia and Chen[1] who implemented a technique to estimate range to an intruder with an EO system by performing a small self-maneuver. The range estimation error in their Kalman filter was noticeably worse when the intruder aircraft was maneuvering.

V. Conclusions and Future Work

This paper analyzed the utility of line-of-sight (LOS) rate measurements for hazard alerting based on a new encounter model developed at Lincoln Laboratory. Our analysis first compared the distributions of near-miss intruder LOS rates prior to their closest approach against the rates of non-hazardous intruder aircraft. These distributions demonstrated that simply using LOS rate measurements for hazard alerting is ineffective. In order to achieve an acceptable correct alert rate (more than 0.95) nearly every intruder is alerted for when a hazard alerting system is based on a threshold LOS rate value. While our partially observable Markov decision process hazard alerting system based on LOS rate measurements demonstrated improved performance, it is still outperformed by a system that alerts based on estimated time of closest approach. A system that uses range and range rate measurements is far more effective at hazard alerting—hazard alerting is simply difficult without accurate range estimates. Our results are restricted to encounters with traffic flying under visual flight rules; conclusions for instrument encounters may be different because instrument flight rules aircraft generally do not accelerate or turn as frequently.

The results presented in this paper highlight the challenges of using LOS rate measurements for hazard alerting. Our partially observable Markov decision process approach to hazard alerting with LOS rate measurements, however, showed promising results that indicate further research may continue to close the current gap between hazard alerting with LOS rate measurements and hazard alerting with range and range rate measurements. Further work will involve applying POMDP solution techniques to finding policies that will maneuver to avoid hazards, not simply alert when a hazard is present.

Acknowledgements

This report is the result of research and development sponsored by the the United States Air Force 303rd Aeronautical Systems Wing (303 AESW/XRX), and the Department of Defense Unmanned Aircraft System Airspace Integration Joint Integrated Product Team.

References

[1]Shakernia, O., Chen, W.-Z., and Raska, M. V. M., "Passive Ranging for UAV Sense and Avoid Applications," *Infotech@Aerospace*, American Institute of Aeronautics and Astronautics, Arlington, VA, September 2005.

[2]Kaelbling, L. P., Littman, M. L., and Cassandra, A. R., "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, Vol. 101, 1998, pp. 99–134.

[3]Kuchar, J. K., "Methodology for Alerting-System Performance Evaluation," *Journal of Guidance, Control, and Dynamics*, Vol. 19, No. 2, 1996, pp. 438–444.

[4]Winder, L. F. and Kuchar, J. K., "Evaluation of Collision Avoidance Maneuvers for Parallel Approach," *Journal of Guidance, Control, and Dynamics*, Vol. 22, No. 6, 1999, pp. 801–807.

[5]Kochenderfer, M. J., Espindle, L. P., Kuchar, J. K., and Griffith, J. D., "A Bayesian Approach to Aircraft Encounter Modeling," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, 2008.

[6]FAA, "Introduction to TCAS II Version 7," November 2000.

[7]Cassandra, A. R., "A survey of POMDP applications," *Working Notes: AAAI Fall Symposium on Planning with Partially Observable Markov Decision Processes*, edited by M. Littman, AAAI, Oct. 1998, pp. 17–24.

[8]Winder, L. F., *Hazard Avoidance Alerting with Markov Decision Processes*, Ph.D. thesis, Massachusetts Institute of Technology, 2004.

[9]Smith, T. and Simmons, R. G., "Point-Based POMDP Algorithms: Improved Analysis and Implementation," *Proceedings of the International Conference on Uncertainty in Artificial Intelligence*, AUAI Press, 2005, pp. 542–547.