JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 20, No. 4, July–August 1997

Prototype Conflict Alerting System for Free Flight

Lee C. Yang* and James K. Kuchar†
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

The development of a prototype alerting system for a conceptual free flight environment is discussed. The alerting logic is based on a probabilistic model of aircraft sensor and trajectory uncertainties that need not be Gaussian distributions. Monte Carlo simulations are used over a range of encounter situations to estimate conflict probability as a function of intruder position, heading, and speed, as determined through a datalink between aircraft. Additionally, the probability of conflict along potential avoidance trajectories is used to indicate whether adequate space is available to resolve a conflict. Intruder intent information, e.g., flight plan, is not included in the model but could be used to reduce the uncertainty in the projected trajectory. Four alert stages are defined based on the probability of conflict and on the avoidance maneuvers that are available to the flight crew. Preliminary results from numerical evaluations and from a piloted simulator study at NASA Ames Research Center are summarized.

Introduction

FUTURE air traffic management concepts such as free flight have been proposed to provide a means by which traffic flow efficiency can be increased.[1] Under free flight, current methods of traffic separation through the use of a rigid airway structure and in-trail spacing would be relaxed. Consequently, aircraft would have more flexibility to follow arbitrary routes in response to changing conditions. To compensate for the loss of airway structure, automated conflict detection and resolution tools would be required to aid pilots and/or ground controllers in ensuring traffic separation.

Because flow efficiency is a driver for free flight, it is desirable that conflicts be resolved using minor course, speed, or altitude changes well before emergency avoidance maneuvers are needed. It is also desirable, given the large number of aircraft in the air, that conflict alerts are only generated when necessary. However, the large amount of uncertainty in the free flight environment makes it difficult to determine how likely a projected conflict is to occur. The result is a tradeoff between alerting early to provide a large safety margin (and also producing unnecessary alerts) vs alerting late to reduce unnecessary alerts (but requiring more aggressive avoidance maneuvers).

Traditionally, alerting systems have been designed through an iterative, evolutionary process.[2–4] After defining alerting thresholds, the performance of the system (in terms of the protection it provides and the unnecessary alert rate) is typically evaluated through simulations of traffic encounters. If collisions or excessive unnecessary alerts occur, the alerting thresholds are modified to improve performance; thus, the performance tradeoffs are generally examined post hoc.

The tradeoff between safety and unnecessary alerts is well known in signal detection problems and alerting systems.[5,6] One recent approach to view the tradeoff is the system operating characteristic (SOC) curve.[7,8] The SOC curve explicitly shows the expected safety level and unnecessary alert rate as a function of the alert threshold setting. The shape of the SOC curve depends on sensor accuracy, uncertainties in the future flight paths of the aircraft, and human performance. Thus, changes in sensors or avoidance strategies can be evaluated by examining their impact on the shape of the SOC curve.

A novel approach is presented to alerting system design in which the performance tradeoffs are directly addressed in order to select alerting thresholds. SOC curves are used to aid in threshold placement, reducing the need for iterative modifications to improve performance. A prototype alerting system was developed using Monte Carlo simulations to assess the probability of a conflict over a range of free-flight traffic encounters. The logic was then exercised in a set of piloted free-flight simulation studies at the NASA Ames Research Center in the fall of 1996. This study examined enroute conflicts and acted as a testbed for the alerting logic presented here.

Received Nov. 1, 1996; presented as Paper 97-0220 at the AIAA 35th Aerospace Sciences Meeting, Reno, NV, Jan. 6–9, 1997; revision received Feb. 21, 1997; accepted for publication Feb. 25, 1997.

*Graduate Research Assistant, Department of Aeronautics and Astronautics, Building 37-435, 77 Massachusetts Avenue. Student Member AIAA.

†Assistant Professor, Department of Aeronautics and Astronautics, Building 33-117, 77 Massachusetts Avenue. Member AIAA.

Background

The use of probability estimation has been explored in conflict analysis before.[7–10] In previous work, Paielli and Erzberger[9] developed a viable analytical solution to determine the probability of a conflict for two aircraft maintaining a straight-ahead course. Their approach used Gaussian uncertainties to model along- and crosstrack error and can be rapidly solved and implemented in real time. If more complex (non-Gaussian) uncertainties (such as aircraft changing course or pilot reaction latency) are modeled, it becomes increasingly difficult to obtain an explicit analytical solution; thus, the use of the Monte Carlo method in this paper. It is not to say that introducing extra parameters leads to more accurate results, but instead, it demonstrates that probability analysis is not restricted to one or two density distributions. Because it is based on Monte Carlo simulation, the approach presented here is more flexible but requires significantly more processing time than the approach in Ref. 9. Uncertainty parameters used in Ref. 9 are incorporated into the prototype logic presented here, and both approaches would produce identical results if aircraft were assumed to continue along a straight, level course.

The work in this paper, however, goes farther than just describing how the calculations were made to determine the probability values. Part of the ongoing research is to determine how this knowledge can be utilized effectively in an actual alerting system. Previous work by Kuchar[7,8] recognized the need to include escape maneuvers as part of the probabilistic alerting analysis. This is necessary to ensure that sufficient maneuvering space is still available to avoid the conflict if required. The avoidance trajectories are evaluated probabilistically to account for uncertainty in flight crew response time and variability in their actions. The most favorable avoidance trajectories can then be explicitly determined based on the likelihood of safely maneuvering out of the conflict.

Methodology

The design of the prototype alerting system was guided in part by NASA requirements for their experiment. Though the concepts developed can be extended to a ground-based system to aid air traffic control (ATC), the prototype logic was tailored to an airborne system where conflicts are primarily resolved on the flight deck. The structure of free flight has not yet been definitively established, but it is presumed that state information from surrounding aircraft (position, speed, and heading) is available through interaircraft datalink

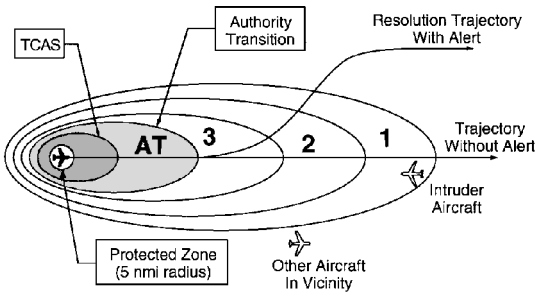

Fig. 1 Multistage alerting concept.

communication such as automatic dependent surveillance-broadcast (ADS-B). The alerting system must, therefore, provide ample warning time so that strategic maneuvers can be examined and coordination between flight crews can be carried out.

To simplify its development, the alerting system described here was designed for one-on-one conflicts during enroute flight. A more complete, operational system would have to be additionally evaluated for its ability to resolve conflicts between more than two aircraft. The aircraft with the alerting system is referred to as the host aircraft; the other aircraft involved in the conflict is termed the intruder. A conflict is defined as a situation in which the intruder enters a protected zone around the host aircraft. Based on current separation standards, the protected zone was defined to be a cylinder 5 n miles in radius and extending 1000 ft above and below the host aircraft.

A multistaged threshold approach was used to provide a series of alerts to indicate trends in conflict hazard. The multistage approach allowed the means of implementing the alert to be tailored to the level of threat. Low-probability threats resulted in relatively passive alerts such as changing the color of a traffic symbol. High-probability, urgent threats produced aural warnings to actively inform the pilots of the conflict.

Figure 1 shows a schematic diagram of the multistage approach. Three stages (marked 1, 2, and 3 in Fig. 1) produced changes in traffic display symbology in the cockpit of the host aircraft. As implemented, the outermost threshold provided a strategic indication of potential threat more than 10 min into the future and up to 200 n miles away. In the NASA 747-400 simulator, a hollow traffic symbol on the map display changed color when the first threshold was exceeded, and the flight crew could begin to coordinate resolution with the other aircraft. If the encounter continued, an additional stage informed the flight crew of the heightening conflict by filling in the traffic symbol. At the third stage (3), an aural alert zone transgression message was provided to the flight crew, indicating that they should take action to resolve the conflict. At this point, there was still ample time to coordinate resolution with other aircraft. If the conflict continued without resolution, an air traffic controller took over authority for conflict resolution at the authority transition (AT) zone.

The current traffic alert and collision avoidance system (TCAS) logic was not modified and was kept in the simulation as an independent, final warning system. But, because the alerting thresholds on the present TCAS 6.04A version are based on limited variables (range and closure rate), TCAS cannot accurately predict whether a conflict will occur beyond a few minutes. TCAS can track traffic within a range of 40 n miles, and its earliest alert can be triggered approximately 1 min before the estimated closest point of approach.[4,11] A new version of TCAS (V7.0) is expected to be available in 1998 and will increase the range to 100 n miles using ADS-B via mode-S to transmit additional position, heading, and vertical speed information.[12] The prototype alerting system can also be expected to utilize the same aircraft state data to estimate future trajectories.

Aircraft Trajectory Model

To determine the probability of conflict, a baseline model of aircraft trajectories was developed. Figure 2 is a pictorial representation of an aircraft in a free-flight environment. The modeled parameters include uncertainty in the current position estimate,

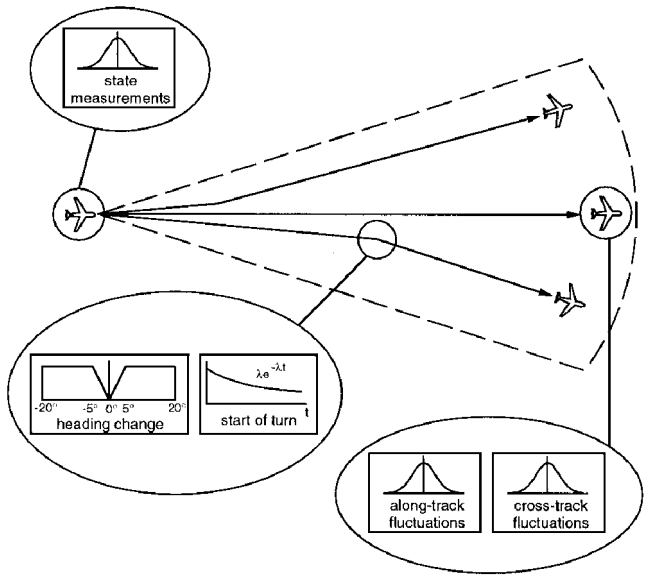

Fig. 2 Probabilistic trajectory model.

future along- and crosstrack position variability, and the potential for and magnitude of course changes.

Figure 3 summarizes the uncertainty parameters used in the baseline trajectory model. Uncertainty in current position results from the accuracy of combined global positioning system and inertial navigation system estimates and is modeled as a normally distributed random variable with standard deviation of 50 m laterally and 30 m vertically. Course drift in the future trajectory is modeled as a 15-kn standard deviation speed fluctuation (along-track error) and a 1 n mile standard deviation crosstrack error. These tracking error values are based on data obtained empirically from traffic by Paielli and Erzberger.[9]

It is assumed that only the current state of the intruder (position, heading, speed) is known on the host aircraft through ADS-B. The intruder's flight plan or target states, e.g., from the flight management system (FMS) or autopilot, are not available to the host aircraft. Accordingly, there is some uncertainty as to whether the intruder will maintain its current course. As an initial estimate, the likelihood of an intruder heading change is modeled as an exponential distribution with a mean rate λ of four turns per hour. This follows from the assumption that heading changes occur in a Poisson manner. When a heading change is made, its magnitude is modeled probabilistically as well. The intruder is equally likely to make heading changes left or right between 5 and 20 deg and is less likely to make turns of less than 5 deg (see Fig. 3).

Similarly, the intruder may change altitude. Altitude changes are also modeled as an exponential distribution with a mean likelihood of four occurrences per hour. When a change in altitude occurs, the intruder is equally likely to climb or descend to any altitude within 10,000 ft of its current altitude.

The host aircraft is assumed to fly a straight trajectory except for the along- and crosstrack described variations. However, to select between alternative conflict resolution options, it is important to also evaluate the reduction in conflict probability that can be achieved if the host aircraft maneuvers. Accordingly, a model of host aircraft resolution maneuvers was also developed. In reaction to a conflict alert, the flight crew response latency is modeled as a probabilistic gamma distribution with a mean of 1 min and a variance such that there is a 95% probability that the response occurs within 2 min. The relatively long latency is intentionally designed to allow time for coordination with other aircraft and/or ATC. Thus, avoidance maneuvers are assumed to have a large time buffer built in. Once initiated, avoidance maneuvers could include turns, altitude changes, or speed changes.

Note that the values of the parameters used in the trajectory model are estimates at this point and are not expected to be completely representative of free flight. Because free flight does not currently exist, it is difficult to predict the probabilistic nature of aircraft trajectories.

	Uncertainty Parameter	Modeled Distribution
Host Aircraft & Intruder Aircraft	Lateral Position Error	Gaussian, σ = 50 m
	Vertical Position Error	Gaussian, σ = 30 m
	Speed Fluctuation (Along-Track Variability)	Gaussian, σ = 15 kts
	Cross-Track Variability	Gaussian, σ = 1 nmi
Intruder Aircraft Only	Heading Change	$4e^{-4t}$; 0, t (hrs); time; -20°, -5°, 0°, 5°, 20°; magnitude
	Altitude Change	$4e^{-4t}$; 0, t (hrs); time; 0, 10,000 ft; magnitude
Host Aircraft Only	Avoidance Response Latency	Gamma; 95%; mean = 1 min.; 0, 1, 2, 3; minutes

Fig. 3 Trajectory model parameters.

However, such a prediction is necessary to estimate the likelihood of conflicts. Even if the values of the parameters are unknown, the impact of changes in the parameters can be evaluated to determine their relative importance. This in turn will help focus future efforts on improving trajectory estimation. For example, the heading and altitude changes in the intruder model can be modified to account for intent information relayed from the intruder's FMS computer. Information such as the next waypoint and the status of the autopilot, e.g., lateral navigation mode, can be used to reduce the probability of a trajectory change (heading, altitude, speed) outside the intended path corridor.

Conflict Analysis

The probability of a conflict [P(C)] is defined as the probability that the intruder will enter the host aircraft's protected zone given that no alert is issued and that the host aircraft maintains its current course and speed. To calculate P(C), the positions of the two aircraft must be projected into the future to determine the likelihood of a protected zone violation. However, an explicit analytical solution incorporating the uncertainty variables listed in Fig. 3 cannot easily be formulated. Instead, Monte Carlo simulations are used.

Given the locations, speeds, and headings of the host and intruder aircraft, the P(C) can be estimated through Monte Carlo simulation. Each Monte Carlo run consists of stepping through the trajectories of both aircraft over time and determining whether a conflict occurs. The trajectories vary randomly with each run according to the distributions from Fig. 3. For instance, in one run the intruder might make a 14-deg course change 1 min into the flight; in another run, the intruder may follow a straight-line path for 30 min. After a certain number of Monte Carlo runs, a count of the number of protected zone intrusions was made. Dividing the number of intrusions by the total number of Monte Carlo runs is then an estimator of P(C).

P(C) was determined through the separate analyses of the horizontal- and vertical-plane situations. A conflict occurs when there are both horizontal and vertical separation violations:

$$P(C) = P(C_{horizontal})\,P(C_{vertical}) \quad (1)$$

In the horizontal plane, the Monte Carlo simulations were performed over a range of state estimates for the intruder and for several host aircraft avoidance maneuvers. The result of each set of Monte Carlo runs is a plot of the probability of a horizontal conflict for the specific situation (intruder position, heading, speed) and host aircraft trajectory (straight ahead or maneuvering). Initial intruder positions were varied over a grid of dimensions 200 n miles on a side, in 1-n mile increments. Two intruder velocities were used along with nine intruder heading angles. Nine different host aircraft trajectories were also examined for each intruder situation. The nine host aircraft trajectories included straight ahead, left and right turns of 10 and 20 deg, and 10 and 50 kn speed increases and decreases. The resulting set of probabilities was then stored in a series of lookup tables indexed by position, heading, and speed.

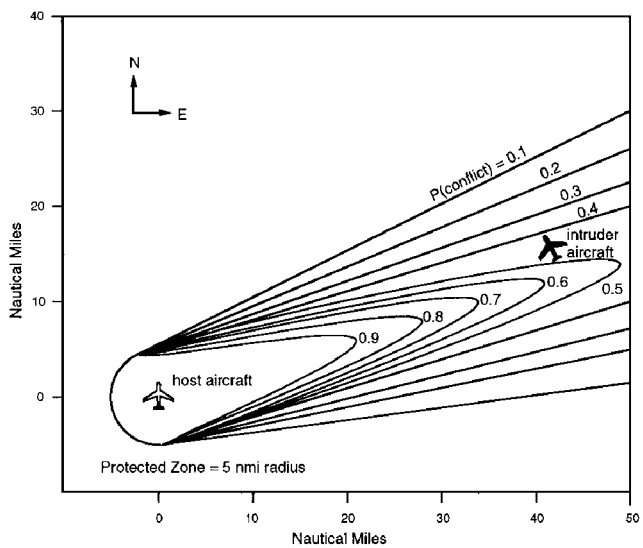

Fig. 4 Example horizontal conflict probability contours.

Figure 4 shows a contour plot of the likelihood of horizontal conflict for a specific encounter situation in which the host aircraft is flying at a heading of 360 deg and an intruder is currently estimated to be flying at a heading of 330 deg. The plot shows actual data based on 10,000 Monte Carlo simulations spaced every 1 n mile. In Fig. 4, the host aircraft is in white at the lower left. The plot shows the conflict probabilities for an intruder aircraft in the surrounding

airspace relative to the host aircraft. For example, an intruder in the position shown in the figure will cause a horizontal conflict in the future with probability 0.45. If the intruder were farther north or east of the host aircraft, this probability would decrease. As the intruder nears the host aircraft, the probability of a conflict will increase if the intruder remains on a collision course. If the intruder changes heading or speed (or if the host aircraft performs an avoidance maneuver), a different contour plot would represent the probability of a conflict.

The vertical conflict probability was obtained by determining the likelihood that the intruder aircraft would follow a vertical path that intersected the protected zone. This was performed through an analytical solution of the vertical probability parameters. Different potential vertical maneuvers of the host aircraft were also evaluated by incorporating the host aircraft's vertical speed into the vertical model.

Alerting Design Tradeoffs

The size of the alert zone affects the performance of the alerting system. If the alert zone is too large, an excessive number of unnecessary alerts will be generated. If the zone is too small, there may not be enough space or time in which to maneuver to avoid a conflict. This tradeoff can be examined using two parameters: the probability of successful alert (SA) and the probability of unnecessary alert (UA).

When an alert is issued, it is defined to be successful if the protected zone is not violated. Thus, the probability of successful alert [P(SA)] is the probability that a conflict does not occur when an avoidance maneuver is performed. Therefore, P(SA) is a function of time and the specific avoidance maneuver that is performed by the host aircraft:

$$\mathrm{P(SA)} = 1 - \mathrm{P(C \mid \text{avoidance maneuver})} \tag{2}$$

An alert is classified as unnecessary if the alert was not required to avoid a protected zone violation. The probability of unnecessary alert [P(UA)] is the probability that a conflict would not have occurred had the host aircraft continued on its current course:

$$\mathrm{P(UA)} = 1 - \mathrm{P(C \mid \text{no avoidance maneuver})} \tag{3}$$

To maximize system performance, it is desirable to maximize P(SA) and minimize P(UA). These goals cannot generally be met simultaneously, and a tradeoff must be managed.

This tradeoff can be visualized using a system operating characteristic (SOC) curve.[7,8] An example is shown in Fig. 5. An SOC curve is a plot of P(SA) for a given avoidance maneuver vs P(UA). Each point on the SOC curve represents an alerting threshold setting.

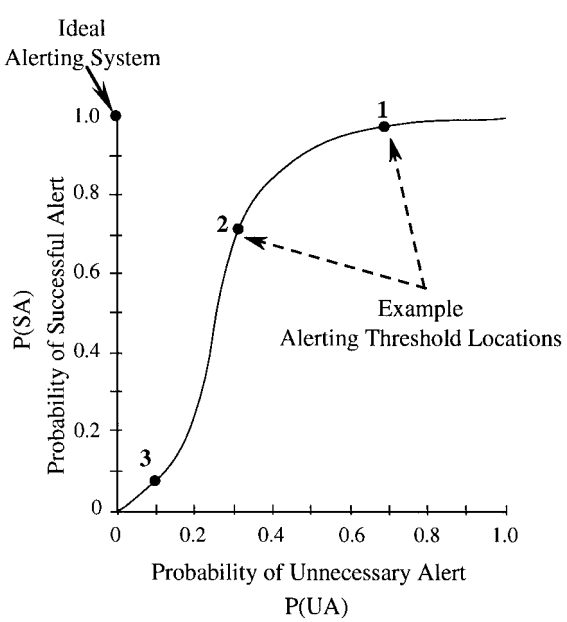

Fig. 5 Example SOC curve.

For example, in Fig. 5, threshold 1 corresponds to a large alert zone: alerts are generated early, resulting in a large value for P(SA) but also a high rate of UAs. As the alert zone size is reduced, the threshold moves along the SOC curve to points 2 and 3. The result is a reduction in UAs but also a reduction in SAs because less time and space are available to perform the avoidance maneuver.

An ideal system would operate in the upper-left-hand corner, where P(UA) is zero and P(SA) is one: ideally, all alerts are necessary and successful. In reality, SOC curves do not reach the ideal operating point. Instead, the threshold must be placed along the curve based on the tradeoff between UAs and SAs.

The shape of the SOC curve is a function of sensor accuracy, the type of avoidance maneuver, operator response latency, and maneuvering aggressiveness. As sensor accuracy is increased or as response time is reduced, for example, the SOC curve will move closer to the ideal operating point.

An SOC curve that lies along the diagonal from the origin [P(SA) = 0, P(UA) = 0] to the upper-right-hand corner [P(SA) = 1, P(UA) = 1] represents a system that is poorly designed. In such a case, an alert is as likely to be successful as unnecessary. This means that alerting is just as likely to produce a conflict as not alerting. Thus, the more that the SOC curve moves away from the diagonal, the better the alerting decision.

Because P(SA) depends directly on the choice of avoidance maneuver, a different SOC curve can be constructed for each maneuver option. The most effective avoidance options can then be identified based on the shape of their SOC curves.

Prototype Alerting Logic

As already described, the prototype system uses four alert stages. The first three stages produce alerts in the cockpit that are intended to aid the flight crew in resolving the conflict before tactical maneuvering is required. At the fourth stage, ATC is notified to issue commands to provide traffic separation. To set the conditions at which these stages are triggered, it is necessary to examine the tradeoffs between P(UA) and P(SA). This requires balancing the likelihood of a conflict against the ability of the host aircraft to avoid a conflict. To do so, five standard conflict resolution maneuvers were considered: 1) left heading change of 30 deg, 2) right heading change of 30 deg, 3) climb or descent at 2000 ft/min, 4) speed increase of 50 kn, and 5) speed decrease of 50 kn.

These maneuvers serve as benchmarks for estimating the ability of the host aircraft to avoid a conflict. When the intruder is far from the host aircraft, any of these five maneuvers could be used to resolve the conflict. As the intruder nears the host aircraft, some of these maneuvers may no longer provide the required separation between aircraft. The premise behind the alerting logic is that if a sufficient number of these maneuvers are still available to the pilot, the alert can be delayed. When the pilot's options begin to disappear, an alert should be issued.

A maneuver was defined to be available to the host aircraft if, by performing the maneuver, the probability of a conflict was reduced to less than 0.05, i.e., P(SA) > 0.95. The five maneuver options just listed included the probabilistic response time described earlier (with a mean latency of 1 min). Thus, when a maneuver was deemed to be not available, safe separation could still be achieved if the pilot reacted more quickly or more aggressively than assumed in the model.

In real time, the logic calculated the number of avoidance maneuvers available N to resolve a conflict with the intruder. This was done using the probability contour data stored in lookup tables for each of the five avoidance maneuvers. By comparing N with P(UA), the appropriate alert stage was defined as shown in Table 1.

The leftmost column of Table 1 shows the probability of a conflict if the host aircraft continues along its current trajectory. This assumes that the intruder's trajectory can be represented by the model discussed earlier. The rightmost column shows P(UA), which as discussed earlier is related to P(C) by Eq. (3). The other columns indicate the defined alert stages as a function of N. Generally, the more options available to the pilot, the lower the alert stage.

For example, if P(UA) is 0.35 and there are two avoidance maneuvers available, then the alert stage is 2. If P(UA) drops below 0.3

Table 1 Alert level classification

P(C | no maneuver)	Number of avoidance maneuvers available N				P(UA)
	None	One	Two	Three or more	
0.0–0.1	——	——	——	——	0.9–1.0
0.1–0.2	1	1	——	——	0.8–0.9
0.2–0.3	1	1	——	——	0.7–0.8
0.3–0.4	2	1	1	——	0.6–0.7
0.4–0.5	2	2	1	1	0.5–0.6
0.5–0.6	3	2	2	1	0.4–0.5
0.6–0.7	3	3	2	2	0.3–0.4
0.7–0.8	AT	3	3	2	0.2–0.3
0.8–0.9	AT	3	3	3	0.1–0.2
0.9–1.0	AT	AT	AT	AT	0.0–0.1

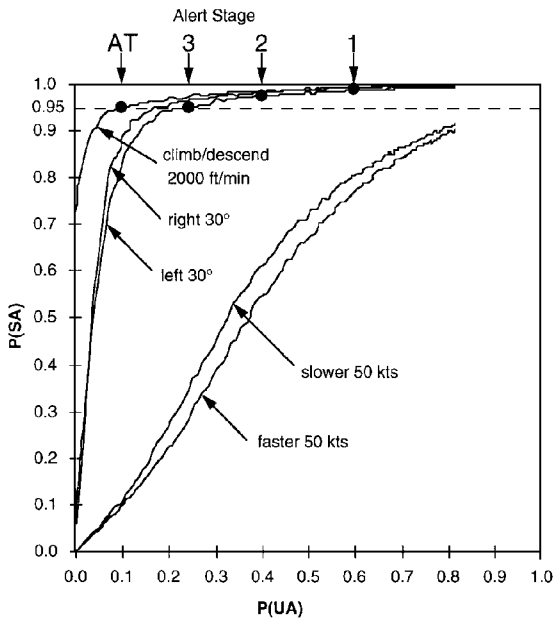

Fig. 6 SOC curve: aircraft on perpendicular tracks.

or if N is reduced to 1, then the alert stage increases to 3. If P(UA) drops below 0.1, then the AT stage is triggered.

Note that because the probability of conflict along different avoidance maneuvers can be estimated, the alerting logic can also be used to determine the magnitude of maneuvering required to resolve a conflict. By interpolating P(SA) between different maneuvering magnitudes, the required action to resolve a conflict with 95% confidence can be determined. For example, if P(SA) for a 10-deg right turn is 0.93 and P(SA) for a 20-deg right turn is 0.97, then a 15-deg turn will result in P(SA) of approximately 0.95. Thus, the probability data can be used both to determine P(SA) when a maneuver is specified and to determine the magnitude of maneuvering that is required to achieve a specified value of P(SA).

To better understand the underlying design process, the thresholds from Table 1 can be mapped into SOC curves. Figure 6 shows SOC curves for two coaltitude aircraft on a collision course along flight paths at right angles to one another. SOC curves corresponding to each of the five resolution maneuver options are shown.

When the intruder is far from the host aircraft, the situation maps into the upper-right-hand corner of the plot: it is likely that a conflict will not actually occur [P(UA) = 1], and it is likely that any avoidance action would resolve the situation [P(SA) for each of the five avoidance maneuvers equals 1]. Data for Fig. 6 were not obtained beyond 200 n miles, and so the SOC curves in the figure do not extend all of the way to the upper-right-hand corner.

As the intruder continues on a collision course, it becomes more clear that a conflict will occur: P(UA) decreases, and the situation moves from right to left along the curves. Thus, P(UA) is related to the distance between aircraft and to the time before closest point of approach. As P(UA) decreases, P(SA) also decreases in differing amounts according to the different SOC curves. The effectiveness of a given maneuver depends on how slowly its P(SA) decreases. When a curve's value of P(SA) drops below 0.95, the corresponding avoidance maneuver is no longer available. Thus, as the situation progresses to the left in Fig. 6, the different avoidance maneuvers become unavailable, in order, from speed changes to turns and finally to climb or descent. Thus, the SOC curves show that for this case, vertical maneuvers are the most effective.

The first maneuvers to become unavailable are the speed change maneuvers, at P(UA) of approximately 0.9. This is because large speed changes are generally required to resolve conflicts in the time scales under consideration.

Until P(UA) drops below approximately 0.25, turns and climb/descent avoidance maneuvers will still provide the required separation. At approximately P(UA) = 0.25, however, a 30-deg left turn maneuver is no longer an option. At approximately P(UA) = 0.2, the 30-deg right turn is also no longer an option. When P(UA) reaches approximately 0.1, the climb/descend options become unavailable.

At a given value of P(UA), N corresponds to the number of SOC curves that have values above P(SA) = 0.95. Figure 6 also shows when the four alert stages are triggered as a function of P(UA). Cross referencing with Table 1, stage 1 is triggered when N is 3 or more and P(UA) drops to 0.6. Stage 2 is triggered when P(UA) drops to 0.4, and stage 3 is triggered when N drops to 2. Finally, the AT stage is triggered when N drops to 0. Although Fig. 6 shows SOC curves for a direct collision between two aircraft on perpendicular flight paths, other geometries produce similar patterns.

The five avoidance maneuvers used here are intended to represent strategic maneuver limits. A large response time (mean = 1 min) is modeled in the avoidance maneuvers (see Fig. 3) and when N is 0, the host aircraft can still maneuver out of the conflict. A more aggressive, tactical maneuver such as a 45-deg-heading turn or a combined climbing turn may still be available when the five assumed strategic maneuvers are not.

Further examination of the SOC curves shows that speed changes make only a limited contribution to the prototype logic. In many cases, a speed change of greater than 50 kn is required for adequate separation with 95% confidence. As can be seen from Fig. 6, the SOC curves for the speed maneuvers deviate only slightly from the diagonal. Thus, it is difficult to provide successful, necessary alerts with speed control alone. Similar difficulties with relying on speed control are mentioned by Krozel et al.,[13] using a much different conflict analysis method based on optimal control theory.

Evaluation

The calculation of the probability of conflict is time consuming due to the large number of required Monte Carlo simulations. Accordingly, the probability contours were stored in lookup tables to be accessed in real time. In operation, the system takes aircraft state data and compares their values against the lookup tables to determine the appropriate alert stage using Table 1. When state values varied between the indices of the lookup tables, the values were linearly interpolated to estimate the probability of conflict.

The alerting logic was evaluated using numerical encounter simulations at the Massachusetts Institute of Technology and also in a human-in-the-loop simulation study at NASA Ames Research Center. These evaluations were not exhaustive but were used to explore several research issues.

As examples, Figs. 7 and 8 show the observed times at which the alert stages were triggered for two different encounter scenarios. Figure 7 shows the same situation described by the SOC curves in Fig. 6: two aircraft on a collision course on perpendicular trajectories. Alert stage 1 is triggered 12.3 min prior to the time of closest point of approach (CPA). Stages 2 and 3 are triggered at approximately 8.5 and 5.8 min to CPA, respectively. ATC is notified to take over authority (at the AT stage) at 3.3 min to CPA. Finally, TCAS produces a traffic advisory (TA) at approximately 45 s and a resolution advisory (RA) at 35 s to CPA.

Figure 8 shows a case in which the two aircraft are not on a direct collision course but will pass within 6 n miles of one another. Stage 1 is triggered 6.5 min before CPA, and stage 2 is triggered 2.2 min before CPA. A TCAS TA is also generated at approximately 30 s before CPA. When the traffic passes the host aircraft, the alert stages gradually decrease. Thus, the logic increases the alert stage as the

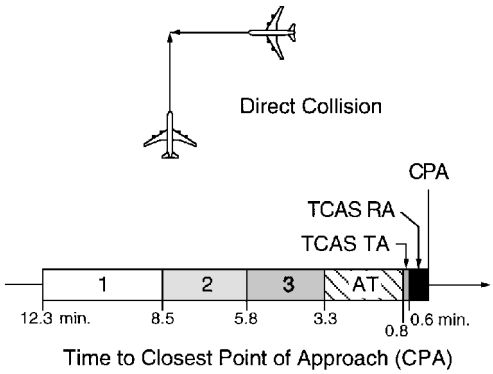

Fig. 7 Alert time line: direct collision, 90-deg crossing angle.

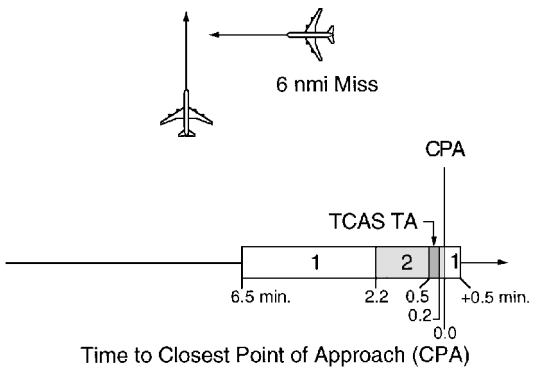

Fig. 8 Alert time line: 6-n mile minimum separation, 90-deg crossing angle.

potential for a conflict rises and reduces the alert stage as it becomes less likely that the intruder could turn and cause a conflict.

At the NASA Ames Research Center, the prototype alerting logic was incorporated in a 747-400 simulator as part of a study of pilot decision-making aids for free flight. In this study, enroute conflicts were scripted to examine pilot response and to exercise the alerting logic.

In operation, the alerting logic was used to trigger the four stages of alerts discussed earlier. Additionally, the probability data were used to determine the magnitude of maneuvering required to resolve conflicts at a specified level of confidence. The pilots in the study were given an interactive tool to explore different maneuvering options. These maneuvers were compared against the probability data to determine whether the conflict would be resolved with 95% confidence. The cockpit display then indicated to the pilot whether the proposed maneuver was likely to be successful.

Preliminary results from the NASA study show that the pilots successfully resolved conflicts without ATC guidance in most cases. AT alert stages were only observed in scenarios where the intruding aircraft was purposely diverted toward the host aircraft at close proximity. However, a more complete analysis is required to more fully evaluate the alerting logic and to determine the potential impact of airborne conflict resolution on air traffic management.

Concluding Remarks

The Monte Carlo approach used has the advantage that it can handle complex, non-Gaussian probability distributions. However, a major limitation is that a significant amount of preprocessing is required before the alerting thresholds can be defined. Once the thresholds are selected, they only apply to the specific probabilistic model used in the Monte Carlo runs. If a change in the model is required, e.g., to examine the effect of varying sensor accuracy, then the Monte Carlo simulations must be rerun using a new model to update the alerting thresholds. In contrast, the analytical solution used in Ref. 9 can be solved in real time but relies on simple Gaussian distributions. Thus, there is a tradeoff between the complexity of the probabilistic model and the ability to estimate probabilities rapidly.

Lookup tables were chosen to codify the alerting thresholds for this system because of their relative simplicity. However, it may be more effective to use neural networks to map the encounter situation directly to the alerting thresholds without the use of lookup tables. Neural nets can be trained off-line using the precalculated data from the Monte Carlo simulations. Once trained, a neural net could represent the alerting thresholds in the real-time system.

Another consideration involves the scope of the conflict. The resolution maneuvers used to develop the alerting logic are based on the immediate problem of avoiding a conflict and do not consider the additional maneuvering required to return to the original flight path. Thus, the logic does not incorporate issues such as increased fuel burn or flight time in the decision to alert. Because the proposed benefits of free flight revolve around efficient traffic flow, it will be necessary to incorporate cost-based considerations into the logic in the future. This can be achieved, for example, by weighing avoidance maneuver options by the additional cost or deviation each option would incur.

Finally, centralized traffic management issues have been ignored. Because, as assumed in this free flight concept, pilots have initial responsibility for traffic separation, ground controllers could have difficulty when suddenly presented with a conflict that was not resolved by the flight crews. Additional conflict detection and resolution aids must be provided for ground controllers to enable them to return to the traffic management loop and handle traffic once they are alerted to a conflict. Alternatively, it may be more appropriate for all conflict detection and resolution activities to be performed on the ground. In either case, the design approach presented could be applied in an air, ground, or mixed mode of operation to develop appropriate alerting thresholds.

Acknowledgments

This research was supported by the NASA Ames Research Center. The authors are especially appreciative of Kevin Corker, Sandra Lozito, Walter Johnson, Barry Sullivan, Joe King, and Paul Soukup at NASA Ames Research Center for their support and input during the development and implementation of the alerting logic.

References

[1]"Final Report of RTCA Task Force 3: Free Flight Implementation," Radio Technical Committee on Aeronautics, Washington, DC, Oct. 1995.

[2]Williamson, T., and Spencer, N. A., "Development and Operation of the Traffic Alert and Collision Avoidance System (TCAS)," *Proceedings of the IEEE*, Vol. 77, No. 11, 1989, pp. 1735–1744.

[3]Miller, C. A., Williamson, T., Walsh, J. A., Nivert, L. J., Anderson, J. L., and Zeitlin, A. D., "Initiatives to Improve TCAS-ATC Compatibility," *Journal of ATC*, Vol. 36, No. 3, 1994, pp. 6–12.

[4]"Minimum Performance Specifications for TCAS Airborne Equipment," Radio Technical Committee on Aeronautics, Document RTCA/ DO-185, Washington, DC, Sept. 1983.

[5]Barkat, M., *Signal Detection and Estimation*, Artech, Boston, 1991, pp. 115–174.

[6]Sheridan, T. B., *Telerobotics, Automation, and Human Supervisory Control*, MIT Press, Cambridge, MA, 1992, pp. 49–60.

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

[8]Kuchar, J. K., "A Unified Methodology for the Evaluation of Hazard Alerting Systems," Ph.D. Dissertation, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, Jan. 1995.

[9]Paielli, R. A., and Erzberger, H., "Conflict Probability Estimation for Free Flight," *Journal of Guidance, Control, and Dynamics*, Vol. 20. No. 3, 1997, pp. 588–596.

[10]Williams, P. R., "Aircraft Collision Avoidance Using Statistical Decision Theory," *Sensors and Sensor Systems for Guidance and Navigation II*, Vol. 1694, Society of Photo-Optical Instrumentation Engineers, Bellingham, WA, 1992, pp. 29–34.

[11]Heuvelink, G. B., and Blom, H. A., "An Alternative Method to Solve a Variational Inequality Applied to an Air Traffic Control Example," *Analysis and Optimization of Systems*, Springer–Verlag, New York, 1988, pp. 617–628.

[12]Klass, P. J., "New TCAS Software to Cut Unneeded Evasive Actions," *Aviation Week and Space Technology*, Vol. 146, No. 4, 1997, pp. 57–59.

[13]Krozel, J., Mueller, T., and Hunter, G., "Free Flight Conflict Detection and Resolution Analysis," AIAA Paper 96-3763, July 1996.