A98-37086

AIAA–98–4242

EVALUATION OF VERTICAL COLLISION AVOIDANCE MANEUVERS FOR PARALLEL APPROACH

Lee F. Winder[*] and James K. Kuchar[†]

Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract

Two candidate evasion maneuvers were evaluated for use with a collision alerting system for independent closely-spaced parallel approaches in instrument conditions. The two maneuvers were a wings-level climb and a climbing turn away from parallel traffic. Pairs of aircraft on parallel approach were simulated using pre-recorded trajectories covering a range of normal approach and blunder scenarios. Each scenario was repeated twice, with the endangered aircraft responding to alerts with either the climb-only or the climbing-turn evasion. The climb-only maneuver is shown to result in 38-times as many collisions as the climbing-turn for nominal alert threshold settings. It is possible to reduce the collision rate by adjusting threshold parameters, but the false alarm rate increases. The climb-only maneuver is shown to be uniformly less safe than the climbing-turn for all parameter combinations. Results are illustrated using System Operating Characteristic curves.

Introduction

As part of the NASA Airborne Information for Lateral Spacing (AILS) program, a cockpit-based alerting system is required for aircraft on independent instrument approaches to closely-spaced parallel runways.[1] The alerting system is designed to prevent collisions in the event that an aircraft blunders from its expected approach path. The current design goal is to enable approaches down to 2500 ft runway spacings, which is well below the current 4300 ft minimum (3400 ft at airports with the Precision Runway Monitor system).[2,3]

The NASA Langley Research Center has developed (and continues to study and modify) a candidate alerting logic for AILS.[1,4-6] As originally conceived, the system requires evading pilots to perform a climbing-turn escape maneuver, combining vertical acceleration with a 45° track angle change away from the approach centerline. Simulations based on this evasion have

shown AILS to improve safety for a variety of blunders, at runway spacings down to 1700 feet.

For reasons of operational simplicity and reduced training costs, a straight-ahead climb ("climb-only") has been advanced as a possible substitute for the climbing-turn. First, such a maneuver is compatible with current missed approach procedures, and with the emergency maneuvers for existing alerting systems such as the Traffic Alert and Collision Avoidance System (TCAS) and Ground Proximity Warning System (GPWS).[7,8] To include a turn as part of a procedural avoidance maneuver would require additional pilot training and expense to airlines. Second, an aircraft turning off of the final approach course is more likely to interfere with other aircraft in the airport vicinity than one performing a straight-ahead climb. Air traffic controllers would require additional procedures to safely handle this.

Analysis was needed to determine the feasibility of this maneuver substitution. In past research, MIT developed a methodology for evaluating the performance of alerting systems in terms of safety (collision rate) and unnecessary alert rate.[6,9] The methodology was previously applied to AILS with the climbing-turn maneuver. That analysis has now been expanded to compare the two candidate evasion maneuvers, testing over a range of threshold settings so that the relative performance potential can be observed. The method and results are the subject of this paper.

AILS Alerting Algorithm

AILS is envisioned as an airborne alerting system similar to TCAS, but specialized for the parallel approach environment. A computer on board each aircraft collects information from sensors, and over datalink from neighboring aircraft. Using this information, the computer decides whether or not to issue an alert based on a worst case assumption and an assumed range of possible aircraft behavior.

---

[*] Graduate Research Assistant
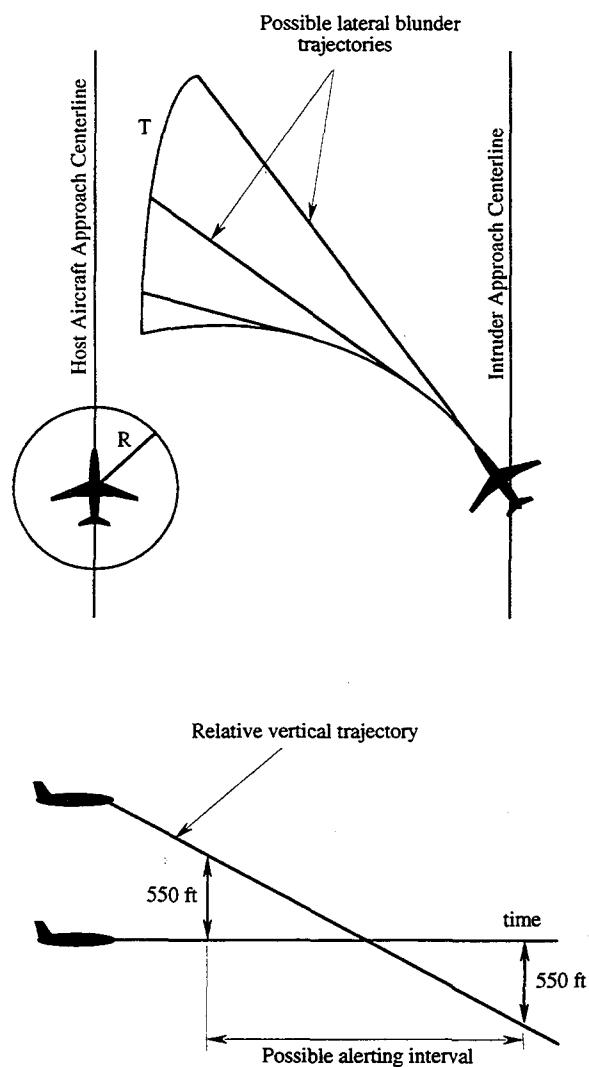[†] Assistant Professor, Member AIAA

838

Fig. 1   AILS Lateral and Vertical Geometry

AILS displays alerts of several levels of urgency on one or both aircraft performing the parallel approach, depending on the nature of the conflict. However, the underlying philosophy is that adequate separation should be ensured even if a blundering aircraft is not responsive to alerts (e.g., due to some mechanical failure).

The full dynamic model used by AILS is too complex to describe here in detail, but the relevant parameters for this analysis are illustrated in Figure 1. In Fig. 1, an encounter situation is shown from the point of view of a normally approaching evader aircraft (on the left). The evader is modeled as maintaining a constant-velocity approach along the extended runway centerline and glide slope. The "intruder" aircraft is modeled as potentially following any of a range of trajectories. The model trajectories include a circular path based on the turn rate measured via datalink, and also a series of cases in

which the intruder rolls out into straight-line flight. The result is a fan of potential trajectories as shown. The intruder's airspeed and vertical velocity are assumed to be constant at all times.

Two parameters, called R and T, are used to define the alerting threshold. If it is possible for the intruder to pass within R horizontal and 550 vertical feet of the host in under T seconds, an alert is issued to the evader. Although the true AILS logic is able to generate a sequence of alerts (each based on a different combination of R and T) of increasing criticality as the intruder closes in, this research focuses on the final breakout alert.

Acceptable values of R and T were determined for several runway spacings by NASA through a trial and error process of blunder simulation. In these simulations, the evading aircraft performed a climbing-turn escape maneuver in response to breakout alerts. For the 2500 foot runway spacing, the values used in simulation tests have been R = 550 ft and T = 13 sec. These values are referred to as the nominal parameter values in the remainder of this paper.

System performance was expected to degrade if the same, nominal R and T values were used with the climb-only maneuver in place of the original climbing turn. This is because the climbing turn generally provides additional separation between aircraft due to the turning component. Thus, the time available in which to escape is different with each type of maneuver, resulting in different optimal values for R and T. The analysis discussed here was designed to determine whether adequate performance is obtainable with the climb-only evasion through adjustment of R and T only, or whether a more complete redesign of the alerting algorithm would be required.

Approach

Alerting logic performance was estimated using a numerical trajectory simulation developed previously for evaluation of an MIT concept logic.[5,6] Intruder trajectories were based on a pre-recorded trajectory set provided by Rockwell-Collins.[4] These were sampled at approximately 2 Hz from piloted flight on a Fokker 70 simulator, and cover a range of behavior: normal approaches; slow constant-rate turns at a 5° bank angle; coordinated heading-changes of 5°, 15°, and 30°; and two types of "fake" blunder, in which the intruder begins a blunder but returns to its proper approach path before crossing that of the threatened aircraft. Separate trajectory data were recorded for calm and turbulent conditions and for airspeeds of 130, 145, and 160 kts. To further expand the variety of possible encounters,

this set of trajectories was used over a series of initial longitudinal separations (within ±1.5 nmi of the threatened aircraft), initial vertical separations (within ±1000 ft), and vertical speeds (ranging from descent along the glide slope, to a 2000 ft/min climb). Lateral spacings were set to a 2500 ft runway spacing. In total there were 36,270 trajectories used in the simulations.

Pairs of aircraft were simulated, one performing an ideal normal approach while another (the intruder) followed each of the blunder or normal approach paths from the aforementioned set. The alerting logic was implemented only for the threatened aircraft, to simulate a case in which the blunderer is unable to respond to alerts.

In separate simulation runs, the threatened aircraft responded to alerts with either the climb-only or climbing-turn avoidance maneuver. The climb-only avoidance maneuver consisted of a 2-second response delay, followed by a 0.25 g pull-up to a 2000 ft/min climb rate, and a 15 knot airspeed increase at 1 kt/sec acceleration. The climbing-turn adds to this a 15°/sec roll (following the delay) to a 30° bank angle, with roll-out at a heading 45° from the approach centerline.

The outcome of each approach was recorded, including (1) whether an alert was generated, (2) whether a collision occurred, and (3) whether an alert was really necessary. Six mutually exclusive categories, listed in Table 1, were used to classify the possible outcomes. A collision was defined to have occurred if separation at any point during an approach was less than 500 ft. An alert was considered necessary if a collision would have occurred without an alert. Thus, an alert in a case where a 501 ft separation would have occurred without the alert is by definition unnecessary. Such a definition is required as a specific performance metric, even though violations of the 500 foot limit do not guarantee collisions in reality, and separations over 500 feet might intuitively merit prevention. A pilot or controller's impression of "necessary" is important, but is subjective and difficult to use analytically.

If an alert was not issued at all during an approach, the run was classified in Table 1 as either a Correct Rejection (if a collision did not occur) or a Missed Detection (if a collision did occur). If an alert was issued, the outcome was placed in one of four categories. An Unnecessary Alert was a case where an alert was not required to prevent a collision, but was issued anyway, and no collision occurred. If an alert triggered a collision when none would have occurred otherwise, the run was classified as an Induced Collision. A Correct Detection occurred when a collision was averted because of an alert. Finally, a Late Alert was a case in which an alert was issued too late to prevent a collision.

To see the dependence of performance on the threshold parameters, separate simulations were performed over a range of R and T values. For this, R was varied from 350 to 750 feet in 100 foot increments, and T was varied from 5 to 25 seconds in 2 second increments. To examine the sensitivity of performance to pilot reaction time, a second series of simulations varied reaction time from 0 to 15 seconds in 1 second increments, with R and T fixed at their nominal values of 550 feet and 13 seconds.

Results

The quantities of actual interest in alerting performance are the rates of collision and false alarm. In terms of the categories of Table 1, collisions are the union of Missed Detections, Induced Collisions, and Late Alerts. False alarms include both Unnecessary Alerts and Induced Collisions. Due to the uncertainties in aircraft trajectories, reducing the collision rate by increasing R or T will result in a corresponding increase in unnecessary alert rate. Because false alarms have deleterious effects of their own on long-term safety (by degrading pilot confidence in the system), choosing thresholds requires a conscious trade-off between collisions and false alarms to attain the best system performance.

Table 1   Outcome Categories

Outcome Category	Alert Issued?	Collision Occurred?	Alert Necessary?
Correct Rejection	No	No	No
Missed Detection	No	Yes	Yes
Unnecessary Alert	Yes	No	No
Induced Collision	Yes	Yes	No
Correct Detection	Yes	No	Yes
Late Alert	Yes	Yes	Yes

Table 2   AILS Performance Summary

(Nominal Threshold Values, 2500 ft runway spacing, 2 sec pilot reaction time)

Maneuver	Correct Rejections (CR)	Missed Detections (MD)	Unnecessary Alerts (UA)	Induced Collisions (IC)	Correct Detections (CD)	Late Alerts (LA)
Climbing-Turn	0.98227	0	0.01257	0.00008	0.00505	0.00003
Climb-Only	0.98227	0	0.00948	0.00317	0.00403	0.00105

The primary question, then, is whether adequate alerting performance is possible for the climb-only evasion maneuver for some combination of R and T. At least a partial answer to this question is obtained by plotting performance metrics as a function of system parameters using simulation output, as is discussed below.

Performance Summary

Table 2 shows a comparison of observed outcome rates between the climbing-turn and climb-only escape maneuvers. These results are for R = 550 ft and T = 13 sec. When climb-only maneuvers are substituted for climbing-turns, there is an approximate 40-times increase in the rate of Induced Collisions, along with a 35-times increase in Late Alerts. This is an overall 38-times increase in collisions.

False Alarm Analysis

Figures 2 and 3 summarize false alarm performance over a range of R and T values. Curves relate the Cumulative Number of False Alarms to the closest approach of the two aircraft had no alert been issued. For example, if the Cumulative Number of False Alarms equals 400 at a Closest Approach value of 2000 ft, then 400 of all False Alarms that occurred were such that a closest approach below 2000 ft would have occurred had there been no alert.

Figure 2 shows data for a constant T of 13 seconds and three values of R, and Figure 3 shows data for R = 550 ft and three values of T. Increasing the value of either T or R results in a greater number of False Alarms at every closest approach distance. The variable slopes of the curves in the plots are due to characteristics of the specific blunder trajectories that were used.

It is useful to distinguish between two types of false alarms: those that occur during actual blunder (or "fake" blunder) scenarios, and those occurring with the intruder on a technically correct approach. The first type of false alarms are of limited importance in system design, in view of the fact that blunders are extremely rare events. Frequent false alarms that disrupt normal approaches, on the other hand, would tend to reduce pilot confidence in the validity of alerts and would negatively impact traffic flow.

To examine the susceptibility of the system to produce false alarms during normal approaches, Figure 4 plots the fraction of normal intruder approaches resulting in false alarms as a function of R and T. Some of these normal approaches were cases in which the intruder was climbing straight-ahead, such as during a standard missed approach procedure. For the R = 350 ft case, normal approach false alarms did not occur until T exceeded 21 seconds. For the maximum value of R of 750 ft, normal approach false alarms began to occur above T = 19 sec. Thus, overall, AILS appears to give
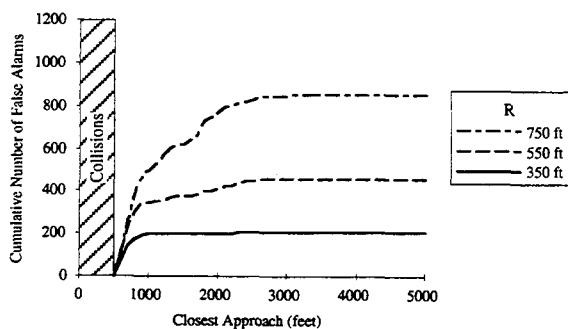

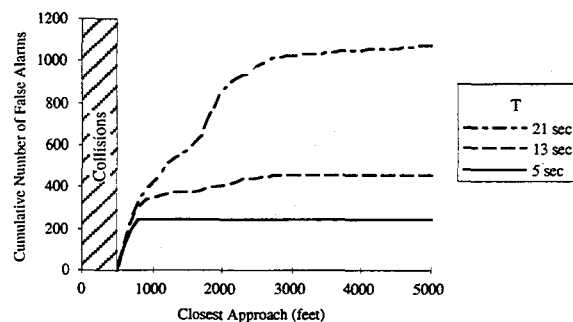**Fig. 2   Cumulative Number of False Alarms** (Varying R at T = 13 sec)

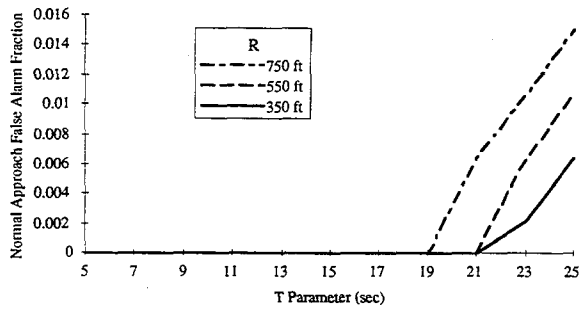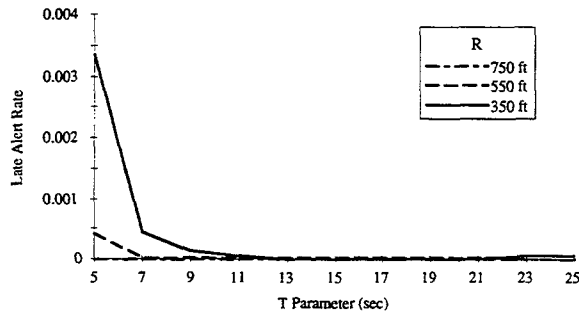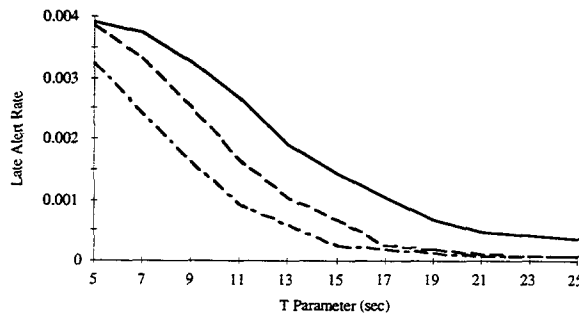**Fig. 3   Cumulative Number of False Alarms** (Varying T at R = 550 ft)

Fig. 4 Normal Approach False Alarm Rate

few "true" unnecessary alerts until the parameter values are increased well beyond nominal settings. However, because normal approaches were only a small fraction of the entire test trajectory set, these bounds must be considered rough estimates. A more thorough study of normal approach false alarms, using a better model of approach behavior, may be needed.

Note that Figures 2 to 4 apply equally to the climbing-turn and climb-only maneuvers, because the curves are functions of the parameters R and T only, and not of the escape maneuver.

Collision Analysis

Figure 5 shows the fraction of all simulated trajectories that resulted in a collision because the system alerted too late (Late Alert) as a function of R, T and the evasion maneuver. Increasing the values of either R or T decreases the rate of collisions. For the climbing-turn, the nominal parameter values are such that Late Alerts occur at a negligible rate (see Table 2) — a condition not equaled for any parameter values with the climb-only maneuver.

Figure 6 shows similar data for the Induced Collision rate. Once again, no values of R and T for the climb-only case attain performance equal to that of the climbing-turn with nominal parameter values.

Note that the induced collision rate peaks for intermediate values of T in the climb-only case. This is due to the fact that induced collisions are indicative of a deficiency in the trajectory model. For an induced collision to occur, the alerting system first fails by alerting when not necessary, and then by not providing adequate separation to prevent a collision. For small values of T, trajectory extrapolation is shortened, and the alerting system is better able to distinguish true

(a) Climbing-Turn

(b) Climb-Only

Fig. 5 Late Alert Rate
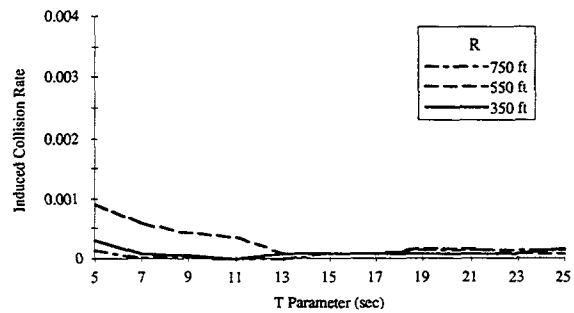
(a) Climbing-Turn
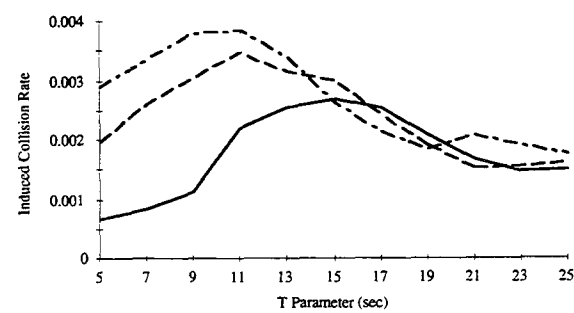
(b) Climb-Only

Fig. 6 Induced Collision Rate

collision threats. However, there is correspondingly less time to escape, resulting in higher late alert rates as shown in Fig. 5. At large values of T, the intruder is far enough away when an alert is issued that adequate separation is likely regardless of the trajectory the intruder follows. Intermediate values of T, however, are such that the trajectory extrapolation is perhaps too long relative to the uncertainties, and induced collisions become more likely.

Response Time Effects

Another important issue is the speed with which a pilot responds to an alert. Both the climbing-turn and climb-only maneuvers were initially based on a 2 second pilot latency. This is an optimistic estimate compared to the 5 seconds allowed by both TCAS and GPWS.[7,8] It is therefore desirable to know how sensitive system performance is to changes in reaction time.

Simulations were run for both the climbing-turn and climb-only evasion using the nominal values of R and T, with varying pilot reaction time. The results are shown in Figure 7, in terms of the fraction of imminent collision blunder scenarios that were correctly resolved. The fraction of imminent collisions avoided, f, can be determined for each maneuver with the expression

$$f = \frac{CD}{MD + CD + LA} \quad (1)$$

where CD, MD, and LA are the Correct Detection, Missed Detection, and Late Alert rates, respectively (c.f. Table 2). For example, for a pilot response delay of 2 sec, these values can be read directly from Table 2. It is clear from Fig. 7 both that the climbing-turn is more robust with respect to reaction time, and that performance of the climb-only maneuver is inferior even for perfect (zero) reaction time.
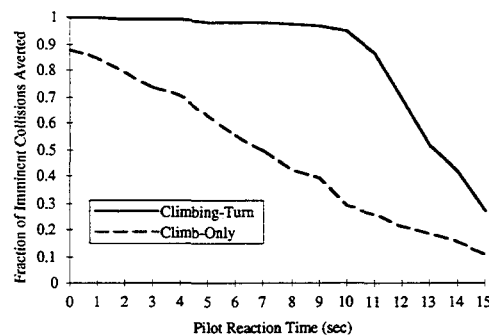

Fig. 7 Imminent Collisions Averted vs. Reaction Time
(R = 550 ft and T = 13 sec)

System Operating Characteristic Curves

System Operating Characteristic (SOC) curves have been devised as a way to directly view the tradeoff between false alarms and collisions.[9] For a given threshold definition and blunder dynamics the alerting system will have certain probabilities of false alarms and of successful system responses (anything not resulting in a collision). An SOC curve is a plot of the functional relationship between the two as a parameter of the alerting system is varied.

For the simulation used in this research it would be unreasonable to interpret the event fractions in Table 2 as the "probabilities" of those events, because blunders make up the vast majority of test trajectories, whereas in actual operation the opposite would be true. Furthermore, it is doubtful that all simulated blunder scenarios are equally likely, as a probabilistic interpretation of the numbers would implicitly assume.

A compromise is to normalize the totals of false alarms and successful alerts by the total number of alerts that were issued. In terms of the variables from Table 2, this can be expressed as:

$$P(FA) = \frac{UA + IC}{IC + UA + CD + LA} \quad (2)$$

$$P(SA) = \frac{UA + CD}{IC + UA + CD + LA} \quad (3)$$

where UA, IC, CD and LA are the fractions of Unnecessary Alerts, Induced Collisions, Correct Detections, and Late Alerts respectively that occurred over all intruder trajectories. The resulting quantities are interpreted as conditional probabilities, given that an alert has been issued. P(FA) as defined here is then the probability that an issued alert is a false alarm. Note that by definition a false alarm is an alert that is issued when a collision would not have occurred had that alert not been generated. The second metric, P(SA), is the probability that an issued alert is successful in avoiding a collision. Thus, 1-P(FA) is the probability that a collision will occur *without* an alert, and 1-P(SA) is the probability that a collision will occur *with* an alert. Plotting P(SA) versus P(FA) produces an SOC curve, such as in those shown in Figure 8.

An ideal alerting system would have P(SA) = 1 and P(FA) = 0, and would therefore operate in the upper-left corner of the plot. Operating points on a diagonal line of slope 1 through the origin represent conditions of no overall benefit. That is, alerting while on the diagonal

(a) Climbing-Turn

(b) Climb-Only

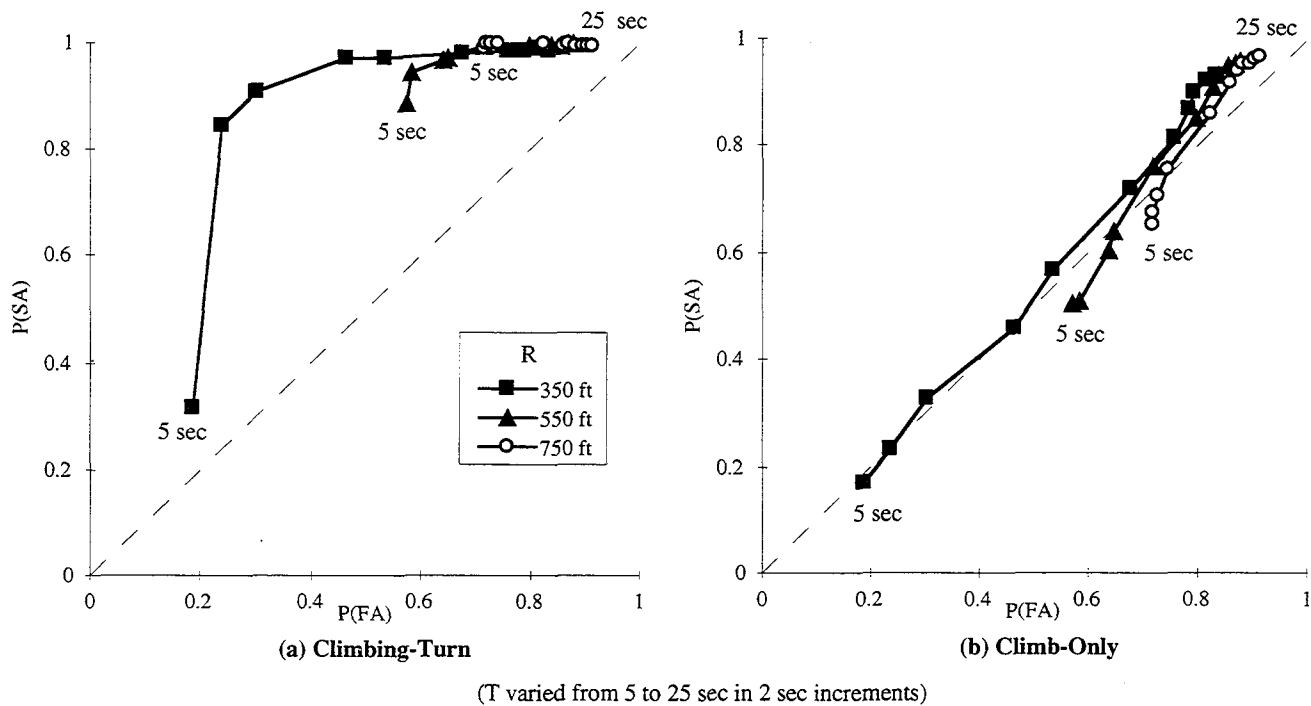(T varied from 5 to 25 sec in 2 sec increments)

Fig. 8 System Operating Characteristic (SOC) Curves

line is equally likely to avoid a collision as not alerting. At points below this diagonal, alerting is more likely to result in a collision than not alerting.

Figure 8 contains SOC curves for AILS with each evasion maneuver as a function of the parameters R and T. For a given value of R, the system operating point will move to the right along each of the curves as the value of T is increased. Thus, increasing T generally increases both P(SA) and P(FA). The benefit of alerting is clearly apparent for the climbing-turn evasion, as all tested combinations of R and T place the operating point above the diagonal. Nominal parameter values place the climbing-turn operating point at P(SA) of approximately 0.994, with P(FA) of approximately 0.713.

In contrast, operating points for the climb-only evasion all lie in the vicinity of the diagonal. A high level of safety is only obtainable for high values of T, and then only along with a high rate of false alarms.

Recall that normal approach false alarms are more important to avoid than false alarms that occur during blunders. Because the P(FA) quantity lumps all false alarms together, it is not necessary to reduce P(FA) to a negligible value, so long as normal approach false alarms do not occur. Normal approach false alarms begin to occur only for large values of T, as was illustrated in Figure 4, and this places an upper bound

on the acceptable T range. According to the SOC plot in Fig. 8b, smaller values of R give higher P(SA) values for a given T when using the climb-only evasion. From Figure 4, smaller R values also allow larger values of T before normal approach false alarms become a problem. Therefore, climb-only safety is maximized for parameter values of approximately R = 350 ft, and T = 21 sec.

Even at this "optimal" operating point, climb-only performance lags behind that of the climbing-turn. Nominal parameter values place P(SA) at 0.994 for the climbing-turn. In contrast, for the climb-only maneuver, the maximum possible value of P(SA) that does not result in false alarms during normal approach is approximately 0.88.

Conclusions

The analysis described above provides some insight into the relative performance of the two evasion maneuvers when used with the current form of the AILS alerting logic. Based on the results, the following conclusions can be made:

1. Over a range of blunder types and flight conditions at the nominal AILS alert threshold parameter values, the climb-only evasion maneuver was observed to produce approximately 38 times more collisions than the turning-climb evasion maneuver.

2. Performance of the climb-only maneuver can be improved by adjusting the alert threshold parameters, but no choice of parameters brings climb-only performance into parity with that of the climbing-turn.

3. Increasing the extrapolation time parameter (T) gives the largest safety benefit, but false alarms during normal approaches begin to occur at values of T greater than approximately 21 sec.

4. The climbing-turn maneuver is less sensitive to pilot reaction time than the climb-only maneuver. The safety level provided by the climb-only maneuver degrades approximately 5 to 10 times more rapidly than the climbing-turn as pilot reaction time is increased.

5. System Operating Characteristic curves show that the climb-only evasion maneuver results in a system that is of little benefit: alerting provides approximately the same level of safety as not alerting. In contrast, with a climbing-turn maneuver, there is a significant safety benefit to producing an alert.

For these reasons, it is believed that the climb-only evasion maneuver will not be adequate to provide sufficient safety at a 2500 ft runway spacing without generating an unacceptable number of false alarms. It must be noted, however, that these results only apply to the current form of AILS alerting logic, which does not modify resolution commands in response to the blundering aircraft's actions. Thus, the evasion maneuvers are performed open-loop. Providing adaptive guidance (such as is used during TCAS Resolution Advisories) may improve safety by strengthening or weakening climb commands as needed.

In any case, the additional separation between aircraft that is provided by a turning maneuver over a straight climb is significant and appears to improve safety by over an order of magnitude. However, specifications for system performance have not yet been defined, and so additional study of the turning-climb maneuver is warranted.

False alarms that occur during normal approaches must be minimized. This requires knowledge of typical lateral deviation, ground track, and turn rate data of aircraft in order to tune the alert threshold parameters. Because normal approaches made up only a small fraction of the total trajectory set used in the analysis described here, further study is needed to more accurately describe the false alarm behavior of AILS.

Acknowledgments

References

[1] Waller, M. C. and C. H. Scanlon, eds., "Proceedings of the NASA Workshop on Flight Deck Centered Parallel Runway Approaches in Instrument Meteorological Conditions", NASA Conference Publication 10191, Hampton, VA, December, 1996.

[2] Federal Aviation Administration, "Precision Runway Monitor Demonstration Report", Document DOT/FAA/RD-91/5, Washington, D.C., February, 1991.

[3] Shank, E. M., and K. M. Hollister, "Precision Runway Monitor", Lincoln Laboratory Journal, Vol. 7, No. 2, MIT Lincoln Laboratory, Lexington, MA, 1994.

[4] Koczo, S., "Coordinated Parallel Runway Approaches", NASA Contractor Report, NAS1-19704 - Task 11, Hampton, VA, May, 1996.

[5] Carpenter, B. D., and J. K. Kuchar, "A Probability-Based Alerting Logic for Aircraft on Parallel Approach", NASA Contractor Report 201685, Hampton, VA, April, 1997.

[6] Kuchar, J. K., and B. D. Carpenter, "Airborne Collision Alerting Logic for Closely-Spaced Parallel Approach", Air Traffic Control Quarterly, Vol. 5, No. 2, 1997.

[7] Radio Technical Committee on Aeronautics (RTCA), Minimum Performance Specifications for TCAS Airborne Equipment., Document No. RTCA/DO-185, Washington, D.C., September, 1983.

[8] Radio Technical Committee on Aeronautics (RTCA), Minimum Performance Standards - Airborne Ground Proximity Warning Equipment, Document No. RTCA/DO-161A, Washington, D.C., May 27, 1976.

[9] Kuchar, J. K., "Methodology for Alerting-System Performance Evaluation", AIAA Journal of Guidance, Control, and Dynamics, Vol. 19, No. 2, March-April, 1996.