GENERALIZED PHILOSOPHY OF ALERTING WITH APPLICATIONS FOR PARALLEL APPROACH COLLISION PREVENTION

by

Lee F. Winder

James K. Kuchar

Department of Aeronautics and Astronautics

Massachusetts Institute of Technology

Cambridge, MA 02139 USA

August, 2000

ICAT-2000-5

Executive Summary

This paper presents the results of research at MIT under NASA grant NAG1-2189 for the period April 1, 1999 to May 31, 2000. The goal of the research was to develop formal guidelines for the design of hazard avoidance systems.

An alerting system is automation designed to reduce the likelihood of undesirable outcomes that are due to rare failures in a human-controlled system. It accomplishes this by monitoring the system, and issuing warning messages to the human operators when thought necessary to head off a problem. On examination of existing and recently proposed logics for alerting it appears that few commonly accepted principles guide the design process. Different logics intended to address the same hazards may take disparate forms and emphasize different aspects of performance, because each reflects the intuitive priorities of a different designer. Because performance must be satisfactory to all users of an alerting system (implying a universal meaning of acceptable performance) and not just one designer, a proposed logic often undergoes significant piecemeal modification before gaining general acceptance. This report is an initial attempt to clarify the common performance goals by which an alerting system is ultimately judged. A better understanding of these goals will hopefully allow designers to reach the final logic in a quicker, more direct and repeatable manner. As a case study, this report compares three alerting logics for collision prevention during independent approaches to parallel runways, and outlines a fourth alternative incorporating elements of the first three, but satisfying stated requirements.

Three existing logics for parallel approach alerting are described (section 3). Each follows from different intuitive principles. The logics are presented as examples of three "philosophies" of alerting system design.

The first philosophy is that in a system with clearly defined normal dynamics, an alert is justified when a clear deviation from normal is observed. This type of thinking is exemplified by the Precision Runway Monitor alerting logic, which issues an alert when an aircraft deviates laterally beyond a threshold distance from its nominal approach path.

Another philosophy stresses justification of alerts through prediction of a specific hazard event. A condition for alerting is that some certainty exists that the alert will not be a false alarm. An example of this type of thinking is the NASA Airborne Information for Lateral Spacing parallel approach logic, which uses an explicit prediction of the future trajectory (when there is no alert) of the approach system to make alerting decisions.

The third philosophy requires that at least a minimum level of safety be ensured for any alerting decision. An alert is forced when the safety level of the available resolution procedure becomes marginal and may become too low with further delay. An example is a probabilistic parallel approach logic developed at MIT. Its alert threshold is based on the notion that alerts should occur when the computed probability of a safe evasion maneuver falls to a minimum acceptable value.

An alerting logic might not be a straightforward application of a single philosophy (section 4). This is because each philosophy focuses on a particular performance metric or metrics (e.g. *false alarm* rate, *hazard event* rate, or *perceived incorrect alert* rate) (section 2), which may not adequately cover the real goals of a particular alerting application. An alerting logic sometimes begins very simply (according to one philosophy), but becomes complex, perhaps needlessly so, as performance issues are addressed later.

An attempt is made in this report to develop a logic for parallel approach collision prevention by considering all performance goals at once (section 5), and then choosing design elements which seem to most directly achieve these goals (section 6). The purpose is not only to shorten development time, but to arrive at approximately the simplest logic that satisfies requirements.

Finally, as an aid to the future analysis of alerting logics based on the normal behavior of an aircraft flying along a planned path, a method is suggested for modeling and simulating this behavior so that performance numbers can be computed (section 7). The method is then used to compute performance metrics for an example threshold based on the first order statistics of the aircraft trajectory model. It is shown that a benefit

(decrease in alerting delay) might be realizable, for a given normal approach alert rate, through an increase in the number of state variables over which the threshold is defined.

Acronyms

AILS	Airborne Information for Lateral Spacing
ATC	Air Traffic Control
DGPS	Differential Global Positioning System
GPS	Global Positioning System
GPWS	Ground Proximity Warning System
IMC	Instrument Meteorological Conditions
INS	Inertial Navigation System
MIT	Massachusetts Institute of Technology
NASA	National Aeronautics and Space Administration
NTZ	No Transgression Zone
PRM	Precision Runway Monitor
TCAS	Traffic Alert and Collision Avoidance System

This paper presents the results of research at MIT under NASA grant NAG1-2189 for the period April 1, 1999 to May 31, 2000. The goal of the research was to develop formal guidelines for the design of hazard avoidance systems. A hazard avoidance system can be described generally as automation that triggers evasive action to prevent rare catastrophes in a human-operated system. Examples in aviation from which insight have been drawn include the existing Traffic Alert and Collision Avoidance System (TCAS), Ground Proximity Warning System (GPWS), and Precision Runway Monitor (PRM), and the proposed Airborne Information for Lateral Spacing (AILS) collision avoidance system for parallel approaches.

1. Introduction

Typically the alerting logic of a hazard avoidance system begins as an intuitive concept, and evolves as inadequacies become apparent through simulation, the input of experts, and actual use. Each development process gives rise to a distinct logic, sometimes dramatically different from others derived for similar applications (e.g. the AILS logic versus the PRM logic, both of which were designed to trigger breakout alerts for aircraft on parallel approach). Of these differences, it is often unclear which are necessary or beneficial due to differences in the applications, and which are the result of subjective choices that became fixed early in the design process. Although each logic is different, consideration of the group of existing logics and their applications reveals general categories of reasoning whose elucidation may simplify and improve the designs of future hazard avoidance systems. Three examples may help to illustrate the different methods.

PRM is a surveillance and collision avoidance system that enables independent approaches in instrument meteorological conditions (IMC) to parallel runways spaced as closely as 3400 ft (Shank & Hollister, 1994). This system employs special air traffic controllers whose purpose is to watch approaching traffic on displays with an alerting capability, and intervene to preclude any collision when a “blunder” occurs. On a PRM controller's display, two adjacent approach paths are shown separated by a strip of forbidden airspace, or “No Transgression Zone” (NTZ) (Fig. 1). If an aircraft is seen to

enter the NTZ (or is predicted to do so within several seconds, depending on display settings), PRM controllers receive an alert requiring communication of corrective maneuver commands to all affected pilots.

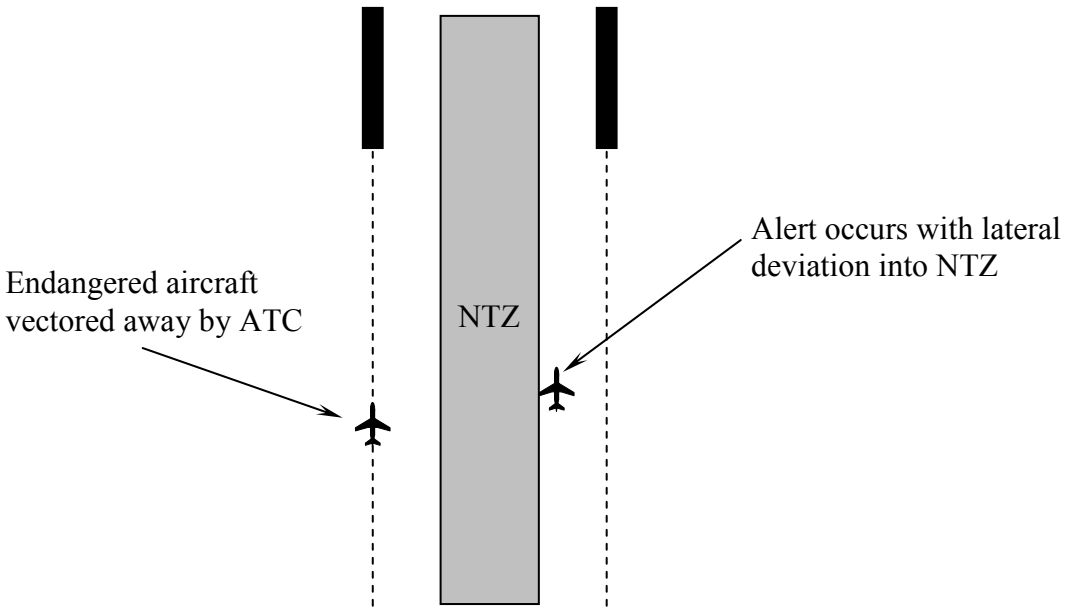

Fig. 1: PRM Alerting Logic Illustrated

The Airborne Information for Lateral Spacing (AILS) logic was developed by NASA Langley with Rockwell-Collins for a proposed cockpit-based parallel approach collision avoidance alerting system (Koczo, 1996; Waller & Scanlon, 1996). It was later modified at Honeywell Technology Center in cooperation with NASA (Samanant & Jackson, 2000). The purpose of AILS was to enable independent IMC approaches to parallel runways spaced more closely than 3400 ft. This was to be accomplished by reducing delays in the blunder detection and alerting process. Radar surveillance and air traffic controller intervention would be replaced by automatic datalink of GPS/INS state data between aircraft, computerized data processing, and alerting aboard each aircraft. After a breakout involving procedural evasion maneuvers, air traffic controllers would intervene to restore normal operation of the system.

The AILS alerting logic is complex, using a combination of approach conformance and trajectory prediction criteria to make alerting decisions (Samanant &

Jackson, 2000). Of importance to this discussion is that the logic can produce an alert whose basis is that one aircraft is specifically threatening another. For such alerts a near collision must be explicitly predicted as shown in Fig. 2. At brief intervals the future trajectories of all aircraft are projected forward in time. Dynamic assumptions for a particular aircraft are dependent on whether or not the aircraft is thought to be blundering from its approach. An aircraft might be judged as blundering based on poor lateral conformance to its approach. Otherwise, each aircraft is assumed in separate cases to be blundering, and the blunder is treated as confirmed when a threshold is crossed with that assumption in effect. To allow for prediction uncertainty, the trajectory model for a blundering aircraft consists of a set of trajectories covering a range of maneuvers. If under the trajectory model a near collision can occur within a limited trajectory projection time T, alerts are generated. The "protection zone" defining a near collision is a volume described by an elliptical area in the horizontal plane and centered about the endangered aircraft, and by bounds on relative altitude. A set of nested alerting thresholds (occurring in a particular sequence as an encounter occurs) are defined for different values of T and protection zone dimensional parameters. Depending on the urgency of the situation, the involved pilots receive either preliminary advisories or final breakout commands from their respective cockpit alerting systems. The breakout procedure involves a 45° turn away from the adjacent centerline and a simultaneous pull-up to a prescribed final climb rate.

An alternative to AILS was developed at the MIT Aeronautical Systems Laboratory (Kuchar & Carpenter, 1997). Breakout alerts are again issued directly to the pilots of aircraft, but are based on the estimated safety level of a procedural evasion maneuver. The metric of safety is the probability of a collision during a procedural turn-with-climb evasion. In an ideal implementation, this metric is recomputed in real time (e.g. via Monte Carlo simulation) at brief intervals (Fig. 3). Note that while the AILS logic involves simulation of non-alert trajectories (that is, the trajectory occurring when there is no alert), the MIT logic simulates post alert trajectories (that is, the trajectory occurring when an alert is issued). As evaluated during its brief development, one aircraft, identified as the host, is modeled as performing a perfect evasion maneuver from its current state, while another is modeled as a potential blunderer, following a variety of

trajectories according to its measured initial state and probabilistic weightings. A collision is approximated as a three dimensional separation of 500 ft or less, and must occur within a limited trajectory projection time. During normal approaches this probability remains virtually zero and an alert is unlikely, provided the runway spacing is large enough relative to normal trajectory deviations from the ideal path. If the probability of a collision reaches 0.001, the evasion is deemed necessary under the reasoning that such risk is marginally acceptable and that the probability of a safe evasion might decrease if there is any further delay before alerting. Otherwise, the alert is deferred to minimize the likelihood of a false alarm.

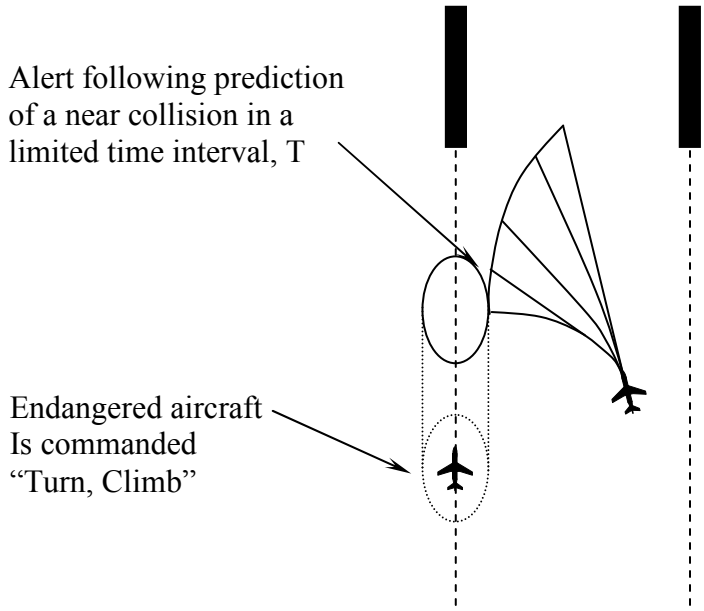

Fig. 2: Partial AILS Predictive Alert Logic

These examples represent three distinct approaches or "philosophies" of decision making that appear to encompass most existing or proposed hazard avoidance logics: alerting when the human-controlled system fails to conform to established procedure (PRM), alerting only if a specific hazardous event may occur if no intervention takes place (AILS), and requiring preservation of options allowing a safe resolution to an alert (MIT). A more detailed discussion of each philosophy will follow. But first some recurrent terms and concepts will be discussed.

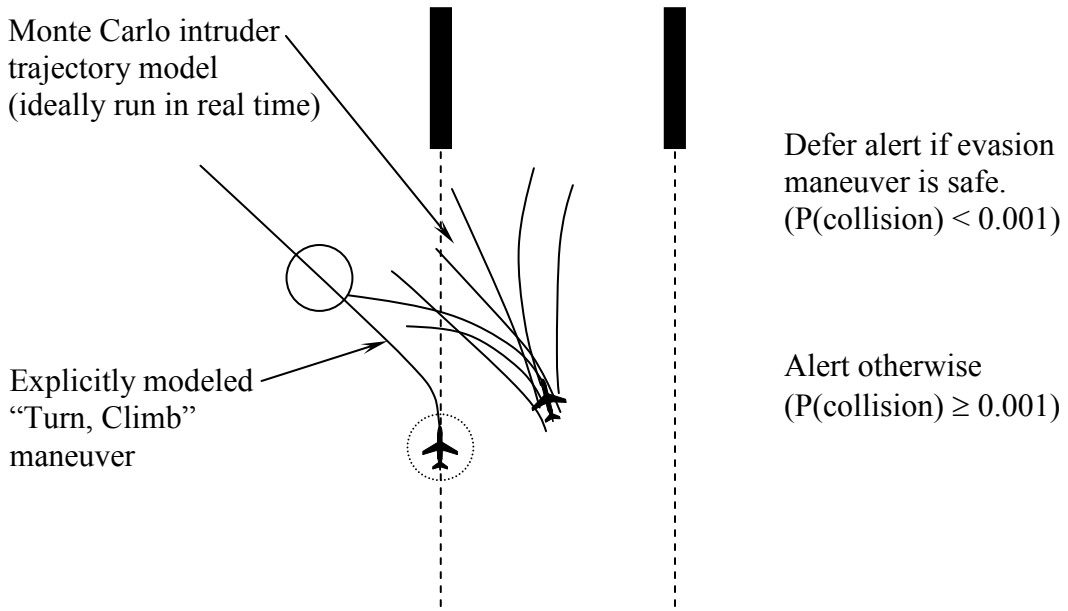

Fig. 3: MIT Parallel Approach Alerting Logic

2. Alerting Background

An alerting system is designed to prevent occurrence of a catastrophe through timely warnings issued to human *operators* within a larger system. There may be several operators simultaneously controlling a system and subject to distinct warning messages for a given alert. In this paper, an *alert* refers to the output event or trajectory (input to the operators) of an alerting system, beginning at a particular time and resulting in altered system dynamics. An example alerting system is the TCAS collision avoidance alerting system for aircraft (Williamson & Spencer, 1989). Because TCAS alert messages are coordinated between aircraft, the TCAS hardware on a single aircraft is not a complete alerting system, but part of a larger alerting system including all suitably equipped aircraft. A TCAS alert can result in distinct messages to different pilots, or no message at all. Also, a TCAS alert generally consists not of a single event but of a sequence of messages delivered to each pilot.

The logic of an alerting system is described in terms of *state variables*. State variables are measurable quantities that aid in describing the state of the larger system of which the alerting system, operators, and environment are parts. They must be

observable by the alerting system. Examples are continuous variables such as position, speed, acceleration and physical dimensions, and discrete variables that might describe different modes or configurations.

In this report a *hazard* refers to a set of system states, any of which is tantamount to a catastrophe, which is assumed to be avoidable through alerting. In this paper there is no such thing as a "soft hazard," or state that is undesirable but not necessarily catastrophic (Kuchar & Hansman, 1995). Note that there may be certain catastrophes that an alerting system is not designed to prevent. For example, in parallel approach alerting, if it is assumed that an aircraft may blunder in a way that makes it unresponsive to cockpit alerts, a terrain collision by that aircraft is not considered as part of the hazard for the purpose of the alerting system design. Flight into terrain or other objects surrounding the airport by an evading aircraft must be considered as part of the hazard, because an alerting system designed in the absence of such events may induce them out of blindness. Thus, one generally should not declare the hazard as a specific type of incident (e.g. mid-air collision between two aircraft) according to the main purpose of the system, and then neglect other types of incidents for the purposes of design.

Some alerting logics employ predictive state *trajectory models* in decision making. Such a model allows the logic to judge the likelihood or possibility of a future event, such as a hazard, given a specific alerting system action. Trajectory models take a number of forms. Here they are divided into three groups: single trajectory prediction, worst case, and probabilistic (Kuchar & Yang, 1997).

The simplest type of prediction is a single trajectory beginning at the current estimated system state (Fig. 4a). This model may be based on knowledge of intent, or lacking that, on a constant highest derivative assumption—that independent state variables remain constant at the most recent measured values (Kuchar & Hansman, 1995). An argument for the latter assumption is that if state variables are known to have continuous and differentiable trajectories, the constant derivative assumption gives a first order approximation that is good for a limited time.

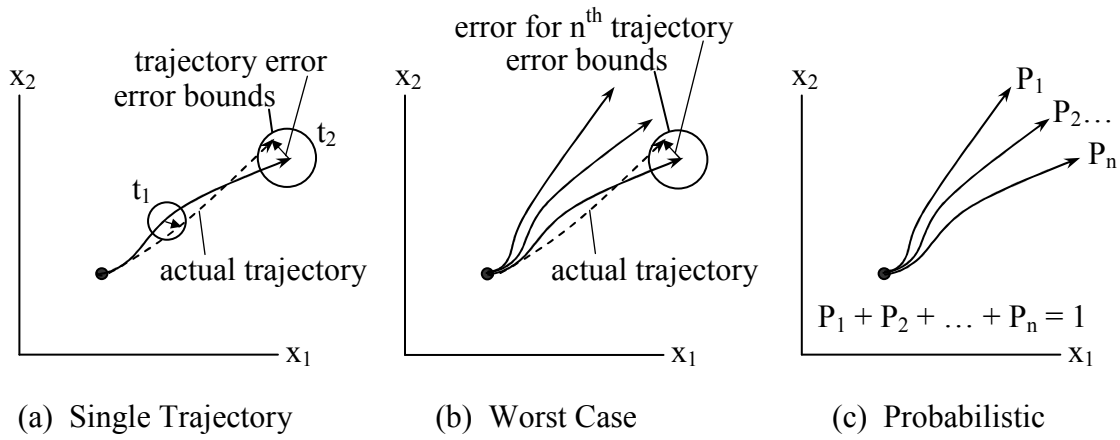

Fig 4: Trajectory Prediction Methods

A single-trajectory prediction model does not include a description of prediction uncertainty. As shown in Fig. 4a, such a model (solid line) is in typically in error at any given point in time with respect to the actual trajectory (dashed line), and the magnitude and direction of error varies with time in the state space. If a constant derivative assumption is made, error magnitude would be expected to increase with time. For the model to be useful, trajectory error should remain within reasonable bounds of magnitude over some interval of time. As an example, error magnitude bounds are drawn as a circle whose radius varies with time. In an alerting system employing a single-trajectory model, trajectory uncertainty might be handled through a consciously conservative hazard description (e.g. defining a collision between two vehicles as a separation below some distance that is large relative to actual vehicle size). The modifications to the hazard definition correspond to the expected bounds on trajectory error.

If uncertainty about the future trajectory is large initially or increases quickly with projection time, multiple trajectories may be used to cover the range of possibilities—that is, to provide a set trajectory predictions such that the actual trajectory lies within specified error bounds of at least one of the model trajectories (Fig. 4b).

If no probabilistic weightings are assigned to its elements, the trajectory set is termed here a *worst case* model. "Worst case" may be a misleading choice of words in that the desired trajectory set might reflect probabilistic or intent knowledge about the

system. In the limit, a worst case model of an aircraft trusted to follow a normal approach might consist of a single trajectory (making it equivalent to a single-trajectory model), even though large deviations from the approach path are dynamically possible. Thus, a worst case model might more descriptively be termed an *unweighted trajectory set* model.

In general the longer the attempted time projection, the more complex the description of the worst case set becomes for a given maximum error, and the more difficult it is to simulate the set in an alerting algorithm. Therefore an issue usually exists in choosing when to cut off trajectories in a worst case prediction model. Even if a particular worst case model can be simulated arbitrarily far into the future, the model may tend to have diminishing value for increasing prediction time due to the increasingly large set of possible system states (i.e. it can become difficult to rule out or guarantee future occurrence of a given event).

If a probability function is defined over a worst case trajectory set (that is, each element is assigned a probability and the probabilities sum to 1), where the set is considered an event space, then the whole is termed a *probabilistic* trajectory model (Fig. 4c). This additional information allows computation of the probability of a particular event (e.g. hazard) occurring within the limited time of the model, whereas when using a worst case model only a statement of whether or not an event is possible can be made.

Different verification requirements apply to probabilistic versus worst case models. Whereas use of a worst case model requires belief that trajectory error lies within acceptable error bounds for an element of the trajectory set, use of a probabilistic model requires belief that computed probabilities are within acceptable error bounds.

Alerting system performance is often quantified in terms of the rates of hazard and false alarm events. These are defined below. In addition, a third event type, the "perceived incorrect alert" is suggested.

A *hazard event* or *hazard encounter* occurs any time the system trajectory encounters (ends in) a hazard state. Thus, the alerting system may have failed to issue needed alerts, issued a late alert, or even induced the hazard through unnecessary alerting.

A *false alarm* occurs if the alerting system issues an alert that is not needed to prevent a hazard event. In real life it may be difficult to say whether an alert that has occurred is a false alarm, because the opportunity to observe the non-alert trajectory of the system is lost when the alert occurs, and there is usually uncertainty in predicting what the system would have done. The frequency or probability of false alarms has sometimes been estimated for a given alerting system by introducing a probabilistic model of the system dynamics, in which the behavior of the system before and after alerts is explicitly defined. Note that a false alarm is not mutually exclusive of a hazard event (i.e. in the case of an induced hazard).

An alerting system action (alert or non-alert) is *perceived incorrect* if an operator believes immediately or in retrospect that a better decision should have been made with available information. This can occur when an operator decides that a false alarm has occurred when there was insufficient risk of a hazard occurring nominally, believes that an alert was necessary but finds commanded maneuvers unsafe, is aware of a past alert that induced a hazardous event, or believes that an alert failed to occur when it was necessary. Thus, such events can in principle range from annoying or disruptive false alarms to disasters blamed on the alerting system. It is important to distinguish the other two alerting event types, false alarms and hazard events, from perceived incorrect alerts. While either of the former two events can also fall into the third category, one can imagine cases where a false alarm, perhaps even a hazard event, is not perceived as an alerting system failure. In addition it may be possible for an outcome that is neither of the first two event types to be considered an alerting system failure. For example, an operator could mistakenly consider an alert an unjustified false alarm when it is not a false alarm at all. The key word in this discussion is *perceived*. Perfect knowledge of the categorization of an event in terms of the first two event categories is insufficient for determining whether it is also of the third category.

Assuming that operators *initially* have high confidence in an alerting system's potential, an accumulation of perceived incorrect alerts can be blamed for any later reduction in operator confidence in an alerting system.

3. Three Philosophies of Alerting Logic Design

Three common philosophies of alerting logic design have been identified and in section 1 were related to existing or proposed systems for parallel approach collision prevention. Following is a more detailed and general description of each philosophy.

3.1 Trajectory Conformance Monitoring

This type of logic uses non-conformance of a system to established procedures as a basis for alerting. For example, Fig. 5 shows a system state with respect to a normal operating region in state space. If the state exits outside the normal operating region an alert is issued. The system exists to prevent occurrences of a hazard, but no explicit prediction of a hazard event is required for triggering an alert. As shown, the normal region is defined so as to be mutually exclusive of the hazard, even though the hazard is not explicitly modeled in the final algorithm. In PRM, for example, as long as both aircraft remain outside the NTZ, the hazard cannot occur. An aircraft entering the NTZ will trigger an alert whether or not it actually threatens another aircraft.

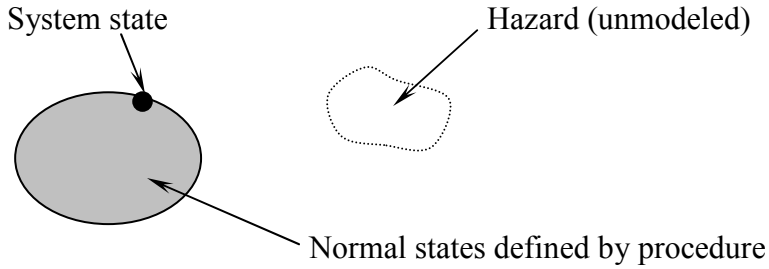

Fig. 5: Trajectory Conformance Alerting

A deviation ("blunder") from the normal procedure is a necessary precursor to a hazard event, so it can be argued that an observed deviation from normal is sufficient reason for an alert and corrective action, provided such a policy does not result in a high

rate of alerts occurring without a blunder. Frequent false alarms during normal system operation would come to be perceived as incorrect by operators, and might in time cause operators to ignore or delay responding to alerts.

In addition to establishing that the non-blunder alert rate is acceptably low, it should be shown that when a blunder does occur there will be an evasive maneuver having an adequate likelihood of success. Such an analysis typically involves a reference dynamic model of the system, and iterative adjustment of the threshold. Because of the dependence of the threshold on the operational procedure, it may be necessary to adjust the procedure itself to achieve performance goals. For example, it was concluded that PRM could be used with parallel runways spaced no less than 3400 feet apart because below this spacing the likelihood of safe resolution of a blunder was too low in simulation studies.

3.2 Nominal Trajectory Hazard Prediction (False Alarm Prevention)

This alerting strategy involves continuous checking for a particular hazard through explicit prediction of the non-alert, or nominal, system trajectory (Fig. 6). For an alert to occur, the hazard event must be predicted. Under this philosophy, the logic avoids alerts that are not clearly justified with respect to the hazard. The hazard is described in terms of a set of state variables composing a state space. The trajectory model, which might be probabilistic, worst case, or a single predicted trajectory, is propagated forward in this state space from the current, measured location. In the Fig. 6 example, the trajectory model is a worst case model, and is predicting that a hazard may be encountered in the future.

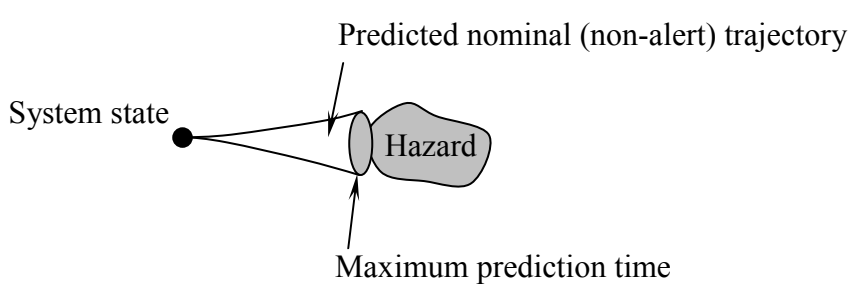

Fig. 6: Nominal Trajectory Prediction Alerting

Due to uncertainty in prediction and the consequent possibility of false alarms, it is insufficient to define the alerting rule as "alert when a hazard is predicted." Typically an additional metric or metrics are required for threshold definition. Possible metrics include the degree of certainty in occurrence of a predicted event (for a probabilistic trajectory model), or the predicted time-to-collision (for worst case or single-trajectory models). The values of threshold parameters must be chosen to satisfy both safety and false alarm goals. In the illustrated example, collision prediction time is the metric and a particular value of this must be selected to define the threshold. Using a reference model of the behavior of the entire human-controlled system (able to describe its dynamics both before and after an alert occurs, and covering all possible initial conditions in state space), optimal threshold parameters are determined, typically through repeated Monte Carlo simulation and adjustment (Yang & Kuchar, 2000).

Whether a hazard is imminent for the nominal trajectory is not a direct indication of whether an evasion maneuver will be safe. In this type of logic, the justification of alerts is inherently stressed over the safety of the alerting decision.

In the course of analysis, the complexity of the logic may increase to cover special cases that were not initially foreseen. This is likely when few state variables are available for measurement or the actual system dynamics are not well understood. An example is the development of TCAS logic for midair collision prevention. This logic began with a simple range rate and time-to-collision prediction model (with adjustable parameters for the threshold prediction time and miss distance) characterized by large trajectory errors, and was eventually augmented with conditional statements and new parameters in order to handle problem scenarios (Drumm, 1996). For example, a situation where two aircraft unknowingly fly parallel at the same speed may be unacceptable, yet trigger no alerts when using a time-to-collision criterion only. To cover such problem scenarios additional checks were added to the logic.

Other examples of logics that use explicit incident prediction as the basis of thresholds are GPWS and AILS.

3.3 Existence of Safe Escape Options (Safety Monitoring)

In general there may be specific completion conditions that must be met in order for a potential incident to be considered resolved, and it is possible to make deferral of alerts conditional on the predicted attainment of such conditions. For example, the MIT logic issues alerts based on knowledge that a collision will probably not occur within a certain period of time following the alert.

This type of logic is superficially similar to the nominal trajectory hazard checking method described in section 3.2 in that it involves a trajectory model. As illustrated in Fig. 7, a hazard event is once again defined in terms of measurable state variables. A trajectory model is used to propagate the system state, but this time under the assumption that an alert has occurred or will occur at a particular time, resulting in escape maneuvers. In general there may be multiple maneuver options (represented by evolving state envelopes—each resembling a horn—in Fig. 7), corresponding to different warning inputs that can be issued to operators. Completion conditions are defined in terms of the evasion trajectory and state variables. As shown, completion conditions may require that the system reach a specific region in state space. In addition, it may be required that the system reach the completion state set within a particular time interval. Finally, the completion state set cannot intersect with the set of hazard states—given that the hazard is a "catastrophe," it cannot be considered as part of a desirable alerting outcome.

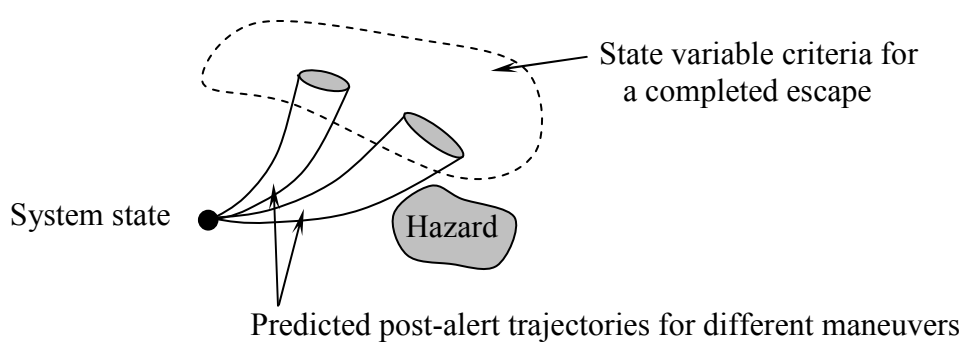

Fig. 7: Ensuring that Safe Options Exist

If a probabilistic trajectory model is to be used, then alerting decisions will be based on the probability of reaching the completion state set within the required time interval. Therefore an additional component of completion is a threshold probability, such that above this an alerting decision is considered "safe." If the probability is equal to the threshold value then safety is "marginal."

If the trajectory model is worst case or a single trajectory, safety requires that all trajectories for an alerting option reach the completion set within a given time interval. Safety is marginal if any one of the trajectories reaches a boundary value of the completion state set or allowed time interval.

According to this philosophy an alert may be deferred as long as an available alerting option is safe. An alert can no longer be deferred when safety becomes marginal. In other words, an alert is considered justified when there may be no safe option remaining at the next alerting opportunity.

In this method safety is fixed at the threshold, resulting in a loss of direct control over false alarms. This is because whether a post-alert maneuver is safe is not a direct indication of whether the nominal system trajectory is safe (i.e. whether an alert will be a false alarm). For example, it may be possible for an evasion option to become marginally safe, triggering an alert, even when no hazard would be encountered on the nominal trajectory.

As part of an alerting option, it may be desired that certain operators within the system continue along predictable nominal courses—either unaware that an alert has occurred, or under advisement not to deviate from the nominal path. In this case the relevant nominal trajectories would be modeled, and their safety monitored. This should not be confused with nominal trajectory-based alerting, because the existence of a hazard along the nominal trajectory is not required for triggering the alert.

For the MIT logic a successful alerting outcome was defined (implicitly) as any case where a collision failed to occur within a fixed time after the alert. More stringent completion conditions could also have been used, such as to require a minimum

separation, divergence rate, heading difference, etc. between aircraft within the limited time interval for which dynamics are modeled. As they are, the MIT conditions do not preclude the occurrence of collisions immediately beyond the prediction time limit, and pilots or air traffic controllers might reasonably object to an alerting system that makes no guarantees about the "resolvability" of the post-alert situation. As shown in Fig. 8, more stringent completion requirements may translate into a smaller set in state space or more limited time requirements, and thus may cause the alerting system to encounter a marginal safety condition earlier than it would have otherwise. So although they increase confidence in the safety of an alerting outcome, more stringent completion requirements can also increase the likelihood of false alarms.

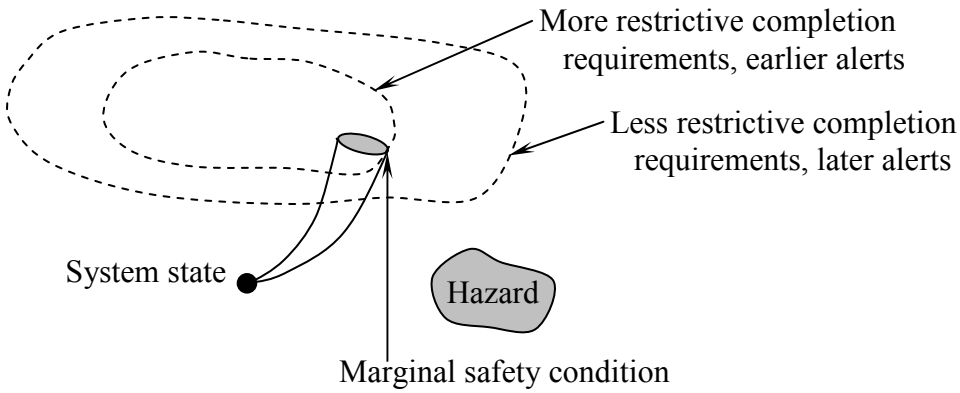

Fig. 8: Effect of Completion Requirements on Alert Deferral

4. Combined Philosophies

An alerting system must be made to satisfy performance goals that are independent of the preferred philosophy of a particular designer. Depending on the philosophy, satisfying performance goals may require extensive modification to the initial design. If important issues are not addressed at first, the resulting performance deficiencies can be eliminated in an *ad hoc* fashion (Kuchar, 1996; Yang & Kuchar, 2000). An adequately functioning logic may be attained in numerous ways, but there is no guarantee that separate logics that are equivalent in performance are also equal in simplicity or understandability.

Each philosophy can be thought of as emphasizing different components of alerting performance. A conformance-based logic has an alert threshold requiring deviation from established normal system dynamics (e.g. operating procedures). It is conceptually simple enough to promote operator belief in the appropriateness of alerts when they occur (though not necessarily in the particular resolution commands chosen), provided it is tuned to minimize normal approach false alarms. It could be said to emphasize minimization of perceived incorrect alerts, but such a logic does not *inherently* ensure that alerts are safe or that they are not false alarms, and successful avoidance of perceived incorrect alerts may require avoiding hazard events or false alarms. A nominal trajectory prediction-based logic provides confidence that an alert is not a false alarm. But based on this trajectory model alone, no direct determination of whether an alert will be safe or how it is perceived can be made. A post-alert trajectory-based logic ensures that alerts occur when they are likely to be successful. But there is no automatic guarantee that such alerts are not false alarms or are perceived as correct.

An example of a logic developed through conventional methods is TCAS. The initial TCAS concept was to issue alerts according to a predicted time-to-collision threshold (with nominal trajectory prediction). The trajectory prediction is made using the approximation that the relative range rate between two aircraft remains constant at its current estimated value. In an attempt to compensate for uncertainty in this single-trajectory prediction, the range defining a mid-air collision is made large relative to aircraft size. This model does not explicitly account for operating procedures of en route flight, but rather employs a more generic prediction with knowingly large uncertainty. In this initial form TCAS exhibits alerting defects that have resulted in various modifications to the logic and its operating environment. To reduce the occurrence of perceived false alarms and other alerts perceived as incorrect in specific situations (a priority of conformance philosophy), threshold parameters have been readjusted and conditional statements have been added. In addition, TCAS has been aided through modification of normal operating procedure. For example, pilots are discouraged from ascending or descending at high vertical rates just before leveling off, in order to prevent TCAS false alarms that can otherwise result (Mellone & Frank, 1993). In large enough numbers, such alerts are likely to reduce pilot conformance to alert commands.

After TCAS deems an alert necessary, evasion maneuver commands are determined by an addition to the logic. By examining a predefined set of maneuver options, this logic seeks cooperative maneuvers that result in at least the minimum acceptable closest approach distance, and therefore no collision if maneuver commands are respected. In this way the safety of a TCAS alert is enhanced beyond what would be possible based on the nominal trajectory prediction if it were arbitrarily paired with a procedural evasion procedure. It should be noted, however, that because the safety of alerts is not a *precondition* of alerting, situations could conceivably arise where no adequate maneuver exists after an alert occurs. So TCAS cannot yet be thought of as fitting the post-alert prediction philosophy—this would require explicit proof that a safe maneuver will always accompany an alert.

The PRM system currently provides adequate safety and normal approach false alarm performance with a simple procedure conformance threshold. The system is operated by air traffic controllers, who ultimately decide which aircraft must be vectored away from approach when a blunder occurs. Proposed future alerting systems (AILS, MIT) require that pilots perform procedural evasion maneuvers when receiving alerts. A problem arises when trying to adapt the PRM alerting logic for use with such procedural maneuvers. An aircraft blundering into the No Transgression Zone could be made to trigger breakout maneuvers by all aircraft on the opposite side of the NTZ. This is probably not the safest action, because it creates a complex traffic problem for air traffic controllers to resolve later. The safest solution may be to break out those aircraft in close proximity to the blunderer, and to allow others to continue on uninterrupted to landing. This minimizes the number of aircraft requiring special recovery guidance. Such a preferred outcome can be stated explicitly in terms of measurable state variables, as described for the post-alert trajectory prediction philosophy. By considering the set of all possible evasion maneuvers for all aircraft in a system, an alerting threshold based initially on a conformance threshold could find a post-alert maneuver solution meeting the explicitly stated requirements.

A strict application of any one of the three discussed philosophies is unlikely to satisfy performance requirements—modifications to each basic concept are needed.

Rather than produce an initial design according to one philosophy and make later adjustments that amount to the addition of properties of the other two philosophies, it may be simpler and more insightful to begin with an approach that combines philosophies.

5. Automatic Alerting for Independent Parallel Approaches

Automated alerting systems have been investigated in recent years as a means to maintain safety for independent IMC parallel approaches at reduced runway separation. The more closely spaced the nominal traffic streams are, the less reasonable it is to expect air traffic controllers to monitor traffic and intervene in the case of an approach "blunder." A blunder refers to a failure in guidance of an aircraft (possibly due to human error, hardware failure, etc.) that could result in large path deviations and an eventual collision with another aircraft. Depending on the suddenness of a deviation and the separation of the parallel runways, there may be too little time from start of a blunder to collision for a controller to detect the blunder, react and communicate with an endangered aircraft, and for the pilot of the endangered aircraft to initiate an evasion maneuver. An automated alerting system may reduce the total delay between blunder initiation and pilot notification, allowing a smaller runway separation for a given level of safety.

The earliest automatic alerting system for parallel approaches was embedded in the controller displays of PRM. PRM is a conservative system that achieves alerting delay reduction through an incremental improvement in surveillance technology and an increase in the number of ATC personnel. It was later suggested that the air traffic controller's role in initial blunder detection could be fully automated and the alerting system implemented as cockpit hardware in a manner similar to TCAS, further reducing delays. This led to the AILS and MIT prototype logics. Both logics seek further performance gains through use of a more complete set of state information. Whereas the PRM alerting system works with horizontal position measurements only, the AILS and MIT logics are designed to make use of direct vertical position, heading, velocity and bank measurements as well.

Before attempting any comparison of logics, some basic assumptions about the approach system and alerting performance requirements will be made clear in the following two sections. These are thought to agree with assumptions stated or implicit in all recent research into independent parallel approach alerting systems, and will not be defended here.

5.1 Assumptions

- The system of concern is a pair of parallel runways operating independently. There is no limit on the number of aircraft on approach to each runway at a given time, though normal in-trail spacing requirements apply.

- The hazard is defined as a mid-air collision or near miss (center of mass separation less than 500 ft) between any two aircraft.

- Any aircraft may commit a blunder. A blunder is defined here as a discrete, randomly occurring transition in aircraft dynamics. It may be characterized by unusually large state deviations from the normal aircraft path, and by a lack of responsiveness to air traffic control or alerting system commands. Such a blunder definition is consistent with assumptions made during evaluation of the AILS logic (Winder & Kuchar, 1999).

- The probability that an aircraft will commit a blunder during an approach is small. Although situations may exist where simultaneous blunders could occur, for simplicity blunders are assumed independent, and the probability of simultaneous blunders by two or more different aircraft is assumed negligibly small.

- A non-blundering aircraft is assumed responsive to air traffic control commands and to automatic alerts, assuming the pilot trusts the alerting system.

- A pilot will trust the alerting system if perceived incorrect alerts are rare.

5.2 Requirements

- Under normal (non-blunder) operation the per-approach probability of a mid-air collision must be negligible. In other words, no system behavior resulting in a catastrophe could reasonably be called normal. In effect, for a collision to occur, a blunder must first occur.

- Air traffic controllers will receive alerts when pilots do. Controllers will intervene at some point after automatic alerts have occurred, and return the system to normal operation, avoiding hazards.

- The needs of air traffic controllers must be provided for. This includes guaranteeing an acceptable post-alert system configuration. For example, the number of aircraft receiving simultaneous breakout maneuver commands should be limited so that controllers are not overwhelmed. Also, aircraft should be adequately separated, not rapidly converging at intervention time, etc.

- The probability of reaching a safe completion must not fall below some minimum level.

- The pilot of a non-blundering aircraft will remain in control during emergency maneuvers.

5.3 Comparative Discussion of the Existing and Proposed Logics

The three alerting logics (PRM, AILS, MIT) will now be discussed and compared in light of the assumptions and requirements just described.

The PRM logic was designed as an aid for ground-based controllers and not as an autonomous cockpit alerting system (which makes direct comparison of it with the later AILS and MIT logics unfair, but it will be done anyway). As such, it has no ability to directly sense that a particular aircraft may be in danger due to a blunderer, or to choose appropriate and distinct maneuvers for each aircraft. Note however, that this requirement could be met with a simple modification, such as to add a longitudinal distance-from-

blunderer criterion for an aircraft to be judged endangered, and to assign a procedural breakout maneuver to such aircraft and none to others. The longitudinal separation criterion is likely similar to what a human air traffic controller would employ in this situation, meaning that the original logic's appeal to human intuition would not necessarily be damaged by this change.

Assuming for a moment that a PRM alert triggered a procedural turn-with-climb evasion maneuver for aircraft near a blunderer, there is still no certainty that responding to the alert ensures any minimum level of safety. Whether safety can truly be guaranteed at alert time depends on the nature of analysis applied to the threshold. In the past, PRM has been evaluated through simulation of a set of entire blunder trajectories and their outcomes with a given alerting system, and calculation of a number representing the overall safety of the approach system with the alerting system (Shank & Hollister, 1994). This number does *not* describe safety for a given alert scenario—rather, it is an average level of safety over all alert scenarios (Yang & Kuchar, 2000). Alternatively, the post-alert safety level could have been evaluated with a post-alert trajectory model initialized at each point on the alerting threshold (i.e. a trajectory model conditional on an given alert state, and therefore containing additional information), to demonstrate that a minimum safety level is met at every threshold point. This type of analysis does meet the safety requirement specified in the assumptions. Note that satisfaction of the safety requirement does not necessarily entail explicit trajectory modeling within the alerting algorithm, only special analysis of the conformance threshold.

PRM has been criticized primarily for two reasons: it employs humans as monitors, a job for which they are notoriously ill suited; and it bases alerting decisions on horizontal position measurements (via radar) and quantities derived therefrom. The proposed AILS and MIT logics eliminate the additional human monitors and incorporate multiple new state variables that are made directly available through datalink technology. It has been argued (Kuchar, 1995) that the addition of state information allows a more accurate prediction of the future trajectory of a system, resulting in improved alerting performance for predictive methods. For example, with a position and velocity estimate one might predict straight-line motion of an aircraft. Given a bank estimate in addition, a

curved trajectory can be predicted. Both AILS and the MIT logic use enhanced state information to construct a predictive trajectory model.

The MIT prototype logic is fundamentally different from PRM in that it uses an explicit trajectory and hazard prediction to make an alerting decision. In an attempt to describe the inherent uncertainties of trajectory prediction it employs a trajectory set and probability function defined over this set. The need for an alert is judged by comparing the estimated probability of a near collision during a breakout maneuver with a threshold value of 0.001. Thus, it conforms to the evasion safety assumption by attempting to ensure a specific minimum safety level at the time of alert. Unfortunately, the method used to achieve this goal is difficult to justify, given that the lack of real world blunder data makes verification of any probabilistic blunder model difficult. Specifically, we have no means of computing error bounds on the hazard probabilities produced by this model so that the appropriateness of the model can be judged relative to other trajectory models.

The MIT logic's definition of a safe evasion may be unsatisfactory to an air traffic controller who must rescue the system after an alert. The logic requires only that a near collision not occur within a limited prediction interval. A controller might desire additional assurance that evasive maneuvers will guarantee a minimum separation or divergence rate (for example) between aircraft at the time of intervention.

Like PRM, the MIT logic is incomplete with respect to the listed assumptions and requirements, but in a different way. It was tested only for the simplified case of exactly two aircraft on adjacent approach paths, with one aircraft the *a priori* blunderer. This greatly simplified the trajectory modeling problem, and eliminated the need to selectively assign evasion maneuvers to multiple non-blundering aircraft. Recall that the desired logic must operate under the assumption of arbitrarily many aircraft, any of which may blunder. It is not clear how, further evolved, the MIT logic would have handled the inevitable relaxation of assumptions.

The AILS logic is unique at this point in judging the difference between a blundering and an endangered aircraft, and providing a distinct alert message

accordingly. An aircraft is judged to be blundering if a worst case trajectory model predicts a near collision within a limited time with another aircraft, under the assumption that the first aircraft is blundering and the endangered aircraft continues a normal approach. AILS may also pre-judge an aircraft the blunderer if it exceeds lateral or vertical bounds on positional deviation from the nominal approach path (i.e. a conformance check).

The AILS logic is structured to prevent false alarms. To accomplish this while simultaneously providing alerts when they are needed, there must be little uncertainty in predictions of the nominal system trajectory. In terms of worst case trajectory modeling (as used by AILS), this means that only a limited error can be tolerated between the actual system trajectory and any element of the worst case nominal trajectory set. Thus, a worst case set could be valid in the sense that at least one element will satisfy error requirements with respect to the actual trajectory, but not suitable for achieving performance requirements due to excessive divergence of the trajectory elements with time. AILS was designed under the assumption that blunder trajectories fit a specific pattern: namely coordinated turns from the approach centerline, possibly followed by rollout into straight-line flight. Assumed vertical motion is along the same lines, involving a brief interval of vertical acceleration coincident with the initial horizontal deviation. Knowing this, more complicated blunders involving multiple heading or vertical rate changes were neglected in the AILS trajectory model. (A few of the more complex maneuvers were considered in evaluation simulations, but were not of a magnitude that could produce a collision with a normally approaching or evading aircraft). Clearly, success of the logic in actual use depends on the validity of assumptions about the form of blunders. The success (acceptability of false alarm and safety performance) of AILS during evaluation simulations is linked to the fact that simulated blunders were exactly those for which the nominal trajectory prediction model had been intended. The assumed blunders are possibilities but there is currently no means of confirming that they are fully representative of those an implemented AILS system would encounter in use.

It is suggested in this report that false alarms are not the culprit in generating mistrust so much as perceived incorrect alerts (including normal approach false alarms). A pilot or air traffic controller may consider an alert correct despite being a likely false alarm if an aircraft has clearly strayed beyond the bounds of normal procedure. Furthermore, false alarms during blunders would be expected to occur at a small rate compared with normal approach false alarms due to the rarity of blunders relative to normal approaches, and thus reduction of this type of alert has little effect on observable false alarm performance. In other words, the AILS design may be stressing minimization of the wrong performance metric.

A type of false alarm that is certainly important to avoid is that occurring during apparently normal system operation. Because the parallel approach system spends the vast majority of time operating normally, such alerts may tend to occur frequently relative to justified alerts—too frequently for operators to believe that the alerts are justified. Such alerts will be perceived as incorrect with time, and later ignored, possibly along with the rare alerts that are justified. In order to minimize such alerts the AILS, MIT and PRM logics must be evaluated using actual or simulated normal approaches, and parameters adjusted.

Unlike the MIT logic, but identically to PRM, an alert from AILS carries no assurance of a safe outcome given a particular system state at the time of the alert. Safety is computed for AILS using a Monte Carlo blunder simulation resulting in overall safety numbers, which as already mentioned, do not apply to individual alert scenarios.

In conclusion, none of the three parallel approach logics considered satisfies all assumptions and requirements that have been put forth in this report. To do so a new logic will be suggested using ideas from the existing logics.

6. Outlining a Logic to Satisfy Requirements

The point of this section is not necessarily to describe a unique logic, but one matching the assumptions and requirements from sections 5.1 and 5.2 with approximately the minimum complexity. Each of the three logics ("philosophies") already considered

satisfies some, but not all items on the list. Desirable properties of each may be borrowed.

The conditions for a successfully completed evasion should be stated explicitly in terms of state variables. Recall that a successful completion in this problem refers to a post-alert system configuration that is acceptable to the air traffic controllers and pilots in a parallel approach system, allowing a comfortable transition back to normal operation after an alert has occurred. Considerations include the locations of fixed hazards or traffic streams in the airport area, the time it takes for controllers to understand the situation after an alert occurs, the number of aircraft that may acceptably be broken out simultaneously, their allowed separations, rates of convergence or relative headings, etc. Clearly, this is a subjective and complicated matter, but it is one that must eventually be thought through in any case.

Due to the requirement that post-alert maneuvers ensure safety of a non-blundering aircraft, the logic must incorporate a means of judging the safety of maneuver options—perhaps including the nominal trajectory (non-alert) option for some aircraft. This implies the need for a post-alert trajectory prediction model similar to that of the MIT logic. Generally this includes both models of the nominal and breakout trajectories for the individual aircraft. One might interpret the need for a nominal as well as a commanded trajectory model for each aircraft as a need to combine attributes of the AILS and MIT alerting philosophies (nominal and post-alert trajectory prediction). But this can be a misleading comparison in that we are not requiring that the alert be *justified* by application of the nominal trajectory model, only that the nominal trajectory be safe if it is to be considered a viable post-alert option for an aircraft. In this way the suggested logic is still purely post-alert prediction-based. The nominal aircraft trajectory is considered here as just another "evasion" option for a particular aircraft in the case that an alert is being issued for the approach system. To repeat, we generally require both a nominal and commanded (where the command could be for a pilot to follow the normal approach) trajectory models for each aircraft so that the safety of each option can be estimated. Multiple options are desired under the reasoning that of a larger set of options,

there is more likely to be an option that is safe at a given time, allowing deferral of an alert.

Recall that an aircraft that has blundered may be unresponsive to alert commands. The aircraft will move according to unknown rules, and it would be pointless to consider a set of maneuver options as though they could be imposed. Instead, the remaining aircraft must be assigned maneuvers that ensure safety without cooperation from the blunderer.

A problem to address is that we have no way of knowing a priori which aircraft, if any, is blundering. Thus, we are faced with several hypotheses—of no blunderer, and of one blunderer that could be any of several aircraft. Each hypothesis corresponds to a distinct reaction by the system to different maneuvers that an alerting system might try to impose. For example, consider a three-aircraft system. There are four hypotheses corresponding to possible realities: no blunder, aircraft 1 blunders, aircraft 2 blunders, and aircraft 3 blunders. In addition, assume there are 8 system maneuver options—2 possible "commands" (one of which may be no command) for each of 3 aircraft. Because of the 4 possible realities, the system may respond in 4 ways to each of the 8 system commands. Under these circumstances, the problem of selecting an option that is safe becomes difficult, or even impossible if further constraints on the system maneuver are introduced, such as to limit the number of aircraft that can be intentionally broken out at once. The problem would be simplified if a particular hypothesis could be verified prior to alerting and maneuver selection.

It is undesirable to issue system alerts when no blunder has taken place. Such normal approach false alarms must be minimized as they can damage an operator's trust in the alerting system. This provides further motivation to select a particular hypothesis prior to alerting, because one of the four possible realities is that in which there is no blunder.

The occurrence of a blunder can be inferred from its effect on aircraft behavior. As mentioned, a blundering aircraft may deviate from a normal approach path. Other than this, there is no characteristic trajectory or pattern that can be used to identify a

blunder. It seems reasonable under these circumstances to judge as blundering any aircraft that deviates enough from a normal path that the likelihood of its being due to normal dynamics is very small. This can be seen as a hypothesis test. It is also the same method by which PRM detects a blunder, or a conformance test.

It is necessary at this point to combine a conformance test with a post-alert trajectory prediction model. Alerts are to be triggered by the conformance test, but the acceptability of the system maneuver associated with the alert is determined by the trajectory model and hazard definition.

To state this differently, the safety of possible system evasion maneuvers is monitored continuously during system operation using the post-alert trajectory model. Prior to an alert, there is no knowledge of whether a blunder has occurred (or of which hypothesis holds). It may be necessary to monitor and preserve safety for escape maneuvers under all hypotheses. This is illustrated in Fig. 8 for a system evolving along a trajectory in state space. Evasion maneuvers under each hypothesis are represented by different trajectory uncertainty envelopes.

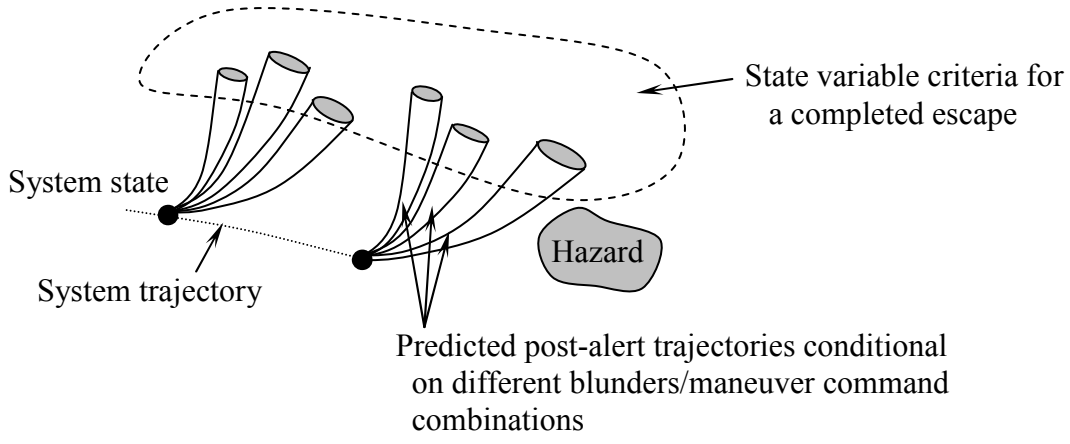

Fig. 8: Generalized Alert Option Monitoring

Recall that according to the post-alert trajectory prediction philosophy an alert must be triggered by the time an evasion option becomes marginally safe. Does this still apply when safety for only one out of several realities becomes marginal? Assuming

there is a non-negligible probability of each blunder hypothesis with respect to the others, the intuitive answer is yes: unless that hypothesis has been ruled out (found improbable) beforehand, the alert should take place. If there is an exact probability associated with each hypothesis, allowing safety to be sacrificed for select hypotheses may be a legitimate option too.

It is desired that alerts be triggered by conformance threshold crossings, and not forced by a near loss of safe options before the conformance threshold is reached. There are multiple reasons for this. One is that a single hypothesis should be selected before an alert occurs, and this only takes place when the conformance threshold is crossed. Another is that the conformance threshold is to be tuned for an acceptable rate of normal approach false alarms under the assumption that alerts do not occur before conformance threshold crossings take place—any additional alerts contribute to an unacceptable increase the normal approach alert rate. Finally, alerting by direct application of the trajectory model would tend to be computationally costly, particularly in light of the numerous maneuver options that may exist. Once a conformance-based threshold has been established, alerting safety along its boundaries can be proven through offline analysis, leaving a relatively simple threshold. In other words, it is desirable to establish that at least marginal safety exists for any reachable point on the conformance threshold for some maneuver option, and that the hazard will not be encountered nominally before the conformance threshold is reached. The possible complexity of demonstrating this may encourage choice of the simplest conformance threshold possible.

To better illustrate these ideas, a combination of a conformance-based threshold with the MIT (post-alert trajectory prediction) logic is shown in Fig. 9 for a 3 aircraft system. The diagram is a simplification showing only the hypothesis under which aircraft A may be blundering. It can be imagined that a similar diagram exists for each of the other blunder hypotheses.

If aircraft A crosses the conformance-based threshold (which has been tuned to an acceptable normal approach alert rate) it is judged to be blundering. The two aircraft on the adjacent approach centerline must then be assigned maneuvers that provide adequate

safety. Each can perform either a turn-with-climb breakout maneuver or continue with the approach. For each of the two aircraft B and C and each maneuver option there is a collision probability contour, such that if aircraft A is on the contour the probability of a collision is 0.001. Inside the contour the probability of a collision is greater than 0.001 and outside it is less. Clearly, aircraft B must choose the turn-with-climb if an alert occurs at the instant shown, or the probability of a collision will be greater than 0.001. Aircraft C could perhaps choose either maneuver (approximately equal safety), but avoiding a breakout might be favored. If completion conditions explicitly specify that only one aircraft can breakout in addition to the blunderer, aircraft C has only one option—to continue the approach.

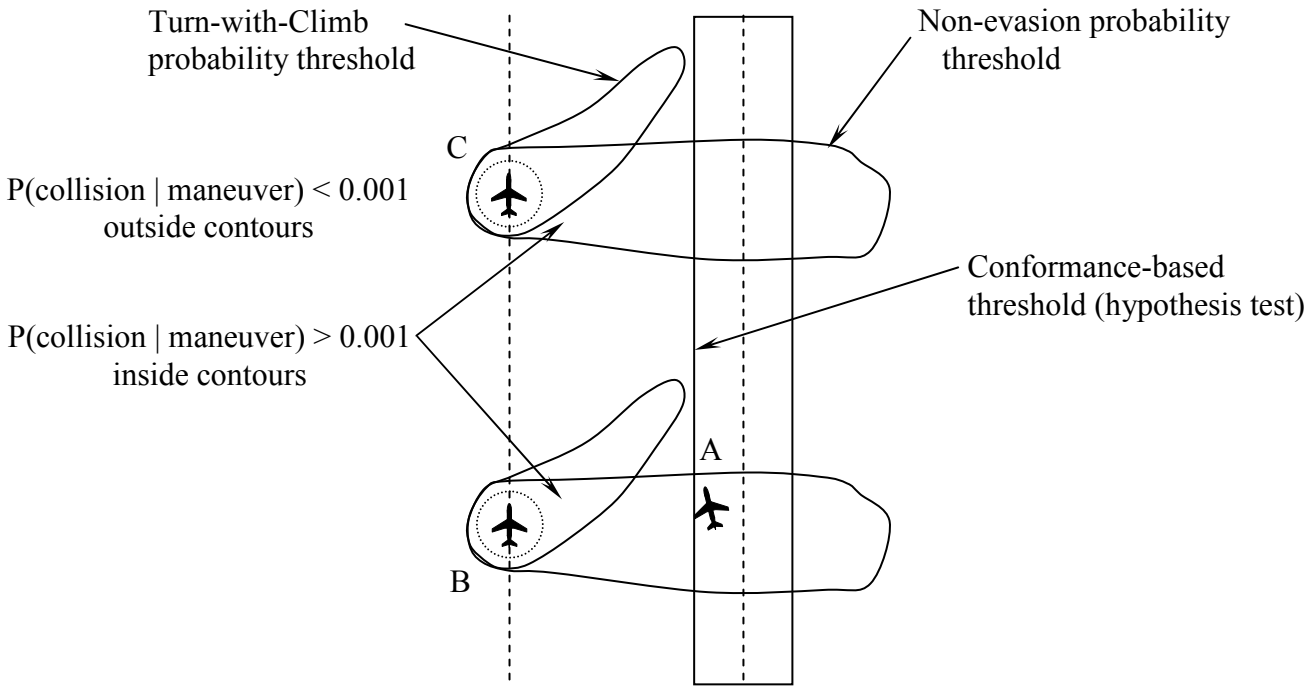

(MIT contours shown for specific heading, bank, and airspeeds)

Fig. 9: Combination Conformance-based and Post-alert Prediction Based Logic

For each of the threatened aircraft, at least one maneuver contour exists which cannot be crossed prior to crossing of the conformance threshold. As described previously, this is important because it may be unreasonable for the illustrated evasion

maneuver commands to be triggered when aircraft A has not been established as the blundering aircraft.

The lack of overlap is also important in that it will allow simplification of the threshold algorithm using the surface defined by the conformance criteria. The MIT thresholds are based on a complex Monte Carlo model that would be difficult or impossible to run in real time, especially if there were several aircraft in the approach system that would require simultaneous monitoring. It may be more convenient to evaluate this trajectory model offline and to implement simplified thresholds known to satisfy the safety requirements, as shown in Fig. 10.

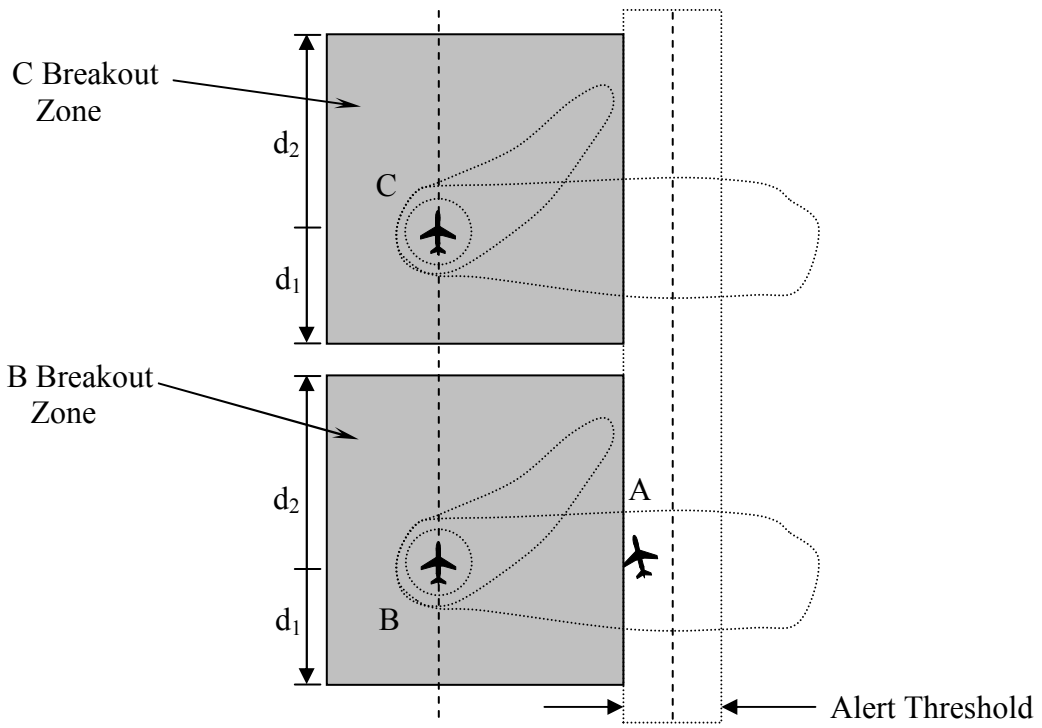

Fig. 10: Simplified Thresholds

In this example an aircraft would receive a breakout command only if an adjacent aircraft were judged a blunderer (after deviating from its approach), and came within a limited longitudinal distance shown as the gray region in Fig. 10. The probability values are no longer evaluated directly. The selection of system evasion maneuvers will have been predetermined for the scenario in which A is judged the blunderer and lies within

specific ranges of longitudinal distance with respect to the other two aircraft in the system (i.e. if d_1 < Longitudinal Distance < d_2 then breakout, else continue with approach).

As apparent in Fig. 11, this particular simplification will tend to result in false alarms that might safely be avoided if the full trajectory model were implemented. The intruder aircraft pictured is in a position where a breakout maneuver could safely be deferred (P(collision) < 0.001), but will be triggered if using the simplified threshold. Conversely, the simplified threshold avoids collision risks that the original logic accepts in trade for reducing false alarms. Realize that the illustrated contour represents a threshold value in a continuum, and the probability of a collision at any point on the simplified threshold is less than the 0.001 value at the trajectory-based threshold. Given the lack of penalty associated with false alarms occurring during blunders, the simplified threshold is making a reasonable trade.

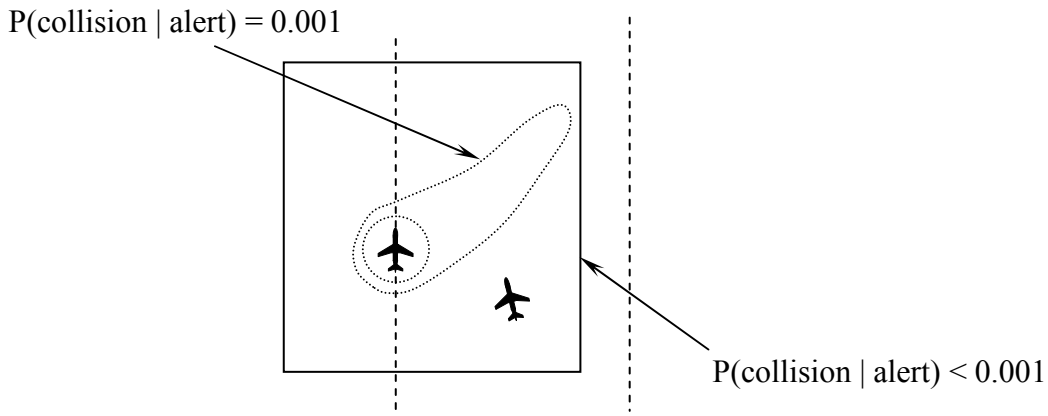

Fig 11: Blunder False Alarm vs. Safety Tradeoff

Because the details of the trajectory model chosen for the MIT logic are open to criticism (e.g. there is no way to compute error bounds on the computed probability values), relaxation of the thresholds through simplification might allay fears that the logic is highly optimized, but with faulty information. The simplification is akin to using a more conservative trajectory model (e.g. worst case), or significantly reducing the acceptable probability of collision over most of the alerting threshold.

The example of Fig. 10 shows a conformance threshold that appears to resemble the PRM threshold in that it checks only for excessive lateral deviations. In general this threshold may be a function of additional variables, such as aircraft heading, bank, and speed. In fact, without this additional information a large runway separation might be required to prevent the 0.001 probability contour for the turn-with-climb case from crossing over the conformance threshold. This is because the lateral extent of the MIT probability contours varies with aircraft heading, bank, and speed, and if considering lateral deviation conformance alone, the minimum distance between the runway centerlines would be determined by the fixed value of the lateral deviation threshold and the maximum value of lateral probability contour extent. Perhaps if the lateral separation threshold were also a function of additional variables, a smaller runway spacing could be allowed without intersection of the two thresholds, as shown in Fig. 12. In the bottom scenario, the conformance threshold has even been relaxed laterally for certain states in order to prevent an increase in the normal approach alert rate, on the grounds that the required safety will not be lost as a result.

This suggestion might not be reasonable. Whether it is depends on assumptions about the form of blunders and on the properties of normal approaches. This issue will be explored further in later sections.

7. Further Work with Conformance-based Thresholds

This section describes in more detail work at MIT into producing and analyzing conformance-based thresholds for parallel approach collision prevention.

A PRM-like alerting threshold based on lateral deviation from the approach path is the simplest example of a conformance-based logic for parallel approach alerting. It was suggested in the previous section that a conformance-based threshold employing state variables in addition to lateral deviation might have advantages over the simpler threshold. For example, a complex trajectory-based threshold might be more effectively simplified as a multi-dimensional conformance threshold than would be possible with a one-dimensional threshold. At this time, however, no attempt has been made to perform this simplification with an actual trajectory-based logic.

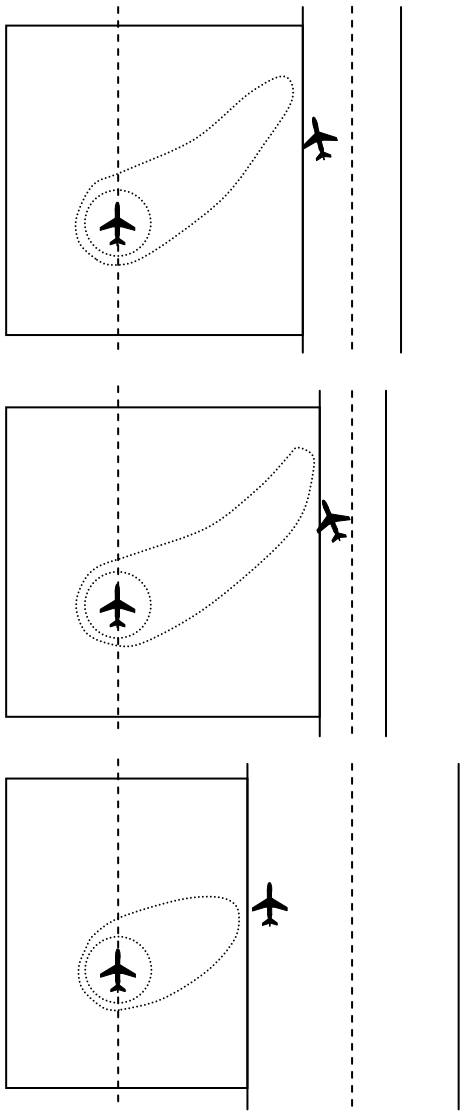

Fig. 12: Conformance Threshold as Function of Additional Variables

The potential value of additional state variables can also be described in terms of their effect on the relative quickness with which blunders are detected. This idea is illustrated using a simple conformance-based logic that was considered at MIT (Fig. 13). The relevant state variables in this case are lateral deviation from the approach centerline, y, heading angle, ψ and roll angle, ϕ. As shown, the threshold consists of a set of independent constant bounds on each state variable. The nominal value of each variable is zero (at the origin of each set of axes). During a normal approach an aircraft's state

wanders, but will tend to lie within a finite distance of the nominal path or state. Bounds are chosen so that a threshold crossing during a normal approach is unlikely. A deviation beyond these bounds is then considered evidence of a blunder, and justification for an alert.

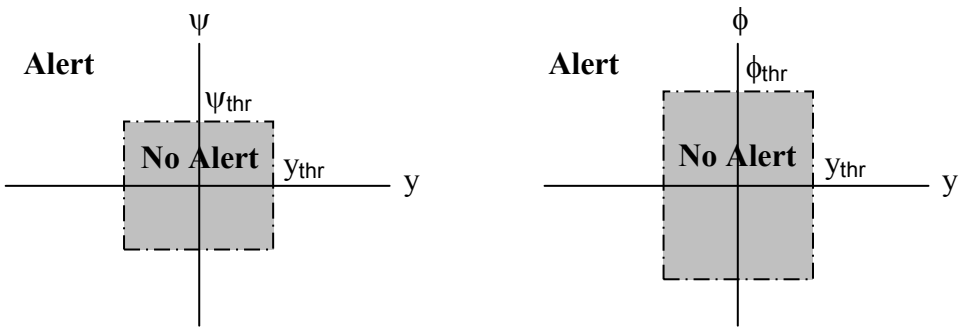

Fig. 13: Prototype Conformance-Based Threshold

Depending on the bounding values, the state of an aircraft may be able to exceed the magnitude bound for any one of the three state variables before exceeding the other two. For example, an aircraft might suddenly bank while at the approach centerline and with zero heading angle, triggering an alert before a significant lateral deviation or heading change has occurred. By nature of the hazard a large lateral deviation is required by at least one aircraft before the hazard (mid-air collision) can occur, and therefore an alert can be guaranteed to occur prior to a hazard event when the threshold is based on the lateral deviation variable alone. However, employing an additional variable or variables as shown may result in earlier detection of the same blunder. This is because both heading and roll are higher order variables than lateral deviation, and can change more quickly when a blunder occurs. Note, however, that there is no guarantee that the higher-order variables will deviate dramatically or quickly enough to trigger an early alert. Consider, for example, the difference between an idealized large heading angle change and slow lateral drift from the approach path, both illustrated in Fig. 14. Normal operating states are shaded, and alerting bounds are represented by dotted lines in the figure. Assume that state variable bounding values have been chosen so that some negligible normal approach alert rate is achieved (i.e. the bounds are large relative to normal deviations in each variable). For the large heading angle change blunder maneuver the first crossing is by the roll angle variable. For the slow drift, it is by the

lateral deviation variable, because roll and heading angle both fail to reach large enough magnitudes during the blunder. Thus, for a given set of bounds, a severe blunder maneuver such as the large heading angle change might be more quickly detectable using a threshold incorporating high order variables. A slow drift is expected to be relatively difficult to detect early due to the need to avoid normal approach alerts. In summary, depending on the magnitude of normal deviations relative to the severity of blunder maneuvers, adding additional state variables to a conformance-based logic may allow quicker detection of blunders.

In the illustrated example note that for neither maneuver is the heading angle threshold crossed before the lateral deviation threshold. Unless there are types of blunder maneuver for which the heading threshold will be the first crossed, there would be no point in including this variable in the logic. With each additional threshold variable a cost is incurred in the form of an increased non-blunder false alarm rate. Thus, it makes sense to minimize the number of variables employed.

An alternative to adding additional variables to a logic in search of quicker alerting is to restrict the normal system behavior to improve blunder detection. For example, reducing the magnitude of normal state variable deviations from the nominal path allows use of a tighter conformance threshold that will result in earlier alerts for any blunder maneuver that can cause a collision. In PRM, replacing ILS approach procedures and technology by more precise GPS-based approaches would presumably allow tightening of the lateral deviation threshold and earlier alerting, even without the addition of state variables to the logic. Requiring autopilot-coupled over hand flown approaches could have a similar effect.

A multi-dimensional conformance threshold need not consist of a set of independent variable bound checks. For example, there might be a strong statistical correlation between two or more measured variables that would make independent variable checking undesirable. As shown in Fig. 15, a threshold recognizing state variable interdependencies might provide earlier alerting with about the same non-blunder alert rate.

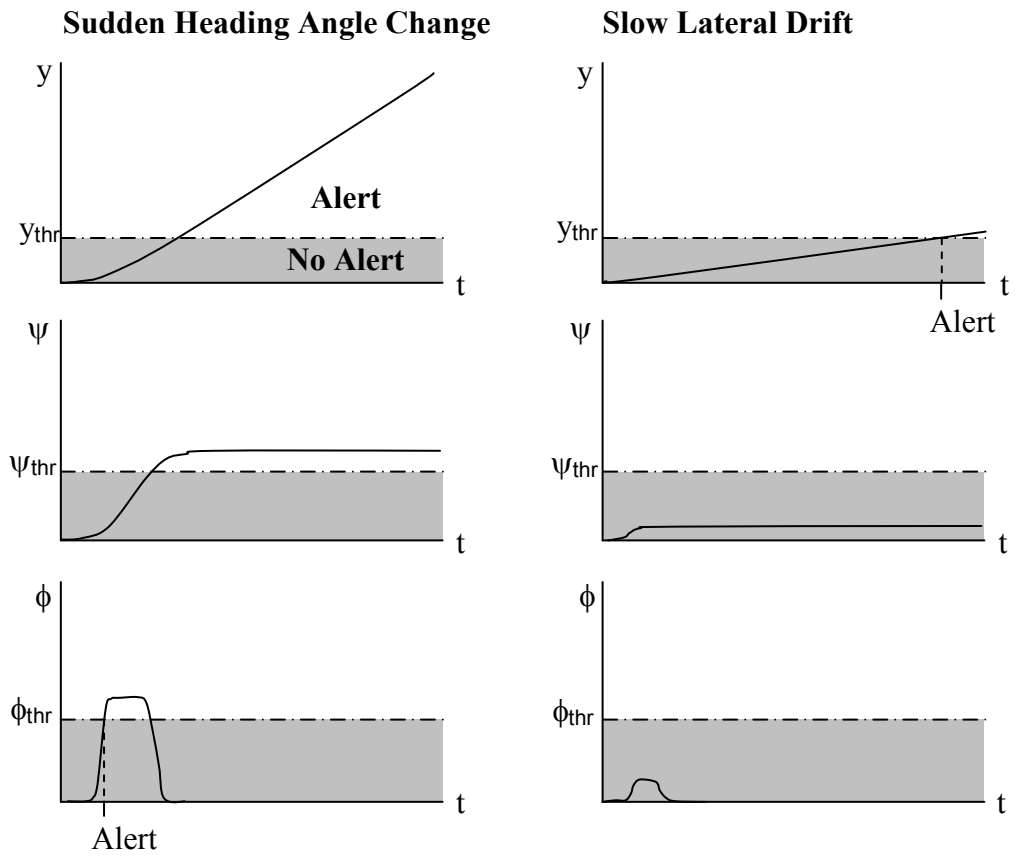

Fig. 14: Benefit of Additional State Variables vs. Blunder Type

If the goal of the alerting system is to ensure some level of safety at the time of an alert, then it may not matter that a particular threshold detects less aggressive blunder maneuvers more slowly than sudden ones. An aircraft drifting slowly off approach will take longer to endanger a neighboring aircraft than one turning sharply, so the alert for the slow drift can be deferred for longer. Based on this point of view and on the need to minimize non-blunder false alarms, the optimal set of state variables to include in a conformance-based threshold is the smallest one allowing a threshold that is compatible with (provides at least the safety of) an initial trajectory-based threshold and which results in a desired rate of non-blunder false alarms. This set will be a subset of the variables in the initial trajectory model.

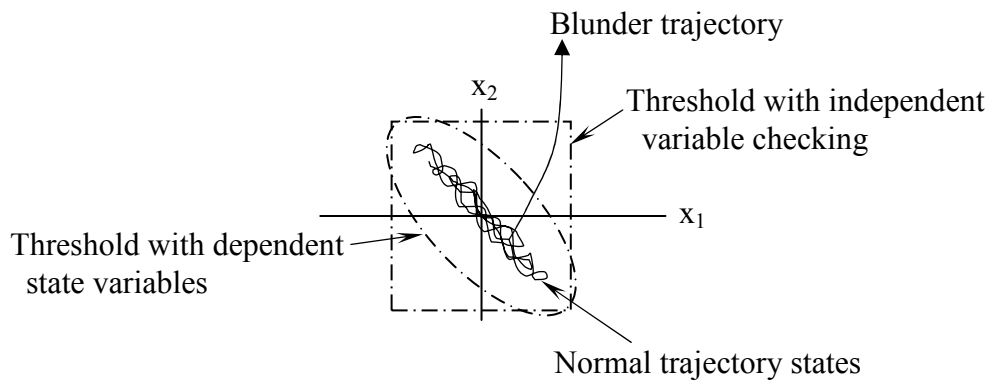

Fig. 15: Conformance Threshold Recognizing Correlated State Variables

7.1 Non-blunder False Alarm Analysis

Much of the discussion in this report is predicated on extensive knowledge of the behavior of a normally operating parallel approach system. Trajectory data for an existing system may be difficult to obtain in large quantities. In addition, any future alerting system for parallel approach collision prevention will likely be designed for an approach system operating under approach guidance technology and procedures that have yet to come into standard use. This makes data even more difficult, if not impossible, to obtain. In this section a method is described for computer generation of random trajectories of an aircraft operating in a future approach system. This tool has been applied to normal approach false alarm analysis for conformance-based thresholds.

Aircraft dynamic models of varying fidelity are commonly available. These range from full nonlinear models incorporating a large set of wind tunnel data to simple models linearized about a particular flight condition. For the current problem it is assumed that an aircraft is established on a straight final approach at constant speed, so that its dynamics are well approximated by a linearized model.

The linearized lateral and vertical approach dynamics of an aircraft (a C-47) were selected for initial experimentation (McRuer et al., 1990). Available state variables include position in three dimensions, velocity, and attitude angles. The control inputs to the aircraft are the aileron, rudder, and elevator angles. As shown in Fig. 16, the aircraft model has been incorporated into a feedback system with a stabilizing controller. The

controller shown has been designed using linear quadratic optimal state feedback methods in order to meet intuitive approach performance criteria, but normally one would attempt to duplicate as closely as possible the dynamics of the existing or planned aircraft/controller system of concern. The intended controller might be a human pilot or an autopilot.

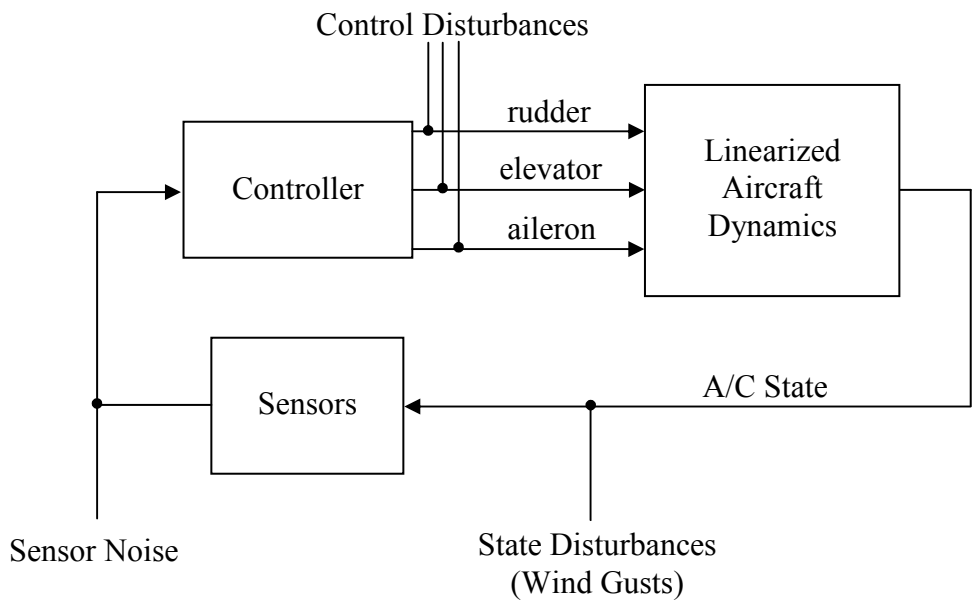

Fig. 16: Random Approach Simulation Design

Random variation of the aircraft state about the nominal approach path can be induced through random disturbance inputs to the system. Possible disturbances include those directly affecting the aircraft state (such as wind gusts), state measurement noise, and random disturbances in the controller outputs. The last item might be included to describe uncertainty inherent in a linear description of human manual control. Each type of disturbance can be approximated as the output of linear filters driven by white noise. For example, wind gusts are currently modeled as low-pass filtered white noise affecting the linear acceleration of the aircraft.

In general no part of the approach model is required to be linear, as this is a simulation and not a controller design problem. However, it has been convenient in the short term to make such assumptions: the controller currently employed is not modeled

after an actual controller, but has been generated through linear control system design methods in order to quickly demonstrate a stable closed-loop system qualitatively resembling an aircraft on approach. For maximum fidelity to a completely defined goal system, a nonlinear simulation may be needed.

A point worth repeating is that random simulation of the normal behavior of an approach system is arguably more reasonable as an analysis technique than random simulation of blunder behavior, as has been attempted in past alerting system analyses (Winder & Kuchar, 1999; Jackson & Samanant, 1999). When operating normally, the approach system should behave according to well-defined dynamic laws and random inputs that can be observed and modeled. The same cannot be said of blunders.

Fig. 17 shows partial data for a 1000 second random trajectory run from the simulation described. Plotted variables include lateral deviation from the approach centerline (y), vertical deviation (z), heading (ψ), roll (ϕ) and pitch (θ). As expected, the aircraft state varies randomly about a nominal approach value, but tends to remain within finite bounds. There is clear correlation between some variables (for example between lateral deviation and bank, and between vertical deviation and pitch) that might affect the choice of threshold shapes if both variables are included in the threshold.

As an example of how the simulation may be used in analysis, performance metrics were computed for an ellipsoidal threshold incorporating different combinations of state variables, and over a range of threshold size. The threshold is based on the first order statistics of the stationary random approach process (that is, the variances of each variable, and covariances between variable pairs at a moment in time.) The process covariance matrix is used to construct a Gaussian function. The alerting threshold is defined as a constant-value surface on this function, which is an ellipsoid in the space of state variables. Note that such a threshold does take into account the correlations between each pair of state variables. No claim is being made that such a threshold is most appropriate for parallel approach alerting—it is merely used as an example.

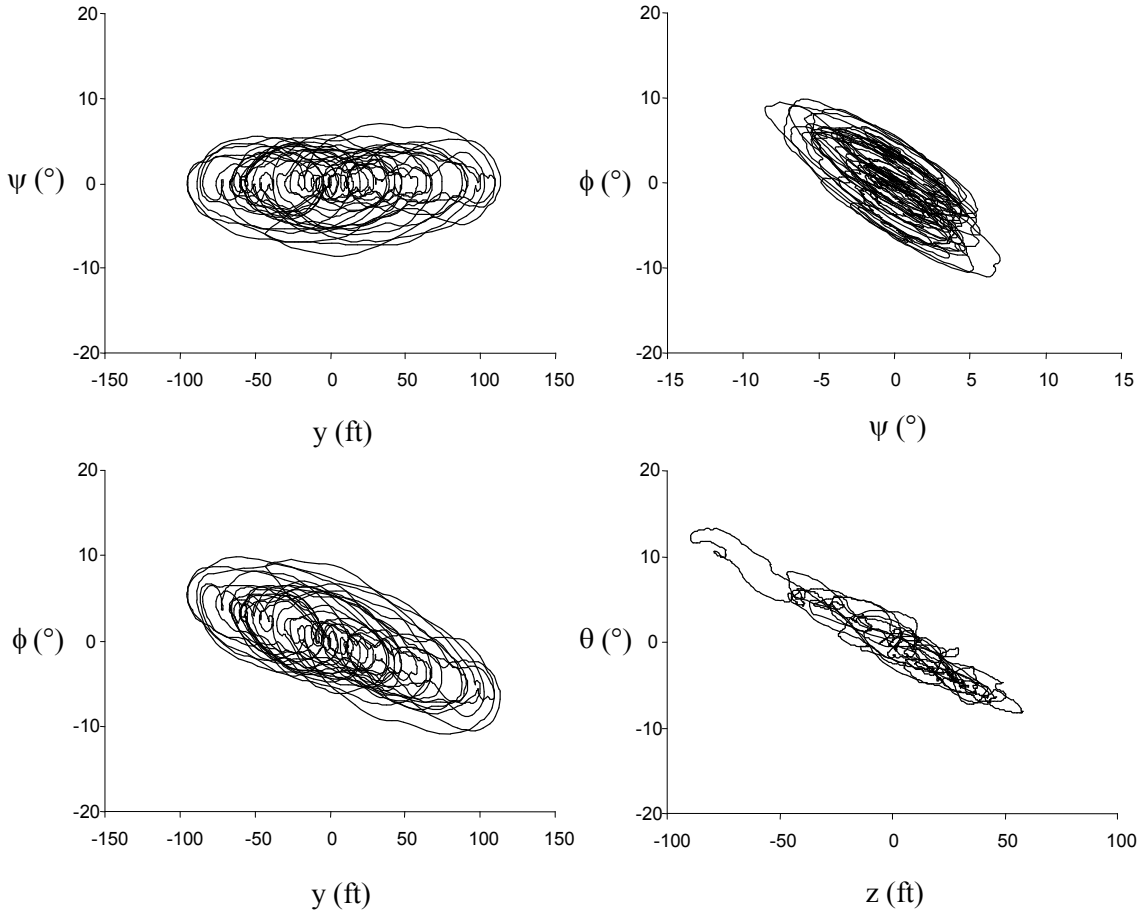

Fig. 17: Sample Random Trajectory (1000 sec)

The normal approach alert performance metric in this example is the mean time-to-alert for a random aircraft trajectory beginning exactly at the nominal approach state.

Another performance metric that is measured for the sake of example is the number of seconds remaining at the time of the alert for an idealized blunder maneuver to result in a 2500 ft deviation from the approach centerline. The blunder maneuver is idealized in that it begins exactly at the nominal approach state, and occurs without random state variations about the average blunder trajectory. Three different blunder maneuvers were simulated, including a 5° constant bank coordinated turn away from the nominal approach, a heading change with roll-out at 30° from the runway centerline, and a heading change with roll-out at 10° from the runway centerline. All are at a constant 145 knots.

Threshold variations included the number of state variables on which the covariance matrix, Gaussian function, and threshold were based (three combinations were considered: y, ψ, ϕ; y, ψ; and y, where y, ψ, and ϕ are lateral deviation, heading angle, and roll angle respectively), and the value of the Gaussian function (or the number of standard deviations) at which the threshold was set. Simulating over these conditions resulted in the curves shown in Fig. 18.

In the plots of Fig. 18 the horizontal axis is for the mean time-to-alert performance metric (normal approach false alarms), and the vertical axis is for the time remaining for a 2500 ft lateral deviation resulting from the blunder. For example, for a 5° constant bank blunder and a threshold based on lateral deviation, heading and roll, the threshold setting resulting in a 1000 second mean time-to-alert for normal approaches results in an alert approximately 42 seconds before the blunderer reaches a 2500 ft lateral deviation.

Notably, it does appear possible to realize a benefit in alerting quickness through the addition of state variables. In this example roll angle is particularly useful, resulting in up to a second of saved time (in the covered mean time-to-alert interval), depending on the blunder type and desired mean time-to-alert. Note, however, that 1 second is still a relatively minor gain when compared to the total time it may take the blunderer to reach an adjacent approach centerline. Heading angle appears to be a relatively ineffective addition to the logic compared to roll, though even it results in some minor improvement over lateral deviation alone for most values of mean time-to-alert.

An issue to point out is that the desired mean time-to-alert may be much larger than the approximately 5000 second maximum attained during these simulations. Assuming that a single approach takes 300 seconds, a 5000 second mean time-to-alert means that an aircraft on normal approach would on average cover a distance of only $5000/300 \cong 17$ approaches before a false alarm occurs. Or assuming that all aircraft have the same approach statistics, approximately one alert would occur for every 17 approaches (6 percent of approaches). Generating data for a mean time-to-alert above the limited range shown in Fig. 18 would require increasingly long simulation runs.

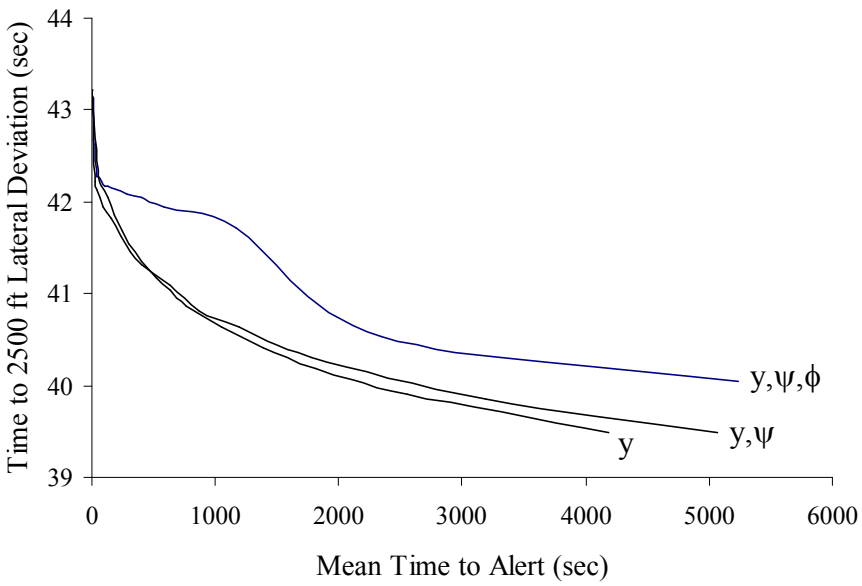

(a) 5° Constant Bank Blunder

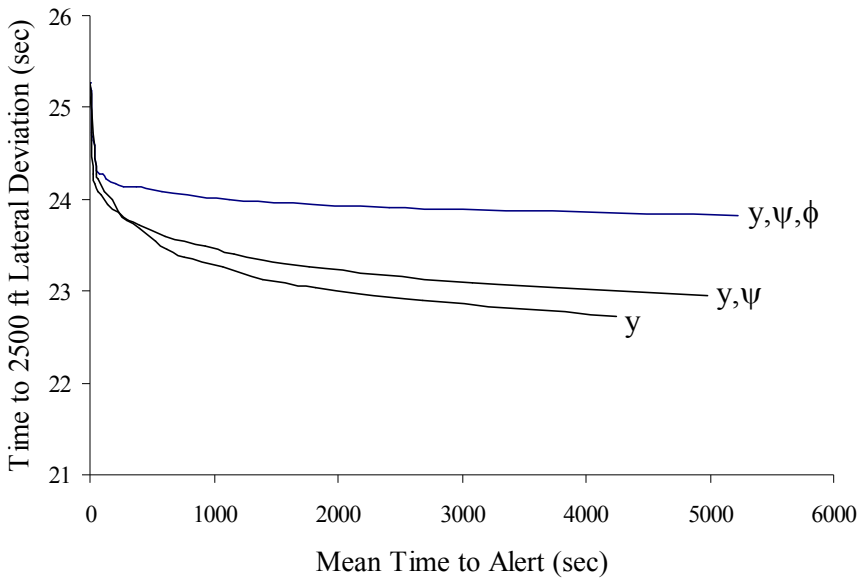

(b) 30° Heading Change Blunder

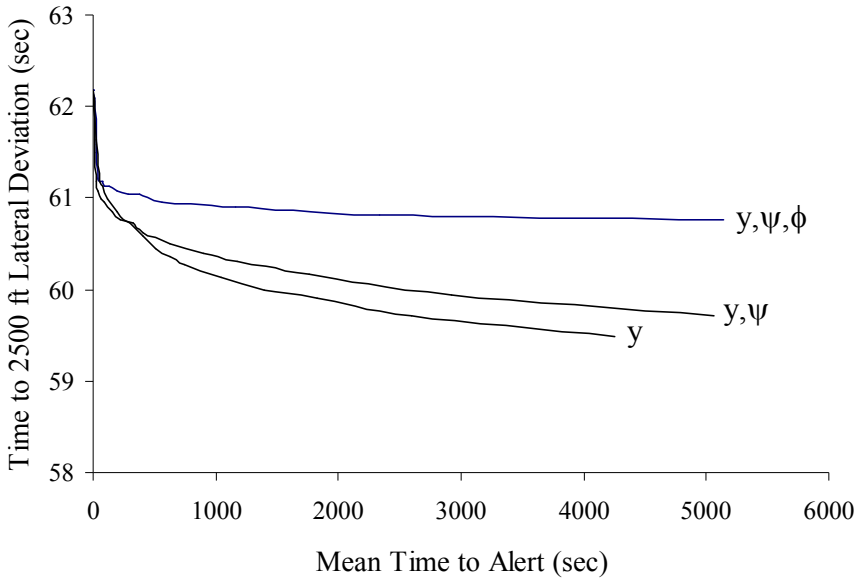

(c) 10° Heading Change Blunder

Fig. 18: Example Analysis of a Conformance-Based Threshold Using Random Trajectory Simulation Technique

References

Drumm, A. C. 1996. Lincoln Laboratory Evaluation of TCAS II Logic Version 6.04a. Volume I. Lincoln Laboratory, MIT. Lexington, MA.

Jackson, Mike & Paul Samanant. 1999. "Analysis of AILS Alerting System." Honeywell Technology Center. July 20.

Koczo, S. Coordinated Parallel Runway Approaches. 1996. NASA Scientific and Technical Information Office. NAS-19704 - Task 11. May.

Kuchar, James K. & Brenda D. Carpenter. 1997. "Airborne Collision Alerting Logic for Closely Spaced Parallel Approach." Air Traffic Control Quarterly. Vol. 5(2). pp. 111-127.

Kuchar, James K. & R. John Hansman. 1995. A Unified Methodology for the Evaluation of Hazard Alerting Systems. Aeronautical Systems Laboratory. MIT. Report No. ASL-95-1. January.

Kuchar, James. K. & Lee C. Yang. 1997. "Survey of Conflict Detection and Resolution Modeling Methods." AIAA-97-3732. AIAA Guidance, Navigation, and Control Conference. New Orleans, LA. August 11-13.

Kuchar, James K. 1996. "Methodology for Alerting System Performance Evaluation." Journal of Guidance, Control and Dynamics. Vol. 19. No. 2. pp. 438-444.

McRuer, Duane, Irving Ashkenas and Dunstan Graham. 1990. Aircraft Dynamics and Automatic Control. Princeton University Press. Princeton, NJ.

Mellone, V. & S. Frank. 1993. "Behavioral Impact of TCAS II on the National Air Traffic Control System." 7th International Symposium on Aviation Psychology. The Ohio State University. April.

Samanant, Paul & Mike Jackson. 2000. Description of the AILS Alerting Algorithm. Honeywell Technology Center for NASA Langley Research Center under NASA order number L-10690. Version 2.0. February 29.

Shank, Eric M. & Katherine M. Hollister. 1994. "Precision Runway Monitor." The Lincoln Laboratory Journal. Vol. 7. No. 22. pp. 329-353.

Waller, Marvin C. & Charles H. Scanlon, editors. 1996. Proceedings of the NASA Workshopon Flight Deck Centered Parallel Runway Approaches in Instrument Meteorological Conditions. NASA Conference Publication 10191. December.

Williamson, T. & N. A. Spencer. 1989. "Development and Operation of the Traffic

Alert and Collision Avoidance System." Proceedings of the IEEE. Vol. 77. No. 11.

Winder, Lee F. & James K. Kuchar. 1999. "Evaluation of Collision Avoidance Maneuvers for Parallel Approach." Journal of Guidance, Control and Dynamics. Vol. 22. No. 6. pp. 801-807.

Yang, Lee C. & James K. Kuchar. 2000. Aircraft Conflict Analysis and Real-Time Conflict Probing Using Probabilistic Trajectory Modeling. International Center for Air Transportation. MIT. Report No. ICAT-2000-2. May.