Approved for Public Release; Distribution Unlimited
Case # 06-1043

SAFETY OF COOPERATIVE COLLISION AVOIDANCE FOR UNMANNED AIRCRAFT

Andrew D. Zeitlin, MITRE Corp./CAASD, McLean VA
Michael P. McLaughlin, Corp./CAASD, McLean VA

Abstract

For Unmanned Aircraft to be routinely used in civil airspace, an effective collision avoidance function is one area deemed essential for safe operation. Like manned aircraft, avoiding collisions with transponder-equipped, or "cooperative" traffic is among the primary hazards.

This paper discusses similarities and differences in the collision avoidance function and the necessity of developing various models of environmental and system components in the collision avoidance functional chain. Potential sensitivities and shortcomings of the TCAS collision avoidance system for unmanned aircraft are discussed.

The analysis method of fast-time simulation can develop a rich sample of collision encounter events from the numerous statistical distributions. This provides an established means to demonstrate system compliance with safety targets, when they are established.

Introduction

Development of Unmanned Aircraft Systems (UAS) is proceeding at a rapid pace, and Government and Industry alike are envisioning numerous uses, aside from battlefield support. There is great pressure to achieve access to the national airspace, but at present no means comparable to those used for manned aircraft are in place. A conspicuous need is to provide protection against midair collision. Manned aircraft utilize a variety of avoidance capabilities, both by the pilot and onboard systems. Some of those aircraft, primarily passenger and cargo carriers, utilize the Traffic Alert and Collision Avoidance System (TCAS II), as a part of their overall safety provision. It is not surprising that TCAS is mentioned as a candidate for use aboard UAS. Concepts and requirements for operating UAS in civil airspace are in the formative stage, and any role for TCAS must be carefully examined within the broader picture of UAS safety.

The UAS safety case needs to address all of the same hazards that are dealt with by manned aircraft. These include not only cooperative targets, but also non-cooperative traffic, terrain and obstacles, severe weather, and terrestrial features. The TCAS system addresses only transponder-equipped (cooperative) traffic, and does so only when airborne, beginning about 500 ft above the ground. For UAS, a system dealing with this traffic would represent only part of a more comprehensive collision avoidance requirement, whose surveillance system might include a suite of sensors for the various hazard types.

This paper presents the work performed to date from the MITRE research program that is examining the suitability of TCAS for use with UAS. The results will be closely coordinated with industry and government activities, particularly within the RTCA SC-203 standards committee Sense & Avoid Working Group.

Considering TCAS Operation for UAS

TCAS was designed prior to UAS technology, so its performance was matched to the operations of manned aircraft which were expected to install it, namely air carriers [1]. This section discusses some aspects of TCAS design that may not be well suited for UAS usage.

Surveillance System

The TCAS surveillance system interrogates nearby transponders using Mode C and Mode S formats on 1030 MHz, and receives replies. It uses antennas mounted on the top and bottom of the aircraft, of which one antenna has direction-finding capability. The transmit power and receiver sensitivity are designed to assure a surveillance range of 14 nm, which is sufficient to establish a timely track when two 600-kt aircraft are approaching head-on, and leave 30 seconds for the Resolution Advisory (RA) and the resulting maneuver. For UAS, where maximum airspeeds should be far less, it should be possible to reduce the surveillance range from the UAS and still detect a target with an equivalent warning time. On the other hand, the warning time parameter is based on a presumed avoidance maneuver of a certain magnitude. If the UAS cannot achieve that maneuver (discussed below), then additional warning time or other measures may need to be taken to provide the equivalent avoidance capability. This would negate any attempts to reduce the surveillance range.

Response to Resolution Advisories

The surveillance system also provides target data to the TCAS Traffic display. In the cockpit, this display presents a graphical representation of nearby traffic, showing the range, bearing, and relative altitude of the traffic. A pilot-operated control may be used to adjust the range around own aircraft that is displayed, and to show only traffic above or below own aircraft (as well as co-altitude). These controls often are used to reduce clutter, such as during a climbing or descending phase of flight, when traffic at other altitudes might obscure the traffic pertinent to collision avoidance. In the UAS application, the traffic information could be linked to the remote pilot. It could be displayed with heading-up orientation as in the manned cockpit, or it could be adjusted to a North-up orientation using a separate navigation input for own heading. If the surveillance range is reduced as suggested above, the situational awareness benefits of the display would be correspondingly reduced.

For manned TCAS operation, the provision of the Traffic display is meant to build confidence in the system so that its RAs are trusted, and it provides part of the traffic situational awareness picture that pilots desire, supplemented by other information such as see-and-avoid and the radio “party line.” A UAS concept should define the intended means of all aspects of safety provision, and may depart from the role TCAS plays in manned cockpits.

TCAS issues RAs only in the vertical dimension. It is able to form tracks on nearby aircraft and estimate the proximity and rates of range and altitude for each one. The TCAS logic determines the minimum vertical maneuver that will achieve a specified miss distance at the closest point of approach. Its projection is based upon an expectation of pilot delay time and acceleration in achieving the indicated vertical rate. Moreover, pilots are trained not to maneuver laterally based upon TCAS traffic information, as the bearing information that it derives for targets is not sufficiently accurate to be trusted for this purpose. For some encounter geometries, it may be difficult to select the correct lateral maneuver from this dynamic display of relative position, especially since pilots are not trained in this skill, as controllers are (for absolute position data).

For the UAS application, if a remote pilot must notice the RA and command an avoidance maneuver, the delays in so doing could be greater than in the manned cockpit. The evaluations of TCAS safety [2] that led to its approval were predicated upon a prompt response, both to initial RAs and even more promptly and vigorously to any subsequent strengthening (for example “Increase Climb”) or sense reversal (Table 1). For UAS, delays could be considerably greater if communication latencies add time before the recognition of the event (for downlink latency) and to maneuvering the aircraft (for uplink command latency). The UAS concept must specify the remote pilot tasking and resulting workload, the displays, the training requirements, and the decision process before a firm assessment of this delay can be made.

Table 1. Expected TCAS RA Response Parameters

RA Type	Delay (s)	Acceleration (g)
Initial	5	0.25
Strengthening	2.5	0.35
Sense Reversal	2.5	0.35

Could the TCAS RA threshold simply be enlarged to compensate for extra delay? This approach might succeed for linear collision encounters, but may not affect the warning time for encounters in which a maneuver causes a late encroachment of the protected volume of airspace. Also, the parameters of TCAS were set so as to provide an acceptable tradeoff between achieving safety versus minimizing nuisance alerts from aircraft that were sufficiently separated. A larger protected volume would tend to declare more threats in the latter category and would unbalance the trade. We expect to look for fruitful logic modifications in later efforts.

Could TCAS RAs be executed autonomously by the UAS vehicle? This option holds some attraction because it eliminates the communication delays to the pilot and back to the vehicle; it eliminates any reliance on the link operating during the short but critical interval between RA and closest approach; and it eliminates the pilot decision delay and human errors. A safety evaluation would need to consider two sources of errors that could lead to hazardous maneuvers. The first is a system fault that leads to an incorrect RA. The second is a barometric altimetry error that "fools" TCAS logic by reporting the wrong altitude for the target aircraft. Altimetry systems, particularly on small aircraft, have a recognized range of inaccuracy. TCAS logic accounts for this error to some extent by the size of its vertical thresholds; but a fundamental tenet of the original TCAS approval was the high probability[1] of the pilot mitigating such a consequence by looking for the target, cued by the traffic display, and rejecting any RA maneuver that seemed clearly wrong. If a UAS design attempts to substitute another sensor that "sees" the target like a pilot, it must either involve the remote pilot in the decision (negating some of the advantages cited above), or develop and implement a decision logic that considers both the TCAS RA and the "visual" sensor input.

Collision Encounter Model

When TCAS was first considered for introduction into the airspace, MITRE constructed a collision encounter model that has become the basis for similar models worldwide. The U.S. model examined moderately close two-aircraft encounters from a variety of terminal airspace locations. The ground radar data was interpolated and smoothed to form tracks, and the qualifying track pairs were characterized in several ways:

- The vertical character of the geometry was expressed in combinations of the two aircraft's profiles, chosen from:
    - level
    - vertical rate
    - vertical rate, leveling off
    - level, changing to vertical rate
- Distribution of vertical rates, including realistic probabilities for transitions from one rate to another
- Distributions of transitions from rate to level and level to rate
- Horizontal and vertical separation at the three-dimensional point of closest approach (CPA)
- Whether a crossing of altitudes occurred during the 30 s before or 10 s after CPA
- The altitude region of the encounter (TCAS logic parameters vary by altitude).

This model has been used with Monte Carlo fast-time simulation to evaluate each version of TCAS collision avoidance logic. For each such evaluation, a large number of encounters are simulated covering all of the geometry classes and vertical separations (Figure 1). Using a surveillance model for measurement noise and a pilot response model of maneuvering in response to an RA, each encounter is repeated, modeling the separation both with and without TCAS. Any maneuvers resulting from a TCAS RA can change the separation at CPA, and these are recorded. The Monte Carlo method repeats the nominal encounter many times, drawing specific values of variables from separate distributions. A large number of repetitions provides a representative sample drawn from the enormous set of possible combinations. Finally, all the results are combined by weighting them over the proportions corresponding to the traffic observed in the airspace.

[1] While visual acquisition would be unlikely to help when flying in Instrument Meteorological Conditions, targets with low-quality altimetry systems rarely fly in those conditions.

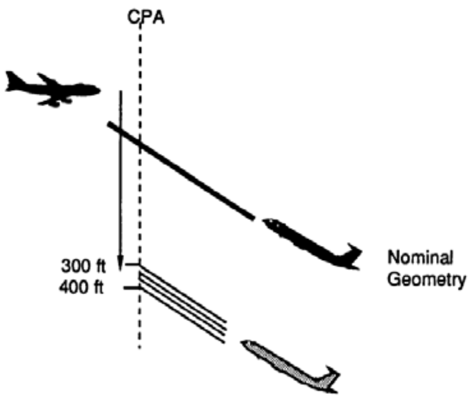

Figure 1. Encounter generation for simulation

The model statistics, drawn from manned operations, do not readily extend to UAS operations. Those would vary in some significant ways:

- Many types of UAS missions have been proposed [3]. Some of these could contain flight profiles very different from conventional manned flight, for example loitering or patrolling a contained area. The altitude mix also could differ from that used for the airspace model, so that encounter probabilities would fall into different altitude bands.
- UAS vehicle flight dynamics span a wide range of values [4]. Some of these differ greatly from manned aircraft, for example in speed, and horizontal or vertical maneuverability. Some UAS exhibit vertical motion, and can hover.

Some UAS vehicles are limited in their climb or descend performance. To illustrate the effect of these limits on collision avoidance, Figure 2 presents the calculated time to vertically maneuver from level flight at 0.25g acceleration (the same expected as a manned aircraft using TCAS), to climb or descend at various maximum vertical rates and achieve a 700 ft vertical separation. This displacement is the value used by TCAS at very high altitudes (i.e., above FL410), and the last value (1500 ft/minute) is the vertical rate expected by the TCAS logic. Although that great a displacement is not required to avoid a collision, the value reflects an allowance for barometric altimetry errors. The figure shows that the maneuver time more than doubles if the vertical rate cannot exceed 600 ft/minute.

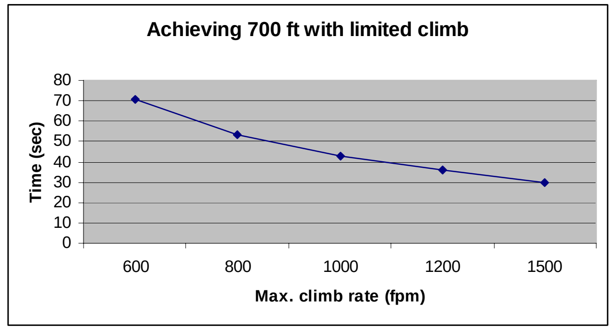

Figure 2. Time to Achieve 700 ft with limited climb performance

When applying the modeling and simulation techniques to UAS operations, the results will only be as good as the underlying assumptions. For fidelity of the results, the particular UAS system concept must be specified and then reflected in the model. Some of the primary issues are:

- Is the UAS remotely piloted in all respects, including collision avoidance, or does it have the ability to maneuver autonomously?
- Are there predictable paths for vehicle failure conditions, such as for a lost control link?
- Is the UAS maneuverability during an avoidance maneuver variable, such as by altitude, weight, temperature, or regime of flight?

The encounter model has served as a basis for TCAS design, since the logic necessarily must involve trades between collision protection and operational acceptability. (As an analogy, consider a home smoke detector. If it could be made far more sensitive to smoke and vapors, it could provide additional warning time in the event of a real fire, at the cost of many false alarms sounding when normal cooking was taking place in the kitchen. That would probably not be acceptable.) For TCAS, the threat detection parameters – warning time, acceleration buffer, and vertical alert and separation thresholds – were chosen in part by testing results for the observed “normal” traffic and ensuring that the nuisance alarm rate was tolerably low. At the same time, the detection parameters had to provide sufficient time to alert, maneuver and avoid a real collision with very high probability. When making similar trades for a UAS, it is not enough to provide the surveillance and alerting for a real collision. The UAS and the manned traffic both need to co-exist without excessive disruption from the alerting system.

The surveillance model must be consistent with the specific system used onboard the UAS. TCAS standards specify its surveillance performance; it is possible that limitations such as vehicle size or power may lead to differences in this function. If different antenna technology is used, variations in field of view could arise. This must be modeled, so that the appropriate encounters (or segments of some encounters) are not incorrectly portrayed as being properly detected by the system. This issue involves the geometric characteristics of UAS-specific mission profiles.

The remote pilot response to a TCAS RA presents an important subject in need of study by human factors experts. The UAS concept for any particular implementation may be among diverse possibilities, including pilot training, experience, workload, displays for both flight dynamics, navigation, mission performance (e.g., surveillance of the ground), communications by data link, and finally, collision avoidance. The suite of displays may not resemble those found on manned aircraft, and in any event, the different situation suggests that there will be differences in piloting performance.

Experience with TCAS in manned aircraft found variations in how pilots used the system and responded to RAs. When it was first introduced, there was a pronounced tendency to over- or undershoot the intended vertical rate until experience was gained making the avoidance maneuver. For UAS pilots, performance could be better or worse depending on many factors: the displays and controls, their experience piloting the UA, and the ability to make the maneuver without the actual onboard visual scene or experiencing the dynamic sensory effects of aircraft motion.

If there is a risk of pilots ignoring the TCAS RA in preference to other information and maneuvering in a different manner, that hazard must be taken very seriously. In encounters between two TCAS-equipped aircraft, the respective RA maneuver senses are coordinated, and safety is contingent upon cooperating by not maneuvering opposite to the displayed sense. A recent evaluation of the TCAS RA Reversal logic [5] showed that non-responding aircraft produce a much higher Risk Ratio (i.e., less safety) than responding ones.

Some concepts envision a single pilot controlling more than one UAS simultaneously. The questions of workload and confusion will need to be studied and compared to the experience for manned aircraft. This could change the probabilities of some hazards and even could create new ones.

The preceding discussion presumed that a remote pilot would respond to RAs. If instead the UAS was able to maneuver autonomously, other issues arise. One of these is informing the pilot of that maneuver so that he can act appropriately, both during and immediately after the maneuver. Another is defining exactly what the pilot can and cannot control during the avoidance maneuver, when and how that reverts to “normal” capability, and how the pilot is informed of these changes. It would be undesirable to have the pilot taking some action that undermines the avoidance. It also would be misleading to have the pilot conclude the vehicle was non-responsive to his inputs due to some failure, when in fact it was correctly performing an autonomous avoidance maneuver.

Simulation Results

The first results examined the sensitivity of TCAS safety to incremental delay. This might be incurred if a remote pilot must respond to RAs, and particularly so if communication delays are significant. Figures 3 and 4 present the Risk Ratio for various incremental delays above those normally modeled for a manned aircraft using TCAS. Figure 3 gives the data for a UAS with TCAS in conflict with a manned TCAS aircraft. Figure 4 shows a UAS with TCAS in conflict with a non-TCAS aircraft. Each data point represents the full set of encounters for the U.S. airspace model, i.e., what was previously a complete safety simulation. These results show that safety is very sensitive to any incremental delay. For example, 5 seconds of additional delay would approximately double the risk in each case.

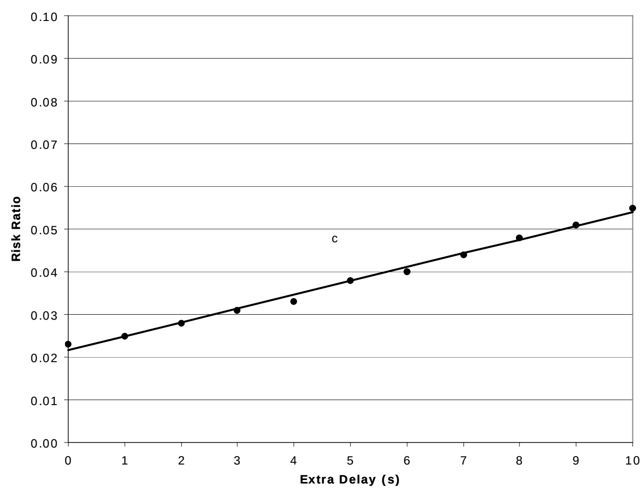

Figure 3. Risk vs. Incremental Delay, UAS vs. TCAS aircraft

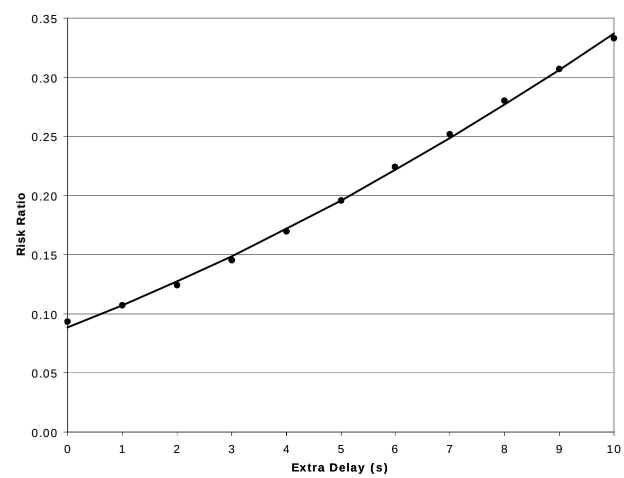

Figure 4. Risk vs. Incremental Delay, UAS vs. non-TCAS aircraft

Future Safety Evaluation

Forthcoming simulations will explore the sensitivities to limited climb and descend performance, and will begin the investigation of some mission characteristics and their effect on collision encounter statistics.

Ultimately, to evaluate the safety of a prospective vehicle, mission, and operational concept, all of the models comprising the simulated inputs, collision avoidance system performance, and pilot and vehicle response, will need to reflect the proposed usage. It has yet to be determined to what extent some of the model parameters can be considered general, and applicable to many UAS. Each safety evaluation could produce different results and lead to custom safety requirements, such as constraints on the operation. The standards activity should try to identify requirements that could assure acceptable levels of safety without undue constraints on desired uses.

Conclusions

While many aspects of piloting manned aircraft are well known, the foreseen scope of UAS operations is broad. The safety evaluations must be comprehensive and specific to the system and certain elements of the operation.

The maneuvering characteristics of some UAS vehicles could severely degrade the safety obtained from TCAS, as its logic is presently designed.

If a remote pilot is to be the means of responding to RAs, the communication link and the pilot response characteristics are both sensitive elements in the safety calculation.

The Monte Carlo safety simulations used to evaluate TCAS performance in an airspace must use models that accurately reflect various aspects of the system. Work must be undertaken to properly characterize sensor performance, human performance, vehicle maneuver dynamics, and encounter characteristics. This last model will depend strongly upon the type of mission profiles to be flown, and the airspace traffic characteristics.

References

[1] RTCA, December 16, 1997, “Minimum Operational Performance Standards for the Traffic Alert and Collision Avoidance System II (TCAS II) Airborne Equipment”, DO-185A, Washington D.C , RTCA, Inc.

[2] McLaughlin, M., June, 1997 , "Safety Study of the Traffic Alert and Collision Avoidance System (TCAS II) – Final Version," MTR97W32, McLean VA , The MITRE Corporation.

[3] DeGarmo, M. and G.M. Nelson, October 2004, "Prospective Unmanned Aerial Vehicle Operations in the Future National Airspace System", ATIO Conference, AIAA.

[4] Weibel, R., and R.J. Hansman, March 2005, "Safety Considerations for Operation of Unmanned Aerial Vehicles in the National Airspace System", Report no. ICAO-2005-1, Cambridge MA, Massachusetts Institute of Technology.

[5] RTCA, November 8, 2005, "Safety Analysis of Proposed Change to TCAS RA Reversal Logic," DO-298, Washington D.C., RTCA Inc.

Disclaimer

Work performed by The MITRE Corporation was produced for the U.S. Government under Contract DTFA01-01-C-00001 and is subject to Federal Aviation Administration Acquisition Management System Clause 3.5-13, Rights In Data-General, Alt. III and Alt. IV (Oct. 1996).

The contents of this document reflect the views of the author and The MITRE Corporation and do not necessarily reflect the views of the FAA or the DOT. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, expressed or implied, concerning the content or accuracy of these views.

25th Digital Avionics Systems Conference

October 15, 2006