

© 2006 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

IMPROVING THE RESOLUTION ADVISORY REVERSAL LOGIC OF THE TRAFFIC ALERT AND COLLISION AVOIDANCE SYSTEM

Andrew D. Zeitlin, MITRE Corp./CAASD, McLean, VA

Thierry Arino, Sofréavia, Toulouse, France

James Kuchar, MIT Lincoln Laboratory, Lexington, MA

Abstract

The Traffic Alert and Collision Avoidance System (TCAS II) is the worldwide standard system for manned aircraft to avoid collisions with airborne transponder-equipped traffic. A safety vulnerability of the collision avoidance logic was reported by European analysts, who also proposed a change to correct it. The safety issue concerns limitations in the ability of TCAS to reverse the sense of a Resolution Advisory (RA) during an encounter. The issue was addressed by a team of experts¹ in the Requirements Working Group (RWG) of RTCA Special Committee 147 [1]. This paper discusses the problem, the metrics and methods used in the analysis, and presents results that quantify the effectiveness of the proposed solution. Finally, recommendations are presented for implementing the change.

Introduction

The latest amendments to TCAS II logic Version 7 were completed in 2000 and it has become the standard version throughout most of the world. Among many other improvements, it provides the capability to reverse the sense of a Resolution Advisory (e.g., from “climb” to “descend”) in a coordinated encounter (with another TCAS-equipped aircraft) when conditions deteriorate. This could happen, for example, if one

pilot failed to follow his displayed RA, or worse, maneuvered in the opposite direction.

The design of the reversal logic [2] involved a number of judgments and assumptions. Operational experience has shown the need to revisit these decisions. The EUROCONTROL ACAS Programme predicted difficulties with the reversal logic, leading to monitoring efforts in Europe during 2001-2005, where the issue was in fact observed. Two accidents occurred, a near midair collision in Japanese airspace, and a midair collision over Überlingen, Germany in 2002. In both, the adverse outcomes might have been avoided if the RA had been reversed. Overall since the end of 2000, eight incidents have been observed in Europe, one in Japanese airspace, and three in United States airspace.

Coordinated Encounters

For purposes of the analysis, two types of scenarios were deemed relevant. The first of these, termed issue SA01a, deals with coordinated encounters between TCAS-equipped aircraft. It involves either a late reversal of the sense of the initial RA, or a failure to reverse the sense of an RA where a reversal could prevent a Near Mid-Air Collision (NMAC). An NMAC is defined as an event in which two aircraft are separated by less than 100 ft vertically and 500 ft horizontally. This issue is referred to as “late reversal RAs or no reversal RAs in coordinated encounters.”

An example of an operationally-realistic encounter that falls into this category involves two converging aircraft flying at the same flight level, as shown in Figure 1. An Air Traffic Control (ATC) instruction just prior to or at the same time as an initial RA could induce a maneuver that is contrary to the sense of the RA for one of the aircraft.

¹ Also making significant contributions to the work were: Christian Aveneau of DSN (France), Kathryn Ciaramella of the Federal Aviation Administration, Hui Men and Nam Phamdo of Johns Hopkins University/Applied Physics Laboratory, Barbara Chludzinski, Ann Drumm, Garrett Harris, Katherine Sinclair and David Spencer of MIT Lincoln Laboratory, Ganghuai Wang of MITRE, Ken Carpenter of QinetiQ (United Kingdom), Hervé Drevillon and Stéphan Chabert of Sofréavia (France). Valuable program direction was provided by Stephen George (FAA) and John Law (EUROCONTROL).




Figure 1. Issue SA01a -- Coordinated Encounter

Uncoordinated encounters

The second category, termed SA01b, involves encounters between one TCAS-equipped and one unequipped aircraft (or a TCAS unit operating in stand-by or TA-only mode). The same Figure applies, except that the aircraft that follows an ATC instruction is not receiving any RA.

Frequency of Incidents

Most of the 12 known SA01 incidents were discovered as a result of monitoring programs being undertaken. One observation of the RWG is that SA01 events have been found whenever a monitoring program to detect them has been operating, and observed SA01 events are not confined to particular airspaces, altitude layers, or airlines.

The EUROCONTROL operational monitoring effort took advantage of the close cooperation between the Safety Issue Rectification Extension (SIRE) team, composed of Sofréavia and DSNAs experts, and operational entities such as the French DSNAs and European airlines. This enabled access to information including airborne recorded data, pilot reports, air traffic controller reports, and radar data.

In the U.S., the production Mode S sensor at MIT Lincoln Laboratory has been used since mid-1994 to record downlinked RA reports and the corresponding surveillance data from aircraft in the Boston airspace. Sensor coverage is approximately 60 nmi, allowing observation of traffic arriving and departing Boston's Logan International Airport and numerous other smaller airports in the New England area as well as en-route traffic inbound and outbound from the North Atlantic Track System.

Analysis of pilots' responses to RAs in this airspace found that the majority of pilots failed to respond promptly and correctly, and a relatively high percentage of pilots moved in a direction opposite to the RA. Figure 2, which shows the vertical rate during the RA vs. the vertical rate prior to the RA for climb sense RAs, illustrates these results. Circles indicate responses that were timely (within 5 seconds); squares indicate responses that were slow. Symbols in the top segment (in green, above the horizontal line and to the left of the diagonal) indicate responses that achieved the intended vertical rate. Symbols in the middle segment (in yellow, lower left triangle) indicate responses that were in the correct direction but did not achieve the desired vertical rate. Symbols below the diagonal (in red, right triangle) indicate responses that were in the wrong direction.

Similar results were found when examining descend sense RAs.




Figure 2. Vertical Rate During vs. Prior to RA

Estimate of Frequency and Criticality

With the close cooperation of European airlines, the SIRE team estimated the frequency of SA01 events in European airspace. Between 1 April 2001 and 31 May 2002, two occurrences of issue SA01a were experienced by aircraft belonging to a major European airline within European airspace, during which time the airline flew 30,190 flight hours per month. Extrapolating this rate across European airspace, an SA01a event would be

expected to occur at an estimated rate of 4.7×10^{-6} per flight hour.

Given the few actual occurrences used for this computation, there is uncertainty in the resulting figure for frequency: there is a 95% confidence that the actual rate lies between 5.7×10^{-7} and 1.7×10^{-5} per flight hour. As there are approximately 12.5 million flight hours annually in Europe, these statistics imply that at least seven SA01a events occur per year in European airspace.

The probability of mid-air collision as a consequence of an SA01a geometry was derived using the dimensions of an NMAC and the data for observed SA01a incidents. The probability of collision given that an SA01a event has occurred is estimated to be 4.6×10^{-3} .

Therefore, with Version 7 TCAS logic, the probability of mid-air collision as a consequence of an SA01a geometry is equal to 4.6×10^{-3} times 4.7×10^{-6} , or 2.2×10^{-8} per flight hour in European airspace. This corresponds to one mid-air collision every four years in Europe, given the total of 12.5 million flight hours per year.

This rate exceeds both the ICAO Target Level of Safety for midair collision and the accepted certification requirements for frequency of a catastrophic event.

Changes to TCAS RA Reversal Logic

The Committee considered many possible reasons for pilot non-compliance to RAs, including following ATC clearances or instructions that differ from the RA, the pilot's confidence in having visual acquisition of conflicting traffic, high workload, and misinterpretation of the TCAS II displays. However, the critical incompatibility arises from the fact that TCAS II selects its RAs expecting pilots to comply with those RAs and is always attempting to model maneuvers to increase separation based on this assumed pilot response. Maneuvers that are not in accordance with RAs effectively defeat the collision avoidance logic.

Two versions of TCAS are presently used in the United States. Version 6.04a allows RA sense reversals only against unequipped (non-TCAS) threats. Version 7 adds reversals against TCAS threats. The main principles are:

- The Mode S address determines which aircraft has priority for tie-breaking
- The RA sense is to be retained as long as its resolution is projected to resolve the conflict. This projection assumes that own TCAS aircraft will follow the RA, but does not assume compatible maneuvering by the other aircraft.
- Only one reversal is permitted during an encounter, except for initial tie-breaking or a multiple-threat encounter.

Change Proposal 112 Enhanced (CP112E) was presented to the Committee by the EUROCONTROL Safety Issue Rectification (SIR) project. It consists of two parts:

- Own aircraft's vertical rate is monitored to detect any non-compliance with the sense of an RA issued to it. If non-compliance is detected, that aircraft's modeling stops assuming compliance, and instead models the observed vertical rate. This feature overcomes the (useless) retention of an RA sense that would work only if followed. The revised modeling is more likely to choose the other sense when it is consistent with the non-complying aircraft's vertical rate. This change is intended to improve performance in SA01a encounters.
- A current requirement that aircraft must have more than a minimal vertical separation in order for a reversal to occur is relaxed. This overcomes a technical limitation in the present logic that was most evident in a same-sense encounter where two co-altitude aircraft maintain similar vertical rates, resulting in an NMAC. This change is intended to improve performance in SA01a and SA01b encounters.

Evaluation Metrics

The principal hazard against which TCAS is evaluated has been widely agreed to be an NMAC. While the true purpose of TCAS is to prevent actual collisions, the NMAC can be defined to represent a precise volume and thus is more straightforward to evaluate. The thresholds defining an NMAC are sufficiently small that it is a reasonable assumption that any separation that does exist is fortuitous: if a

collision does not occur, this is merely by chance. It is generally reckoned that there is a one in ten chance that an NMAC could be a collision.

A simulation, using an accepted airspace encounter model, can evaluate defined encounter situations in a more practical manner than using radar data or pilot reports. Millions of situations are generated randomly during the course of a safety study, and are archived so that each situation can be reused in a different, controlled condition. Metrics are based on compiling the statistics of various types of outcomes (e.g., whether NMACs occur) under different conditions (e.g., with or without TCAS equipage or different pilot behavior). In some cases, it is necessary to track separate pairs (or even triplets) of outcomes for every situation (e.g., when an NMAC occurs with TCAS but does not occur without TCAS). This is necessary to differentiate between several important types of NMAC, termed resolved, unresolved, and induced.

A *resolved* NMAC is an outcome in which TCAS prevents an NMAC that would occur without TCAS or with another version of TCAS. An *unresolved* NMAC corresponds to an outcome where there is an NMAC both with and without TCAS. An *induced* NMAC is an outcome in which an NMAC occurs with TCAS, when none would have occurred without TCAS or with another version of TCAS. It is important, however, that the terms “unresolved” and “induced” be interpreted carefully. Even if TCAS increases separation compared to a non-TCAS condition, there may still be some probability of induced NMAC due to altimetry errors; blame should not be placed solely on TCAS.

The core analysis of CP112E focused on five airspace encounter models: a U.S. model [3], an ICAO standard model [4], a European model [5], an SA01a encounter model developed by the SIRE team [6], and a model used to stress-test TCAS [7].

One key condition affecting the performance of TCAS is the pilot response to RAs. Three pilot response models in particular have been used in the analyses. First, the standard pilot response model [4] uses a 5 s delay to an initial RA and 0.25 g acceleration until the target vertical rate is achieved by the aircraft. A non-response model is also used, in which the pilot simply disregards any TCAS RA information and continues to do what the encounter

model specified. Finally, a slow pilot response model is also used. In the case of Lincoln Laboratory and European analysis, the slow response included a 9 s delay and a vertical rate limit of 500 ft/min when responding to RAs; the MITRE analysis applied the 9 s response delay without the vertical rate limit.

SA01a events can be generated by simulating TCAS-TCAS encounters with one pilot responding to RAs and the other pilot not responding to RAs. The non-responding pilot follows the trajectory prescribed by the encounter model, which in some cases may be contrary to a concurrent RA and result in an SA01a geometry. SA01b encounters can be generated by simulating TCAS-unequipped encounters: some fraction of these encounters involve the unequipped aircraft following a trajectory that results in SA01b events.

Risk Ratio

Risk ratio represents the change in risk in an airspace between a condition in which there is some specified mix of TCAS equipage relative to a condition in which no aircraft have TCAS. A risk ratio less than one indicates a corresponding reduction in NMAC probability. Its computation requires performing two simulation runs, one under each TCAS condition, over identical encounter situations. Risk ratio is valid only for the specific conditions tested, including the airspace encounter model, intruder equipage, and pilot response model.

$$\text{Risk Ratio} = \frac{P(\text{NMAC}) \text{ with TCAS}}{P(\text{NMAC}) \text{ without TCAS}}$$

Risk ratio is composed of two effects: the degree to which TCAS cannot resolve pre-existing NMACs and the degree to which TCAS induces NMACs. While risk ratio is a concise metric to assess overall safety benefit by equipping with TCAS, it does not provide visibility into the relative frequency of outcomes such as unresolved NMACs, induced NMACs, or NMACs that involved reversals.

Induced Risk Rate

The induced risk rate is the absolute rate (or frequency) of induced NMAC events over a defined period of time (e.g., per flight hour, per year, or per

simulated encounter situation) and over a defined region (e.g., an entire airspace or sector).

Induced risk rate directly measures how often induced NMAC events occur per unit time or per set of encounter situations. If the unit time and encounter sets are held constant between test conditions, the metric aids in making direct comparisons of induced risk due to TCAS. This metric requires computing separate outcome pairs (with TCAS and without TCAS) for every encounter scenario.

One advantage of induced risk rate is that it is not sensitive to the baseline NMAC risk without TCAS, as risk ratio is. However, an estimate of flight hours (or years) per close encounter situation is required.

Status-Quo and Upgrade Risk Rates

If it is assumed that TCAS is not upgraded with CP112E, the status-quo risk rate represents the rate with which an NMAC will occur that would have been prevented had TCAS been upgraded with CP112E.

Alternatively, if it is assumed that TCAS is upgraded with CP112E, the upgrade risk rate represents the rate with which an NMAC will occur that would have been prevented had TCAS not been upgraded with CP112E.

When status-quo risk rate is greater than the upgrade risk rate, there is a safety benefit by upgrading with CP112E. Status-quo and upgrade risk rates are computed using the rate of joint NMAC events with Version 7 and with CP112E. Both metrics require an accurate estimate of the number of flight-hours per encounter situation.

Evaluation Results

Five independent evaluations of CP112E were performed using complementary evaluation tools. Models of the U.S. and European airspaces attempt to reproduce accurate proportions of relevant close operational encounters. The ICAO model is a blend of the preceding two airspaces' statistics. The SA01a model concentrates the characteristics of susceptible encounters to generate more instances of this relatively rare event.

For each model, multiple runs were performed to evaluate various combinations of TCAS equipage, altitude reporting quantization (25 or 100 ft), and pilot maneuvering behavior.

Details of each evaluation and more comprehensive results are reported in [1].

U. S. Airspace Model

The evaluation by MITRE Corp. used the same Monte Carlo simulation and the same model of U.S. airspace as was used to evaluate Version 7 logic [3]. Each run evaluated 1,086,000 encounters. The primary evaluation metric was the risk ratio. Supplementary metrics were used to measure operational acceptability. A substantial improvement provided by CP112E ("CP") was found in comparison to Version 7 ("V7") for the non-responding cases. These runs, of which a portion of encounters exemplify the contrary response situation of SA01a, confirm the success of the change proposal. Table 1 presents the risk ratios for TCAS-TCAS encounters with 100 ft altitude quantization, for various combinations of equipage and response to RAs.

Table 1. Risk Ratios for TCAS-TCAS

Respond to RAs	V7-V7	V7-CP	CP-CP
Both Aircraft	2.3%	2.3%	2.3%
One Aircraft	9.7%	9.4%	8.1%

The encounters in the U.S. database also were simulated to evaluate the operational acceptability of CP112E relative to Version 7. For responding cases, virtually no increase in reversal rate was observed between Version 7 and CP112E. Some increases in reversal rates did occur when using CP112E in non-responding cases, as would be expected. These remain at a very small fraction of the total RAs. No evidence is seen that the change proposal causes any adverse operational effect.

Table 2. Risk Ratios for TCAS vs. TCAS encounters, standard vs. standard or standard vs. non-response.

Encounter Model					
	European			ICAO	
Respond to RAs	V7-V7	V7-CP	CP-CP	V7-V7	CP-CP
Both Aircraft	2.00%	2.00%	2.00%	1.00%	1.00%
One Aircraft	23.10%	21.70%	20.20%	42.20%	32.00%

U.S. Simulation of European, ICAO, and SA01a Models

MIT Lincoln Laboratory performed an independent set of simulations to evaluate CP112E. Three encounter models were used: the European model was used to generate 1,000,000 encounters in each run; the ICAO model was used to generate 1,140,000 encounter scenarios in each run, and the SA01a model was used to generate 250,000 encounters in each run.

In TCAS vs. TCAS encounters with both aircraft following standard responses (Table 2), when comparing Version 7 to CP112E there was an increase in risk ratio that is too small to appear in the table. There was a reduction in risk ratio from Version 7 to CP112E in the no-response condition (where only one aircraft responded to its RA), demonstrating to some degree the benefit from CP112E in SA01a encounters. Performance in slow response conditions (not shown) showed little difference between Version 7 and CP112E. Mixed equipage encounters (Version 7 vs. CP112E, tested only in standard-standard and standard-no response conditions) had risk ratios in between the corresponding same-equipage conditions.

For TCAS vs. unequipped encounters (Table 3), CP112E reduced risk ratio in the standard response condition. These results primarily demonstrate the achieved benefit from CP112E in SA01b encounters. There was a smaller benefit from CP112E when slow responses were used.

Induced risk rate was also computed assuming a baseline (non-TCAS) NMAC rate of 3×10^{-7} per

flight hour [8]. The induced risk rates ranged over two orders of magnitude, with the smallest rate at approximately 1.4×10^{-9} per flight hour for the TCAS vs. TCAS standard response case. Variation in induced risk rate between Version 7 and CP112E was similar to the variation in risk ratios, indicating that CP112E had the largest impact on induced risk. It is also worth noting that the induced risk rate when the intruder did not respond to TCAS was approximately 40 times larger than the rate when the intruder did respond to TCAS, underscoring the importance of following RAs. If both aircraft responded slowly, induced risk rate was approximately 100 times larger than if both aircraft responded promptly.

Table 3. Risk Ratios for TCAS vs. Unequipped encounters, standard response

Encounter Model			
European		ICAO	
V7	CP	V7	CP
23.1%	22.1%	28.6%	26.2%

The relative reduction in NMAC probability was also computed for the SA01a encounter model. CP112E was most effective in altitude layers 2-5, where it reduced NMAC probability by approximately 50% relative to Version 7. Performance in layer 1 was still positive, but CP112E only reduced NMAC probability by approximately 24%. A key limitation at lower altitudes is that RAs are issued only 15-20 s before closest point of approach, and so there is little

opportunity for CP112E to observe non-compliance in order to trigger an RA reversal. When all layers were weighted together, the overall reduction in NMAC probability was 46%. Benefit from CP112E in slow response cases was significantly smaller than in standard response conditions, with an overall reduction in NMAC rate of approximately 10%.

Figure 3 shows the status-quo and upgrade risk curves as a function of the airspace RA response rate for the European model (ICAO model results are similar). When all aircraft follow RAs (RA response rate = 1), upgrade risk rate is larger than status-quo risk rate, though both are small ($\sim 10^{-10}$ per flight hour and $\sim 10^{-13}$ per flight hour, respectively). When there is some degree of RA non-conformance, the status-quo risk quickly rises to become approximately seven times greater than the upgrade risk. Status-quo risk can be as large as 5.2×10^{-9} per flight-hour; the upgrade risk never exceeds 1×10^{-9} per flight-hour. Crossover between status-quo and upgrade risk curves occurs at approximately 99% RA compliance, a rate that does not appear realistic based on the monitoring results.

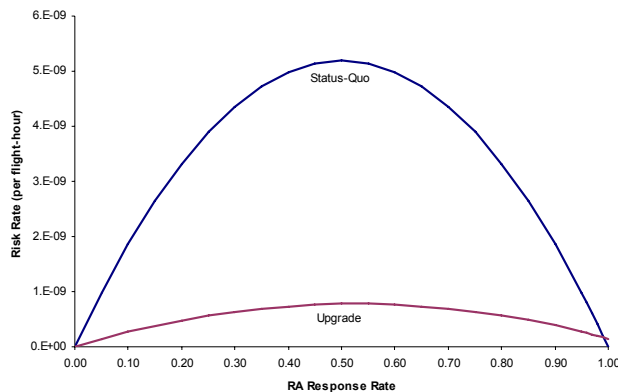


Figure 3. Status-quo vs. Upgrade Risk

To summarize, the Lincoln Laboratory evaluations showed that CP112E provides a significant reduction in collision risk in the SA01a and SA01b conditions for which it was designed. In TCAS vs. TCAS SA01a encounters where only one aircraft responds promptly to its RAs, CP112E cuts the probability of collision nearly in half compared to Version 7. In the European and ICAO encounter models, CP112E reduces collision risk to approximately 75%-85% of the Version 7 level when only one aircraft responds promptly to its

RAs. In TCAS vs. unequipped encounters (SA01b), CP112E provides a less-significant but still positive reduction in collision risk (reducing risk to approximately 92%-95% of the Version 7 level).

Decision risk analysis supports upgrading to CP112E. From a decision-making standpoint, choosing to remain with Version 7 carries the risk that a collision will occur that CP112E would have prevented. Alternately, choosing to upgrade with CP112E carries the risk that a collision will occur that Version 7 would have prevented. If pilots always promptly follow their RAs, these risks are very small (on the order of 10^{-10} per flight hour) and there is no incentive to upgrade with CP112E. However, it is clear that pilots do not always follow their RAs, in which case the data show that the decision risk associated with remaining with Version 7 is approximately seven times larger than the risk associated with upgrading with CP112E. Given the supporting monitoring data that show a significant rate of late, weak, or no compliance with RAs, it appears that there is a strong incentive to upgrade with CP112E.

European Simulations

Key safety metrics were computed, both for Version 7 and CP112E, on a large number of scenarios to assess the performance of CP112E over a broad spectrum of TCAS equipage and pilot response combinations. These scenarios were all investigated on three different encounter models (ICAO, European, and US-like) in order to mitigate the specificities of each model and to develop a complete picture. Additionally, key metrics were also computed on a specific safety encounter model dedicated to the SA01a issue.

Table 4 presents results using the various models, for TCAS-TCAS encounters with one aircraft not responding to RAs. The different simulations that have been run clearly show that on the SA01a issue, CP112E significantly outperforms Version 7: CP112E reduces the risk of NMAC by up to 36%, induces up to 86% fewer NMACs, and generally provides greater vertical separation. These results are confirmed by the status-quo and upgrade risk metrics, which show a clear benefit in upgrading from Version 7 to CP112E.

Table 4. Key Metrics, Responding/Non-Responding Scenario

Encounter model	CAS logic	Risk ratio	Induced risk rate $\times 10^{-9}$
US-like	V7	12.6%	29.9
	CP112E	9.1%	19.6
European	V7	19.0%	45.1
	CP112E	16.8%	38.7
ICAO	V7	11.0%	22.7
	CP112E	7.0%	12.2
SA01a	V7	15.9%	15.9
	CP112E	7.2%	10.8

Regarding scenarios with unequipped threats, CP112E also provides noticeable benefits over Version 7. CP112E reduces the risk of NMAC by up to 6%, induces 13% fewer NMACs, and also provides greater vertical separation than Version 7 in similar situations. These results are also confirmed by the status-quo and upgrade risk metrics, which show that there is a benefit in upgrading from Version 7 to CP112E.

Fast Time Encounter Generator “Stress Testing”

As was done for previous logic versions, the FAA William J. Hughes Technical Center and MIT Lincoln Laboratory performed logic “Stress Testing” using fast-time encounter generation. Aircraft parameters (e.g., planned vertical separation at CPA, vertical speed, and acceleration) used in these encounters are designed to span and exceed the typical values observed in the airspace. Aircraft maneuvers are timed to generate worst-case situations for TCAS in order to test the performance limits of the system. This stress testing provides an assessment of the strengths and weaknesses of the logic without representing any particular airspace.

The Technical Center simulated 5 million encounters across vertical profiles, quantization levels, and pilot response conditions. Several

thousand of these produced differences between Version 7 and CP112E. Lincoln Laboratory performed detailed analysis of NMAC events that resulted, looking in particular for patterns that could identify deficiencies in performance. The results of both efforts confirmed that the Change Proposal improves logic performance, and that no geometry classes indicate a performance deficiency for the Reversal logic.

Code Evaluation

A code evaluation effort was performed, independent of the various simulations, to gain a full understanding of the single-threat reversal logic in Version 7 and CP112E and to ensure that the proposed CP112E pseudo-code is consistent with its high-level design principles. This work was primarily performed by committee members from Johns Hopkins University Applied Physics Laboratory and MIT Lincoln Laboratory.

The manual code evaluation effort identified a few significant logical errors, verified conformation to the MOPS format, and polished the proposed pseudo-code changes. It also conducted coupling analysis confirming that CP112E's impact is confined to the single-threat reversal logic and affects other CAS functionalities only through the Version 7 reversal logic. Two by-products of the evaluation effort are a CP112E design principles report [1] and four decision tables [1]. The design principles report is a thorough explanation of CP112E logic changes with traceability to the pseudo-code. It also includes a set of logic-flow diagrams and a set of data-flow diagrams illustrating the coupling analysis. The decision tables offer a concise high-level summary of the threat declaration, sense selection, Version 7 reversal logic, and CP112E reversal logic, highlighting all of the main decision criteria at a glance.

Conclusions and Recommendations

The evaluation of CP112E led to the following conclusions and recommendations [1]:

Conclusions

- A significant safety vulnerability has been identified in the TCAS Version 7 RA reversal logic. The vulnerability is termed SA01, and encompasses a class of TCAS-TCAS and TCAS-unequipped encounters where a necessary RA reversal is not issued in a timely manner (if at all).
- Based on observations in European airspace, it is estimated that SA01 events occur at a frequency of 4.7×10^{-6} per flight hour, corresponding to an estimated mid-air collision rate due to SA01 of 2.2×10^{-8} per flight hour. The U.S. monitoring indicates a rate consistent with European airspace. This risk is unacceptable because the observed frequency of SA01 exceeds that which is tolerated for catastrophic hazards.
- The CP112E change will provide a significant improvement in performance for SA01 events. For example, analyses of SA01a events indicate that CP112E would reduce the collision rate due to SA01 to 30 to 50% of the rate with Version 7. CP112E achieves this reduction by revising several assumptions made for the Version 7 logic reflecting operational experience.
- The evaluation shows the greatest improvement where all aircraft carry CP112E. Improvement is even seen for airspace in which some aircraft carry CP112E while others carry other versions. No problems of interoperability between versions have been found.
- Side effects and performance degradations are minimal for CP112E and are considered acceptable compared to the collision risk with current versions of TCAS II. No RA reversal logic can be expected to be perfect, given inherent limitations such as altitude tracking lag and variable pilot response. The evaluation effort compared CP112E to existing versions of TCAS using several complementary methods and airspace models. The evidence strongly indicates that the benefits of CP112E outweigh its limitations.

Recommendations

- It is recommended that FAA and international authorities commence work towards regulatory action that would expedite implementation of the revised logic, as now defined. Safety would be improved as soon as the change can be installed in the TCAS fleet.

Regulatory measures could include issuance of Airworthiness Directives, requirements to enhance pilot and controller training so as to minimize the occurrence of the observed problems, and mandatory equipage of the change by specific dates for both reverse and forward fit.
- It is recommended that RTCA proceed with a revision to TCAS MOPS based on the CP112E change to the RA reversal logic.
- It is recommended that airspace monitoring be expanded to assess the performance of TCAS in the changing airspace.
- It is recommended that resources of expertise in TCAS technical analysis and supporting tools for simulation be sustained.

As a result of the evaluation, SC-147 is commencing work on a revision to the TCAS MOPS that would revise the RA Reversal Logic. FAA and EASA, in association with EUROCONTROL, are investigating the associated rulemaking measures. ICAO may also revisit the ACAS SARPS.

References

- [1] DO-298, 2005, "Safety Analysis of Proposed Change to TCAS RA Reversal Logic," RTCA, Washington D.C.
- [2] Chabert, Stéphan, July 2, 2003, "Reversal Logic of TCAS II Version 7.0 Overview," EUROCONTROL, SIR/WP2/04/W., Toulouse, France.
- [3] McLaughlin, Michael P., June 1997, "Safety Study of the Traffic Alert and Collision Avoidance System (TCAS II) Final Version", MTR 97W32, The MITRE Corporation, McLean, VA.
- [4] ICAO, 1998, "ICAO Standards and Recommended Practices – Annex 10, Volume IV,

Surveillance, Radar and Collision Avoidance Systems”, Montreal, Canada.

[5] CENA/Sofréavia and QinetiQ, March 2002, “European encounter model – Specifications and probability tables”, ACASA/WP1/186/D, Toulouse, France.

[6] Chabert, Stéphan, 1 April 2005, “Safety Encounter Model Focused on Issue SA01a”, SIRE/WP2/21/D, Toulouse, France.

[7] Choyce, Thomas A., Kathryn M. Ciaramella, October, 2000, “Test and Evaluation of TCAS II Logic Version 7”, Federal Aviation Administration, Atlantic City, NJ.

[8] Carpenter, Ken, 28 June 2000, “NMAC Rate”, ACASA/WP1.1.7/115/W Malvern, UK.

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.

Work by Lincoln Laboratory was sponsored by the Federal Aviation Administration under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the United States Government.

Approved for Public Release; Distribution Unlimited.
Case Number: 06-0988

*25th Digital Avionics Systems Conference
October 15, 2006*