AERONAUTICAL SURVEILLANCE PANEL

(Montréal, Canada, May 1-5, 2006)

Evaluation of Proposed Changes to the ACAS Modified Tau Calculation

(Prepared by James K. Kuchar)

Summary

Modified tau is a parameter computed by ACAS to estimate the earliest time at which a collision could occur should an intruder aircraft accelerate toward the own aircraft. A concern with the modified tau calculation has been raised in a class of encounters where intruders are already close and converging slowly. In these problem cases, ACAS may induce a Near Mid-Air Collision by generating RAs with inappropriate timing or initial sense or failing to reverse sense when necessary. Performance in some problem encounters is greatly improved when using several proposed changes to the modified tau equations. These changes are outside CP112E, which focuses only on RA reversals. Although changes to modified tau resolve some problem encounters, aggregate risk-ratio results do not support implementing the existing proposals. There remains a concern about mid-air collision risk due to vulnerability in the existing modified tau equations, yet a robust solution to the problem has not been developed.

References:

1. K. M. Carpenter, “The ACAS Critical Interval”, SICASP/WG2 WP2/723, Brussels, Belgium, 9 October 1998.
2. N. Phamdo, H. Men, and F. Deal, “A Critical Examination of the Critical Interval”, RTCA SC-147 RWG Presentation, Washington, DC, 27-29 September 2005.
3. S. Chabert, “SIR Final Report”, SIR/WP3/20/D, 16 July 2004.
4. K. M. Carpenter, “Comments on Critical Interval”, RTCA SC-147 RWG, 10 January 2006.
5. H. Drévillon, “CP112E Pseudo-Code Baseline,” EUROCONTROL, SIRE/WP2/50/W – Version 1.1, 19 September 2005.
6. ICAO, Section 4.4 of Annex 10 – Aeronautical Telecommunications, Volume IV, 1998
7. T. Miquel and K. Rigotti, “European Encounter Model” Parts 1 and 2, ACASA/WP1.1/186, December 4, 2001.
8. H. Drévillon, “Evaluation of the Change to the Critical Interval Calculation on the New European Encounter Model”, SIRE+/WP1/04/W, 10 February 2006.

This work is sponsored by the Air Force 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 U.S. Government.

1. Background on the Problem

1.1. The so-called 'critical interval' defines a window of time in which a collision might occur with an intruder being tracked by ACAS. This interval extends between "true tau", computed as $\tau = -r / \dot{r}$, and "modified tau", computed as $\tau_{\text{mod}} = -(r^2 - DMOD^2)/(r\dot{r})$ as shown in Appendix A. A constraint is also placed on modified tau such that if $r <$ $DMOD$, $\tau_{\text{mod}} = 0$. True tau is an estimate of the time to collision, assuming that the closure rate between aircraft remains constant. Modified tau is always shorter (less) than true tau, and was based on an assumption that the intruder could accelerate toward the own aircraft at a sustained 1/3 g. The intent, then, was to protect against acceleration between the two aircraft by estimating a lower-bound on the time to collision.

1.2. When determining whether to issue an RA, separation is first-estimated by projecting linear trajectories from own and intruder aircraft (Figure 1). Two vertical separations between these projections are measured, one at the start and one at end of the critical interval. The smaller separation value is used to determine whether an RA is required. Should the two aircraft cross in altitude between the start and end of the critical interval, the projected separation is defined to be zero.

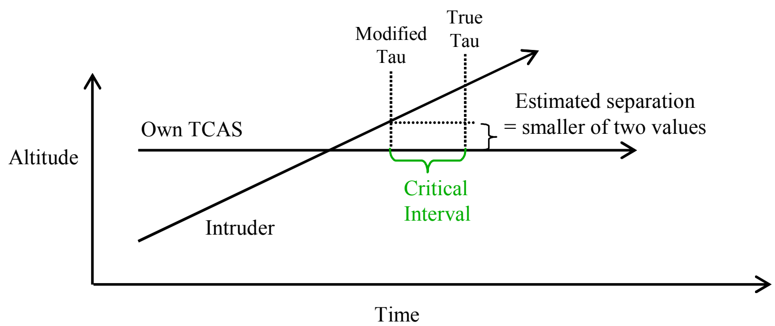

Figure 1: Critical Interval and Separation Estimates

1.3. Once ACAS determines that an RA should be issued, it then makes additional projections of climb and descent maneuvers. The critical interval is also used here to estimate vertical separation along these trajectories, though the details are somewhat more complex than discussed in Section 1.2. The key point is that a change in the duration of the critical interval can produce a change in the predicted separation, which in turn may affect RA timing, initial sense, or reversals.

1.4. A change to the modified tau calculation was proposed by Ken Carpenter in 1998 [1]. The proposal noted that the existing modified tau is an inaccurate simplification of the actual quartic equation needed to compute time to collision assuming constant acceleration between two aircraft. The use of a different, quadratic solution was proposed as an alternative. This quadratic solution relied on an alternate assumption in which acceleration increased as the inverse of the time to collision.

1.5. The 1998 proposed change to the modified tau calculation was rejected (as recommended by its author) due to concerns over the delay such a change would impose on the schedule of equipping aircraft with ACAS. Part of the concern was that changing the modified tau calculation could interact in a negative way with other elements of the ACAS logic, and so a more widespread analysis and redesign might be required.

1.6. In 2005, as part of the RTCA SC-147 Requirements Working Group (RWG) evaluation of CP112E, concerns over the critical interval calculation were again raised. An analysis by Johns Hopkins University Applied Physics Laboratory (JHU APL) proposed an exact solution to the quartic equation for modified tau (reproduced in Appendix A) [2]. The exact solution is somewhat more computationally intensive than the existing modified tau equation, relying on one cube-root and several square-root operations. The APL analysis also highlighted the extent to which the existing modified tau varies from the exact value computed using the quartic equation. The existing modified tau could underestimate the time to collision by more than 20 seconds in certain low closure-rate situations, leading to incorrect estimates of vertical separation and incorrect RA timing or sense selection. Figure 2 shows one example (sensitivity level 7 with $\dot{r}$=-500 ft/s) where the existing critical interval at a range of 1.1 nmi is approximately 13 s, whereas the critical interval using the exact modified tau equation is only 0.25 s.

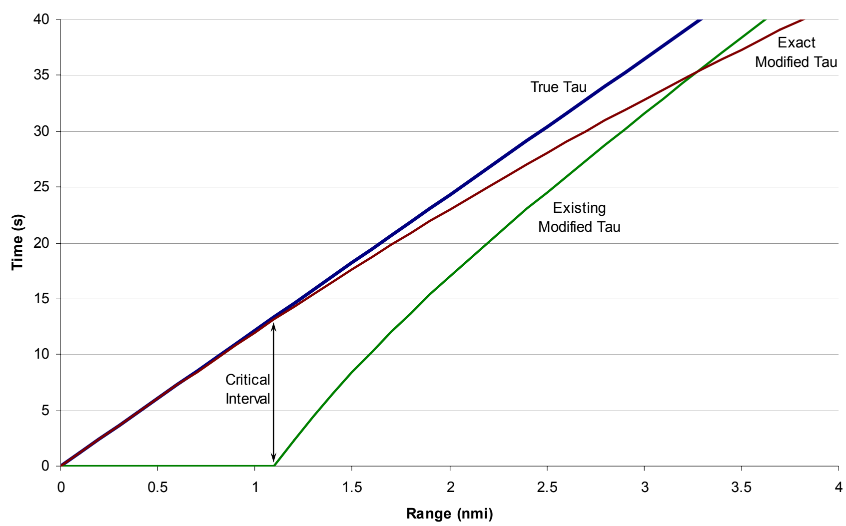

Figure 2: True Tau, Existing Modified Tau, and APL Exact Modified Tau
The critical interval spans the time between modified tau and true tau
(ACAS Sensitivity Level 7, 500 ft/s range rate)

1.7. The APL modified tau is called the exact modified tau in this paper because it represents the exact (as opposed to approximated) solution to the quartic equation assuming 1/3 g acceleration. It should be noted, however, that a constant 1/3 g acceleration is only one of many possible assumptions, and so even the 'exact' modified tau may not be the most accurate or representative.

1.8. Also in 2005, Lincoln Laboratory had noted a certain set of problem encounters in its ACAS simulations in which an initial RA sense was not reversed against an unequipped intruder in a vertical chase, resulting in Near Mid-Air Collisions (NMACs). Although this general type of encounter had been previously identified (SA01b) by Sofréavia and DSNA during the development of CP112E [3], the Lincoln encounters were not remedied by CP112E and were traced specifically to the modified tau calculation.

1.9. The main concern raised by the Lincoln encounters was that in vertical-rate encounters with an intruder closing slowly in range (i.e., on nearly parallel tracks), the existing modified tau would be significantly smaller than the actual time to closest point of approach. As shown in Figure 3, this could lead ACAS to project the aircraft as crossing in altitude during the critical interval when in reality one aircraft might be clearly above the other during the critical interval.

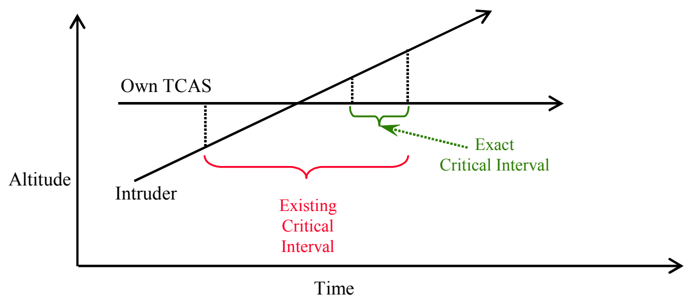

Figure 3: Potential Effect of Existing vs. Exact Modified Tau on Estimated Separation

1.10. A significant proportion of the Lincoln problem encounters were successfully resolved when re-simulated using the APL or 1998 modified tau equations, suggesting that repairing the modified tau calculation could improve the robustness of ACAS. For example, Figure 4 shows an encounter using the existing modified tau calculation where a non-crossing descend RA is issued that leads to a vertical chase and an NMAC at the Closest Point of Approach (CPA). With the exact modified tau calculation (right side of Figure 4), a more effective crossing climb RA is issued instead.

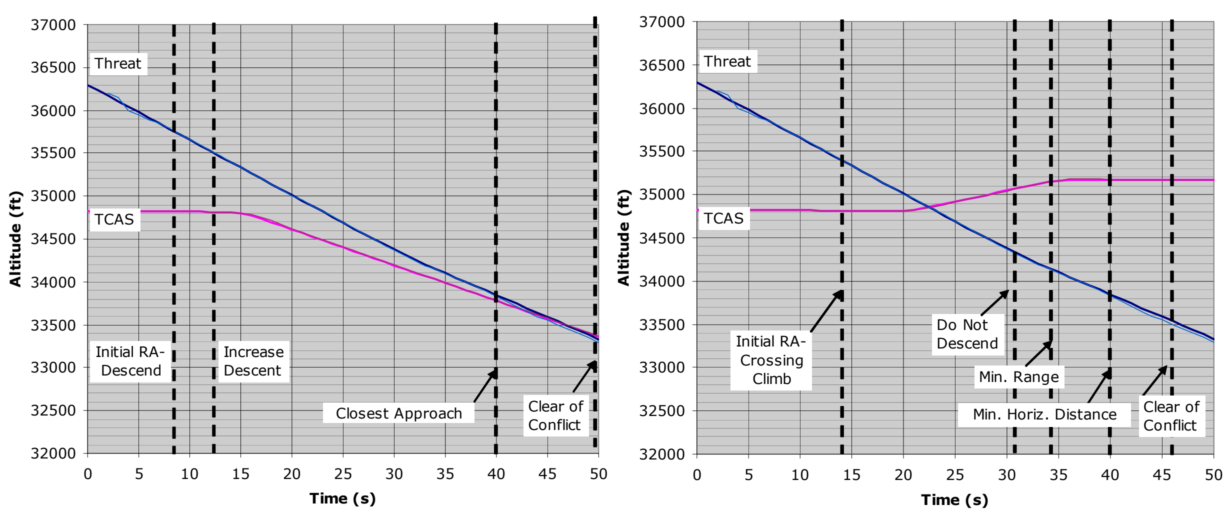

Existing Modified Tau | Exact Modified Tau

Figure 4: Example Problem Encounter

1.11. It should be noted that the existing modified tau calculation appears to be acceptable when making calculations up until the point of crossing the range threat boundary. Concerns with modified tau are restricted to cases in which threat declaration is delayed until well after the range threat boundary is crossed. This may occur, for example, when waiting for the altitude threat criteria to be met or in the case of a pop-up threat at close range. Modified tau may also be a problem when projected separation calculations are made well inside the range threat boundary in order to determine if a change in RA strength of an RA reversal needs to be performed.

1.12. To summarize the situation up to the end of 2005, a problem with the existing modified tau method had been identified in which RAs might be issued at an improper time, with an improper sense, and/or might not be reversed when appropriate. The existing modified tau calculation could lead to using a projection time much shorter than it should be when predicting vertical separation. CP112E was targeted at a different set of problems in the ACAS logic and did not directly affect the issues related to modified tau. Two solutions to modified tau had been proposed, each of which appeared to resolve these problems in initial analyses.

1.13. Between December 2005 and January 2006, RTCA SC-147 RWG undertook a more extensive but rapid analysis of the proposed changes to the modified tau calculation by running a large set of encounters using three airspace models: 1980-1990s U.S. airspace, ACAS Analysis (ACASA) European airspace, and a new ACAS Safety Analysis Post-RVSM Project model (ASARP). Simulation facilities at MITRE, Lincoln Laboratory, and Sofréavia and DSNA were used to analyze each of these three models, respectively. The goal was to determine whether a change to modified tau might be warranted without delaying the implementation of CP112E.

1.14. In early January 2006, Ken Carpenter proposed one more solution for modified tau that used a variation on his 1998 concept. In the 2006 proposal, acceleration was modeled as increasing in inverse proportion to the time to collision raised to the 3/2 power, resulting in the pseudocode implementation shown in Appendix A [4]. The goal was a solution that provided the general characteristics of the exact modified tau calculation proposed by APL but was simpler to calculate.

1.15. Lincoln Laboratory evaluated the performance of the existing modified tau calculation, the exact APL proposal, and the 2006 proposal by Ken Carpenter using the European encounter model. The remainder of this information paper describes the Lincoln Laboratory analysis and its results.

2. Analysis Method

2.1. The Lincoln Laboratory Collision Avoidance System Safety Assessment Tool (CASSATT) was used to compute risk ratios as a function of (1) critical interval calculation method, (2) altitude layer, (3) aircraft equipage, and (4) pilot response. TCAS Version 7, including the baseline CP112E algorithms, was used [5].

2.2. CASSATT ran pairwise encounters using 1 s updates for ACAS and 0.1 s timesteps for aircraft dynamics and pilot response. The ACAS units on the two aircraft were modeled to update with a 0.5 s time offset between them. ICAO-specified range, bearing, and altimetry errors were included [6]. ACAS-equipped intruders were assumed to have 25 ft altimeter quantization; non-ACAS intruders were assumed to have 100 ft quantization. A total of 200,000 encounters were simulated for each critical interval / equipage / response condition at each of five altitude layers using the ACASA European encounter model specifications [7]. Altitude layer weights for the ACASA model are shown in Table 1.

Table 1: ACASA Model Altitude Layer Weights [7]

Altitude Layer	Weight
1	0.392525
2	0.328396
3	0.153016
4	0.121628
5	0.004414

2.3. Three critical interval (modified tau) calculation methods were examined:

(1) “Baseline”: current method used in TCAS Version 7 + CP112E baseline
(2) “Rising accel”: Ken Carpenter’s proposal of 10 January 2006 using powers of 3/2
(3) “Exact”: exact computation method proposed by APL

3. Risk Ratio Results

ACAS – unequipped encounters

3.1. Figure 5 shows the risk ratio results for the ACAS vs. unequipped case. Reading across the Figure, each European model altitude layer (1-5) is presented separately, all to the left of a gray dividing line. Data for a different form of altitude layer 5 is then shown (labeled “5 – later RAs”; this will be discussed below), and the final, rightmost data set is the overall weighted risk ratio based on the 5 altitude layers from the left side of the Figure. Three bars are shown at each altitude layer, one each for the baseline, rising accel, and exact modified tau methods.

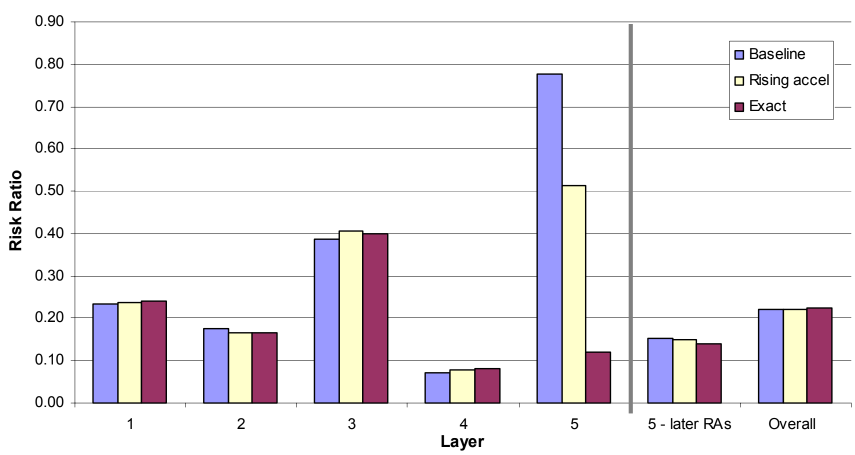

Figure 5: ACAS-Unequipped

3.2. Figure 5 shows generally mixed results. In layers 1, 3, and 4, risk ratio is increased when using a method other than the existing modified tau baseline. Layers 2 and 5 show improvement with the rising accel or exact modified tau methods. Overall, risk ratio degrades slightly when using one of the proposed modified tau methods.

3.3. Behavior in layer 5 is clearly different from the other layers, with a significantly higher baseline risk ratio. Many of the layer 5 encounters involve slow-closure scenarios that begin with the aircraft within or close to DMOD. In the baseline condition, these encounters often produce RAs very quickly and degrade into a prolonged vertical chase with small resultant Vertical Miss Distance (VMD).

3.4. Risk ratio in layer 5 is greatly improved when using the rising accel method and is especially improved with the exact modified tau method. 70% of the exact method "saves" in layer 5 occurred because the sense of the initial RA was changed. A different initial sense prevented the situation from evolving into a vertical chase and instead resulted in an altitude crossing.

3.5. Figure 6 shows VMD (with TCAS) vs. the time of the initial RA for 10,000 encounters in layer 5, for the baseline and exact cases in TCAS-unequipped encounters. CPA occurs at t = 40 s. The baseline plot (left) shows that every small-VMD case occurs with early RA times (between 5 – 10 s, or equivalently, 30 – 35 s before CPA). It is these early RA encounters that are causing the large baseline risk ratios in layer 5 of Figure 4. Changing from baseline to exact modified tau removes most of these small-VMD cases.

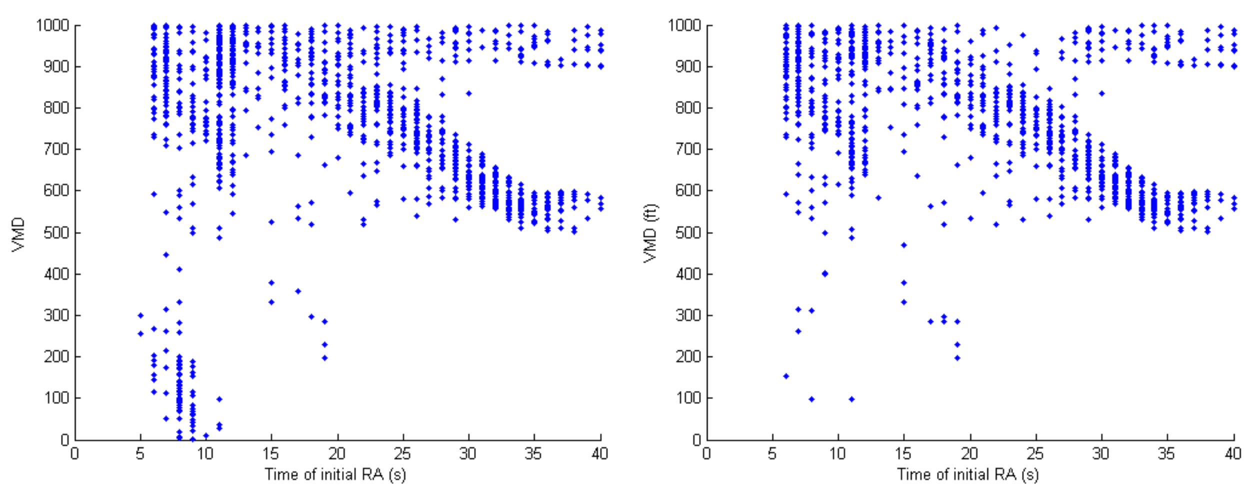

Baseline *Exact*

Figure 6: Baseline and Exact Modified Tau VMD vs. RA Time

(Layer 5, TCAS-unequipped; time measured from the start of the simulation)

3.6. Slow-closure encounter geometries only produce significant problems in layer 5 simply because these geometries are common in layer 5 and are rare in lower layers. Replicating the layer 5 encounter geometries at a lower altitude produces similarly large risk ratios in the baseline condition, and similarly improved risk ratios when using the proposed changes to modified tau. Thus, the issue is not due only to the large protected volume at high altitude. TCAS is vulnerable to these types of slow-closure encounters (regardless of whether they are operationally realistic), but slow-closure encounters are too rare to produce a measurable impact outside of layer 5. Layer 5 is only weighted by 0.004 toward the overall risk ratio, so the effect of these problem encounters on overall risk ratio is also diluted. Slow-closure problem encounters are also rare in any layer of the ICAO encounter model.

3.7. One additional note is that because simulations begin only 40 seconds before CPA and because the trajectories are designed to miss laterally by at most 500 ft, a slow-closure encounter necessarily starts with aircraft very close together. "Slow closure" is synonymous with "low initial separation". It may be valuable to examine slow-closure encounters that begin significantly earlier than 40 seconds before CPA (thus giving TCAS more time before the aircraft come close together), though this may interfere with the realism of vertical profiles and horizontal accelerations used during encounter generation. An initial analysis that extended the start time to greater than 40 s before CPA showed some improvement in baseline TCAS performance, but low-VMD problems remained.

3.8. To remove the effect of slow-closure encounters that produce early RAs, the portion of Figure 5 labeled "5 – later RAs" shows data from only those encounters in layer 5 in which RAs were issued more than 10 seconds into the simulation (or equivalently, less than 30 seconds from CPA). These data show risk ratios more in line with the other altitude layers, and show a more modest effect of changes to the modified tau calculation (but still, both proposed methods outperform the existing baseline).

ACAS – ACAS with non-responding intruder

3.9. The modified tau methods in ACAS-ACAS non-responding cases (Figure 7) only show improvement in risk ratio in layer 5. As with the unequipped case, layer 5 produces very large risk ratios (all greater than 1!) unless the early RAs are excluded.

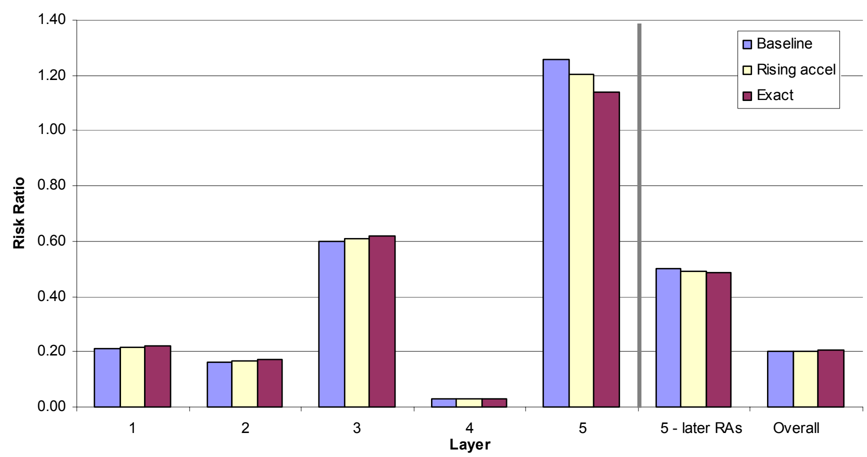

Figure 7: ACAS-ACAS Non-Responding Intruder

ACAS – ACAS both aircraft follow standard response

3.10. Results for ACAS-ACAS standard responses (Figure 8) show essentially no benefits from alternate modified tau calculations at any altitude layer. Note also that excluding the early RAs in this condition results in a higher risk ratio: ACAS-ACAS standard response cases with early RAs are handled robustly by the logic.

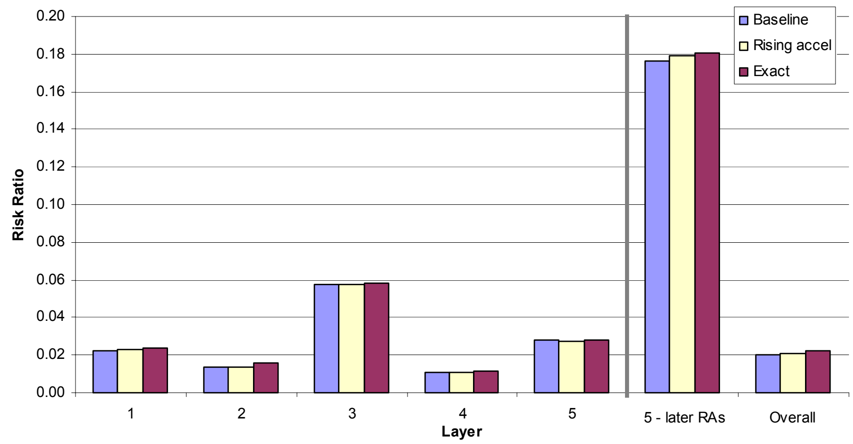

Figure 8: ACAS-ACAS Standard Responses

4. Other Analyses

4.1. All of the encounters from the Lincoln simulations were categorized by altitude layer and ACAS performance with the existing vs. proposed modified tau methods. Sets of characteristic encounters were then defined to help focus analysis on typical geometries that caused problems for either or both of the existing or proposed modified tau equations. The exact modified tau equations both improved and degraded performance across encounters at each altitude layer, though the overall trend was a slight degradation of performance in the lower layers. Detailed analyses of specific example encounters helped to describe how modified tau affected ACAS performance, but this analysis was not by itself sufficient to suggest a new robust solution to the problem.

4.2. An analysis was performed to examine the effect of modified tau on crossing RA rate. The exact method was found to produce approximately 10% more crossing RAs than there were with the baseline method, regardless of altitude layer. Some of these crossing RAs occur in layer 5 when the exact method successfully resolves an encounter that would have been an NMAC due to vertical chase under the baseline method. Additional details behind this behavior have not been investigated.

4.3. The use of a horizontal miss distance (HMD) estimate (from the range/bearing and parabolic trackers) to compute modified tau was briefly investigated. A simplistic implementation of HMD for modified tau produced worse performance than the other modified tau proposals. It is not known whether this is due to limitations in the HMD filters (which were not designed for this purpose) or due to more fundamental issues.

4.4. MITRE CAASD analyses using the U.S. encounter model and Sofréavia and DSNA analyses using the ACASA and ASARP models were in general agreement with the

Lincoln Laboratory results. All analyses showed degradations in risk ratio when using the exact modified tau method. As with the Lincoln results, a small improvement in risk ratio was observed with the exact modified tau equations in the ASARP model in ACAS-unequipped encounters in RVSM layers. The Sofréavia and DSNA authors note, however, that ACAS-unequipped encounters are currently unlikely at high altitudes [8].

5. Conclusions

5.1. A class of problem encounters has been identified involving low-initial-separation slow-closure intruders that almost immediately produce RAs. These encounters (if the intruder is unequipped or non-responding) are poorly-handled by the existing modified tau calculation (often due to vertical chase), and are resolved much more effectively by several proposed modified tau methods.

5.2. The frequency of exposure to these problem encounters is thought to be exaggerated in the Lincoln Laboratory results. The ACASA model that was used generated a larger proportion of nearly-parallel track encounters at high altitude than would be the case from a more uniform distribution of approach angles. There have not been any identified modified tau problem events in actual operations, although as we have learned from the SA01 analysis, operational problems may exist that are not detected until focused monitoring is conducted. The new ASARP model developed by Sofréavia and DSNA does not appear to have the same concentration of slow-closure encounters in the highest altitude layer as the Lincoln Laboratory version of the ACASA model [8]. Given that the ASARP model incorporates several improvements over the earlier ACASA model (including RVSM data), it is likely that performance assessed using the ASARP model is more representative of the current airspace. The exposure to problem encounters is therefore expected to be smaller than the ACASA model predicts. The development of an updated U.S. airspace model is strongly encouraged so that the exposure to critical interval risk or other issues can be assessed more accurately.

5.3. With the exception of the slow-closure problem encounters, there does not appear to be a benefit from using any of the proposed modified tau methods – risk ratios are generally degraded by doing so.

5.4. RTCA SC-147 RWG has recommended that implementation of CP112E not be delayed in order to carry out work on modified tau. It appears that more extensive logic changes would be needed beyond simply changing the modified tau equations. Due to the complexity of the CAS logic, the development of an acceptable solution could span several years of work (perhaps on the order of the CP112E effort).

5.5. However, it should also be emphasized that a vulnerability in TCAS does in fact exist. Although it appears unlikely, a slow-closure encounter at close range could occur that results in a collision induced by TCAS because of the current modified tau calculation. A probability for such an event has not been estimated.

5.6. The ASP is invited to note these findings and to consider placing additional attention on remedying the modified tau vulnerability as soon as an opportunity presents itself.

Appendix A: Modified Tau Equations

In the DO-185A pseudocode, `T1` represents modified tau.

A.1 Current implementation (DO-185A):

```
IF (ITF.R GT DMOD)
      THEN T1 = -(ITF.R–((DMOD**2)/ITF.R))/RDTEMP;
      ELSE T1 = P.MINTAU;
```

A.2 APL exact implementation [2]:

```
IF (ITF.R LE 0)
      THEN T1 = P.MINTAU;
      ELSE XX = SQRT(27)*(RDTEMP**2)/(8*(32.2/3)*ITF.R);
           YY – (SQRT(XX**2+1)-XX)**(0.33333);
           ZZ = SQRT(0.33333)*(1/YY-YY);
           T1 = SQRT(ITF.R/(32.2/3))*(SQRT(2*SQRT(ZZ**2+1)-ZZ)-SQRT(ZZ));
```

A.3 Rising accel implementation [4]:

```
T1 = (ITF.R**2)/(-ITF.R*RDTEMP+(DMOD**2)/TRTHR);
```