MET'2015
return

MET'2015 Croatia
Trogir, October 4 - 7, 2015
12th International Conference
Modern Electrified Transport

THE TRAFFIC ALERT AND COLLISION AVOIDANCE SYSTEM

Petar Obradović(1), Slobodan Obrenović(2)
Bosnia and Herzegovina Air Navigation Services Agency(1), Zagreb, Croatia (2)

Abstract

A collision between aircraft is one of the most sudden and catastrophic transportation accidents imaginable. These tragic events are rarely survivable – hundreds of people may die as the two aircraft are destroyed. Some airborne systems have been developed and are currently in use to prevent mid-air collisions. This article focuses on the widely fielded, crucial technology called the Traffic Alert and Collision Avoidance System (TCAS). TCAS has had extraordinary success in reducing the risk of mid-air collisions. Now mandated on all large transport aircraft, TCAS has been in operation for more than two decades and has prevented several catastrophic accidents. TCAS is a unique decision support system in the sense that it has been widely deployed (on more than 25,000 aircraft worldwide) and is continuously exposed to a high-tempo, complex air traffic system. TCAS is the product of carefully balancing and integrating sensor characteristics, tracker and aircraft dynamics, maneuver coordination, operational constraints, and human factors in time-critical situations. Missed or late threat detections can lead to collisions, and false alarms may cause pilots to lose trust in the system and ignore alerts, underscoring the need for a robust system design.

Introduction

Over the years, air traffic has continued to increase. The developments of modern air traffic control systems have made it possible to cope with this increase, whilst maintaining the necessary levels of safety. The risk of collisions is mitigated by pilots exercising the “see and avoid” principal and staying away from other aircraft and by ground based Air Traffic Control (ATC) which is responsible for keeping aircraft separated. Despite technical advances in ATC systems, there are cases when the separation provision fails due to a human or technical error. Any separation provision failures may result in an increased risk of a mid-air collision.

To compensate for any limitations of “see and avoid” and ATC performance, an airborne collision avoidance system, acting as a last resort, has been considered from the 1950s. In 1955, the use of the slant range was proposed between aircraft divided by the rate of closure or range rate for collision avoidance algorithms, i.e. time rather than distance, to the Closest Point of Approach (CPA). Today’s airborne collision avoidance system is based on this concept [1].

In 1956, the collision between two airliners, over the Grand Canyon in the USA, prompted both the airlines and the aviation authorities to advance the development of an airborne collision avoidance system. It was determined in the early 1960s that, due to technical limitations, the development could not be progressed beyond the overall concept.

During the late 1960s and early 1970s, several manufacturers developed prototype aircraft collision avoidance systems. Although these systems functioned properly during staged aircraft encounter testing, it was concluded that in normal airline operations, these systems would generate a high rate of unnecessary alerts in dense terminal areas. This problem would have undermined the credibility of the system with the flight crews.

In the mid-1970s, the Beacon Collision Avoidance System (BCAS) was developed. BCAS used reply data from the Air Traffic Control Radar Beacon System (ATCRBS) transponders to determine an intruder’s range and altitude.

In 1978, the collision between a light aircraft and an airliner over San Diego, California led the US Federal Aviation Administration (FAA) to initiate, three years later, the development of TCAS (Traffic Alert and Collision Avoidance System) utilizing the basic BCAS design for interrogation and tracking with some additional capabilities.

Despite the terrifying prospect of a mid-air collision, aviation travel is incredibly safe. A person who flew continuously on a jet transport aircraft in today’s environment could expect to survive more than 11,000 years of travel before becoming the victim of a mid-air collision. This accomplishment has only recently been realized. The number of hours flown annually by jet transport aircraft has more than quadrupled since 1970, but the rate of mid-air collisions over that period of time has dropped by an order of magnitude. The result is that today we can expect one mid-air collision every 100 million flight hours. Such an exceptional safety level was achieved through advances in air traffic surveillance technology and relentless attention to improving operational procedures. TCAS is one component of a multi-layered defense against mid-air collisions. The structure of airspace and operational procedures provide the first strategic layer of protection. Traffic flows are organized along airways at segregated altitudes to aid air traffic controllers in managing aircraft and predicting potential conflicts well before problems arise. Aircraft are normally kept three to five miles apart laterally or 1000 ft vertically, to provide sufficient safety margins. Air traffic control ensures that separation minima are not violated by issuing tactical commands (including altitude restrictions and heading change vectors) to the pilots in response to nearby traffic. Should these nominal traffic separation processes fail, the TCAS system aids pilots in visually acquiring potential threats and, if necessary, provides last-minute collision avoidance guidance directly to the flight crew.

It is obviously imperative that TCAS alert the flight crew early enough that evasive action can be taken. But it is also important that TCAS does not alert unnecessarily. Collision avoidance alerts represent high-stress, time-critical interruptions to normal flight operations. These interruptions, in addition to distracting the aircraft's crew, may lead to unnecessary maneuvering that disrupts the efficient flow of traffic and may over time also cause pilots to distrust the automation. Monitoring and safety assessments led to a series of changes resulting in the latest international version of TCAS – referred to as Version 7.1, or the Airborne Collision Avoidance System (ACAS). Starting in January 2003, the International Civil Aviation Organization mandated the use of ACAS worldwide for all turbine-powered aircraft with passenger capacity of more than 30 or with maximum take-off weight exceeding 15,000 kg. In January 2005, that mandate was extended to cover aircraft with more than 19 passenger seats or maximum take-off weight of more than 5700 kg [2].

How TCAS works

TCAS processes are organized into several elements, as shown in Figure 1. First, surveillance sensors collect state information about the intruder aircraft (e.g., its relative position and velocity) and pass the information to a set of algorithms to determine whether a collision threat exists. If a threat is identified, a second set of threat-resolution algorithms determines an appropriate response. If the intruder aircraft also has TCAS, the response is coordinated through a data link to ensure that each aircraft maneuvers in a compatible direction. Collision avoidance maneuvers generated and displayed by TCAS are treated as advisories to flight crews, who then take manual control of the aircraft and maneuver accordingly. Pilots are trained to follow TCAS advisories unless doing so would jeopardize safety. The following sections provide more detail on the methods used to perform surveillance, threat detection, and threat resolution.

Surveillance

Surveillance of the air traffic environment is based on air-to-air interrogations broadcast once per second from antennae on the TCAS aircraft using the same frequency (1030 MHz) and waveform as ground-based air traffic control sensors [3]. Transponders on nearby intruder aircraft receive these interrogations and send replies at 1090 MHz. Two types of transponders are currently in use: Mode S transponders, which have a unique 24-bit identifier, or Mode S address, and older Air Traffic Control Radar Beacon System (ATCRBS) transponders, which do not have unique addressing capability.

Fig. 1. Elements of TCAS processes

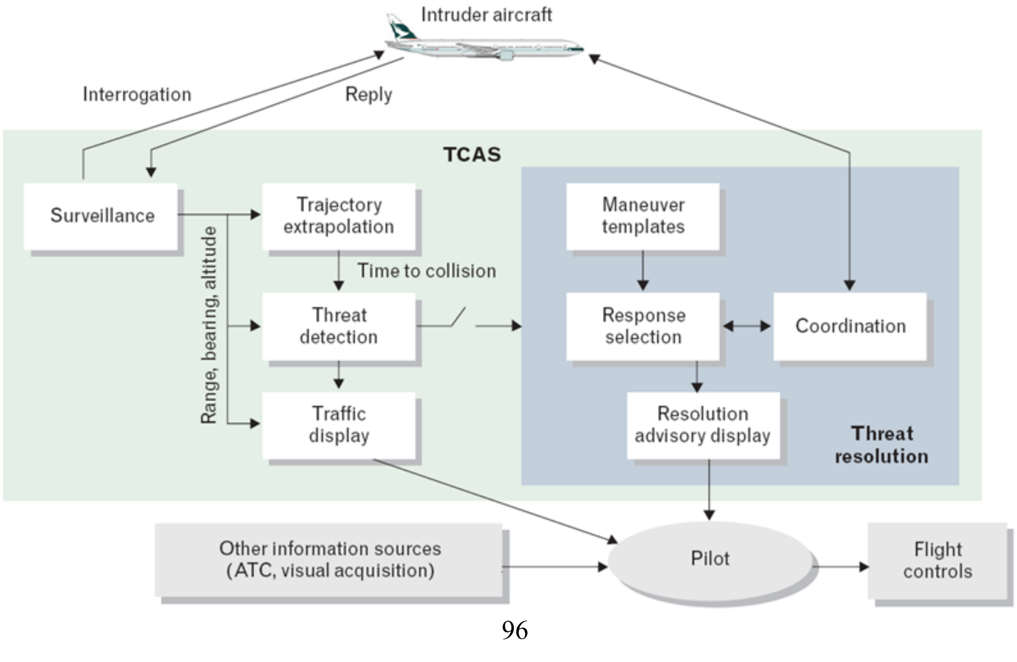

To track ATCRBS intruders, TCAS transmits "ATCRBS-only all-call" interrogations once per second. All ATCRBS aircraft in a region around the TCAS aircraft reply. In contrast, Mode S – equipped intruders are tracked with a selective interrogation once per second directed at that specific intruder; only that one aircraft replies. Selective interrogation reduces the likelihood of garbled or overlapping replies, and also reduces frequency congestion at 1030/1090 MHz. Replies from most ATCRBS and all Mode S transponders contain the intruder's current altitude above sea level. TCAS computes slant range on the basis of the round-trip time of the signal and estimates the bearing to the intruder by using a four-element directional antenna. Alpha-beta and non-linear filters are used to update range, bearing, and altitude estimates as well as to estimate range rate and relative-altitude rate. Mode S transponders also provide additional data-link capabilities. All aircraft with TCAS are equipped with Mode S transponders so that this data link can coordinate collision avoidance maneuvers.

One of the most difficult challenges in the development of TCAS is balancing the surveillance requirements of TCAS and air traffic control ground sensors – in particular, managing their shared use of the 1030/1090 MHz frequencies. As the density of TCAS equipped aircraft grows, transponders in an airspace are interrogated by more and more TCAS units. As a result, transponders now devote more of their time to responding to TCAS and less of their time responding to ground interrogations. Because of concerns about frequency congestion, TCAS uses interference-limiting algorithms to reduce competition between TCAS and ground sensors. Each second, TCAS determines the number and distribution of other TCAS units in its vicinity. With that information, TCAS can reduce its maximum transmit power (i.e., reduce its surveillance range) – limiting the impact on the victim transponders and, in turn, on the ground sensors.

National and international requirements in this area are quite strict. Interference limiting is intended to ensure that for any given transponder, no more than 2% of its available time is consumed in communications with all nearby TCAS units. Because TCAS requires a minimum surveillance range to provide adequate collision avoidance protection, however, a limit is imposed on how much the TCAS transmit power can be reduced. As a result, it is possible for a transponder to exceed the 2% utilization figure in high-density airspace. Transponder utilization due to TCAS has been the focus of worldwide monitoring, and monitoring results continue to motivate the development of innovative TCAS surveillance techniques. Many such techniques were developed for Version 7, including using Mode S interrogation schemes that are different for distant, non-threatening intruders than for potential threats, and transmitting sequences of variable-power ATCRBS interrogations to reduce garble, or overlap, among concentrations of ATCRBS intruders. In addition, standards are nearing completion for TCAS Hybrid Surveillance. This is a new technique that allows TCAS to make use of passive (Automatic Dependent Surveillance – Broadcast, or ADS-B) transmissions, thereby reducing TCAS interrogation rates. Two other issues affect the ability of TCAS to track intruders. First, some older transponders do not report altitude information when interrogated. TCAS can not generate collision avoidance commands against these threats. (Large aircraft, aircraft flying in the vicinity of large airports, and aircraft flying above 10,000 ft are required to be equipped with altitude-reporting transponders.) Second, aircraft without a functioning transponder can not be detected or tracked by TCAS at all. Some small aircraft, such as gliders or ultralights, may not carry any electronic equipment or transponders. Pilots therefore must take the responsibility to see and avoid such traffic.

Threat detection and display

TCAS's complex threat-detection algorithms begin by classifying intruders into one of four discrete levels [6]. To project an aircraft's position into the future, the system performs a simple linear extrapolation based on the aircraft's estimated current velocity. The algorithm then uses several key metrics to decide whether an intruder is a threat, including the estimated vertical and slant range separations between aircraft. Another parameter, called *tau*, represents the time until the closest point of approach between aircraft. A display in the cockpit depicts nearby aircraft, indicating their range, bearing, and relative altitude; an arrow indicates whether the intruder is climbing or descending. Such traffic display information aids the pilot when attempting to visually acquire traffic out the windscreen. Distant, non-threatening aircraft appear as hollow diamond icons. Should the intruder close within certain lateral and vertical limits, the icon changes to a solid diamond, alerting the flight crew that traffic is proximate, but is not yet a threat. If a collision is predicted to occur within the next 20 to 48 seconds (depending on altitude), TCAS issues a Traffic Advisory (TA) in the cockpit [4]. This advisory comes in the form of a spoken message, "traffic, traffic." The traffic icon also changes into a solid yellow circle. The TA alerts the pilot to the potential threat so that the pilot can search visually for the intruder and communicate with ATC about the situation. A TA also serves as a preparatory cue in case maneuvering becomes required.

If the situation worsens, a Resolution Advisory (RA) warning is issued 15 to 35 seconds before collision (again depending on altitude). The RA includes an aural command such as "climb, climb" and a graphical display of the target vertical rate for the aircraft. A pilot receiving an RA should disengage the autopilot and manually control the aircraft to achieve the recommended vertical rate. Figure 2 shows both the TA and RA displays.

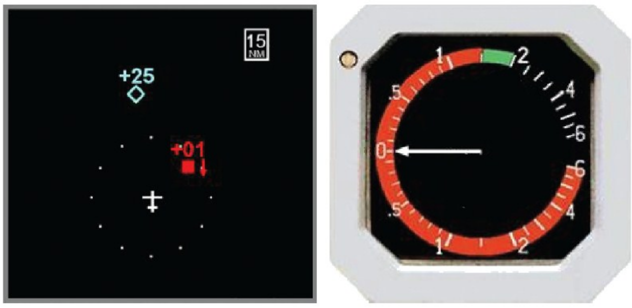

Figure 2. TCAS TA display (left) and RA display (right)

Threat resolution

Once the criteria for issuing an RA have been met, TCAS's threat-resolution algorithms determine what maneuver is appropriate to avoid a collision. First, the algorithm decides the vertical sense of the maneuver – that is, whether the aircraft needs to climb or to descend. Second, the system figures the strength of the RA – that is, how rapidly the aircraft needs to change its altitude. TCAS works only in the vertical direction; it does not select turning maneuvers, because bearing accuracy is generally not sufficient to determine whether a turn to the left or right is appropriate. Figure 3 shows a simplification of the sense-selection process. In general, two maneuver templates are examined: one based on a climb, and one based on a descent. Each template assumes a 5 sec delay before a response begins, followed by a 0.25 g vertical acceleration until reaching a target vertical rate of 1500 ft/min. In the meantime, the intruder aircraft is assumed to continue in a straight line at its current vertical rate. The TCAS algorithm selects the maneuver sense providing the largest separation at the predicted closest point of approach. In the situation shown in Figure 3, TCAS would on the basis of these criteria advise the aircraft to descend.

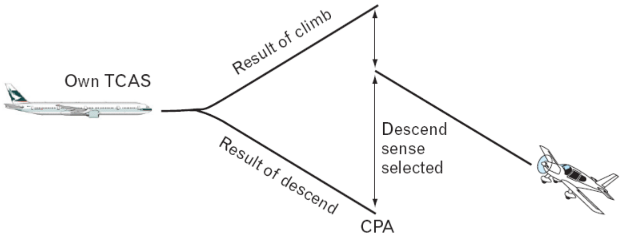

Fig. 3. TCAS algorithm selects the maneuver that provides the largest separation at the predicted Closest Point of Approach (CPA). In the scenario shown here, the correct maneuver would be to descend.

If the intruder is also TCAS equipped, the sense of the RA is coordinated through the Mode S data link to ensure that both aircraft do not select the same vertical sense. Should both aircraft simultaneously select the same sense – say, both select a climb RA – the aircraft with the lower numerical-valued Mode S address has priority and will continue to display its climb RA. The aircraft with the higher Mode S address will then reverse its sense and display a descend RA. Once the sense has been selected, the strength of the RA maneuver is determined by using additional maneuver templates (Figure 4). Each template again assumes a 5 sec delay, followed by a 0.25 g acceleration to reach the target vertical rate. TCAS selects the template that requires the smallest vertical-rate change that achieves at least a certain minimum separation. In the example shown in Figure 5, the TCAS aircraft is currently descending at a rate of 1000 ft/min when an RA is issued. Five maneuver templates are examined, with each template corresponding to a different target vertical rate. The minimum-strength maneuver that would provide the required vertical separation of at least 400 ft would be to reduce the descent rate to 500 ft/min; the pilot would receive an aural message stating that instruction. Descent rates exceeding 500 ft/min would appear in red on the RA display. Note that in Figure 4 if the intruder were 100 ft higher, then the selected RA would instead be "don't descend." If the intruder were another 100 ft higher still, the selected RA would be "climb."

Due to TCAS's 1 Hz update rate and filtering lags, its estimates may lag the actual situation during periods of sudden acceleration. This lag may in turn lead to an inappropriate RA sense or strength. To help alleviate this problem, TCAS refrains from issuing an RA, if there are large uncertainties about the intruder's track.

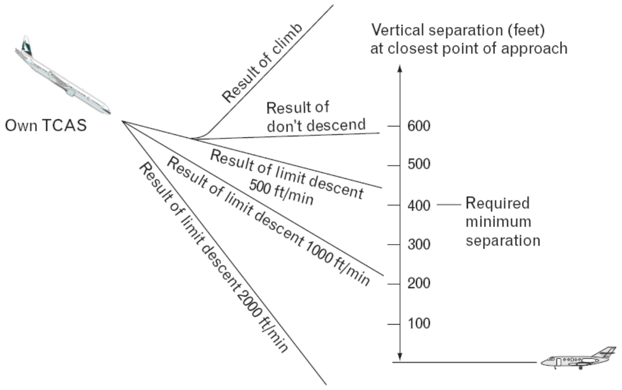

Fig. 4. Once TCAS determines whether to advise an aircraft to climb or to descend, it calculates the speed at which the aircraft must maneuver to avoid collision. TCAS selects the template that requires the smallest change in vertical rate that achieves the required separation.

TCAS also includes algorithms that monitor the evolution of the encounter and, if necessary, issue a modified RA. The strength of an RA can be increased – for example, changing from "don't descend" to "climb" (target rate of 1500 ft/min) to "increase climb" (target rate of 2500 ft/min). Under certain conditions, if it becomes clear that the situation is continuing to degrade, TCAS can even reverse the sense of the RA, from climb to descend, or vice versa. Coordination of this reversal with a TCAS-equipped intruder aircraft will also be performed through the Mode S data link. Sense reversal is especially challenging because only a few seconds may remain before collision. Any latencies involved in pilot and aircraft response could result in an out-of-phase response that further reduces separation.

Lessons from a disaster

On the night of 1 July 2002, a Boeing B-757 operated by the cargo carrier DHL collided with a Russian Tu-154 passenger jet at 34,940 ft over the small town of Überlingen, Germany. The accident destroyed both aircraft and killed all 71 crew members and passengers aboard the two aircraft. What was especially troubling about this accident is that both aircraft were equipped with TCAS. As with most aviation accidents, a string of events occurred leading up to the collision. First, the nominal separation standards between aircraft were lost through a combination of problems and errors at the air traffic control facility monitoring the aircraft. As a result, the two aircraft were on a collision course much closer together than is normal while cruising at 36,000 ft.

Figure 5 schematically summarizes the event. Forty three seconds before the collision, ATC instructed the Russian aircraft to descend because of the traffic conflict. Before the controller finished his verbal instruction, however, TCAS on the Russian aircraft issued an RA advising the pilot to climb. A coordinated descend RA was issued on the DHL aircraft at the same time. The DHL pilots followed their RA and began to descend. The Russian flight crew followed the ATC instruction and also descended. Shortly thereafter, the RAs on each aircraft were strengthened to “increase climb” on the Russian aircraft and “increase descent” on the DHL aircraft. About 35 seconds after the TCAS RAs were issued, the aircraft collided.

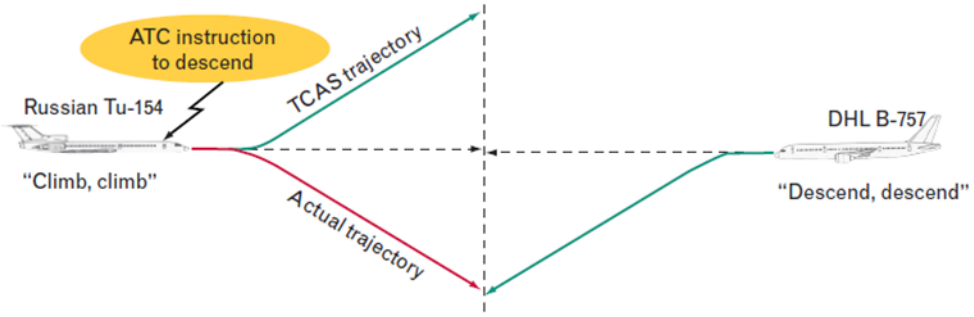

Fig. 5. The Überlingen mid-air collision occurred after the Russian pilot decided to heed the air traffic control instruction to descend rather than the TCAS advisory to climb

One of the immediate causes for the accident, as described in the German accident report, was the fact that the Russian flight crew chose to follow the ATC clearance to descend rather than follow the TCAS RA to climb [5]. The Russians’ choice to maneuver opposite to the RA defeated the coordination logic in TCAS. An advisory system like TCAS can not prevent an accident, if the pilots don’t follow the system’s advice. The DHL crew, however, did follow the TCAS RA and yet they still collided. The question thus arises: why didn’t TCAS reverse the sense of the RAs when the situation continued to degrade? Had it done so, the Russian aircraft would have received a descend RA, which presumably it would have followed, since the crew had already decided to descend in response to the ATC clearance. The DHL aircraft would have received a climb RA, which it likewise would have presumably followed, since its crew had obeyed the original RA. This is not to say that a reversal is always a good idea, however. In many encounters, a reversal would reduce separation and increase the risk of a collision. Because of sensor limitations and filtering lags, it turns out to be quite difficult to trigger reversals when they are needed while avoiding them when they are not needed.

A closer examination of the reversal logic revealed several areas in which earlier design assumptions proved inadequate in situations when one aircraft maneuvers opposite to its RA. In order for a RA reversal to be issued, the Version 7 threat logic requires four basic conditions to be satisfied; these conditions are illustrated in Figure 6. First, a reversal will be triggered only by the aircraft with priority – that is, the aircraft with the lower Mode S address. If the aircraft has a higher Mode S address than the intruder, the RA sense will be reversed only when directed to do so by the priority aircraft through the data link. Second, the maneuver templates projecting the situation into the future need to predict that insufficient separation between aircraft will occur unless a sense reversal is issued. Third, a maneuver template projecting the response to a reversed-sense RA needs to predict adequate separation between aircraft. Fourth, the two aircraft in danger of colliding must be separated by at least 100 ft vertically. (This last condition is intended to prevent reversals from occurring just as aircraft cross in altitude.)

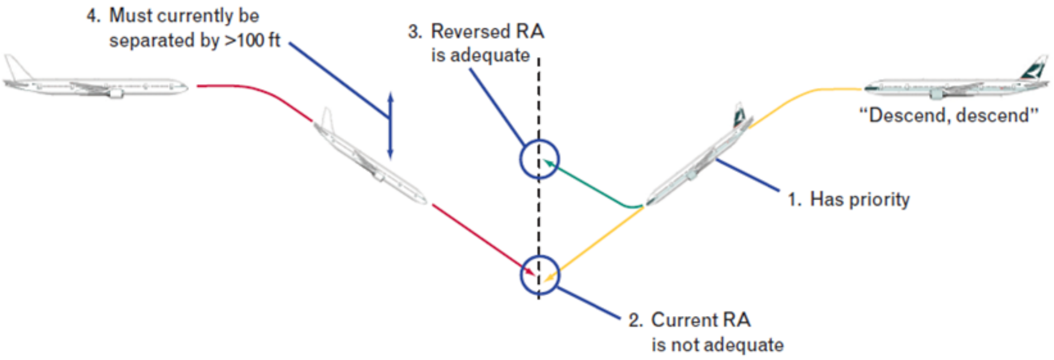

Fig. 6. In order for TCAS to reverse its maneuver instruction – e.g., from "descend" to "climb" – four conditions must hold

A closer look at the Überlingen accident, as shown in Figure 7, reveals why TCAS did not issue an RA reversal. Responsibility for triggering the reversal rested with the Russian aircraft, which had a lower Mode S address. The Russian aircraft was operating under an active climb RA. The climb-RA maneuver template predicted adequate separation between aircraft, at least until the final few seconds. Therefore, TCAS did not issue an RA reversal. Since the Russian aircraft was not actually following the climb maneuver, of course, the template's predictions were invalid. What is startling, however, is that even if the DHL aircraft had the lower Mode S address (and therefore priority), the aircraft still probably would have collided. In the hypothetical case in which the DHL aircraft had priority, three of the four conditions required to trigger a reversal, as shown in Figure 6, would have held: the DHL aircraft would have had priority; the DHL aircraft's descend RA would have shown that a collision was still predicted; and the projection of a reversal-climb RA would have predicted adequate separation. However, both aircraft remained within 100 ft vertically of each other throughout the encounter, and so this fourth criterion for permitting a reversal still would not have been met.

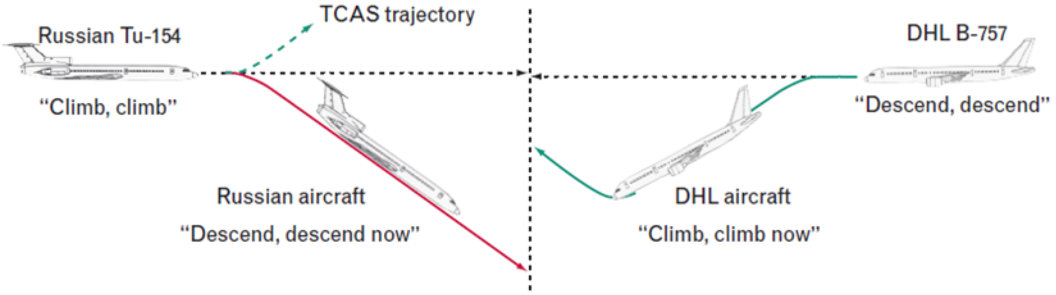

Fig. 7. The Überlingen accident might have been averted, if TCAS had issued an RA reversal as shown. Responsibility for triggering the reversal rested with the Russian aircraft, which had priority and which was operating under a "climb" RA. But until the final few seconds, the climb RA maneuver template predicted adequate separation between aircraft. Therefore, TCAS did not issue an RA reversal. Since the Russian aircraft was not actually following the climb maneuver, but rather the air traffic control instruction to descend, the template's predictions were tragically invalid.

To reduce the risk of this type of collision, researchers funded by the European Organization for the Safety of Air Navigation (Eurocontrol), have proposed a change to the TCAS threat logic. Eurocontrol's proposal aims to improve reversal performance in encounters in which both aircraft become involved in a so-called vertical chase, as occurred at Überlingen. The proposal includes two major components. First, when using maneuver templates, TCAS would no longer assume that the TCAS aircraft would follow its RA. Instead, TCAS would check the recent vertical motion of the aircraft. If this motion is not compatible with the RA that had been issued, then TCAS would revert to models using the aircraft's current vertical rate instead of its predicted motion in response to the RA. Second, the proposal would eliminate the 100 ft separation requirement, allowing TCAS to reverse sense in vertical-chase situations. The combination of these changes would have produced RA reversals in the Überlingen accident – no matter which aircraft had priority.

Conclusion

TCAS represents a clear success story in aviation safety. Its successful design was achieved through detailed consideration of sensor characteristics and the coupled dynamic interactions among pilots, air traffic controllers, and aircraft. The result is a fine balance that provides sufficient time to take action and that minimizes alert rates. As the Überlingen accident shows, however, safety can not be taken for granted, and areas of improvement will always exist in systems that rely on integrating humans and automation for information processing and decision making.

The real challenge lies in integrating new collision avoidance technologies with the existing systems and procedures. The Überlingen accident demonstrated the catastrophic outcome that can result from dissonance between two different decision makers in a time-critical situation: namely, an air traffic controller's decision to request a descent and TCAS's Resolution Advisory to climb. While this specific problem is being solved by improving pilot training to comply with RAs and refining the TCAS algorithms, related problems are likely to surface as unmanned aircraft and enhanced collision avoidance technologies mix. Ensuring compatible operation also extends well beyond TCAS or aviation to many integrated sensing and decision support system applications.

References

[1] *ACAS Guide: Airborne Collision Avoidance Systems (incorporating TCAS II version 7.0 & 7.1 and introduction to ACAS X)*, Eurocontrol, April 2015

[2] *Airborne Collision Avoidance System (ACAS) Manual, Doc 9863,* International Civil Aviation Organization (ICAO), First Edition, 2006

[3] *Annex 10 to the Convention on International Civil Aviation – Aeronautical Telecommunications, Volume IV (Surveillance Radar and Collision Avoidance Systems)*, Second Edition, July 1998

[4] *Introduction to TCAS II Version 7.1,* U.S. Department of Transportation, Federal Aviation Administration, 28 Feb. 2011

[5] *Investigation Report AX001-1-2/02,* Bundesstelle für Flugunfalluntersuchung (BFU, German Federal Bureau of Aircraft Accident Investigation), Braunschweig, Germany, May 2004

[6] *Minimum Operational Performance Standards for Traffic Alert and Collision Avoidance System II (TCAS II) Airborne Equipment*, RTCA/DO-185A, Washington, D.C., 16 Dec. 1997