What Should Aviation Safety Incidents Teach Us?

William S. Greenwell John C. Knight

Department of Computer Science
University of Virginia
151 Engineer's Way, P.O. Box 400740
Charlottesville, VA 22904-4740, USA
{greenwell,knight}@cs.virginia.edu

Abstract. Accidents and incidents involving safety-critical software systems often provide lessons to the systems' users and designers, to industry, and to the software engineering community at large. Proper identification and documentation of these lessons is critical in order to prevent the recurrence of an untoward event. In this paper we examine two commercial aviation incidents involving failures of safety-critical software systems. Based on our analysis of the incidents and the official investigations that followed, we conclude that the aviation community is missing important lessons regarding safety-critical software systems, especially concerning the broad role these systems play in preserving the safety of commercial air travel. This is primarily because incidents involving such systems are not being investigated and documented with sufficient rigor to identify these lessons and disseminate them throughout the aviation community effectively.

1 Introduction

Safety-critical software systems are becoming increasingly ubiquitous in our society. In the realm of commercial aviation, these systems are used both on board aircraft and at air traffic control facilities to assist pilots in operating their aircraft safely and air traffic controllers in managing the national airspace system (NAS) in a safe and efficient manner. Either in their capacity to control potentially hazardous operations or to advise human controllers via warnings and guidance when danger is present, we rely on these systems to function in a dependable fashion and not threaten our safety. If such a system falls short of its dependability requirements, the consequences could be catastrophic, and lives or property could be put at risk. Therefore, incidents involving safety-critical systems are serious occurrences. Whether or not an incident results in a catastrophe, it indicates a weakness in the systems involved and underscores the need for improving the affected systems to prevent future occurrences that could have more severe consequences. How we investigate incidents in which a safety-critical system failed to function as intended might determine whether lives or property are affected in the future by the same system behavior.

In this paper we describe two commercial aviation incidents involving safety-critical software systems. The first incident involved the failure of a ground-based warning system that contributed to an accident with extensive casualties. The second concerned a failure of an onboard collision avoidance system that caused two aircraft to nearly collide, jeopardizing the lives of over 400 passengers and crew members. In

our opinion, the official investigations of these incidents did not examine the software systems involved with sufficient rigor, and consequently crucial lessons in software engineering were not documented or acted upon. We extract new lessons for the aviation community relating to the design of safety-critical software systems and the manner in which incident investigations are conducted when software systems are involved.

The remainder of this paper is organized as follows. Sections 2 and 3 present each of the incidents we selected as case studies for this research and review each incident separately. Section 4 contains a common lesson to the aviation community on the design and maintenance of safety-critical systems along with our recommendations and conclusions.

2 Korean Air Flight 801

On August 6, 1997 at about 1:42am Guam local time, Korean Air flight 801, a Boeing 747-300, crashed into Nimitz Hill, Guam while attempting a nonprecision approach to runway 6L at A.B. Won Guam International Airport. Of the 254 persons on board, 237 of which were passengers, only 23 passengers and 3 flight attendants survived. The National Transportation Safety Board (NTSB) investigated the accident and classified the crash as a *controlled-flight-into-terrain*, or CFIT, accident. During its investigation, the NTSB found that a ground-based minimum safe altitude warning system (MSAW), designed to alert air traffic controllers of aircraft flying too low, had been inhibited. In its final report [1], the NTSB concluded that the crash was largely due to pilot error, but also noted:

> Contributing to the accident was the Federal Aviation Administration's (FAA) intentional inhibition of the minimum safe altitude warning system (MSAW) at Guam and the agency's failure to adequately manage the system.

Despite its finding that the inhibition of the MSAW system at Guam was a contributory factor, the NTSB did not issue any safety recommendations to the FAA pertaining to the MSAW system in response to this accident.

2.1 The Incident

Korean Air flight 801 crashed during its final approach to runway 6L at Guam International Airport while operating under instrument flight rules (IFR). At the time of the accident, the FAA had issued a Notice To Airmen (NOTAM) for the Guam airport stating that the runway 6L glideslope was out of service, meaning that pilots were not to rely on the glideslope signal when landing at Guam. The flight crew of flight 801 received this notice both prior to departure and again from air traffic control during their approach to Guam. When the glideslope is unavailable, it is still possible to perform a nonprecision or localizer-only ILS approach. The nonprecision approach procedure is published alongside the precision approach as a sequence of step-down altitudes. In lieu of a glideslope, pilots make a series of intermediate descents using these step-down altitude fixes.

Postaccident analysis of radar data indicates that flight 801 began a premature descent on its nonprecision approach and violated the 2,000 step-down clearance. The

aircraft proceeded on a steady descent, violating the 1,440 step-down clearance before impacting terrain approximately 3.3 nm short of the runway threshold. In its report, the NTSB concluded, "the captain lost awareness of flight 801's position on the [ILS] localizer-only approach to runway 6L at Guam International Airport and improperly descended below the intermediate approach altitudes...which was causal to the accident." Thus, the NTSB classified the accident as a controlled flight into terrain (CFIT).

The FAA's Commercial Aviation Safety Team (CAST) cited controlled flight into terrain as "the leading cause of fatal commercial air accidents worldwide" [3]. It defines a CFIT accident as one in which "a fully qualified and certified crew flies a properly working airplane into the ground, water, or obstacles with no apparent awareness by the pilots." Under its own initiatives and in response to Safety Recommendations from the NTSB, the FAA has adopted numerous systems and procedures designed to reduce the frequency of CFIT-induced accidents. In the cockpit, the Instrument Landing System (ILS), comprised of the localizer and glideslope, marker beacons, and special runway lighting, provides precision guidance to the flight crew as the aircraft makes its final approach. Approach plates specify procedures for becoming established on the ILS approach as well as backup approach procedures in case the ILS approach is unavailable. In addition, an onboard Ground Proximity Warning System (GPWS) gives aural altitude callouts as the aircraft descends and features a special callout when the aircraft reaches its decision height or minimum descent altitude (MDA). Lastly, the other members of the flight crew, typically the copilot and possibly the flight engineer, monitor the pilot's approach and may challenge it if they sense trouble. On the ground, the MSAW system alerts air traffic controllers to low-flying aircraft so that they can contact the flight crews and advise them accordingly.

Under normal circumstances, each of these measures—the ILS components, the flight crew following approach-plate procedures with onboard instruments, and the MSAW system—serves as a *barrier* against CFIT. While individual systems might fail occasionally, an accident will be prevented if just one of the systems above functions as intended. For a CFIT-induced accident to occur, *all* of these barriers must fail to prevent the accident, and typically the probability of such a catastrophic failure is extremely small provided the systems fail randomly and independently of each other.

2.2 MSAW System Overview

The MSAW system is a ground-based system that alerts controllers visually and aurally when an IFR-tracked flight descends below, or is predicted to descend below, a predetermined minimum safe altitude (MSA). The system itself is entirely software, relying on existing radar hardware to assist controllers in detecting low-flying aircraft.

Each MSAW installation operates with a terrain database and configuration information that are tailored to the airport at which the installation is running. The system identifies low-flying aircraft by employing two monitoring techniques. General monitoring tracks all aircraft operating within the MSAW service area. For each aircraft, the system reads the maximum terrain elevation for the region in which the aircraft is operating from the terrain database and applies a 500 foot margin to determine the MSA for that region (although the margin value can be adjusted). If the aircraft has violated its MSA, the system raises an alert to air traffic controllers. Approach path

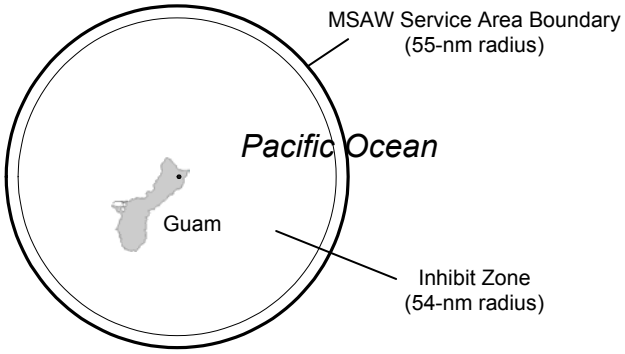

Fig. 1. The Guam MSAW inhibit zone (not drawn to scale)

monitoring, the second technique, tracks aircraft operating within specially designated rectangular regions called capture boxes where aircraft typically perform final approach maneuvers. Inside each capture box, the MSAW system simulates a glideslope descent path to determine whether an aircraft on final approach has descended, or is projected to descend, below the desired path.

After its introduction, some MSAW-equipped sites were plagued by frequent nuisance warnings generated by the system. These warnings were typically triggered by aircraft that had just taken off or were about to land. In order to reduce the frequency of nuisance warnings, site adaptation managers could request that inhibit zones be added to the configuration information for their airports.

According to the NTSB report, the Guam MSAW system was installed in 1990 and originally configured to monitor an area within a 55-nautical mile (nm) radius of the Guam radar. In March 1993, Guam air traffic managers, in conjunction with the FAA's Western-Pacific Region office and the FAA Technical Center, prepared new site adaptation parameters for the Guam MSAW system that included a 54-nm inhibit zone centered at the Guam radar site as illustrated in Figure 1. According to NTSB investigators, this change was "neither a fluke nor a malfunction but rather was an intentional adaptation change for the purpose of eliminating numerous nuisance low altitude alerts," and was put in place "for temporary use until a better solution to the problem of nuisance alarms could be found" [4]. According to testimony from the acting manager of the National Field Support Division (NFSD) of the FAA Technical Center at the time of the accident, this change effectively reduced MSAW processing to a 1-nm ring located between 54 and 55 nm from the radar facility as shown in Figure 1. No general or approach path warnings would be generated by the system for aircraft operating within the inhibit zone [2].

The new system became operational in February 1995. In July of the same year an FAA inspector conducted a biannual evaluation of the Guam facility and noted the inhibition of the MSAW system as an "informational" item, but did not recommend any follow-up action be taken. In April 1996 the FAA installed new MSAW software at Guam to update the terrain database; however this software also contained the 54-nm inhibit zone. This software remained in operation through the time of the accident. The FAA conducted another facility evaluation of Guam in May 1997, but this inspec-

tion failed to note the MSAW inhibition entirely.

2.3 Postincident Actions

After the accident, the FAA and NTSB investigators conducted a simulation of Korean Air flight 801's final approach with the MSAW inhibit zone removed. The simulation indicated that, without the inhibit zone, the MSAW system would have generated visual and aural low altitude warnings for KA 801 64 seconds prior to impact. The NTSB concluded in its report that this would have been sufficient time for air traffic controllers to notify KA 801 and for the flight crew to take remedial action.

On August 15, 1997, the FAA announced that it had begun a review of MSAW systems nationwide as a precautionary measure. Of the 192 in-service systems, the FAA found two that were configured improperly, and reported that these systems were corrected and recertified. In addition, FAA flight inspections of 23 ATC facilities uncovered a previously unknown inhibit zone at Florence, SC. In response to these findings, the FAA developed policy to "require that MSAW be flight checked and ground certified as part of the commissioning process for a new radar and periodically thereafter" [1]. The FAA also conducted a fact-finding review of 10 ATC towers to assess controllers' knowledge of the MSAW system. The review found that most controllers possessed only a cursory knowledge of the system and gave inconsistent answers when asked about who had the authority to adapt MSAW parameters and how daily MSAW testing should be conducted if the system had been inhibited.

The review recommended that, among other things, (a) uniform site adaptation parameters should be established for MSAW operation, (b) periodic evaluations of MSAW systems should be conducted "to ensure system integrity and reliability," and (c) configuration management of MSAW software should be appropriately documented and centrally controlled. In an October 1997 briefing to the NTSB, the FAA also presented new guidelines for certifying and maintaining MSAW systems to establish "strict management oversight and control" over MSAW operations [1].

2.4 Analysis

The MSAW system was developed to address scenarios in which the onboard barriers designed to prevent CFIT accidents fail. This is precisely what happened on August 6, 1997 over Guam. The glideslope for runway 6L was out of service, and the captain lost awareness of the aircraft's position on final approach. Although the onboard Ground Proximity Warning System (GPWS) gave aural altitude callouts to the flight crew as the aircraft descended and an additional callout when the aircraft reached its MDA, Cockpit voice recorder transcripts indicate these callouts were largely ignored by the flight crew, possibly because traditional GPWS systems were known to generate nuisance messages over Guam. Poor illumination surrounding Nimitz Hill made it difficult for the captain to verify his approach visually. Lastly, the copilot and flight engineer failed to challenge the captain's approach in time to save the aircraft.

In its final report, the NTSB concluded that "the FAA's quality assurance for the minimum safe altitude warning system was inadequate, and the agency's intentional inhibition of that system contributed to the flight 801 accident" [1]. The NTSB did not, however, identify the underlying problems in the FAA's quality assurance process or recommend changes to the FAA's maintenance programs for MSAW or its other soft-

ware systems.

Clearly the FAA's quality assurance of MSAW was inadequate. The underlying problem with the manner in which the FAA maintained the MSAW installations at its 193 ARTS IIA and ARTS III facilities is that it allowed changes to be made to the system without examining the effect those changes would have on the safety case the system was designed to address. They had taken a trial-and-error approach to adapting MSAW site parameters, allowing sites to make configuration changes at their discretion to reduce or eliminate nuisance warnings with little oversight from the AOS or the ATO. Moreover, the FAA provided individual sites with no instructions for making configuration changes or guidance for reducing nuisance warnings while minimizing the extent of inhibit zones. These are merely symptoms of a deeper problem, however, and as Leveson notes, "If we only patch the symptoms and ignore the deeper underlying cause of one accident, we are unlikely to have much effect on future accidents" [5].

The MSAW system stands as the only ground-based barrier against CFIT-induced accidents, aside from the vigilance of air traffic controllers. While the FAA does not, in general, regard ground-based software systems as safety-critical, in his testimony at the NTSB hearing, the acting manager of the NFSD classified the MSAW system as a "safety-critical item" [2], just as most people would.

2.5 Related Incidents

Dulles International Airport, 1994. On June 18, 1994, a Transportes Aereos Ejecutivos, S.A. Learjet crashed on final approach to runway 1R at Dulles International Airport approximately 0.8 nm short of the runway. During its investigation of the accident, the NTSB found two discrepancies in the site adaptation variables used by the Dulles MSAW installation. These discrepancies caused the system to model the location of the threshold for runway 1R incorrectly and to apply the wrong MDA for aircraft subject to approach path monitoring. While the NTSB did not find these discrepancies to be causal to the accident, on November 21, 1994 the NTSB issued the following Safety Recommendation to the FAA:

> Conduct a complete national review of all environments using MSAW systems. This review should address all user-defined site variables for the MSAW programs that control general terrain warnings, as well as runway capture boxes, to ensure compliance with prescribed procedures [6].

The FAA responded that it would undertake such a review, and on January 26, 1996 reported that the review had been completed. The 1995 facility inspection of Guam in which the MSAW inhibition was cited as an "informational item" was undertaken during the review period, even though no corrections were made.

Houston Intercontinental Airport, 1998. Four years later, on January 13, 1998, a Learjet crashed 2.3 nm short of runway 26 while on final approach to Houston Intercontinental Airport. Investigators determined that the MDA specified in the site adaptation parameters for the Houston MSAW installation was incorrect. The MSAW system was configured to use an MDA of 100 feet above ground when the actual MDA was 402 feet above ground.

The configuration error in the Houston MSAW installation was the same error that

had been made at Dulles four years earlier. This type of error should have been detected and fixed during the FAA's national MSAW review campaign. Moreover, the Houston accident occurred after the FAA had implemented the recertification programs and uniform site adaptation standards it had proposed in response to the Korean Air flight 801 accident.

2.6 Lessons Learned

Accidents occur because of complex sequences of events and intricate combinations of circumstances. It is clear that many things could have prevented this accident. The NTSB report blames three factors—the flight crew, the lack of operation of the glideslope, and the FAA's inhibition of the MSAW system's service area. Presumably, changes were made based on the first two, and we have documented the changes that were made as a result of the third. After looking at all the evidence about this accident that is available to us, however, we conclude that the lessons learned from this accident were far short of what they should have been. Two additional prominent problems should have been identified and additional significant corrective actions taken.

Lesson 1—Configuration Management. Korean Air flight 801 crashed into Nimitz Hill, Guam in part because of the manner in which the FAA made software changes to the MSAW system. By allowing each of the 193 MSAW-equipped air traffic control facilities to modify their MSAW installations at their discretion without guidance or review, the FAA allowed the system to be modified without regard to how the modifications might affect the system's ability to detect low-flying aircraft.

The MSAW system at Guam was a barrier designed to help prevent CFIT accidents. As such, it was a component of an overall system that included all of the barriers designed to prevent the hazard of flying below a safe altitude. Prevention of the hazard could have been achieved by any one of the barriers provided that particular one was perfect in its operation. None of them were. The goal of preventing the hazard was to be achieved by accepting that no barrier was perfect and providing several. Thus, the MSAW system's functionality was an integral part of the analysis of the overall system's safety. This does not mean that the system itself has to be ultra-dependable. It means that the system's dependability when coupled with that of the other barriers reduces the probability of an accident to acceptable levels.

The crucial lesson here is that all aspects of a system's software configuration might be essential parts of maintaining system safety, and that the initially deployed configuration and any changes to a software system must be undertaken only in concert with a comprehensive safety analysis. The importance of this lesson is underscored by the fact that there were two similar incidents at Dulles and Houston airports.

Lesson 2—Human Error. The second lesson that should have been learned from this accident concerns human error. Human error in the maintenance of software in safety-related systems is likely just as it is in the operation of those systems, yet software in a safety-critical system is an integral component that cannot be changed without detailed analysis of the impact of the change.

Complementing the first lesson noted above, the research community needs to examine the complex circumstances that are present in widely deployed safety-related

software systems and develop techniques to verify properties that are crucial to safety. For example, requiring that the software (including all data and configuration files) for some particular system not be changed in the field unless the change is accompanied by suitable verification activities, re-establishment of safety properties, and compatibility checks with other software components is essential.

3 British Airways Flight 027

On June 28, 1999, British Airways flight 027, a Boeing 747 carrying 419 passengers and crew members en route to Hong Kong, China, and another Boeing 747 operated by Korean Air Cargo nearly collided over a remote region of Chinese airspace. At their closest point of approach, the two aircraft passed within 600 feet of each other, and the British Airways copilot later recounted that his windshield was consumed by the fuselage of the other jet. No injuries resulted from the incident and both aircraft arrived at their destinations. If the two aircraft had collided, however, it is likely that none of the persons aboard either aircraft would have survived [7].

3.1 The Incident

Prior to the incident, the two aircraft were cruising in opposite directions along the same airway with 2,000 feet of vertical separation. The British Airways passenger flight was cruising at 33,500 feet and the Korean Air Cargo jet at 31,500 feet. The Korean Air jet was flying in a cloud, preventing the pilots from visually identifying each other's aircraft. Both aircraft were equipped with a collision avoidance system known as the Traffic Alert and Collision Avoidance System, or TCAS.

The TCAS unit installed on the Korean Air jet indicated traffic 400 feet below and approaching head on and shortly thereafter instructed the pilot to climb to avoid the oncoming traffic. In reality, there were no other aircraft in the vicinity of the Korean Air jet except for the British Airways flight 2,000 feet *above*, and the TCAS unit's indication and climb instruction were erroneous. The pilot had no way of knowing this, however, as he was operating in a region of airspace without air traffic control service and the cloud layer severely limited his visibility, and thus he followed the climb instruction issued by TCAS. The Korean Air pilot reported that the vertical separation between his aircraft and the phantom aircraft indicated by TCAS decreased to zero before increasing, and before reaching zero TCAS instructed him to increase his rate of climb. The pilot complied and pitched his aircraft further, unknowingly placing it on a collision course with British Airways flight 027, which was now closing in rapidly from above as shown in Figure 2 [7, 8].

As the Korean Air Cargo jet was making its climb, the crew of the British Airways passenger flight reported nothing unusual in their cockpit. Their TCAS display indicated traffic approaching head-on but still flying safely 2,000 feet below their own aircraft. Then, the TCAS unit suddenly issued a descend instruction and showed the traffic now approaching from only a few hundred feet below. The flight crew began to comply with the instruction and pitched the nose down just before seeing the Korean Air jet emerge from the cloud layer below right in front of them. The two aircraft darted past one another separated only by 600 feet, well below the minimum separa-

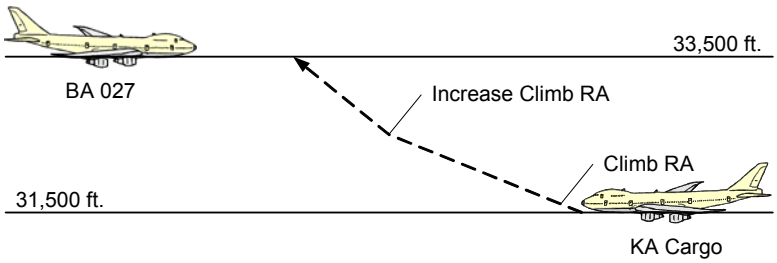

Fig. 2. British Airways flight 027 incident sequence

tion limits for commercial aircraft. The entire incident lasted about 35 seconds [10].

3.2 TCAS Overview

The Traffic Alert and Collision Avoidance System, or TCAS, is an onboard system designed to alert pilots of approaching traffic and provide guidance to avoid traffic conflicts and maintain proper aircraft separation [9]. TCAS detects and tracks surrounding aircraft using technology similar to that used by air traffic control radars to track aircraft from the ground. The system uses a radio transceiver to broadcast an interrogation signal via a directional antenna. Any nearby aircraft that are equipped with transponders will detect the signal and "squawk" back a reply containing information such as the aircraft's altitude, heading, airspeed, and vertical rate of climb or descent.

TCAS receives flight data for the aircraft on which it is operating from two independent air data sources. These data are passed into a comparator where they are averaged before being sent to the TCAS logic unit as illustrated in Figure 3. If the comparator detects that the variance in the inputs from the air data sources is too large, it raises an error signal that prompts TCAS to shut down and print an error message on the primary flight display. This design allows the system to detect but not tolerate disagreement between the air data sources or a failure of one of the sources. The comparator on the TCAS unit installed in the Korean Air Cargo jet featured an Enable line that if set to one would cause the comparator to function normally and if set to zero would

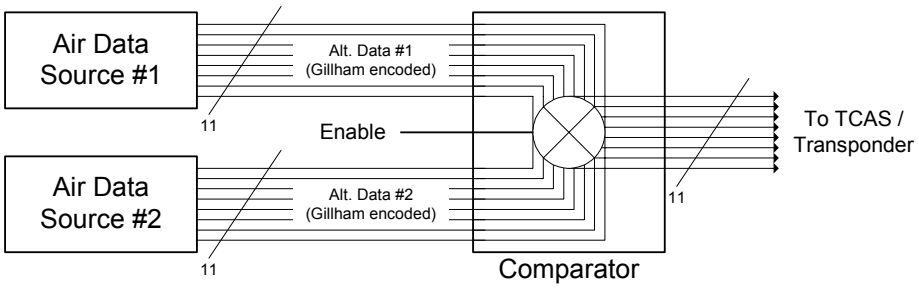

Fig. 3. Simplified schematic of the air data comparator

cause it to function as a pass-through between one of the air data sources and TCAS [7].

The air data sources report altitudes using an 11-bit binary encoding scheme known as Gillham code. Altitude data is sent using this encoding both to the transponder where it is transmitted as part of the transponder's interrogation reply and to the comparator where the data is averaged and forwarded to TCAS. Neither Gillham code nor the transponder protocol employ any error detection or correction mechanisms to verify the integrity of the data. TCAS compensates for this by maintaining histories of the transponder returns from each tracked aircraft that it compares with new returns as they arrive. If a new return contains a fluctuation that is atypical of the performance capabilities of jet aircraft, such as a sudden change in altitude, TCAS assumes that the return is faulty and discards it. TCAS will continue to discard faulty returns for up to one minute from the detection of the original faulty return, at which point it resumes processing the returns normally irrespective of whether the fluctuation has disappeared.

TCAS is one of three mechanisms in the air traffic system designed to help maintain proper aircraft separation and prevent midair collisions. Air traffic control (ATC) is the primary line of defense, and air traffic controllers can resolve traffic conflicts long before pilots or TCAS are even aware of them. TCAS is the secondary system and only reports conflicts when they are projected to occur within one minute in order to give ATC time to resolve the conflict first. Visual identification is the last defense mechanism.

3.3 Postincident Actions

In response to the incident, the UK's Civil Aviation Authority (CAA) and British Airways each conducted investigations to explain the behavior of the Korean Air jet and determine why the TCAS unit onboard the British Airways flight failed to issue an advisory to the flight crew until moments before the two aircraft reached their closest point of approach. An inspection of the TCAS unit installed on the British Airways jet did not detect any problems. When the TCAS unit aboard the Korean Air jet was inspected, however, investigators found that circuitry related to TCAS function had been damaged in two locations. In the first, part of the data line used to send the pressure altitude reading from one of the air data sources had been damaged, resulting in a bit-stuck-at-one error on the line. In the second, a pin on the Enable line to the comparator had been pushed back, causing it to short open, thereby disabling the comparator. The result of this problem was that faulty altitude values were allowed to pass through to the TCAS logic unit unchecked.

Although the air data source was sending the correct altitude value, the bit-stuck-at-one error on the data line caused TCAS to receive a value containing a one-bit discrepancy that corresponded to a 2,400-foot difference in altitude according to Gillham code [7]. Thus, the TCAS unit aboard the Korean Air jet believed it was flying at 33,900 feet, placing it 400 feet above the British Airways jet. According to the separation rules used by TCAS, this created a conflict between the two aircraft, and since TCAS believed it was the one on top, it issued a climb resolution advisory (RA) to the pilot. As the pilot executed the instruction and the aircraft's altitude began to increase,

the altitude value reported to TCAS also changed; however the one-bit error caused TCAS to think the aircraft was actually descending, decreasing the separation between it and the intruder. This led the system to revise its RA and instruct the pilot to climb faster, placing the two aircraft on what was actually a near-collision course.

Just as the error on the data line was causing incorrect altitude values to be sent to TCAS, it was sending the same incorrect readings to the British Airways jet where the fault tolerance mechanism discussed earlier detected the sudden altitude jump and began discarding the erroneous returns. This prevented the TCAS unit aboard the British Airways jet from issuing a false RA, but it also meant that the crew of flight 027 was unaware that the Korean Air jet below was climbing toward them. This continued until moments before the closest point of approach, when the TCAS unit finally started processing the returns again and issued a descend RA.

With the assistance of Korean Air, the CAA determined that the damage to the Korean Air Cargo jet's TCAS unit occurred during maintenance to the aircraft's avionics systems. Upon concluding its investigation, the CAA issued an airworthiness directive requiring air carriers using Gillham code to check the altitude values being transmitted by the transponder throughout the operational envelope of the aircraft and to periodically inspect the comparator unit to ensure that it is functioning properly. The CAA also notified other European aviation regulatory agencies and the FAA of the problems it found as well as manufacturers of transponder and TCAS equipment, and it issued a recommendation to aircraft operators urging them to consider using other encoding schemes for transmitting altitude data instead of Gillham code. At the end of its report, the CAA noted, "This incident shows the effects that secondary failures can have on primary systems fitted to aircraft today. Regardless of the integrity of the collision avoidance system, it shows that relatively minor faults in the interfacing system can still contribute to a serious safety risk" [7]. Indeed, safety is a systems issue, and the fact that one subsystem has high "integrity" does not imply that the resulting system will as well.

3.4 Analysis

TCAS Design Issues. The follow-up actions taken by the CAA focused on the maintenance issues that caused the damage to the TCAS system aboard the Korean air jet and those that allowed it to operate in such a state. While these issues are important, serious design issues also exist in TCAS, at least in the models aboard the incident aircraft. The transponder protocol does not employ any error detection or correction mechanisms to verify the integrity of the data, and even simple transmission faults that could be detected by employing parity checking or cyclic redundancy checks (CRCs) can pass through to TCAS unnoticed by the transceiver hardware. TCAS attempts to compensate for this by examining the history of each tracked aircraft to check for erroneous returns, but this is by no means a complete solution.

The second issue pertains to the design of the comparator used to verify the data gathered from the two air data sources. The purpose of receiving data from two separate sources is to stabilize the air data TCAS receives and reduce the likelihood that faulty data is allowed to pass into TCAS undetected. The dependability increase achieved through this fault detection mechanism is defeated, however, if the compara-

tor itself introduces vulnerabilities into the system. In the design used by the TCAS unit aboard the Korean Air jet, the Enable line to the comparator presented such a vulnerability.

Incident Investigation. The result of the CAA's investigation into the flight 027 incident was a three-page report briefly describing the incident and the investigation, a paragraph documenting the analysis, and summary lists of the actions taken, conclusions, and recommendations [7]. The investigation was fairly informal and conducted with the assistance of British Airways and Korean Air officials. This pales in comparison to the formal investigations that are launched in response to accidents involving loss of life, injury, or substantial damage to property and the voluminous reports they produce. This is not to single out the UK's Civil Aviation Authority, however, as the practice is shared by investigative agencies worldwide. British Airways conducted a more detailed investigation of the incident, but has not officially released the report of its investigation or findings to the public.

3.5 Related Incident

The British Airways incident over China in June 1999 followed a similar incident that occurred between two aircraft in January 1998 over Hawaii. One aircraft's TCAS unit issued a false traffic advisory because an air data computer had malfunctioned and was reporting the aircraft's altitude as 1,500 feet higher than its actual position. Fortunately, air traffic controllers happened to notice a discrepancy between the aircraft's altitude as reported by its transponder and that reported by the flight crew and were able to defuse the situation before it escalated further.

The Australian Transport Safety Bureau (ATSB) launched an investigation into the incident. When the British Airways incident occurred 15 months later, ATSB investigators saw the similarities between the two incidents and obtained a copy of British Airways' findings. The findings went beyond recommending better maintenance and addressed the design issues highlighted in this paper along with other issues such as human factors. British Airways investigators recommended changes to the TCAS design and displays, and also advised that Gillham code be abandoned in favor of more robust solutions for transmitting altitude data. Had such an investigation been performed in a timely fashion following the January 1998 incident over Hawaii, the near-collision over China might have been prevented.

3.6 Lessons Learned

British Airways and the CAA presented lessons and recommendations for the improvement of TCAS, transponder systems, and policies for maintaining and inspecting these systems. Once again, after looking at all the evidence about this accident that is available to us, we conclude that the lessons learned from this accident were far short of what they should have been. Two additional prominent problems should have been identified and additional significant corrective actions taken.

Lesson 1—Incident Classification. The first lesson is that classification schemes in which investigative resources are allocated to accidents and incidents based on their associated losses de-emphasize the importance of incidents with no losses even though these incidents might still have important lessons to be learned. Incidents provide

opportunities to improve the affected systems without the consequences associated with accidents, and investigators should seize upon these opportunities to prevent similar sequences of events from occurring in the future, possibly with more dire consequences. Many accidents have been preceded by similar incidents, and lives and property could have been spared if the problems contributing to those accidents had been addressed when they manifested themselves earlier.

Lesson 2—Criticality of Design Faults. The second lesson is that when an accident or incident occurs involving a safety-critical computing system, investigators must pay particular attention, systematically and comprehensively, to possible system design faults. Because design faults are difficult to understand, attempting to compensate for the design faults in a system by trying to prevent the conditions that trigger the faults from recurring rather than correcting the faults themselves is a strategy that is unlikely to succeed.

In this incident, attention was paid correctly to TCAS system maintenance. But the TCAS design faults that permitted the effects of the maintenance damage to go unnoticed are clearly items that should have been recognized and acted on by regulating authorities.

4 Conclusions

The systems examined in this paper are each part of much larger systems designed to enhance the safety of commercial air travel. The MSAW system is part of the FAA's program to prevent CFIT accidents, and TCAS plays a significant role in reducing the likelihood of mid-air collisions. In both of the incidents described in this paper, the systems involved were viewed as if they were isolated; ample consideration was not given to the roles these systems played in the overall systems of which they were part.

When changes are made to a safety-critical system, the original safety analysis of the system is invalidated and must be performed again to ensure the system is still compliant with its original safety requirements. Moreover, when a new safety system is to be added to an existing set of barriers, that system must be examined to ensure that it will not adversely affect safety through faulty operation. These lessons are not new to the safety engineering community, but their importance in safety-critical, software-intensive systems is not fully appreciated.

The two incidents described in this paper illustrate the need for more comprehensive investigations of incidents involving safety-critical software systems. Both the NTSB and the UK CAA successfully determined the sequence of events leading to the Korean Air and British Airways incidents; however investigators missed important lessons when they interpreted this information, and as a result their recommendations were incomplete. When an incident involves a software system, that system's development and maintenance histories, safety analysis and operational context should be investigated as a fundamental component of the investigation just as human actions, aircraft maintenance records, and organizational policies and regulations are. Doing so could reveal additional lessons to system developers and managers that would otherwise remain hidden.

Finally, we note that the need to consider subtle differences in event sequences

reemphasizes the problems associated with investigating accidents and incidents differently. The Korean Air flight 801 accident received much more investigative and public attention than did the British Airways flight 027 incident, even though the latter could have developed easily into a tragedy with twice the number of casualties. All accidents begin as incidents, and luck might be the only factor preventing an incident from developing into a catastrophe. From an investigative perspective, the lessons to be learned from an incident and its related accident are equally important since these lessons usually focus on preventing the incident rather than mitigating the extent of the loss. This is especially important in the context of safety-critical software systems, where design faults are shared by all instances of a particular system. If an incident reveals the presence of a design fault, investigators have an opportunity to develop recommendations to prevent the fault from manifesting itself again in other installations of the same system, possibly with more severe consequences.

Acknowledgements

It is a pleasure to thank Michael Holloway for his suggestion that important lessons are sometimes not learned in accident investigations. We also thank Peter Ladkin and Chris Johnson for their insightful reviews of an earlier version of this paper. This work was funded in part by NASA Langley Research Center under grants numbered NAG-1-2290 and NAG-1-02103 and in part by Microsoft.

References

1. National Transportation Safety Board. *Controlled Flight Into Terrain, Korean Air Flight 801, Boeing 747-300, HL7486, Nimitz Hill, Guam, August 6, 1997*. Aircraft Accident Report NTSB/AAR-00/01. Washington, DC.
2. National Transportation Safety Board. *Public Hearing in Connection With the Investigation of Aircraft Accident, Korean Air Flight 801, B-747-300, Agana, Guam, August 6, 1997*. 24 March 1998. Honolulu, Hawaii.
3. Federal Aviation Administration. “Fact Sheet: CAST Accomplishments: Civil Aviation Controlled Flight Into Terrain (CFIT).” 26 March 2001. Washington, DC.
4. National Transportation Safety Board. “Guam ARTS-11A MSAW Chronology.” *Korean Air Flight 801, B-747-300, Agana, Guam, August 6, 1997, Public Hearing Exhibit List*. docket no. SA-517, exhibit no. 3-U. 17 October 1997. Washington, DC.
5. Leveson, N.G. *Safeware: System Safety and Computers*. Reading, MA: Addison Wesley. ‘95.
6. National Transportation Safety Board. *Controlled Collision with Terrain Transportes Aereos Ejecutivos, S.A. (TAESA) Learjet 25D, XA-BBA Dulles International Airport Chantilly, Virginia June 18, 1994*. Aircraft Accident Report NTSB/AAR-95/02. Washington, DC.
7. U. K. Civil Aviation Authority. “Hazardous Loss of Separation Between Two Aircraft Over Chinese Airspace.” Doc Ref KMH/Pap/059, issue 1. 28 October 1999. London, U. K.
8. Carley, William M. “Wires Crossed: Flawed Safety Device In Jets Gets Blamed For a Near Catastrophe.” *Wall Street Journal.* 12 October 1999, eastern ed.: A1.
9. MITRE Corp. “TCAS: Traffic Alert and Collision Avoidance System.” 31 December 1998. <http://www.mitre.org/pubs/showcase/tcas/tcas.html>
10. Australian Transport Safety Bureau. “Safety Deficiencies: Errors in Traffic Alert and Collision Avoidance Systems.” Output no. R19990156. 9 Sept. 1999. Canberra City, Australia.