

# Risk-Based Classification of Incidents

William S. Greenwell    John C. Knight    Elisabeth A. Strunk

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

**Abstract.** As the penetration of software into safety-critical systems progresses, accidents and incidents involving software in such systems will inevitably become more frequent. Identifying lessons from these occurrences and applying them to existing and future systems is essential if these incidents are to be prevented in the future. Unfortunately, investigative agencies do not have the resources to fully investigate every incident under their jurisdiction and domains of expertise and thus must prioritize certain undesired events and allocate resources accordingly. In the aviation community, most investigative agencies prioritize undesired events based on the severity of their associated losses, allocating more resources to those involving extensive aircraft damage or injury to passengers. We argue that this scheme is inappropriate because it undervalues incidents whose recurrence could have a high potential for loss while overvaluing fairly straightforward accidents involving accepted risks. We then suggest a new strategy for prioritizing incidents based on the risk arising from incident recurrence.

## 1 Introduction

By their very nature, commercial aviation accidents demand our attention. Major accidents can create spectacular scenes of carnage and destruction that threaten public confidence in commercial air travel. At the very least, accidents remind us that, while very safe, there is still some risk in commercial air travel, and they often force engineers and regulators to rethink their safety analyses and add additional safeguards to the air transit system. It is out of a desire to improve safety and prevent the recurrence of tragedy that society demands investigations into accidents in an attempt to learn as many lessons from them as possible.

It is well known that incidents in which no loss is incurred are as valuable as accidents in their ability to teach lessons [2]. Despite this, incidents rarely command the attention that accidents do, and we argue in this paper that this is a serious imbalance with possibly serious consequences. We discuss two commercial aviation events involving safety-critical software systems in which the failure of a safety-critical software system contributed to the occurrence of the event. The first event resulted in a crash with hundreds of fatalities. The second event did not develop into an accident, although the failure of the system involved led to a near-collision between two large commercial aircraft. After summarizing the events, we show that the first event received a much more rigorous investigation than

the second, even though the later could have resulted in almost twice the number of fatalities. We then suggest an alternative incident classification scheme that we claim will more appropriately match investigative resources to events whose recurrence would likely have catastrophic consequences.

Both of the events that we discuss in this paper could have been prevented in many ways. However, the need for change in incident classification is illustrated very clearly by the fact that both events were *preceded* by similar incidents that indicated the possibility of a systemic problem [3].

## 2 Investigative Resources

Unfortunately, most investigative agencies simply do not have the resources to fully investigate every aviation-related undesired event that occurs within their jurisdiction and must prioritize certain events when allocating investigative resources. Agencies typically prioritize accidents according to the severity of their associated losses. For example, the National Transportation Safety Board (NTSB) classifies an accident as “major” if the accident results in the destruction of a commercial aircraft, multiple fatalities, or one fatality and substantial damage to a commercial aircraft. According to NTSB statistics, 74 major accidents occurred between 1983-2002 compared to 581 accidents receiving less severe designations [5]. NTSB investigators use a special operating manual when investigating major accidents that guides them in collecting evidence, holding public hearings, and preparing final reports [6]. Reports are typically reserved for major accidents; synopses are prepared for less severe accidents and then stored in a database.

The NTSB is not unique in employing these procedures. Although many other agencies worldwide do not limit their investigations to accidents as the NTSB does, most distinguish between accidents and “serious incidents,” including the U.K. Air Accidents Investigation Branch (AAIB), the French Bureau d'Enquêtes et d'Analyses (BEA), the German Federal Bureau of Aircraft Accidents Investigation (BFU), the Accident Investigation Board of Finland (AIB), the Australian Transport Safety Bureau (ATSB), the Taiwanese Aviation Safety Council (ASC), and others. While the definition of “accident” is typically clear, the term “serious incident” is often not well-defined. The AAIB and BFU offer guidelines that give examples of serious incidents, but admit that these guidelines are not comprehensive. The ATSB uses a five-category system to classify accidents and incidents, but the criteria for categorizing an occurrence are subjective. The Canadian Transport Safety Board (TSB) does not actually distinguish between accidents and incidents but labels both types of events as “occurrences.” They classify and investigate occurrences based on “whether the investigation is likely to lead to reduced risk to persons, property, or the environment” [9]. This is similar to the scheme we propose; however their criteria are still quite subjective.

The effect of allocating resources to accident and incident investigations based on the severity of their associated losses is that less severe accidents might receive only a small amount of attention from investigators, and incidents might not be investigated at all. However, many major *accidents* are preceded by similar *incidents*, in which it was only by coincidence that a loss did not occur. This is particularly important in the context of safety-critical software systems because design faults present in such systems can manifest themselves with unpredictable consequences. If the systems control hazardous opera-

tions, they might bring direct harm to passengers or crew. Alternatively, if the systems provide advice or warnings to pilots, they might raise false alerts or issue erroneous guidance to pilots, who could inadvertently jeopardize safety by acting on this information.

### **3 Event Descriptions**

To illustrate the distinction between the way in which accidents and incidents are investigated, we examine the investigations conducted following a major accident and a major incident. We begin in this section with brief descriptions of the events and then present details of the event investigations.

#### **3.1 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 [7], 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.”

We discuss in detail how the MSAW system at Guam contributed to the accident elsewhere [3]. Essentially, the system had been disabled years before the accident in order to eliminate nuisance low-altitude warnings. Prior to the accident, the FAA had received multiple warnings that MSAW systems were being configured improperly. These included a safety recommendation from the NTSB issued in response to a previous accident urging the FAA to verify the MSAW configurations at each of its air traffic control facilities as well as an evaluation of the Guam facility that noted its MSAW inhibition. After the Korean Air flight 801 accident, the FAA developed a comprehensive program to manage its MSAW installations, but continued to be plagued by accidents in which MSAW configuration errors were cited as contributory factors.

The NTSB began its investigation into the Korean Air flight 801 accident immediately after the crash. The Board adopted its final report, a 212-page document, on January 13, 2000. The report contains 134 pages of factual information pertaining to the accident and 37 pages of analysis. The investigation yielded 36 findings and a set of 15 recommendations mostly addressed to the FAA. During the investigation, the NTSB held a three-day public hearing into the accident in which officials from the FAA, Korean Air, the government of Guam, and other organizations gave testimony. The transcript from this hearing spans approximately 430 pages [8].

### **3.2 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 [10].

Prior to the incident, the two aircraft were travelling in opposite directions along the same airway with a safe margin of 2,000 feet of vertical separation. The British Airways passenger flight was flying above the Korean Air Cargo flight. The incident sequence began when a collision avoidance system onboard the Korean Air Cargo flight malfunctioned and mistakenly determined the aircraft's altitude to be 2,400 feet higher than its true altitude. This caused the system to believe that a traffic conflict existed between the two aircraft, which prompted it to erroneously instruct the Korean Air pilot to climb in order to avoid the conflict. Because no air traffic control service was available in the region of airspace in which the aircraft were operating and meteorological conditions prevented the pilots from visually identifying each other's aircraft, the Korean Air pilot had no reason to question the collision avoidance system's instruction and thus complied. This placed the two aircraft on a collision course that neither flight crew detected until moments before the aircraft reached their closest point of approach. British Airways officials later noted that it was only by coincidence that the two aircraft avoided each other and that they would have likely collided had they been using more precise navigation systems [1].

With the assistance of Korean Air, the CAA determined that the malfunction in the Korean Air Cargo jet's collision avoidance system was caused by damage inflicted during maintenance to the aircraft's avionics systems. Upon concluding its investigation, the CAA issued an airworthiness directive requiring air carriers using similar systems to periodically conduct inspections to ensure the systems are using correct altitude values. The CAA also notified other European aviation regulatory agencies and the FAA of the problems it found, as well as equipment manufacturers, and it issued a recommendation to aircraft operators urging them to consider using more robust schemes for handling altitude data.

The U.K. Civil Aviation Authority (CAA) and British Airways each conducted their own investigations into the incident. The CAA's report does not indicate when its investigation into the incident began; however the report is dated October 28, 1999, suggesting that the investigation lasted at most four months. The report is three pages long and includes eight paragraphs of factual information spanning two pages and a single paragraph of analysis. It contains a single conclusion and three recommendations directed at operators and equipment manufacturers. No public hearing was held in response to this incident. British Airways prepared a more detailed report on the incident, but that report has not been officially released to the public.

### **3.3 Event Comparison**

Clearly, the Korean Air flight 801 accident received a more rigorous investigation than did

the British Airways flight 027 incident. In order to help quantify the extent of the difference, we have summarized data from the events and their investigations in Table 1. On the basis of loss, the near-collision involving British Airways 027 had no casualties compared to a 90% fatality rate in the Korean Air 801 accident. In addition, neither of the Boeing 747s involved in the near-collision sustained any damage from the incident, whereas the 747 involved in the Guam accident was destroyed. Examining loss alone, the Korean Air accident over Guam appears far more important than the near-collision over China, and thus the large discrepancy in output from the two investigations might seem warranted.

Comparing these events solely on the basis of loss is deceiving, however, as the British Airways incident could have easily developed into an accident with almost twice the number of fatalities as the Korean Air flight 801 crash in Guam. As British Airways officials noted, it was entirely by luck that the British Airways passenger flight and the Korean Air Cargo flight did not collide. By the time the Korean Air pilot inadvertently placed his aircraft on a collision course with British Airways flight 027, all of the barriers designed to prevent midair collisions had been defeated. If the incident sequence were to recur with similar aircraft, a variation in wind direction or in navigational precision could lead to a much more dire outcome and almost certainly would have if the incident aircraft had been using the Global Positioning System (GPS) navigation systems in widespread use today. Under the accident classification schemes employed by most investigative agencies, this catastrophic outcome would be necessary for a major investigation to be undertaken, even though the findings and recommendations would likely be the same as if an equally rigorous investigation had been conducted into the incident alone. This should

	Korean Air 801	British Airways 027
Classification	Accident	Incident
Persons On Board	254	419
Fatalities	228	0
Injuries, Serious	26	0
Injuries, Minor	0	0
Total Casualties	254	0
Aircraft Damage	Destroyed	None
Investigation Length (months)	30	4
Final Report Length (pages)	212	3
Factual Information (pages)	134	2
Analysis (pages)	37	1
Findings / Conclusions	36	1
Recommendations	15	3

**Table 1: Comparison of Korean Air flight 801 and British Airways flight 027**

not be the case. New classification schemes are necessary in order to better allocate investigative resources to incidents whose recurrence could have more severe consequences.

In reviewing this comparison, one might argue that the vast difference between the Korean Air and British Airways events was not necessarily because of their associated losses but rather due to the fact that different agencies investigated each event. Had both events been investigated by the NTSB or CAA, the figures might have matched more closely. Because the NTSB does not investigate incidents, however, had the near-collision involving British Airways 027 occurred in U.S. airspace, it is unlikely the NTSB would have issued any report on it. Similarly, if the Korean Air flight 801 accident had occurred in British airspace, it would have been investigated not by the CAA but by the AAIB, whose formal reports are similar to the NTSB's final reports in structure and length.

#### **4 Risk-Based Classification of Incidents**

The term "incident" can be defined in a variety of ways but typically involves the failure of a network of barriers designed to protect a system from one or more hazards. An incident becomes an accident when it is coupled with a loss event such as a crash or collision in which damage or casualties are incurred. It is often the case that luck determines whether an incident develops into an accident and, if so, what the extent of the loss will be.

When investigating accidents, investigators can issue recommendations aimed at preventing the associated incident or at mitigating the severity of the loss, and they usually do both. While attempting to mitigate loss given the occurrence of an incident can help to reduce the severity of accidents, some degree of loss is almost always inevitable. On the other hand, if the incident itself is prevented, it cannot develop into an accident and thus no loss will occur. Therefore, recommendations aimed at preventing incident recurrences are likely to be more effective in preventing future losses. Indeed, 13 of the 15 recommendations issued by the NTSB in response to the Korean Air flight 801 accident were aimed at preventing the recurrence of incidents in which aircraft descend below safe altitudes during final approach. Only two focused on mitigating losses by suggesting improvements to Guam's emergency response units.

Given that accidents begin as incidents and that incident prevention should be the focus of investigations, incidents are opportunities for investigators to identify problems and suggest safety improvements without the losses associated with accidents. Accident classification schemes based on loss should be de-emphasized because loss alone is a poor indicator of an incident's potential for learning new lessons and preventing future incidents. New schemes should be adopted in which resources are allocated to incident investigations based on the risk associated with the recurrence of each incident. To this end, we propose the fundamentals for such a scheme.

Risk is defined as the probability that an event will occur multiplied by the anticipated cost derived from the occurrence of the event. When an incident occurs, it suggests the presence of a deficiency in the safety systems involved that, if not corrected, could lead to one or more recurrences of the incident. A useful measure of the importance of an incident, therefore, is the total risk that society faces if nothing is done to prevent recurrences.

The total risk of such a recurrence is given by the formula:

$$\text{Total Risk} = P[\text{Incident Recurrence}] \times E[\text{Cost}] \times (\text{Number of Susceptible Systems})$$

In this formula,  $P[\text{Incident Recurrence}]$  is the probability of a recurrence of the incident for an individual system and  $E[\text{Cost}]$  is the anticipated cost associated with a recurrence. The three terms comprising this equation follow one's intuition in prioritizing incidents. Clearly, those that have relatively high probabilities of recurrence with high expected costs warrant significant investigation, particularly if numerous systems are already deployed that might also be susceptible to the incident. Likewise, an incident with a small probability of recurrence, a low expected cost, or for which there are only a handful of susceptible systems might warrant only a minor investigation.

As an example of the use of total risk, consider the incident involving British Airways flight 027. It is very difficult to estimate the probability of recurrence but not impossible. The rates of failure of the relevant hardware components are probably known as is the rate of undetected damage occurring during maintenance. The cost of such an incident were it to result in an accident would be very high since there would be considerable loss of life and equipment. The Number of Susceptible Systems is also likely to be very high because of the prevalent use of TCAS. Thus, a rough estimate of the total risk could be calculated quickly and used as an indicator of the significance of the incident. As investigative actions proceeded, the estimate could be refined so as to permit a rational decision as to when to terminate investigation.

A second important use of the concept of total risk is to guide the actions taken following an investigation. If the total risk is high, then the follow-up actions should have a high probability of reducing the total risk to an acceptable level. Many options are available to investigative and regulatory agencies and they need to be used carefully. At one extreme, there is the option of grounding the fleet and at the other there is the option of no action. In between, there are a variety of possibilities including required inspections, required equipment replacement, required equipment redesign, and so on. There are also options about how quickly any action should occur. Selection among options is a difficult activity if there is no effective rating mechanism for the seriousness of an incident.

Using BA 027 as an example once again, the actions taken following the incident were insufficient and fragmented despite the fact that the total risk by our estimation was very high. Upon concluding its investigation, the CAA issued an airworthiness directive requiring air carriers using similar equipment to check and periodically inspect the equipment 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 since that was part of the problem. The CAA's recommendations did not require mandatory changes and the probability that they would reduce total risk to an acceptable level was small. More importantly, the report by British Airways contains useful insights about the incident yet it has not been made public nor led to appropriate general recommendations.

The notion of total risk is a starting point for a metric that will allow investigators to better assess the importance of incidents and allocate investigative resources accordingly. By assessing incidents based on the risks of future losses from their recurrence rather than

their immediate losses, investigators can be more proactive in detecting safety problems before they contribute to accidents involving casualties or damage to aircraft.

A lot of work remains to be done before this metric can be put into practice. Because incidents are rare occurrences, estimating their probabilities is difficult. Moreover, a model of cost will be needed to assess the likely loss associated with an incident that takes into account fatalities, serious and minor injuries, and damage to aircraft and other property.

Once these challenges are overcome, a technique for quickly assessing the probability of recurrence and expected cost for a given incident will need to be developed in order to compute the relative risk of a future recurrence. Investigators will need to be able to apply the technique early in the investigative process, meaning that most of the details pertaining to an incident will still be unavailable when the technique is applied and some details might be inaccurate.

## **5 Conclusions**

Commercial aviation accidents are serious occurrences that demand public investigations in order to correct safety problems and prevent future losses. Incidents are also important, however, since they often present the same opportunities to identify new lessons without the losses associated with accidents.

Current accident classification schemes used by investigative agencies to allocate resources to investigations place too great an emphasis on the immediate loss from an accident and as a result undervalue the importance of incidents with no losses. Consequently, incidents suggesting the presence of serious safety problems in onboard and ground-based systems are often ignored or not investigated with sufficient rigor to uncover these problems, which if left uncorrected could contribute to future incidents with more tragic outcomes. This dilemma was illustrated by the vast difference in the investigations conducted into the Korean Air flight 801 and British Airways flight 027 incidents, despite the observation that the later could have easily developed into a major accident with almost twice the number of casualties as the Korean Air crash into Guam.

To mitigate this problem, investigators should reconsider the practice of classifying incidents based on their losses, and instead classify them based on the risk of future losses. Adopting risk-based schemes will allow investigators to be more proactive and address safety problems before they contribute to accidents with extensive casualties. For risk-based classification schemes to be useful, techniques will have to be developed for investigators to quickly assess the risk level of incidents early in the investigative process so that they can allocate resources accordingly.

Finally, we note that the need for change in incident classification is indicated strongly by the occurrence of incidents and accidents closely related to and preceding the ones we have used for illustration. The Korean Air flight 801 accident followed a similar incident in 1994 that also involved a mis-configured MSAW system in which 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 [3]. The British Airways flight 027 incident followed a similar incident that also involved TCAS processing incorrect altitude data that occurred between two aircraft in January 1998 over Hawaii [3].

## 6 References

- [1] 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.
- [2] Department of Transportation, Federal Aviation Administration, "Aircraft Accident and Incident Notification, Investigation, and Reporting", Order 8020.11B (August 2000) <[http://www2.faa.gov/avr/aai/8020\\_11b.pdf](http://www2.faa.gov/avr/aai/8020_11b.pdf)>
- [3] Greenwell, William S. and Knight, John C. "What Should Aviation Safety Incidents Teach Us?" Submitted to: SAFECOMP 2003, The 22nd International Conference on Computer Safety, Reliability and Security, Edinburgh, Scotland (September 2003).
- [4] National Transportation Safety Board. "Accidents, Fatalities, and Rates, 2002 Preliminary Statistics, U.S. Aviation." <<http://www.nts.gov/aviation/Table1.htm>>
- [5] National Transportation Safety Board. "Accidents and Accident Rates by NTSB Classification, 1983 through 2002, for U.S. Air Carriers Operating Under 14 CFR 121." <<http://www.nts.gov/aviation/Table2.htm>>
- [6] National Transportation Safety Board. *Aviation Investigation Manual: Major Team Investigations*.
- [7] 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.
- [8] 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.
- [9] Transportation Safety Board of Canada. "Investigation Process." (18 September 2002). <[http://www.tsb.gc.ca/en/investigation\\_process/what\\_we\\_do.asp](http://www.tsb.gc.ca/en/investigation_process/what_we_do.asp)>
- [10] 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.