Safety Science, Volume 43, Issue 9, November 2005, Pages 715-738.

To be published in Safety Science
Post Referees Review
@ 19 November 2004

REDUCING MID-AIR COLLISION RISK IN CONTROLLED AIRSPACE: LESSONS FROM HAZARDOUS INCIDENTS

Peter Brooker

Cranfield University

ABSTRACT

The collection and analysis data on hazardous air traffic management (ATM) incidents is an important task. Expert judgement about such incidents needs to be carried out within a systematic and consistent safety framework. The mark of the genuine safety expert is to be able to ask the right questions concerning potential accidents.

Hazards and risks are not 'facts' or 'events' that 'exist', but rather judgements made about conditional futures and their consequences. A hazardous situation is one in which the outcome was not 'system controlled', with some potential outcomes having significant negative consequences". System controls in this sense cover all the means by which the system is held stable (= defended) against the potential negative consequences.

The ATM system can be (over-) simplified to consist of three structural system layers acting as the system controls: Planning (pre-operational), Operation (the flight in progress), and Alert (the ground and air protection enabled by conflict alert systems, on which the controller/pilot will act). A hazardous event is one in which a high degree of conflict between aircraft is observed plus a low confidence that the remaining system layers would generally provide the necessary corrective action.

1. INTRODUCTION

Aviation's safety track record demonstrates that it is an industry whose people are focused on continuous improvement. The prime safety goal of the air traffic management (ATM) of en route commercial flights is to reduce the risk of mid-air collisions. Safety has improved to such an extent that collisions are now rare, so collecting data on hazardous ATM Incidents has therefore always been seen as an important task. This incident data, collected consistently, can be viewed as a key indicator of the 'health' of the ATM system. Incidents can provide insights into the frequency of known error and failure types, and also enable new types to be detected. Analyses of the processes and characteristics of incidents provide insights

into potential system design weaknesses. Moreover, planned changes to the system to improve safety can be tested against such incidents to demonstrate that risk is being reduced. [NB: 'risk' is often used in safety analyses as a combination of frequency of occurrence and its severity: in the following, the accidents always have the same severity – an accident with many fatalities.]

A large variety of crucial – and intrinsically difficult – safety questions can be asked about ATM incidents:

> Which incidents should be judged the most important to ATM system safety?
>
> Which incidents give most guidance about potential accidents?
>
> In what ways should incidents be categorised and analysed to help pinpoint key safety issues?
>
> Are 'minor' incidents of any safety importance?
>
> How should the relevant importance of different incidents be assessed or weighted to provide a true picture of the health of the ATM safety system?

This paper attempts to make a start – no more than that – in answering these kinds of questions. Interpretation of incident data needs a systematic and consistent safety framework. Without such a framework, it is difficult to gain the desired insights into system design weaknesses that have the potential for accidents. There is little point in collecting and categorising a great deal of data, unless it is analysed in a way that systematically reveals the safety lessons that help to reduce the likelihood of potential future accidents.

The aim here is to try to ensure that expert judgement about ATM incidents can be carried out within a systematic and consistent safety framework, rather than producing formulaic prescriptions. The focus is <u>how</u> the <u>system</u> can be improved rather than the whys of causes. System jargon is avoided: the approach is mainly through an 'ordinary language' analysis of safety terms and logic, plus some metaphors to describe the key features of a safe ATM system.

The following text consists of six sections:

2. The Nature of Hazardous ATM Events
3. ATM System Layers and Risk Probabilities
4. Airproxes, Hazards and System Layers
5. Discussion
6. Conclusions

2. THE NATURE OF HAZARDOUS ATM EVENTS

There are several different – and complementary – ATM incident systems in use in the UK. Airproxes derive from pilot and controller reports, and are the province of the UK Airprox Board [UKAB] (1999-). The Civil Aviation Authority's (CAA) Safety Regulation Group (SRG) has a Mandatory Occurrence Reporting scheme, which covers all kinds of aviation incident not just ATM-related ones (CAA, 2003). National

air traffic Services Ltd (NATS) records inter alia, data on Short Term Conflict Alert (STCA) and Separation Monitoring Function (SMF) monitoring (eg NATS (2004)).

What actually is an ATM incident? To start this analysis, consider a frequently used phrase about ATM incidents: 'hazardous'. A dictionary definition is on the lines of:

> hazardous: "anything which might cause an accident, create danger, etc"

This is a complex statement. It has three important elements:

> "might" – this is a future conditional statement, referring to a possible future
>
> "cause" – some chain of related events is being considered
>
> "accident/danger" – there are significant negative consequences

Thus, when something is referred to as being hazardous, this is actually a statement about something that will happen (or is the process of happening) that could have significant negative penalties to the participant(s) or other individuals.

Some examples help to illustrate these different elements (these will also be used in the further development of the concept). To walk a tightrope would be hazardous because an inexperienced person would be likely to fall of it (but not so for an experienced and healthy circus performer). Driving a car on an icy road around a corner could be hazardous because the driver's knowledge about the road's friction would not be as good as in normal conditions, and so the driver might fail to steer the vehicle safely. It would be hazardous for a pilot to adopt a markedly non-standard phraseology when communicating with a controller because this could lead to a misunderstanding about the route the aircraft should follow, and hence possibly produce a flight path conflicting with other aircraft.

Thus, the word hazardous cannot properly be used about events that have already occurred. There is no conditionality about past events: the tightrope was walked, the car was driven, etc. What people are often commenting on is the hazard involved in some kind of similar event in the future. Thus, walking a tightrope is a hazardous thing for people to do because the next person to attempt it may fall off – or perhaps the person after that, etc. Sometimes, cars that are to be driven on icy roads will lead to an accident. Sometimes, there will be occasions when poor voice discipline by aircrew can lead to a mid-air collision.

Not all future things presently viewed as hazardous will actually turn out to have negative consequences – some non-acrobats will be able to walk the tightrope, etc. The likelihood of the negative consequences will depend on factors such as the participant's skills, the environment and safety mechanisms. Thus, the tightrope walking will be much more difficult for the average 90-year-old non-acrobat, and the pilot's poor communication may not lead to danger if the controller detects the choice of the wrong course. Sometimes, the negative consequences are avoided by a deterministic (ie definitely present) measure, eg the traffic police might have set a special speed limit for the icy road and are monitoring every driver. In other instances (eg the controller monitoring the aircraft), there will be probabilistic elements involved, so that, on some proportion of occasions, the negative penalties will occur and the rest of the occasions it will not.

This suggests a definition of a past 'hazardous situation':

> hazardous situation: "one in which the outcome was not 'system controlled', with some potential outcomes having significant negative consequences"

The phrase 'system controlled' means something like:

> system controlled: "the ability to determine the outcome against reasonably foreseen changes and variations of system parameters, such as the abilities of the participant(s), the environment (in the largest sense), and the safety mechanisms in place.

System controls in this sense cover all the means – the effective feedbacks – by which the system is held stable (= defended) against the potential negative consequences: designers, pilots, controllers, software engineers, etc – not just the operational controller. What is preventing an unsafe system state from persisting?

A failure of system control covers both of these kinds of mistake, where the mechanisms make the situation worse, and when they essentially fail to intervene (eg a conflict alert system could fail either by putting the aircraft into more danger or by not alerting). The new phrase here is 'reasonably foreseen'. This means that the assessment of hazard is not to be carried out against 'unreasonable' system parameters. What is unreasonable is a matter for debate and convention. In the example of the driver on the icy road, if the driver had a great deal of experience of driving in such conditions, then it might be reasonable for him or her to assume that the road friction was no worse than he had previously encountered. Phrased another way, to say that a past situation was hazardous implies that if some of the system parameters applicable during that situation had been 'slightly different' or 'reasonably perturbed' then significant negative consequences would have resulted. Something was hazardous because of what might have happened.

This idea, that system control is a key thought, has already been adopted in safety assessment of operational ATM systems. To quote (ARIBAa):

> "In general, a (sub)process is named *controllable* if it has the property that its behaviour can always be steered from an undesirable sequence of events to another, more desirable, one. Opposite to controllable (sub)processes are (sub)processes that have the property that once the initial state is known, the future evolution of the (sub)process is completely determined. This type of (sub)process is named *autonomou*s."

But is this the correct understanding of a concept such as 'hazardous incident'? This 'ordinary language' approach seems to be consistent with the dictionary definitions of related words – Appendix A. The definitions in Appendix A share the underlying ingredients detected in the word 'hazardous'. Thus, the future conditional tense for is implied by phrases such as 'likely to', 'possibility', and 'depending on chance'. Hazards and risks are not 'actual' things – 'facts' or 'events' that somehow 'exist' – but rather judgements made about conditional futures and their consequences, given a lack of information about current system parameters and further events. This

judgement (or perception or opinion) is about the degree of possibility of some unpleasant state of things that may come into existence at some future time.

There is no such thing as ‘actual risk’ unless it is interpreted in this way: combining the definitions above only produces sense if there is this kind of interpretation. [A simple test is to try to explain a putative conceptual phrase such as ‘actual risk’ to someone – what kind of thing would it consistently describe? The logical implications of terms such as ‘hazard’ have generated considerable interest by safety and computing researchers, eg Ladkin (1998).]

Thus, the conclusion is that, whenever attempts are made to classify incidents in terms of risk, these should be 'what if' exercises. The central message is that hazards and risks are judgements, which implies the need to put in place a framework and processes that ensure that the experts are asked to make the most valuable judgements in safety terms. The implication of this to ATM incident analysis is picked up again in latter sections: first, a new description of the ATM system is needed in the next section.

3. ATM SYSTEM LAYERS AND RISK PROBABILITIES

For an ATM System, the ‘system control’, as described in the previous section, can be the responsibility of a new concept: ‘system layers’. The present ATM system has evolved over the decades. It has several distinct components in its operational and technical concept. In roughly their date order of introduction, these are.

Controllers and pilots – people are an integral part of the whole system.

Formal Safety Rules – for the control of traffic, including the minimum separation to be permitted between aircraft.

Radio Telephony – communication between controllers and pilots (only introduced shortly before World War II).

Controlled Airspace – the creation of sectors, volumes of airspace, each handled by a controller team, with a small number of routes; civil commercial traffic is separated from general aviation and military aircraft.

Flight Progress Strips – paper strips, generated by the flight plan computer and kept on plastic holders in ordered racks, which are used to record the details of a flight – being changed to electronic versions.

Radar – processed Secondary Surveillance Radar (SSR) is now used, with the displayed aircraft symbols complemented by callsign and height information, passed down from aircraft transponders. [Not is use in oceanic airspace – ATC is supplied with position reports from the on-board navigation systems.]

Computer Processing – of radar and flight data.

High Quality Aircraft Navigation – progressively improved from VOR/DME to Inertial Navigation Systems through to satellite-based aids such GPS.

STCA – the computer processing system has the facility for analysing SSR tracks to predict if aircraft might come into close proximity in the near future and, if they do, warn the controller by flashing a message on his radar screen.

Airborne Collision Avoidance System ACAS – On board collision avoidance system based on detection of other aircraft in the vicinity carrying SSR transponders. These tell the pilot of nearby traffic – TA (Traffic Advisory) – and aircraft coming into conflict – RA (Resolution Advisory). RAs tell the pilot to climb or descend as appropriate to take it out of risk. [NB: ACAS is the generic term – used here. TCAS – Traffic Alert and Collision Avoidance System – is a commercially available version of ACAS.]

For present purposes, these can be grouped as follows:

Safe Route Design	**Formal Safety Rules** **Controlled Airspace**
Technical Infrastructure	**Radio Telephony** **Radar** **High Quality Aircraft Navigation** **Computer Processing** **Flight Progress Strips**
Planning	**Controller**
Flying	**Pilot**
Controlling	**Controller**
Ground Protection	**STCA** **Controller and Pilot**
Air Protection	**ACAS** **Pilot**

Figure 1. ATM System Components - illustrative

Figure 2 then shows, in a very abstract and simplified fashion, the dynamical transit of a typical flight – in ATM system terms – through these structural system layers. The three pre-operational – Planning – layers have been grouped together because they are highly related, eg separation minima depend on suitable equipment being available, while the controller has to work within the safety constraints using the equipment. The picture is very simplified: for example, flight scheduling and flow management should be included in Planning. The Operation Layers cover the activities of pilots and controllers while the flight is in progress. The Alert Layer is the ground and air protection enabled by STCA and ACAS, on which the controller/pilot will act. The assumption that a flight goes through these layers in one sequence is too simple; in reality, there are complex feedback loops.

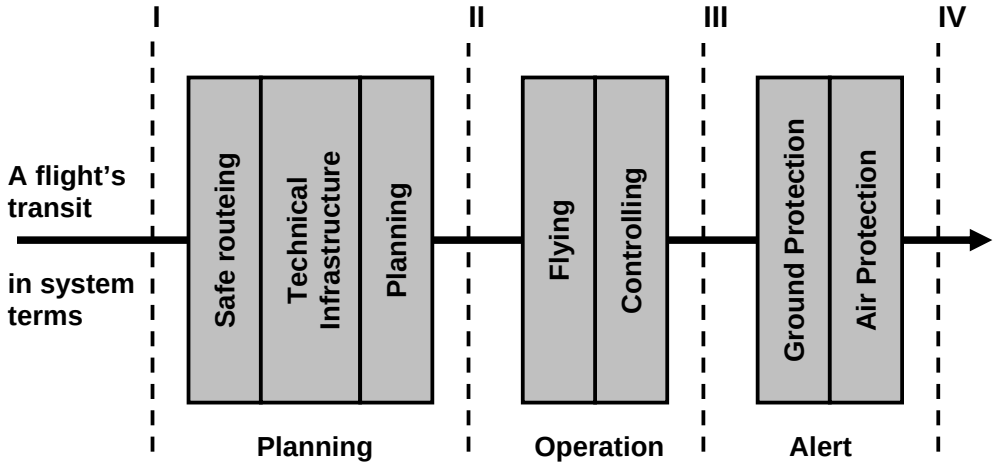

Figure 2. The ATM system layers – highly simplified, without 'loops'

But what do these layers actually accomplish in terms of system risks? The answer is that they act systematically to reduce mid-air collision risk by 'shuffling the risk pack'. The purpose of the system layers is to change the statistical distribution of risks. A formal structure is imposed by the Planning Layers, next the Operation Layer should then eliminate inherent conflicts (but note that the Planning Layer does not produce conflict free paths from departure to landing), and then the Alert Layer warns the controller and pilot about impending conflicts.

Figure 3 illustrates the risk structure – the frequencies of flights having particular 'degree of conflict' (D_c) values – of a large number of flights being processed by these layers, measured at two of the 'slice' points I to IV in the transit shown in Figure 2. [The term 'degree of conflict' is used because of similarities with related phrases adopted by researchers assessing STCA and ACAS.]

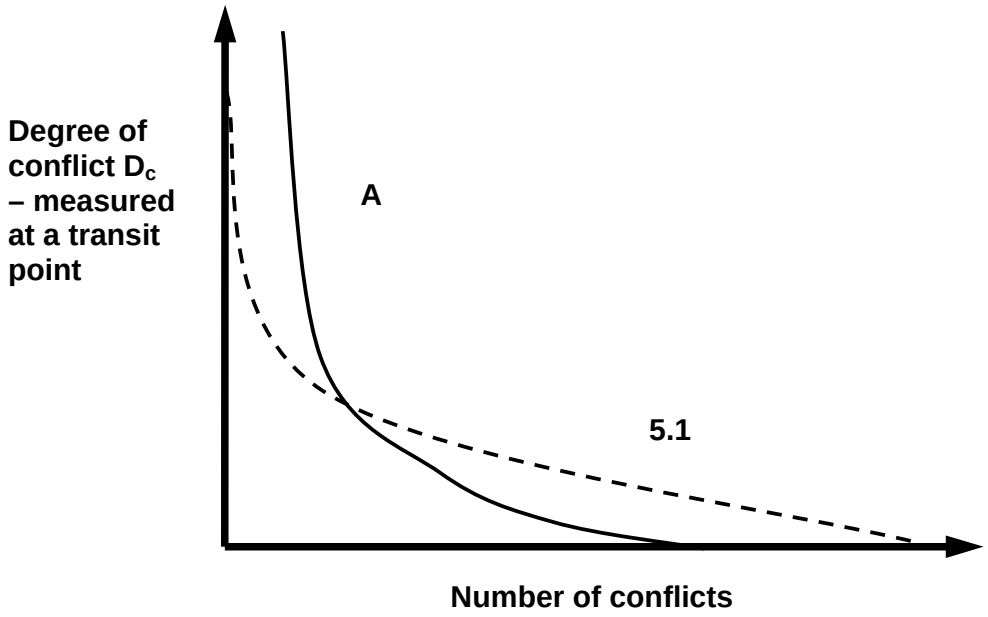

Figure 3. Statistical distribution of D_c in a large population of flights

The full line curve A shows a comparatively risky set of flights: some have very high values – off the scale – and the number of flights with zero D_c value is not that high. The dashed curve B is a much better profile: there is a maximum D_c value and most flights at near to zero D_c value. The job of the ATM system layers is to improve type A distributions so that they become type Bs: fewer high degrees of conflict and more of the lower degrees of conflict).

The 'starting population' of flights in Figure 3 could be all those that operated on a given day in a given airspace. But most of the pairs of flights fly at different times in different geographical areas. As these aircraft are too far apart in time and/or space ever to pose a risk to each other, they are excluded from the analysis – they have 'zero degree of conflict' (or most precisely, 'negligible' in any practical sense). The bulk of flights are to be found in the very long tail near to this 'zero degree of conflict'. (For a given traffic pattern, there is an upper limit to the number of conflicts: it is certainly bounded above – by all aircraft being instantaneously in conflict with all other aircraft.) Those significant conflicts that exist are comparatively few in number, and are generally confined to those flights operating at about the same time and in particularly heavily loaded airspace sectors.

On what kind of scale is this, so far vaguely described, degree of conflict D_c to be measured? It needs to indicate that the aircraft flightpaths, as known at the time the distribution is measured, will lead to the aircraft being in close proximity. It is a measure of hazard potential. Note that this measurement takes no account of the action of the remaining system layers. The choice of a scale for D_c is essentially arbitrary: Figure 3 shows a choice in which the full line curve is unbounded (eg for high values of D_c could correspond to the inverse of the projected closest point of approach (CPA) distance).

If the statistical distribution is rescaled by using something like a logarithmic transformation of the number of flights – so as to compress the horizontal scale (indicated by the zigzag on the horizontal axis) – then the action of the layers looks like Figure 4.

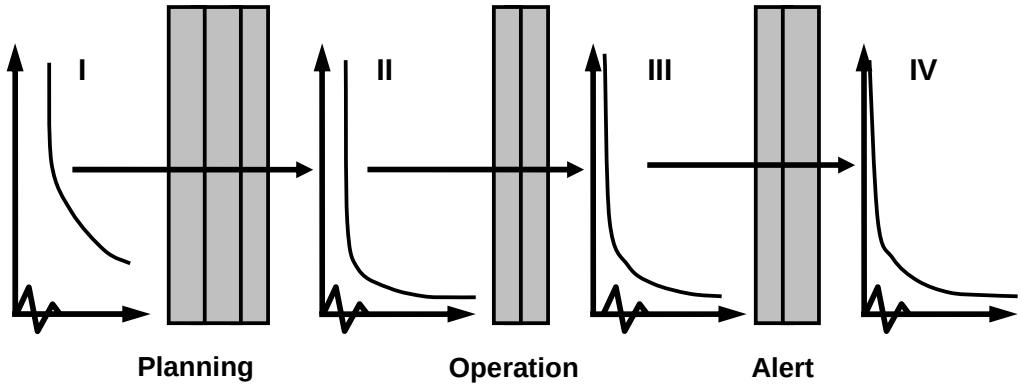

Figure 4. System layers changing the shape of the Conflict Distribution

Each of the Layers acts to reduce the frequency of the high conflict events – they transform the conflict pattern from the initial I to the final IV. A beneficial change then increases the rate of descent at the left and pushes the intercept on the $D_c = 0$ axis to the right.

How do the system layers work? The language of System Theory as applied to safety distinguishes between tightly versus loosely coupled systems (Weick, 1976; Perrow, 1984):

> The sub-components of a tightly coupled system have immediate impacts on each other. Tightly coupled systems can survive failures, if that failure has been anticipated and provided for in the original design. Thus, tightly coupled systems must therefore be designed to anticipate all the failure modes and providing safety features for continued operation and recovery. These kinds of designs can usually be modelled quantitatively, and their performance can be validated against what happens in the real world, eg by studying the kinds of accidents, incidents, failures and errors that occur.
>
> Loosely coupled systems have flexibility in the timing, nature or intensity of responses. They accommodate failures through adaptive responses. There is some 'play' in the (negative) feedback loops—a little over-correction, then some under correction. Such systems are adaptable and error tolerant, but can have long reaction times. Loosely coupled designs use much more complex information sources – eg through human visualization and situational awareness, so they tend to be open and continually interacting with the outside environment.

> In summary, in safety terms, loosely coupled systems can accommodate shocks, failures, and pressures for change without destabilization, while tightly coupled systems respond more rapidly to perturbations – but the response can be disastrous.

The ATM system layers have elements of both types of coupling. The Planning and Alert Layers tend to be more tightly coupled, because their action should be 'programmed'. The Operation Layer tends to be more loosely coupled, because pilots and controllers make strategic and tactical decisions. Their decision-making does however reflect their training, so they will tend to do similar things in similar situations and (eg) highly skilled pilots will make fewer strategic/decision errors (Wiegmann and Shappell, 2001).

However, each layer acts to transforms the distribution in a probabilistic fashion, rather than a deterministic one in which specific events necessarily follow particular causes. A particular aircraft configuration or routeing should be changed by the layers into a different one, which probably has a lower degree of conflict. But, on some occasions, the pilot or controller will choose to do something that does not improve the D_c value. A collision can occur if the system is in the hazardous state and the remaining system layers do not correct this. Thus, on a small proportion of cases, eg when an aircraft has just manoeuvred considerably and the aircraft pair is suddenly very close, ACAS might recommend an inappropriate course of action.

Note the difference between this kind of system metaphor and other kinds of system perspectives, such as those of Bird (1974), Rasmussen (1990) and Reason (1990). The approach here is foreshadowed in Rasmussen: "The causal tree found by an accident analysis is only a record of one past case, not a model of the involved relational structure…We should be fighting types, not individual tokens…" In Reason's descriptive models of system safety defences, the probabilistic aspects are modelled by the number and size of holes in defensive layers (eg Shappell and Wiegmann, 2001). The metaphor used here has the advantage that it shows clearly how the degree of conflict can reduce and increase during the sequence of an aircraft's transit through the ATM system. Human error, in the widest sense, remains a major element in this causal chain, as it does for aviation in general (Courteney and Newman, 2003).

Philosophical models of 'Probabilistic Causation' have a long history, eg Sosa and Tooley (1993). Examples of relevant aviation-related work are by Ladkin (2000), which provides a different 'logic-orientated' perspective, and by Greenwell et al (2003). There is a huge, multidisciplinary literature on this topic.

ATM layered concepts are the subject of research current interest. A very thought-provoking paper is by Graham et al (2003), which inter alia proposes the uses of a loop picture to explain the nature of information flow and decision-making between ATM layers. Graham et al trace back the idea of a layered approach to Villiers (1968). Interestingly, the French title of Villiers' paper uses the phrase "la méthode des filtres", thus giving filtering as a useful metaphor.

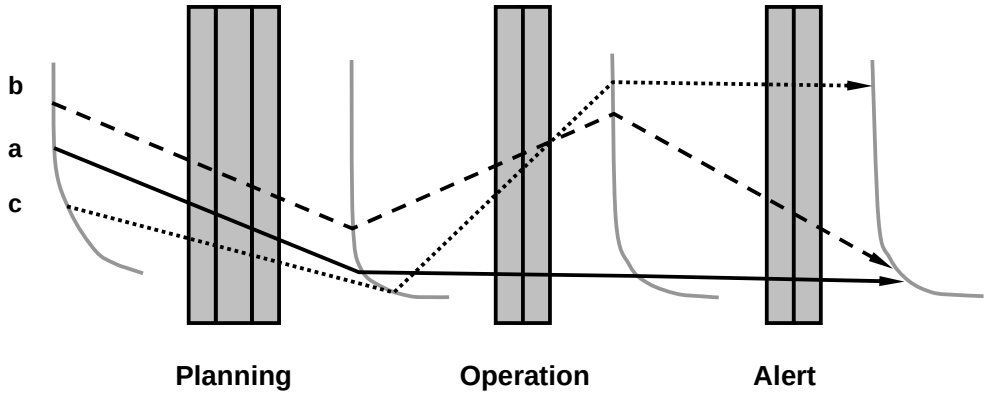

Figure 5. Examples of transition paths

Figure 5 shows a 'cut-down' version of Figure 4 with some examples of conflict level transformations:

a [full line] A comparatively high conflict level is reduced markedly by the planning Layers. The conflict level then hardly changes: the Operation Layers do not increase the D_c, the Alert Layers are not required, and so the end result is a low conflict level.

b [dashed line] This path shows a very high conflict level reduced to a moderate conflict level by the Planning Layers. Then some kind of Operation problem increases the D_c, sufficient the Alert Layer to come into action and reduce the conflict markedly.

c – [dotted line] In this transition path the low D_c at the initial stage is reduced even further by the Planning Layers. But the Operation Layers then increase the conflict level considerably. The Alert Layers turn out to be ineffective, so the D_c is very high.

The key point is that the D_c value can rise and fall through the sequence of system layers. There are obviously many different possibilities, eg these transition paths do not show a case in which the Planning Layers increase the risk. These could be very significant, given that the route and airspace structure is probably more likely to increase the number of aircraft on conflicting paths. This also serves to highlights an example of interaction between the layers. The main purpose of route structures for ATM could be considered to be helping controllers to identify conflicts (ie by forcing crossing points at particular locations etc), and hence improving the effectiveness of the Operational layer.

It is one thing describing such a probabilistic process as a metaphor, but this needs to be turned into something that is quantitative. For present purposes, all that is required is a simple categorisation of 'errors' made by the system layers. These are of two types so, using analogous language to statistical hypothesis testing:

Type 1 error – failing to reduce the D_c when it is high (α)

Type 2 error – increasing the D_c when it is low (β)

Next, it is necessary to categorise the D_c distributions in two bands: hazard potential and safe:

> The 'hazard potential' band comprises flights for which there is a high degree of risk.
>
> Aircraft in the 'safe' band are not in any risk now – but could be if the further safety layers did not act appropriately.

The proportions of hazard potential and safe flights are h and s, with their sum adding to unity. (Obviously, a much more complex banding is possible, with h_1, h_2, h_3…representing a finer-structured grouping, but for present purposes the two bands are sufficient.) Thus, probabilities α and β are respectively the proportion of situations with hazard potential that remain in this state and the proportion of safe instances that move to having hazard potential.

As an aside, it is possible to view the system layer process as a Markov Chain (eg Cox and Miller, 1977) or as a Bayesian network (eg Neil et al, 2003). These offer new ways of thinking – in rational and quantitative terms – about ATM safety and the actions of the system layers. This kind of analytical approach (compare Wiegmann and Shappell, 2003) may be more productive for modelling than (eg) Reason's descriptive ideas.

4. AIRPROXES, HAZARDS AND SYSTEM LAYERS

The material of the previous two sections can now be used to analyse Airproxes. . Airproxes are chosen here because they are publicly available and well documented – Airprox statistics are often publicised as a 'gold standard' of UK ATM safety. It must be stressed that other safety incident data – eg the CAA's MORS data, and NATS SMF data and other databases would no doubt be at least equally valid ATM safety incident sources.

Airprox Category	Description
CAT A - Risk of Collision:	The risk classification of an aircraft proximity in which serious risk of collision has existed.
CAT B - Safety Not Assured:	The risk classification of an aircraft proximity in which the safety of the aircraft may have been compromised.
CAT C - No risk of Collision:	The risk classification of an aircraft proximity in which no risk of collision has existed.
Risk Not determined:	The risk classification of an aircraft proximity in which insufficient information was available to determine the risk involved or inconclusive or conflicting evidence precluded such determination.

Figure 6. ICAO Doc 4444 AIRPROX Severity Scheme (eg Eurocontrol, 2002)

Figure 6 shows the internal (ICAO) definitions for Airprox categories (eg Eurocontrol SRC, 2002). Note the use of the word 'existed', ie at some point there was a perception/judgement of hazard. Next, Figure 7 illustrates a generic Airprox.

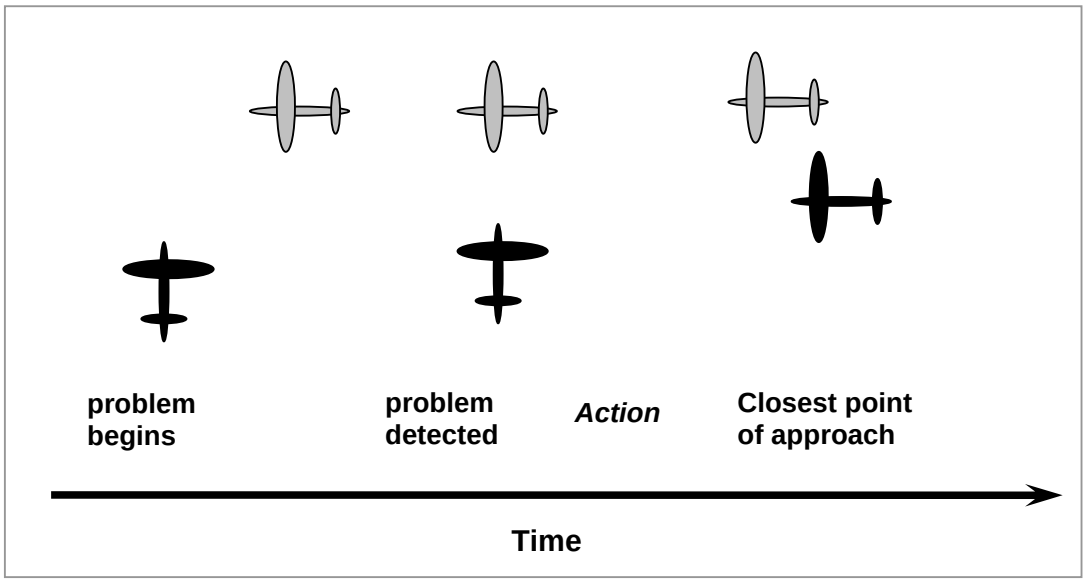

Figure 7. Sequence for a possible Airprox – see text for explanation

Figure 7 illustrates the time sequence for a possible Airprox – the black aircraft might collide with the grey one: three 'snapshots' in time are shown. A problem of some kind is the precursor of a possible Airprox – it could be a miscommunication and a failure to note a relevant piece of information. This problem could lead to a mid-air collision within a few minutes if nothing further is done. At some point, this problem begins, but this is probably not immediately detected by the aircrew or ATC. Later, potential consequences of the problem are detected: this could be at any point from the initiation of the problem to an automatic alert to the pilot/controller. Action then takes place. This action might be a new instruction to the pilot, a manoeuvre by the pilot, or just close monitoring by ATC (if the aircrafts' projected flightpaths do not in fact lead to close proximity, then new instructions would not be necessary – or even desirable). Finally, the two aircraft reach their CPA. This whole process is an Airprox.

Note that the reporting of such an occurrence as a potential Airprox need not be a statement by the reportee that the situation was 'hazardous'. It need only be a view by the reportee that 'this kind of aircraft configuration should not have occurred if the rules are followed properly'. To be hazardous in the terms used here, the occurrence must satisfy:

> On the information available and without further action/intervention, the flightpaths, allowing for typical variabilities, would produce a close proximity between the aircraft; and
>
> In these circumstances, the remaining system layers would not prevent the dangerous CPA.

Expressed in a different way:

> A hazardous event is one in which there is a high degree of conflict plus a low confidence that the remaining system layers would generally provide the necessary corrective action.

One test is to judge the probability that the remaining system layers would deal with the conflict in a straightforward fashion, in other words estimate the effects of the remaining α and β parameters. What is the probability that the remaining system layers will fail to resolve this high degree of conflict? Turning the question around, if it had not been detected, is it highly probable that the system layers would have converted to 'safe'? The expert judgement/estimate that has to be made about an ATM incident (such as an Airprox) is of the chance that the remaining system layers will not resolve the situation safely. This assessment might be made using something like the ARIBA (1999b) approach (which attempted to build an accident risk tolerability/acceptability matrix for air traffic operations on HSE (UK Health and Safety Executive, 1992 and 1999) lines).

5. DISCUSSION

This discussion section consists of seven largely independent sub-sections, mainly exploring some of the questions raised in the Introduction.

5.1 Real-Life Analyses

5.2 Which incidents should be judged the most important to ATM system safety?

5.3 Are 'minor' incidents of any importance?

5.4 The Value of the System Layer Concept

5.5 Examples of Airproxes

5.6 The Right Lessons from Mid-air Collisions?

5.7 NATS Safety Significant Events

5.1 Real-Life Analyses

What are really the questions about ATM incidents that need to be answered by safety managers and regulators? What kinds of action should they take? How can these people demonstrate that they have done their job properly? Certainly, one way of knowing what one should do is to contemplate the consequences of not doing certain things. This implies that they must act reasonably and rationally – but they must be active rather than passive (ie ATM incidents must be analysed from the viewpoint of 'looking for trouble').

However, the public's and air traveller's concerns must surely be focused on the safety levels achieved in the real world, rather than on speculations. Why should they care about 'what might have been in some theoretical alternative universe' – composed of 'what ifs' – in which things had been 'slightly different'? So is it valid or sensible to examine these 'what ifs'?

The answer is that it is necessary for ATM incident analysts to think in this kind of way. This is not some nugatory or over-sophisticated 'icing on the cake'. The mark of the genuine safety expert is to be able to ask the right questions about incidents concerning potential accidents. The messages from an incident do not simply 'announce themselves' to the analyst: they provide the raw material for his or her productive thoughts about system safety.

When aviation systems fail, they often do so fail in apparently complex and unpredictable ways (Amalberti, 2000; Wickens, 2003): the aviation expert is someone who perceives the fundamental causal and system factors. Such people are rare individuals – the kind of structured thinking outlined here is intended to supplement their expertise.

[If an expert had been able to predict the failure mode of the Comet aircraft – that metal fatigue concentrated at the corners of the aircraft's windows would probably cause catastrophic crashes (three in 1953) – then the UK's commercial aviation industry would have followed a very different path. One aspect of the tragedy that is relevant here was the absence of an earlier related minor incident that could have warned the designers of this kind of failure mode. How much worse would a catastrophe be viewed by relatives and the public if it could have reasonably been prevented by learning from an earlier incident's characteristics?]

One way of recognising the need to examine 'what ifs', is to envisage an ATM system covering a large area. For simplicity, assume a constant amount of traffic every year, that the underlying 'safety culture' and regulatory framework is unchanging. The nature of ATM incidents in such a scenario will vary in an infinite number of ways: there will be changes in aircraft types and their flight departure times, runway usage, meteorological conditions, choices made by controllers, etc, will lead to different kinds of incidents occurring at different times from year to year. However, the 'pattern' of these incidents – the causal nature of these incidents and the action of the system layers – will be much the same from year to year; the observed D_c values are effectively a 'sample' from this much larger population of potentially feasible conflict patterns.

The expert is the person who can comprehend the underlying patterns revealed by incidents. He or she can then rationally extrapolate those observed over many years to identify the most probable accident in a coming year.

5.2 Which incidents should be judged the most important to ATM system safety?

The most important incidents to ATM system safety are surely those in which only the Alert layer prevents a collision, because these represent failures of the previous system layers. Was the incident reported subsequent to an Alert? If so, the protection against mid-collision could have depended on the chance relation between the aircraft trajectories relied – the lack of such an event could have been wholly fortuitous. The analyst could therefore ask: "If one of the aircraft departure times had been (say) up to 30 seconds different, would there have been a collision?"

There has been important work on the performance of the Alert layer (ie which relate to its α and β parameters). As part of ICAO panel studies, Harrison (1993), reported calculations and simulation results that:

> For every 100 critical Airproxes and under a variety of assumptions about traffic awareness and pilot responses (eg neglecting the use that a pilot might make of TA information prior to an RA – a very cautious assumption):
>
> > 94 will be resolved safely
> >
> > 6 will not be resolved
> >
> > 4 new Airproxes will be induced because non-critical encounters are converted into critical ones

McLaughlin (1999) has produced some current estimates on the benefits of ACAS.

Hale and Law (1989) provides an examination Simultaneous Operation of Conflict Alert and ACAS II in UK En-Route Airspace.

5.3 Are 'minor' incidents of any importance?

ATM incidents do not need to be deemed hazardous to be informative or to generate action. The decision-maker needs to have in mind something like the HSE's ALARP philosophy (Appendix B; HSE, 1992/1999). Given that the system is very safe, specific safety management measures – changes to procedures, equipment, software – should be implemented as long as such is reasonably practicable. This generally implies some kind of analysis of costs and benefits. To give an example (edited text from NATS, 2004):

> "In November 2002, an Airprox involving a Virgin 747 and a Delta 767, in which track data blocks for the two aircraft were inadvertently swapped on-screen, which was subsequently assessed by the UKAB as category C (no risk of collision). Immediately after this particular incident, an instruction was issued to controllers reminding them of the correct procedure to follow when individual track data blocks are re-positioned, in order to prevent a repeat of these events. In 2004, NATS changed its software, further improving the legibility of track data blocks. The change ensures that, whenever a data block is moved, it will be automatically linked (by way of a strut on the screen) to the aircraft target to which it belongs."

Thus, key questions are: “Is there evidence of a systematic design flaw? Can it be corrected without disproportionate cost? Are international agreements an issue?”

5.4 The Value of the System Layer Concept

The safety of the system is the product of the effectiveness of the system layers. The questions are obvious ones. Did they operate as they were planned? Did a failure in the earlier system layers produce an aircraft pair configuration and/or circumstances that might not generally be corrected by the remaining layers?

If the incident is a consequence of a known category of system layer failure, is the frequency of this type of failure (= the α value) increasing over time? If so, what is being done to reduce the rate by ‘tweaking’ the system layers to improve their capture of such failures? Actions/decisions include such things as additional training for pilots or controllers, patches to airspace designs, tailoring of STCA/ACAS, etc.

It is not sufficient to focus on the extent to which separation minima are breached, or the nature of actions by the pilot/controller. These are vital issues, but they must be seen in the context of the system control provided by the system layers. A low CPA may represent little hazard if, in those particular circumstances, both pilots could see the other aircraft visually and on their ACAS displays. A larger CPA for an event not detected by either pilots or controller, and subject to a very late ACAS alert because of manoeuvring aircraft, is a much more serious matter in system safety terms. The pilot/controller may have taken no action because there was insufficient time for it to be effective of because they judged that it might make a bad situation even worse. An STCA alert and/or a pilot/controller action represent the system controls – the system layers – functioning effectively, rather than a necessarily hazardous situation.

It almost goes without saying that monitoring the frequency of failures – the α and β parameters – is vitally important. For example, the ATM provider might monitor ‘Safety Separation Breaches’, composed of incidents in which separation was significantly breached (rather than a minor infringement by a fraction of a nm) or in which the controller had to act on a STCA alert (Brooker, 2004).

Another way in which this kind of model is useful is to consider how the different layers vary in different types of ATM operation. For instance, in Oceanic operations the operational layer is comparatively weak compared with radar control, but the planning layer is generally very effective in reducing the degree of conflict – but see the next sub-section.

5.5 Examples of Airproxes

It is worth examining half a dozen examples of Airproxes to back up the arguments of the preceding text. None of these incidents was rated as deserving the (highest) Airprox category A. But to what extent do they provide important lessons about potential accidents? Some of these particular incidents were chosen because the UKAB text used or implied the word ‘fortuitous’. The following text – Figure 8 – is of

course heavily summarised from the original material [NB: 127/99 is incident number 127 in 1999, etc and the aircraft are A and B]:

Airprox 127/99	Recorded Separation was 1100 ft	
Context & Issues	**Warnings**	**1.7.1.1.1.1 Event, Monitoring & Intervention**
Extremely busy period, issues about display of data blocks	Outside STCA parameters. ACAS RAs to both aircraft	Controller had issued a descent clearance that would have led aircraft A to descend through the level of aircraft B, which he had inadvertently not taken into account.

Airprox 29/00	Recorded Separation at CPA was 0·6 NM and 600 ft.	
Context & Issues	**Warnings**	**Event, Monitoring & Intervention**
Flight deck procedure error, fatigue	ACAS RA both to A and B. STCA activated before ACAS passed to ATC	Aircraft descended below its cleared level. Neither pilot saw the other aircraft. ATC occupied with other traffic, did not spot high descent rate. Board agreed that any separation was to a large degree fortuitous…with different geometry …could have been considerably more serious.

Airprox 54/00	Recorded Separation at CPA was 12 NM	
Context & Issues	**Warnings**	**Event, Monitoring & Intervention**
In oceanic airspace, period of data processing (FDPS) manual reversion	No warnings	Neither pilot saw the other aircraft. …the SCACC Domestic Controller…detect[ed] the conflict from a routine scanning of her radar display.…an ATCO member said that, in his opinion, separation in this case was fortuitous. had the encounter taken place further west, outside the cover of domestic radar, the confliction would have remained undetected and the outcome might have been more serious…

Airprox 145/00	Recorded Separation was 0 45 NM and 400 ft	
Context & Issues	**Warnings**	**Event, Monitoring & Intervention**
Previous military fighter formation, severe weather, ATC preoccupied with other conflict, workload	ACAS RA to A only. ATC at airport not STCA equipped. LATCC STCA alert	ATC did not maintain standard separation between aircraft A and B. Pilot A, because of weather, saw B only after the ACAS alert. Pilot B heard the ACAS alert on RT and saw A subsequently. ATC was not aware of conflict until advised of ACAS alert.

Airprox 164/03	Recorded Separation was 3.7 NM and 500 ft	
Context & Issues	**Warnings**	**1.7.1.1.1.2 Event, Monitoring & Intervention**
No apparent ATC causal factors, non-standard phraseology	STCA alerted after the aircraft had received an ACAS RA.	Aircraft A crew descended below their cleared level into conflict with aircraft B.

Airprox 184/03	Recorded Separation was 3.4 NM and 600 ft	
Context & Issues	**Warnings**	**1.7.1.1.1.3 Event, Monitoring & Intervention**

Controller with an inexperienced trainee, combined sector – 'busy' but within capabilities'.	STCA activated and shortly afterwards an ACAS RA climb was issued.	The aircraft B crew read back the wrong heading and level instructions, which went undetected by the controller. The controller said he had no reason to doubt that the aircraft would not comply with the issued clearance.

Figure 8. Edited text from Airprox Reports

Each of these incidents exhibits the potential for a more serious incident or even an accident. Most of them show the ATM system getting into an operational state in which recovery was not assured by subsequent system defensive processes, or where the failure/error occurred at such a late stage that the remaining safety defensive layers had a large element of chance. Each of them should therefore be marked as having more safety significance than incidents in which (say) STCA gave a very early warning and the controller was able to resolve the problem before any kind of ACAS warning was given.

Several of the incidents show well-known Human Factors aspects. Airprox 127/99 is an example of a failure of prospective memory (eg Loft et al, 2003). Airprox 29/00 is an example of an aircrew procedure error with fatigue as a factor. Airprox 145/00 is a loss of the necessary situation awareness (eg Endsley et al, 2003), given a complex set of circumstances. Communications procedures are well known to be a source of potential problems, eg Airprox 184/03.

The oceanic incident, Airprox 54/00 is perhaps an example of a (currently?) inherent ATM system design limitation. The North Atlantic region ATM system does not have radar coverage, hence it does not have STCA available. The necessary safety is delivered by the use of large, essentially procedural separation minima, Mach Number flying techniques, and positional reports every 10 degrees of longitude. ACAS is a critical safety defensive layer.

Some defensive system layers do not function for certain types of manoeuvre. If the aircrew or the controller mistakenly climb or descend an aircraft, then a potential conflict may not be picked up in time for STCA to alert, so the hazard is reduced by a combination of ACAS and the pilot's see-and-avoid, eg Airproxes 164/2003 and 127/99.

Equipment failures and mal-functions are comparatively rare these days compared with Human Factors issues (in the broadest sense). Airprox 54/00 notes a case of FDPS manual reversion, which could affect considerably the oceanic safety layers. Airprox 222/02, referred to in the previous sub-section, arose because of a confusing display of overlapping track data blocks and aircraft symbols. This led the controller to confuse the relative positions of the two aircraft so that one was descended into conflict with the other.
A caveat needs to be entered about the degree to which the safety system layers can be fully understood solely from this kind of incident data. In particular, incidents rarely give much insight into the working of the Planning layer. The Planning layer is not designed to eliminate all conflicts (except, as already noted, in procedural control in oceanic systems), so the fact that aircraft may be in conflict is not an incident until

the Operational layer fails to detect and resolve the conflict. This implies the need to collect other precursor information. How many conflicts of different types are being solved routinely by the existing safety system?

5.6 The Right Lessons from Mid-air Collisions?

European mid-air collisions in controlled airspace are rare. Their characteristics can be analysed in exactly the same way to incidents. On 1 July 2002, two ACAS-equipped aircraft collided over the Swiss-German border at Überlingen. One important feature of the accident was that the flight crew of one aircraft did not follow the ACAS alert, but followed instead the ATC instruction. Guidance material now stresses that pilots should follow alerts and that controllers should not attempt to modify the flight path of an aircraft responding to an alert.

The full official report on this tragedy has recently been published (BFU, 2004). There have already been attempts to analyse the causal factors involved, eg Nunes and Laursen (2004) identify six 'contributing factors': Single Man Operations, Downgraded Radar [STCA], Dual Frequency Responsibility, Phone System, ACAS, Corporate Culture.

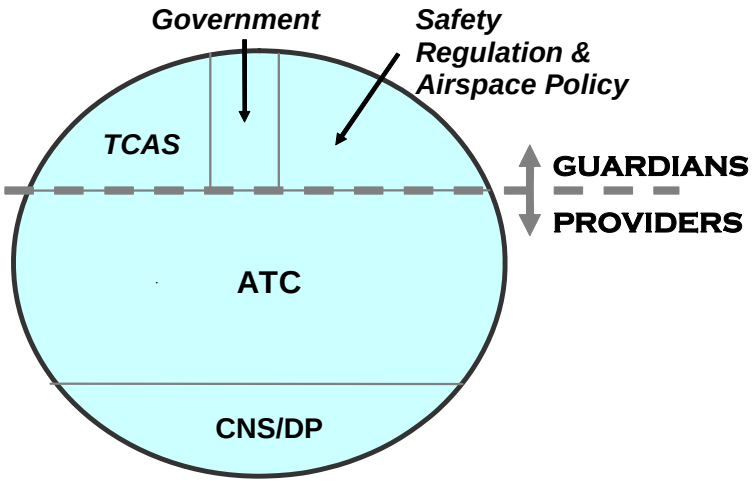

Figure 9. Components of ATM System safety

One clear message from Überlingen is that it is not only the providers who must take responsibility for preventing accidents. In Figure 9, ATC is air traffic control, CNS is the communications, navigation and surveillance systems, DP is all the data processing and information flows. These can be termed system 'Guardians', as distinct from the various service providers. The Government sets up the regulatory regime; the airspace policy has to recognise the needs to accommodate all users: commercial flights, military flights and general aviation. The safety regulator has to be confident that everything works safely: "To require enterprises to take proper account of the hazards to which they expose people". The regulator cannot always believe what it is told by or reads in documents by the service provider. The

regulator's job does not end there, for example, it must ensure that the system delivers the right kinds of safety-related training.

This multiplicity of factors indicates the unusual nature of the Überlingen accident. A picture of the mid-air collision on the lines of Figure 5 would show a very high conflict level at every point in the diagram, with the final failure to act correctly on the ACAS warning raising the conflict level to the actual collision event. But if the accident had not occurred – if the aircraft had, by providence, missed by several hundred feet – would the same kinds of international actions have been taken by regulators and providers? These questions just make an even stronger case for international learning from the collection and thorough analysis of incidents. This learning has to identify those incidents that reveal inherent system control flaws or regulatory gaps/inconsistencies and offer ways of dealing with them.

5.7 NATS Safety Significant Events

Airprox classifications represent an example of what might be termed a 'Board of Inquiry' perspective, exemplified in UK aviation by the work of the AAIB (2004), which examines accidents and incidents. The AAIB's remit is stated as:

> "The fundamental purpose of investigating accidents is to determine the circumstances and causes of the accident…It is not to apportion blame or liability"

This appears straightforward – but the hidden text is:

> "with a view to the preservation of life and the avoidance of accidents in the future",

which actually coincides with a central message here about asking rational and professional 'what if' questions.

Safety managers and analysts working for ATC providers do attempt this kind of thing. An example is NATS' work on the Safety Significant Event (SSE) scheme, which has been under development since the mid-1990s [NB: the author participated in the early work]. Eurocontrol has recently given support to this kind of approach for categorising incidents (eg Eurocontrol, 2004).

SSEs in a radar control environment are defined relative to bands:

***Band 1**: Separation ≤ 66% of the prescribed separation.*

***Band 2**: Separation > 66% of the prescribed separation.*

SSEP: A Band 2 incident, which involved aircraft losing separation or potentially losing separation with another aircraft where there was a possible ATC error.

SSE4: A Band 1 incident which was detected and resolved in the most effective and timely manner by the controller who was providing the service when the incident occurred, and no systems failures or procedures affected the resolution.

SSE3: A Band 1 event which was detected and resolved by ATC but:

- it was not resolved by the controller who was providing the service when the event

was initiated; or

o it was detected by colleague warning, STCA or pilot query; or

o it was not resolved in either a timely manner; or

o it was not resolved in an effective manner; or

o systems or procedures failures affected the resolution.

SSE2: A Band 1 event, which was resolved by the pilot/other.

SSE1: A Band 1 event for which no timely/effective pilot/other action was taken to resolve the event (providence) or there was a high risk that the action taken would not have been successful.

Figure 10. SSE classification re 'Loss of separation in a radar control environment'

Figure 10 shows how incident data on breaches of separation is classified according to the SSE scheme. The similarity to the kinds of ideas explored here is in the use of words such as 'providence' and 'action taken would not have been successful'.

Recent discussions by the author with NATS experts suggest that the SSE scheme has been judged very successful in safety management terms, in particular that it has helped the organisation to focus on key safety issues. Unfortunately, there has been little published by NATS in the safety literature on how the scheme is used. Searches on the web and on academic electronic databases (in particular the British Library's Electronic Table of Contents 'zetoc') reveal very few open-literature technical documents about the use of SSEs. One that is very much worth mentioning is the work by Neil et al (2003), which addresses the same kind of problem as that examined here.

6. CONCLUSIONS

The aim is to try to ensure that analysis and expert judgement about ATM incidents can be carried out within a systematic and consistent safety framework. Hazards and risks are not 'facts' or 'events' that 'exist', but rather judgements made about conditional futures and their consequences. This judgement (or perception or opinion) is about the degree of possibility of some unpleasant state of things that may come into existence at some future time. A hazardous situation is not one in which aircraft happened to be close, but rather one in which the outcome was not 'system controlled', with some potential outcomes having significant negative consequences. System controls in this sense cover all the means by which the system is held stable (= defended) against the potential negative consequences.

The ATM system can be (over-) simplified to consist of three structural system layers acting as the system controls: Planning (pre-operational), Operation (the flight in progress), and Alert (the ground and air protection enabled by STCA and ACAS, on which the controller/pilot will act). Each Layer acts to reduce the frequency of high conflict events in a probabilistic fashion. ATM safety improvements correspond to monitoring and acting upon two probabilities α and β for each system layer: these are respectively the proportion of hazardous situations that remain hazardous and the proportion of safe instances that become hazardous.

A hazardous event is one in which a high degree of conflict between aircraft is observed plus a low confidence that the remaining system layers would generally provide the necessary corrective action. The expert judgement/estimate that has to be made about the hazardousness of an ATM incident (such as an Airprox) is of the likelihood that the remaining system layers will not resolve the situation safely.

It is necessary for ATM incident analysts to think in this kind of way – and indeed for their managers to ensure that they have enough time to think. It is not enough just for the analysts to be intelligent, knowledgeable and energetic. Just 'keeping to the facts' could fail to anticipate the future. The mark of the genuine safety expert is to be able to ask the right questions concerning potential accidents – to show imagination and system insight. The messages from an incident do not simply 'announce themselves' to the analyst. However, they do provide the raw material for productive thoughts and wise judgements about system safety. Thus, ATM safety analysts need first to get data to determine the general effectiveness of the different safety layers, and then to be able to determine the differences in layer performance for different varieties of event.

Returning to the kinds of specific safety questions that can be asked about ATM incidents, some answers can be attempted:

> *Which incidents should be judged the most important to ATM system safety?* The most important ones are those in which only the final stages of the Alert layer prevented a collision, because this represents a failure of the previous system layers. The protection against mid-collision could have depended on the chance – fortuitous – relation between the aircraft trajectories relied. "If one of the aircraft departure times had been (say) up to 30 seconds different, would there have been a collision?" Would there have been enough time to analyse, decide and act successfully?
>
> *Which incidents give most guidance about potential accidents?* The crucial incidents are those in which the system got into a state in which recovery was not assured by subsequent system defensive processes, or where the failure/error occurred at such a late stage that the remaining safety defensive layers had a large element of chance. Was an accident prevented by design or by chance? Were, for example, the system defensive layers there (but did not function to specification), weakened (through other elements in the system environment), or absent (perhaps because of regulatory flaws/gaps)?
>
> *In what ways should incidents be categorised and analysed to help pinpoint key safety issues?* Additional categorisation schemes probably do not add much to safety improvement! Describing, in ever greater detail, the symptoms of a patient's illness is much less important than finding a cure for that illness. The important thing is to highlight flaws in system controllability by using something like the transition path picture; and to identify from this what might be possible solutions (or to recognise openly the inherent limits of present ATM system design, technology and operation).

Are 'minor' incidents of any safety importance? Yes. Incidents that are minor, in terms of the 'Board of Inquiry view' of the actual event, can be very informative or even generate decisive action. Is there evidence of a systematic design flaw? Can this flaw be corrected without disproportionate cost? Are international agreements an issue? Does the incident reveal readily correctable flaws in regulatory instructions or training?

How should the relevant importance of different incidents be assessed or weighted to provide a true picture of the health of the ATM safety system? The vital need is international learning from incident reports. ATM systems in developed countries use much the same equipment and operating concepts, so that 'ATM health' is (at least) that of the European system. This learning has to identify those incidents that reveal inherent system control flaws or regulatory gaps/inconsistencies – and offer ways of dealing with them. This safety evolution has to be an international process, not a national one. Safety experts must use all the information they can get to improve the ATM system, which strongly supports the need for international openness about safety data and its analysis.

The key safety management question, for both Providers and Guardians, to bear in mind is: "If Überlingen had been a severe incident not an accident, would the same safety lessons have been learned or pursued so vigorously? Would it have been put to one side as a 'unique event'?"

ACKNOWLEDGEMENTS

This work was in part supported by a research grant by the Civil Aviation Authority's Safety Regulation Group (SRG). I would like to thank SRG staff; Ian Parker, the Head of NATS Safety Management Development for updating me on work in NATS; and both he and Mike Shorthose of Helios Technology Ltd for their comments on earlier drafts. I would like to thank the Director of the UKAB and his colleagues for their help. I would also like to thank the referees for their insightful, indeed motivational, comments.

REFERENCES

AAIB [UK Air Accidents Investigation Branch} (2004). Website. http://www.dft.gov.uk/stellent/groups/dft_control/documents/contentservertemplate/dft_index.hcst?n=5161&l=1

Amalberti, R. (2000). The paradoxes of almost totally safe transportation systems. Safety Science, 1, 1-16.

ARIBA [ATM system safety criticality Raises Issues in Balancing Actors responsibility] (1999a). WP4 Final Report: Human Operators Controllability of ATM Safety. European Commission Project ARIBA/NLR/WP4/FR. http://www.nlr.nl/public/hosted-sites/ariba/rapport4/index.htm

ARIBA (1999b) WP6 Final Report Part II: Safety Cases for a new ATM operation. European Commission Project ARIBA/NLR/WP6/FR-II. http://www.nlr.nl/public/hosted-sites/ariba/rapport6/part2/title.htm.

Baumgartner, M. (2003). One safe sky for Europe – A revolution in European ATM. The Controller, July, 8-12.

BFU [German Federal Bureau of Aircraft Accidents Investigation] (2004). Investigation Report 'Überlingen Mid-air collision'. AX001-1-2/02. http://www.bfu-web.de/berichte/02_ax001efr.pdf

Bird, F. (1974). Management Guide to Loss Control. Institute Press, Atlanta, Georgia.

Brooker, P. (2002). Future Air Traffic Management: Quantitative En Route Safety Assessment Part 2 – New Approaches. Journal of the Institute of Navigation 55(3), 363-379.

Brooker, P. (2004). Why the Eurocontrol Safety Regulation Commission Policy on Safety Nets and Risk Assessment is Wrong. Journal of the Institute of Navigation 57(2), 231-243. (in press).

CAA [Civil Aviation Authority] (2003). The Mandatory Occurrence Reporting Scheme. CAP 382. CAA, London. http://www.caa.co.uk/docs/33/CAP382.pdf

Chambers (2001). Chambers 21st. Century Dictionary. Chambers Harrap, Edinburgh. http://www.chambersharrap.co.uk/chambers/chref/chref.py/main?query=hazard&title=21st.

Courteney, H. and Newman, T. (2003). Taming Human Error with a Systems Approach. FSF/IASS Conference, Washington, USA.

Cox, D. R. and Miller, H. D. (1977). The Theory of Stochastic Processes. Chapman & Hall, London.

Endsley, M.R., Bolte, B. and Jones, D.G. (2003). Designing for Situation Awareness. Taylor and Francis, London.

Eurocontrol SRC (1999). ESARR 2 Guidance to ATM Safety Regulators: Severity Classification Scheme for Safety Occurrences in ATM. EAM 2/GUI 1. Eurocontrol, Brussels. http://www.eurocontrol.int/src/documents/deliverables/esarr2_awareness_package/eam2gui1e10ri.pdf

Eurocontrol SRC (2002). ESARR 2 Mapping between the Eurocontrol Severity Classification Scheme & the ICAO Airprox Severity Scheme. EAM 2/GUI 3. Eurocontrol, Brussels.

Eurocontrol (2004). Integra Safety Metrics. http://www.eurocontrol.int/care/integra/safety_metric.htm

Graham, R., Hoffman, E., Pusch, C. and Zeghal K. (2003). Absolute versus Relative Navigation: Theoretical Considerations from an ATM Perspective. 5th Eurocontrol/FAA ATM R&D Seminar, Budapest, Hungary

Greenwell, W. S., Knight, J. C. and Strunk, E. A. (2003). Risk-based Classification of Incidents. Second Workshop on the Investigation and Reporting of Incidents and Accidents (IRIA 2003). http://shemesh.larc.nasa.gov/iria03/p03-greenwell.pdf

Hale, S. and Law, M. (1989). Simultaneous Operation of Conflict Alert and ACAS II in UK En-Route Airspace. DORA Report 8914, CAA.

Harrison, D. (1993). Results of ACAS II Safety Analysis. ICAO Secondary Surveillance Radar Improvements and Collision Avoidance Systems Panel (SICASP/5)

HSE [Health and Safety Executive] (1992 and 1999). The Tolerability of Risk from Nuclear Power Stations. HMSO; Reducing Risks, Protecting People. HSE Books.

JAA [Joint Airworthiness Authorities] (2000). Advisory Joint Material relating to JAR 25 Large Aeroplanes. AMJ 25.1309. Change 15. Joint Airworthiness Authorities.

Ladkin, P. B. (1998). Notes on the Foundations of System Safety and Risk. RVS-Bk-00-01. RVS Group, Faculty of Technology, University of Bielefeld. http://www.rvs.uni-bielefeld.de/publications/books/safetyNotes.pdf.

Ladkin, P. B. (2000). Causal Analysis of Aircraft Accidents Computer Safety, Reliability and Security, Proceedings of the 19th International Conference, SAFECOMP 2000, Lecture Notes in Computer Science No. 1943, Springer-Verlag, 2000. At: http://www.rvs.uni-bielefeld.de/cms/publications/records

Loft, S., Humphreys, M. and Neal, A. (2004). Prospective memory in air traffic control. In, Edkins, G. and Pfister, P. (Eds.), Innovation and consolidation in aviation. Aldershot, UK Ashgate.

McLaughlin, M. (1999). Predicting the Effect of TCAS II on Safety. Air Traffic Control Quarterly, 7(1), 1-18.

National Air Traffic Services Ltd [NATS] (2004). Safety Website. http://www.nats.co.uk/about/safety.html

NATS (2004). Press notice, NATS responds to media comments on 2002 airprox. http://www.nats.co.uk/news/index.html

Neil, M., Malcolm, B. and Shaw, R. (2003). Modelling an Air Traffic Control Environment Using Bayesian Belief Networks. 21st International System Safety Conference, Ottawa, Ontario, Canada.

Nunes, A. & Laursen, T. (2004). Identifying the factors that led to the Ueberlingen mid-air collision: implications for overall system safety. Proceedings of the 48th Annual48th Annual Chapter Meeting of the Human Factors and Ergonomics Society, New Orleans, LA, USA. http://www.aviation.uiuc.edu/UnitsHFD/conference/humfac04/nuneslaur.pdf.

Perrow, C. (1984). Normal Accidents: Living with High-Risk Technologies. Basic Books, New York.

Rasmussen, J. (1990). Human error and the problem of causality in analysis of accidents. Philosophical Transactions of the Royal Society B 327, 449-462.

Reason, J. (1990). Human Error. Cambridge University Press, Cambridge UK.

Review of the General Concept of Separation Panel (RGCSP) (1995). Working Group A Meeting: Summary of Discussions and Conclusions. (1995). ICAO.

Shappell, S. A. and Wiegmann, D. (2001). Applying Reason: The Human Factors Analysis and Classification System (HFACS). Human Factors and Aerospace Safety 1, 59-86.

Sosa, E. and Tooley, M. (1993) editors. Causation. Oxford Readings in Philosophy. Oxford University Press, Oxford.

UK Airprox Board. (1999 onwards - biannual). Analysis of Airprox in UK Airspace. www.ukab.org.uk.

Villiers, J. (1968). Perspectives for Air Traffic Control for Advanced Phases of Automation - the Method of Layers (in French: Perspectives pour le contröle de la circulation aérienne dans les phases avancées d'automatisation - la méthode des filtres), Navigation n° 61, January.

Weick, K.E. (1976). Educational organizations as loosely coupled systems. Administrative Science Quarterly, 21(1), 1-19.

Wickens, C. D. (2001). Attention to Safety and the Psychology of Surprise. 11th International Symposium on Aviation Psychology. Columbus, Ohio. http://www.aviation.uiuc.edu/UnitsHFD/conference/Osukeynote01.pdf

Wiegmann, D. and Shappell, S. A. (2001); Applying the Human Factors Analysis and Classification System (HFACS) to the analysis of commercial aviation accident data. Paper at the 11th International Symposium on Aviation Psychology. Ohio State University, Columbus.

Wiegmann, D. and Shappell, S. A. (2003). A Human Error Approach to Aviation Accident Analysis: The Human Factors Analysis and Classification System. Ashgate Publishing Company.

APPENDIX A

WORD DEFINITIONS RELEVANT TO HAZARDOUS EVENTS

Figure A1 sets out a group of word definitions relevant to hazardous events. These are extracted from a well-established dictionary (Chambers, 2001); the deleted text is irrelevant material (eg that a metal box used to store valuables is a 'safe'). [The words selected here also include some relating to factual evidence – 'actual' and 'existing', which will be used later here.]

Word	Definition – Relevant Extracts
chance *noun*	**1** the way that things happen unplanned and unforeseen. **2** fate or luck; fortune. **3** an unforeseen and unexpected occurrence…
danger *noun*	**1** a situation or state in which someone or something may suffer harm, an injury or a loss… **2** something that may cause harm, injury or loss. **3** a possibility of something unpleasant happening. [from French *dangier* power, therefore 'power to harm']
dangerous *adjective*	likely or able to cause harm or injury
hazard *noun*	**1** a risk of harm or danger. **2** something which is likely to cause harm or danger… **4** chance; accident
hazardous *adjective*	**1** very risky; dangerous. **2** depending on chance; uncertain
risk *noun*	**1** the chance or possibility of suffering loss, injury, damage, etc; danger. **2** someone or something likely to cause loss, injury, damage, etc…
safe *adjective*	**1** free from danger or harm. **2** unharmed. **3** giving protection from danger or harm; secure… **4** not dangerous or harmful… **5** involving no risk of loss; assured… **7** cautious
safety *noun*	**1** the quality or condition of being safe…
actual *adjective*	**1** existing as fact; real. **2** not imagined, estimated or guessed. **3** current; present. [from French *actuel*, meaning 'demonstrated by one's actions']
exist *verb* **(existed, existing)**	**1** to be, especially to be present in the real world or universe rather than in story or imagination. **2** to occur or be found…
straightforward adj*ective*	**1** without difficulties or complications; simple…

Figure A1. Dictionary Definitions of Hazard-related terms (Chambers (2001)

APPENDIX B

HSE RISK ASSESSMENT

The Health and Safety Executive (HSE, 1992, 1999) version of risk assessment is illustrated in Figure B1

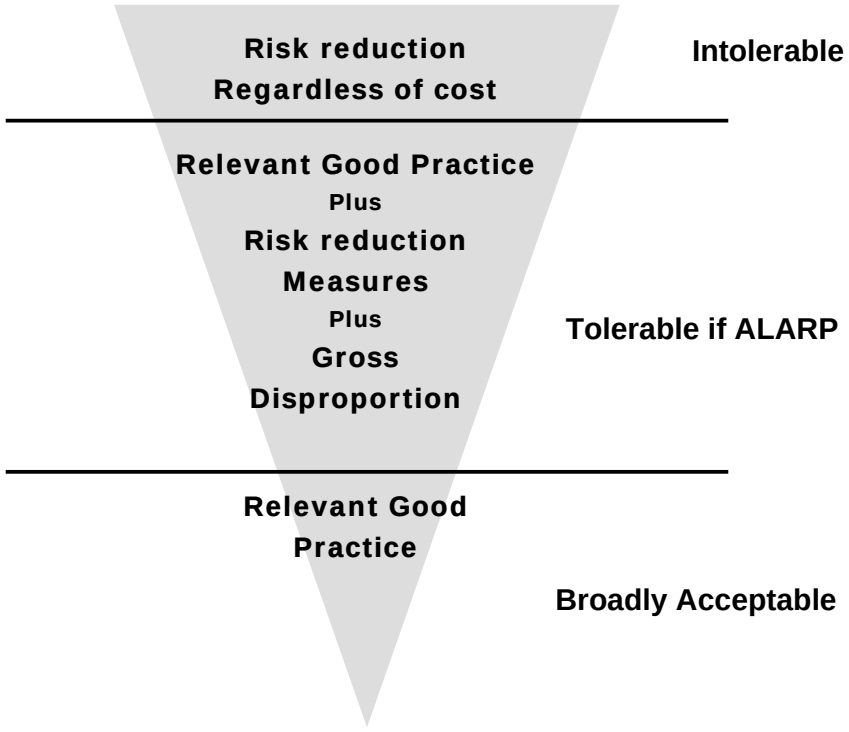

Figure B1. ALARP Approach (taken from HSE)

Risk is classified by the HSE as being in one of three categories: intolerable, tolerable if ALARP, and broadly acceptable ('negligible' in some variants) (Figure B1). Note that the boundary lines between the risk categories negligible, tolerable, and intolerable need to be specified; they are not automatically set.

A checklist of (simplified) HSE definitions is:

ALARP principle The principle that no risk in the tolerability region can be accepted unless reduced 'As Low As Reasonably Practicable'.

broadly acceptable risk A risk which is generally acceptable without further reduction.

intolerable risk A risk which cannot be accepted and must be reduced.

tolerability region A region of risk which is neither high enough to be

unacceptable nor low enough to be broadly acceptable. Risks in this region must be reduced ALARP.

The decision processes are:

- If a system's risk falls into the intolerable category, then action must be taken to redress this. If this is not possible, the system should be halted or not implemented.
- If a system's risk falls into the tolerable category, it must be proven that it is low as reasonably practicable within that region for the system to be considered acceptable. Thus, showing a system is ALARP means demonstrating that any further risk reduction in the tolerable zone is either impracticable or 'grossly disproportionate' (ie it can be shown that the cost of the measure is far in excess of any benefit to be gained).
- If a system's risk falls into the negligible category, no action is required other than monitoring to ensure that the negligible risk is maintained.

An examination of 'what people actually do in aviation' suggests that the message is that the regulators nearly always operate in the 'tolerable if ALARP' region – it cannot be 'intolerable' because flights would therefore have to stop or being tightly constrained. It cannot be broadly acceptable because there are always incidents and accidents that the public expects regulators to examine to see if things can be tightened up. In the middle region, the regulators usually deal with procedural/organizational changes, which are often of low cost (ie tend to involve neither substantial infrastructure capital investment nor dramatically changed operational concepts).

Thus, in the ALARP region, specific safety management measures should be defined (eg safety monitoring, safety improvement projects, etc.) as long as such is reasonably practicable.