THE JOURNAL OF NAVIGATION (2002), **55**, 363–379. © The Royal Institute of Navigation
DOI: 10.1017/S037346330200187X Printed in the United Kingdom

Future Air Traffic Management: Quantitative En Route Safety Assessment Part 2 – New Approaches

Peter Brooker

(*Cranfield University*)

This is the second of two papers on Quantitative Safety Assessment – vital to the successful introduction of future Air Traffic Management systems. The focus is en route European controlled commercial traffic, particularly the mid-air collision risk. Part 2 develops soundly based and practical methods for safety assessment. The objective is to determine the key questions and the best ways to answer them. Aspects covered include lessons from Hazard Analysis and Airproxes together with 'realistic' risk budgeting. Two abstract concepts are introduced: *Position Integrity* and *Reasonable Intent* (essentially the need to be on the 'right' flight path), and their implications for risk calculations are discussed.

KEY WORDS

1. Air Traffic Control. 2. Safety. 3. Modelling.

1. INTRODUCTION. Part 1 of this paper explained the need for improving the Quantitative Safety Assessments of future Air Traffic Management (ATM) systems. It reviewed the techniques that have been used, and identified their strengths and weaknesses. It demonstrated that the traditional techniques are in some cases becoming markedly less appropriate for the newer generations of ATM system.

This part moves beyond criticism and aims to set out a positive approach for the future ATM system based on a thorough grasp of the likely key safety issues. The objective here is to develop appropriate models for safety assessment of future ATM systems. The concern is not with final answers, that is the responsibility of service providers and regulators, but to determine what are the key questions that need answers and what sorts of analyses are required to deal with them.

To reiterate points made in Part 1, the focus here is on European commercial traffic in controlled en route airspace, although other types of ATM system, in particular oceanic, are touched upon. Many of the references used here are UK ones, mainly the Civil Aviation Authority (CAA) and the National Air Traffic Services (NATS – prior to its becoming a subsidiary company of the CAA in the late 1990s). It should be stressed that a great deal of the work in this area has been carried out in international fora in partnership with other countries' representatives.

There are six sections in the following text:

(a) Lessons from Hazard Analysis – sets out some introductory material on Hazard Analysis.

(b) ATM Safety Structure – describes ATM safety in the context of Hazard Analysis.
(c) Lessons from Airproxes – examines a recent sample of commercial Airproxes in terms of the safety defences in operation.
(d) 'Realistic' Risk Budgeting – develops a risk budget breakdown for mid-air collisions using two abstract concepts of *Position Integrity* and *Reasonable Intent*, the latter being the need to be on the 'right' flight path.
(e) Consequences of *Reasonable Intent* Failures – carries out a stylised risk calculation for *Reasonable Intent* failures, and indicates its limitations.
(f) Discussion and Conclusions – discusses the lessons learned and lists some conclusions about safety modelling.

2. LESSONS FROM HAZARD ANALYSIS. The first step in trying to develop new methods for quantitative en route safety assessment of ATM systems is to see what can be learnt from the research literature on risk assessment and its practical application. To begin with a quote from Profit (1995): 'Hazard Analysis is based on prediction, and the accuracy of its results is dependent on the correct identification of all significant (*sic*) potential hazards, the assumptions made about the (usually new) operating environment and on the accuracy of the data analysed.' Profit (1995) also notes, in connection with Fault Tree Analysis techniques, the importance of estimating the probability of initiating risk events and, moreover, the need for 'completeness of the tree in terms of identification of all the contributory lower events'.

'Completeness' is a key problem in safety analysis generally. Mid-air collisions are very rare, so how can an analyst be confident that the most significant causal factors have been properly covered? Hazardous incidents are obviously a key ingredient – in Reason's (1990) phrase, 'They are free safety lessons.' A risk analysis that did not learn the lessons from incidents, in particular Airproxes and Mandatory Occurrence Report data, would be a poor one.

However, incident data is not guaranteed to provide all the information necessary about the performance of new systems. It provides necessary information but is not always sufficient. There may be new modes of error and the designers may have 'blind spots' about significant failure modes. Perhaps the standard example is the sinking of the ocean liner Titanic in 1912 (Encyclopaedia Britannica, 2000). It had a double-bottomed hull divided into 16 compartments, presumed watertight. The liner would stay afloat if four of these were flooded, and it was therefore believed unsinkable. Unfortunately, at least five of the compartments were ruptured when the liner collided with an iceberg. They filled with water and, because the compartments were not capped, they filled each succeeding compartment until the bow was below the waterline. In effect, the compartments acted like a tilted ice cube tray filling with water.

The Titanic catastrophe arose because of a failure to envisage this type of accident. Information on previous accidents and model tests had not uncovered this causal chain – i.e. there was a failure to ensure completeness. This oversimplifies the safety lessons from the Titanic. For example, there were only half enough lifeboats for the people aboard – although 'unsinkable', these had to be in place to meet UK Government regulations – and, in the event, only about half of these were used, because of crew ineptitude and passenger panic. With a better safety culture,

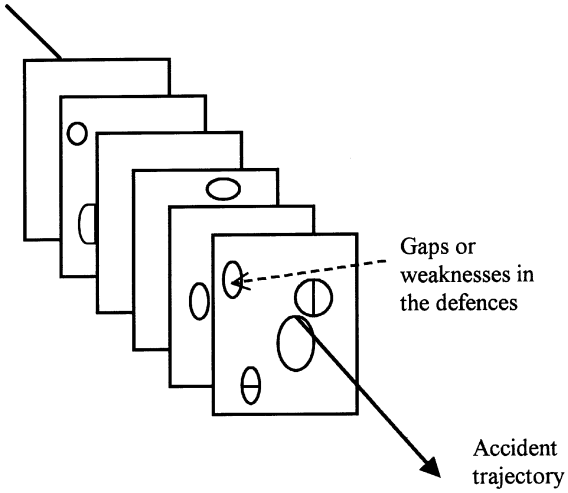

Figure 1. An event involving the complete penetration of a safety system's defences, barriers, and safeguards (from Maurino *et al*., 1995).

including a focus on training, plus the right number of lifeboats, there might have been far fewer casualties.

Completeness and accuracy of probability estimates are particularly difficult when incorporating Human Reliability Analysis (HRA) – trying to put into risk analyses quantitative estimates of the types and magnitudes of contributing factors involving human error. Reason (1990) has carried out a considerable body of work into these methodological problems.

A dramatic illustration of the importance of human factors in risk assessment is that of Chernobyl. Fremlin (1987) was unfortunate enough to estimate nuclear power production risks just before the catastrophe at the nuclear plant at Chernobyl. The accident there was caused by a combination of operator faults and design defects. Operators broke several safety rules, failed to inform the safety managers of their proposed actions, and disconnected all the automatic reactor shut-down devices and the emergency core-cooling system. Chernobyl was obviously an extreme failure of 'safety culture'. In the present context, it shows that HRA calculations can be very dependent on assumptions about 'reasonable' behaviour being extrapolated into the future. It also serves as a lesson in the importance of rigorous defensive mechanisms in ensuring safety, and the consequences of 'cutting corners'.

3. ATM SAFETY STRUCTURE. Figure 1, taken from Maurino *et al*. (1995), illustrates the idea of an accident trajectory. This is a penetration of all the system safety defences, barriers, and safeguards, which could bring uncontrolled hazards into contact with potential victims. The defensive layers range from human monitoring to warning systems (such as STCA and TCAS) to the procedural rules (e.g. separation standards). Gaps in the defensive layers can occur for a variety of reasons:

(a) Undiscovered shortcomings in the defences, which only surface in particular circumstances.
(b) Temporary weaknesses, e.g. during maintenance work.

(c) Unsafe acts by operators, both violations of safety procedures and their disabling to 'improve' system performance.

A very small proportion of unsafe acts actually lead to accidents. This is because of the protection that is offered by the other defensive layers and engineering redundancies.

What are the ATM defensive barriers? UK ATM is a well-defended system. On the illustration in Figure 1, an accident trajectory would require holes to be precisely lined up in each one of the defensive layers. Thus, normally, pilots and controllers notice that something is wrong; Short Term Conflict Alert (STCA) and Traffic and Collision Avoidance System (TCAS) alert them to significant problems and monitoring reveals that an earlier decision was not one that should have been implemented. However, defensive layers are not all 'independent' – STCA and TCAS both depend on SSR transponder carriage, for example.

The present ATC system has several distinct components in its operational concept and hence its safety defences. These have evolved over the last few decades; and are shown below in roughly their date order of introduction.

(a) Controllers and pilots – people are an *integral* part of the whole system.
(b) Formal Rules – for the control of traffic, including the minimum separation permitted between aircraft.
(c) Radio Telephony – communication between controllers and pilots (only introduced shortly before World War II).
(d) 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.
(e) 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.
(f) 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.
(g) Computer Processing – of radar and flight data.
(h) High Quality Aircraft Navigation – progressively improved from VOR/DME to Inertial Navigation Systems through to satellite-based aids, GPS/GLONASS.
(i) Short Term Conflict Alert (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.
(j) Traffic alert and Collision Avoidance System (TCAS) – 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 the aircraft out of risk.

To reiterate, the present concept has evolved over the decades. To some extent, it is 'overlaid', in that new technology has been added on to the previous concept,

rather than being a clean sheet redesign. The system might be said to be 'backwards compatible', in that new functionality generally tends to be able to carry out both the tasks of the previous generation plus some new ones – thus primary radar was followed first by secondary surveillance radar and then by its variant 'monopulse' radar.

These features, and many others, add up to give a system that has to work, day in and day out, delivering safety, capacity, and quality. This system has evolved through the decisions made by the ATC development managers.

However, it may well be the wrong mental model to conceive of the ATM safety system as being constructed upwards from 'foundations'. An alternative approach would be to emphasize the *coherency* of the system structure. The system is in reality made up of interconnecting parts, and it derives its strength from just these interconnections. Each aspect of the safety structure is both supported by and supports other parts. It may be that, as the system evolves over time, some interconnections become much less crucial to the whole – an example would be primary radar, which has become a much less frequently used tool by civil controllers. Thus, some elements can – and should – be replaced if other parts of the system fulfil at least the same tasks. If this were not to be done, then the system could retain unnecessary 'appendices' – in the medical use of the word.

4. LESSONS FROM AIRPROXES. As noted earlier, learning the lessons from Airproxes and other safety incidents is not sufficient to prevent all future types of accident – but it does offer some *necessary* tests. Therefore, as a starting point, the Airprox Report data for the year 2000 (UK Airprox Board, 2000 and 2001) were examined, to help *illustrate* the safety issues. Airprox data is chosen here because it is publicly available and well documented, and indeed Airprox statistics are widely publicised as a 'gold standard' of UK ATM safety. The 2000 data covers the last complete calendar year. It must be stressed that other safety incident data, e.g. the CAA's Mandatory Occurrence Reports, and NATS Separation Monitoring Function data and Safety Significant Events database, would be at least equally valid ATM safety data sources.

Brooker (2002) lists summarised Airprox data for those incidents that involved solely Commercial Air Transport aircraft ('CAT' – covers scheduled/non-scheduled passenger flights in airliners and helicopters, plus cargo flights). All categories of Airprox from A to C are included, although C is 'no risk of collision'. This subset of civil traffic has been chosen because general aviation traffic frequently operates outside controlled airspace and is thus less relevant to the present concern. The information in Brooker (2002) is not official Airprox Board data – it may contain errors of interpretation and omissions.

Many of the inferences about the nature and causes of Airproxes are already made in the Airprox Board Reports. The emphasis here, however, is on three types of defensive barrier that the ATM system now includes:

(a) Human Monitoring and Intervention. Did the pilots see the other aircraft (prior to any TCAS alert or ATC instruction)? Was the controller aware of the conflict prior to any alarms?
(b) Warnings. Did TCAS alert in the aircraft – Traffic Advisory (TA) or Resolution Advisory (RA)? Did the controller get a STCA alert?

(c) Procedures, Rules & Structures. Were there special factors, e.g. extreme weather? Were there procedural problems or confusions? Did equipment fail? Were there human factors issues, etc?

The 'Cause' shown in the summarised information against each numbered incident is taken directly from the Airprox Board Report. Aircraft and controllers have been de-identified. In a few cases, the text has been edited down.

The three types of barrier are illustrated in Figure 2 – developed from Figure 1:

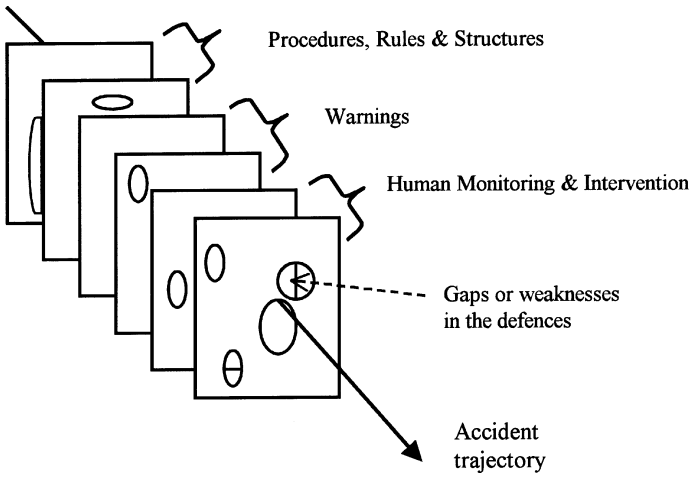

Figure 2. The ATM safety system's defences, barriers, and safeguards.

Some conclusions from the incidents summarised in Brooker (2002) are:

(a) There are *no* incidents recorded in this sample that arise because of a radar accuracy or resolution problem, or from normal altimetry operation. This is not to say that these do not exist but, on the statistical evidence here, they would be a causal factor in only a small proportion of Airproxes.
(b) TCAS and STCA provide alerts in the great majority of cases. One case is noted where ATC might have expected an earlier STCA warning, and in one other there was a possible TCAS malfunction.
(c) There were no instances where the pilots or controllers expressed strong concerns about disruption by warning systems.
(d) Either one or both pilots saw the other aircraft in more than half the incidents before there was any TCAS alert.
(e) In around two thirds of incidents, ATC were actively monitoring the aircraft before there was a warning alert.
(f) Most of the Airproxes were caused by failures in the 'procedures, rules and structures' category by both pilots and controllers. The following are some common phrases from the reports: 'pilot misunderstood the ATC instruction'; 'ATC and pilot procedure errors'; 'controller distraction when sector split'; 'pilots' poor RT discipline'; 'pilot procedure for altimetry in error'; 'ATC memory slip'.

https://doi.org/10.1017/S037346330200187X Published online by Cambridge University Press

Past Airprox Board reports list as the top four causal factors: 'did not separate/poor judgement by controllers', level busts, 'did not pass/late passing traffic information', poor coordination by controllers.

5. 'REALISTIC' RISK BUDGETING. Given the discussions in Part 1 about the differences between past safety analyses and the current – and likely future – risks in the ATM system, how should the system be analysed? Given the problems of completeness in risk modelling, what can be done to ensure that all potentially significant failure modes are covered? Some ideas on how to answer these questions are set out in the following.

To begin, some degree of abstraction is needed. A mid-air collision arises – obviously – when two aircraft collide. What ways are there in which this could happen? Two concepts are introduced: *Position Integrity* and *Reasonable Intent*, whose meanings are as follows:

(a) *Position Integrity*: the system has this when positional equipment is functioning 'normally' – when the errors on radar, GPS, altimetry, measurements are not extreme, when displays work properly, when signals are not corrupted or lost, etc.
(b) *Reasonable Intent*: this is an inference that would usually be made 'after the event': did the controller do what a competent controller would have considered a reasonable (albeit perhaps not perfect) course of action; did the pilot do something that other pilots would have judged decent practice (albeit perhaps not the ideal decisions)? It thus includes misjudgements and blunders as normally understood. It is primarily a human factors issue.

The next step in using these concepts to think about collisions is to bring them together in Figure 3 below. The question is asked: 'Did that concept apply?' This

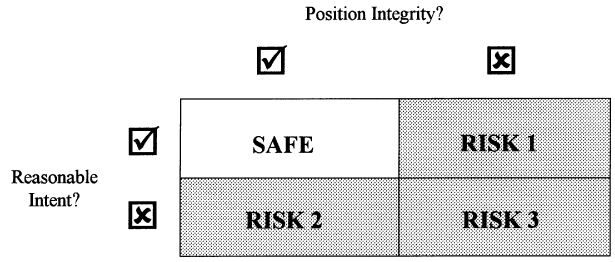

Figure 3. Failures of *Position Integrity* and *Reasonable Intent* – the possibilities.

gives the two-by-two table below, which includes *all* the logical possibilities. This completeness is crucial in risk budgeting. There is one 'safe' box in the table – but three boxes showing 'risk'. Examples of the three would be:

(a) Risk 1 – An aircraft has an altimetry error of 1000 ft.
(b) Risk 2 – Pilot misunderstands clearance, descending to wrong flight level.
(c) Risk 3 – SSR transponder fails *and* controller unable to communicate.

Risk 1 is typical of the sort of event resulting from the loss of planned separation that the Reich model and the radar separations standards work investigated, arising from equipment failure, which would be 'attributable to the loss of correctly

https://doi.org/10.1017/S037346330200187X Published online by Cambridge University Press

established separation'. Risk 2 is typical of observed Airproxes, as summarised in the earlier section: generated by human error in the widest sense. In crude terms, Risk 1 is 'wrong place on right flight path' and Risk 2 is 'right place on wrong flight path'.

Risk 3 would usually be described as a complex system failure, with multiple equipment and human problems happening at the same time. Safety regulators would usually see Risk 3 events as indicating 'Category 1' hazards (Profit, 1995). This is, for example, why the time period of any loss of communications is a critical safety requirement.

Risk 1 and 2 events are 'first order', in that they reflect what are essentially single failures, whereas Risk 3 events would be expected to be the product of multiple problems and specific emergencies rather than the coincidence of just the two types of failure. Risk 3 events of this type are well covered by safety regulation – as evidenced in Profit (1995) – and are therefore not discussed further here.

It is reasonable to assume that collision risks are roughly proportional to the frequency of Airproxes (the next section has a discussion on this point). On this basis, Risk 2 events would be far more frequent than Risk 1 events. In risk budgeting terms, Figure 4 below shows one way of partitioning the risks; the percentages used are

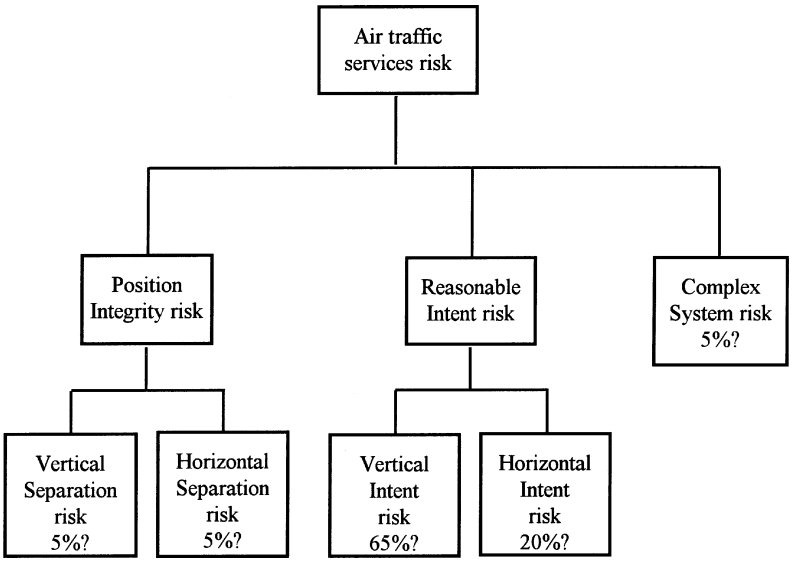

Figure 4. One way of partitioning the en route risk budget, based on Airprox proportions.

illustrative but 'reasonable'. It should be noted that the Reich model described in Part 1 sometimes required *Reasonable Intent* failures to be included in calculations as *Position Integrity* errors. An example was the so-called 'ATC loop error', by which aircraft flew on the wrong track. This could arise from an error in ATC instructions or through errors in communication between aircraft and ATC.

6. CONSEQUENCES OF *REASONABLE INTENT* FAILURES. The discussion on 'Realistic' Risk Budgeting above uses the observed relative proportions of different types of Airprox to suggest that comparatively large proportions of the

Air Traffic Services risk budget should be allocated to the two types of *Reasonable Intent* risk. The next questions are: 'how can these risks be estimated? How can the safety barrier concept be used in such calculations?'

One point that must be made at the outset is that these risks cannot be calculated by a similar means to the *Position Integrity* risks. With these 'equipment' risks, it is possible to analyse relevant data of what are essentially engineering observations, as sketched in Part 1. *Reasonable Intent* risks are, by their very nature, deeply embedded in human functions and performance. As noted earlier, this is studied by Human Reliability Analysis (HRA), which is not yet a fully developed science. Dougherty (1990) and Hollnagel *et seq*. (1997) serve as examples of the continuing debate on the difficulties of integrating human operator elements into quantitative safety modelling. To quote Hollnagel (1997c): 'The HRA community was seriously shaken in 1990 by Dougherty's (1990) concise exposure of the lack of substance in the practically used HRA approaches, and has not yet recovered from that.'

HRA is a complex blend of human error theory, human factors engineering, cognitive psychology, and probabilistic analysis. Much of the development work has been focused on nuclear power stations (but see Fremlin (1987)) and chemical process plants; there appears to be comparatively little published work on aviation or ATM.

No examination of future ATM systems will be able to avoid detailed scrutiny of HRA aspects, so the problems of the application of HRA's application to ATM will need to be solved. This might well be a formidable undertaking: controllers' tasks appear to be much less structured, i.e. with more discretion to determine solutions – than are nuclear workers'. However, the present aim is to set out a context for these sorts of studies rather than to try to answer all the questions at one attempt.

The methods adopted below are simple and technically stylised, but probably cautious in a safety sense. The test for models having greater complexity is the extent to which their mechanisms and interconnections can be *validated* quantitatively. In the following, the 'Procedures, Rules & Structures' defensive barriers will be treated as essentially a 'black box'.

6.1. *Nature of Collision Risks*. The first step is to examine the nature of the collision risk posed through failure of *Reasonable Intent*, ignoring for the moment warning systems and controller/pilot actions to prevent an accident. There are several mathematical models in the literature for estimating the likelihood of near and actual mid-air collisions. The method used here is a simplified version of that set out in May (1971).

The first point to make is that airspace today is probably not as three-dimensional as it once was viewed (Lee, 1976). Most modern jets are optimised to fly at around the tropopause, say 35,000 ft, so FL350 and the neighbouring flight levels are very popular. En route flight occupies a large proportion of most flights. Figure 5 illustrates two flight levels, each of which is densely packed with aircraft. Improved altimetry means that aircraft fly in narrow bands around the flight level. Moek *et al*. (1993) gives some data on vertical errors recorded by various categories of aircraft. For long/medium range types, the standard deviation was 85 feet; and it was 155 feet for medium/short range types; some aircraft were 'rogues' with deviations exceeding 300 feet; these rogues occurred at a rate of 0·17% and 1·66% respectively.

Since the research by Moek *et al*. (1993) was published, there have been strenuous efforts on the part of airlines and ATC providers, as part of the programme of work to reduce vertical separation above FL 290 to 1000 feet, to improve and monitor

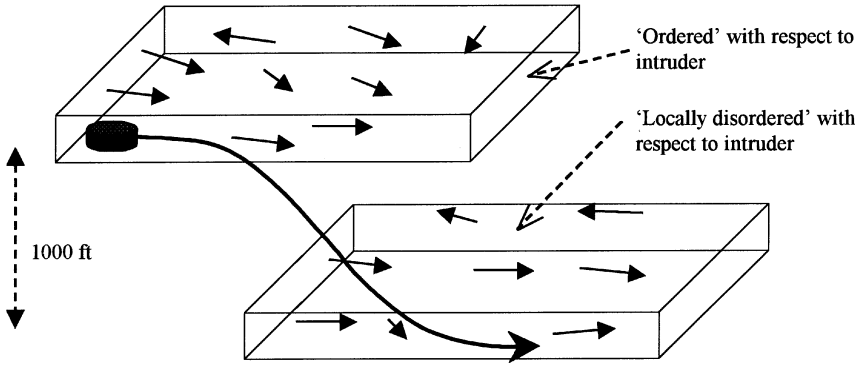

Figure 5. Descent by one flight level.

vertical performance (Moek *et al.* (1993) discusses these efforts). Figure 5 therefore shows aircraft at the two levels as being distinct and well separated.

The next point to note is the difference between the two 'flight level cuboids' when an aircraft moves – without *Reasonable Intent* – from one to the other. Figure 5 shows such a descent by aircraft 1 – represented by a disc, its path being shown as the heavy curved arrow. The other aircraft are shown by short arrows indicating their velocity vectors.

The descending aircraft does not face abnormal risk in the higher flight level, because air traffic control has 'ordered' the upper level traffic. In other words, the controllers have aimed to set up the aircraft positions and velocities so that conflicts do not occur for some period ahead. Controllers may be coordinating the flight paths of particular aircraft that could soon be in conflict. They may recognise that they need to do certain things to ensure that particular aircraft pairs are conflict-free when they are passed on to their next controller. But the point is that, at any given time, aircraft at the upper level have a high degree of ordering as regards their position and velocity vectors.

In contrast, the aircraft at the lower level are *not* positionally 'ordered' with respect to the intruder. In no sense were their flight paths constructed to avoid conflicts with that aircraft. It may be that their velocity vectors are statistically correlated with that of aircraft 1, because there would be expected, at a given time and in particular geographical areas, to be a flow of aircraft in one general direction, but their positions at the time that aircraft 1 reaches the lower level will *not* have been designed to be conflict-free. Thus, the intruder has essentially moved into a 'locally disordered' layer, i.e. what is essentially 'statistically random' traffic with respect to the intruder.

What are the relative risks in the intruder's descent? This is difficult to estimate accurately because the path of the aircraft and the time it spends in the different vertical zones are in practice quite variable. However, a simple cautious estimate can be made, on the following chain of argument:

(a) Clearly, the main collision risk from the intruder is when it is in the bottom flight level.
(b) The risks in the initial flight level are relatively less than this because of the ordering of traffic.
(c) The risks in the middle zone are less than in the bottom flight level, because there are very few aircraft – even if they are similarly 'disordered'.

(d) The total risk during the descent must therefore be less than if the intruder had spent the whole descent time in the 'locally disordered' flight level.
(e) Thus, the total risk of collision is at most equal to that for an aircraft in a narrow 'slice' of random traffic at the destination flight level.

That simplifies the Vertical Intent case – what about the Horizontal Intent risks? The concept of ordered airspace helps to resolve this. An examination of the Airproxes in Brooker (2002) reveals that in most of the more hazardous instances the controller has lost awareness of an aircraft or of potential conflicts between aircraft. In essence, therefore, the degree of 'orderliness' of the system has been reduced markedly for the intruder – the aircraft would be intruding into near random traffic. Again, as for the vertical case, there may be some positive statistical correlation of velocity vectors, e.g. because of traffic flows, but the positions of some of the other aircraft could be relatively 'disordered', i.e. essentially random. It is unlikely that aircraft at about the same flight level in a local area would be flying in opposite directions. Given the relationship between aircraft relative velocity and risk set out below, the risk of collision from a failure of Horizontal Intent would be at most equal to that for an aircraft that had intruded into a 'slice' of random traffic.

The discussion above is of a general logical nature rather than a detailed model embodying specific assumptions. It should therefore be a robust set of arguments. However, it is possible to think of circumstances that do not readily fit the logic; these are discussed later in this section.

6.2. *An Illustrative Risk Calculation for En Route ATM.* May (1971) derives the statistical expectation of a collision for an aircraft, represented by a disc of radius R and height H, that spends a time period T in an airspace volume A filled with N similar-sized aircraft; and all aircraft have velocity V. N divided by A is the density D of aircraft.

The statistical expectation of a collision is:

$$\frac{2H \times 4R \times T \times N \times (\text{Average}V_r)}{A}.$$

'Average V_r' is the magnitude of the average relative velocity between the intruder aircraft and the N aircraft. For random traffic, i.e. with all values of the angle θ between the aircraft velocities being equally likely (May, 1971), $V_r = 4V/\pi$. Were all the N aircraft to be flying in the opposite direction to aircraft 1, then V_r would equal 2V. In practice, there would tend to be flows along routes, so that the velocities would be positively correlated, and θ would on average be a small angle, with V_r itself being small in comparison with V.

Cautious values of the aircraft dimensions would be (Harrison and Moek, 1992):

$$H = 0{\cdot}010 \text{ nm} = \text{about } 60 \text{ feet};$$
$$R = 0{\cdot}017 \text{ nm} = \text{about } 100 \text{ feet}.$$

The figures correspond to large jets, i.e. those on oceanic routes. However, the other parameters in the expression (T and D) are more difficult to estimate. There are several ways to do this; the method used in the following relies on inferences from Airprox data.

Brooker (2002) lists the CAT/CAT Airproxes in 2000. Incidents in oceanic airspace, those with aircraft in the final landing or take off phases, and one involving

a helicopter are marked by italic text. Excluding these 5, there were 30 Airproxes for UK airspace under radar control. In 2000, there were 1·389 million flying hours in UK airspace (UK Airprox Board, 2000 and 2001). Thus, the rate of Airproxes is around $2{\cdot}2 \times 10^{-5}$ per flying hour (NB: for a better estimate this should be scaled down to en route hours). With an environment incorporating TCAS and STCA, the extent of under-reporting of CAT/CAT Airproxes is generally believed to be small.

Given the argument above, that the positions of air traffic in the 'intruded' flight zones is essentially random as far as the intruding aircraft is concerned, the likelihood of a collision compared to an Airprox will be in proportion to their relative dimensions. This is because the calculations in May (1971) follow through in exactly the same way for an Airprox as they do for a collision – the values of T, N, A and (Average V_r) would be the same. This assumes that a separation infringement corresponding to an Airprox has consistent effective dimensions, having a similar disc shape to the aircraft.

The calculation on the lines of May (1971) for an intruder in random traffic does not distinguish between *Position Integrity* and *Reasonable Intent* conflicts. Thus, the same argument, scaling the risk of a mid-air collision and an Airprox according to the aircraft dimensions, works for both types of conflict. This supports the argument earlier that the collision risk is generally proportional to the Airprox rate. (NB: the absolute level of risk would be lower for *Positional Integrity* failures in same-direction track systems because of the lower relative velocity factor.)

As approximations to Airprox dimensions H_a and R_a, analogous to H and R respectively, 1990s NATS Separation Monitoring Function (SMF) research studies generally used 'encounters' of aircraft pairs within 2 nm and 600 feet (below FL290) and 1600 feet (above FL290). Using this as a guide, the following dimensions are chosen:

$$H_a = 0{\cdot}050 \text{ nm} = \text{about } 300 \text{ feet},$$
$$R_a = 1 \text{ nm}.$$

Thus the statistically expected ratio of Airproxes to collisions is:

$$\frac{H_a \times R_a}{H \times R} = 0{\cdot}050 \times 1/0{\cdot}010 \times 0{\cdot}017 = 294.$$

On this basis, the statistically expected rate of a mid-air collision posed through failure of *Reasonable Intent* would be $2{\cdot}2 \times 10^{-5}/294 = 7{\cdot}5 \times 10^{-8}$.

As already noted, this calculation does not include the defensive barrier effects of automatic warning systems and controller/pilot action (including the effectiveness of see-and-avoid). Research has been carried out into this subject, but little of it is in the published literature; Datta and Oliver (1992) is one of the few examples, and notes that 'there is a great shortage of relevant data'. Datta and Oliver (1992) assume (*sic*) that pilots make a wrong manoeuvre on about 2% of occasions, and refers to some earlier work by the Mitre Corporation, primarily focused on TCAS II design, to support this assumption.

Also very relevant is Hale and Law (1989), which analysed simultaneous STCA and TCAS performance in UK en route airspace. It showed that the two systems worked effectively, in that all genuinely 'serious' encounters found in a *post hoc* analysis were detected by both systems.

Given this limited statistical evidence, it is difficult to estimate what safety gain there is from the combination of warning systems and human intervention. It is certainly better than a factor of 10; the data from Airprox reports could be consistent with a factor of 100, and it could be even higher. For present purposes, a factor of 20 is assumed – noting that this implies that up to 95% of mid-air collisions would be prevented. This would give an expected risk of a mid-air collision posed through failure of *Reasonable Intent* as $7{\cdot}5 \times 10^{-8}/20 = 3{\cdot}75 \times 10^{-9}$. As one mid-air collision equals 2 fatal aircraft accidents, this gives the accident rate as $7{\cdot}5 \times 10^{-9}$. This rate is to be compared to the 1993 safety target (for 2010) of 0·1 fatal aircraft accidents per 10 million aircraft en route flying hours, i.e. $1{\cdot}0 \times 10^{-8}$. For *Reasonable Intent* accidents, this has to be scaled down, according to the 'Realistic' Risk Budgeting proposal (Figure 4) by about 85% (i.e. 65% + 20%), i.e. to $8{\cdot}5 \times 10^{-9}$.

6.3. *Flaws in the Collision Risk Calculation*. Thus, on this simple calculation, the en route risk meets the target. The obvious question is to what degree the calculation is robust. There is one immediate flaw in that the analysis only covers CAT aircraft rather than the full range of controlled airspace users. Airproxes quite frequently occur between CAT aircraft and general aviation aircraft, which would tend to increase the estimated risks. These types were excluded because not all are fitted with TCAS, and incidents tend to occur at the boundaries of controlled airspace.

A much more important issue is the extent to which the 'intruded airspace' traffic is truly 'random' – locally disordered with respect to the intruder. There are some obvious cases when this might not be so. Thus, Airprox reports give examples where an aircraft at one flight level would be likely to have aircraft nearby at the adjacent level. One category is incidents in holds – Incident 29 (Airprox Board reference number 176) is an example. Here, an aircraft was descended to the level of the next aircraft in the hold; there was a TCAS TA; and ATC issued avoiding action. (NB: the degree of hazard was assessed by the Airprox Board as 'C' – no risk of collision.)

This is an instance where the safety barriers might not have operated according to the exemplary calculation sketched earlier in this section. There was less protection from 'random traffic' but the defensive barriers offered by the warning system and human monitoring worked very well. In particular, the pilot showed good situational awareness in spotting the other aircraft on TCAS and acquiring it visually.

This type of incident shows the importance of action to reduce 'height busts' in those circumstances where there might be traffic at the adjacent flight levels. The CAA Aviation Safety Review (2000) gives some statistical information on them and programmes of work to reduce their incidence. For comparison, there were 319 height busts by CAT aircraft in 2000 (source: Civil Aviation Authority).

Another example is a loss of 'orderliness' occurring at a flight level through a trainee deciding on an inappropriate course of action. The trainee might not have the experience to judge accurately enough the effects of relative speeds or aircraft manoeuvrability. In such a case, the trainee would seek help from the controller mentor, with the various warning systems alerting pilots and controllers to hazardous circumstances. In essence, the mentor provides an extra defensive layer.

Where traffic is not 'random', safety protection may be being provided through low relative velocities, which offer added time to detect proximity, by special rules and procedures – to lessen the likelihood of intent failures – or by controller mentors in the case of trainees' errors.

7. DISCUSSION AND CONCLUSIONS. Part 1 reviewed ATM safety assessment for en route ATC, and the preceding text has then examined a different framework for supporting quantitative aspects of such assessments. This section discusses the lessons learned regarding future ATM system development and sets out some conclusions.

Traditional techniques used are in some cases becoming less appropriate for the coming generations of ATM system. Once the answers to yesterday's key questions are resolved by model analysis, there is a choice to be made about the safety improvement path; is it best to refine those models or to tackle new types of problem? The argument here is that it is essential to create models that match the patterns shown in, for example, Airprox reports; these models will then provide useful information about the types of safety barrier that should be incorporated to deliver the necessary safety.

The behaviour of a model – logical structures, sets of equations or computer simulations – can give valuable insights, often of a quantitative nature, into the performance of the real system. But the model is not always a good approximation to present and future systems. There is always a temptation to model that which can be modelled and put the rest aside 'temporarily'; thus producing a distorted picture of the true safety situation. 'Completeness' is, however, essential to safety assessment and risk budgeting for ATM system design.

The review sketched out in Part 1 shows how useful the Reich model has been, initially for the North Atlantic and most recently for the reduction of vertical separation above FL290, the latter being one of the most thorough and worthwhile pieces of research and analysis into ATM safety. However, the model did not prove very useful for problems involving en route radar-based control. The Reich model has proved its value, but it is a 'snare and delusion' to believe that it or any simple 'equipment-based' model will be appropriate in the safety assessment of future ATM systems.

It is no longer sufficient to model equipment performance and to focus on *Position Integrity*, the main concern of the Reich model (altimetry seems to be the sole exception). The limited statistical analysis of Airproxes carried out above shows that the concept of *Reasonable Intent* is crucial. The misjudgements, blunders, and other human factors/cognitive failures that lead to an aircraft taking the 'wrong' flightpath are currently the main cause of Commercial Air Transport Airproxes, and hence *prima facie* correspond to the greatest mid-air collision risk. It is not sufficient to engage in 'large-scale' thinking about ATM systems; it is necessary to think about the nature of individual hazardous events such as Airproxes.

Thus, on the evidence of Airproxes, the risk budget for en route ATM has to be re-thought. *Positional Integrity* risks, such as those caused by radar inaccuracy and altimetry errors, probably contribute a small proportion of the budget, as do those where there is effectively a loss of air traffic control over a limited period. The main contributors arise from failures of *Reasonable Intent*. Risk budgets should match the proportion of Airproxes of the different types of risk. However, this is not a totally 'foolproof' approach; attention has to be paid to exposure to and protection from complex system failures, and to identifying any varieties of risk that might not show early warnings through Airproxes.

7.1. *Future ATM System Quantitative Safety Assessment: Key Points.* The simple model of ATM risk sketched out above, using Reason's ideas of defensive

safety barriers, indicates that the ATM system can meet the very demanding safety targets that are being placed upon it. However, even such a simple model indicates where the safety defences might be at risk in future systems development. The following paragraphs draw out some of the explicit and implicit points made above.

7.1.1. *Safety Targets*. Some of the problems of en route collision risk safety targets were identified in the mid 1990s (Davies and Sharpe, 1993). Statistically, mid-air collisions are a much smaller proportion of aircraft accidents than in earlier decades. This may well be because of the addition of new safety defences, such as TCAS and STCA, on top of the considerable improvements in display technology, radar performance and data processing. Whatever the reasons, the en route safety target, derived from projecting forward overall safety improvements, is now an extremely demanding design target. A mid-air collision risk of 5×10^{-9} per aircraft flying hour applied to the current UK annual CAT flying hours of $1{\cdot}389 \times 10^{6}$ would imply a collision in the UK at a rate of about 1 in 150 years – longer than any individual's lifetime.

7.1.2. *Separation Standards*. It has been argued above that the methods for estimating the risks 'attributable to the loss of correctly established separation' are extremely cautious, and that the evidence from Airproxes is that these types of *Position Integrity* risk are much less important than those arising from *Reasonable Intent* failures. However, the present values of separation standards *might* still be appropriate because of the corresponding time they allow controllers to resolve potential conflicts, i.e. to some extent they perform a 'fail safe' function in route planning and tactical control. To improve flight efficiency, these standards might be set at much smaller distances. However, if controllers operated to those smaller values, then there would, for example, be a much higher frequency of TCAS and STCA advisories, which would change the nature of ATC tasks dramatically. The key safety model would then need to be one that estimated the frequency of these alerts and the implications for acceptable controller workload. Any air/ground datalinking of intent information in a future ATM system (or the putting in place of some sort of 'governor mechanism' on manoeuvres) would be an important beneficial factor.

7.1.3. *Procedures, Rules & Structures*. These defensive barriers have been treated in the modelling above as essentially a 'black box', but they actually incorporate complex human reliability structures. Major changes to responsibilities or procedures could produce equally extensive changes to the likelihood of certain types of Airprox, so elements of the black box would need to be modelled.

7.1.4. *Defensive Barrier Philosophy*. The defensive barrier philosophy used here for analysing ATM system risks derives from research work on human errors (Reason, 1990). It is a rational way of examining system risk and, when supported by quantitative estimates, helps to indicate the kinds of event defences that need to be strengthened. Human reliability analysis techniques provide systematic ways of carrying out this quantification, but they do not appear to have been used very often in this area.

7.1.5. *Risk Budgeting*. A new framework for thinking about mid-air collision risk budgeting has been suggested. This uses the concepts of *Positional Integrity* and *Reasonable Intent*. In conjunction with Airprox statistical analyses, these shift the focus away from equipment errors. The vital need is to ensure that risk budgeting is a 'complete' exercise and that attention is not over-focused on the more readily modelled types of risk.

7.1.6. *Use of Airprox Information*. Airproxes have been used to indicate what would be the most likely 'generic' causes of a mid-air collision. Other safety incident data could equally well be used. It is, however, important to ensure that too much reliance is not placed on analyses of particular kinds of incident rather than on types of failure (e.g. of *Reasonable Intent*); it is the generic nature of system barriers that can deliver their greatest benefits. Otherwise, there can be the 'Titanic risk', i.e. where the lack of any experience of similar non-catastrophic events produces over-confidence about the scope of potential hazards.

7.1.7. *Use of MORS data*. Airprox data has been used here because it is published and widely available. It must be stressed that MORS (Mandatory Occurrence Reporting Scheme – CAA (1996)) data is also a valuable resource on ATM incidents. In particular, it can be used to test the Airprox and risk budget arguments set out above that radar-related *Position Integrity* failures are infrequent (as a proportion of flying hours). As noted in Part 1, vertical errors and their time duration are already being monitored as part of the introduction of reduced separation above FL290.

7.1.8. '*Randomness*' *as a Safety Barrier*. A degree of 'randomness' safety protection is provided when the air traffic into which an aircraft intrudes is 'locally disordered', i.e. its position is not statistically correlated with that of the intruder. Thus, the low density of traffic helps reduce risks. This is not always fully the case in the present system, so extra defences have to be incorporated, e.g. a mentor when a trainee is handling traffic, and special procedures at holds. A future ATM system that, for example, introduced significantly higher local traffic densities or increased tactical methods of separating aircraft could reduce the extent of randomness protection.

7.1.9. *Warning Systems*. A decade ago, warning systems such as TCAS and STCA would have been seen as 'last ditch' safety bonuses. Today, it can be argued that they are vital pieces of safety equipment, integral to the delivery of the very demanding safety targets for en route ATM. An examination of the Airprox reports summarised in Brooker (2002) shows that a TCAS warning is frequently the first indicator of potential hazard. Perhaps even more relevant are the occasions when a pilot detects potential problems by monitoring the traffic information on the TCAS display, i.e. before even a TA occurs. Thus, TCAS is additionally being used as a situational awareness tool. In system terms, TCAS and STCA are very effective safety protection barriers, not just aids to pilot and controller vigilance.

7.1.10. *Communications, Navigation and Surveillance* (*CNS*). For reasons of space, very little has been said specifically above about CNS, but these technologies do raise a whole range of new issues. In particular, high performance from all these systems is essential if all the defensive safety barriers are to function effectively. Communications is well recognised (e.g. Profit, 1995) as key to safe ATC. The present use of SSR is crucial: it has transformed radar control, it underpins STCA, and SSR transponders equally underpin TCAS. Navigational improvements have led to reductions in the likelihood of 'gross' errors and hence enabled reduced separation standards.

7.2. *General Conclusions*. The aim has been to show what kinds of ATM modelling and safety analysis would be fruitful in future. Conceptually simple models have served the aviation industry well in the past. The next generation will necessitate more complex modelling, and assessments will need to take more account of both human factors aspects and the effective usage of automatic warning systems.

Comparatively little of the work in this area is being published in the open research literature – it tends to be documented in ICAO papers, conference proceedings and ATC providers' internal reports. This needs to change. Given the paramount importance that the industry attaches to safety, it is essential that the merits of different approaches for safety modelling and assessment are publicly tested.

ACKNOWLEDGEMENTS

This work was supported by a research grant by the Civil Aviation Authority's Safety Regulation Group (SRG). I would like to thank SRG staff and Dr David Harrison of National Air Traffic Services for their comments on earlier drafts. Thanks are also due to Ms Ros Howell and the Airprox Board staff for their help with Airprox data.

REFERENCES

Brooker, P. (2002). Summarised CAT/CAT Airprox Data for 2000. *Cranfield Research Report PB/1/3/02*, Cranfield University.

CAA (1996). Mandatory Occurrence Reporting Scheme, *CAP 382*, CAA.

CAA (1998). Air Traffic Services Safety Requirements. *CAP 670*, CAA.

CAA (2000). Aviation Safety Review: 1990–1999. *CAP 701*, CAA.

Datta, K. and Oliver, R. M. (1992). A model to predict mid-air and near-mid-air collisions. *Journal of Forecasting*, **11**, 207.

Davies, E. H. and Sharpe, A. G. (1993). Review of the Target Level of Safety for NAT MNPS Airspace. *CS Report 9301*, NATS.

Dougherty, E. M. (1990). Human reliability analysis – where shouldst thou turn? *Reliability Engineering and System Safety*, **29**, 283.

Encyclopaedia Britannica Inc. (2000). Titanic. Britannica.co.uk, CD 2000 Version, London.

Fremlin, J. H. (1987). *Power production: what are the risks?* Oxford University Press, Oxford.

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. and Moek, G. (1992). European studies to investigate the feasibility of using 1000 ft vertical separation minima above FL 290: Part II – Precision data analysis and collision risk assessment. This *Journal*, **45**, 91.

Hollnagel, E. (1997). (a) Reliability analysis and operator modelling. *Reliability Engineering and System Safety*, **52**, 327, 1996. Plus follow-up articles in the *same journal*: Lydell B. (1997). (b) A practitioner's view on the state of HRA methodology. **55**, 257; Hollnagel E. (1997). (c) Reply to …, **55**, 261.

Lee, L. (1976). [Author of 'Three-dimensional darkness: The World of the Airline Pilot' – out of print], personal communication.

Maurino, D. E., Reason, J., Johnston, N. and Lee, R. B. (1995). *Beyond Aviation Human Factors*. Ashgate Publishing, Aldershot UK.

May, G. T. A. (1971). A method for predicting the number of near mid-air collisions in a defined airspace. This *Journal*, **24**, 204.

Moek, G., ten Have, J. M., Harrison, D. and Cox, M. E. (1993). European studies to investigate the feasibility of using 1000 ft vertical separation minima above FL 290: Part III – Further results and overall conclusions. This *Journal*, **46**, 245.

Profit, R. (1995). Systematic safety management in the Air Traffic Services. *Euromoney*, London.

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

UK Airprox Board. (2000 and 2001). Analysis of Airprox in UK Airspace. *Reports Numbers 4 and 5*.