17th October 2007

AIR TRAFFIC SAFETY: CONTINUED EVOLUTION OR A NEW PARADIGM?

Peter Brooker

Cranfield University

Abstract

The context here is Transport Risk Management. Is the philosophy of Air Traffic Safety different from other modes of transport? – yes, in many ways, it is. The focus is on Air Traffic Management (ATM), covering (eg) air traffic control and airspace structures, which is the part of the aviation system that is most likely to be developed through new paradigms. The primary goal of the ATM system is to control accident risk. ATM safety has improved over the decades for many reasons, from better equipment to additional safety defences. But ATM safety targets, improving on current performance, are now extremely demanding. What are the past and current methodologies for ATM risk assessment; and will they work effectively for the kinds of future systems that people are now imagining and planning? The title contrasts 'Continued Evolution' and a 'New Paradigm'. How will system designers/operators assure safety with traffic growth and operational/technical changes that are more than continued evolution from the current system? What are the design implications for 'new paradigms', such as the USA's 'Next Generation Air Transportation System' (NextGen) and Europe's Single European Sky ATM Research Programme (SESAR)? Achieving and proving safety for NextGen and SESAR is an enormously tough challenge. For example, it will need to cover system resilience, human/automation issues, software/hardware performance/ground/air protection systems. There will be a need for confidence building programmes regarding system design/resilience, eg Human-in-the-Loop simulations with 'seeded errors'.

1. Introduction

The context here is Transport Risk Management. Is the philosophy of Air Traffic Safety different from other modes of transport? Yes, in many ways, it is. The title contrasts 'Continued Evolution' and a 'New Paradigm'. Is the use of these two clichés appropriate? What are the past and current methodologies used for air traffic risk assessment? Will they work effectively for the kinds of future systems that people are now imagining and planning?

The focus here is on Air Traffic Management (ATM), which is the part of the aviation system that is most likely to be developed through new paradigms. This is a definition of ATM from an ICAO – International Civil Aviation Organization – document.

> "Air Traffic Management is the dynamic and integrated management of air traffic and airspace, safely, economically and efficiently, through the provision of facilities and seamless services, in collaboration with all partners."

Note the emphasis on partners: airlines and airport operators play a big role. The following list shows some of the features of ATM. .

- Safety
- Air Traffic Control – ATC
- Airspace design & routes
- Technology

Most of the following deals with commercial, passenger and freight carrying, aviation in developed countries.

The text is divided into a number of sections; the abbreviations are explained later:

2. History and Technical Background
3. Decision-making on Safety Improvements
4. ATM Safety – Systems Theory
5. Tightly-coupled Sub-systems: Separation Minima
6. Loosely-coupled Sub-systems: STCA and TCAS Decisions
7. Future ATM Systems
8. New Paradigms
9. Safety of New ATM Paradigms
10. Key Safety Assessment Messages and Ideas

The first two thirds of these Sections use aviation material to draw out general and interesting lessons, eg about decision-making, data, risk models, as well as lessons for aviation itself. In the final third, the focus is more specifically on new ATM paradigms – which are known as Nextgen and SESAR.

History and Technical Background

Where to start? How did things evolve? The obvious place is the first flight, which the USA would have us believe took place on 17th December 1903 by the Wright brothers. Some people do not believe that the Wrights actually achieved powered flight. The 2003 re-enactment of the flight failed totally.

But Orville Wright certainly has a first in that he piloted an aircraft in 1908 that crashed and caused the first passenger death: September 10th, Fort Myer, Virginia. The first passenger to die was Thomas Selfridge. He happened to be the heaviest passenger so far carried, at 175 pounds (12½ stone, about 80 kilograms): very heavy for those days – Orville weighed 30 pounds less.

Orville Wright was the pilot; but he was also the crash investigator (the cause was a stress crack in the propeller); and he was the aircraft engineer who designed the improved aircraft. There were no flying licences; there was no accident investigation; there was no regulatory action; there was no government inquiry.

The first airliner mid-air collision was in 1922. It happened in Beauvais, north of Paris. The aircraft were on the London-Paris route flying towards each other in bad weather. Figure 1 shows the aircraft types involved – dramatic advances on the Wright Brothers early machines. The mid-air collision was reported in Flight magazine, but only at the bottom of page 7. There were rather a lot of fatal accidents in those days. This is a key point: aviation safety culture has changed, because society's safety culture generally has changed dramatically. What might happen today? – presumably some combination of accident investigation/report, public inquiry, and criminal/civil prosecution? [Try looking through old copies of Flight to find them.]

It took only a matter of months to solve the London-Paris safety problem. It was simple. Safety in the air then depended on looking out of the cockpit window. So, pilots were instructed to keep to the left of the various roads and rail lines that they used to navigate their way between the cities. This was an operational rather than a technical solution. It involved the airline companies and government departments. But a crucial point is that it was an Anglo-French solution – an agreement between States. Aviation safety generally requires States to agree on safety developments.

During the 1920s and 1930s there were technical and organisational changes that improved operating efficiency and safety. Radio telephony began to be used. There was 'Wireless Traffic Control' by 'control officers'. The airspace around the major aerodromes – Croydon was the biggest UK aerodrome – started to have restrictions placed upon its use ('control zones').

The Second World War led to huge changes in aviation: many of them fed into civil use. The present ATM system has several distinct components in its operational concepts and technology infrastructure – Figure 2. Air traffic controllers are important decision-makers. They communicate through radiotelephony; they use flight plans agreed with pilots; they monitor highly processed secondary surveillance radar – SSR – data. These data flows are embedded in 'safety structures', eg with well-defined controlled airspace and formal rules for control such as the minimum separation permitted between aircraft.

Navigation has developed enormously from the war-time systems. The system has moved from point source beacons to satellite-based aids, mainly GPS, which are incredibly accurate. In the UK, there are ground-based short-term conflict alert – STCA – systems available to warn controllers of aircraft coming into close proximity. Commercial aircraft now carry a Traffic alert and Collision Avoidance System – TCAS – which actually tell pilots what to do to avoid another aircraft that has lost separation.

The crucial change, in terms of both its current benefits and potential for system development, is SSR. The controller sees displayed aircraft symbols, callsign and height information, which have been passed down from aircraft transponders. This innovation was a huge step, because it meant that operational information could be passed from the aircraft to the ground system. It is a huge potential step for the future because the aircraft effectively has the equivalent of a telephone number and wireless connection: information can be sent from air to ground and vice versa. In the technical jargon, this is 'datalink'– but it is basically text-messaging between all the active aviation parties.

What safety performance does aviation deliver? There are plenty of statistics one can examine. Figure 3 shows worldwide fatalities for people travelling in airlines over the last sixty years. It does not include the effects of terrorism or the deaths of third parties on the ground. The long term trend of total fatalities is slowly downwards; over the last twenty years it has varied between about 400 and 1800 people, averaging about 1100 a year. Over the sixty years, air travel has expanded dramatically, typically doubling every decade. Mid-air collisions are a small fraction of the total: the last one in Europe – at Überlingen – was five years ago.

Figure 4 shows a comparison of the trend in British – not worldwide – accidents from Evans (2005). Again using figures for Great Britain, the 1995–2004 average rates of fatality per billion passenger kilometres, across the passenger carrying modes, shows that air has a good relative historical record – Table 1 – compared to the other modes. This safety picture is not universally endorsed, eg compare the most recent (2007) Rail Safety & Standards Board.

3. Decision-making on Safety Improvements

But what has produced the progressive improvements in safety? And can the system go on improving safely? These are actually two very different questions. The first is answered by examining a variety of important safety changes as they actually happened. But before that, it is necessary to cover some aspects of ATM safety decision-making and the systems theory underpinning ATM safety.

Formal theories of safety appraisal – particularly decisions about safety in the context of the best use of resources – play a part in ATM decision-making, but different safety decisions have different contexts (Evans, 2005). It is important to keep asking three questions about safety decision-making.

> First, what motivates the need to make decisions (other than decide 'Do nothing')? Why is there a need for the people in charge to make a decision? Why did they decide that 'Do nothing' was not the best option?
>
> Second, what are the analytical/political decision processes that are followed? What data has to be gathered? What quantitative assessments have to be constructed? What comparisons have to be made? What are the actual decisions? Who has to be convinced? If there is a 'Go' or 'No Go' choice, is 'Further work' an option?

> Third, what are the resource implications? What are the constraints? Who has to be persuaded or instructed to spend money on new equipment/training?

Moreover, when asking these questions, a cynic might keep some things in mind. The first is formally termed 'Post hoc Rationalisation'. This means re-writing history to make things appear better: eg a decision that was politically forced is translated into one that was really the process of taking wise decisions. The second is 'Public Relations', for a variety of motives. The media are known to be generally enthusiastic about exaggerating bad things, so occasionally the story that is told to them – and hence to the public – is a sugar-coated version of reality. And sometimes people manage to convince themselves that they were not really responsible for something that went wrong: 'Mistakes were made (but not by me)'.

Many safety-related decisions are not in fact 'big decisions'. The system performs well because of the sum of all the 'small' engineering and operational decisions made by aviation professionals. People want to make things safer. Professionals learn from mistakes – the UK aviation industry is generally very good at keeping data on hazardous incidents.

Thus safety, as well as effectiveness, improves over time by what might be called 'Good Systems Engineering'. This includes elements such as:

> rigorous planning of procedures for design inspection and review
>
> quality assurance based on a wide range of targeted tests
>
> continuous evolution by adaptation of products already in widespread use
>
> deliberate over-engineering

This particular set of good practices actually comes from a famous paper 'How did Software get so Reliable without Proof?' (Hoare, 1996) about software development. Thirty years ago, quite a lot of people believed that aviation would suffer very badly from computer hardware and software failures – but it has not happened in aviation or other industries to the extent they feared (but see CSTB, 2007). Could future ATM systems be more susceptible to such problems?

4. ATM Safety – Systems Theory

This leads on to the systems theory underpinning ATM safety. A useful way of thinking about potential ATM accidents is to construct two broad categories according to the kind of technological and human system structures that are being employed to ensure safety. These are called tightly- and loosely-coupled. These terms originated with Weick (1976), and were subsequently used by Perrow (1984) to analyse accidents.

Perrow in fact defined two important dimensions: interactive complexity and loose/tight coupling:

> Interactive complexity refers to the presence of unfamiliar or unplanned and unexpected sequences of events in a system, either not visible or not immediately comprehensible.

> A tightly-coupled system is highly interdependent, with each part of the system being tightly linked to many other parts, so a change in one part can rapidly affect the status of other parts. So tightly-coupled systems respond quickly to perturbations—but this response may be disastrous.
>
> Loosely-coupled systems have less tight or fewer links between their parts, so they are able to absorb failures or unplanned behaviour without destabilization

A tightly-coupled design generally uses traditional engineering methods, with bits of electronic kit, aircraft construction, software, etc. Tightly-coupled systems can survive failures, but only if that kind of failure has been anticipated and provided for in the original design. Designers of tightly-coupled systems must therefore invest effort and thought into anticipating failure modes and providing safety features to permit survival and recovery. In contrast, loosely-coupled systems tend to accommodate failures through adaptive responses.

A loosely-coupled system allows some 'play' in the system stabilizing (negative) feedback loops—a little over-correction, followed by some under-correction. Loose systems are more adaptable, have more tolerance for error, but can have much longer reaction times. If what happens in one part has little impact on another part, or if everything happens slowly, eg on the scale of human thinking times, the system is not tightly-coupled. Loosely-coupled systems tend to be open and continually interacting with the outside environment.

It would be dangerous to construct a safety-critical system with both interactive complexity and tight coupling. In such systems, an apparently trivial incident can potentially cascade in unpredictable ways that cannot be remedied, and hence produce severe consequences. But ATM does not fall into this category (eg Marais et al, 2004). In general, much of ATM system design is deliberately de-coupled in order to increase safety. Thus, large minimum separations are required between aircraft, so that mistakes by pilots or controllers can be remedied; hence loosely-coupled. This is in a system containing independent and engineering-redundant safety defensive layers.

It must also be noted that system safety performance necessarily depends on organisational safety performance. Sorensen (2002) notes:

> "There is a widespread belief that safety culture is an important contributor to safe operations…The commonly accepted attributes of safety culture include good organizational communications, good organizational learning, and senior management commitment to safety. Safety culture may be particularly important in reducing latent errors in complex, well-defended systems."

Safety culture aspects are seen as increasingly important for European ATM, eg Eurocontrol (2006).

Some sub-systems of the ATM System are designed to be tightly-coupled. The key element is that the range of expressed 'failure modes' is comparatively limited and well defined. Thus, the sub-system acts in a 'programmable' or routine fashion (with specific designated functions). These kinds of approximately tightly-coupled systems would include navigation of well defined route systems, altimetry and instrument landing systems. In such cases, Human Factors 'failures' can be sufficiently regular in nature to permit a simple accident model to be used. For example, it may be possible to measure the frequency of a straightforward error in inputting flight level or North Atlantic track data into an aircraft computer. The operation of these kinds of tightly-coupled designs can therefore usually be modelled quantitatively.

In contrast, loosely-coupled ATM sub-systems would include the pilot flying the aircraft away from airport runways and ATC/pilot interactions in sectors. Loosely-coupled ATM designs use much more complex information sources. For example, the controller's job requires visualization and situational awareness skills.

The ATM system's safety defensive layers combine a variety of tightly-coupled and loosely-coupled sub-systems – Figure 5. Together, they act systematically to reduce mid-air collision risk. The purpose of the system layers is to reduce the 'end product risk'. Thus, each defensive layer should scale down the probability of a potentially hazardous situation. This leads to two kinds of qualitative/quantitative safety model:

> Tightly-coupled models — Accident risk is a function of specific failures, e.g., gross navigational errors or a restricted set of Human Factor failures occurring comparatively regularly. Risk can be numerically quantified in terms of a limited number of key failure modes using collision risk models.
>
> Loosely-coupled models — Safety is provided through a structure of defensive layers: risks occur if these layers perform poorly and do not filter out potentially hazardous situations. In some circumstances, risk can be roughly numerically quantified, based on past defensive layer performance.

'Collision risk model' – CRM – here means analytical frameworks on which are 'hung' empirical and statistical data about rates at which errors occur, recovery mechanisms, and failure probabilities. These models are effective because they involve changing one or two sub-system parameters, and use probabilities of gross error events: they do not usually involve major changes to controller or pilot tasks. The earliest models were constructed to deal with aircraft simply relying on accurate navigation to avoid other aircraft. The collision risk could be estimated as the product of the frequency of a gross navigational error, ie bringing the aircraft across another flightpath, and the probability that this crossing aircraft would then pass very close to the aircraft on that flightpath (Reich, 1966).

5. Tightly-coupled Sub-systems: Separation Minima

It is easy to list some successful (tightly-coupled) CRMs (details of the references are given in Brooker (2006a)):

> Navigation beacon defined routes (ICAO, 1976)
>
> Longitudinal North Atlantic separation (Brooker and Lloyd, 1978)

North Atlantic Track System (Brooker and White, 1979)

Radar separation (Sharpe, 1991)

Precision Runway Monitor (FAA, 1991)

Vertical separation (Harrison and Moek, 1992)

Area navigation parallel routes (DNV, 1997, 2003)

The end result of these studies is that various separation standards – separation minima is the formal phrase – were reduced. A very important recent example – implemented several years after the risk sums were done – was in fact the reduction in the vertical separation standard from 2000 feet to 1000 feet for aircraft flying above 29,000 feet. In most cases, the change examined is of a single operational parameter.

But *why* were these separation minima reduced? What was the motivation for this? The main motivation is to deliver operational benefits from investment in improved equipment. Better navigational kit on the aircraft or new radars on the ground enable the elimination of some of the 'deliberate over-engineering' referred to earlier. Thus, the aim is to introduce new technologies and ways of operating in what can reasonably be justified as safe system improvements. This is very different from the situation where new technology is introduced simply to achieve some safety benefit. (These 'safety benefit cases' do exist in ATM: the obvious examples are STCA and TCAS, which are examined later.)

So what would be safe system improvements? Figure 6 shows the key arguments. The first question asks about the acceptable risk in order to derive the Safety Target – the abbreviation TLS, target level of safety', is often used. The two components are the way that risk is to be measured and a *value judgement* about acceptability. But somewhere in the process there has to be a value judgement about accidents and deaths, eg about the rate of safety improvement.

Establishing Safety Targets requires some quantitative measure of risk. For ATM, the metric used in the UK has been 'fatal aircraft accidents per 10^7 aircraft flying hours'. The choice of the number of aircraft accidents rather than the number of deaths was made because ATC handles aircraft rather than individual passengers. (A constant safety rate would correspond to an increasing number of people killed each year.) A fatal accident is one in which at least one person in the aircraft is killed. Aircraft flying hours matches 'exposure' to the ATC service, which is provided over the duration of the flight.

The next step is to determine the safety target to be used in system design, against which prediction of the effects of system changes can be compared. The method chosen is to extrapolate safety performance – which historically has improved over time – to some 'design year' in the future. Figure 7 illustrates the method – the statistically fitted trend line is of negative exponential or similar form, although tending to flatten out in recent decades. Boeing Commercial Airplane (2007) shows the flatness of the statistical trends and the tightness of collision risk targets. (For example, over the last decade, the Boeing statistics show there were 89 fatal aircraft accidents involving commercial jets, of which two resulted from mid-air collision.)

The trend line method essentially maintains current overall trends towards safety performance. It must be stressed that the target is based on a progressive improvement on achieved performance rather than some 'absolute' figure.

The next step is to move from aircraft accidents in general to targets for mid-air collisions. This process is 'risk budgeting' – essentially the setting of minimum design targets for the contributory types of accident. The actual allocations of risk between the contributors generally bear some relationship to the present estimated rates. Air traffic services risks are a fraction of the total aviation risk. En route collision risk is then a fraction of the air traffic services risk.

An example of a safety target is the ICAO figure of '1.5×10^{-8} fatal aircraft accidents per flying hour' as the rate corresponding to mid-air collisions – for any reason and in any spatial dimension – in en route flight in controlled airspace. This is a target for *total system design* to an assured level of safety, ie all (*sic*) types of failure, mechanical, procedural and human, which generate a risk of collision will be accounted for.

The second question in Figure 6 asks about the actual risk that would follow a change. To quantify the future risk level, it is first necessary to develop a sufficiently comprehensive model of the processes and factors that contribute to this risk level. The key mechanisms that generate risk consequences need to be established. The significant causal factors and the associated risk probabilities must all be examined, albeit that not all can be quantified.

The model then has to be used to predict risks. *Data* on the risk mechanisms – equipment failure, human beings' failure rates – needs to be input into the model. These sorts of data may sometimes be immediately available, more often measurement exercises will be needed, but in key areas it may well be difficult or impossible to collect data sufficient to establish adequate statistical confidence.

The third component is validation. Is the model OK? Is the data input OK? How can its accuracy be tested? These can be exceedingly difficult tasks. Much of the problem is with the extremely tight targets that are now placed on aviation safety.

Validation is inherently a major problem. Historically – 40 or 50 years ago – the focus was on equipment failure modes, so that low risk rates relied upon equipment redundancy and monitoring. Today – and presumably in the future too – the focus is much more on abnormal events, in which human factors (HF) can play a major part.

6. Loosely-coupled Sub-systems: STCA and TCAS Decisions

Tightly-coupled sub-systems can potentially be modelled with some confidence, if certain conditions (mainly regarding data) are met, ie accident rates can be expressed in terms of sub-system failure rates. But what about loosely-coupled parts of the total system? The system relies on safety defences in depth—a multiplicity of formal, technical and human safety defensive layers—to deliver the necessary safety. But they cannot usually be modelled quantitatively with great precision. So

the safety target and modelling arguments will not work: the problem is that it is not possible to estimate the safety consequences of loosely-coupled elements with precision (Brooker, 2006a). There are just too many options, too large a potential for adaptive response and flexibility, too many probabilities to estimate, and not enough 'accurate' data available. It is actually easy to do the sums by standard hazard analysis techniques, but the problem then is assessing how good the sums are; what kind of precision can be attached to risk estimates?

Simply making lots of 'cautious' assumptions generally tends to produce over-pessimistic risk estimates, and hence is of little value for safety decision-makers. Does 'expert judgement' solve the problem? To quote Moray (1990) (in response to Dougherty (1990)):

> "The use of 'expert judgement' is a polite name for 'expert guesses', and we do not have data to validate the accuracy of the guesses."

Safety for loosely-coupled operational sub-systems is improved – purportedly ''to meet safety targets'' – by an on-going process of safety feedback plus the introduction of additional safety-related defensive layers and engineering redundancies, eg STCA, TCAS, error-free aircraft Flight Management System databases, etc. The key thing to ensure safety is that the ATM safety layers work effectively enough together to produce the necessary corrective action. For example, there is a need to focus attention on circumstances and geometries when STCA and TCAS do not provide large amounts of extra protection or when the geometries/velocities mean that they induce risk.

But how did anyone decide to introduce loosely-coupled operational sub-systems like STCA and TCAS? Different countries have different decision processes – the following is the UK experience.

First, STCA in the UK: in the UK's National Air Traffic Services (NATS) version of STCA, a computer system continually monitors secondary surveillance radar (SSR) data and alerts air traffic controllers if it detects a situation where two aircraft are in danger of approaching too close to one other. STCA is concerned with potential conflicts in projected flight paths (Figure 8 – the overlapping discs for cautiously projected flightpaths). The goal is to provide a warning – with special symbols on the controller's radar display – around 90 to 120 seconds before the Closest Point of Approach (CPA) of the two aircraft. This gives them time to redirect the aircraft if they judge it necessary. STCA alerts do not imply specific mandatory action by the controller. He or she is presented with the extra information as part of the normal air traffic control task.

The algorithms in the STCA computer software are specifically tailored for the varieties of airspace and separation rules. The STCA software contains a large number of parameters, whose values have to be fixed by extensive safety testing.

STCA does not 'know' the intentions of the pilots or air traffic controllers who may be aware of a potential conflict and already be taking measures to avoid it. As STCA must make cautious predictions, there are necessarily nuisance alerts as well as genuine alerts. There is a trade-off between genuine and nuisance alerts: if the

software eliminated all the nuisance alerts, then it would also fail to identify many genuine alerts. But if there were too many nuisance alerts, it would be difficult to maintain the controllers' confidence in STCA. The right balance has to be struck.

So was STCA introduced through some rational process involving safety targets and benefits? No, it was not. Twenty years ago, there were some serious incidents involving aircraft coming too close together, at much less than the separation minimum. These incidents were known then as Airmisses, now as Airproxes. One involved a blunder by a controller sequencing traffic from two of the holding stacks north of Heathrow: the aircraft flew towards each other.

The Chairman of the Civil Aviation Authority (CAA) and the head of NATS had a conversation about the Airmiss. It went on the lines:

> "Controllers sometimes make these very bad blunders. What's stopping a mid-air collision?"
>
> "ATC Supervisor plus chance."
>
> "That's a guarantee?"
>
> "No."
>
> "There must be something else."
>
> "We could try implementing the STCA in the ATC computers."
>
> "You do that - and soon."

The NATS STCA system became operational for parts of UK en route airspace in 1988. An essential ingredient for its introduction was the bringing on stream of a new generation of secondary surveillance radars with much better accuracy and performance. There had been some in-house research going on into STCA, but it had not been top-priority for implementation until the Chairman had made his views very clear. The research, development and implementation costs included the optimisation of the software (eg re nuisance alerts), design of screen displays and interfaces, production of training packages, and individual STCA training briefings and simulations for several hundred controllers – costing several (but not tens of) million pounds in total.

TCAS actually works on very similar principles to STCA. TCAS is an aircraft system using SSR transponder signals to provide advice to the pilot on potential conflicting aircraft that are also equipped with SSR transponders. It operates independently of ground-based equipment. TCAS produces Traffic Alerts (TAs) and vertical Resolution Advisories (RA): Figure 9 shows TAs in yellow and RAs in red.

Based on the horizontal and vertical closing rates, TCAS calculates dynamic protective volumes around its aircraft. If the closing intruder is assessed as a threat, then a TCAS system proposes an RA to the pilot as a Vertical Avoidance Manoeuvre. The system can coordinate its RA with the intruder aircraft, if it can generate an RA, so that the manoeuvres are complementary. Corrective RAs require the pilot to change the flightpath of the aircraft; preventive RAs require the pilot to keep the aircraft on that flightpath. TCAS's RAs are generated much nearer

to the predicted CPA – Closest Point of Approach – than are STCA alerts. Typical threshold times are between 15 and 35 seconds before predicted CPA, ie they are much closer to the CPA than STCA alerts to controllers.

So was TCAS introduced through some rational process involving safety targets and benefits? No, it was not. It followed a political decision by the USA. One of the main causes was the Cerritos, California, mid-air collision in 1986. An Aeromexico DC-9 with 64 passengers collided with a private Cessna aircraft carrying a family of three. The DC-9 crashed into a neighborhood and destroyed18 homes and killed 15 people on the ground.

MIT (2007) provides some history of mid-air collisions in the USA. It comments that "mid-air collisions have the effect of raising public awareness and causing a great deal of interest and pressure from the press and the public to 'do something'. Early versions of TCAS had in fact been under development for at least a decade before the Cerritos accident (Williamson and Spencer, 1989).

In December 1987, the Congress of the USA enacted Public Law 100-223, which 'requires the administrator of the Federal Aviation Administration (FAA) to complete the development of the Traffic Alert and Collision Avoidance System (TCAS II) and ensure that it is installed on all airplanes with more than 30 passenger seats by December 1991'. Thus, TCAS has been mandatory in USA airspace since 1991. It became mandatory in Europe in 2000, and there has been an ICAO world-wide mandate operating from 2003. ICAO concluded that the use of TCAS would reduce markedly the risk of collision. ICAO recognised that TCAS is not a panacea: it cannot resolve all possible collisions; it may increase some risks of collision.

It took quite a long time to move from the initial USA Public Law to world-wide introduction. Airlines were in fact buying TCAS kit well before it was mandated. Was there some kind of safety cost benefit analysis? There were some attempts to do this, but the focus on what is termed the 'Risk Ratio':

> *Risk Ratio*: the net improvement in safety arising from TCAS implementation. If the probability of a Near Mid-Air Collision is P(NMAC), then the ratio of P(NMAC) when TCAS is used to P(NMAC) without TCAS is the Risk Ratio. Properly-used TCAS will successfully resolve most potential mid-air collisions, but some will not be resolved; and an additional fraction will be 'induced' (as some non-critical encounters are converted into critical ones).

Carpenter (2004) is a good source for the history of this work. Risk Ratios depend on the kinds of conflicts that occur. As there are very few mid-air collisions, potential future conflicts have to be estimated by simulating realistic aircraft encounters. These conflicts use real traffic events and then vary their parameters realistically, eg seeing the effects of putting aircraft closer to each other when some manoeuvre occurs. The key point is that the Risk Ratio is essentially an experimental measurement of projected performance, not a simple output from purely mathematical calculations or computer science. The development work through States and ICAO over the decades has focused on improving the Risk Ratio and reducing the induced risks.

Interestingly, the final stages of the USA TCAS mandate did involve considerations of the kind envisaged by Evans (2005). The final mandate (FAA, 2003) was specified in terms of airplane weight and performance characteristics, and had the consequence of covering larger cargo aircraft. [NB: TCAS installation costs of the order of €200,000 per aircraft and the aircraft operator incurs significant ongoing maintenance costs.] The justification for the inclusion of cargo aircraft included a safety cost benefit analysis that estimated the frequency of different kinds of mid-air collision, the costs of installation, the costs of lost aircraft, plus a valuation of 'avoided deaths'. To quote: "…it is assumed that a midair collision will result in fatalities for all passengers and crew, rather than some percentage attributed to various classifications of injuries. The value per averted fatality is estimated to be $3.0 million." (This figure is taken from USA Department of Transportation recommendations.) Thus, while the initial TCAS policy decision was political, the final details incorporated formal risk analyses and safety costings: the weighing of costs and benefits did affect the installation decision for cargo aircraft.

7. Future ATM Systems

The introduction of TCAS was essentially the last safety layer in the current ATM system. The system is expected to continue evolving in future years, with safety lessons being learnt and engineering systems being improved. So why are 'new paradigms' being suggested? What indeed is meant by a new paradigm? Paradigm was used originally in the History of Science to refer to a theoretical framework. Researchers in many different fields now often see themselves as developing new paradigms. Paradigm is a buzzword or a key concept, depending on one's attitude to adapted words (paradigm actually derives from a Greek word meaning 'to compare').

For ATM systems, the main current safety paradigm – the prevailing view of things – is that air traffic controllers work to prevent mid-air collisions. A new paradigm would move substantial parts of this control workload to either or both aircrew and computer assistance, usually requiring a considerable enhancement of the data available for decision-making. Why move the workload? The need to do this reflects peoples' views that the present controller-based paradigm is in one or several ways 'reaching its limits', so its continued evolution will not solve future problems effectively.

So, people have big ideas about what the future ATM system could look like. If people in the aviation industry are asked what are the issues for future ATM, the kind of ranked list one gets is on the lines of:

Improving Safety

Environmental Impact

Meeting traffic demand

Reducing costs of ATC provision

Reducing costs of ATC constraints (eg delays, excess route length)

If you are going to make changes, it is vital to develop a clear picture of the intended end-product. How to establish the characteristics of that picture is the problem.

Answering technical questions is important in getting to that picture – but it is vital to try to work out what really are the full range of questions that need to be answered. It is essential to frame the 'Key Tests' that will enable the full picture of the future ATM system to be established. In analysing a potential major system change, it is important to try to break down the problems to be resolved into several distinct types of issue. A simple formal division into five key test areas, not in any particular order of importance, is:

Safety Credibility

Technological Feasibility

Operational Concept

Benefits and Costs

Transition Path

Given the title here, the focus is on some aspects of just three of these: Safety Credibility, Operational Concept, and Transition Path.

Safety is always the top priority in the introduction of new aviation systems. Key questions would be:

What sorts of quantitative tests would be needed to prove that a new system is sufficiently safe?

What safety management performance is necessary before the safety regulator would be convinced?

Does the new system pass the safety regulator's tests?

The present system relies for safety on a complex combination of hardware, software and 'peopleware'. It has a great deal of engineering-type redundancy built into it. The assurance that the present system delivers safety is largely because it has built up layers of defensive safety barriers, often independent of each other.

Even 'simple' sub-system changes, from a whole system point of view, require extensive data collection and analysis before they can be accepted. This would be a much harder job with linked and integrated communications, navigation and surveillance systems. Automation routes require, as a foundation, a complete logical structure for ATC decision-making – against all eventualities. Then the hardware and the software design would have to be demonstrated as providing the necessary reliability and safety critical integrity.

The phrase Operational Concept means no more than a clear picture of <u>how</u> the future system would operate:

Who uses what information to do what things?

What are the responsibilities, and on whom do they rest?

How are decisions made?

The present operational concept has evolved over the decades. In safety terms, it is sequentially ‘overlaid’ on to the immediately previous concept, rather than being a clean sheet redesign. Many in commercial aviation believe that the current system is too rigid, with flights not always being able to use preferred routeings and profiles (often termed ‘free routeing’).

There has to be a realistic transition path from the present system. Does the transition path make sense in both safety and business terms? The international dimension is a major factor. Formal processes involved in international development, testing and certification are lengthy. TCAS and the major reductions in separation minima took 10-15 years to reach substantial operational implementation.

At the heart of the present Operational Concept are the controller and the work that has to be done in handling aircraft. Control workload is a multi-dimensional concept encompassing both the difficulty of tasks and the effort – physical and mental – that has to be brought to bear, plus a personal dimension – Figure 10.

A common characteristic of the different avenues for reducing controller workload is that they must somehow produce more time for the controller to make necessary decisions. The aim is not to make theoretically "optimal" decisions, but to enable the controller to make good decisions, but most important to prevent bad decisions from ever being made. The controller has to act in real time, so the systems designers have to make sure there is as much of this available to the controller as is necessary for the critical tasks.

One crucial question is always the extent to which a system improvement changes the air traffic controller’s tasks. Airspace capacity is in many instances – very much so in Europe – the leading constraint on traffic throughput. This capacity is largely determined by the acceptable workload on sector controllers. To gain capacity in a block of airspace, the controller tasks have to be made ‘easier’ – by computer assistance – or some tasks have to be eliminated – Figure 11. These eliminated tasks could be automated or transferred to the flight deck, ie to the pilot. In a complementary fashion, to prevent the pilot from being over-burdened, aircrew tasks need to be eliminated or cut down. Hence, some pilot voice communication would be replaced by electronic data communication. But tasks added to the flight deck need to have all the on-board systems to perform them safely.

Controller and pilot errors are by far the main ‘causes’, using the word vaguely, of hazardous ATM incidents. A vital tool in learning about ATC safety is data from Airproxes. An Airprox is formally defined as ‘a situation in which, in the opinion of a pilot or a controller, the distance between aircraft as well as their relative positions and speed have been such that the safety of the aircraft involved was or may have been compromised’.

Figure 12 shows some statistics – unofficial – from recent UK Airprox Board (UKAB) Reports. The aim was to find any with ‘root causes’ beyond immediate human error/misjudgement etc. Thus, the Airspace category focuses on incidents with airspace design and/or operational procedures that raise questions about ways of de-conflicting traffic features. If these issues could be satisfactorily dealt with, then this

would prevent reoccurrences of this kind of Airprox in the particular locality. 'Incorrect readback' means that the pilot incorrectly read back ATC instructions and this was undetected by the controller. The three 'Technical' incidents are those where equipment was involved. They were consequences of sudden cabin decompression, the misreading of a navigational chart, and the failure of flightdeck procedures to detect an incorrect setting on the flight computer.

So, for a markedly different ATM system to deliver safety, the entire system (people, equipment, procedures) is designed to help prevent human error and capture the inevitable errors before they result in a collision. To quote Kim Cardosi: "People make mistakes – even the most intelligent, well-trained, conscientious, and well-intentioned people make mistakes." In the 2007 Wimbledon Final, Roger Federer Double-faulted or made an Unforced Error roughly every six minutes. The average USA physician kills two people by accident during his/her career (Dekker, 2006).

Some examples of controller and pilot mistakes, from Airprox Data (Brooker, 2005c), are:

> "Controller did not monitor aircraft 1's progress – he was bandboxing and had been concentrating on traffic situation elsewhere.
>
> ATC occupied with other traffic, did not spot high descent rate.
>
> Controller gave '*erroneously and essentially unforced descent instruction*' to aircraft 2.
>
> The aircraft 2 crew read back the wrong heading and level instructions, which went undetected by the controller."

Future systems designs have to prevent these kinds of errors happen or to deal with them safely.

8. New Paradigms

So what is being proposed? SESAR is Europe's 'Single European Sky Air traffic Research system'. NextGen is the USA's 'Next Generation Air Transport System' [previously known as NGATS]. SESAR and NextGen are developments targeted at post 2020. Neither of them is fully developed, so their current descriptions still include different options for achieving safe and cost-effective systems. In particular, the degree to which control tasks are transferred to the aircraft/aircrew and to ground automation (and the transitional steps involved) are questions which will need to be definitively answered. At present, a range of options is still being explored.

The common SESAR and NextGen vision is to integrate and implement new technologies to improve air traffic management (ATM) performance – a 'new paradigm'. SESAR and NextGen combine increased automation with new procedures to achieve safety, economic, capacity, environmental, and security benefits. The technical systems do not have to be identical, but must have aligned requirements for equipment standards and technical interoperability.

Figure 13 shows a very simplified picture of NextGen. It creates a 'cooperative surveillance' model for civil aircraft operations, where aircraft are constantly transmitting their position, flight path intent, and other useful aircraft parameters (ADS-B – Automatic Dependent Surveillance-Broadcast). Both expect aircraft position to be determined using a satellite navigation constellation, such as GPS or Galileo.

The basis for planning and executing system operations is an aircraft 4D trajectory, which is the aircraft path, three space dimensions plus time, from gate-to-gate, including the path along the ground at the airport. Data on the planned and actual trajectories are exchanged between air and ground. Digital satellite communication constellations report positions to ground facilities. All other communication is through this constellation as well. So, gate-to-gate 4-D trajectories are broadcast and, if necessary, so are voice communications. Thus, Communications, Navigation and Surveillance functions are much less ground-based than the current system. The operating concept for NextGen is summarized in Figure 14. Figure 15 is summarised from an early 2007 description of SESAR concepts (EC, 2007). Very recently, the SESAR Consortium has issued its 162-page ATM Target Concept (SESAR Consortium, 2007), which inter alia shows the complexity of the concept. It notes that:

> "The ATM Target Concept is not about one size/one solution fits all; it offers different concept features which can be tailored to the specific local needs to meet the local performance objectives and their evolution in the life time of SESAR."

SESAR currently appears to be more general than NextGen. The main difference shown in Figures 14 and 15 is that NextGen additionally uses automation through the Evaluator function to analyze these trajectories to ensure aircraft remain at safe distances from one another. The Evaluator uses the extra information and communication power to enable safe and efficient decisions to be centralised but system-wide.

Regarding transferring some tasks to pilots (independent autonomous decision-making), the terms ABS-B and ASAS are widely used. ADS-B is 'Automatic Dependent Surveillance – Broadcast' and ASAS is 'Airborne Separation Assistance System'. They are inter-related. Some brief explanations.

> ADS-B: aircraft periodically broadcast details of their position, altitude, velocity and other flight data via a digital data link.
>
> ASAS can transfer the responsibility of maintaining separation between two aircraft from ground ATC to the airborne side when feasible.
>
> > Aircraft transmit their position via a data link.
> >
> > These data are presented on the traffic display in the cockpit.
> >
> > ASAS alerts and advises the crew.
> >
> > Separation maintenance is the responsibility of the cockpit crew.

There are lots of different interpretations and schemes for the use of ASAS, some tactical, others more strategic. ASAS's safety depends on retaining large separation

minima (*inter alia* restricting traffic density), and introducing more sophisticated conflict detection algorithms and cockpit displays/automatic warning systems. So they are inter-related concepts – but basically rely on airborne text messaging.

A key jargon phrase in NextGen and SESAR discussions is 'Net-centric': 'A networked collection of capabilities that empower end-users to pull the information they require from any available source, with minimal latency to support the mission at hand'. An internet-like network carries a common real-time information set. This 'net-enabled information' has to be accessible, usable, and secure for all ATM system decision-making parties. Information is 'pushed' to known users, and available to be 'pulled' by new users. Everything operating in the system is part of NextGen. Aircraft are mobile 'nodes' within a larger information network. Aircraft both use and provide information, but are also able to route messages/information sent from another aircraft or ground source.

Will NextGen and SESAR be implemented in the way that people are currently envisaging? The investment side of things is a major challenge: even early cost estimates for the systems are in the tens of billions of Euro/dollars, and stakeholders will need to be convinced that the benefits outweigh the costs.

9. Safety of New ATM Paradigms

But safety is the paramount challenge. The final system must be safe, and so must all the transitional stages. The problems are with the ideas behind designing a safe system and those of proving it to be safe. If the system safety defences change their nature considerably, then where will the evidence come from to substantiate claims of a safe system? Political decisions, eg re TCAS, had a great impact on change in the past, but could politicians really be expected to be the main safety decision-makers on the acceptability of large-scale 'new paradigm' system changes?

This leads to a long list of both generic and specific questions:

- Can a safe 'new paradigm' ATM system be designed?
- Can an ATM system be implemented in safe stages?
- Can the interim and final ATM systems be proved safe?
    - Resilience against extreme events?
    - Human Factors?
    - Emergent Properties
    - Software/Hardware?
    - Safety philosophy?

NextGen has the potential to be very safe. Most current accidents and serious incidents are caused by a 'lack of reasonable intent' rather than equipment failures or software errors. Reasonable intent essentially means that aircraft are committed to sensible flightpaths. But sometimes a pilot decides to climb or descend for no strong reason; an urgent problem may hold the controller's attention to the detriment of new developing problems; airport staff can drive onto an active runway they believed was shut etc. These precursors of accidents arise out of workload, miscommunication, and lack of up-to-date information. NextGen could potentially eliminate these precursors by the common knowledge of 4D trajectories, safety checks through the Evaluator, and machine communication rather than voice messages.

But there are few publications as yet on the safety of NextGen. One very interesting one is by Andrews et al (2005): it is mainly intended to highlight where further work is needed. Two issues about extreme technical events they note are:

> Aircraft fly 4D trajectories using their Flight Management Systems (FMS). But sometimes the flightpath will not conform to the specified trajectory, eg because of engine failure, extreme turbulence, FMS performance limits. This 'control fault non-conformance' is a key fault type. What improvements in FMS performance will be needed?
>
> Ground computer systems covering large areas may fail or be shut down to respond to anomalous events ('outages'). The suggestion is that structured recovery planning can handle this 'troubling question'.

NextGen and SESAR must successfully deal with Human Factors issues. But this work is very much in its early stages. A sample of issues from Sheridan (2006) and Sheridan et al (2006) is:

- Who (human) or what (computer) has authority at different stages of flight
- What network information would be 'pushed' (mandatory display to human operators), what would need to be 'pulled' (explicitly requested), etc
- Problems of robustness, reliability and operator trust in computer decision support tools/control
- Control instabilities resulting from closed-loop time delays, eg due to ATC time-sharing of attention
- Operator error in 'automation mental modelling' and situation awareness of what the automation has done, is doing, or will do

An unknown is the extent of 'emergent properties', sketched in Figure 16: useful references are Johnson (2006) and Chalmers (2006). A good – and very relevant example here – of a weak emergent property is found in mobile phone text messaging. It started as a message service, allowing operators to inform all their own customers about things such as problems with the network. SMS (Short Messaging Service) was not initially meant to communicate from consumer to consumer. But Texting took off when it found its natural/full markets: teenagers attracted to pre-paid phones; *and* when cell phone users could send SMS to

someone on a different operator. Initially, some networks did not even charge for SMS.

An expert, defined as somebody who was very bright and really knowledgeable about a system, would see unexpected occurrences much less frequently than a routine performer. But there would be no way that the expert could *prove* there would be no emergent properties when systems were changed significantly. Human experts, even Human Factors (HF) experts, are not omniscient. So, it seems very unlikely that even the most expert system modellers and HF experts could guarantee that their understanding of the final ATM system would encompass all potential emergent behaviours. NextGen and SESAR are likely to exhibit emergent properties because the responsibilities of intelligent people in the ATM system are changed considerably *and* there are new tools for them to use – and adapt.

Software and hardware become much more important in the new paradigm systems, because they are much more tightly-coupled and safety-critical in decision-making than the present ATM system. But it is a fantasy to believe that either software or hardware can be proved 'correct' (whatever that may mean – eg see Cohn (1989), MacKenzie (1991)). Modern thinking (eg Thomas, 2004) warns that the safety critical software industry 'falls far short of the standards expected from a mature industry developing very complex and highly critical systems', in particular:

> "…it is impractical to have sound evidence that a system has achieved a pfh [probability of dangerous failure per hour [pfh] of 10^{-5} or lower and that the safety assurance of safety-related systems is therefore inevitably a matter of judgement.
>
> To show that some system met the targets for SIL 1 [Safety Integrity Level 1] (10^{-6} <pfh<10^{-5}) would involve testing the system continuously for more than ten years, under operational conditions, with no unsafe failures and no modifications to the system (Littlewood and Strigini, 1993)."

'Safety Philosophy' is a very hard problem at the heart of the safety assessment process for a new paradigm. To recall, the safety target approach works well for sub-systems where the changes mainly affect simple parameters of sub-systems and there is the opportunity to understand and collect data on a limited number of error modes. There is no good evidence that the same approach works for much more loosely-coupled systems including a variety of human interactions and responses. But ICAO and Eurocontrol currently have some strange safety regulation policies.

The best example of these is that Andrews et al analysis of NextGen systems get two orders of risk reduction by taking account of the safety benefits from TSAFE (an intent-based short-term conflict alert system) and TCAS (airborne collision avoidance). But ICAO and Eurocontrol safety policy would consider aspects of these calculations invalid, because currently the policy is that a system has to be 'proven safe' without the use of these aids. It is not obvious why such a 'policy' is needed, given aviation's tradition of rational safety testing. There has to be a willingness to recognise the fact that STCA and TCAS are intrinsic to the ATM system delivering its current high safety performance, rather than considering them as little more than 'optional extras'.

Figure 17 sets out the rationale for Ground and Air Protection layers being included in hazard analysis. Future ATM System Safety depends crucially on both these kinds of protection. To ignore them in risk assessment is to prevent improvements that will reduce aviation deaths as traffic grows.

What *process* dangers could there be? The answer is you get *Stasis.* This was a word used by the ancient Greeks to mean many different things: civil war, arguments between factions, 'a stoppage'. Today, it generally means a cessation of progress or change. It would be worse than that, because one consequence would be a great number of safety and human factors analysts producing increasingly elegant mathematical models and unverifiable complex calculations – but which do not convince decision-makers about practical 'Go/No go' choices. The danger to the Safety Innovation Process is that there is a great deal of highly intellectual activity, with lots of entertaining conferences and seminars, but no conclusive outputs. Compare 'The Glass Bead Game' (Figure 18).

10. Key Safety Assessment Messages and Ideas

So what are the key messages for Safety Assessment? What really does matter? The tough challenge is to turn a new paradigm into something real. But what has often worked in the past is unlikely to work generally in the future. Tightly-coupled models, such as CRMs, correspond to simple mechanisms and few assumptions, so validation is not a major issue for them, but loosely-coupled models produce results of largely unknown precision. The development of NextGen and SESAR will involve much more than tightly-coupled system changes. NextGen and SESAR change the nature of the ATM system considerably, so emergent properties will add to both the complexity of modelling and to an increased lack of confidence in quantitative risk predictions.

Some key safety lessons from the previous text are:

- Achieving and proving safety for NextGen/SESAR is an enormously tough challenge
- It will not be done by employing current patchwork of methods focused on tightly-coupled sub-systems
- Decision-makers need rational, evidence-based and realistic modelling for Risk Assessment of the Total ATM system

- Risk Assessment needs to cover system resilience, human/automation issues, software/hardware performance/ground/air protection systems
- Safety decision-making process must be systematically set out very early on
- Confidence building programmes re system design/resilience
- Openness to peer review

These are all very simple and obvious, but what would be the consequences?

Brooker (2007c) offers some ideas for 'Safety Innovation Strategies' that would meet these needs:

> Implement the final system in phases and prove each phase safe and resilient.
>
> *'Quantitative safety filte*r': test if the estimated safety level of each new ATM system is 'about the same' as the present ATM system. To prove that it is actually better than the present system is likely to be impossible: the problems are the difficulty of making precise estimates of very small risks, and saying with confidence that one group of risks is less than another group. Try to use CRM-type models wherever possible, and use data from 'Human-in-the-Loop' (HITL) simulations. [HITLs, often called real time simulations, put controllers and pilots into accurately simulated ATC environments. A typical large-scale ATC centre HITL is essentially a mock-up of a control room, and the controllers carry out the same tasks that they would for real traffic. Aircrew HITL are usually cockpit simulators, with the most advanced kinds being able to move the cockpit around in an imitation of climbs, turns, etc.]
>
> *'Seeded Errors Resilience':* HITL simulations can potentially be valuable in testing how *resilient* the system is to errors, blunders, etc. This mode of testing can be called *seeded errors*. The analogy here is a tank with 'impenetrable' armour. What actually sells tanks is that tanks in simulated real operations come out operationally unscathed from an attack by a large range of missiles and antitank weapons. Realistic HITL simulations that show resilience to them build confidence in the real system's resilience. So test that each new ATM system phase proves resilient in HITL simulations to a wide variety of 'incentivised' novel system challenges. As a start, the HITL simulation needs to be seeded with key features observed in past accidents and stressed further by increased traffic levels. Incentivising a group of challengers might mean paying them significant amounts of money for identifying potential system weaknesses. HITL simulations must be seeded with unusual, but realistic, failure modes.
>
> *'Emergent Properties Detection':* If unsafe emergent properties and new kinds of HF error are to be detected/corrected quickly in operation, there needs to be increased emphasis on automatic detection systems. Novel 'un-thought of' HF failure modes and general emergent properties must be detected early thorough much more rigorous monitoring than at present. Eg, if aircrew take on separation responsibility, so controllers do not routinely monitor flightpaths,

then need to have cockpit systems monitoring automatically for separation breaches & retain evidence of aircrew actions & equipment performance

These ideas for Safety Innovation Strategies would be a high cost process – but probably highly cost-effective option in terms of efficiently impementing safe systems.

Acknowledgement

I am very grateful to Professor Andrew Evans for his tough questioning and constructive criticisms of my draft texts.

BIBLIOGRAPHY

Adams, N. (2004). The Flyers: In Search of Wilbur and Orville Wright. Three Rivers Press.

Andrews, J. W., Welch, J. D., and Erzberger, H. (2005). Safety Analysis for Advanced Separation Concepts. 6th USA/Europe Air Traffic Management R&D Seminar, Baltimore. http://atmseminar.eurocontrol.fr/all-seminars/6th-usa-europe-atm-2005-r-d-seminar/paper_120/attachment_download/file

Boeing Commercial Airplanes (2007). Statistical Summary of Commercial Jet Airplane Accidents Worldwide Operations1959-2006. http://www.boeing.com/news/techissues/pdf/statsum.pdf

Brooker, P. (2002a). Future Air Traffic Management – Passing the Key Tests. Aeronautical Journal, 106(1058), 211-215.

Brooker, P. (2002b). Future Air Traffic Management: Quantitative En Route Safety Assessment Part 1 – Review of Present Methods. Journal of Navigation, 55(2), 197-211.

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

Brooker, P. (2003a). Control Workload, Airspace Capacity and Future Systems. Human Factors and Aerospace Safety, 3(1), 1-23.

Brooker, P. (2003b). Future Air Traffic Management: Strategy and Control Philosophy. Aeronautical Journal, 107(1076), 589-598.

Brooker, P. (2004a). Airborne separation assurance systems: towards a work programme to prove safety. Safety Science, 42(8), 723–754.

Brooker, P. (2004b). Consistent and up-to-date aviation safety targets. Aeronautical Journal (July), 345–356.

Brooker, P. (2004c). 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.

Brooker, P. (2005a). Airborne collision avoidance systems and Air Traffic Management safety. Journal of the Institute of Navigation, 58(1), 1–16.

Brooker, P. (2005b). Reducing mid-air collision risk in controlled airspace: Lessons from hazardous incidents. Safety Science, 43(9), 715–738.

Brooker, P. (2005c). STCA, TCAS, Airproxes and Collision Risk. Journal of the Institute of Navigation, 58(3), 389–404.

Brooker, P. (2006a). Air Traffic Management Accident Risk Part 1: The Limits of Realistic Modelling. Safety Science, 44(5), 419-450.

Brooker, P. (2006b). Air Traffic Management Accident Risk. Part 2: Repairing the Deficiencies of ESARR4. Safety Science, 44(7), 629-655.

Brooker, P. (2007a). Air Traffic Management Innovation: The Risks of *Stasis.* Air Traffic Technology International 2007.

Brooker, P. (2007b). NGATS: Strategy and Challenges. Navigation News, Jan/Feb, 12-15.

Brooker, P. (2007c). Air Traffic Management Safety Challenges. Keynote Address, 2nd IET System Safety Conference.

Brooker, P. (2008a). SESAR and NextGen: Investing in New Paradigms**.** To appear in the Journal of the Institute of Navigation.

Brooker, P. (2008b). The Überlingen Accident: Macro-Level Safety Lessons. To appear in Safety Science.

Carpenter, K. (2004). ACAS Safety Study Summary. Qinetiq Report for Eurocontrol. Available from: http://www.eurocontrol.int/msa/gallery/content/public/documents/Safety/ACAS%20Safety%20Studies.pdf

Chalmers, D. J. (2006). Strong and Weak Emergence. In The Re-emergence of Emergence. (Eds Clayton, P. and Davies, P.) Oxford University Press. http://consc.net/papers/emergence.pdf

Cohn, A. (1989). The Notion of Proof in Hardware Verification. Journal of Automated Reasoning, 5(2), 127-139.

CSTB [Computer Science and Telecommunications Board] (2007). Software for Dependable Systems: Sufficient Evidence? National Academies Press. (Eds Jackson, D., Thomas, M. and Millett, L. I.) Pre-publication copy. http://books.nap.edu/catalog.php?record_id=11923

Dekker, S. W. A. (2006). Doctors Are More Dangerous Than Gun Owners: A Rejoinder to Error Counting. Technical Report 2006-01 Lund University School of Aviation. http://www.lusa.lu.se/upload/Trafikflyghogskolan/TR2006-01_RejoindertoErrorCounting.pdf

Dougherty, E. M. (1990). Human reliability analysis – where shouldst thou turn? Reliability Engineering and System Safety, 29, 283-299.

EC [European Commission] (2007). State of progress with the project to implement the new generation European air traffic management system (SESAR) 2nd March 2007. http://ec.europa.eu/transport/air_portal/sesame/doc/0315_comm_sesar_en.pdf

Eurocontrol (2006). Understanding Safety Culture in Air Traffic Management. Safety Domain, DAP/SAF. http://www.eurocontrol.int/safesky/gallery/content/public/SafetyDomainSept06.pdf

Evans, A. W. (2005). Safety Appraisal Criteria. The Royal Academy of Engineering and Lloyd's Register Lecture on Risk Management.

FAA (2006). Wilbur and Orville Wright. FAA Wright Brothers Curriculum material. http://www.faa.gov/education_research/education/educator_resources/curriculum/wright_brothers/media/wbtext.pdf

FAA [Federal Aviation Administration] (2003). Collision Avoidance Systems; Final Rule. Department of Transportation Part III 14 CFR Parts 121, 125, and 129. Available from: http://a257.g.akamaitech.net/7/257/2422/14mar20010800/edocket.access.gpo.gov/2003/pdf/03-7653.pdf

Flight (1922). The Air Disaster. XIV (15), 210.

Hesse, H. (1943/2002). The Glass Bead Game: (*Magister Ludi*). English Translation. Picador.

Hoare, C. A. R. (1996). "The Role of Formal Techniques: Past, Current and Future or How Did Software Get so Reliable without Proof?" Proceedings of the 18th International Conference on Software Engineering (ICSE '96), IEEE.

Johnson, C. W. (2006). What are Emergent Properties and How do They Affect the Engineering of Complex Systems? Reliability Engineering and System Safety, 91(12), 1475-1481. http://www.dcs.gla.ac.uk/~johnson/papers/emergence.pdf

JPDO [Joint Planning and Development Office] (2007). Concept of Operations for the Next Generation Air Transportation System. *Draft 5,* 28th February, 2007, Version 1.2. Available from http://techhangar.jpdo.aero/

Littlewood, B. and Strigini, L. (1993). Validation of Ultra-High Dependability for Software-based Systems Communications of the ACM, 36(11), 69-80.

MacKenzie, D. (1991). The Fangs of the VIPER. Nature, 352, 8 August, 467-468.

Marais, K., Dulac, N., Leveson, N. (2004). Beyond normal accidents and high reliability organizations: The need for an alternative approach to safety in complex systems. In: Engineering Systems Division Symposium. MIT, Cambridge, MA. http://sunnyday.mit.edu/papers/hro.pdf.

Marais, K. and Weigel, A. L. (2006). Encouraging and Ensuring Successful Technology Transition in Civil Aviation. MIT ESD-WP-2006-07. http://esd.mit.edu/wps/esd-wp-2006-07.pdf

MIT (2007). The Story of Mode S: An Air Traffic Control Data Link Technology. http://mit.edu/6.933/www/Fall2000/mode-s/collisions.html#1986#1986

Moray, N. (1990). Dougherty's Dilemma and the One-sidedness of Human Reliability Analysis Reliability Engineering and System Safety, 29, 337-344.

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

Rail Safety & Standards Board (2007). Annual Safety Performance Report 2006. http://www.rssb.co.uk/pdf/reports/ASPR_2006.pdf

Reich, P. G. (1966). Analysis of long-range air traffic systems: separation standards. Journal of Navigation, 19, 88–98, 169–193 and 331–347.

Reuters UK (2007). Tennis-Match statistics for Federer v Nadal. July 8th. http://uk.reuters.com/article/tennisNews/idUKL0840015820070708

SESAR Consortium (2007). The ATM Target Concept. DLM-0612-001-02-00. SESAR Definition Phase, Deliverable 3. 27th September 2007. http://www.eurocontrol.int/sesar/gallery/content/public/docs/DLM-0612-001-02-00.pdf

Sheridan, T. B. (2006). Next Generation Air Transportation System: Human-Automation Interaction and Organizational Risks. Resilience Engineering Symposium**,** Juan les Pins, France, November 8-10, 2006. http://www.resilience-engineering.org/REPapers/Sheridan_R.pdf

Sheridan, T. B., Corker, K. and Nadler, E. (2006). Report on a Workshop on Human-Automation Interaction in NGATS**.** DOT-VNTSC-NASA-06-02. http://www.volpe.dot.gov/hf/docs/worshop-hai-sheridan.doc

Sorensen, J. N. (2002). Safety culture: a survey of the state-of-the-art. Reliability Engineering & System Safety, 76(2), 189-204.

SDG [Steer Davies Gleave] (2005). SESAME CBA and Governance. Assessment of options, benefits and associated costs of the SESAME Programme for the definition of the future air traffic management system: Final Report. http://ec.europa.eu/transport/air_portal/traffic_management/sesame/doc/2005_06_24_sesame_final_report.pdf

Thomas, M. (2004). Engineering Judgement. 9th Australian Workshop on Safety Related Programmable Systems (SCS'04). Brisbane. Conferences in Research and Practice in Information Technology, Vol. 47 (Cant, T., Ed.). http://crpit.com/confpapers/CRPITV47Thomas.pdf

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

Williamson, T. and Spencer, N. A. (1989). Development and Operation of the Traffic Alert and Collision Avoidance System (TCAS) Proceedings of the IEEE, 77(11) 1735-1744.

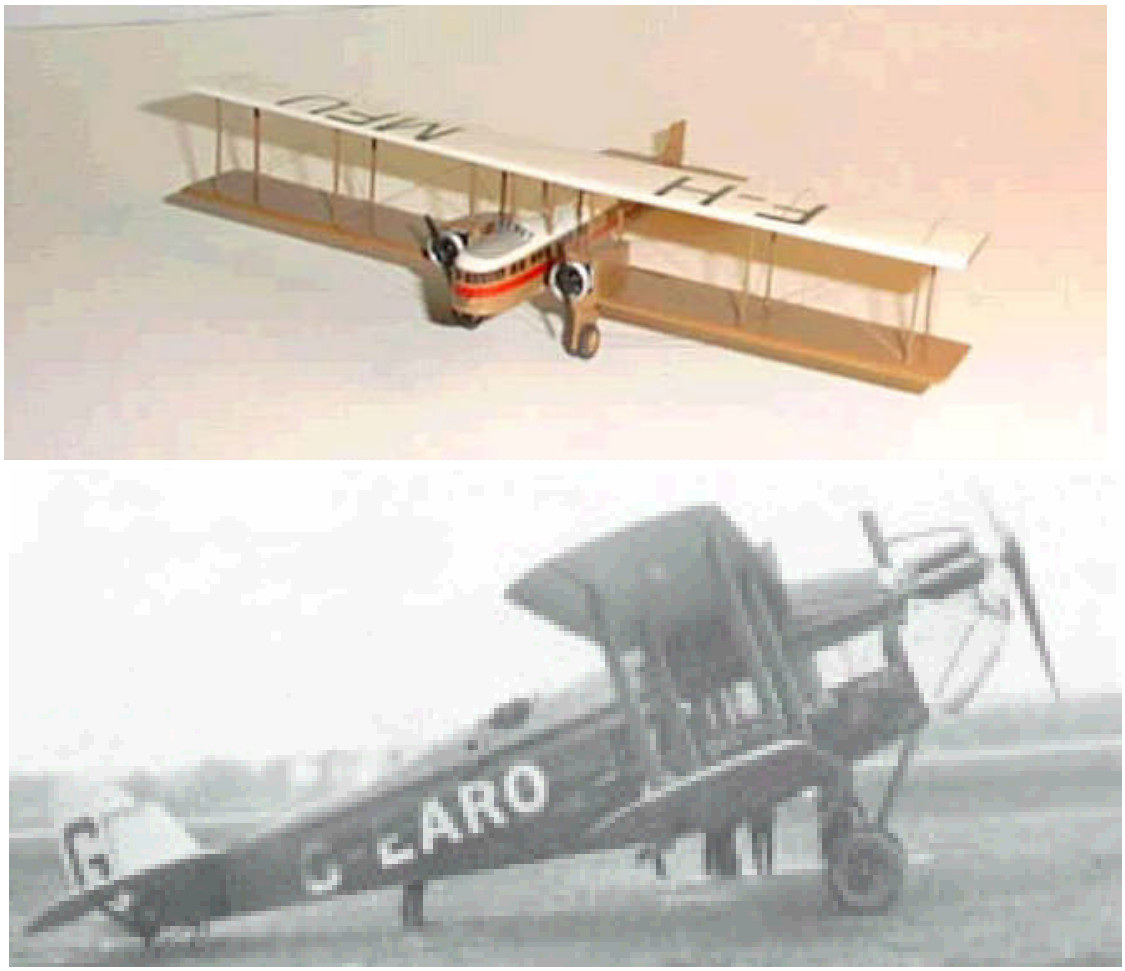

Figure 1. The First Mid-Air Collision – aircraft types involved

- Controllers and pilots – people are an *integral* part of the whole system
- Formal Rules for the control of traffic, eg minimum separations allowed between aircraft
- Radio Telephony
- Controlled Airspace – broken down into sectors handled by controller teams
- Flight Progress Information – flight plan computing
- Radar – processed SSR: displayed aircraft symbols, callsign and height information, passed down from aircraft transponders
- Computer Processing of radar and flight data.
- High Quality Aircraft Navigation – Point source beacons to INS [Inertial Navigation Systems] through to satellite-based aids
- Conflict Alert (STCA) – the computer processing system can analyse SSR tracks to predict if aircraft might come into close proximity and warn the controller by radar screen messages
- Traffic alert and Collision Avoidance System TCAS – on board collision avoidance system based on detection of other aircraft in the vicinity carrying SSR transponders

Figure 2. The Current System's Evolved Safety Defences

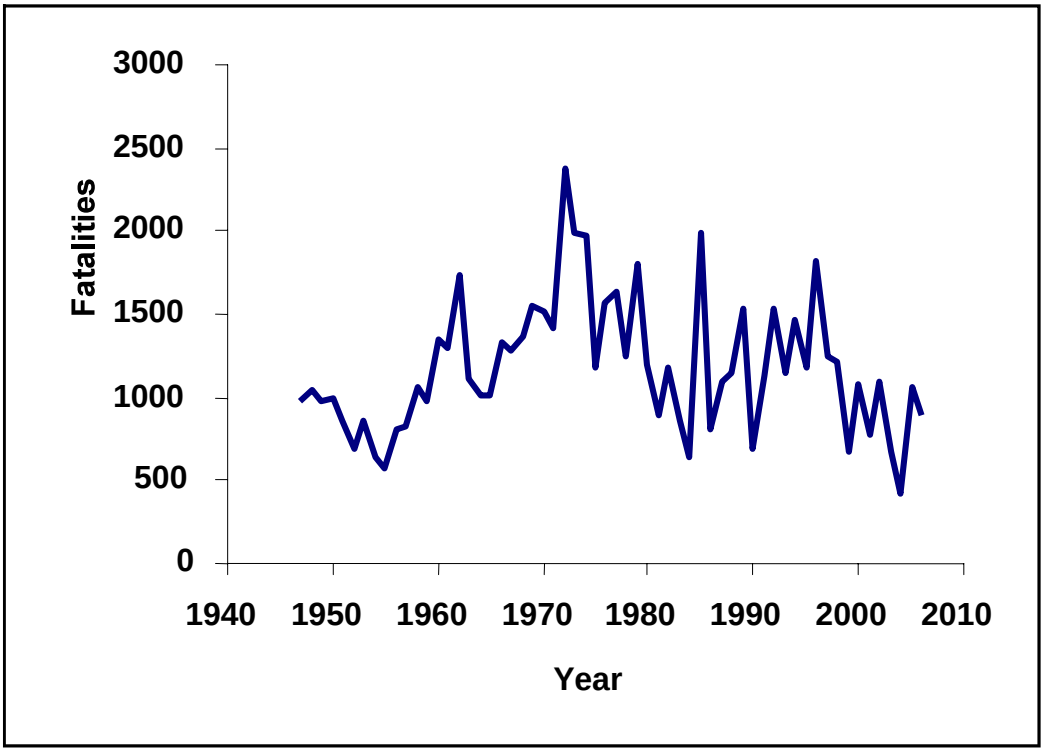

Figure 3. Worldwide Airline Fatalities 1947-2006. Taken from Airline Safety Network Statistics, http://aviation-safety.net/statistics/period/

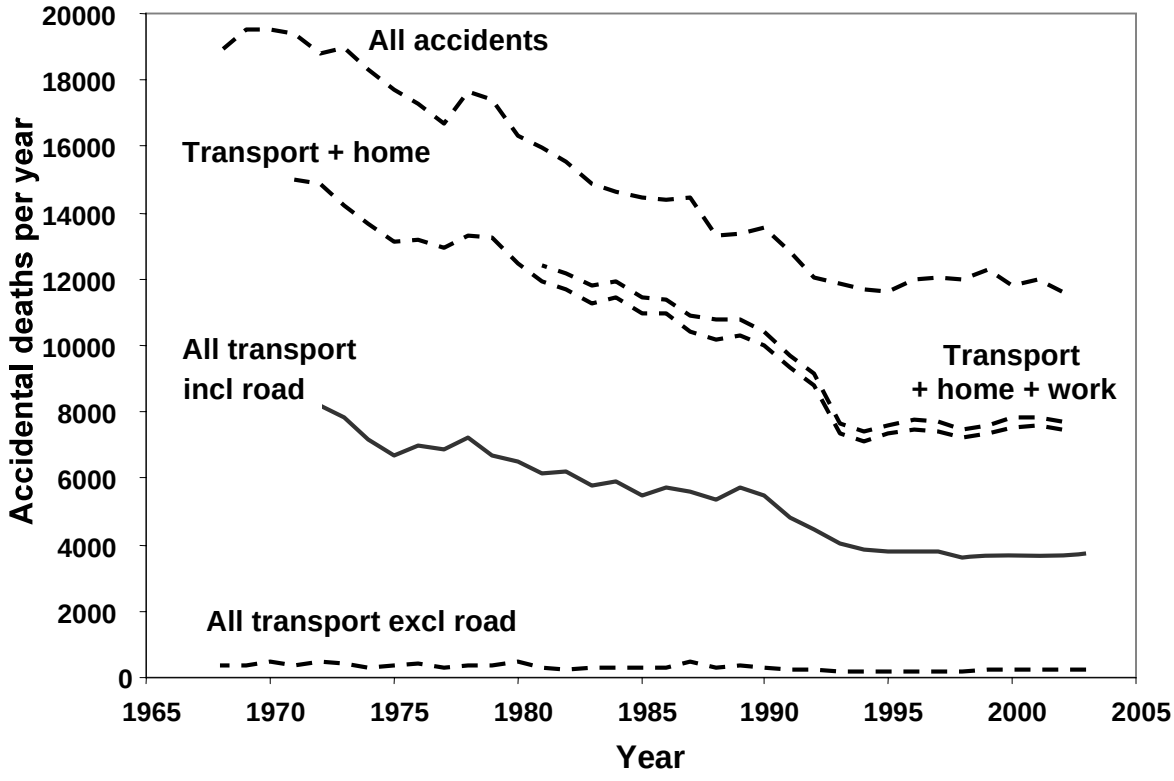

Figure 4: Accidental deaths in Great Britain 1967–2003. Source Evans (2005).

Mode	Fatalities per billion Passenger km
Air	0.00
Water	0.3
Bus/coach	0.3
Rail	0.4
Car	2.8

Table 1. Passenger fatality rates by mode in Great Britain 1995-2004 Average. (Source: Department for Transport 'Road Casualties 2005 Great Britain', 2006)

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

Figure 5. ATM Safety System Layers

What is the acceptable risk?

➔ X (the Safety Target)

⇒ measurement + value judgement

What would be (*sic*) the actual risk?

➔ Y (the estimated risk)

⇒ modelling + prediction +validation

If Y ≤ X then decision is 'go ahead'

Figure 6. ATM Decisions – Safety Target Philosophy

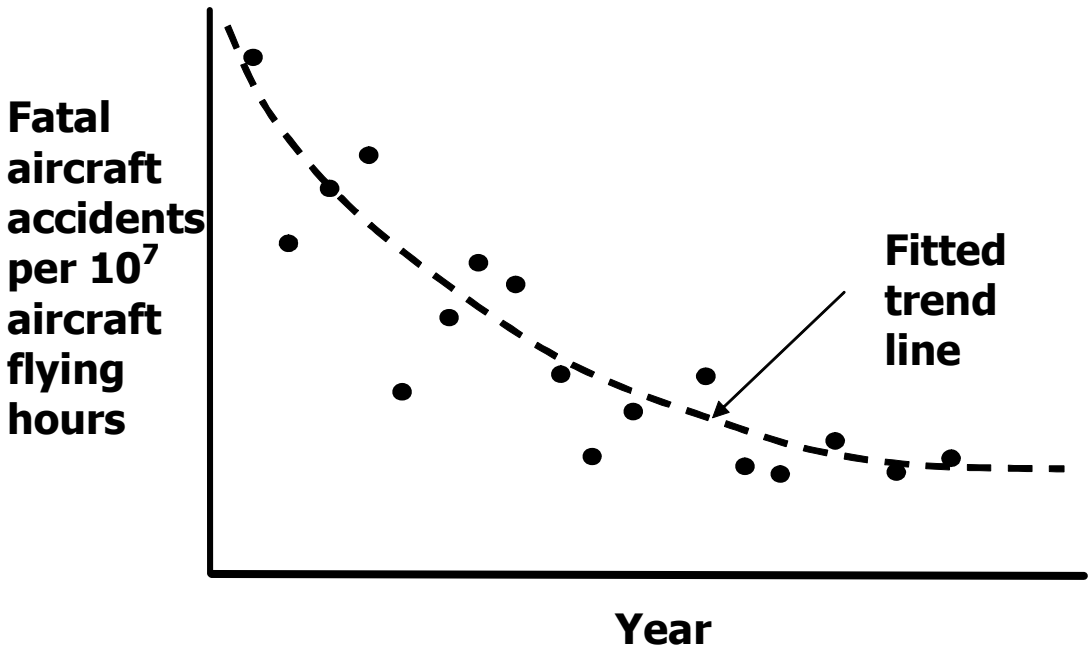

Figure 7. Fatal aircraft accident rate data and trend – *illustrative*

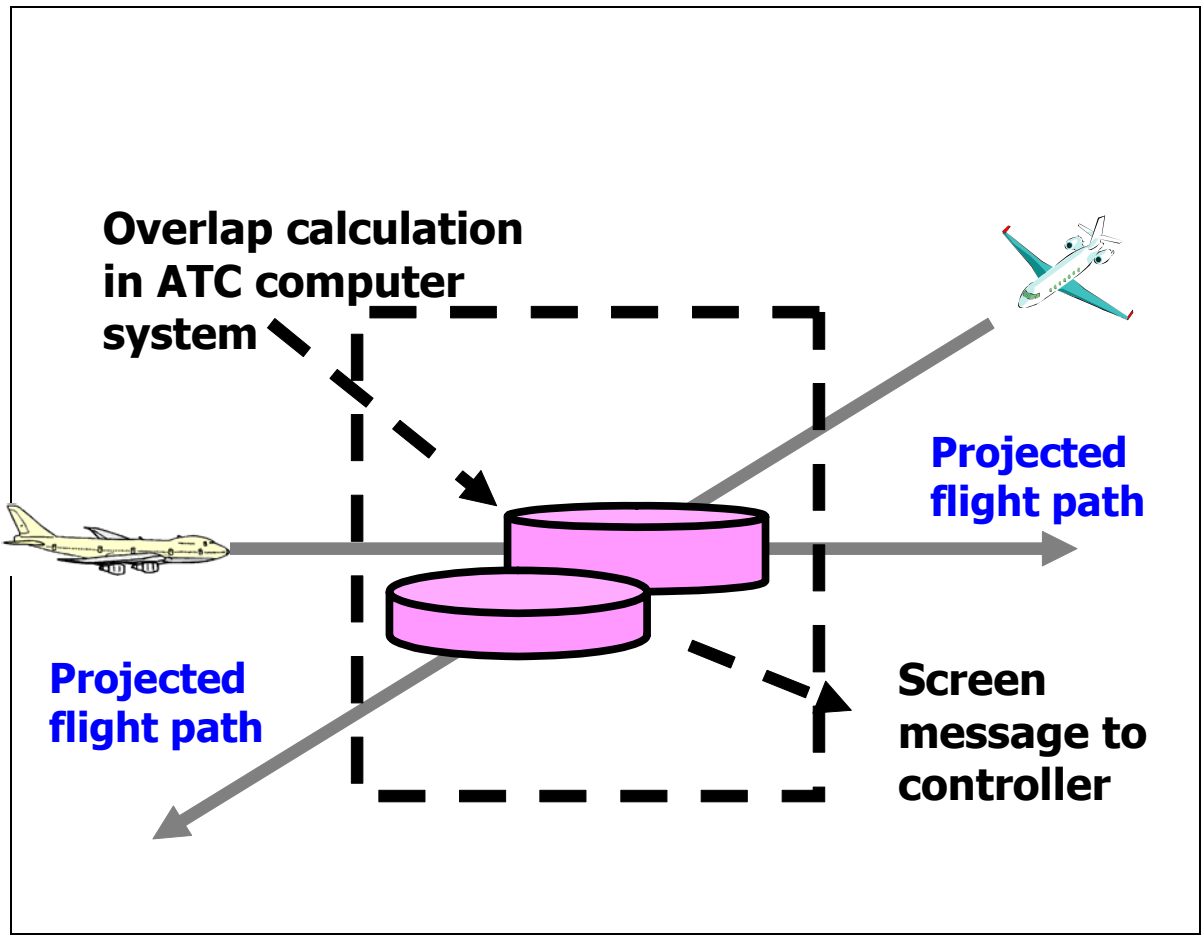

Figure 8. Short Term Conflict Alert - STCA

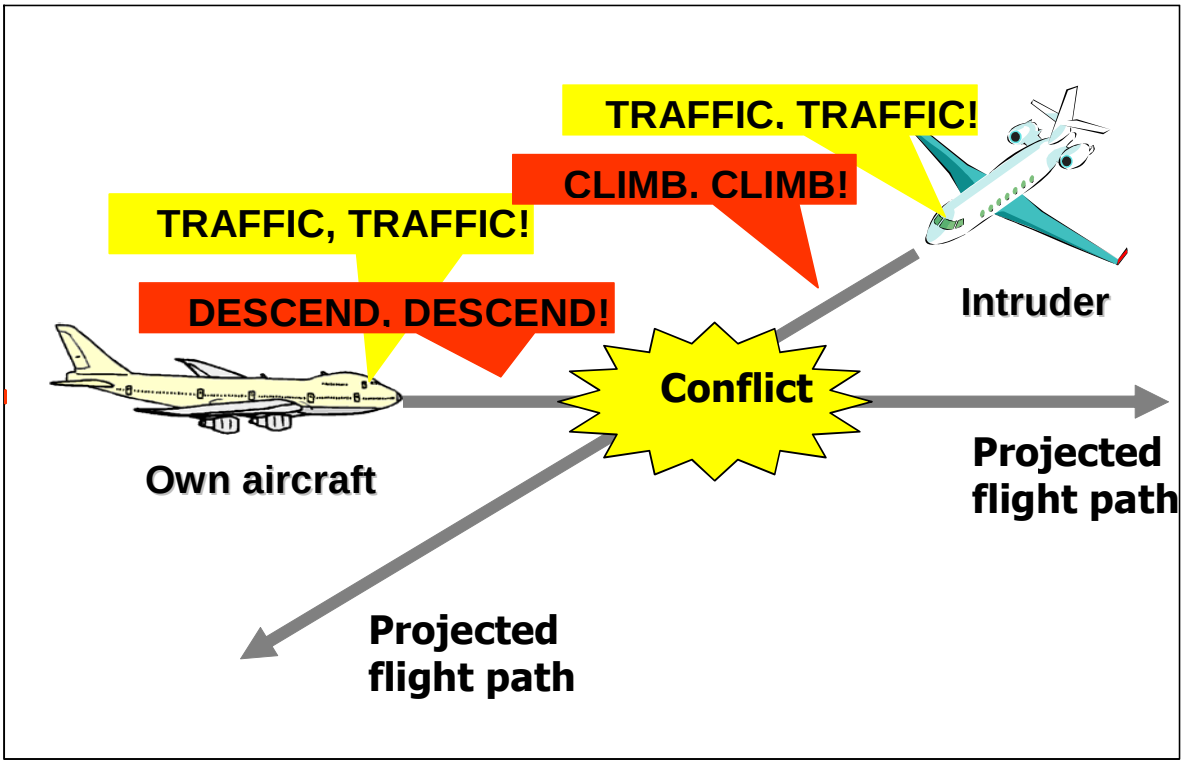

Figure 9. Traffic Alert and Collision Avoidance System - TCAS

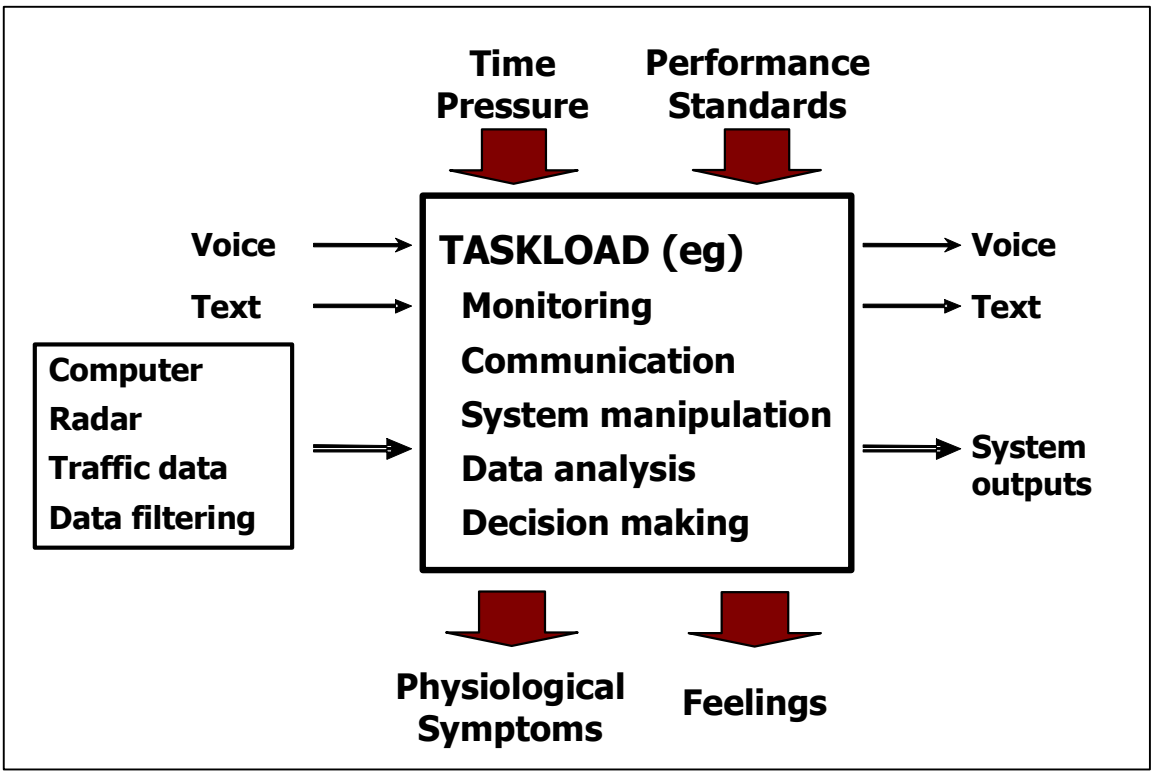

Figure 10. Aspects of Control Workload

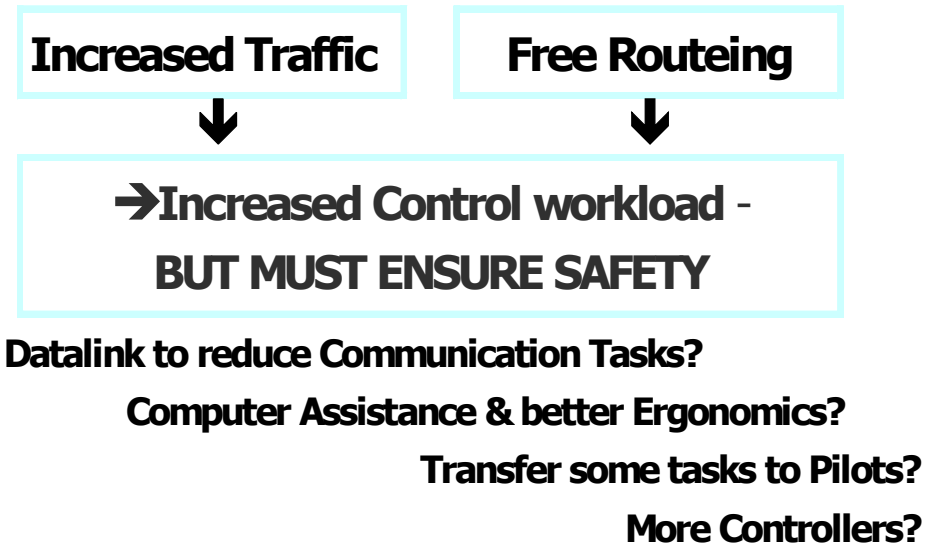

Figure 11. Control Workload Problem

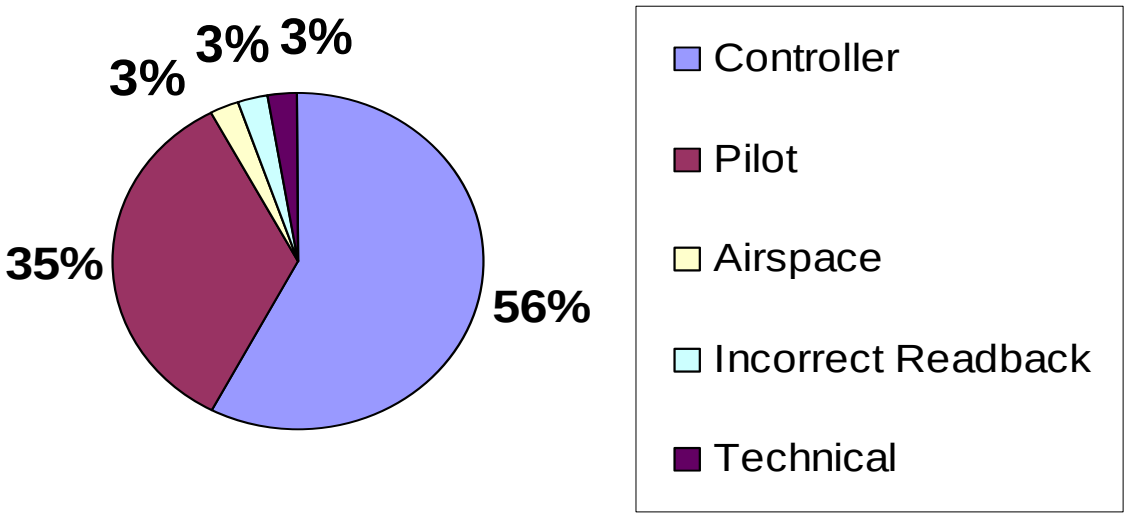

Figure 12. Airprox 'Causes'

From a sample of 117 recent UK Airproxes involving commercial flights. Eliminated from data set: Airproxes in 'uncontrolled' airspace (Class F/G), military zones, North Sea; military aircraft, parachutist, balloons, sighting reports.

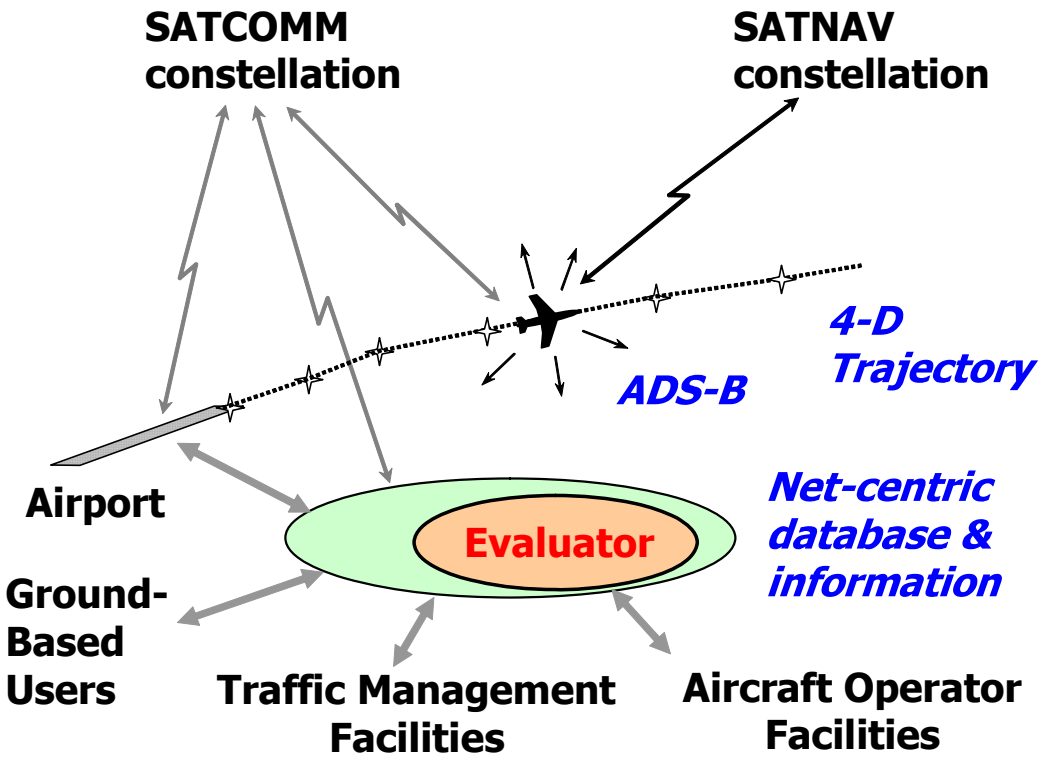

Figure 13. NextGen

- User-selected flexible 4D trajectories from gate-to-gate
- Agreed-upon trajectory contracts flown with approved variances
- Variances in flightpaths managed by exception
- Seamless airspace across current or projected boundaries, so no need to distinguish between airports/terminal area/en route airspaces/boundaries
- Network centric system-wide perspective with information shared by all users: datalink – *ie text-messaging between all parties* – is vital

Figure 14. NextGen Operating Concept

- **Operations based on better forecasting**

  Change from *reactive ATM* to *anticipatory ATM* – to reduce operational pressure on human operators

- **Better anticipation of problems**

  Collaborative decision-making procedures – stakeholders share and negotiate relevant information

  Merge the different "trajectory" representations into a single one established by the on-board computers

  Accurate monitoring of the scheduled 'trajectory' by means of extremely accurate satellite navigation

- **Efficient telecommunications network**

  Network of ground-to-air data links to enable accurate 'trajectory" information exchanges

  All stakeholders to have effective and simultaneous access to flight information status

- **Optimisation of the use of airports**

  'Smooth' approaches to reduce noise/gaseous emissions

  Better forecasting and detection of turbulence phenomena

- **Increased automation of air traffic control tools to assist operators**

  Share workload between controller and pilot

  Trajectory negotiation planning and support tools

  Cockpit tools to visualise surrounding traffic

Figure 15. Some SESAR Features (extracted from EC, 2007)

An Emergent property of a complex system is a behaviour that surprises its designers

Strong emergent properties occur if the system exhibits behaviours that cannot be identified through functional decomposition – the system is more than the sum of its component parts.

Weak emergence properties are unexpected simply because of the degree of difficulty that the designer has in deducing them from his/her understanding of the component parts. *(Compare 'Interactive complexity' concept for system coupling.)*

Figure 16. Emergent properties

1. All systematically applied safety defences should be considered as full parts of the integrated ATM safety system
2. This includes STCA and TCAS
3. Hazard analysis calculations incorporating STCA and TCAS provide a measure of the true risk potential in the real world
4. Excluding them puts an extra burden on risk estimation: the calculations will tend to be even more cautious – and hence more pessimistic about the value of new concepts
5. This is *backward-looking*: it retards the introduction of acceptably safe systems embodying novel operational concepts – it has become more difficult to *prove* their safety

Figure 17. The rationale for Ground and Air Protection layers being included in hazard analysis.

The Glass Bead Game (*Das Glasperlenspiel*) by Hermann Hesse

Elegant intellectual activity, using up bright people's time, but of no practical value

The game's exact nature remains elusive

There are only allusions about the precise rules of the game

The rules are so sophisticated that they are not easy to imagine

Figure 18. The Glass Bead Game