Air Traffic Technology International 2007 Paper
Draft Version @ 25th September 2006
Peter Brooker

AIR TRAFFIC MANAGEMENT INNOVATION: THE RISKS OF *STASIS*

INTRODUCTION

Stasis is 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. ATM in Europe is in danger of being in stasis, because current ATM safety regulation policies are tending to make it more and more difficult to innovate, to introduce new technologies and ways of operating.

The following essentially highlights the key findings of a group of published research papers analysing a variety of problems with ATM safety regulation policies. These policies mainly derive from the Eurocontrol Safety Regulation Commission (SRC), but also from ICAO. The case studies discussed are: the role of ground-based safety nets (Short Term Conflict Alert – STCA); air-based safety nets – Airborne Collision Avoidance Systems (ACAS); and risk assessment and mitigation in ATM (SRC's Safety Regulatory Requirement Number 4 – 'ESARR4'). These policies were no doubt developed with good intentions but, in quality-management jargon, they are not 'fit for purpose'.

CHECKLIST OF POSSIBLE SAFETY POLICY PRINCIPLES

A safety policy must add significant worthwhile and valid content to decision-making. If it does not, then what is its value? Who needs unnecessary unproductive bureaucracy and fruitless activity in ATM safety management and regulation? It can add value in several ways. The kinds of things it can do are:

- Clarify decision-making processes. It can show the process by which decisions can be made rationally. It can give examples of how to set out logical and comprehensive safety arguments. It can correct misunderstandings.
- Learn lessons from past accidents and incidents. It would be a poor policy that did not learn from the past. If certain types of human error produces significant risks, then the policy would show what steps are need to correct or mitigate them. If certain kinds of failure mode had been ignored in the past, then it would instruct the reader to include them in hazard analyses.
- Generate safe system improvements. This must be the most important feature. A good safety policy is one that considerably reduces the risks of an accident in future and/or permits safe innovations. If, for example, technology offers ways of dramatically reducing fuel costs and/or carbon emissions, then ways should surely be found for using it safely.

Make sense organisationally. Safety policies may have complex implications on safety culture and ATC provider organisations. The ownership of policies must be clear and sensible. This ownership should include incentives for better safety performance. Does the policy help?

Recognise the reality of ATC. Controllers are people. People make mistakes. Therefore controllers make mistakes.

A rigorous process needs to be followed. Draft safety policies need to be exposed to scrutiny by the full range of professional criticism, with all the key source material underpinning regulations being in the public domain. What assumptions have been made? Are safety responsibilities clear, complete and comprehensive? Could there be sensible arguments against this openness and vigilance?

When there is a feasible, preferably quantitative, way of doing something, then it should be explained clearly, with real-life examples. Are there arguments against it? When building a house, would guidance material from Mademoiselle M., who has built many houses and who knows what difficulties can be faced in reality, be better than help from Mr P., who has never built a house in his whole life, but who has read introductory books on house building?

CASE STUDY: STCA

Current SRC policy says that changes to the ATM system (including safety minima) must be demonstrated through risk assessments to meet the Target Level of Safety (TLS) without needing to take safety nets such as STCA into account. The TLS here is a design hurdle, a quantified risk level that a system should – i.e. be designed to – deliver. A TLS covers all aviation-related causes, but does not usually attempt to cover the consequences of terrorism or criminal behaviour. This policy of excluding STCA is wrong and because it does not build rationally and consistently from ATM's firm foundations of TLS and hazard analysis.

It is worthwhile setting out some UK background. The UK's National Air Traffic Services (NATS) introduced STCA progressively from the late 1980s. STCA functionality had been in its radar data processing system for several years previously but had not been implemented. This was mainly because of inadequate secondary radar performance.

Until the mid-1980s, NATS' senior managers believed that the UK ATC system was acceptably safe, even with rapid traffic growth and pressure on operational staff. Several serious Airproxes then occurred and received wide publicity. The chairman of the UK CAA, effectively the safety regulator, asked the head of NATS to give guarantees that the UK system was safe and that the ATC-related causes of the Airproxes would quickly be eliminated.

NATS could not give such guarantees and therefore decided to introduce STCA. The key policy points here are:

- STCA was not a desirable add-on to an already acceptably safe ATC system. STCA was necessary for NATS to deliver the level of ATC safety required by the CAA.
- STCA could offer a near-guarantee of detecting the ATC-related causes for the observed Airproxes in sufficient time for controller action.

- NATS hazard analyses showed that STCA would improve safety, except in situations where there were wider ATC operational and/or regulatory failures (compare the Überlingen mid-air collision).

The UK and NATS are not unique in these kinds of decisions.

In system safety analyses, STCA appears in exactly the same way as help from colleagues. STCA is effectively an extra colleague, dedicated to one supporting task – extrapolating radar tracks. Thus STCA is part of the ground ATC system. The effects of STCA in reducing deaths from mid-air collisions should therefore be fully included in safety assessments, hazard analysis and safety audits. Figure 1 sets out a list of questions and answers about this 'Rational STCA Policy'.

CASE STUDY: ACAS

Present ICAO ACAS Policy says that ATC systems and procedures must be 'acceptably safe' – i.e. meet TLSs or similar requirements – without considering the effects of the ACAS safeguard. This policy is invalid. The core problem with this policy is:

- As a first step, ATC safety targets for future systems generally assume that the current system is acceptably safe, and then add in the challenge of a sizeable factor for continued improvements or to meet the challenge of increased traffic
- But there are significant numbers of incidents – potential accidents – in the current system that are resolved only by the use of ACAS. Current ATC performance, and hence the baseline for safety targets, therefore relies on the safety gains delivered by ACAS. If the current system were to be proved acceptably safe 'without ACAS', it would need to be demonstrated that the great majority of these kinds of incidents were either not safety significant or that other system defences would have successfully resolved them. ICAO has not offered such a safety demonstration for the current European system, with its dense, complex ATC-controlled airspaces.
- The present ICAO policy therefore demands that the future system 'without ACAS' will deliver markedly higher safety performance than the current system 'with ACAS' is just about managing to achieve. Where is a reasoned justification for this extra constraint? This is like saying that, just to qualify to enter the next Olympic Games 100 metres competition, you will need to break the world record whilst wearing winter clothing: certainly a highly cautious performance target, but probably leading to zero competitors taking part – the stasis effect.

Tables 1 and 2 illustrate this core problem; the data is from UK Airproxes with ACAS Resolution Advisories (RAs). In Table 1 the controller did not detect the conflict and STCA did not trigger. In Table 2 STCA triggered but the controller did not act on the alert.

The miss distances H and V at the closest point of approach in Tables 1 and 2 were measurements taken after the pilots' action in following the RA. Miss distances, had ACAS not existed, would need to be estimated from the individual incident reports. They would generally have been much less than those achieved with the benefit of ACAS.

The incidents in Tables 1 and 2 had safety significance, given the breaches of separation minima. In several of the incidents the report does not suggest that the ATC team detected the potential close approach. These failures often seem to be from states of impaired situation

awareness in a controller's mind. Controllers need to remember to perform activities in the future. Successful completion of an intended action in the future depends on a type of remembering called 'prospective memory'. What produces these problems, and how can they be prevented from happening or additional protection be provided?

The safety defence that made the difference – that delivered acceptable safety – was ACAS. If the present system is supposed to be acceptably safe without ACAS, then where is the demonstration that the kinds of incidents in Tables 1 and 2 would have been successfully resolved in such as case?

Thus ACAS should be viewed as a fully integrated part of the ATM safety system, not some kind of supplement. The effects of ACAS in reducing the risks of mid-air collision should be fully included in safety assessments, hazard analysis and safety audits. Figure 1 sets out a list of questions and answers about this Rational ACAS Policy, based on an analysis of ICAO documents.

CASE STUDY: ESARR4

The SRC's ESARR4 and its supporting material are long and complex documents, so this is a sketch of the issues involved. The main point is that the onus is on the SRC to demonstrate, by reference to past successful work documented in the policy/research literature, that the ESARR4 approach actually delivers the goods. Much of ESARR4 is quite reasonable – concerned with sensible good housekeeping and safety process/documentation issues. The problems arise in respect to strategic safety and hazard analysis issues: the focus is very much on these critical pieces of text.

ESARR4 is an SRC Requirement. An ideal Requirement would presumably look something like Figure 1. The diagram focuses on the word 'specified'. The data needed to do the job is clearly set out and it should be processed in well-defined ways. The data would be a description of the ATM subsystem and its relevant performance. The process would be a logical set of calculations and well-defined assessment procedures. But a critical reading shows that ESARR4 does not deliver. It does not produce answers usable by prudent decision makers. It is in practice more complex and less helpful than the ideal of Figure 3 and more like Figure 4. The answer that comes out at the end is usually an indecisive one, leading to a need for more data – which often does not exist – or for some kind of expert judgement in order to make a decision. Excessive use of expert judgement is open to many criticisms. Consistency may just be consistently wrong: differences in viewpoint may be expressions of people's pessimism or optimism. It is important to restrict expert judgement to topics and events that have been experienced by operational experts.

Some specific practical and technical flaws of ESARR4 include:

- Lack of clarity about responsibilities for ATM safety (e.g. accidents attributable to safety regulatory decisions or practices);
- Not based on practically proven techniques. Surely ATM safety must build on successful best practice?
- Reliance on 'mechanical' probabilistic safety assessment;
- Ignoring risk assessment techniques successfully used in ATM, which are well documented in other Eurocontrol publications;

- Not distinguishing between different types of ATM subsystem (e.g. *Weick*'s tightly coupled and loosely coupled);
- Unstructured use of expert judgement;
- Careless use of 'worst case'.

Thus ESARR4 leaves it to the service provider to make a safety case but provides no examples about how this might reasonably be done. This is very poor guidance material. Examples need to be given, covering the whole range of real-life potential problems – of how such tasks are to be accomplished. ESARR4 must show how big, real-life ATM safety problems have been tackled successfully.

Does ESARR4 tend to suggest that safety regulators have a limited role – almost a fire-and-forget mentality? It is surely not sufficient for a safety regulatory body to deal in abstractions – a well-founded learning process for service providers is needed. The regulator must surely act professionally, only setting regulations that have a good chance of being achieved through reasonable efforts. To quote a former CAA colleague: "Its high-level principles are good but the detail [of ESARR4] is seriously flawed. It certainly seems odd that a piece of regulation can be produced for which there is no agreed method of compliance – and possibly no practical means."

In general the claims ESARR4 etc make for the methodologies proposed are overstated – not supported by sound evidence from real-world hazard analysis. This obviously generates subsequent problems with the intended practical application of ESARR4. Continued use of this document could therefore misallocate scarce safety resources, divert attention from real safe system improvements, waste regulators' and managers' time and delay safe system innovation. The most important underlying change would be a refocusing on practical safety assessment based on methods that have already demonstrated their merits, including references to realistic specimen risk calculations.

CONCLUSIONS

All systematically applied safety defences should be considered as full parts of the integrated ATM safety system. Hazard analysis calculations incorporating STCA and ACAS provide a measure of the true risk potential in the real world. Excluding them has no rational basis and puts a huge extra burden on risk estimation: the calculations will tend to be very over-cautious – and hence much more pessimistic – about the value of new concepts. This is backward looking: it retards the introduction of acceptably safe systems embodying novel operational concepts, because it is more difficult to prove their safety. To ignore STCA and ACAS in risk assessment is to produce innovatory stasis, i.e. prevent safe system improvements.

All ATM safety system design/operations participants have a duty of care in that they apply reasonable, skilful and proper use of relevant evidence, with intelligence, foresight and scrutiny. Regulators have some measure of responsibility for safe ATM system design/operations. It is not enough to promulgate abstract regulations in the office. Regulatory guidance material needs big, real-life examples. Examples for risk assessment policy and guidance must cover the whole range of potential problems, including professional analyses of relevant incident data, and practical-orientated safety assessment based on methods that have already demonstrated their merits. Poor guidance material, without clarity about compliance, produces indecision – and hence stasis.

Safety analysts and aviation decision makers usually have to make decisions based on evidence that is – inherently – incomplete. Ideal hazard assessments will not be practically achievable in every case, so decision making requires the cooperative hard work – indeed wisdom – that multi-disciplinary ICAO panels can demonstrate. Openness about data analysis, assumptions and reasoning, coupled with peer review and rational responses to critical challenge, are vital components in these processes.

BIBLIOGRAPHY

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

Brooker, P. (2004). Why the Eurocontrol Safety Regulation Commission policy on safety nets and risk assessment is wrong. Journal of Navigation 57(2), 231-243.

Brooker, P. (2005). Airborne Collision Avoidance Systems and Air Traffic Management Safety. Journal of Navigation 58(1), 1-16.

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

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

Eurocontrol SRC (2001). Risk assessment and mitigation in ATM, Eurocontrol Safety Regulatory Requirement ESARR4, Edition 1.0., Eurocontrol, Brussels.

Eurocontrol SRC (2003). Explanatory material on ESARR4 requirements. EAM 4 / GUI 1. Edition 1.0, Eurocontrol, Brussels.

Eurocontrol SRC (2003). SRC Policy 2: Use of safety nets in risk assessment & mitigation in ATM. Edition: 1.0 dated 28 April. Eurocontrol, Brussels.

ICAO ANC11 (2003). Eleventh Air Navigation Conference Montreal, 22 September to 3 October 2003. The Role of collision avoidance in future ATM systems (presented by the secretariat) ANConf.11.WP.034.1.en.wpd AN-Conf/11-WP/34.

Nunes, A. and Laursen, T. (2004). Identifying the factors that led to the Ueberlingen mid-air collision: implications for overall system safety. Proceedings of the 48th annual chapter meeting of the Human Factors and Ergonomics Society, New Orleans, LA, USA.

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

Question	Answer
Couldn't the ATC safety regulator just mandate STCA?	No. What safety basis would the regulators have for mandating? The regulator would need to possess the skills to write/enforce the rules for (e.g.) appropriate optimisation of parameters covering all the individual STCA applications.
Didn't the original derivations of TLSs *exclude* safety nets from risk assessment	No. TLSs were derived as system design targets against all the risks affecting aircraft and passengers. Thus risk calculations would potentially include *all* mitigating factors, from controller monitoring and intervention to STCA.
Is it really necessary to include STCA in hazard analyses?	Yes, because this measures the true risk potential in the real world. Excluding STCA retards the introduction of acceptably safe systems embodying novel operational concepts. This would be because it would become more difficult to *prove* their safety. *Safety policies should generate safe system improvements.*
Is STCA really part of the ATC system?	Yes. STCA is enabled through software in the main ATC computer. It is carefully designed into the controller's display. Safety regulators expect detailed guidance on STCA in the ATC provider's and regulator's formal safety documents. UK Airprox reports show that controllers *in practice* view STCA as simply one of their tools.
STCA can fail, so it should not be relied upon.	No. ATC system parameters and procedures are generally based on properly functioning system components. If (e.g.) STCA *or* radar *or* communications fails, then heightened risks must be mitigated by (e.g.) extra staff, flightpath restrictions, flow regulation. Compare Überlingen accident – *Nunes and Laursen.*
Surely the present policy will not cause any harm?	No. If STCA were seen as a safety add-on, then what is the incentive for its Europe-wide introduction? What incentives would they have for optimising parameters to reduce unnecessary warnings? STCA must be optimised for controllers' usage. *Safety policies should generate safe system improvements.*
Surely STCA should be viewed as a safety 'add-on'?	No. The UK CAA and NATS judged STCA as *necessary* to achieve acceptable ATC safety, not an optional feature.
Surely the ATC system is acceptably safe without STCA?	No. Where is there comprehensive quantitative evidence to support such a view for dense, complex ATC operations?
The safety benefits from STCA are only 'marginal'	No. Analyses of Airproxes demonstrate the practical importance of STCA in preventing very serious incidents, and show that it is integrated into ATC to ensure safety.

Figure 1. Questions and answers on a rational STCA policy

Question	Answer
Aren't the safety benefits from ACAS 'marginal'?	No, see the evidence above. If ACAS did not exist, then the continued achievement of current safety levels would probably require extra staff (for monitoring), flightpath restrictions (eg to reduce climb/decent frequencies), additional flow regulation (to reduce further controller workload peaks), and restrictions on non-commercial traffic (to prevent airspace incursion problems).
Didn't the original derivations of ATC safety targets (TLSs) exclude safety nets from risk assessment?	No, see Figure 1
Is it really necessary to include ACAS in hazard analyses?	Yes, see Figure 1.
People measure ATC's performance excluding ACAS.	No. The public's concern is with the safety level achieved in the real world, not with what it might have been in some theoretical universe – composed of 'what ifs' – in which ACAS did not exist.
Separation minima guarantee safety**.**	No. The use of separation minima alone provides no guarantee that the risk of collision will be acceptably safe. The use of separation minima, STCA and ACAS alerts are safety barriers combining to provide statistical assurance of safety rather than absolute guarantees.
Surely ACAS performs erratically and has side effects?	No. ACAS produces fewer side effects for ATC, in terms of serious incidents, than cautiously estimated when first introduced. ACAS certainly generates some nuisance and false alerts, but they do not frequently lead to hazardous incidents – the real measure of what matters in safety terms. ACAS improves safety, except if there are wider ATC operational and/or regulatory failures (compare Überlingen mid-air collision). *Safety policies should learn lessons from past accidents and incidents.*

Figure 2. Questions and answers on a rational ACAS policy

Figure 3. An ideal Requirement structure

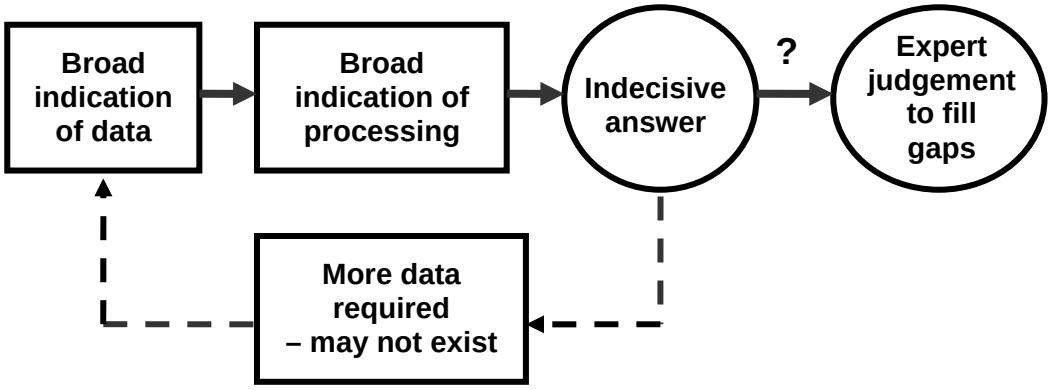

Figure 4. How ESARR4 fails

Airprox number	Summary	H Nm	V feet
1999127	Controller had issued a descent clearance that would have led aircraft 1 to descend through the level of aircraft 2, which he had inadvertently not taken into account.	0	1100
1999200	Controller did not take aircraft 1 into account when he descended aircraft 2.	4.5	400
2000018	Aircraft were being vectored from N and S to line up on the Heathrow ILS. Controller did not ensure standard separation.	1.2	300
2000032	Controller gave 'erroneously and essentially unforced descent instruction' to aircraft 2.	2.0	700
2000055	Apparently anomalous RA [NB: TCAS II Version 7.0 not in use, so incident not now relevant]. Might have been malfunction or misheard aural warning – 'reduce climb' for 'climb'?	2.5	500
2000126	Controller did not ensure standard separation between the two aircraft. STCA did not alert because of geometry of the situation.	1.1	700
2001069	Controller allowed aircraft 1 to climb to the level that he had cleared aircraft 2 to fly at, without coordination.	2.8	700

Table 1. UK 'Outside STCA parameters' Airproxes

Airprox number	Summary	H Nm	V feet
1999221	Mentor did not detect that instruction by trainee put both aircraft at same level without standard separation. Mentor issued avoiding action at time when pilots were responding to RAs.	1.0	700
2002112	The control team did not ensure that aircraft 1 was coordinated with the neighbouring Control Centre. STCA was dismissed as a 'nuisance warning' – aircraft 2 assumed to be climbing to level 1000ft below aircraft 1.	1.6	400
2003075	Controller dispensed with vertical separation without ensuring lateral separation. Controller had seen the STCA alert but did not consider it a problem, as he believed that aircraft 1 would safely descend through the level of aircraft 2.	1.4	700
2003164	Aircraft 1 crew descended below their cleared level into conflict with aircraft 2. STCA alerted after aircraft 1 had received a TCAS RA.	3.7	500
2003169	A sighting report, given that aircraft 1 was in TMA airspace and the other in Class G airspace – deemed separated (according to the CAA regulations) when STCA activated.	0.6	500
2003184	The aircraft 2 crew read back the wrong heading and level instructions, which went undetected by the controller. The controller said he had no reason to doubt that the aircraft would not comply with the issued clearance. STCA activated and shortly afterwards a TCAS RA climb was issued.	3.4	600

Table 2. 'STCA Warned but not actioned by controller' Airproxes

(a) Selection extracted from UK Airprox Reports, 1999-2003. The first four digits of the number is the year, the last three is the identification. **H** is the radar recorded horizontal distance at closest approach; **V** is the corresponding vertical distance.
(b) In Table1, the STCA safety defensive layer could not provide protection; in Table 2, STCA alerted but was not acted upon.