THE JOURNAL OF NAVIGATION (2004), **57**, 39–51. © The Royal Institute of Navigation
DOI: 10.1017/S037346330300256X Printed in the United Kingdom

Radar Inaccuracies and Mid-Air Collision Risk: Part 2
En Route Radar Separation Minima

Peter Brooker

(*Cranfield University*)

A review of safety targets for en route ATC radar separation suggests that the existing target level of safety (TLS) is over-cautious. If risk budgeting principles are followed consistently, a 'radar TLS' of $1{\cdot}0 \times 10^{-9}$ fatal aircraft accidents per flying hour is appropriate. This rate is consistent with Joint Aviation Authorities (JAA) guidance on system failure conditions leading to catastrophic accidents. Dynamic and static calculations using published data are compared. The new methodology shows where there are problems with the traditional static calculations, and how to improve the estimation. A further improvement introduces a simple robust model of the controller's decision processes. The focus is not on describing what controllers would generally do, but on setting criteria based on what they could not reasonably be expected to do. This additional ingredient into the calculation adds realism and ensures that attention is focused on hazardous correlated errors. Focused data collection would be an essential component of new risk estimates. The key information required would be on radar performance and the nature and frequency of use of radar separation, including the relative velocities for proximate events at closest point of approach and the frequency of correlated gross errors (through a conditional probability factor). If this factor is not properly taken into account, then the data collection and analysis could be inefficient.

KEY WORDS

1. Mid-air collision. 2. Radar separation minima. 3. TLS.

1. INTRODUCTION. Part 1 presented a 'dynamic' methodology for the estimation of mid-air collision risk arising from radar inaccuracy. Part 2 uses this methodology to compare static and dynamic calculations, and then examines the nature of radar errors as seen by controllers. First, Section 2 reviews safety targets for en route ATC radar separation minima. Section 3 then sets out some examples of dynamic and static calculations using published data, to show where there are problems with the traditional static calculations and how the new methodology improves the estimation. Section 4 further develops the methodology by introducing a simple model of the controller's decision processes, starting this analysis from the International Civil Aviation Organisation (ICAO) regulations governing radar separation minima. The focus is on how correlated errors could mislead the controller. Conclusions are in Section 5.

2. SAFETY TARGETS FOR EN ROUTE ATC RADAR SEPARATION.

A target level of safety (TLS) is a key concept in aviation safety. Appendix A gives some background on risk metrics for a TLS. The TLS is a design hurdle, a quantified risk level, measured as some sort of accident rate, which a system should deliver – i.e. in planning, design and actual performance. The main motivation is to answer questions about separation minima between aircraft and the design of route structures, although it should be stressed that the TLS relates to the risk of an accident due to *all* causes. There have been several much more recent reviews of TLSs, by Davies and Sharpe (1993), by Eurocontrol's Safety Regulation Commission (2000) and as part of joint studies by the FAA and Eurocontrol (1998). The TLS is a quantified risk level (measured as an accident rate) that a system should – i.e. be designed to – deliver, usually as a proportion of fatal accidents per so many flying hours (or airport movements when that is more appropriate). Most of the practical problems are with the proper estimation of the safety level that is or would be achieved. There is an Actual Level of Safety (ALS) being achieved in the system under examination. The key question is then how is this to be calculated with sufficient accuracy to be confident that the $\text{ALS} < \text{TLS}$?

TLSs appropriate for accidents arising from mid-air collisions in controlled airspace are usually derived by taking historical accident rates – which show a progressive reduction over time – and extrapolating forward, i.e. an 'improvement factor' is applied. The original centre of attention was on commercial passenger jet flights – in fact the statistical analyses were restricted to scheduled services. This basic method has continued to be used to the present day. The TLS measures the rate of fatal aircraft accidents, i.e. accidents in which at least one person in the aircraft was killed, per so many aircraft flying hours. [NB: in these calculations, it is invariably assumed that a mid-air collision would produce fatalities in both aircraft, i.e. two fatal accidents.] The current ICAO (RGCSP, 1995) figure is $1{\cdot}5 \times 10^{-8}$ fatal aircraft accidents per flying hour for mid-air collisions – of any type for all causes – for en route flight in controlled airspace. Given an equal allocation of the 'risk budget' (Profit, 1995) in three dimensions (i.e. $0{\cdot}5 \times 10^{-8}$ per x, y and z dimension), the horizontal (i.e. summing the x and y dimensions) TLS for mid-air collisions for en route flight in controlled airspace is $1{\cdot}0 \times 10^{-8}$ fatal aircraft accidents per flying hour.

This TLS is different from that adopted in RSSWG (1982), Sharpe (1991) and related Civil Aviation Authority (CAA) and National Air Traffic Services (NATS) documents (e.g. Greenwood *et al.*, 1994). These take the horizontal TLS to be the same as the vertical TLS [NB: for *all* causes, not just aircraft systems affecting height-keeping], i.e. $0{\cdot}5 \times 10^{-8}$ fatal aircraft accidents per flying hour. With the benefit of hindsight, this does not appear valid. Radar separation involves control of aircraft pairs in potentially all orientations in the xy 'playing field', and, in real-life, controllers use both lateral and longitudinal separations between aircraft (Part 1 gave examples), so the horizontal TLS should be two thirds the total TLS, i.e. $1{\cdot}0 \times 10^{-8}$. If one were to take vertical plus horizontal risk budgets as just twice the vertical, i.e. $1{\cdot}0 \times 10^{-8}$, then the question is where has a third of the total (all causes) mid-air collision risk budget in the TLS disappeared to? Note that Sharpe (1991) focused on azimuth errors: his paragraph 5.1.8 comments on the infrequency of gross range errors and that the focus of analysis is on a 'worst case' scenario of azimuth separation. This focus on azimuth errors may have translated into 'wasting' the range error risk budget (presumably a third of the total TLS), as the estimated collision risk for radar range errors would be much less than the amount deemed to be 'allowable'.

The TLS for horizontal separation next has to be partitioned against the different sources of risk, in particular to arrive at a value for a 'radar minimum in en route control TLS' (RTLS), i.e. the accident rate target appropriate for radar inaccuracy. In early work, a third of the total was allocated to radar inaccuracy, but later a figure of 10% was used (Sharpe, 1991). It was emphasized that this was a somewhat arbitrary figure, designed to reflect the increased accuracy of radar systems and hence the much reduced likelihood that radar errors could produce a mid-air collision. Recent work examining Airproxes (Brooker, 2002a and b, UK Airprox Board, 1999 et seq) confirms that the frequency of near-collisions is largely a consequence of pilot/controller-related issues, crudely expressed as 'right place on wrong flight path'. If the factor of 10% is used, the RTLS is $1{\cdot}0 \times 10^{-9}$ fatal aircraft accidents per flying hour (compare $5{\cdot}0 \times 10^{-10}$, i.e. half as much, in Sharpe, 1991). Is the 10% figure appropriate? In particular, given the Airprox evidence (Brooker, 2002b), which would generally be expected to match the proportion of accidents in the very long run, should an even smaller percentage be used?

One obvious problem with, say, varying the partitioning to match directly Airprox proportions, is that radar systems designers could be penalised for their over-achievements and successes. Very low rates of accident (or rather incident) would effectively reduce the risk design budget. For example, if the radar inaccuracy budget corresponded to 10 incidents and just one occurred, then radar accuracy would have to be made 10 times more stringent in risk terms. This type of TLS problem has already been recognised (Davies and Sharpe, 1993) in respect of mid-air collisions, which have reduced markedly as a proportion of total accidents compared to the statistics in the mid 1970s. Davies and Sharpe (1993) noted that there would be a large variability in the mid-air collision proportion of the total number of accidents because of the statistical fluctuations in the small number of collisions, even over quite long periods. But they also commented: 'To take a hypothetical case, if no collisions occurred during the period in question, it would be impossible to determine a sensible target level of safety for design purposes on such a basis.' They therefore retained the collision proportion derived from previous work, and noted that 'for system design purposes, any new target level of safety will be at least as demanding (in relative terms) as the previous version.'

Is there a different way of deriving the TLS that validates this approach and hence making it possible to justify the 10% factor? Ways of thinking about these kinds of issues have in fact been developed through the European Commission DGVII ARIBA (1999) project. ARIBA stands for 'ATM system safety criticality Raises Issues in Balancing Actors responsibility'. This project addressed certification in ATM services and, inter alias, attempted to build an 'accident risk tolerability matrix' for air traffic operations. To begin with, it is necessary to set out a simplified version of the JAA classification criteria adopted in JAR AMJ 25.1309 for tolerability criteria for any failure condition of a technical system. A 'failure condition' is defined by:

> "Failure condition: A failure condition is defined at the level of each system by its effects on the functioning of that system. It is characterised by its effects on other systems and on the whole system. All single failures and combinations of failures … which have the same effects on the system under consideration are grouped in the same failure condition."

Table 1. Definitions of frequency levels.
(From JAR AMJ 25.1309 'acceptable' numerical frequency ranges for each flight hour.)

Description	Estimate of Frequency	Frequency per aircraft flight hour
Probable	Anticipated to occur one or more times during the entire operational life of each aeroplane.	More than 10^{-5}
Remote	Unlikely to occur to each aeroplane during its total operational life but which may occur several times when considering the total operational life of a number of aeroplanes of the type.	Between 10^{-7} and 10^{-5}
Extremely remote	Unlikely to occur when considering the total operational life of all aeroplanes of the type, but nevertheless, has to be considered as being possible.	Between 10^{-9} and 10^{-7}
Extremely improbable	So unlikely that they are not anticipated to occur during the entire operational life of all aeroplanes of one type.	Less than 10^{-9}

First, each failure condition is classified according to its severity. The JAA qualitative definition of severity relevant to a mid-air collision is:

> Catastrophic: Failure conditions which would prevent continued safe flight and landing.

For each failure condition a classification of its frequency or probability of occurrence is next given. Qualitative definitions of probability according to the JAA standard are given in Table 1, along with 'acceptable' numerical frequency ranges for each flight hour. JAR AMJ 25.1309 (Section 4: Background) essentially assumes that there are 100 possible failure conditions on an aircraft that can prevent continued safe flight and landing. This assumption is similar to assuming that there are 100 safety-critical systems on each aircraft (there are around 70 safety-critical systems on the A320). The allowable probability of a catastrophic accident from a system failure (as distinct from a human error, weather and other natural environmental occurrences) is taken by JAA to be '1 in 10 million hours'. This is then apportioned equally among the 100 failure conditions. This gives a maximum permissible frequency of occurrence of each catastrophic failure condition of one per thousand million hours of flight, and hence a maximum acceptable probability of catastrophic failure of 10^{-9} per flight hour.

The European Commission ARIBA (1999) project has attempted to build an accident risk tolerability matrix for air traffic operations on HSE (UK Health and Safety Executive, 1992 & 1999) lines. The key point is that UK industry safety assessments usually use the HSE studies and guidelines about 'tolerable' and 'acceptable risk', and these are probably the most well developed decision making frameworks regarding the control of risk in Europe. A checklist of (simplified) HSE definitions is:

- ALARP Principle. The principle that no risk in the tolerability region can be accepted unless reduced 'As Low As Reasonably Practicable'.
- Intolerable Risk. A risk that cannot be accepted and must be reduced.
- Tolerability Region. A region of risk which is neither high enough to be unacceptable nor low enough to be broadly acceptable. Risks in this region must be reduced "as low as reasonably practicable (i.e. the ALARP principle)".

Table 2. Failure condition tolerability matrix adapted from JAR 25.1309 (Annex A).

Probability level (frequency per flight hour)	Severity			
	Minor	Major	Hazardous	Catastrophic
Probable: $>10^{-5}$	Tolerable	Intolerable	Intolerable	Intolerable
Remote: 10^{-7}–10^{-5}	Negligible	Tolerable	Intolerable	Intolerable
Extremely remote: 10^{-9}–10^{-7}	Negligible	Negligible	Tolerable	Intolerable
Extremely improbable: $<10^{-9}$	Negligible	Negligible	Negligible	Tolerable

The JAR AMJ 25.1309 guidance then allows failure conditions with the combinations of severity and frequency shown in Table 2 (definitions of the terms used are set out in ARIBA (1999)). [NB: it must be stressed that the words 'Intolerable', 'Tolerable' and 'Negligible' are as suggested by ARIBA, not the JAA.] The key point from the JAA and ARIBA documents (see also Eurocontrol SRC, 2000) is that 10^{-9} is taken as an acceptable/tolerable risk rate for *any* aircraft system failure condition, albeit that safety managers should endeavour to reduce such types of failure by reasonable means. So the obvious question is: "Does a gross radar error in the aircraft's 'environment' correspond in safety terms to such a failure condition?" As the JAA material deals with *all* kinds of system failures leading to catastrophic accidents, it is reasonable to argue that a gross radar error would indeed fall into this category, so that the 10^{-9} figure would be a 'RTLS ceiling'. [NB: the match between the two types of requirements is discussed in Eurocontrol SRC (2000), Section 3.1.]

To summarise this section: two approaches have been used, the first based on the current model, making it consistent with safety budgeting but retaining the assumption of the 10% factor; the second method has examined JAA and ARIBA reference material and argued that a gross radar error would fall into the category of a system failure condition leading to a catastrophic accident. Both methods lead to a TLS of 1×10^{-9}, which is higher than assumed in previous work on the subject. Note that in mid-air collision risk work, it is always assumed that every collision constitutes two fatal accidents, so the collision risk target is $0{\cdot}5\times10^{-9}$ collisions per flight hour.

3. EXAMPLE OF DYNAMIC AND STATIC CALCULATIONS.

The equations from Part 1 (see List of Symbols at end of paper) are:

$$P_z = P_z(0)\left\{1+\left(\frac{w}{u}\right)\left(\frac{G}{H}\right)\right\} \tag{1}$$

$$C(R) \approx p\left\{1+\left(\frac{\mu}{\lambda}\right)^2\right\}\mu e^{-\mu R} \tag{2}$$

$$P_y(R) = 2G\left\{1+\frac{v}{u}\right\}C(R) \tag{3}$$

$$Rk = P_z \sum RF(R)P_y(R) \tag{4}$$

How do the static and dynamic results compare, noting that Rk is to be compared with the radar TLS derived in the previous section?

Table 3. Frequency of use of small separations in area with 5 nm radar separation minimum. (Adapted from Table 1 of Sharpe (1991), 5904 flight hours in controlled airspace within 120 nm of Debden Radar with 5 nm applicable *and* maximum vertical separation of 500 feet.)

Horizontal separation	Number of aircraft pairs		
	Level flight	Transition	Aggregate
<5·0	19	363	382
5·0–5·5	3	234	237
5·5–6·0	1	303	304
6·0–6·5	10	326	336
6·5–7·0	10	410	420
7·0–7·5	8	430	438
7·5–8·0	5	406	411
Total	56	2472	2528

Table 4. How the $C(R)$ scaling factor reduces the relative risks. (Derived from Sharpe (1991), Figure 10.)

Horizontal separation (nm)	Aggregate number	Scaling factor for $C(R)$	Products of columns 2 and 3
<5·0	382	1	382
5·0–5·5	237	1	237
5·5–6·0	304	0·2249	69
6·0–6·5	336	0·0505	17
6·5–7·0	420	0·0114	5
7·0–7·5	438	0·0026	2
7·5–8·0	411	0·0006	1
Total	2528		713

It was noted in Part 1 that $C(R)$ is a rapidly declining function of R, so the calculation of Rk will be dominated by values of R planned to be at and near to S. Table 3 is taken from Sharpe (1991). There obviously has to be a 'health warning' about this data, which is used here purely to illustrate the calculation: in over a decade the density of traffic and the frequency of use of the separation minimum will have tended to increase – although sector and routeing re-design may well have countered such trends to some extent. It should also be noted that aircraft pairs separated by markedly less than 5 nm should be eliminated from the calculations. A significant loss of separation of this nature would generally reflect a failure by the pilot(s) and/or controller(s) rather than a risk caused by radar inaccuracy. Aircraft pairs with a CPA of just under 5 nm would however be included in the risk calculations (being taken as at 5 nm), as they might reflect typical scatter in radar positional accuracy – 'normal technical inaccuracies'. Aircraft pairs with a much lower achieved separation as a consequence of a radar error would in fact be excluded from this data set, as it is the distribution of *planned* separations that is of interest. The statistical distribution of distances below 5 nm is not examined in Sharpe (1991), so it is assumed that the aircraft pairs are indeed all close to 5 nm, rather than being eliminated from the calculations. The $P_z \sum RF(R) P_y(R)$ factor in equation (4) is calculated in Table 4, again using the Sharpe (1991) data. The weighted number of small separations, i.e. adjusted to give the risk for incidents at the separation minimum S, is 713. This is to be compared with the 'very conservative' value of 2528 assuming that all small separations were actually at the separation minimum. The present value is 0·282 times the earlier one. Note that this estimate does not attempt to adjust for the typical scatter in radar positional accuracy, i.e. these achieved separations are not modified to estimate planned separations. Order of magnitude calculations indicate that such adjustments would be more than countered by the cautious assumption that all the data in each range cell is counted as being at the lower end of the cell, e.g. data in the cell 5·0–5·5 nm is taken to be at 5 nm. The height of the collision box in Sharpe (1991) is 110 feet, giving a probability of overlap for transitioning traffic of 0·11. Sharpe uses a $P_z(0)$ value of 0·40. Sharpe then estimates P_z by weighting the level flight and transition flights by 5% and 95% respectively – a cautious estimate given the data in Table 4. [NB: Sharpe's calculation is correct but slightly cryptic, because he does not show a diagram on the lines of Figure 8 in Part 1.] P_z is thus estimated by Sharpe at 0·125. Still ignoring the $\left\{1+\left(\frac{w}{u}\right)\left(\frac{G}{H}\right)\right\}$ factor, i.e. calculating an 'effective $P_z(0)$', the present methodology, using the recent NATSPG value for $P_z(0)$, gives:

$$\text{Effective } P_z(0) = (\pm 0{\cdot}05)\,.\,(0{\cdot}48) + (0{\cdot}95)\,.\,(0{\cdot}11) = 0{\cdot}129$$

This is 1·03 times the earlier value.

Gathering together the results above into equations (1) to (4) gives:

$$\text{New estimate of risk} = \text{Standard Risk} \times 1{\cdot}03 \times 0{\cdot}282 \times (\text{dynamic factor}) \tag{5}$$

where the factor in brackets is:

$$\text{dynamic factor} = \left\{1+\left(\frac{w}{u}\right)\left(\frac{G}{H}\right)\right\} \times \left\{1+\frac{v}{u}\right\} \tag{6}$$

Thus, to estimate collision risk it is necessary to be able to estimate the values of u, v and w. In fact, it is necessary to estimate appropriately weighted averages for the velocity ratios $\frac{w}{u}$ and $\frac{v}{u}$. Weighting, rather than just inputting average values for the velocities, is necessary because slowly passing aircraft will have a low value of u and passing-behind manoeuvres would have much larger values; thus the ratios could potentially be very variable.

The two questions to address are therefore: what are the values of u, v and w at observed passings, and what how would they differ in hazardous incidents? The answer to the first question requires measurements of traffic data. This is not an intrinsically difficult exercise – existing measurement programmes to gather data on passing aircraft could readily be expanded to collect this velocity data. The second question is much more difficult, because it requires an understanding of the nature of near mid-air collisions as compared to normal passing configurations. The next section puts forward strong arguments that a mid-air collision would most probably occur in circumstances when the controller might – conceivably – be seriously misled by correlated gross errors. In such circumstances, the velocity data at passing might be expected to be similar to those in normal passing configurations.

4. CONTROLLERS' ABILITY TO INTERPRET DISPLAYED INFORMATION.

The analysis and calculations above have largely followed the logic and processes in earlier UK documents – the standard UK methodology. However, it is important to ask whether the right questions are really being posed.

The underlying philosophical problem is that inaccuracies in radar position cannot 'of themselves' cause mid-air collisions: it is the decisions made by controllers based on faulty radar information that might cause them. But it cannot be appropriate to construct detailed logical/human factor models of controller methods and cognitive processes in such circumstances – for how would such models be verified empirically? The need is therefore for something that is simple, robust and cautious – and which relies on measurable quantities. What could reasonably be assumed about the controller's ability to interpret displayed information? Are there valid 'absolute minimum' assumptions that will generate useful results?

Is the collision mechanism assumed in the above calculations the right one? There are good reasons for believing that it is not sufficient. In particular, Eurocontrol SRC (1998) notes the particular importance of cases of correlated position errors, where the 'apparent separation' of aircraft from radar plot information could be much larger than the real separation. Eurocontrol SRC (1998) notes that these correlated position errors can arise from 'partial loss of resolution' and 'multi-path' (distortion of the radar beam due to the characteristics of the terrain surrounding the radar site – but this would generally have a very limited effect on the separation applied, as it would tend to affect adjacent aircraft plots in the same way). But where are the consequences of correlated errors reflected in the calculations above? These kinds of problems were recognised in the earliest work, e.g. RSSWG (1982):

> "This static or snapshot analysis is artificial … . The controller will in reality base his judgement of their actual separation (or their likely future separation at the anticipated point of passing) not only on the current displayed positions but on their track histories."

The starting point for the radar separation minimum has to be the ICAO rules (ICAO Doc 4444 Part VI, para. 7.4.3):

> "… the radar separation minimum shall be prescribed by the appropriate ATS authority according to the capability of the particular system to accurately identify the aircraft position in relation to the centre of the radar position symbol".

So what do – or more relevantly what *should*-the words 'accurately identify' and 'capability' really mean in this context? What operational definition of accuracy would be most appropriate for the radar system designer and controller? The reasoned deductions in the following paragraphs are not based on describing what controllers would generally do, but rather on setting criteria based on what they could *not* reasonably be expected to do.

Consider the illustration in Figure 1, which shows the plan view of aircraft B passing aircraft A. The frame of reference fixed in A is used, i.e. A is at rest and B has the relative velocity of the two aircraft. The controller observes B as a sequence of radar plots – B1 to B6 are shown here. In this scenario, B1, B2, B4, B5 and B6 are very accurate radar plots: the true track is a straight line through these plots. Unfortunately, one plot, B3, is badly affected by gross radar errors, and it appears shifted in the y-axis from the true position of the aircraft (shown as a 'ghost' plot). If the controller were to make plans and take decisions on the basis of two successive radar plots – and nothing else – then there would be a choice of false tracks, as shown. Acting upon them would in both cases result in the need for a very sharp manoeuvre of the B aircraft, i.e. changing its direction flight quite dramatically. The controller would also have to ignore the apparently very large speed of the aircraft.

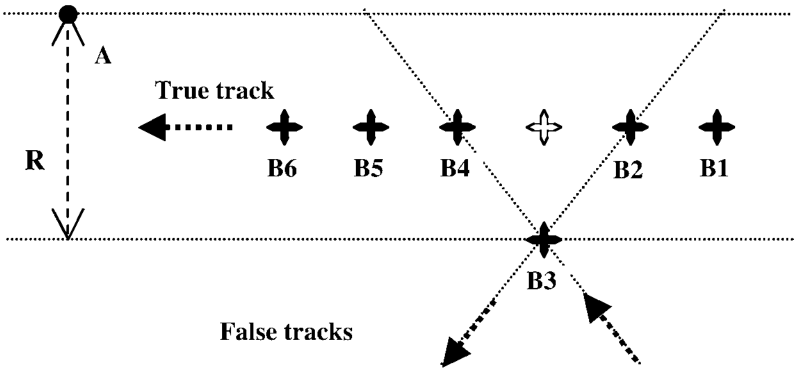

Figure 1. Controller potentially misled by single gross radar error. (Plan view in A's frame of reference.)

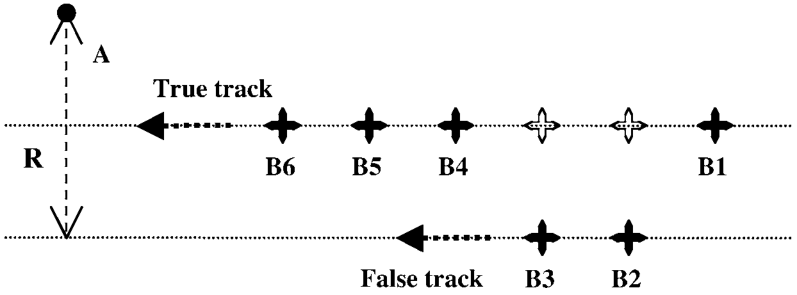

Figure 2. Controller potentially misled by double gross radar error. (Plan view in A's frame of reference.)

For example, if (to make the sums easy) plots were displayed at 5-second intervals then a 5+ nm change from one plot to the next would imply an apparent speed of 3600 knots.

Next consider Figure 2. In this scenario, B1, B4, B5 and B6 are again very accurate radar plots: the true track is again a straight line through these plots. Unfortunately, on this occasion, two plots, B2 and B3, are badly affected by gross radar errors and they appear shifted in the y-axis from the true position of the aircraft – shown as 'ghost' plots. If the controller were to visualise/estimate future aircraft positions *and* base planning decisions solely on the basis of these two plots, i.e. immaterial of flight plan information and past trail dots, then the judgement could be that B was on the false track shown. Thus, in such a case the belief would be that B would pass A at a distance R, greater than the separation minimum. In truth, B would follow its true track and pass A at only about half the aimed-for separation, according to the diagram. Hence, in this instance there would not be a mid-air collision but there would have been a substantial breach of the separation minimum. But a larger error in the B2 and B3 positions could have produced an accident. Note that the relative aircraft velocities in such a collision would largely be due to navigational 'flight technical errors', i.e. with the pilots endeavouring to keep to their tracks. This picture implies that w and v would usually be small velocities, while u would depend on the kinds of aircraft passings observed in practice (see Section 2 of Part 1). So Figure 2 is an illustration of how a mid-air collision might 'just conceivably' arise. It shows that the controller would need to have been misled by two plots with large errors rather than a single rogue plot. In practice, individual gross errors would probably be 'obvious' to the controller, and hence would be disregarded in decision-making. A controller would surely not make reasoned decisions on the likely path of an aircraft by observing a single plot – for this could never provide any information on the direction and speed of the aircraft. Hence a minimum of two plots is needed for a controller to make any kind of judgement on the likely flight path.

The convolution integral calculation in Part 1 essentially supposes that the controller is aware of the direction of the aircraft's velocity vector, but is then misled by a single gross positional error into believing that aircraft B will pass aircraft A with at least the minimum separation. But if, through intense work pressure or other extreme circumstances, a controller has little time to scan plots, and hence were to make decisions in such a way, would this really constitute 'positive radar control' as envisaged in MATS Part 1? Could such controller decision-making properly be termed 'reasoned' or the 'achievement of planned separation.'

These illustrations pose fundamental questions about the limits of 'reasonableness' for a controller's operational practices. The radar minimum should surely not be determined on the basis that a controller acts irrationally or on wholly insufficient evidence. The argument in the previous paragraphs is that the controller's assessment and decision-making in this context can be based on the *assured* achievement of some specific – albeit very low – performance level. This is not to assert that a controller would actually behave in such a way in practice, as the behaviour described above would be verging on the absurd for a professional controller, but that this determines a limiting boundary for the extremes of reasonable decision-making by a controller. The simplest possibility is to take the 'minimum' controller performance when using radar to be that the controller makes no decisions unless at least (sic) two proximate aircraft plots have been viewed. By 'proximate' is meant two plots during a normal sequence of five (?) trail dots – some of which may be missing. This means that the controller could, in theory, estimate the aircraft flight path – or at least its velocity vector – from these rogue plots. In these circumstances, there are two decision criteria that the controller effectively checks that:

- The first plot has to be subject to a (relative) gross error.
- *And* a proximate plot has to have an error of a comparable size.

It is stressed that the controller carries out *both* these tests before putting the aircraft on passing flight paths. This effectively implies that the convolution integral (and hence the lateral probability $P_y(R)$ and in turn the collision risk) would need to be multiplied by a conditional probability – 'probability of a gross error given that a proximate plot has a gross error'. Introducing this additional ingredient into the calculation would therefore add realism *and* ensure that attention is focused on the frequency of serious correlated errors. Note that these are independent decisions by the controller: the distribution of relative errors $C(R)$ does not determine, nor is it determined by, the conditional probability of proximate gross errors. [NB: if the conditional probability is near unity, then this implies localised regions with gross errors.]

In terms of the statistical fitting of data, the cautious approach would be to fit the tails of the probability distribution to 'raw data', i.e. without trying to exclude 'rogue data plots'. This would then be multiplied by the observed conditional probability for proximate gross errors. Thus, if only 5%, taken just for illustration, of gross errors had a proximate gross error, the tail of the lateral probability $P_y(R)$, and hence the collision risk, would have to be scaled down by a factor of 20. This conditional probability factor has considerable consequences as regards the size and nature of radar data collections. If it is not properly taken into account, then the data collection could be unnecessarily large.

5. CONCLUSIONS.

The derivation of the en route radar minima is important both in itself, i.e. ensuring that the minima are acceptably safe, and because of its 'reverse path' – a particular separation minimum necessarily feeds back into requirements on radar design criteria, data processing and performance. The estimation of the radar separation minimum has been significantly improved by the use of a new 'dynamic' methodology, as set out in Part 1. Specific conclusions from a review of safety targets and comparisons of dynamic and static calculations using work here are that:

i. A review of safety targets for en route ATC radar separation suggests that the existing TLS is over-cautious. If risk budgeting principles are followed consistently, a 'radar minimum TLS' of $1{\cdot}0 \times 10^{-9}$ fatal aircraft accidents per flying hour is appropriate. This rate is consistent with JAA material dealing with system failure conditions leading to catastrophic accidents.

ii. Dynamic and static calculations using published data are compared. The standard static model is actually a special case of the dynamic methodology, but it must be stressed that some elements in the standard model have been shown not to be cautious assumptions. The new methodology shows clearly where there are problems with the traditional static calculations, and how to improve the estimation.

iii. The new methodology can be further developed by introducing a simple model of the controller's decision processes. This methodology starts with an analysis of the ICAO regulations governing radar separation minima and then considers how correlated errors could mislead the controller. The focus is not on describing what controllers would generally do, but on setting criteria based on what they could *not* reasonably be expected to do. Introducing this additional ingredient into the calculation adds realism and ensures that attention is focused on correlated gross errors – recognised as particularly hazardous. This model needs development and validation.

Focussed data collection would be an essential component of new risk estimates. The key information required would be on radar performance and the nature and frequency of use of radar separation, including the relative velocities for proximate events at closest point of approach and the frequency of correlated gross errors (through a conditional probability factor). If this latter factor is not properly taken into account, then the data collection and analysis could be inefficient.

ACKNOWLEDGEMENTS

This work arose out of discussions with Dr. Mike Shorthose of Helios Technology Ltd. I would like to thank him for perceptive comments on an earlier draft. I would also like to acknowledge the major steps forward made by my former colleagues Rob Hunter and Andrew Sharpe in the earlier work on this subject.

APPENDIX A

RISK METRICS FOR TLS

What is an appropriate Risk Metric for the risk of mid-air collision? Aviation risk can be measured in terms of the chance of an adverse event per unit of activity. There are a number of possible metrics that may conceivably be used to describe risk. These include (and this is obviously not exhaustive):

number of accidents number of fatal accidents number of fatalities	**per**	year aircraft flying hour aircraft mile passenger passenger hour passenger mile stage flight passenger journey

Thus, given sufficient data, the same 'risk situation' could be expressed by a variety of metrics, such as fatal accidents per flight hour, expected number of accidents per calendar year, and fatalities per en route passenger hour.

The controlled airspace TLS is not framed in terms of passenger fatalities. It was chosen to be measured in 'fatal aircraft accidents', i.e. accidents in which at least one person in the aircraft was killed, per so many aircraft flying hours. This rate covers collisions between aircraft – for any reason and in any spatial dimension – in en route flight in controlled airspace. 'En route' has the technical meaning of the phase of flight occurring between the end of the initial climb and the beginning of the initial approach. Brooker and Ingham (1977) explains why this metric was chosen. The choice of 'fatal aircraft accidents' was because the focus was on the performance of ATC, which deals with aircraft *not* the passengers they contain. So a collision involving two small aircraft is just as serious an ATC system failure as one with two large aircraft. Aircraft flying hours was chosen because it was a measure of the time aircraft were under control – the 'exposure to ATC'. Several reasons were given why other possible metrics would be unsatisfactory given this context, for example, the sole use of metrics such as the probability of an accident per year give no indication, by themselves, of the individual's exposure to the risk when taking a flight.

Hence, the TLS is not framed in terms of passenger fatalities because of the concentration on ATC performance. Brooker and Ingham (1977) comments that rates based on (e.g.) the number of fatalities, i.e. by multiplying the TLS by the average number of fatalities per accident would be little different in safety assessment terms. However, TLS calculations based on particular assumed sizes of aircraft would have to be re-examined should aircraft average sizes increase significantly – as indeed they have done for commercial aircraft over much of the last 30 years. The choice of a risk metric depends upon the type of risk one is attempting to measure. It is the end-user of the risk metric who should determine the most appropriate choice for the metric. It is surely very reasonable that a regulatory, standard-setting safety authority would consider the number of (fatal) accidents per mile/departure/flight hour to be the most appropriate metric of ATC safety.

References in this Appendix are listed in the main text.

REFERENCES

ARIBA (1999). WP6 Final Report Part II: Safety Cases for a new ATM operation. http://www.nlr.nl/public/hosted-sites/ariba/rapport6/part2/index.htm

Brooker, P. (2002a). Future Air Traffic Management: Quantitative En Route Safety Assessment Part 2 – New Approaches. This *Journal* **55**, 363–379.

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

Brooker, P. and Ingham, T. (1977). *Target Levels of Safety for Controlled Airspace*. CAA Paper 77002. CAA, London.

Davies, E. H. and Sharpe, A. G. (1993). *Review of the target level of safety for NAT MNPS airspace*. CS Report 9301. Civil Aviation Authority, London.

Eurocontrol SRC [Eurocontrol Safety Regulation Commission] (1998). *Guidelines for the Application of the ECAC Radar Separation Minima*. ASM.ET1.ST18.1000 – REP – 01.00, Eurocontrol, Brussels. http://www.eurocontrol.int/eatmp/library/documents/ECAC_Radar_Sep_Min.pdf

Eurocontrol SRC [Eurocontrol Safety Regulation Commission] (2000). *Safety minima study: review of existing standards and practices*, SRC Doc 1, Eurocontrol, Brussels. http://www.eurocontrol.int/src/index.html

FAA/Eurocontrol (1998). A Concept Paper for Separation Safety Modeling: An FAA/Eurocontrol Co-operative Effort on Air Traffic Modeling for Separation Standards. http://www.faa.gov/asd/ia-or/pdf/cpcomplete.pdf

Greenwood, J., Fearn, T. and Pitts, S. M. (1994). *Criteria for the accuracy of Radar required to support a given Separation Minimum*. CS Report 9449, Civil Aviation Authority, London.

HSE – Health and Safety Executive (1992). *The Tolerability of Risk from Nuclear Power Stations*. HMSO, London.

HSE – Health and Safety Executive (1999). *Reducing Risks, Protecting People*. HSE Books, London.

ICAO Doc 4444 (PANS/RAC), Part VI (Radar Services).

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

MATS Part 1 (2002). CAP 493 Manual of Air Traffic Services ('MATS') Part 1, Civil Aviation Authority, London. http://www.caa.co.uk/docs/33/CAP493__Part1.pdf

NATSPG – North Atlantic Systems Planning Group (2002). *Summary of Discussions and Conclusions of the North Atlantic Systems Planning Group*. NATSPG/38.

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

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

RSSWG (1982). *Radar Separation Standards Working Group Report: Volume One*. CAA Paper 82014, Civil Aviation Authority. London.

Sharpe, A. G. (1991). *Application of the 5 nm radar standard separation at ranges up to 160 nm from Claxby, Debden and Pease Pottage SSRs*. CAA Paper 91013, Civil Aviation Authority, London.

UK Airprox Board (1999 onwards – biannual). Analysis of Airprox in UK Airspace.

LIST OF SYMBOLS

u, v, w	relative velocity of aircraft B (in an encounter with aircraft A) in the x, y and z directions at the planned passing point
$P_z(0)$	probability of vertical overlap for two aircraft at the same altitude
P_z	probability of vertical overlap with the extended collision box, for aircraft either in level flight or transitioning through the altitude
$P_y(R)$	probability of an overlap of for a lateral deviation R
G, H	size of box containing aircraft – G is a side, H its height. The 'collision box' has twice these dimensions. The 'extended collision box' is the collision box with dimensions extended by velocity-dependent factors
p, λ, μ	parameters characterising the radar data errors according to the probability density expression $f(y) = \frac{1}{2}\{(1-p)\lambda e^{-\lambda\|y\|} + p\mu e^{-\mu\|y\|}\}$
$C(y)$	the convolution integral $\int_{-\infty}^{\infty} f(\xi) f(y-\xi) d\xi$ (a probability density function)
$F(R)$	frequency of planned radar separation of value R