International Civil Aviation Organization

WORKING PAPER

SASP-WG/WHL/21-WP/29
26/10/2012

SEPARATION AND AIRSPACE SAFETY PANEL (SASP) MEETING OF THE WORKING GROUP OF THE WHOLE

TWENTY FIRST MEETING

Seattle, United States, 29 October to 9 November 2012

Agenda Item 2: En-route separation minima and procedures - horizontal

Longitudinal Collision Risk Estimates under the NAT RLongSM Operational Trial

(Presented by Mr. Harry Daly)

(Prepared by Ms. Jacky Civil)

SUMMARY

At SASP-WG/WHL/20 in May 2012 the Mathematicians' Sub Group (MSG) reviewed WP/29 and WP/25 which detailed a longitudinal collision risk assessment in support of the North Atlantic (NAT) Reduced Longitudinal Separation Minimum (RLongSM) operational trial, performed using data collected prior to the initiation of the trial. The review concluded that the collision risk estimates should only be used in support of an operational trial, but not for a final collision risk assessment to support the transition of an operational trial into full operational use. For that, a collision risk assessment based on sufficient data collected during the operational trial was required. The UK accepted an action to perform the requested collision risk assessment for the NAT region.

To address this action, this paper details a longitudinal collision risk assessment based on 12 months of data collected from the Shanwick Oceanic Control Area (OCA) during the NAT RLongSM operational trial. It concludes that the longitudinal collision risk meets the Target Level of Safety (TLS), and therefore that the RLongSM procedure in the NAT region is acceptably safe.

The meeting is invited to endorse RLongSM as a globally applicable procedure via an amendment to ICAO Doc 4444.

1 Introduction

1.1 Background

At SASP-WG/WHL/20 in May 2012 the Mathematicians' Sub Group (MSG) reviewed WP/29 [Ref 1] and WP/25 [Ref 2] which detailed a longitudinal collision risk assessment in support of the North Atlantic (NAT) Reduced Longitudinal Separation Minimum (RLongSM) operational trial, performed using data collected prior to the initiation of the trial. The review concluded that the collision risk estimates should only be used in support of an operational trial, but not for a final collision risk assessment to support the transition of an operational trial into full operational use. For that, a collision risk assessment based on sufficient data collected during the operational trial was required. The UK accepted an action to perform the requested collision risk assessment for the NAT region.

To address this action, this paper details a longitudinal collision risk assessment based on 12 months of data collected from the Shanwick Oceanic Control Area (OCA) during the NAT RLongSM operational trial.

1.2 Longitudinal Reich Collision Risk Model

The collision risk analysis will be performed using the Longitudinal Reich Collision Risk Model (CRM). This can be written as:

$$N_{ax} = 2P_y(0)P_z(0)\left(\frac{\overline{|\dot{x}(m)|}}{2\lambda_x} + \frac{\overline{|\dot{y}_0|}}{2\lambda_y} + \frac{\overline{|\dot{z}_0|}}{2\lambda_z}\right)\frac{2\lambda_x}{\overline{|\dot{x}(m)|}}\frac{1}{T}\sum_s E(s)Q(s)$$

Where N_{ax} is the collision risk expressed as number of fatal accidents per flight hour. The meaning and derivation of each model parameter will be described in later sections, and a full summary of the adopted values for each parameter is given in Table 4-1.

2 Data

The data used as input to the collision risk analysis is from aircraft transiting the Shanwick region of the North Atlantic (NAT) during the first year of the RLongSM operational trial. This covers the period starting 1st April 2011 and ending 31st March 2012 (Note: the NAT RLongSM operational trial became active on 30th March 2011 and was available continuously during this period). The entire year was used in order to maximise the available data. The reason for not using a larger dataset incorporating the Summer months of 2012 was so that the weather patterns in Summer months would not be over-represented in the analysis.

2.1 Data Processing

2.1.1 ADS-C Reports

The specific information of interest was extracted from daily Shanwick Automated Air Traffic System (SAATS) audit files which contain ADS-C position reports and SAATS generated forecasts of future waypoint times. The post-operational data extraction failed for 26 daily audit files and in those cases data for the entire day was excluded. Those days were: 11, 12 April 2011; 4, 10, 13 September 2011; 8, 9, 15, 16, 22, 26, 27, 29 October 2011; 12, 14, 21, 22, 27, 28 November 2011; 18 December 2011; 6, 31 January 2012; 15, 16 February 2012; 10, 19 March 2012.

The data used from each ADS-C report (waypoint and periodic) consists of the following:

- Date;
- Time;
- Aircraft Identification (Callsign and Tail Registration);
- Position in lat/long;
- Flight Level;
- SAATS forecast time at the next waypoint.

2.1.2 Periodic Reporting Data

For the purpose of the analysis we require data for an 18 minute reporting period, without any interim waypoint or demand reports. To achieve this, pairs of successive periodic reports were extracted, together with the previous and next waypoint reports. A single reporting period transit was therefore represented by the data for four ADS-C reports; two waypoint reports bounding two periodic reports. It should be noted that many flights had two complete reporting periods between 20W and 30W (in either direction). In those cases each reporting period was considered to be a unique reporting period transit [see Section 3.3].

For each 18 minute reporting period transit, the following information was calculated:

- Local Route (i.e. a unique identifier for previous and next waypoints);
- Distance travelled between periodic reports (Great Circle Distance);
- Ground Speed between periodic reports.

2.1.3 Paired Periodic Reporting Data

Pairs of Leader and Follower aircraft were generated using the following rules:

- Leader and Follower travel on the same date;
- Leader and Follower are on the same Local Route (previous and next waypoints);
- Leader and Follower maintain the same Flight Level for the duration of the waypoint to waypoint transit;
- The actual time separation between Leader and Follower at the first waypoint is less than or equal to 20 minutes;
- The reporting periods of the Leader and Follower overlap in time;

- The Leader's first periodic report is earlier than the Follower's first periodic report. [The purpose of this criterion is discussed in Section 3.3]

Given the rules above, in some cases the Follower aircraft of one pair is also the Leader aircraft of a different pair.

2.1.4 Potential Bias

For unavoidable reasons, the paired periodic reporting data resulting from the processing described above may not be fully representative of RLongSM-capable aircraft behaviour within Shanwick. A specific flight will only be in the dataset if it has another aircraft following on the same track shortly afterwards. This will bias the analysis towards the behaviour of flights on the Organised Track Structure (OTS) during busy times of day. This is reasonable however, since it arguably represents the flights most at risk of collision due to erosion of longitudinal separation.

The requirements of the analysis also dictate that whole reporting periods between waypoint reports should be identified. This may result in an unbalanced representation of aircraft ground speeds within the analysis. For example, very fast aircraft with strong tail winds may only have a single full reporting period between 20W and 30W, whereas slower aircraft may have two full reporting periods between the same waypoints.

Finally, the requirement that both aircraft in a pair maintain level flight during the waypoint to waypoint transit may mean that any specific aircraft behaviour or risks associated with changes of altitude are not being captured.

2.2 Data Description

The data processing described in Section 2.1 results in 14,184 unique pairs of Leader and Follower reporting period transits, consisting of data from 19,429 unique flights. During this period it was estimated that a total of approximately 180,000 RLongSM-capable flights transited the Shanwick OCA. The number of flights estimated to be gaining benefit from RLongSM within Shanwick was 1,941 during the 12-month period.

The following sections characterise the data in terms of potential factors of interest such as direction, flight level, speed etc.

2.2.1 Direction

Table 2-1 shows the breakdown of the data by direction of travel and which pair of waypoints bound the reporting period.

Transit (Longitudinal Co-ordinates)	Number of reporting period pairs
15-20 (Westbound)	1337
20-15 (Eastbound)	12
20-30 (Westbound)	5721
30-20 (Eastbound)	7114
TOTAL	**14,184**

Table 2-1

2.2.2 Flight Level

Figure 2-1 shows the distribution of flights by Flight Level.

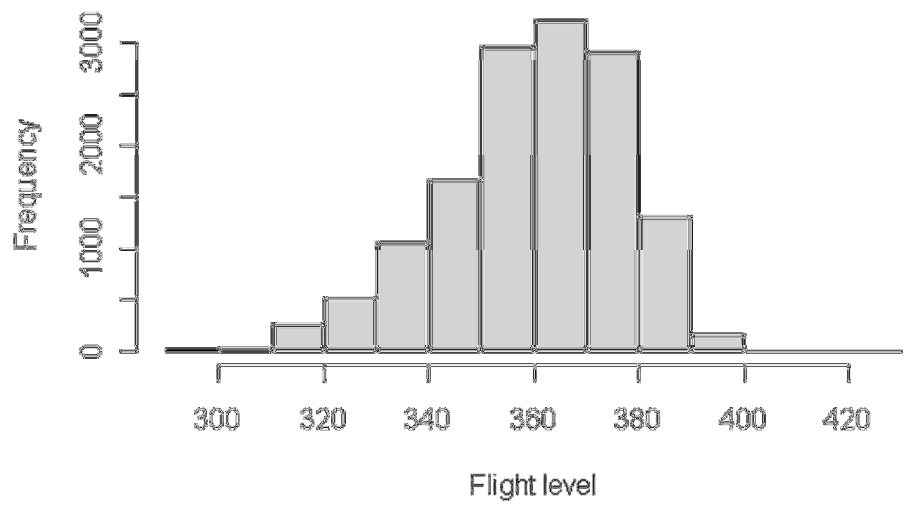

Figure 2-1: Histogram of pairs by flight level

2.2.3 Ground Speed

Figure 2-2 shows the distribution of flights by calculated ground speed (expressed in Nm/min), and Figure 2-3 shows a scatterplot of paired Leader and Follower calculated ground speeds. A high correlation between Leader and Follower speeds can be observed, which is expected of flights which are closely separated.

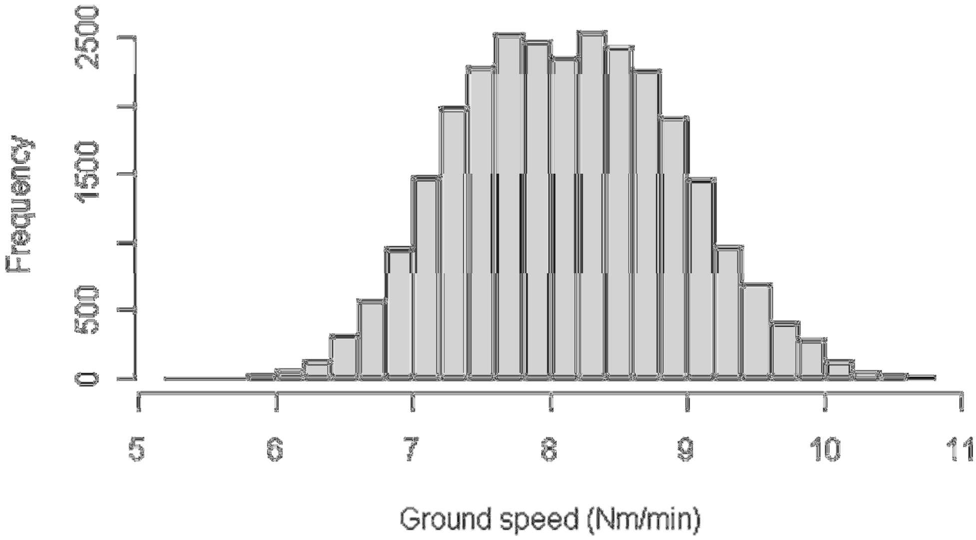

Figure 2-2: Histogram of Leader and Follower ground speeds

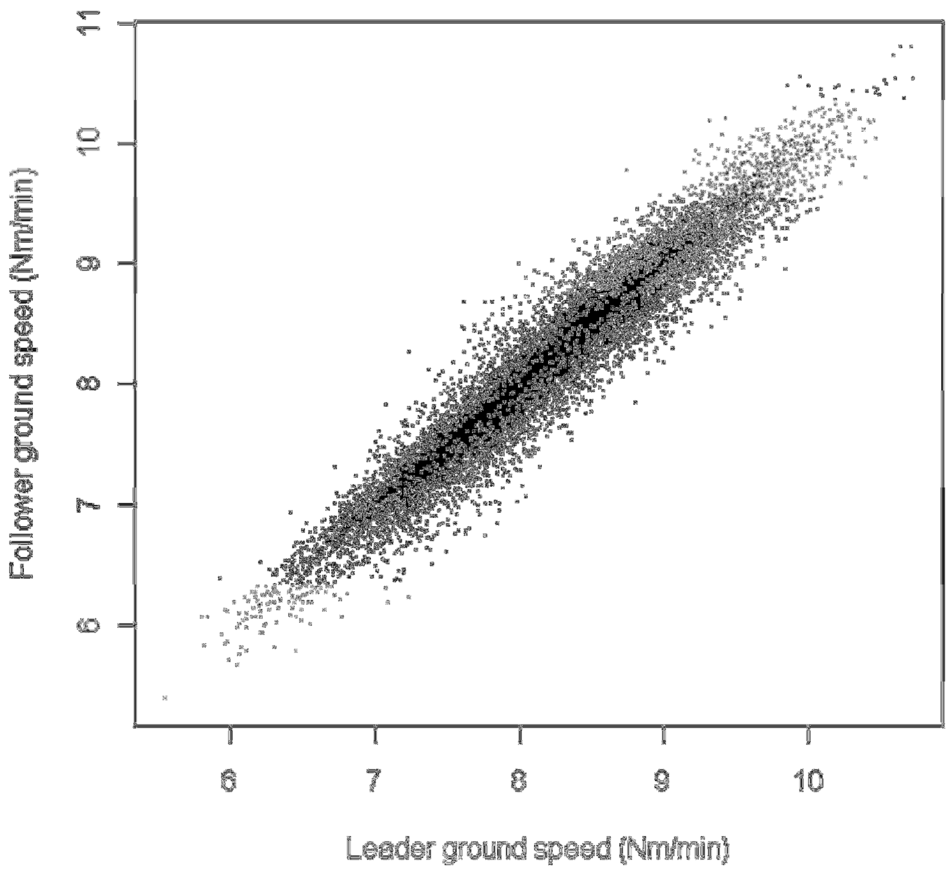

Figure 2-3: Paired Leader and Follower ground speeds

3 Methodology

This section describes the difficulties associated with this type of analysis, along with the methodology that has been used to overcome them.

3.1 Introduction

The data described in Section 2 was gathered for the purpose of estimating the two longitudinal Reich CRM components, $E(s)$ and $Q(s)$, where $E(s)$ is the distribution of separations and $Q(s)$ is the probability of losing s or greater minutes of separation during the at-risk period. There are two possible ways to define this pair of distributions:

1. $E(s)$ is based on the actual separations at the start of the at-risk period and $Q(s)$ is based on the actual change in separation experienced during the at-risk period;

2. $E(s)$ is based on the intended separations at the end of the at-risk period, forecast at the beginning of the period, and $Q(s)$ is based on the error in the forecast separations at the end of the at-risk period.

Under the first option, any catch-up or pull-away due to speed differences that the controller is aware of and willing to permit will be included within both components. This causes a problem since a catch-up would only be allowed when the initial separation between a pair of aircraft is large. Therefore, under this option, $E(s)$ and $Q(s)$ could not be considered to be independent of each other and should be estimated as a joint distribution. For this reason, the second option is preferable.

In practice, the available data cannot be used directly to determine the two parameters using either option. The difficulty is illustrated in Figure 3-1. The aircraft provide their periodic reports asynchronously, so neither the distance nor time separation is known at any of the periodic reporting points and the components for option 1 cannot be directly estimated from the data. Also, conformance checking at the periodic reports is performed by generating an updated forecast of the time that the aircraft is expected at the next waypoint and comparing it to the previous forecast. SAATS does not explicitly forecast where the aircraft will be at the next periodic reporting point, so a forecast separation at any reporting point does not exist and the components for option 2 cannot be estimated from the data.

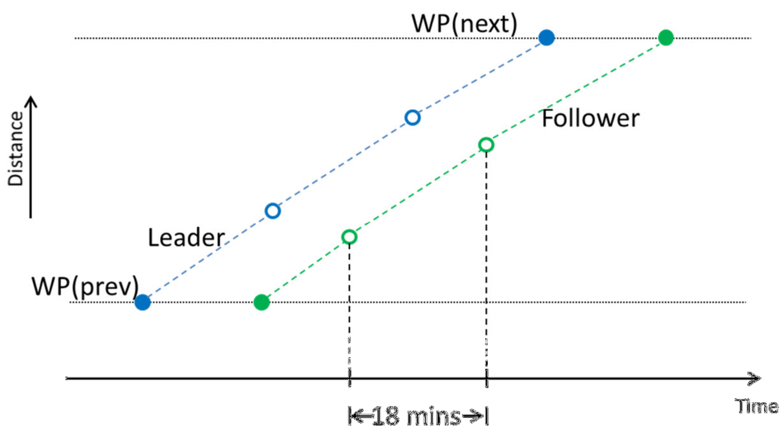

Figure 3-1: Illustration of a typical aircraft pair reporting at waypoints (filled circles) and periodic reporting points (empty circles)

There are two further difficulties for the analysis with respect to the definition of the at-risk period. In prior work a conservative assumption has been made that the aircraft provide their periodic position reports at the same time (synchronous reporting). The at-risk period can then unambiguously be considered as the reporting period plus an additional time allowance for the controller to understand the conflict, decide on a resolving action, communicate it to the pilot and for the resolving action to be initiated (Conflict Resolution Delay (CRD)).

Since aircraft report asynchronously in practice there will always be an interim Leader's periodic report between any pair of Follower's periodic reports. Thus the 18 minute reporting period for the Follower cannot be considered as continuously at-risk since there is an opportunity for the controller to receive and act on partial information regarding a potential loss of separation at the Leader's interim periodic report. Similarly, no information exists within the data as to the impact of the additional time allowance for the CRD.

3.2 Derivation of E(s) and Q(s)

At each periodic reporting point the current separation and forecast separations at future periodic reports are unknown. However, an updated forecast for the separation at the next waypoint can be calculated from the available data. Figure 3-2 illustrates the initial forecast separation at the next waypoint, s_1, based on the forecast times at the next waypoint for both Leader and Follower aircraft at the first periodic reporting point. This value, s_1, will be used to derive the distribution $E(s)$.

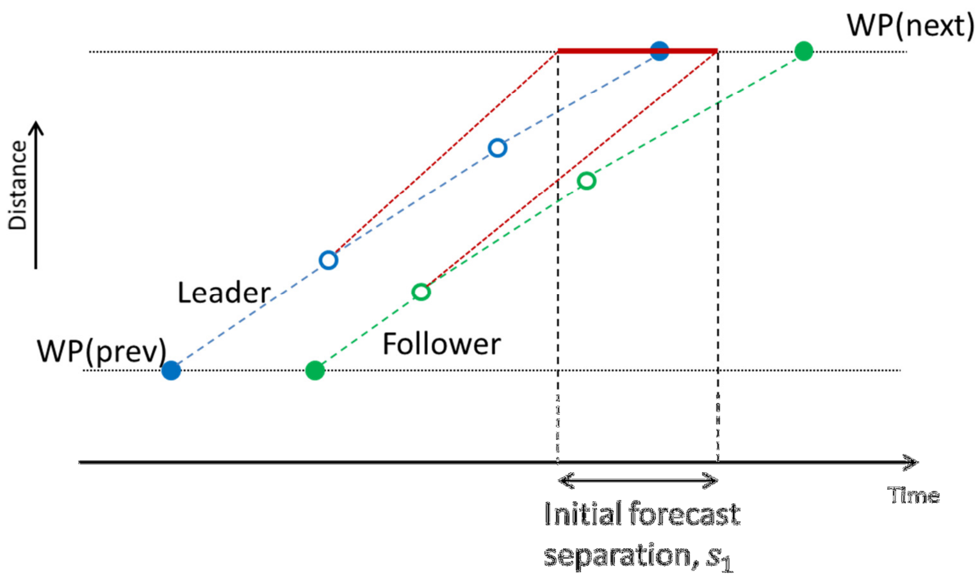

Figure 3-2: Initial forecast separation at the next waypoint

Figure 3-3 shows the final forecast separation at the next waypoint, s_2, based on forecasts made at the second periodic reporting point. The change between the two separation forecasts, $s_2 - s_1$, will serve as a proxy for the change in separation over the reporting period, and will be used as the basis for the distribution $Q(s)$.

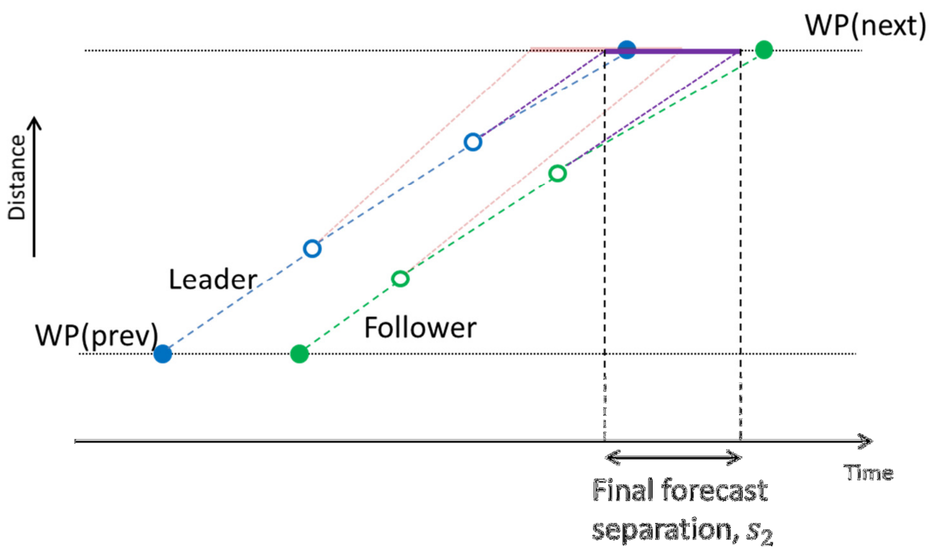

Figure 3-3: Final forecast separation at the next waypoint

One limitation of this approach is that in theory we are not assessing the risk of collision within the reporting period. In an extreme example, if the final forecast separation was equal to zero, this wouldn't necessarily indicate that all separation was lost within the reporting period, but instead that it was forecast to be lost at the next waypoint. This is discussed further in Section 3.4.

3.3 Reporting Period

Figure 3-4 illustrates a typical pattern of multiple reporting periods between 20° West and 30° West. In this circumstance there are three possible sets of periodic reports that could be considered as distinct paired reporting periods:

Set 1 - $\{PR_L(1), PR_L(2), PR_F(1), PR_F(2)\}$;

Set 2 - $\{PR_L(2), PR_L(3), PR_F(1), PR_F(2)\}$;

Set 3 - $\{PR_L(2), PR_L(3), PR_F(2), PR_F(3)\}$.

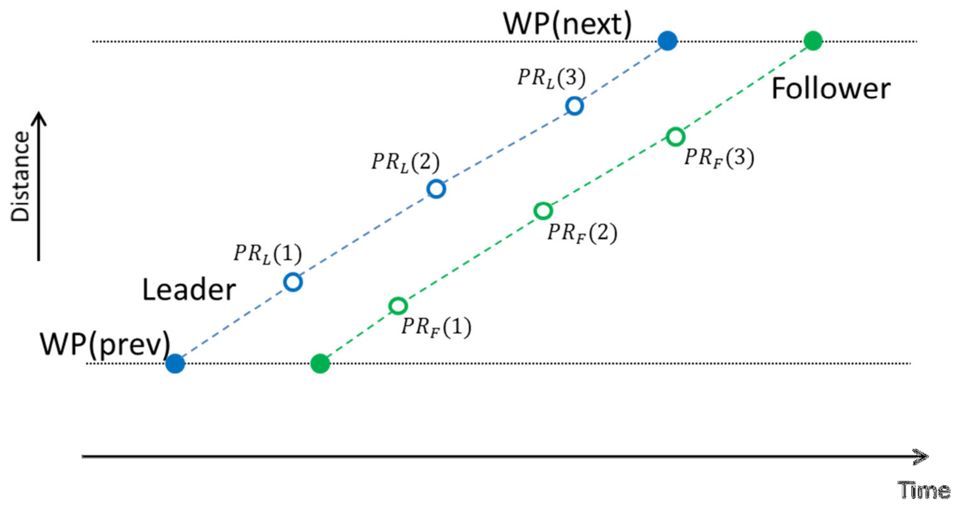

Figure 3-4: Typical periodic reporting for slower aircraft pairs

However, Set 2 contains only data which is duplicated from either Set 1 or Set 3. Therefore, Set 2 paired reporting periods are excluded from the analysis. This is accomplished by requiring that the Leader's first periodic report is earlier than the Follower's first periodic report (as specified in Section 2.1.3).

The reporting period is therefore explicitly defined as the time between the Follower's periodic reports. The Leader's periodic reports contribute to the calculation of forecast separation at the next waypoint (as in Figure 3-3), but for simplicity we ignore the possibility of the controller initiating an intervention triggered by the Leader's report. This is conservative, effectively equivalent to the assumption of synchronous reporting.

3.4 Conflict Resolution Delay

3.4.1 No Adjustment Option

Section 3.2 mentioned one limitation of the analysis being the possibility that the risk of collision is not being assessed within the reporting period. Instead, it could be considered that what is being assessed is the risk of a collision occurring between the first periodic report and the following waypoint report. If this is the case, the additional time between the second periodic report and the waypoint report could serve as a proxy for the Conflict Resolution Delay (referred to here as pseudo-CRD), with no further adjustment required to the analysis.

To have confidence that this would be sufficient to account for CRD, we should be able to observe that the variance of the derived Q(s) is larger, and thus the probability of extreme separation erosions is higher, when there is a longer additional time period following the reporting period before the next waypoint.

Figure 3-5 shows the observed distributions of separation gain/loss for five minute bands of pseudo-CRD, calculated as the difference between the time of the second periodic report and the next waypoint report, averaged across Leader and Follower. By eye there does not appear to be a consistent increase in variance or extreme observations when the pseudo-CRD is greater. This is confirmed by Table 3-1, which shows some increase in variance with higher levels of pseudo-CRD, but not sufficient evidence to give confidence that no adjustment to the analysis for CRD is required. Instead, the results in Figure 3-5 and Table 3-1 appear to imply that the change in forecast separation during a reporting period is primarily capturing the actual change in separation during the reporting period.

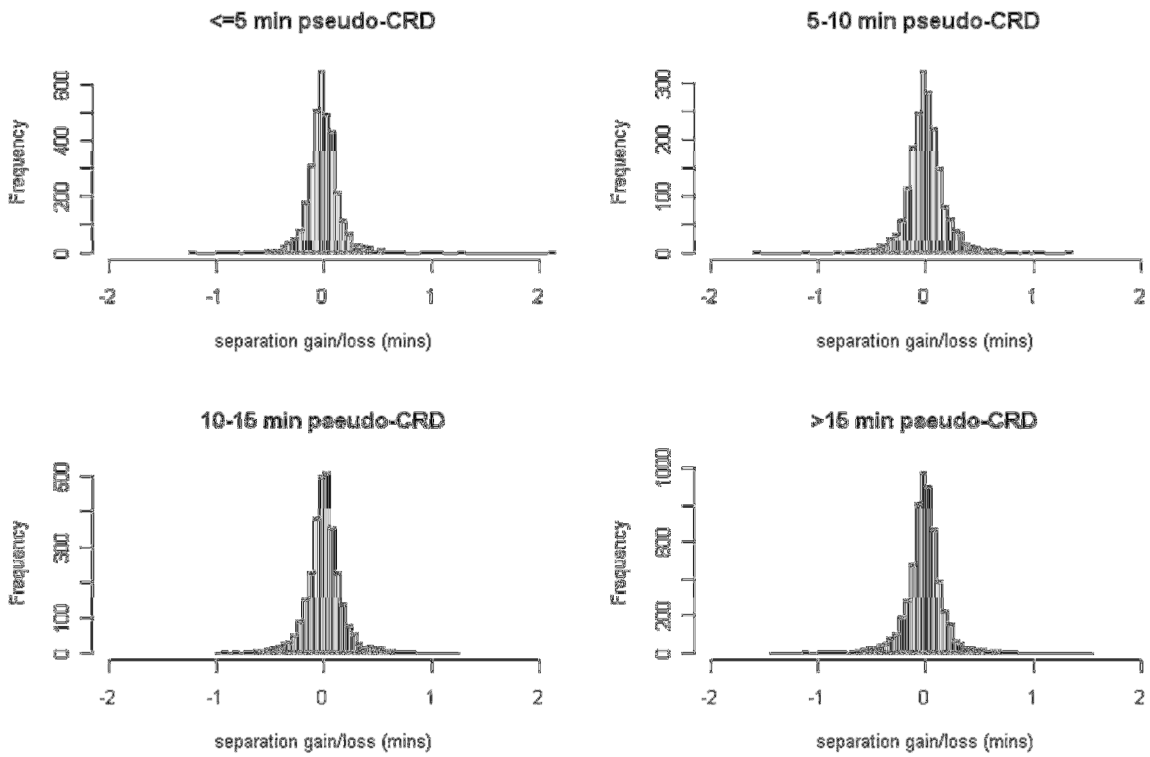

Figure 3-5: Observed separation gain/loss distributions at different ranges of pseudo Conflict Resolution Delay

Pseudo-CRD (minutes)	Separation gain/loss variance	Percentage of separation gain/loss greater than 1 minute
<5	0.027	0.21%
5-10	0.035	0.35%
10-15	0.033	0.07%
>15	0.038	0.29%

Table 3-1: Variance of separation gain/loss distributions and percentage of extreme observations, by pseudo-CRD

3.4.2 Scaling

To account for the additional duration of the at-risk period due to the CRD, it is necessary to scale upwards the variance of the separation gain/loss distribution derived from the 18 minute reporting period.

Prior work [Ref 1] explored in depth the related problem of scaling down the variance observed during waypoint to waypoint transits to give an appropriate estimate for the variance under shorter periodic reporting periods. This work was necessary since a strictly linear scaling (e.g. assuming that halving the period would halve the standard deviation of the separation gain/loss distribution) would not have been a conservative assumption. The outcome of this work was a multiple component auto-regressive model which (simplifying) scaled the variance according to the formula:

$$Var(gain/loss|t = T_2) = \frac{T_2^2}{T_1^2} Var(gain/loss\ |t = T_1) \frac{F(\alpha, T_2)}{F(\alpha, T_1)},$$

$$F(\alpha, T) = \frac{1}{T}\left\{\frac{1+\alpha}{1-\alpha} - \frac{2\alpha(1-\alpha^T)}{(1-\alpha)^2 T}\right\}$$

where α is the auto-correlation term, $T_1$is the time period over which the separation gain/loss variance is known and T_2 is the time period over which the variance is needed.

For this work, using plausible values for the time periods and auto-correlation parameter this would give an additional scale factor due to the ratio of $F(\alpha, T)$ terms in the range of approximately 0.95-0.99. For simplicity, and as a conservative assumption, we will assume this additional scale factor is equal to one. Therefore the standard deviation of the separation gain/loss distribution will be scaled upwards proportionately to the additional time due to CRD.

4 Derived Distributions and Parameter Estimates

4.1 E(s)

As described in Section 3.2 the distribution of separations, $E(s)$, will be based on the forecast separation at the next waypoint, where the forecast is made at the first periodic reporting point. Figure 4-1 shows the empirical distribution of forecast separations from 12 months of Shanwick RLongSM operational trial data.

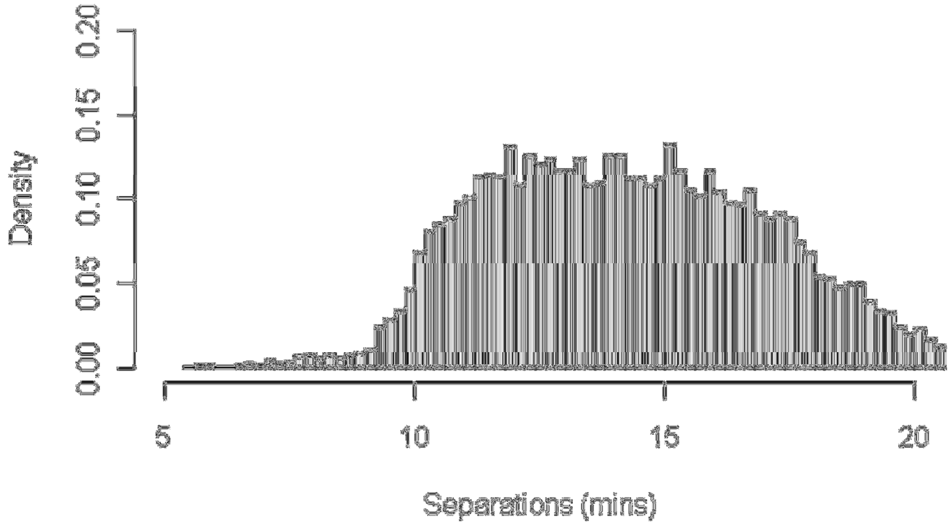

Figure 4-1: Empirical distribution of forecast separations at the next waypoint

It can be seen that despite longitudinal separations of less than 10 minutes being available to these pairs of aircraft, in the majority of cases the applied separations remain greater than this. Where RLongSM has been applied, it is generally to permit only a small erosion of separation below the previous rule of 10 minutes. The smallest forecast separation within this dataset was 5 minutes, 29 seconds, which is in line with regular controller practice to allow a buffer separation above the minimum. It should be noted that the data was selected for pairs of aircraft with initial separations at the first waypoint of less than 20 minutes, so greater longitudinal separations than those shown here do frequently occur in the Shanwick airspace.

It is appropriate that a distribution close to the observed empirical distribution should be fitted for the purpose of the baseline analysis. However, it is necessary to also assess the impact of theoretical distributions that represent possible future scenarios with greater take-up of RLongSM and increased traffic levels.

4.1.1 Baseline

A Gamma distribution is used for $E(s)$ as it is a flexible distribution form with varied shapes dependent on the parameters. It also allows easy integration of $E(s) \times Q(s)$ when $Q(s)$ is modelled using the Laplace distribution (see Section 4.3).

The Gamma distribution is naturally bounded below by zero. However, the lower limit of separations that would be permitted without intervention is five minutes. Therefore, we denote the forecast separation s as $s = 5 + x$, where x is considered the buffer separation above the minimum and has a Gamma distribution. The Gamma(k, θ) distribution has shape parameter k, scale parameter θ, and is written:

$$f(x) = \frac{x^{k-1}\exp(-\frac{x}{\theta})}{\Gamma(\mathrm{k})\theta^k}, x \geq 0$$

Figure 4-2 shows the fit of the Gamma(6, 1.5) distribution for x chosen for the baseline analysis. While this distribution does not provide a perfect fit to the data, it is broadly in line with the shape of the empirical distribution and is moderately conservative.

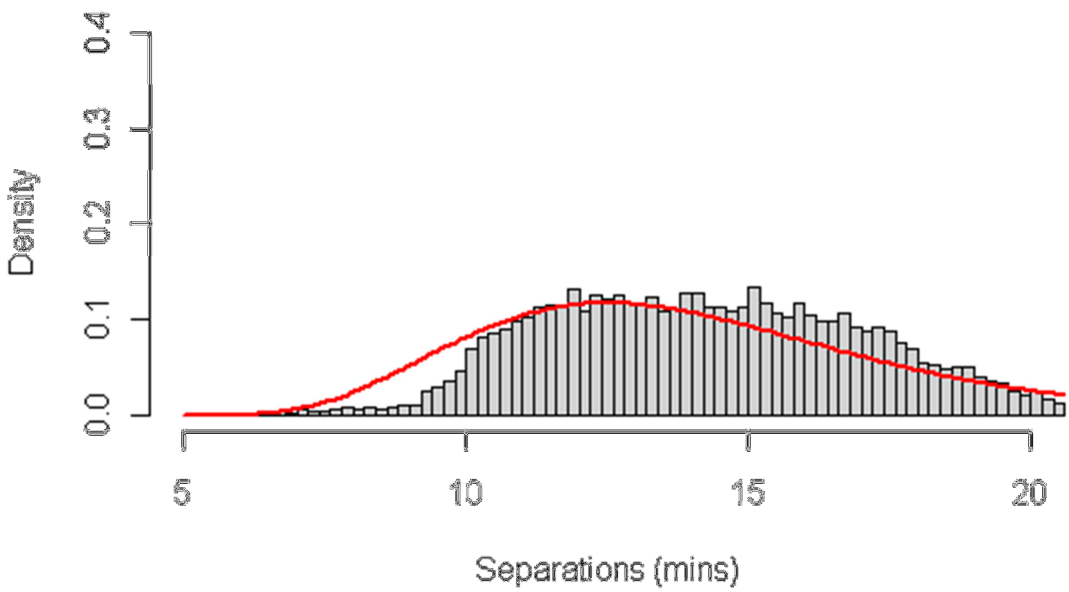

Figure 4-2: Histogram of forecast separations with Gamma(6, 1.5) distribution in red

4.1.2 Variations for Sensitivity Analysis

In addition to the baseline distribution representing our best knowledge of the current situation in the NAT, it is also necessary to assess the impact of theoretical distributions that represent possible future scenarios with greater take-up of RLongSM and increased traffic levels. Figure 4-3 shows a Gamma(5, 1.4) distribution that represents a plausible increase in RLongSM usage. Figure 4-4 shows a Gamma(2, 1) distribution that represents a much more extreme increase in RLongSM usage. This version of the distribution was used in the earlier collision risk assessment work that was used to support the introduction of the operational trial [Ref 2].

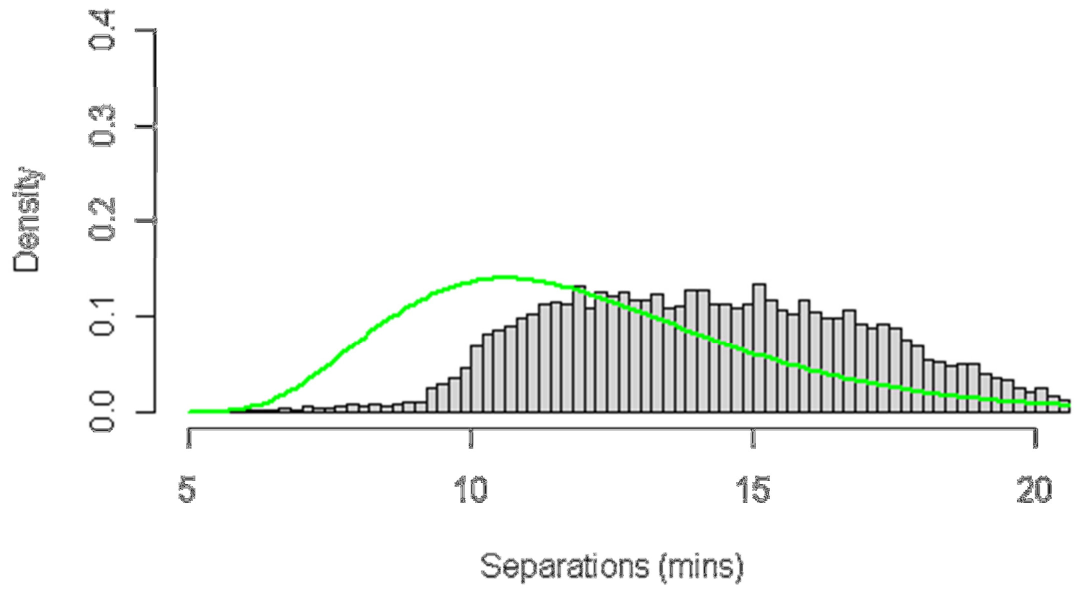

Figure 4-3: Gamma(5, 1.4) distribution representing plausible future increase in RLongSM usage

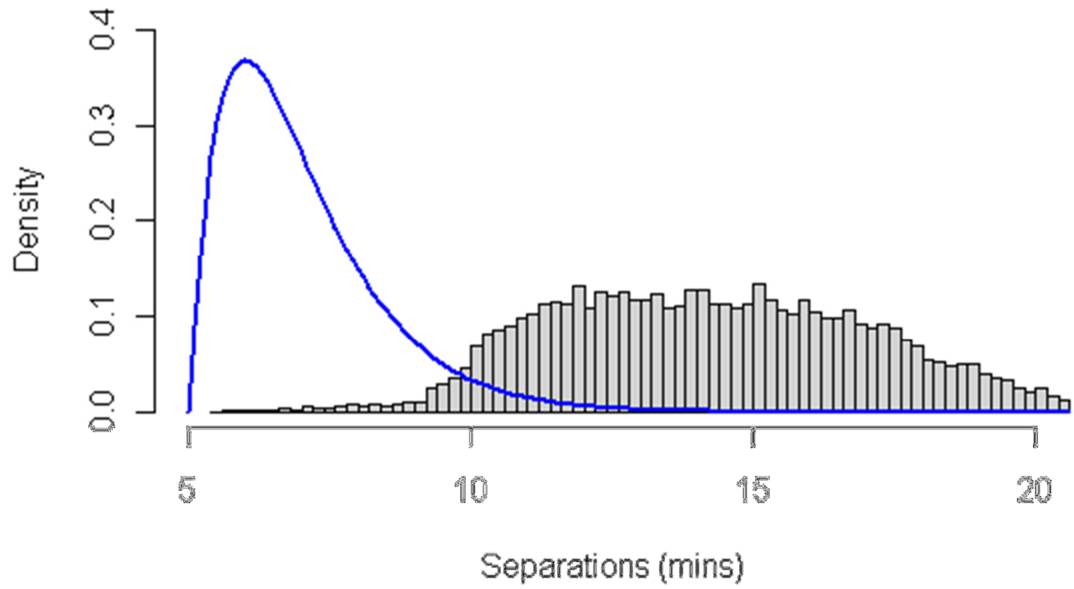

Figure 4-4: Gamma(2, 1) distribution representing extreme future increase in RLongSM usage

4.2 Q(s)

The probability of losing s or greater minutes of separation during the at-risk period, $Q(s)$, will be based on the difference between forecast separations at the next waypoint made at the two successive Follower's periodic reports. Figure 4-5 shows the empirical distribution of the forecast separation gain/loss based on 12 months of Shanwick RLongSM operational trial data.

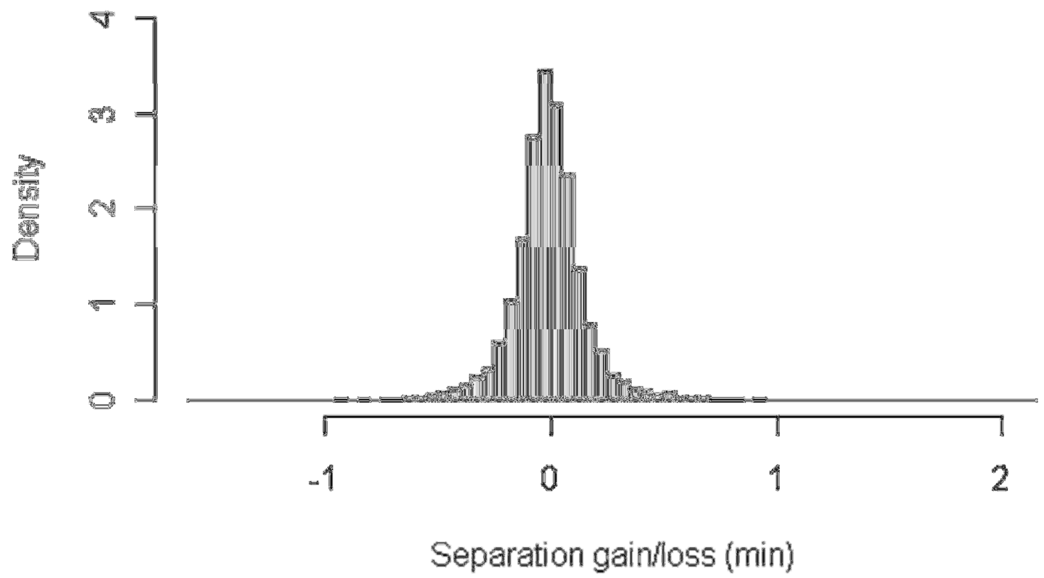

Figure 4-5: Empirical distribution for Q(s)

The distribution closely resembles a Laplace distribution (double-exponential distribution), which is commonly used to model this kind of data. The distribution is symmetric, mostly centred close to zero and has heavy tails. The largest observed separation loss was 1 minute 34 seconds [Note: this is erosion from intended longitudinal separation, not a Loss of Separation], and the largest observed separation gain was 2 minutes 8 seconds.

The derivation of an appropriate $Q(s)$ distribution is arguably the most important part of longitudinal collision risk modelling, since variations in the variance of the distribution can change the eventual collision risk estimate by several orders of magnitude.

4.2.1 Baseline

4.2.1.1 Simple Distribution

The best fit Laplace distribution (zero-centred with parameter b=0.12) is shown in Figure 4-6. The top two plots show the overall fit of the theoretical distribution to the empirical distribution as a histogram and in a quantile-quantile plot. The lower two plots show the fit at the extreme tails of the empirical distribution.

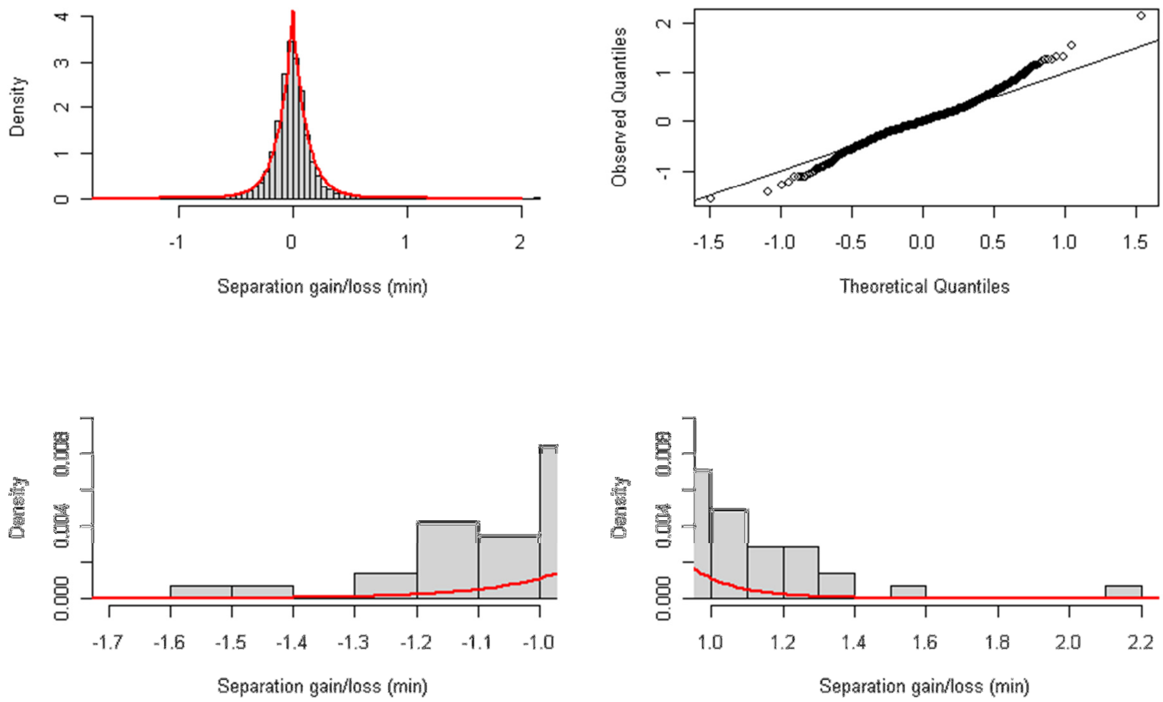

Figure 4-6: Plots showing the goodness of fit of the Laplace(0, 0.12) distribution

While the Laplace(0, 0.12) distribution fits the body of the data very well, the tails are clearly being under-represented. Since the purpose of fitting the theoretical distribution is to extrapolate the probability of separation losses more extreme than those observed, the poor fit of the tails renders this distribution unsuitable.

4.2.1.2 Mixture Distribution

In theory, it can be considered that the observed errors come from two distinct sources: technical errors which are random variations due to unavoidable causes such as the effect of weather, and operational errors due to mistakes by the pilot or controller, which will often result in a larger variation than can be observed by technical error alone. Thus the overall separation gain/loss distribution can be modelled as a mixture of two distributions.

Figure 4-7 shows the goodness of fit of a mixture distribution for the separation gain/loss. Both components of the mixture are Laplace distributions, with 89% of observations assumed to be drawn from a Laplace(0, 0.11) distribution and 11% drawn from a Laplace(0, 0.27) distribution. The parameters of the mixture distribution were derived using WinBUGS software to perform a Markov Chain Monte Carlo (MCMC) model fitting. The software derives posterior probability distributions for the parameters, from which the medians were selected to give point estimates for the model.

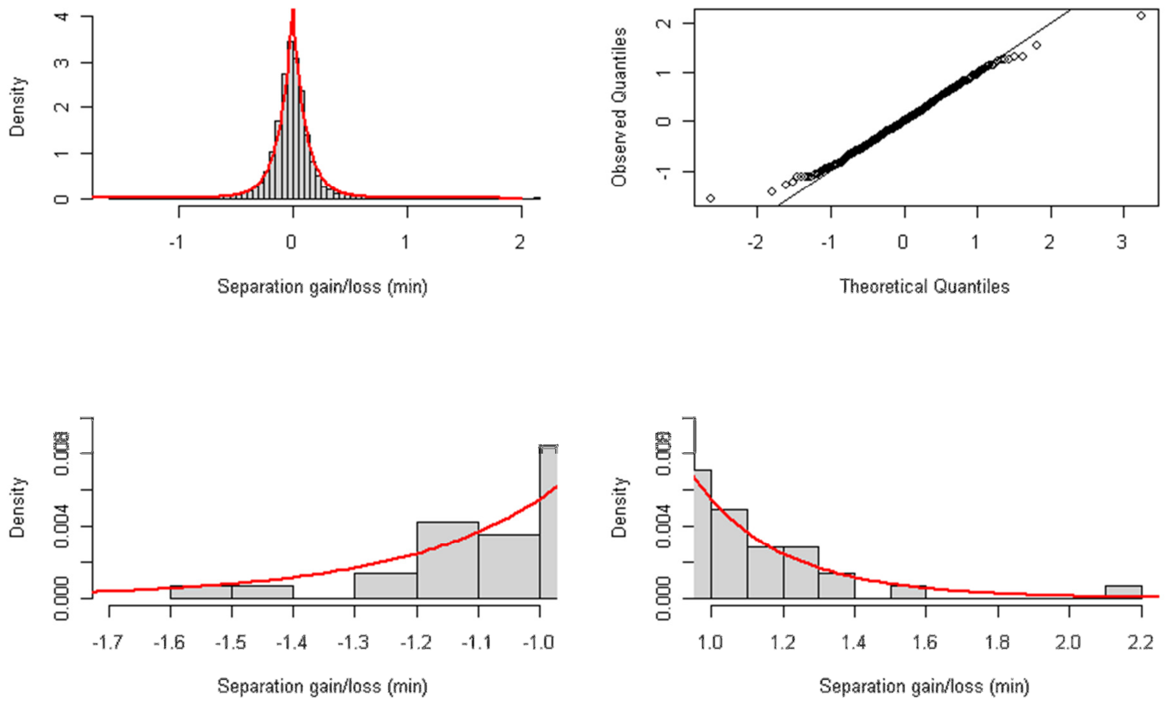

Figure 4-7: Plots showing the goodness of fit of the optimised mixture distribution with 89% of observations from Laplace(0, 0.11) and 11% of observations from Laplace(0, 0.27)

This mixture distribution appears to fit well both the body and tails of the empirical distribution, although from the quantile-quantile plot it can be seen that there is a conservative fit to the extreme tails. It should be noted that there are other heavy-tailed distributions or combinations of distributions that might also fit the data well. It should also be noted that the mixture component with the larger variance is heavily dependent on the few observations in the tails, and is therefore subject to substantial uncertainty. However, the Laplace distribution has an analytically tractable functional form and there is little indication that it is inappropriate in this case, and so for convenience the Laplace mixture distribution in Figure 4-7 will be used for the baseline analysis.

4.2.1.3 Derivation of Q(s)

The fitted mixture distribution describes the probability of a pair of aircraft gaining or losing exactly x minutes of separation during the reporting period. If the parameters of the two Laplace distributions are b_1 and b_2, with proportion p of observations drawn from the first distribution, then the mixture can be written:

$$f(x|p, b_1, b_2) = \frac{p}{2b_1}\exp\left(-\frac{|x|}{b_1}\right) + \frac{(1-p)}{2b_2}\exp\left(-\frac{|x|}{b_2}\right)$$

The probability of losing s or greater minutes of separation can then be derived as:

$$Q(s) = \int_s^\infty f(x|p, b_1, b_2)dx = \frac{p}{2}\exp\left(-\frac{s}{b_1}\right) + \frac{(1-p)}{2}\exp\left(-\frac{s}{b_2}\right)$$

4.2.2 <u>Variations for Sensitivity Analysis</u>

Figure 4-8 shows a mixture distribution with a more extreme distribution component for the tails. The parameters were selected using the upper 97.5% percentiles of the posterior probability distribution derived using MCMC.

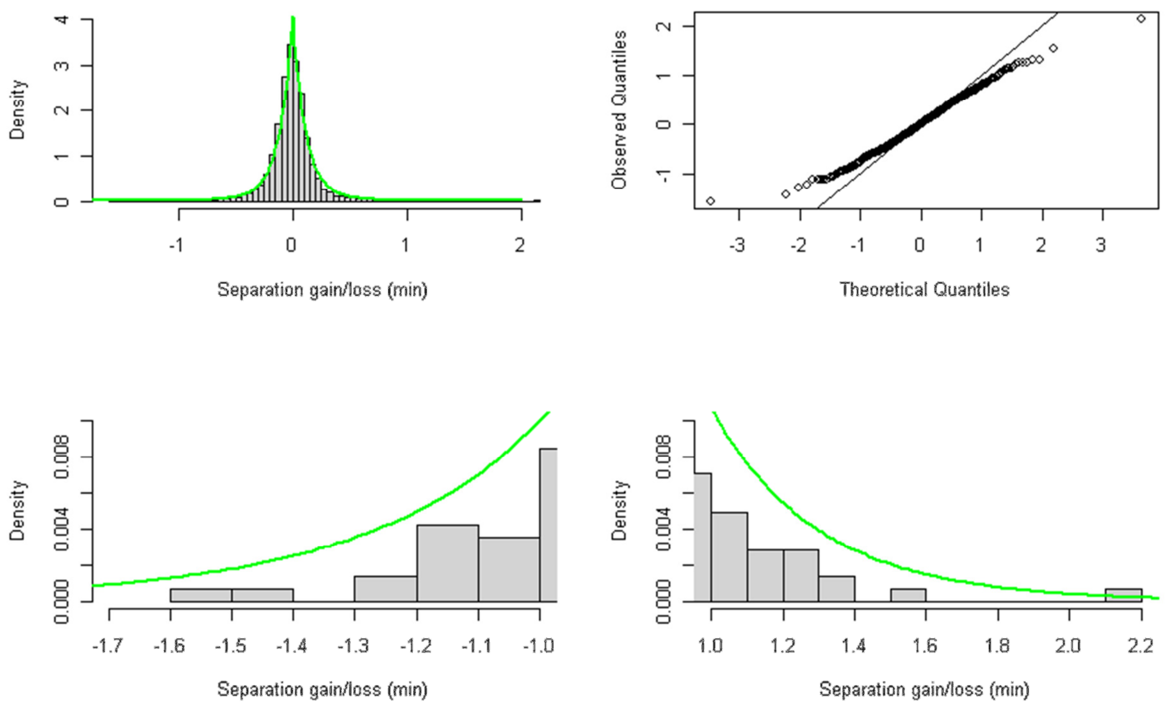

Figure 4-8: Plots showing the goodness of fit of the optimised mixture distribution with 84% of observations from Laplace(0, 0.11) and 16% of observations from Laplace(0, 0.31)

The single Laplace distribution shown in Figure 4-9 will also be used in the sensitivity analysis to show the effect of modelling the data using a simpler form of the distribution. The distribution has been manually fitted to ensure adequate tail coverage.

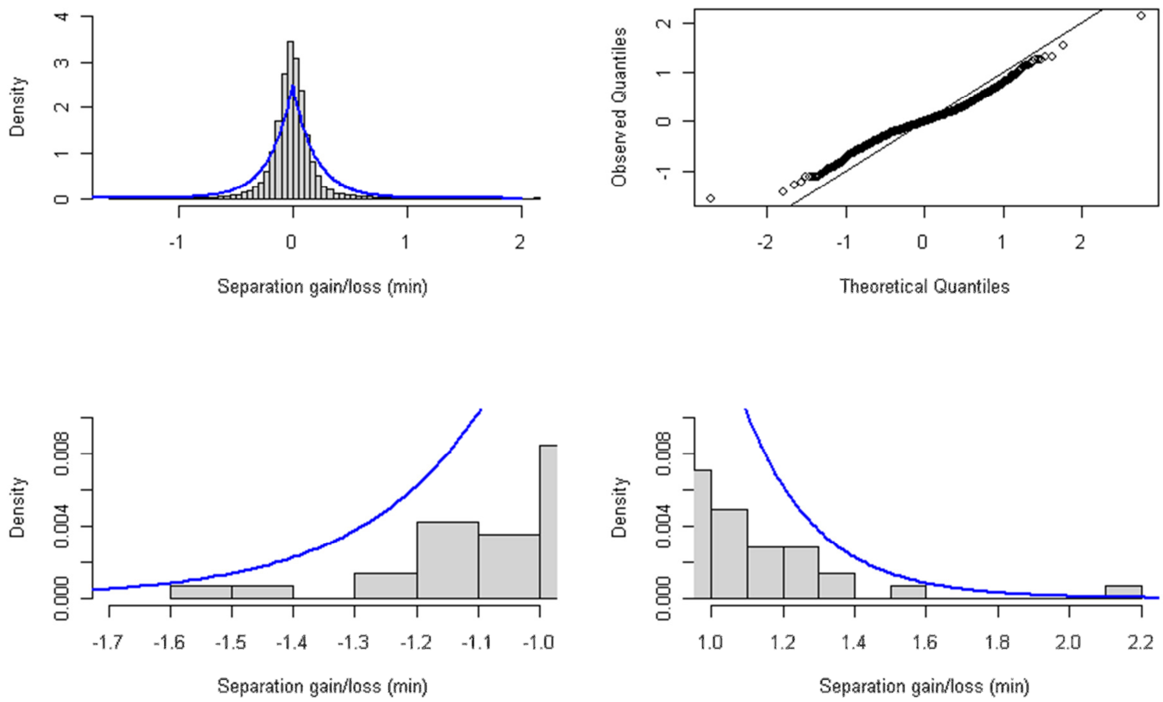

Figure 4-9: Plots showing the goodness of fit of a single Laplace(0, 0.2) distribution

4.3 Probability of Overlap

The two distributions $E(s)$ and $Q(s)$ are integrated within the longitudinal collision risk model to calculate the probability that two aircraft will be in longitudinal overlap by the end of the at-risk period. Given a Gamma(k, θ) form for $E(s)$ and a p ×Laplace(0, b_1) + $(1-p)$ ×Laplace(0, b_2) form for $Q(s)$, this integral can easily be calculated as:

$$\Pr(overlap) = \int_5^\infty E(s)Q(s)ds$$

$$= \int_0^\infty \frac{x^{k-1}\exp(-\frac{x}{\theta})}{\Gamma(\mathrm{k})\theta^k} \times \left(\frac{p}{2}\exp\left(-\frac{5+x}{b_1}\right) + \frac{(1-p)}{2}\exp\left(-\frac{5+x}{b_2}\right)\right) dx$$

$$= \frac{p}{2\left(1+\frac{\theta}{b_1}\right)^k}\exp\left(-\frac{5}{b_1}\right) + \frac{(1-p)}{2\left(1+\frac{\theta}{b_2}\right)^k}\exp\left(-\frac{5}{b_2}\right)$$

4.4 Conflict Resolution Delay

4.4.1 Baseline

In previous work [Ref 2] a value of 9 minutes was used for the average Conflict Resolution Delay (CRD), denoted τ. In discussion of the analysis at SASP-WG/WHL/20 this value was strongly criticised for being too conservative and a more appropriate value of 4 minutes was recommended, based on work presented in [Ref 3]. This recommendation was accepted and a CRD value of 4 minutes is used for the baseline analysis.

4.4.2 Variations for Sensitivity Analysis

The previously used value of 9 minutes for the CRD will be retained for the purpose of a sensitivity analysis.

4.5 At-Risk Period

The total at-risk period T (hours) comprises the reporting period plus the CRD. Therefore:

$$T = \frac{18 + \tau}{60}$$

4.6 Longitudinal Closing Speed

The average longitudinal closing speed (knots) given minimum longitudinal separation of 5 minutes can conservatively be derived as:

$$\overline{|\dot{x}(5)|} = \frac{5d}{60T}$$

where d is the average aircraft speed in knots and T is the total at-risk period in hours. A value of 480 knots will be used for the average aircraft speed in this analysis.

4.7 $P_y(0), P_z(0)$

The probability that two aircraft which are on the same track are in lateral overlap, $P_y(0)$, has a value of 0.1172 currently adopted for use in the NAT. Similarly, the probability that two aircraft which are nominally at the same level are in vertical overlap, $P_z(0)$, has an adopted value of 0.48. These values are based on the average behaviour of all aircraft operational in the NAT [Ref 4].

Since this analysis is limited to aircraft which are RLongSM-capable, and thus FANS-equipped, it can be assumed that their average navigation performance is better than that implied by these adopted values. An improved estimate specifically for FANS-equipped aircraft is not available, however in discussion at SASP/WG-WHL/20 it was suggested that a value of 0.6 could be more appropriate for these parameters. This value will therefore be adopted for the baseline analysis. To be prudent a value of 1.0 will additionally be used for both probabilities for the sensitivity analysis.

4.8 Additional Parameters

The parameters for aircraft dimensions and average absolute relative cross track and vertical speed will take the values currently adopted for use in the NAT [see Table 4-1] [Ref 4].

4.9 Parameters Summary

Parameter	Description	Baseline Value	Sensitivity Analysis Values
CRD	Conflict Resolution Delay (note: not a direct model parameter, but affects other components)	4 mins	9 mins
λ_x	Average aircraft length	0.03108Nm	-
λ_y	Average aircraft wing-span	0.02846Nm	-
λ_z	Average aircraft height (includes, incorrectly, the undercarriage)	0.00892Nm	-
$\overline{\lvert\dot{y}_0\rvert}$	Average absolute relative cross track speed for aircraft nominally on the same track	5kts	-
$\overline{\lvert\dot{z}_0\rvert}$	Average absolute relative vertical speed for an aircraft pair that have lost all vertical separation	1.5kts	-
$\overline{\lvert\dot{x}(5)\rvert}$	Average absolute longitudinal closing speed between an aircraft pair with 5 minutes intended separation, given that a longitudinal overlap event occurs during the at-risk period	109.1kts (4 mins CRD)	88.9kts (9 mins CRD)
$P_y(0)$	Probability that two aircraft which are on the same track are in lateral overlap	0.6	1.0
$P_z(0)$	Probability that two aircraft which are nominally at the same level are in vertical overlap	0.6	1.0
T	Average at-risk period	0.37 hours (4 mins CRD)	0.45 hours (9 mins CRD)
$E(s)$	Distribution of planned separations at the end of the at-risk period	Gamma(6, 1.5)	[1] Gamma(5, 1.4), [2] Gamma(2, 1)
$Q(s)$	Probability distribution of a loss of planned separation of s or greater minutes by the end of the at-risk period	p=0.89, b_1=0.13 (b_1=0.11 scaled for 4 mins CRD), b_2=0.33 (b_2=0.27 scaled for 4 mins CRD)	[1] p=0.84, b_1=0.13 (b_1=0.11 scaled for 4 mins CRD), b_2=0.38 (b_2=0.31 scaled for 4 mins CRD), [2] p=1, b_1=0.24 (b_1=0.2 scaled for 4 mins CRD), b_2=0

Table 4-1**: All parameters values adopted within this analysis**

5 Collision Risk Estimates

5.1 Baseline Calculation

Using the baseline parameters described in Table 4-1 within the longitudinal collision risk model, the collision risk expressed as number of fatal accidents per flight hour is estimated to be:

$$N_{ax} = 1.07 \times 10^{-12}$$

This represents our best estimate for the current longitudinal collision risk in the Shanwick region of the NAT due to erosion of longitudinal separation between RLongSM-capable aircraft, reporting periodically with an interval of 18 minutes and a minimum longitudinal separation of 5 minutes.

The currently adopted Target Level of Safety (TLS) in the NAT for the longitudinal dimension is 5×10^{-9} fatal accidents per flight hour. Our collision risk estimate is substantially below this, with less than 1/1000th of the maximum risk it is required to be demonstrated.

5.2 Sensitivity Analysis Results

Each of the variations for sensitivity analysis described in Table 4-1 were assessed in turn. The collision risk estimates in each case are presented in Table 5-1.

Sensitivity Analysis Variation	Collision Risk Estimate
CRD = 9 mins	3.96×10^{-11}
$P_y(0)$ = 1.0	1.79×10^{-12}
$P_z(0)$ = 1.0	1.79×10^{-12}
$E(s)$ ~ Gamma(5, 1.4) [Plausible future increase in RLongSM usage]	7.87×10^{-12}
$E(s)$ ~ Gamma(2, 1) [Extreme future increase in RLongSM usage]	1.92×10^{-9}
$Q(s)$: p=0.84, b_1=0.13 (b_1=0.11 scaled for 4 mins CRD), b_2=0.38 (b_2=0.31 scaled for 4 mins CRD) [Mixture distribution with extreme tails]	2.15×10^{-11}
$Q(s)$: p=1, b_1=0.24 (b_1=0.2 scaled for 4 mins CRD), b_2=0 [Single Laplace distribution with extreme tails]	1.07×10^{-14}

Table 5-1: Sensitivity analysis results

It can be seen that the choices of CRD, $E(s)$ and $Q(s)$ can substantially alter the resulting collision risk estimates. In particular, a choice of $E(s)$ where the majority of separations are assumed to be close to the minimum of 5 minutes gives a collision risk estimate close to the longitudinal TLS. This was assumed to be an extreme distribution, unlikely to be realised within the NAT so this may not be cause for concern.

The choice of functional form for $Q(s)$ (i.e. single distribution versus mixture distribution) is seen to make a large difference. This strongly indicates how important the fitting of the theoretical distribution to the tails of the empirical distribution is.

6 Conclusions of NAT RLongSM Longitudinal Collision Risk Analysis

The analysis has been conducted on an entire year's data from pairs of RLongSM-capable aircraft transiting the Shanwick region of the NAT during the RLongSM operational trial. All model parameters have been selected to represent our current best knowledge of aircraft behaviour in the NAT. The result of this analysis is a longitudinal collision risk estimate of 1.07×10^{-12} fatal accidents per flight hour, which is substantially below the longitudinal TLS of 5×10^{-9} fatal accidents per flight hour.

The sensitivity analysis does show that variations in key model parameters can have a large effect on the collision risk estimates. However, for each variation tested, the estimated collision risk was still below the TLS.

On the basis of this analysis we conclude that the application of RLongSM in the Shanwick region of the NAT is acceptably safe.

7 Recommendations

The meeting is invited to note and review the results presented which indicate that the estimated longitudinal collision risk under RLongSM operations in the NAT (based on 1 year of data from the Shanwick region) is within the longitudinal TLS of 5 x 10^{-9} fatal accidents per flight hour. This indicates that the RLongSM procedure in the NAT region is acceptably safe.

The meeting is invited to determine whether this analysis meets the requirements for a collision risk assessment to support the transition of a regional operational trial into full operational use. If this analysis meets the requirements and can thus be considered a proof of concept of the RLongSM procedure, the meeting is invited to endorse RLongSM as a globally applicable procedure via an amendment to PANS-ATM (ICAO Doc 4444).

8 References

1. SASP-WG/WHL/20-WP/29, *Estimating the Longitudinal Separation Loss Distribution under Periodic Reporting using Waypoint Reporting Data*
2. SASP-WG/WHL/20-WP/25, *Collision Risk Estimates under Reduced Longitudinal Separation Minimum Operations*
3. SASP-WG/WHL/4-WP09, *A general collision risk model based on reliability theory*
4. North Atlantic MNPS Quick Reference Guide, (Last Revised 29 September 2011), Edited by the Rapporteur of the Mathematicians' Working Group