Safety Regulation Group

Civil Aviation Authority

CAA PAPER 2003/2

DGPS Guidance for Helicopter Approaches to Offshore Platforms - Follow On Studies

www.caa.co.uk

Safety Regulation Group

CAA PAPER 2003/2

DGPS Guidance for Helicopter Approaches to Offshore Platforms - Follow On Studies

Important Note

The CAA has made many of the documents that it publishes available electronically (in addition to traditional printed format). Where practical, the opportunity has been taken to incorporate a clearer revised appearance to the documents. Any significant changes to the content of this document will be shown in the Explanatory Note. If no such changes are indicated the material contained in this document, although different in appearance to the previously printed version, is unchanged. Further information about these changes and the latest version of documents can be found at www.caa.co.uk.

9 June 2003

ISBN 0 86039 923 0

Enquiries regarding the content of this publication should be addressed to:
Research Management Department, Safety Regulation Group, Civil Aviation Authority, Aviation House, Gatwick Airport South, West Sussex, RH6 0YR.

The latest version of this document is available in electronic format at www.caa.co.uk, where you may also register for e-mail notification of amendments.

Printed copies and amendment services are available from: Documedia Solutions Ltd., 37 Windsor Street, Cheltenham, Glos., GL52 2DG.

List of Effective Pages

Study	Annex	Page	Date	Study	Annex	Page	Date
Study 3		8	9 June 2003				
Study 3		9	9 June 2003				
Study 3		10	9 June 2003				
Study 3		11	9 June 2003				
Study 3		12	9 June 2003				
Study 3		13	9 June 2003				
Study 3		14	9 June 2003				
Study 3		15	9 June 2003				
Study 3		16	9 June 2003				
Study 3		17	9 June 2003				
Study 3		18	9 June 2003				
Study 3		19	9 June 2003				
Study 3		20	9 June 2003				
Study 3		21	9 June 2003				
Study 3		22	9 June 2003				
Study 3		23	9 June 2003				
Study 3		24	9 June 2003				
Study 3		25	9 June 2003				
Study 3		26	9 June 2003				
Study 3		27	9 June 2003				
Study 3		28	9 June 2003				
Study 3		29	9 June 2003				
Study 3		30	9 June 2003				
Study 3		31	9 June 2003				
Study 3		32	9 June 2003				
Study 3		33	9 June 2003				
Study 3		34	9 June 2003				
Study 3		35	9 June 2003				
Study 3		36	9 June 2003				
Study 3		37	9 June 2003				
Study 3		38	9 June 2003				
Study 3		39	9 June 2003				
Study 3		40	9 June 2003				
Study 3		41	9 June 2003				
Study 3		42	9 June 2003				
Study 3		43	9 June 2003				
Study 3		44	9 June 2003				
Study 3		45	9 June 2003				
Study 3		46	9 June 2003				
Study 3		47	9 June 2003				
Study 3		48	9 June 2003				
Study 3	Annex 1	1	9 June 2003				
Study 3	Annex 1	2	9 June 2003				
Study 3	Annex 1	3	9 June 2003				
Study 3	Annex 1	4	9 June 2003				
Study 3	Annex 2	1	9 June 2003				
Study 3	Annex 3	1	9 June 2003				
Study 3	Annex 3	2	9 June 2003				

Contents

Differential GPS Guidance for Helicopter Approaches to Offshore Platforms - Follow-On Studies

General Foreword

The research reported in this paper was funded by the Safety Regulation Group of the UK Civil Aviation Authority, and was performed by Cranfield Aerospace Ltd (CAe) and Lambourne Navigation Ltd (LNL). The work comprises three follow-on studies to the flight trials of DGPS guidance for helicopter approaches to offshore platforms reported in CAA Paper 2000/5 which, in turn, was instigated in response to the findings of the Helicopter Human Factors Working Group reported in CAA Paper 87007 (Recommendations 4.1.1 and 4.2.1). The Helicopter Human Factors Working Group was formed in response to Recommendation 1 of the Report of the Helicopter Airworthiness Review Panel (HARP Report - CAP 491).

- The first study, undertaken by CAe, comprises an analysis of a number of the largest horizontal errors observed in the data from the earlier flight trials. The initial analysis of this data reported in CAA Paper 2000/5 indicated that, although the DGPS errors were generally relatively low, there were a significant number of occurrences of much larger errors. It was considered appropriate to attempt to understand the causes of these errors in order that they could be either prevented, or mitigated in an operational system.
- The second study, also performed by CAe, investigates the difficulties encountered with the MF correction system during the earlier flight trials. It was considered that a non-local area correction system could represent an attractive solution to providing differential corrections, particularly an existing marine MF service such as that employed for the trials. The problems encountered during the trials, however, would first need to be understood and addressed.
- The third study, performed by LNL, investigates the effects of single satellite unavailabilities on the results obtained from the earlier flight trials. The study focused on GPS performance in terms of availability, precision and receiver tracking. Other studies have indicated that satellite reliability is such that unavailabilities must be considered, but it would have been impractical to explore their effects during the trials.

Overall, the CAA concurs with the results of the studies which will all be taken into account in regulating the use of GPS in offshore helicopter operations. Although not investigated in any of the three studies reported in this paper, it is expected that the costs and logistics of installing and maintaining differential correction stations on offshore platforms will render local area systems unattractive to Industry; the results of the first study serve to cast further doubt on their viability for offshore applications. The marine MF-based wider area system employed for the flight trials, however, was not without problems as illustrated in the second study. Nevertheless, significant improvements have been made to this service since the flight trials, and a satellite-based correction system is also likely to become available in the near term. Perhaps the most significant finding of all the work reported in this paper is the poor satellite tracking performance observed during the flight trials, and identified in the third study. As a consequence, further experimental work addressing the effects of helicopter rotors on GPS reception has been commissioned by CAA and has recently been completed. This work will be published in a separate CAA paper.

Study 1

Analysis of Large DGPS Errors

Contents

Study 1 Analysis of Large DGPS Errors

Study 1 Analysis of Large DGPS Errors

1 Introduction

During 1996 a series of flight trials was undertaken in the North Sea to examine the use of Differential Global Positioning System (DGPS) equipment as an approach aid for offshore installations. The flight trials programme was undertaken by the Flight Systems and Measurement Laboratories (now incorporated into Cranfield Aerospace Ltd) of the College of Aeronautics, Cranfield University in the role of prime contractor on behalf of the UK Civil Aviation Authority. The trials were performed using a Sikorsky S76C helicopter operated by Bond Helicopters Ltd. The aircraft was fitted with an experimental DGPS installation which was complemented by additional recording equipment sited at fixed locations.

In the course of seven test flights totalling 36 hours, over 70 predefined manoeuvres were performed at a set of four offshore production platforms with differing topside layouts. At each platform, approach trajectories and guidance presentations based upon the use of DGPS data were evaluated by the trials team which comprised representatives from CAA, Bond and Cranfield.

The trials installation allowed a comparison to be made between alternative sources of differential corrections: Medium Frequency “MF” corrections received from onshore marine radiobeacons; and Ultra High Frequency “UHF” corrections received from a dedicated DGPS reference station positioned on the target platform. It was also possible to perform a comparison between GPS receivers produced by two different manufacturers (Navstar and Trimble).

Post-flight processing of the data recorded during each trial enabled an assessment to be made of the performance of the real-time airborne DGPS equipment. Full details of these results are contained in the project Final Report, which has been published as CAA Paper 2000/5 (ref 1).

Volume 2 of ref 1 includes a statistical analysis of the horizontal (two-dimensional) error characteristics of each receiver, covering the periods where operation in differentially corrected mode was achieved. This analysis demonstrates that, although these errors were relatively low for the majority of the data samples, there were also a significant number of occurrences of considerably larger errors.

In the case of the receiver supplied with differential corrections from the UHF source, the 95% confidence limit for the horizontal accuracy was 16.7m. However, the maximum error observed was 123.1m (i.e. nearly an order of magnitude greater) and the 99% confidence limit was 48.5m. For the identical GPS receiver which was supplied with the MF corrections, the corresponding figures were as follows: 95% confidence limit 6.5m; maximum error 17.7m; and 99% limit 11.4m.

This report presents the results of an additional, more detailed, analysis of a number of time periods during which some of the largest horizontal errors were observed to occur. Data relating to the operation and performance of the GPS receivers is presented in graphical form, along with suggestions regarding the likely reasons for the reduced receiver accuracy.

2 References

1 CAA Paper 2000/5, "DGPS Guidance for Helicopter Approaches to Offshore Platforms", CAA, 2000.

2 CA/CSG/7067, "Analysis of MF Correction Datalink Performance", CAe, 2001.

3 CoA-FS-96-426, "Report on Test Flight 4", Cranfield, 1997.

4 CoA-FS-96-427, "Report on Test Flight 5", Cranfield, 1997.

3 Abbreviations

CAA	UK Civil Aviation Authority
CAe	Cranfield Aerospace Ltd
Bond	Bond Helicopters Ltd
Cranfield	Cranfield University, Cranfield Aerospace Ltd
DGPS	Differential Global Positioning System
ft	Foot
GD1	Identifier for MF-corrected Navstar GPS Navigation data
GD2	Identifier for UHF-corrected Navstar GPS Navigation data
GD3	Identifier for MF-corrected Trimble GPS Navigation data
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HDOP	Horizontal Dilution of Precision
Hz	Hertz
kt	Knot
L-band	Region of electromagnetic spectrum around 1.5GHz
L1	GPS Link 1 Fequency (1575.42 MHz)
m	Metre
MAP	Missed Approach Point
MF	Medium Frequency
min	Minute
Navstar	Navstar Systems Ltd
nm	Nautical mile
OTruth	Aircraft truth position history using onshore system as reference
PRC	Pseudorange Correction
PRN	Pseudo-Random Noise (GPS Satellite Identifier)
PTruth	Aircraft truth position history using platform system as reference
ref	Reference
RRC	Range-Rate Correction
s	Second
SA	Selective Availability
std dev	Standard deviation
Trimble	Trimble Navigation Ltd
UHF	Ultra High Frequency
UK	United Kingdom
UTC	Universal Time Co-ordinated
o	Degree
^{o}T	Degrees True
2-D	Two dimensional
3-D	Three dimensional

4 Background

The statistical results, presented in section 6 of volume 2 of ref 1, were derived from a comparison of the position data output by the relevant DGPS receivers with the "OTruth" truth position history. Periods during which the DGPS receivers reverted, for whatever reason, to non-differential operating mode were excluded from the analysis.

The OTruth truth position data was derived using a commercial post-processing software package, which combined GPS data from a reference receiver on the aircraft with similar data recorded at a fixed onshore location. A second truth position computation, which employed reference data from a fixed platform-based receiver in place of the onshore system, was also generated and was termed "PTruth".

The post-processing technique employed the GPS carrier phase observable on the L1 frequency to improve the solution accuracy, and to provide increased immunity to multipath effects. As described in section 4.5 of volume 2 of ref 1, a ground-based trial demonstrated that the system was capable of providing positions whose horizontal accuracy was better than 1.5m. Although no similar confirmation of the system accuracy could be achieved in flight, it was noted that a comparison between the separate post-processed OTruth and PTruth solutions (and specifically, the occurrence of a discrepancy of greater than one metre between them) allowed identification of the periods where the truth data might be considered to be suspect.

The data in Table 1 below is based upon ref 1 (volume 2 table 38) and provides a summary of the statistics of the 2-D (horizontal) DGPS errors encountered during the flight trials. The identifiers "GD1", "GD2" and "GD3" refer to the three DGPS units on the trials aircraft, which differed both by manufacturer (GD1 and GD2 were Navstar XR5-M12 receivers, whereas GD3 was a modified Trimble TNL-2100); and by differential correction source employed (GD1 and GD3 used MF corrections, and GD2 used UHF corrections).

Table 1 Summary of Statistics of GD1, GD2 and GD3 2-D Error (metres)

	GD1 receiver (MF Navstar)	GD2 receiver (UHF Navstar)	GD3 receiver (MF Trimble)	GD1 plus GD3 (Both MF units)
Samples	25831	29444	7853	33684
Std dev	2.0	8.0	2.5	2.2
50%	3.0	5.1	4.0	3.2
95%	6.5	16.7	8.5	7.0
99%	11.4	48.5	14.4	13.6
Maximum	17.7	123.1	21.0	21.0

The data in the Table is based upon those samples where the receivers operated in either 3-D differential mode or 2-D differential mode. The former can be considered the normal operating mode of the receivers, and the latter as a reversionary mode (based upon assumptions regarding the aircraft altitude) which was sometimes employed when insufficient satellites were available to provide a 3-D solution. All samples where the receivers were unable to apply differential corrections were excluded from the analysis.

In order to determine when the largest errors had occurred the samples for each receiver were re-sorted into descending order, beginning with the sample associated with the maximum error. Each sample was identified by its flight number and UTC time. Examination of the entries at the start of the re-ordered list then revealed that the largest error values were not evenly spaced throughout the trials dataset but, instead, were contained within a smaller number of time windows. Each time window typically contained a number of (often consecutive) high error data samples. When the data was plotted in the time domain, it became apparent that there were a number of different error distribution "shapes" associated with the occurrence of the largest errors.

Following consultation with CAA, the representative series of time windows detailed in Table 2 was selected for analysis. The start and end times for each window were chosen, largely for convenience, to fall on integer minute boundaries.

Table 2 Selected Time Windows for Analysis

Time window	Flight number	Location	Date	Time period (UTC)	Worst receiver	Worst error (m)
A	7	Beatrice C	31/10/96	14:38 to 14:40	GD1	17.7
B	6	Buchan A	26/09/96	15:01 to 15:03	GD1	17.1
C	4	Piper B	30/07/96	13:53 to 13:56	GD2	123.1
D	5	Tartan A	01/08/96	13:31 to 13:36	GD2	95.5
E	5	Tartan A	01/08/96	13:21 to 13:24	GD2	66.5

Included within time windows C to E are the 220 largest 2-D error samples from the GD2 receiver. They range from the worst error of 123.1m down to 53.4m. These latter samples are also, by virtue of the fact that the GD2 errors were considerably larger than those for the other GPS receivers, the 220 worst errors observed for all three systems during the trials. This data corresponds to about 0.3% of the total dataset.

Time windows A and B include four out of the five largest GD1 samples (including the worst error of 17.7m, plus three 17.1m samples) observed during the flight trials. Although the maximum errors statistics are lower than those applicable to the GD2 data, it was considered important to gain an insight into the error mechanisms which were at work when the MF correction source was utilised: in particular, to determine whether there were any significant differences compared with the use of the UHF source.

No time periods relating to the worst GD3 errors were selected for analysis. Although the largest error statistics for GD3 were slightly greater than those for the GD1 receiver, it transpired that the data in question all related to a period of largely non-differential operation on Flight 6, where the receiver entered differential mode for only four seconds. The fact that such a short section of data was involved, combined with the fact that the causes of the reversion to non-differential mode have been identified elsewhere (ref 2), suggested that it would be more appropriate to perform an additional analysis upon some of the GD1 data.

5 Analysis of Selected Time Windows

5.1 Time Window A: Flight 7, 14:38 to 14:40 UTC

5.1.1 Manoeuvre details

This time window covers a two minute period commencing shortly before the start of approach number 7 on Flight 7 at the Beatrice C platform, and ending midway through the approach. At the start of the window, the aircraft was in level flight in a left turn at a range of 3.6nm from the platform and an altitude of approximately 800ft. The descent was commenced at time 14:39:28, and at the end of the window the aircraft was passing through 420ft at 1.4nm range from the platform.

5.1.2 Truth data availability

Satisfactory truth data was available throughout this time window. The OTruth and PTruth solutions were in agreement to better than ±1m in each axis.

5.1.3 GD1 position accuracy

Figure 1 shows the horizontal and 3-D position errors for the GD1 (MF-corrected Navstar) receiver. The receiver was operating in 3-D non differential mode prior to time 14:39:00, at which time it returned to 3-D differentially corrected mode. The horizontal error for the first sample following entry into differential mode was 17.1m, which was the largest such error observed with the GD1 receiver during the trials programme. One second later the error reduced to 12.1m, and it then fell to less than 4.5m for the remainder of the time window.

Examination of the data for the period during which the receiver operated without the benefit of differential corrections, reveals that the error was steadily reducing for nearly a minute before a sudden increase, beginning at time 14:38:57 and covering the final three samples prior to returning to differential operation, was observed.

Figure 2 depicts the individual GD1 position error components in each axis (i.e. latitude, longitude, and altitude). Although these three error components are not independent of each other, this presentation allows the effect of the error in positioning terms to be determined more clearly.

The figure reveals that the trend of the error starting at time 14:38:57 was in a southeasterly and downward direction. Once differential operation was regained at 14:39:00, the error components appear to have exhibited a decay characteristic over the following three samples.

5.1.4 GD1 velocity components

Figure 3 shows the velocity components, in each axis, obtained from the GD1 receiver. The velocity samples appear to be continuous, with no evidence of any form of hiatus at the time that the disturbances were observed on the position output.

5.1.5 GD1 along-track and cross-track position errors

Figure 4 depicts the result of resolving the GD1 latitude and longitude error components along, and at right angles to, the aircraft velocity vector to obtain along-track and cross-track position errors. This reveals that the position error following return to differential operation at time 14:39:00 was largely in the cross-track sense, and would have erroneously indicated that the aircraft was to the right (by up to 17m) of the intended approach track.

Examination of the data for Flight 7 Approach 3 in volume 3 of ref 1 reveals that the additional effect of this error on the indications displayed to the pilots was negligible (there is evidence for a very small step change on the localiser deviation plot at just over 5000m from the MAP), owing to the low cross-track sensitivity at this range from the platform.

5.1.6 GD1 estimated accuracy parameter

The Estimated Horizontal Accuracy parameter ("nknteh") from the GD1 receiver is plotted in Figure 5. The colour coding on the plot also indicates the point at which the receiver operating mode changed. The following statement was provided by the receiver manufacturer regarding how this parameter is calculated:

> "This variable is calculated by an algorithm that takes into effect the satellite geometry (hdop), the accuracy of individual satellites (level of SA) and the corrections applied to compensate for the atmospheric effects (tropospheric and ionospheric delays), but does not include any effects from multipath signals."

It may be observed that the value of this parameter more than doubled at time 14:38:59 (i.e. two seconds later than the start of the increase in the non differential position error), from 48m to 104m. It then reduced to less than twenty metres in a series of steps, commencing at time 14:39:02 (by which time the horizontal position had reduced to under five metres). This suggests that, on this occasion, examination of the estimate accuracy parameter from the Navstar receiver would have provided an indication of a reduced precision solution.

5.1.7 GD1 satellite usage

Figure 6 shows the satellites, identified by their PRN numbers, which were employed by the GD1 receiver to generate its position solution. Although the receiver employed five satellites throughout the time window, it may be observed that there was a change in the combination used around time 14:39:00, with PRN 16 being replaced by PRN 26 for a five second period which appears to correlate with the occurrence of the increased position error.

It was considered possible that the reason for this temporary change in the satellite combination might have been due to the receiver having lost lock on PRN 16. Although tracking information from the GD1 receiver was not recorded during the trials, it was possible to obtain the equivalent data from the aircraft truth reference Navstar receiver. Since the latter was identical to, and operated from a common antenna as, the GD1 receiver it was concluded that any satellite tracking problems might be expected to affect both units equally. Examination of the truth receiver tracking data confirmed that PRN 16 was lost from the list of satellites tracked for a five second period, commencing at time 14:38:57. This corresponds directly with the period during which the GD1 receiver switched to PRN 26.

The approximate positions of the visible satellites are shown in polar plot form in Figure 7, as calculated by a commercial software program using almanac data which was downloaded from one of the Navstar receivers following the flight. Comparison with the recorded attitude data reveals that, during the period when PRN 16 was temporarily lost, the satellite was within the azimuth sector 050° to 070° and at an elevation of between 30° and 40° relative to the aircraft. This is not a direction from which the satellite signals were known to have been subject to reception difficulties.

5.1.8 Differential corrections.

The GD1 receiver operated using differential corrections derived from the MF system, and Figure 8 shows the correction data which was received on the aircraft from this source during the time window. No correction information was available during the first 58 seconds of the time window owing to intermittent reception problems affecting the MF equipment, which have been the subject of a separate report (ref 2).

In addition to the basic pseudorange corrections (PRC), each of which represents the error in a specific satellite range measurement as determined by the differential reference station, the transmitted data also included range-rate corrections (RRC). The latter are intended to represent the rate of change of the PRC information, to allow the user receiver to perform an extrapolation between successive correction messages.

Since the MF source only provided an updated correction message every few seconds, the possibility was identified that any problems in the RRC data could result in an rapidly increasing error, proportional to the age of the corrections, in the computed differential position. Figure 8 has therefore been arranged to depict both the individual PRC measurements themselves (in the form of the circular dots), and the associated rate-of-change RRC information (the trend lines emanating from the dots), for each satellite included in the received messages. Examination of the plot reveals no evidence of any problems (e.g. step changes) in the transmitted differential corrections, and also suggests that the range-rate data provided a satisfactory prediction of the forward trend on each measurement.

The first of the correction messages was received on the aircraft at time 14:38:59, i.e. in the second prior to the GD1 receiver entering differential mode. There is no evidence of any significant differences between this correction data and the subsequent messages from the MF station.

5.1.9 **GD2 and GD3 operation.**

Figure 9 depicts the position error components in each axis for the GD2 (UHF-corrected Navstar) receiver. There is no evidence of this receiver having been affected by any form of hiatus around time 14:39:00: the solution remained in 3-D differential mode throughout, and the maximum horizontal position error was 8.6m. Regular correction updates were being received by the unit, normally once per second (the maximum interval was 5s). No changes were observed in the combination of satellites employed by the GD2 receiver: the set remained fixed as PRNs 2, 10, 17, 19 and 27. It is therefore not possible to determine whether this receiver also lost lock on PRN 16.

Figure 10 depicts the position error components for the GD3 (MF-corrected Trimble) receiver. This unit, which was operating from the same correction source as the GD1 Navstar receiver, operated in 3-D non differential mode up until time 14:39:24 when it changed to 3-D differentially corrected mode. Thereafter, the maximum horizontal position error was 2.4m. The reasons why the GD3 unit appeared to require an additional 24 seconds to enter differential mode, compared to the GD1 receiver, are discussed in ref 2.

Although four satellites were in use by the GD3 receiver at all times, the unit was observed to make several changes of satellite selection, particularly during the period when it was computing a non-differential solution. This may explain the various step changes in the position errors associated with the latter.

From time 14:38:58 until the end of the time window, satellite PRN 16 was employed in the GD3 position computation, despite the fact that at least two of the Navstar receivers appeared to lose lock on the satellite during a portion of this period.

PRN 26 was in use by the GD3 receiver for seven seconds starting at time 14:38:58.

5.1.10 **Discussion**

Any change in satellite combination when operating without the benefit of differential corrections will normally cause a step change in a GPS position solution, particularly with Selective Availability (SA) in operation as was the case during these flight trials.

It is almost certainly coincidental that the GD1 receiver apparently lost lock on PRN 16, forcing the receiver to change to a different combination of satellites in which PRN 26 replaced PRN 16, only a second or so before the first valid MF correction message was received. Reception of PRN 16 was then regained a few seconds later, allowing the GD1 receiver to revert to its original satellite combination after only two seconds in differential mode.

There are three potential reasons as to why the first two samples from the GD1 receiver (at times 14:39:00 and 14:39:01) may have exhibited an increased position error.

The first hypothesis is that the increased error was the consequence of a deterioration in Dilution of Precision (DOP), which led the receiver to initiate a change in the set of satellites used. Detailed information on the satellite selection algorithms employed by the receiver is not available: in particular, the criteria which determine whether four or five satellite measurements are to be employed in the position solution.

To explore this hypothesis, the CAA GNSS Receiver Performance Simulator was used to calculate the five-satellite horizontal DOP (HDOP) associated with the two different sets in question (PRNs 10, 16, 17, 19 and 27; and PRNs 10, 17, 19, 26 and 27). In addition, the four-satellite HDOP value was also computed for each of the ten possible subsets of these PRNs.

The results reveal no evidence for any sudden change in HDOP with time. For the set of satellites (and its four-PRN subsets) which was in use prior to and following the transition, the HDOP remained less than 6. This suggests that a forced changed in the satellite set due to a sudden increase in DOP is unlikely to have been the cause.

A second hypothesis is that the increased error resulted from a DOP deterioriation which was a consequence, and not the cause, of the change in satellite usage. Evidence from the truth receiver does suggest that the GD1 receiver may well have lost lock on PRN 16 at time 14:38:57, necessitating a change to a satellite set involving PRN 26. The associated HDOP was indeed higher (1.4 increasing to 3.2 for the five-satellite computation; and to as much as 11.0 for one of the four-satellite subsets), so this hypothesis cannot be entirely dismissed.

The third hypothesis is that, due to some form of filtering process within the Navstar receiver, the effect of the step change on the immediately preceding non-differential data was not removed instantaneously but instead involved a couple of seconds' worth of "recovery time" which resulted in an exponential-type decay of the error.

On this occasion, it so happens that the effect of the increased error on the GD1 solution was negligible due to the fact that the aircraft was operating at a significant range from the platform, where the sensitivity of the cross track deviation information presented to the pilots was low. However this could not always be assumed to be the case: instead, it will be necessary to consider what the effect of such an occurrence might be if it were to happen when the indications were at their most sensitive.

The need to improve the reliability of the differential correction source has already been identified (making the assumption that, in the current SA-off environment, an external augmentation source will still be required). The above analysis suggests that it may be necessary, when considering the system failure cases, to pay particular attention to the receiver behaviour when transitioning between the different operating modes.

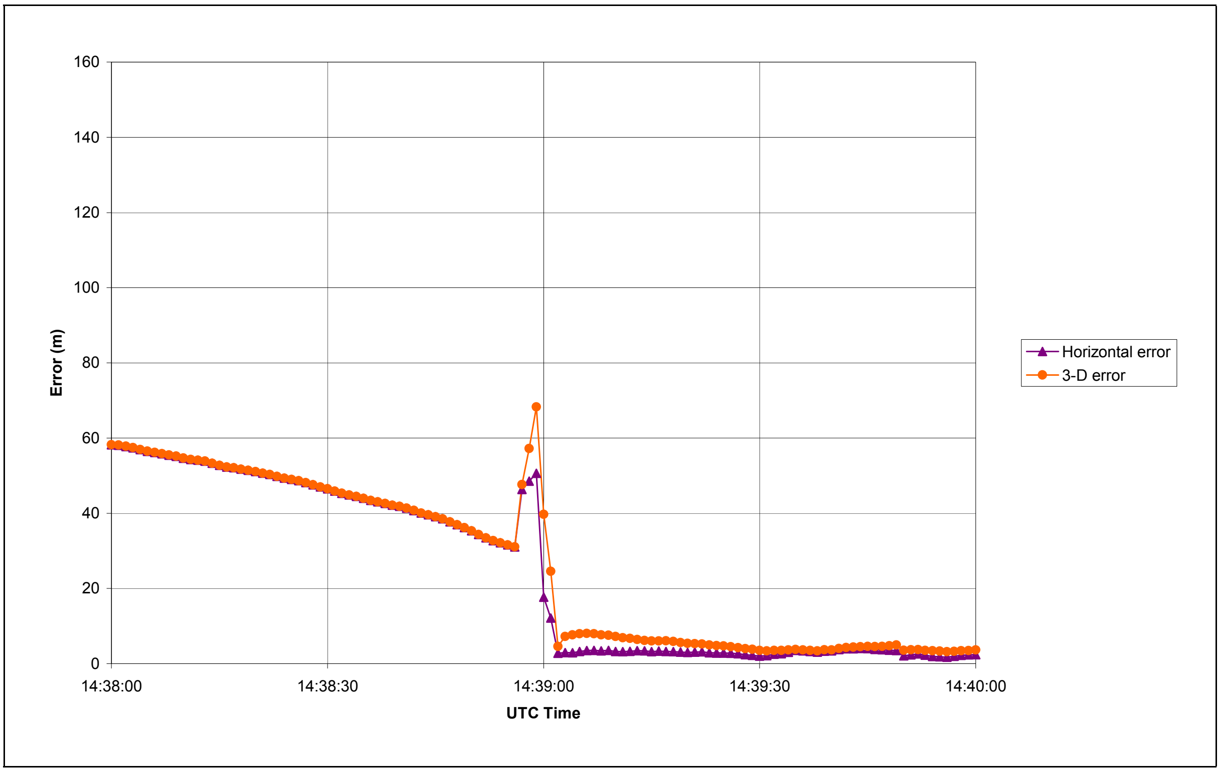

Figure 1 GD1 Position Errors for Window A

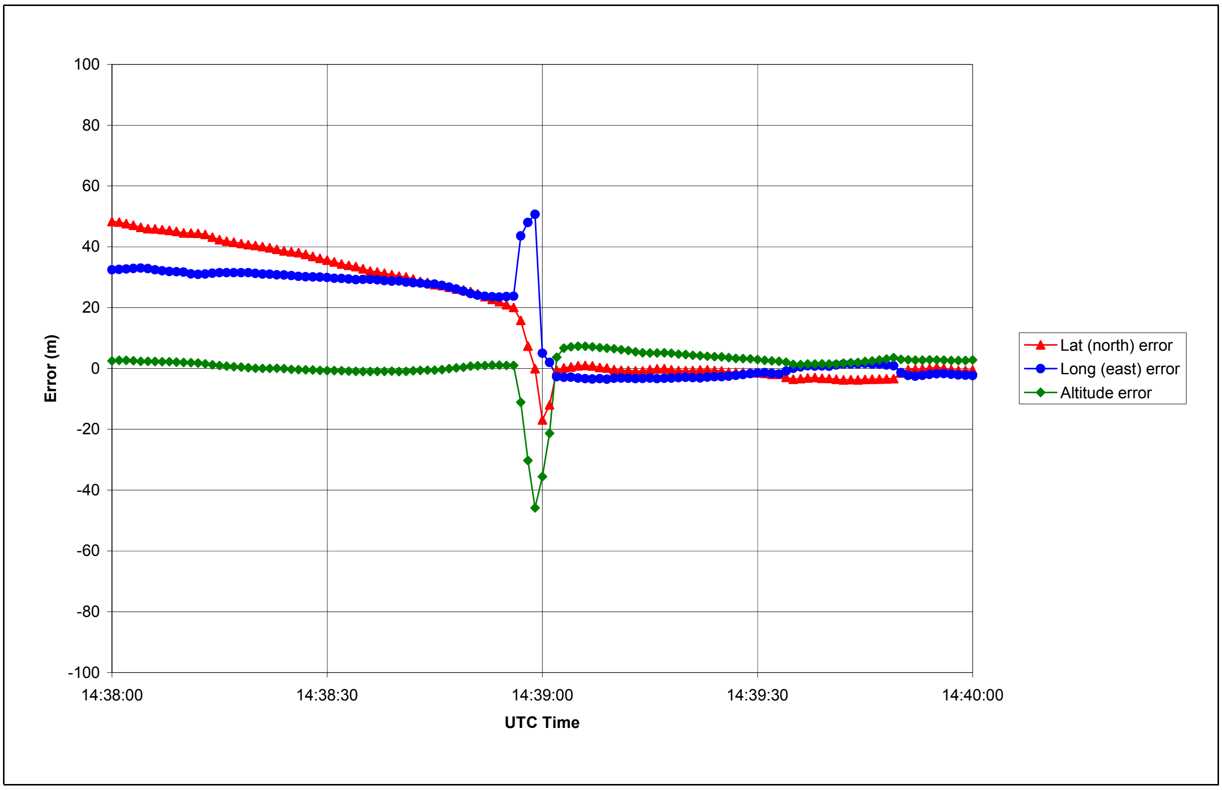

Figure 2 GD1 Position Error Components for Window A

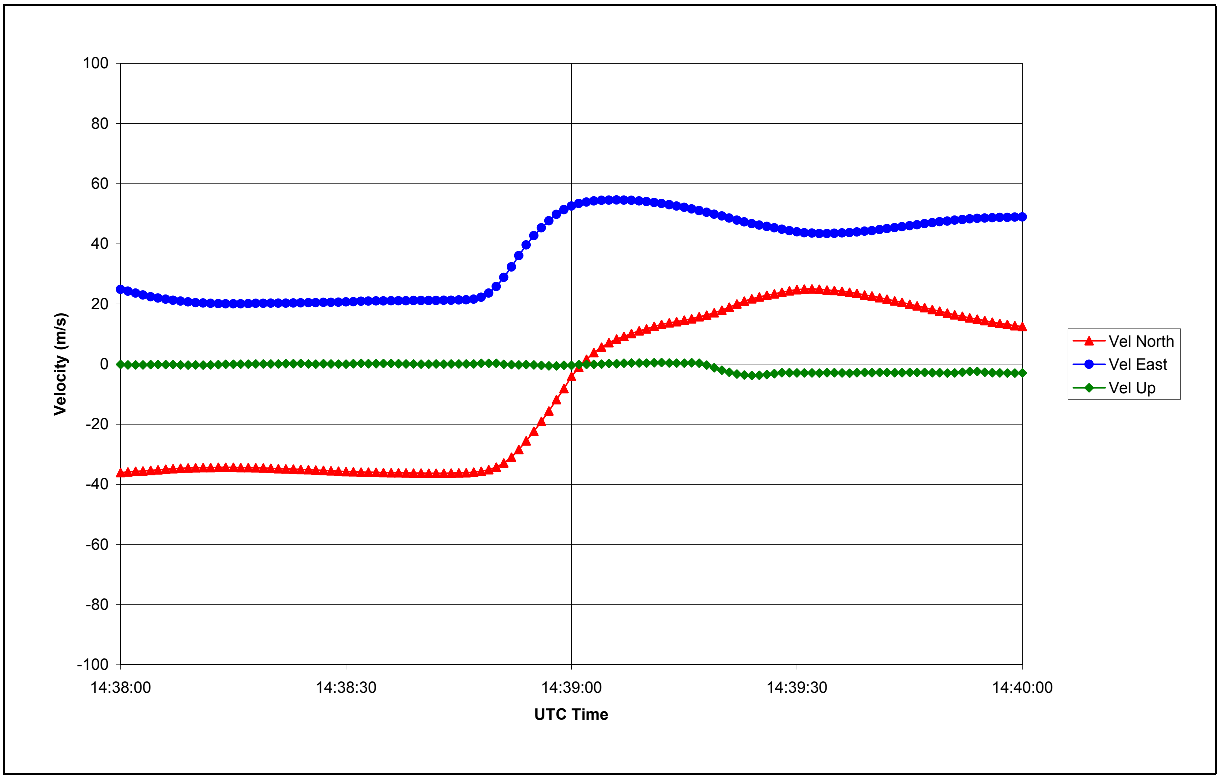

Figure 3 GD1 Velocity Components for Window A

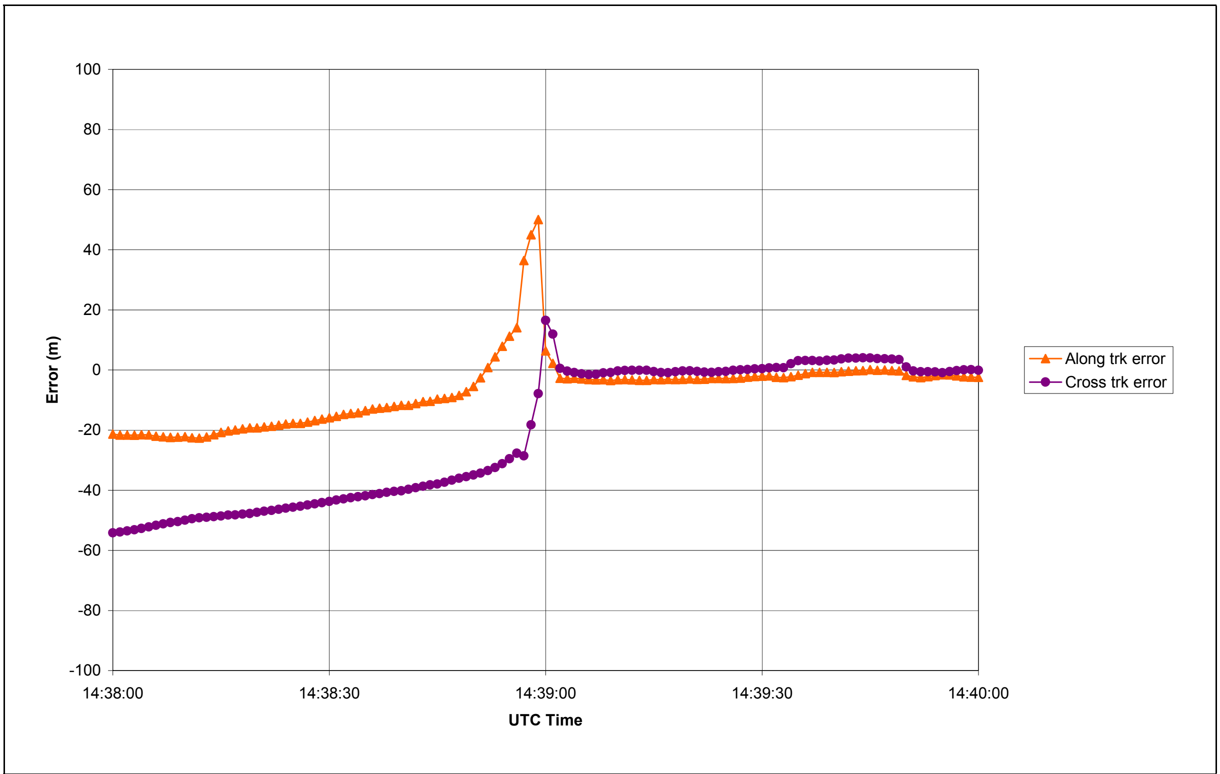

Figure 4 GD1 Along-Track and Cross-Track Position Errors for Window A

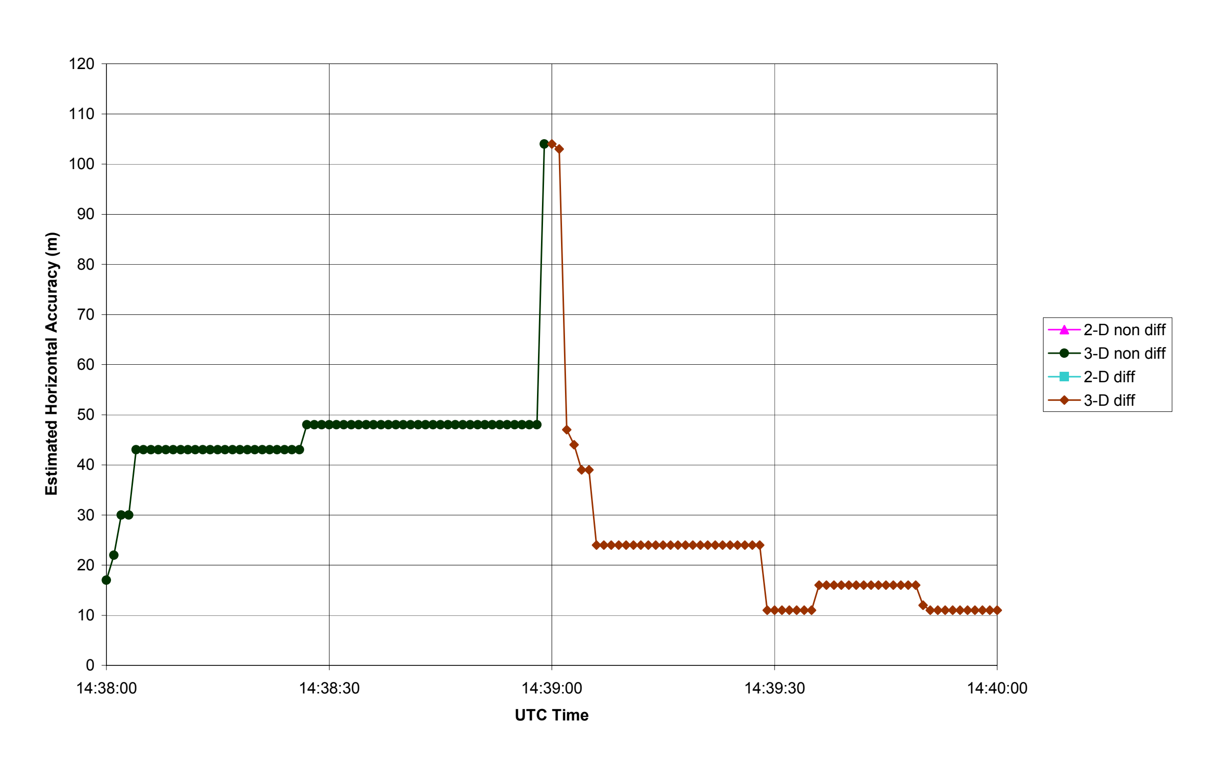

Figure 5 GD1 Estimated Horizontal Accuracy and Operating Mode for Window A

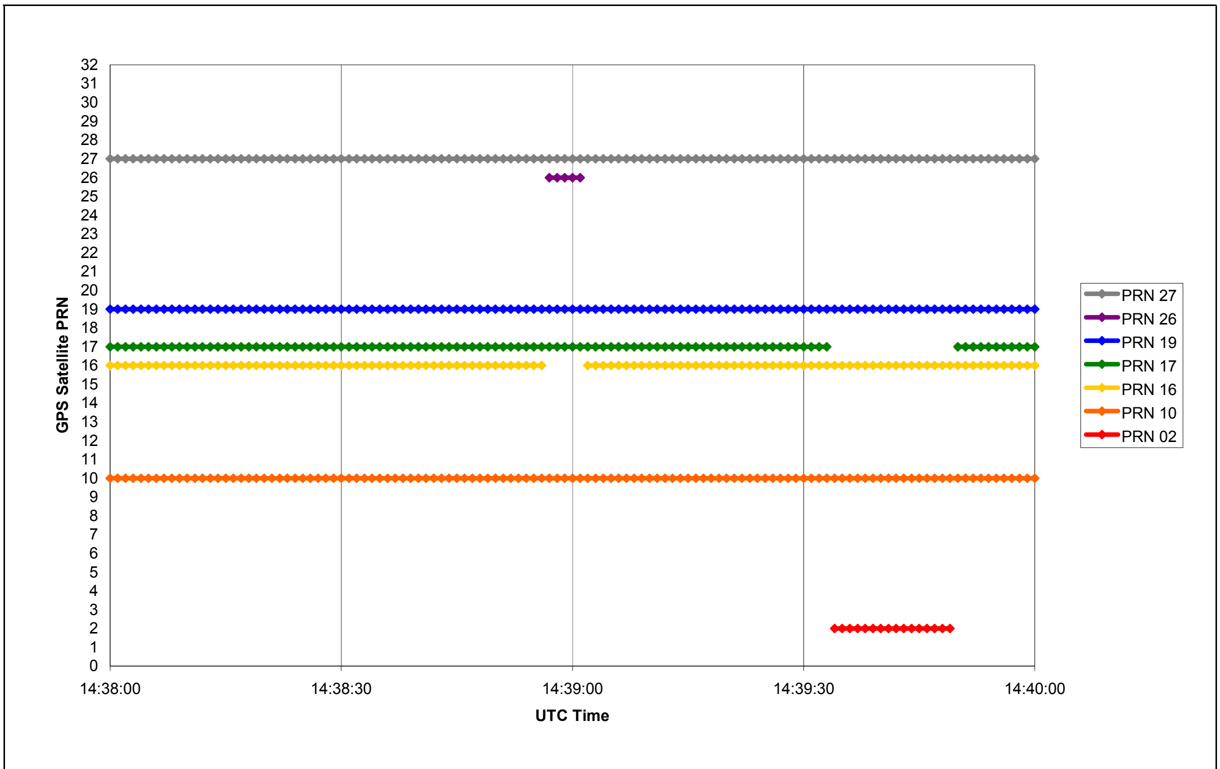

Figure 6 GD1 Satellite Usage for Window A

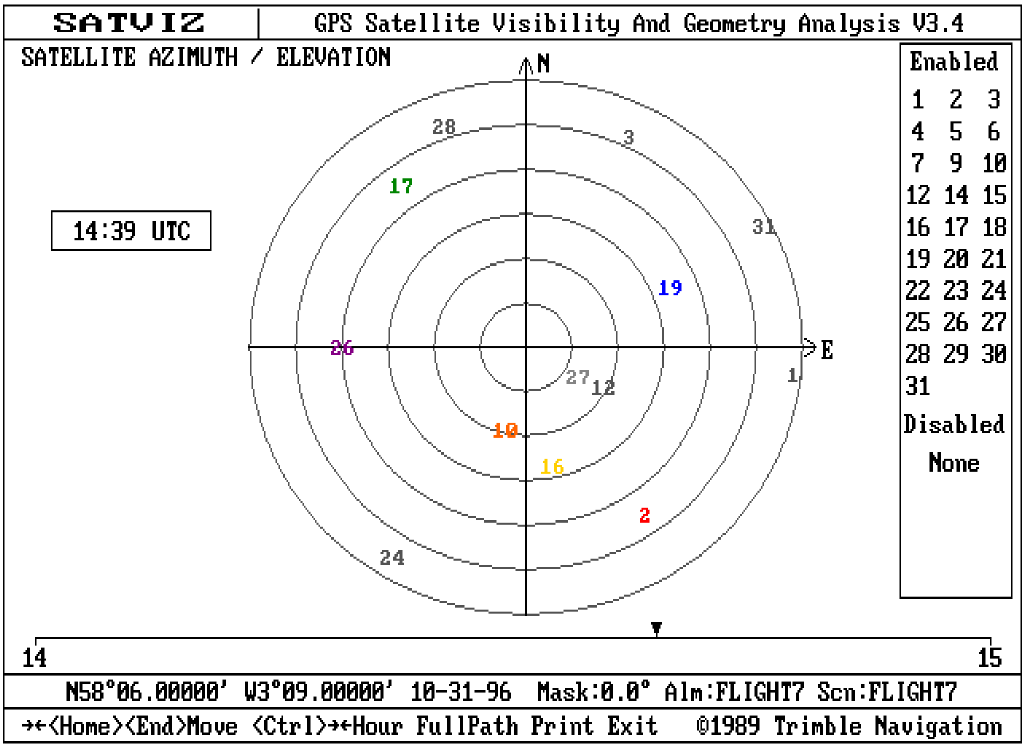

Figure 7 Satellite Azimuth and Elevation for Window A

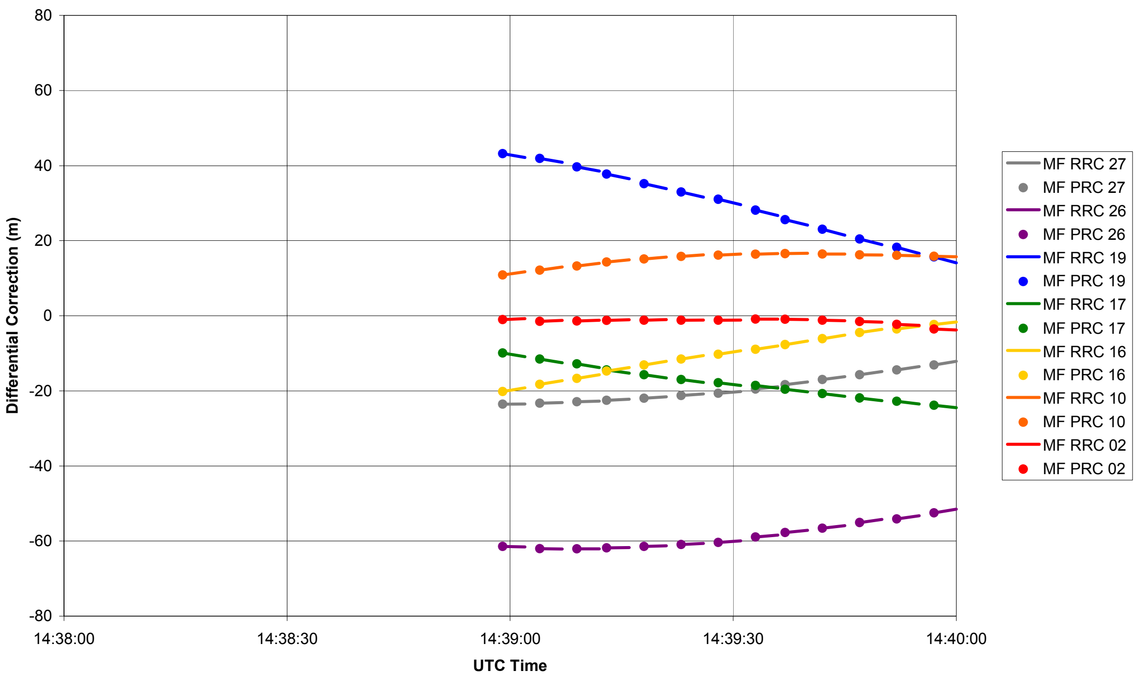

Figure 8 MF Differential Corrections for Window A

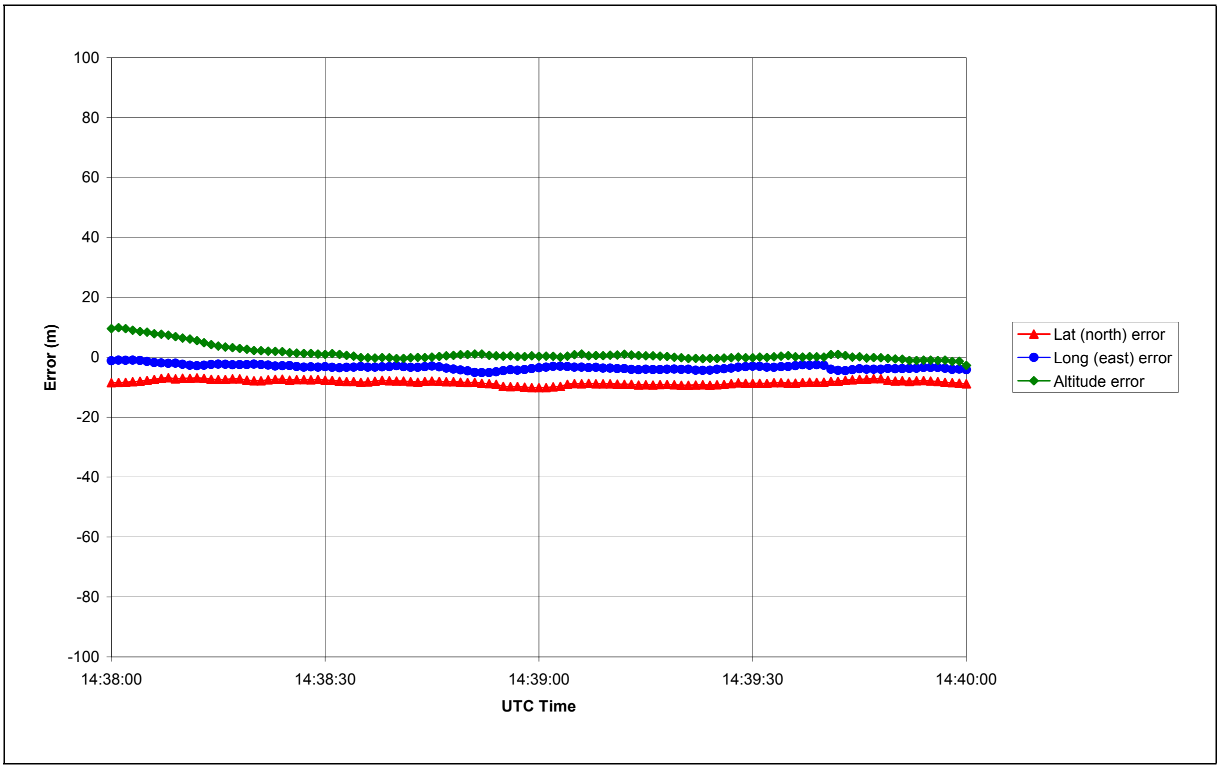

Figure 9 GD2 Position Error Components for Window A

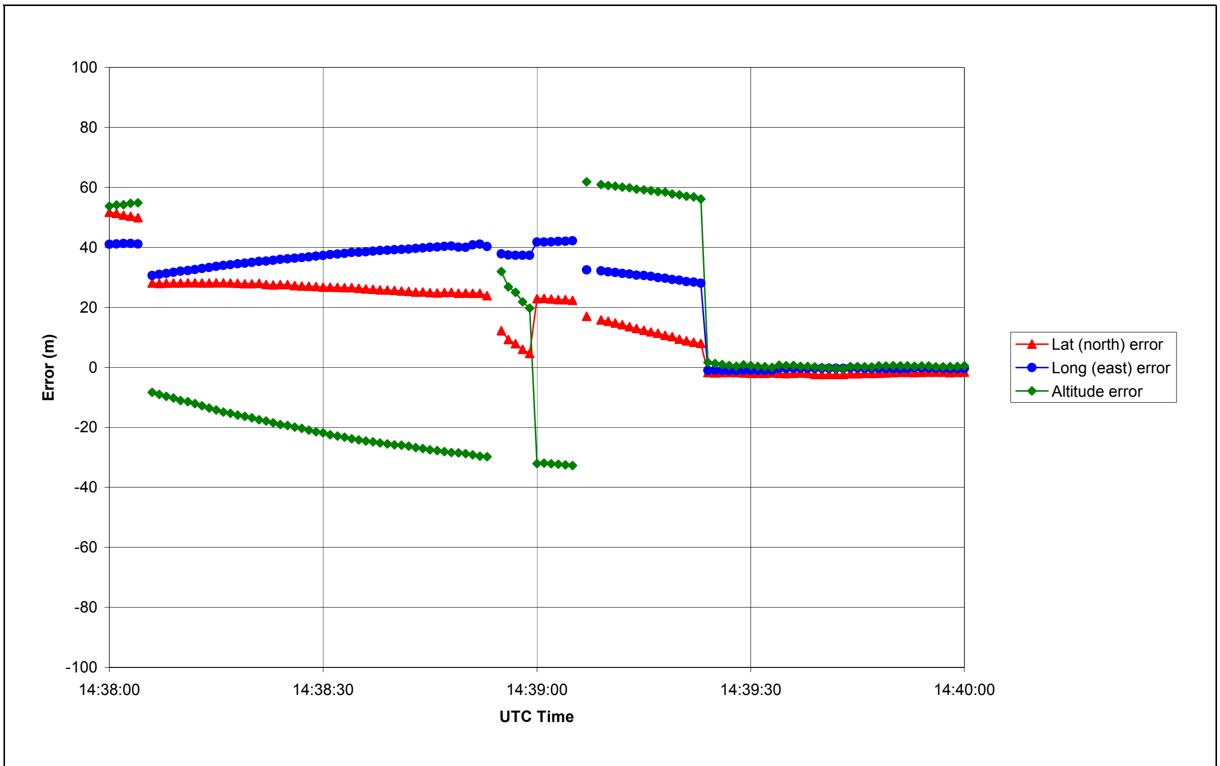

Figure 10 GD3 Position Error Components for Window A

5.2 Time Window B: Flight 6, 15:01 to 15:03 UTC

5.2.1 Manoeuvre details

This time window covers a two minute period ending shortly before the start of approach number 11 on Flight 6 at the Buchan A platform. At time 15:01:21, the aircraft was to the northwest of the platform, on a heading of 110°T and at an altitude of approximately 800ft. It then turned left (with a maximum roll attitude of around 25°), maintaining this altitude, onto a southerly heading to position for the approach. The minimum and maximum ranges from the platform were 3.9nm and 4.6nm, respectively.

Owing to recording problems, no data is available for the first twenty seconds of this time window. Consequently, the first available sample relates to time 15:01:21.

5.2.2 Truth data availability

Truth data only became available from time 15:01:21 onwards and it is possible that its accuracy during this time window was reduced: the maximum difference between the OTruth and PTruth solutions was in excess of 10 metres.

Some of this discrepancy may have been due to the fact that the platform reference receiver was positioned on a semi-submersible (floating) platform, with the result that the system’s position did not remain fixed. Although this problem would only have affected the PTruth data; it is not possible to conclusively determine whether any errors lie in the OTruth solution, in the PTruth solution, or in both. As in ref 1, all of the analysis below has been based upon the use of the OTruth data.

5.2.3 GD1 position accuracy

Figure 11 shows the horizontal and 3-D position errors for the GD1 (MF-corrected Navstar) receiver. The receiver was operating in 3-D differentially corrected mode throughout. There are three samples for which the horizontal error was 17.1m.

Figure 12 depicts the individual GD1 position error components. This shows that the error was in a southwesterly and upward direction, and remained broadly constant for the duration of the time window.

5.2.4 GD1 velocity components

Figure 13 shows the velocity components obtained from the GD1 receiver. The velocity values appear to have been continuous, with no evidence for any disturbances.

5.2.5 GD1 along-track and cross-track position errors

Figure 14 depicts the along-track and cross-track components of the GD1 horizontal error, both of which varied within the range ±15m (approximately) as the aircraft track changed.

5.2.6 GD1 estimated accuracy parameter

The Estimated Horizontal Accuracy parameter from the GD1 receiver is plotted in Figure 15. It may be observed that this parameter was indicating a value of 11m throughout the window, and that the receiver was operating in 3-D differential mode at all times.

5.2.7 GD1 satellite usage

Figure 16 shows the satellites which were employed by the GD1 receiver to generate its position solution.

The plot reveals that the receiver was employing five satellites throughout, although the combination can be observed to have changed on six different occasions during the period when the aircraft was undergoing the greatest heading changes. There is no evidence for any variation in the position error which correlates with these changes in satellite usage.

The satellite positions are shown in the polar plot in Figure 17.

5.2.8 Differential corrections

Figure 18 shows the correction data which was received on the aircraft from the MF source during the time window. Both the pseudorange correction (PRC) and range-rate correction (RRC) information is included.

The plot indicates that the correction data was continuous, with no missing messages, and that the RRC data would have allowed the receiver to apply the appropriate extrapolation between successive updates.

5.2.9 GD2 and GD3 operation

Figures 19 and 20 depict the position error components for the GD2 (UHF-corrected Navstar) and GD3 (MF-corrected Trimble) receivers. Both units operated in 3-D differential mode throughout, using the maximum possible number of satellites (five for the Navstar, four for the Trimble).

The GD2 receiver was receiving regular differential correction updates (maximum interval 4s); and it must be assumed that this statement also applies to the GD3 receiver since it remained in differential mode.

A comparison between Figures 12, 19, and 20 indicates that the magnitude and time variation of the error components from each of the three receivers were very similar.

5.2.10 Discussion

The GD1 data provides no obvious explanation for the occurrence of the larger position errors. In particular, it appears that the receiver was operating under stable conditions, with plenty of satellites visible and with regular differential correction updates.

The fact that the GD1 error components were observed to be very similar (both in magnitude, and in their variation with time) with those for the GD2 and GD3 receivers suggests that some form of common factor may have been at work, to affect the position error computations for all three receivers. The fact that several changes occurred to the receiver satellite selections, combined with the diversity in the differential correction sources, tends to eliminate the possibility that this common factor was associated with the receivers themselves. However it is not possible to make such a definitive statement regarding the antenna system or the L-band reception environment.

The alternative hypothesis is that it was the truth solution itself which was providing a substantial contribution to the individual receiver “errors”, particularly in view of the fact that a significant discrepancy existed between the OTruth and PTruth solutions. However, in the absence of any other independent check upon the accuracy of the truth solution while the aircraft was airborne, it is not possible to directly quantify what the contribution of the OTruth data to the GD1 “error” may have been.

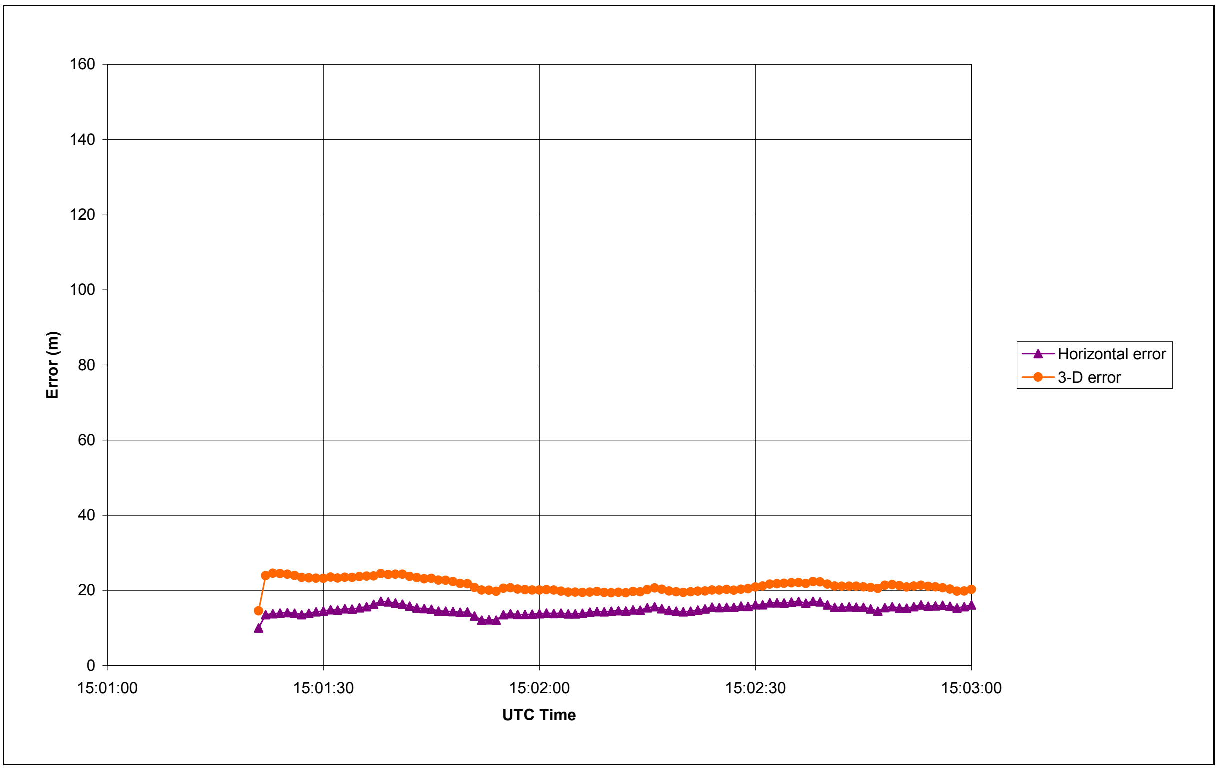

Figure 11 GD1 Position Errors for Window B

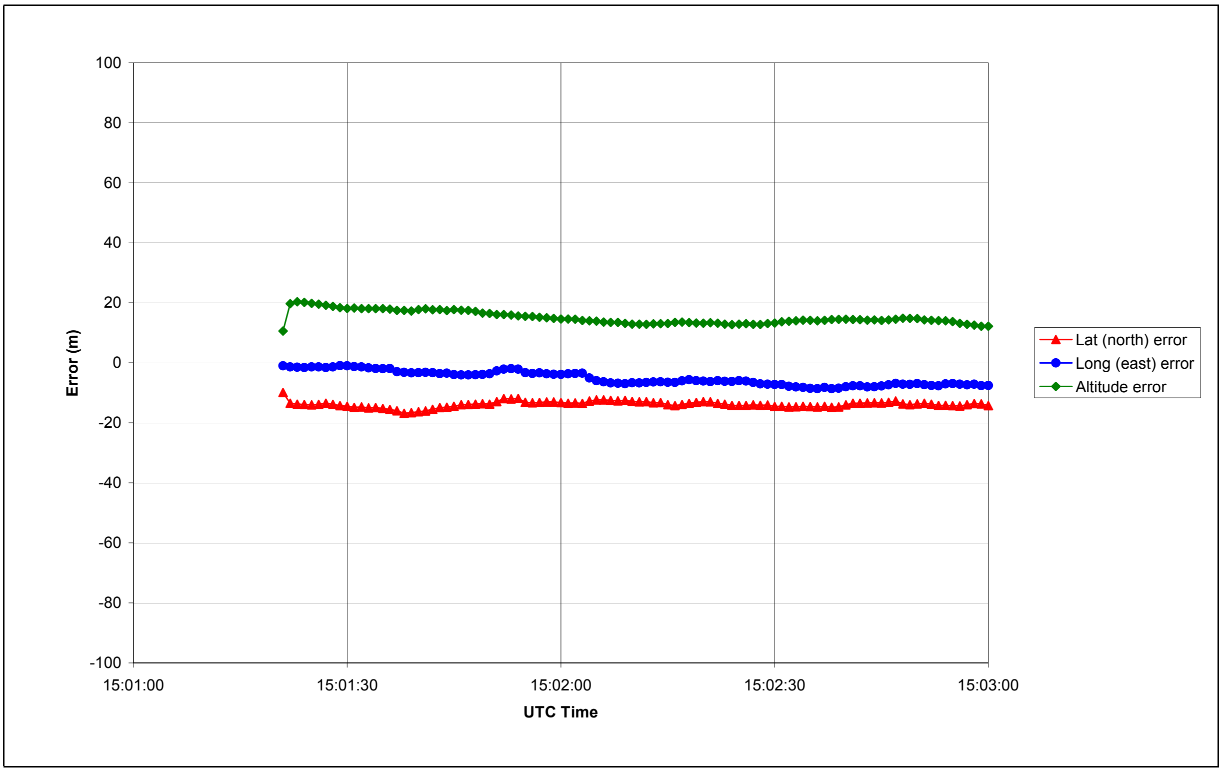

Figure 12 GD1 Position Error Components for Window B

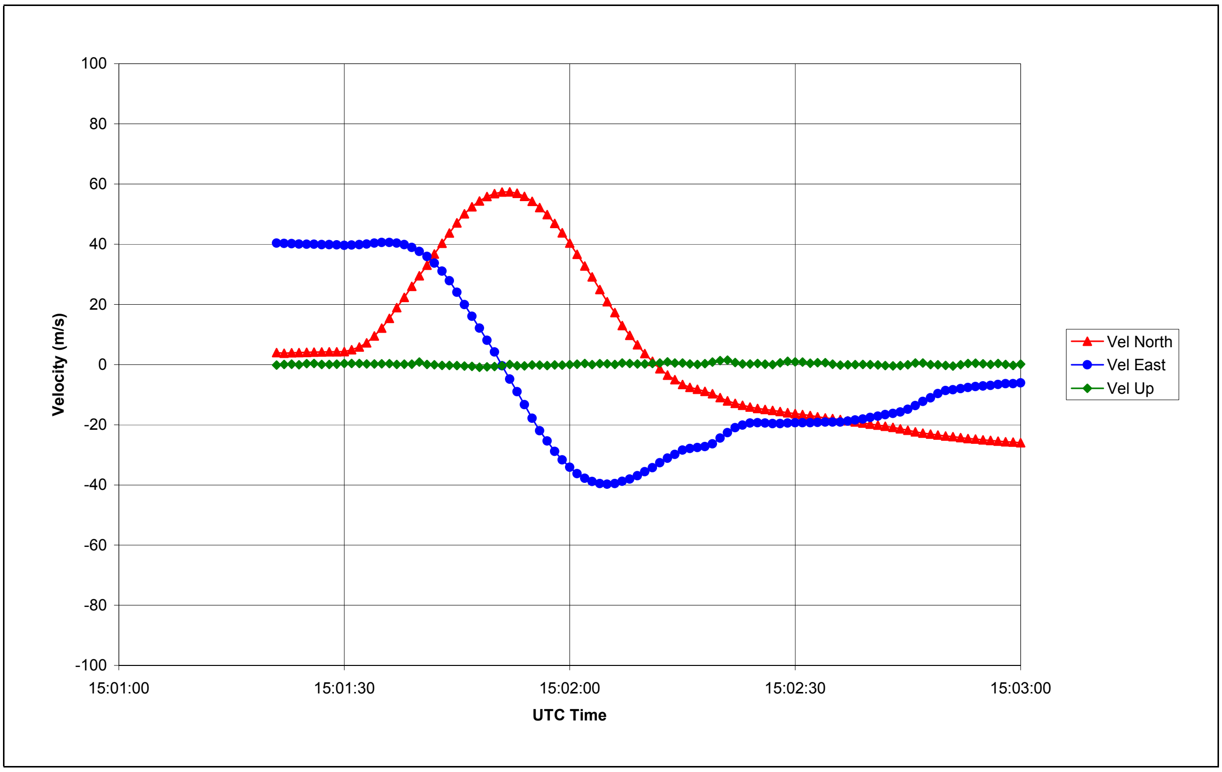

Figure 13 GD1 Velocity Components for Window B

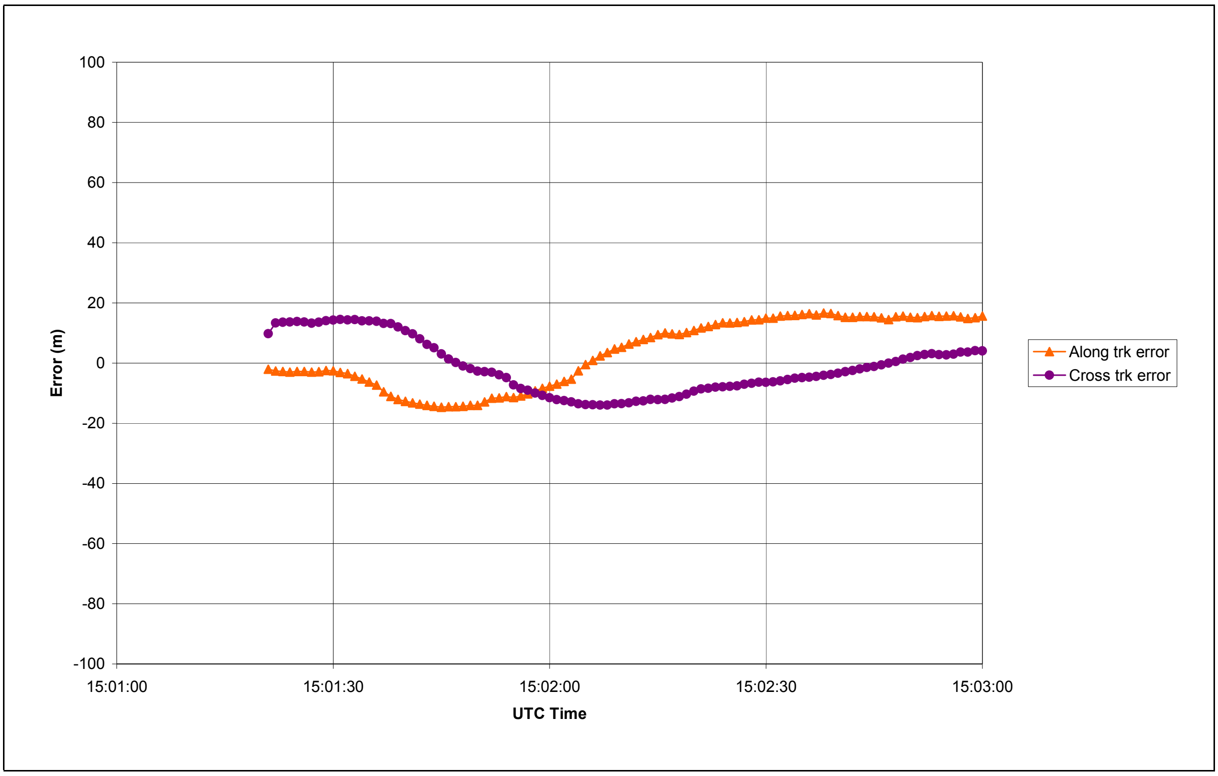

Figure 14 GD1 Along-Track and Cross-Track Position Errors for Window B

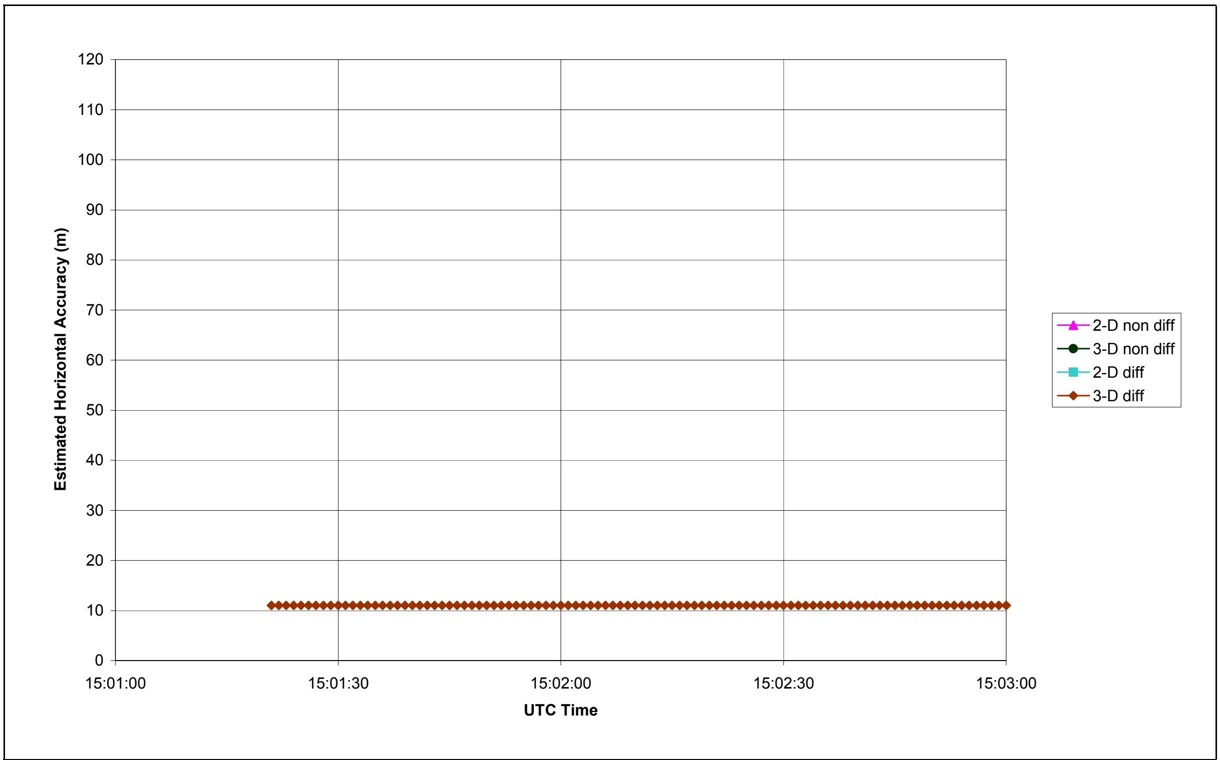

Figure 15 GD1 Estimated Horizontal Accuracy and Operating Mode for Window B

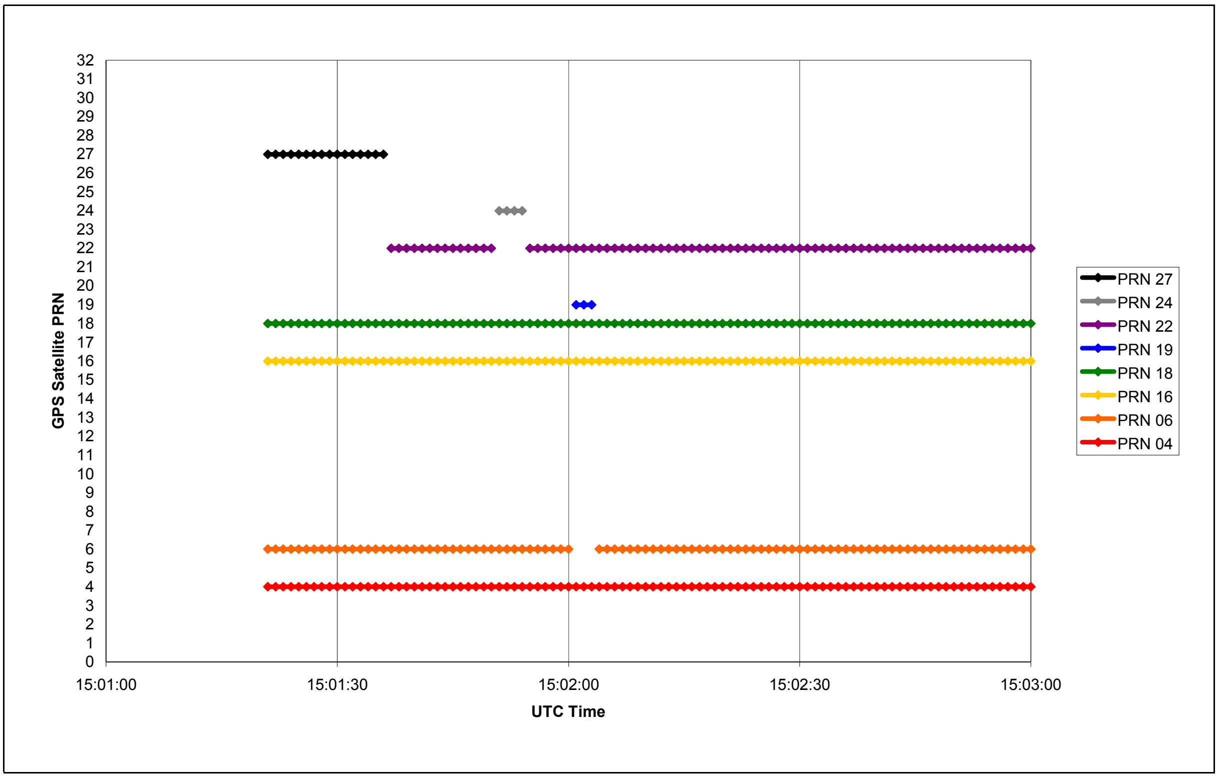

Figure 16 GD1 Satellite Usage for Window B

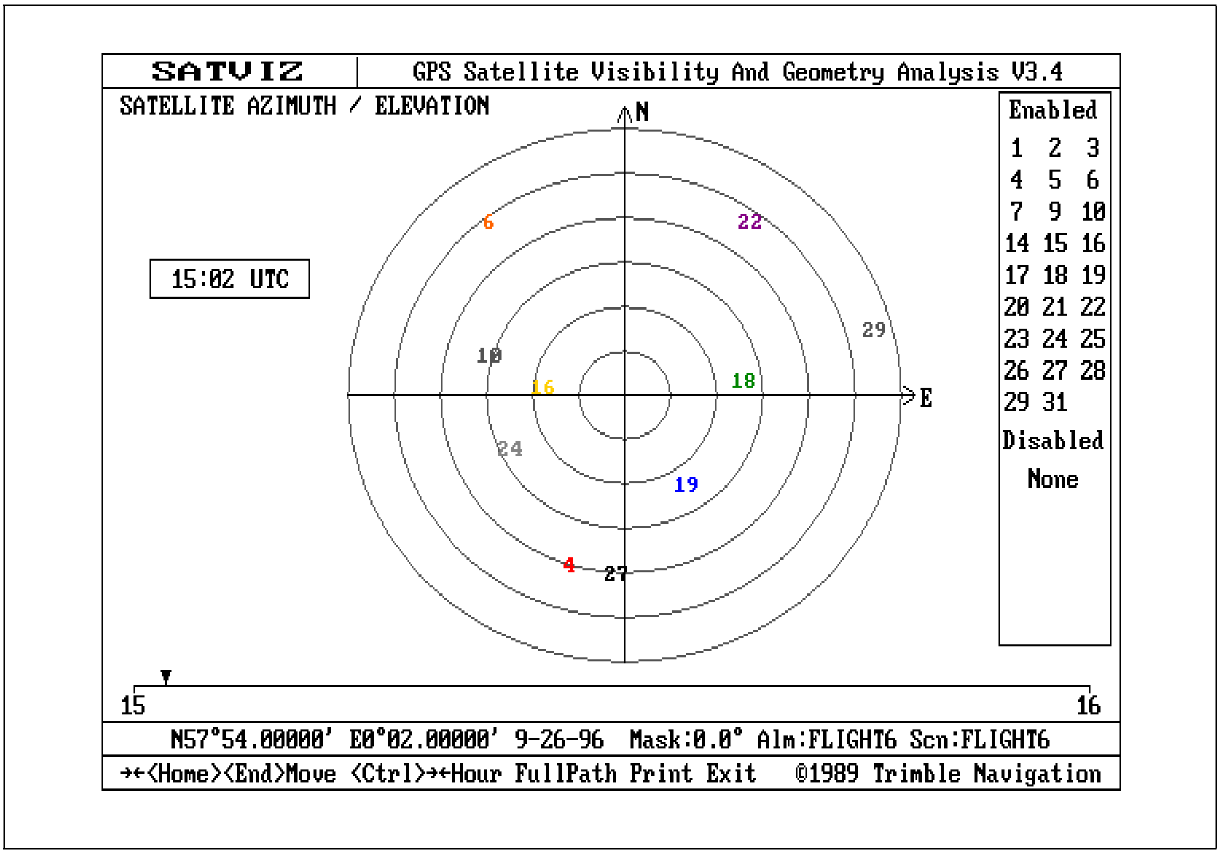

Figure 17 Satellite Azimuth and Elevation for Window B

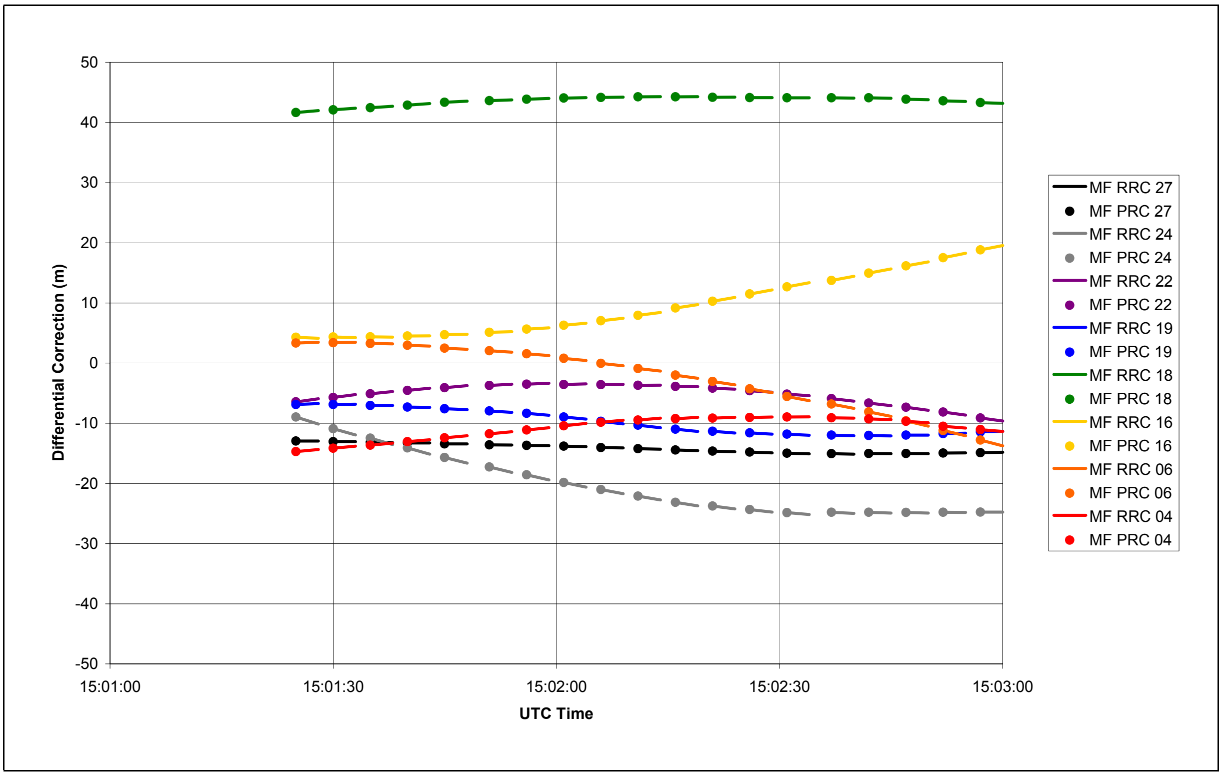

Figure 18 MF Differential Corrections for Window B

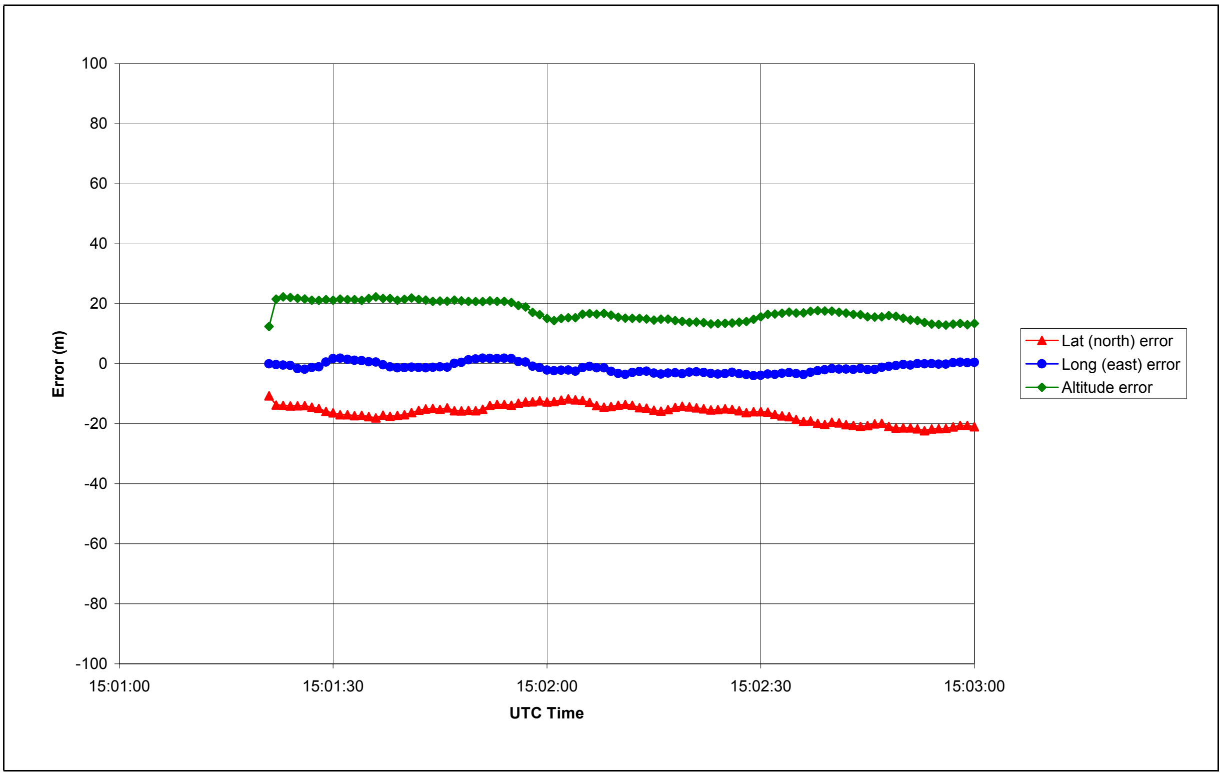

Figure 19 GD2 Position Error Components for Window B

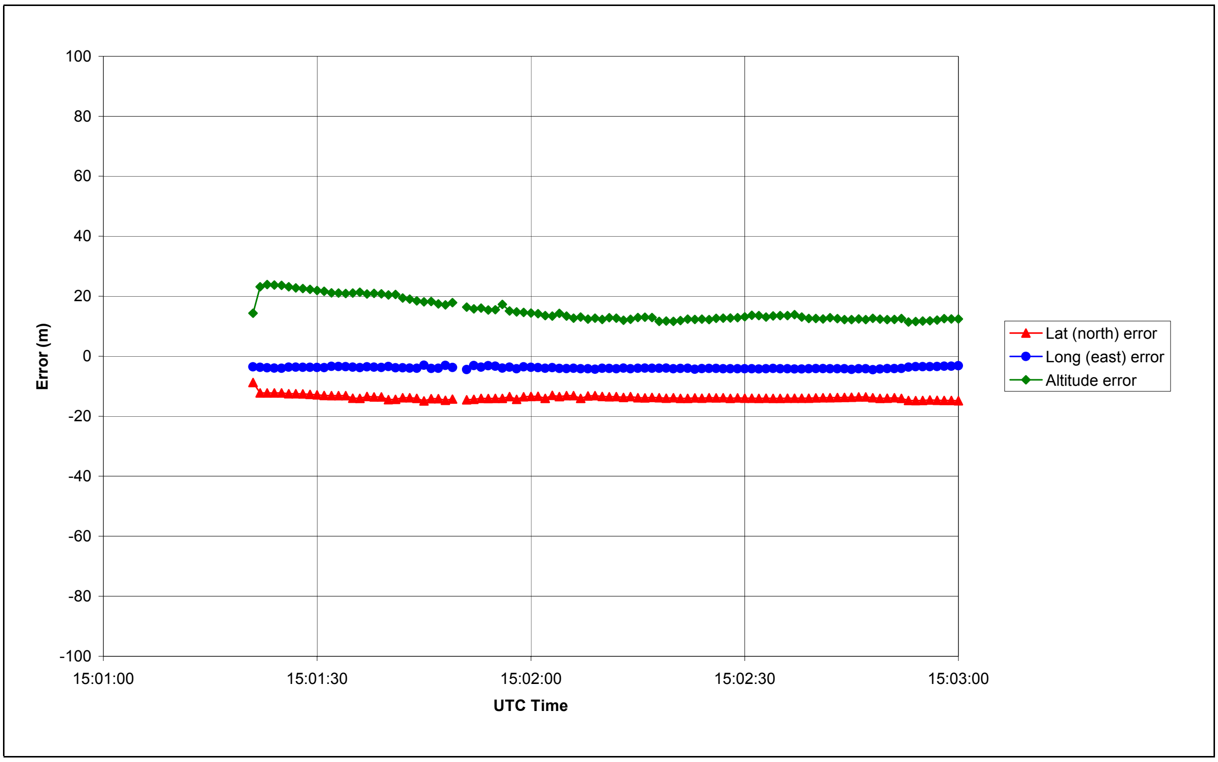

Figure 20 GD3 Position Error Components for Window B

5.3 Time Window C: Flight 4, 13:53 to 13:56 UTC

5.3.1 Manoeuvre details

This time window covers a three minute period commencing midway through approach number 3 on Flight 4 at the Piper B platform, and terminating during the go-around manoeuvre which followed the approach. At the start of the window, the aircraft was 1.7nm from the platform descending through 660ft altitude. At time 13:53:38 the aircraft transitioned to level flight at an altitude of approximately 200ft, and it passed within 175m of the platform at time 13:53:59. The go-around was commenced at time 13:54:08 and the aircraft climbed straight ahead. At the end of the time window it was 3.6nm from the platform passing 1400ft in the climb.

The approach was performed at high speed in a downwind direction, with the result that the rate of descent was observed by the pilot to be “too high” (ref 1 volume 3). Assuming a ground speed of 100kt, the rate of change of altitude during the 3.5° approach and go-around segments would have been approximately 600ft/min (or just over 3m/s).

5.3.2 Truth data availability

Satisfactory truth data was available throughout this time window. The OTruth and PTruth solutions were in agreement to better than ±1m in each axis.

5.3.3 GD2 position accuracy

Figure 21 shows the horizontal and 3-D position errors for the GD2 (UHF-corrected Navstar) receiver. Between times 13:53:23 and 13:55:08, the receiver operating mode changed to 2-D differentially corrected on six occasions. For the remainder of the window, the unit operated in 3-D differential mode.

The plot demonstrates that the position errors underwent a series of six ramp-style excursions, which correlate with the times when the receiver was operating in 2-D mode. In each case, the horizontal position error appears to have increased at a broadly constant rate (typically increasing by between 6m and 9m per second), before rapidly reducing to a value below 20m when the receiver returned to 3-D mode.

On the first excursion, the horizontal error reached 123.1m (the worst error sample observed during the trials programme) at time 13:53:41, after which the error reduced slightly over the following six seconds prior to the receiver re-entering 3-D mode.

On the subsequent excursions, the maximum errors were 117.2m, 72.8m, 45.0m, 18.2m and 99.8m. The errors increased continuously on these excursions, with the maximum value being reached on the final sample before 3-D operation was resumed.

On all but the second excursion, the error began increasing as soon as the receiver reverted to 2-D mode. For the second excursion, it appears that the onset of the rapid increase was delayed by fifteen seconds until around time 13:54:08.

Figure 22 depicts the individual GD2 position error components. This reveals that the direction of the first excursion (northwesterly and upward) was opposite to that of the subsequent ones (southeasterly and downward). It is perhaps significant that the aircraft altitude was reducing during the first of these excursions, but was increasing during the subsequent ones. It is also apparent that the relationship between the ratio between the magnitudes of the three error components remained approximately constant during the various excursions (e.g. the ratio of the latitude to the longitude component was around 1.2:1).

5.3.4 GD2 velocity components

Figure 23 shows the velocity components obtained from the GD2 receiver. Although the underlying velocity components appear to have been fairly stable, there is evidence for a series of step changes which correlate with the start and end of the various position error excursions (those at the end of the excursions are, in general, more pronounced) and with the changes in the receiver operating mode.

The direction in which each velocity component changed at each “step” appears to be the same as the rate of change of the associated position component. It can also be observed that the value of the vertical velocity component as computed by the receiver was zero during the times when the receiver operated in 2-D mode, despite the fact that the aircraft altitude was changing.

5.3.5 GD2 along-track and cross-track position errors

Figure 24 depicts the along-track and cross-track components of the GD2 horizontal error. This reveals that the errors were principally in an along-track sense (up to ±120m), with the cross-track error component remaining within the range +37m to -22m.

Since the guidance indications during this approach were being derived from the output of the GD1 receiver, the presence of the excursions on the GD2 data was not apparent to the pilots.

5.3.6 GD2 estimated accuracy parameter

The GD2 Estimated Horizontal Accuracy parameter, and the receiver operating mode, are plotted in Figure 25. The accuracy parameter was observed to increase, from around 15m up to 30m, when operating in 2-D differential mode.

Clearly, the algorithm used to compute this parameter was underestimating the extent of the errors by a considerable margin.

5.3.7 GD2 satellite usage

Figure 26 shows the satellites which were employed by the GD2 receiver to generate its position solution. Examination of the plot reveals that only three satellites (PRNs 1, 14 and 15) were used in the solution during the periods where the receiver reverted to 2-D mode.

Figure 27, the polar plot showing the satellite positions in elevation and azimuth, demonstrates that the projected positions of the three satellites appear to have been almost in a straight line. This is a situation (with the satellites close to being co-planar in three dimensions) which can often result in a reduced precision solution due to poor satellite geometry.

The commercial satellite prediction software’s estimate for a three-satellite equivalent to HDOP (although it is not known precisely how the software calculates this) was between 10.3 and 11.4 over the analysis time window, suggesting that the geometry with these three satellites was poor. The introduction of an additional satellite (PRN 29) into the calculation resulted in the HDOP falling below 4.

The tracking data for the truth reference receiver was examined to determine the number of satellites which this unit had been tracking. This revealed that a minimum of four satellites had been tracked throughout the time window, with PRNs 1, 7, 14 and 15 available throughout. The presence of PRN 7 in the tracking list, and its geometry relative to the aircraft (the satellite elevation was 33°, and its azimuth would have placed it slightly to the left of the aircraft nose), suggested that this satellite might have been expected to feature in the set used by the GD2 receiver. It is therefore somewhat surprising that this satellite does not appear at all in Figure 26.

5.3.8 Differential corrections

Figure 28 shows the pseudo range correction (PRC) data which was transmitted by the UHF reference station on the platform. It should be noted that the transmission of a correction message did not necessarily guarantee that it had been received on the aircraft (there is a subtle difference compared to the MF correction data, since the recorded information for the latter related directly to the corrections received on the aircraft rather than being recorded at the transmitter). However, examination of the GD2 data confirms that the interval between successive correction updates was normally either 1s or 2s; with the longest such interval being 4s, suggesting that there was not a significant problem with correction reception.

To allow a comparison to be made with the corrections received from the onshore station, the plot also includes the MF PRC data for those satellites for which corrections were available from both sources. Since the MF message rate was lower than that for the UHF system (which provided corrections every second), the individual MF messages are denoted by the diamond symbols which are joined by straight lines.

Unfortunately, a software fault, resulting from a coding error, was present on this flight and affected the precision of some of the recorded differential correction data. The fault affected approximately one-third of the satellites for which corrections were being generated by the two systems, and did not necessarily affect the same satellite PRNs in the two cases since this was dependent upon the order in which the satellites appeared in the correction messages.

For those satellite corrections which were affected by the fault, the effect was that the resolution of much of the recorded PRC data changed from 0.02m to 5.12m (i.e. by a factor of 256). The affected data has, where possible, been included in Figure 28 but is denoted by an asterisk: the effect of the reduced resolution is clearly visible in the figure. The remaining portion of the data has had to be eliminated from the figure to avoid confusion.

Two facts are immediately apparent from a consideration of Figure 28. The first is that satellite PRN 07 was absent from all of the UHF correction messages, which provides the explanation as to why the GD2 receiver did not use this satellite in its differentially-corrected solution. The second observation is that there is clear evidence for the presence of sinusoidal variations, with a period of several tens of seconds, on the transmitted UHF corrections for the majority of the satellites (PRNs 2, 14, 29, 31 and conceivably also 15). These variations are absent from the corresponding MF correction data and it is believed that they were due to L-band multipath effects, resulting from the presence of metallic platform structures close to the UHF reference station.

The greatest discrepancy between the UHF and MF corrections appears to have been for PRN 29, where the amplitude of the variation approached ±20m. However, for the three satellites (PRNs 1, 14 and 15) associated with the large position excursions, the discrepancy appears to have been somewhat lower (less than ±10m).

5.3.9 GD1 operation

Figure 29 depicts the position error components for the GD1 (MF-corrected Navstar) receiver. There is no evidence of any significant disturbances on the receiver output : the largest horizontal error was 7.4m, at time 13:55:33.

The GD1 receiver remained in 3-D differentially corrected mode throughout, with a minimum of four satellites employed in the solution (PRNs 1, 7, 14, and 15 were used throughout the time window; along with periodic use of PRN 2, 29 and 31). An MF differential correction was being received regularly, with a maximum interval between successive messages of 6s.

No data for the GD3 (MF-corrected Trimble) receiver is available, since the unit was not operational on this flight.

5.3.10 **Discussion**

A preliminary analysis of the events described above was included in the original flight test report for this trial (ref 3). In particular, the absence of a UHF correction for PRN 7 was identified in the report, in spite of the fact that the satellite should have been clearly visible. Evidence for the latter statement is provided by the fact that signals from this satellite were being received by the platform truth reference receiver, which operated from a common antenna.

It was also suggested in ref 3 that the corrections for PRN 29 were likely to have been affected by multipath reflections from the platform structure, an effect which can be clearly observed on Figure 28. The report identified that the largest errors were observed whilst the GD2 receiver was operating in 2-D mode with only three satellites, but did not reach any conclusions regarding the precise mechanism involved.

However, with the benefit of the additional analysis presented above, it is possible to suggest that the errors were due to the following series of factors:

a) The satellite geometry resulted in a lower than average number of visible satellites at "reliable" elevations: for example, there were only four satellites with an elevation in excess of 30°.

b) For substantial portions of the approach, the Navstar receivers on the aircraft were only able to maintain lock on these four highest elevation satellites (PRNs 1, 7, 14 and 15).

c) For reasons which remain unknown, although they are probably connected with the presence of an L-band multipath environment on the platform, the UHF correction station failed to generate corrections for one of these four satellites (PRN 7).

d) In order to continue operating in differential mode, the GD2 receiver was forced to revert to a three satellite (2-D solution) using PRNs 1, 14 and 15.

e) This particular combination of satellites provided poor geometry, with the prediction software computing a three-satellite HDOP equivalent in excess of 10.

f) No external altitude source to the GPS receiver was available, with the result that the unit was obliged to rely on its last known altitude estimate (and an implicit assumption that the latter continued to remain valid) to provide a fourth "measurement" for the computation of horizontal position.

g) Unfortunately the aircraft altitude was not constant, but was decreasing (at a rate of approximately 3m/s) prior to time 13:53:40, and then began increasing at a similar rate after time 13:54:08.

h) The resulting altitude measurment "error" propagated directly into the receiver's horizontal position solution, with an effect which was almost certainly magnified by the poor satellite geometry. This resulted in an increasing horizontal error at a rate of up to 9m/s, and with a direction which varied according to whether the aircraft was descending or climbing.

i) On each occasion where the receiver regained lock on a fourth satellite, it returned to 3-D operating mode and was thereby able to recompute an accurate altitude solution. This caused the horizontal error at the start of each subsequent period of 2-D operation to be reset, resulting in the "ramp up / fly back" effect visible on the position error plots.

j) The presence of a limited amount (less than ±10m) of multipath corruption on the UHF differential corrections for PRNs 1, 14 and 15 may have provided an additional error component in the GD2 solution, but this was probably not the primary cause of the large errors. In particular, the inclusion of an additional satellite for which the multipath component was larger (PRN 29, up to ±20m of multipath) into the calculation resulted in an improvement, and not a degradation, of the position solution.

k) Confirmation that the aircraft itself was probably not being affected by multipath is provided by the fact that there were no significant errors in the position solution (which included PRNs 1, 14 and 15) generated by the GD1 receiver.

l) Receiver clock drifts provide a alternative hypothesis for the increasing horizontal position errors in 2-D mode. However, the strong correlation with the changes in aircraft altitude suggest the latter as being the most likely cause.

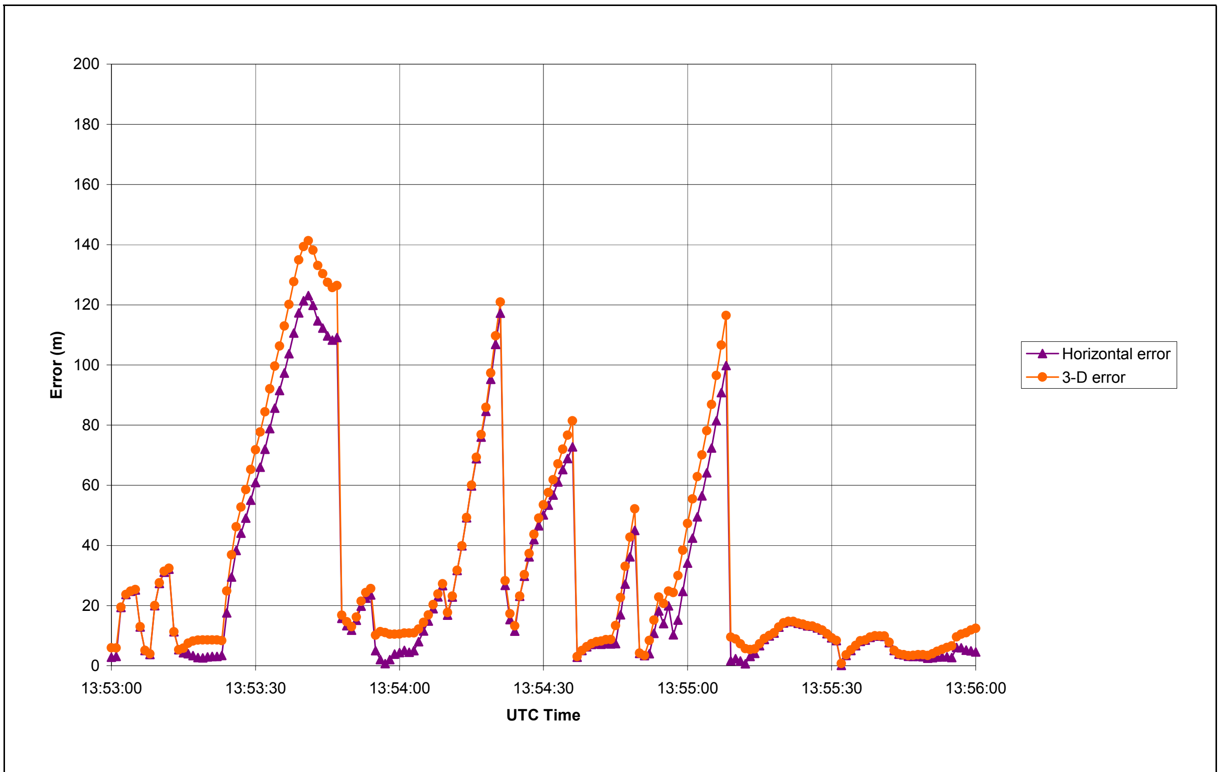

Figure 21 GD2 Position Errors for Window C

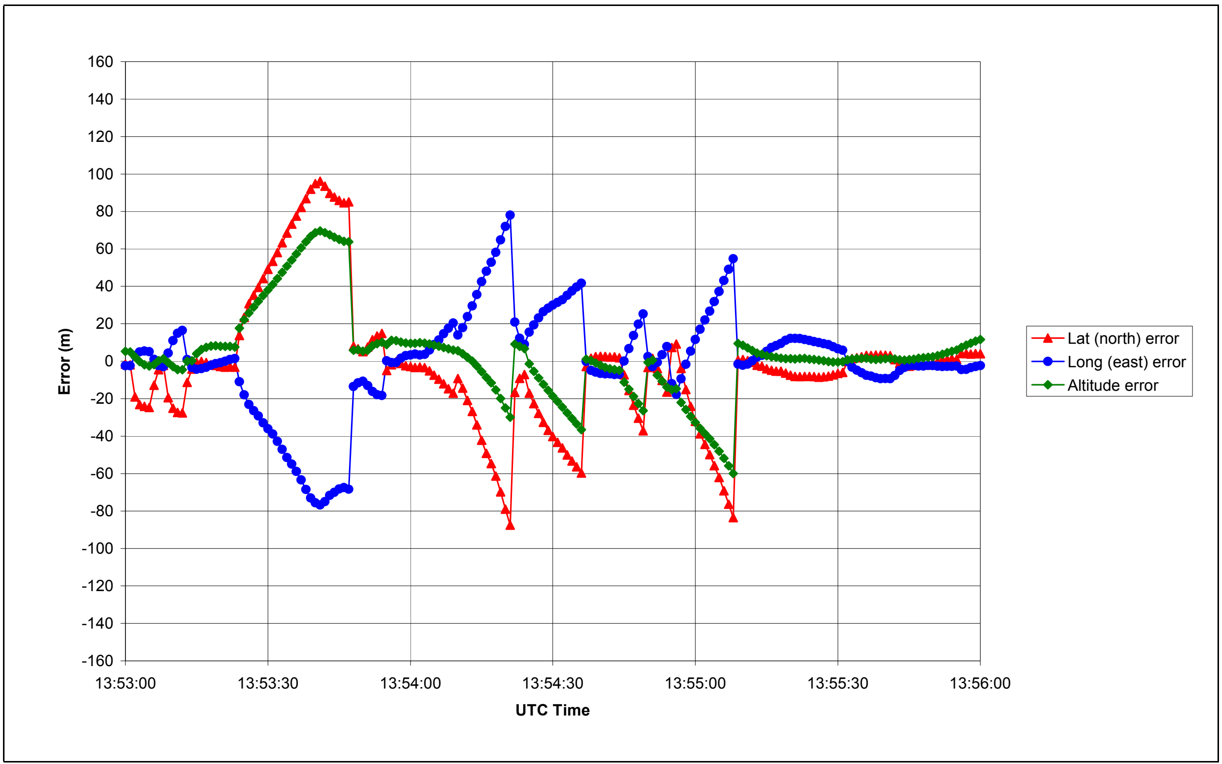

Figure 22 GD2 Position Error Components for Window C

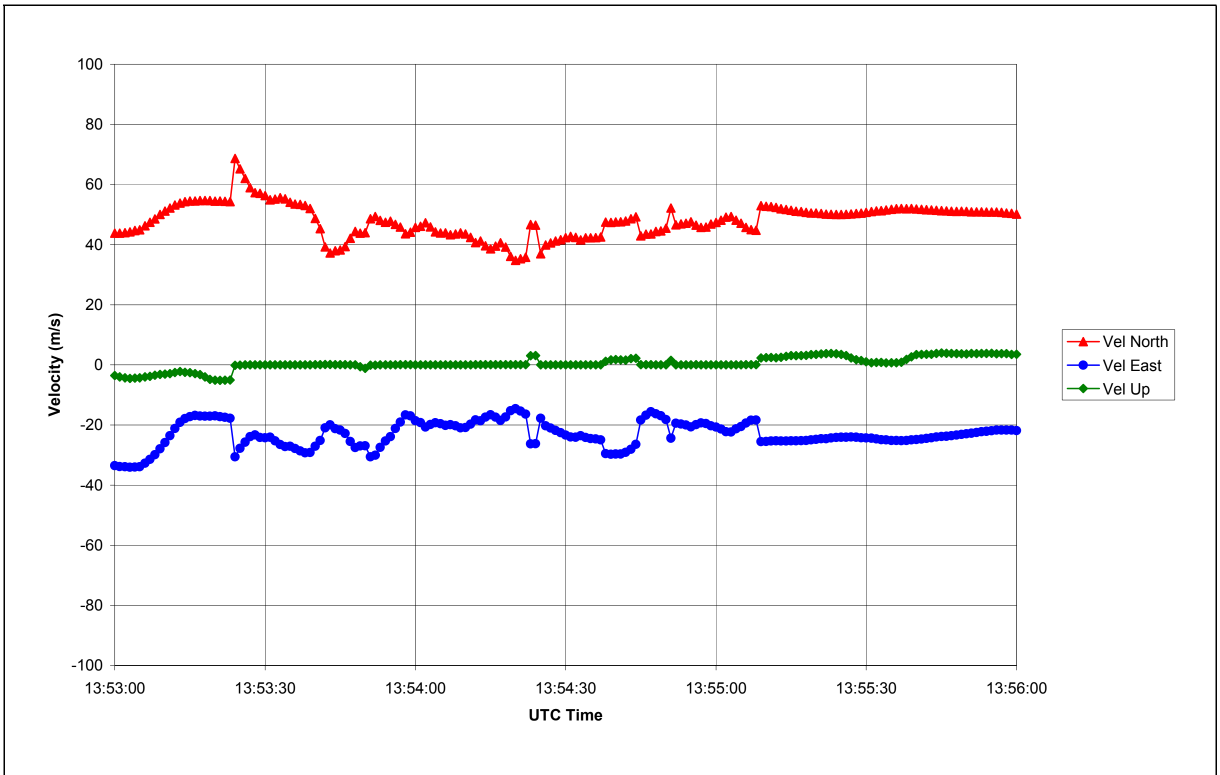

Figure 23 GD2 Velocity Components for Window C

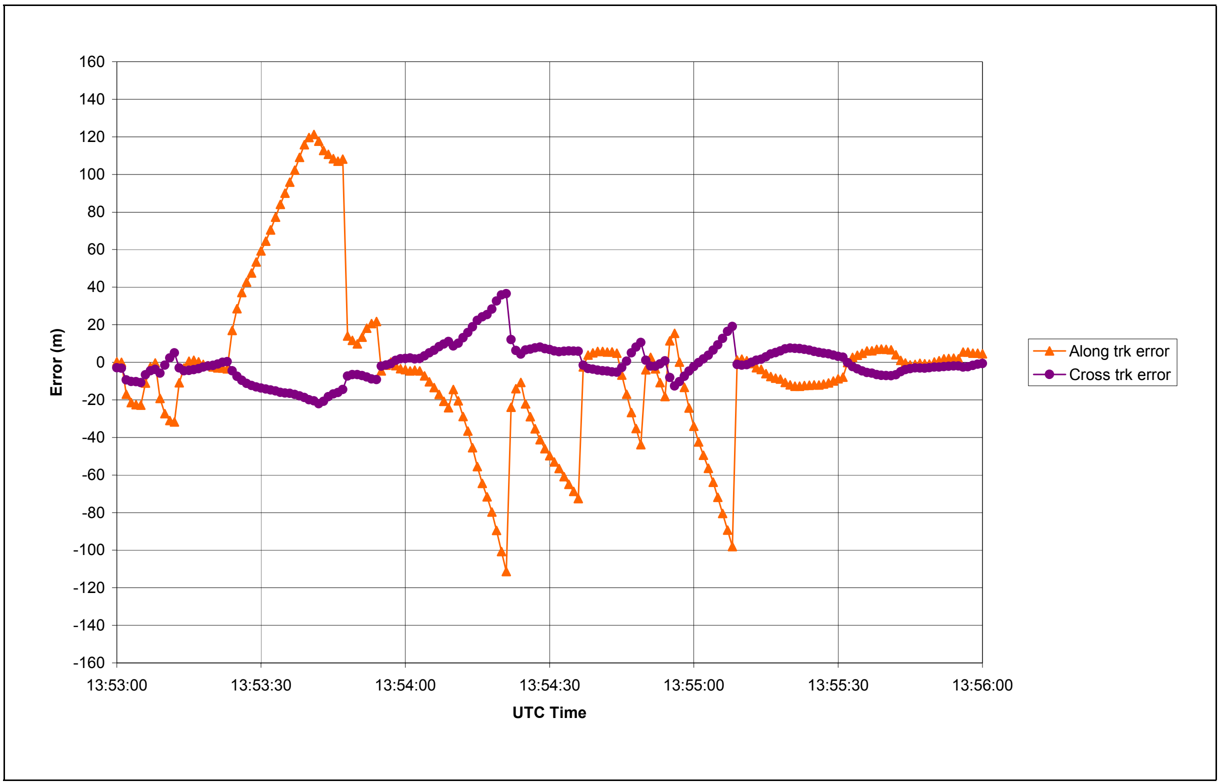

Figure 24 GD2 Along-Track and Cross-Track Position Errors for Window C

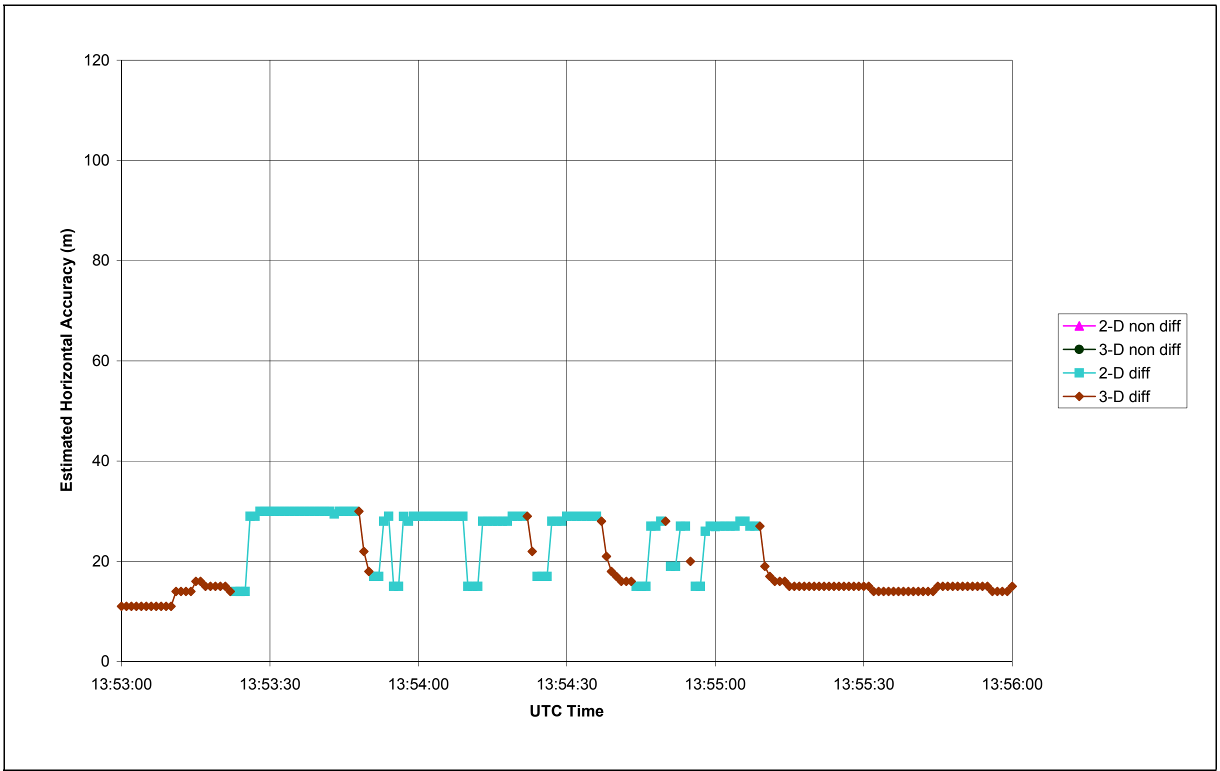

Figure 25 GD2 Estimated Horizontal Accuracy and Operating Mode for Window C

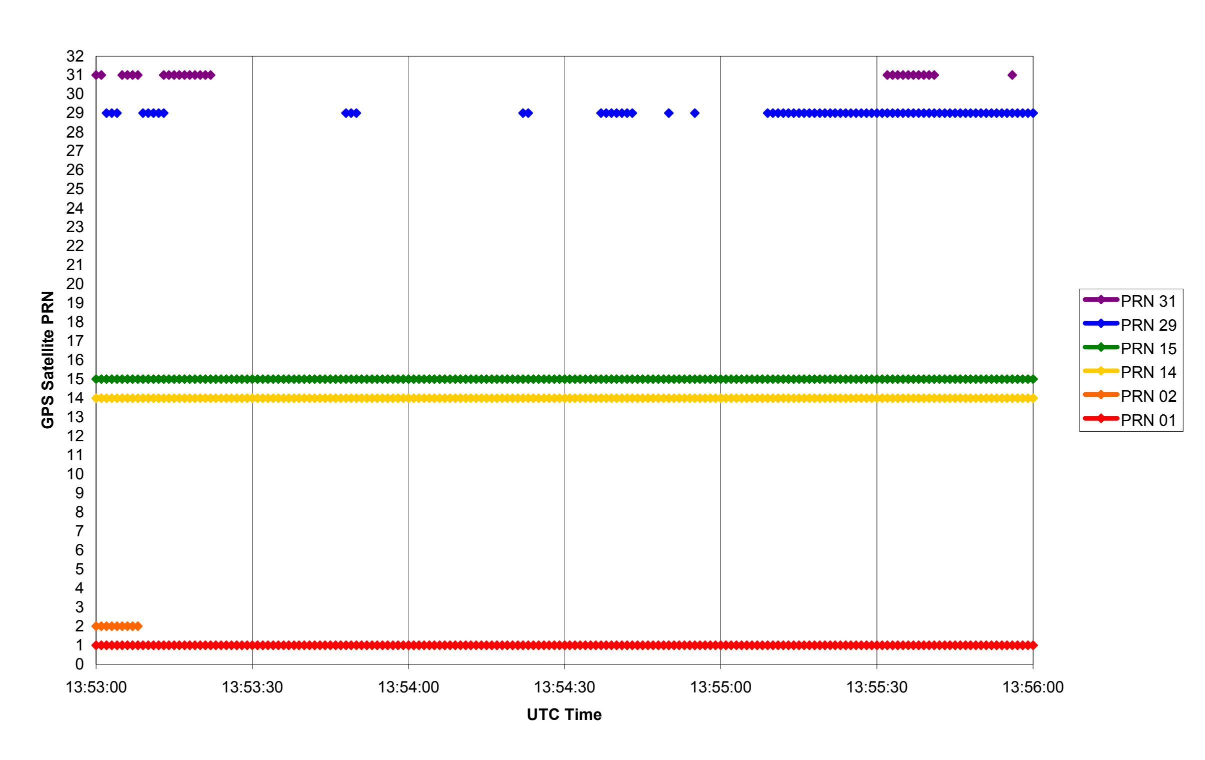

Figure 26 GD2 Satellite Usage for Window C

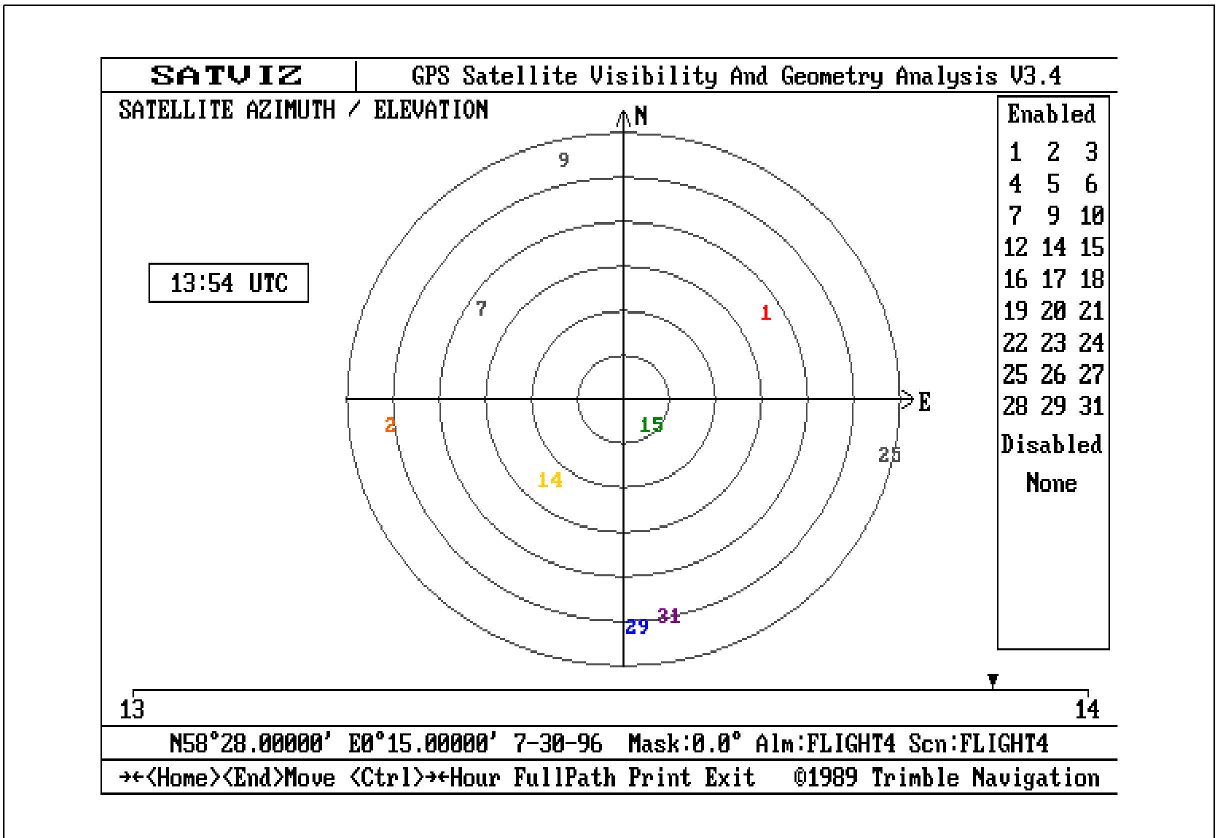

Figure 27 Satellite Azimuth and Elevation for Window C

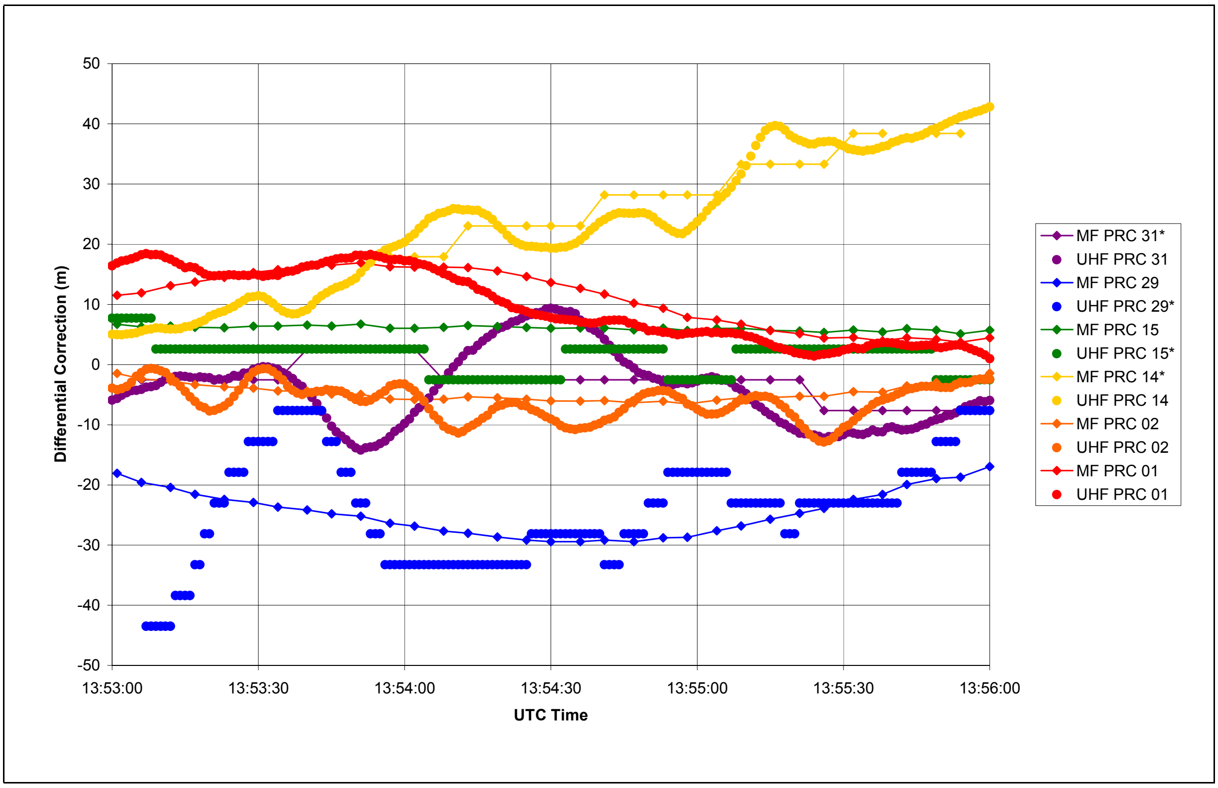

Figure 28 UHF and MF Differential Corrections for Window C (asterisk denotes reduced precision corrections due to software fault)

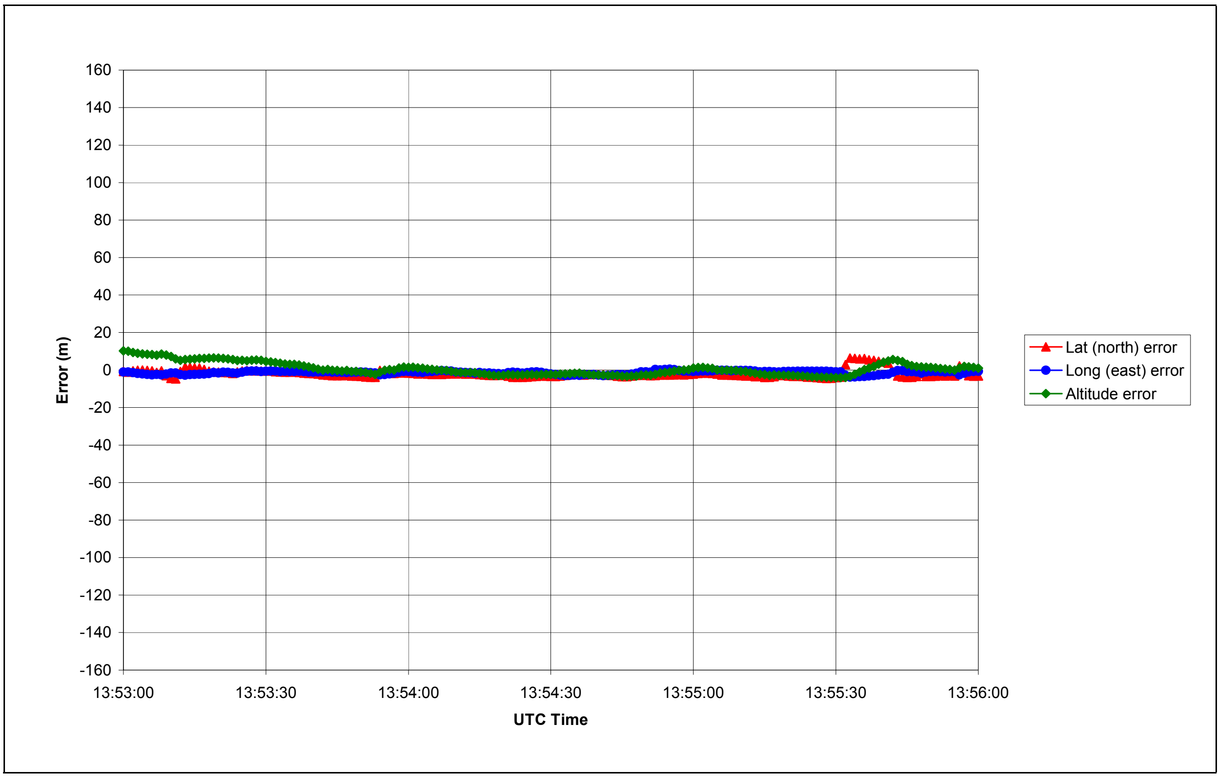

Figure 29 GD1 Position Error Components for Window C

5.4 Time Window D: Flight 5, 13:31 to 13:36 UTC

5.4.1 Manoeuvre details

This time window covers a five minute period commencing midway through approach number 11 on Flight 5 at the Tartan A platform, and terminating during the go-around manoeuvre which followed the approach. At the start of the window, the aircraft was 3.6nm from the platform in level flight at approximately 800ft altitude. The descent was commenced at time 13:32:10, and at time 13:34:28 the aircraft had completed the transition to level flight at an altitude of approximately 200ft, passing within 230m of the platform at time 13:35:14. The go-around was commenced at time 13:35:25 and the aircraft entered a left turn as it climbed. At the end of the time window it was 0.4nm from the platform passing through 650ft.

5.4.2 Truth data availability

Satisfactory truth data was available throughout this time window. The OTruth and PTruth solutions were in agreement to better than ±1m in each axis.

5.4.3 GD2 position accuracy

Figure 30 shows the horizontal and 3-D position errors for the GD2 (UHF-corrected Navstar) receiver. The receiver operating mode was 3-D differentially corrected up until time 13:33:02 where it changed to 2-D differential. During the following thirty seconds it switched several times between 2-D differential and 3-D differential mode, then remained in 2-D differential mode until time 13:35:34. The mode then returned to 3-D differential for the remainder of the time window.

Examination of the plot reveals that there were two oscillations of the position error during the initial period of 3-D differential operation. In the case of the second oscillation, the horizontal error increased to a maximum of 95.5m. Following these oscillations, the errors fell to around 20m before increasing (but more gradually), to a maximum of 63.2m whilst the receiver was operating in 2-D mode. On the final return to 3-D operation, the errors fell below 20m once again and no further oscillations were observed.

Figure 31 depicts the individual GD2 position error components. This reveals evidence for a sinusoidal variation of the latitude and longitude components, with period of around 1.5 minutes, throughout the time window. The altitude component was observed to vary initially in a sinusoidal manner (peaking at 153.2m) during the period of 3-D operation, then changed to a ramp-type effect while the receiver had reverted to 2-D mode.

Changes in the rate of change of the altitude error component were observed to occur at times which correlated with changes in the aircraft altitude profile. In particular, prior to time 13:34:28 the aircraft was descending and the altitude error was increasing. The aircraft then flew level (with the altitude error remaining approximately constant) until time 13:35:25, after which the aircraft climbed (and the altitude error reduced). This suggests that a similar effect to that which occurred during Window C may have been at work.

5.4.4 GD2 velocity components

Figure 32 shows the velocity components obtained from the GD2 receiver. Prior to the initiation of the go-around, the velocity components appear to have been stable, with the exception of a very small step change at time 13:33:03 (which correlates with the first reversion to 2-D mode). Whilst operating in 2-D mode, the velocity component remained within ±1m/s of zero.

5.4.5 GD2 along-track and cross-track position errors

Figure 33 depicts the along-track and cross-track components of the GD2 horizontal error. This reveals that the along-track error components were slightly larger (between -20m and +80m) than those in the cross-track sense (-25m to +50m). Since the guidance indications during this approach were being derived from the output of the GD1 receiver, the presence of the excursions on the GD2 data was not apparent to the pilots.

5.4.6 GD2 estimated accuracy parameter

The GD2 Estimated Horizontal Accuracy parameter, and the receiver operating mode, are plotted in Figure 34. A steady reduction in the accuracy parameter, from an initial 43m down to less than 15m, was observed. Lower values were indicated in 2-D mode than in 3-D mode. It would appear that the variations in the horizontal accuracy parameter appeared to have borne little relation to the position errors which were actually occurring.

5.4.7 GD2 satellite usage

Figure 35 shows the satellites which were employed by the GD2 receiver to generate its position solution. Initially four satellites were available (PRNs 2, 7, 14 and 15), and reversion to 2-D mode correlated with the loss of PRN 2 from the solution. At time 13:35:34, not only did PRN 2 reappear in the solution but a fifth satellite (PRN 31) also became available. This event correlated with the final return to 3-D differential mode, following which no further large oscillations occurred.

The satellite elevation and azimuth positions are shown on the polar plot in Figure 36. This reveals that PRN 2, the satellite whose loss from the position solution caused the reversion to 2-D operation, was at relatively low elevation (18°) and would have been almost directly ahead of the aircraft.

Satellite 31 was at 19° elevation on the left-hand side of the aircraft, where the presence of the tail rotor was known (ref 1 volume 2) to give rise to reception difficulties. When the aircraft began turning to the left, measurements from this satellite began to be included in the solution.

Examination of Figure 36 suggests that PRN 1 (at an elevation of 36°) should also have been available for inclusion in the position solution. Examination of the tracking data for the truth reference receiver confirms that PRN 1 was indeed available to this latter unit (and also that the tracking situation in respect of PRN 2 and PRN 31 was as suggested above).

A geometry calculation using the satellite prediction software reveals that the HDOP associated with a four satellite solution (i.e. PRNS 2, 7, 14 and 15) would have been very large during this time window: initially 25.6, then reducing to a value of 9.1. When PRN 2 was removed from the calculation to determine a three-satellite solution, the resulting three-satellite HDOP was much lower (around 2.5). This may explain why the value of the estimated accuracy parameter was less when operating in 2-D mode. The addition of a fifth satellite (PRN 31) into the computation almost certainly resulted in a further improvement to the geometry. Although it was not possible to compute the HDOP for an overdetermined five-satellite solution using the satellite prediction software, it was demonstrated that an HDOP as low as 2.4 could be achieved using one of the available subsets of four satellites (PRN 2, 7, 14 and 31).

5.4.8 Differential corrections

Figure 37 shows the pseudo range correction (PRC) data transmitted by the UHF system, overlaid with the identical information received from the MF correction system for the same set of satellites. The maximum interval between successive UHF corrections received by the GD2 receiver was 3s. This flight was also affected by the recording fault described in section 5.3.8, and the affected correction measurements have been denoted by an asterisk.

The reason for the non-appearance of PRN 1 in the set of satellites used by the GD2 receiver in its navigation solution, is explained by the absence of this satellite from the transmitted UHF corrections. Consideration of the relative geometry of the UHF reference station, satellite PRN 1 (azimuth 061°, elevation 36°), and the platform structure suggests that the satellite signals were probably not being directly masked by the latter at the time in question. It is possible, however, that some form of masking may well have occurred during the thirty minute period which preceded the events described above.

Examination of the correction plot also reveals evidence for the presence of significant, and mostly oscillatory, discrepancies between the corresponding correction measurements generated by the UHF and MF systems. In many cases it is difficult to estimate the extent of the discrepancies due to the recording problems, but there is evidence for discrepancies of perhaps ±15m on PRN 7, and up to 30m on PRN 2 (the error on the latter is more akin to a constant bias). The period associated with these oscillations appears to have typically been in the region of between 1 and 1.5 minutes, which is compatible with the period of the oscillations which were observed in the individual position error components.

5.4.9 GD1 operation

Figure 38 depicts the position error components for the GD1 (MF-corrected Navstar) receiver. The behaviour of this receiver was clearly benign throughout the time window; the maximum horizontal error was 6.7m. The GD1 receiver remained in 3-D differential mode throughout. The maximum interval between MF correction updates was 15s, due to the two missed messages around time 13:35:20 which can be observed on Figure 37. Examination of the list of satellites employed in the GD1 solution reveals that PRNs 1, 7, 14 and 15 were in use throughout, plus PRN 2 during the periods where it was being included in the tracking data.

No data for the GD3 (MF-corrected Trimble) receiver is available, since the unit was not operational on this flight.

5.4.10 Discussion

It would appear that several factors are likely to have been at work. The absence of PRN 1 from the UHF correction messages probably explains why the GD2 receiver employed a combination of satellites (PRNs 2, 7, 14 and 15) with very poor geometry. Associated with the poor satellite geometry was the presence of errors, of up to around ±15m, on the UHF corrections for the satellites employed in the solution. These errors are most likely to have been due to the L-band multipath environment on the platform, as was suggested in the flight test report for this trial (ref 4). The combination of the correction errors and the poor satellite geometry (high HDOP), probably resulted in the large and oscillatory position error characteristics during the period where the receiver was operating in 3-D differential mode.

The loss of one of the four satellites then forced the receiver into 2-D mode. Somewhat counter-intuitively, this resulted in a reduction in the HDOP implying that the geometry was improved (this may result from the different calculation method associated with the generation of a dilution of precision value for a three satellite solution). However, the fact that the aircraft altitude was changing while operating in 2-D mode, then introduced another error source resulting from the assumption made within the receiver that the altitude remained constant. This probably caused the altitude error characteristic to exhibit a ramp behaviour, and the horizontal error components to be affected both by the ramp effect and by the continued presence of multipath corruption on the differential corrections.

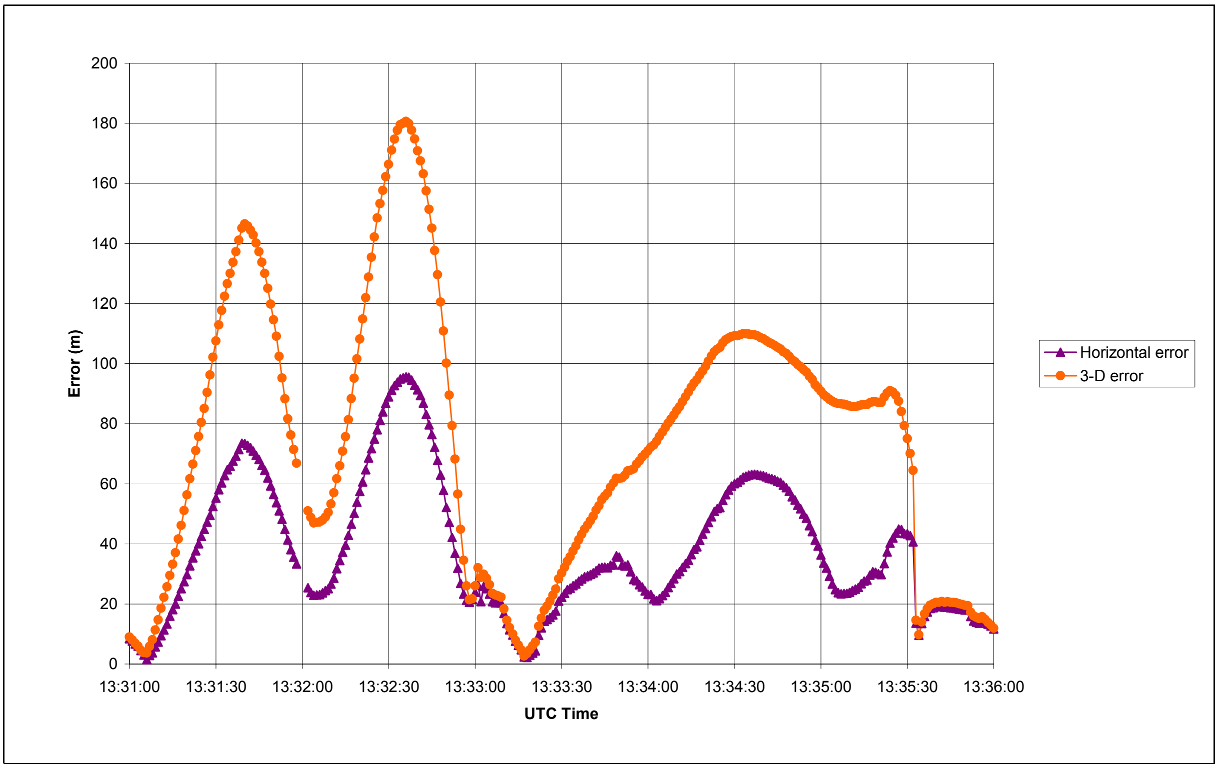

Figure 30 GD2 Position Errors for Window D

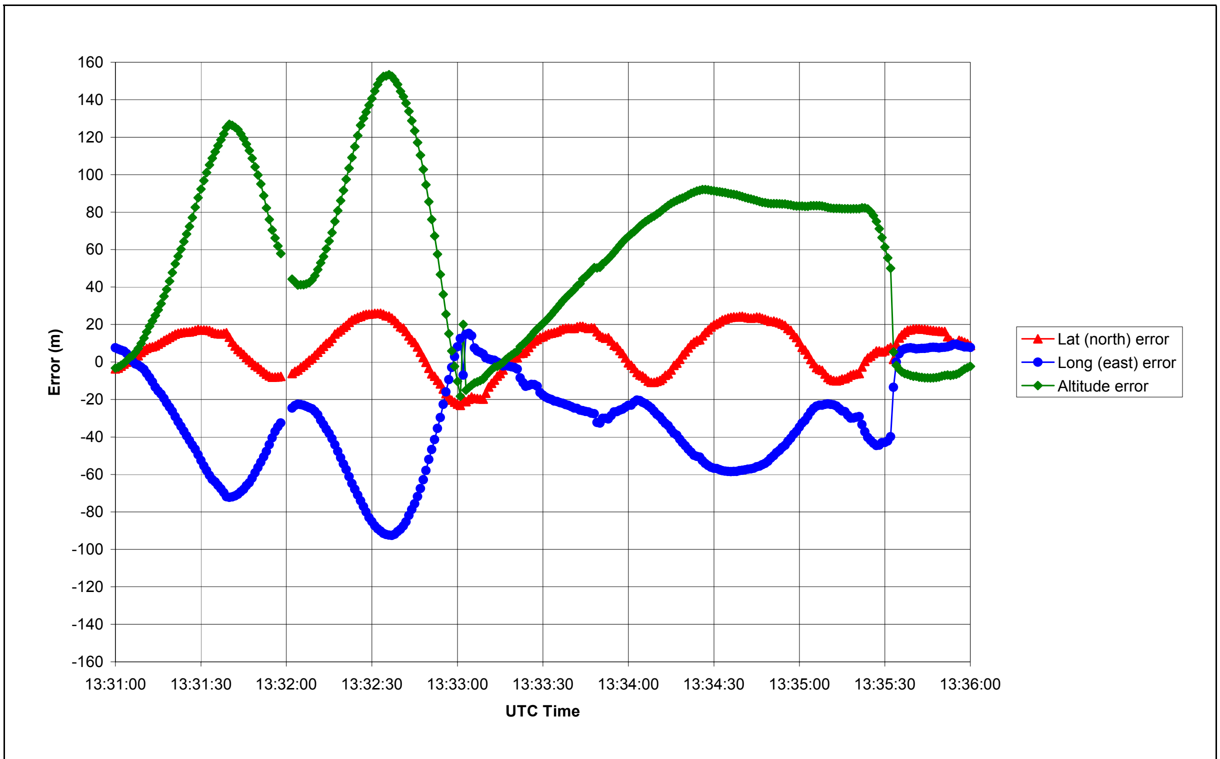

Figure 31 GD2 Position Error Components for Window D

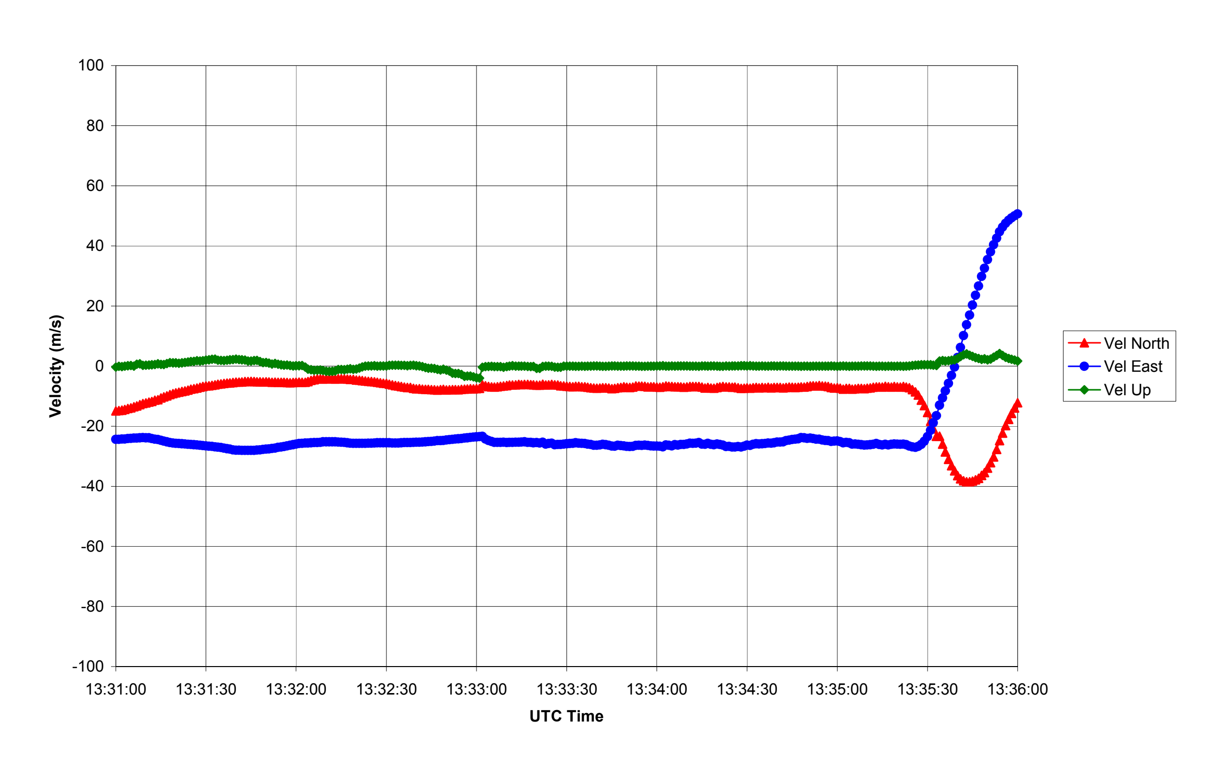

Figure 32 GD2 Velocity Components for Window D

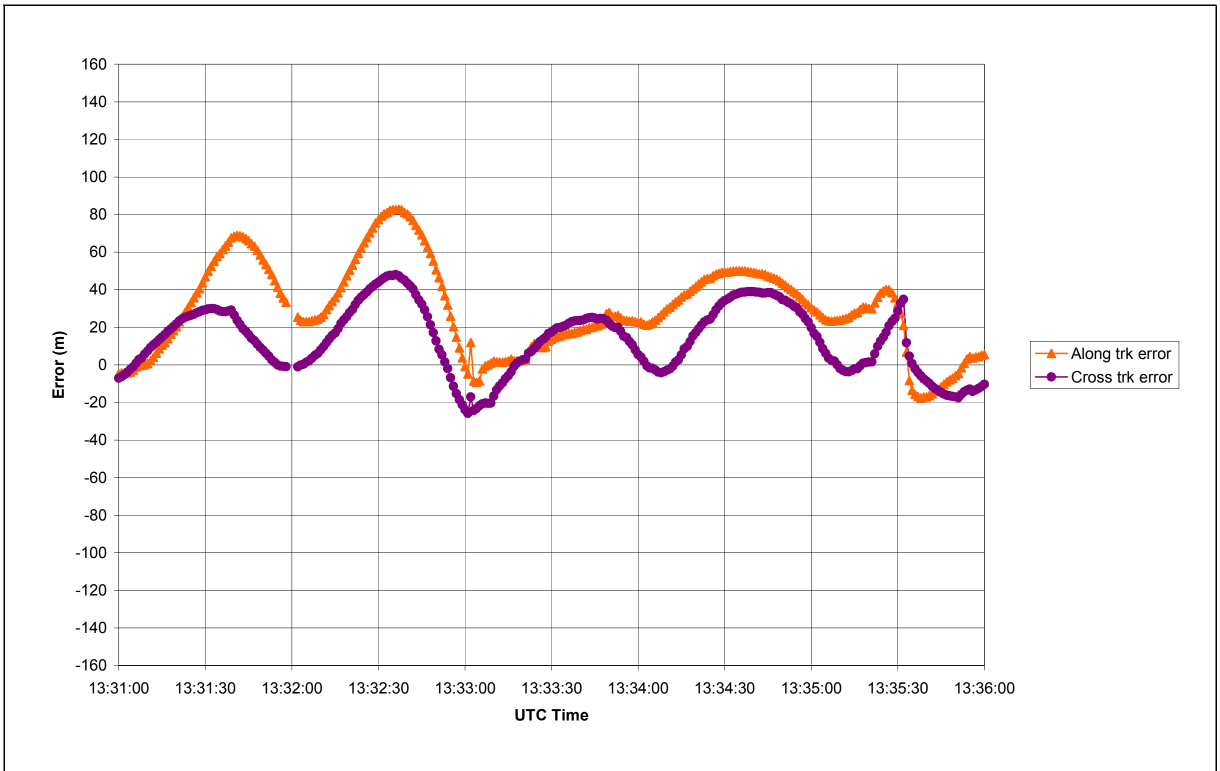

Figure 33 GD2 Along-Track and Cross-Track Position Errors for Window D

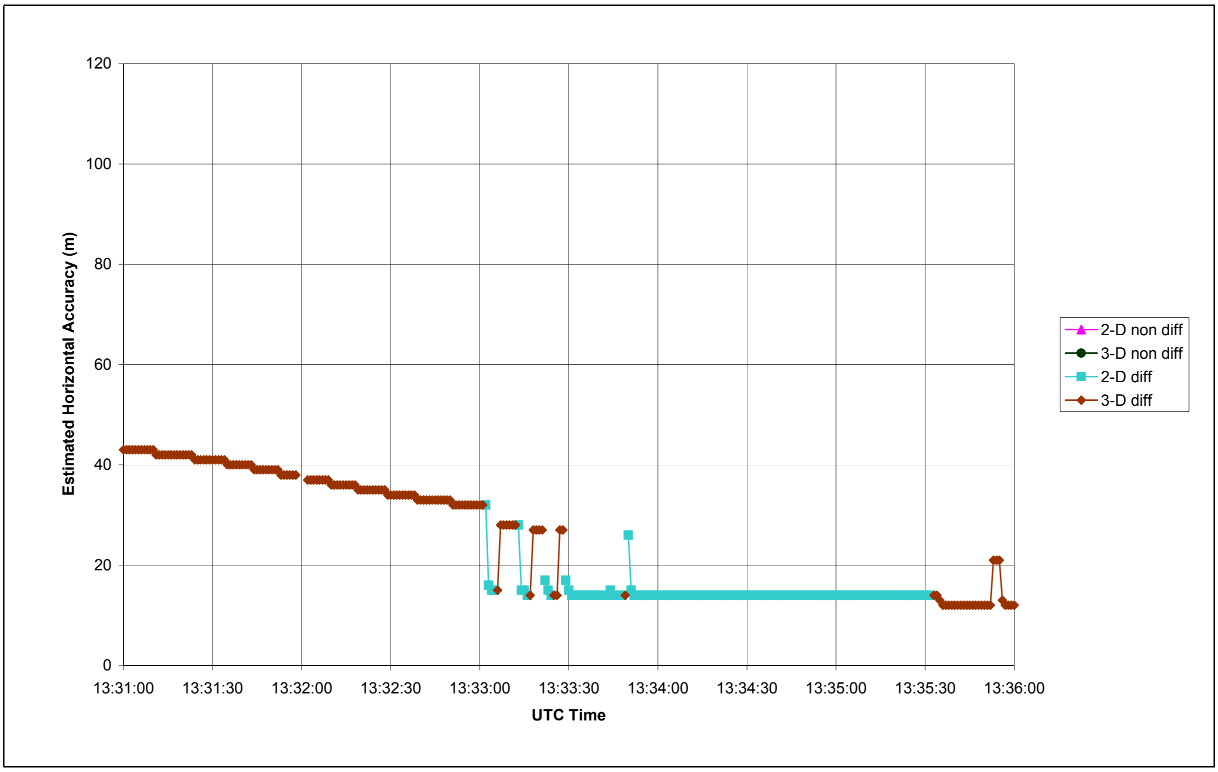

Figure 34 GD2 Estimated Horizontal Accuracy and Operating Mode for Window D

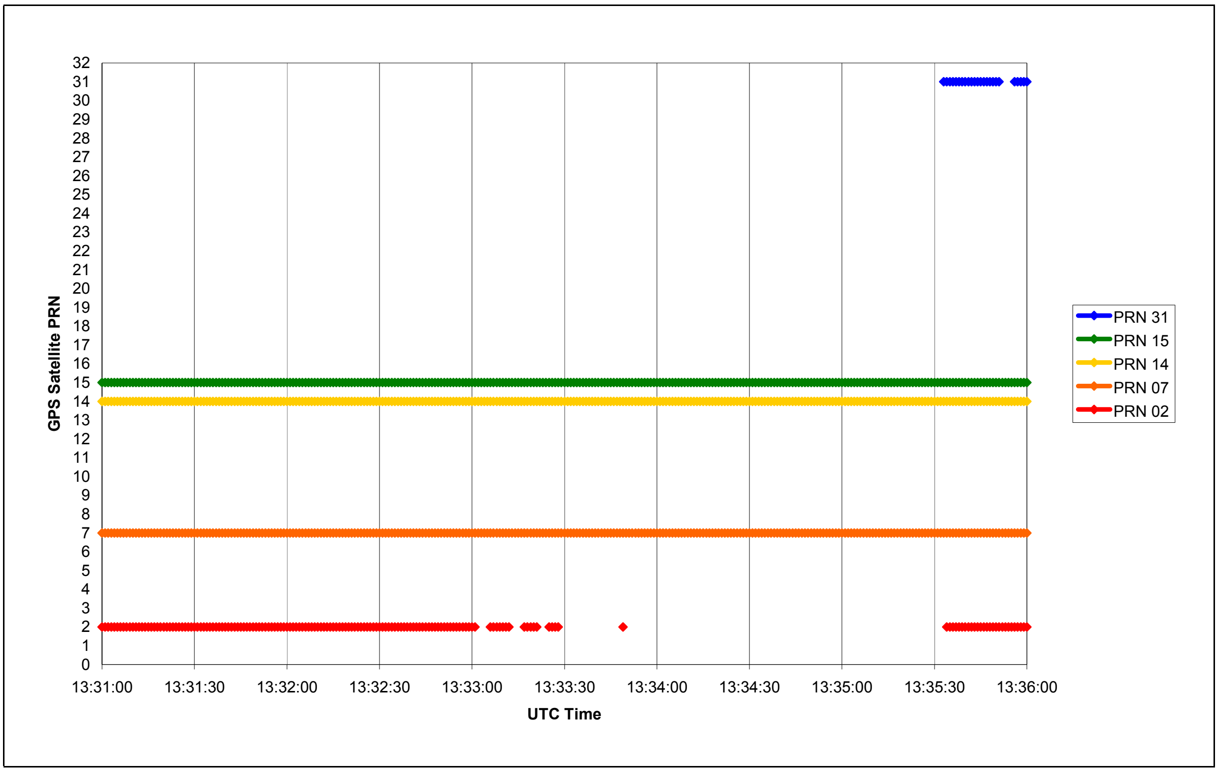

Figure 35 GD2 Satellite Usage for Window D

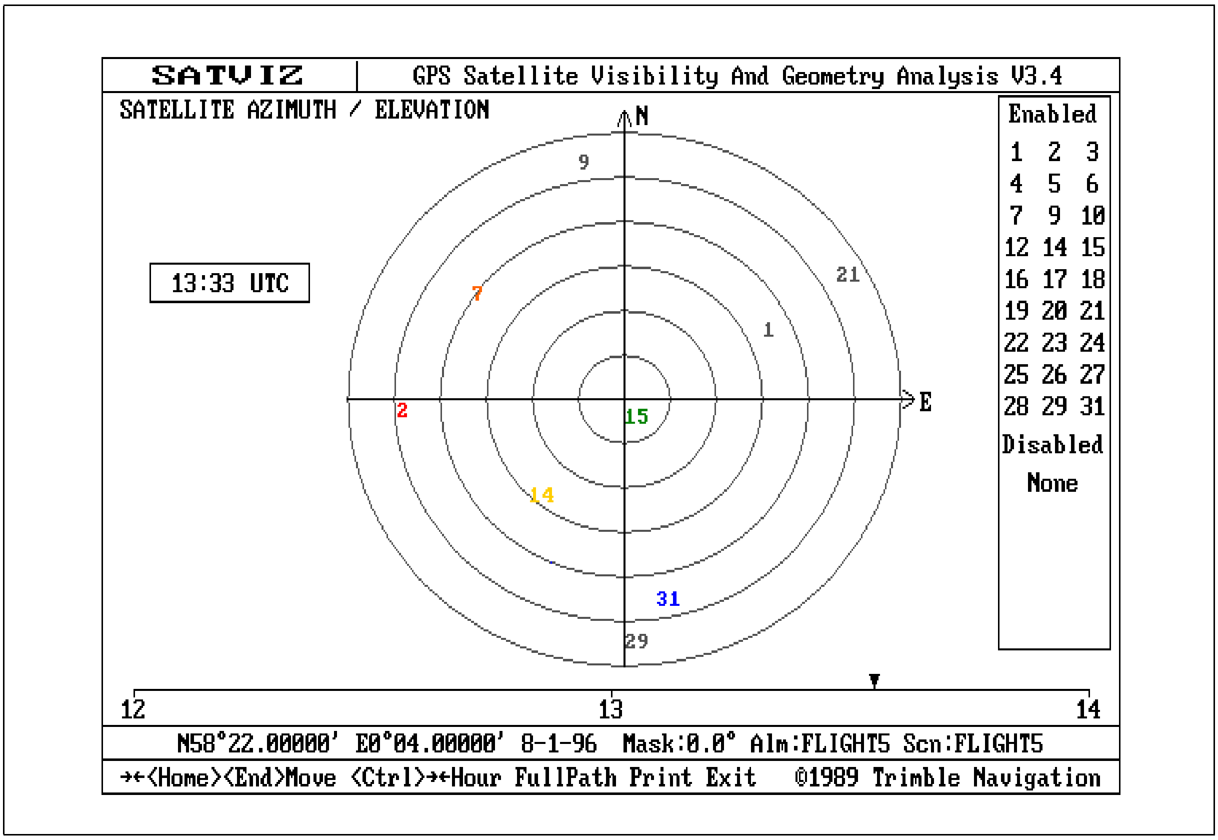

Figure 36 Satellite Azimuth and Elevation for Window D

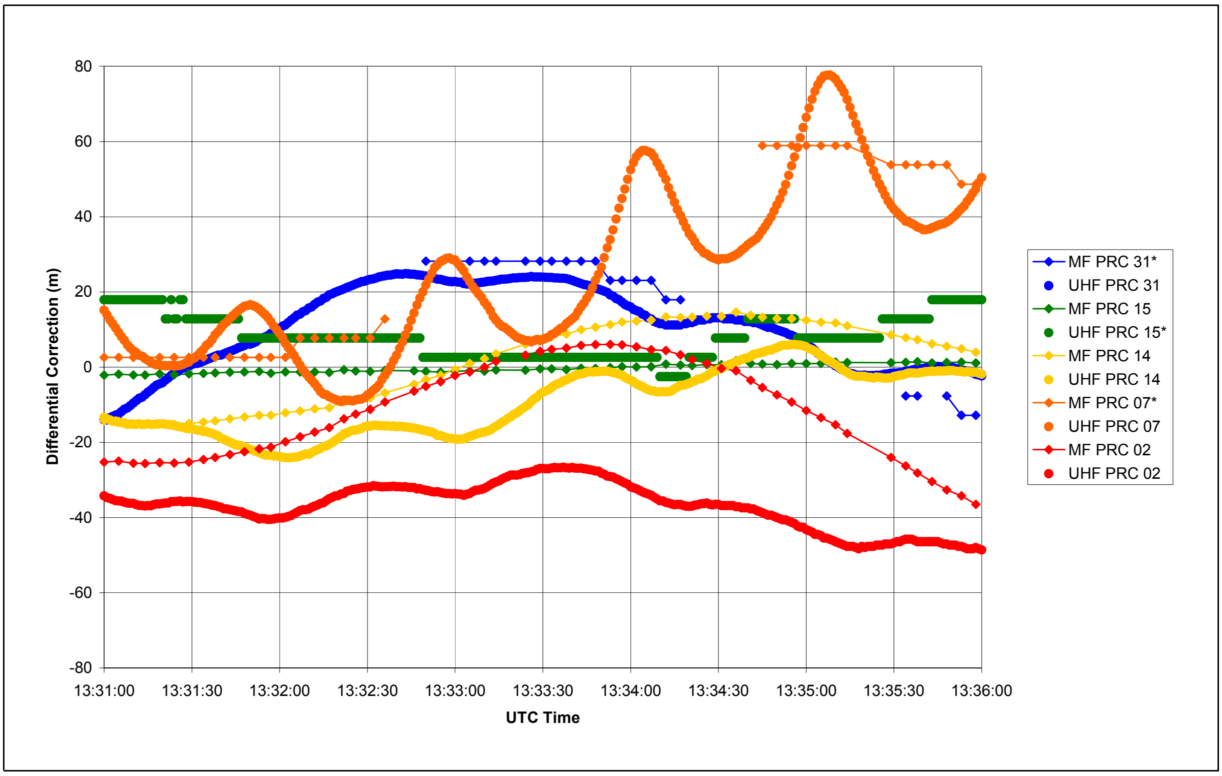

Figure 37 UHF and MF Differential Corrections for Window D
(asterisk denotes reduced precision corrections due to software fault)

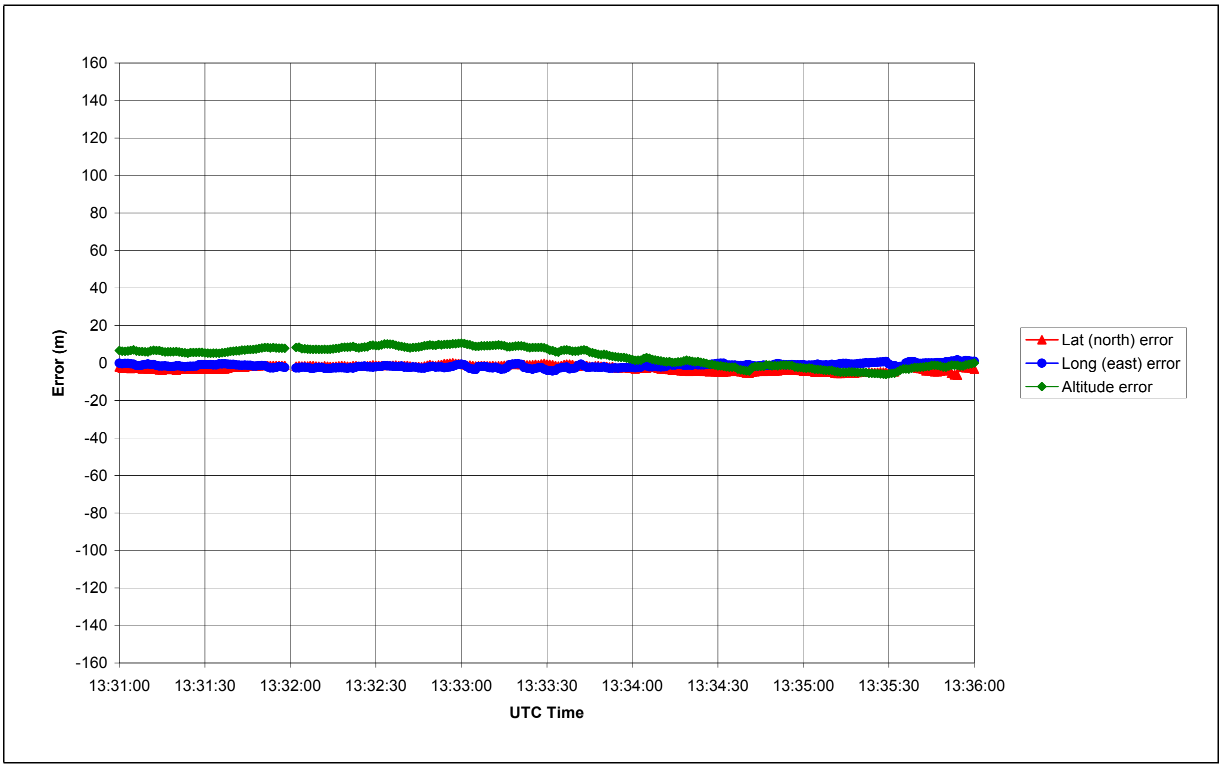

Figure 38 GD1 Position Error Components for Window D

5.5 Time Window E: Flight 5, 13:21 to 13:24 UTC

5.5.1 Manoeuvre details

This time window covers a three minute period commencing at the start of approach number 10 on Flight 5 at the Tartan A platform, and terminating during the go-around manoeuvre which followed the approach. At the start of the window, the aircraft was 1.9nm from the platform descending through 620ft altitude. The aircraft completed the transition to level flight at time 13:22:21 at an altitude of approximately 200ft, and passed within 250m of the platform at time 13:23:03. The go-around was commenced at time 13:23:12 and the aircraft entered a left turn as it climbed. At the end of the time window it was 0.6nm from the platform passing 1000ft.

5.5.2 Truth data availability

Satisfactory truth data was available throughout this time window. The OTruth and PTruth solutions were in agreement to better than ±1m in each axis.

5.5.3 GD2 position accuracy

Figure 39 shows the horizontal and 3-D position errors for the GD2 (UHF-corrected Navstar) receiver. The receiver operating mode was 3-D differentially corrected up until time 13:21:22 where it changed to 2-D differential. During the following two minutes it switched several times between 2-D differential and 3-D differential mode, but was predominantly in 2-D mode. From time 13:23:27 onwards until the end of the window, 3-D mode was regained.

Examination of the plot reveals that the position error increased steadily, beginning around the time that the receiver first switched to 2-D mode. The rate of increase slowed at around time 13:22:10 and the maximum horizontal error was 66.5m at time 13:23:02. After time 13:23:15, the errors reduced rapidly just prior to 3-D operation being regained at 13:23:27.

Figure 40 depicts the individual GD2 position error components. This reveals that the excursion was principally in the longitude and altitude components, with the time variation similar in each case. Comparison with the aircraft altitude profile (from volume 3 of ref 1) confirms evidence for a correlation with the GD1 altitude error component: it is even possible to distinguish the variations in aircraft height, around time 13:22:40, during the level segment. This suggests that these errors were due to the operation of the receiver's internal altitude hold algorithm on entry into 2-D mode, similar to the situation which had been observed during time windows C and D.

5.5.4 GD2 velocity components

Figure 41 shows the velocity components obtained from the GD2 receiver. The velocity components appear to have been stable and continuous with little evidence for any step changes.

5.5.5 GD2 along-track and cross-track position errors

Figure 42 depicts the along-track and cross-track components of the GD2 horizontal error. This reveals that the most of the error was in the along-track sense: up to around +60m, compared to a maximum of just over +30m in the cross-track component. Since the guidance indications during this approach were being derived from the output of the GD1 receiver, the presence of the excursion on the GD2 data was not apparent to the pilots.

5.5.6 GD2 estimated accuracy parameter

The GD2 Estimated Horizontal Accuracy parameter, and the receiver operating mode, are plotted in Figure 43. In general, the value of this parameter was observed to be lower (typically around 15m) during the periods of 2-D differential operation, as opposed to around 30m when operating in 3-D differential mode. The exception was after time 13:23:27, were the value on reversion to 3-D mode remained below 14m. Once again, the value of the horizontal accuracy parameter appears to have borne little relation to the position errors which were actually observed.

5.5.7 GD2 satellite usage

Figure 44 shows the satellites which were employed by the GD2 receiver to generate its position solution. Initially four satellites were available (PRNs 2, 7, 14 and 15), and reversion to 2-D mode correlated with the loss of PRN 2 from the solution. At time 13:23:27, not only did PRN 2 reappear in the solution but a fifth satellite (PRN 31) also became available. This pattern is very similar to that which occurred during time window D. The two approaches in question were undertaken on identical tracks about ten minutes apart.

The satellite elevation and azimuth positions are shown on the polar plot in Figure 45: it can be seen that the positions were very similar to those relating to window D (Figure 36). Once again, it would have been expected that PRN 1 would have been present in the GD2 computation: examination of the tracking data for the truth reference receiver confirms that this satellite was available to the latter unit throughout the time window. As with time window D, the HDOP associated with the four satellite solution with PRN 2, 7, 14 and 15 was considerably higher (between 10.8 and 17.9) than that for a three satellite solution with PRN 2 eliminated (around 2.8). Similarly, the introduction of PRN 31 was demonstrated to allow the HDOP to be reduced to 2.5.

5.5.8 Differential corrections

Figure 46 shows the pseudo range correction (PRC) data transmitted by the UHF system, overlaid with the identical information received from the MF correction system for the same set of satellites. The maximum interval between successive UHF corrections received by the GD2 receiver was 3s. The correction measurements affected by the recording fault described in section 5.3.8 have been denoted by an asterisk.

As with time window D, corrections for PRN 1 were absent from the transmitted UHF differential corrections, which explains why this satellite did not figure in the set used by the GD2 receiver.

There is also evidence for some significant discrepancies (assumed to be due to the degraded L-band reception environment on the platform) between the two sets of corrections. The discrepancies appear to have approached 30m in the case of some of the satellites, although a precise comparison is not possible owing to the effects of the recording fault. There is perhaps less evidence for any oscillatory variation with time.

5.5.9 GD1 operation

Figure 47 depicts the position error components for the GD1 (MF-corrected Navstar) receiver. The largest horizontal error was 5.0m, with the receiver maintaining 3-D differential operation throughout the time period. The maximum interval between MF corrections was 14s. The GD1 receiver was operating using PRNs 1, 7, 14 and 15 throughout, together with either PRN 2 and PRN 31 during periods where one or other of these satellites was also tracked.

No data for the GD3 (MF-corrected Trimble) receiver is available, since the unit was not operational on this flight.

5.5.10 Discussion

This time window proved to be very similar to the previous one (Window D), with many of the same causal factors present. These include:

a) PRN 1 not included in UHF correction messages;

b) Poor geometry in 3-D mode using the available differentially corrected satellites;

c) Aircraft altitude variations when operating in 2-D mode;

d) Errors (assumed to be due to L-band multipath) in the correction measurements.

However, on this occasion there was no evidence for any significant oscillatory variation in the position error components, possibly due to the absence of any such variation from the multipath-corrupted differential corrections. As a result, the position error excursion was similar to that which occurred during the second half of time window D (i.e. an error which gradually increased to a “plateau” corresponding to the period of level flight close to the platform, and then reduced when the go-around manoeuvre was commenced).

Another curiosity is the fact that, although the GD2 receiver returned to 3-D mode on several short occasions during the position error excursion, it would appear that this did not cause the altitude error (and hence the horizontal error) at the start of each subsequent period of 2-D operation to be reset. In this respect, the behaviour differed from that which had been observed during time windows C and D.

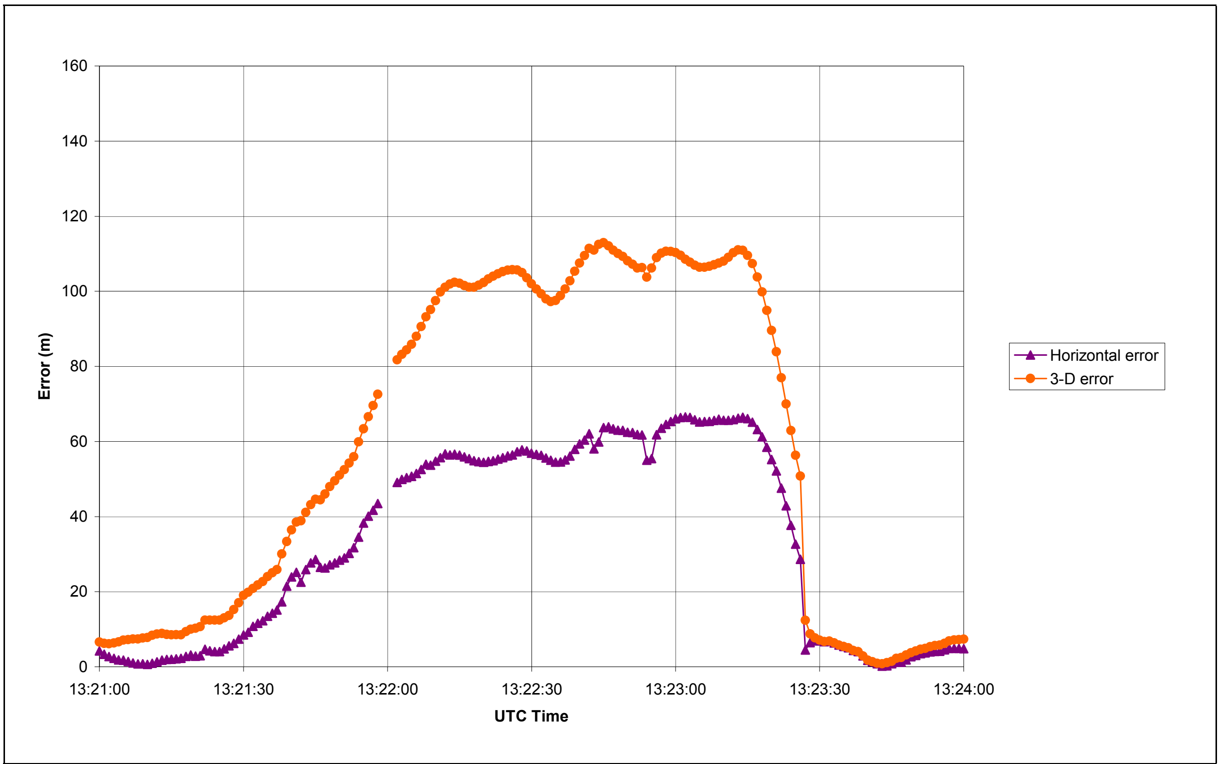

Figure 39 GD2 Position Errors for Window E

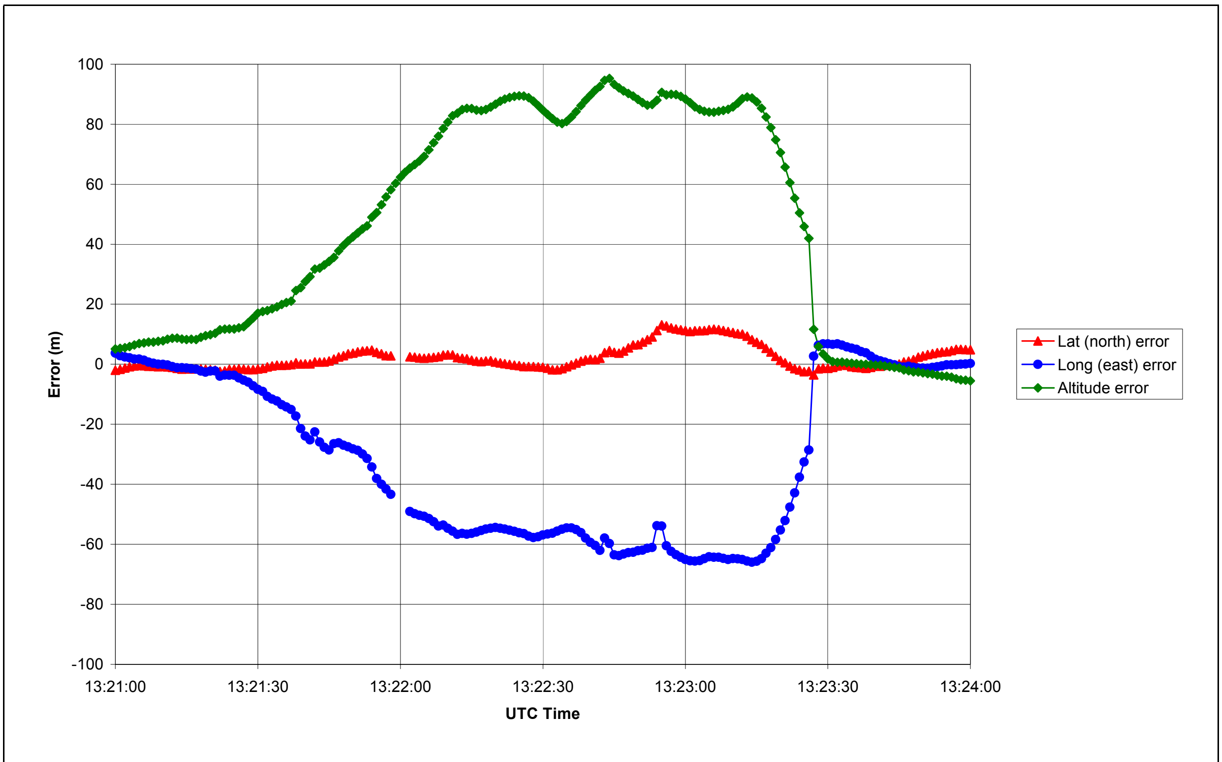

Figure 40 GD2 Position Error Components for Window E

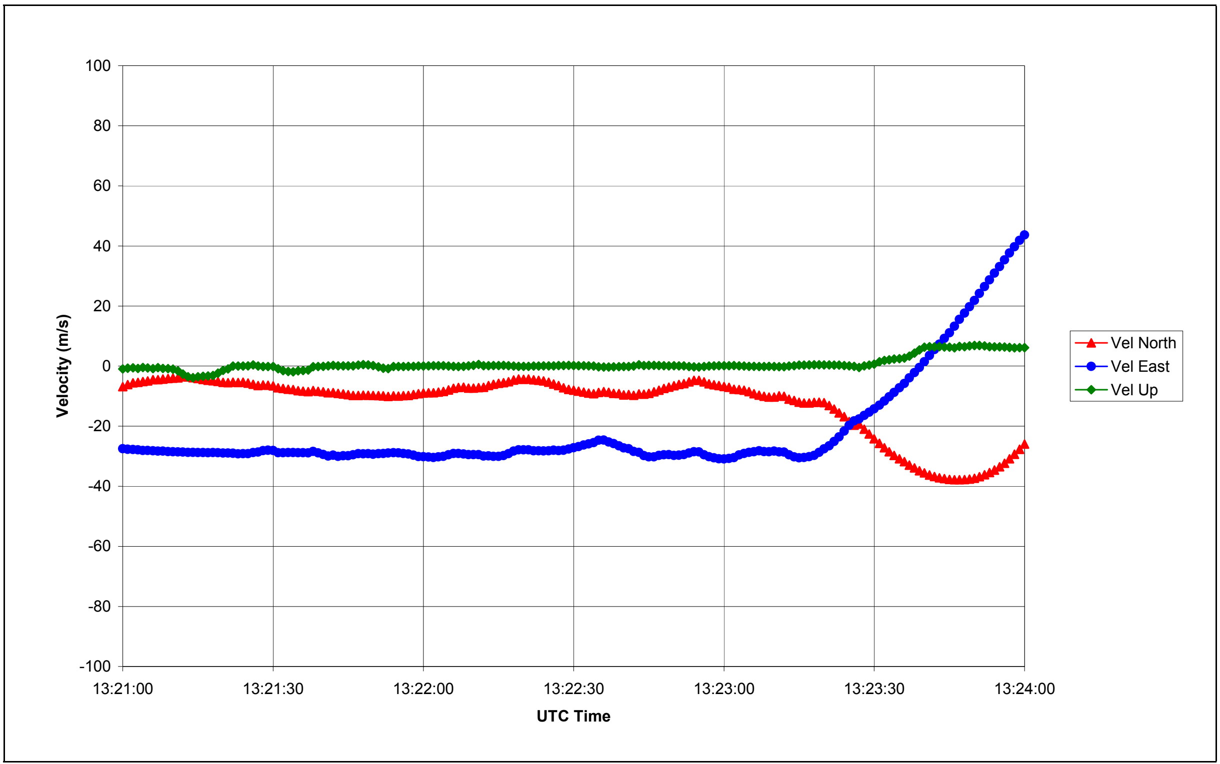

Figure 41 GD2 Velocity Components for Window E

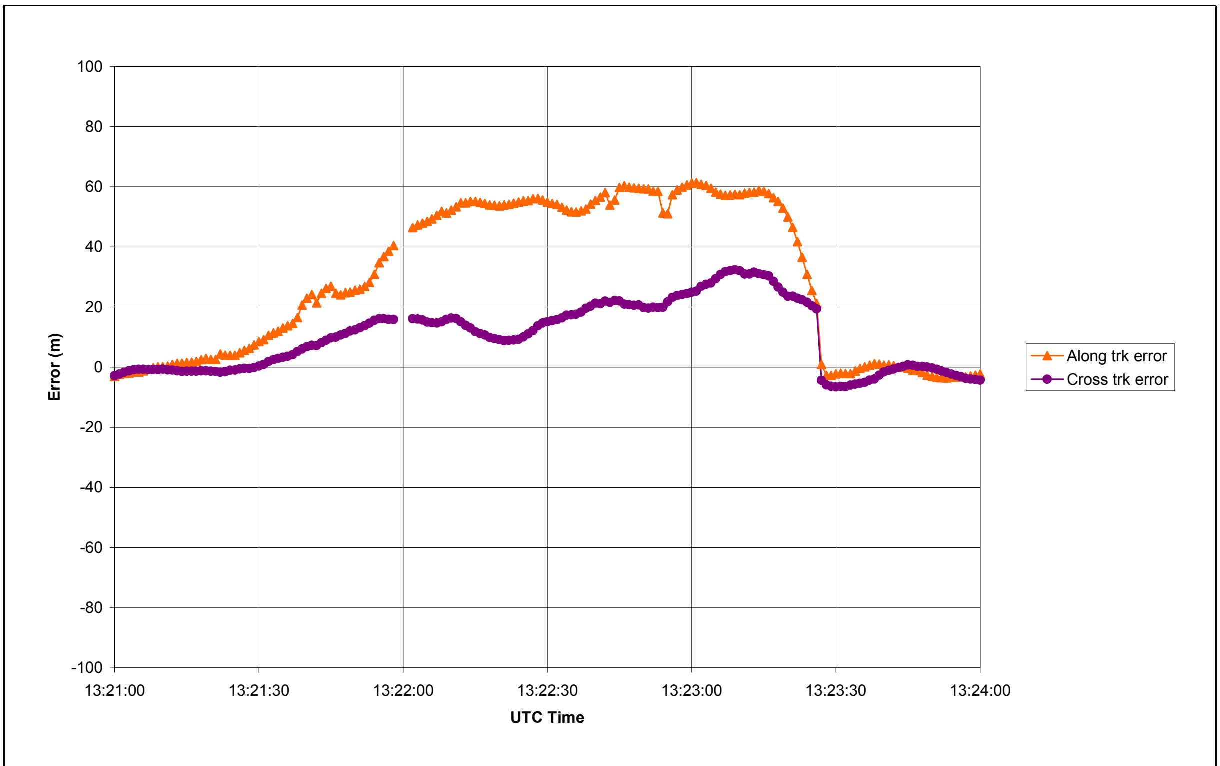

Figure 42 GD2 Along-Track and Cross-Track Position Errors for Window E

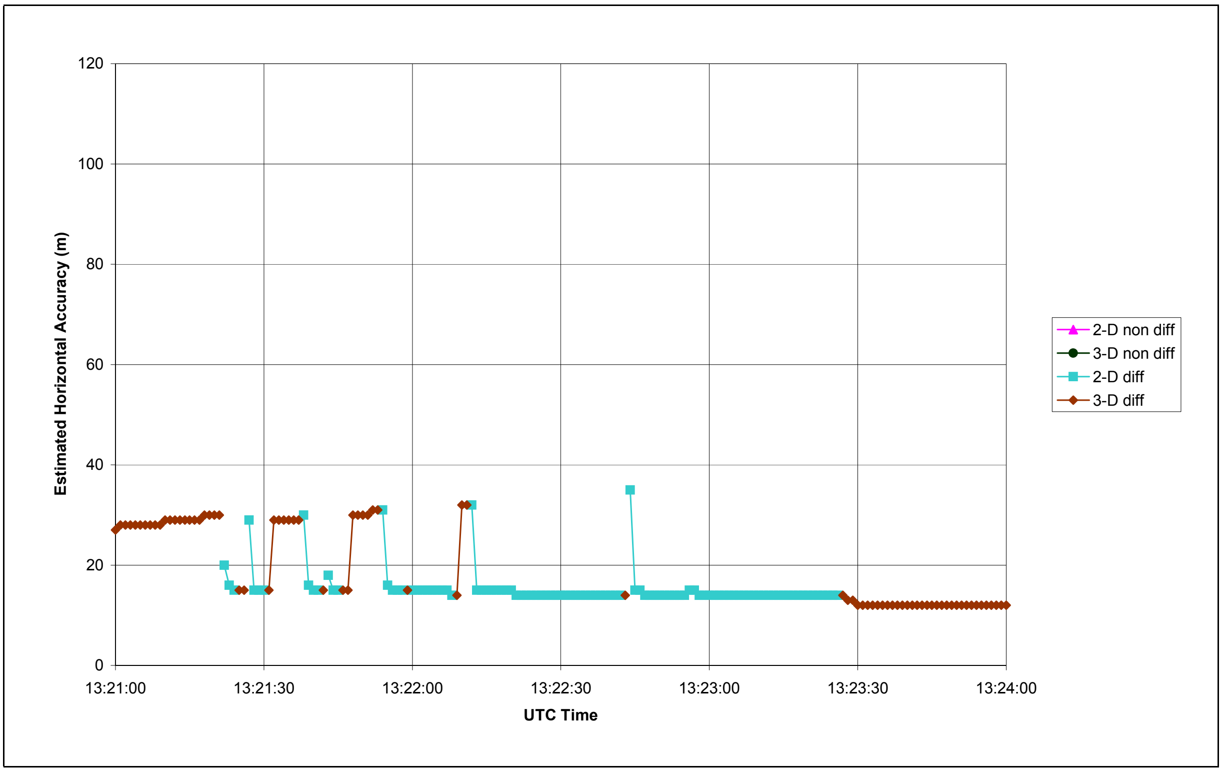

Figure 43 GD2 Estimated Horizontal Accuracy and Operating Mode for Window E

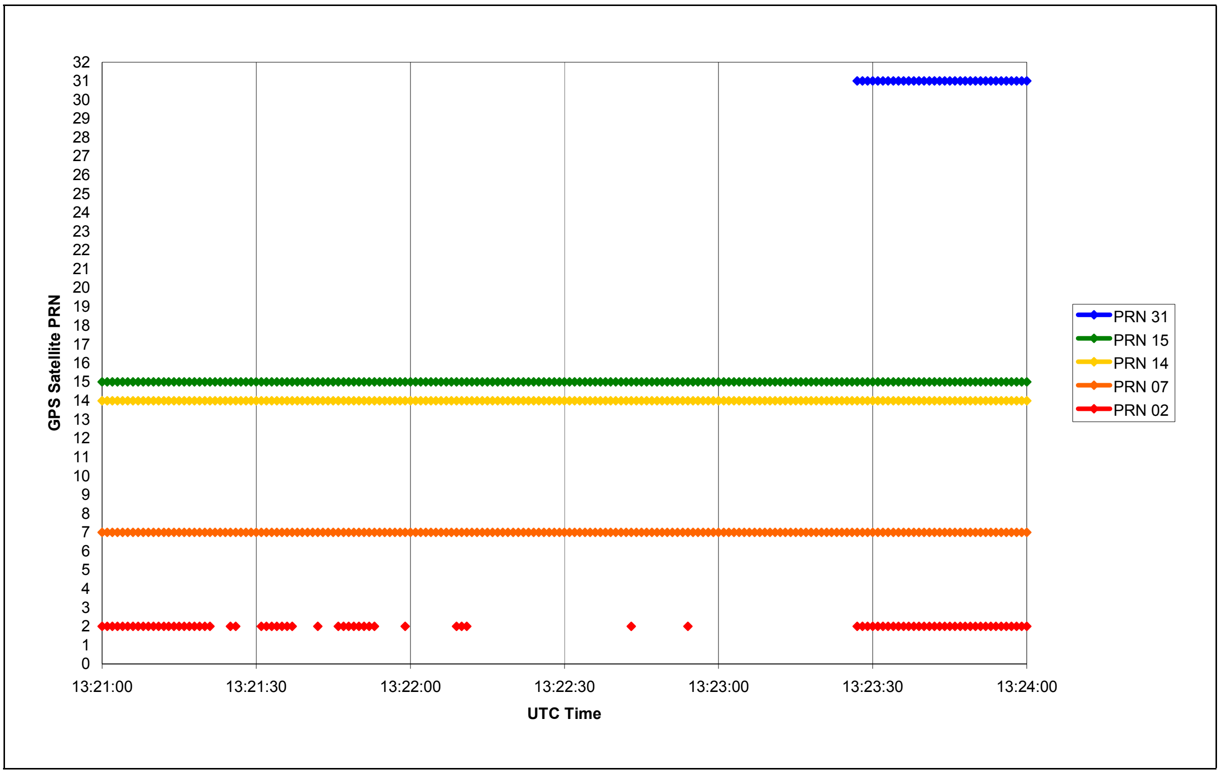

Figure 44 GD2 Satellite Usage for Window E

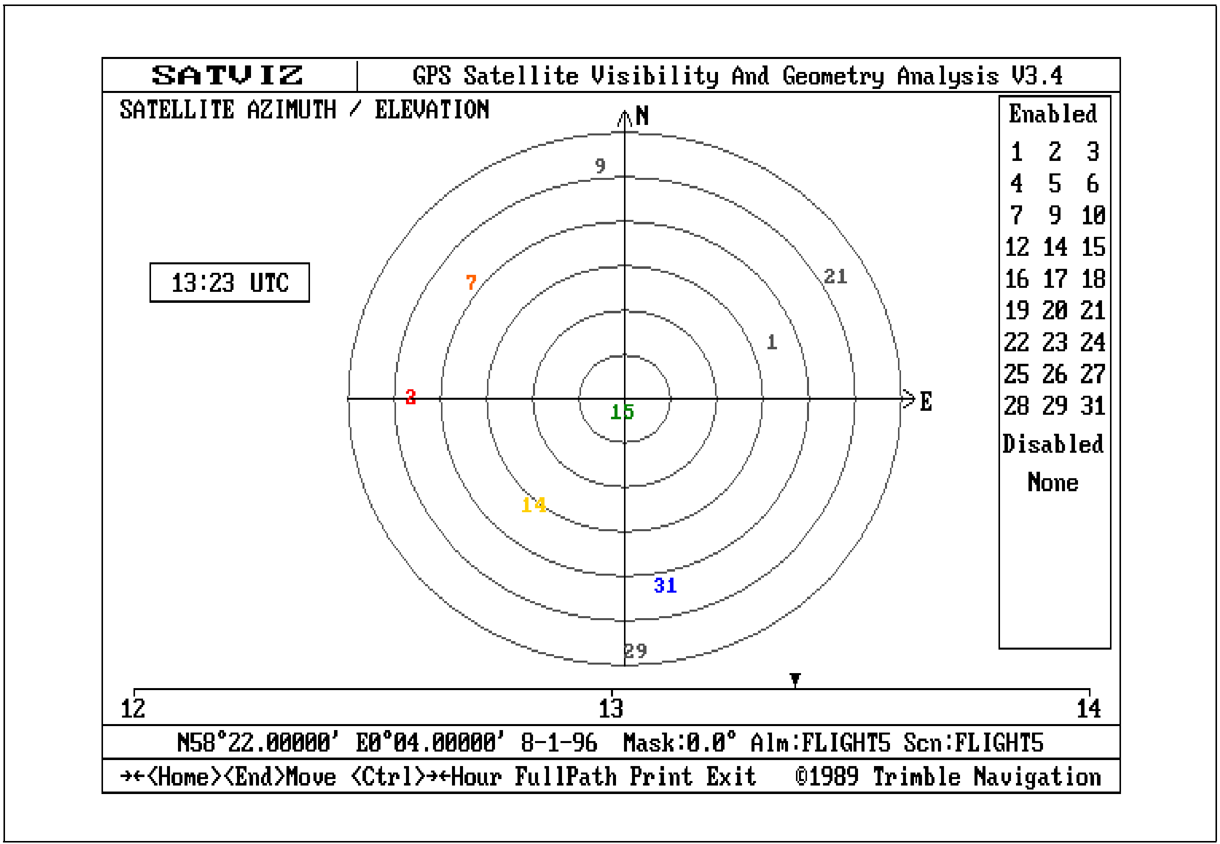

Figure 45 Satellite Azimuth and Elevation for Window E

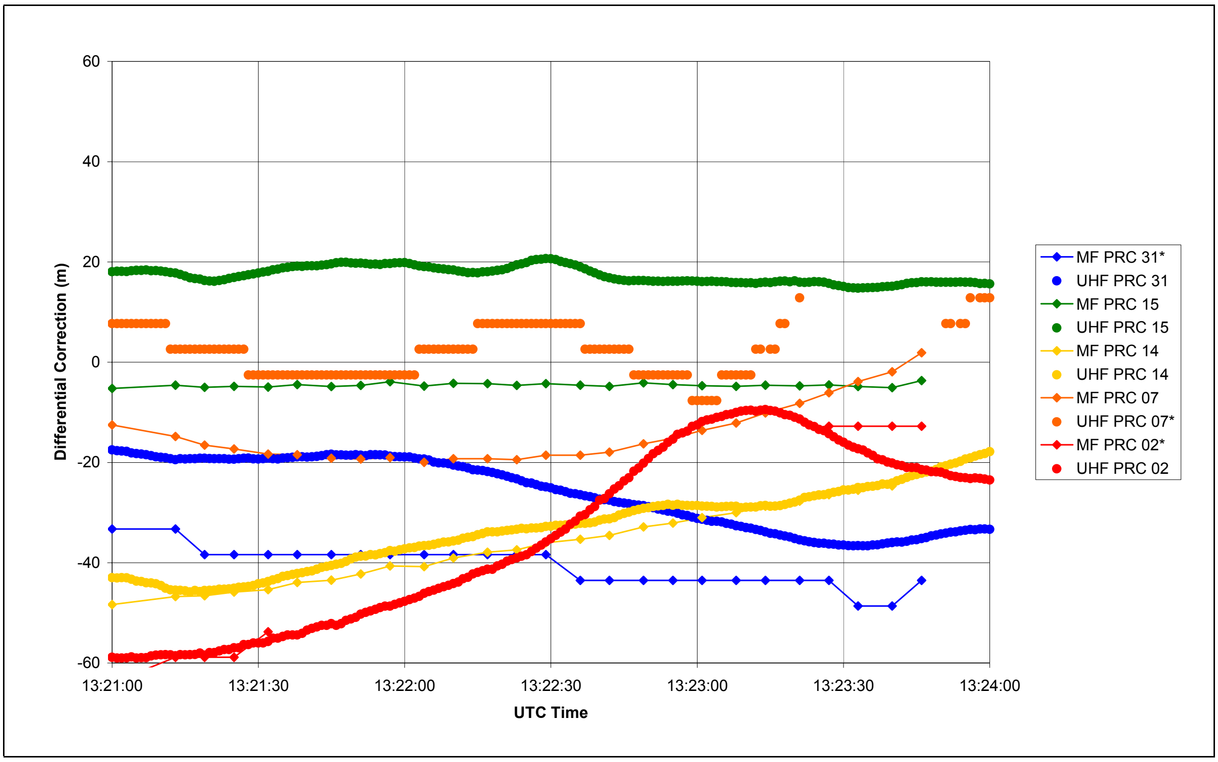

Figure 46 UHF and MF Differential Corrections for Window E
(asterisk denotes reduced precision corrections due to software fault)

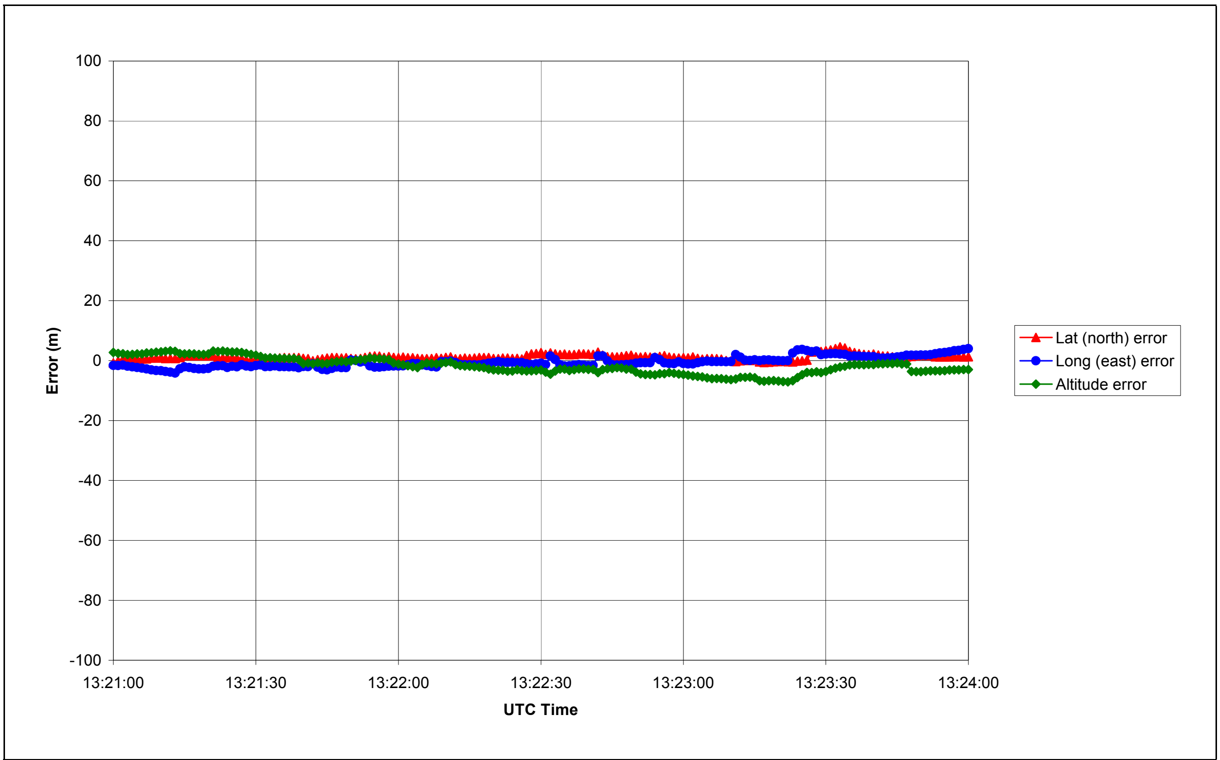

Figure 47 GD1 Position Error Components for Window E

6 Conclusions

Causes of the Large Errors

The following factors were identified as the most likely reasons for the occurrence of the largest DGPS horizontal position errors which were observed during the trials programme.

(1) The differential correction station on the offshore platform sometimes failed to generate measurements for inclusion in the transmitted differential corrections for satellites at a high elevation angle. This resulted in a reduction in the number of differentially corrected pseudorange measurements which were available for inclusion in the airborne receiver's position solution.

(2) Discrepancies of up to 30m were observed to be present between the pseudorange corrections transmitted by the platform reference station, and the corresponding measurements from an onshore station. These discrepancies were attributed to an unsatisfactory L-band reception environment on the offshore platforms.

(3) The airborne receiver had no external source of altitude information to provide assistance when it was operating with only three satellite measurements in 2-D positioning mode. During flight phases where the aircraft was climbing or descending, an invalid assumption that the altitude was constant caused the associated height error to propagate into the horizontal position solution.

(4) A transient position error, lasting less than three seconds, was observed to occur when a DGPS receiver re-entered differential operating mode following a period of non-differential operation.

Since several of these causal factors were sometimes present concurrently, it is not necessarily possible to identify the maximum positioning error which resulted from each one. However, it is possible to state as follows:

(5) A combination of causes (1) and (2) led to cross-track errors of up to 50m and along-track errors in excess of 80m.

(6) A combination of causes (1) and (3), plus a lesser contribution from (2), resulted in cross-track errors of up to 40m, and along-track errors of up to 120m.

(7) Cause (4) resulted in a maximum cross-track error of 17m, and an along-track error of 6m.

Implications of the Results.

(8) The results confirm that the L-band reception environment at the reference station locations on the offshore platforms was unsatisfactory, presumably due to the presence of the platform structure.

Although relocation of the reference antenna to the highest point of the platform would probably have avoided these problems, siting an antenna in this position is probably impractical for operational reasons.

(9) Consideration must be given in the design of any system reliant on differential correction measurements, to ensuring that the possibility of satellite measurements being absent from the transmitted data is taken into account when predicting the overall system accuracy, availability and reliability.

(10) The effect of the problems associated with reversion to 2-D mode could be eliminated either by treating this as an alarm condition, or by introducing an external altitude source such as barometric height. In the latter case, careful consideration would need to be given to the error characteristics, latency, update rate and failure modes of the external source of information.

(11) It was discovered that degraded truth data was probably employed in the calculation of the position errors associated with a portion of the trials dataset, resulting in pessimistic assumptions about the associated DGPS position errors.

(12) In future, it would be preferable if a clearer indication could be provided for the validity, or otherwise, of the truth solution. Ideally a single unambiguous validity flag, associated with a defined probability limit, should be available.

(13) It was demonstrated that the Estimated Horizontal Accuracy parameter output by a DGPS receiver was not capable of reliably identifying the presence of position errors due to all of the different sources described above.

(14) Accuracy estimates of this nature should therefore be treated with extreme caution, particularly when full details of the algorithm employed to calculate the parameter have not been provided by the receiver manufacturer. It is essential that the pilot is not provided with hazardously misleading information.

Future Work

(15) The data presentation developed for this study allowed direct examination of the difference between the contents of the MF and the UHF correction streams. An extension of this technique could usefully be used to investigate the L-band characteristics of any future candidate site for a differential correction station antenna, or the relative merits of two or more such sites. Some form of monitoring programme along these lines would be essential if a platform based location was to be considered.

Study 2

Analysis of MF Correction Datalink Performance

Contents

Study 2 Analysis of MF Correction Datalink Performance

Study 2 Analysis of MF Correction Datalink Performance

1 Introduction

During 1996 a series of flight trials was undertaken in the North Sea to examine the use of Differential Global Positioning System (DGPS) equipment as an approach aid for offshore installations. The flight trials programme was undertaken by the Flight Systems and Measurement Laboratories (now incorporated into Cranfield Aerospace Ltd) of the College of Aeronautics, Cranfield University in the role of prime contractor on behalf of the UK Civil Aviation Authority. The trials were performed using a Sikorsky S76C helicopter operated by Bond Helicopters Ltd. The aircraft was fitted with an experimental DGPS installation which was complemented by additional recording equipment sited at fixed locations.

In the course of seven test flights totalling 36 hours, over 70 predefined manoeuvres were performed at a set of four offshore production platforms with differing topside layouts. At each platform, approach trajectories and guidance presentations based upon the use of DGPS data were evaluated by the trials team which comprised representatives from CAA, Bond and Cranfield.

The trials installation allowed a comparison to be made between alternative sources of differential corrections: Medium Frequency “MF” corrections received from onshore marine radiobeacons; and Ultra High Frequency “UHF” corrections received from a dedicated DGPS reference station positioned on the target platform. It was also possible to perform a comparison between GPS receivers produced by two different manufacturers (Navstar and Trimble).

Post-flight processing of the data recorded during each trial enabled an assessment to be made of the performance of the real-time airborne DGPS equipment. Full details of these results are contained in the project Final Report, which has been published as CAA Paper 2000/5 (ref 1).

As described in volume 2 (section 7.1) of ref 1, various difficulties were encountered with the MF correction system over the course of the trials programme. These problems included the loss of reliable correction signals for portions of the flights, necessitating the selection of alternative beacons in order to re-establish differential GPS operation. Differences were also identified between the performance of the Navstar and Trimble receivers when operating with an unreliable correction source.

Although ref 1 included some suggestions as to why these problems had been experienced with the MF corrections during the trials, the report stated that additional tests would be required in order to explore the situation further. This report presents the results of some additional experimentation which was undertaken using the trials equipment, and of an analysis of the variations which were observed in the performance of the MF correction system during the flight trials.

2 References

1 CAA Paper 2000/5, "DGPS Guidance for Helicopter Approaches to Offshore Platforms", CAA, 2000.

2 "RTCM Recommended Standards for Differential Navstar GPS Service", Radio Technical Commission for Maritime Services, 1998.

3 875-0007-001, "MBX-2/00 Marine Beacon Receiver Programming Manual", Communication Systems International, 1995.

4 ICD-GPS-200, "NAVSTAR GPS Space Segment and Navigation User Interface", ARINC Research Corporation, 1997.

5 "Global Positioning System Standard Positioning Service Signal Specification", DoD, 1995.

6 LNL/C06/058/R, "Further Analysis of the Flight Trials Data – Effects of Satellite Unavailability and Analysis of Receiver Tracking Performance", Lambourne Navigation Limited, 2001.

3 Abbreviations

ARINC	Aeronautical Radio, Inc
Bond	Bond Helicopters Ltd
CAA	UK Civil Aviation Authority
CAe	Cranfield Aerospace Ltd
Cranfield	Cranfield University, Cranfield Aerospace Ltd
dB	Decibel
DGPS	Differential Global Positioning System
DoD	Department of Defense
EMC	Electromagnetic Compatability
ft	Foot
FTE	Flight Test Engineer
GBAS	Ground Based Augmentation System
GD1	Identifier for MF-corrected Navstar GPS Navigation data
GD2	Identifier for UHF-corrected Navstar GPS Navigation data
GD3	Identifier for MF-corrected Trimble GPS Navigation data
GPS	Global Positioning System
H-field	Magnetic field
Hz	Hertz
IMC	Instrument Meteorological Conditions
kt	Knot
L-band	Region of electromagnetic spectrum around 1.5GHz
m	Metre
MF	Medium Frequency
min	Minute
MSK	Minimum Shift Keying
Navstar	Navstar Systems Ltd
NDB	Non-Directional Beacon
nm	Nautical mile
ref	Reference
rms	Root Mean Square
RS232	Electronics Industry Association Recommended Standard 232
RTCM	Radio Technical Commission for Maritime Services
RTCM-SC104	RTCM Special Committee Number 104
s	Second
SA	Selective Availability
SBAS	Satellite Based Augmentation System
SNR	Signal-to-Noise Ratio
SS	Signal Strength
Trimble	Trimble Navigation Ltd
TSO	Technical Standard Order
UHF	Ultra High Frequency
UK	United Kingdom
UTC	Universal Time Co-ordinated
V	Volt
o	Degree
^{o}T	Degrees True
2-D	Two dimensional
3-D	Three dimensional

4 Background

Of the three onshore differential correction stations employed during the trials, two were located in the UK and the third in Norway. The two UK stations, at Girdle Ness and Sumburgh, were operated by a commercial service provider and transmitted encrypted signals to allow costs to be recovered via licence fees (although this arrangement has now been superseded by a free access MF correction service). By contrast, the signals from the Norwegian station at Utsira were unencrypted.

The difficulties experienced on the DGPS flight trials with the MF correction system were described in detail in section 7.1 of volume 2 of ref 1. Successful results were, in general, achieved for much of the trials programme by using the two UK stations. However on each of the final two flights, it proved necessary to switch to the Utsira station (despite its considerably greater range from the platforms) in order to achieve satisfactory differential GPS operation.

The fact was recognised that, in all of these cases, operations were taking place at distances considerably greater than the published operating ranges for the marine NDB transmitters used to broadcast the correction signals. Although satisfactory reports had been received from other aircraft operators at these distances, the possibility could not therefore be discounted that (for example) interference from other sources in the same frequency band had contributed to the reception difficulties. Discussions with the operator of the UK correction stations had revealed that no operational difficulties (such as a reduced power level, or a signal-in-space problem) were known to have existed during the times that the flight trials were undertaken.

The possibility of the reception problems (in particular, those which were observed on the final two flights) having resulted from some kind of deficiency with the receiver equipment, or from its installation on the trials aircraft, could not be discounted, particularly as the trials equipment was being removed between flight trials to allow the aircraft to re-enter normal revenue service. Accordingly, various ground tests were undertaken at the time of the trials in an attempt to determine whether any obvious problem (e.g. a faulty antenna or receiver) could be identified. No conclusive evidence for any problem could be determined, although some of the possibilities (one example would be an intermittent connection which only manifested itself in flight) could not be explored in the time available. The possibility was also identified that the aircraft may have been operating in conditions conducive to the formation of precipitation static, which could have affected the MF antenna installation.

As described in ref 1, the output of the MF datalink receiver was supplied to two of the DGPS units on the trials aircraft: one of the Navstar XR5-M12 receivers (the output from which was termed “GD1”), and the Trimble TNL-2100 receiver (“GD3”). This arrangement allowed both receivers to operate from a common L-band antenna system, and with identical differential corrections sourced from the onshore MF reference station. It therefore allowed a direct comparison to be made between the performance of the two models of receiver.

Loss of the correction signal from the ground station was demonstrated to cause both of these receivers to revert, following a short time-out period, to operation in non-differential mode (i.e. one in which no corrections were applied to any of the satellite measurements). Although the receivers continued to output a navigation solution, reversion to this stand-alone GPS mode resulted in a reduction in positioning accuracy, the effect of which was particularly significant as the Selective Availability (SA) degradation was in operation during the trials.

The analysis of the receiver accuracy contained in ref 1 considered only those time periods where the relevant receivers operated in full differential mode. All of the "non-differential" samples were excluded from the analysis, with the result that the extent of the available statistical data for the MF-corrected receivers was smaller than that for the equivalent UHF-corrected ("GD2") unit.

The Trimble receiver was, for reasons not directly related to the trials themselves, only functional on the final two test flights. These, as it happens, were the flights on which the availability of the MF differential corrections was considerably reduced; however no direct correlation is believed to exist between these two facts.

From a consideration of these final two flights, the ratio of the availability of the differential-corrected position solution from the MF-corrected Navstar receiver, compared to that from the UHF-corrected Navstar receiver (for which, in general, drop-outs into non-differential mode were not observed) was calculated to be 74.7%. By contrast, the corresponding availability ratio for the MF-corrected Navstar solution on the two earlier flights (numbers 4 and 5) for which a full statistical analysis was performed, was calculated to be 96.8%. The significant difference between these two figures highlights the extent to which the operation of the MF-corrected sub-system had been relatively trouble-free on the earlier flights.

A similar calculation for the Trimble receiver on the final two flights results in an availability ratio of 65.1%, confirming the statement in ref 1 that this receiver appeared to require a longer recovery period than the Navstar unit before returning to differential operation following a drop-out. It was postulated in the report that this might be due to differences between the processing performed by the two units' internal software programs.

5 Differential Correction Decoding

5.1 General

The differential correction signals from the onshore MF reference stations were transmitted in the form of a continuous Minimum Shift Keying (MSK) encoded bit stream, at a rate of 100 bits per second, on a side band of the main radiobeacon carrier frequency. The format used to encode the correction signals as a bit stream is defined in ref 2, and is known as the RTCM-SC104 (colloquially, "RTCM") format. The document does not, however, represent a single fixed transmission standard, but merely specifies a series of recommended message formats which can be employed for the transmission of GPS range and range-rate corrections, and associated data, from a reference station to a mobile user.

The function of the MF datalink receiver on the trials aircraft was to demodulate the received MSK sequence, and to output it as a byte-orientated character stream in RS232 format (an electrical interface standard), following the "User Equipment Interface" recommendations defined in ref 2.

In the case of the two UK radiobeacon stations employed during the trials, the use of an encrypted signal at the time of the trials introduces an additional complication. Although no detailed information regarding the nature of the encryption is available, the receiver was assumed to operate to reconstitute the original unencrypted data, with the result that the data appearing at its RS232 output would be indistinguishable (as far as the user equipment was concerned) from that associated with a conventional correction station.

The following statements are included by the manufacturer of the MBX-2 MF datalink receiver in the associated documentation (ref 3):

> "The MBX-2 does not pre-process the output data, and will always output a stream of characters. If the receiver is not locked onto a valid beacon transmitter, these characters will be random; when locked, the output will be valid correction errors in RTCM format.
>
> Internally, the MBX-2 performs complete RTCM messages synchronization and parity checking in order to provide a true "Signal Lock" indication, and to develop message quality statistics as auxiliary output messages."

The "Signal Lock" output referred to in the above text is an indicator on the receiver front panel, which was not monitored during the flight trials owing to the fact that the trials equipment was remotely located in the aircraft baggage bay.

The "message quality statistics" are assumed to refer to a series of parameters which can be accessed via the receiver's front panel user interface (again, these could not be accessed during the trials), but which in some cases could also be accessed remotely via an RS232 interface. Two of these parameters (Signal Strength and Signal-to-Noise Ratio) were recorded during the test flights and an analysis of these recordings is included in section 6 of this report.

The reference to "parity checking" relates to the use of a strong parity algorithm as an integral part of the RTCM standard. The parity algorithm in question is identical to that employed for the GPS transmissions themselves. To quote from ref 2:

> "A strong parity algorithm is required to detect errors in the data, preventing the use of erroneous corrections that could affect user safety."
>
> "The GPS parity algorithm is a known and proven algorithm with which the users are familiar."

Details of the GPS parity algorithm are contained in, for example, ref 4 and ref 5. A technique known as a Hamming Code is employed which, in theory, allows a one-bit error within a word[1] to be detected and corrected; and a two-bit error within a word to be detected (but not corrected). Both ref 4 and ref 5 include an example flow chart which is, presumably, intended to enable a receiver designer to implement the necessary parity algorithm. However, in both cases the flow chart only provides for the detection of an error within a word (there are two outcomes: "Parity check fails" and "Parity check passes"), without any mention of the steps which would be necessary in order to correct a one-bit error.

In the absence of a satisfactory input signal to the MF datalink receiver (whether due to low signal level, interference, installation-dependent factors, or for any other reason); the manufacturer's statement indicates that the receiver would continue to output a random stream of characters. In this situation, reliance would be placed upon the correct operation of the parity detection algorithm in the subsequent user equipment (in this case, the Navstar and Trimble DGPS receivers) to reject the incoming data.

In marginal reception conditions, it would be expected that the output of the MF receiver would consist of a mixture of valid RTCM data and of random noise bits (which might periodically include "missing bit" or "extra bit" errors). Again, reliance would be placed upon the parity algorithm to reject and/or correct the incoming data stream in any subsequent equipment.

1. Each word in an RTCM message is 30 bits long, including six parity bits. The length of a message is variable, but for example a Type 1 "Differential GPS Corrections" message containing data for nine satellites would be seventeen words (510 bits) long.

Considering the case of a parity failure which affects only one particular word within an RTCM message, then no guidance appears to be contained in any of the documents referred to above as to how this situation should be dealt with. Instead, it can only be assumed that these matters are expected to be left to the discretion of the systems engineer and/or software designer. For example, one approach to this problem might be for the software to attempt to decode and make use of the data contained within the remainder of the message in question (and to ignore only the words which are affected by the parity failure). This technique is potentially fraught with pitfalls, particularly if the software has been poorly specified, coded or tested. The more cautious approach would be for the software to ignore the whole of the message affected by the parity error, and to await the beginning of the next message before resuming processing. However, owing to the variable message length, it might not always be possible for the receiver to determine where the next message began, particularly if multiple bit errors were present in the received data.

5.2 Bench Tests

In the absence of any specific information on the subject from the respective manufacturers, it was considered useful to explore how the Navstar and Trimble receivers would respond when corrupted data was supplied at their RTCM correction inputs. Accordingly, a series of bench tests was undertaken during August 2001 using the DGPS Trials Pallet which had been installed in the aircraft.

To provide a controlled source of differential correction messages, the MF datalink receiver was replaced by a Navstar GPS receiver which was configured to output RTCM Type 1 messages. This arrangement was very similar to that which had been employed to generate the UHF corrections for the flight trials. The output from this RTCM source was routed to the inputs of the GD1 (Navstar) and GD3 (Trimble) receivers, thus replicating precisely the arrangement which had been employed on the flight trials for the MF corrections. However, a facility was added to enable the content of the correction messages to be modified under controlled conditions using a microprocessor-based interface unit, which allowed simulated “errors” to be introduced into the data stream.

The purpose of the trial was to explore the ability of the Navstar and Trimble receivers to operate in the presence of corrupted differential corrections and, in particular, to determine the conditions under which they reverted to the non-differential operating mode. No particular attention was paid to the accuracy of their navigation solution outputs, other than to perform a simple “reasonableness check” to confirm that they were generating an appropriate latitude/longitude position.

Both the Navstar and Trimble receivers provided a continuous indication of the current operating mode, both in terms of the fix type (no fix, 2-D, or 3-D), and of the DGPS mode (non-differential, or differential). During the flight trials, this information had been recorded via the “gpsfix” parameter (Table 10 in volume 1 of ref 1 refers). For the Navstar receiver, additional information regarding the status of the RTCM correction input was available in the form of the “difage” parameter, which indicated (to a resolution of one second) the length of time since the last valid correction message was received. Unfortunately this information was not available from the Trimble receiver, for which the only available indication of a loss of differential input was reversion to non-differential operating mode on expiry of the time-out period.

With all of the GPS receivers arranged to operate from a common roof-mounted antenna (which was known to provide a satisfactory L-band signal for bench testing), and the equipment configured as described above, the tests described in sections 5.2.1 to 5.2.8 were undertaken.

5.2.1 Normal operation

The reference receiver was configured to output a new Type 1 correction message every second, and the operation of the Navstar and Trimble units was monitored for several minutes to confirm that they were operating correctly in 3-D differential mode.

The value of the "difage" indication from the Navstar receiver was observed to be permanently set at one, confirming that a new correction message was being received and decoded every second.

5.2.2 No L-band signal to Navstar receiver

The antenna input to the Navstar receiver was temporarily disconnected.

This resulted in the receiver losing lock on all satellites and reverting to the "no fix" mode, but with the "difage" parameter continuing to indicate a value of one second. This test confirmed that the Navstar "difage" parameter (indicating the number of seconds since the last correction message was received) still provided the correct indications, even if the receiver was incapable of navigating due to loss of the L-band satellite signal(s). It was therefore assumed that the internal process which decoded the correction messages operated independently of the satellite tracking sub-system.

5.2.3 Reduced correction message rate

The rate at which the reference receiver transmitted a Type 1 correction was temporarily reduced, from once per second to once every five seconds.

Both the Navstar and Trimble units continued to operate in 3-D differential mode. The Navstar "difage" parameter was observed to count up on successive seconds from one to five, before returning to one on the next update following correction transmission. This test provided confirmation that the operation of the Navstar "difage" parameter was correctly understood.

As described in ref 1 (volume 2 section 7), the MF corrections employed on the flight trials were only updated approximately every five or six seconds. This was a function of the time necessary to transmit a complete Type 1 message at the low data rate (100 bits per second) employed for the MF signals, the precise update rate varying according to the number of satellites for which corrections were present in the RTCM messages. In spite of the fact that it would have been more representative to retain the 5 second update rate for the remainder of the bench tests, it was decided to revert to the use of a 1Hz update rate as this enabled closer monitoring of the "difage" parameter in some of the later tests.

5.2.4 Corrections disabled

The transmission of the Type 1 corrections was temporarily disabled.

The Navstar "difage" parameter was observed to count upwards from 1 at each second. After approximately 30 seconds, both the Navstar and Trimble receivers were observed to revert to non-differential mode. This test was repeated several times to determine whether both receivers were applying a consistent time-out period. This revealed that the elapsed time between loss of corrections, and reversion to non-differential mode, was 30±1 seconds for both units.

5.2.5 Single one-bit errors.

The effect upon receiver operation of a simulated single bit error within a correction message was investigated.

This test was performed by selecting, at random, a particular character byte within a message, and a bit position within that byte. The affected bit was then inverted (to simulate a reception error) before the resulting message was transmitted to the

Navstar and Trimble receivers. Subsquent correction messages in the sequence were not modified.

From a series of 50 such tests, it was observed that on each occasion, the Navstar "difage" parameter increased from one to two, for a single second immediately following transmission of the erroneous message. Both the Navstar and Trimble receivers remained in differential mode throughout.

The increase in the "difage" parameter suggests that the Navstar receiver had rejected the incoming message rather than correcting the erroneous bit. There was no reversion to stand-alone navigation mode, presumably since the receiver was capable of retaining the previous correction information for up to 30s.

No conclusions could be drawn regarding the operation of the Trimble receiver, since it does not generate an equivalent to the "difage" parameter.

5.2.6 **Single two-bit errors.**

The previous test was repeated, but with errors introduced into two randomly selected bits within the affected byte.

The results, from a series of 25 tests, were identical to those obtained when a one-bit error was employed. The only conclusion which could be drawn from this test was that the Navstar receiver appeared to also reject messages which contained two erroneous data bits within a word.

5.2.7 **Multiple one-bit errors**

The simulated one-bit error test was repeated, but with the error introduced into a randomly selected byte within each successive Type 1 message.

From a series of tests, it was observed that, as soon as the error injection was enabled, the Navstar "difage" parameter began counting upwards from one at each second. After thirty seconds, both the Navstar and Trimble receivers reverted to non-differential mode. When the error injection was disabled, both receivers were observed to return to differential mode (and the Navstar "difage" parameter to one) within two seconds.

It was concluded that neither the Navstar nor the Trimble receiver were making use of the error-correcting capabilities of the Hamming parity algorithm. Instead, both units appeared to be rejecting the complete RTCM message in response to a single erroneous bit within any of the data words.

5.2.8 **Multiple "extra bit" errors**

Owing to the character-oriented nature of the RS232 format, it was not readily possible to introduce only a single "extra bit" error (e.g. to simulate a temporary loss of bit synchronisation within the MF receiver) into the RTCM message stream.

Instead, a null character byte was inserted at a randomly chosen position within each Type 1 message, thereby injecting eight additional data bits at this point. Although not particularly representative of the type of reception problem which might occur in reality, it was felt that this would at least allow a test to be performed of the receivers' re-synchronisation capabilities.

When the error injection was enabled, the "difage" parameter began counting upwards and both receivers reverted to non-differential mode after thirty seconds, confirming that the incoming messages were being rejected. Once the error injection was disabled, the Navstar receiver reverted to differential operation within between one and three seconds, similar to the results obtained following the multiple one-bit error test.

For the Trimble receiver, however, the characteristics of this time delay were very different. From a series of thirteen separate tests, the times taken to regain differential operation were as follows:

12s, 48s, 5s, 2s, 13s, 11s, 1s, 2s, 30s, 22s, 6s, 2s, 3s (average time = 12s)

The implication of this result is that, if the RTCM correction messages being supplied to the Trimble receiver have been subjected to certain forms of data corruption, it would appear to be possible for the unit to enter a state in which it is incapable of returning to differential operation for a considerable time period (e.g. of up to nearly one minute) even if a series of uncorrupted correction messages has subsequently been provided to it.

Owing to the importance of this result, this test was repeated with the correction update rate reduced to one every five seconds, so as to be more representative of the arrangement used during the flight trials. From a series of nine separate tests, the times taken for the Trimble receiver to regain differential operation were:

11s, 5s, 53s, 5s, 5s, 10s, 9s, 17s, 5s (average time = 13s)

It was noted that another occurrence of a delay approaching one minute was observed. The reduction in the incoming message rate by a factor of five did not, however, appear to result in any significant increase in the total delays which were observed.

The corresponding results for the Navstar receiver revealed that, in each case, differential operation was regained within five seconds (i.e. following receipt of the next uncorrupted RTCM message). The contrast between these two sets of results provides confirmation for the hypothesis in ref 1 to the effect that the two receivers processed the incoming RTCM corrections in different ways, and that this may well provide the explanation as to why the availability of a differentially corrected position solution was considerably lower in the case of the Trimble unit.

5.3 Manufacturer’s Comments

Enquiries of the manufacturers of the Navstar receiver (or strictly, of their successor company which has inherited the responsibility for supporting their range of GPS equipment) confirmed that the decision not to implement any error-correcting scheme at the RTCM correction input had been intentional. It was also stated that this facility was not utilised when decoding the GPS satellite transmissions themselves.

The manufacturer’s response indicated that a single bit error was unlikely to occur in practice, unlike a multiple burst of errors would could be caused by noise or a jamming source. It was not clear, however, whether or not a separate analysis had been undertaken to consider the bit error characteristics at both the GPS transmission frequency (L-band), and at frequencies commonly used for differential corrections (in this case MF, although the receiver is compatible with any RTCM format correction source).

It did not prove possible to obtain a response to a similar query from the manufacturers of the Trimble equipment.

6 MF Datalink Performance

This section presents the results of an analysis of the performance of the MF correction system as a function of the aircraft’s range from the transmitting station.

The data which was recorded on the aircraft during the flight trials is defined in ref 1 (volume 1, section 7), and included an “RTCM Correction Message” derived from the

MF system. Unfortunately, this information did not comprise the raw RS232 data as output by the MF datalink receiver, but instead resulted from the application of a decoding algorithm (similar to that which would be employed within a GPS receiver) to the RTCM format data. The performance of the algorithm and, in particular, its behaviour in response to the presence of corrupted data, would therefore be dependent upon the robustness of the underlying software design.

Although all of the details of this decoding algorithm and of the structure of the software routines within which it was implemented were available, it was felt that an analysis of the reliability which with which the MF corrections had been received should (in the absence of any “raw data” recording) ideally be based upon the use of a proven GPS receiver to decode the corrections. As demonstrated by the results presented in section 5, the bench tests had shown that the Navstar receiver appeared to provide a satisfactory indication (via the “difage” parameter) as to when an uncorrupted RTCM message was received. Accordingly, it was decided to employ the “difage” data contained in the GD1 recorded data from the MF-corrected Navstar receiver to provide a measure of how accurately the corrections had been received on the trials aircraft.

This analysis was undertaken by dividing the recorded trials data into a series of one-minute periods. For each period, the number of occasions that the “difage” parameter was reset to a lower value (signalling the receipt of a new Type 1 message) was determined using an automated software routine. This provided an indication (shown as “Messages/min” on the subsequent plots) as to the number of error-free RTCM correction messages received on the aircraft within that one minute period.

Unfortunately the interval between successive Type 1 messages transmitted by the reference station was not constant throughout the trials, but varied according to the number of satellites for which corrections were included in each message. For example, if corrections for ten satellites were being generated, then the interval between messages would theoretically be 5.7s which corresponds to 10.5 messages/min. If corrections were only available for six satellites, then the interval would be 3.6s, or 16.7 messages/min. These figures correspond to the minimum and maximum number of satellites present in the correction messages, for the approaches described in ref 6.

Although the use of some form of normalising technique to adjust the “Messages/min” results for the effect of the number of satellites in the Type 1 message was considered, it was decided that it would be preferable to employ the non-normalised figures. In any event, the most significant time periods would be those during which a low number of correction messages were received, and these are the values for which any normalisation technique would have the least impact.

Instead, a decision was taken to employ a threshold value of seven messages per minute, below which the performance of the MF datalink could be considered to be “unacceptable”. This value was determined by assuming an average value of eight for the number of satellites in each correction message. This implied a theoretical transmission rate of 12.5 messages/min. Since it was known that both receivers were capable of operating for up to thirty seconds with stale correction data, it could be assumed that the reception of seven or more uncorrupted messages in each minute (minimum transmission time 34s) would enable the receivers to operate continuously in differential mode.

Two other performance parameters, obtained directly from the MF datalink receiver, were recorded during the flight trials. These comprised Signal Strength (“SS”) and Signal-to-Noise Ratio (“SNR”).

It was stated in ref 1 that the units for the Signal Strength parameter were "dBµV/m" (i.e. decibels relative to a field strength of 1 microvolt per metre), this information having been obtained from the receiver documentation (ref 3). However, it was subsequently pointed out that the receiver would only be capable of measuring the field strength if the "effective height" of the antenna was known. A series of e-mail communications with the receiver manufacturer revealed that the units stated in the manual are incorrect, and that the Signal Strength parameter actually refers to the voltage at the receiver input (i.e. the units quoted in ref 1 should have been "dBµV"). The manufacturer's representative recommended that this parameter should only be employed in a relative sense (i.e. the higher the figure, the stronger the input signal).

No information was available from the manufacturer as to how the Signal-to-Noise Ratio (SNR) parameter was computed. The manufacturer also pointed out that the H-field loop antenna employed on the trials aircraft has been obsolete for several years and that, as a consequence, no additional information on its characteristics could be provided.

In spite of this lack of detailed information on the SS and SNR parameters, it was considered that they might well prove a useful additional source of data. Accordingly, the average SS and SNR was calculated for each of the one minute periods described above, and this information was plotted along with the "Messages/min" statistic as a function of the range in nautical miles from the MF transmitting station using a common vertical axis.

The aim has been to present all of the recorded data from each of the flight trials (with the exception of Flight 1, which was beset with data recording problems), including those periods where the aircraft was not airborne. Reference lines are included on the plots to indicate the positions of Aberdeen (Dyce) airport and of the other onshore and offshore locations at which the aircraft landed.

Separate plots have been prepared for each of the MF stations which were employed during a particular flight. It should be noted that on a number of occasions during the trials, a number of different beacons was selected by the FTE in quick succession in order to assess which of them provided the best reception. To avoid confusion, all such periods were excluded from the analysis and appear as "gaps" in the plotted data.

The nature of the test flights was such that the distribution of the results plotted below is not uniform. Instead, the concentration of data points is greatest in the vicinity of the platforms where the DGPS approaches and other manoeuvres were undertaken. There are also a large number of samples associated with those periods where the aircraft was stationary on the platform helideck: it would be expected that this might modify the MF reception performance owing to the proximity of the deck surface to the aircraft tail boom, below which the MF antenna was located.

The greatest uniformity of results was obtained for the transit flights to and from the platforms, during which the aircraft track, speed and altitude generally remained constant. Details of the aircraft altitude during each transit flight are provided below the plots.

Except where otherwise stated below, and in the immediate vicinity of the onshore aerodromes, the signal path from the MF station to the aircraft was entirely over the sea.

6.1 Girdle Ness, Flight 2

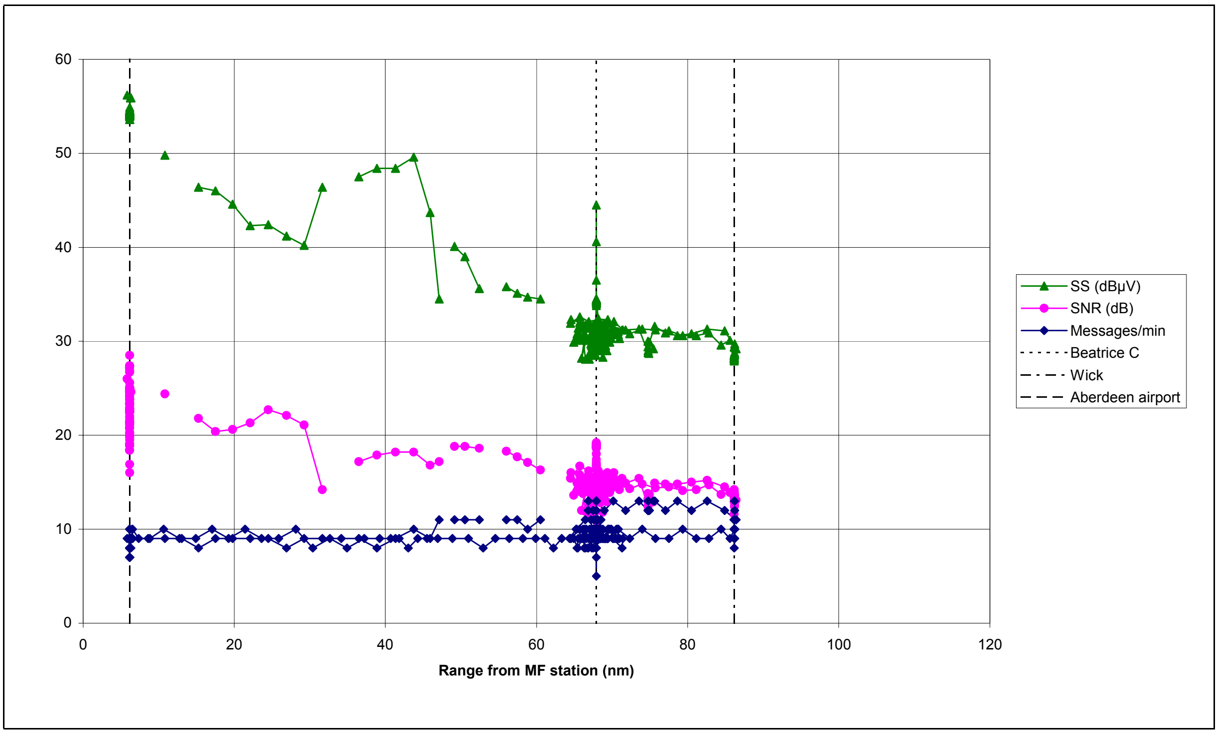

Figure 1 MF Performance on Flight 2, using Girdle Ness

6.1.1 Flight details

The Girdle Ness MF station was employed for the whole of Flight 2. The outbound transit from Aberdeen to the Beatrice C platform was undertaken at an altitude of between 1500ft and 1900ft. Owing to an equipment configuration problem, SS and SNR data was not recorded during this portion of the flight. The intermediate transits between Beatrice C and Wick were undertaken at an altitude between 400ft and 1200ft, and included various forms of aircraft manoeuvre.

The return transit to Aberdeen was undertaken at various altitudes up to 1700ft, with the aircraft descending to 200ft during two experimental approach manoeuvres which were performed over the sea at ranges of between 60nm and 47nm from Girdle Ness. The first 35nm of the MF signal path from Girdle Ness was over land, with the remainder being over the sea.

6.1.2 Comments

There is little evidence for any variation with range of the rate at which the error-free correction messages were received. With the exception of one short period whilst the aircraft was stationary on the platform, the rate remained above the 7 messages/ minute threshold throughout the flight and the performance can therefore be considered acceptable.

The limited amount of SS/SNR data allows few conclusions to be drawn, beyond the fact that the underlying trend of both parameters was to reduce as the range from the transmitter increased. The different slopes on the two plots suggest that the two parameters did not necessarily decrease at the same rate.

6.2 Girdle Ness, Flight 3 (Outbound)

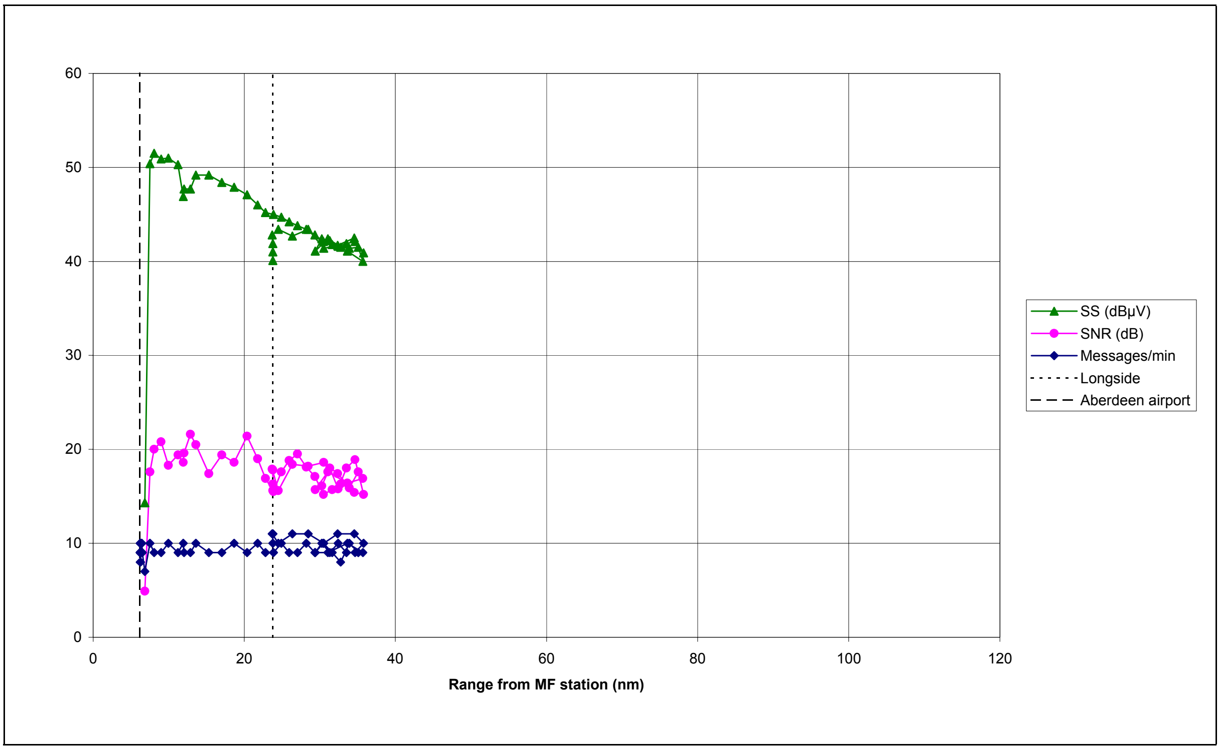

Figure 2 Performance on Flight 3 (Outbound), using Girdle Ness

6.2.1 Flight details.

This data comprises the outbound transit from Aberdeen to Peterhead (Longside) aerodrome. The Girdle Ness MF station was employed throughout this sortie.

The majority of the transit was undertaken at an altitude of 1500ft. A series of four experimental approaches, involving descents to 200ft, was performed over the sea at ranges of between 15nm and 36nm from Girdle Ness.

6.2.2 Comments

Reception of MF correction messages was acceptable throughout this portion of the flight.

A steady decrease can be observed in the SS as the range increases, and there is also limited evidence for a decrease in the SNR (although there appears to be a greater amount of random variation in the latter parameter).

6.3 Girdle Ness, Flight 3 (Inbound)

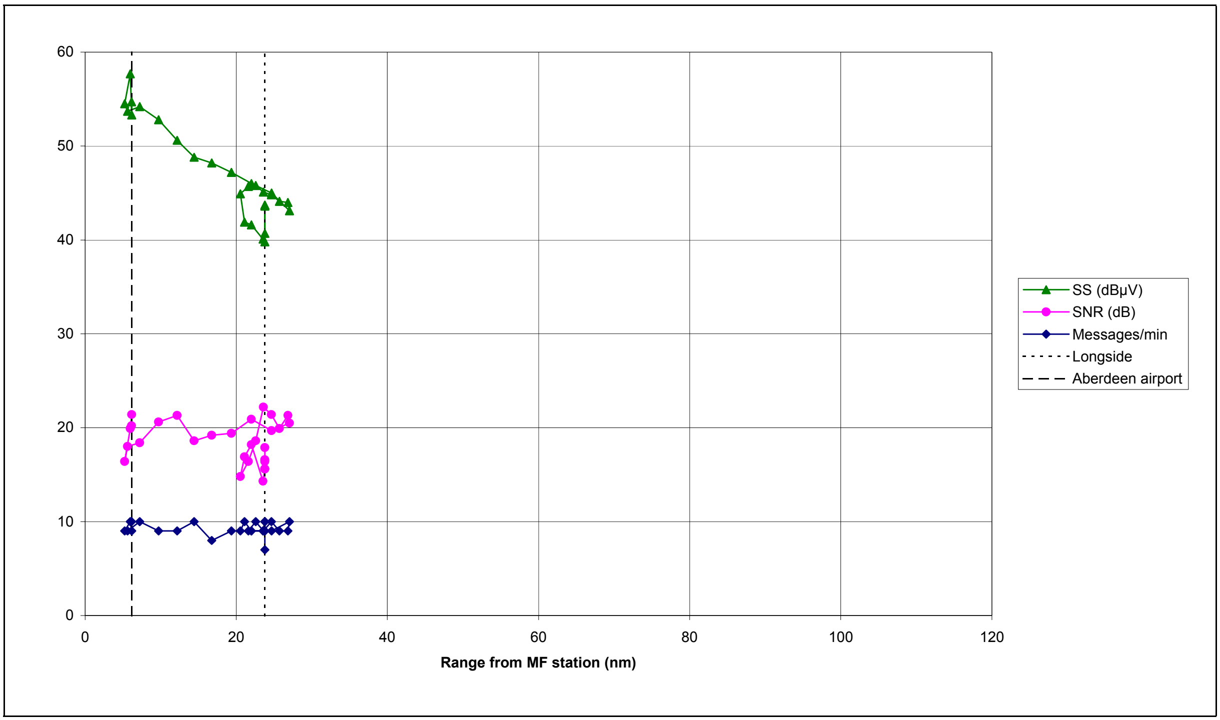

Figure 3 MF Performance on Flight 3 (Inbound), using Girdle Ness

6.3.1 Flight details

This data comprises the return transit from Longside to Aberdeen and has been presented separately from that relating to the outbound flight, due to the fact that the equipment configuration had been modified for the ground manoeuvres at Longside before reverting to the original setup. The Girdle Ness MF station was employed throughout this sortie.

The majority of this transit was undertaken at an altitude of betwen 1200ft and 1500ft. One experimental approach, involving a descent to 200ft, was performed over the sea at ranges of between 22nm and 26nm from Girdle Ness.

6.3.2 Comments.

The correction message rate was again acceptable throughout this portion of the flight, and the characteristics of the SS and SNR parameters are similar to those observed on the outbound sortie.

6.4 Girdle Ness, Flight 4

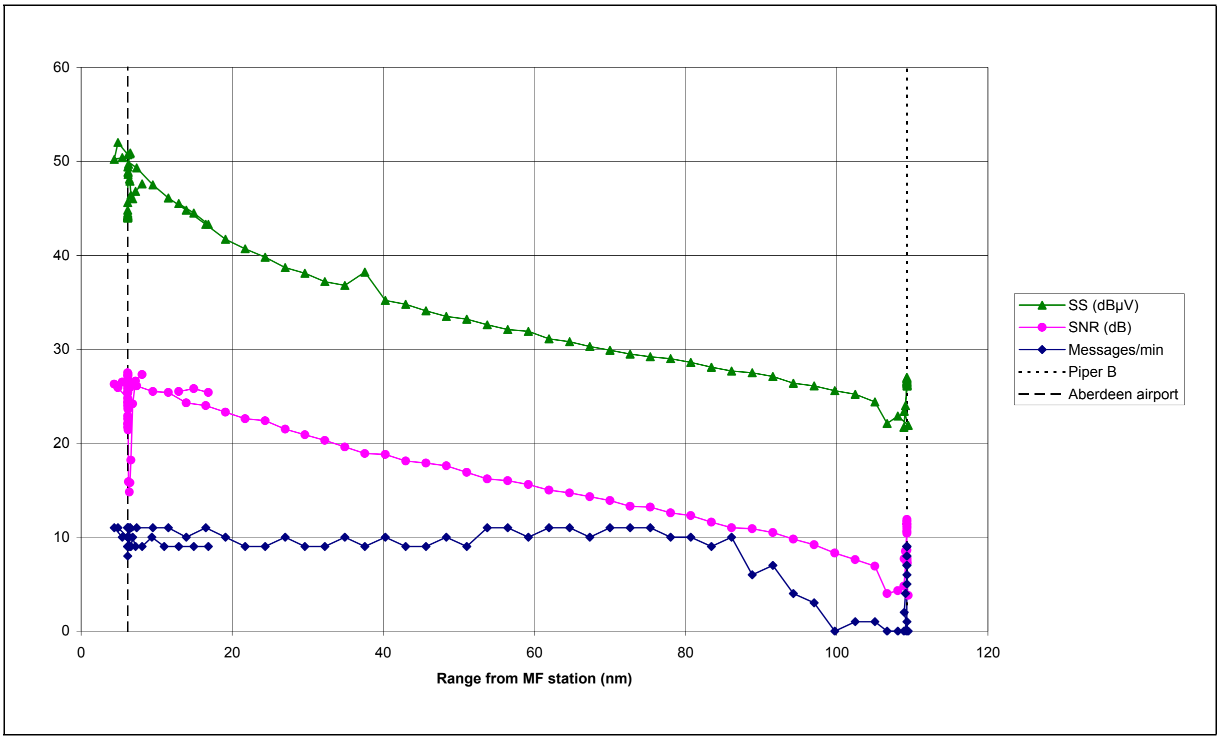

Figure 4 MF Performance on Flight 4, using Girdle Ness

6.4.1 Flight details

This data comprises the outbound transit from Aberdeen to the Piper B platform, plus a very small portion of the return flight (within 17nm of Girdle Ness), during which the Girdle Ness MF station was employed.

The outbound transit was undertaken at an altitude of 4900ft, with the descent to the platform commenced at a range of 100nm from Girdle Ness. On arrival at the platform, it was observed after a short time that reception of the Girdle Ness signals had becoming unreliable, and so the Sumburgh station was selected instead for the majority of the remainder of the flight. The associated data is shown in section 6.8.

6.4.2 Comments

This was the first flight on which an unacceptably low correction message rate was observed during a significant portion of the recording.

The correction message rate fell below the 7 messages/minute threshold at a range of approximately 90nm from Girdle Ness, and generally remained unacceptable beyond this distance (except for a portion of the time that the aircraft was stationary on the platform).

Both the SS and SNR exhibit a steady decrease with range. With the exception of an anomalous sample at approximately 38nm, calculations have shown that the signal strength parameter follows the expected inverse square law relationship (to within ±1dBµV) between ranges of 10nm and 100nm from Girdle Ness.

At 90nm range, the point at which correction reception became unacceptable, the SS and SNR were approximately 27dBµV and 11dB, respectively. These figures conceivably represent thresholds below which it is not possible for the receiver to continue to operate, although more evidence would be required to confirm this assumption.

The possibility was recognised that the one-minute averaging process used to generate the SS and SNR values might be masking any time dependent behaviour in the data (such as interfering transmissions from a pulsed source). To confirm whether this was the case, two short sections of the SS and SNR data (one relating to a time period at 20nm range where the correction reception was acceptable; and a second to a period at 100nm range where it was not) were plotted against time, together with indications as to when each new correction message had been received.

This data, shown in Figure 5 and Figure 6, suggests that there is no evidence for any detailed correlation between the instants at which error-free corrections were received, and the values of the two performance parameters. The time domain noise characteristics of the SS and SNR appear to be similar in both figures, suggesting that there is little evidence for any time-varying interference being present at the higher range.

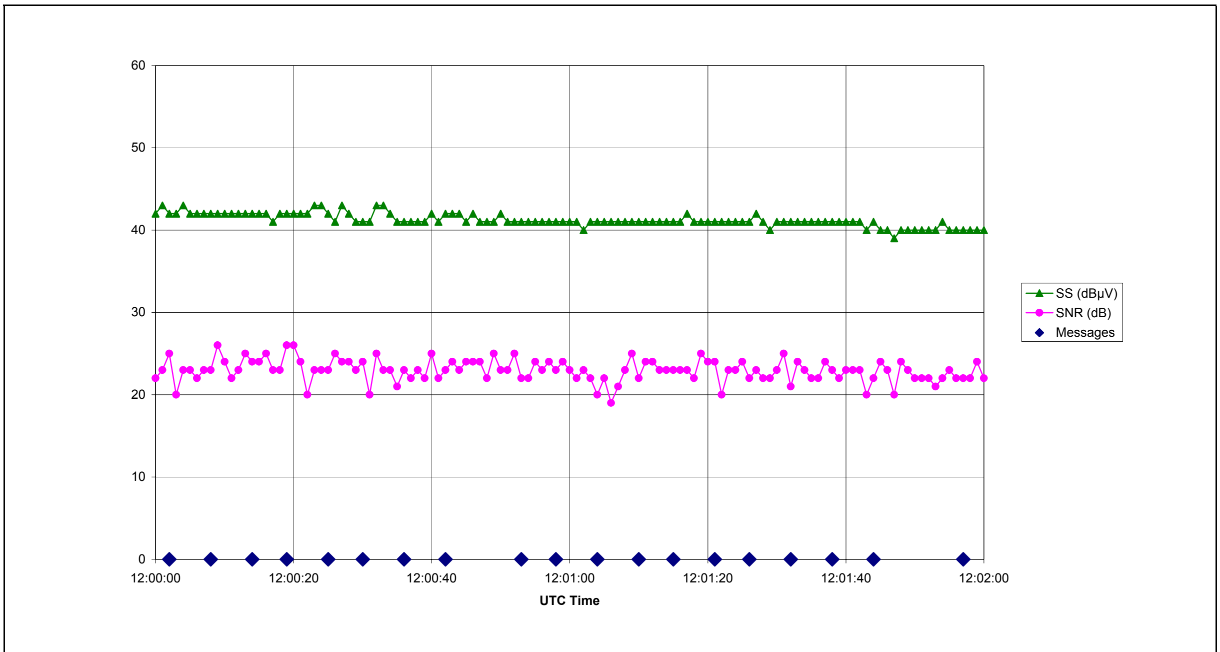

Figure 5 MF Performance on Flight 4, at 20nm Range from Girdle Ness

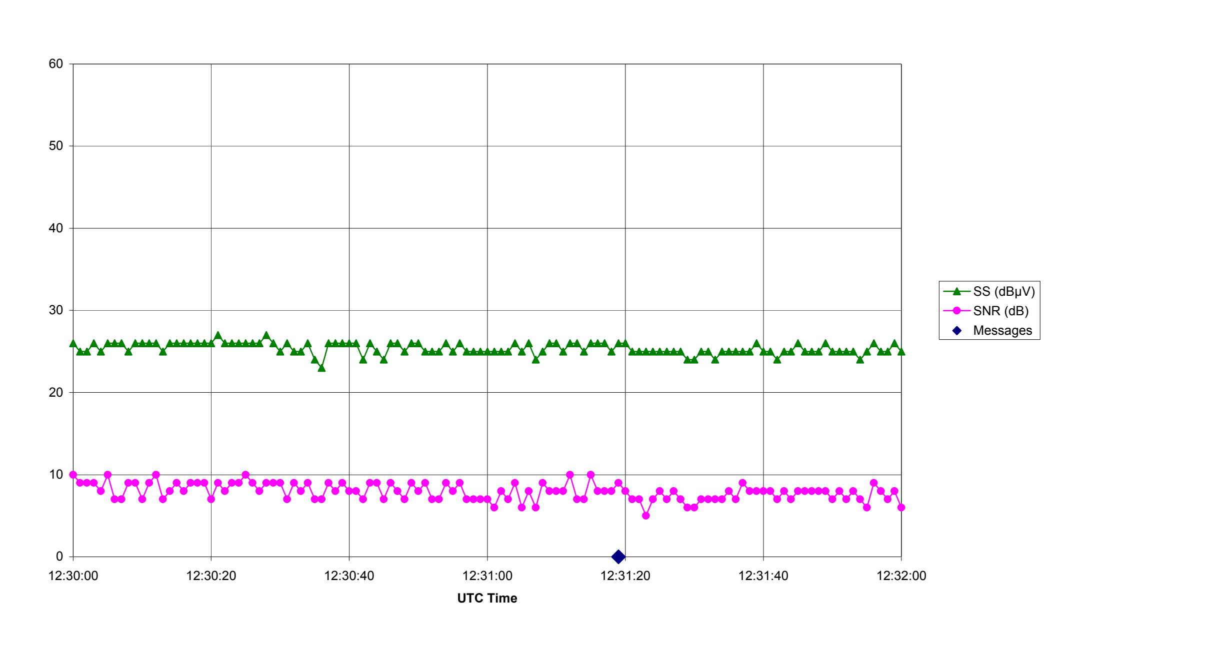

Figure 6 MF Performance on Flight 4, at 100nm Range from Girdle Ness

6.5 Girdle Ness, Flight 5

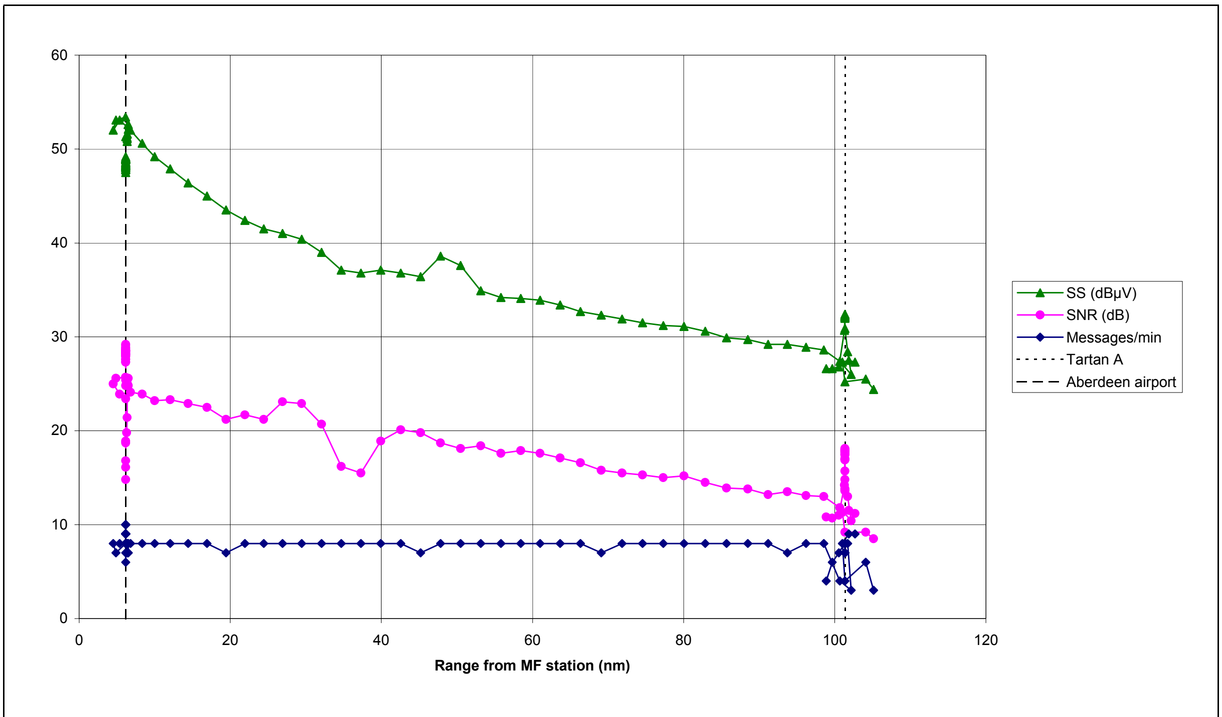

Figure 7 MF Performance on Flight 5, using Girdle Ness

6.5.1 Flight details

This data comprises the outbound transit from Aberdeen to the Tartan A platform, plus a short portion (six minutes' worth) of data during the course of the manoeuvres at the platform, for which the Girdle Ness MF station was employed. The outbound transit was undertaken at an altitude of approximately 4500ft, with the descent to the platform commenced at a range of 91nm from Girdle Ness.

Upon lifting from the platform ready to perform the first of the trials manoeuvres, it was observed that reception of the Girdle Ness signals had become unreliable. Apart from the short period described above, the Sumburgh station was left selected for the remainder of the flight and the associated results are presented in section 6.9.

6.5.2 Comments

There is little available data relating to the period where the message rate became unacceptable. The problems only occurred at a range of around 100nm from Girdle Ness, slightly greater than the 90nm figure on the previous flight.

A comparison between Figure 4 and Figure 5 reveals that the signal strength indication was slightly higher on this later flight, by around 2dB irrespective of range. Beyond 40nm range, the SNR was also slightly higher on this flight (e.g. about 4dB greater at 100nm). This is despite the fact that the outbound tracks and altitudes from Aberdeen on both sorties were very similar.

Application of the estimated SS and SNR thresholds determined on the previous flight (27dBμV and 11dB, respectively) reveals that there is a degree of correlation with the point at around 100nm where the reception became unacceptable.

6.6 Girdle Ness, Flight 6

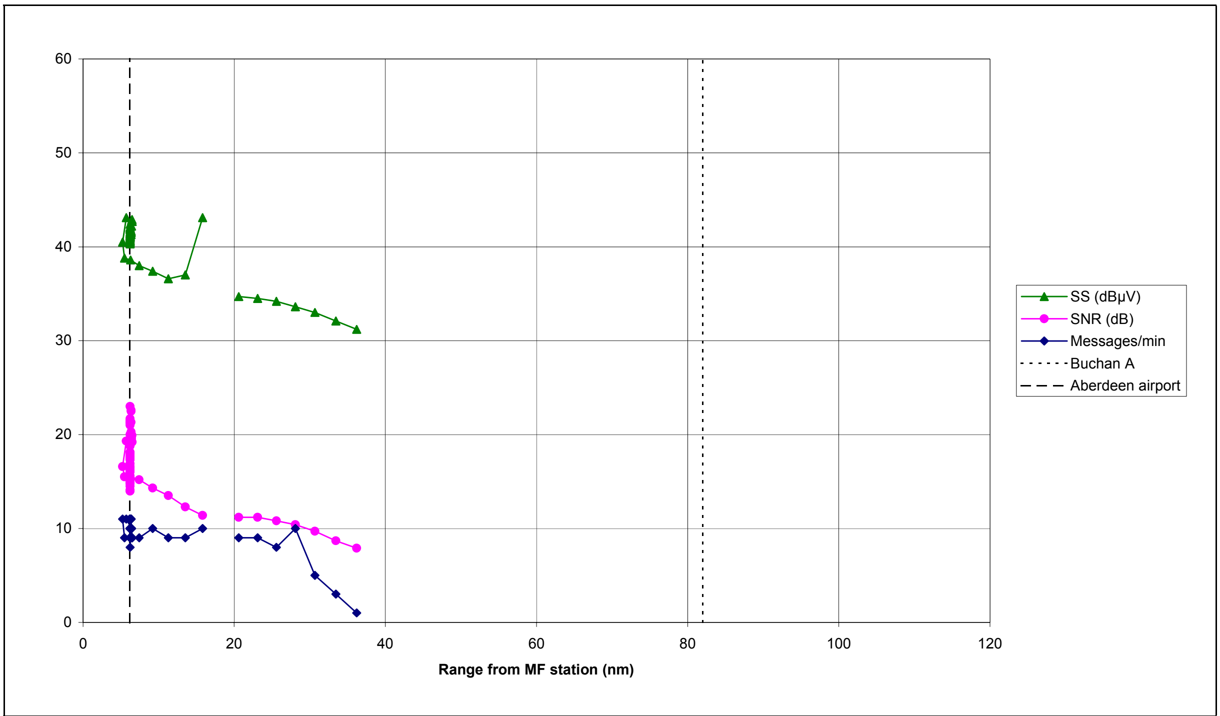

Figure 8 MF Performance on Flight 6, using Girdle Ness

6.6.1 Flight details.

This data comprises the first part of the outbound transit from Aberdeen towards the Buchan A platform, undertaken at an altitude of approximately 4700ft in IMC, during which the Girdle Ness MF station was employed. It was soon identified that reliable reception of the Girdle Ness signals was not being achieved. Attempts to receive signals from Sumburgh, or from any other UK beacon, during the outbound flight proved unsuccessful. Once the aircraft had landed on the platform, the Utsira station was selected for the remainder of the flight (section 6.10). This could not be achieved remotely, but had to be undertaken via the MF receiver's front panel in the aircraft baggage bay.

6.6.2 Comments

Comparison with the data for the previous two flights reveals that both the SS and SNR were significantly lower (by perhaps around 10dB in both cases) on this flight, even during the period prior to departure from Aberdeen.

In view of this fact, it is probably not surprising that the correction message rate became unacceptable at a range of only around 30nm from Girdle Ness. Once again, the associated SNR was approximately 11dB, although the signal strength was higher than the threshold value of 27dBµV suggested in section 6.4.2. This indicates that it is probably the former parameter which is the more significant in determining when the correction rate is likely to become unacceptable.

In the absence of any definite evidence that the Girdle Ness station itself was operating at reduced power, it might be concluded that the reduction in performance was most likely to have been equipment related (e.g. a fault in the receiver or antenna, or because of increased interference from another on-board source). It is considered particularly significant that degradation was observed during the period where the aircraft was on the ground at Aberdeen, suggesting that it is less likely for any in-flight factor to have been the sole cause.

6.7 Girdle Ness, Flight 7

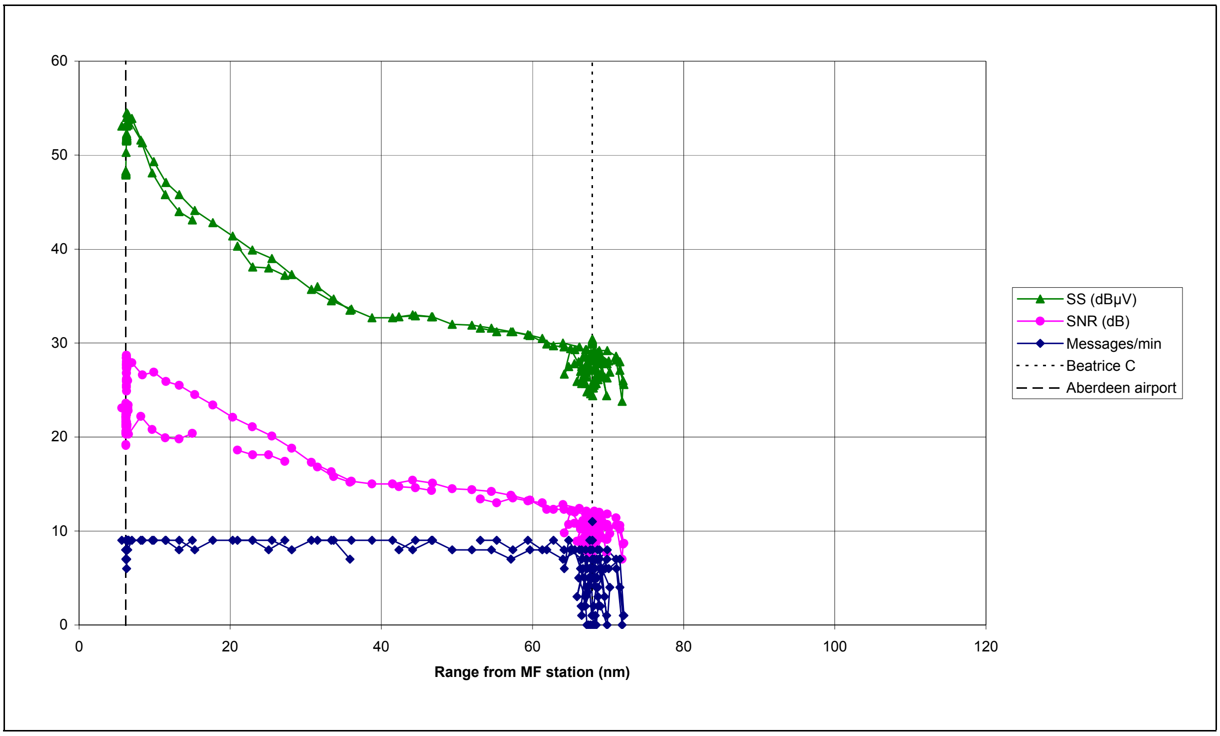

Figure 9 MF Performance on Flight 7, using Girdle Ness

6.7.1 Flight details

This data comprises the whole of the first sortie to and from the Beatrice A and Beatrice C platforms, during which the Girdle Ness MF station was employed. The outbound transit (the intermittent data on the plots) was undertaken at an altitude of approximately 4000ft, and the return transit (continuous data) was undertaken at 3000ft.

The first 35nm of the MF signal path from Girdle Ness was over land, with the remainder being over the sea. Reliable reception of the Girdle Ness signals proved difficult during this sortie and it did not prove possible to receive signals from

Sumburgh, in spite of several attempts (including during the outbound transit, hence the periodic gaps in the data). On returning to Aberdeen to prepare for the second sortie, the Utsira beacon was selected for the remainder of the flight, the results for which are presented in section 6.11.

6.7.2 Comments

Although the Girdle Ness reception difficulties were not as bad as had been experienced on Flight 6, it still did not prove possible to maintain adequate reception throughout the manoeuvres at the platform. By contrast, Flight 2 had been undertaken at the same location without any difficulties being observed.

A comparison between Figure 2 and Figure 9 reveals that the signal strength in the vicinity of the platform was slightly lower (by perhaps around 4dB at 60nm range) on the later flight, although close to Aberdeen the results were very similar. There is certainly no evidence for the large reduction in signal strength which was seen on Flight 6.

The SNR near the platform was also lower (by perhaps 4dB at 60nm range), and during the manoeuvres it frequently fell below the notional 11dB threshold which was previously identified. However, there is less evidence for any significant reduction in SNR close to Aberdeen, with the results certainly having much more in common with those from Flights 1-5 than with those from Flight 6.

6.8 Sumburgh, Flight 4

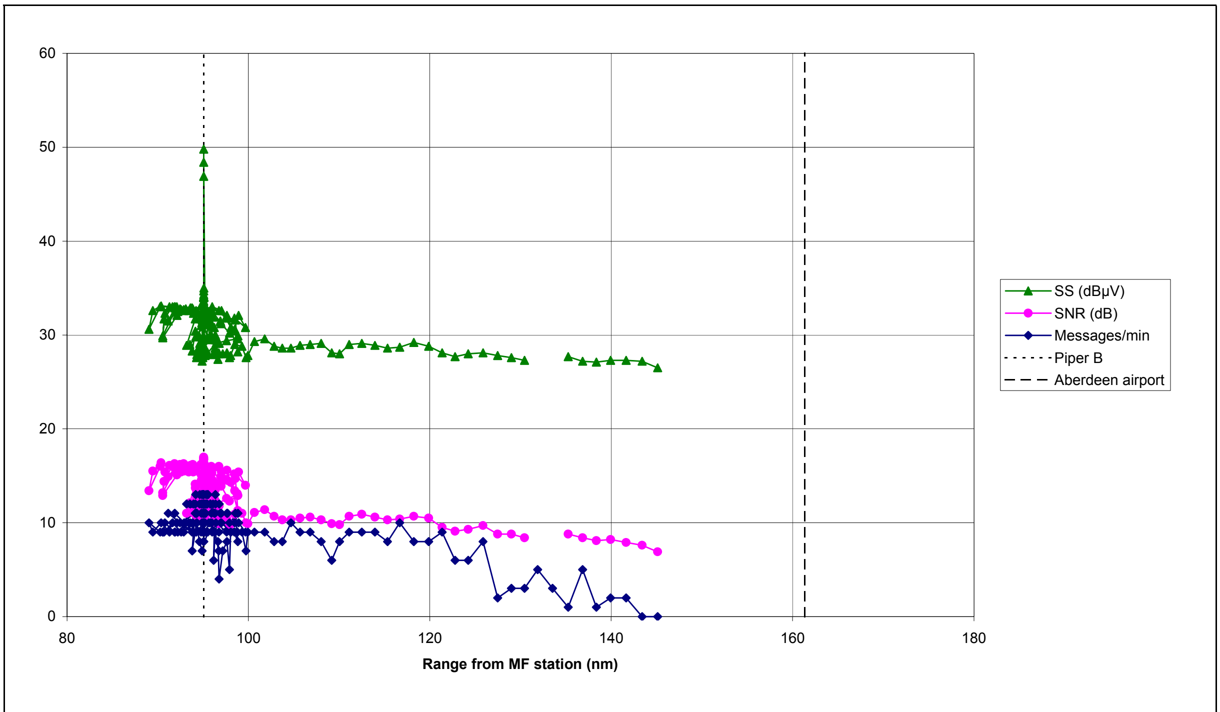

Figure 10 MF Performance on Flight 4, using Sumburgh station

6.8.1 Flight details

This data comprises the approaches and manoeuvres at the Piper B platform, and the subsequent return transit to Aberdeen at an altitude of 3800ft, during which the Sumburgh MF station was employed. The Girdle Ness station was reselected on the return transit once the loss of reliable reception from Sumburgh had been recognised.

6.8.2 Comments

The correction message rate at the platform using Sumburgh was generally acceptable, with the exception of a few short periods (less than five minutes in total).

Comparison with the Girdle Ness plots (in particular, Figures 4 and 5) reveals that the SS and SNR figures appear to have been slightly higher at equivalent ranges from the two stations (e.g. the signal strength measured at 100nm from Girdle Ness was 26dBµV, whereas for Sumburgh it was around 29dBµV).

Although detailed information on the output power of the two correction stations was not available, the published operating range for the Sumburgh MF beacon was higher, at 56nm rather than 40nm (ref 1 volume 2 section 7.1), which may explain why better reception was achieved using Sumburgh. In the MF band, atmospheric noise represents the dominant noise source and would have been very similar at the two frequencies involved (Girdle Ness 311.5kHz, Sumburgh 304.5kHz). However, the possibility of interference from a co-channel source cannot be completely eliminated.

The reason for the occasional larger (up to nearly 50dBµV) values of SS, observed while the aircraft was stationary on the platform, is not clear. It is conceivable that they arose from the operation of a very localised interfering source, however there appears to have been no corresponding increase in the SNR.

6.9 Sumburgh, Flight 5

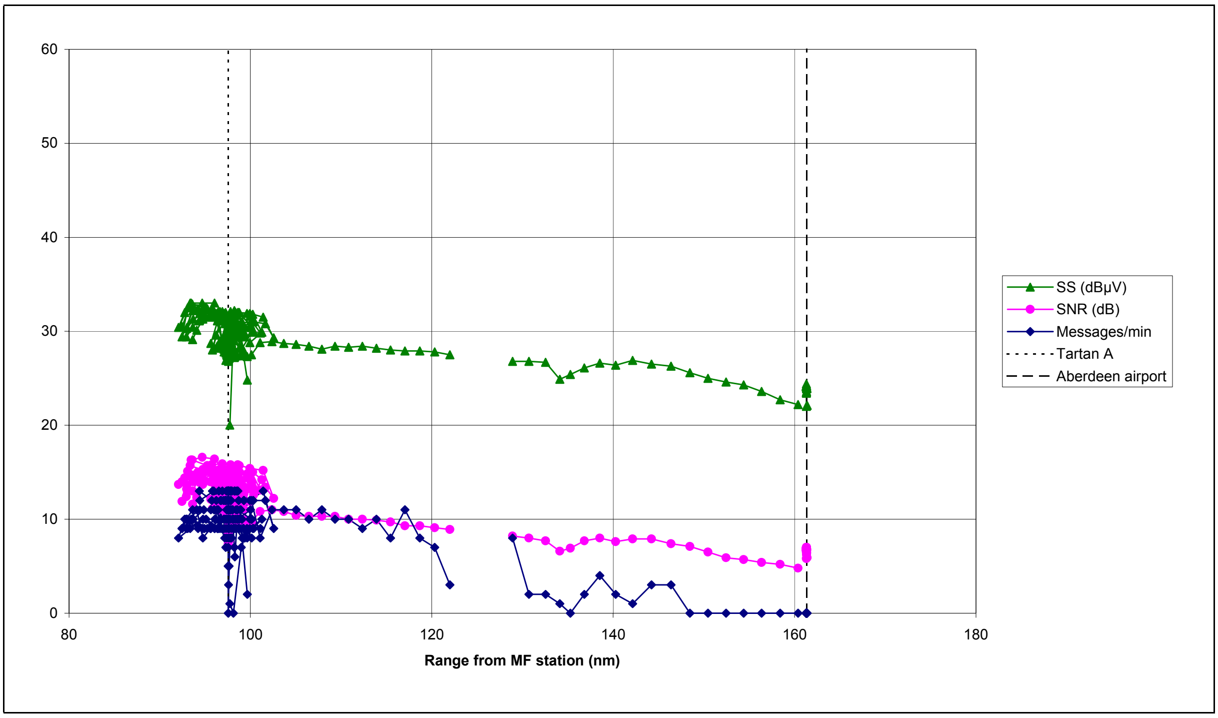

Figure 11 MF Performance on Flight 5, using Sumburgh station

6.9.1 Flight details.

This data includes the majority of the approaches and manoeuvres at the Tartan A platform, and the subsequent return transit to Aberdeen at an altitude of 3800ft, during which the Sumburgh MF station was employed.

6.9.2 Comments

These results are very similar to those which were obtained using Sumburgh on Flight 4.

The correction message rate again became unacceptable at a range of approximately 125nm from Sumburgh, where the SNR was around 9dB.

6.10 Utsira, Flight 6

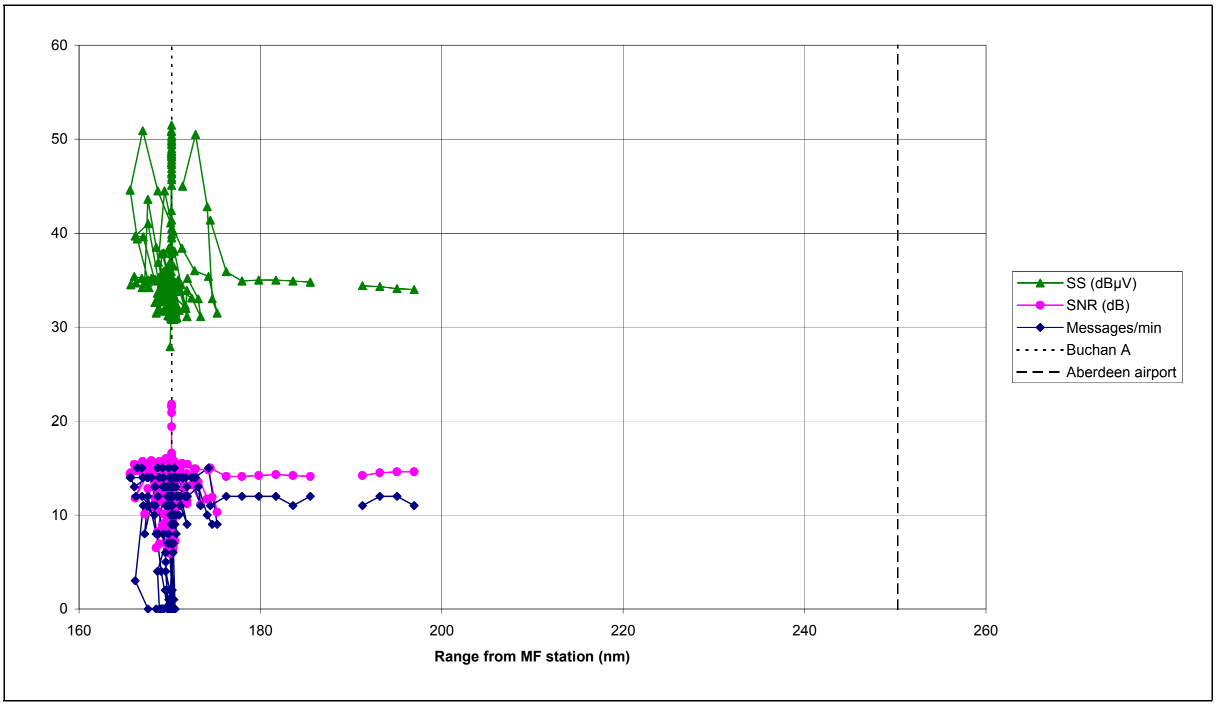

Figure 12 MF Performance on Flight 6, using Utsira station

6.10.1 Flight details.

This data includes the approaches and manoeuvres at the Buchan A platform, and the subsequent return transit to Aberdeen at an altitude of approximately 2000ft, during which the Utsira MF station was employed.

6.10.2 Comments

The range from the transmitting station (approximately 170nm at the platform) was considerably higher than that at which satisfactory results had previously been obtained using the two UK stations.

Large variations were observed in the correction message rate, and in the SS and SNR parameters, whilst operating in the vicinity of the platform. The correction rate was below the acceptability threshold for a significant proportion of this period.

The signal strength was generally well above 30dBμV, and the SNR (ignoring some samples taken whilst stationary on the platform) ranged between 6dB and 16dB. These figures are considerably higher than would have been expected at this range on the basis of the Girdle Ness and Sumburgh data (the nearest available comparison is with Figure 11, which indicates that the Sumburgh signal strength at 160nm range was around 22dBμV, and the SNR around 5dB: well below the threshold for an acceptable message rate).

Although the published operating range for the Utsira beacon (60nm) was comparable with that of Sumburgh (56nm), these results suggest that the output power of the transmitter was considerably higher. The fact that the Utsira transmissions were not encrypted may also be a factor. It is perhaps unfortunate that the Utsira station was apparently deselected during the return transit, particularly as the correction rate still appeared to be acceptable.

6.11 Utsira, Flight 7

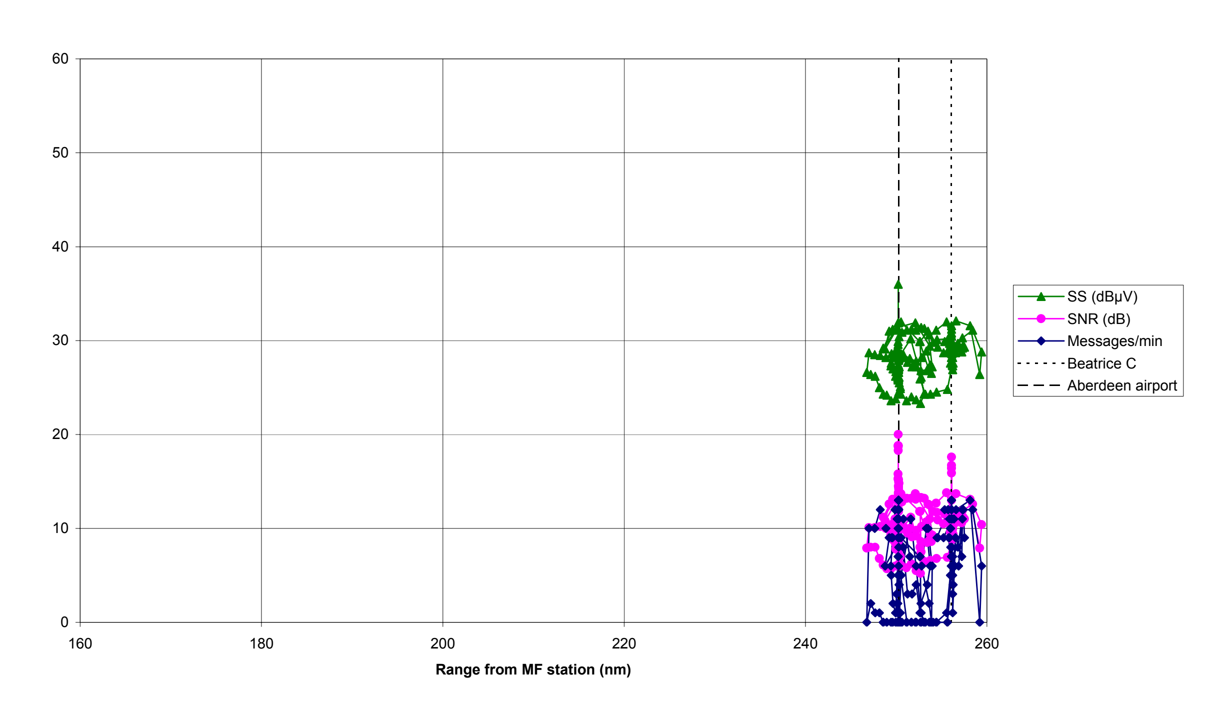

Figure 13 MF Performance on Flight 7, using Utsira station

6.11.1 Flight details

This data comprises the whole of the second sortie to and from the Beatrice A and Beatrice C platforms, during which the Utsira MF station was employed. Both the outbound and return transits were undertaken at an altitude of approximately 1500ft.

6.11.2 Comments

The extended range from Utsira (in excess of 250nm) is such that, on the basis of the results using the two UK beacons, it might be considered very surprising that any correction messages were received.

Once again, large variations were observed in all three of the plotted parameters. Although it is difficult to perform a direct comparison with Figure 12 owing to the different ranges involved, it can be stated that the results are not dissimilar from those observed when using Utsira on the preceding flight.

6.12 Variation with Aircraft Heading

A common factor in all of the above plots of the MF performance parameters was the observation that although the results obtained during the outbound and return transit portions of the flights were very smooth (exhibiting a clear correlation against range from the transmitter), a degree of additional scatter developed once the aircraft began manoeuvring around the platform.

Consideration was given as to how these two portions of the flights had differed, and as to how these differences might have affected the MF reception. For example, it was recognised that the manoeuvres had encompassed changes in aircraft altitude, heading, attitude and speed; whereas these parameters had generally remained constant whilst in transit to and from the platforms.

In addition, the possibility of localised interference sources either on, or close to, the platforms could not be discounted as a mechanism for poorer MF performance. At several of the platforms, in particular, the existing MF NDB transmitter had been in operation during portions of the trials.

From an analysis of a portion of the recorded data, it was discovered that there was clear evidence for a correlation between the SS and SNR parameters and the aircraft heading.

For example, Figure 14 is a plot of these two parameters against true heading, for the period in Flight 2 during which manoeuvres (approaches and orbits) were undertaken at the Beatrice C platform. No averaging has been performed: instead, each individual sample is plotted.

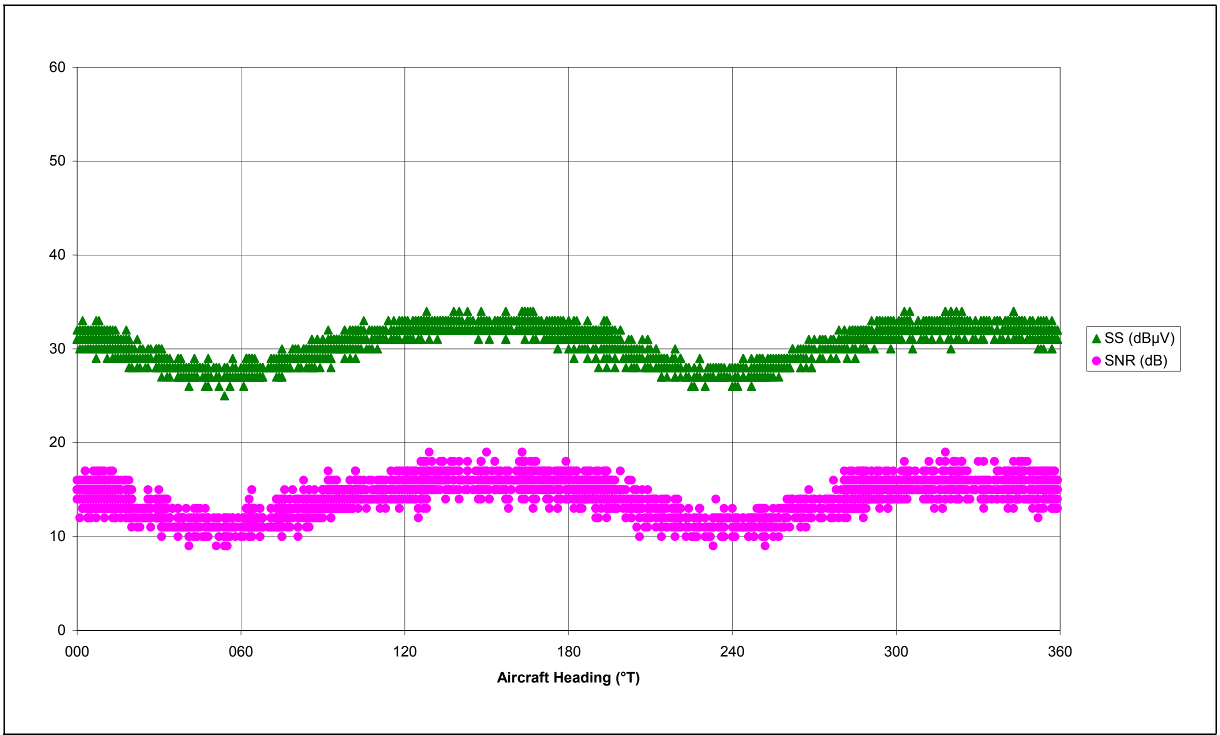

Figure 14 MF Performance against Heading on Flight 2, using Girdle Ness

It is apparent from the figure that a sinusoidal variation is present on both the SS and the SNR data, with a peak-to-peak amplitude of around 5dB for SS, and 6dB for SNR. The two maxima of this variation can be observed to be located at headings of around 150°T and 330°T, and the minima are at approximately 060°T and 240°T.

These maxima were then shown to coincide with the relative geometry of the transmitter and the platform: the bearing of Girdle Ness from Beatrice is approximately 150°T (reciprocal 330°T). This implies that the MF reception was at its best when the aircraft was heading either directly to or away from the transmitter, and that it was at its worst when flying at right angles to it.

This provides evidence for the reception pattern of the MF reception antenna having been non-uniform, with two main lobes oriented along the aircraft longitudinal (fore/aft) axis. Since the antenna was mounted beneath the rearward extension of the aircraft tail boom, this is perhaps not a totally unexpected result, although the extent of the variation is very large.

A consequence of the proximity of the Girdle Ness station to Aberdeen airport was that the aircraft heading during the transit flights to and from each of the platforms would have coincided with one of the maxima of the antenna "polar diagram". This provides an explanation as to why, even though the MF reception from Girdle Ness might have been satisfactory whilst proceeding outbound towards a platform, problems were sometimes encountered once the manoeuvres were commenced upon arrival.

Examination of Figure 4 demonstrates that, if it is assumed that a heading change through 90° would cause the measured SNR to reduce by 6dB, then this would have the same effect as increasing the range to the transmitter by approximately 30nm. When operating in the vicinity of the platform, it is clear that it is this heading dependence which would dominate: if the SNR was already sufficiently marginal, then reception would be lost when flying on certain headings.

Although a full analysis of the "polar diagram" effect and its impact upon the results at each platform is beyond the scope of this study, some additional confirmation of its existence was obtained by plotting the Utsira data from Flight 7 as a function of heading (Figure 15).

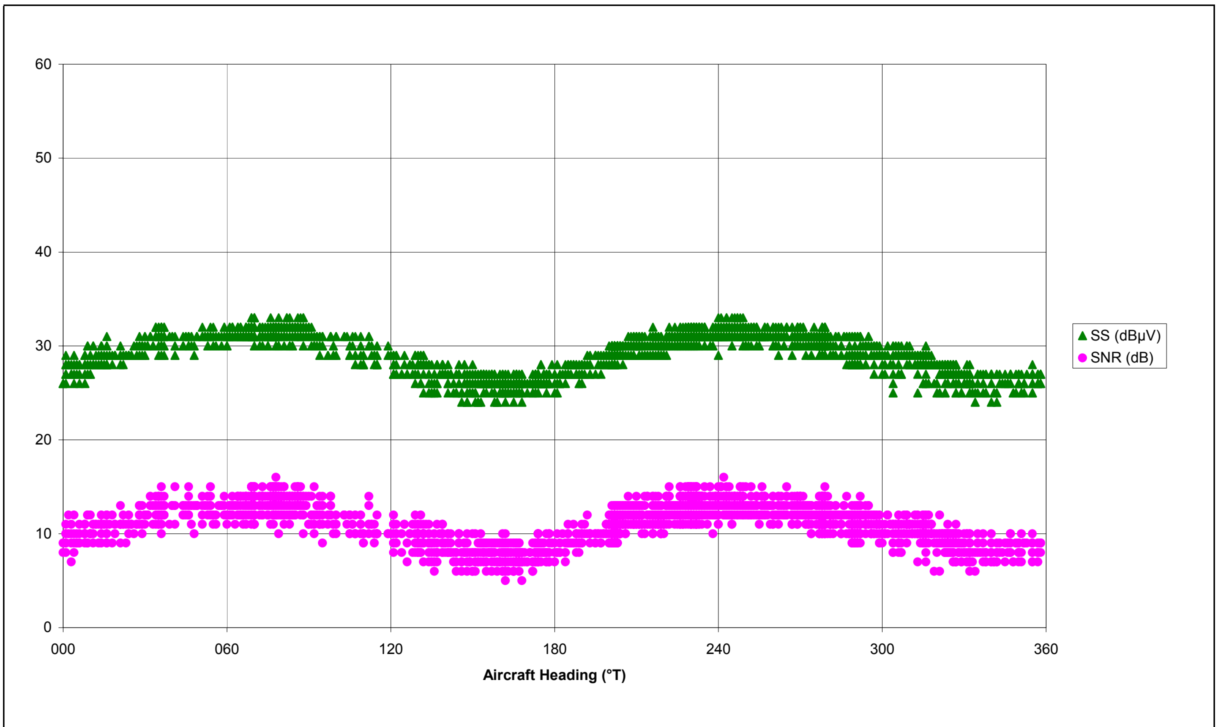

Figure 15 MF Performance against Heading on Flight 7, using Utsira

A sinusoidal effect is again present, with amplitudes similar to those observed on the Girdle Ness data. The maxima of the pattern are at headings of around 080°T and 240°T, which correspond with the bearing of Utsira from Beatrice (approximately 250°T).

An explanation is still required as to why the MF reception was observed to become significantly worse on the final two test flights, compared to the first five. Fortuitously, it is possible to undertake a direct comparison between a flight in each group, owing to the fact that the Beatrice platform was employed on both Flight 2 and Flight 7. On both occasions, data using the Girdle Ness station is available.

Figure 16 shows the results of plotting the Flight 7 Girdle Ness data against heading, which can be compared with the equivalent Flight 2 data shown in Figure 14.

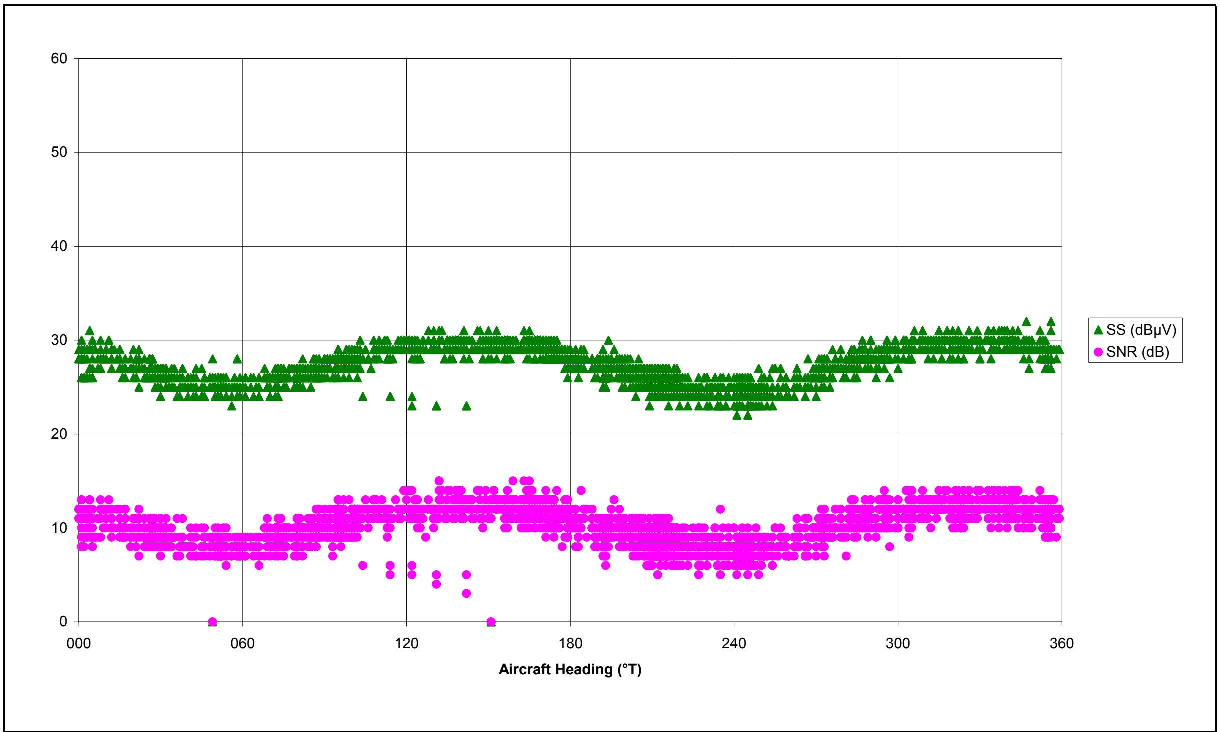

Figure 16 MF Performance against Heading on Flight 7, using Girdle Ness

A comparison between the two plots reveals that, although the maxima and minima of the SS and SNR are in the same positions, and the amplitudes of the sinusoidal variations are similar, the underlying mean values were lower on Flight 7.

It was estimated that the SS was on average around 3dB lower, and the SNR about 4dB lower. These values are comparable with those which were estimated (section 6.7.2) for the associated transit flights, suggesting that the "polar diagram" effect was no worse on the later flight, but that some other mechanism must instead have been at work.

The possibility that the transmitter output power was reduced on the later sortie must be considered. Although a response at the time of the trials from the operator of the stations had stated that this was unlikely to be the case, it is believed that it is not now possible to confirm this fact with 100% certainty. The service has, in any event, since been superseded by a series of centrally funded, free access correction stations which, although in many cases located in the same positions, employ different equipment and a revised frequency allocation plan.

Alternatively, there is the possibility that the performance of the airborne equipment had become degraded in some way. It so happens that the MF receiver had been changed between flights 6 and 7 (for reasons unconnected with its performance: the power supply on the original unit had developed a fault during bench testing). Since poorer performance was observed on both of these sorties, it is considered unlikely that a fault with the receiver unit itself would have been the principal cause.

The remaining possibilities include a problem with the antenna unit, or a cable fault (e.g. a poor contact or an impedance mismatch). The antenna was removed from the aircraft between trials sorties and could potentially have been damaged in storage. Unfortunately the relevant items are now believed to have been discarded by the operator, and cannot be recovered for further testing.

Another possibility is that interference was being generated from another on-board source (anecdotal evidence from the operator was that the S76C avionic systems were known to be prone to EMC and static problems, and that these were not always repeatable). Certainly, the data from Flight 6 (section 6.6.2) suggests that the problems were already present whilst on the ground at Aberdeen before departure.

7 Discussion

During the flight trials programme described in ref 1, the sole purpose of the differential corrections was to provide a mechanism which improved the positioning accuracy of the airborne GPS equipment. Since the trials were undertaken whilst SA was active, the resultant increase in accuracy was significant. A horizontal positioning accuracy of 7m at the 95% level was demonstrated when using the MF corrections.

Correction signals from a reference station can also be used to provide integrity as well as accuracy: this provides a mechanism whereby failure of the signal-in-space can be identified and signalled to the user equipment. A term commonly applied to refer to this concept is “augmentation” (e.g. GBAS: Ground Based Augmentation System, SBAS: Satellite Based Augmentation System).

The loss, whether temporary or permanent, of the signals from the augmentation system will inevitably result in a loss of integrity; and would also normally cause a reduction in accuracy. It would be expected that this situation would be signalled to the pilot in the form of a navigation system failure warning: the resultant action which he would need to take would be dependent upon the phase of flight and upon the relevant operational procedures. For example, during the approach phase it could result in the execution of a missed approach.

During the offshore flight trials, loss of the MF correction signals caused the GPS receiver to revert to non-differential positioning mode. Although this situation was communicated to the pilots, the frequency with which these failures occurred meant that it was necessary to continue to display valid lateral and vertical guidance indications derived from the non-differential solution, in order to avoid prejudicing the trials programme. Due to the presence of SA, reversion to non-differential mode frequently resulted in large “step changes” in the deviation indications presented to the pilots. Since the indications still remained “flyable” in this situation, these problems were considered to represent a nuisance rather than a hazard. However, it was made clear by the pilots that this situation would not be acceptable in a normal operational environment.

A commercial set of GPS equipment, rather than units specifically intended for civil aviation applications, was employed for the trials (the Trimble receiver, although based upon a “TSO” compliant aviation model, had undergone a special modification to permit its use with an RTCM differential correction source). Likewise, the MF corrections were obtained from a commercial service provider whose primary market was the marine survey industry, and employed a correction format (RTCM) which was based upon a marine rather than an aviation standard.

Section 5 of this report has demonstrated that reception of a single corrupted data bit, anywhere within an RTCM correction message, was observed to result in both the Navstar and Trimble GPS receivers rejecting the entire message. Although the encoding scheme used to transmit the corrections potentially offered the ability to correct for single bit errors, this capability is not mentioned in the standard industry documents. Since the Type 1 RTCM messages are quite long, typically around 500 bits per message, the bit error rate at which reception difficulties would be expected is actually very low (e.g. a bit error rate of around 0.2% would result in an average of one incorrect bit per message). Implementation of the error-correcting scheme would have allowed the system to continue to operate in the presence of a higher bit error rate.

It was also demonstrated in section 5 that the recovery capability of the Trimble GPS receiver following reception of a corrupted RTCM message was considerably worse than that of the Navstar unit. This may explain why the availability of the differentially-corrected position samples was lower for the Trimble receiver.

In the UK, a centrally funded MF correction service has now superseded the system which was employed for the flight trials. The new service incorporates a form of integrity monitoring which is intended to detect failures in the correction broadcast, although this was not specifically designed to meet aviation requirements.

The new service is understood to transmit correction messages using the RTCM Type 9 format (ref 2) which, in effect, involves the substitution of a series of shorter messages in place of a single, long, Type 1 transmission. This may improve the ability of GPS receivers to operate in the presence of corrupted data, although this assumption would have to be tested experimentally.

Both the Navstar and Trimble GPS receivers were demonstrated in section 5 to apply a thirty second time-out between failure of the differential correction input and reversion to non-differential mode. Analysis may show that this represents an excessively long time-to-alarm if the external signal is needed to provide integrity, although there is potentially a trade-off to be made against system availability, particularly when operating in an environment where the correction signal is prone to transmission errors.

Section 6 of this report considered how the MF correction reception was observed to vary during the flight trials, as a function of range from the transmitter and of aircraft heading.

Clear evidence was demonstrated for the aircraft antenna reception pattern having been significantly non-uniform. When operating at extended range from the transmission station, this caused the availability of the differential-corrected position solution to vary according to the aircraft heading. For example, a heading change through 90° was shown to have an effect similar to that of a 30nm increase in range. It is conceivable that the use of an alternative antenna, or of a different installation arrangement on the aircraft, might reduce this “polar diagram” effect to a certain extent. However it is considered unlikely that a completely uniform reception pattern could be achieved, particularly in view of the very limited antenna installation options on the offshore helicopter fleet. Careful analysis and flight testing of any future operational system using similar technology would therefore be needed: for example, satisfactory performance during a transit flight will not guarantee that problems will not be encountered once manoeuvres are commenced.

Although the availability of the MF correction messages was generally observed to reduce as the distance from the transmission station was increased, the results were not always repeatable between different flights (in particular, there is a some evidence for the performance of the airborne equipment having become degraded between the first five flights and the final two, although it was not possible to determine the precise nature of the problem).

It is therefore likely that an increase in the transmitter output power, and/or an increase in the sensitivity of the airborne equipment, might have improved the correction reception at the platforms. However, due to the lack of repeatability in the results, and to the fact that it was not possible to isolate all of the other potential interference sources (such as distant MF transmitters operating on similar frequencies, or equipment operating at offshore locations) there can be no guarantee that this would necessarily solve the problem.

Precipitation static was identified in ref 1 as another potential source of interference to the aircraft MF installations. Since no additional analysis or experimentation regarding this aspect has been undertaken, it cannot be eliminated as a reason for the reception problems, and could conceivably explain some of the variations between different sorties.

All of the flight trials were undertaken during daylight hours. As a result, it is not possible to draw any conclusions regarding the performance of the MF equipment at night, when atmospheric effects can modify the signal propagation characteristics in this band.

8 Conclusions

Summary of Conclusions

(1) The reception performance of the airborne MF equipment exhibited considerable directional variation. The effect of this variation was observed to dominate the performance results when operating in marginal reception conditions.

(2) It was demonstrated that neither the Navstar nor the Trimble receiver appeared to make use of the error-correction capabilities associated with the RTCM correction parity algorithm. As a consequence, the availability of the differential solution from both units was more susceptible to transmission errors than would otherwise have been the case.

(3) It was observed that the Trimble receiver sometimes required a lengthy recovery period, of up to one minute, before it would re-enter differential mode following reception of corrupted correction messages. Similar behaviour was not observed with the Navstar receiver.

(4) Both the Navstar and Trimble receivers were shown to be capable of operating using "stale" differential correction data for a period of up to thirty seconds, which improved the availability of the differential-corrected solutions in marginal reception conditions. However, the absence of these updates could result in a degradation in accuracy with time, due to the extrapolation effect of the range-rate corrections, and would also remove the source of integrity information.

Implications of the Results.

(5) The results were obtained during 1996 and are not necessarily applicable to the present situation. For example, various changes which have been made to the UK MF correction service in the intervening period, including the installation of new equipment, the removal of encryption, and the introduction of a new frequency allocation plan.

(6) Improvements have also been made to the MF datalink receiver and antenna technology during the period since the trials were undertaken. It is therefore conceivable that better results could be obtained by using the latest equipment.

(7) An increase in the transmitter power would probably have improved the MF correction signal reception at the locations employed on the flight trials. However, it is not possible to predict the effect of such an increase from the existing dataset with any certainty.

(8) The likelihood of a strong directional variation in reception performance must be taken into consideration when defining the certification requirements for any airborne MF correction receiver installations.

(9) If the datalink is necessary to ensure the integrity of the navigation solution, then the use of a thirty second timeout period, which could represent a substantial portion of an approach, might prove to be excessively long. However, a shorter timeout might reduce the probability of maintaining a continuous differential-corrected solution.

(10) An excessively long timeout period could also affect the solution accuracy due to the reliance on extrapolated range-rate corrections, particularly if Selective Availability was ever reactivated.

Future Work

(11) The performance of the current MF technology could best be determined by undertaking a further data collection exercise, using state-of-the-art equipment, over a wider geographical area, and in a range of weather and day/night conditions.

(12) These experimental results should ideally be compared against a theoretical model of the MF signal propagation characteristics to confirm whether the predicted performance is being achieved.

(13) In order to define the timeout periods associated with the generation of an alarm condition, an analysis must be performed to consider the probability and impact of the augmentation signal having been affected by errors in transmission. The recovery time following the detection of faulty data must not be excessive.

(14) Before any such analysis can be undertaken, supporting data is required which is based upon the empirical monitoring of the signal-in-space and bit error characteristics of a representative augmentation signal.

(15) Strong consideration must also be given to the environment in which the equipment will operate: for example, it must not be assumed that the results of fixed-wing tests will necessarily be applicable to helicopter operations.

Study 3

Effects of Satellite Unavailability and Analysis of Receiver Tracking Performance

Contents

Study 3 Effects of Satellite Unavailability and Analysis of Receiver Tracking Performance

Study 3 Effects of Satellite Unavailability and Analysis of Receiver Tracking Performance

Executive Summary

During 1996 a series of flight tests was undertaken to explore the viability of using GPS for helicopter approach guidance to offshore platforms. The satellite constellation had up to 26 satellites. Since the US policy for GPS promises only 24 satellites, and since satellites sometimes fail, the results from the flight tests were optimistic. This document reports on a study that explored the sensitivity those results would have to reductions in the satellite constellation. The study focused on performance as quantified by availability, precision and receiver tracking.

The primary findings were that:

i) availability of an integrity-assured position solution would be seriously compromised if there was a loss of a single satellite,

ii) in those cases where a solution was available the precision was not seriously affected,

iii) availability was lost because the receiver tracked fewer satellites than expected during the flights,

iv) this problem appeared to be linked to an interference effect from the tail rotor.

There are two safety implications from the findings:

a) reliance on GPS alone for helicopter approach guidance is unsafe,

b) predictive RAIM for helicopter GPS performance will give a very optimistic result.

It is recommended that research into this area must focus on the potential causes of the loss of tracking performance and on the potential for techniques and operational procedures to alleviate the effect.

Glossary

A number of terms have the following specific meaning when used in this document.

Architecture

The term describes some of the aspects that determine the GPS receiver's performance. It includes:

a) the number of satellites tracked and able to contribute a measurement of pseudo range to the navigation solution,

b) the selection criteria if only a subset of the available satellites is used,

c) the extent of integrity checking,

d) any external aiding,

e) the manner in which different measurements are combined,

f) any smoothing in the position solution.

Section 2.3 discusses the architecture that was selected by CAA and Section 3.6 indicates how the availability results would be altered if different architectures were used.

FD RAIM

This term has two uses.

a) It is a general term for Fault Detection Receiver Autonomous Integrity Monitoring, implying that the software is designed to detect when there is a fault in one of the satellite measurements but (usually) cannot determine which satellite is associated with that fault. Being a general term, there are several techniques for achieving such fault detection, most of these rely on some form of consistency checking in an over-specified solution.

b) In the Simulator (see later in this glossary) it has a specific meaning relating to an architecture in which five satellites are used in five groups of four satellites. For each group PDOP is computed, these are added and the sum tested against that for other groups of five. The value for the group with the lowest sum is used, subject to it being below a pre-set threshold. Within the selected five satellites the position is computed for each of the five subsets and the mean is taken to be the correct position. An alternative implementation is to weight the five individual solutions proportionally to the inverse of the DOP. The spread of the five solutions is used as the integrity indicator.

DOP

The Dilution Of Precision describes the geometry involved in the computation of position and is akin to the angle of cut in conventional direction finding. DOP is computed using as a starting point the angles measured at the receiver between the three axes of an orthogonal co-ordinate system and the line of sight to the satellites. Matrix mathematics is used to manipulate the cosine of these angles leading to four components:

a) NDOP, the DOP in the north direction,

b) EDOP, the DOP in the east direction,

c) VDOP, the DOP in the vertical direction and

d) TDOP, the DOP in the time domain.

Combining NDOP and EDOP leads to HDOP, the DOP in the horizontal plane.

Combining HDOP and VDOP leads to PDOP, the position DOP or DOP in 3-dimensions.

Combining TDOP and PDOP leads to GDOP, the Geometric DOP.

In all cases these "combinations" are vector additions. An example of these additions may be found in Section 3.1 of this report.

PDOP4sum and HDOP4sum

PDOP4sum is the sum of the PDOPs for the five groups of four satellites as used in the Simulator FD RAIM architecture. It operates in three dimensions. HDOP4sum refers to the sum of the five values of HDOP associated with the same five groups of four satellites if the user is only concerned with the horizontal plane.

Limiting HDOP4sum

This is the maximum value permitted within any one sample for the best HDOP4sum before the receiver declares a sample to be not available.

Protection Radius

This is the radius of a circle centred on the true position within which circle will be found a given proportion of those of the samples which have the limiting HDOP4sum. The proportion that is most commonly used in this report is 95%. Referring to Figure 1, the Protection Radius is related to all points within the green curve. (All figures appear after Tables 1 to 5 and before the annexes.)

2-D error

For a **single sample** this is the difference between the true and computed positions. Generally, a group of samples is studied together, and the error cited for **that group** is the radius of a circle centred on the true position within which circle will be found a given proportion of the available samples. The proportion that is most commonly used in this report is 95%. Referring to Figure 1, this error relates to all points that are within the red curve and are at the same time left of the line marked HDOP4limit. In some cases it is convenient to show the radius as a function of the proportion: this is referred to as the cumulative error probability distribution. In some cases it is convenient to refer to the proportion, not out of all available samples, but out of all samples whether available or not. In all cases it is the magnitude of the error that is considered, not the direction. Thus, if two samples have errors of 10 metres north and 12 metres south respectively, this report considers only the values of 10 and 12, giving a mean of 11, rather than 1 metres south which would be the mean if the direction had been considered.

Mask angle

The range measurements for satellites with a low elevation angle sometimes have larger than average errors due to a combination of multipath effects and perturbations from the troposphere and ionosphere. Manufacturers therefore choose an elevation threshold: satellites with elevation angle below that threshold are ignored in the receiver software. This threshold is known as the "mask angle".

Available

In the case of the specific architecture CAA selected for this study (defined in Section 2.3), the system is said to be "available" if the receiver tracks 5 or more satellites and the geometry they form with each other and with the receiver is such that the HDOP4sum is less than a limiting HDOP4sum of just over 75. The background to this value is presented in Section 2.4.

Availability

This is a ratio (expressed in percentage terms) defined as A = 100*K/L

where

A	availability
L	the number of samples taken
K	the number of these samples which were available.

Thus, availability is a quantity not exceeding 100%. The associated term "un-availability" is given by (100 – A). There is no distinction in this report between un-availabilities that are predictable (as in predictive RAIM) and un-availabilities that are not predictable. Section 3.1 describes how the availability was computed in this report.

Realism factor

A correction factor which allows for the fact that, within a receiver, the number of satellites that is available to contribute to the navigation solution may be less than expected from a knowledge of elevation angle, satellite health and the receiver's advertised tracking ability. Numerically the factor may be expressed as

F	=	BC

where

B	=	Σ (T_r) summing from r = 1 to r = n
C	=	Σ (Er) summing from r = 1 to r = n
T_r	=	the number of satellites actually tracked for the r^{th} sample
E_r	=	the number of satellites expected to be tracked for the r^{th} sample
n	=	the total number of samples.

Thus F is a quantity not exceeding 1.0. Section 5.2 of this report discusses in detail the way the factor was worked out. Since a receiver sometimes tracks a satellite intermittently it can be useful to use this Realism Factor also to describe the degree of intermittency. In those cases the Realism Factor is worked out for each sample and factors B and C are used differently as follows:

B	=	the number of satellites actually tracked for that sample,
C	=	the number of satellites expected to be tracked for that sample.

Simulator

This refers to the CAA/Nortel GNSS Receiver Performance Simulator. An overview of the Simulator is presented in Section 1.4.

Sample

A single sample is one group of data that relates to the same situation. Thus the "sample" is made up of one latitude, one longitude, one height and one epoch of time, these being the four unknowns that may be solved when four pseudo ranges are provided.

As flown

This term implies that the satellites available for the analysis were those which the recording equipment showed to be available during the flight trials. Thus it does not imply the receiver architecture as flown since for all the analysis the receiver architecture was that selected by CAA.

One satellite off

This term implies that the satellites available for the analysis was one less than that which the recording equipment showed to be available during the flight trials.

Acronyms

BNR	A company in Harlow Essex known as BNR
CAA	UK Civil Aviation Authority
DGPS	Differential GPS
DOP	Dilution Of Precision
FD RAIM	Fault Detection RAIM
FOC	Full Operational Capability
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
LNL	Lambourne Navigation Limited
MF	Medium Frequency (in this report approximately 300 kHz)
MRE	Measured Range Errors
MTBF	Mean Time Between Failures
MTTR	Mean Time To Repair
NANU	Notice Advisory to Navstar Users
NATS	National Air Traffic Services
PC	Personal Computer
RAIM	Receiver Autonomous Integrity Monitoring
TSO	Technical Standard Order
UHF	Ultra High Frequency

1 Introduction

1.1 General

The UK CAA has undertaken a flight trials programme in the North Sea to investigate the use of Differential GPS (DGPS) for instrument approaches to offshore platforms. The flight trials were conducted during 1996, using a chartered S76C helicopter equipped with a DGPS trials installation designed by Cranfield Aerospace Ltd. The airborne system included four GPS receivers, MF and UHF data-link receivers, and an acquisition and processing unit. The processing unit reformatted DGPS data for transmission to the helicopter's area navigation system (RNAV-2) and generated guidance information for display on the cockpit instruments. Two methods of providing differential corrections were adopted for the trials: existing shore-based differential corrections transmitted by marine beacon, and platform-based differential corrections transmitted by a platform station purpose-built for the trial. This enabled the advantages and disadvantages of each system to be assessed. The GPS "carrier phase" positioning technique, via post-processing, was adopted as the "truth" system against which the performance of the real-time DGPS equipment was compared.

In total, seven test flights were conducted involving data-gathering exercises at several offshore platforms, representing low, medium and high multipath environments. The current weather radar approach pattern was flown using DGPS guidance and alternative "DGPS approach" trajectories were also investigated. Further details were given in [1]. The primary output from these trials was data on the flyability of the system and data on multipath. The output included some other data that were secondary for the Flight Trials Programme, but were used for the work described in this current report. Section 2.1.3 describes these secondary data in detail.

1.2 GPS constellation performance

During 1997 Lambourne Navigation Limited (LNL) performed a study for NATS into the satellite unusabilities as notified through the NANU system, covering the first 20 months after GPS was declared to have achieved Full Operational Capability (FOC). In this context the term "unusable" was the term used in the NANUs. The findings included the following:

a) Number of unusabilities (all satellites)	114
b) Mean Time Between Unusabilities (per satellite)	3,140.9 hours
c) Mean Time To Repair	16.2 hours
d) Probability that in a group of 8 visible satellites at least one satellite is having an outage	0.04120

The period analysed was fairly close to the period covered by the flight tests. Preliminary analysis of the satellite performance during 1999 and 2000 suggested that such unusabilities continued to be a feature of the constellation. However, the precise numbers may have altered. Nevertheless it was natural to explore the consequence on off-shore operations of such incidents. Apart from such unusability events, there is also the possibility that the GPS satellite constellation may in the future have fewer satellites. It was these two possibilities that lay behind the study described in this report.

1.3 Methods of investigation

Computer simulations form an important part of evaluation of GPS for helicopter approaches. Their strengths include that:

a) it is possible to explore the consequence of changes in the satellite constellation,

b) different receiver architectures may be examined,

c) one can examine the performance changes second by second,

d) it is possible to repeat a simulated test with controlled conditions.

Weaknesses of computer simulations include that in modelling real-world situations some assumptions inevitably have to be made, and errors in those assumptions lead to errors in the results.

On the other hand, flight tests deal with the actual environment which include:

a) electrical noise,

b) reflections and aircraft shielding which may be modelled only imperfectly,

c) equipment that is imperfect,

d) wind and pilot performance which interfere with the desired flight path.

However, flight trials are costly, and when an unexplained phenomenon is observed it is difficult to repeat the situation in a manner that allows the phenomenon to be examined in detail.

It follows from this that the most cost-effective way to evaluate DGPS for helicopter approaches to offshore platforms is through a combination of simulations and flight trials, in a manner which utilises the strengths of the two techniques and takes into account the associated weaknesses.

1.4 The CAA/Nortel GNSS receiver performance simulator

For several years a GNSS Receiver Performance Simulator has been under development by BNR Ltd (now Nortel Networks). This simulator is a suite of software modules that may run on a standard PC. It allows different satellite constellations and different receiver architectures to be explored for an aircraft positioned at a specific location or moving along specified tracks. It allows different error models to be incorporated and is a powerful tool in that, given the above inputs, it produces as its output a variety of statistical parameters that define the system performance. The simulator development by Nortel is largely documented in a Requirements Specification and a Design Specification. Formal tests are described in a Test Specification. This GNSS Receiver Performance Simulator was developed under funding by CAA. In the following it will be referred to simply as the “Simulator”. Section 2.2 of this report provides details of those Simulator aspects that were particularly important for this study.

1.5 Study objectives

The overall purpose of the work reported in this document was to estimate the impact on performance of a particular receiver architecture if one of the satellites that had been usable during the flight trials had, in fact, been unusable or not there. Later in this document this is referred to as the “one satellite off” condition. This was compared to the “as flown” condition in which the satellites were as used during the flight trials.

The work, therefore, entailed two translations:

a) the performance obtained during the flight trials using an actual receiver and an actual constellation was translated to the performance that would have been experienced with the same constellation but using the CAA-selected receiver architecture,

b) the performance as at (a) above was translated to the performance that would have been experienced with the CAA-selected receiver architecture but using the less complete constellation.

The project objective was to quantify the performance in each case by availability and 2-D precision.

The proposal envisaged an 11 month programme. However, an opportunity to present some results at a conference arranged by the US Institute of Navigation in January 2000 arose. It was felt to be valuable to have some reaction to the results and also to check whether, on the international scene, there were other developments of importance to this simulation work. Therefore, to make use of this opportunity, a limited set of data was provided by Cranfield Aerospace. The data was analysed and the paper [2] presented. A major conclusion in the paper was that the number of satellites usable to the navigation computations in the receiver was below that expected from considerations of the mask angle alone. The only reaction to the paper was that both in the US and in Australia it had been found that, in general, the number of satellites that a GPS receiver was able to use was in excess of that expected. Thus it would appear that the UK had identified a problem that needed to be considered further. Quantification of the tracking performance was therefore added to the project objectives.

NOTE: When pursuing the reasons for the US/Australian findings it appeared that that particular discrepancy arose because one or two of the receiver manufacturers actually used a lower mask angle than they had declared.

2 Background

2.1 The flight trials data forming the input to this study

2.1.1 Criteria for approach selection

Cranfield Aerospace had recorded data from a large number of approaches to five different off-shore platforms, focusing on four of these. These approaches were grouped in seven flights. To constrain the cost of the subsequent analysis it was decided that:

a) this analysis should be limited to ten of the approaches,

b) it should be limited to a study of approaches that had relied on MF corrections,

c) for each approach it should explore only 4 different cases of a satellite being unusable.

The approaches had to be selected in a way that provided a good cross-section of the conditions that were experienced during the flight trials. The selection was therefore based on a number of criteria.

Criterion 1 required that the approaches should be selected from different platforms.

This enabled different geographic locations and different platform characteristics to be sampled.

Criterion 2 required that the number of usable satellites should be different.

The number of usable and visible satellites influences the system availability. At the start of the flight test program there were 24 good satellites, at the end of the programme the number had risen to 26. However, since performance is controlled both by the number of satellites and their location within the orbital planes, it does not follow that the receiver's performance towards the end of the programme should necessarily be better than at the start of the programme.

Criterion 3 required that the approaches should be from different portions of the satellite repeat cycle such that there was a spread of satellite subset characteristics.

This nominal repeat cycle in GPS is 0.99727 days, which is four minutes short of a day. The available flights covered the period from GPS week 852 to GPS week 877, a span of 25 weeks. During this period the constellation had moved $4_{minutes} * 7_{days} * 25_{weeks}$ giving a slip of 700 minutes such that any constellation condition at the end of the period would arise some 11 hours 40 minutes earlier than at the beginning. This is a nominal slip: in reality the orbital period sometimes differs from this. The study therefore required an examination of the way this repetition was at work during the trial so that approaches which avoided the repetition could be used. The timing of each of the flights within this repeat cycle is illustrated in Figure 2 in which the ellipses indicate the portions of the flights from which approaches needed to be selected in order to meet this criteria. Two of these ellipses span two flights indicating that the criteria would be met whether the approach was selected from the end of one flight or from the beginning of the other.

Criterion 4 required approaches to be selected for which the data recorded had good quality for the MF corrections.

This ensured that the receiver had not stopped tracking a satellite due to the absence of the correction. Corrections were sometimes missing, the cause was uncertain but became the subject of a separate study by Cranfield Aerospace.

Criterion 5 required that the approaches to be selected should be as long as possible in duration.

This gave a maximum number of samples.

2.1.2 The approaches selected

Based on the above the following were selected:

a) Flight 2, approach 3.

b) Flight 2, approach 10.

c) Flight 4, approach 1.

d) Flight 4, approach 5.

e) Flight 5, approach 1.

f) Flight 5, approach 11.

g) Flight 6, approach 3.

h) Flight 6, approach 11.

i) Flight 7, approach 3.

j) Flight 7, approach 7.

The numbering scheme is based on that used in the Cranfield Aerospace documentation.

2.1.3 The data available for these approaches

Section 1.1 referred to "primary" and "secondary" data from the flight trials. These data are identified in Figure 3 which shows those elements of the test flight equipment that were most relevant to this study. In that figure each box is numbered in the lower left corner.

Cranfield Aerospace's Recording and Analysis System (box 4) received inputs from a GPS receiver (box 2) which was differentially corrected by MF corrections emitted by coastal reference stations and received by a MF receiver (box 1). Apart from the Primary Outputs (boxes 8 and 9) referred to in Section 1 the equipment also provided data that was secondary to the purpose of the trials, but nevertheless used for this study. This comprised:

a) For each of the approaches a computer file (box 5) known as "mfcorr.txt" containing the differential corrections received and decoded at the aircraft. Each correction was tagged with the relevant GPS time and showed clearly the satellites to which the corrections applied.

b) For each of the approaches a computer file (box 6) known as "tracked.txt" listing the satellites that were tracked by the airborne GPS receiver. This data was also tagged with GPS time.

The Cranfield analysis included:

c) 2-D error statistics (box 7) for the combined programme and relating to the MF corrected position data.

Apart from the boxes cited here, the flight trials installation included various other equipment such as a UHF receiver and other GPS receivers referred to in Section 1.1. Those items of equipment were not used in this study and are therefore not referred to further.

2.2 The CAA/Nortel simulator

The Simulator referred to in Section 1.4 has a number of operating modes. These include:

a) taking samples at random in space and time within a region geographically bounded,

b) taking samples at a fixed geographic point but at regular time intervals,

c) taking samples along a track at regular time intervals.

Central to all modes of operation there is a main loop that is exercised repeatedly, once per sample, until the simulation is halted. This loop is illustrated in Figure 4 in which all boxes are numbered in the lower left corner. In this loop the inputs comprise latitude, longitude, height and time that are considered to be "truth" (box 2) for that sample. The positions of the satellites at that time are computed from the Keplerian parameters contained in an almanac (box 1) which also gives the health of the satellites. The satellite elevations are then computed relative to the "truth" position. Those satellites that are below a pre-set mask angle (box 3) are then excluded leaving a list (box 4) of satellites that are visible to the receiver and healthy.

It is possible for the Simulator operator to inhibit (box 5) one or more of these satellites for the duration of a Simulator run. It is also possible to cause satellites to be ignored intermittently to simulate temporary failures (box 6) in the space segment. The rate at which satellites are ignored may be controlled by specifying the satellite Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR).

These two aspects, represented by boxes 5 and 6 in Figure 4, modify the Satellite Visibility List to a Satellite Tracking List (box 7) that represents the satellites that are then available to the “Performance Algorithm” (box 9). As indicated in Figure 4 this box also receives an input from the Range error model (box 10). The performance algorithm delivers availability data and 2-D error data for statistical purposes (box 12), and the situation for each sample may be followed on a monitor (box 11). A text file is generated (box 8) providing a record for every sample: this details the GPS time and the errors for that sample plus various other data that are of no concern here.

Additional details of the performance algorithm are illustrated in Figure 5. From the Satellite Tracking List a subset of satellites is selected (box 22) based on the algorithm of the chosen receiver architecture. This generally includes an analysis of the geometries expressed through the Dilution Of Precision (DOP) (box 21). The DOP for one subset is compared to that for all the other possible subsets and the best of these is compared to a pre-set threshold above which a subset is considered to be unavailable for navigation purposes.

For each of the satellites that are so selected the true distance between satellites and receiver are computed (box 23). To each of these distances an error is added (box 24) from the chosen error model, and a time error is also incorporated such that erroneous pseudo ranges result. These are then presented to the algorithm that estimates receiver position (box 25) from a set of pseudo ranges. The difference between this computed position and the “truth” position is then calculated and represents the 2-D error (box 26) for that particular sample. In DGPS mode the error model provides range errors selected at random from a population with a Gaussian distribution and controlled standard deviation. The range error for any satellite is independent of the range errors for all other satellites. The randomness is controlled from a random number seed that is under operator control. However, when this seed has been used once, a new seed that is used for subsequent random selections is generated. Only the first seed is selected by the operator.

When all required samples have been examined a statistical analysis is performed that leads to the cumulative probability distributions for 2-D, 3-D and height errors.

A related aspect of the Simulator development was a software module referred to as the Measured Range Error (MRE) Model. The input data were to be the ephemeris, the pseudo ranges and the true position, all of which were measured by the aircraft receiver during the flight tests. The Simulator module would then translate the data into measured range errors which, in turn, would be used in the Simulator instead of the theoretical range error models that are normally used. These could then be used to explore a variety of conditions. However, while planning this module, Nortel indicated that they wished to withdraw from the Simulator project. The task was then altered to one in which Nortel would develop and deliver a specification for the module, but the module itself would not be a **formal** deliverable. Nevertheless, Nortel decided to develop the module, carry out some limited tests and deliver it informally.

While the specification delivered is a useful start, the tool itself must be considered as incomplete until it has been tested further. This is in sharp contrast to other parts of the Simulator, particularly the main loop referred to above, that may be considered to have achieved a high degree of maturity. Since Nortel was unable to provide long term support for the Simulator and, in particular, dealing with problems that might emerge with the MRE module, the use of this module was effectively curtailed.

The Simulator has been developed over a considerable time period and currently exists in different versions, two of which have been used in this work. Their characteristics differ slightly, particularly in terms of the range of MTBF, MTTR, DOPs permitted and the format of some of the output data. However, the principles are the same. If a reader seeks to reproduce the simulations discussed he should first consult with the CAA about the version differences.

2.3 The selected receiver architecture

One of the benefits of the Simulator is that it allows a number of different receiver architectures to be examined and compared. The term "architecture" is defined in the Glossary earlier in this document.

The architecture selected for this study was one having the following characteristics:

i) there should be no external aiding from barometric altimeter, inertial system or clock,

ii) an element of Receiver Autonomous Integrity Monitoring should be provided to the extent of providing an integrity assured solution (FD RAIM), therefore 5 satellites would need to be usable and tracked,

iii) the geometry should be good enough to provide a horizontal protection radius of 100 metres at the 95% level of confidence,

iv) the position estimate should be based on the un-weighted average of the 5 estimates possible from 5 satellites,

v) each sample should have its estimate derived without any regard to the positions estimated from prior samples (i.e. no smoothing),

vi) the mask angle should be fixed.

This particular architecture was chosen since it was considered to be the least attractive of the known alternatives for providing FD RAIM. In particular, it was considered that some improvement would result if the outputs were smoothed, if each solution was weighted according to its DOP, or if all satellites in view were used simultaneously rather than the subset of five satellites providing the best geometry. The precision improvement possible with these other alternatives was estimated to be of the order of some 10%, the improvement in availability much less. Thus, it was considered that whatever the results achieved with the chosen architecture, the alternatives would be slightly better, and some estimate of the improvement in availability would be possible at the end of the study.

All the measurements performed during the flight trials had been recorded using a receiver which claimed to be able to track 11 satellites simultaneously. Since there were never more than this number visible during the trials, the recording equipment was in fact an all-in-view installation from which satellites could be selected for analysis by the Simulator to represent the appropriate architecture.

There are therefore two different receiver architectures involved in this study:

a) the architecture for the receiver actually flown on the aircraft,

b) the architecture for the theoretical receiver whose performance was under investigation.

It is the Simulator that is being used to translate the performance from the former to the latter of these two.

2.4 DOP and protection radius

In GPS, as in most navigation systems, the 2-D errors vary. Many errors are too small to be of concern to aviation, but some errors are large enough to present a potential hazard. It would be useful if the pilot could be warned when the error is in the second category i.e. that it exceeds a pre-determined value. This would give him the option of disregarding the navigation output (relying, for instance, on dead reckoning while the condition lasts), or possibly allow other users in the area to be warned. This would, in effect, protect the aircraft against proceeding outside a circle centred on desired position and with radius equal to that pre-determined value. In this paper this pre-determined value is referred to as the “protection radius”. The difficulty is, however, that the airborne receiver does not know what the instantaneous error is. However, the receiver does know the DOP it is working with and the question therefore arises as to whether the DOP can somehow be used to estimate the 2-D error.

For any given sample the 2-D error is controlled by two factors that are independent of each other. These are:

a) the geometry as quantified by the HDOP, and

b) the combination of errors associated with each individual satellite range measurement.

For any given DOP all of the range errors can be low, in which case the 2-D error will be small. For the same DOP all of the range errors can be high. This probably leads to a much larger 2-D error. More generally, there may be a mixture of some low range errors and some larger range errors. Because the 2-D error is heavily influenced by the range errors it follows that for any single sample the 2-D error is not proportional to DOP. However, there is a probabilistic relationship in that if the mean were taken of many 2-D errors all sharing a low DOP, that value would be smaller than if a mean were taken of many 2-D errors all sharing a high DOP.

This probabilistic relationship may be expressed in its simplest form through the expression

$$E_{2\text{-}D} = \sigma * HDOP * k \qquad (1)$$

In which

HDOP is the instantaneous HDOP for the architecture employed,

$E_{2\text{-}D}$ is the 2-D error which contains a proportion “p” of all samples having that particular instantaneous HDOP,

σ is the range error at the one standard deviation level. In DGPS this s describes the residual error after differential correction.

k is a constant which enables $E_{2\text{-}D}$ to be cited at the level of confidence associated with “p”. The value of “k” depends on the selected receiver architecture and on the statistical distribution that describes the errors associated with the particular DOP.

CAA selected 0.95 as the preferred value for the probability “p”.

There was no obvious way of determining the form of the statistical distribution of the probabilistic relationship from theoretical considerations, nor was it necessary for this particular study. All that was required was to establish the 2-D error that contained 95% of the 2-D errors for each of several values of the HDOP4sum.

To this end the simulator was set up to run a sequence of time/position combinations known to contain a wide range of values of HDOP4sum. For each value of HDOP4sum the 2-D error was extracted. The run was then repeated a number of times changing only the random number seed such that different 2-D errors arose for the same values of HDOP4sum. Figure 6 plots the situation after 10 different values of HDOP4sum had been explored with 10 different random number seeds. Thus, for each of these DOPs (plotted along the horizontal axis) there are 10 values of 2-D error (plotted vertically). This sequence was continued until 100 different seeds had been used, giving 100 values of 2-D error in each case. The process was also extended to 18 different values of HDOP4sum giving a more even distribution along the horizontal axis. For each of these 18 DOPs that value of 2-D error which contained 95% of the errors was then extracted. These values were plotted in blue in the diagram shown in Figure 7. When these points are joined together they form a very jagged line. Therefore, the points were submitted to a least squares analysis such that the best straight line was derived. In Figure 7 this is plotted in red. This line has the form

$$E_{2\text{-}D} = a * \text{HDOP4sum} + b. \qquad (2)$$

It is more general than (1) since it does not presume that the line passes through the point given by

$$(E_{2\text{-}D}, \text{HDOP4sum}) = (0, 0).$$

For these simulation runs the almanac was that valid on one of the flight days and the location was taken from that off-shore area where the flights took place. At the time of this analysis the residual range errors after differential corrections were not known, the simulations used a value of 10 metres for σ and the analysis was later refined when an opportunity arose for determination of the actual range errors. This led to

$$E_{2\text{-}D} = 1.506 * \text{HDOP4sum} - 13.626 \qquad (3)$$

such that a 95% protection radius of 100 metres is achieved by disallowing geometries for which HDOP4sum exceeds 75.45.

While this value is the one that is used in the analysis it should not be concluded that the 100 metres at 95% is precisely the value that always corresponds to the HDOP4sum of 75.45. There are several reasons for this:

a) It relates only to situations in which the value of σ is as during the flight trials.

b) It assumes the same σ on all satellites and that is not always a good assumption.

c) Only 100 different random number seeds were used.

d) The resulting plot gave an un-even line which was smoothed out by means of a best straight line fit analysis. There is not strictly any proof that the line should be straight, and an alternative approach might have been to explore more data points close to the critical limit (say from a DOP of 70 to a DOP of 80) and then to ignore the other data points.

3 Availabilities

3.1 Method followed to determine the availability

As satellites rise and set the DOP varies with time, and the Simulator monitor shows this variation. However, it was observed that the GPS receiver used during the trials was not tracking all of the satellites that were visible and healthy, despite there being less than 11 such satellites and the receiver being advertised as being able to track

11. It was also observed that some of the differential corrections were absent from the MF data stream. To ensure that the Simulator's performance accurately reflected the performance that would be experienced with the selected architecture, it was therefore necessary to exclude the satellites that were absent from the GPS receiver's tracking list and the satellites that were absent from the MF correction list from the simulation process. To this end the time-tagged MF corrections (box 5 in Figure 3) and the time-tagged satellite tracking list (box 6 in Figure 3) were merged manually to provide one time-tagged list of satellites usable to the DGPS operation. Satellites absent from this list were then excluded using the Simulator's satellite inhibit feature (box 5 in Figure 4).

Each approach was then divided into a number of "sections" for which there was a stable set of satellites present in this merged list. Such a section comprised one or more samples. As soon as a new sample used a different subset of satellites a new section began. An example of this may be seen in Table A1.1 of Annex 1. There the first 11 samples used satellites PRN01, 04, 05, 07, 14, 15, and 29. These samples comprised data Section 1. PRN05 was then absent for the next 5 samples that represented Section 2 after which PRN05 re-appeared again for Section 3 etc.

For each approach the situation was then explored using the Simulator which was run a number of times, each run comprising the sections sharing a common set of satellites. These runs used as input parameters the almanac for the week of the flight, the time as derived from the time-tagged Cranfield files, and the mask angle known to feature in the receiver employed for the flight trials. The satellite MTBF/MTTR feature (box 6 in Figure 4) was turned off such that intermittent unusabilities did not arise within the simulation.

Ideally one would explore the five satellite combinations ("as flown" plus four different satellites "off") for each sample. This would be extremely time-consuming since an approach typically contained 200 samples. During an approach an aircraft altered its position only by a few miles, and over such short distances the changes in satellite visibility and geometry would be insignificant. It was therefore felt appropriate to approximate the aircraft position to that of the platform throughout the approach.

Figure 8 illustrates how the various boxes from Figures 3 and 4 were used together for this determination of availability.

The ten approaches identified in Section 2.1.2 were examined one at a time. Annex 1 shows an example of how the analysis proceeded for Flight 4 Approach 5. Table A1.1 illustrates the list of sections which was used in two ways:

a) It was first used to examine the performance of the selected receiver architecture assuming all the satellites that featured in the above list were operating correctly. This was called the "as flown" condition. The Simulator was set up for the platform location, stepping forward in time for the duration of the selected approach and allowing those satellites to be used by the Simulator. For each sample the DOP was extracted and the sample was considered "available" if it had 5 usable satellites and the HDOP4sum was below the threshold of 75.45. This value of DOP had been determined as the value that corresponded to a horizontal protection radius of 100 metres at the 95% level of probability. Section 2.4 describes how this threshold was determined. In Annex 1, Table 2 illustrates the number of satellites usable (column 5) and the DOP extracted (column 6). Accordingly, for each sample it is shown whether it was available (termed "good" in the Table) or not available (here termed "bad" in the Table). This information is provided in columns 7 and 8.

b) The list was then used to examine the performance resulting if a single satellite among those in the list was not usable. This was called the “one satellite off” condition. This was achieved by using the Simulator’s inhibit satellite feature (box 5 in Figure 4). Again the DOP was extracted for each sample and the sample was again considered “available” if the HDOP4sum was below the threshold of 75.45. This was performed with four different satellites set unusable. In Annex 1, Table A1.3 shows the working in one of these cases, namely with PRN07 turned off.

Unfortunately the Simulator does not have the facility to directly analyse CAA’s chosen architecture as this requires access to HDOP4sum and this is not available as an output. However, PDOP4sum is available and HDOP4sum is related to PDOP4sum through the relationships:

$$(HDOP4)2 = (PDOP4)2 - (VDOP4)2 \qquad (4)$$

$$HDOP4sum = HDOP41 + HDOP42 + HDOP43 + HDOP44 + HDOP45 \qquad (5)$$

$$PDOP4sum = PDOP4_1 + PDOP4_2 + PDOP4_3 + PDOP4_4 + PDOP4_5 \qquad (6)$$

The subscripts denote the various combinations of four satellites possible in a subset of five satellites. From these relationships it follows that HDOP4sum is never greater than PDOP4sum. It could therefore be concluded that if PDOP4sum (the value which is normally computed by the Simulator in the FD RAIM architecture and displayed on the Simulator monitor) was less than 75.45, then HDOP4sum (the value required by the analysis) would also be less than 75.45. In those cases where PDOP4sum was shown to exceed 75.45, a more detailed analysis was required to find the five constituent components in equation (5) above. This analysis required HDOP4 to be determined for each of the 5 subsets of 5 satellites.

When the “as flown” performance had been determined, the situation was similarly explored with additional single satellites set unusable. The satellites selected were in each case those four having the lowest PRNs featuring in both files of the Cranfield data.

3.2 Magnitude of the availability

The overall results were:

For the “as flown” condition

(a)	Number of available samples	2136	
(b)	Number of samples not available	132	
(c)	Total number of samples	2268	
(d)	Availability	94.18%	(100*2136/2268)

For the “one satellite off” condition

(e)	Number of available samples	7649	
(f)	Number of samples not available	1423	
(g)	Total number of samples	9072	
(h)	Availability	84.31%	(100*7649/9072)

The low availability is mainly caused by Flight 5 Approach 11 for which there was an unusually low number of satellites being tracked (see further under section 3.3 below). Note that the "one satellite off" condition was explored for 4 different satellites. The result is that the "one satellite off" condition has 4 times as many samples as the "as flown" condition. Thus 2268*4 = 9072.

3.3 Variations between the approaches

Tables 1 and 2 provide the availability per approach, Table 1 describing the "as flown" condition and Table 2 describing the "one satellite off" condition. The Tables show that the availability for Flight 5 Approach 11 was only 57.89% as flown and only 0.82% with one satellite off. In all other cases the availability was 100% or close to 100%. The variation between the approaches demonstrates clearly the importance of ensuring that any analysis such as this must include a wide variety of situations, thus justifying the wide range of criteria outlined in Section 2.1.1. It needs to be stated that since Flight 5 Approach 11 gave such a very poor result, the input data were submitted to a special scrutiny by Cranfield Aerospace who confirmed that there were no signs of any faults with the data or with the recording of that data.

3.4 Duration of the unavailabilities

An analysis of the duration of each period of unavailability was carried out. This showed that the majority of the periods had a short duration, but that there were a few periods with a duration of the same order as the duration of the approach.

The unavailability durations fell into three groups:

a) There were 41 cases of an unavailability period with a duration of between 1 and 7 seconds,

b) There were 3 cases of an unavailability period with a duration of between 33 and 38 seconds,

c) There were 4 cases of an unavailability period with a duration of between 300 and 302 seconds.

The total of 48 cases of a continuous unavailability is too small a sample size to conclude anything useful about the statistical distribution of the unavailability duration. However, while the shorter durations may represent little more than a nuisance, the longer duration cases would seriously reduce the value of the system.

3.5 The cause of the unavailabilities

Table 3 shows the split between those cases where the unavailability was due to "too few satellites" and those cases where it was due to bad geometry. It was clearly the number of satellites that was the dominant cause, accounting for 1343 cases out of 1423. This corresponds to about 94%.

3.6 The effect of altering the selected receiver architecture.

The study was to focus on one particular architecture and one particular protection radius. However, the following observations could be made about the effect if one of these was altered:

i) Altering the receiver architecture.

Five satellites are needed for all unaided types of FD RAIM. This applies whether the selection criteria is GDOP, PDOP, HDOP or VDOP, whether or not a weighted approach is used or sequential smoothing allowed. Five satellites are also required for an all-in-view approach. The 1343 cases of "too few satellites" as shown in Table 3 would therefore still pose a problem. The extent to which the 80 cases of "bad geometry" might change with a different architecture would require additional analysis. The "too few satellites" cases would be overcome by aiding (altitude or clock), though the resulting geometry would need to be checked. It would probably also be overcome by an HDOP approach using 4 satellites if the integrity protection relied on the differential corrections. However, one would need to check that the geometry of the 4 satellites was adequate.

ii) Altering the protection radius.

Alteration of the horizontal error to be protected will affect the 80 cases of bad geometry. The largest HDOP4sum encountered was 310.6. Using equation (3) in Section 2.4 this value of DOP corresponds to a horizontal protection error of 454.1 metres. If the required protection error was set at 455 metres there would therefore be no "bad" cases, other than those caused by "too few satellites". Increasing the protection error beyond this value would not improve the availability further.

4 Precision

4.1 Method followed to determine the precision

The 2-D errors clearly depended upon the range error model. The test flight results had shown that during the flights:

i) the contribution to the 2-D error from airframe-generated multipath was generally less than 1 metre,

ii) the contribution to the 2-D error from platform-generated multipath when the corrections came from the MF coastal station (rather than from the platform mounted reference station) was in most cases not measurable.

The level of systematic errors was therefore sufficiently small to justify the use of a range error model that comprised range errors taken at random from a population with Gaussian distribution and controlled standard deviation. To determine this standard deviation, the simulator was run for the region and time of the tests providing a very high sample size and using a first guess standard deviation. This generated a population of 2-D errors from which the 95% and 50% errors were extracted. These values were compared with the 2-D errors actually determined during the flights (box 7 in Figure 3). The guessed standard deviation was then scaled so that the simulated 95% and 50% 2-D errors gave a good fit to those measured. This yielded a standard deviation for the range errors of 4.07 metres. It was noted that this value was somewhat higher than had been expected. Various investigations were carried out and it was concluded that, apart from components (i) and (ii) above, there was a contribution from latency in the MF corrections (possibly from weather), some interference, in addition to general receiver noise. The 4.07 metres was therefore the value accepted as the best estimate. In these simulations the receiver architecture was an all-in-view type since that corresponded most closely with the receiver actually used.

In the manner outlined above, the 2-D error statistics from the flight trials generated the range error model from which range errors were fed to the performance algorithm. This is illustrated in Figure 9.

Reference has been made earlier to the output file (box 8 in Figure 4) generated within the main Simulator loop. This was not used for the determination of either availability or the Realism Factor. However, it was used when establishing the precision.

Since a particular approach was represented by sections from different simulation runs, the analysis of precision required that the error data also be extracted from those sections. Two Pascal routines were written to extract the relevant data. This required the following actions:

a) selection of the correct files,

b) selection of the appropriate records within these files,

c) selection of the latitude and longitude errors within the record,

d) forming the 2-D error from them (Pythagoras).

These errors were then placed in a histogram of errors having a constant bin width in metres. When all errors had been extracted the cumulative error distribution was determined. This error distribution was copied to an Excel worksheet and the Excel "Chart Wizard" was then used to provide the plot. Figures 10 – 12 were derived in this way. They are discussed in the next section.

4.2 Overall results

All the individual errors from the 10 approaches were analysed together and the cumulative probability distributions presented in Figures 10 and 11. It is particularly Figure 10 that illustrates the deterioration when one of the satellites is set unusable. The key statistical parameters were:

Out of the 2136 available samples in the "as flown" condition, 95% had an error less than 15 metres. Out of the 7649 available samples in the "one satellite off" condition 95% had an error less than 16 metres. Although this seems like a small deterioration it should be remembered that the availability in the "one satellite off" condition was much smaller than in the "as flown" condition.

As noted in Section 3.2, the availability was less than 95%. Thus the 95% level could not be determined for either condition except in terms of "out of the available samples". A better way of expressing the deterioration is therefore to use the error as an independent variable and the percentage as the dependent variable. This gave the following:

i) An error of 10 metres contained 79.7% of all samples in the "as flown" condition; this fell to 69.3% in the "one satellite off" condition.

ii) An error of 20 metres contained 92.4% of all samples in the "as flown" condition; this fell to 82.5% in the "one satellite off" condition.

iii) An error of 30 metres contained 93.9% of all samples in the "as flown" condition; this fell to 84.0% in the "one satellite off" condition.

iv) An error of 40 metres contained 94.2% of all samples in the "as flown" condition; this fell to 84.2% in the "one satellite off" condition.

v) The largest error found in the "as flown" condition was 39.5 metres; the largest error found in the "one satellite off" condition was 62.6 metres.

Figure 12 plots the cumulative probability error curve up to 46 metres which is the error beyond which there are only 10 samples. Plotting the curve beyond this point is considered to be unreliable.

5 The Realism Factor

5.1 Introduction

Performance evaluations using simulations generally assume that a receiver will "use" all satellites that:

a) are set “healthy”,

b) have an elevation angle above the receiver’s mask angle.

In this connection “use” means that for all-in-view receivers the pseudo range measurement is incorporated into the navigation solution. For those receivers which elect to employ only a subset of the satellites, “use” means that the receiver will incorporate that pseudo range measurement if the satellite satisfies the geometric selection criteria.

Analysis of the flight trials data showed that the above assumption gave an optimistic result. This is illustrated in Annex 3. Figure A3.1 shows that PRN19 was being tracked intermittently. However, the elevation angle as measured at the platform was around 19 degrees and the MF corrections were always present for the satellite. There were many cases in the recorded data of such intermittent tracking, or indeed no tracking at all, even though the satellite was high in the sky and the almanac declared the satellite to be usable.

To obtain a more realistic estimate of performance one must therefore, somehow, incorporate a correction factor which, in this report, is described as the “Realism Factor”. This is the probability that a healthy satellite above the mask angle is usable when computing position or integrity. The Simulator was not designed to allow for this. However, the Simulator has a “random satellite failure” mechanism (box 6 in Figure 4) which could be utilised.

Let

$MTBF_a$ = The actual Mean Time Between Failures per satellite.

$MTTR_a$ = The actual Mean Time To Repair a satellite

$MTBF_s$ = The Mean Time Between Failures per satellite input to the Simulator

$MTTR_s$ = The Mean Time To Repair input to the Simulator

F_r = The realism factor.

These may be related through the relationship

$$MTTRs/ MTBFs = (1- F_r) + MTTR_a / MTBF_a \qquad (7)$$

Thus, if all satellites are always present in the tracking list this gives $F_r = 1$ and $(1- F_r) = 0$ such that the MTTR and MTBF values that are input to the simulations will be the actual ones applying to the satellites.

As F_r deviates from 1 the value of $(1- F_r)$ increases in importance. Suppose the satellite MTBF and MTTR were 10,000 hours and 10 hours respectively. This would represent a considerable improvement relative to the values during the months immediately following declaration of FOC as referred to in Section 1.2. The ratio MTTR/MTBF would then be 0.001. Therefore if the realism factor fell from unity to 0.999, the importance of this factor would equal that of the satellite failures. This shows it to be a very important parameter. If F_r were to fall to 0.9, its importance would totally dominate the importance of failures in the space segment. The following sections now explore this factor in light of the flight trials data recorded.

At the time of the flight trials, Cranfield Aerospace set the elevation mask angle in the tracking receiver to 5 degrees. For this study it was the Simulator that was used to determine whether a particular satellite at a particular instant was below or above such a mask angle. If a satellite is several degrees above or below the mask, then this will be clear. However, the Simulator displays angles to an integer value and this displayed value cannot be used to determine the satellite elevation relative to the mask in those cases where the elevation is close to the mask. However, apart from extracting the current elevation angle from the displayed value, one can also derive information from the presence or absence of a satellite in the list of satellites used by the Simulator, as displayed on the Simulator monitor (box 11 in Figure 4). This is because the Simulator will not display satellites that are below the mask angle.

5.2 **Method followed to determine the Realism Factor**

The time-tagged satellite tracking list (box 6 in Figure 3) was compared with the satellite visibility list (box 4 in Figure 4) except that in both cases those satellites that were absent from the time-tagged MF correction list (box 5 in Figure 3) were ignored. The differences were extracted and the Realism Factor computed according to the definition given in the Glossary earlier in this report. Figure 13 illustrates how the different modules were used.

Annex 3 illustrates the method using Flight 2 approach 10 as the example. This approach started 33 seconds after 1426 UTC and ended 8 seconds after 1429 UTC providing 156 samples. There were MF corrections for 7 satellites, these being PRNs 17, 19, 21, 22, 23, 28 and 31. All of these were sometimes tracked, but sometimes one or more were missing. Table A3.1 of Annex 3 shows the situation for the early and late parts of the approach. For example at 26 minutes and 34 seconds PRN21 and PRN31 were missing such that only 5 out of 7 satellites featured in the tracking list. The last column of Table A3.1 shows for each sample the ratio of satellites actually tracked to those expected.

At the start of this report the Realism Factor was defined numerically as

F = B/C

where

B = Σ (T_r) summing from r = 1 to r = n

C = Σ (Er) summing from r = 1 to r = n

From the Annex one has

B = 982 and

C = 1092 such that

F = 0.899 for that approach.

5.3 Magnitude of the Realism Factor and variation between the approaches

The results were obtained for the individual approaches and these were then added to give an overall result. The individual and overall results were:

Flight	Approach	Tracked	Expected	Factor
2	3	1264	1484	0.852
2	10	982	1092	0.899
4	1	1528	1936	0.789
4	5	1912	2235	0.855
5	1	1737	2073	0.838
5	11	1394	1824	0.764
6	3	1166	1260	0.925
6	11	2378	2439	0.975
7	3	1874	2260	0.829
7	7	879	1008	0.872
Overall		**15114**	**17611**	

This gives an overall Realism Factor of 15114/17611 = 0.858

5.4 An estimate of the consequence of this Realism Factor on availability

To estimate the effect this Realism Factor would have on RAIM availability the Simulator was used as illustrated in Figure 14. Two simulations were run for which the following input conditions were common:

a) Sampling was done with random time and random location within the world wide region.

b) The almanac was GPS_RTCA.547, that which page 15 of the TSO [3] demands shall be used for all evaluations. It has 21 satellites.

c) Evaluation was for day 0 in week 547 which is the one referred to in [3].

d) No barometric aiding.

e) Mask angle at 7.5 degrees. Page 29 of [3] would allow lower angles "provided the applicant develops acceptable test conditions and supporting analysis to substantiate use of the desired mask angle." It is not clear to LNL to what extent manufacturers have done that yet.

f) Number of samples = 100,000.

g) Receiver architecture was all-in-view based on HDOP. This architecture is assumed by LNL to be the one that most commonly will be used in the future.

For these simulations the inhibit feature was not used but the MTBF/MTTR feature was used to cover both satellite failures and the Realism Factor. It was assumed that the MTBF and MTTR values referred to in the introduction would improve as new satellite designs are implemented. Just what values would be realistic for the future is not known, but for the purpose of this analysis it was assumed that 10,000 hours for the MTBF and 10 hours for the MTTR might be achievable. For random time and place simulations the precise numbers do not matter, but the ratio between them does. Let:

$MTBF_a$ = The actual Mean Time Between Failures per satellite.

$MTTR_a$ = The actual Mean Time To Repair a satellite

$MTBF_s$ = The Mean Time Between Failures per satellite input to the Simulator

$MTTR_s$ = The Mean Time To Repair input to the Simulator

F_r = The Realism Factor.

These are related through the relationship

$$MTTR_s / MTBF_s = (1- F_r) + MTTR_a / MTBF_a$$

Substituting the realism factor of 0.858 as found in Section 5.3 into this relationship gives:

$$MTTR_s / MTBF_s = (1- F_r) + MTTR_a / MTBF_a$$

$$= 0.142 + 0.001$$

$$0.143$$

This shows that the situation is totally dominated by the tracking problems in the receiver, rather than unavailabilities in the space segment.

The permitted boundaries of the MTBF and MTTR in the particular version of the Simulator used for this exercise allows this to be achieved with MTBF = 1000 hours and MTTR = 143 hours. When ignoring the Realism Factor such that only the space failures are considered MTTR = 1 hour. (This is a factor of 10 times better to match the 10 times reduction in MTBF required to get the simulation ratio right.)

The results of this comparison yielded the following results:

a) The simulation run ignoring the Realism Factor (i.e. assuming the receiver tracked all the satellites) gave a RAIM availability of 99.58% .

b) The simulation run including the Realism Factor gave a RAIM availability of 85.06%.

This is a significant drop in availability. While the TSO [3] covers only Supplemental Navigation Equipment for which a reduced availability could be said to be only moderately important, such a reduction for Sole Means or Primary Means Navigation Equipment would be very serious. It is therefore clear that there is a real need to understand the nature of the phenomenon that gives rise to the reduced availability to ensure that it is not carried into Sole Means systems.

5.5 Directionality of the low Realism Factors

It is relevant to consider whether the worst cases of loss of satellites are linked to a particular direction relative to the helicopter. Figure 116 of [5] illustrates that a similar effect had been observed by Cranfield Aerospace when the Carrier to Noise Ratio had been plotted as a function of direction. It was therefore decided to plot the location of the satellites missing in the tracking list on a sky plot. "Missing" means here those satellites that were tracked for less than 10% of the time within an approach. There were five cases within this category, these were:

PRN24 in flight 2 approach 03 was tracked for 2.5% of time
PRN21 in flight 4 approach 01 was tracked for 3.3% of time
PRN25 in flight 4 approach 05 was tracked for 0.0% of time
PRN31 in flight 5 approach 11 was tracked for 1.3% of time
PRN14 in flight 7 approach 03 was tracked for 0.0% of time

A typical approach was conducted on a nominally fixed track and lasted only a few minutes. During such a short period the direction to any given satellite relative to true north would have varied by only a few degrees. For the purpose of such a sky plot it was therefore considered reasonable to assume the satellite position to be constant within an approach. Furthermore, although the approach track was nominally along a straight line the wind and the pilot's limitations in track maintenance would have led to some aircraft heading and attitude alterations. These two considerations justify marking the satellite's position during the approach as a single point, rather than as a line.

For each of these 5 combinations the Simulator was used to compute the elevation angle and the azimuth angle relative to true north. The track for each approach was computed as the vector from the latitude and longitude of the first sample to the latitude and longitude of the last sample using the truth file generated during the flight trials. The azimuth relative to north was then transformed into azimuth relative to the aircraft's nose. Figure 15 shows the resulting sky plot.

The figure shows that the 5 cases are all located within a 37 degree arc of azimuth. This is the shortest of the two arcs bounded by the green lines in Figure 15. If the cause of the difficulty was random, then the probability of one specific satellite being in this particular sector would be 0.103 (i.e. 37/360); the probability that 5 specific satellites are in this sector = $0.103^5 = 1.15*10^{-5}$. This is clearly a very unlikely situation. Comparison may again be made with Figure 116 of [5] which shows a similar difficult sector on the port side of the aircraft. However, in comparing the information provided by Figure 116 of [5] with Figure 15 of this current document it must be borne in mind that the former indicates signal to noise loss of the satellites that remained tracked, whereas Figure 15 indicates the whereabouts of satellites that were not tracked. One is therefore justified in claiming that this has a systematic cause. Since the tail rotor was on the side of the tail with the GPS aerial, this rotor must be considered a prime candidate for the cause. The maximum elevation angle for these five cases was 20 degrees.

Figure 15 may be compared with Figure 16 which shows the sky plot for all satellites being used in the 10 approaches. The satellites with higher values of Realism Factor are clearly distributed much more randomly.

5.6 The Realism Factor and predictive RAIM

Predictive RAIM is sometimes used to determine whether it is safe for an aircraft to fly with GPS as a navigation sensor. Predictive RAIM assumes that all satellites that are set healthy in space and which are above the receiver's mask angle will be usable by the receiver's navigation calculation algorithm. The above analysis has shown this to be an optimistic assumption. The conclusions from such predictive RAIM will therefore also be optimistic. Therefore this raises the question: how safe is predictive RAIM? It is beyond the scope of this study to answer this question.

5.7 Procedures

Where the number of satellites a receiver actually tracks is reduced so much that the receiver is unable to provide a position output one needs to look for either technical solutions or for procedures that can maintain safety. Section 3.6 indicated potential technical solutions. The following are some procedures that may have potential.

a) Checks before take-off.

It may be possible for a helicopter operator, or the Air Traffic Service provider to procure a simulator that can explore the likely performance along the proposed track. In effect this is what predictive RAIM is intended to do. It would, however, need to be adapted to incorporate a two-stage mask angle which can deal with the sector defined in Section 5.5 above. However, such a simulator may be unacceptably costly to a helicopter operator. Ensuring safety would require the procedure to explore the track in advance of flight. Such an exploration would identify what time periods had to be avoided. Clearly, this procedure cannot readily deal with unexpected extra satellite outages, but this difficulty is shared with ordinary predictive RAIM. Delaying (or advancing) take-off time may be operationally undesirable and impractical.

b) Transits dead reckoning.

Many of the outages identified during this analysis were of short duration. A simple procedure is therefore that if this loss arises unexpectedly during a transit then the aircraft simply coasts through the periods with dead reckoning on the assumption that the outage will last only a few minutes after which the situation will recover. A subset of this would be to arrange that the receiver should not inhibit display of anything, but show position and declare there is no integrity check.

c) Transit Zig-zag tracks

An alternative may be for the aircraft to alter its heading by an amount sufficient to bring one of the “lost” satellites out of the difficult azimuth sector. This, however, takes the aircraft away from its intended track and the extent to which ATC would be able to clear such a deviation clearly depends upon other traffic. Where constant radials are flown (as out of or into Aberdeen) Zig-zag tracks may be more practical at higher range from land where the distance between radials is greater.

d) On approach

If this problem arises on approach to a platform, it may be appropriate to abandon the approach and to wait a few minutes until the satellites have adequately altered their position in the sky, or until a different satellite has risen in the sky, then to try again. An alternative may be to approach the platform in a slightly different direction if the wind strength and direction allows.

6 Sensitivity Considerations

Both the availability figures and the Realism Factors have been derived to a precision of several digits. This precision reflects the underlying conditions and assumptions. It is appropriate, however, to consider factors that influence how representative the results will be in future equipment and installations. The following paragraphs outline some of the relevant aspects.

a) The receivers used.

When the flight test programme was planned it had been hoped to use TSO-approved receivers. However, SRG accepted Cranfield’s proposal to use receivers that were not so approved. This was largely because it provided several advantages such as comparing UHF corrections from the platform with MF corrections from the shore. There was, in any case, no TSO specification for differential operation using corrections derived largely in a marine environment. It therefore cannot be said with certainty that the specific figures for availability and for the Realism Factor that were derived with the particular receivers will automatically apply in a future TSO-controlled environment. Nor is LNL aware of any attempts to include a requirement to check, or to limit this Realism Factor in the TSO.

b) The platforms to which the approaches were made.

There are several hundred platforms in the North Sea: only five of these were “flown” and the majority of the data was derived from four of them. They were specifically selected to be representative in terms of their physical layout and the likely effect on 2-D errors. Nevertheless it cannot be said with certainty that the same numerical results would be found with other platforms.

c) The direction of the approaches and the time of day.

Since it appears that the problem arises when there are few satellites usable and one or more of them is in a particular sector, it follows that a different approach direction at that time might present a different situation for the receiver. Also, the sector that posed a problem at one part of the day might not do so at a different time of day.

d) The locations where the flights took place.

These were all in the Northern North Sea in a common area. It should not be assumed that the performance in the areas flown will copy elsewhere. However, to minimise any uncertainty in this regard CAA placed a study contract which explored geographic variations. At the time those studies concluded that the geographic variations were insignificant.

e) The aircraft flown and the antenna.

There are several aircraft types flying in the North Sea and only one was “tested”. It cannot be stated with certainty that the figures derived are necessarily representative of what would be experienced with other aircraft. A particular uncertainty here is the effect of rotor interference. If some of the satellite losses arise from this effect, then it may well be that a different result would be obtained if any of the following were altered:

- the helicopter type if it had a different rotor speed or rotor blade construction,
- the antenna location (for instance if it was moved to the aircraft nose where the main rotor could have a different effect),
- the antenna type and mounting method (for instance some L-band antennas have a cylindrical shape and are mounted on a post; this could be mounted on the tail such that the L-band antenna was above the tail rotor).

See further under i) below.

f) The satellite constellation.

The constellation changes on a long time scale as satellites are withdrawn and some new ones are launched. It also changes from hour to hour as short-term problems arise and are dealt with. During the 1996 flight trials the number of healthy satellites in the constellation varied from 24 in the case of flights 2 to 5, to 26 in the case of flight 7. During the early part of 2001 (GPS weeks 71, 72 and 73), the number of healthy satellites in the constellation was 28. Comparing the constellation during the early flight trails with that during week 71, say, one would therefore intuitively think that today’s constellation would be less likely to cause a problem. That assumes that the cause is linked to the tail rotor rather than a processing limitation. However, the US Federal Radionavigation Plan for GPS only promises 24 satellites [6].

A special case of a change in the satellite constellation is if a receiver becomes capable of tracking both GPS and the proposed European satellite navigation system Galileo. The increased number of satellites would render the airborne installation less sensitive to such blockage.

g) Selective Availability

During the flight trials Selective Availability was "on". It is now "off". The extent to which this might affect the results is not known. There will probably be an impact on the residual error after correction, though the magnitude of the change is uncertain. If the residual error after correction is reduced, that raises the limiting HDOP4sum and therefore improves the availability. It is unlikely that the Realism Factor would be affected.

h) The MF corrections.

During the flight trials the MF corrections arose from a private company's reference station. The reference stations are now under the control of the General Lighthouse Authorities, the reference stations come from a different manufacturer, the transmitted messages are different, there is a new frequency plan and the advertised operational ranges quoted are different. The extent to which this might affect the situation is not known.

i) Aircraft installation and data sourcing.

The aircraft carried several GPS receivers all sharing a common L-band antenna. While this is unusual for an operational situation it should not affect the results in the manner that was observed here. The "tracked.txt" file was extracted from one receiver while the DGPS position was derived from another receiver. Since the receivers were identical, again the results should not be affected. Nevertheless, in both cases an uncertainty remains.

j) Health data discrepancies

The health of a satellite is described in the data stream from the satellites in several places. The data stream is partitioned into 25 sections referred to as "frames" and each frame is made up of 5 sub-frames that are numbered 1 to 5. In sub-frames 4 and 5 a satellite describes all the satellites in the constellation, giving information about their location and health (sometimes collectively called "almanac"). These data are transmitted once every 12.5 minutes. It is typically used by a receiver to determine which satellites it should try to acquire. In sub-frame 1 a satellite describes its own health and these data are transmitted every 30 seconds. It is used typically to detect a health deterioration that might require the receiver to drop the satellite. It takes time for the satellite system ground segment to formulate and to upload a new message for sub frames 4 and 5. Therefore, it could well be that the sub frame 4 and 5 data indicate that the satellite health is good while the sub frame 1 data indicate that the satellite health is bad. This might explain some of the cases where the receivers flown during the flight trials tracked fewer satellites than was expected from consideration of the almanac data. However, since the probability of a malfunction in a satellite is very low, this type of discrepancy cannot explain more than a tiny proportion of the cases identified in this report. In any case, it would be the almanac health data that would be used in performance predictions such as predictive RAIM. It is therefore entirely reasonable to rely on the almanac data in computing both availability and the Realism Factor as has been done in this report.

k) The basic project philosophy.

The philosophy has been that if one of the satellites tracked had in fact become unusable, only the other satellites that were already tracked would remain available. This would presumably be the case if the cause of the absence was a poor signal to noise ratio. It may be, however, that if one of the usable satellites was withdrawn, the receiver would be forced into a renewed search mode. If this proved successful it would invalidate the basic assumption in the project philosophy.

Despite these sensitivities it is felt appropriate to draw the conclusions presented in the following section.

7 Conclusions

The study has concluded that:

7.1 On availability

a) As flown, the availability was 94.2% during the approaches analysed. This "availability" refers to a position solution integrity assured by FD RAIM based on HDOP4sum. The solution was associated with a protection radius of 100 metres in 2-D and assessed at the 95% level of confidence.

b) When single satellites were ignored in the simulator (simulating a satellite failure or a less complete constellation) this availability dropped to 84.3%, averaged over the satellites and averaged over the ten approaches.

c) For these availabilities the number of samples examined in the case of one satellite ignored was 9072, this being the number of receiver output messages multiplied by four (the number of different satellites ignored). The number of samples for the "as flown" condition was 2268.

d) Cases of non-availability were dominated by conditions with too few satellites, rather than by poor geometry. Changing the RAIM architecture, therefore, will have limited benefit unless aiding is allowed.

e) The "not available" cases had a widely varying duration, sometimes lasting for the entire approach.

7.2 On precision

a) Expressing precision as a percentage of the available samples only, the 95% 2-D error was 15 metres "as flown" and 16 metres in the "one satellite off" condition.

b) Expressing precision as a percentage of all samples whether available or not the 95% values could not be established since the availability was less than 95% in both cases. However, an error of up to 30 metres contained 93.9% of all samples in the "as flown" condition; this fell to 84% in the "one satellite off" condition.

c) The percentages in (a) and (b) above were all tied to a protection radius of 100 metres at 95% confidence.

7.3 The Realism Factor

a) The Realism Factor was 0.858 (the probability that the receiver tracked what it was expected to track).

b) For the assessment of the Realism Factor the number of samples was 17611.

c) This low Realism Factor can be expected to have a significant impact on performance tested in accordance with the constellation referred to in the TSO specification.

d) Since predictive RAIM is based on the assumption that the Realism Factor is unity, the findings in this study suggest that predictive RAIM as currently used will provide an optimistic performance prediction. The safety of predictive RAIM must therefore be questioned.

e) Low Realism Factor was associated with a direction in space which had an elevation angle below 20 degrees and had an azimuth width of 37 degrees in the port direction. Since the aircraft's tail rotor is to the side of the tail, it is concluded that this rotor is at least part of the cause of the low Realism Factor.

7.4 Procedural aspects

a) Some operational procedures might be employed to alleviate the problems of low availability.

7.5 Overall conclusion

a) Exclusive reliance on GPS for helicopter navigation is unsafe.

b) Predictive RAIM for helicopter operations will give optimistic results.

c) Operational procedures have the potential to compensate for some of the lack of safety, but only if those procedures fully take into account the nature of the weaknesses of GPS for this application.

d) Design changes like the incorporation of aiding may offer some compensation, though the resulting geometry and its impact on precision would need to be checked. Transferring the responsibility for integrity verification from RAIM to that of the differential link is also a candidate technical solution.

8 Recommendations

a) The study has identified that there was a substantial difference between the satellites that were tracked and the satellites that were expected to be tracked from a basis of almanac data, mask angle and published receiver characteristics. There is therefore an urgent need for a research programme to examine the **cause** of the low Realism Factor and potential technical **solutions**. The effect of the tail rotor should form the initial focus of this work.

b) Consideration should be given to the extent to which operational **procedures** can adequately overcome the safety weaknesses the study has identified in relation to the use of GPS receivers on helicopters.

c) The GPS receiver **certification process** needs to ensure that the receiver's tracking performance is adequately defined in the specification. It must include tests to demonstrate that the actual tracking performance is consistent with the specification, not only as a stand-alone unit, but also when installed on the aircraft concerned and with any rotors (propellers in the case of fixed wing aircraft) working at normal speed. It must also include tests to demonstrate that any predictive RAIM built into the receiver is based on a tracking performance that is consistent with the performance that applies under realistic operating conditions.

d) While the tail rotor looks the most likely explanation for the unexpectedly poor tracking performance observed, **other explanations** should not be ignored. Papers that are freely available in the literature [4] have referred to cases where a GPS receiver tracked far fewer satellites than expected even when installed at a fixed site away from any rotor that could influence the results. Therefore, work should also be carried out to compare the satellites tracked with the satellites expected in GPS receivers installed on fixed wing aircraft.

9 References

1 A North Sea Trial To Investigate The Use Of Differential GPS For Instrument Approaches To Offshore Platforms. K.M. Dodson and J.R.A. Stevens. 23rd European Rotorcraft Forum. Dresden, Germany, September 1997.

2 The use of GPS Simulations to Enhance the Value of Flight Tests when evaluating DGPS as an Approach Aid to Off-Shore Platforms. R Johannessen and DA Howson, The proceedings of the 2000 National Technical Meeting, Anaheim CA, January 26-28, 2000. US Institute of Navigation.

3 Technical Standard Order TSO-C129a, Airborne Supplemental Navigation Equipment using the Global Positioning System (GPS).

4 Ten Million Data Points from TSO-approved aviation GPS receivers: Results of analysis and applications to Design and Use in Aviation. Navigation, Spring 2000. Vol 47, No 1.

5 DGPS Guidance for Helicopter Approaches to Offshore Platforms. CAA Paper 2000/5.

6 1999 Federal Radionavigation Plan, section 3.2.1B.

Tables

Table 1 Individual approaches and the overall result.Availability for the "as flown" condition.

Flight/Approach	Sample count		Availability
	Available	Not available	
Col 1	Col 2	Col 3	Col 4
2/3	202	0	100.00%
2/10	156	0	100.00%
4/1	242	0	100.00%
4/5	283	0	100.00%
5/1	274	4	98.56%
5/11	176	128	57.89%
6/3	180	0	100.00%
6/11	271	0	100.00%
7/3	226	0	100.00%
7/7	126	0	100.00%
Overall	2136	132	94.18%

Table 2 Availability for the "one satellite off" condition. Individual approaches and the overall result.

Flight/ Approach	Sample count when one of the satellites is ignored										Avail
	Not PRNA		Not PRN B		Not PRN C		Not PRN D		Overall		
	Good	Bad	Good	Bad	Good	Bad	Good	Bad	Good	Bad	%
Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8	Col 9	Col 10	Col 11	Col 12
2/3	199	3	199	3	190	12	199	3	787	21	97.40
2/10	152	4	155	1	153	3	152	4	612	12	98.08
4/1	235	7	235	7	235	7	242	0	947	21	97.83
4/5	283	0	281	2	283	0	279	4	1126	6	99.47
5/1	274	4	274	4	233	45	268	10	1049	63	94.33
5/11	2	302	4	300	2	302	2	302	10	1206	0.82
6/3	139	41	176	4	139	41	172	8	626	94	86.94
6/11	271	0	271	0	271	0	271	0	1084	0	100.0
7/3	226	0	226	0	226	0	226	0	904	0	100.0
7/7	126	0	126	0	126	0	126	0	504	0	100.0
Total	1907	361	1947	321	1858	410	1937	331	7649	1423	84.31
Total	2268		2268		2268		2268		9072		-
Availability %	84.08		85.85		81.92		85.41		84.31		-

NOTE:

1 "PRNA", "PRNB", "PRNC" and "PRND" are the four satellites that, one at a time, were set "not usable" in the simulations. This nomenclature was used since the actual numbers were different in different approaches.

2 "Good" = available, "Bad" = Not available.

Table 3 Causes of non-availability with one satellite off.

Flight/approach	Too few satellites	Bad geometry	Total
Col 1	**Col 2**	**Col 3**	**Col 4**
2/3	9	12	21
2/10	12	0	12
4/1	21	0	21
4/5	4	2	6
5/1	63	0	63
5/11	1206	0	1206
6/3	28	66	94
6/11	0	0	0
7/3	0	0	0
7/7	0	0	0
Total	1343	80	1423

Table 4 Summary of 2-D error (m) at 95% level of confidence for all approaches. Out of all available samples.

Condition		2-D error (metres)		No of available samples	
Flight	**Approach**	**As flown**	**One satellite off**	**As flown**	**One satellite off**
2	3	14	24	202	787
2	10	14	18	156	612
4	1	9	12	242	947
4	5	18	18	283	1126
5	1	12	13	274	1049
5	11	26	-	176	10
6	3	11	14	180	626
6	11	10	10	271	1084
7	3	12	14	226	904
7	7	11	12	126	504
Overall		15	16	2136	7649

Table 5 Realism Factors

Flight	Approach	Tracked	Expected	Factor
2	3	1,264	1,484	0.852
2	10	982	1,092	0.899
4	1	1,528	1,936	0.789
4	5	1,912	2,235	0.855
5	1	1,737	2,073	0.838
5	11	1,394	1,824	0.764
6	3	1,166	1,260	0.925
6	11	2,378	2,439	0.975
7	3	1,874	2,260	0.829
7	7	879	1,008	0.872
Totals		15,114	17,611	0.858

Figures

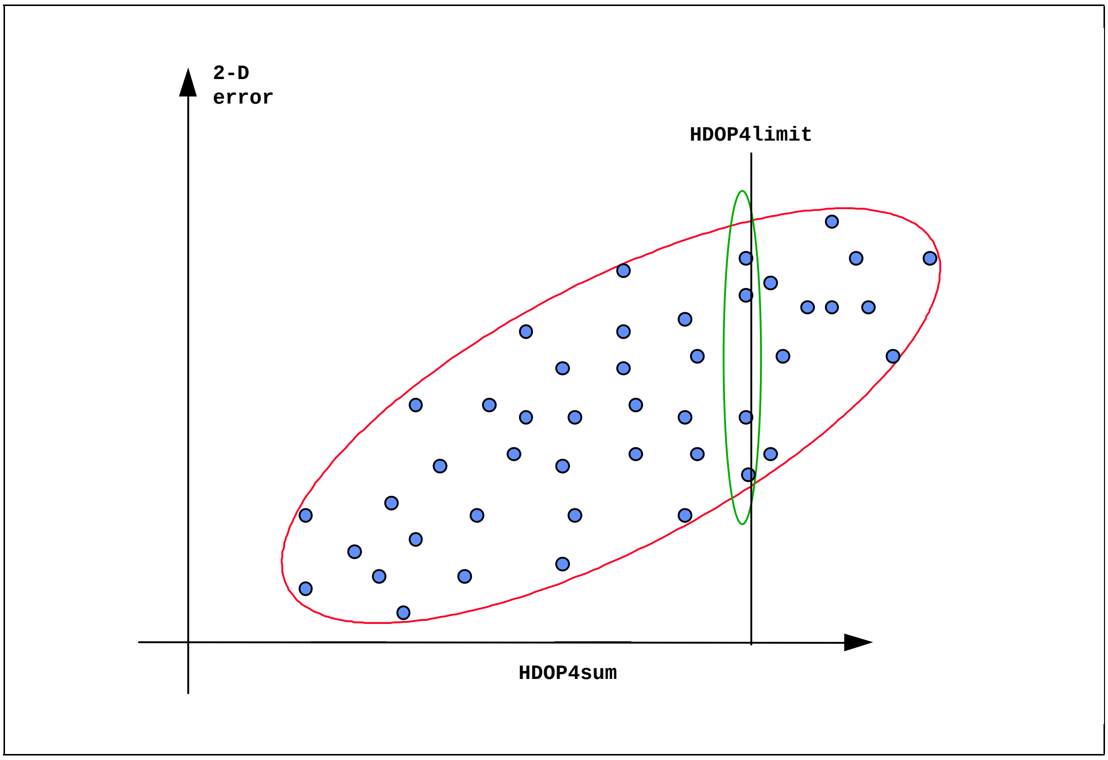

Figure 1 Scatter diagram showing a plot of error against dilution of precision. Points within the green curve are used to determine the protection radius. Points within the red curve are used to determine the general 2-D error distribution.

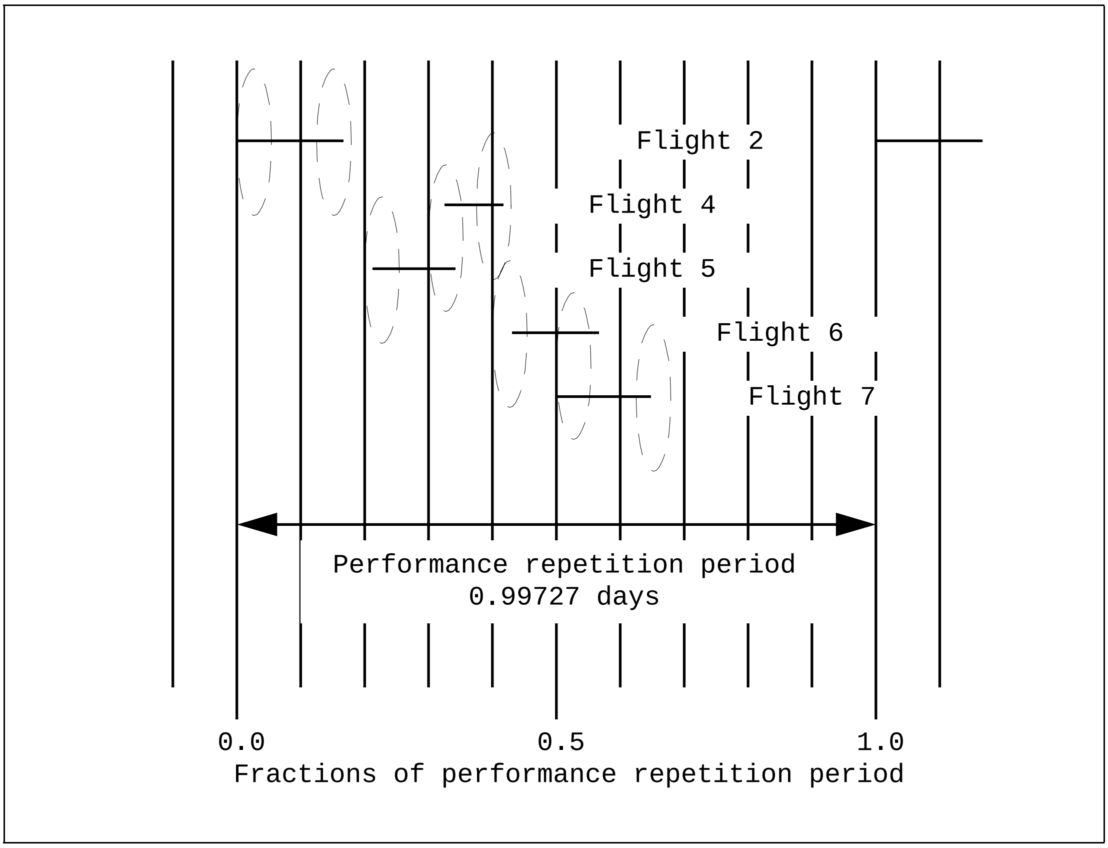

Figure 2 Location of flights within the performance repetition period, relative to the start of flight 2. Ellipses indicate the portions from which the approaches were selected.

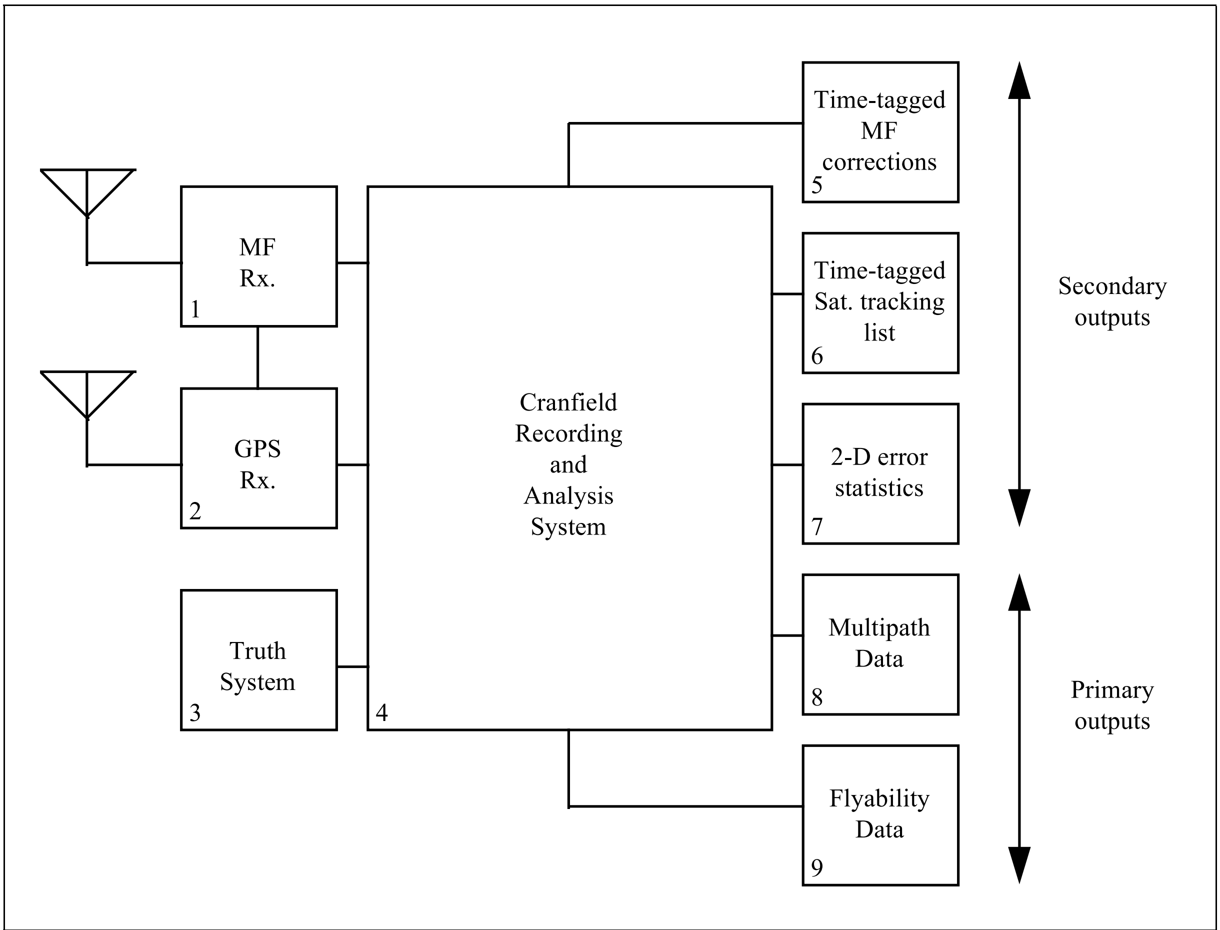

Figure 3 Arrangements for the Cranfield Aerospace Flight Test Data

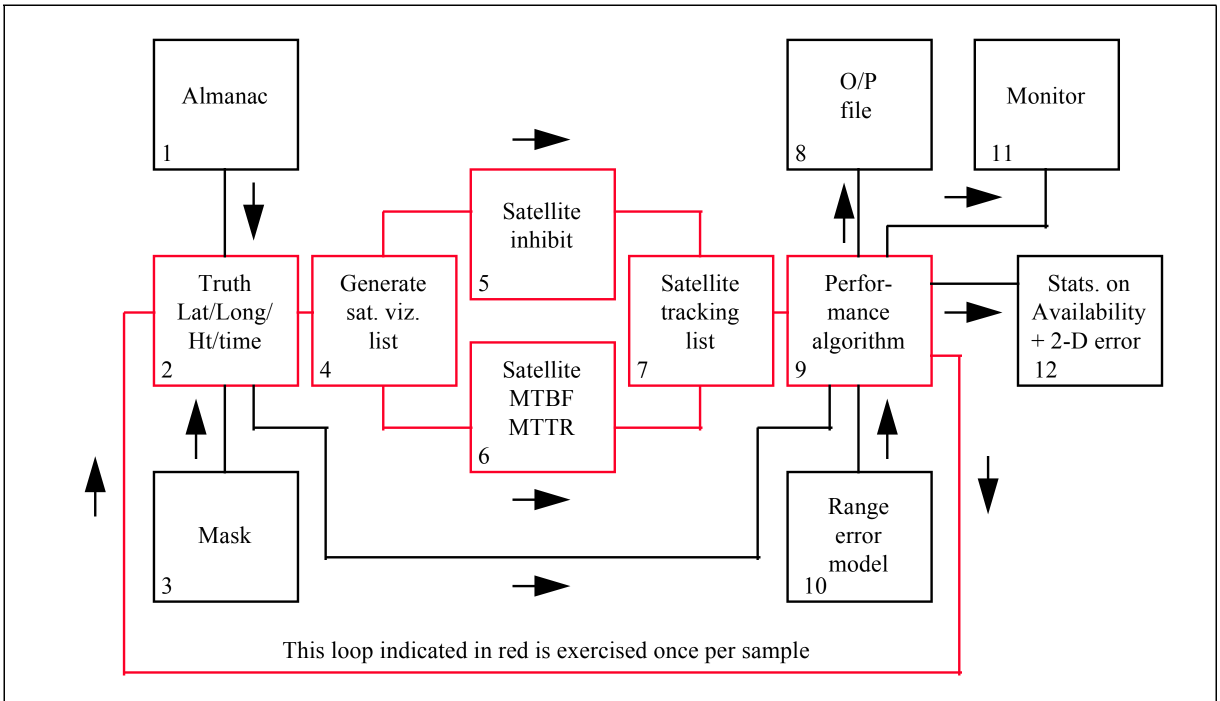

Figure 4 Main simulation loop.

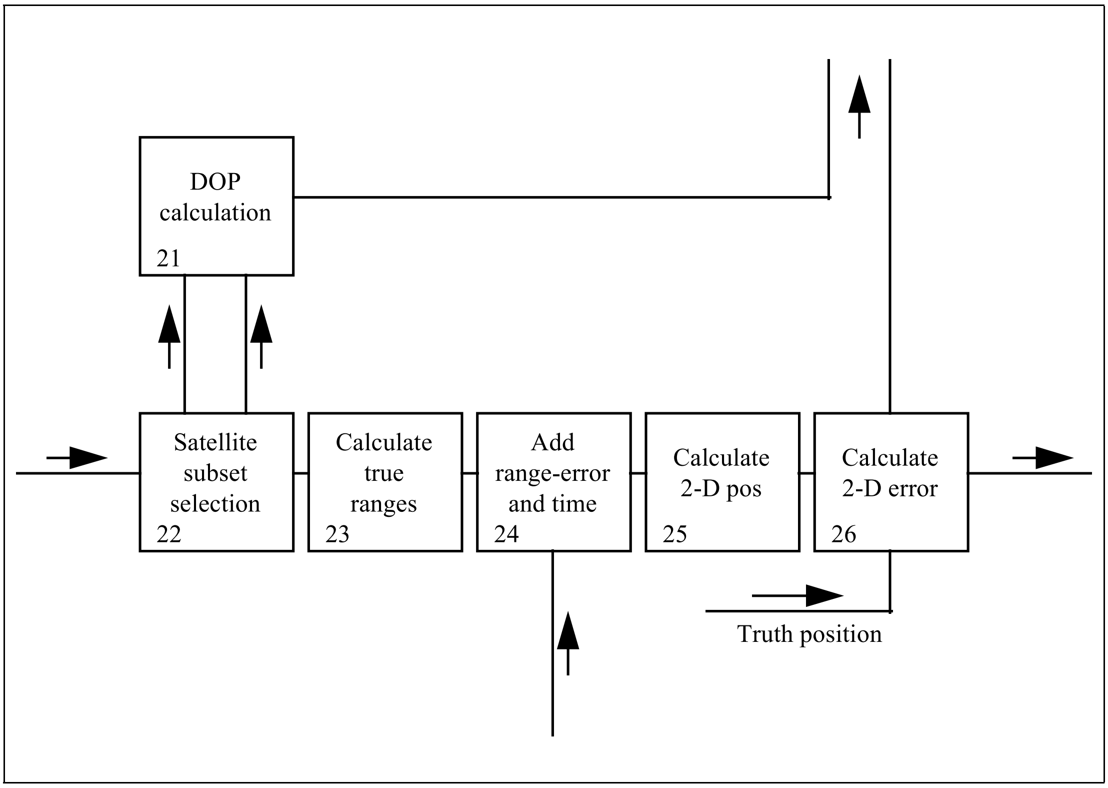

Figure 5 Performance algorithm

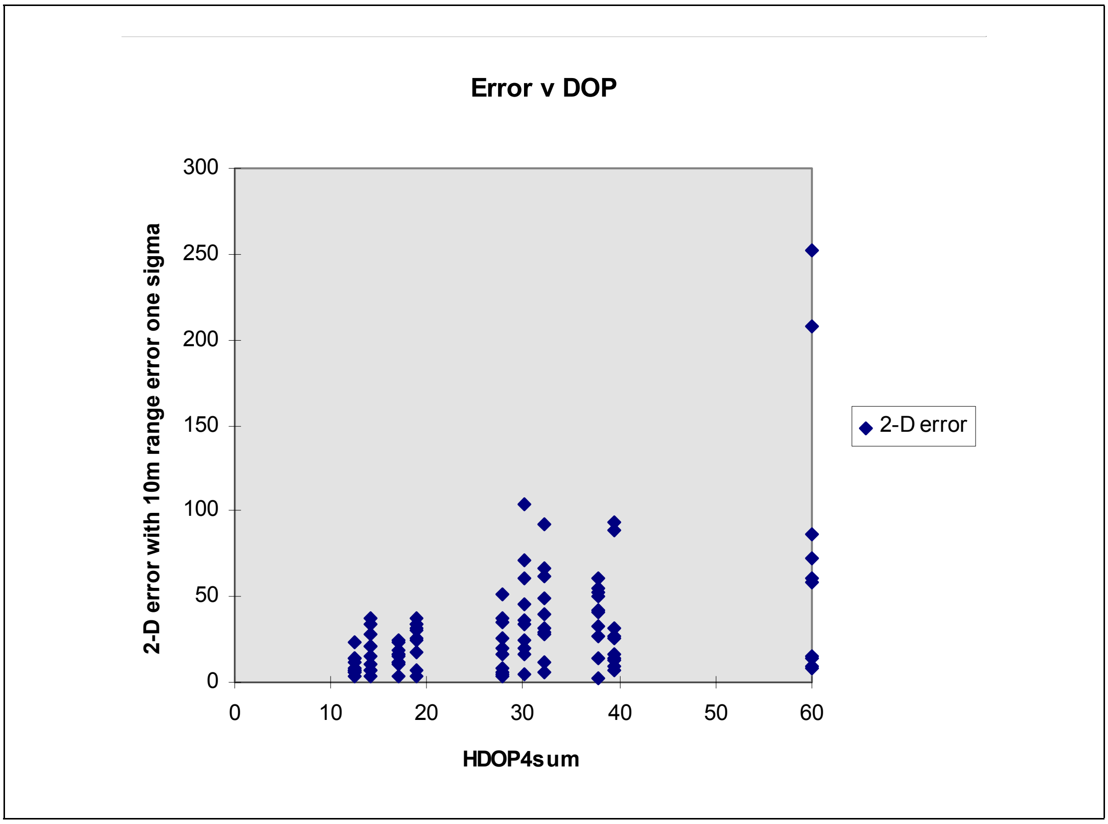

Figure 6 Scatter plot showing 2-D error in metres v. HDOP4sum.

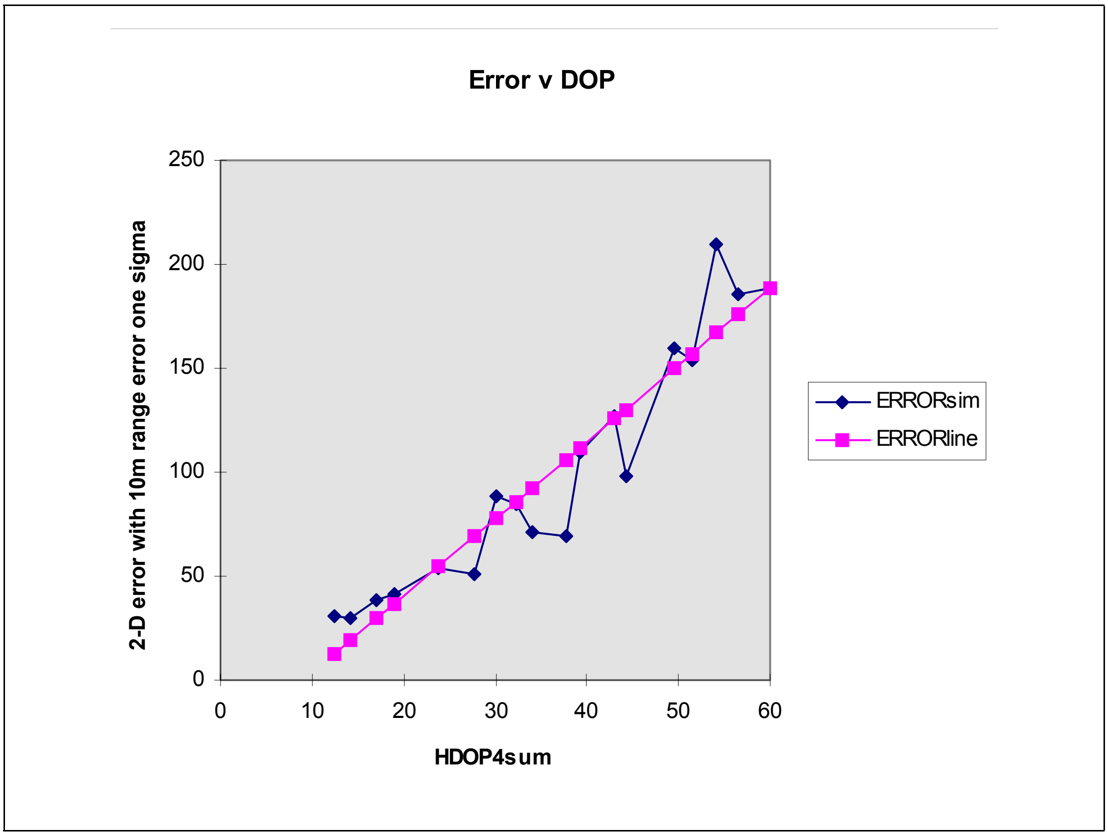

Figure 7 95% 2-D error v HDOP4sum in FD RAIM un-weighted.
ERRORsim is from the Simulator
ERRORline is the corresponding least squares fit straight line.
100 simulations, residual range error 10 metres one σ.
Number of DOP values = 18.

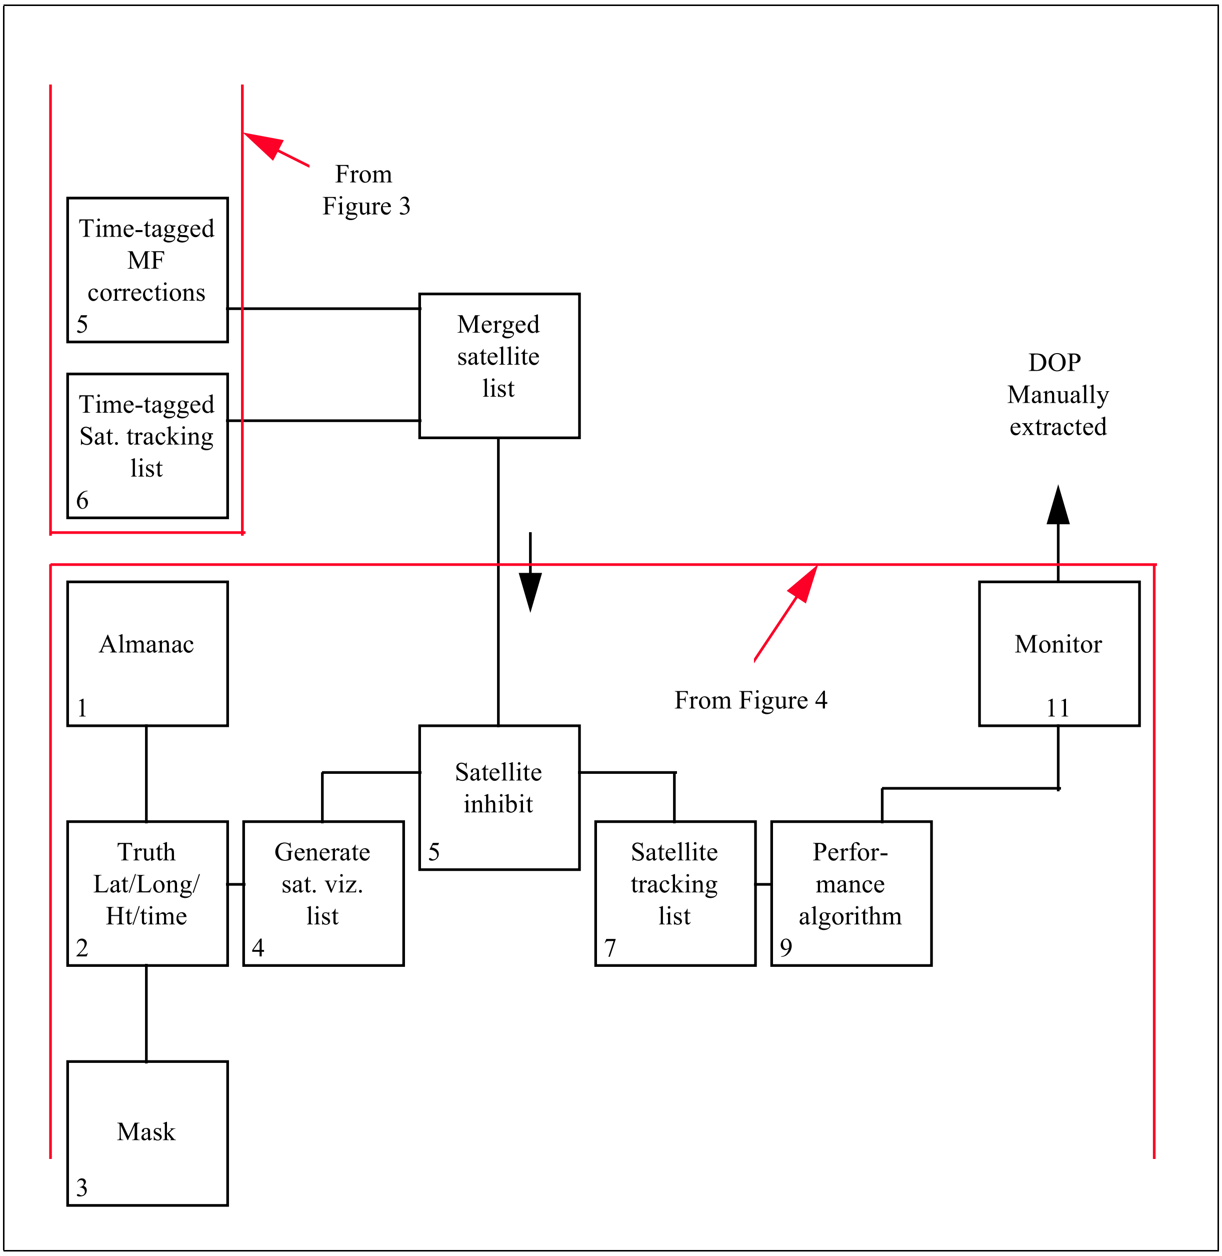

Figure 8 Determination of availability.

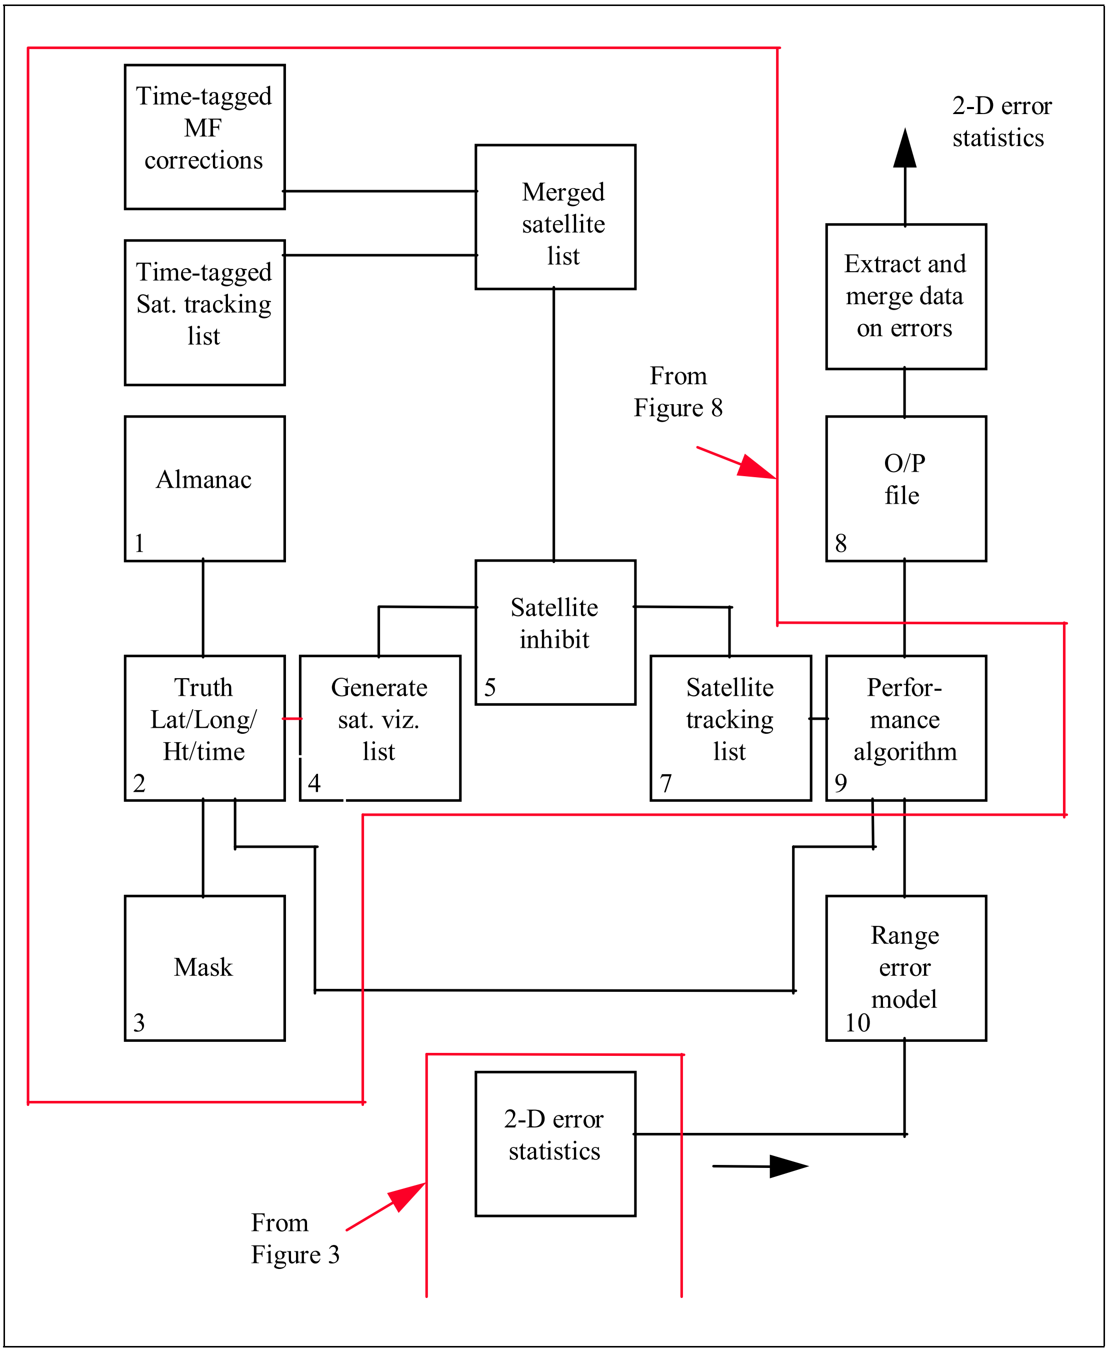

Figure 9 Determination of precision

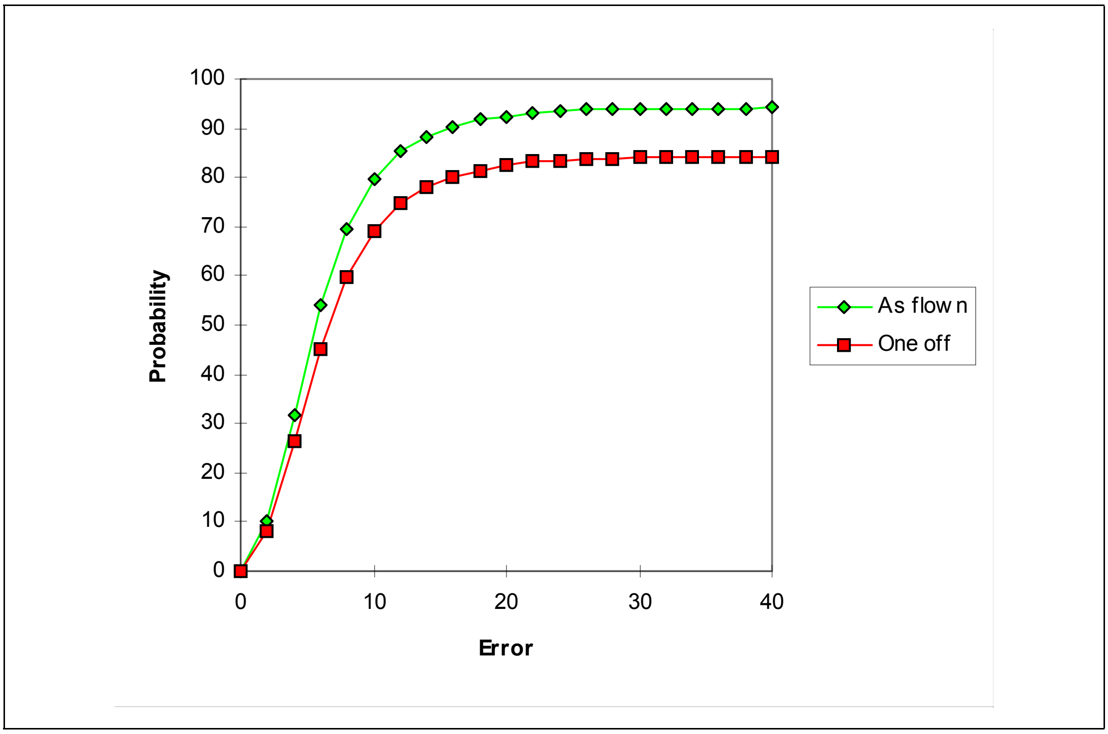

Figure 10 Probability (%) that 2-D error (m) is less than value shown. All ten approaches combined, % based on all samples irrespective of availability.

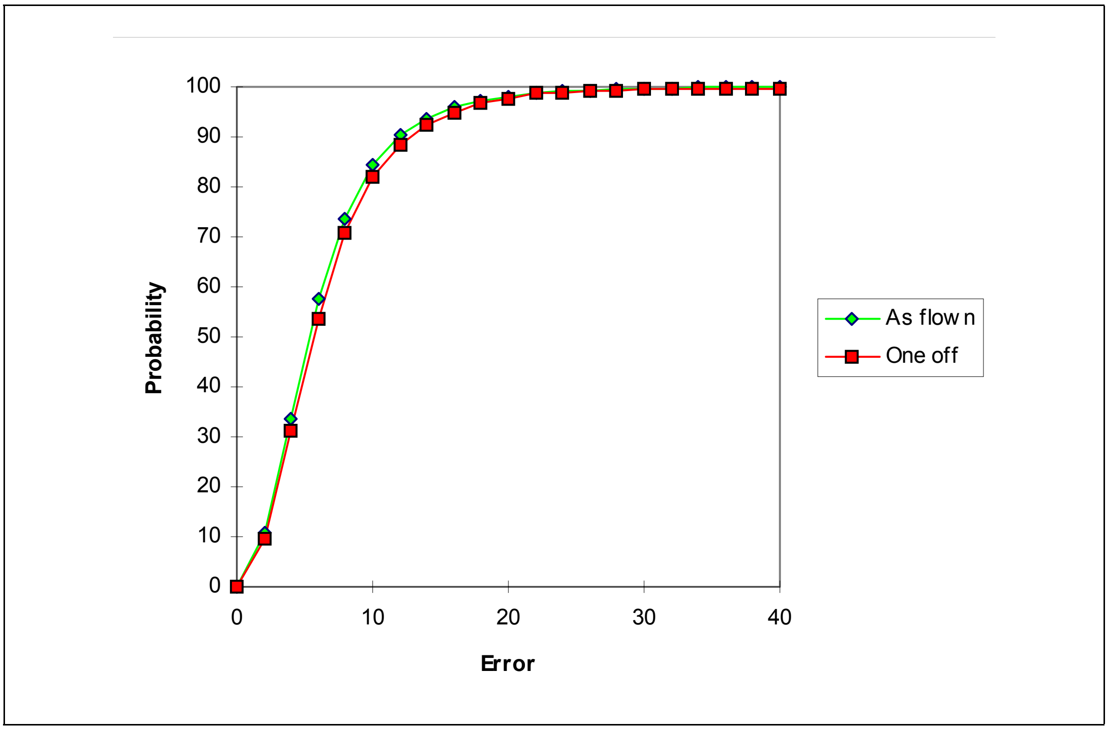

Figure 11 Probability (%) that 2-D error (m) is less than value shown. All ten approaches combined, % based on available samples only.

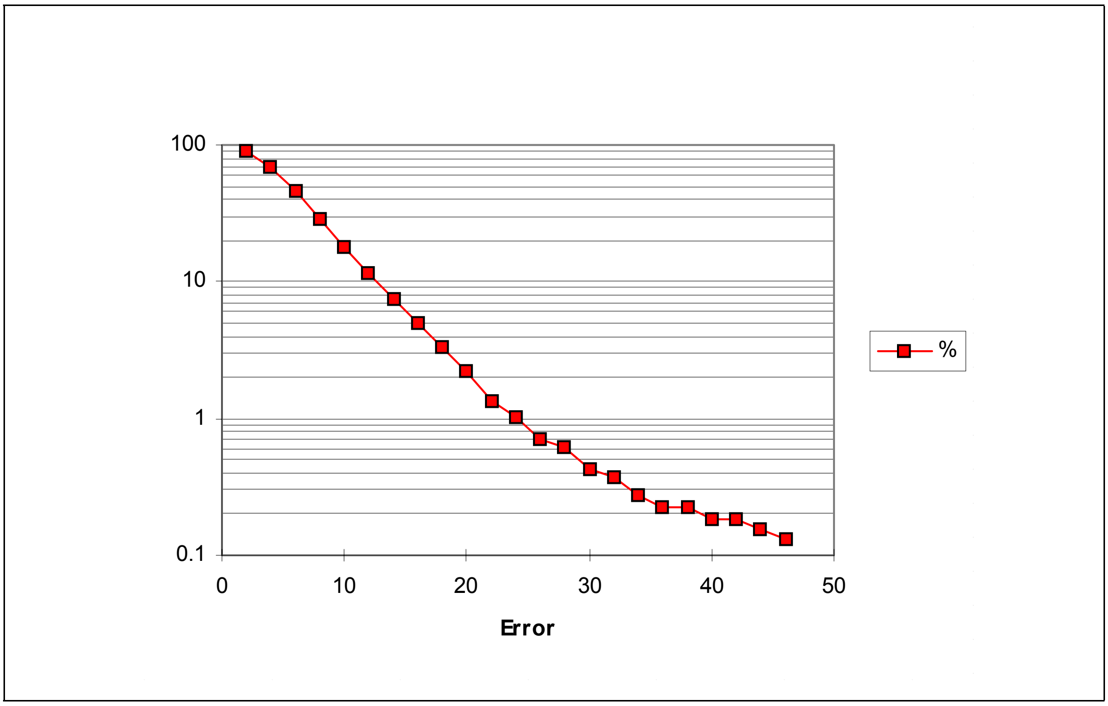

Figure 12 Probability (%) that 2-D error (m) exceeds the value shown by the horizontal axis, % based on the available samples only, "one satellite off" condition.

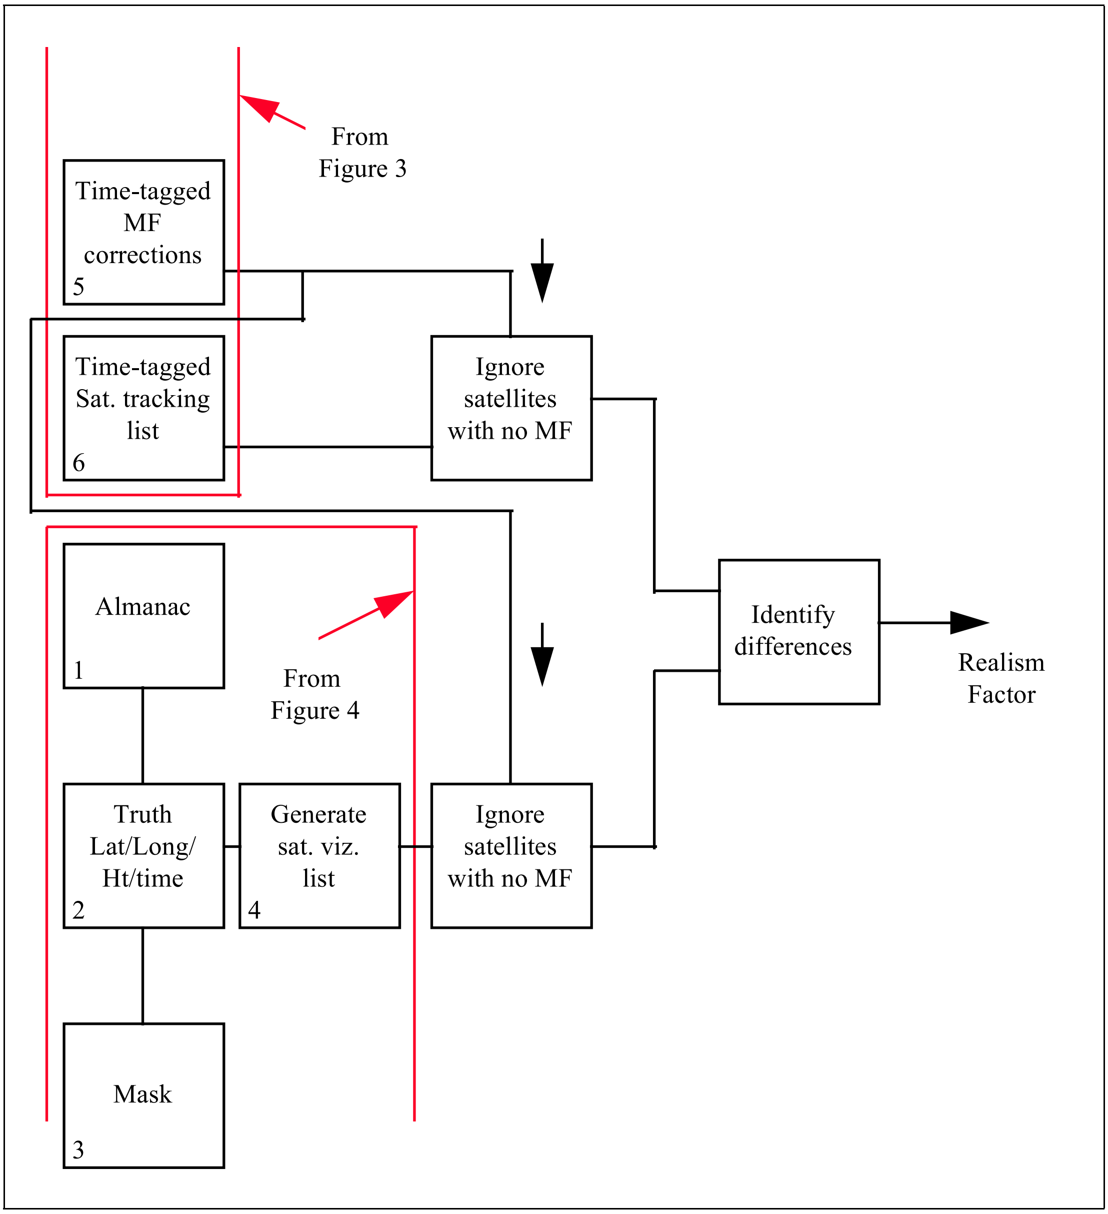

Figure 13 Determination of the Realism Factor

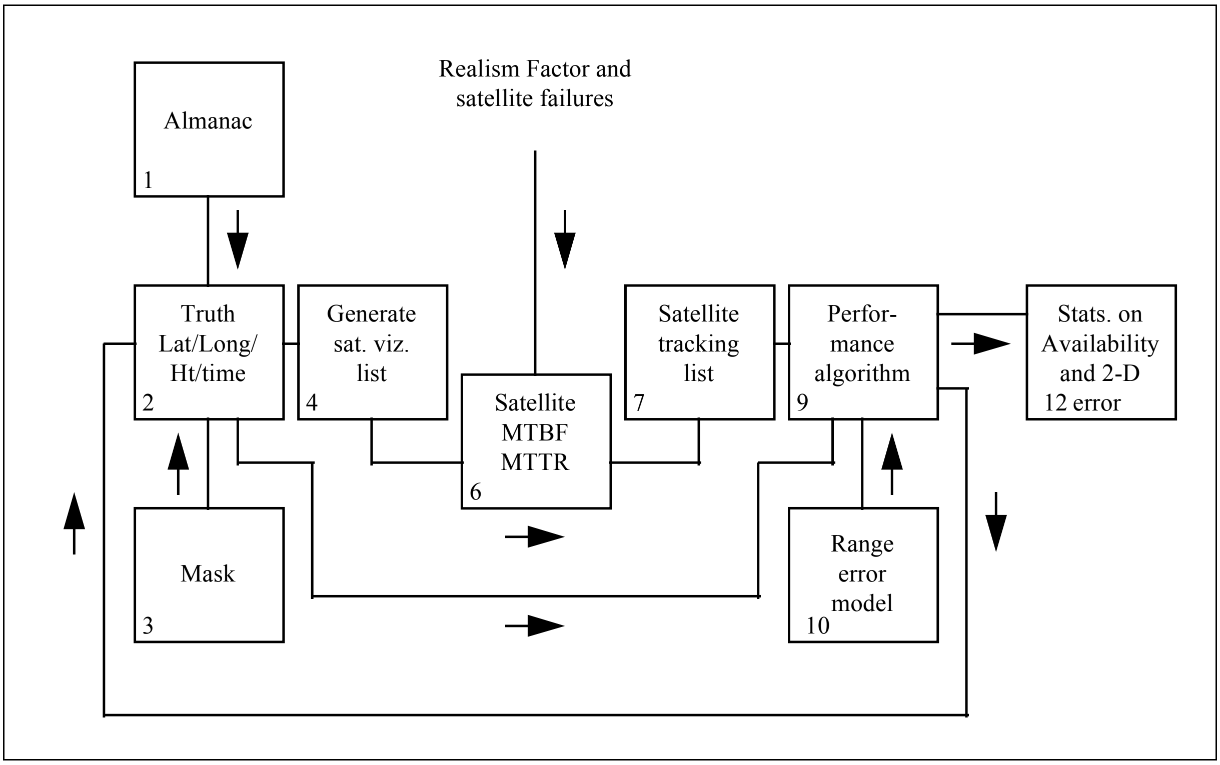

Figure 14 Determination of the effect of the low Realism Factor

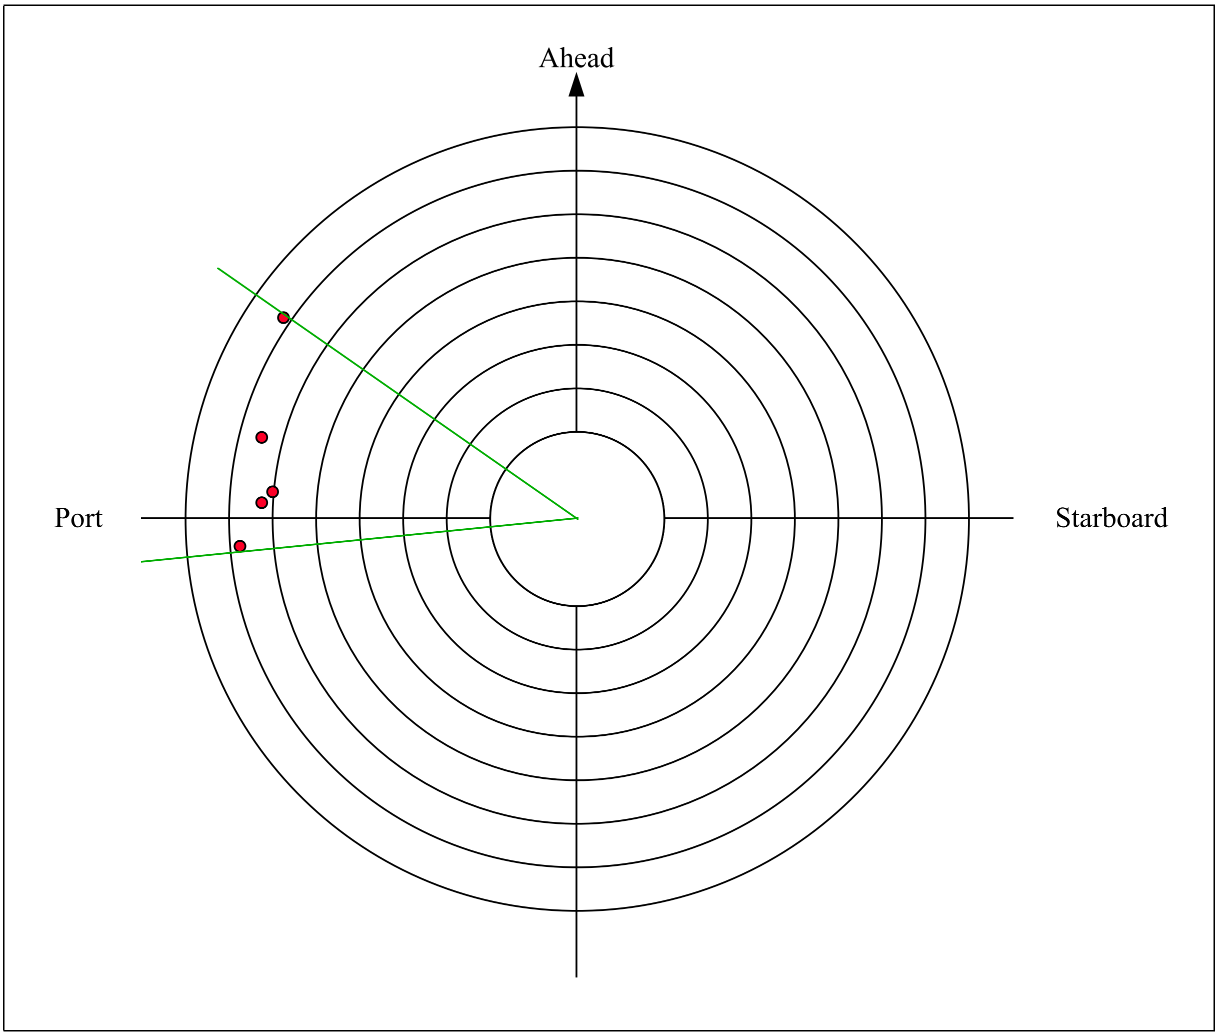

Figure 15 Sky plot for the 5 cases of lowest Realism Factor.
(The outer ring represents zero degree elevation, the centre represents 90 degree elevation.)

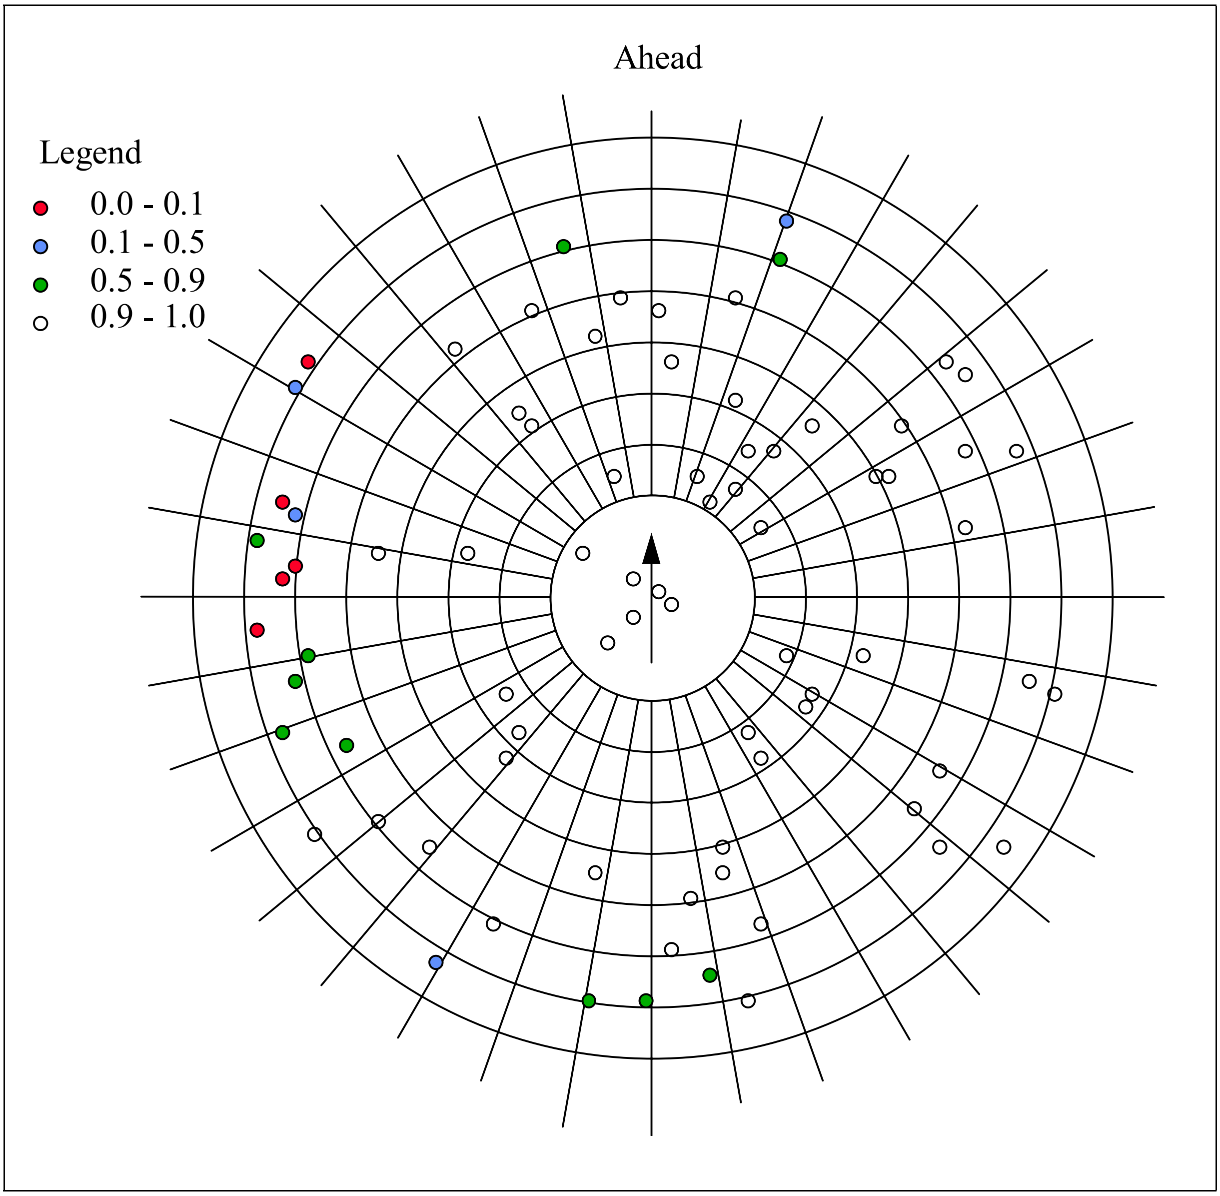

Figure 16 Sky plot for all the Realism Factors.
(The outer ring represents zero degree elevation, the centre represents 90 degree elevation.)

Annex 1 Estimating the availability for flight 4 approach 5.

This annex illustrates the key processes used to determine the availability.

Table A1.1 shows how the data recorded during an approach was divided into 24 sections. The approach started at time 15:01:15 and was completed at time 15:05:57. At a GPS data rate of 1 Hz this provided a total of 283 samples. During the first section of the approach the receiver tracked satellites PRNs 01, 04, 05, 07, 14, 15 and 29 such that there were 7 satellites usable. The receiver then dropped PRN05 from its list and continued without that satellite for 5 samples until at time 15:01:31 it was again included. Such changes arose several times during the approach. PRNs 04, 07, 14, 15 and 29 were used continuously but PRNs 01 and 05 were used intermittently.

Within the approach sections 1, 3, 5 and 7 plus a number of later sections the same satellites were used. These sections were therefore grouped together within a single simulation run which Table A1.2 shows was run number 896. Similarly, other sections were grouped together such that the "as flown" analysis required four different simulation runs. The same Table shows the worst DOP for each section. In the case of Table A1.2 the DOP cited in column 6 is the PDOP4sum as displayed by the Simulator monitor (box 11 in Figure 4). Table A1.2 shows that for all sections the DOP was below the threshold and all sections had 5 or more satellites: therefore all samples satisfied the criterion for "available". The terms "good" and "bad" are used to save some space in the Table headings.

The "one satellite off" situation was then analysed for 4 different satellites off, one at a time. Table A1.3 summarises one of these situations, namely with PRN07 set unusable. Approach sections with a common subset of satellites were again grouped together. One of these groups comprised the samples at times 15:04:11 and 15:04:13 for which there were only four satellites usable. These two samples were therefore unavailable on account of there being less than 5 satellites. For two of the groups the PDOP4sum was always less than 75.45 and HDOP4sum would therefore also be less. However, for the last group the displayed PDOP4sum was above 75.45 and a detailed analysis was then carried out using the Simulator's HDOP4 architecture five different times, each time setting one of the five satellites unavailable. The five values of HDOP4 were then added to give HDOP4sum. For two of the samples this exceeded the threshold and for the rest it was below.

The entries in each of column 7 and column 8 were added separately, giving 279 available samples and 4 unavailable. The availability for this approach with this particular satellite set unavailable was therefore 100*279/283 = 98.59%.

In Tables A1.2 and A1.3 column 3 is used to show which satellites were used for the simulation. In order to achieve this some satellites (shown in column 2) had to be inhibited, this was achieved by the Simulator's satellite inhibit feature (box 5 in Figure 4).

Table A1.1 Measurements availability history.
Flight 4 approach 5.

Section	From	To	Samples	PRNs tracked by both	N
Col 1	**Col 2**	**Col 3**	**Col 4**	**Col 5**	**Col 6**
1	15:01:15	15:01:25	11	01-04-05-07-14-15-29	7
2	15:01:26	15:01:30	5	01-04- -07-14-15-29	6
3	15:01:31	15:03:11	101	01-04-05-07-14-15-29	7
4	15:03:12	15:03:40	29	01-04- -07-14-15-29	6
5	15:03:41	15:04:04	24	01-04-05-07-14-15-29	7
6	15:04:05	15:04:07	3	01-04- -07-14-15-29	6
7	15:04:08	15:04:08	1	01-04-05-07-14-15-29	7
8	15:04:09	15:04:10	2	01-04- -07-14-15-29	6
9	15:04:11	15:04:11	1	-04- -07-14-15-29	5
10	15:04:12	15:04:12	1	-04-05-07-14-15-29	6
11	15:04:13	15:04:13	1	-04- -07-14-15-29	5
12	15:04:14	15:04:16	3	01-04- -07-14-15-29	6
13	15:04:17	15:04:41	25	01-04-05-07-14-15-29	7
14	15:04:42	15:04:44	3	01-04- -07-14-15-29	6
15	15:04:45	15:04:50	6	01-04-05-07-14-15-29	7
16	15:04:51	15:04:53	3	01-04- -07-14-15-29	6
17	15:04:54	15:05:12	19	01-04-05-07-14-15-29	7
18	15:05:13	15:05:16	4	01-04- -07-14-15-29	6
19	15:05:17	15:05:17	1	01-04-05-07-14-15-29	7
20	15:05:18	15:05:23	6	01-04- -07-14-15-29	6
21	15:05:24	15:05:39	16	01-04-05-07-14-15-29	7
22	15:05:40	15:05:44	5	-04-05-07-14-15-29	6
23	15:05:45	15:05:56	12	01-04-05-07-14-15-29	7
24	15:05:57	15:05:57	1	-04-05-07-14-15-29	6
Total	15:01:15	15:05:57	283		

Table A1.2 Simulation conditions for the "as flown" analysis. Flight 4 approach 5.

Run	PRNs		Period	N	DOP	Samples	
	Inhibited	**Used**				**Good**	**Bad**
Col 1	**Col 2**	**Col 3**	**Col 4**	**Col 5**	**Col 6**	**Col 7**	**Col 8**
896	18-25	01-04-05-07-14-15-29	15:01:15 – 15:01:25	7	24.9	11	0
			15:01:31 – 15:03:11		25.9	101	0
			15:03:41 – 15:04:04		26.1	24	0
			15:04:08 – 15:04:08		26.1	1	0
			15:04:17 – 15:04:41		26.2	25	0
			15:04:45 – 15:04:50		26.2	6	0
			15:04:54 – 15:05:12		26.3	19	0
			15:05:17 – 15:05:17		26.4	1	0
			15:05:24 – 15:05:39		26.5	16	0
			15:05:45 – 15:05:56		26.6	12	0
897	05-18-25	01-04-07-14-15-29	15:01:26 – 15:01:30	6	25.9	5	0
			15:03:12 – 15:03:40		26.0	29	0
			15:04:05 – 15:04:07		26.1	3	0
			15:04:09 – 15:04:10		26.1	2	0
			15:04:14 – 15:04:16		26.1	3	0
			15:04:42 – 15:04:44		26.2	3	0
			15:04:51 – 15:04:53		26.3	3	0
			15:05:13 – 15:05:16		26.4	4	0
			15:05:18 – 15:05:23		26.4	6	0
898	01-05-18-25	04-07-14-15-29	15:04:11 – 15:04:11	5	26.1	1	0
			15:04:13 – 15:04:13		26.1	1	0
898	01-18-25	04-05-07-14-15-29	15:04:12 – 15:04:12	6	26.1	1	0
			15:05:40 – 15:05:44		26.5	5	0
			15:05:57 – 15:05:57		26.6	1	0
Totals						283	0
						283	
Availability						100.00%	

Table A1.3 Simulation conditions with PRN07 ignored.
Flight 4 approach 5.

Run	PRNs		Period	N	DOP	Samples	
	Inhibited	Used				Good	Bad
Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8
907	07-18-25	01-04-05-14-15-29	15:01:15 – 15:01:25	6	28.9	11	0
			15:01:31 – 15:03:11		30.6	101	0
			15:03:41 – 15:04:04		31.6	24	0
			15:04:08 – 15:04:08		31.7	1	0
			15:04:17 – 15:04:41		32.0	25	0
			15:04:45 – 15:04:50		31.9	6	0
			15:04:54 – 15:05:12		31.9	19	0
			15:05:17 – 15:05:17		31.9	1	0
			15:05:24 – 15:05:39		31.9	16	0
			15:05:45 – 15:05:56		31.9	11	0
908	05-07-18-25	01-04-14-15-29	15:01:26 – 15:01:30	5	29.8	5	0
			15:03:12 – 15:03:40		34.5	29	0
			15:04:05 – 15:04:07		35.9	3	0
			15:04:09 – 15:04:10		36.1	2	0
			15:04:14 – 15:04:16		36.4	3	0
			15:04:42 – 15:04:44		38.1	3	0
			15:04:51 – 15:04:53		38.8	3	0
			15:05:13 – 15:05:16		40.5	4	0
			15:05:18 – 15:05:23		41.1	6	0
N/A	01-05-07-18-25	04-14-15-29	15:04:11 – 15:04:11	4	-	0	1
			15:04:13 – 15:04:13		-	0	1
909	01-07-18-25	04-05-14-15-29	15:04:12 – 15:04:12	5	See annex text	1	0
			15:05:40 – 15:05:44			4	1
			15:05:57 – 15:05:57			0	1
Totals						279	4
						283	
Availability						98.59%	

Annex 2 Estimating the precision for flight 4 approach 5.

This annex illustrates the format of the output file (box 8 in Figure 4) produced by the Simulator. It is **part of** the file produced during simulation run 897. The first parameter is the GPS week number (864), the second parameter is the GPS time in seconds within the week. After the gap the first two parameters are the latitude and longitude errors in metres. All the other parameters may be ignored for the purpose of this analysis. The records that are in bold are some of the records that had to be used. Thus the first block shown comprises 227045 to 227047 seconds corresponding to times 15:04:05 – 15:04:07. The next block comprises 15:04:09 – 15:04:10 etc. In addition to the blocks highlighted in this annex, there are some blocks earlier and some blocks later than those shown such that, in all, there are nine groups to be included from the data generated during simulator run number 897. This may be checked against column 4 in Annex 1 Table A1.2.

864	227043.00	+058.460333000	+000.249167000	+00360.00	0.00	5.69	3.26	11.27	26.08	6	1
864	227044.00	+058.460333000	+000.249167000	+00360.00	0.00	-6.38	-1.89	2.52	26.08	6	1
864	**227045.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**8.89**	**8.06**	**-9.50**	**26.09**	**6**	**1**
864	**227046.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**-1.11**	**1.43**	**0.42**	**26.09**	**6**	**1**
864	**227047.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**9.48**	**3.18**	**-2.82**	**26.09**	**6**	**1**
864	227048.00	+058.460333000	+000.249167000	+00360.00	0.00	4.41	6.18	-21.50	26.09	6	1
864	**227049.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**3.56**	**5.02**	**-0.08**	**26.10**	**6**	**1**
864	**227050.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**3.06**	**3.30**	**-11.01**	**26.10**	**6**	**1**
864	227051.00	+058.460333000	+000.249167000	+00360.00	0.00	4.55	0.32	7.74	26.10	6	1
864	227052.00	+058.460333000	+000.249167000	+00360.00	0.00	3.47	4.63	-1.32	26.11	6	1
864	227053.00	+058.460333000	+000.249167000	+00360.00	0.00	8.33	5.24	14.60	26.11	6	1
864	**227054.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**-3.70**	**-5.94**	**-8.42**	**26.11**	**6**	**1**
864	**227055.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**1.76**	**-9.45**	**19.97**	**26.12**	**6**	**1**
864	**227056.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**-13.82**	**-5.51**	**-14.98**	**26.12**	**6**	**1**
864	227057.00	+058.460333000	+000.249167000	+00360.00	0.00	5.43	0.38	6.97	26.12	6	1
864	227058.00	+058.460333000	+000.249167000	+00360.00	0.00	-7.15	-8.47	18.11	26.13	6	1
864	227059.00	+058.460333000	+000.249167000	+00360.00	0.00	2.13	5.39	-22.61	26.13	6	1
864	227060.00	+058.460333000	+000.249167000	+00360.00	0.00	2.79	2.88	4.12	26.13	6	1
864	227061.00	+058.460333000	+000.249167000	+00360.00	0.00	4.07	3.47	-14.07	26.14	6	1
864	227062.00	+058.460333000	+000.249167000	+00360.00	0.00	-8.08	-10.29	5.00	26.14	6	1
864	227063.00	+058.460333000	+000.249167000	+00360.00	0.00	18.46	11.86	-8.39	26.14	6	1
864	227064.00	+058.460333000	+000.249167000	+00360.00	0.00	6.28	0.37	5.19	26.15	6	1
864	227065.00	+058.460333000	+000.249167000	+00360.00	0.00	2.72	0.28	4.33	26.15	6	1
864	227066.00	+058.460333000	+000.249167000	+00360.00	0.00	-5.75	-5.34	-1.75	26.15	6	1
864	227067.00	+058.460333000	+000.249167000	+00360.00	0.00	9.12	4.06	-15.16	26.16	6	1
864	227068.00	+058.460333000	+000.249167000	+00360.00	0.00	0.34	3.56	-3.39	26.16	6	1
864	227069.00	+058.460333000	+000.249167000	+00360.00	0.00	1.25	4.85	-16.58	26.16	6	1
864	227070.00	+058.460333000	+000.249167000	+00360.00	0.00	-6.64	-6.93	1.64	26.17	6	1
864	227071.00	+058.460333000	+000.249167000	+00360.00	0.00	7.56	4.31	6.99	26.17	6	1
864	227072.00	+058.460333000	+000.249167000	+00360.00	0.00	11.08	3.19	7.52	26.17	6	1
864	227073.00	+058.460333000	+000.249167000	+00360.00	0.00	5.96	1.63	7.38	26.18	6	1
864	227074.00	+058.460333000	+000.249167000	+00360.00	0.00	-2.11	4.74	-2.14	26.18	6	1
864	227075.00	+058.460333000	+000.249167000	+00360.00	0.00	6.44	3.74	-3.91	26.18	6	1
864	227076.00	+058.460333000	+000.249167000	+00360.00	0.00	-0.28	-2.44	-3.23	26.19	6	1
864	227077.00	+058.460333000	+000.249167000	+00360.00	0.00	15.28	12.75	11.63	26.19	6	1
864	227078.00	+058.460333000	+000.249167000	+00360.00	0.00	5.28	3.71	-8.53	26.19	6	1
864	227079.00	+058.460333000	+000.249167000	+00360.00	0.00	2.67	-1.54	14.36	26.20	6	1
864	227080.00	+058.460333000	+000.249167000	+00360.00	0.00	-24.19	-8.27	-15.52	26.20	6	1
864	227081.00	+058.460333000	+000.249167000	+00360.00	0.00	13.63	5.29	-8.10	26.21	6	1
864	**227082.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**-2.86**	**-1.72**	**-6.00**	**26.21**	**6**	**1**
864	**227083.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**-6.53**	**-7.58**	**23.57**	**26.21**	**6**	**1**
864	**227084.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**6.81**	**-3.43**	**-2.44**	**26.22**	**6**	**1**
864	227085.00	+058.460333000	+000.249167000	+00360.00	0.00	9.12	4.90	-10.08	26.22	6	1
864	227086.00	+058.460333000	+000.249167000	+00360.00	0.00	-5.90	-0.84	-0.48	26.22	6	1
864	227087.00	+058.460333000	+000.249167000	+00360.00	0.00	-11.98	-10.44	-0.52	26.23	6	1
864	227088.00	+058.460333000	+000.249167000	+00360.00	0.00	11.06	6.68	4.85	26.23	6	1
864	227089.00	+058.460333000	+000.249167000	+00360.00	0.00	-22.25	-7.78	-1.05	26.24	6	1
864	227090.00	+058.460333000	+000.249167000	+00360.00	0.00	-2.80	0.49	4.72	26.24	6	1
864	**227091.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**3.68**	**0.60**	**5.58**	**26.24**	**6**	**1**
864	**227092.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**2.51**	**4.70**	**-5.64**	**26.25**	**6**	**1**
864	**227093.00**	**+058.460333000**	**+000.249167000**	**+00360.00**	**0.00**	**-2.75**	**-3.00**	**-4.54**	**26.25**	**6**	**1**
864	227094.00	+058.460333000	+000.249167000	+00360.00	0.00	-4.29	0.63	17.81	26.26	6	1

Annex 3 Estimating the Realism Factor

This Annex illustrates how the Realism Factor was computed in the case of Flight 2 Approach 10.

Figure A3.1 is an extract from the time-tagged satellite tracking list (box 6 in Figure 3) recorded during the flight. The first numeric line shows that on 5^{th} August 1996 at 1428 hours and 2.9894782 seconds the receiver tracked 7 satellites, these were PRNs 17, 28, 21, 23, 03, 31 and 22.

Table A3.1 shows for part of the approach the comparison between the satellites that were expected, based on the simulations, and the satellites that actually featured in the tracking file.

The Glossary section of this report defined the Realism Factor as

$$F = B/C$$

where

$$B = \Sigma\ (T_r) \quad \text{summing from } r = 1 \text{ to } r = n$$

$$C = \Sigma\ (E_r) \quad \text{summing from } r = 1 \text{ to } r = n$$

Table A3.1 shows for the r^{th} sample the PRNs expected to be tracked, the number expected (E_r), the PRNs actually tracked and the number actually tracked (T_r). In each case the PRNs listed are only those for which there was a valid MF correction. Thus, since there was no correction for PRN03 (the reason for this was unknown) in computing the Realism Factor PRN03 was ignored even though this featured in the receiver's actual tracking list.

Figure A3.1 Extract from the tracking file for flight 2, approach 10

DATE			HH	MM	SECONDS		N	PRNs							
96	5	8	14	28	2.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	3.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	4.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	5.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	6.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	7.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	8.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	9.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	10.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	11.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	12.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	13.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	14.9894782	0	8	17	28	21	19	23	3	31	22
96	5	8	14	28	15.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	16.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	17.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	18.9894782	0	7	17	28	21		23	3	31	22
96	5	8	14	28	19.9894782	0	7	17	28	21		23	3	31	22

Table A3.1 Comparison between expected and actual tracking

r	Min	Sec	Expected		Actual		Factor
			PRNs	E_r	PRNs	T_r	
1	26	33	17 19 21 22 23 28 31	7	17 19 22 23 28 31	6	0.857
2	26	34	17 19 21 22 23 28 31	7	17 19 22 23 28	5	0.714
3	26	35	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
4	26	36	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
5	26	37	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
6	26	38	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
7	26	39	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
8	26	40	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
9	26	41	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
10	26	42	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
11	26	43	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
12	26	44	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
13	26	45	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
14	26	46	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
15	26	47	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
16	26	48	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
17	26	49	17 19 21 22 23 28 31	7	17 19 21 22 23 28 31	7	1.000
18	26	50	17 19 21 22 23 28 31	7	17 19 21 22 23 28 31	7	1.000
19	26	51	17 19 21 22 23 28 31	7	17 19 21 22 23 28 31	7	1.000
20	26	52	17 19 21 22 23 28 31	7	17 19 21 22 23 28 31	7	1.000
150	29	02	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
151	29	03	17 19 21 22 23 28 31	7	17 21 22 23 28 31	6	0.857
152	29	04	17 19 21 22 23 28 31	7	17 21 22 23 28	5	0.714
153	29	05	17 19 21 22 23 28 31	7	17 21 22 23 28	5	0.714
154	29	06	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
155	29	07	17 19 21 22 23 28 31	7	17 19 21 22 23 28	6	0.857
156	29	08	17 19 21 22 23 28 31	7	17 21 22 23 28	5	0.714
Totals				1092		982	

Note that for all samples the time was between 1400 and 1500 hours, thus the first sample was at 14h26m33s etc.