

4000126063/18/NL/MP
NAVISP-EL3-001



D3a EGNOS and Beacons Final Report

11th March 2020

MarRINav

MARITIME RESILIENCE AND INTEGRITY OF NAVIGATION



MarRINav is a project delivered on behalf of the European Space Agency



MarRINav – Maritime Resilience and Integrity in Navigation Work Package 2 EGNOS and Beacons Final Report

Version	Date	Author	Reason for change
1.1	25.11.2019	C Hill	Initial submission
2.1	11.03.2020	C Hill	Final submission

© NLA International Limited 2020

The copyright in this document is vested in NLA International Limited.

This document may only be reproduced in whole or in part, or stored in a retrieval system, or transmitted in any form, or by any means electronic, mechanical, photocopying or otherwise, either with the prior permission of NLA International Limited or in accordance with the terms of ESA Contract No. 4000126063/18/NL/MP.



Document Information

Client	ESA
Project Title	MarRINav – Maritime Resilience and Integrity in Navigation 4000126063/18/NL/MP NAVISP-EL3-001
Deliverable Number	D3a
Report Title	D3a EGNOS and Beacons Final Report
Report Version	V2.0
Report Version Date	11 th March 2020
Lead Author(s)	Name: Dr Chris Hill Organisation: University of Nottingham
Contributing Author (s)	Name: Dr Lukasz K Bonenberg Organisation: University of Nottingham
Project Manager	Richard Greaves richard.greaves@nlaltd.co.uk NLA International Ltd
Circulation	1. Client 2. Project Files
File Name	20 03 11 D3a EGNOS and Beacons Final Report v2.1.docx
File Location	Googledrive, Dropbox



Executive Summary

The analysis presented in this report has attempted to answer two related questions:

1. Does EGNOS satisfy the maritime requirement for 10m (95%) positioning accuracy, with 99.8% availability and 99.97% continuity over 15 mins? [‘IMO A.915 - Revised Maritime Policy and Requirements for Future Global Navigation Satellite System (GNSS)’, IMO, Resolution A.915(22), Jan. 2002.
2. If not, do IALA beacons have a beneficial role when EGNOS fails to meet the required accuracy?

To address these questions, a number of simulations, using the Iguassu SBAS Simulator Version 2 (SSv2) have been performed. To ensure that the results from SSv2 are realistic, an extensive test and calibration exercise was carried out prior to this work. This is documented in the D2 deliverable from this project. The result of this exercise was a configuration for SSv2 which enables performance predictions that statistically match the real performance of EGNOS as determined from the post-processing of sample data sets. The results of this exercise are summarised in the first part of this report.

With a properly calibrated and configured simulation tool, performance simulations have been performed to assess whether EGNOS meets the maritime requirement for availability of 10m (95%). Assessment of continuity has not been performed due to the limitations of a deterministic simulation and a 24-hour constellation repeat period. Continuity is closely related to availability: any occurrence of a predicted positioning error of greater than 10m is an ‘availability event’, which in turn would trigger ‘continuity events’. A single availability event in a 24-hour simulation will affect a 15-minute continuity requirement for at least $15/(24*60) = 1\%$ of the time, hence exceeding the 99.97% continuity requirement. This analysis has therefore not attempted to determine continuity from the accuracy statistics, but has identified various availability events and has then attempted to determine the origins of these events.

The analysis has found that the identified availability events are not caused by:

- The geometrical weakness of the GPS constellation (represented through DOP).
- Exclusion of satellites from the EGNOS positioning solution due to inadequate RIMS tracking.

The most likely cause of the availability events appears to be the exclusion of satellite measurements from the positioning solution due to a lack of ionospheric corrections from the EGNOS ionospheric grid. Since EGNOS is currently a service that caters for single frequency receivers, users are required to correct their single frequency measurements using an ionospheric correction interpolated from a grid that is determined by the RIMS network



and broadcast as part of the EGNOS data message. If a user's measurement passes through a part of the ionospheric grid that is not adequately monitored by the EGNOS RIMS network, then no correction is possible, and the affected measurement must be excluded from the positioning solution. This will result in poorer DOP and probably a poorer positioning solution. We have identified through simulation and processing of historical data, a case that appears to repeat each day (appearing 4 minutes earlier each time); this is investigated specifically on days 075 and 134 of 2019.

Monitoring of the EGNOS ionospheric grid relies on a dense distribution of ionospheric pierce points (IPPs) generated by the RIMS network. A dense network of RIMS locations, each generating a scatter of IPPs within a certain elevation-dependent radius of the RIMS location, will produce a dense scatter of IPPs throughout the coverage area. On the periphery of the RIMS network the density of the IPPs generated by the RIMS network will decrease, to the point that grid points beyond this coverage area will not be adequately monitored. However, it is still possible that users' measurements will pass through these inadequately monitored parts of the network and will consequently have to be excluded from the positioning solution. In the area of interest, this is most likely to happen in the North West part of the British Isles.

Having identified this as the likely cause of the availability events observed in the simulations, this report then addresses whether the use of IALA beacons can solve these availability events, i.e. whether they can keep the accuracy to better than 10m. Two approaches were investigated:

1. Does conventional beacon-based DGPS provide a positioning solution with a better accuracy than EGNOS positioning in the cases where EGNOS would require satellite measurements to be excluded due to poor ionospheric grid monitoring?
2. Could the inclusion of additional RIMS locations (possibly close to the infrastructure of the existing DGPS beacon sites) solve the inadequate monitoring of the ionospheric grid?

To investigate whether conventional DGPS can meet the 10m requirement, we have processed historic measurement data from a period covering the EGNOS availability events. The processing has shown that the spike in performance is present also in the DGPS positioning results. This was an unexpected result, and suggests that uncorrelated measurement errors are present in the user receiver measurements and/or the reference receiver measurements. A full analysis of these errors has not been possible to date. The spike in performance is below the 10m accuracy requirement, so although this needs further investigation the main conclusion is that DGPS can indeed provide the required accuracy even at a significant distance (several hundred km) from the beacon.

The inclusion of additional RIMS locations has proven to be particularly effective at improving the monitoring of the EGNOS ionospheric grid, to the point that a few, e.g. three, additional RIMS in the UK would fill those holes in the ionospheric grid that were observed in these analyses.



In addition to the repeating spike in predicted performance, testing also revealed a case where a spike above 10m in the actual performance was not predicted by the 95% performance prediction. Whilst this single event was within the 5% of the statistically expected cases, a brief investigation revealed that this seems to have been caused by a combination of a reduced number of available satellites (above a 15° elevation mask) and the probable down-weighting of one or more satellites due to EGNOS model weighting. The nature and cause of the down-weighting remains to be explored.

This report also includes a section in which a specific detail of the simulations was investigated. The exercise to calibrate and configure SSv2 identified that the optimum calibration to suit the north of the UK resulted in slightly pessimistic simulation in the south of the UK. This suggested that the EGNOS corrections used by receivers in the south of the UK were of a better quality than those used by receivers in the north of the UK. Although not strictly necessary for the simulations, an analysis of the factors affecting the accuracy of the EGNOS ephemeris and clock corrections was undertaken. The concept of IDOP (inverse DOP, between one satellite and multiple RIMS locations) was explored in an attempt to assess whether this is a better indicator of ephemeris and clock accuracy than a simple count of the number of RIMS tracking a given satellite ('tracking depth'). The analysis has shown some subtle differences between IDOP and tracking depth but has not yet been conclusive as to whether IDOP is a better metric.

The report has identified areas where further investigation should be carried out

- It is clear that the elevation mask angle applied to the EGNOS positioning solutions has a significant effect on the performance. The trade-off between improved satellite geometry and increased measurement errors, arising from variations in the user receiver elevation mask, should be investigated, to determine whether the exclusion of satellites due to ionospheric monitoring can be mitigated by the inclusion of measurements with slightly higher measurement errors. This should help to inform the choice of a standardised user receiver elevation mask.
- The validity of the SSv2 results should be further investigated by comparing the modelled performance with performance predicted from historic EMS data.
- To investigate whether there is correlation between the computed error in the EGNOS ephemeris and clock corrections (versus the Precise Ephemeris) and either IDOP or the RIMS tracking depth.



Contents

Table of Contents

DOCUMENT INFORMATION	4
EXECUTIVE SUMMARY	5
CONTENTS	8
GLOSSARY	9
1 INTRODUCTION	11
2 SUMMARY OF TOOLS VALIDATION	13
3 PART 1: VARIATIONS IN THE QUALITY OF EGNOS CORRECTIONS	16
4 PART 2: PERFORMANCE OF EGNOS FOR MARITIME OPERATIONS	24
4.1 BASIC SATELLITE GEOMETRY	27
4.2 SATELLITE EXCLUSION	29
4.2.1 <i>Exclusion due to inadequate RIMS tracking</i>	29
4.2.2 <i>Exclusion due to poorly monitored Ionosphere Grid</i>	30
4.2.3 <i>Impact of Ionospheric Grid Monitoring further from RIMS network</i>	36
4.3 EGNOS CORRECTION WEIGHTING.....	38
5 BENEFITS OF IALA BEACONS	48
5.1 BEACON DGPS PERFORMANCE	49
5.2 ADDITIONAL RIMS LOCATIONS	54
6 CONCLUSIONS AND RECOMMENDATIONS	61



Glossary

3G	Third-generation Cell-phone Technology
AL	Alert Limit
A-RAIM	Advanced Receiver Autonomous Integrity Monitoring
C/N0	Carrier-to-noise (spectral Density) Ratio
C1	GPS Code L1 Observation
CDF	Cumulative Distribution Function
CEP	Circular Error Probable
CNMP	Code-noise And Multipath
CORS	Continuously Operating reference Station
DGNSS	Differential GNSS
DGPS	Differential GPS, herein used as a short form of ‘marine beacon DGPS’
DOP	Dilution Of Precision
EC	European Commission
EDAS	EGNOS Data Access Service
EEZ	Exclusive Economic Zone
EGNOS	European Geostationary Navigation Overlay Service
EMRF	European Maritime Radionavigation Forum
EMS	EGNOS Message Server (archive of historic transmitted EGNOS messages)
EPFS	electronic Position Fixing Systems
ESA	European Space Agency
ESSP	European Satellite Services Provider
FLT	SBAS Fast And Long-term Correction Messages
GAGT	GLA Augmented-GNSS Toolset
GEO	Geostationary Satellite
GIVE	Grid Ionospheric Vertical Error
GLA	General Lighthouse Authority
GLONASS	Globalnaya Navigazionnaya Sputnikovaya Sistema
GNSS	Global Navigation Satellite System
GPRS	General Packet Radio Services, Second-generation Cell-phone Technology GPS
GSA	European Global Satellite System Agency
HAL	Horizontal Alert Limit
HDOP	Horizontal Dilution Of Precision
HMI	Hazardously Misleading Information
HNSE	Horizontal Navigation Standard Error
HNSP	Horizontal Navigation Standard Precision
HPL	Horizontal Protection Level
IALA	International Association Of Aids To Navigation And Lighthouse Authorities
IBPL	Isotropy-based Protection Level
ICD	Interface Control Document
IDOP	Inverse
IGP	Ionospheric Grid Point
IGS	International GNSS Service



IMO	International Maritime Organization
IONEX	Real World Ionospheric Measurements
IPP	Ionospheric Pierce Point
L1	GPS Upper L-band Transmission Frequency (1575.42 Mhz)
LoS	Line Of Sight
MHSS	Multiple-hypothesis Solution-separation
MOPS	Minimum Operational Performance Standards
MP	Multipath
M-RAIM	Maritime – Receiver Autonomous Integrity Monitoring
MSR	Multi-system Receiver
NAVISPS	Navigation Innovation And Support Programme
NGI	Nottingham Geospatial Institute
NLoS	Non Line-of-sight
NOAA	National Oceanic And Atmospheric Administration, United States Department Of Commerce
PDF	Probability Distribution Function
PE	Position Error
PL	Protection Level
PPM	Parts Per Million
PVT	Position Navigation Time
R95	Horizontal 95th Percentile Errors
RAIM	Receiver Autonomous Integrity Monitoring
RIMS	Receiver Integrity Monitor Stations
RINEX	Receiver Independent Exchange Format
RMS	Root Mean Square
RNP	Required Navigation Performance
RS	Reference Station
RTCM	Radio Technical Commission Maritime
SBAS	Satellite-based Augmentation System
SOLAS	Safety Of Life At Sea
SPP	Single Point Positioning
SSv2	SBAS Simulator V2
TLE	Two-line Element Set Files
TEC	Total Electron Count (a measure of ionospheric disturbance)
TSO-RAIM	Technical Standard Order – Receiver Autonomous Integrity Monitoring
TTA	Time To Alarm
UERE	User Equivalent Range Error
UDRE	User Differential Range Error
UIVE	User Ionospheric Vertical Error
UK	United Kingdom
WAAS	Wide Area Augmentation System
WP	Work Package
ZTD	Zenith Tropospheric Delay



1 Introduction

This report presents the results of a study to determine whether the accuracy and availability of EGNOS meet maritime requirements. The study has focussed on accuracy and has studied in particular the origins of occasions when the predicted 95% accuracy spikes above 10m.

The analyses have relied on simulations, using the Iguassu SBAS Simulator Version 2 (SSv2). To ensure that the results from SSv2 are realistic, an extensive test and calibration exercise was carried out prior to this work. This is documented in the D2 deliverable from this project. The result of this exercise was a configuration for SSv2 which enables performance predictions that statistically match the real performance of EGNOS as determined from the post-processing of sample data sets. The results of this exercise are summarised in Section 2 of this report.

Section 3 presents the results of a short study into variations in the apparent quality of EGNOS corrections across the UK. Using a consistent simulation model, variations in the comparison with the results from historic data were shown in the SSv2 calibration exercise, suggesting that the EGNOS corrections are of a higher quality in the south of the UK. The study in this section investigated the concept of ‘inverse DOP’ (IDOP, between one satellite and multiple RIMS locations) to assess whether the geometrical strength of the RIMS network was a better indicator of the quality of orbit and clock corrections than a simple count of the number of RIMS (the so-called tracking depth). This analysis has shown some subtle differences between IDOP and tracking depth, which may be sufficient to account for the apparent variations in the quality of the corrections. However, a complete analysis of the accuracy of the corrections (e.g. versus post-mission precise values) and their correlation with IDOP and tracking depth, has not been possible in the time available, so this remains an open question.

Section 4 shows the results of analyses which identify a repeating deficiency in the performance of EGNOS for a station in the north of the UK (Stornoway, STOR). Predicted performance is shown to spike above 10m on a daily repeatable basis, 4 minutes earlier each day in accordance with the shifting of the GPS constellation. The causes of this spike are investigated, from the strength of the basic GPS constellation to possible exclusion of satellites due to tracking deficiency or inadequate EGNOS ionospheric grid monitoring. The analysis appears to identify a lack of monitoring in certain parts of the ionospheric grid as the main cause of this spike, which particularly affects users in the north of the UK; Stornoway is affected, whereas Margate, on the south coast, does not suffer from the same deficiency. This section also includes a short analysis of a spike in real performance that is not predicted by SSv2, concluding that it is probably due to high uncertainties in the EGNOS corrections, leading to one or more satellites being down-weighted.

Section 5 of the report investigates whether the use of IALA beacons can mitigate the accuracy spikes observed in the EGNOS positioning results. This section considers first the conventional use of the beacons to provide DGPS corrections, and shows that DGPS positioning is able to perform within the 10m accuracy requirement during the period when



the EGNOS accuracy spikes. The analysis then considers the hypothetical case where additional EGNOS RIMS are created close to some beacon locations. It shows that additional RIMS locations would be able to improve the depth of monitoring of the EGNOS ionospheric grid, such that the exclusion of satellites due to unmonitored ionospheric grid points would no longer be necessary.

The report is concluded in Section 6.



2 Summary of Tools Validation

The purpose of the validation of the software tools that was presented in D2 was to find the appropriate configuration options such that predictions made with the tools are statistically representative of the performance that is obtained from real data.

SSv2 aims to predict the performance of EGNOS for aviation, and as such it intentionally overestimates the errors in the EGNOS corrections in order to overbound the random errors and biases in those corrections, to ensure that the resulting position solutions have aviation levels of integrity. SSv2 was developed to provide estimates of the errors in the EGNOS corrections that match the error estimates broadcast by EGNOS. For the majority of the time, EGNOS performs better than indicated by these high integrity aviation models. The task in D2 was to unpick that overbounding and find ways to configure the models so that the predicted positioning performance is more representative of typical EGNOS performance. This meant scaling down some of the estimates of the errors in the EGNOS corrections.

The expectation was that SSv2 overbounds the error estimates by a factor of 2.5 to 3, and indeed it was found that by scaling down SSv2's model of σ_{UDRE} (the model that predicts the accuracy of satellite orbits and clocks) by a factor of 3 (from 0.91m to 0.3m), the predicted positioning performance matched well with the performance observed with historic measurement data from a number of locations in the UK and under varying atmospheric conditions. Figure 1 to Figure 4 below show the relationship between different values of σ_{UDRE} and the percentage of position solutions meeting the 95% position errors predicted by SSv2. Since the aim is to configure the software so that 95% of position errors meet the predicted 95% position errors, the task was to find a σ_{UDRE} value that corresponds to where the curves cross the horizontal red 95% line. On each chart, curves are shown for three separate days representing different ionospheric conditions:

- Day 075 = 16th March 2019, classified as the 4th most ionospherically 'disturbed' day of the month (but not particularly disturbed in an absolute sense).
- Day 107 = 17th April, an ionospherically quiet day, classified as the 4th quietest day of the month.
- Day 135 = 14th May 2019, an ionospherically disturbed day, classified as the most 'disturbed' day of the month.

The charts for STOR (Stornoway, northern Scotland, 58° 12' N, 6° 23' W), KIRW (Kirkwall, northern Scotland, 58° 57' N, 2° 54' W) and NGB2 (Nottingham, English midlands, 52° 57' N, 1° 11' W) all show that a value of 0.3m for σ_{UDRE} gives the desired result for all three days, and a further test of sensitivity to this value demonstrated in D2 that the small spread of crossing points was not a concern.



The chart for MART (Margate, English south coast, 51° 24', 1° 23' E) shows that a smaller value for σ_{UDRE} is required, in the range of 0.1m to 0.25m. This suggests that the position solutions derived from real EGNOS corrections were generally of even better quality compared to the predictions than those from STOR, KIRW and NGB2. Since the predictions already take account of satellite geometry, this better performance of EGNOS-derived positions at MART suggests that the EGNOS corrections themselves are of a better quality for the receiver at MART than for the receivers at STOR, KIRW and NGB2. Section 3 of this report explores some possible reasons for this.

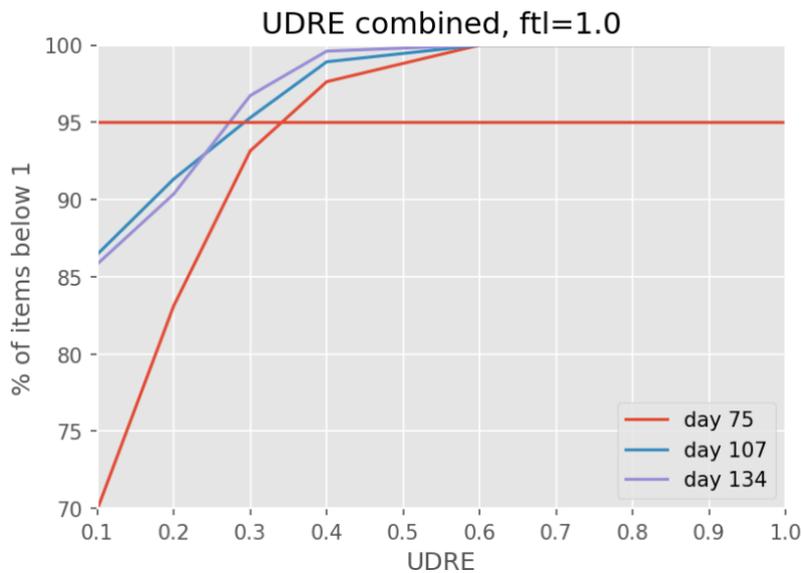


Figure 1: Effect of σ_{UDRE} on error bounding (STOR)

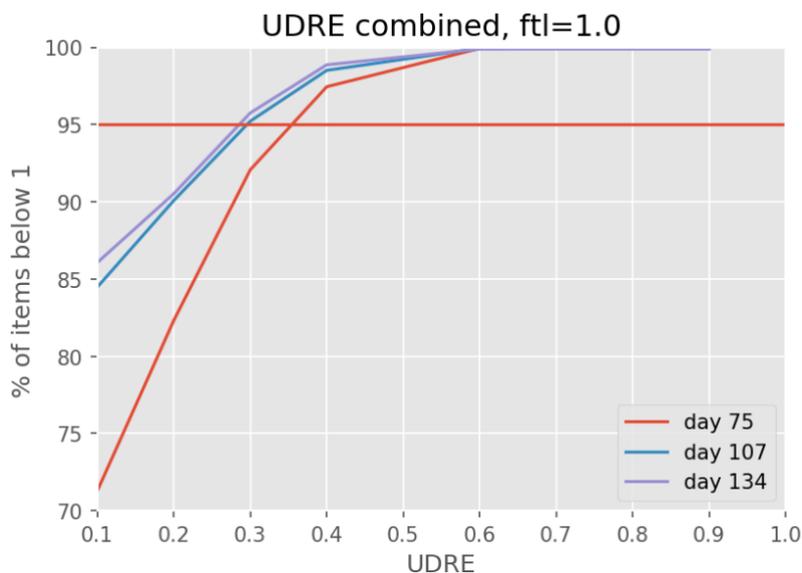


Figure 2: Effect of σ_{UDRE} on error bounding (KIRW)



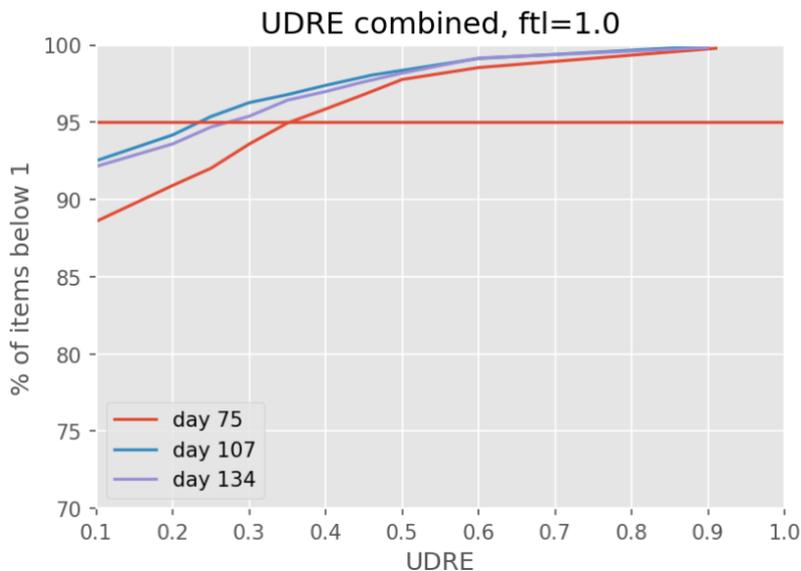


Figure 3 Effect of σ_{UDRE} on error bounding (NGB2)

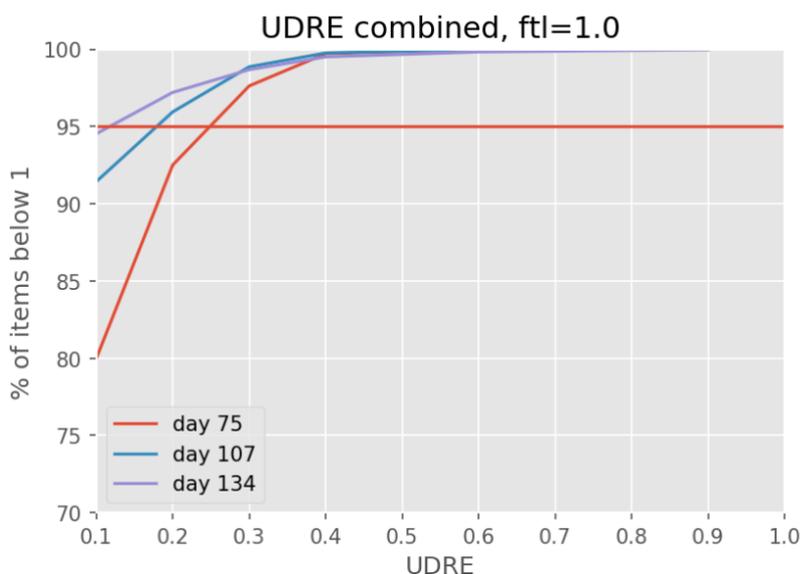


Figure 4: Effect of σ_{UDRE} on error bounding (MART)

The result of this validation and calibration is that SSv2 can be configured to produce realistic estimates of actual EGNOS performance. With a σ_{UDRE} value of 0.3 the results at STOR, KIRK (northern UK) and NGB2 will be representative, but the results at MART (southern UK) may be slightly pessimistic. This is preferable to having realistic predictions in the south and optimistic predictions in the north. The model is effectively tuned best for the areas of interest to this project.

With the SSv2 models configured accordingly, the tool can now be used to predict the performance of EGNOS over a given service area, and the predicted performance can be assessed against maritime requirements. This is covered in section 4.



3 Part 1: Variations in the Quality of EGNOS Corrections

The model in SSv2 which determines the predicted error in the orbit and clock corrections (σ_{UDRE}) uses a simple look-up table based on the number of RIMS observing the satellite. With 8 or more RIMS observing the satellite the model predicts that σ_{UDRE} will be 0.91m (Figure 5).

RIMS	σ_{UDRE} [m]
0	NM
1	NM
2	NM
3	4.56
4	4.56
5	2.27
6	1.37
7	1.14
>= 8	0.91

NM = not monitored

Figure 5: Configuration of σ_{UDRE} within SSv2.

It was shown in D2 that, due to the large number of RIMS and their tracking down to lower elevations than a typical user elevation mask, very few satellites visible in or around the UK will ever be monitored by fewer than 8 RIMS (see Figure 6 for example).



RIMS analysis for 134_A_UdreDfre

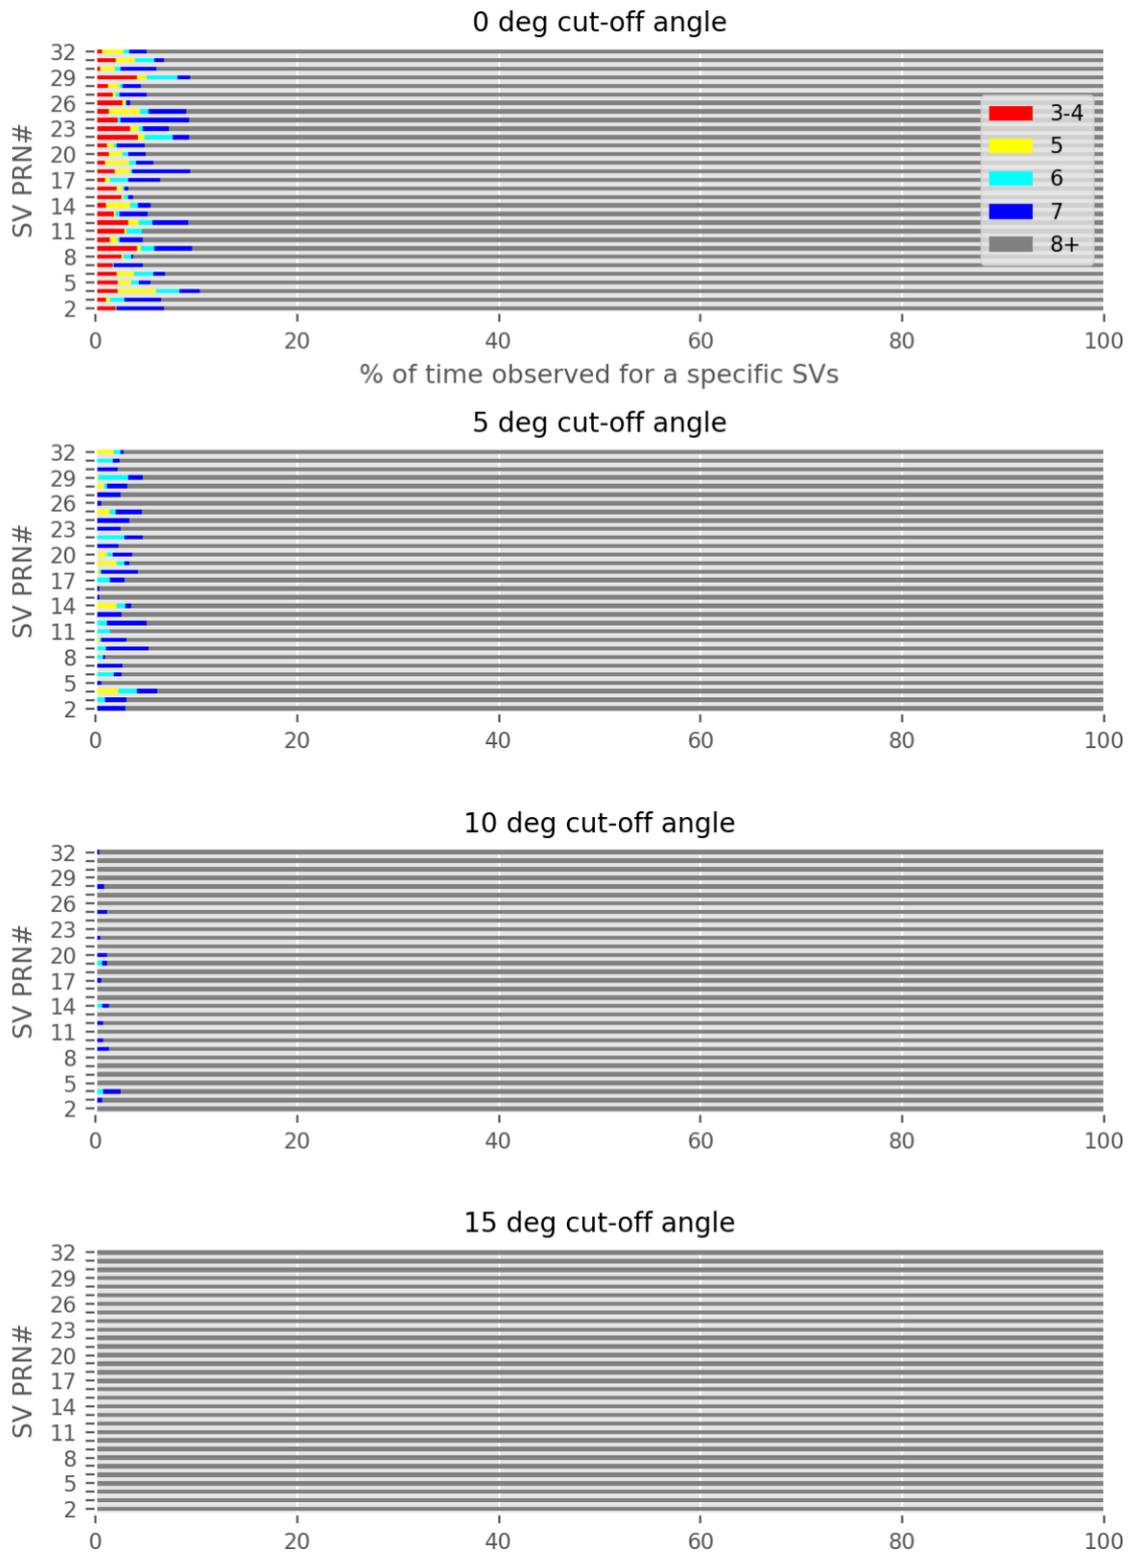


Figure 6: RIMS tracking depth for satellites visible from 65 °N, 12°W



To analyse this further we have mapped the visibility of RIMS stations from anywhere on the surface of the globe. For this exercise we have used the 38 RIMS stations that Ssv2 defines as it's default configuration. These are detailed in the table below. Whilst this network may not accurately reflect the current operational network, the conclusions regarding tracking depth will not be significantly affected by small differences in the network.

RIMS ID	Latitude (°)	Longitude (°)
ACR	38.51	-28.62
AGA	30.50	-9.66
ALB	57.10	9.90
ALX	31.18	29.95
ATH	37.88	23.73
BRN	52.32	13.25
CNR	27.95	-15.38
CRK	51.85	-8.50
CTN	37.47	15.07
DJA	33.87	10.77
EGI	65.28	-14.40
GLG	55.70	-4.10
GOL	39.63	32.80
GVL	60.67	17.13
HBK	-25.88	27.70
JAN	70.92	-8.72
KIR	69.68	29.92
KOU	5.17	-52.68
LAP	61.53	27.55
LPA	28.62	-17.76
LSB	38.78	-9.13
MAD	32.75	-16.70
MLG	36.68	-4.52
MON	46.07	-64.78
NKC	18.07	-15.95
PAR	48.83	2.33
PDM	39.57	2.73
RKK	64.13	-21.93
ROM	41.80	12.58
SDC	42.92	-8.42
SOF	42.80	23.42
SVA	78.12	15.40
SWA	50.88	-1.28
TLS	43.42	1.50
TRD	63.45	10.90
TRO	69.67	18.95
WRS	52.22	21.07
ZUR	47.45	8.57

Table 1: Default Ssv2 RIMS locations, used to investigate Tracking Depth



Figure 7 shows contours of visibility from an altitude of 20,200 km, representing typical GPS satellite altitudes. The red crosses in the figure represent the RIMS locations. As can be seen, a satellite anywhere over Europe will be seen by nearly all of the available RIMS, and it is not until it is located in the orange/red area that a satellite would be tracked by fewer than 8 RIMS.

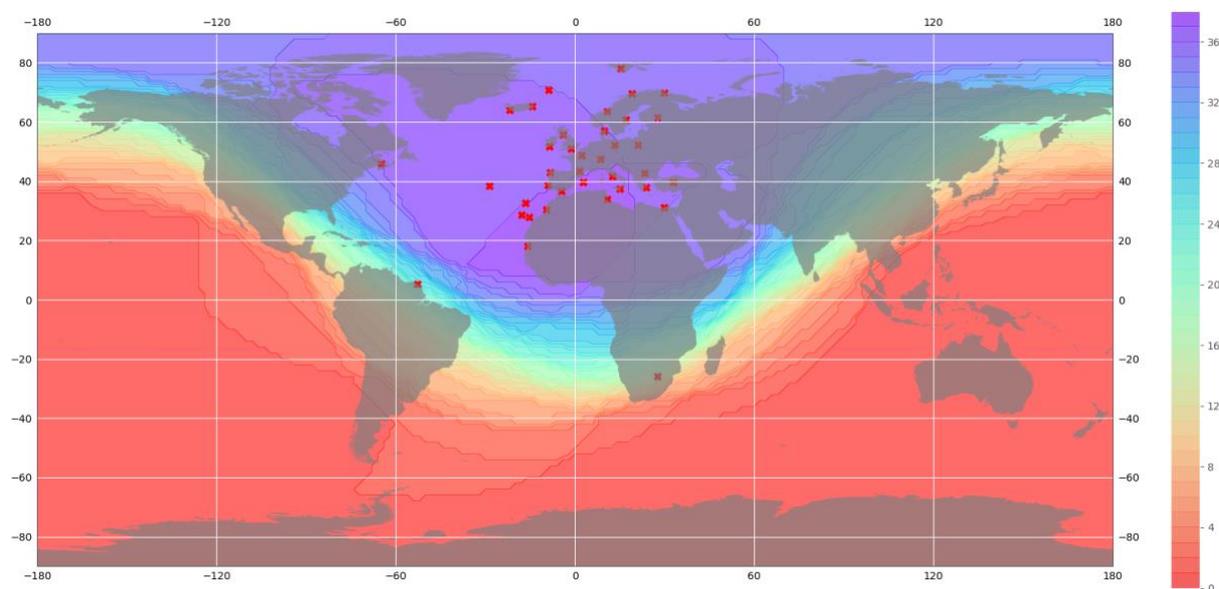


Figure 7: Global RIMS tracking depth

The SSv2 model of σ_{UDRE} will assign a uniform orbit and clock error of 0.91m for all satellites inside the 8 RIMS contour. From a purely geometrical consideration, in the same way that a simple count of the number of satellites that a user receiver can see is not a good indicator of the quality of the user's positioning, it seems unlikely that a simple count of the number of visible RIMS will properly reflect the quality of the satellite's orbit and clock estimation.

In the case of user positioning, the concept of DOP (Dilution of Precision) is a better indicator of the positioning quality, so we have attempted to apply the same concept, inverted, to investigate the orbit and clock quality. Again, assuming a typical satellite altitude of 20,200 km we have computed a Dilution of Precision figure based on the geometry between one satellite and all of the visible RIMS locations. In Figure 8 we have plotted a contour map of this 'Inverse DOP' (IDOP), for comparison with the simple tracking depth presented in Figure 7.

Firstly, the IDOP values illustrated in Figure 8 are very much higher than those typically experienced between one user and multiple satellites. This is because the geometry between one satellite and multiple RIMS is generally much poorer. Conventional (user positioning) DOP can be reasonably well visualised as being inversely proportional to the volume of the polyhedron with the user at one vertex and satellites at the other vertices. For a ground-based receiver with satellites well distributed in azimuth and elevation, the volume of such a polyhedron is typically very large. But for the polyhedron defined by one satellite and multiple RIMS, the volume is much smaller, since the RIMS are tightly clustered within a



relatively small region of the Earth’s surface, so their distribution in azimuth and elevation from the point of view of the satellite is poor.

The second point to note is that, although the shape of the IDOP contours is similar to the shape of the tracking depth contours, there are some important differences. The influence of the non-European RIMS can be observed as contour ‘walls’, where IDOP changes rapidly.

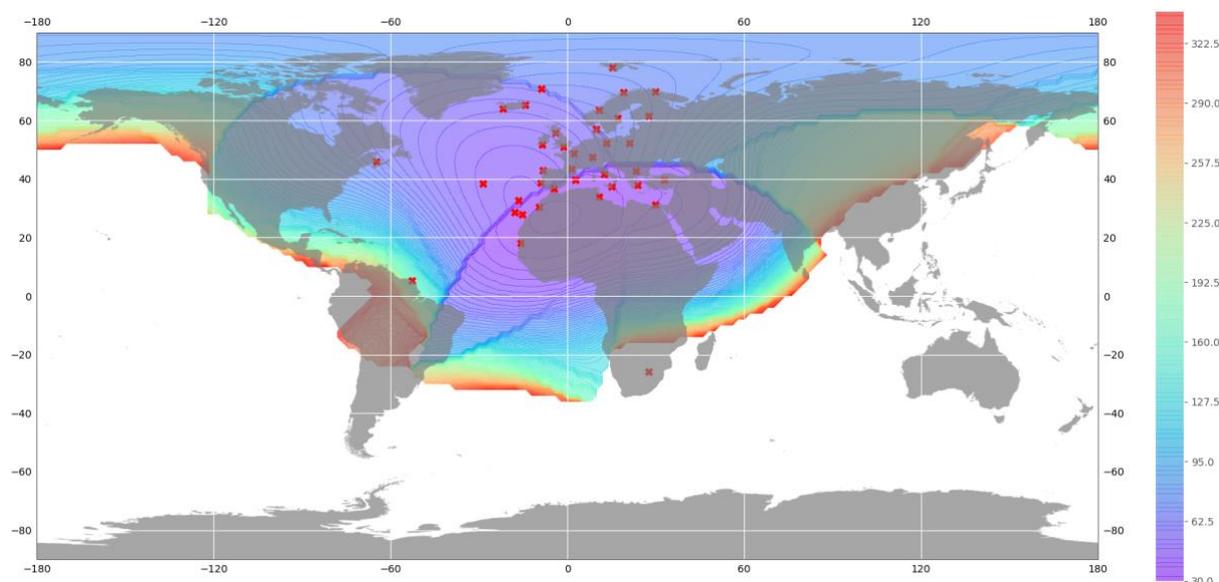


Figure 8: Global ‘Inverse DOP’

In an effort to determine whether the differences between RIMS coverage depth and IDOP can help to explain observed performance differences between STOR and MART, we have investigated the visible satellites from STOR and MART for a snapshot in time.

Figure 9 shows the tracking depth contour map from Figure 7 overlaid with the positions of the satellites visible from MART at the snapshot epoch. Figure 9 is also overlaid with the red visibility ‘circle’, representing the limits of the possible locations of satellites for them to be visible from MART above 10° elevation (the ‘circle’ does not show as a circle on this projection – it would be a perfect circle if viewed on the spherical Earth). It demonstrates that any satellite tracked by 8 or fewer RIMS (the orange/red zone of the contour map) would be outside the visibility circle for MART, i.e. not visible from MART.

Figure 10 shows the IDOP contour map overlaid with the same satellite locations and the same visibility circle. Most satellites visible from MART have an IDOP of better than 100, but there are areas where this rises to ~250.

For comparison, Figure 11 and Figure 12 show the same plots for STOR. Many of the satellites visible from STOR are common to those visible from MART. However, satellite 31, observed by ~24 RIMS, is not visible from STOR, whereas satellite 7, observed by ~12 RIMS is. In the SSv2 model all of these satellites would be assigned a σ_{UDRE} value of 0.91m.



Referring to the IDOP plots, satellite 31, observed only by MART, has an IDOP value of ~ 150 , whereas satellite 7, observed only by STOR, has an IDOP value above 300.

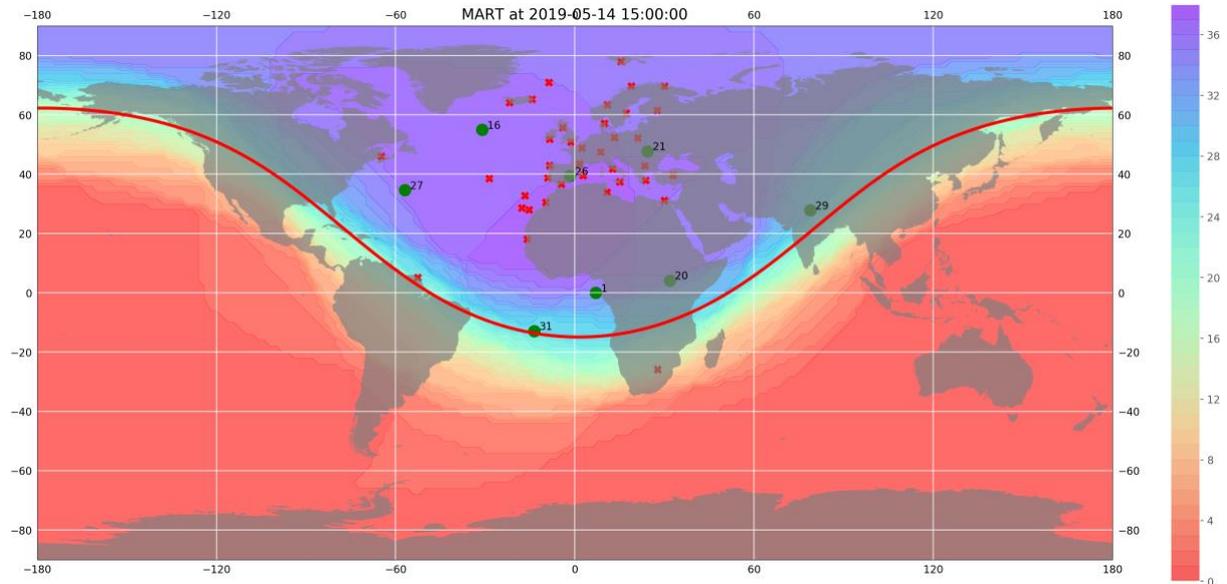


Figure 9: RIMS tracking depth for MART at 15:00, 14th May 2019.

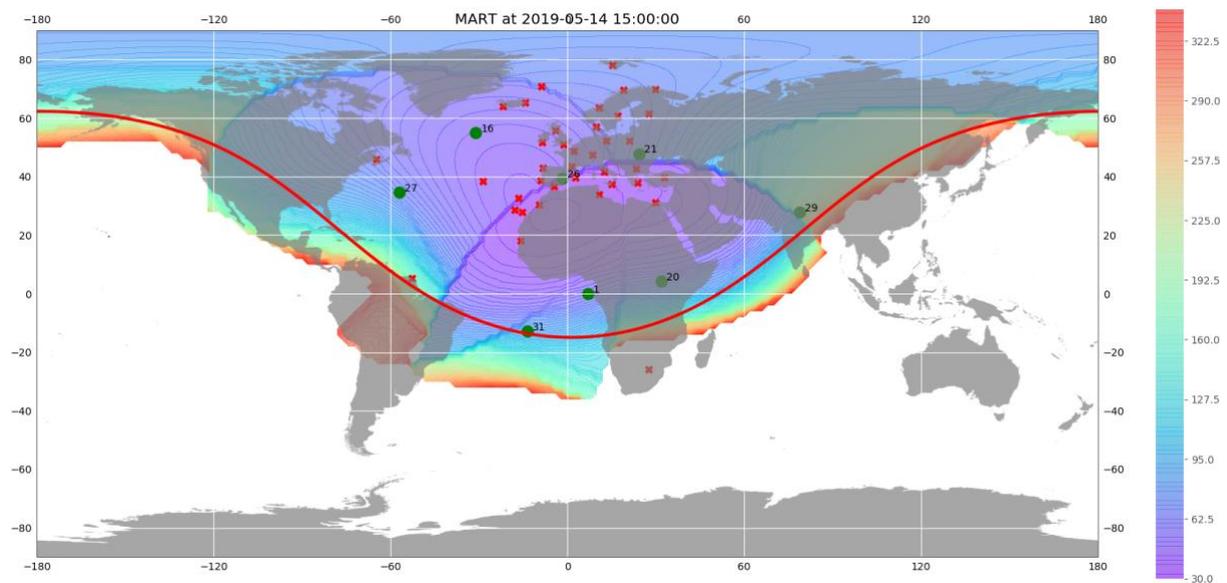


Figure 10: IDOP for MART at 07:45, 17th April 2019.



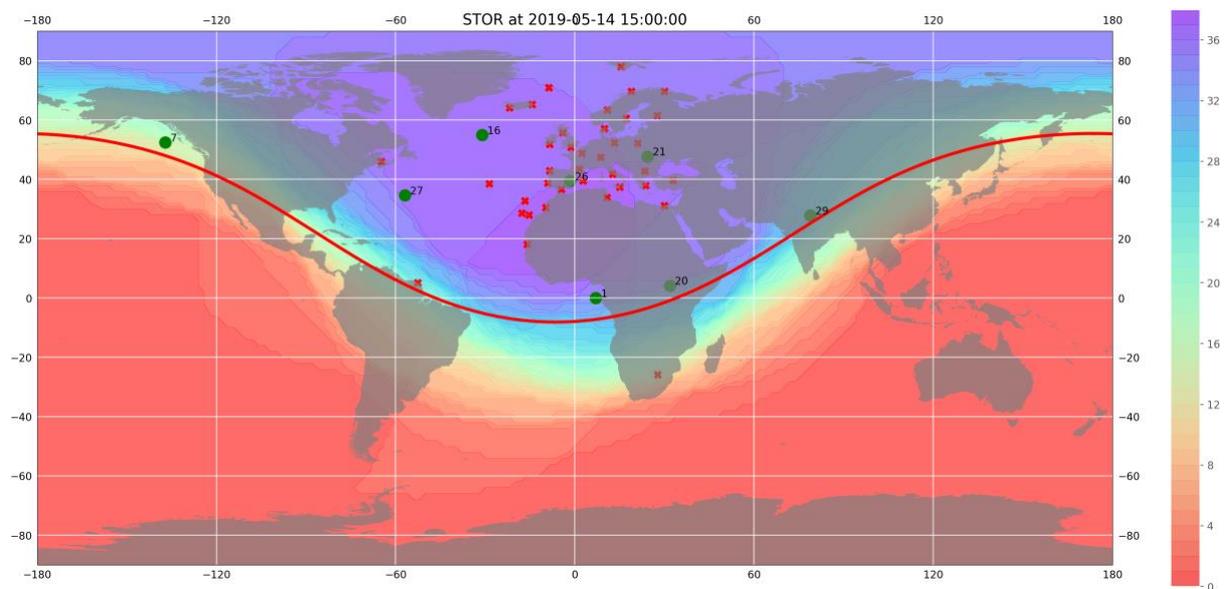


Figure 11: RIMS tracking depth for STOR at 15:00, 14th May 2019.

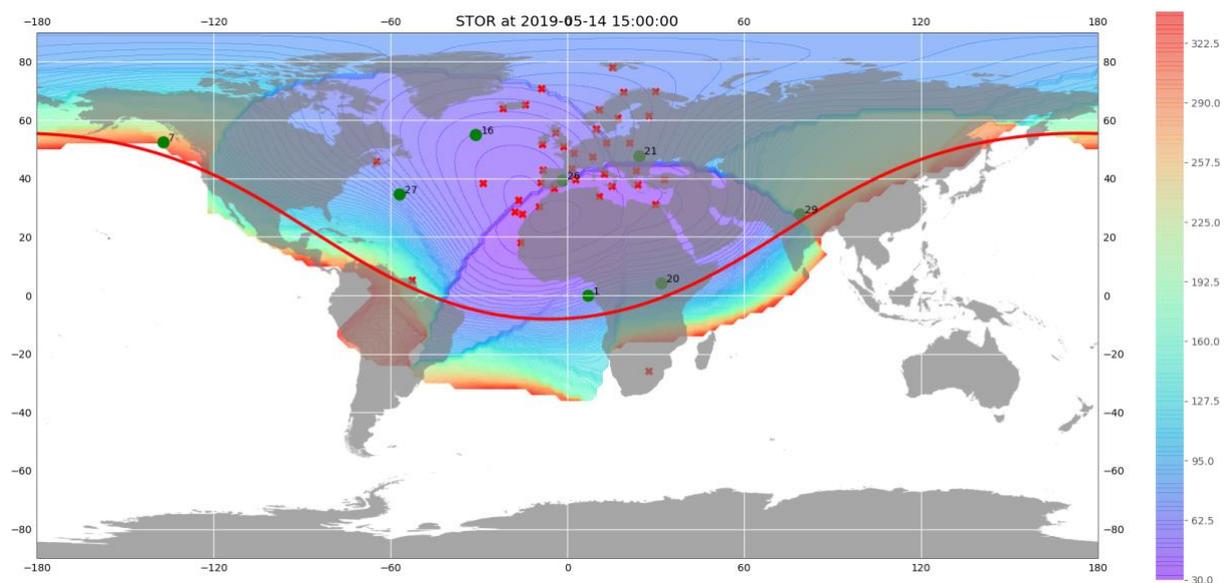


Figure 12: IDOP for STOR at 07:45, 17th April 2019.

The orbit and clock estimation algorithm in EGNOS uses a dynamic model that includes tracking from long periods, so instantaneous indicators such as the number of visible RIMS or their inverse DOP are unlikely to be fully representative of the accuracy of that process. Nevertheless, it does appear that the pessimistic performance predictions for MART (or conversely the optimistic predictions for STOR) could be due to the algorithm not distinguishing adequately between satellites that are well tracked and those that are not so well tracked. Whether the tracking depth is a suitable metric for ‘well tracked’, or whether IDOP provides a better metric, is not conclusively shown from these tests, however.



Further analysis could investigate whether there is correlation between the computed error in the EGNOS ephemeris and clock corrections (versus the Precise Ephemeris) and either IDOP or the RIMS tracking depth. The nature of any such correlation could be used to improve, if necessary, the σ_{UDRE} prediction model, such that predicted receiver performance better matches real performance.



4 Part 2: Performance of EGNOS for Maritime Operations

With the SSv2 tool calibrated to produce realistic results, especially at northern UK sites, we have used it to generate 95% accuracy plots for STOR and MART for some sample 24-hour periods, as well as availability maps covering the same 24-hour periods. These plots assume a user elevation mask of 15°. This is a typical elevation mask employed by user receivers, especially in survey applications, since it eliminates the lowest elevation observations which are likely to be more contaminated by errors from the atmosphere and multipath. However, it has not been determined whether this is typical of a maritime and/or EGNOS receiver. There is no doubt that receiver performance will vary as a function of the elevation mask that is chosen. A lower elevation mask will allow more satellite measurements to be included in the position solution, which will always have a beneficial effect on satellite geometry (DOP). However, these lower elevation measurements are also likely to have higher multipath errors, due to signal reflections from buildings (for land-based receivers) or vessel infrastructure (for marine receivers). Furthermore, low elevation measurements will also have greater atmospheric errors (both ionospheric and tropospheric), and the corrections (in the case of the ionosphere) will have greater uncertainties. If those uncertainties are used to weight the individual measurements, then the benefits to DOP are no longer as clear-cut: a lower weight means that the measurement will contribute less, not only to the combined errors but also to the overall geometrical strength. One such probable example is explored in section 4.3. So, it is possible that the results of an analysis using a 10° elevation mask will show different individual results from those presented above.

We first show time series plots of 95% accuracy computed at STOR and MART on the same three days that were used in the tool's validation. Each plot includes a red horizontal line to highlight the 10m Alert Limit so that events where this limit is breached can be seen clearly.

Each plot shows the predicted 95% accuracy from SSv2 simulations (the blue 'HNSP' line) and the real instantaneous errors (difference between computed solution and known truth) in positions computed by gLab from EMS data files (the red 'HNSE' line). Since the SSv2 predictions assume a 95% probability, the expectation is that 95% of the real errors (red lines) are bounded by the blue lines. It can be seen that there are a few occasions where the real position error noticeably exceeds the predicted error; this is expected since those real errors are from the 5% of the distribution that is not bounded by the 95% probability.

It is also evident that spikes in the real position error usually correspond to spikes in the predicted error. However, the magnitude of the spikes in the real positioning errors is rarely as large as the simulations predict. Again, this is due to the probability distribution. The simulations aim to predict the value which will bound 95% of the position errors, but at any given moment the actual measurement errors have a 95% probability of combining to give a position error less than this predicted bound. Nevertheless, it is the occasions when the predicted error exceeds 10m that should concern us here, since it is those occasions when measurement errors *could* combine to give an unacceptable position error – so-called



‘availability events’. It is clear that there are several occasions in the plots for STOR where the predicted 95% error exceeds 10m, but none in the plots for MART.

There is an interesting event that occurs at STOR at ~07:45 on day 075, ~05:30 on day 107 and ~03:45 on day 134. The fact that the GPS constellation advances by 4 minutes per day confirms that this is in fact the same ‘event’ (075 to 107 = 32 days = ~128 minutes; 107 to 134 = 27 days = ~108 minutes). The same event shows as no more than a slight degradation in the predicted performance at MART. This event therefore forms the focus of the following analysis.

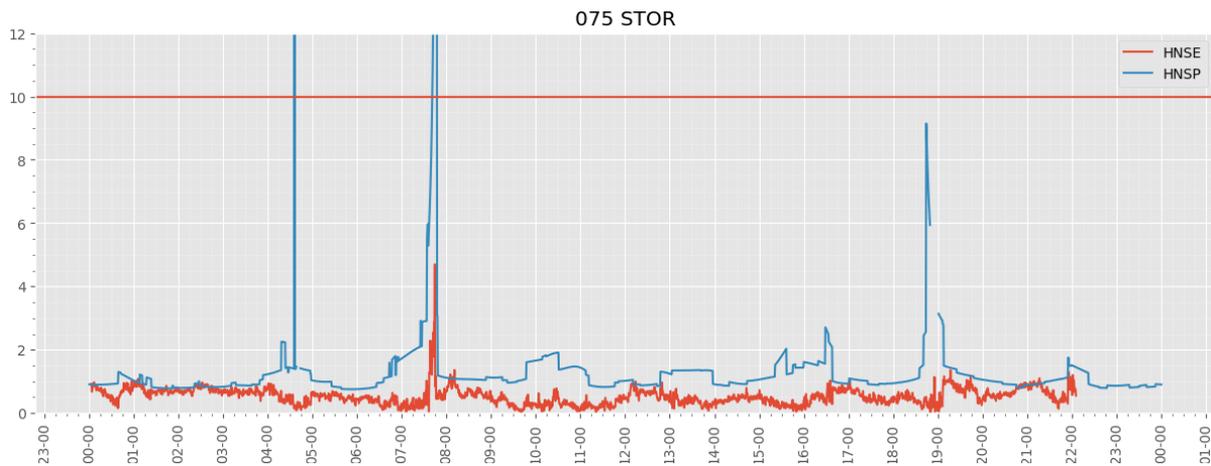


Figure 13: 95% accuracy plot for STOR on 16/03/19.

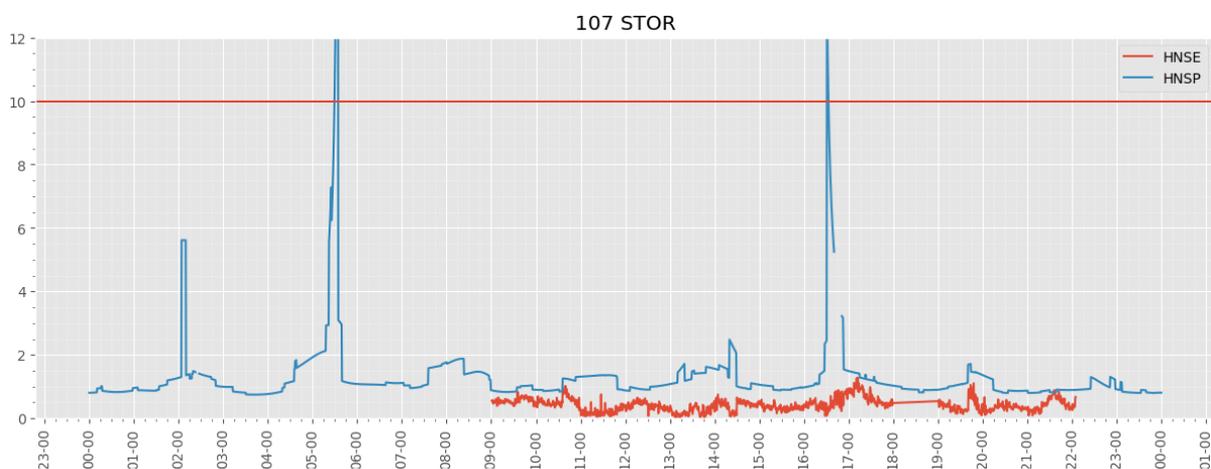


Figure 14: 95% accuracy plot for STOR on 17/04/19.



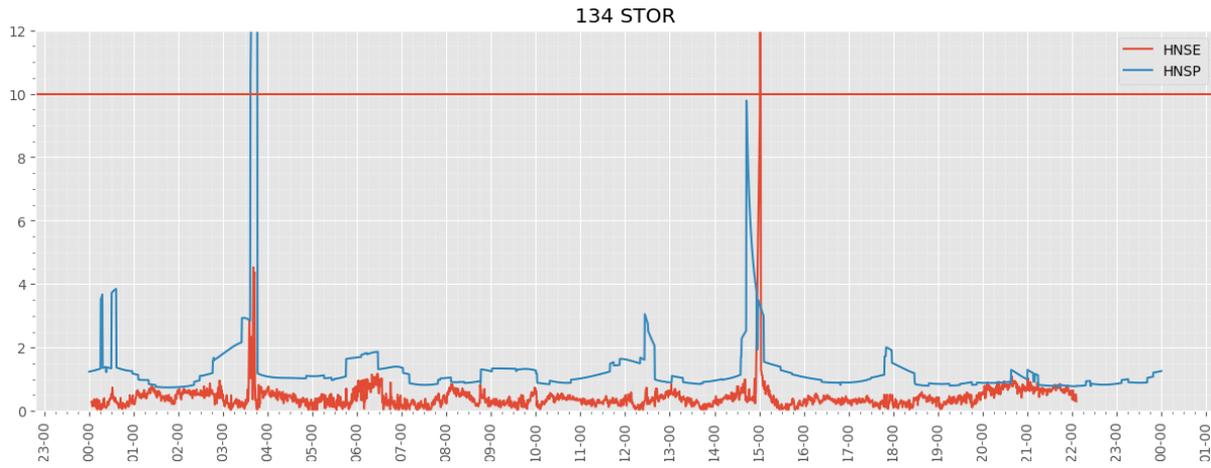


Figure 15: 95% accuracy plot for STOR on 14/05/19.

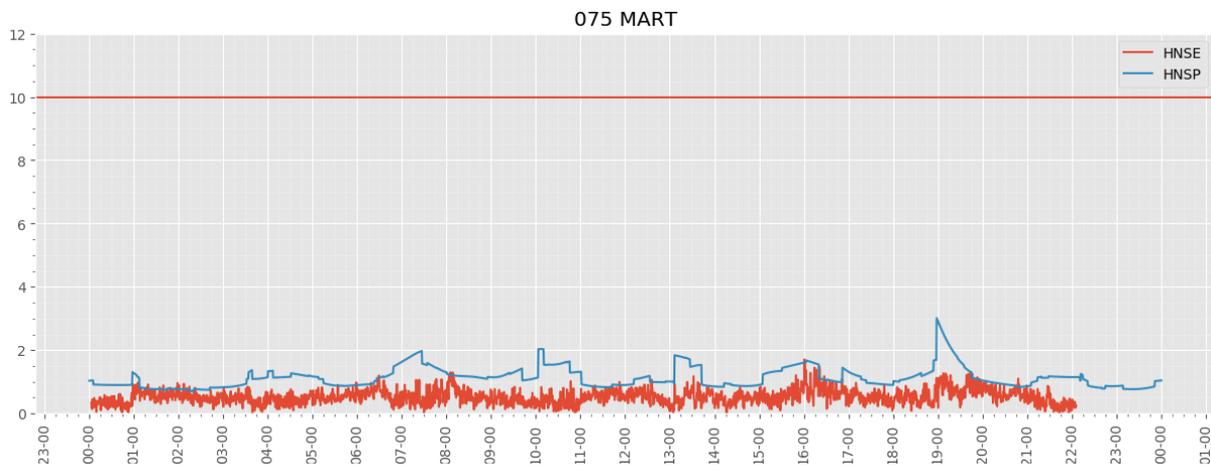


Figure 16: 95% accuracy plot for MART on 16/03/19.

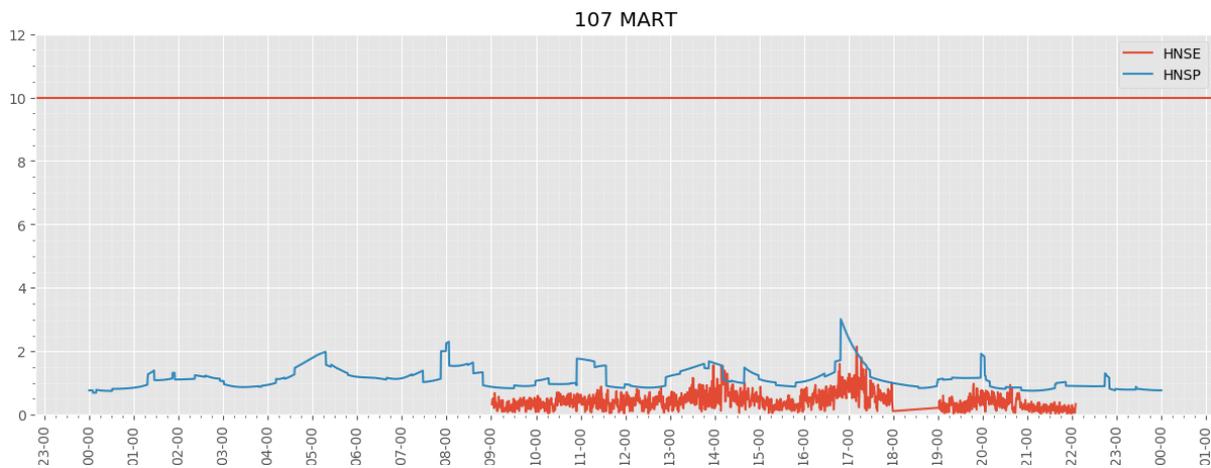


Figure 17: 95% accuracy plot for MART on 17/04/19.



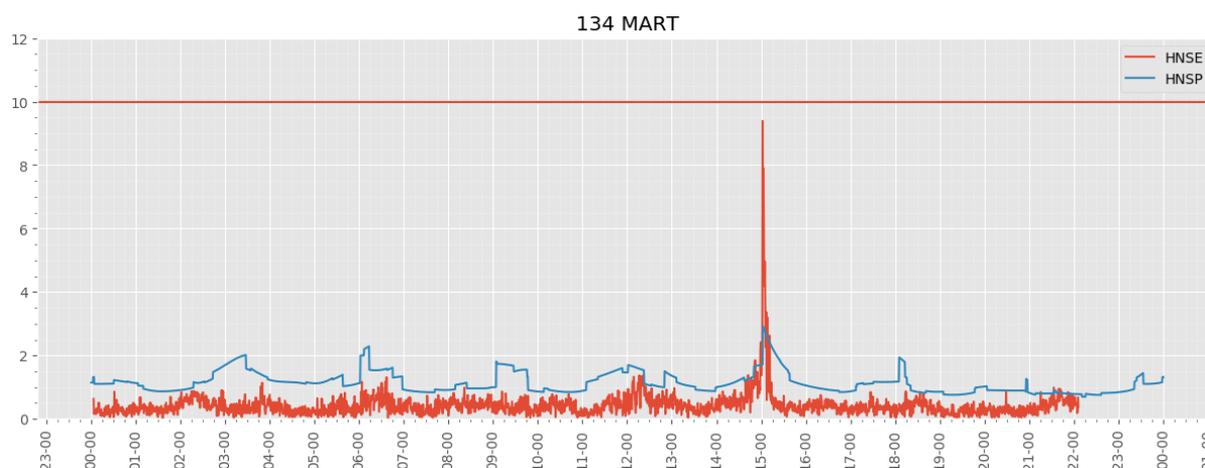


Figure 18: 95% accuracy plot for MART on 14/05/19.

The quality of the user position solution is a function of both the geometry of the satellites with respect to the user location (usually represented by Dilution of Precision, DOP), and the quality of the measurements (represented in simulation by the measurement standard error, σ). Poor positioning performance can result even with good measurements if the geometry is poor enough, or with good geometry if the measurement quality is poor.

So, the possible reasons for the predicted user error exceeding 10m are:

1. Poor geometry of visible satellites.
2. Poor geometry caused by the exclusion of some satellites due to ionosphere grid or RIMS tracking limitations (i.e. satellites that are ‘not monitored’, or ionospheric pierce points that pass through an area of the ionospheric grid that is ‘not monitored’).
3. Poor quality corrections which result in low measurement weight. With a sufficiently low weight the impact will be similar to excluding the satellite completely, and the effective DOP will suffer accordingly.

The following sections aim to explore these issues in more detail.

4.1 Basic Satellite Geometry

SSv2 has a tool to generate a grid of DOP values for all visible satellites. Figure 19 shows the instantaneous GDOP values at 03:45:00 on 14/05/19, plotted as a contour map between 80°N and 15°N, and 35°W and 45°E. The figure demonstrates that the GPS constellation is sufficiently dense that GDOP does not exceed ~3 anywhere within this coverage area at the given epoch (except in the far north), and for most of the area the GDOP value is below ~2.



As the satellites move in their orbits, their geometry with respect to users on the ground varies, so DOP varies continuously. When the worst case GDOP values over a 24-hour period (from 2880 30-second snapshots) are plotted as a contour map (Figure 20), the typical latitude banding of such plots becomes apparent. The band of higher worst-case DOP at high latitudes is clearly visible. The better performance at mid latitudes is caused by the convergence of the GPS satellite orbits at their maximum latitude of 55° together with visibility of some satellites that can be seen on the other side of the globe, beyond the north pole.

Figure 19 shows that the availability event observed at STOR is not caused by abnormally poor geometry in the visible satellites. Nevertheless, to achieve a positioning accuracy worse than 10m with a DOP value of, say, 2 would require average measurement errors of only 5m ($10\text{m}/2$), so even with such uniformly good DOP values it is still possible to achieve a poor position. Alternatively, the exclusion of one or more satellites from the visible constellation, due to 'not monitored' conditions (number of RIMS, or Iono grid points), could lead to worse DOP for the remaining satellites.

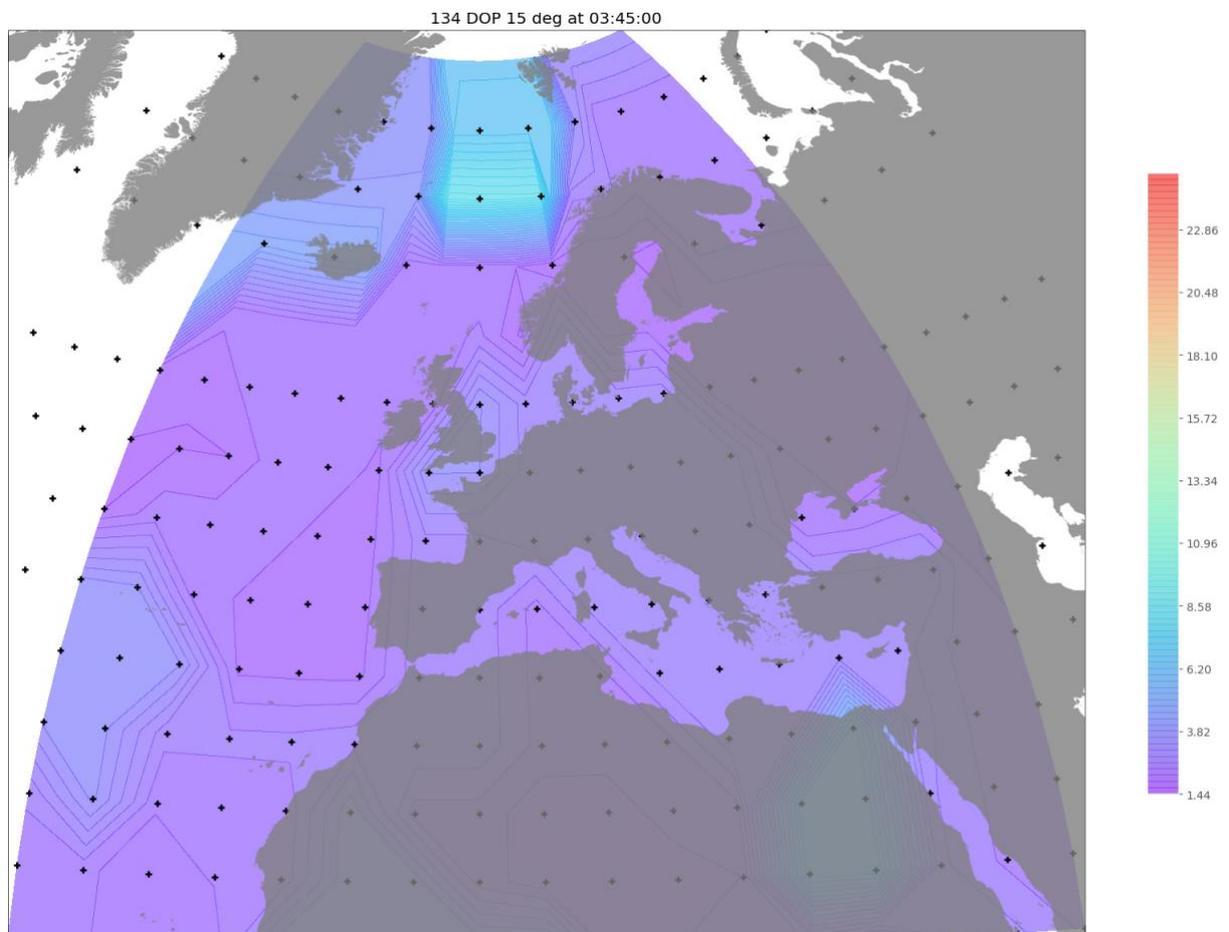


Figure 19: GDOP of all visible satellites above 15° elevation mask at 03:30:00 on 14/05/19.



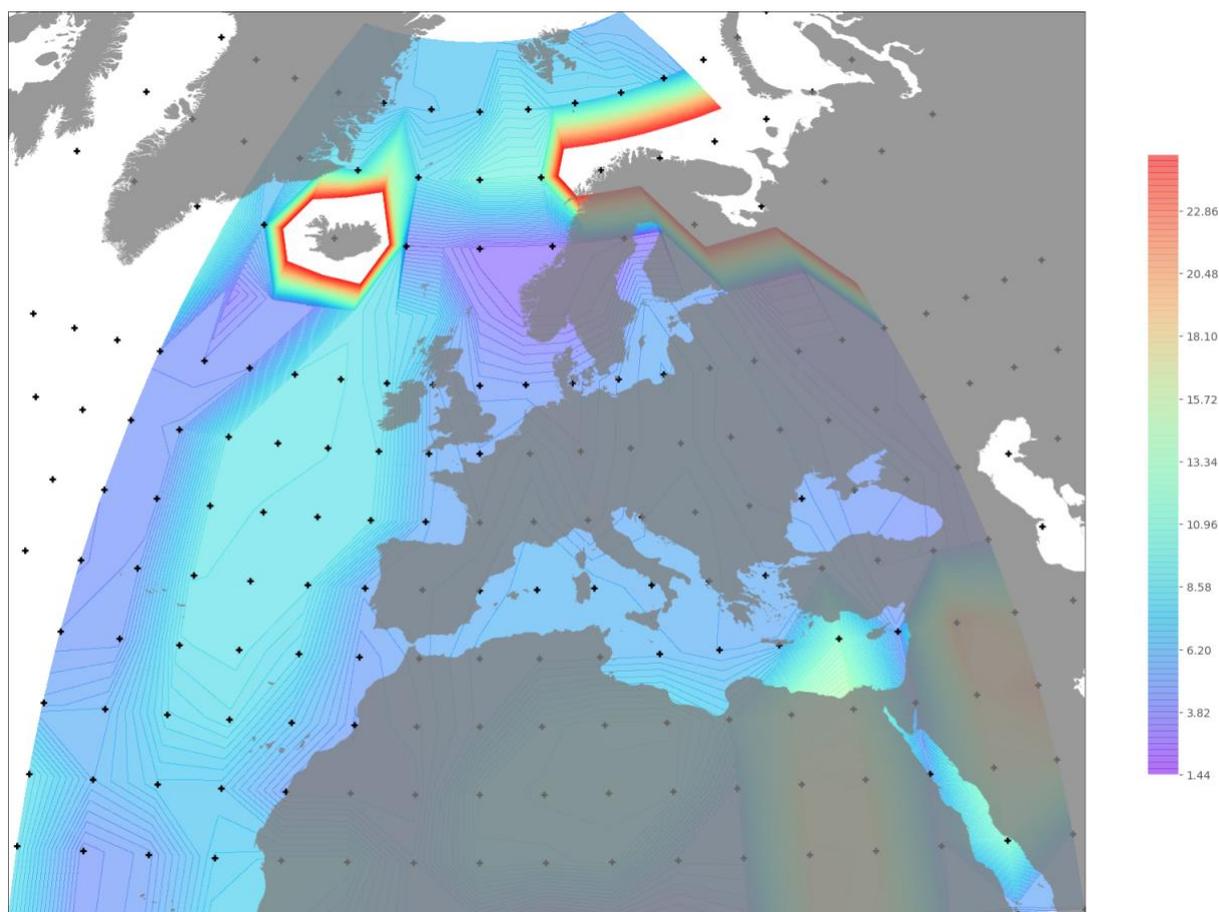


Figure 20: Worst case GDOP of all visible satellites above 15° elevation mask throughout the day on 14/05/19.

4.2 Satellite Exclusion

4.2.1 Exclusion due to inadequate RIMS tracking

EGNOS may determine that some satellites within the user's constellation are 'not monitored' and therefore should not be used. There are two principal ways that this condition may arise.

1. EGNOS requires every satellite to be monitored by a certain minimum number of RIMS in order to ensure that the satellite's status is reliably monitored (i.e. so that errors in the measurements from a single RIMS can be excluded) and that the ephemeris and clock corrections can be accurately determined.
2. If the satellite measurement passes through a part of the EGNOS ionospheric grid that is not well determined by the RIMS network, a reliable ionospheric correction will not be possible, and the measurement may be marked as 'not monitored'.



Concerning the first point, Figure 5 shows that SSv2 requires the satellite to be visible to a minimum of 3 RIMS before it will assign a numerical value to σ_{UDRE} – any fewer and the satellite is labelled as ‘NM’ (not monitored).

However, given the extent of the RIMS network, satellites visible to a user above a typical elevation mask (e.g. 10 to 15 degrees) will rarely be seen by fewer than 3 RIMS. This is illustrated by Figure 7 and Figure 6. Figure 7 shows a map of global tracking depth, and it can be seen that the contour for a depth of 3 RIMS covers a significant portion of the globe, encompassing the vast majority of locations where satellites would be visible to an EGNOS user. Figure 6 confirms this and shows that for the location selected for that figure (65°N, 12°W) no satellites above a 5 degree elevation mask are seen by fewer than 5 RIMS.

Figure 11 confirms that it is not possible for a satellite visible from STOR above a 10-degree elevation mask to be tracked by fewer than 3 RIMS, so this cannot be the cause of the observed availability events.

4.2.2 Exclusion due to poorly monitored Ionosphere Grid

The RIMS network is used to monitor the state of the ionosphere and determine values of Total Electron Count (TEC) (with associated error estimates) at fixed grid points (Ionospheric Grid Points, IGPs). A single frequency user is required to determine an ionospheric correction (and an estimate of the error in that correction) by interpolating from those grid points to the coordinates of each Ionospheric Pierce Point (IPP), i.e. the points at which the lines of sight between the user and the visible satellite pass through the altitude which corresponds to the notional altitude of the ionosphere. The ability of the RIMS network to determine the corrections and error estimates at the grid points is determined by the distribution of the points at which the ionosphere is ‘sampled’ by the measurements made by each RIMS.

These sampling points are the IPPs of the observations from the RIMS network. Each observation from a RIMS to a satellite generates one IPP, which is used to give an observation of the TEC at that IPP. Each RIMS observes multiple satellites at different azimuths and elevations, so the IPPs are distributed around the RIMS location. For the purposes of the ionosphere grid the ionosphere is considered to be located at an altitude of 350km. From simple geometry, a satellite with an elevation angle of 45 degrees from a given RIMS will generate an IPP that is 350km from that RIMS location. Higher elevation satellites will have IPPs closer to the RIMS location, and lower elevation satellites will generate IPPs further away. At 10 degrees elevation the IPP will be over 1200km from the RIMS location. So, a dense network of RIMS locations, each generating a scatter of IPPs, will produce a dense scatter of IPPs throughout the coverage area. EGNOS then uses an algorithm to combine these IPPs into values at the fixed grid points. For a given grid point to have a reliable TEC value and a reliable estimate of the error in that TEC value, it is required that the grid point should be determined from at least a minimum number of nearby IPPs.

The definition of that ‘minimum number’ and ‘nearby’, as well as the algorithm for combining the IPPs, are part of the EGNOS ground segment function. SSv2 includes models which aim to mimic the ground segment algorithms in this respect. For example, Figure 21 below shows the values assigned to the errors in the ionospheric grid points as a function of the number of



RIMS IPPs that contribute to each grid point. The error is reduced as more IPPs contribute, but below 4 IPPs the grid point is considered ‘not monitored’ (NM), meaning that a value for the correction is not possible. A consequence of this is that any user measurements whose IPPs are close to such a ‘not monitored’ grid point may not be able to interpolate a correction value from the grid, so will have to be excluded from the user’s position solution.

Geometric model - GIVE	
IPP	σ_{GIVE} [m]
0	NM
1	NM
2	NM
3	NM
4	0.91
5	0.82
6	0.73
7	0.73
8	0.64
≥ 9	0.54

NM = not monitored

Figure 21: SSv2 Ionospheric Grid error settings

In fact, a single ‘not monitored’ grid point is not normally sufficient to lead to the exclusion of the user’s observation, because several strategies exist for interpolating between other nearby grid points should a single nearby grid point be ‘not monitored’. The strategies vary depending on the latitude of the user’s IPP, because the (constant) longitude separation of the EGNOS grid points results in decreasing distance with increasing latitude. From (gAGE 2002), the strategies can be summarised as follows:

There are four distinct grid regions:

- First look for surrounding square cell
- Else seek surrounding triangular cell
- If neither available for $5^\circ \times 5^\circ$ look at $10^\circ \times 10^\circ$ (square cell then triangular cell)
- From 75°N to 85°N interpolate using virtual IGPs
- No correction possible if not surrounded

This is illustrated in Figure 22.



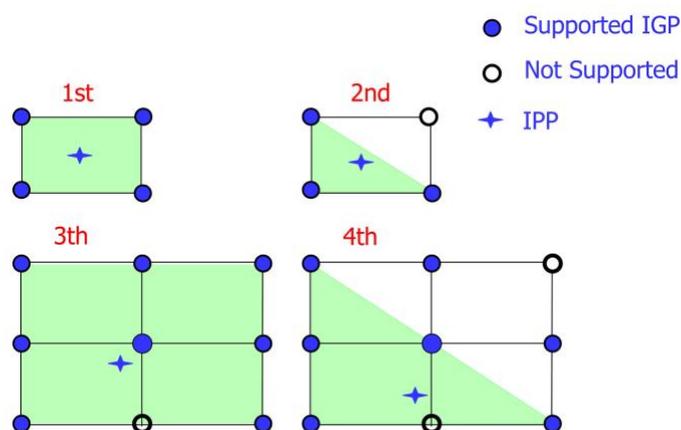


Figure 22: Illustration of IPP interpolation algorithm under various ‘not monitored/not supported’ conditions. (Ref: EGNOS Tutorial, gAGE, Universitat Politècnica de Catalunya)

The above figure illustrates that one or two ‘not monitored’ grid points in the vicinity of the user’s IPP will not necessarily result in the satellite measurement being excluded, since there is usually a way to interpolate between other nearby grid points. However, if the IPP is not surrounded by monitored grid points at either 5° or 10° spacings then no interpolation is possible, and the measurement must be excluded.

Using IPPs computed by SSv2, we have generated contour maps based on a weighted combination of the IPPs contributing to each Ionospheric Grid Point. SSv2 was used to output a value at each grid point based on an internal algorithm that closely matches the EGNOS IGP determination algorithm. It uses a weighted combination of the IPPs with three circles of increasing radius (320km, 400km and 480km) around each IGP and estimate the quality of those corrections by counting a total number of IPP around IGP within 600km radius. We then simply contoured these output values.

In addition, we have overlaid the monitoring status of the IGPs as determined from a simple count of the number of IPPs within a 480km radius of each IGP. Green dots indicate IGPs that are considered to be properly monitored (≥ 4 IPPs), while black crosses indicate ‘not monitored’ IGPs (≤ 3 IPPs). It is clear from each of these plots that although the simple count is only an approximation of the weighted combination algorithm, the contour information agrees closely with the monitoring status determined from the simple IPP count, i.e. the black crosses are always coincident with the red parts of the contour map, indicating few IPPs.

Note that the contour data is not available above latitude 65°N, due to a limitation with the output from SSv2.

In order to investigate whether unmonitored IGPs are the source of the spikes in performance at STOR, we have also overlaid these ionospheric grid point plots with the specific IPPs of the satellites visible from STOR at a few key epochs. These are shown as blue dots with SV numbers.



Firstly we show the plot for a nominal epoch where EGNOS performance is better than 10m.

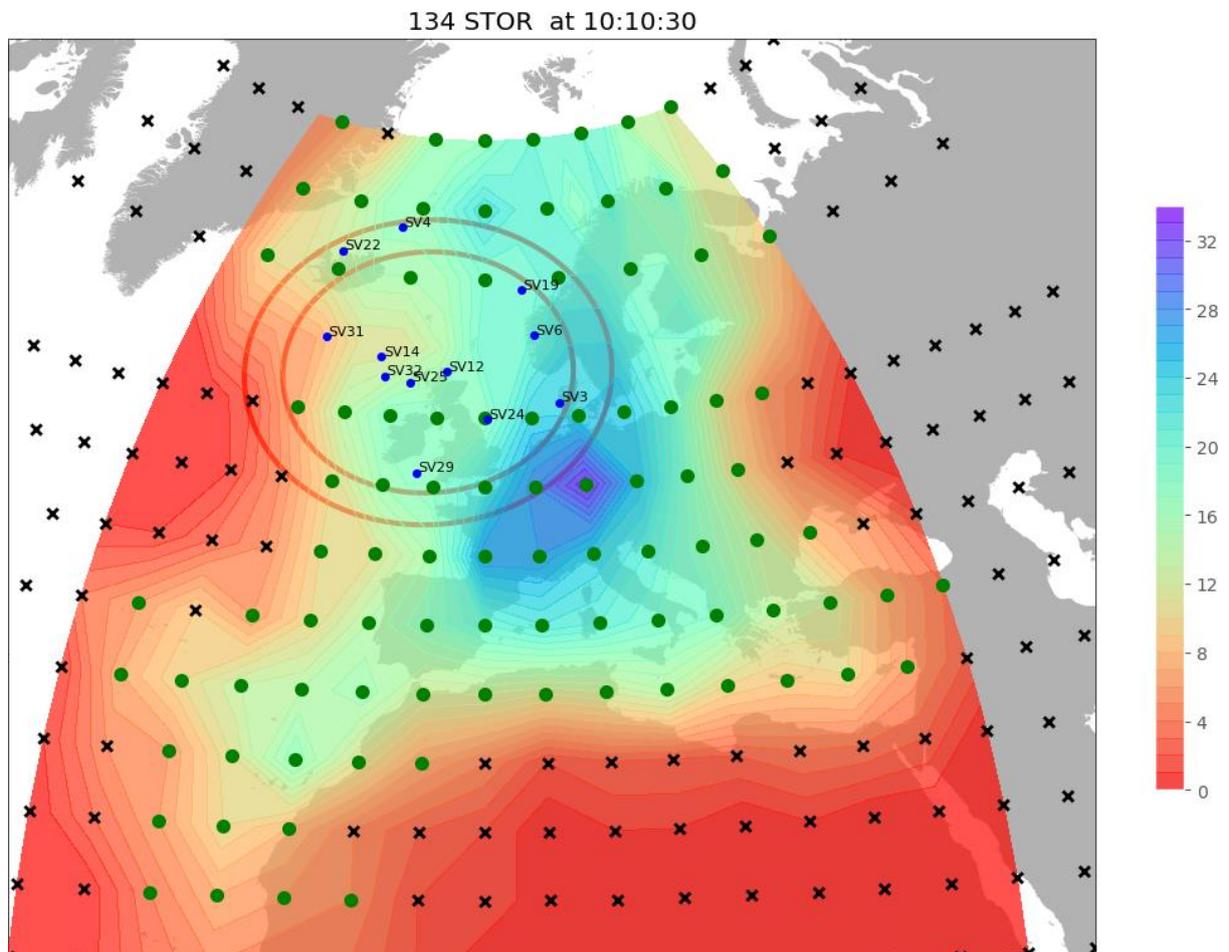


Figure 23: IGP and STOR IPPs for 10:10:30 on 14/05/19

From this plot it is clear that the IPPs of all satellites visible from STOR pass through parts of the EGNOS ionospheric grid that are properly monitored, i.e. every IPP is surrounded by green dots representing monitored IGPs, so an ionospheric correction can be interpolated for every satellite measurement.

Figure 24 shows the status at 03:45:00 on 14/05/19 and demonstrates a different situation. It shows that the IPP for SV15 is in a part of the EGNOS ionospheric grid that is not properly monitored. This IPP is not surrounded at either 5° or 10° by properly monitored IGPs, so no ionospheric correction can be interpolated for this observation. It is assumed that this measurement would therefore be excluded from the EGNOS positioning solution, and the remaining measurements will lead to a worse DOP. This figure is also overlaid with the visibility circles for both 15° and 10° elevation masks, showing the extent of the ionosphere (at an assumed 350km) that is visible above those elevations from STOR. It can be seen that



the lower elevation mask captures a couple of extra satellites, SV9 and SV28, compared to the higher mask angle.

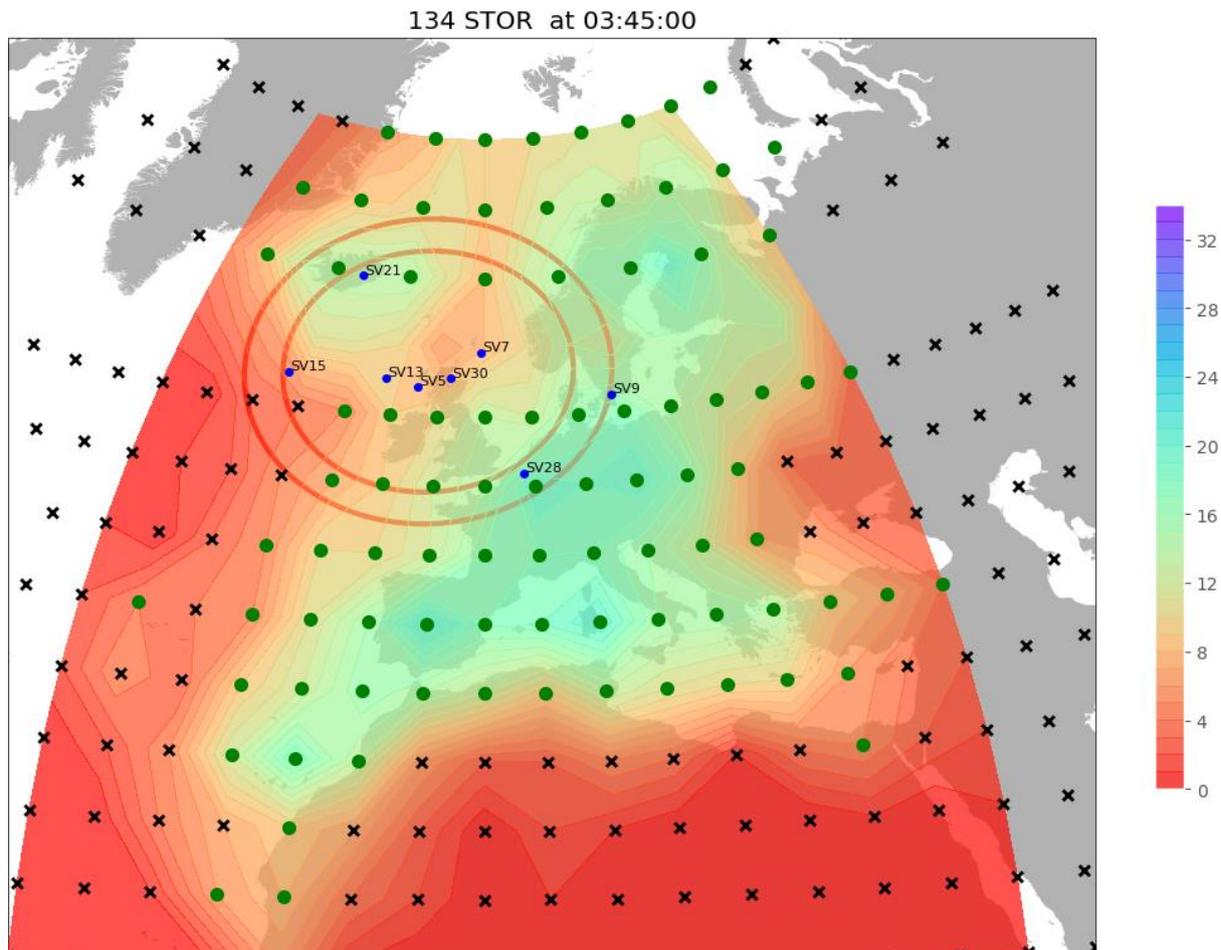


Figure 24: IGP and STOR IPPs for 03:45:00 on 14/05/19

The above plots illustrate one very specific scenario where a spike in the performance of EGNOS is due to inadequate monitoring of the ionosphere. Figure 25 illustrates the visibility circles around MART at the same epoch as Figure 24. It shows that the more southerly location of MART leads to IPPs that are not in the same unmonitored part of the ionospheric grid. This explains why the HNSP plots in Figure 16 to Figure 18 do not show the spikes in predicted accuracy that are seen in the STOR plots.



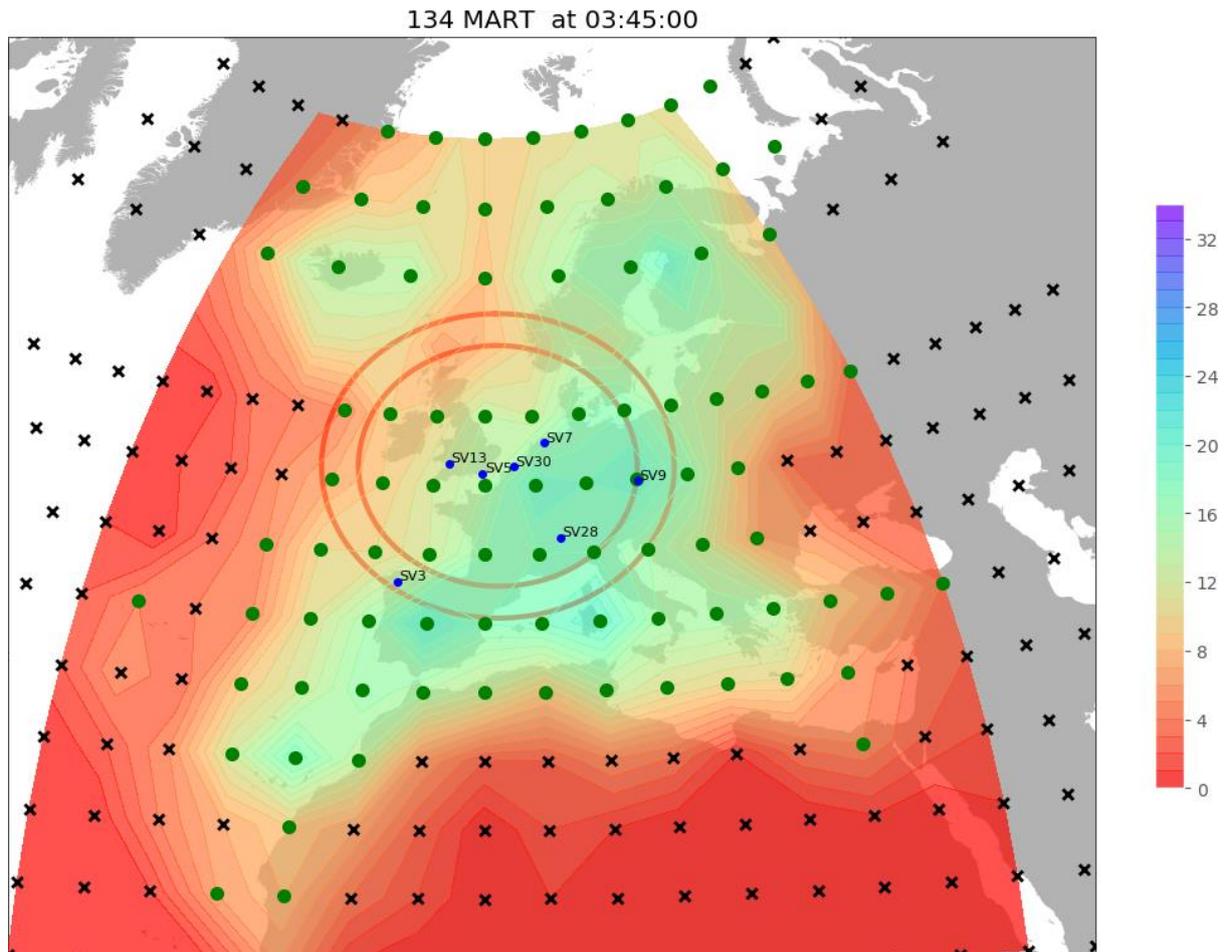


Figure 25: Ionosphere visibility from MART

While the analysis is intended to mimic the performance of the EGNOS system as closely as possible, it does not take into consideration historical EGNOS corrections. A preliminary analysis of EMS corrections for STOR shows it is in agreement with gLab data (presented in Figure 15 and reproduced below). For visualisation purposes the EMS accuracy prediction (shown added in purple in Figure 26) was scaled down to 1/3 of the original value, to approximately account for the overbounding nature of the EGNOS error estimates. This is not intended to be a conclusive demonstration, merely an illustration that the SSv2 models perform similarly to historic EMS data in this example. Further study in this area is suggested.



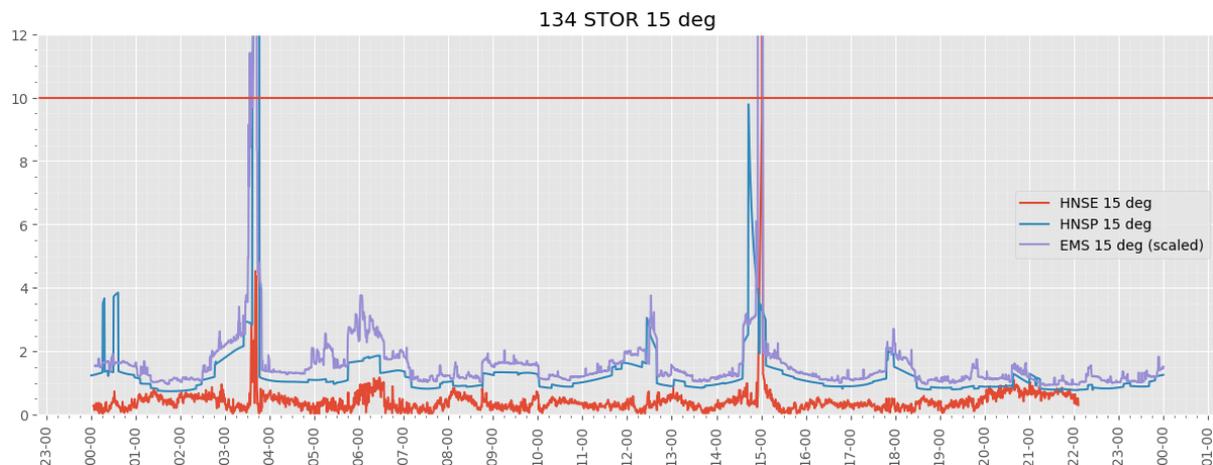


Figure 26: 95% accuracy plot for STOR on 14/05/19 with EMS based accuracy estimation.

4.2.3 Impact of Ionospheric Grid Monitoring further from RIMS network

The above examples illustrate the impact of reduced IGP monitoring towards the periphery of the EGNOS coverage area. The sites chosen (STOR and MART) are nevertheless land-based sites, and it is clear that it is possible for a vessel to be within UK or Irish waters and to be much further from the core of the EGNOS coverage area, in locations where the Ionospheric grid is substantially less well monitored. In the following example we consider the impact on a user at the approximate western limit of the Irish EEZ, at 55°N, 15°W.



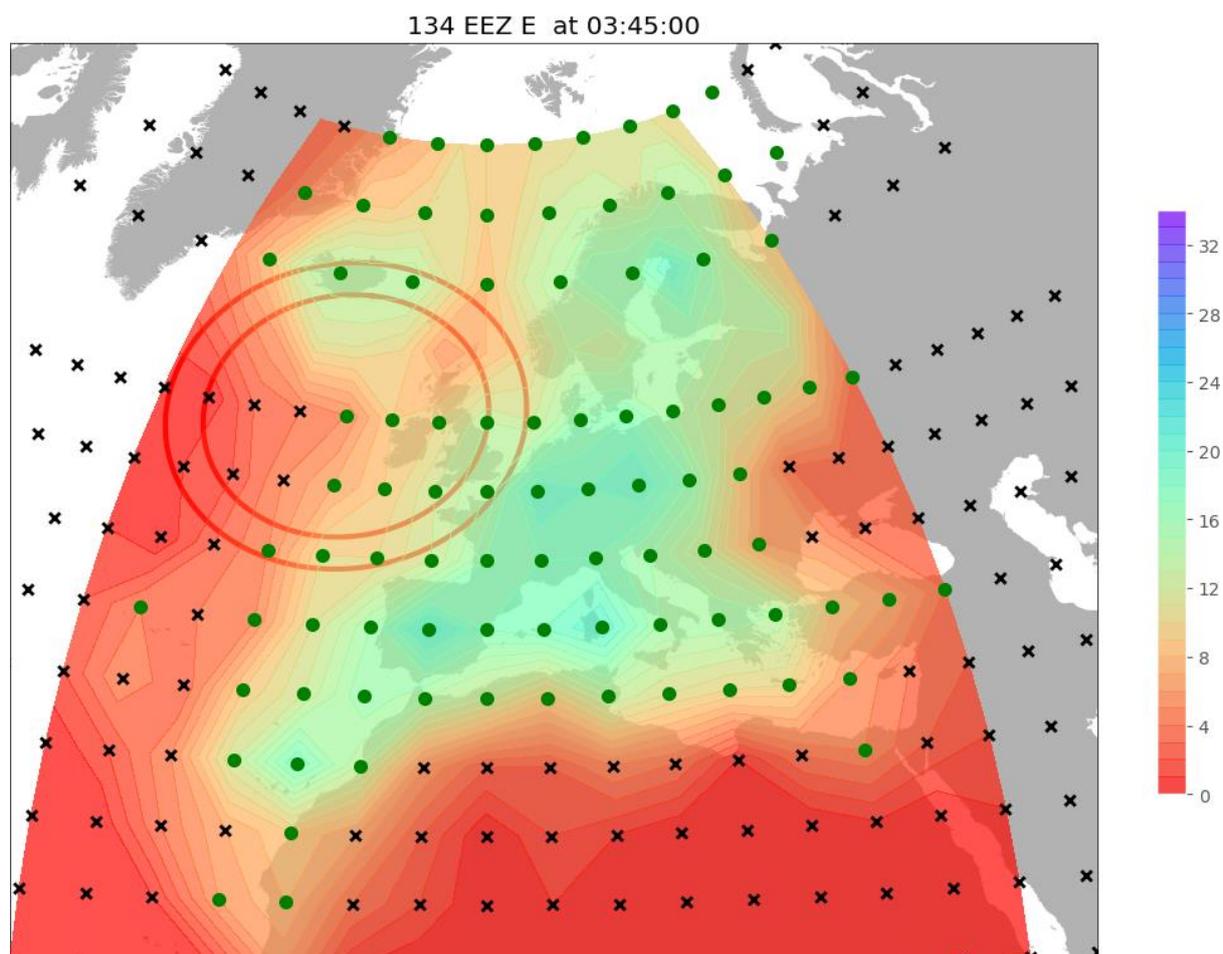


Figure 27: Ionosphere visibility from 55°N, 15°W

It is clear that a significant proportion of the ionospheric grid, where pierce points from this user could pass, is not monitored. If satellites in that part of the grid are excluded, it is likely that the geometry of the remaining satellites will be significantly poorer. To illustrate this, we show a plot of the predicted 95% accuracy for a user at this location.

Since we do not have real receiver measurement data for a user at this location, we show only the predicted 95% accuracy. The blue line shows the prediction based on a user elevation mask of 15° and demonstrates that there are numerous periods when the predicted 95% accuracy is worse than 10m. The likely explanation for this is the exclusion of satellites due to their pierce points being within a poorly monitored section of the ionospheric grid. To explore whether a lower elevation mask would improve the performance, we also show the predicted 95% accuracy with a mask angle of 10°. It is clear that performance is improved, and that there is a reduction in the number of ‘availability events’, but there are still several spikes where the predicted performance is worse than 10m, and numerous occasions when the performance is in the 6-8m range.



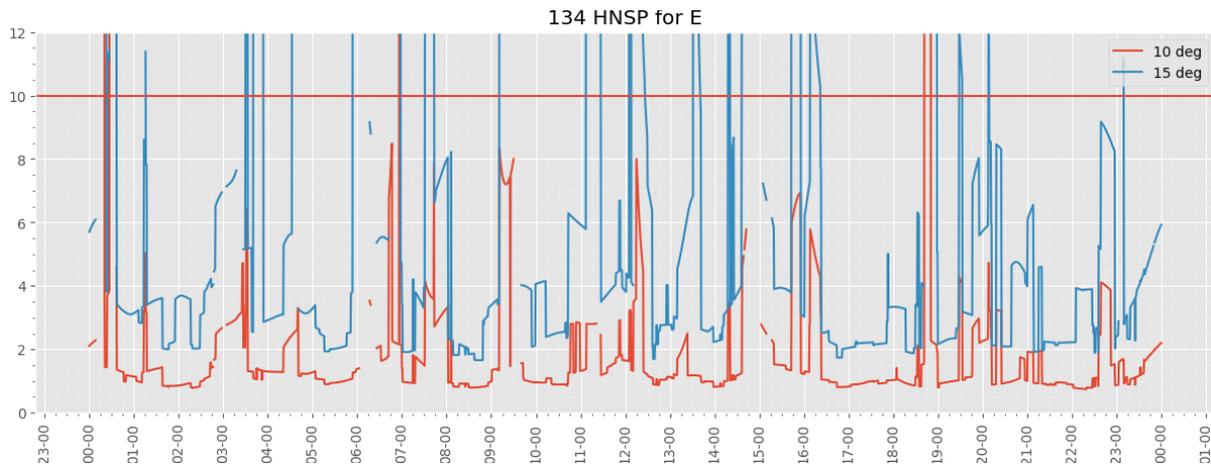


Figure 28: Predicted 95% accuracy for a user at 55°N, 15°W

4.3 EGNOS Correction Weighting

The aim of this section is to investigate the spike in EGNOS positioning performance observed at STOR at 15:00 on 14/05/19 (seen earlier in Figure 15, repeated here as Figure 29).

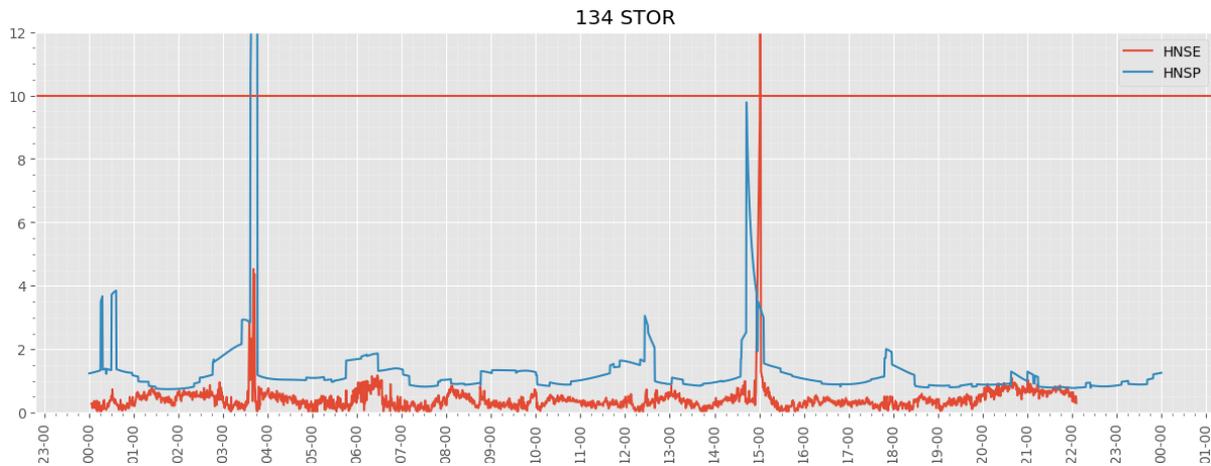


Figure 29: 95% accuracy plot for STOR on 14/05/19.

We first check whether the same spike is present in the uncorrected single point positioning (SPP) solution. Figure 30 below shows the EGNOS positioning results in more detail, while Figure 31 shows the equivalent single point positioning results. It is clear that the spike at 15:00 is not present in the SPP solution, which indicates that some feature of the EGNOS solution is responsible for the spike.

We next check whether the elevation mask angle has an impact on the spike. Figure 32 to Figure 35 show the EGNOS and SPP results for 10° and 5°. It is apparent that the spike is considerably diminished in the EGNOS solutions, and remains absent in the SPP solutions.



It should be noted that, contrary to the example under investigation here (15:00), other spikes are present in the SPP solutions, which are not present in the EGNOS solutions, indicating that the EGNOS corrections have successfully accounted for weaknesses in the SPP solutions, either by correcting measurement errors, excluding poor measurements, or giving them a low weight.

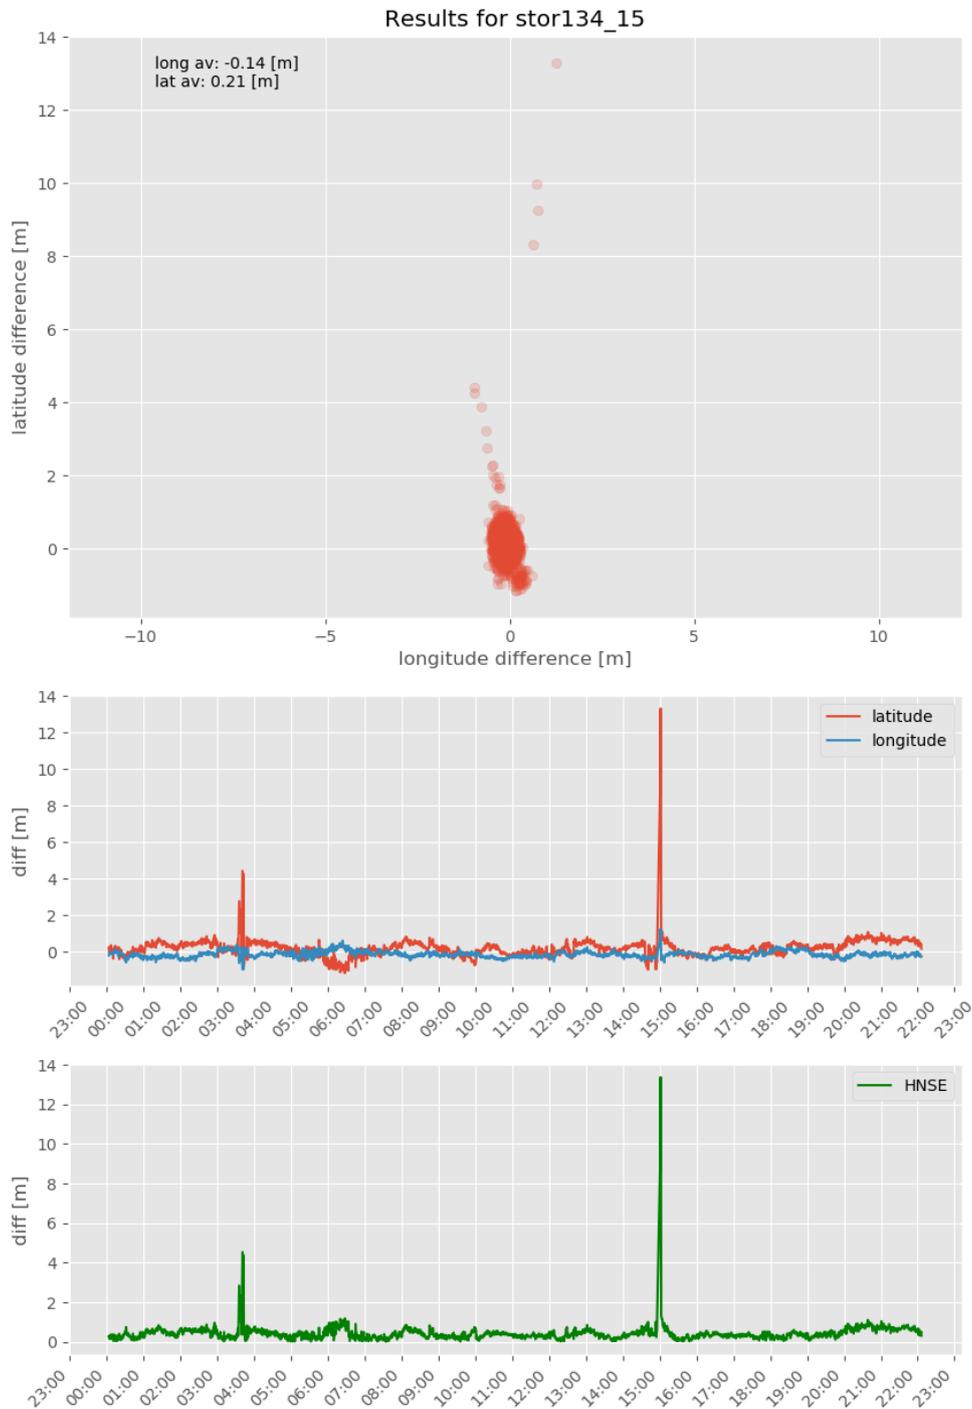


Figure 30: EGNOS Positioning Performance with 15° Elevation Mask



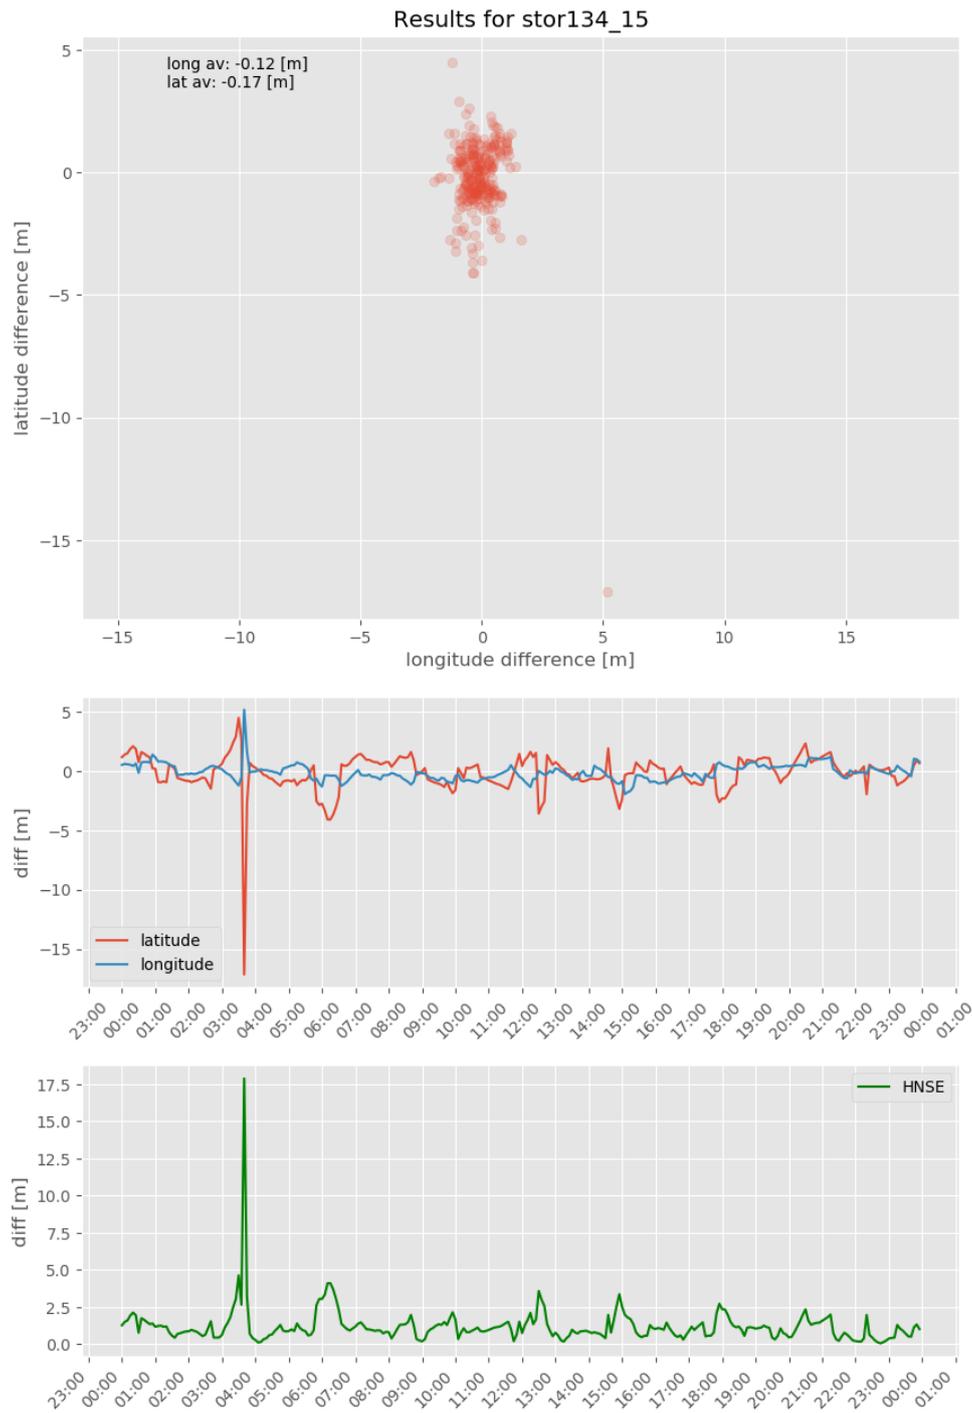


Figure 31: Single Point Positioning Performance with 15° Elevation Mask



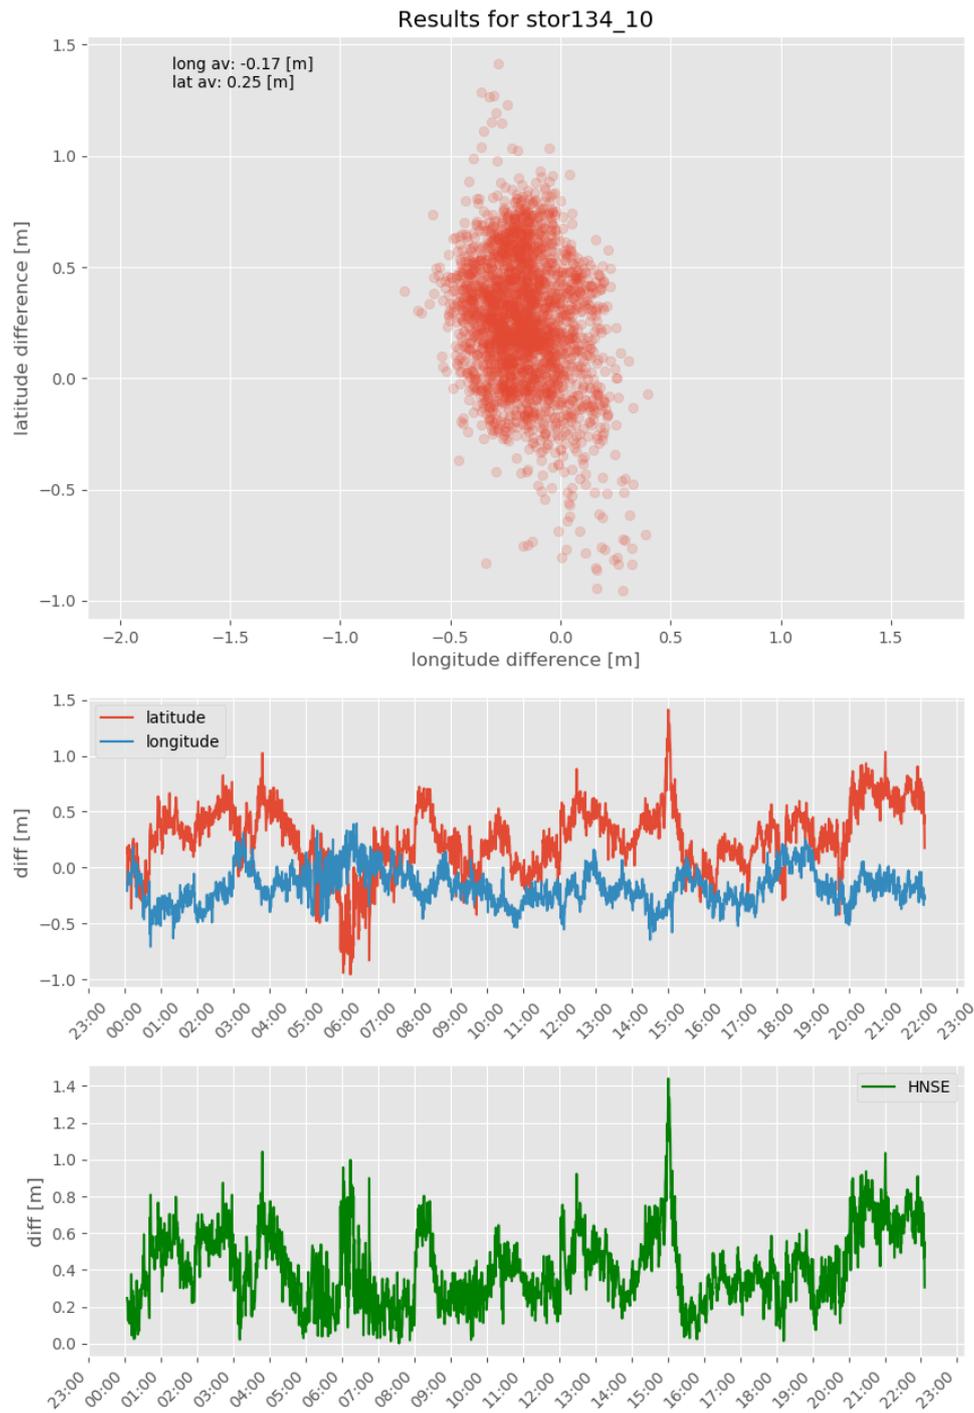


Figure 32 EGNOS Positioning Performance with 10° Elevation Mask



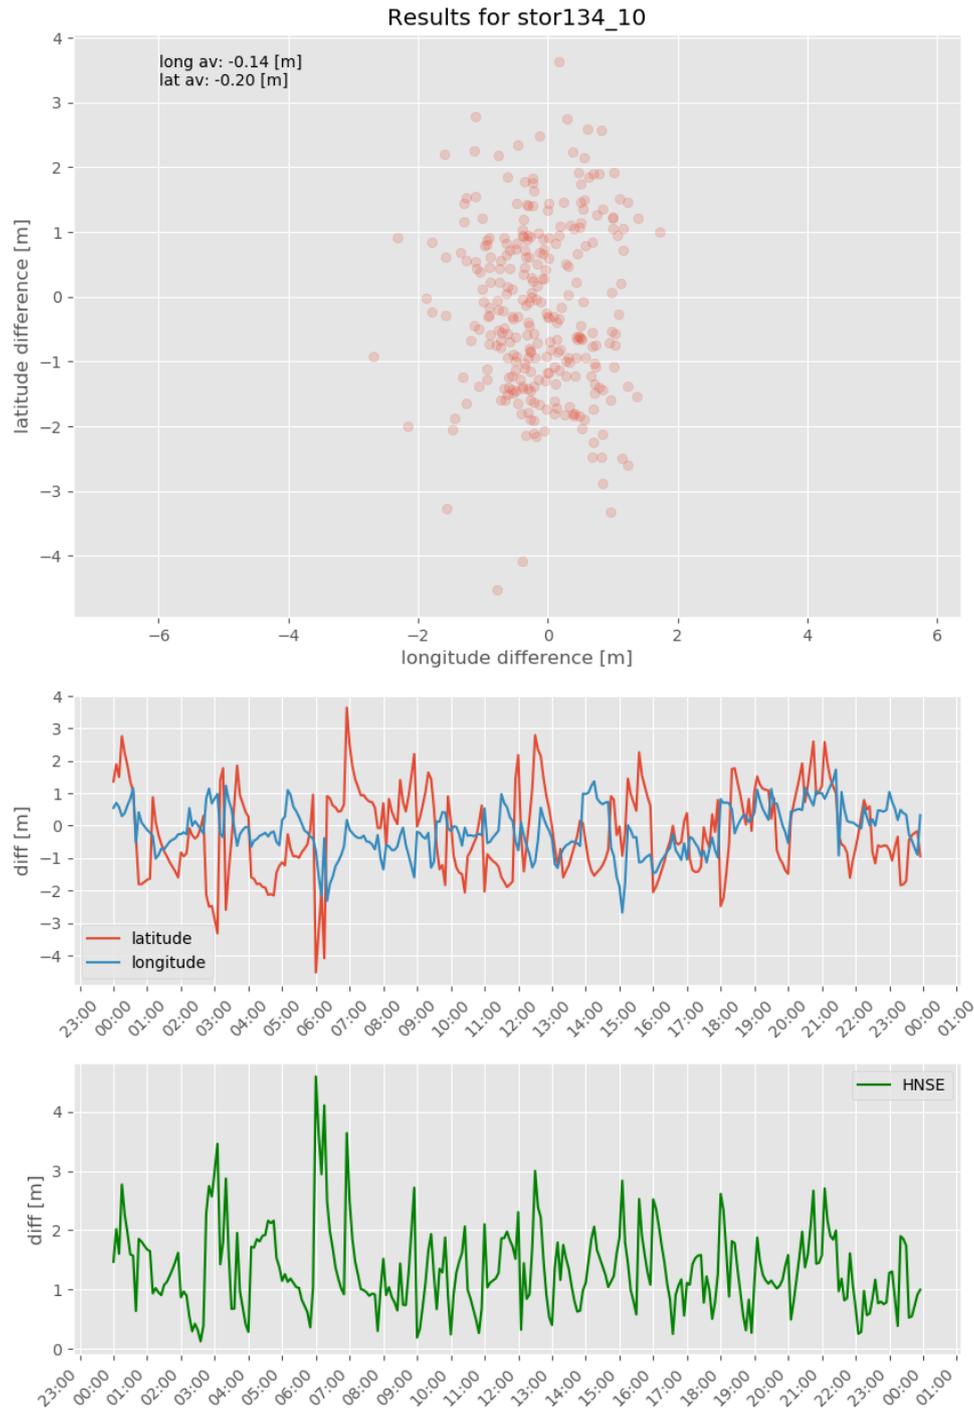


Figure 33 Single Point Positioning Performance with 10° Elevation Mask



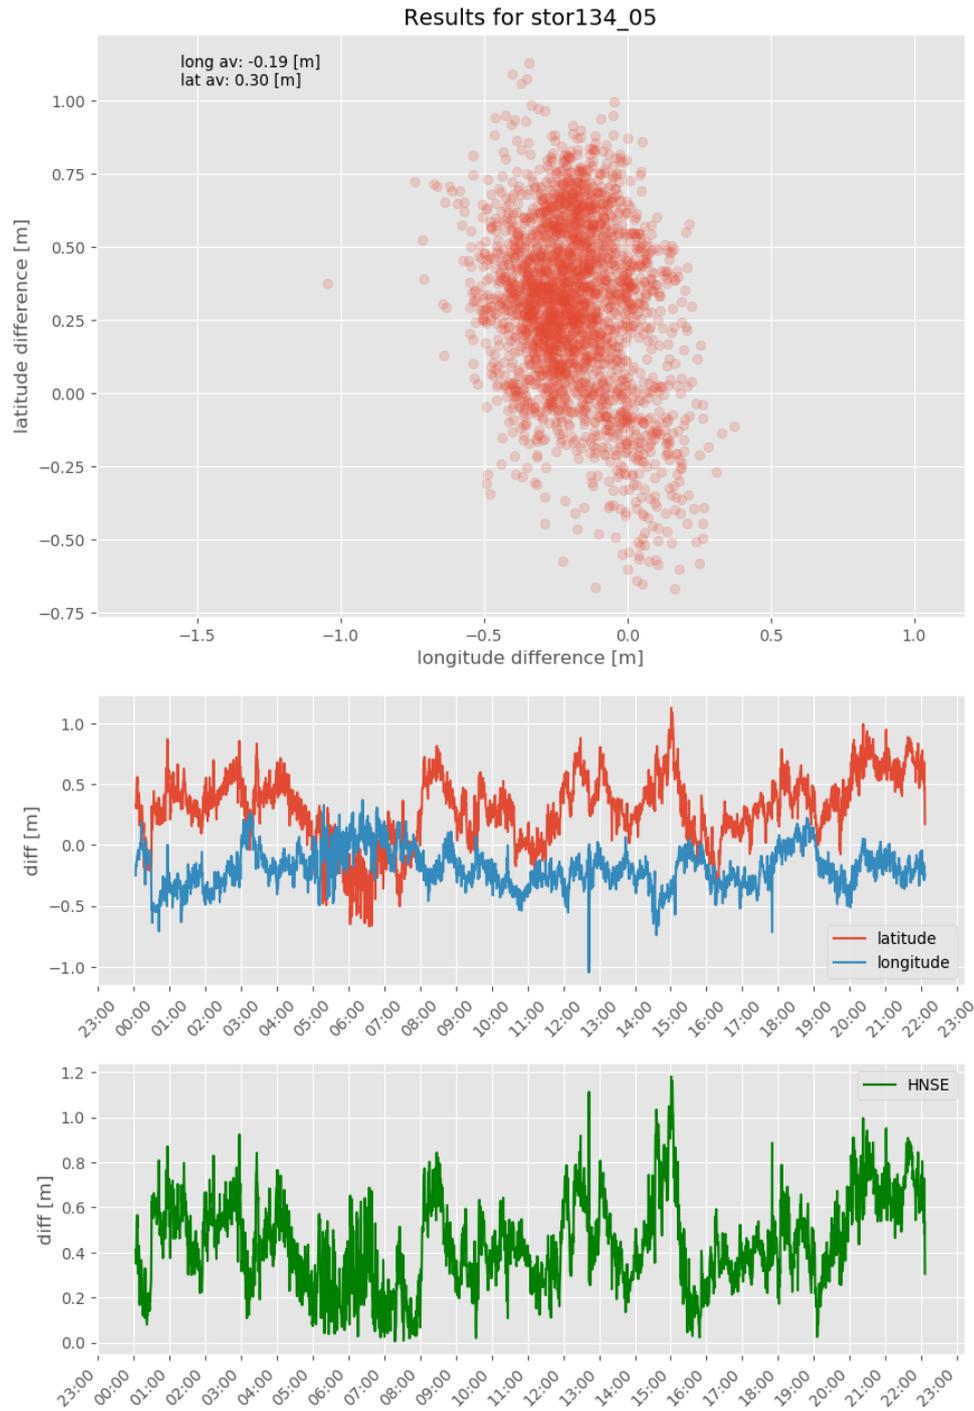


Figure 34: EGNOS Positioning Performance with 5° Elevation Mask





Figure 35: Single Point Positioning Performance with 5° Elevation Mask

The differences between the SPP and EGNOS solutions at the epoch under investigation suggest that the satellite geometry may be different between these two solutions, possibly due to the exclusion of one or more satellites by the EGNOS corrections. We next investigate the HDOP for the two types of solution, as output by gLAB. Figure 36 shows the EGNOS HDOP values for the three different elevation mask angles, and Figure 47 shows the same plot for the SPP solutions. It is clear that the EGNOS 15° solution exhibits a spike in HDOP at 15:00



which is not present in the SPP solution. The expectation therefore was that the EGNOS processing has excluded one or more satellites from the solution.

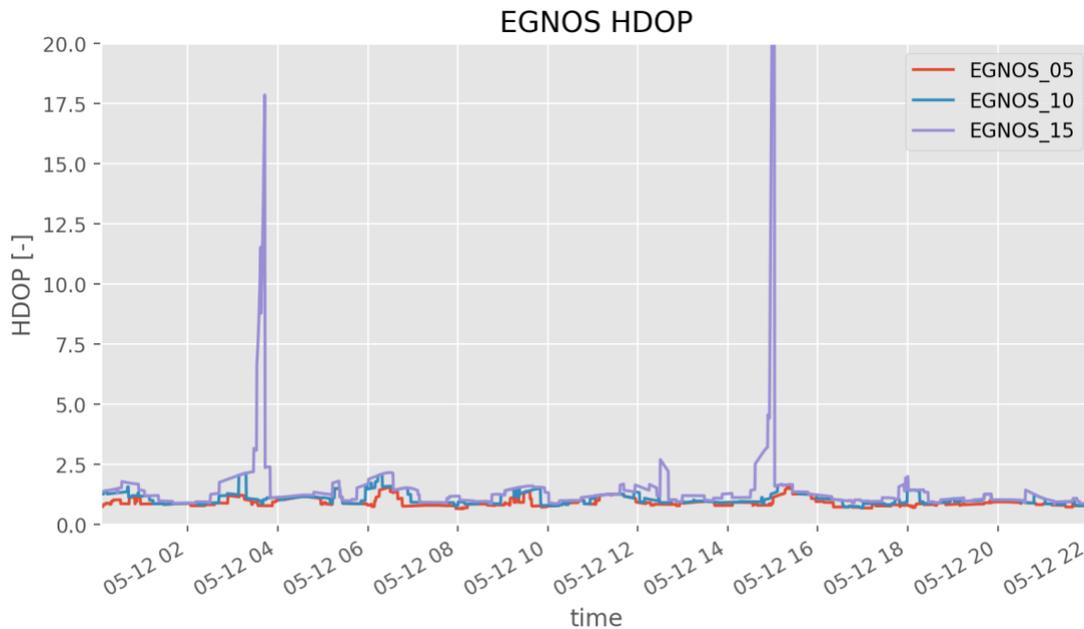


Figure 36: gLAB HDOP for EGNOS solutions at 5°, 10° and 15°

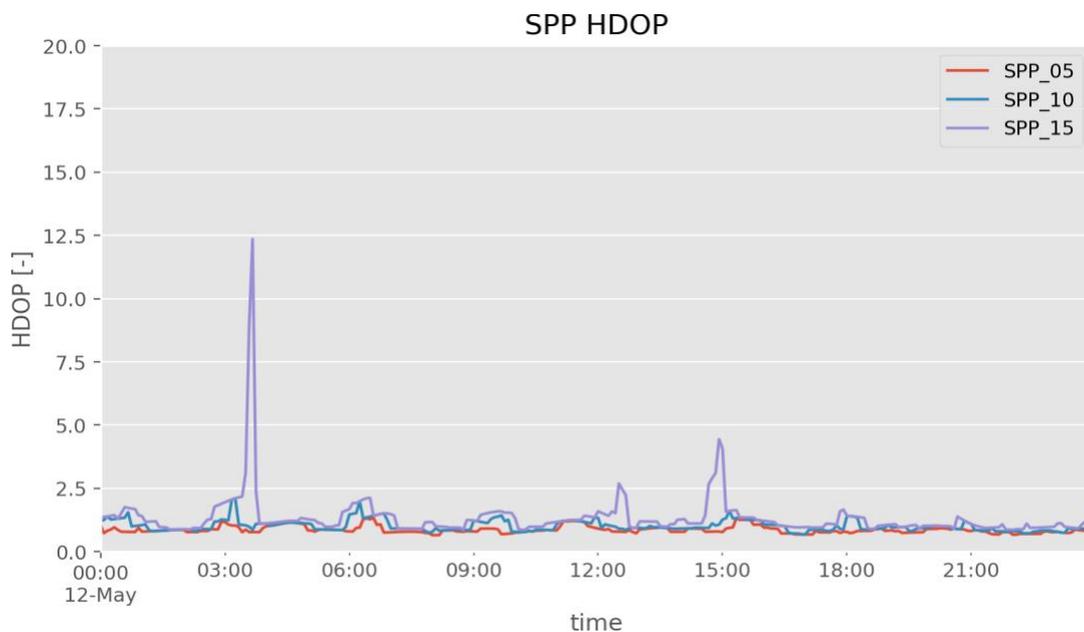


Figure 37: gLAB HDOP for Single Point Positioning solutions at 5°, 10° and 15°

Figure 38 and Figure 39 therefore show the number of satellites used in the gLAB solutions for EGNOS and SPP respectively. It is clear that the number of satellites used in the solution drops to 5 momentarily at 15:00 in both the EGNOS and SPP solutions, and since only the mask angle can be responsible for satellite exclusion in the SPP case, the implication is that the *same* 5 satellites are present in both solutions, and that they must have the same geometry.



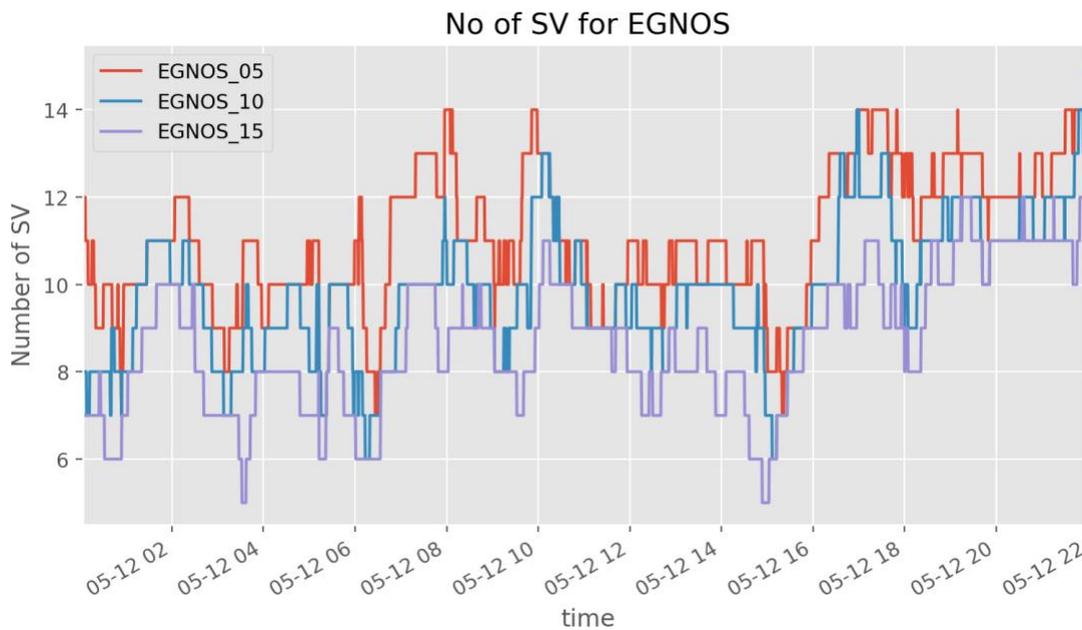


Figure 38: Number of used satellites in gLAB EGNOS solutions at 5°, 10° and 15°

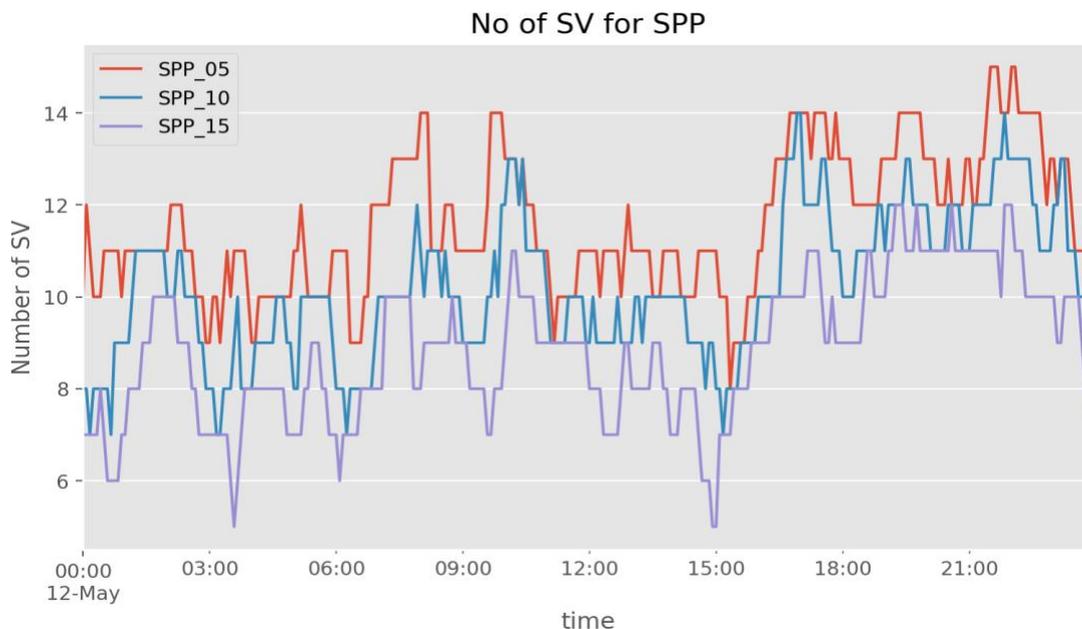


Figure 39: Number of used satellites in gLAB Single Point Positioning solutions at 5°, 10° and 15°

We therefore conclude that the HDOP values output by gLAB must be *weighted* HDOP, i.e. the effect of measurement weighting must have been included. It is likely that gLAB uses an elevation dependent weighting function in the SPP solutions. In the EGNOS solutions, however, measurement weights will be obtained from the EMS data files.

We can therefore conclude that the EMS messages at 15:00 have down-weighted one or more of the 5 satellites visible above the 15° elevation mask, resulting in poorer geometry, and it is this poorer geometry that is responsible for the spike in performance. Whilst the reason for



the down-weighting is not explored here, it is clear from the equivalent SPP results that the measurement(s) in question do not contain abnormally high measurement errors, which implies that the down-weighting has arisen due to EGNOS modelling (e.g. insufficient tracking depth, or, more likely, increased IGP uncertainties).

With the lower elevation mask angles, the minimum number of satellites visible at 15:00 increases from 5 to 7 or 8, so the impact of the down-weighted measurement(s) among the subset of 5 is diminished. Nevertheless, although the performance spike is diminished with the lower elevation mask angles, it is still present, and in fact seems to be correlated to increased positioning error for approximately one hour before this event and 30 minutes after. The measurement down-weighting therefore appears to be related to a satellite that is present in the solution for an extended duration.



5 Benefits of IALA Beacons

The foregoing discussion has demonstrated that there are occasions where EGNOS performance does not meet the 10m accuracy requirement for maritime navigation. Time has not permitted a comprehensive study of a long period of time to explore the frequency of these events or their impact on overall accuracy and availability performance, but the identification of at least one such repeating occurrence has allowed a study of the causes of this occurrence, which have led to some general conclusions. It has been shown that the lack of an ionospheric grid correction for some satellites, leading to a reduced constellation with poorer DOP, is the most likely explanation for the spike in performance that causes such an ‘availability event’. Some satellites visible from STOR have pierce points that are within a poorly monitored part of the EGNOS ionospheric grid.

This poor monitoring is a function of the instantaneous location of the satellites in the constellation, and the distribution of the RIMS sites, which together define the distribution of the IPPs that are used to determine the ionosphere grid.

Poor monitoring is not necessarily a function of the state of activity within the ionosphere. Higher or lower ionospheric activity will not affect the distribution of the IPPs, but will affect the uncertainty in the TEC values assigned to each grid point. Higher uncertainty will lead to a lower weight being assigned to measurements with pierce points in that vicinity, which in turn will lead to poorer ‘weighted’ DOP: a very high uncertainty would result in a very low weight, which can have a similar effect to removing the satellite altogether.

In either case, the ability of the RIMS network to determine good quality TEC values at each point in the ionosphere grid is key to obtaining good positioning performance.

The following analysis will attempt to determine whether IALA beacons can benefit in situations when the EGNOS ionospheric grid is not adequately monitored.

We will demonstrate that:

- Beacon-based DGPS performance during the period around the identified availability event remains better than 10m at all times, even at very long distances between the user and the beacon.
- Implementing additional RIMS locations (here considered at the locations close to the infrastructure of the existing beacons) would improve the monitoring of the EGNOS ionospheric grid to the point where no satellites would be excluded due to poor monitoring.



5.1 Beacon DGPS Performance

To assess the performance of DGPS at varying distances from the beacon, the measurement data from STOR, on 14th May 2019 has been processed by gLab in single difference mode (to simulate DGPS positioning), using a 15° elevation mask, against a number of CORS sites with increasing distances from STOR. The CORS sites used, together with their distances from STOR are as follows:

ULLO: Ullapool, 79km, OBAN: Oban, 203km, CAML: Campbelltown, 313km, KIRK: Kirkwall, 403km.



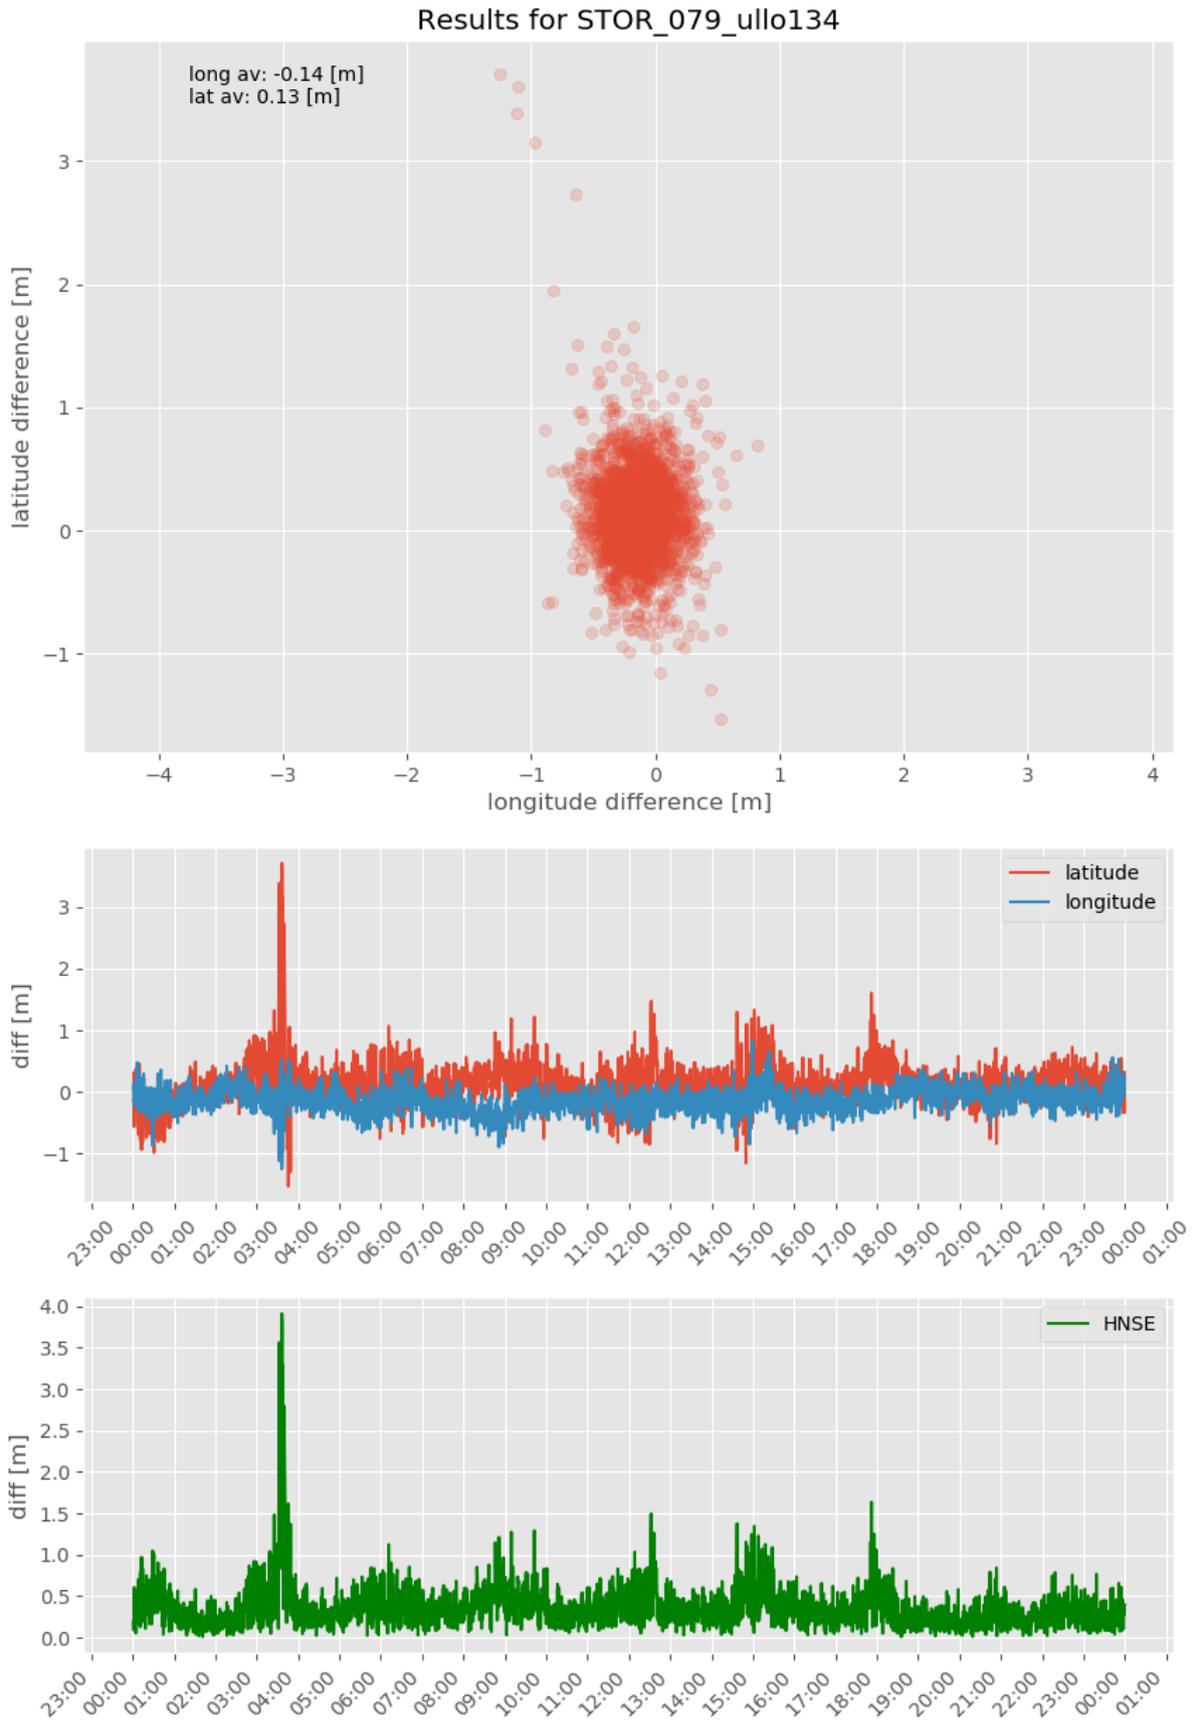


Figure 40: DGPS positioning results for STOR from ULLO on 14/05/19



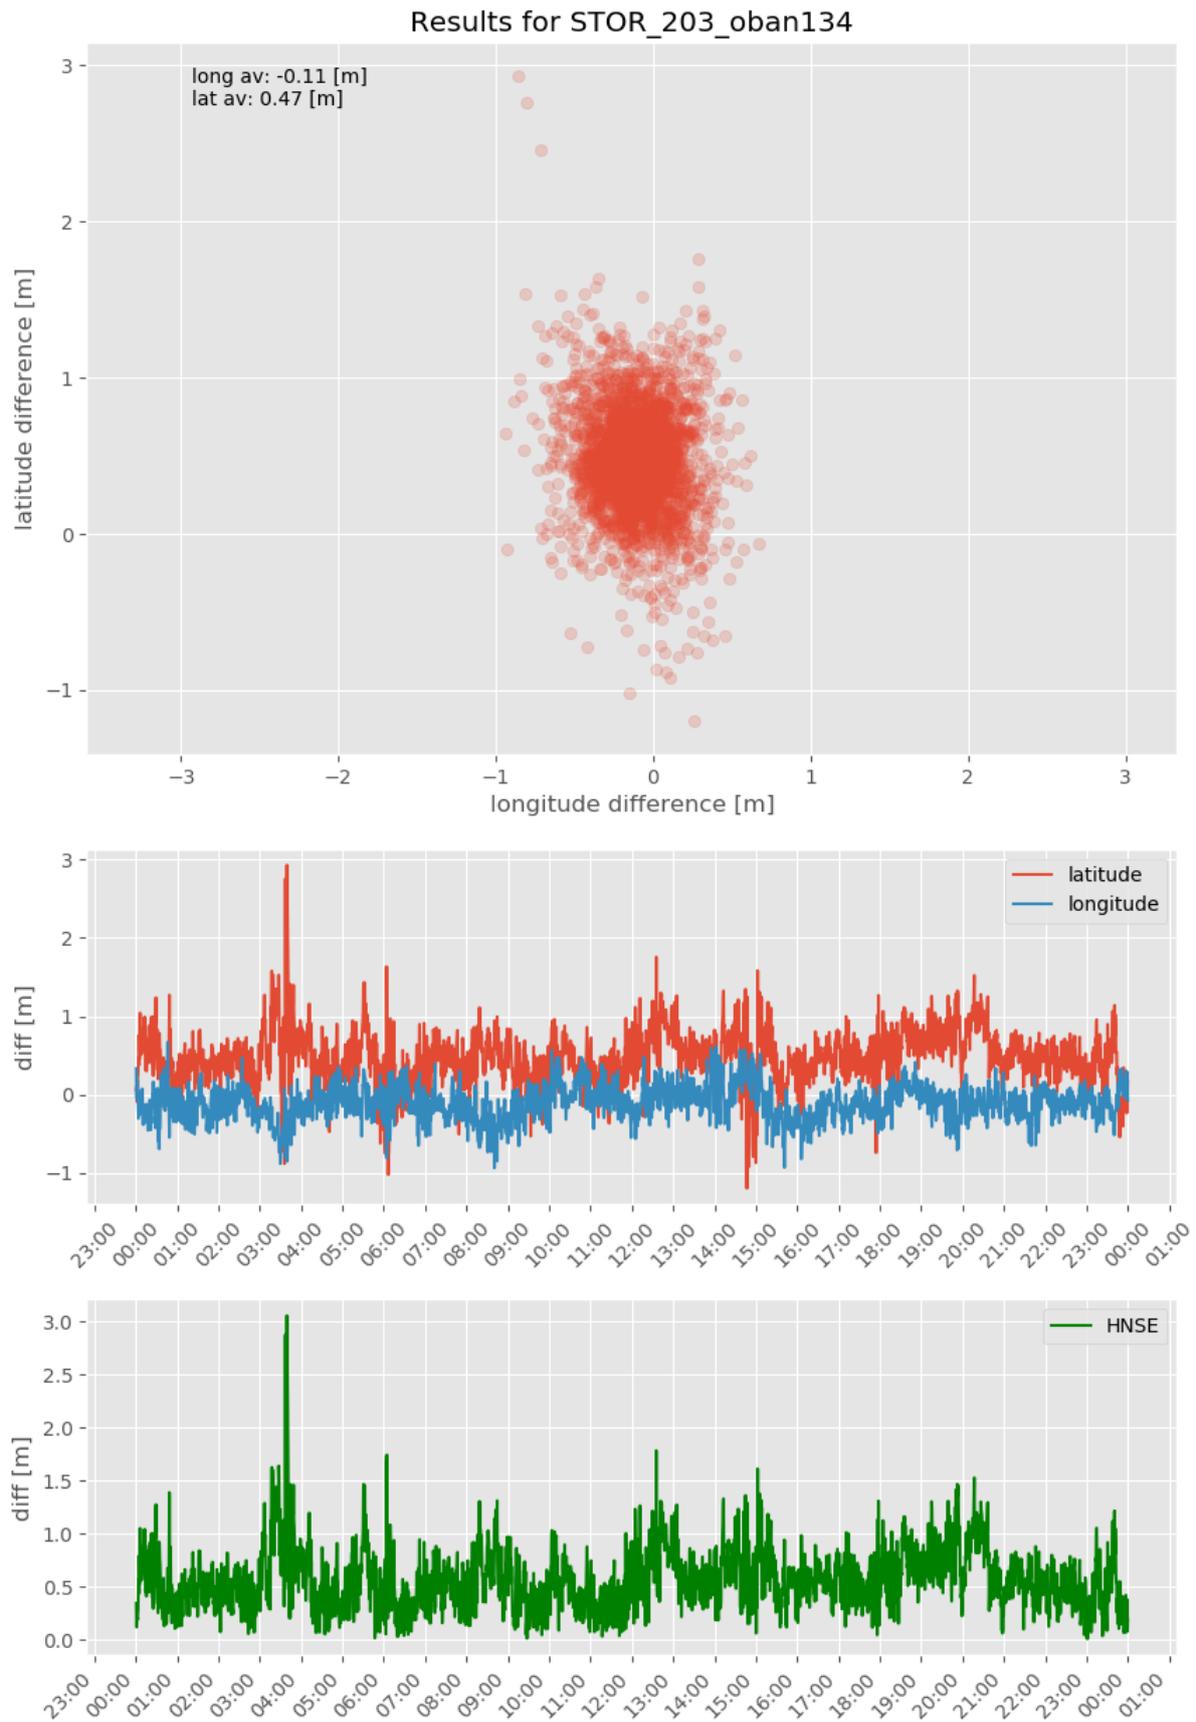


Figure 41: DGPS positioning results for STOR from OBAN on 14/05/19





Figure 42: DGPS positioning results for STOR from CAML on 14/05/19





Figure 43: DGPS positioning results for STOR from KIRK on 14/05/19



The above figures demonstrate that the spike in positioning performance at 03:45 is of a similar magnitude to the spike observed in the EGNOS positioning (~4-5m, Figure 15) using any CORS site up to approximately 300km. Using the closest sites (ULLO and OBAN) the DGPS positioning error is smaller than in the case of the EGNOS results, but the spike in the position error is still present. This suggests that there is a significant decorrelation between the errors in the STOR measurements and those from the reference (CORS) measurements. No analysis of these measurement errors has been possible in the time available, so the only conclusion possible at this stage is that close range DGPS has performed better than EGNOS during this event. A fuller analysis would investigate individual single difference measurements to attempt to determine which measurements are responsible for the spike in position error, and whether there is a correlation with, say, the relative IPP locations from STOR and the reference stations

5.2 Additional RIMS Locations

Since the ‘monitored/not monitored’ status of each IGP is based on a weighted count of the number of RIMS IPPs falling within a given radius of the IGP, we have explored whether additional RIMS locations would provide sufficient extra IPPs to upgrade the ‘not monitored’ IGPs to ‘monitored’ status. We first show, in Figure 45 and Figure 46 the impact, at 07:33 on day 075, of including the 5 additional EGNOS RIMS (at locations plotted in Figure 44):

- Mizen Head
- Loop Head
- Tory Island
- Butt of Lewis
- Sumburgh Head



Figure 44: Additional RIMS locations used to supplement existing RIMS network



Figure 45 shows the monitoring status of the EGNOS ionospheric grid using the default RIMS network, both in contour form and in the form of the individual IGP status (based on the 480km radius approximation). It is overlaid with both the IPPs from STOR at the specified epoch (03:45:00 on 17/5/19) and with the 10° and 15° ionosphere visibility circles around STOR. The figure demonstrates that a few of the IGP status have the status of 'not monitored' at this epoch, and further, as already illustrated, that satellite SV15 is not surrounded by 'monitored' IGP, so it would not be possible to interpolate valid TEC values and uncertainties for this satellite.

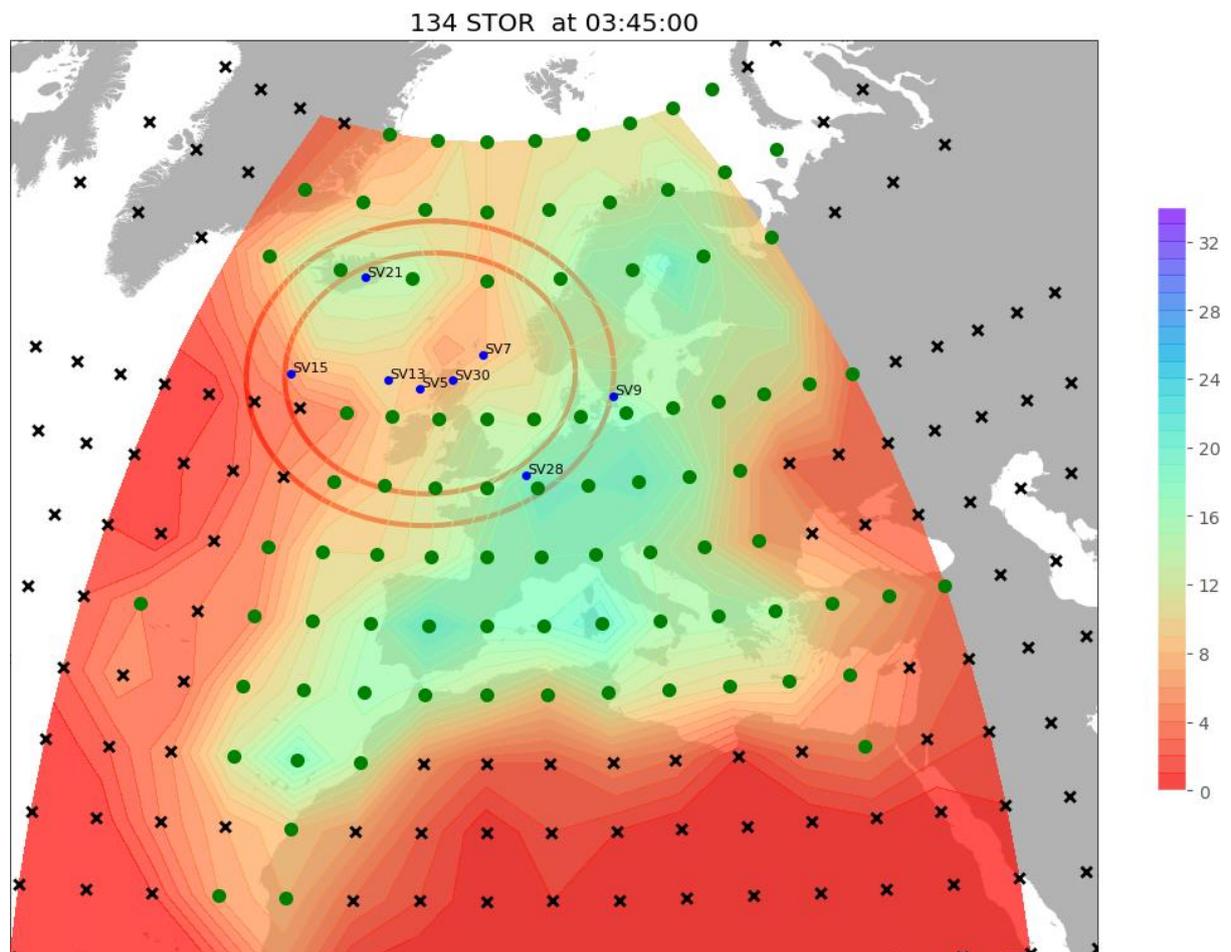


Figure 45: IGP status and STOR IPPs for 03:45:00 on 17/05/19

Figure 46 shows the same plot but generated by assuming that the five additional RIMS at locations listed above were included. It is clear that the area of the green part of the contour map is increased, particularly in the waters around the UK, and that several of the previously black 'not monitored' IGP status are now green ('monitored'). It demonstrates that the additional RIMS have generated additional IPPs in the previously poorly monitored section of the ionospheric grid, such that all the IGP status are now adequately monitored. As a result, satellite SV15, which would have been excluded previously, would now be able to determine a valid ionospheric correction and would be included in the position solution at this epoch.



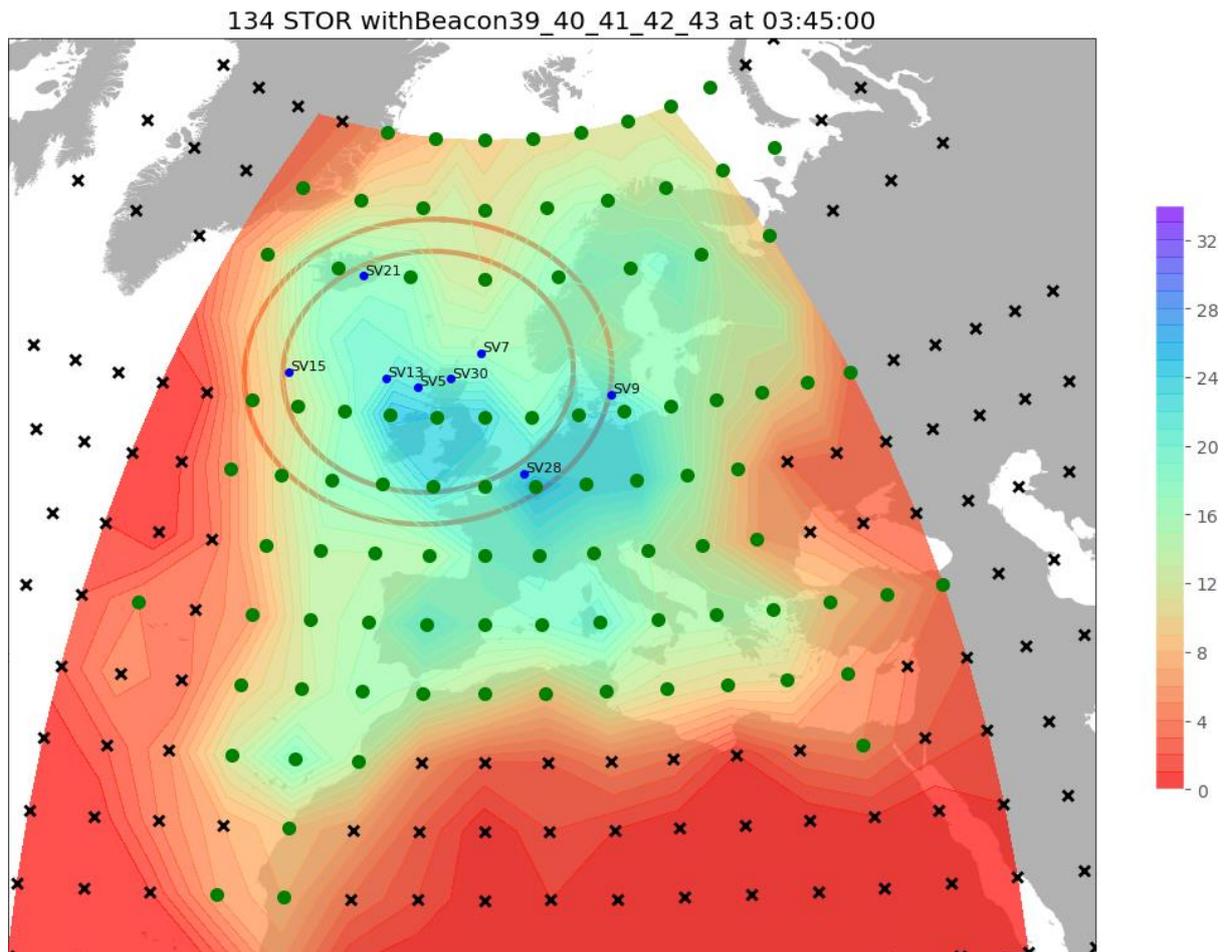


Figure 46: IGPs and STOR IPPs for 03:45:00 on 17/05/19, with 5 additional RIMS locations

The improvement in the monitoring of the ionospheric grid around STOR is due to the increase in the number of RIMS IPPs. The geometry of the additional RIMS locations is not critical in this respect. One additional RIMS location will generate one additional IPP per satellite. RIMS locations that are close together will generate IPPs that are also close together and will therefore boost the number of IPPs contributing to individual IGPs, so the improvement is a function of the density of the RIMS locations, rather than the extent of their geographic coverage. To illustrate this, we now show the impact of fewer than 5 additional RIMS locations. We demonstrate this by incrementally adding the additional RIMS one by one.

The following figures demonstrate that the major deficiency with the ionospheric grid monitoring requires just three additional RIMS locations. With three additional RIMS, the ionospheric grid is sufficiently monitored to mean that no satellites need to be excluded from STOR due to missing ionospheric corrections at the chosen epoch.



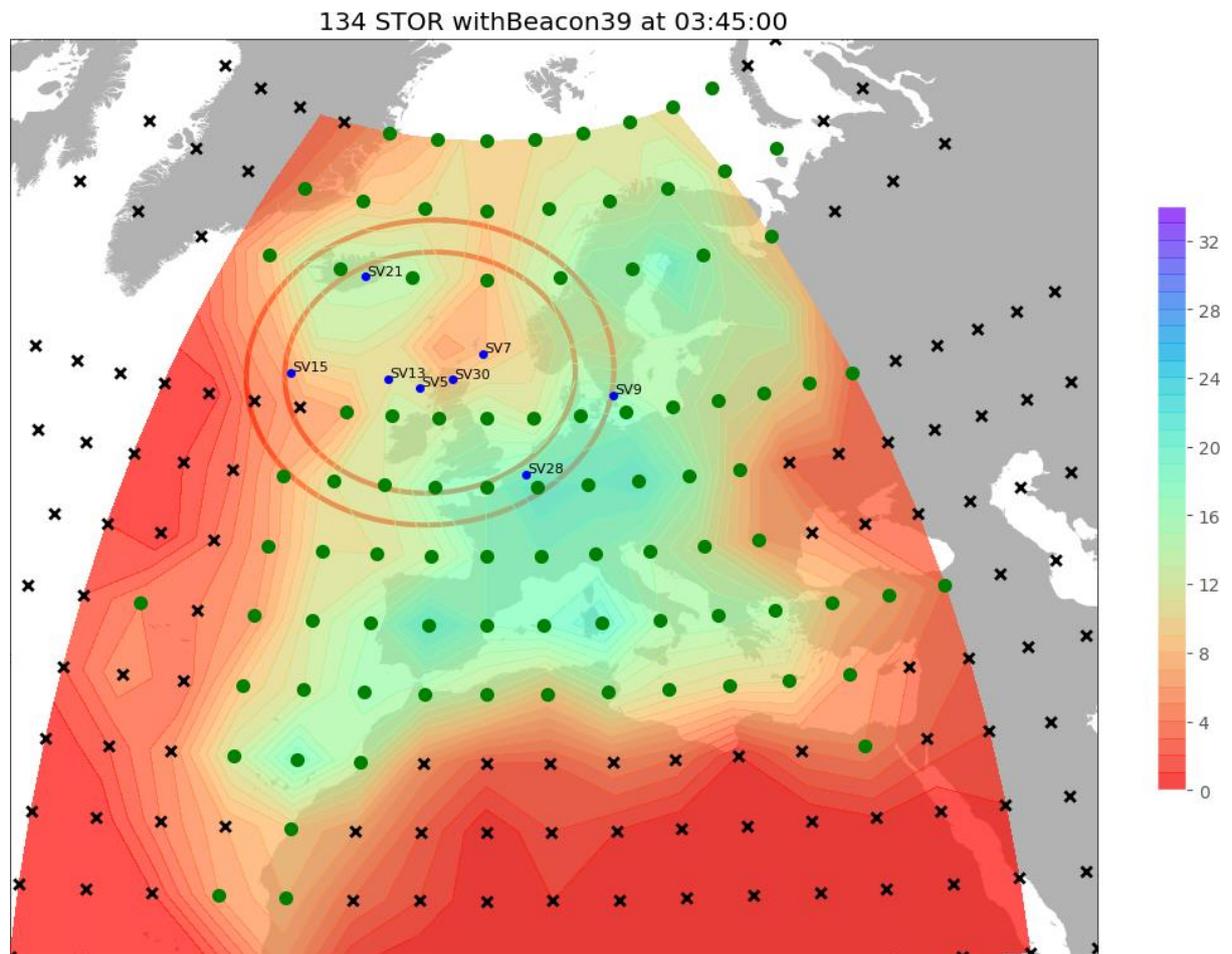


Figure 47: IGPs and STOR IPPs for 03:45:00 on 17/05/19, with one additional RIMS location



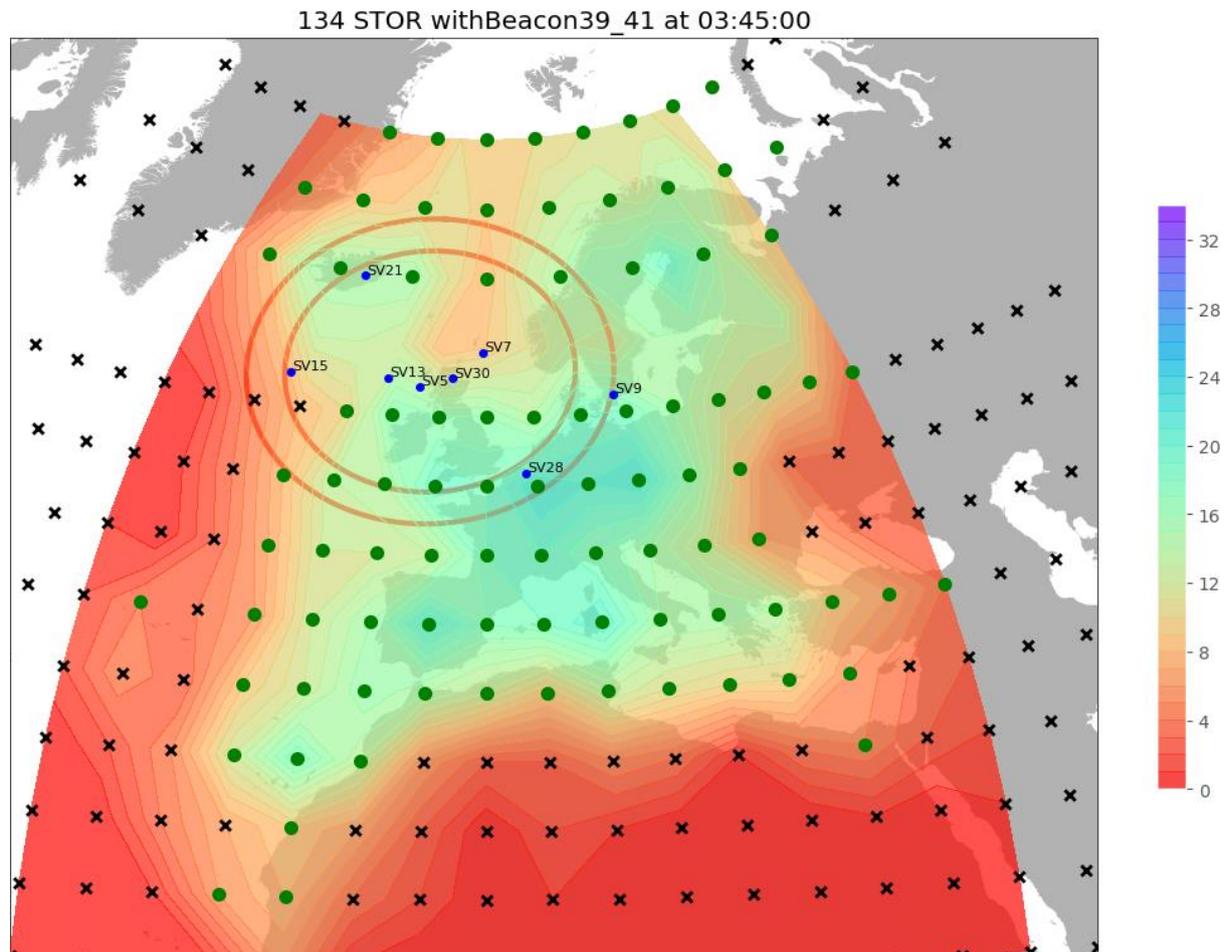
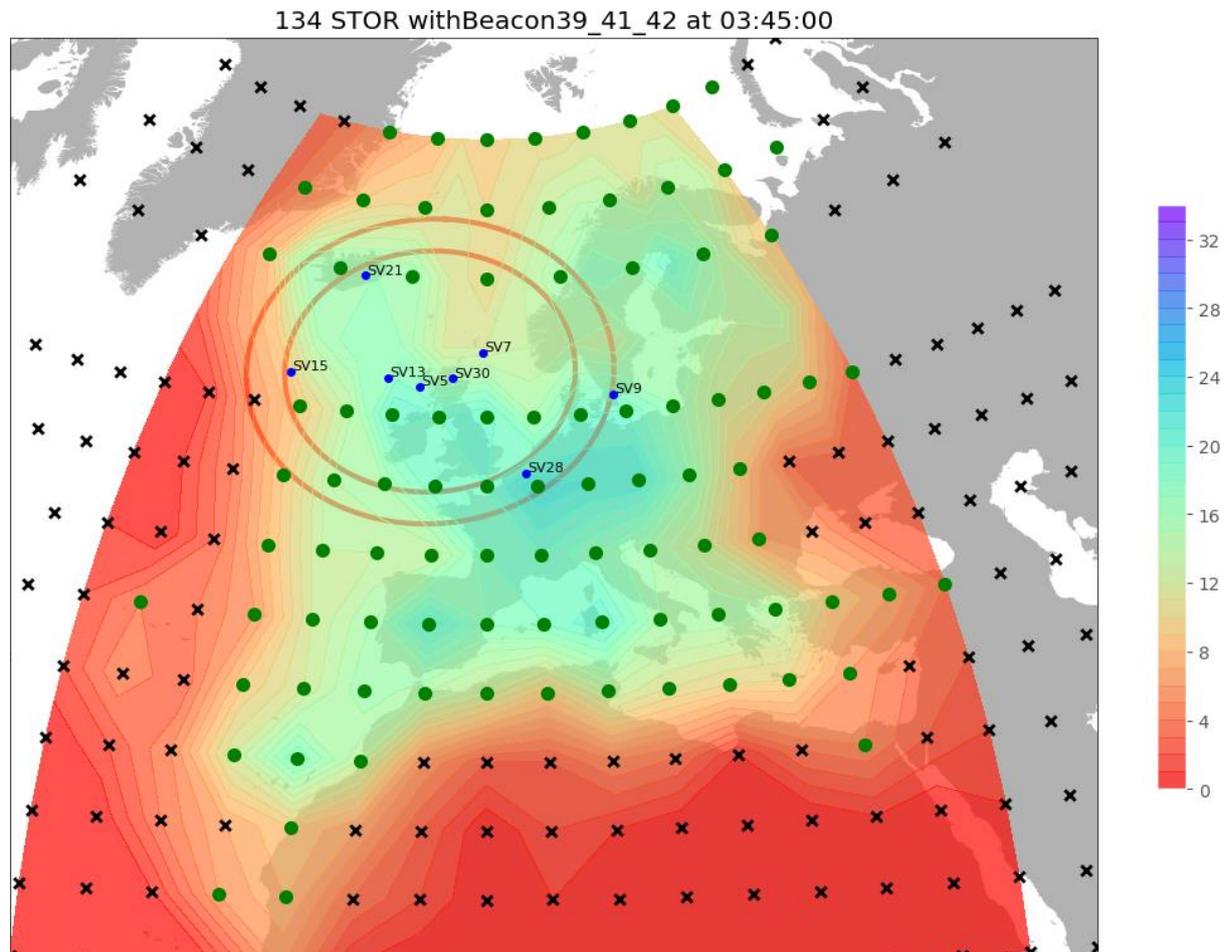
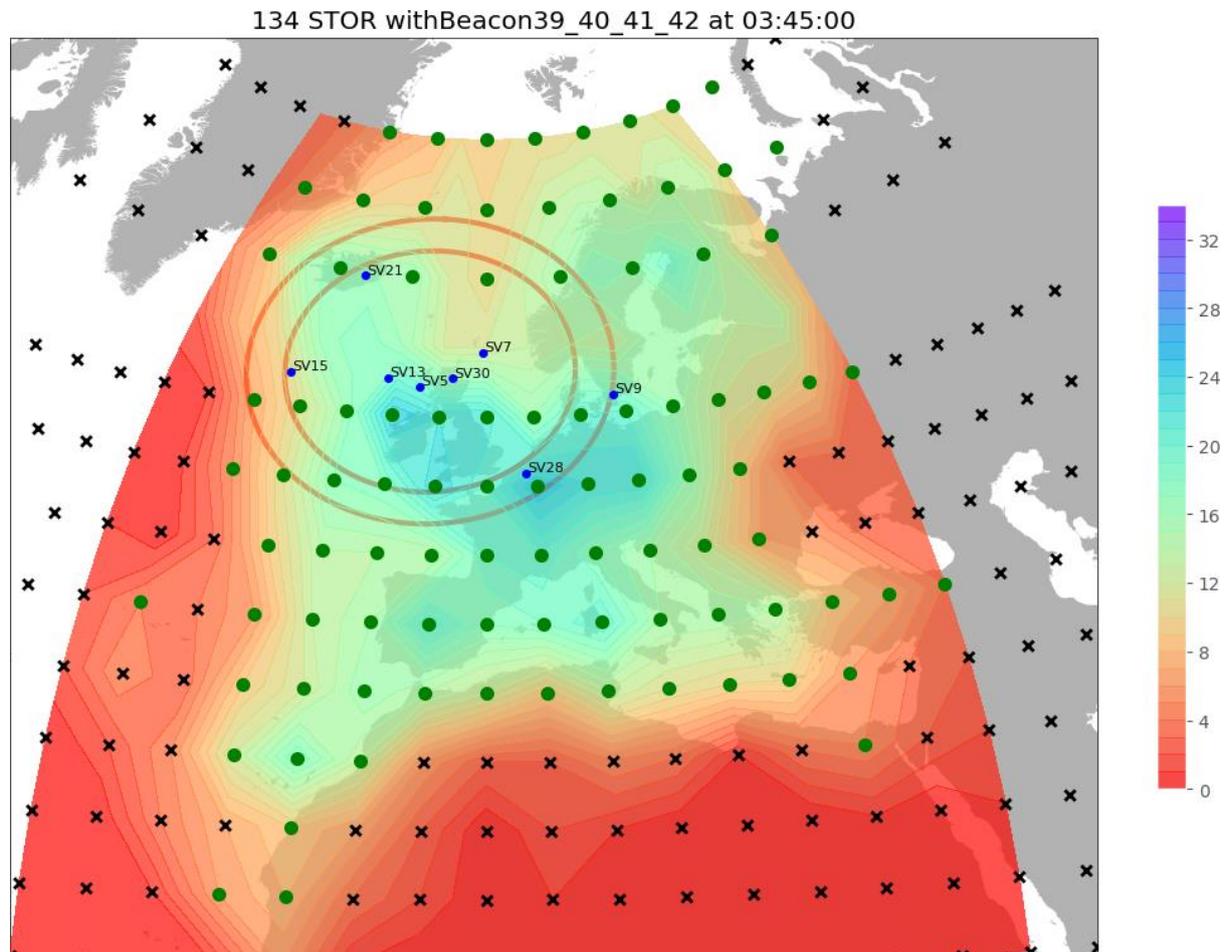


Figure 48: IGPs and STOR IPPs for 03:45:00 on 17/05/19, with two additional RIMS locations







6 Conclusions and Recommendations

The foregoing analysis has demonstrated that predicted EGNOS performance, using a calibrated simulation tool (SSv2), does not always meet the requirement for 10m (95%) accuracy. Spikes in performance contribute ‘availability events’. With the resolution available from a deterministic simulation tool, e.g. 30 second samples over a 24-hour period, it is not useful, or meaningful, to attempt to quantify continuity based on these availability events.

However, an investigation of the causes of these spikes in performance has shown that the exclusion of some satellite measurements from a user’s EGNOS position solution, due to the inability of the RIMS network to provide ionospheric corrections for those measurements, is their most likely cause. Exclusion of these satellites will result in worse DOP for the remaining satellites, which will probably result in a worse positioning solution. It will certainly result in a poorer 95% accuracy prediction. At locations further from the ‘core’ of the EGNOS RIMS network, the monitoring of the EGNOS ionospheric grid degrades, and it has been shown that the impact of satellite exclusion at the approximate western limit of the Irish EEZ is much greater than at coastal sites in the UK, with many more availability events being predicted.

These availability events can be mitigated by conventional beacon DGPS, where satellite exclusion due to a poorly monitored EGNOS ionospheric grid cannot be a factor – DGPS corrections for all common view satellites will be available, which will maintain full DOP, unless satellites are determined to be out of bounds. Historic data analysis has shown that DGPS can provide accuracy better than 10m during these predicted availability events.

They can also be mitigated by densifying the RIMS network in the area where the user receivers are affected by poorly monitored ionospheric grid points. The analysis has shown that the addition of just 3 extra RIMS at suitable north western locations could improve the monitoring of the ionospheric grid in the area where it affects users in the UK and Irish waters of their EEZs.

The findings presented in this report are based on a few test sites, and it is not possible to make recommendations based on these limited results. Nevertheless, they have highlighted an issue which could merit further investigation. Satellite exclusion due to inadequate ionospheric grid monitoring appears to lead to availability events, because the remaining satellites in the user’s local constellation have reduced geometrical strength. Geometrical strength is very closely linked to the elevation mask employed by the user receiver. There is likely to be a trade-off between the improved geometrical strength that comes from a lower elevation mask angle and the greater atmospheric and multipath errors, leading to lower measurement weights, that will be present in low elevation measurements. It is suggested that this should be studied in more detail. In addition, an investigation of the correlation between historic EGNOS corrections, from EMS files, and the SSv2 model predictions, should be undertaken, to corroborate the examples where predicted 95% position error exceeded the 10m threshold.



This report has identified areas where further investigation should be carried out:

- It is clear that the elevation mask angle applied to the EGNOS positioning solutions has a significant effect on the performance. The trade-off between improved satellite geometry and increased measurement errors, arising from variations in the user receiver elevation mask, should be investigated, to determine whether the exclusion of satellites due to ionospheric monitoring can be mitigated by the inclusion of measurements with slightly higher measurement errors. This should help to inform the choice of a standardised user receiver elevation mask.
- The validity of the SSv2 results should be further investigated by comparing the modelled performance with performance predicted from historic EMS data.
- To investigate whether there is correlation between the computed error in the EGNOS ephemeris and clock corrections (versus the Precise Ephemeris) and either IDOP or the RIMS tracking depth.



Richard Greaves, Project Manager NLA

Richard.greaves@nlaltd.co.uk



Registered office: Mill House Farm, Blackthorn Hill, Blackthorn, Oxfordshire OX25 1TJ
Registered in England: 10801372

