

An assessment of GNSS radio occultation data produced by Spire

March 23, 2020

Neill E. Bowler

Weather science technical report 640

This work was supported by funding from the Public Weather Service (PWS) at the Met Office, UK.

The dataset produced by Spire was supported under an ESA ARTES grant.

Abstract

Over recent years there has been increasing interest from commercial companies to produce observations of the atmosphere using GNSS-RO techniques (global navigation satellite system - radio occultation). Spire has achieved this by launching a large number of cubesats that operate in a constellation to observe the atmosphere. Over a period of 18 months these data have been assessed in comparison to current GNSS-RO data.

The differences between the observations and a short-range forecast (the innovations) indicate that the data are closely comparable to currently operational GNSS-RO data below 30km. Above this height the data are noisier than current satellites. Below 30km the data often have a smaller standard deviation than other observations, which can be achieved by using increased vertical smoothing. Although the vertical correlation length-scales for Spire data are larger than for the current data, they do not seem excessively large.

Given the good performance of the innovation statistics for the Spire data impact trials were run in the Met Office's NWP system. Tests were run assimilating data from the current observational network, and with no GNSS-RO data, adding data from Spire in various proportions, and replacing the data from the Metop-C satellite with data from Spire. It was found that there is a substantial forecast benefit from assimilating Spire data. In accordance with previous work the forecast impact of increasing the data volume is roughly proportional to the logarithm of the total amount of GNSS-RO data assimilated. Theoretical arguments suggest a different relationship with the number of observations, but these do not fit the results well.

1 Introduction

Numerical weather prediction (NWP) has brought great benefit to society over many years. NWP forecasts have been gradually improving in quality over time (Bauer *et al.*, 2015). This is due to improved numerical models, increased numbers of observations and improved methods for assimilating those observations.

Radio occultation (RO) observations from global navigation satellite systems (GNSSs) provide one source of observations to the global observing system. These measurements use the time delay in receiving signals from GNSS satellites at a low-earth orbiting (LEO) satellite, caused by refraction of these signals in the earth's atmosphere. As the LEO orbits the earth the GNSS satellite will appear to rise above or set below the earth's horizon, depending on the relative motion of the satellites. Thus a sequence of observations provides a profile through the earth's atmosphere — each such profile is referred to as an occultation (for more information, see for instance Kursinski *et al.*, 1997; Anthes, 2011).

A large increase in the number of GNSS-RO observations happened when the COSMIC-1 constellation was launched in April 2006 (Anthes *et al.*, 2008). Although this constellation has lasted much longer than expected, over time the number of observations received from COSMIC-1 has slowly decreased as the satellites have encountered operational problems. To some extent these observations have been replaced by GNSS-RO observations from the Metop and FY-3 satellites (see Figure 1 of Gleisner, 2018). Nonetheless the number of occultations being assimilated operationally at the time of writing remains less than it was soon after the launch of COSMIC-1.

The expected impact of increasing the number of GNSS-RO observations on NWP was studied by Harnisch *et al.* (2013) using an ensemble of data assimilations. In that study synthetic observations were created, and the impact of the observations were estimated by examining the change in the spread of the ensemble. They simulated 128,000 occultations per day and the effect of using certain fractions of these observations was considered. They found that 16,000 occultations per day gave approximately half the reduction in the ensemble spread of using the full set of observations. Therefore they recommended that 16,000–20,000 globally distributed occultations per day be considered as a minimum requirement for the global observing system. At the time of writing approximately 3,600 occultations are being made each day, clearly far short of the target recommended by Harnisch *et al.* (2013). One should also note that the COSMIC-2 constellation of equatorial satellites was launched in June 2019. These will more than double the number of occultations being received and will provide very high-quality observations. However, they are not yet operationally available to most NWP centres in near-real time. The second part of the COSMIC-2 constellation, with polar orbiting satellites was cancelled for budgetary reasons.

The Coordination Group for Meteorological Satellites (CGMS) of the World Meteorological Organisation (WMO) provides recommendations on the number of observations that should be made each day by the various observation platforms. They recommend that 6,000 globally distributed occultations should be a baseline value (Coordination Group for Meteorological Satellites, 2018), and that 20,000 occultations should be the target (Coordination Group for Meteorological Satellites, 2017).

For various reasons commercial companies have decided to launch satellites to make GNSS-RO observations. Three such companies are known to this author: Spire, GeoOptics and PlanetIQ. Of these companies Spire have made most progress and have a constellation of small satellites in orbit that are making a large number of GNSS-RO observations in near-real time (NRT). The assessment of those observations is the subject of this report. Spire's constellation is made up of small satellites (cube-sats) that are 30x10x10 cm in size. Each satellite has a forward- and rearward-pointing antenna for taking radio-occultation measurements, and can receive measurements from multiple GNSS constellations. Due to its small size and low power constraints there are limitations on the accuracy of the bending angle measurements at high levels. More information on Spire's satellites can be found at Masters (2019).

In addition to providing accurate observations of the atmosphere, GNSS-RO performs a secondary role as an anchor to the data assimilation scheme (Eyre, 2016). The bias in the observations is very small, and

therefore they are assimilated without bias correction by the data assimilation system. In turn this means that the bias correction applied to other observations cannot drift too far, thus helping to anchor the data assimilation system.

One factor that affects the number of observations used is the time taken for the observations to be produced and transmitted to the Met Office. For global NWP the Met Office runs the data assimilation twice every six hours. The main run is used to produce the forecasts that are presented to customers and is run at approximately 2:42 after the main data time. An update run is performed at 6:19 after the main data time, and is run to include any observations that may have missed the main run cut-off. It provides the background forecast for the subsequent main and update runs.

Taking the 1200 UTC cycle as an example, the main and update runs will seek to assimilate any observations that are taken between 0900 and 1500 UTC. The main run is made at approximately 1442 UTC, so will not use any observations taken after this time. Therefore, observations made between 1442 and 1500 UTC cannot be included in this run. In practice very few observations made after 1400 UTC will arrive in time to be used. The update run is performed at approximately 1819 UTC. Therefore an observation made at 1459 UTC will be included if it arrives within 3:20, but an observation made at 0901 UTC can arrive after 9:17 and will still be used in the update run. This variation in the amount of delay that is acceptable to the system means that the utility of an observation is approximately proportional to the speed with which it arrives. In order to compare the number of observations across the different tests the average of the number of observations assimilated by both the main and update runs will be used, since both types of run influence the forecast quality.

This report is intended to assess the quality of data from Spire against the Met Office NWP system, and to investigate the benefit of assimilating many more GNSS-RO observations than are currently operational. No comment is made on whether these observations should be purchased by the global weather community, or the relative benefits of publically or privately funded missions. Those are controversial questions that will need to be discussed by international bodies, such as the WMO.

2 Assessment of innovation statistics

The most straightforward way to assess any new source of observations is to calculate the innovation statistics for the observations. The innovation is defined as

$$d_b^o = y - H(x_b) \quad (1)$$

where y is a vector of the observed values, x_b is the short-range forecast from the most recent analysis of the atmospheric state and H is the observation operator. For this assessment the update run of the Met Office's operational global NWP system is used to provide the short-range forecast. The mean and standard deviation of the innovations, normalised by the value of the pseudo-observation from the background, is often used as a measure of observation quality for RO.

The observation operator H used at the Met Office is a 1-dimensional model that calculates refractivities and bending angles from a vertical column extracted from the NWP model (Healy and Thépaut, 2006; Buontempo *et al.*, 2008; Rennie, 2010). A set of intermediate pseudo-levels (Burrows *et al.*, 2014) is used to reduce the biases that can be caused by interpolation between model levels. The motion of the satellites is accounted for by considering the drift of the tangent point when extracting the model data (Burrows, 2014).

The figures are typically accumulated after quality-control (QC) has been applied. The principal QC steps are the flagging that is made by the data provider, a check on the difference between the background field and the observed value, and whether a 1-dimensional retrieval to calculate an analysis profile con-

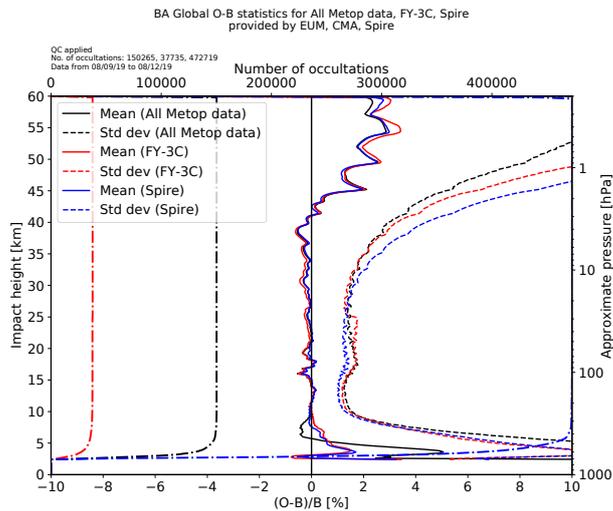


Figure 1: Number of observations, and mean and standard deviation of the bending angle normalised innovations for the Metop satellites, FY-3C and Spire.

verges sufficiently quickly (see [Rennie \(2010\)](#) for more details). The graphs shown are based on data between 0300 UTC on 8th September 2019 and 2100 UTC on 8th December 2019. On 20th September 2019 there were issues with the archive at the Met Office, so observations between 2100 UTC on 19th September and 0300 UTC on 20th September are not used in these plots.

Figure 1 shows the number of observations, the mean and standard deviation of the bending angle innovations normalised by the background bending angle, after QC has been applied. There are many more observations from Spire than from either the three Metop satellites or from FY-3C. Above 10km the mean innovation from the three data sources is very similar. There are subtle differences in the mean innovation between 10km and 40km, with Spire generally having slightly higher values of the mean innovation than FY-3C or the Metop satellites. At least in part these appear to be due to differences in latitudinal sampling of the observations, since the data from Metop and Spire is similar when examining similar graphs for restricted latitude ranges. Below 10km the data from FY-3C and Spire are relatively similar, and a substantial bias in the data from Metop is evident.

As expected the observations from Spire have larger standard deviations above approximately 30km. The Metop satellites have particularly low standard deviations in this region. Between 10km and 25km the Spire observations have the smallest standard deviation of the three data sets. Below 10km FY-3C has a similar standard deviation to the Spire data, which is lower than that from the Metop satellites.

Below 30km the differences in standard deviation between the data sets is partially explained by differences in the amount of vertical smoothing that is applied to the bending angle observations. Using greater vertical smoothing decreases the standard deviation of the innovations, but this comes at the expense of increasing vertical correlation in the innovation statistics. A vertical profile which has substantial vertical correlations will contain less information than one whose errors are uncorrelated. Figure 2 shows the vertical correlations of the bending angle innovations. Ideally this would be close to a diagonal matrix, with vertical correlations only deriving from the background-error covariances. Figure 2 shows that both the Metop and Spire systems have relatively clean vertical correlations, without large correlations away from the diagonal. The core of positive correlations near the diagonal is wider for Spire than for Metop, which is consistent with the smaller standard deviations being due to greater smoothing. There is a weak pattern of positive and negative bands above 35km in the Spire data — it is not clear what causes this. The Metop data has a band of off-diagonal positive correlations above 40km — this is also poorly understood.

There are some oscillations in the vertical correlations for Spire at around 16–20km impact height. The

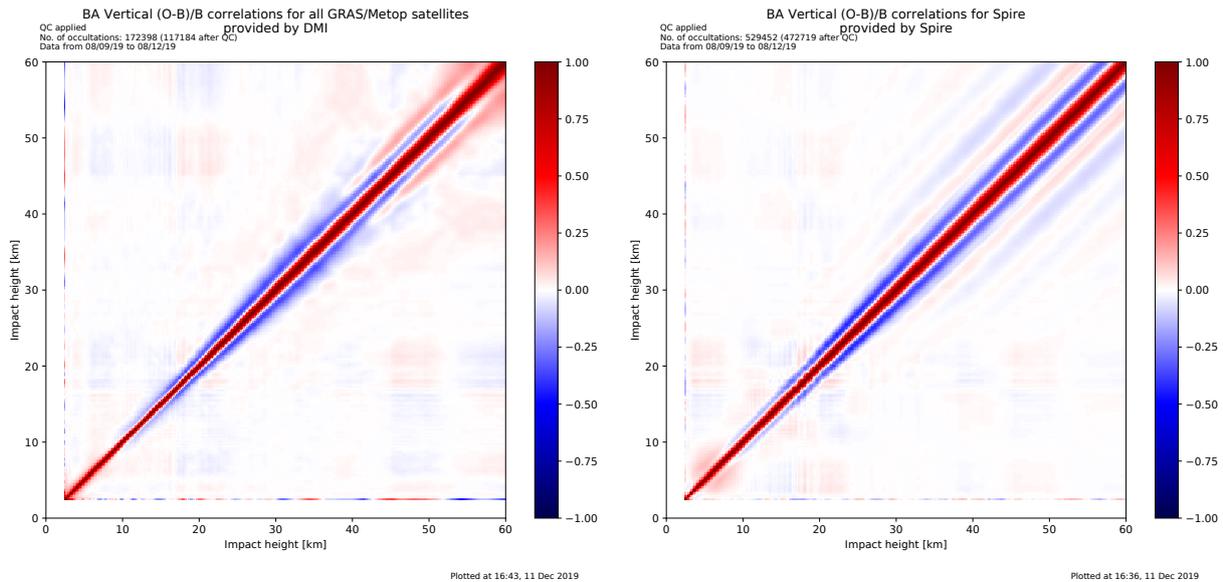


Figure 2: Vertical correlations of bending angle innovations for all Metop observations (left) and Spire observations (right).

separation of the maxima in these oscillations is approximately the same as the separation of the vertical levels in the Met Office model. Therefore they might be related to the interpolation of data between model levels. It is not clear why these oscillations are peculiar to data from Spire. Spire provides its bending angle observations on a regular vertical grid, with 200m separation between each level. By contrast EUMETSAT and the Chinese Meteorological Administration (CMA) provide observations on a vertical grid whose levels are closely spaced near the surface and further apart at higher altitudes. On the other hand UCAR use an approximately regular grid, similar to Spire.

Figure 3 shows the global bending angle innovation statistics from Spire, separated by GNSS classification. The largest fraction of the data comes from the GPS network, followed by GLONASS and Galileo. QZSS provides the least data, since this is a regional system with few satellites. There are subtle differences in the mean behaviour; these are largely due to the different latitudinal sampling with the different GNSS constellations (see Figure 4). GLONASS provides many more observations in the tropical region, whereas GPS and Galileo provide more observations in the extra-tropical regions. The most notable difference is that the data from Galileo have a slightly larger mean innovation in the tropical troposphere than for the other GNSS constellations, see Figure 5. This bias has been a subject of active research for Spire, and was substantially reduced by a change made in late September (D Masters, personal communication) but has not yet been reduced to zero. It is also apparent that some of the older satellites in the constellation have larger biases in the tropical troposphere than the newer satellites (not shown). Whether the evolution of the observing constellation is viewed as a positive or negative feature may depend on the particular user in question. Typically users in NWP applications value regular improvements in the quality of an observation source as it can lead to forecast improvements. On the other hand, users studying the climate typically value stability in order to understand trends in the data.

The differences in the standard deviation are also affected by latitudinal sampling, but there are also notable differences between the constellations. Observations using GPS and QZSS generally have the smallest innovation standard deviation at high levels of the GNSS constellations, and Galileo is the largest. Processing Galileo signals has proved challenging due to the higher chip-rate (the speed of the data encoding) used by these satellites (D Masters, personal communication). In the core region for GNSS-RO observations (10km to 30km) the difference in standard deviation between the GNSS constellations is small

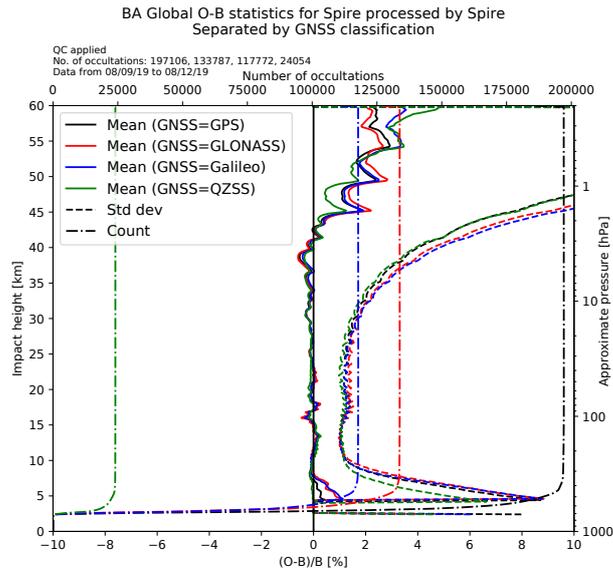


Figure 3: Global bending angle innovation statistics from Spire, separated by GNSS classification.

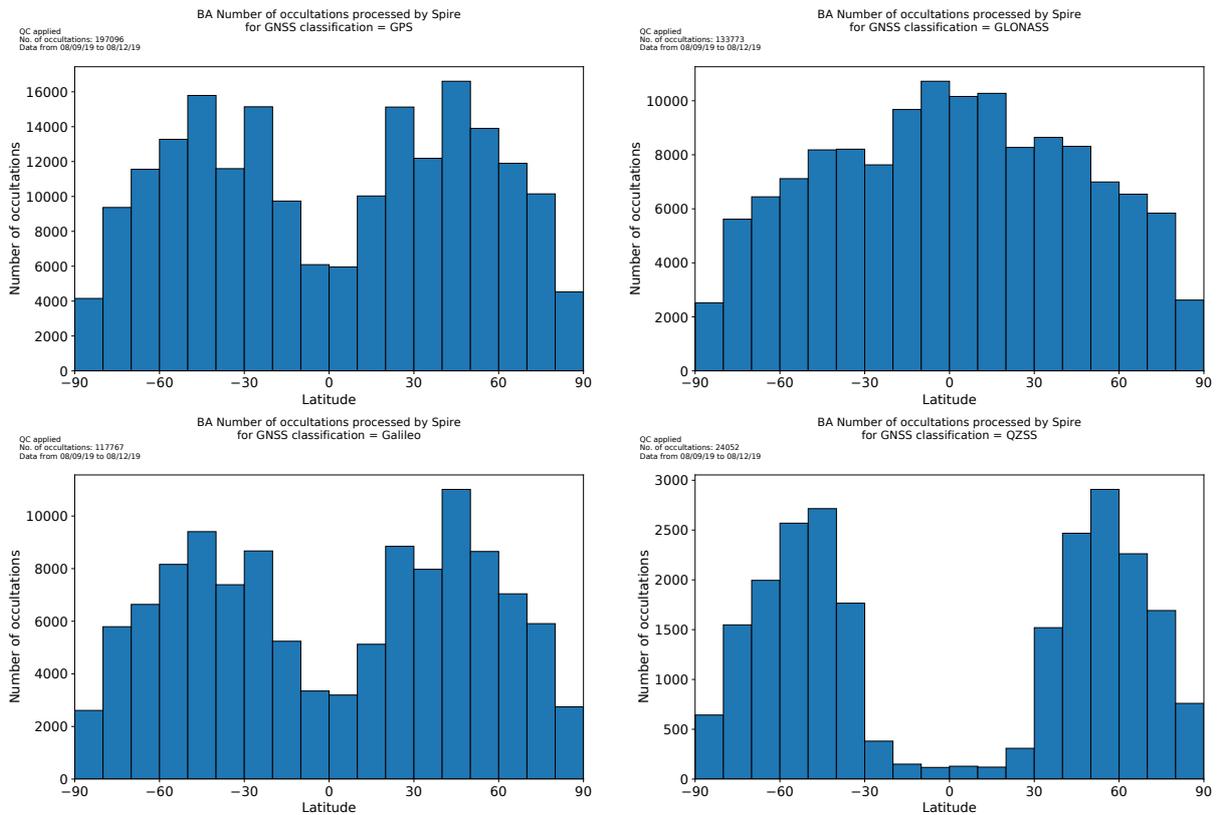


Figure 4: The number of occultations received from Spire plotted against latitude. Each graph shows the distribution for a given GNSS classification.

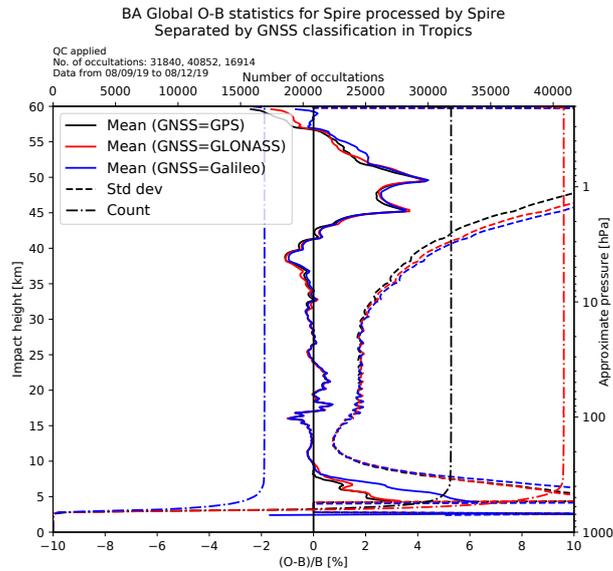


Figure 5: Bending angle innovation statistics from Spire in the tropical region, separated by GNSS classification.

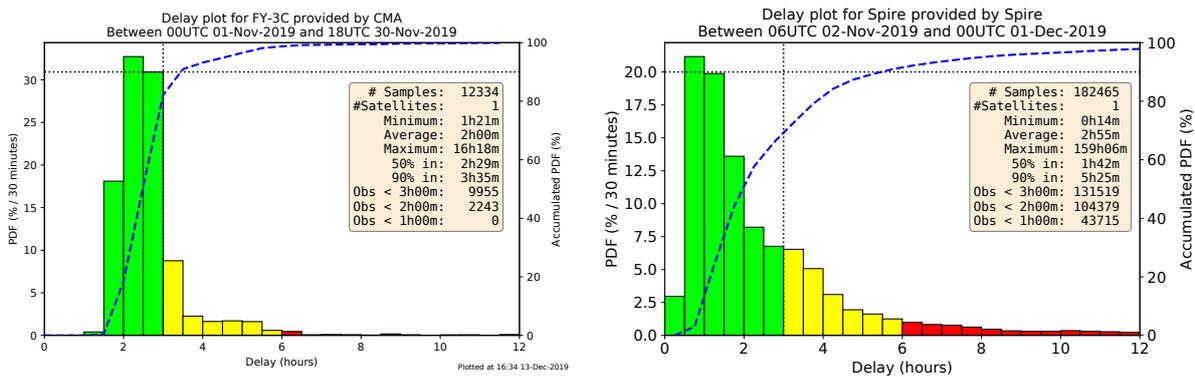


Figure 6: (left) Time taken between an FY-3C observation being made and it being available in the Met Office data-base. (right) Time between an observation being made and the data being available on Spire AWS servers, for the first processing of the observations. Data calculated for the month of November 2019.

and it is difficult to conclude that one observation group is better than any other. In the troposphere QZSS has smaller innovation standard deviations than the other constellations. In part this is due to QZSS providing fewer observations in the tropics than the other platforms, but also due to the innovation standard deviations being smaller in the northern-hemisphere extra-tropics (not shown).

Assessing the time taken for an observation to be available is more difficult for Spire data than for the operational data sources. The precise orbit information for each satellite can be updated after a given observation is made. When this occurs Spire reprocess the occultation using the updated orbit information. The updated occultation replaces the original occultation on the Amazon Web Services (AWS) servers, which increases the apparent delay in creating the observations. Figure 6 shows the time taken between the observation being made and being available for FY-3C and for the first processing of each occultation from Spire. Since the times given for Spire are based on the time-stamp of the file on the AWS servers, one should add the time taken to retrieve the observations to the Met Office, which is typically between 7 and 21 minutes, depending on the amount of data that is being extracted.

Figure 7 shows the delay histogram for the final time that each occultation is processed. As can be

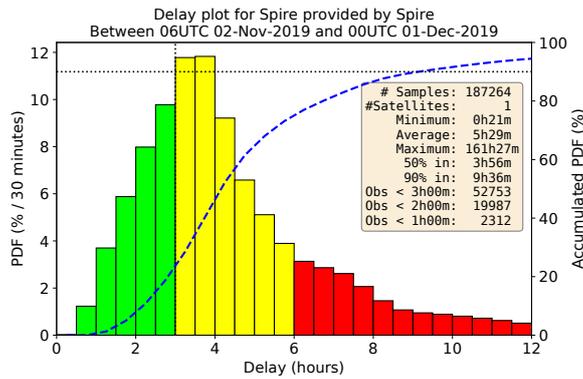


Figure 7: Time between an observation being made and the data being available on Spire AWS servers, for the final reprocessing of the observations. Data calculated for the month of November 2019.

Source	Main	Update
Metop	71%	100%
FY-3C	52%	95%
Spire	56%	92%

Table 1: Estimates of the fraction of observations are available in time for the main and update runs of the Met Office NWP system.

seen, the additional delay created for reprocessing the observations can be substantial. Validation of the first and final processing of each occultation against Met Office short-range forecasts (not shown) indicates that the benefit from the reprocessing is small, and limited to altitudes above 35km.

Another way to assess the timeliness of the data is what fraction of the observations meet the cut-off times for the main and update runs. Table 1 provides an estimate of the fraction of observations that make cut-off for the Met Office’s main and update runs. Data from the Metop satellites arrive very quickly and hence the percentages for these satellites are very high. Data from FY-3C take longer to arrive, and just over 50% arrives in time for the main run. The arrival times from Spire are more diverse, which means that much of the data arrives very quickly, but some of the data fail to arrive in time for the update run.

This analysis has shown that the mean and standard deviation of the innovations for Spire observations is similar to that of other satellites below 35km. Above this height the observations from Spire are noisier than the operational platforms. The observations using different GNSS constellations sample different regions of the globe, and are of different quality due to differences in the transmitted signals. The timeliness of the data received is similar to that from FY-3C, although the distribution of times is somewhat different. Spire’s reprocessing of the observations makes the timeliness assessment more complicated, but this is a minor detail since the quality of the data is only marginally affected by this reprocessing.

3 Assimilation experiments

An important incentive for gathering more GNSS-RO data is to improve numerical weather prediction (NWP). Data assimilation systems for NWP ingest large numbers of observations in order to calculate an accurate estimate of the current state of the atmosphere. The Met Office uses a hybrid four-dimensional variational method (4D-Var) to perform the data assimilation [Rawlins *et al.* \(2007\)](#); [Clayton *et al.* \(2013\)](#); [Bowler *et al.* \(2017\)](#).

Observations are assimilated from a wide variety of sources, including radiosondes, surface weather stations, aircraft and many satellite platforms. To mimic the data cut-off used by operational systems

observations from Spire were extracted three times every six hours from the AWS servers. The first two extractions were started at 2:42 and 6:19 after the central time in the data assimilation window, to mimic the main and update runs. The final extraction was run at 24h after the central time, with a view to receiving all the GNSS-RO observations that were made. At various times during the experimental period there were problems that led to the data not being extracted at the given time. All the issues were related to the experiment being run in a non-operational configuration. Fortunately staff at Spire were able to provide the observations that would have been extracted by the main and update runs had the problems not been encountered.

To test the impact of data from Spire in the NWP assimilation system a number of experiments were run between 8th September and 8th December 2019. These used a low-resolution system that aims to mimic the operational NWP system as closely as possible. The forecast model was run at N320 resolution (640x480 grid-points on a latitude-longitude grid) and the data assimilation was run at N108 (216x162 grid-points) and N216 resolution (432x324 grid-points). A hybrid 4D-Var data assimilation scheme was used, as described by Clayton *et al.* (2013). The NWP model is the GA6 / GL8 configuration as described by Walters *et al.* (2017). The first week of data was removed from the verification to avoid transient effects at the start of the period. The “control” against which all other experiments will be judged is a run of the NWP system with all currently operational observational data. This also includes GNSS-RO observations from FY-3D and KOMPSAT-5 as they began to be used operationally on 10th December 2019. It also includes the revised observation errors that were developed in Bowler (2019).

3.1 Removing all GNSS-RO data

The first experiment is to remove all the GNSS-RO observations, to understand the impact that the existing observation network has. Figure 8 demonstrates that the existing RO network has a very large impact on the quality of the forecast. The area of the triangles shows the size of the change relative to the control of the the root mean square difference (RMSD) between the forecast and ECMWF analyses. If the triangle fills the box in which it sits, then this corresponds to a 20% change in the RMSD (or greater). The verification is performed against ECMWF analyses, as this covers the globe evenly and is independent of the Met Office modelling system. Verification against observations shows a similar pattern, but with smaller impacts overall — one example scorecard is shown in section 3.4. If a change in the RMSD is judged to be statistically significant by a two-sided Wilcoxon signed-rank test at the 5% level, then the area surrounding the triangle is shaded. It is based on the difference in RMSD for each forecast cycle (12h apart) which are treated as independent data points. The largest degradations in forecast quality are for temperature at 100hPa and 50hPa, as might be expected. There are a few quantities where the forecast appears to be improved by removing all RO data. Temperature at 250hPa in the tropics appears improved due to a reduction in forecast bias. The bias is also improved by the introduction of Spire data (see below). Screen temperature in the northern-hemisphere extra-tropics is also degraded, in terms of bias and standard deviation. It is not clear why this is.

As will be shown below, the degradation caused by removing existing observations is much larger than the benefit from adding additional observations. Harnisch *et al.* (2013) found that the increase in ensemble spread caused by removing all GNSS-RO observations was approximately the same as the decrease in ensemble spread from increasing the total number of occultations to 8000 per day in the northern hemisphere and tropics, and to between 16,000 and 32,000 per day in the southern hemisphere. The results below indicate that degradation in the Met Office’s system for losing all GNSS-RO observations is larger than any reductions in RMSD seen from adding observations from Spire. The role of GNSS-RO in anchoring the data assimilation was not tested by Harnisch *et al.* (2013), and this may be the cause of the increased impact seen here. It is also likely that verification against ECMWF analyses overestimates the impact of the denial of GNSS-RO observations in the short range, since the same observations are used

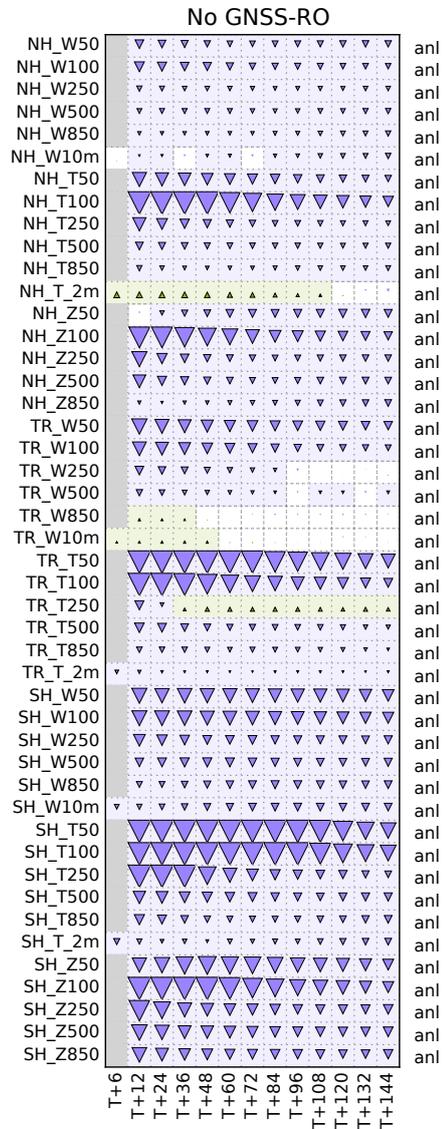


Figure 8: Forecast verification scorecard using the current observation network without any GNSS-RO data. The area of the triangles is proportional to the change in RMSD, relative to the control system measured against ECMWF analyses. Green (purple-blue) triangles indicate better (worse) performance for the experiment. The y-axis denotes the forecast variable being considered. NH, TR and SH denote the northern hemisphere > 20 degrees, tropics and southern hemisphere < -20 degrees, respectively. The letter following the underscore indicates the weather parameter, with W, T and Z signifying vector wind field, temperature and geopotential height, respectively. The numbers after this indicate the height of the observation in hPa, except for 2m and 10m that are heights above the surface. The x-axis shows the forecast lead time in hours.

in creating the ECMWF analyses. This only applies to observations before the verification time, so the effect would diminish rapidly with increasing forecast lead time. However, the conclusion that the existing observations provide a much larger benefit than supplementing the observation network remains when performing the verification against observations (not shown).

3.2 Adding data from Spire

When assimilating Spire observations, observation uncertainties from the COSMIC-1 satellites were used. The COSMIC-1 constellation performed well, and the innovation statistics from Spire are reasonably close to these data. In addition the minimum error for a bending angle observation was raised from 3 μrad to 6 μrad , since we noted in Figure 1 that the observations from Spire have a larger standard deviation than other satellites above 30km.

Figure 9 shows the impact from assimilating the observations from Spire that were available in near-real time (that is, by the time that the main and update runs usually occur) in addition to all the data that were used in the control. The prevalence of green upward-pointing triangles on this graph demonstrates that there is a substantial benefit from the use of these observations. The main negative element in this verification is screen temperature in the northern hemisphere extra-tropics at shorter lead times. Given that this quantity was improved by removing all GNSS-RO observations this result does not seem to be specific to Spire data.

Figure 10 shows a further experiment that uses 50% of the near-real-time observations from Spire to test whether we see 50% of the benefit when using 50% of the data. The 50% of observations to use were randomly selected from the available observations each cycle. Therefore they have similar data coverage to the full set used in Figure 9, and this is not the same as using the observations from half of the satellites, which would have a less regular data coverage. The overall pattern of the verification is very similar to that seen for Figure 9 with positive impact for most variables. The negative impacts are also seen in very similar locations. The overall impact is slightly larger than 50% of the impact from the full NRT set of observations, as might be expected since the second 50% of the observations are added to an already well-observed system. This conclusion may be affected by the metric used in the verification, as shown by Eyre and Weston (2014).

Figure 11 shows the impact from assimilating all the Spire observations that were available after a delay of 24h. This experiment demonstrates the impact that would be seen if all the observations were delivered in time for the main run (including those for times that are after the main run cutoff). Thus the main and update runs use the same number of observations from Spire, and this number is larger than the number of observations seen by the update run in the NRT experiments. All the other observations use the standard data cut-offs. The impact in this experiment is approximately 30% larger than the experiment assimilating NRT observations.

Another way to assess the impact of assimilating new data is to consider the change in the RMS innovation for other observation types. Figure 12 shows the change in the RMS innovation for the CrIS instrument on JPSS0 (NPP) and the IASI instrument on Metop-B. These are both hyper-spectral IR sounders, but their results give a representative picture of the overall changes. As seen for the CrIS instrument, the general pattern is that the standard deviation of the innovations have been reduced by the introduction of the extra data. This is to be expected since the forecasts at short lead times have been improved. For the IASI instrument, however, there is a group of channels where the standard deviation of the innovation has been increased.

The channels where degradation in the innovations are seen are channels that are sensitive to temperature in the stratosphere and upper troposphere. None of the verification results (against radiosondes) have degradations in the short-range temperature verification (not shown). So it is not clear why this degradation is apparent for IASI, and mostly absent for CrIS, which is a very similar instrument.

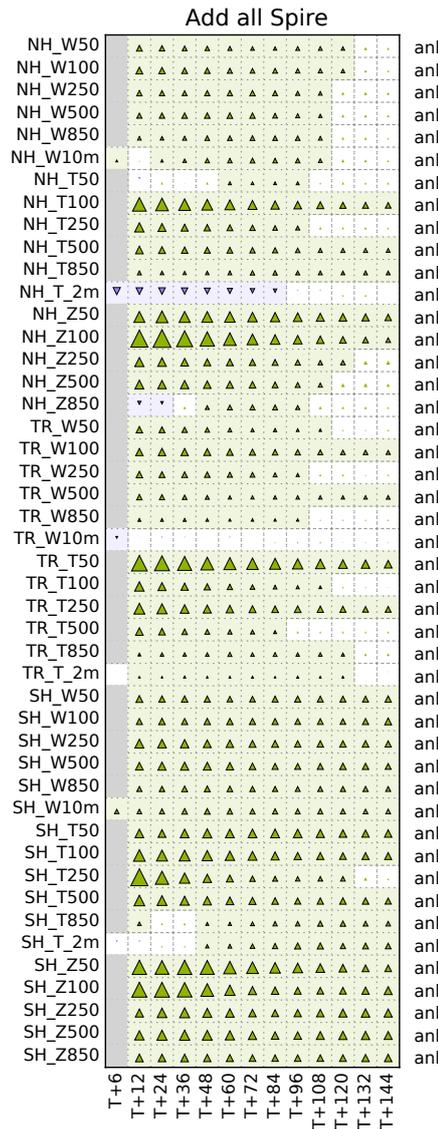


Figure 11: Forecast verification scorecard using all the observations from Spire in addition to the current observation network. Figure format as Figure 8.

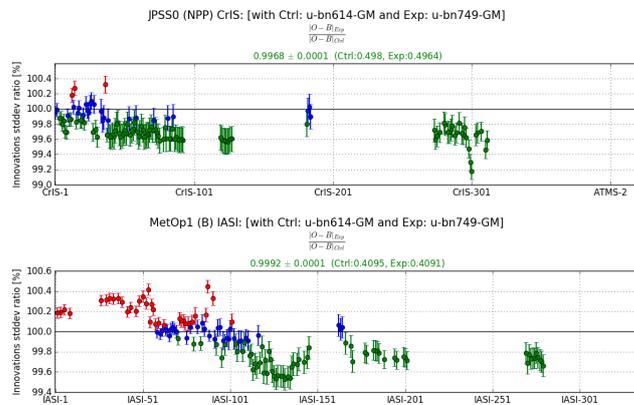


Figure 12: Change in the RMS innovation between the experiment assimilating all the observations from Spire and the control. The innovations are calculated over the period for the CrIS instrument on JPSS0 (NPP) (top) and the IASI instrument on Metop-B (bottom).

The experiments assimilating varying amounts of data from Spire have shown substantial positive impacts across many variables. The tests using varying numbers of observations have indicated that greater benefit is seen with a greater number of observations. There are negative impacts in screen temperature in the northern-hemisphere extra-tropics. This mirrors an improvement seen in this quantity when removing all GNSS-RO data. This negative impact is also seen in the short range when verifying against observations (see section 4).

3.3 Comparison with Metop-C

By launching a large number of small satellites Spire is aiming to provide a very large number of observations. These observations are expected to be of lower quality than those from larger satellites. The GRAS instrument on Metop, for instance, uses an ultra-stable oscillator and the receiver exhibits very low thermal noise (C Marquardt, personal communication) which means that its observations are particularly accurate at high levels. The ultra-stable oscillator has relatively high power consumption, which means it is not suitable for use on a small satellite. The low noise from Metop can be seen in the standard deviations above 30km in Figure 1. However, for NWP the performance at lower levels is more important than at higher levels. Therefore it is interesting to compare the benefit seen in NWP between observations from Metop with those from Spire.

Figure 13 shows the impact from the removal of observations from Metop-C. The forecast scores are widely worse without these data, as shown by the blue downward-pointing triangles across the scorecard.

Figure 14 shows the impact of removing the observations from Metop-C and including an equivalent number of observations from Spire. The number of observations added has been chosen to give an approximately equal amount of data assimilated. The difference in the assimilated data volume is approximately 0.16%. The change in performance for replacing the data source is mixed, with some small decreases and increases in RMSD for some variables. Most of the changes in the forecast skill are not statistically significant. Therefore although we see that observations from Spire and Metop-C are of similar benefit to the forecast scores, it is not possible to conclude which are more beneficial.

The Met Office uses quality control (QC) flags generated by the Danish Meteorological Institute (DMI) that remove around 30% of the data from Metop. Therefore a greater amount of raw data from Metop is used in this comparison. This is likely to change in the coming year due to processing changes expected by EUMETSAT. The observations from Spire have been selected by making a random selection from the full set of observations. Therefore the data from Spire are coming from a large number of satellites, but observations from Metop-C are taken from a single satellite. Many of the Spire satellites are in a very similar orbit, meaning that it is not clear how important this factor is.

We have seen that removing the observations from Metop-C brings a clear degradation to the overall forecast performance. Replacing that with an equivalent amount of data from Spire gives equivalent performance for NWP.

3.4 Extrapolation of impact

The scorecards shown in this report summarise the change in the NWP forecast performance across a range of variables. It is common practice to summarise the change in performance into a single number by finding the average percentage change in RMSD across all the variables and lead times. Here we consider how this overall impact varies with the number of GNSS-RO observations assimilated. Figure 15 shows the overall impact (measured as the percentage reduction in RMSD across a range of variables against ECMWF analyses and observations) as a function of the total number of observations. A linear best-fit for observations and ECMWF analysis is also shown on the graph, including the equation that defines this line. Looking back at the results of Harnisch *et al.* (2013) it appears that these results also fit closely to a

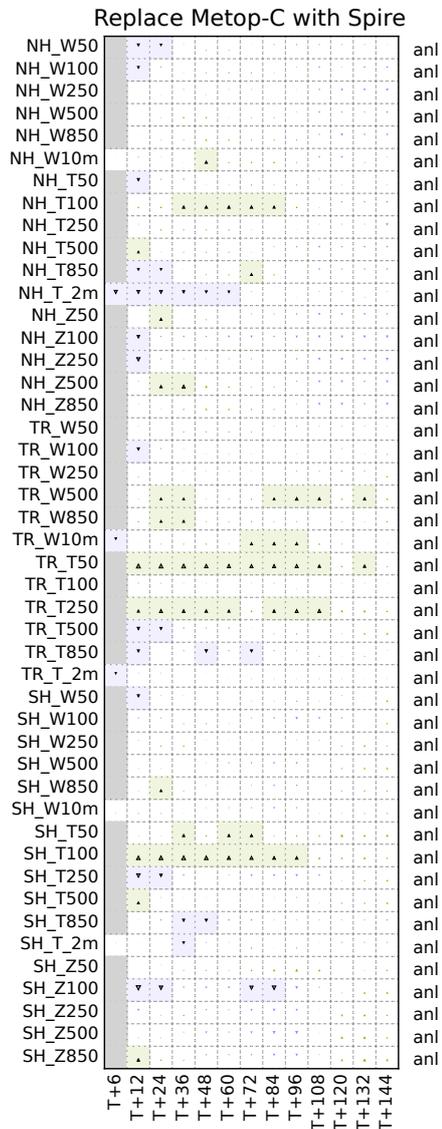


Figure 14: Forecast verification scorecard using the current observation network, minus Metop-C and adding an equivalent amount of Spire data. Figure format as Figure 8.

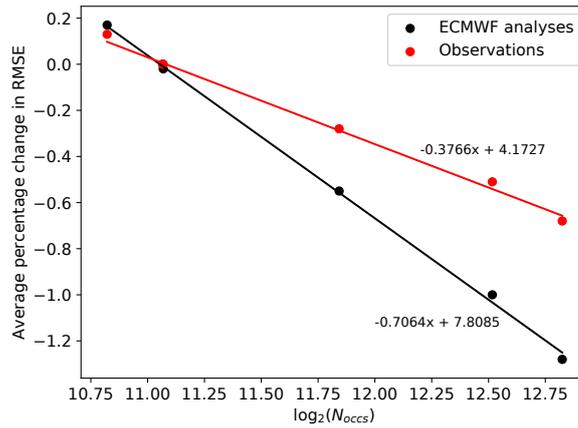


Figure 15: Overall impact of assimilating various numbers of GNSS-RO observations compared with the control. A linear fit against the logarithm of the number of occultations per day is shown. The control against which changes are measured is at $x = 11.07$, with points to the left of this representing the experiment removing Metop-C observations.

logarithmic dependence. Therefore it seems that the results here are somewhat consistent with those of [Harnisch *et al.* \(2013\)](#).

[Harnisch *et al.* \(2013\)](#) used a simple fitting function based on the square root of the number of observations to approximate the ensemble spread for 128,000 profiles, based on the results for 32,000 and 64,000 profiles. [Eyre and Weston \(2014\)](#) showed that the analysis “accuracy” (the inverse of the analysis-error covariance) should scale with the inverse of the number of observations, although note that this is the total analysis “accuracy” not the change in RMSE shown here. Therefore we note that the logarithmic fit given in Figure 15 is an empirical relationship and is not supported by a theoretical understanding. Various attempts have been made to fit the change in the RMSD with a power of the number of occultations, but these were unsuccessful.

The linear best-fit allows us to estimate the expected impact from introducing additional observations. From these we would expect that a doubling of the total number of GNSS-RO observations would lead to a 0.71% reduction in the RMSD as measured against ECMWF analysis and a 0.38% reduction against observations. Obviously, such numbers are purely estimates and should be treated with a healthy degree of skepticism. Given that the reduction in the RMSD for adding observations to the NWP system decreases with the number of observations that are present, it may be the case that estimates of the impact per observation (such as FSOI ([Cardinali, 2009](#))) need to consider this effect.

4 Verification against observations and with standard deviation

So far the verification has focused on comparisons with ECMWF analyses, since these are highly accurate and evenly distributed around the globe. Impacts against observations are generally similar, but smaller in magnitude. Figure 16 shows the scorecard for assimilating all Spire observations that are available in NRT verified against observations. Comparing this with Figure 9 shows that the pattern of the change in RMSD is very similar, but smaller in magnitude. Two possible reasons for the difference in magnitude of impact between the two verifications are as follows. Firstly, that GNSS-RO observations are likely to have greater impact over the oceans, since there are fewer conventional observations and their benefit in bias correcting satellite radiance observations will also be concentrated over regions where radiances are used more. Secondly the noise in a radiosonde observation is typically larger than that in the ECMWF analyses, and this reduces the apparent improvement from an increase in the forecast skill.

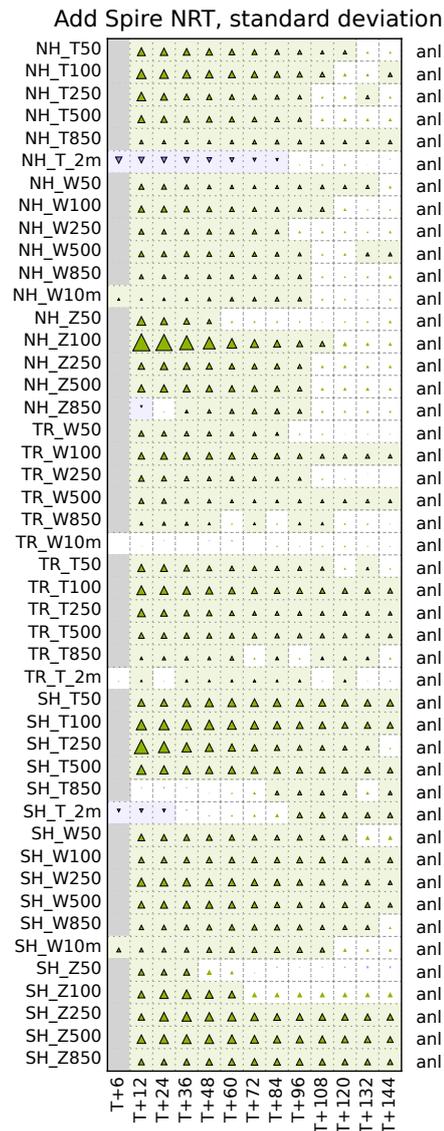


Figure 17: Forecast verification scorecard using real-time observations from Spire in addition to the current observation network. Figure format as Figure 8, but verification using the standard deviation instead of the RMSE.

The improvement in RMSD can be a result of changes in the forecast bias as well as reducing the random errors of the forecast. In order to help separate these factors Figure 17 shows the change in the forecast standard deviation when adding near-real-time observations from Spire. This should be compared with Figure 9. The overall pattern of improvement is rather similar, although many of the impacts are smaller with the use of standard deviation. Since GNSS-RO data generally reduces the forecast biases it should be expected that the change in RMSD is benefitting from the reduction in the biases.

Earlier in this paper it was noted that the RMSD of forecasts of 250 hPa temperature in the tropics was reduced when removing all GNSS-RO observations from the assimilation (Figure 8). Figure 18 shows the standard deviation and bias of forecasts of this quantity as a function of lead time. This shows that removing all GNSS-RO observations reduces the bias of the forecasts, and this leads to the reduction in the RMSD measured against ECMWF analyses. Similar behaviour is seen against radiosondes. At the same time removing these observations increases the standard deviation, but it is the bias change that is dominant in the RMSD calculation. Interestingly, adding the near-real-time observations from Spire also reduces the

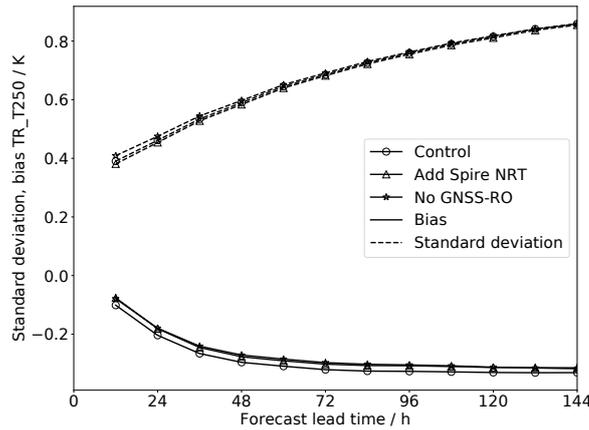


Figure 18: Standard deviation and bias of forecasts of 250 hPa temperature measured against ECMWF analyses as a function of lead time.

bias as well as reducing the standard deviation. Thus, both removing and adding GNSS-RO observations improves the forecast bias — this is not well understood. Given that ECMWF analyses and radiosonde observations may have biases there may be interactions between these data sources.

5 Conclusion

GNSS-RO observations have been produced by Spire, and its quality has been evaluated. The amount of data provided more than doubles the number of observations compared to that which is currently used operationally, although noting that the COSMIC-2 constellation will also do this, and will soon be operationally available in NRT. The statistics of the innovations are similar to those from established satellites. The mean innovation is very similar, with hardly any difference between Spire and FY-3C. Above 30km the standard deviation of Spire innovations is larger than other satellites. Below 30km the standard deviation is similar to, or smaller than other satellites. This is due to differences in the vertical smoothing of the observations, as seen in the vertical correlation of the innovations. The vertical correlations for Spire data are larger than other satellites, but not excessively so.

Observations from the different GNSS constellations perform in a broadly similar manner. Many of the differences that are seen are due to different latitudinal sampling of the observations. In addition to this it appears that observations from GLONASS and Galileo have slightly larger standard deviations above 30km, and observations from Galileo have larger biases in the tropical troposphere. This bias has been reduced, and some of the remaining problems are due to the presence of some of the older satellites in the Spire constellation.

The time taken to receive data from Spire is similar to that from operational satellites, this assessment is complicated by Spire’s reprocessing of the observations. Although much of the data is quick to arrive, some observations take much longer to progress to the AWS servers. Improving the timeliness of the data supply is an ongoing development by Spire (D Masters, personal communication). A similar percentage of data makes the Met Office’s cut-off for the main and update runs to that for data for the FY-3C satellite.

Tests with the Met Office’s NWP system demonstrate the substantial benefits to forecast quality that GNSS-RO data brings. Assimilating observations from Spire in addition to the current operational network brings substantial benefits to forecast performance. These benefits are seen for almost all forecast variables and lead times. However, the reductions in the RMSD are much smaller than the detriments seen from removing all operational GNSS-RO observations, despite this being a smaller number of observations.

Removing data from Metop-C and at the same time including an equivalent number of observations from Spire leaves the system approximately unchanged. Therefore we conclude that the two data sources are of similar quality. There are factors which make direct comparison of quality difficult. The Met Office uses aggressive quality control on observations from the Metop satellites meaning that more raw data from Metop-C is used in this comparison. In addition the sub-selection of Spire observations was random, leading to an even distribution of observations across the globe.

The reduction in the RMSD when assimilating additional GNSS-RO observations is seen to be approximately proportional to the logarithm of the total number of observations assimilated. Although not supported by theoretical arguments, this is consistent with the results of Harnisch *et al.* (2013). The impact of GNSS-RO data is seen to be larger when measured against ECMWF analyses than when measured against observations. Doubling the total number of observations assimilated decreases the RMSD by approximately 0.71% (ECMWF analyses) or 0.38% (observations).

These results have looked at the performance of GNSS-RO observations from one commercial company (Spire) when using one NWP system (the Met Office's). It would be beneficial to extend this study to compare multiple sources of additional observations using multiple NWP systems to improve the confidence in these conclusions.

Acknowledgements

This work was facilitated by a considerable amount of work from people employed by Spire. In particular the author has had very helpful interactions with Dallas Masters, Vladimir Irisov and Timothy Duly. However, it should be noted that this work has been conducted independently by the Met Office, and the conclusions in this paper have not been directly influenced by Spire employees in any way. Sean Healy (ECMWF) provided some very useful feedback on the results and their interpretation.

References

- Anthes RA. 2011. Exploring Earth's atmosphere with radio occultation: contributions to weather, climate and space weather. *Atmospheric Measurement Techniques* **4**: 1077–1103. doi: [10.5194/amt-4-1077-2011](https://doi.org/10.5194/amt-4-1077-2011).
- Anthes RA, Bernhardt PA, Chen Y, Cucurull L, Dymond KF, Ector D, Healy SB, Ho SP, Hunt DC, Kuo YH, Liu H, Manning K, McCormick C, Meehan TK, Randel WJ, Rocken C, Schreiner WS, Sokolovskiy SV, Syndergaard S, Thompson DC, Trenberth KE, Wee TK, Yen NL, Zeng Z. 2008. The COSMOC/FORMOSAT-3 - Mission early results. *Bulletin of the American Meteorological Society* **89**: 313–333. doi: [10.1175/BAMS-89-3-313](https://doi.org/10.1175/BAMS-89-3-313).
- Bauer P, Thorpe A, Brunet G. 2015. The quiet revolution of numerical weather prediction. *Nature* **525**: 47–55. doi: [10.1038/nature14956](https://doi.org/10.1038/nature14956).
- Bowler NE. 2019. Revised GNSS-RO observation uncertainties in the Met Office NWP system. *submitted to QJR Meteorol Soc*.
- Bowler NE, Clayton AM, Jardak M, Jerney PM, Lorenc AC, Wlasak MA, Barker DM, Inverarity GW, Swinbank R. 2017. The effect of improved ensemble covariances on hybrid variational data assimilation. *QJR Meteorol Soc* **143**: 785–797. doi: [10.1002/qj.2964](https://doi.org/10.1002/qj.2964).
- Buontempo C, Jupp A, Rennie M. 2008. Operational NWP assimilation of GPS radio occultation data. *Atmospheric Science Letters* **9**: 129–133. doi: [10.1002/asl.173](https://doi.org/10.1002/asl.173).
- Burrows CP. 2014. Accounting for the tangent point drift in the assimilation of gpsro data at the met office. Satellite applications technical memo 14, Met Office.

- Burrows CP, Healy SB, Culverwell ID. 2014. Improving the bias characteristics of the ROPP refractivity and bending angle operators. *Atmospheric Measurement Techniques* **7**: 3445–3458. doi: [10.5194/amt-7-3445-2014](https://doi.org/10.5194/amt-7-3445-2014).
- Cardinali C. 2009. Monitoring the observation impact on the short-range forecast. *Quarterly Journal of the Royal Meteorological Society* **135**: 239–250. doi: [10.1002/qj.366](https://doi.org/10.1002/qj.366).
- Clayton AM, Lorenc AC, Barker DM. 2013. Operational implementation of a hybrid ensemble/4D-Var global data assimilation system at the Met Office. *QJR Meteorol Soc* **139**: 1445–1461. doi: [10.1002/qj.2054](https://doi.org/10.1002/qj.2054).
- Coordination Group for Meteorological Satellites. 2017. Report of the 45th Meeting of the Coordination Group for Meteorological Satellites. URL <https://www.cgms-info.org/documents/CGMS-45-FullReport.pdf>.
- Coordination Group for Meteorological Satellites. 2018. CGMS baseline: Sustained contributions to the Global Observing System. URL https://www.cgms-info.org/documents/CGMS_baseline_-_Sustained_contributions_to_the_Global_Observing_System.pdf.
- Eyre JR. 2016. Observation bias correction schemes in data assimilation systems: a theoretical study of some of their properties. *Quarterly Journal of the Royal Meteorological Society* **142**: 2284–2291. doi: [10.1002/qj.2819](https://doi.org/10.1002/qj.2819).
- Eyre JR, Weston PP. 2014. The impact of the temporal spacing of observations on analysis errors in an idealised data assimilation system. *Quarterly Journal of the Royal Meteorological Society* **140**: 1441–1452. doi: [10.1002/qj.2227](https://doi.org/10.1002/qj.2227).
- Gleisner H. 2018. Validation report: Reprocessed level 3 gridded data. Technical Report v1.2, ROM SAF, URL http://www.romsaf.org/product_documents/romsaf_vr_grd_rep.pdf.
- Harnisch F, Healy SB, Bauer P, English SJ. 2013. Scaling of GNSS Radio Occultation Impact with Observation Number Using an Ensemble of Data Assimilations. *Mon Weather Rev* **141**: 4395–4413. doi: [10.1175/MWR-D-13-00098.1](https://doi.org/10.1175/MWR-D-13-00098.1).
- Healy S, Thepaut J. 2006. Assimilation experiments with CHAMP GPS radio occultation measurements. *QJR Meteorol Soc* **132**: 605–623. doi: [10.1256/qj.04.182](https://doi.org/10.1256/qj.04.182).
- Kursinski E, Hajj G, Schofield J, Linfield R, Hardy K. 1997. Observing Earth's atmosphere with radio occultation measurements using the Global Positioning System. *Journal of Geophysical Research - Atmospheres* **102**: 23 429–23 465. doi: [10.1029/97JD01569](https://doi.org/10.1029/97JD01569).
- Masters D. 2019. Status and plans for Spire's growing commercial constellation of GNSS science cubesats. URL https://www.romsaf.org/romsaf-irowg-2019/en/open/1570201296.44bbb0d2c4dd6f896d0761e696cb5cc8.pdf/Masters__2019-IR0WG-Spire-R0-Masters-Status+and+Plans+for+Spires+Growing+Commercial+Constellation+of+GNSS+Science+CubeSats-Thu.pdf. Talk given at IROWG 2019.
- Rawlins F, Ballard SP, Bovis KJ, Clayton AM, Li D, Inverarity GW, Lorenc AC, Payne TJ. 2007. The Met Office global four-dimensional variational data assimilation scheme. *QJR Meteorol Soc* **133**: 347–362. doi: [10.1002/qj.32](https://doi.org/10.1002/qj.32).
- Rennie MP. 2010. The impact of GPS radio occultation assimilation at the Met Office. *QJR Meteorol Soc* **136**: 116–131. doi: [10.1002/qj.521](https://doi.org/10.1002/qj.521).
- Walters D, Boutle I, Brooks M, Melvin T, Stratton R, Vosper S, Wells H, Williams K, NigelWood, Allen T, Bushell A, Copsey D, Earnshaw P, Edwards J, Gross M, Hardiman S, Harris C, Heming J, Klingaman N, Levine R, Manners J, Martin G, Milton S, Mittermaier M, Morcrette C, Riddick T, Roberts M, Sanchez C, Selwood P, Stirling A, Smith C, Suri D, Tennant W, Vidale PL, JonathanWilkinson, Willett M, Woolnough S, Xavier P. 2017. The Met Office Unified Model Global Atmosphere 6.0/6.1 and JULES Global Land 6.0/6.1 configurations. *Geoscientific Model Development* **10**: 1487–1520. doi: [10.5194/gmd-10-1487-2017](https://doi.org/10.5194/gmd-10-1487-2017).

Met Office
FitzRoy Road
Exeter
Devon
EX1 3PB
United Kingdom