


The separation of noise and signal components in Doppler RADAR returns

M. Dixon and J.C. Hubbert

National Center for Atmospheric Research, Boulder, Colorado, USA, dixon@ucar.edu

25 June 2012



Mike Dixon

1. Introduction

Knowing the correct noise value in a Doppler radar return is essential for (a) computing moments with good data quality (Ivic and Torres 2010), (b) optionally censoring (i.e. setting to missing) data which contains noise only and (c) contributing to a data quality metric (Friedrich et. al 2006, Osrodka et. al 2010). The receiver noise, however well calibrated a radar may be, will drift over time. Clutter, weather and water vapor emit radiation at all wavelengths, and these emissions will add to the thermal noise. This is especially problematic at shorter wavelengths, such as Ka-band and W-band. Radar moments in noise-only regions have well-known statistical properties. For stationary radars (such as vertically pointing instruments) it is possible to compute these statistics from a single gate over time. For a scanning radar we need to consider the statistics from a number of adjacent gates, substituting variability in space for variability in time. In this paper we present a method to identify noise regions in data from a scanning radar, utilizing the known behavior of returned power and radial velocity in noise. We show that the method is applicable to radars over a range of frequencies, from S-band to Ka-band. We demonstrate that the method can robustly identify noise regions, allowing us to compute the noise on a ray-by-ray basis. The radar moments for individual beams can therefore be effectively adjusted.

The sensitivity of the radar to these emissions depends on the wavelength of the radar. Fig. 1, below, shows an example of increased noise in a PPI scan from ground clutter targets at S-band – see the yellow ellipse.




Fig. 1: S-band low level scan showing increased noise around 300 degrees, caused by ground clutter.

Similarly Fig. 2 shows increased noise at lower elevation angles for a Ka-band RHI scan.




Fig. 2: Ka-band RHI showing increased noise in the lower levels due to thermal emissions from storms and water vapor

Both of these cases come from the SPol-Ka radar during the 2011/2012 DYNAMO field project in the Maldives Islands.

Recent work on this topic by Ivic and Torres (2011) makes use of the raw I/Q time series to analyze the behavior of power vs. range. Their algorithm has a number of steps, the first of which identifies regions along a ray with almost constant power levels, since that is a property of noise regions. They then follow up with techniques designed to eliminate false detections, and finally they compute the noise using gates that have been identified as containing only noise. They verify their results by comparing with known noise power computed at gates at long range.

In this paper we take an approach which shares some of the ideas from Ivic and Torres, but which uses moments computed over a dwell rather than the raw time series. As in Ivic and Torres, we make use of the fact that power is rather constant in noise regions. In addition, we rely on the random nature of radial velocity estimates in noise regions. We draw on experience with a clutter detection algorithm that makes use of fuzzy logic to combine information from a number of features into a single metric, which we then use to identify noise. Having identified the noise gates, we compute the noise by averaging the powers at all of those gates, provided there are a sufficient number of them in a single ray (> 50). We also optionally censor the moments at noise gates to improve compression in the stored data files.

2. Noise identification method

The noise identification method makes use of the fuzzy logic approach used successfully in the Clutter Mitigation

Decision algorithm (Hubbert et. al 2009). The first step is to compute a number of so-called *feature fields*, which are designed to indicate either the presence or absence of the phenomenon under study. We then use *interest maps* to convert the feature fields into *interest fields*, each of which has a value bounded between 0 and 1, inclusively. These interest fields are then combined into a single *decision field*. A threshold is applied to this decision field, yielding either a true or false decision for the presence of the phenomenon.

We use three feature fields in this algorithm, which are designed to be independent of each other to the greatest extent possible.

To demonstrate the method, we use an S-band RHI case. Fig. 3 shows the reflectivity field, while Fig. 4 shows the signal-to-noise ratio (SNR) for this case. There are weather and clear-air echoes below about 9 km, and generally noise only above that height.




Fig. 3: S-band reflectivity for RHI example case




Fig. 4: S-band SNR for RHI example case.

The red ellipses highlight regions in which the SNR is low but there is a coherent velocity field (See Fig. 5).

3. Feature fields

3.1 Standard deviation of phase in range (PHASE_SDEV)

When some signal is present at a gate, the radial velocity has values that show ‘coherence’ in range – in other words, adjacent gates have somewhat similar values for velocity, and it is somewhat smoothly varying. In noise, however, the velocity signature appears random. Fig. 5 below shows the radial velocity for this case. The random nature of the

velocity in the noise region is readily observable. The red ellipses highlight regions in which the SNR is rather low – less than -6 dB, and yet the velocity field still shows coherence. This demonstrates how sensitive the velocity field is to low SNR coherent signals and thus velocity makes a good indicator of noise-only regions.




Fig. 5: RHI at S-band, showing the random nature of velocity in noise.

In our method, we use the Doppler phase rather than the velocity itself for detecting noise. Velocity is the phase scaled by the Nyquist velocity, and the phase varies from -180 degrees to +180 degrees. We postulate that the phase is random in regions of pure thermal noise. For a uniformly-distributed random variable, the expected value of the variance can be shown to be equal to the square of the value range divided by 12. For random phase in degrees, the expected value of the standard deviation is:

$$\sigma = \sqrt{360^2 / 12} \approx 104 \text{ deg}$$

We compute the standard deviation over a series of gates in range, typically 9. There are, of course, gradients in the phase that will contribute to its standard deviation. Therefore, we first remove the any trend by fitting a straight line and computing the residuals – i.e. the difference between the original values and the fitted line. The phase standard deviation (PHASE_SDEV) is computed using the residuals.

Fig. 6 shows the phase field, and Fig. 7 the PHASE_SDEV field.




Fig. 6: Phase in degrees




Fig. 7: PHASE_SDEV – standard deviation of phase in range, in degrees, over 9 gates

In regions with signal, PHASE_SDEV has values generally below 20, while in noise the values are much higher, above 70 or so, on the order of the theoretical value of 104 for random phase. Therefore, this shows promise as a good discriminator between signal and noise.

3.2 Standard deviation of power in range (DBM_SDEV)

As shown by Ivic and Torres (2011), in regions of noise the power does not vary significantly with range. In contrast, in echoes, the variability in power is considerable. We capture this behavior by computing the standard deviation of power, in log (dB) units, over a number of range gates, typically 9. Fig. 8 shows this feature field, which we refer to as DBM_SDEV. Low values, less than 0.5, are indicative of noise, and values greater than 0.75 indicate the presence of signal.




Fig. 8: DBM_SDEV – standard deviation of power, in dB units, over 9 gates

3.3 Mean of normalized coherent power – NCP_MEAN

The so-called NCP (normalized coherent power) is defined as the magnitude of the lag1 covariance divided by the lag0 power. It gives an indication of how ‘coherent’ the signal is, in other words how predictable the phase is from one sample to the next. NCP is unit-less and ranges from 0 to 1. In noise NCP is low, and in signal it tends to be high. NCP is inversely related to spectrum width, so it must be used carefully. If turbulence is present, or the sample volume is large (as at long range), the spectrum width will be high and NCP will be low, even if signal is present.

To improve the reliability of this feature field, we compute a mean over a number of gates, typically 9. We refer to this field as NCP_MEAN – see Fig. 9. Values below 0.15 are indicative of noise.




Fig. 9: NCP_MEAN – mean of normalized coherent power, over 9 gates

4. Interest maps

In order to combine the feature fields, we need to remap their values into normalized ‘interest’ values ranging from 0 to 1. We do this by applying an interest map, which is a simple piece-wise linear mapping function. Fig. 10 shows the details of these mappings. The maps show the transition from 0 to 1, or view versa, as the feature field values vary. On either side of the transition, the interest value is constant at either 0 or 1.

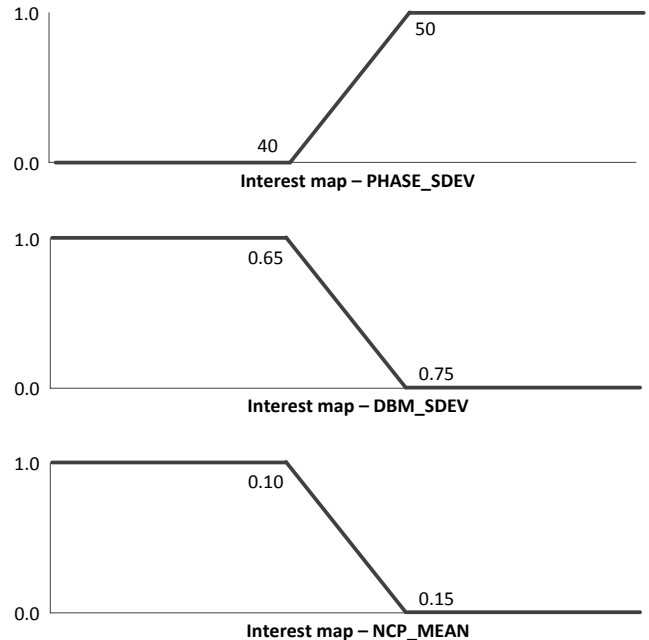


Fig. 10: Interest maps for converting feature fields to interest fields

5. Finding the noise locations

The next step is to combine the interest fields, as follows:

- for each of the three feature fields, compute the interest by applying the interest mapping translation

- multiply each interest value by a suitable weight
- compute the sum of the weighted interest values, divided by the sum of the weights. This yields a normalized combined interest value.

We then apply a threshold to the combined interest value, in order to deduce whether signal is present or if there is only noise.

Table 1 lists the he weights and thresholds we applied for the S-band and Ka-band cases.

	S-band	Ka-band
Weight for PHASE_SDEV	1.0	1.0
Weight for DBM_SDEV	1.0	1.0
Weight for NCP_MEAN	0.65	1.0
Interest threshold for noise	0.65	0.9

Table 1: interest map weights and thresholds

Fig. 11 shows the result for the S-band RHI case – below 9 km we have mostly signal, and above 9 km mostly noise.

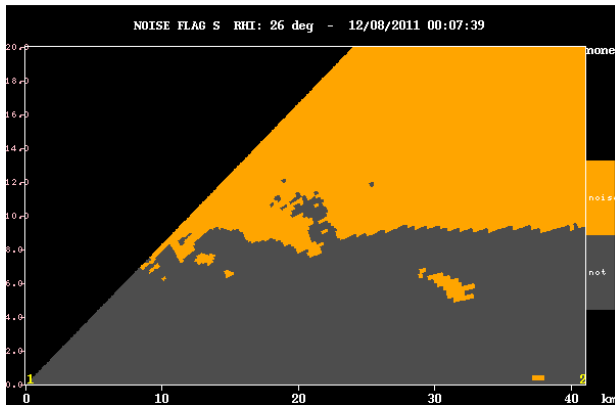


Fig. 11: flag indicating the presence of noise only

The advantage of using a number of feature fields is robustness. Errors which occur in a single field are mitigated by information from the other fields. Specifically, it is helpful to have at least 3 fields, as we do in this case. Using only 2 fields can be problematic because we then need to decide which of the 2 fields should take precedence in the event of a tie. Having a third field allows us to break the tie.

6. Ka-band case

In general, the weights and thresholds applied should be very similar for radars of different frequencies, and would follow those used for the S-band. In this particular Ka-band case, the radar had problems with phase locking, making the PHASE_SDEV field less reliable that it ideally would be. Therefore, we decided to weight the fields equally, and set a high threshold for determining the presence of noise only.

The following figures show the result of applying the method to data from a Ka-band radar in PPI mode at 0.5 degrees elevation.




Fig. 12: Ka-band DBZ, 0.5 degrees elevation




Fig. 13: Ka-band PHASE_SDEV




Fig. 14: Ka-band DBM_SDEV




Fig. 15: Ka-band NCP_MEAN

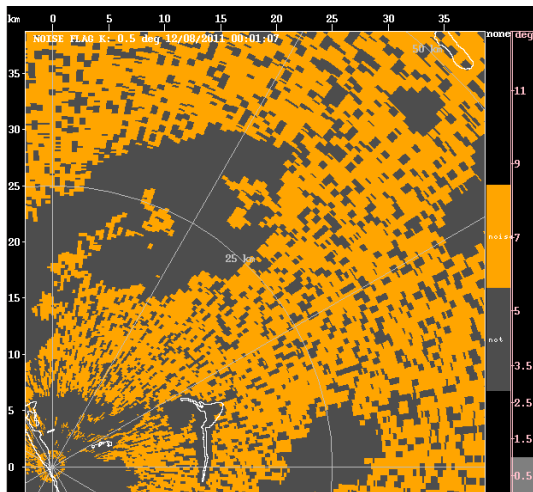


Fig. 16: Ka-band noise flag

7. Applying the method for noise estimation and censoring

7.1 Computing the noise power for each ray

After identifying the gates at which only noise is present, the noise for each ray can be determined by computing the mean power for the gates that have only noise. To do so requires, of course, that a certain minimum number of gates – say 50 – have noise only, so that a good estimate of the mean noise power can be computed. If too few gates are available, the mean noise estimate will be unreliable.

There are a number of approaches to handling the case in which the number of noise gates is too low: (a) we can use the calibrated noise value; (b) we can use a running estimate of the noise from previous rays; or (c) we can keep an estimate of mean noise for rays by azimuth and elevation, and use the value for the azimuth and elevation closest to the actual ray.

We implemented method (c), by creating a table of elevation and azimuth and keeping a running estimate of the noise for each element in that table. If the number of gates available for computing noise at a gate fall below say, 50, we use the estimated value from the table instead.

7.2 Censoring using noise gates

Censoring data at gates that contain only noise is a useful technique for making the ‘good’ data easier to visualize, and for compressing the data to keep the file sizes small. Once we have identified the gates at which only noise is present, we can simply censor the data fields at those gates, setting the value to ‘missing’.

7.3 Results

The figures below show the result of applying this method to a low-level PPI at S-band. Figures 17 and 18 show the reflectivity and phase for a 0.5 degree PPI, with no censoring applied.




Fig. 17: S-band DBZ PPI at 0.5 degrees. No censoring is applied.

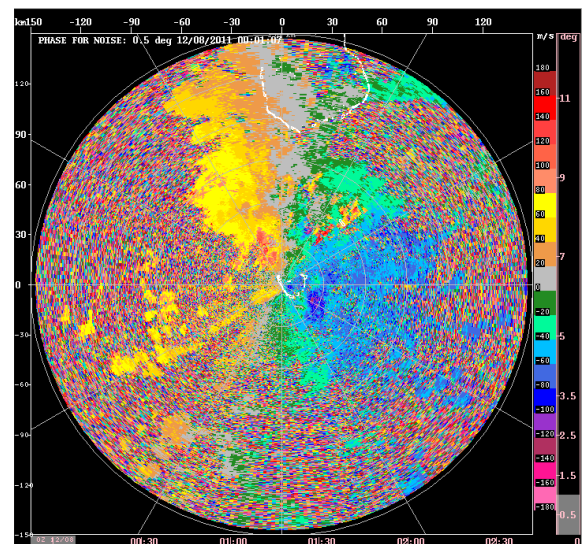


Fig. 18: S-band phase. No censoring is applied.

Fig. 19 shows the gates that are flagged by this algorithm as having noise only. Fig. 20 shows the ‘noise-bias’ field, computed as the ray-estimated mean noise minus the calibrated noise value. This shows an increase in mean noise in the NW quadrant due to ground clutter. Figures 21 and 22 show the reflectivity and velocity fields, after censoring of the noise-only gates.

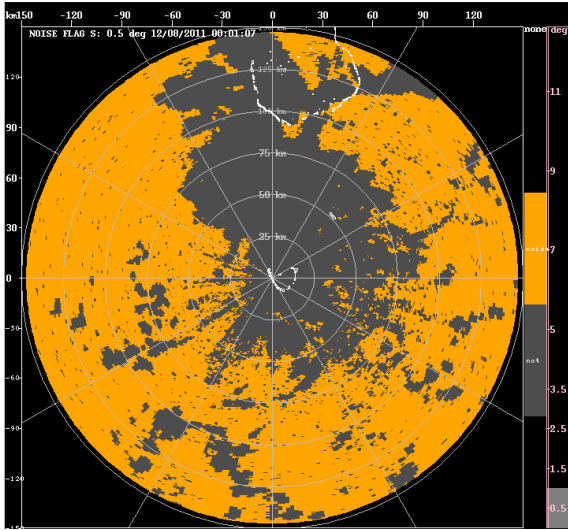


Fig. 19: Noise flag for S-band 0.5 degree PPI




Fig. 22: S-band velocity, at 0.5 degrees, after censoring

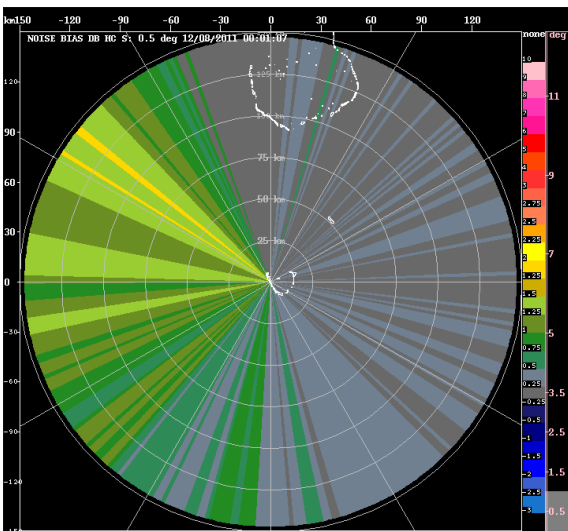


Fig. 20: Noise bias for 0.5 degree PPI. This is the mean noise, estimated for each ray, minus the calibrated noise value.

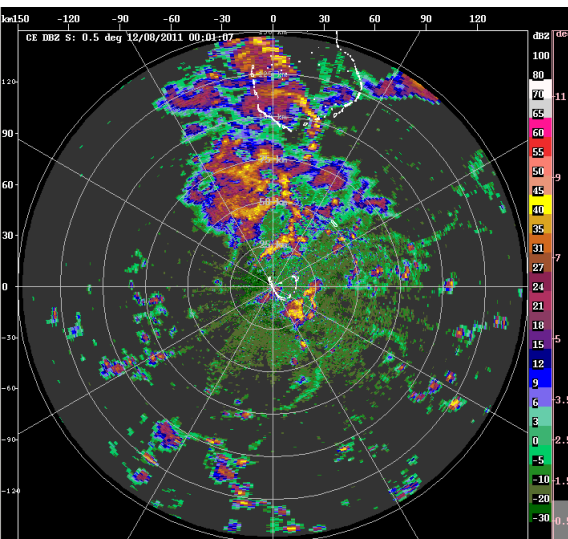


Fig. 21: S-band DBZ at 0.5 degrees, after censoring

8. Summary and conclusions

We augment previously published techniques to develop a method for identifying radar range gates at which only noise is present. The method makes use of feature fields and fuzzy logic to combine the information from a number of fields into a single decision field. Accurate estimates of the thermal background noise can significantly improve low SNR radar estimates.

The method has proved effective in identifying noise in both S-band and Ka-band radars.

Acknowledgments

The National Center for Atmospheric Research is sponsored by the National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Friedrich K., Hagen M. and Einfalt T., 2006: A quality control concept for radar reflectivity, polarimetric parameters, and Doppler velocity. *Jtech* volume 23, July 2006, 865 - 887.

Hubbert J. C., Dixon M and Ellis S. M., 2009: Weather Radar Ground Clutter. Part II. Real-time Identification and Filtering. *6th Journal of Atmospheric and Oceanic Technology*, 26, 1181 – 1197.

Ivic I. R., Torres S. M., 2010: Online determination of noise level in weather radars. *6th European Conference on RADAR in Meteorology and Hydrology*.30 - 36.

Ivic I. R., Torres S. M., 2011: Online determination of noise level in weather radars. *27th Conference on Interactive Information Processing Systems (IIPS), AMS, January 2011*, 369 – 376.

Osrodka K., Szturc J., Jurczyk A., Michelson D., Haase G., Peura M., 2010: Data quality in the BALTRAD processing chain. *ERAD 2010 - The sixth European conference on radar in meteorology and hydrology*.