

A time-series method to identify and correct range sidelobes in meteorological radar data

Article

Published Version

Westbrook, C. D. ORCID: <https://orcid.org/0000-0002-2889-8815> and Nicol, J. C. (2013) A time-series method to identify and correct range sidelobes in meteorological radar data. *Journal of Atmospheric and Oceanic Technology*, 30 (10). pp. 2417-2424. ISSN 1520-0426 doi: 10.1175/JTECH-D-13-00063.1 Available at <https://centaur.reading.ac.uk/34921/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1175/JTECH-D-13-00063.1>

Publisher: American Meteorological Society

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

A Time-Series Method to Identify and Correct Range Sidelobes in Meteorological Radar Data

C. D. WESTBROOK

Department of Meteorology, University of Reading, Reading, United Kingdom

J. C. NICOL

National Centre for Atmospheric Science, Department of Meteorology, University of Reading, Reading, United Kingdom

(Manuscript received 27 February 2013, in final form 4 July 2013)

ABSTRACT

The use of pulse compression techniques to improve the sensitivity of meteorological radars has become increasingly common in recent years. An unavoidable side effect of such techniques is the formation of “range sidelobes,” leading to the spreading of information across several range gates. These artifacts are particularly troublesome in regions where there is a sharp gradient in the power backscattered to the antenna as a function of range.

In this article a simple method for identifying and correcting range sidelobe artifacts is presented. The method makes use of the fact that meteorological targets produce an echo that fluctuates at random, and that this echo, like a fingerprint, is unique to each range gate. By cross correlating the echo time series from pairs of gates, therefore, information that has spread from one gate into another can be identified, and hence regions of contamination can be flagged. In addition it is shown that the correlation coefficients contain quantitative information about the fraction of power leaked from one range gate to another, and a simple algorithm to correct the corrupted reflectivity profile is proposed.

1. Introduction

Pulse compression is a popular method to increase the sensitivity and/or range resolution of meteorological radars. In recent years pulse compression has found extensive use in millimeter cloud radars (e.g., Kollias et al. 2007; Moran et al. 1998) in order to detect the weak echoes associated with thin stratocumulus or cirrus clouds, where the particles are small. Pulse compression is commonly used by mesosphere–stratosphere–troposphere wind profilers (e.g., Hooper et al. 2008), and it has also been proposed as a means to increase the number of independent samples for rain radars with short dwell times (Mudukutore et al. 1998).

A conventional pulsed radar has a sensitivity that is limited by the peak power output of the transmitter multiplied by the length of the pulse (in addition to other

factors, such as antenna size, dwell time, etc.). Longer pulses lead to higher sensitivity, but also to poorer range resolution. Pulse compression attempts to improve the sensitivity of a radar while maintaining high range resolution. This is achieved by transmitting a long pulse that has extra information encoded into it on time scales corresponding to the desired range resolution. The echo is then decoded using a matched filter. The encoded information takes the form of either phase or frequency modulation—see Farnett and Stevens (1990) for a review of the various implementations that are possible.

A side effect of pulse compression is the formation of range sidelobes, where echoes from a given range leak into neighboring range gates. In essence this occurs where there is not enough information encoded into the transmitted pulse to uniquely decode the reflected long pulse into the desired short-range resolution, or where that encoded information has been corrupted by the motion of the particles being probed.

Because these sidelobes are typically much smaller in magnitude than the echo from which they originate, in regions where the power backscattered to the antenna is

Corresponding author address: Chris Westbrook, Department of Meteorology, University of Reading, P.O. Box 243, Earley Gate, Reading RG6 6BB, United Kingdom.
E-mail: c.d.westbrook@reading.ac.uk

quite uniform with range, the sidelobes have little effect on quantities of interest, such as the radar reflectivity or Doppler velocity spectrum. However, where there are sharp gradients as a function of range, as is common in clouds and precipitation, the influence of sidelobes can be significant and problematic for quantitative interpretation of the measurements. Such artifacts are particularly important for dual-wavelength techniques, where even small biases can lead to large retrieval errors (Hogan et al. 2005).

At present there is no objective method to identify range sidelobe artifacts, or to correct the reflectivity data for their influence. Some radars interleave short uncompressed pulses between the long compressed ones, allowing a cross check at some gates; however, since the aim of pulse compression is to detect echoes that cannot be detected using an uncompressed pulse at the same range resolution, artifacts cannot be diagnosed this way in many cases. Empirical methods of flagging affected data, based on identification of situations likely to give rise to artifacts (e.g., where there are sharp gradients in backscatter with range), have been developed (e.g., Moran et al. 1998); however, it is highly desirable to develop more rigorous techniques.

In this article we develop a new idea to identify and correct for the presence of range sidelobes. We show that the cross correlation between echo time series sampled at pairs of range gates contains information on the occurrence of sidelobes. This is first used to visualize the locations of sidelobe artifacts in a drizzling stratus cloud profile, and then to flag the gates corrupted by sidelobes objectively. Next, we show that there is a quantitative link between the correlation coefficients and the leakage of power from one range gate to another, and we use this fact to develop a tentative algorithm for correcting a reflectivity profile corrupted by range sidelobes. We test these ideas on a second example (a midlevel ice cloud) and conclude with a brief discussion and directions for future work.

2. Method

The essential idea that we will utilize in this study is that meteorological targets produce an echo that fluctuates from pulse to pulse (Marshall and Hitschfeld 1953; Wallace 1953). This fluctuation occurs as particles move relative to one another on scales of order one-quarter of the wavelength, leading to waves that may add constructively or destructively at the radar antenna. Since this reshuffling occurs on scales that are very small compared to the range resolution of the radar (a factor $\sim 10^4$ difference in scale for the radar in section 3), this leads to the expectation that each range gate will sample

a different fluctuating echo to any other range gate (in other words, each time series is *unique*). This means that for a pair of range gates $i \neq j$, we anticipate that the correlation coefficient ρ_{ij} between the time series samples at those two gates should be zero, since the time series are uncorrelated with one another. In practice this is only true in the limit where the length of the time series is much longer than the decorrelation time of the echo, and for a finite time series this means that $|\rho_{ij}|$ will be slightly greater than zero. However, we expect that it should be rather small for many radar configurations, and this expectation is verified observationally in sections 3 and 4.

Where range sidelobes are present, we should expect $|\rho_{i \neq j}|$ to be significantly greater than zero. This is because information has leaked from a range gate where the echo is strong to another range gate where the echo is weak, giving rise to a correlation between the two gates. To detect range sidelobe artifacts, therefore, we simply calculate $|\rho_{ij}|$ and look for values significantly above zero.

3. Example 1: Drizzling stratus

We illustrate the method outlined above using data from the National Centre for Atmospheric Science (NCAS) 35-GHz cloud radar at the Chilbolton Observatory in the United Kingdom. We are interested in the general method rather than the specifics of the pulse compression scheme used, and so we describe the details of the compression in brief only.

Two complementary 10-bit binary phase codes are transmitted in sequence. Complementary codes have the advantageous property that the range sidelobes produced by each code theoretically cancel each other out when the decoded signals are summed. In practice (and in the example shown here) this is not always achieved because the assumption that the echo is unchanged between the first and second coded pulse is frequently violated by meteorological scatterers moving relative to the antenna, and relative to one another, by a significant fraction of the Nyquist velocity (Wakasugi and Fukao 1985); hence, range sidelobes are produced. The length of the coded pulses is $4 \mu\text{s}$, leading to a compressed range resolution of 60 m for our 10-bit code. Coded pulses are transmitted every 0.2 ms. The data are oversampled in range at intervals of 30 m.

As well as long coded pulses, short $0.4\text{-}\mu\text{s}$ pulses with no compression are interleaved between them. Again, the range resolution is 60 m oversampled to 30 m, and the pulses are transmitted every 0.2 ms. The coded measurements are approximately 13 dB more sensitive than the uncoded pulse data as a consequence of the longer pulse, coherent averaging of the two complementary




FIG. 1. Radar reflectivity measured through a drizzling stratus cloud over a 2-h period while the radar was dwelling at vertical. (a) Data collected using a simple uncoded pulse. (b) Data collected using a 10-bit complementary-pair coded pulse (see text). Note the weak echoes above ≈ 2000 m in (b) that are absent from (a): these are suspected to be range sidelobes. There is a small dead time after transmission of the pulse during which no echoes are received: for the uncoded data in (a), this blind zone is 200 m; for the coded data in (b), this is augmented by the length of the long coded pulse leading to a larger blind zone of 800 m. The dashed line indicates the vertical profile that will be investigated in more detail in sections 3a, 3b, and 4.

coded pulses, and the oversampling in range by the receiver.

On 28 December 2012, a layer of drizzling stratus cloud was observed over much of southern England. Figure 1 shows a 2-h period of radar reflectivity data measured using the radar at the Chilbolton Observatory while dwelling at vertical. Figure 1a shows the reflectivity measured using the uncoded pulses, while Fig. 1b shows the results measured using the coded (i.e., compressed) pulses. While both methods yield similar reflectivity fields at lower levels, there are weak echoes in the coded data above ≈ 2 km that are not present in the uncoded data. This behavior has been observed quite frequently in drizzling boundary layer clouds, and we strongly suspect that these are range sidelobes resulting from the pulse compression.

In what follows we will test this hypothesis objectively using a 0.5-s sample of echo time series data collected at 2200 UTC (indicated by the dashed line in Fig. 1). This corresponds to 2048 coded and uncoded pulses (note, however, since pairs of coded pulses are combined, this leads to a coded pulse time series that is only 1024 points in length). Figure 2 shows the specific profile being considered in detail: here, the signal-to-noise ratio (SNR)¹ is shown as a function of range for both coded and uncoded pulses. Note the SNR drops to 0 dB for the

uncoded pulse at 1940 m; meanwhile, the coded pulse detects a significant echo up to 2360 m. Also shown for reference is the uncoded SNR with a 13-dB offset (dashed line). At ranges < 1700 m the profiles are almost identical. Between 1700 and 1940 m, however, the gradient of the SNR with range is steep (~ 6 dB per 100 m) and here the coded pulse overestimates the reflectivity relative to the uncoded data: at 1820 m this difference is 1 dB, while at 1880 m the difference is 4 dB, suggesting that range sidelobes may be affecting these higher-altitude gates.

a. Correlations between pairs of time series

The correlation coefficient ρ_{ij} between the complex samples at each possible pair of range gates was computed for both coded and uncoded pulses. Note that in an operational algorithm, one would only need to compare gates within one code length from each other—here, we show all possible correlations for completeness. Figure 3 shows the magnitude of the correlation coefficient (grayscale) as a function of range on the ordinate and abscissa. Note that the diagram is symmetrical, since $|\rho_{ij}| = \rho_{ji}$. In Fig. 3a, the uncoded echo time series are, as expected, uncorrelated with one another and values of $|\rho_{ij}|$ are small (mean value of $|\rho| = 0.05$ in regions of $\text{SNR} > 5$ dB, i.e., where the sample is dominated by meteorological echoes rather than noise). The exception to this rule is the diagonal elements of the figure, where $i = j$ and one is simply correlating the time series at that gate with itself ($|\rho_{ii}| = 1$). Note also that because of the oversampling to 30-m bins, there is some correlation between neighboring bins where a detectable

¹ We define $\text{SNR} = (P - \mu)/\sigma$, where P is the received power, μ is the mean noise power sampled in empty (noise dominated) range gates, and σ is the standard deviation of that noise from gate to gate.




FIG. 2. (a) The vertical profile of the SNR through the cloud. Squares are for the complementary-pair coded pulse (see text). Circles show the same profile measured using a simple uncoded pulse. The dashed line is simply the uncoded profile shifted by 13 dB to aid comparison with the coded data (which is 13 dB more sensitive). (b) The sidelobe flag derived from the simple identification algorithm described in the text. A value of one indicates significant range sidelobe artifacts are likely at that range gate.

meteorological echo is present (800–1900 m in the figure). At ranges where noise is dominating the echo (>1900 m), this correlation between neighboring gates disappears, and $|\rho|$ is very close to zero, demonstrating that the receiver noise is uncorrelated from gate to gate.

Figure 3b shows the same correlation coefficients for the coded pulses. The influence of the range sidelobes is immediately obvious in this figure, with correlation coefficients significantly higher than zero present in the off-diagonal elements (peak values of $|\rho_{ij}|$ in this case were ≈ 0.7). The regions of correlation are very clearly defined, running parallel to the diagonal but offset from it by ≈ 100 and ≈ 400 – 600 m (indicated with arrows in one-half of the figure). This is the behavior that we were led to anticipate from range sidelobe artifacts in the discussion in section 2. We therefore identify these areas of correlation as range sidelobes produced by the strong drizzle echoes between the 1400- and 1800-m range masking the much weaker (or absent) echoes above. These well-defined features demonstrate that the correlation coefficients can provide a robust indicator for the presence of sidelobes, which can be exploited to identify affected gates.

b. A simple flag for sidelobe artifacts

Having graphically identified those pairs of gates where information has leaked from one gate to another, we suggest a rudimentary algorithm to flag corrupted data. The implementation of a fully developed operational algorithm to flag and remove range sidelobe artifacts is beyond the scope of this initial study. However, we have experimented with some simple approaches. The most obvious idea is simply to flag any range gate where there is a significant correlation with another gate

separated from it by less than the length of the coded pulse. However, in the example presented here, this would remove much of the data at ranges 1400–1800 m, even though the agreement between the coded and uncoded SNR profiles here is very close.

A simple refinement is to identify “which way” the information is likely to be flowing. Specifically, we assume that weak signals (or background noise) are masked by sidelobes from much stronger signals, and not vice versa. Given a correlation coefficient between a pair of gates above some threshold value $|\rho|_{\text{crit}}$, one then seeks the weaker SNR among the two gates, and flags those data as affected, while leaving the stronger signal unflagged. The result of this simple algorithm is shown in Fig. 2b alongside the SNR profile. Here, we have chosen $|\rho|_{\text{crit}} = 0.25$, although the results were not found to be sensitive to this choice of threshold, and identical flagging was obtained using $|\rho|_{\text{crit}} = 0.15$ and $|\rho|_{\text{crit}} = 0.35$. Data at ranges between 1850 and 2400 m have been flagged as corrupt, while the data at lower ranges have not. This is an encouraging result. The algorithm has removed the spurious echoes above 1900 m, where there was in fact almost certainly no cloud present. It has also identified a few pixels near cloud top, where the gradient in the SNR is steep and where we had noted discrepancies between coded and uncoded results. This gives us optimism that this relatively simple algorithm can be used to identify sidelobe artifacts.

4. Correcting the corrupted reflectivity profile

To correct the corrupted reflectivity profile for the influence of range sidelobes, a quantitative link is




FIG. 3. Correlation coefficients computed for complex time series at each pair of range gates between 800 and 2600 m. The absolute magnitude of the correlation coefficient is shown here (grayscale). (a) Results for a simple uncoded pulse. (b) Results for a 10-bit complementary-pair coded pulse (see text). Note the off-diagonal regions of high correlation in (b): these are correlations arising from range sidelobes. The diagrams are symmetrical since $|\rho_{ij}| = |\rho_{ji}|$.

needed between the correlation coefficients measured in section 3a and the amount of power being leaked from a range gate with a strong echo to another range gate with a weak echo. Consider the complex samples V of the echo measured at range gate i :

$$V_i^m = V_i^t + \sum_{j=i-CL}^{i+CL} f_{ji} V_j^t, \quad (1)$$

where V_i^t is the true echo at range gate i and V_i^m is the (potentially corrupted) measurement. The second term on the right-hand side represents the contributions from any sidelobes originating in other range gates within one

code length CL of gate i (for our 35-GHz radar, CL = 10). The factors f_{ji} represent the sidelobes themselves; specifically, they are the (possibly complex) fraction of the true echo at gate j that has leaked into the measured echo at gate i . We assume in this analysis that $|f_{ji}| \ll 1$.

The cross-correlation coefficient between a pair of measured time series at gates k and i is

$$\rho_{ki} = \frac{\sum_{\text{time}} V_k^m (V_i^m)^*}{\left(\sum_{\text{time}} |V_k^m|^2 \sum_{\text{time}} |V_i^m|^2 \right)^{0.5}}. \quad (2)$$

Substituting Eq. (1) into the numerator, expanding out, and neglecting terms of order f^2 , one obtains three terms remaining in the numerator N :

$$N = \sum_{\text{time}} V_k^t (V_i^t)^* + \sum_{\text{time}} \left[(V_i^t)^* \sum_{j=k-CL}^{k+CL} f_{jk} V_j^t + V_k^t \sum_{j=i-CL}^{i+CL} (f_{ji} V_j^t)^* \right]. \quad (3)$$

The first term approaches zero for a sufficiently long time series, since V_k^t and V_i^t are uncorrelated. All of the terms in the first sum over j are approximately equal to zero for the same reason, except when $j = i$. Finally, all of the terms in the second sum over j are equal to zero, except when $j = k$. This yields the following result:

$$\rho_{ki} = \frac{\sum_{\text{time}} f_{ik} |V_i^t|^2 + f_{ki}^* |V_k^t|^2}{\left(\sum_{\text{time}} |V_k^m|^2 \sum_{\text{time}} |V_i^m|^2 \right)^{0.5}}. \quad (4)$$

Now we recognize that since $|f| \ll 1$, we need only concern ourselves with the scenario where one gate k contains a strong echo and the second gate i contains a much weaker echo (i.e., $|V_k^t|^2 \gg |V_i^t|^2$ and $|V_k^t|^2 \approx |V_k^m|^2$). We may then simplify further to obtain the magnitude of the correlation coefficient (as measured in Fig. 3):

$$|\rho_{ki}| = \left| \frac{\sum_{\text{time}} f_{ki} |V_k^m|^2}{\left(\sum_{\text{time}} |V_k^m|^2 \sum_{\text{time}} |V_i^m|^2 \right)^{0.5}} \right| \quad (5)$$

or equivalently

$$|\rho_{ki}|^2 = |f_{ki}|^2 \times \left(\frac{\sum_{\text{time}} |V_k^m|^2}{\sum_{\text{time}} |V_i^m|^2} \right), \quad (6)$$

where

$$\langle f_{ki} \rangle = \frac{\sum_{\text{time}} f_{ki} |V_k^t|^2}{\sum_{\text{time}} |V_k^t|^2} \quad (7)$$

is the time-averaged leakage factor, weighted by the power of the echo at gate k .

Equation (6) reveals how the correlation coefficient is quantitatively related to the sidelobe leakage from one gate to another. On the basis of this, we now propose a simple algorithm to determine these time-averaged leakage factors $|f_{ki}|$ and we use these to correct the profile of SNR (and hence the radar reflectivity). Note that it is the power measured by the receiver that is to be corrected: range-weighted quantities such as radar reflectivity should only be computed after the correction is complete.

The algorithm is extremely simple. At each gate i , one identifies other range gates k within one code length that have a signal-to-noise ratio significantly higher than that at i (here, we use 3 dB as a threshold). If we take the amount of power leaked from gate k to gate i to be $|\langle f_{ki} \rangle|^2 \times \text{SNR}_k$, then Eq. (6) shows that this is equal to $|\rho_{ki}|^2 \times \text{SNR}_i$. In other words, $|\rho_{ki}|^2$ represents the fraction of the echo power at gate i that is introduced from the stronger return at gate k . Given this, we can use our measured correlation coefficient $|\rho_{ki}|$ to compute the contribution to the SNR at gate i that is due to leakage from gate k . We subtract this number from the corrupted SNR value at gate i for all k within one code length. This provides us with a corrected SNR value for that gate. We then move on and repeat the process for all other range gates, yielding a complete corrected SNR profile.

Figure 2a shows the result of this correction process for our drizzling stratus example [solid line with the asterisks (*)]. The spurious echoes above 2 km have now been completely removed. In addition the reflectivity profile at the top of the cloud (≈ 1.8 –2-km range) is now in close agreement with the uncoded profile, showing a sharp decrease in reflectivity with range. This result is extremely encouraging, giving us a profile that matches the uncoded data and provides a correct cloud-top height. This indicates that the procedure outlined above is a sensible one, and that the analysis above is a valid approximation to the problem.

One source of error in this approach is that it is likely to slightly overestimate the amount of power that should be subtracted from the corrupted gates. This is because even in the absence of sidelobes, $|\rho_{ki}| > 0$ because of the finite decorrelation time of the echo relative to the total length of the time series. This “noise” in $|\rho_{ki}|^2$ can be

significant, and is additive when sidelobes are present, leading to excess power being subtracted from a corrupted gate. In the gates above 2 km in our example above (where we should have subtracted 100% of the signal), our algorithm actually subtracted between $\approx 100\%$ and 130% in many of the gates. In the figure we have set these negative points to a value of $\text{SNR} = 0$ dB. This issue can be ameliorated by using a longer time series, which will yield a more accurate correction since the noise in $|\rho_{ki}|^2$ will be reduced.

A second possible source of error is the form of the retrieved correction factor itself. We are estimating the leakage factor f_{ki} for the complex samples, weighted by the power in gate k on each pulse, and averaged over the time series. However, the leakage of *power* from one gate to another is proportional to $|f_{ki}|^2$, and therefore it is in fact the quantity

$$\langle |f_{ki}|^2 \rangle = \frac{\sum_{\text{time}} |f_{ki}|^2 |V_k^t|^2}{\sum_{\text{time}} |V_k^t|^2}, \quad (8)$$

which we seek in order to correct the reflectivity profile. In the correction procedure above we have approximated this by $|\langle f_{ki} \rangle|^2$. This is reasonable if f_{ki} does not vary greatly over the time series. The agreement between the corrected profile and the uncoded profile in Fig. 2a is evidence that this approximation is indeed an acceptable one, at least for our radar setup.

5. Example 2: The base of a thick midlevel ice cloud

To illustrate the idea further, we now briefly present a second example profile, this time collected in a deep midlevel ice cloud sampled at 1500 UTC 23 December 2012. The experimental setup is identical to section 3. Figure 4a shows the uncoded SNR profile: note the sharp gradient close to the cloud base between 1.7- and 2-km range, as the ice particles fall into dry air below and evaporate. The coded profile is consistent with the uncoded measurements at ranges above 1.75 km; however, between ≈ 1.65 and 1.75 km the coded and uncoded measurements disagree, and between 1.45 and 1.65 km (where the uncoded signal is dominated by noise) the coded profile contains a significant echo, which we suspect to be the result of range sidelobes.

Figure 4b shows the correlation values for all pairs of range gates between 1 and 2.5 km using the uncoded time series. As in the stratus case, we observe $|\rho_{i \neq j}| \approx 0$. Figure 4c shows the same for the coded time series—this time two lines of significant correlation ($|\rho_{i \neq j}| \approx 0.5$) are present, and they show that the time series at range gates between 1.45 and 1.75 are highly correlated with the




FIG. 4. (a) The signal-to-noise profile at the base of a deep ice cloud using coded and uncoded pulses. Also shown is the uncoded profile with a 13-dB offset, and a coded profile that has been corrected using the algorithm described in the text. The magnitude of the correlation coefficient (grayscale) for pairs of range gates for (b) uncoded and (c) coded pulses. (d) The result of the algorithm described in section 3b identifying sidelobe-corrupted data (value of 1) and unaffected data (value of 0).

time series sampled approximately 400 and 550 m higher up the profile.

Using the same algorithm described in section 3b, the areas of likely contamination were automatically flagged, and this is shown in Fig. 4d. The algorithm has diagnosed all gates between 1450 and 1650 m as corrupt, and this is consistent with the region where we suspected that no cloud was present. It has also flagged the region between 1.65 and 1.75 km, where the uncoded and coded profiles did not match.

Finally, we have corrected the coded SNR profile using the procedure in section 4. The corrected profile is shown in Fig. 4a. As in the stratus case, the corrected profile shows very close agreement with the uncoded profile, and it has again removed all of the signals in gates where we suspected there was no cloud, leading to a correct cloud-base height. This time between 90% and 110% of the echo was removed in gates where we inferred that there was no cloud present. There is one

range bin at 1500 m where the corrected SNR is +7 dB, and where we believe cloud is absent. However, this single pixel would be easily removed with a simple speckle filter (used as part of our standard processing). These results encourage us further that our correction methodology is a reasonable one.

6. Conclusions and discussion

We have shown how pulse-to-pulse fluctuations can be used to diagnose the presence of range sidelobe artifacts associated with pulse compression, and have illustrated the approach using data collected from a drizzling stratus cloud and a thick midlevel ice cloud with the 35-GHz cloud radar at the Chilbolton Observatory. We have suggested a simple algorithm to flag the sidelobe-affected data. We have also shown how the correlation coefficients are quantitatively linked to the amount of the power in a particular range gate that

originates as the result of a sidelobe from a second range gate, and hence we have formulated a simple correction algorithm that appears to perform well on the two examples presented here.

While the example data presented here were collected using complementary phase-coded pulses, the methodology is quite general and ought to be applicable to other pulse compression techniques, provided that a decoded time series of echo amplitudes is available that can be cross correlated. We emphasize that although we have shown simultaneous uncoded measurements in Figs. 1–4 for the purposes of comparison, it is only the compressed pulse data that are used in the flagging and correction algorithms.

One outstanding issue is the small spurious correlations that are produced when the decorrelation of the signal is slow relative to the length of time series (i.e., the number of independent samples is small). This makes it more difficult to identify some of the less prominent sidelobe artifacts, and leads to an overcorrection of the reflectivity profile, which we would like to minimize, or at least quantify. The magnitude of these correlations will be different depending on the meteorological situation. In cirrus clouds where the Doppler spectrum is very narrow for a vertically pointing radar and the decorrelation of the echo is concurrently slow, we have measured the mean and standard deviation of $|\rho_{i \neq j}|$ of ≈ 0.1 using uncoded pulses. This leads to a few pairs of range gates where the correlation is as large as 0.3 despite the absence of sidelobes. The most logical way to improve this situation is to make use of the longest possible time series, taking into account the desired time resolution of the flagged/corrected data. We are currently studying the possibility of setting a variable $|\rho_{ij}|$ threshold in our flagging algorithm that is diagnosed based on the decorrelation time (or Doppler spectrum width) of the samples in question.

A second outstanding issue is the neglect of terms of order f^2 in the derivation of Eq. (3). The neglected terms are of the form $f_{ik}f_{li}^*|V_l^t|^2$. Physically, these terms correspond to a correlation between gates k and i being introduced via sidelobes from a third gate l that is leaking power into both k and i . All three gates must lie within one code length of each other. These terms become significant in the scenario where the meteorological echoes at both k and i are comparable and weak relative to the echo at gate l . This does not introduce any difficulties in the flagging algorithm, but it does emphasize the need to only apply the correction algorithm to gate pairs where $|V_k^t|^2 \gg |V_l^t|^2$ (in the profiles shown, a 3-dB threshold was used, but depending on the magnitude of f , a higher threshold might prove more robust). We plan to investigate the role of these higher-order correlations

theoretically via computer simulation (see future work below).

Another source of range sidelobes (and hence correlation) between range gates may occur in the absence of pulse compression, simply as a result of the finite bandwidth of the receiver and finite duration of the transmitted pulse, leading to a slight spreading of information across neighboring range gates (e.g., Doviak and Zrnić 1984; Nicol and Illingworth 2013). However, in section 3a we noted that where the echo is dominated by noise, the time series were uncorrelated from one range bin to the next. This is evidence that the receiver filter effect is not a significant source of correlation, at least for the 35-GHz radar used here.

Unlike the reflectivity profile, we have not attempted to correct the Doppler information. This is because Doppler velocity depends on the change in phase between pairs of pulses, and hence a correction of the echoes on a pulse-by-pulse basis would be required. The correlation coefficient, on the other hand, only describes a time-averaged leakage factor $\langle f_{ki} \rangle$. In section 4 we suggested that f_{ki} might be approximately constant in time ($f_{ki} \approx \langle f_{ki} \rangle$). However, the Doppler information is much more sensitive to the accuracy of this approximation, since any overcorrection or undercorrection on individual pulse amplitudes will effectively introduce spurious power into the Doppler spectrum. More analysis is needed to investigate this aspect.

So far we have only applied our method to a vertically pointing cloud radar. It is interesting to consider whether the same approach could be applied to scanning radars. For scanning cloud radars, or research radars that can scan slowly, it should be possible to apply the present methodology. However, for operational weather radars the dwell time per ray is usually very short, while the decorrelation of the echo is relatively slow because of the longer wavelength: for example, Illingworth (2004) suggests that one might expect only around 35 independent samples when scanning in rainfall at low elevations. The vertical dwells analyzed in sections 3 and 5 had a spectral width of $\approx 0.3 \text{ m s}^{-1}$, corresponding to 120 independent samples. According to Fisher (1921), the error on a correlation coefficient estimated from n independent samples is approximately proportional to $1/\sqrt{n-3}$ when the populations from which those samples are drawn are uncorrelated. This means that we should expect ρ_{ij} to be a factor of ~ 2 noisier for operational weather radar data than for our examples in sections 3 and 5. As explained in section 4, range sidelobe corrections are likely to be overestimated because of this extra noise. As a result we anticipate that the identification and correction of less prominent sidelobe artifacts will be more difficult for operational weather radars.

Future work will focus on the application of the flagging and correction procedures developed here to other meteorological scenarios to assess how well they perform. We also hope to use the methodology to characterize the sidelobes of our radar setup using different code and pulse repetition frequency options in order to best minimize the formation of sidelobes for the NCAS 35-GHz radar at the Chilbolton Observatory. Idealized simulations using synthetic echoes are also being performed to forward model the expected correlation coefficients that we will then measure using the real radar system.

Acknowledgments. We thank the staff at the Chilbolton Observatory for operating and maintaining the 35-GHz radar. Some of this work has been supported by the NCAS Facility for Ground-Based Atmospheric Measurements (FGAM).

REFERENCES

- Doviak, R. J., and D. S. Zrnić, 1984: *Doppler Radar and Weather Observations*. Academic Press, 458 pp.
- Farnett, E. C., and G. H. Stevens, 1990: Pulse compression radar. *Radar Handbook*, 2nd ed. M. I. Skolnik, Ed., McGraw-Hill, 10.1–10.39.
- Fisher, R. A., 1921: On the “probable error” of a coefficient of correlation deduced from a small sample. *Metron*, **1**, 3–32.
- Hogan, R. J., N. Gaussiat, and A. J. Illingworth, 2005: Stratocumulus liquid water content from dual-wavelength radar. *J. Atmos. Oceanic Technol.*, **22**, 1207–1218.
- Hooper, D. A., J. Nash, T. Oakley, and M. Turp, 2008: Validation of a new signal processing scheme for the MST radar at Aberystwyth. *Ann. Geophys.*, **26**, 3253–3268.
- Illingworth, A. J., 2004: Improved precipitation rates and data quality by using polarimetric measurements. *Weather Radar: Principles and Advanced Applications*, P. Meischner, Ed., Springer, 130–166.
- Kollias, P., E. E. Clothiaux, M. A. Miller, B. A. Albrecht, G. L. Stephens, and T. P. Ackerman, 2007: Millimeter-wavelength radars: New frontier in atmospheric cloud and precipitation research. *Bull. Amer. Meteor. Soc.*, **88**, 1608–1624.
- Marshall, J. S., and W. Hitschfeld, 1953: Interpretation of the fluctuating echo from randomly distributed scatterers. I. *Can. J. Phys.*, **31**, 962–994.
- Moran, K. P., B. E. Martner, M. J. Post, R. A. Kropfli, D. C. Welsh, and K. B. Widener, 1998: An unattended cloud-profiling radar for use in climate research. *Bull. Amer. Meteor. Soc.*, **79**, 443–455.
- Mudukutore, A. S., V. Chandrasekar, and R. J. Keeler, 1998: Pulse compression for weather radars. *IEEE Trans. Geosci. Remote Sens.*, **36**, 125–142.
- Nicol, J. C., and A. J. Illingworth, 2013: The effect of phase-correlated returns and spatial smoothing on the accuracy of radar reflectivity retrievals. *J. Atmos. Oceanic Technol.*, **30**, 22–39.
- Wakasugi, K., and S. Fukao, 1985: Sidelobe properties of a complementary code used in MST radar observations. *IEEE Trans. Geosci. Remote Sens.*, **GE-23**, 57–59.
- Wallace, P. R., 1953: Interpretation of the fluctuating echo from randomly distributed scatterers. II. *Can. J. Phys.*, **31**, 962–994.