The effect of phase-correlated returns and spatial smoothing on the accuracy of radar refractivity retrievals

Article

Published Version

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-12-00077.1

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

The Effect of Phase-Correlated Returns and Spatial Smoothing on the Accuracy of Radar Refractivity Retrievals

J. C. NICOL

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

A. J. ILLINGWORTH

University of Reading, Reading, United Kingdom

ABSTRACT

Radar refractivity retrievals have the potential to accurately capture near-surface humidity fields from the phase change of ground clutter returns. In practice, phase changes are very noisy and the required smoothing will diminish large radial phase change gradients, leading to severe underestimates of large refractivity changes (ΔN). To mitigate this, the mean refractivity change over the field ($\langle \Delta N \rangle_{\text{field}}$) must be subtracted prior to smoothing. However, both observations and simulations indicate that highly correlated returns (e.g., when single targets straddle neighboring gates) result in underestimates of $\langle \Delta N \rangle_{\text{field}}$ when pulse-pair processing is used. This may contribute to reported differences of up to 30 N units between surface observations and retrievals. This effect can be avoided if $\langle \Delta N \rangle_{\text{field}}$ is estimated using a linear least squares fit to azimuthally averaged phase changes. Nevertheless, subsequent smoothing of the phase changes will still tend to diminish the all-important spatial perturbations in retrieved refractivity relative to $\langle \Delta N \rangle_{\text{field}}$; an iterative estimation approach may be required. The uncertainty in the target location within the range gate leads to additional phase noise proportional to ΔN, pulse length, and radar frequency. The use of short pulse lengths is recommended, not only to reduce this noise but to increase both the maximum detectable refractivity change and the number of suitable targets. Retrievals of refractivity fields must allow for large ΔN relative to an earlier reference field. This should be achievable for short pulses at S band, but phase noise due to target motion may prevent this at C band, while at X band even the retrieval of ΔN over shorter periods may at times be impossible.

1. Introduction

Radar refractivity retrieval promises to provide valuable insights into the dynamic variability of near-surface water vapor. Changes in the refractive index (n) of the atmosphere near the earth's surface are dominated by humidity changes, particularly during summer in temperate latitudes. Despite growing interest in their use, greater emphasis needs to be placed on quality control and understanding sources of error if the full potential of refractivity retrievals is to be achieved. Retrievals use the phase change between two different plan position indicator (PPI) radar scans from stationary targets (ground clutter). Initially described in Fabry et al. (1997), refractivity changes can be derived from differences in the measured phase change between pairs of stationary targets on the same azimuth (couplets). The mean refractivity change (ΔN), where $N = (n - 1) \times 10^6$, between two targets (A and B) along the same azimuth at ranges of r_A and r_B, respectively, is related to the measured phase changes ($\Delta \phi_A$ and $\Delta \phi_B$):

$$\Delta N = \frac{c}{4\pi f_{\text{Tx}}} 10^6 \frac{\Delta \phi_B - \Delta \phi_A}{r_B - r_A}. \quad (1)$$

Here, f_{Tx} is the transmitted frequency and c is the speed of light in a vacuum. At 20°C, $\Delta N = 1$ equates approximately to a 1% change in relative humidity. At S-, C-, and X-band wavelengths ($\sim$10, $\sim$5, and $\sim$3 cm), the rates of phase change with respect to range are approximately

Corresponding author address: Dr. John Nicol, National Centre for Atmospheric Science, Dept. of Meteorology, University of Reading, Whiteknights, P.O. Box 243, Reading RG6 6BB, United Kingdom.
E-mail: j.c.nicol@reading.ac.uk

$7°$, $13°$, and $23°$ km^{-1} when $\Delta N = 1$, respectively. Clearly, difficulties will increasingly occur at C and even more so at X band because of aliasing. This occurs when the *phase change difference* $(\Delta\phi_B - \Delta\phi_A)$ exceeds $\pm 180°$ over the distance $(r_B - r_A)$ (henceforth, this quotient is the *phase change gradient*). Phase change differences are usually calculated by "pulse-pair processing" (e.g., Skolnik 1990, p. 23.15) although changes are estimated between adjacent range gates rather than successive pulses as for Doppler velocity. Aliasing occurs when $|\Delta N|$ is greater than the maximum unambiguous refractivity change $(|\Delta N|_{\text{folding}})$, given by (2); this effect is minimized when contiguous gates are used, as Δr is some multiple of the range-gate spacing (Δr_{gate}) and is minimized when $\Delta r = \Delta r_{\text{gate}}$:

$$|\Delta N|_{\text{folding}} = \frac{c}{4 f_{\text{Tx}} \Delta r} 10^6. \qquad (2)$$

The majority of published work comparing refractivity retrievals with surface observations has been at S band, for example, with the McGill radar (Fabry et al. 1997; Fabry 2004) and with S-band dual-polarization Doppler radar (S-Pol; Weckwerth et al. 2005) during the International H$_2$O Project (IHOP). The range resolution and range-gate separation for both these radars is 150 m, hence aliasing occurs beyond $|\Delta N|_{\text{folding}} \approx 200$ (assuming no noise in the phase changes), whereas for the 250-m gate separation of the Next Generation Weather Radar (NEXRAD) Weather Surveillance Radar-1988 Doppler (WSR-88D) radars (Bodine et al. 2010, 2011), aliasing occurs beyond $|\Delta N|_{\text{folding}} \approx 120$. Since the seasonal (and often annual) range of N will often be less than ± 120, if a "reference period" can be found when N is approximately constant over the clutter domain (i.e., the ground clutter coverage at a particular elevation), then one might expect reliable retrievals to be obtained throughout the year relative to N_{ref}, the *reference refractivity field*. Aliasing occurs beyond $|\Delta N|_{\text{folding}} \approx 44$ for the operational C-band radar data used in this study.

In practice, the observed phases are very noisy, so some spatial smoothing of phase change measurements is necessary. However, smoothing with respect to range tends to reduce large phase change gradients and underestimates of ΔN can result. For this reason, $\langle\Delta N\rangle_{\text{field}}$, the *field-averaged refractivity change*, is typically subtracted from all the raw phase change observations prior to smoothing. Though as we shall see later, this often results in an underestimate of the variability of retrieved refractivity fields even when $\langle\Delta N\rangle_{\text{field}}$ is accurately estimated.

In this paper, we show that the phase of returns from neighboring range gates often displays significant correlations, indicating that the returns from the two gates are not independent. If the same target is responsible for returns at adjacent gates, then the relative phases will be constant with time, and the implied radial gradient of the phase changes with time will tend to zero. When pulse-pair processing is used on the raw phase change data, underestimates of $\langle\Delta N\rangle_{\text{field}}$ can result and refractivity retrievals will be biased toward N_{ref}. We shall see that noise in the observed phase changes increases these biases, which can be significant, in particular as $|\Delta N|$ becomes an appreciable fraction of $|\Delta N|_{\text{folding}}$. A vastly improved performance is demonstrated using an azimuthally averaged linear fit rather than the pulse-pair method in the estimation of $\langle\Delta N\rangle_{\text{field}}$.

In section 2, a discussion of the origin and magnitude of phase change noise is presented with an overview of refractivity retrievals from the literature. Section 3 explains how correlated phase measurements in ground clutter result from the finite pulse length and the filter response of the radar receiver. Data from one of the operational U.K. C-band weather radars is used in section 4 to demonstrate the correlation of phases between adjacent range gates and the resultant bias in the derived $\langle\Delta N\rangle_{\text{field}}$. In section 5, physically based simulations of ground clutter incorporating a realistic range-weighting function are used to quantify these biases in estimating $\langle\Delta N\rangle_{\text{field}}$. In section 6, we demonstrate that the smoothing applied to phase change measurements also tends to bias estimates of ΔN, using smoothing kernels from the literature as examples. The possible relevance of correlated returns and spatial smoothing for previous published refractivity observations is discussed in section 7. In practice, the linear fit approach still leads to underestimates as $\langle\Delta N\rangle_{\text{field}}$ approaches $|\Delta N|_{\text{folding}}$; so it is suggested in the final section 8 that the use of shorter pulses would enable larger values of $|\Delta N|$ to be accurately retrieved.

2. Overview of phase noise in observations and refractivity retrieval algorithms

a. Sources of random phase errors

Phase measurements of ground clutter targets at each range gate are obtained by averaging the returns over many pulses as the radar scans in azimuth. In practice, phase change measurements are very noisy. The primary sources of noise are due to target motion, target height variability, and the unknown location of the target within the range gate, and are discussed below. In general, phase change noise is inversely proportional to wavelength and is therefore more of a problem for refractivity retrievals with short-wavelength radars. Phase errors related to transmitter and local oscillator frequency drifts are

restricted to magnetrons and are not considered here. Although data from a magnetron radar are considered in this paper, we limit our analyses to times when the frequency was constant. Propagation phase shifts in rainfall are negligible except in very heavy rainfall and are not discussed.

1) TARGET MOTION

The most obvious source of phase noise is due to target displacement. For instance, a target displacement of just 2.8 mm will result in a phase error of 36° at C-band wavelengths and about 20° and 60° at S-band and X-band wavelengths, respectively. When considering scan-to-scan ΔN, the phase *change* errors will be larger by a factor $\sqrt{2}$, in contrast to comparisons with a reference period when many scans are averaged. Target motion noise can be mitigated to some extent by a judicious selection of suitable targets. Fabry (2004) evaluated a target reliability index using the temporal coherence (ratio of the lag-1 and lag-0 time correlation) from a series of scans throughout a "calibration period" during which time the refractivity near to the surface was constant. In addition, real-time observations such as signal-to-noise ratio (SNR), velocity, and spectral width have been used to derive a target quality index. The minimum of these indexes was then used to create a weighting function in the retrieval algorithm. Another approach (e.g., Nicol et al. 2012a) is, as explained in section 3b, to use the Power Ratio (a real-time measure of the pulse-to-pulse variability) to identify sufficiently stationary targets.

2) TARGET HEIGHT UNCERTAINTY

Park and Fabry (2010) considered the effect of target height variability combined with changes in the vertical gradient of refractive index (dn/dh) and showed that, particularly at close ranges, the resulting phase change noise increased almost linearly in proportion to local target height variability, range, and dn/dh. For a local target height variability of 10 m at a range of 20 km, the root-mean-squared (rms) phase change errors due to this effect would be about 15°, 27°, and 47° at S band, C band, and X band, respectively [Park and Fabry 2010, their Eq. (9)], assuming a change in dn/dh of 20 ppm km^{-1}. For the example given, both the range and magnitude of these changes are quite conservative; in certain situations this can become a significant source of error. Relative to dn/dh in low levels of the standard atmosphere (-40 ppm km^{-1}), the greatest changes are likely in ducting conditions ($dn/dh \leq -157$ ppm km^{-1}).

3) TARGET LOCATION UNCERTAINTY

Refractivity retrieval based on (1) assumes that each clutter target is located at the center of the range gate, but in reality their positions within the gate will be random, which introduces a random phase change noise (Nicol et al. 2012b, manuscript submitted to *J. Atmos. Oceanic Technol.,* hereafter NIDK). For a 300-m gate at C band, NIDK established that the typical distance of the target from the gate center is 150 m; this 300-m difference in the length of the actual two-way propagation path from that assumed introduces random phase changes of 2° when $\Delta N = 1$. We will refer to this phase change noise, which scales linearly with pulse length, radar frequency, and ΔN, as "target location phase noise." For the U.K. C-band radars, $\Delta N \approx 20$ leads to $\sigma_{\Delta\phi} \approx 40°$; whereas at S band with 250-m gate separation (e.g., NEXRAD radars), the same phase change noise ($\sigma_{\Delta\phi} \approx 40°$) would result for $\Delta N \approx 50$ and for $\Delta N \approx 80$ with a 150-m gate separation (e.g., S-Pol).

4) PHASE NOISE OBSERVATIONS

The combination of these random errors for any given target can rather easily result in large total errors. The magnitude of the noise in observed phase changes is rarely reported, but Park and Fabry (2010) estimated approximately 70° noise for targets within a range of 10 km for scans separated by almost 1 month (their case 2). The change in dn/dh was estimated to be 15 ppm km^{-1}, and assuming a target height variability of 10 m, only a small fraction ($\sim 5°$) of total noise ($\sim 70°$) could be attributed to changes in dn/dh. As $\langle \Delta N \rangle_{\text{field}}$ was negligible (0.5 ppm), target location uncertainty would have been insignificant and the majority of the 70° noise was most likely due to target motion. For simplicity, we shall henceforth refer to all phase noise not resulting from target location uncertainty as "target motion phase noise," even though other sources of noise such as that due to target height uncertainty will also contribute to a given measurement.

Noise of this magnitude ($\sim 70°$) would scale to 125° and 220° at C and X bands, respectively. Our experience over several years at C band in the United Kingdom is that, even using the Power Ratio (see section 3b) to identify suitable stationary targets, the phase change noise for scans only 5 min apart is typically in the range from 20° to 90° for targets out to the 30-km range. The phase change noise increases with time and often exceeds 100° for changes over several hours, presumably because of sources of noise other than target motion, prohibiting reliable retrievals. This implies that, whereas the use of a suitable reference field to calculate refractivity fields over subsequent weeks or months may be possible at S band, phase noise and aliasing make such an approach very difficult at shorter wavelengths.

b. Outline of refractivity retrieval calculations

The refractivity algorithm is now described assuming that a judicious selection of targets has been made. To obtain robust refractivity estimates, a large number of phase change gradients must be averaged or spatial smoothing of the phase changes is required before these gradients are estimated. Smoothing with respect to range is equivalent to low-pass filtering of the measured phase change field. More weight is then given to low-frequency radial fluctuations (small ΔN) than high-frequency radial fluctuations (large ΔN), so ΔN may be underestimated. To avoid this, $\langle \Delta N \rangle_{\text{field}}$ is typically calculated using pulse-pair processing and its effect subtracted prior to any smoothing. Park and Fabry (2010) used a least squares fit to the azimuthally averaged phase changes as a function of range to infer $\langle \Delta N \rangle_{\text{field}}$ out to 40 km, though they did not explain why this was preferred to pulse-pair processing. The phase change corresponding to $\langle \Delta N \rangle_{\text{field}}$ as a function of range is then subtracted from the original phase changes before smoothing. Local refractivity perturbations are estimated from phase change gradients in small regions (typically several kilometers squared) using pulse-pair processing. Finally, $\langle \Delta N \rangle_{\text{field}}$ is added to these perturbations to derive the total ΔN. The smoothing then acts as a bandpass filter centered on the estimated $\langle \Delta N \rangle_{\text{field}}$ and ΔN may be biased toward the estimated $\langle \Delta N \rangle_{\text{field}}$, particularly in regions where large local refractivity perturbations occur.

3. Phase correlations in ground clutter

a. Origins of phase correlations in ground clutter

The returned radar signal may be modeled by the convolution of a distribution of targets with a range-weighting function and a beam-weighting function in azimuth of the transmitted pulse. Our model is based on Hubbert et al. (2009), who found from NEXRAD data that the signal variability of ground clutter returns is most accurately modeled by a single clutter target dominating each range gate. The range-weighting function for radar measurements depends on the duration and shape of the transmitted pulse and the filter response of the receiver. The range-weighting function (in terms of amplitude) is expressed here assuming a rectangular pulse and a Gaussian filter transfer function (Doviak and Zrnić 1993):

$$|W(x)| = [\text{erf}(x + b) - \text{erf}(x - b)]/2, \quad (3)$$

where $b = B_6 \tau \pi / 4\sqrt{\ln 2}$, $a = \pi/2\sqrt{\ln 2}$, and $x = (2aB_6/c)(r_0 - r)$.

Here, B_6 is the 6-dB receiver bandwidth, τ is the pulse duration, r is the range, and the weighting function is maximum at the range r_0. This range-weighting function is depicted in Figs. 1a,b for a point target located at the center of the range gate and located halfway between adjacent range-gate centers, respectively. The corresponding range-weighting functions calculated for the Met Office operational weather radars are shown in Figs. 1c,d.

The Gaussian filter transfer function corresponds to an infinite propagation delay through the receiver filters and is an excellent approximation around the peak responses, although poor at the tails (Doviak and Zrnić 1979). In modern radar systems, digital filters are often used in the receiver chain, which allows a great deal of flexibility in the filter design. For most applications, the range-gate spacing is matched to the nominal range resolution in relation to the pulse duration. However, these filters are designed with meteorological targets rather than ground clutter targets in mind; the secondary lobes in the range-weighting functions (in Figs. 1c,d) have little influence on meteorological echoes, but may not be appropriate for the extreme reflectivity gradients typical of ground clutter and lead to the return from a single dominant clutter target appearing in several adjacent range gates.

From Figs. 1c,d, we may expect absolute reflectivity differences from a single dominant clutter target of approximately 20 dB on average from one gate to the next. The corresponding reflectivity differences across two and three range gates are approximately 30 and 45 dB, respectively. These figures agree well with observations to be presented later (end of section 3). For both dominant and isolated clutter targets, when no competing returns exist in adjacent gates, the measured phase will be roughly the same at adjacent gates and the correlation between the gates will be close to unity.

b. Phase correlations in ground clutter observed with an operational weather radar

1) RADAR DESCRIPTION

The C-band radar at Cobbacombe Cross in southwest England (details in Table 1) is operated by the Met Office as part of the operational radar network and executes plan position indicator (PPI) scans at various elevations that are repeated every 5 min, but we shall only consider PPIs at the lowest operational elevation (0°). Measurements are obtained for each 1° in azimuth corresponding to about 44 transmitted pulses. Using complex notation ($v = I + iQ$), where I and Q are the in-phase and quadrature voltages, respectively, the mean signal from N_P consecutive pulses is given by

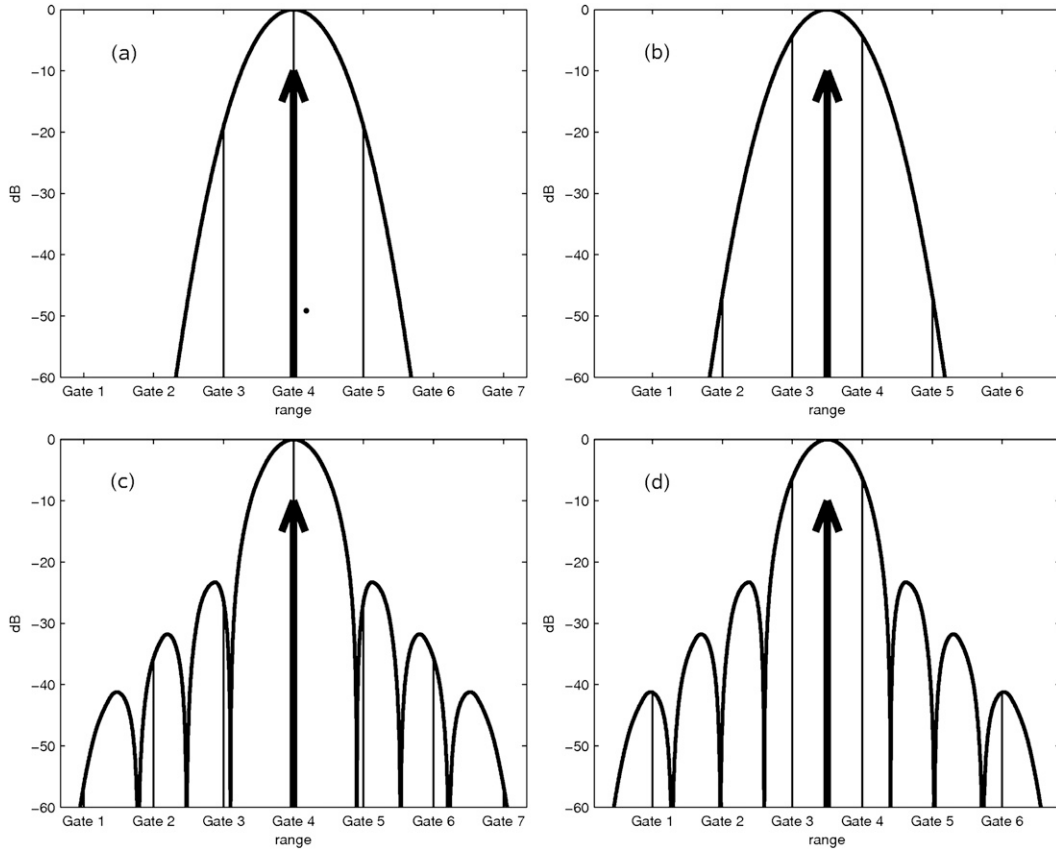$$\overline{V} = \sum_{j=1}^{N_P} v_j \Big/ N_P. \quad (4)$$

FIG. 1. Range-weighting function in terms of power (dB) corresponding to a point target located (a),(c) in the center of range gate 4 and (b),(d) exactly halfway between the centers of range gates 3 and 4. The range-weighting function in (a) and (b) is derived from (2) with a Gaussian filter transfer function assuming a bandwidth–pulse duration product of unity, and in (c) and (d) is derived from the actual digital filtering applied for radars of the operational U.K. weather radar network. Compared to the Gaussian filter transfer function, less power is spread into adjacent gates though more power is distributed into more distant range gates.

The Power Ratio (PR) provides a measure of the signal variability during the averaging process (i.e., per degree in azimuth) as expressed below and is a useful means of identifying suitable targets for refractivity in real time. PR = 1 indicates a perfectly constant signal in both phase and amplitude. The variability of either phase or amplitude results in lesser values and the mean value approaches $1/N_P$ for uncorrelated Gaussian noise:

$$\mathrm{PR} = \frac{|\overline{V}|^2}{\sum\limits_{j=1}^{N_P} |v_j|^2 \Big/ N_P}. \tag{5}$$

Empirical tests show that targets with PR > 0.7 are sufficiently stationary (Nicol et al. 2012a) to be used in refractivity retrievals. Simulations with Gaussian-distributed phase noise and a single target with a constant backscattering cross section indicate that this

threshold corresponds to rms pulse-to-pulse fluctuations of 34°, assuming that measurements are made as the radar scans past the target within a 1° half-power Gaussian beamwidth. This results in a standard error of 5° if 44 pulses with uncorrelated phase noise are averaged for each degree in azimuth, implying that the corresponding phase change noise for a target with PR = 0.7 would be only 7°. This phase change noise appears to contradict the observed phase change noise of

TABLE 1. Technical specifications of U.K. weather radars and operational parameters for the low-elevation scan.

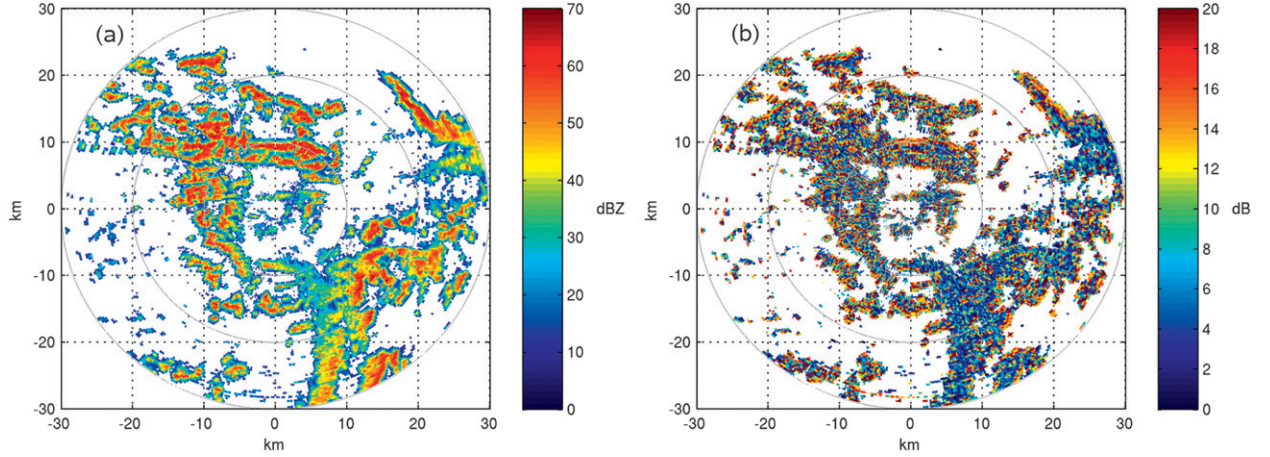|	
Frequency	5.6 GHz
Wavelength	5.4 cm
PRF	300 Hz
Antenna scan rate	$7.2°\ \mathrm{s}^{-1}$
Pulse length	2 μs
Range resolution	300 m
Pulses per degree	44

FIG. 2. (a) Dry weather example of the mean reflectivity (dBZ) and (b) the absolute radial reflectivity change ($|\Delta$dB$Z|$) between adjacent range gates for the ground clutter field averaged over 288 consecutive scans in 24 h (7 May 2008) for the radar at Cobbacombe Cross.

20°–90° between scans separated by 5 min reported in section 2a even though these simulations only account for target motion. However, motion-induced phase errors would typically be correlated from pulse to pulse (e.g., targets swaying slowly relative to the interpulse period) and the standard error in the mean phase will therefore be larger. So, without knowing the pulse-to-pulse phase correlation, it is not possible to interpret PR directly in terms of the standard deviation of the phase noise.

Although magnetron transmitters are prone to drift in frequency with ambient temperature changes (e.g., Skolnik 1990), only times when the transmitter frequency is essentially unchanged are considered in this paper, replicating measurements made by radars with klystron transmitters. The effect of frequency changes on phase measurements is beyond the scope of this work but has been treated in NIDK.

Prior to each PPI, the local oscillator (LO) frequencies in the receiver chain are digitally requested to match the transmitter frequency to maintain a well-centered intermediate frequency. This is primarily achieved by putting a numerically controlled oscillator through occasional adjustment in steps of 100 kHz of the digitally requested stable local oscillator (STALO). Both the real-time measurement of the transmitter frequency and the selection of local oscillator frequencies are achieved with a very high degree of accuracy (<0.25 ppm; NIDK). The high-resolution automated frequency control (AFC) of this system implies that the same phase will be observed for returns from a single target that dominates returns at adjacent range gates. This is identical to measurements made by radars with klystron transmitters if the sum of the LO frequencies is equal to the transmitter frequency. Provided that the difference between this sum of the LO

frequencies and the transmitter frequency does not change in time, the phase change difference between adjacent range gates for highly correlated targets will always equal zero. While both the LO and transmitter frequencies are constant for klystron radar systems, this continuity may be broken if the radar is turned off, as identical frequencies are not guaranteed when the system is restarted. Refractivity retrievals bridging any breaks in the continuity of operation may exhibit both refractivity biases due to LO frequency changes and increased phase change noise due to transmitter frequency changes (NIDK).

2) GROUND CLUTTER CHARACTERISTICS

An example of the dry-weather ground clutter reflectivity (dBZ) field averaged over 288 consecutive scans in 24 h for the radar at Cobbacombe is shown in Fig. 2a. The corresponding mean absolute radial reflectivity change ($|\Delta$dB$Z|$) between adjacent range gates is shown in Fig. 2b, where the maximum absolute changes approach 20 dB as predicted from Figs. 1c,d. As expected, the radial gradients are largest surrounding targets with high reflectivities and at the edges of the clutter field with respect to range.

The phase correlation may be derived from the covariance of the mean complex radar voltage (V) reconstructed from the measured phase (ϕ) and reflectivity (Z) at each range gate, where j and k are the range and azimuth indices, respectively:

$$V_{j,k} = \sqrt{Z_{j,k}}\, e^{-i\phi_{j,k}}. \qquad (6)$$

Strictly speaking, the received power should be used rather than the range-corrected reflectivity in (6),
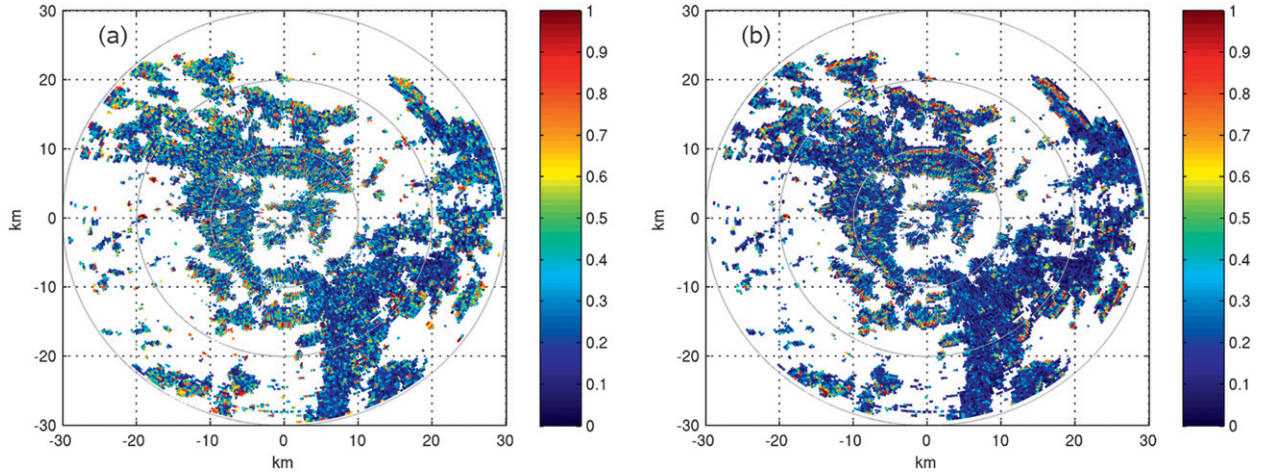
FIG. 3. Phase correlations from returns at a separation of (a) one gate and (b) two gates on the same azimuth using 288 consecutive scans over 24 h (7 May 2008).

although this distinction is not important since we are considering differences over a few range gates at most. The phase correlation of returns separated by m range gates along the same azimuth, calculated over N_S PPI scans, is expressed here:

$$|\overline{R}_{j,k}| = \frac{\left| \sum\limits_{n=1}^{N_S} V_{j,k} V_{j-m,k}^* \right|}{\sqrt{\sum\limits_{n=1}^{N_S} Z_{j,k} \sum\limits_{n=1}^{N_S} Z_{j-m,k}}}. \quad (7)$$

Correlations have been estimated using 1-, 2-, 3-, and 4-gate separations for all cluttered gates with mean reflectivity (dBZ) > 10 using 288 consecutive scans (every 5 min) in 24 h (7 May 2008). The corresponding field-average correlations are 0.33, 0.26, 0.17, and 0.12, respectively, for the various gate separations. Correlations calculated with 1- and 2-gate separations are shown in Figs. 3a,b. The correlations have some structure within the ground clutter field and while the correlation generally decreases going from a 1-gate to 2-gate separation, some locations exhibit persistently high or even higher correlations. Very similar correlation distributions have been obtained from other dry days during summer 2008; this is consistent with the highly correlated target couplets being from a single target and hence unaffected by target motion and refractivity changes.

In Figs. 4a–d, the 1-, 2-, 3-, and 4-gate phase correlations, respectively, are plotted against the mean absolute reflectivity change ($|\Delta\mathrm{dB}Z|$) across the range-gate couplets, again using 288 scans within 24 h (7 May 2008). Hence, Fig. 4a shows the correlations from Fig. 3a

plotted versus the 1-gate absolute reflectivity changes from Fig. 2b. The strong relationship between correlation and the absolute reflectivity gradient for some targets is quite apparent. The highest correlations tend to occur in the vicinity of strong targets, presumably when there are no or only weak targets in the neighboring range gates.

One may note that the correlations greater than 0.5 generally correspond to absolute radial reflectivity differences of approximately 20, 30, and 45 dB for 1-, 2-, and 3-gate separations, respectively. This is consistent with the actual range-weighting function shown in Figs. 1c,d. Based on this, one might expect absolute radial reflectivity differences greater than 60 dB with a 4-gate separation, but these are not observed, considering the minimum threshold for clutter (dB$Z_{\min}$ = 10) and that the ground clutter returns do not exceed about dBZ = 70. In the analysis presented in section 4, the subjectively chosen thresholds (0.5, 0.4, 0.35, and 0.3) indicated in Figs. 4a–d are used to separate the majority of ground clutter couplets with relatively low correlations from the highly correlated couplets associated with steep reflectivity gradients that, as we shall see, can result in refractivity biases. Decreasing thresholds have been used because the correlation between independent targets decreases with increasing gate separations as a result of the diurnal cycle of refractivity changes during the 24-h period considered.

4. Observed biases in field-averaged refractivity changes at C band

In this section, $\langle \Delta N \rangle_{\mathrm{field}}$ between two times separated by 8 h were estimated out to the 10-km range using
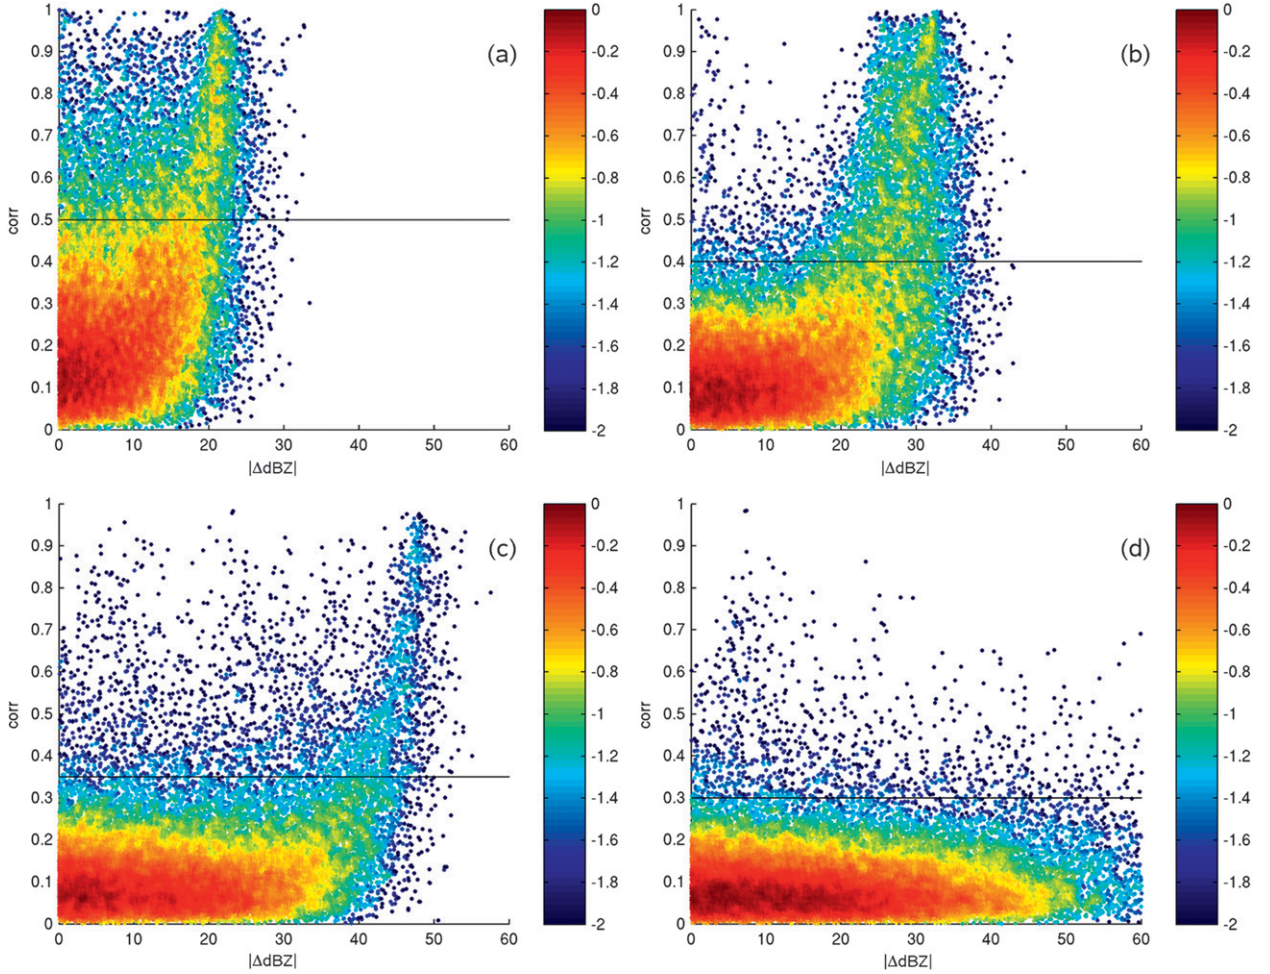
FIG. 4. The (a) 1-, (b) 2-, (c) 3-, and (d) 4-gate phase correlations plotted against the modulus of the mean reflectivity change ($|\Delta dBZ|$) in the ground clutter field averaged over 288 consecutive scans from 7 May 2008. Subjectively chosen thresholds are indicated in each case to isolate the most highly correlated couplets for the analysis in section 4. Number density (color) is indicated using a logarithmic scale relative to the maximum.

pulse-pair processing with the aim of quantifying any biases introduced by using couplets with different gate spacing and also the effect of including and excluding correlated targets. The true $\langle \Delta N \rangle_{\text{field}}$ was assumed to be equal to that estimated using a "least squares fit" method, derived from a linear fit of all phase changes averaged over all azimuths. Although we do not demonstrate that this approach generally leads to unbiased estimates in the presence of phase-correlated returns until section 5b, this approach resulted in better agreement with measurements from synoptic weather stations.

To avoid spurious effects due to drifts in the magnetron frequency, two periods were carefully chosen when both the magnetron frequency and the local oscillator were essentially unchanged. Although the transmitter frequency changes have no direct effect on the retrieved

refractivity, the associated LO frequency changes bias the estimated ΔN in proportion to the LO frequency changes in parts per million (NIDK). Identifying occasions when significant $\langle \Delta N \rangle_{\text{field}}$ have occurred while the magnetron frequency remains constant is extremely difficult, as both can be caused by temperature changes. However, one such occasion has been identified between about 2300 UTC 21 May 2008 and 0700 UTC 22 May 2008. This period appears to be primarily associated with a change in air mass introducing more humid air. Surface observations indicated that the humidity gradually rose from around 70% RH to near saturation while the temperature remained largely unchanged. Two sets of 8 scans separated by 5 min from 2230 to 2305 UTC 21 May 2008 and from 0625 to 0700 UTC 22 May 2008 were used to estimate $\langle \Delta N \rangle_{\text{field}}$ based on the raw phase change measurements out to a range of just

TABLE 2. Mean estimated $\langle\Delta N\rangle_{field}$ with errors based on two sets of 8 scans (64 scan pairs) between 2230 and 2305 UTC 21 May 2008 and 0625 and 0700 UTC 22 May 2008 using pulse-pair processing for various gate separations within 10 km of the radar. The true $\langle\Delta N\rangle_{field}$ was estimated to be $+6.8 \pm 0.2$ N units using a linear best fit to phase changes averaged over all azimuths within 10 km.

	1-gate	2-gate	3-gate	4-gate
$\langle\Delta N\rangle_{field} \pm \sigma_{\Delta N}$ (all targets)	3.9 ±0.3	4.7 ±0.2	6.0 ±0.3	6.7 ±0.3
$\langle\Delta N\rangle_{field} \pm \sigma_{\Delta N}$ (excluding highly correlated couplets)	5.5 ±0.4	6.5 ±0.3	6.5 ±0.3	6.8 ±0.3
$\langle\Delta N\rangle_{field} \pm \sigma_{\Delta N}$ (only highly correlated couplets)	1.6 ±0.3	0.8 ±0.2	2.5 ±0.3	3.3 ±0.7
Total number of couplets	3673	3098	2701	2533
Correlation thresholds	0.5	0.4	0.35	0.3
% of highly correlated couplets	23.3	15.6	7.0	3.1

10 km. The rms frequency change between the 64 pairs of scans was 2.26 kHz or 0.4 ppm, which corresponds to a refractivity uncertainty of 0.4 N units. There are effectively only 8 independent pairs of scans between the two periods, so the frequency changes should introduce a mean error of just 0.14 ppm into $\langle\Delta N\rangle_{field}$.

The hourly synoptic observations (between 2300 and 0700 UTC) at two stations, one 20 km to the northwest and the other 20 km to the southwest of the radar, indicated gradual increases of 6.5 and 9.2 N units, respectively, throughout the period, so the average increase was 7.85 N units. Estimates of $\langle\Delta N\rangle_{field}$ using a linear least squares fit to the phase changes averaged over all azimuths as a function of range were limited to 10 km. Targets beyond 10 km were not used because visual inspection of the azimuthally averaged phase changes indicated that reliable unfolding could no longer be achieved. Least squares estimates are prone to large errors when proper unfolding cannot be achieved. For this reason, such approaches are not well-suited to local ΔN estimates (i.e., with a limited number of observations) though they work well for $\langle\Delta N\rangle_{field}$ when averaging phase changes at close ranges over all azimuths. The least squares estimates indicate a mean increase of 6.8 N units between 2245 and 0640 UTC, which is not significantly different from the mean change of 7.85 N units from the two surface stations. As already stated, this increase of 6.8 N units is taken as truth when calculating pulse-pair biases.

The two sets of 8 scans provided 64 estimates of $\langle\Delta N\rangle_{field}$ between about 2245 and 0640 UTC. These 64 estimates were used to estimate the mean $\langle\Delta N\rangle_{field}$ as a function of the range-gate separation used in the pulse-pair processing; the pulse-pair technique has been used in previously published studies. The standard error has also been determined considering 8 independent pairs of scans. The mean and standard error were estimated using gate steps from 1 to 4, initially including all stationary target couplets having PR > 0.7 and dBZ_{min} > 10. In addition, the phase correlation thresholds in Figs.

4a–d were used to exclude (and isolate) highly correlated target couplets as a function of the gate separation. The $\langle\Delta N\rangle_{field}$ estimates are shown in Table 2, where the average total number of couplets, the correlation thresholds, and the percentage of "highly correlated" couplets in each category are also included.

Table 2 confirms that, when all stationary targets are used for a 1-gate and 2-gate separation, there is an underestimate of about 43% and 32%, respectively, in $\langle\Delta N\rangle_{field}$, but for a 4-gate separation the estimate is very close to the true value. For contiguous gates, 23% of gates are highly correlated returns and their removal reduces the bias from 43% to 19%. However, for 2-gate separation only 16% of the gates are highly correlated and their removal reduces the bias from 32% to less than 4%. This demonstrates that pulse-pair processing of raw phase changes for the widely used 1-gate separation can result in significant underestimates of ΔN when correlated returns are present, even when such changes are much less than $|\Delta N|_{folding}$. There is some scope for reducing the bias in $\langle\Delta N\rangle_{field}$ by increasing the minimum clutter threshold; for example, the 43% bias with dBZ_{min} = 10 reduces to 38% and 20% with dBZ_{min} = 20 and 30, respectively. However, increasing the clutter thresholds leads to greater errors in estimating $\langle\Delta N\rangle_{field}$, and as we shall later see, biases in $\langle\Delta N\rangle_{field}$ may be avoided completely using a least squares fit method.

5. Simulated biases in field-averaged refractivity changes at C band

The observations in section 3 revealed that phase-correlated returns from clutter are significant. In this section, physically based simulations of these returns are used to quantify the biases in $\langle\Delta N\rangle_{field}$ estimated using pulse-pair processing of raw phase change measurements and also using a least squares approach; the predicted biases are then compared with the observed biases from section 4.

a. Simulation procedure

The simulated ground clutter is based on the observed values from dry-weather clutter (dB$Z > 10$) within 30 km of the radar at Cobbacombe, averaged over a 24-h period (7 May 2008) shown in Fig. 2a and are constructed assuming a single target at each range gate. The observed distribution of ground clutter reflectivities has not been used directly, as these measurements have already been smoothed by the real range-weighting function. Instead, while the spatial coverage was maintained, the reflectivities at cluttered range gates were randomly redistributed. Hence, the reflectivities in neighboring gates prior to smoothing by the range-weighting function are independent. This random redistribution in some respects mitigates the radial smoothing due to the real range-weighting function in modeling the underlying ground clutter field.

Each target was assigned a random location within the range gate and the range-weighting function corresponding to the actual digital filtering applied in the U.K. operational weather radars (shown in Figs. 1c,d) was applied. To model the effects of the finite beamwidth, the "raw" complex signals were then smoothed by the range-weighting function and in azimuth by a Gaussian smoothing function with a 3-dB width of 1°. The process was then repeated with the same distribution of targets but adding both range-dependent phase changes due to the prescribed uniform $\langle \Delta N \rangle_{\text{field}}$ and phase change noise ($\sigma_{\Delta\phi}$). Phase change noise has been represented as an independent, zero-mean Gaussian random variable (with standard deviations ranging up to 70° in 10° steps in the various simulations). Although phase change noise was only added at the later time, this is equivalent to phase noise (σ_ϕ) added at each time, where $\sigma_\phi = \sigma_{\Delta\phi}/\sqrt{2}$. Although this additional phase change noise is later referred to as "target motion phase change noise," it also includes contributions from other sources such as target height uncertainty and changes in dn/dh. Target location phase noise is implicitly included in all simulations; it is the only source of noise in the simulations when none is added explicitly (i.e., when $\sigma_{\Delta\phi} = 0$).

Pulse-pair processing was used to estimate $\langle \Delta N \rangle_{\text{field}}$ with the standard 1-gate separation along with larger separations up to 4 gates. In addition, $\langle \Delta N \rangle_{\text{field}}$ was also derived using a linear least squares fit to phase changes averaged over all azimuths. The mean and standard deviation of the estimated $\langle \Delta N \rangle_{\text{field}}$ were calculated from 1000 realizations. The mean biases derived from these simulations are more robust than those estimated from observations because of the large number of realizations and also the large number of target couplets in each realization. The simulations include all clutter targets out to 30 km (cf. 10 km for the observations), and as only suitably stationary target couplets (passing the PR threshold) were considered from the observations, the number of target couplets was about 5–6 times larger in the simulations than in the observations. An additional difference is that all targets have the same prescribed phase error characteristics in the simulations while the error characteristics in the observations are unknown and will presumably vary significantly from target to target.

b. Simulation results for pulse-pair and least squares estimates

1) PULSE-PAIR ESTIMATES

Initially, we shall consider C-band wavelengths and a range resolution (and range-gate spacing) of 300 m corresponding to the operational weather radar data analyzed in this paper, for which $|\Delta N|_{\text{folding}} \approx 44$. Histograms of ΔN estimated from each individual 1-gate couplet are shown in Figs. 5a–d from single simulations with a uniform modeled $\langle \Delta N \rangle_{\text{field}} = 10$, 20, 30, and 40, respectively. No additional target motion phase change noise has been added, so only the effect of target location phase noise is included.

The peak at $\Delta N = 0$ corresponds to highly correlated couplets that tend to bias estimated $\langle \Delta N \rangle_{\text{field}}$. The distribution of ΔN from less-correlated targets becomes progressively broader as the modeled $\langle \Delta N \rangle_{\text{field}}$ increases, as expected from the discussion on phase noise in section 2a. This demonstrates that the target location within the range gate not only determines the spreading of power into adjacent gates via the range-weighting function but also introduces phase change noise (proportional to ΔN), even when none has been explicitly added in these simulations.

In contrast to these results, simulations performed with a rectangular range-weighting function (corresponding to a rectangular transmitted pulse matched to the gate length and no receiver filtering) confirmed that no refractivity biases occur. In this case, returns from adjacent gates are uncorrelated and the histograms of estimated changes are symmetrical about the modeled $\langle \Delta N \rangle_{\text{field}}$ and the peak at $\Delta N = 0$ is absent, regardless of the phase noise present (results not shown). However, a more realistic range-weighting function introduces the peak at $\Delta N = 0$, which leads to biases in the simulated refractivity retrievals as shown in Table 3a; calculations using pulse-pair processing have been made for $\langle \Delta N \rangle_{\text{field}}$ up to 40 in steps of 5 N units. The first column is for target location phase noise only with the bold numbers for the specific cases in Fig. 5, and subsequent columns show the effect of additional target motion phase change noise.
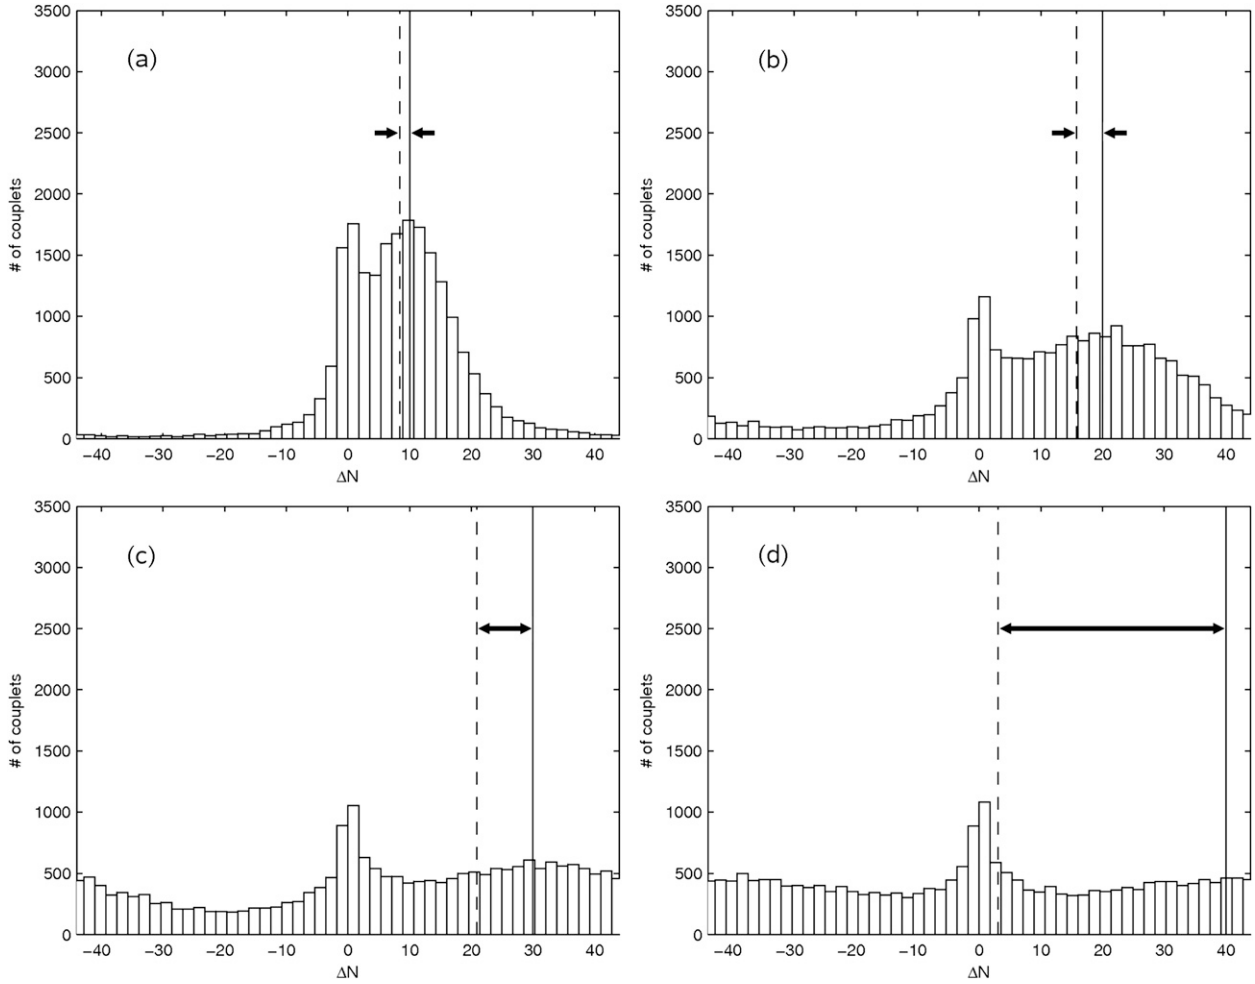
FIG. 5. Histograms of ΔN from each individual 1-gate couplet from single simulations with uniform modeled $\langle \Delta N \rangle_{\text{field}}$ = (a) 10, (b) 20, (c) 30, and (d) 40 (solid vertical lines). These simulations incorporate only phase change noise due to random target location within each gate and no other sources of phase noise. The corresponding pulse-pair estimates of $\langle \Delta N \rangle_{\text{field}}$ were (a) 8.4, (b) 15.5, (c) 21.1, and (d) 2.2 (dashed vertical lines); the biases are indicated by arrows. The peak at $\Delta N = 0$ only occurs in simulations incorporating a realistic range-weighting function and is associated with highly correlated target couplets.

Small biases exist even for small $\langle \Delta N \rangle_{\text{field}}$ and with minimal phase errors. However, the biases become rapidly larger as both the modeled $\langle \Delta N \rangle_{\text{field}}$ and the phase noise increase so that the estimated $\langle \Delta N \rangle_{\text{field}}$ reaches a maximum, even when the true $\langle \Delta N \rangle_{\text{field}}$ is well below $|N|_{\text{folding}}$, and the $\langle \Delta N \rangle_{\text{field}}$ estimate subsequently falls to zero as $|\Delta N|_{\text{folding}}$ is approached. This suggests that it becomes impossible to retrieve ΔN values approaching $|\Delta N|_{\text{folding}}$ even when target motion phase noise is absent. It would be practically impossible to anticipate these biases, as they are not only proportional to the quantity that is to be measured (ΔN) but also to the unknown magnitude of the target motion phase noise. Table 3b shows the results of simulations for pulse-pair calculations with 2-, 3- and 4-gate separations, reducing $|\Delta N|_{\text{folding}}$ to 22, 15, and 11, respectively.

The results indicate that although the range of modeled $\langle \Delta N \rangle_{\text{field}}$ has been reduced, the biases are much smaller.

2) LEAST SQUARES ESTIMATES

As an alternative to pulse-pair processing for estimating the $\langle \Delta N \rangle_{\text{field}}$, we shall now consider estimates using a least squares fit to the phase changes averaged over all azimuths out to a range of 30 km. This requires that the azimuthally averaged phase changes are corrected for aliasing when changes greater than 180° are observed from one range gate to the next. The results using this approach (shown in Table 3c) indicate that these estimates are generally unbiased, even when correlated targets are considered. The biases that start to occur as the $\langle \Delta N \rangle_{\text{field}}$ approaches $|\Delta N|_{\text{folding}}$ result

TABLE 3. Mean estimated $\langle\Delta N\rangle_{\text{field}}$ (to 30 km) using pulse-pair and least squares approaches from simulations as functions of ΔN and phase change noise ($\sigma_{\Delta\phi}$). These simulations are for C-band wavelengths with a range-gate separation of 300 m and the standard errors are less than the precision shown in all cases. Values in bold correspond to the examples in Fig. 5.

(a) Pulse-pair processing (1-gate separation)							
$\sigma_{\Delta\phi}$ ΔN 1-gate 0	10	20	30	40	50	60	70
5 4.3	4.3	4.1	3.9	3.6	3.2	2.7	2.1
10 **8.4**	8.3	8.0	7.5	6.9	6.0	5.0	3.8
15 12.1	12.0	11.5	10.8	9.6	8.3	6.6	4.9
20 **15.5**	15.4	14.6	13.4	11.8	9.8	7.5	5.2
25 18.6	18.4	17.3	15.6	13.1	10.1	7.1	4.6
30 **21.1**	20.6	19.1	15.9	12.2	8.3	5.2	3.1
35 20.9	19.8	17.2	11.3	7.3	4.0	2.5	1.3
40 **2.2**	0.8	1.1	0.2	0.1	0.0	0.0	0.0

(b) Pulse-pair processing (2-, 3-, and 4-gate separation)							
2-gate							
5 4.6	4.6	4.6	4.5	4.3	4.1	3.9	3.5
10 9.2	9.2	9.1	9.0	8.7	8.3	7.7	6.7
15 14.0	13.9	13.8	13.6	13.2	12.7	11.7	9.9
20 19.1	19.1	19.0	18.9	18.6	18.2	17.0	13.6
3-gate							
5 4.8	4.8	4.8	4.7	4.7	4.6	4.5	4.3
10 9.6	9.6	9.6	9.6	9.5	9.4	9.2	8.9
4-gate							
5 4.9	4.9	4.9	4.9	4.8	4.8	4.8	4.8
10 9.8	9.8	9.8	9.8	9.8	9.8	9.8	9.7

(c) Least squares linear fit to phase data averaged over all azimuths							
$\sigma_{\Delta\phi}$ ΔN 0	10	20	30	40	50	60	70
5 5.0	5.0	5.0	5.0	5.0	5.0	5.0	5.0
10 10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
15 15.0	15.0	15.0	15.0	15.0	15.0	15.0	15.0
20 20.0	20.0	20.0	20.0	20.0	20.0	20.0	19.9
25 25.0	25.0	25.0	25.0	25.0	25.0	24.9	24.4
30 30.0	29.9	29.9	29.9	29.8	29.4	28.4	26.4
35 33.6	33.7	33.6	33.1	32.0	30.1	27.0	22.7
40 25.5	25.7	23.7	23.0	20.7	17.4	14.6	11.5

from incorrect dealiasing of the azimuthally averaged phase changes rather than from phase correlations across individual couplets. It is for this reason that the "least squares" approach is not well-suited to ΔN estimates over small areas (e.g., several kilometers), as the number of targets averaged at each azimuth is relatively small and successful dealiasing becomes much more difficult.

c. Comparisons with observations (case study: 21–22 May 2008)

For this case study, $\langle\Delta N\rangle_{\text{field}} = 6.8$ was derived in section 4 using an azimuthally averaged least squares fit that should be unbiased, but Table 2 showed that the

TABLE 4. Mean estimated $\langle\Delta N\rangle_{\text{field}}$ (to 30 km) using pulse-pair processing for various gate separations from simulations as a function of phase change noise ($\sigma_{\Delta\phi}$). These simulations are for C-band wavelengths with a 300-m range-gate separation and $\langle\Delta N\rangle_{\text{field}} = 6.8$, corresponding to the observations in Table 2. The standard errors are less than the precision shown in all cases.

$\sigma_{\Delta\phi}$	0	10	20	30	40	50	60	70
1-gate	5.8	5.8	5.6	5.3	4.9	4.2	3.6	2.8
2-gate	6.3	6.3	6.2	6.1	5.9	5.6	5.2	4.6
3-gate	6.5	6.5	6.5	6.5	6.4	6.3	6.1	5.8
4-gate	6.6	6.6	6.6	6.6	6.6	6.6	6.6	6.5

changes were biased low when pulse-pair processing was used. Simulated $\langle\Delta N\rangle_{\text{field}}$ are presented in Table 4 for pulse-pair calculations using gate separations from 1 to 4 and additional target motion phase change noise, again with standard deviations ranging up to 70° in 10° steps. The simulations with additional target motion noise of 60° compare well with the observations in Table 2, confirming that phase-correlated targets can result in significant biases in estimates of $\langle\Delta N\rangle_{\text{field}}$ using pulse-pair processing.

6. The effect of spatially smoothed phase changes on refractivity biases at S band

Spatial smoothing of raw phase change data is required to reduce the influence of phase change noise, but smoothing with respect to range essentially acts as a low-pass filter on phase change measurements, tending to bias ΔN estimates toward zero. To avoid this it is common practice to subtract the effect of $\langle\Delta N\rangle_{\text{field}}$ from the raw phase change data before smoothing is applied, though this can result in biases toward the estimated $\langle\Delta N\rangle_{\text{field}}$. In this section we discuss 1) the form of smoothing kernels that have been used, and then their effect on various simulations; 2) simulations of $\langle\Delta N\rangle_{\text{field}}$ derived from pulse-pair processing following spatial smoothing of phase changes; 3) as in 2), but the spatial smoothing is done after subtraction of the $\langle\Delta N\rangle_{\text{field}}$ derived from pulse-pair processing; and 4) as in 3), but the subtracted $\langle\Delta N\rangle_{\text{field}}$ is derived from an azimuthally averaged least squares fit. In all cases, the simulated value of $\langle\Delta N\rangle_{\text{field}}$ is uniform over the domain; the effect of smoothing on our ability to retrieve the spatial perturbations of ΔN will be discussed in section 8.

a. Implementation of the smoothing kernels

We shall consider two categories of radars and the refractivity algorithms that have been applied in the literature: 1) NEXRAD WSR-88Ds (e.g., Bodine et al. 2011) with a range resolution of 235 m using the refractivity algorithm from Cheong et al. (2008), and 2) S-Pol (e.g.,
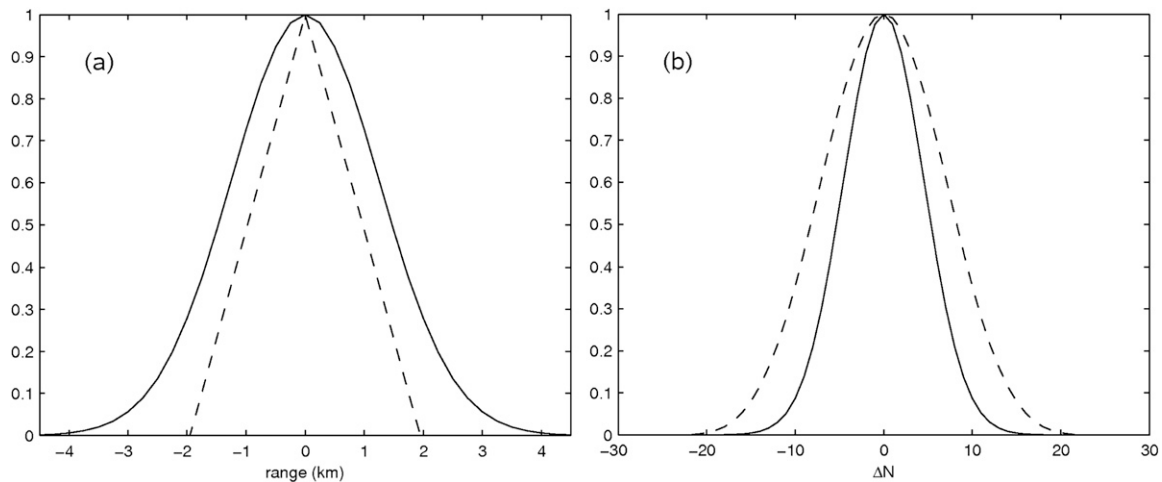
FIG. 6. (a) The smoothing kernels with respect to range as used in the C08 (solid line) and F04 (dashed line) algorithms; (b) the corresponding frequency responses as a function of ΔN.

Weckwerth et al. 2005) and the McGill radar (e.g., Fabry 2004), both with range resolutions of 150 m using the refractivity algorithm described in Fabry (2004). These two categories will henceforth be referred to simply as the C08 and F04, respectively. Apart from the differences in range resolution, there are some subtle differences between the two retrieval algorithms, including the spatial smoothing kernel applied to the phase change data.

The C08 smoothing kernel consists of a two-dimensional (2D) Gaussian filter with a physical width of 2.5 km. Following the definition in Cheong et al. (2008), the physical width is equivalent to twice the standard deviation of the Gaussian function. The F04 smoothing kernel consists of a 2D triangular function with a base of 4 km. For simplicity, we shall only consider the kernels in one dimension (i.e., radially), as this determines the biasing effect. Each smoothing kernel can be characterized by a frequency response (though in the spatial rather than the time domain), which may be represented as a function of ΔN (proportional to the phase change gradient). These two smoothing kernels are shown as a function of range in Fig. 6a and their frequency responses are shown in Fig. 6b in terms of ΔN (all normalized with a peak value = 1). The width of these smoothing kernels at the 0.5 level (i.e., ~3 km and ~2 km for C08 and F04, respectively) is approximately 13 range gates in each case and much greater than the radial extent of phase-correlated returns. Characterizing the width of the frequency responses the same way results in half-widths of ~5 and ~8 N units.

The effect of the smoothing kernels may be considered equivalent to the multiplication of their frequency responses (Fig. 6b) with the underlying ΔN (on a couplet-by-couplet basis) as depicted, for example, in

Figs. 5a–d (ignoring the smoothing of the beam in azimuth). If $\langle \Delta N \rangle_{\text{field}}$ is not subtracted, smoothing acts as a low-pass filter. When $\langle \Delta N \rangle_{\text{field}}$ is subtracted before smoothing, it acts as a bandpass filter centered at $\langle \Delta N \rangle_{\text{field}}$. While smoothing largely negates the effect of phase-correlated returns (which is confirmed by the simulations that follow), it tends to bias refractivity estimates toward the estimated $\langle \Delta N \rangle_{\text{field}}$.

b. Simulations of smoothing with no mean-field change subtraction

Simulations similar to those in section 5 were performed, though now for S-band wavelengths and assuming a Gaussian receiver filter frequency response; the particular details about the receiver filters of these radars are not readily available and this approximation is likely to be sufficient concerning returns in adjacent range gates (Doviak and Zrnić 1993). The range-weighting function is then described by (3) and shown in Figs. 1a,b. Spatial smoothing is also applied now to the normalized (i.e., unit length) phase change data prior to pulse-pair processing. Hence, these simulations combine the effects of the smoothing kernel and phase-correlated returns.

Results are shown in Table 5 for C08 and F04 with modeled $\langle \Delta N \rangle_{\text{field}}$ of up to 20 in steps of 5 N units, again with additional target motion phase change noise ranging up to 70°. Results from simulations using a rectangular range-weighting function matched to the range-gate spacing (resulting in no phase correlations) were almost identical to those presented here. This confirms that the smoothing kernel determines these biases and that biases due to phase-correlated returns are not significant after smoothing. This is not surprising, as the extent of the smoothing kernels with respect to range is several

TABLE 5. Mean estimated $\langle\Delta N\rangle_{field}$ (to 30 km) using (1 gate) pulse-pair processing from simulations as functions of ΔN and phase change noise ($\sigma_{\Delta\phi}$). These simulations are for S-band wavelengths using smoothing kernels from previous publications: C08 as applied to the NEXRAD WSR-88D radars and F04 as applied to the S-Pol and McGill radars. The standard errors are less than the precision shown in all cases.

$\sigma_{\Delta\phi}$ ΔN C08	0	10	20	30	40	50	60	70
5	4.3	4.2	4.2	4.2	4.2	4.1	4.0	3.8
10	7.2	7.2	7.1	6.9	6.6	6.1	5.5	4.7
15	5.3	5.1	4.7	4.2	3.5	2.7	2.0	1.4
20	2.3	2.1	1.8	1.4	1.0	0.7	0.5	0.3
F04								
5	4.3	4.3	4.2	4.2	4.2	4.2	4.1	4.0
10	8.1	8.1	8.1	8.0	7.9	7.8	7.5	7.2
15	10.4	10.4	10.2	10.0	9.7	9.1	8.4	7.3
20	7.5	7.4	7.1	6.6	5.8	4.9	3.9	2.9

times larger than the extent of phase correlations. As expected, biases are larger for C08 than F04 because of the larger smoothing kernel.

The results presented in Table 5 correspond to the biases that would occur if $\langle\Delta N\rangle_{field}$ were not removed prior to phase change smoothing, or alternatively, if $\langle\Delta N\rangle_{field}$ were close to zero. Typically, $\langle\Delta N\rangle_{field}$ is subtracted prior to smoothing so these results indicate the biases for local perturbations in ΔN relative to $\langle\Delta N\rangle_{field}$. This implies that retrieved refractivity fields will often be significantly smoother than the true fields, as the retrieved fields will tend to be biased toward $\langle\Delta N\rangle_{field}$. These effects become greater as the radial extent of the smoothing kernel increases and again as phase change noise increases.

Considering target motion phase change noise of 30° as an example and $\Delta N = 5$, 10, 15, and 20 relative to $\langle\Delta N\rangle_{field}$ would be expected to result in quite unacceptable biases of −16%, −31%, −72%, and −93% for C08 and −16%, −20%, −33%, and −68% for F04, respectively. As a rough approximation, biases greater than 2 N units tend to occur when local ΔN perturbations (relative to $\langle\Delta N\rangle_{field}$) exceed the frequency response half-widths of the smoothing kernels as previously defined (i.e., ~5 and ~8 N units for C08 and F04, respectively). These simulations indicate the importance of restricting the radial extent of smoothing kernels applied to phase change measurements, as otherwise refractivity retrievals may exhibit dramatically less spatial variability than the actual fields. There is a trade-off between this effect and the fact that the smoothing must sufficiently reduce phase change noise to obtain reliable retrievals. However, the smoothing kernel does not need to be of equal length in range and azimuth.

c. Simulations of smoothing after subtracting $\langle\Delta N\rangle_{field}$ derived using pulse-pair processing

To mimic typical retrieval algorithms, the effect of the estimated $\langle\Delta N\rangle_{field}$ is now subtracted from the raw phase change data prior to smoothing and the mean ΔN is estimated relative to $\langle\Delta N\rangle_{field}$, all using pulse-pair processing. Finally, $\langle\Delta N\rangle_{field}$ is added to these estimates to quantify the total bias in estimating $\langle\Delta N\rangle_{field}$, incorporating the two-stage approach typical of refractivity retrievals.

Simulations have been performed for C08 and F04 and the total estimated ΔN for modeled $\langle\Delta N\rangle_{field}$ up to 70 are shown in Table 6a. The improvement compared to Table 5 is striking, confirming why the two-stage approach of retrievals is required. For a target motion phase change noise of 30° and $\Delta N = 20$, the biases with mean-field subtraction are only 3.5% for both C08 and F04, rather than 93% and 68%, respectively, without. However, for the same phase change noise and $\Delta N = 60$, the resulting biases are almost 30 N units (50%) for C08 though only 7 N units (12%) for F04, primarily because of the initial underestimates of $\langle\Delta N\rangle_{field}$ due to phase-correlated returns.

d. Simulations of smoothing after subtracting $\langle\Delta N\rangle_{field}$ derived using least squares

Table 6b shows that a remarkable improvement in performance is achieved if the subtracted $\langle\Delta N\rangle_{field}$ is derived from a least squares fit to the azimuthally averaged phase changes. The advantage of this approach is that phase correlation artifacts associated with spreading targets are effectively negated. From Table 6b, we see that accurate mean values of $\langle\Delta N\rangle_{field}$ can now be retrieved. For a target motion phase change noise of 30° and $\Delta N = 60$, the biases are now only 0.2 N units for both C08 and F04.

7. Discussion of previously published refractivity results

When $\langle\Delta N\rangle_{field}$ is estimated using pulse-pair processing, phase-correlated returns tend to bias estimated ΔN toward zero and consequently refractivity retrievals will be biased toward N_{ref}. This may explain some of the extreme refractivity differences, often around 20 N units and at times as large as 30 N units, reported by Bodine et al. (2011) concerning a NEXRAD (KTLX) in comparison with Mesonet surface observations in central Oklahoma and also the refractivity "shifts" using reference maps made at different times of the day.

Bodine et al. (2011) attributed these large discrepancies to differences in the height of the radar targets

TABLE 6. Mean estimated total $\langle \Delta N \rangle_{\text{field}}$ (to 30 km) from simulations using spatial smoothing of phase changes after the subtraction of the initial $\langle \Delta N \rangle_{\text{field}}$ estimate as functions of ΔN and phase change noise ($\sigma_{\Delta \phi}$). These simulations are for S-band wavelengths using smoothing kernels from previous publications: C08 as applied to the NEXRAD WSR-88D radars (250-m gate spacing; $|\Delta N|_{\text{folding}} \approx 120$) and F04 as applied to the S-Pol and McGill radars (150-m gate spacing; $|\Delta N|_{\text{folding}} \approx 200$). The standard errors are less than the precision shown in all cases.

(a) Pulse-pair processing (1-gate)							
$\sigma_{\Delta \phi}$ ΔN C08	0	10	20	30	40	50	60	70
10	9.8	9.8	9.8	9.7	9.6	9.4	9.0	8.3
20	19.7	19.6	19.5	19.3	18.8	17.7	14.8	10.2
30	29.4	29.3	29.0	28.3	26.3	20.9	14.5	9.8
40	38.8	38.6	37.7	34.7	27.6	20.7	15.3	10.5
50	46.7	45.7	42.1	34.1	26.9	20.2	13.6	9.0
60	46.2	44.2	38.4	31.4	22.5	15.1	9.6	5.7
70	36.1	32.6	23.9	16.8	10.6	6.5	3.7	2.2
F04								
10	9.8	9.8	9.8	9.7	9.6	9.4	9.1	8.7
20	19.7	19.6	19.5	19.3	19.0	18.5	17.6	15.4
30	29.5	29.4	29.2	28.8	28.2	26.8	22.8	15.0
40	39.2	39.1	38.8	38.2	36.6	31.6	21.6	14.7
50	48.8	48.7	48.2	46.7	42.1	31.1	23.0	17.2
60	58.3	58.0	56.9	53.2	42.3	33.4	26.2	17.7
70	67.3	66.7	63.9	54.6	44.4	36.2	25.7	16.9

(b) Least squares linear fit to phase data averaged over all azimuths							
$\sigma_{\Delta \phi}$ ΔN C08	0	10	20	30	40	50	60	70
10	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
20	19.9	19.9	19.9	19.9	19.9	19.9	19.9	19.9
30	29.9	29.9	29.9	29.9	29.9	29.9	29.9	29.9
40	39.9	39.9	39.9	39.9	39.9	39.9	39.9	39.9
50	49.8	49.8	49.8	49.8	49.8	49.8	49.8	49.8
60	59.8	59.8	59.8	59.8	59.8	59.8	59.7	59.0
70	69.7	69.7	69.7	69.6	69.3	68.3	65.5	56.9
F04								
10	10.0	10.0	10.0	10.0	10.0	10.0	10.0	10.0
20	19.9	19.9	19.9	19.9	19.9	19.9	19.9	19.9
30	29.9	29.9	29.9	29.9	29.9	29.9	29.9	29.9
40	39.9	39.9	39.9	39.9	39.9	39.9	39.9	39.9
50	49.8	49.8	49.8	49.8	49.8	49.8	49.8	49.8
60	59.8	59.8	59.8	59.8	59.8	59.8	59.8	59.8
70	69.8	69.8	69.8	69.8	69.8	69.8	69.8	69.8

relative to the 2.5-m Mesonet observations, as dn/dh between 2.5- and 9-m observations at one of these stations at times exceeded -1000 ppm km^{-1}. These discrepancies often occurred at all Mesonet stations in the radar domain simultaneously and for longer periods than the maximum gradients were observed (e.g., on 29 September in their Fig. 8). Park and Fabry (2010) estimated that the vast majority of ground targets in the

Great Plains were less than 10 m in height based on the spatial extent of ground clutter and propagation calculations. However, discrepancies of 30 N units reported by Bodine et al. would require that the average effective target height (depending on the illumination by the radar beam) approached or exceeded 30 m. Based on the results of simulations in Table 6a, biases of 20 N units could be explained by phase-correlated returns for any of the following combinations: $\Delta N = 40$ and $\sigma_{\Delta \phi} \approx 50°$, $\Delta N = 50$ and $\sigma_{\Delta \phi} \approx 35°$, or $\Delta N = 60$ and $\sigma_{\Delta \phi} \approx 20°$. It is worth noting that the phase change noise due to target height uncertainty is likely to be significant when very large dn/dh occurs, exacerbating the biases presented in this paper. It appears that the biases due to phase-correlated returns could have contributed significantly to the extreme refractivity differences with respect to Mesonet observations reported by Bodine et al.

Such large refractivity biases were not observed at the McGill Radar Observatory in Montreal (Fabry 2004) or with S-Pol in the Oklahoma Panhandle during IHOP (Weckwerth et al. 2005). The simulations presented here suggest that biases are typically smaller for these radars, primarily because of their higher range resolution and the corresponding reduction in the radial extent of phase-correlated returns. In addition, conditions in central Oklahoma are typically warmer and more humid, so both diurnal refractivity changes and those over longer periods are generally larger.

In addition, refractivity retrievals during IHOP were improved by using one of two calibration periods (F. Fabry 2012, personal communication), one typical of conditions by day ($N_{\text{ref}} = 264$) and the other by night ($N_{\text{ref}} = 298$). The improvements were attributed to a reduction in phase change noise due to target height uncertainty by reducing differences in the vertical gradient of refractivity. Although the partition between the two calibration periods was more closely related to differences in dn/dh rather than in refractivity, the use of these two calibration periods is likely to have reduced ΔN between measured fields and the reference times, thereby reducing the bias toward N_{ref}. Other contributing factors could be due to differences in the filter response in the radar receivers or possibly the treatment of weak ground clutter targets in the processing algorithms. For example, it is not clear if the exclusion of sidelobes in F04, which may have serendipitously resulted in the removal of range sidelobes (phase-correlated returns), was included in C08.

8. Conclusions

Significant phase correlations in ground clutter returns may be observed typically associated with large

radial reflectivity gradients surrounding strong "point" targets and at the edges of regions of ground clutter. The magnitude and extent of these correlations is largely determined by the filtering in the receiver: broader filters reduce the phase correlations. The filters are usually tuned to cope with precipitation echoes and could be adjusted to better accept the large change in magnitude of neighboring clutter echoes without degrading their performance for precipitation. Alternatively, dedicated refractivity scans could apply very broad filtering, as clutter returns are typically very strong so the reduction of thermal noise is not important.

Radar refractivity retrievals rely on returns from independent targets to estimate ΔN. The presence of correlated returns will bias estimates of $\langle\Delta N\rangle_{\text{field}}$ toward zero using the standard approach of pulse-pair processing of the unsmoothed phase changes. Large biases were observed (-43%) using data from a C-band weather radar. The simulations suggest that little or no bias would have occurred at S band with such changes (~ 7 N units) because of the wavelength dependence of the sensitivity to both refractivity changes and phase change noise. However, significant biases appear to be evident in published S-band results when much larger refractivity changes have occurred.

Observed phase changes are often extremely noisy and must be smoothed if reliable refractivity fields are to be retrieved. It is likely that target motion often contributes significantly to this noise, though target height variability combined with changes in the vertical gradient of refractivity may be significant at times, particularly at farther ranges. The uncertainty in the precise location of the target within the range gate can also lead to random errors proportional to radar frequency, pulse length (gate spacing), and ΔN; this results in $2°$ phase change noise for $\Delta N = 1$ with a 2-μs pulse (300-m gate) at C band, though just $0.8°$ and $0.5°$ phase change noise for $\Delta N = 1$ with a 1.57-μs pulse (250-m gate; e.g., NEXRAD) and a 1-μs pulse (150-m gate; e.g., S-Pol), respectively, at S band. Smoothing is typically done after the effect of $\langle\Delta N\rangle_{\text{field}}$ has been subtracted from the raw phase change measurements to avoid severely biasing estimated ΔN toward zero. However, smoothing now tends to bias estimated ΔN toward $\langle\Delta N\rangle_{\text{field}}$. If $\langle\Delta N\rangle_{\text{field}}$ is underestimated, the final estimates of ΔN after smoothing may also be underestimated, though to a lesser degree.

Physically based simulations indicate that significant biases occur as ΔN approaches $|\Delta N|_{\text{folding}}$ and as phase change noise increases; these biases are exacerbated by the radial extent of both the range-weighting function and the smoothing kernel. As a result, biases are most severe at short wavelengths and with long pulse lengths.

The simulated biases in $\langle\Delta N\rangle_{\text{field}}$ at S band will depend somewhat on the minimum clutter reflectivity threshold and may not be precise for several other reasons: 1) the underlying clutter field may not be well-represented by the random redistribution of clutter targets used in these simulations, 2) the Gaussian frequency response does not accurately describe the real range-weighting function away from the middle of the pulse, and 3) the transmitted pulse is typically not perfectly rectangular. Since reasonable agreement was found between simulations and observations at C band (for which the real range-weighting function was available), it seems that the first of these reasons has not significantly affected the results, while the later two assumptions should only have a small quantitative impact.

As an alternative to estimating $\langle\Delta N\rangle_{\text{field}}$ using pulse-pair processing, unbiased estimates of $\langle\Delta N\rangle_{\text{field}}$ can generally be achieved with a least squares fit to azimuthally averaged phase change measurements. For example, assuming a true $\langle\Delta N\rangle_{\text{field}} = 60$ and a plausible target motion phase change noise of $30°$, the simulations indicate that biases of 50% (or 30 N units) using pulse-pair processing could be completely avoided with this approach. The least squares approach requires that the azimuthally averaged phase change measurements can be accurately corrected for aliasing. This is only possible for sufficiently robust data (i.e., a large degree of averaging is required), so it is well-suited for estimating $\langle\Delta N\rangle_{\text{field}}$ provided that the maximum distance from the radar is limited. However, this approach is not well-suited for estimating ΔN over small spatial regions (e.g., 4 km $\times$ 4 km), as large errors will occur when accurate dealiasing is not possible.

Refractivity retrievals require estimates of the local perturbations in the refractivity field relative to the mean change: these perturbations are best retrieved using pulse-pair processing. Even when $\langle\Delta N\rangle_{\text{field}}$ is accurately estimated (or more trivially equal to zero), the smoothing of the phase change field will tend to bias refractivity perturbation estimates toward $\langle\Delta N\rangle_{\text{field}}$ and retrievals will exhibit significantly less spatial variability than the true fields. Phase-correlated returns (which may extend over several range gates) have little effect once smoothing has been applied to the raw phase changes as the smoothing kernel typically extends over several kilometers. To reduce these biases, it may be necessary to lessen the extent of smoothing with respect to range, though this will tend to result in noisier retrievals. The best solution is to use a shorter pulse for refractivity retrievals, which has several benefits:

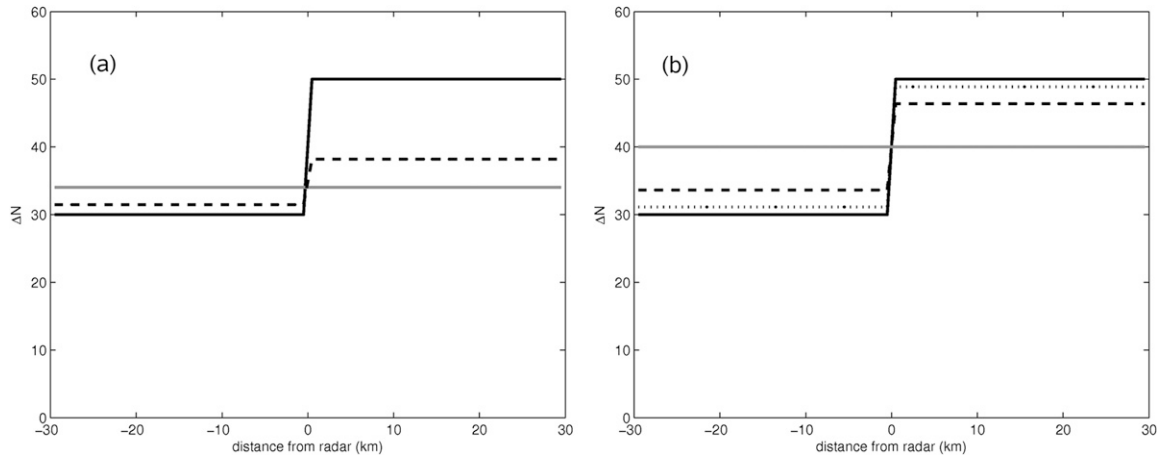1) The greater number of independent clutter targets in a given region will allow the radial extent of the

FIG. 7. Idealized refractivity change transect (black solid lines) of 30 N units to the west and 50 N units to the east of the radar. The mean refractivity change (gray lines) is (a) biased when estimated using pulse-pair processing, though (b) unbiased when the least squares approach is used. The retrieved transects (black dashed lines) are biased toward the mean refractivity change because of the smoothing of phase changes. These biases may be significantly reduced by a second iteration (black dotted line) in (b). Further details can be found in the text.

smoothing kernel to be reduced; this would reduce biases in the retrieved perturbations in the refractivity field.

2) Since $|\Delta N|_{\text{folding}}$ will increase, accurate dealiasing of the azimuthally averaged phase changes for the least squares approach would be possible for larger $\langle\Delta N\rangle_{\text{field}}$. This makes it much easier to use a reference field of constant N obtained weeks or months prior to the observations.

3) The shorter gate length will reduce the phase noise associated with target location uncertainty and thus improve data quality. Target motion phase noise will be unaffected; however, with more targets, the criteria for accepting suitable targets (e.g., Power Ratio) can be made more stringent.

The use of the shorter pulse, combined with subtraction of $\langle\Delta N\rangle_{\text{field}}$ using a least squares fit to the azimuthally averaged phase changes (rather than the pulse-pair technique used at present), should efficiently mitigate phase-correlated returns, leading to much more accurate retrievals both when $\langle\Delta N\rangle_{\text{field}}$ and local perturbations in the refractivity field are large.

Phase change noise due to uncertainty in the target location is proportional to pulse length, radar frequency, and ΔN; it can be significant when ΔN is large, which may occur when comparisons are made with a reference refractivity field. Phase change noise due to target motion increases with radar frequency; at S band the effect is usually tolerable but becomes problematic at C band and may be prohibitive at X band. Considering other sources of phase noise, the use of a constant reference

refractivity field obtained weeks or months earlier may only be achievable at S band. Performance at L band would be better, whereas at C band, retrievals of ΔN may be limited to a few hours and even less at X band. Care must always be taken to ensure that the STALO and transmitter frequencies have not changed since the reference field was measured; this could occur if the radars were temporally switched off.

The problem of dealing with very inhomogeneous ΔN remains a challenge. For example, an idealized transect with $\Delta N = 30$ to the west of the radar and $\Delta N = 50$ to the east is depicted in Fig. 7. In Fig. 7a, $\langle\Delta N\rangle_{\text{field}}$, is estimated to be 34 N units using pulse-pair processing rather than 40 N units (gray lines), which is correctly estimated using the least squares approach (Fig. 7b). Spatial perturbations (estimated with pulse-pair processing) relative to $\langle\Delta N\rangle_{\text{field}}$ will be underestimated because of the smoothing of phase changes (dashed black lines). The biases represented in Fig. 7 correspond to the simulations using the C08 algorithm with 30° additional phase change noise (shown in Tables 5, 6a). To mitigate these biases, an iterative solution may be required. First, for the initial (standard) retrieval, $\langle\Delta N\rangle_{\text{field}}$ is considered the "best-guess" field and the corresponding uniform radial phase change gradient is subtracted before smoothing is applied. For subsequent iterations, the previous retrieval of the spatially variable ΔN is used as the best guess (as opposed to a constant $\langle\Delta N\rangle_{\text{field}}$) and the corresponding radial phase change gradient (now spatially variable) is subtracted prior to smoothing. While biases toward $\langle\Delta N\rangle_{\text{field}}$ may occur in the initial retrieval because of the smoothing of phase changes,

biases should diminish progressively in the second (dotted line in Fig. 7b) and subsequent iterations as the true local perturbations relative to the best-guess field diminish and convergence with the true field is achieved.

Targets such as communication towers and power poles, while previously proposed as the best ground echoes for refractivity retrievals (Fabry et al. 1997), may in fact provide phase-correlated returns. The best ground echoes would be more homogeneous in terms of reflectivity spatially, providing independent returns between adjacent range gates. Although the spatial smoothing applied to phase change measurements effectively negates the biases in ΔN perturbations relative to $\langle \Delta N \rangle_{field}$ because of the phase correlations, greater phase change noise will effectively occur in regions with a significant density of such targets. While highly correlated targets may be censored by increasing the minimum clutter reflectivity threshold, this tends to discard a similar proportion of "good" (weakly correlated) targets. It is relatively simple to identify highly correlated couplets directly and a better approach would be to exclude (significantly) weaker returns in highly correlated couplets from refractivity retrievals.

REFERENCES

Bodine, D., P. L. Heinselman, B. L. Cheong, R. D. Palmer, and D. S. Michaud, 2010: A case study on the impact of moisture variability on convection initiation using radar refractivity retrievals. *J. Appl. Meteor. Climatol.,* **49,** 1766–1778.

——, and Coauthors, 2011: Understanding radar refractivity: Sources of uncertainty. *J. Appl. Meteor. Climatol.,* **50,** 2543–2560.

Cheong, B. L., R. D. Palmer, C. D. Curtis, T.-Y. Yu, D. S. Zrnić, and D. Forsyth, 2008: Refractivity retrieval using a phased array radar: First results and potential for multi-function operation. *IEEE Trans. Geosci. Remote Sens.,* **46,** 2527–2537.

Doviak, R. J., and D. Zrnić, 1979: Receiver bandwidth effect on reflectivity and Doppler velocity estimates. *J. Appl. Meteor.,* **18,** 69–76.

——, and ——, 1993: *Doppler Radar and Weather Observations.* 2nd ed. Dover Publications, 562 pp.

Fabry, F., 2004: Meteorological value of ground target measurements by radar. *J. Atmos. Oceanic Technol.,* **21,** 560–573.

——, C. Frush, I. Zawadski, and A. Kilambi, 1997: On the extraction of near-surface index of refraction using radar phase measurements from ground targets. *J. Atmos. Oceanic Technol.,* **14,** 978–987.

Hubbert, J. C., M. Dixon, S. M. Ellis, and G. Meymaris, 2009: Weather radar ground clutter. Part I: Identification, modeling, and simulation. *J. Atmos. Oceanic Technol.,* **26,** 1165–1180.

Nicol, J. C., K. Bartholemew, T. Darlington, A. J. Illingworth, and M. Kitchen, 2012a: Operational radar refractivity retrieval for numerical weather prediction. *Weather Radar and Hydrology,* R. J. Moore, S. J. Cole, and A. J. Illingworth, Eds., IAHS Publ. 351, IAHS Press, 348–351.

Park, S., and F. Fabry, 2010: Simulation and interpretation of the phase data used by the radar refractivity retrieval algorithm. *J. Atmos. Oceanic Technol.,* **27,** 1286–1301.

Skolnik, M., 1990: *Radar Handbook.* 2nd ed. McGraw-Hill Inc., 1200 pp.

Weckwerth, T. M., C. R. Pettet, F. Fabry, S. J. Park, M. A. LeMone, and J. W. Wilson, 2005: Radar refractivity retrieval: Validation and application to short-term forecasting. *J. Appl. Meteor.,* **44,** 285–300.