



Royal Netherlands
Meteorological Institute
*Ministry of Infrastructure and the
Environment*

Doppler Clutter Removal on KNMI Weather Radars

H. Leijnse, H. Beekhuis, and I. Holleman

De Bilt, 2016 | Technical report; TR-355

Doppler Clutter Removal on KNMI Weather Radars

Hidde Leijnse, Hans Beekhuis, and Iwan Holleman

Technical Report, KNMI TR-355,
January 5, 2016

Contents

1	Introduction	5
2	Methods for clutter removal	7
2.1	KNMI's statistical scheme	7
2.1.1	Clutter signal processing	7
2.1.2	Clutter flag processing	9
2.2	Doppler clutter removal	11
2.2.1	Time-domain and frequency-domain filtering	14
2.3	Speckle filtering	15
3	Operational radar data	17
3.1	Volume coverage pattern	17
3.2	Available data for comparison	18
4	Comparison of clutter removal schemes	21
4.1	Analyses of two cases	21
4.1.1	Severe anaprop clutter on 30 March 2014	22
4.1.2	Rainfall event on 7 April 2014	25
4.2	Objective verification using cloud masks	28
4.3	IQ-replay	30
5	Summary and next steps	33
	References	37

Chapter 1

Introduction

To support the KNMI mission the institute operates a weather radar network. The radar products of the network are essential for several tasks of KNMI, such as monitoring the meteorological conditions, performing research, and issuing forecasts and alerts for hazardous weather to the Dutch society. The radar network consists of two C-band Doppler weather radars which cover the whole of the Netherlands. These radars, together with radar data from neighboring countries, are primary input to the radar product processor that in turn generates various products and disseminates them to several groups of users. These radar products are considered primary geophysical meteorological data.

In 2014 KNMI started a project called “Radar Sensor Replacement” (RASER) in which both weather radars will be renewed. In the RASER project the following goals will be achieved:

- An enhanced quality of the radar observations and products by using polarimetric weather radars.
- An increased update frequency of the primary radar products to once every 2.5 minutes.
- A maximum assured availability by replacing the two radar sensors with modern low-maintenance radars.
- A reduced maintenance effort for KNMI staff by outsourcing most of the maintenance of these radars.

During 2015 the tendering procedure and contract negotiations were completed and the actual installation of the new C-band polarimetric weather radars is foreseen in 2016.

As part the Digital Radar Upgrade (DRUP) of the KNMI weathers in 2006 the statistical clutter removal scheme (Wessels and Beekhuis, 1992, 1994) was reviewed by Holleman and Beekhuis (2005). The clutter removal scheme was transformed from a method for 2-dimensional Cartesian images to a method for 3-dimensional polar data. Currently the lowest elevation of the operational KNMI volume scan is filtered using this transformed statistical clutter removal scheme while all other elevations are filtered using a standard Doppler-based scheme. It should be noted, however, that the lowest elevation is the most important one for the standard radar surveillance product and for radar-based quantitative precipitation estimation. In this Technical Report the performances of the statistical clutter removal scheme and the standard Doppler-based scheme are compared for the lowest radar elevations. If the Doppler scheme performs at least as good as the statistical scheme, there is no need to port the non-standard statistical clutter removal scheme of KNMI to the new weather radars and this would lead to a considerable simplification of the radar sensor replacement project.

In this Technical Report the outcome of the comparison between the statistical and Doppler clutter removal schemes and the further evaluation of the Doppler scheme are described. It is concluded that the standard Doppler-based scheme performs somewhat better than the statistical scheme and thus it is recommended to use Doppler clutter filtering for all elevations of the KNMI volume scan. The outline of the remaining of the report is as follows:

- In chapter 2 descriptions of the KNMI statistical clutter removal scheme and the Doppler-based clutter removal and filtering schemes are given. Furthermore the one-dimensional and two-dimension speckle filters, which are standard features on most radar signal processors, are briefly introduced.
- The characteristics of the current 5-minute volume scan of the KNMI operational Doppler weather radars are presented in Chapter 3.
- In chapter 4 a comparison of the statistical and Doppler schemes based on two cases, a case with severe anomalous propagation ('anaprop') and a case of intense rainfall is described. Furthermore an objective comparison of the two clutter removal schemes using the cloud mask product from Meteosat, and an analysis of different types of Doppler filters based on IQ data are presented.
- In the last chapter the conclusions are summarized and recommendations for the next steps are made.

Chapter 2

Methods for clutter removal

In this chapter the KNMI statistical clutter removal scheme and the standard Doppler-based clutter removal schemes are described in more detail. Furthermore the one-dimensional and two-dimension speckle filters, which are standard features on most radar signal processors, are briefly introduced.

2.1 KNMI's statistical scheme

In the early 90s, a stepwise procedure for the rejection of (anomalous propagation) clutter was developed for the C-band weather radars of KNMI (Wessels and Beekhuis, 1992, 1994). This procedure is based on distinguishing between the inherently fluctuating Rayleigh-scattered precipitation signals and the relatively stable ground clutter signals (Aoyagi, 1983). A dynamical clutter map is constructed from the “fluctuation” flags using a spatial averaging procedure and a decision-making model. Anomalous propagation clutter over land is removed almost completely, while the system is partly (40%) effective in removing sea clutter. Holleman and Beekhuis (2005) modernized the KNMI clutter removal scheme in view of the digital radar upgrade. A brief description is given below.

2.1.1 Clutter signal processing

During the clutter signal processing the fluctuations of the received power (echoes) within each processed range bin are analyzed. In Figure 2.1 examples of the standard deviation spectra of the raw power samples for areas with precipitation and clutter are shown. The precipitation spectrum peaks around a standard deviation of 4.2 dB and the clutter spectrum peaks around 2.0 dB. The echoes from precipitation areas are caused by Rayleigh scatter-

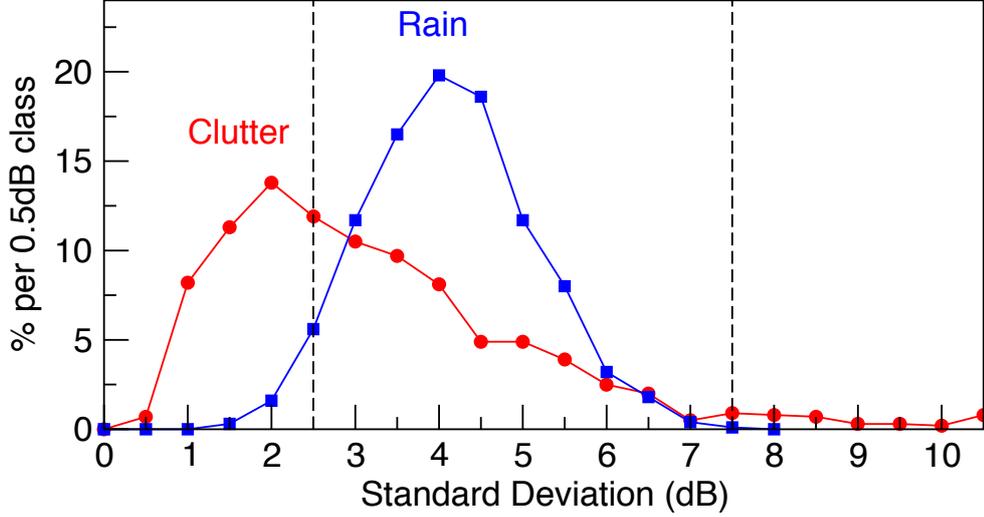


Figure 2.1: Examples of standard deviation spectra for precipitation and anomalous propagation clutter. The vertical dashed lines mark the thresholds that are used to make the distinction between precipitation and clutter in the first step of the algorithm (Figure taken from Wessels and Beekhuis (1994)).

ing and for this scattering process from random targets a standard deviation of 5.5 dB is expected from theory. The observed standard deviation is somewhat lower which is probably caused by the (small) correlation of the samples in azimuthal direction. The vertical dashed lines mark the thresholds that are used to make the distinction between precipitation and clutter signals. Evidently there is a large overlap between the standard deviation spectra for precipitation and clutter, and thus the separation of precipitation and clutter signals based on this criterion is not good enough: about 45% of the clutter signals will be detected (Wessels and Beekhuis, 1994). Information from neighboring samples is, therefore, needed to improve the performance.

The samples of the received power in dBm calculated from the linear In-phase (I) and Quadrature-phase (Q) channels and corrected for the background noise are the input for the clutter signal processing. From these raw power samples the 5-point running average P_5 is calculated. The standard deviation of the power samples σ_p is calculated from the actual values $P(r, \alpha)$ and the 5-point running average:

$$\sigma_p \equiv \sqrt{\frac{1}{N_a} \sum_{\alpha=0}^{N_a-1} [P(r, \alpha) - P_5(r, \alpha)]^2} \quad (2.1)$$

Table 2.1: Default values of the parameters for the statistical clutter signal processing.

Parameter	Value [Unit]
PRF	250 Hz
Azimuthal speed	18 deg/s
Minimum std. dev.	2.5 dB
Maximum std. dev.	7.5 dB
Range bin size	1 km \times 1 deg
Samples per bin	8
Clutter threshold	3 or 4

where r and α indicate the range and azimuth, respectively, of the processed range sample and N_a is the number of samples in azimuthal direction. With a PRF of 250 Hz and an azimuth speed of 18 degrees/s, about 14 independent raw power samples per 1 degree are available for calculation of the standard deviation. Using fixed minimum threshold σ_{min} and maximum threshold σ_{max} , the clutter flag for a certain range sample is set depending on the observed standard deviation (see Figure 2.1).

The clutter flag for the processed range bin is based on the flags for the underlying range samples. Up to 8 samples per processed range bin can be obtained and therefore the clutter flag $C(R, A)$ for a bin at range R and azimuth A is set according to:

$$C(R, A) = \begin{cases} 0 & \text{if } \sum_{\rho=0}^7 c(r + \rho, A) < C_t \\ 1 & \text{else} \end{cases} \quad (2.2)$$

where C_t represents the clutter threshold which is either 3 or 4.

2.1.2 Clutter flag processing

In Figure 2.2 a schematic view of the operational spatial averaging procedure of the reflectivity and clutter flags is presented. The reflectivity flags for the processed range bins N_R are set when the observed reflectivity exceeds a certain reflectivity threshold dBZ_t (default 1 dBZ). The clutter flags T_R are copied from Equation 2.2 but are only considered when the corresponding reflectivity flag is set. The central range bin and its two neighbors are marked in black and almost black, respectively, in the figure, and the inner and outer areas are colored with dark gray and light gray, respectively. The outer area is measuring $\Delta R_{out} \times \Delta A_{out}$ and the inner area is measuring $\Delta R_{in} \times \Delta A_{in}$.

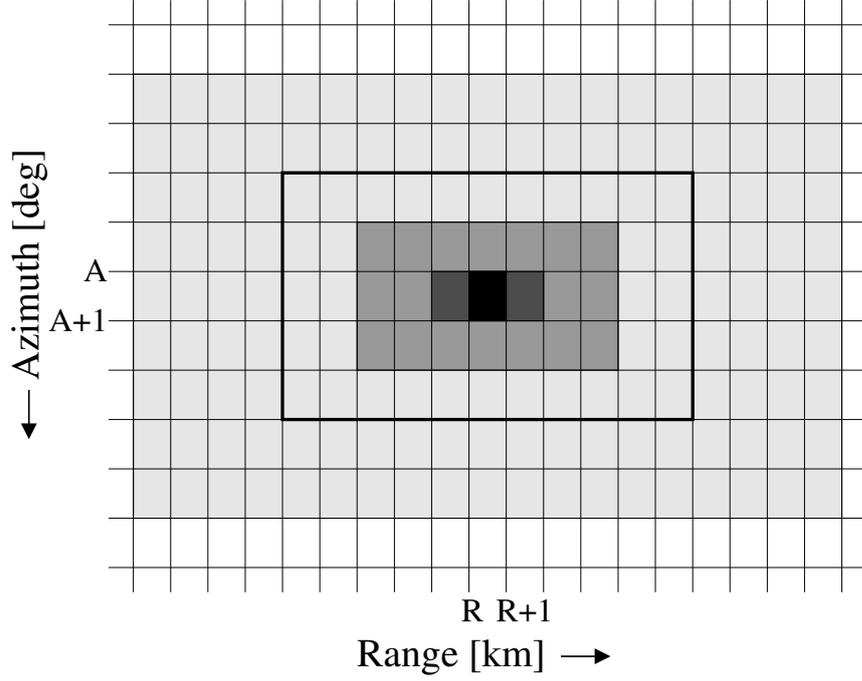


Figure 2.2: Schematic view of the operational spatial averaging procedure of the reflectivity and clutter is presented. The central range bin and the two neighbors are marked in black and almost black, respectively, and the inner and outer areas are colored with dark gray and light gray, respectively. The bold square indicates the area where the horizontal fluctuations of the reflectivity pattern is analyzed.

The default values are listed in Table 2.2. The reflectivity and clutter flags for the range bins are summed into N and T , approximating a Bell-shape. The sum of the reflectivity flags N is made over the inner and outer areas around each processed range bin:

$$N = \sum_{\rho=-\delta R_o}^{\delta R_o} \sum_{\alpha=-\delta A_o}^{\delta A_o} N_R(\rho, \alpha) + \sum_{\rho=-\delta R_i}^{\delta R_i} \sum_{\alpha=-\delta A_i}^{\delta A_i} N_R(\rho, \alpha) \quad (2.3)$$

where the summation limits $\delta R_{i,o}$ and $\delta A_{i,o}$ are defined by:

$$\delta A_{i,o,s} = (\Delta A_{in,out,std} - 1)/2 \quad (2.4)$$

$$\delta R_{i,o,s} = (\Delta R_{in,out,std} - 1)/2. \quad (2.5)$$

In the same way, the sum of the clutter flags T is made over the inner and outer areas around each processed range bin:

$$T = \sum_{\rho=-\delta R_o}^{\delta R_o} \sum_{\alpha=-\delta A_o}^{\delta A_o} T_R(\rho, \alpha) + \sum_{\rho=-\delta R_i}^{\delta R_i} \sum_{\alpha=-\delta A_i}^{\delta A_i} T_R(\rho, \alpha). \quad (2.6)$$

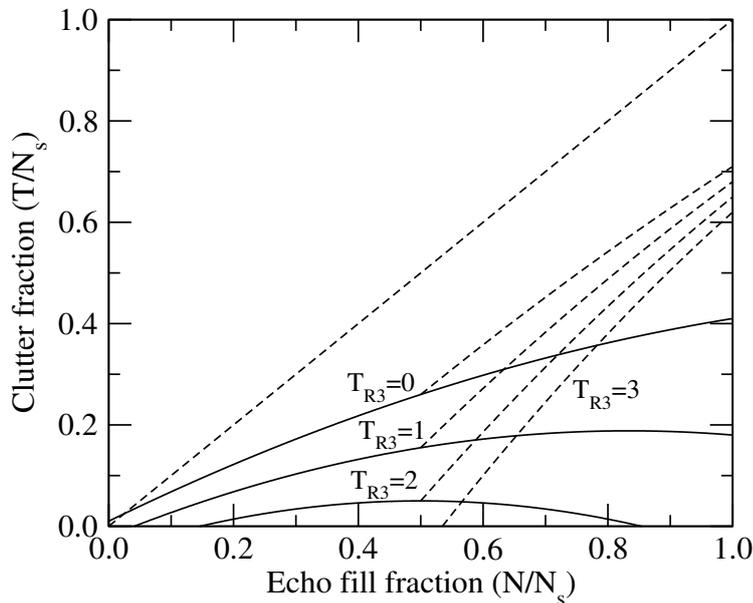


Figure 2.3: The threshold curves of the current clutter decision making model are displayed in this figure. Four different threshold curves, one for each possible value of the triple clutter flag of the central bin and its two neighbors T_{R3} , are presented.

In the above summations for N and T , a periodic boundary condition in the summations over α is applied again and the last value of ρ is assumed to be valid beyond the natural limits (zero and maximum range). The maximum possible value of N and T is denoted as N_s .

The threshold curves of the decision-making model are shown in Figure 2.3. The underlying formula and parameters can be found in Holleman and Beekhuis (2005). The final clutter flag for the corresponding processed range bin is set when the clutter fraction T/N_s (on y-axis) for a certain echo-fill fraction N/N_s (on x-axis) is higher than the threshold curve. The clutter flags T_R of the two neighboring range bins (see Figure 2.2) are added to obtain a triple clutter flag $T_{R3} = 0, 1, 2, 3$ for the central range bin. Depending on the value of this triple flag the appropriate threshold curve is selected and the decision clutter yes or no is made.

2.2 Doppler clutter removal

The power spectrum, often called the Doppler spectrum, contains the information on the distribution of radial velocities of the meteorological scatterers

Table 2.2: Default values of parameters for the proposed clutter flag processing. The dimensions of the areas are given in number of processed range bins ($1 \text{ km} \times 1 \text{ deg}$).

Parameter	Value	Parameter	Value
ΔA_{in}	3	ΔR_{in}	7
ΔA_{out}	9	ΔR_{out}	19
ΔA_{std}	5	ΔR_{std}	11
dBZ_t	1 dBZ	η_t	0.5
σ_t	3.5 dBZ		

in the power-weighted measurement volume. From the shape of the Doppler spectrum, the hydrometeor parameters, most notably the mean radial velocity and spectral width, can be deduced. In Figure 2.4 an example of a Doppler spectrum from hydrometeor scatterers is depicted (Doviak and Zrnić, 1993). The unambiguous velocity interval of this Doppler spectrum is $\pm 28.5 \text{ m/s}$ and the received power for each velocity bin is given in decibel relative to the peak power. The Doppler spectrum has a clear maximum around a velocity of 15 m/s which is equal to the mean radial velocity. The spectral width of 2.2 m/s is determined from a Gaussian fit to the power spectrum. In addition it is evident from the figure that the noise level is about 25 dB below the peak power. Ground clutter, i.e., signal from fixed non-hydrometeor targets, can easily be recognized in a Doppler spectrum. An example of a Doppler spectrum of ground clutter is depicted in Figure 2.5 which again is taken from Doviak and Zrnić (1993). It is obvious that ground clutter produces a narrow peak centered around zero velocity. The spectral width of this clutter peak is only 0.45 m/s and thus it is a factor of four smaller than that of the hydrometeor peak.

For a Doppler weather radar at least three types of contributions to the power spectrum can be distinguished: noise, ground clutter, and hydrometeor signal. The relative importance of each of these contributions depends on the range from the radar and the actual meteorological circumstances. In order to process the received Doppler signal, assumptions on the spectral shape of these contributions have to be made. For white noise the power spectrum $S(f)$ is constant and equal to:

$$S(f) \equiv T_n N_0 \quad (2.7)$$

$$R(n) = \begin{cases} N_0 & \text{for } n = 0 \\ 0 & \text{for } n \neq 0 \end{cases} \quad (2.8)$$

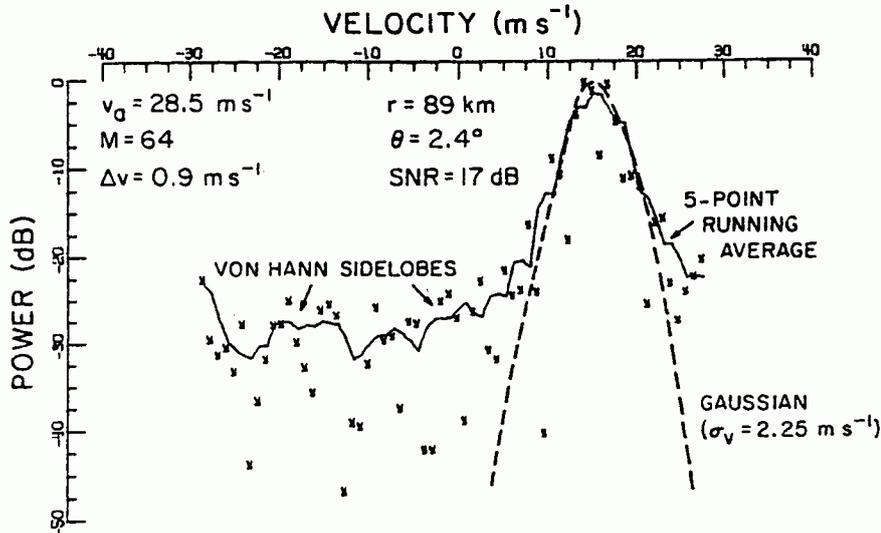


Figure 2.4: A typical Doppler spectrum for a resolution volume within a thunderstorm. This spectrum is obtained from a Fourier transform using 64 samples and a Von Hann window. This figure is taken from Doviak and Zrnić (1993).

where T is the pulse repetition time and N_0 the noise power. The corresponding autocorrelation function $R(n)$ for white noise is obtained via an inverse discrete Fourier transform. The autocorrelation function is zero for $n \neq 0$ reflecting that white noise is completely uncorrelated. The power spectrum of a ground clutter signal can be approximated by a Gaussian function centered around zero frequency:

$$S(f) \equiv \frac{C_0}{\sigma_c \sqrt{2\pi}} \exp(-f^2/(2\sigma_c^2)) \quad (2.9)$$

$$R(n) = C_0 \exp(-(2\pi n T \sigma_c)^2/2) \quad (2.10)$$

In these equations for the power spectrum and the autocorrelation function, C_0 and σ_c represent the power and the spectral width, respectively, of the ground clutter signal. For a hydrometeor scattering signal, the power spectrum is well approximated by a Gaussian function centered on the mean Doppler frequency shift f_d which is directly related to the mean radial velocity of the scattering hydrometeors (Doviak and Zrnić, 1993). The power spectrum and autocorrelation function for the hydrometeor signal are described by:

$$S(f) \equiv \frac{S_0}{\sigma \sqrt{2\pi}} \exp(-(f - f_d)^2/(2\sigma^2)) \quad (2.11)$$

$$R(n) = S_0 \exp(i2\pi n T f_d) \exp(-(2\pi n T \sigma)^2/2) \quad (2.12)$$

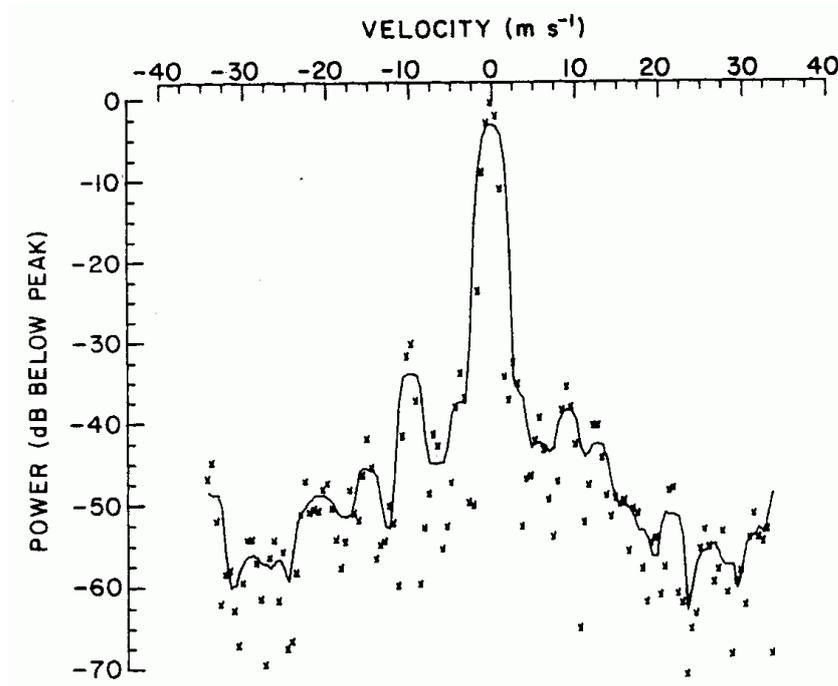


Figure 2.5: A Doppler spectrum of ground clutter for an antenna scanning at 10 degree/s. The Doppler processing is similar that of the data in Figure 2.4. The ground clutter peak is centered at zero velocity and its spectral width is only 0.45 m/s. This figure is taken from Doviak and Zrnić (1993).

where S_0 is the scattered power and σ is the spectral width. Any Doppler spectrum can be described by a linear combination of the preceding contributions. As a matter of fact this also holds true for the autocorrelation functions, because the conversion between the power spectrum and the autocorrelation is linear.

2.2.1 Time-domain and frequency-domain filtering

It has been shown in Figure 2.5 that ground clutter gives rise to a narrow peak around zero frequency in the Doppler spectrum. This low frequency component can be removed from the Doppler signal by applying a steep high-pass filter. Digital filtering of the Doppler signal can either be done in the frequency domain or in the time domain. For the latter the filtered signal $\tilde{Z}(n)$ is given by:

$$\tilde{Z}(n) = \sum_{m=0}^M b_m Z(n-m) - \sum_{m=1}^M c_m \tilde{Z}(n-m) \quad (2.13)$$

where M is the order of the time-domain filter and $Z(n)$ the unfiltered time-series. It is evident that the filtered signal is a linear combination of the (un)filtered signals at previous time steps. When all coefficients c_m are zero the filter is a so-called finite impulse response filter, and else it is a so-called infinite impulse response filter. More information on digital filtering in the time domain can be found in Press et al. (1992). The filter characteristics and function depend on the choice of coefficients b_m and c_m . It is important to note that, in contrast to frequency domain filtering, time domain filtering does not allow for reconstruction of the power spectrum around zero frequency which can lead to loss of weather signal. Moreover, time domain filters require a settling time after e.g. a change in Pulse Repetition Frequency and therefore the first pulses can not be used.

Nowadays all radar signal processors have sufficient computational power to perform frequency domain filtering. For this the timeseries of pulses at a certain range is transformed to a spectrum using a discrete frequency domain filter, then the intensity around zero frequency (potentially due to clutter) is removed, and finally the weather spectrum is reconstructed using interpolation. After back-transformation the radar quantities, like power, reflectivity, radial velocity, etc., can be estimated using the standard procedures. A minimum number samples per range bin is required for a proper functioning of the frequency domain clutter filters, typically at least 32 samples.

Usually the ratio between the received power before and after filtering is monitored and when this ratio is too high, indicating that the spectrum is dominated by clutter, the corresponding range bin value is rejected. The maximum clutter correction ratio is typically in the range between 10 and 20 dB, implying that Doppler radar data with 90 to 99% clutter power can still be reconstructed.

2.3 Speckle filtering

All radar signal processors have the option to perform one-dimensional speckle filtering on the volume data and many of them also offer a two-dimensional speckle filtering. The term 'speckle' is used for isolated range bins with valid data that are surrounded by bins with no data. A speckle filter removes these isolated range bins by setting them to 'no data' and thus cleans the raw radar volume data. The most common implementation is a one-dimensional speckle filter which only considers the radial direction ('rays') and checks whether a range bin with valid data has a one or two valid (before and/or after) neighbors.

More advanced two-dimensional speckle filters evaluate the number of

valid neighbors and thus decide to remove the valid range bin or not. The precise implementation of such a two-dimensional speckle filter depends on the brand of the radar signal processor.

Chapter 3

Operational radar data

In this chapter the volume coverage pattern of the operational radar scans of KNMI and the available data for the comparison of the clutter removal schemes are described.

3.1 Volume coverage pattern

Originally the KNMI radars performed three interlaced volume scans with different characteristics: a 4-elevation reflectivity scan which was repeated every 5 minutes, a 14-elevation reflectivity scan every 15 minutes, and a 3-elevation Doppler scan which was repeated every 15 minutes. The settings for each elevation within a volume scan were fixed due to limitations of the old radar scan controller. The Digital Radar Upgrade (DRUP) in 2006 made it possible to define the settings for each elevation in a volume scan independently. Thus a single volume coverage pattern with optimized settings for all applications was developed and this was actually one of the major goals of the DRUP project. Since then the operational scanning of the KNMI weather radars generates a 14-elevation volume dataset every 5 minutes. Table 3.1 lists the settings of the current volume coverage pattern of the operational weather radar scan of KNMI. Apart from the scan elevation, antenna rotation speed, high and low pulse repetition frequency (PRF), and the unambiguous velocity also the selected maximum range, range bin length, and clutter removal method are listed. The rationale for the scan settings is that with increasing elevation a shorter radar range is required as the tropopause is reached sooner. As a consequence higher pulse repetition frequencies with higher unambiguous velocities and higher antenna rotation speeds can be used. In this way the space-velocity coverage and recording time of the radar volume scan is optimized.

Table 3.1: Settings of the current volume coverage pattern of the operational radar scan of KNMI. For each elevation the applied clutter removal method is indicated: statistical (S), Doppler time domain (T), or Doppler frequency domain (F).

N	Elev [deg]	Speed [deg/s]	PRF h/1 [Hz]	Velocity [m/s]	Range [km]	Bin [km]	Clutter
1	0.3	18	250	3.3	320	1.0	S
2	0.4	18	600/450	24.0	240	1.0	T
3	0.8	18	600/450	24.0	240	1.0	T
4	1.1	18	600/450	24.0	240	1.0	T
5	2.0	18	600/450	24.0	240	1.0	T
6	3.0	24	800/600	32.0	170	0.5	T
7	4.5	24	800/600	32.0	170	0.5	T
8	6.0	30	1000/750	40.0	150	0.5	T
9	8.0	30	1000/750	40.0	150	0.5	T
10	10.0	36	1200/900	48.0	120	0.5	T
11	12.0	36	1200/900	48.0	120	0.5	T
12	15.0	36	1200/900	48.0	120	0.5	T
13	20.0	36	1200/900	48.0	120	0.5	T
14	25.0	36	1200/900	48.0	120	0.5	T

In Figure 3.1 a graphical representation of the volume coverage of the KNMI radar scan is given. It is evident that the elevations are distributed evenly so that the range-height plane is covered rather homogeneously. The width of the radar beam, about 1 degree for the KNMI radars, is not visualized in the figure but makes that the low elevations are actually overlapping. This is in particular true for the two lowest elevations which are only 0.1 degree apart.

3.2 Available data for comparison

The volumetric data produced during the operational scanning of the KNMI radars in De Bilt and in Den Helder are archived. As the two lowest elevations are only 0.1 degree apart and recorded within 25 seconds for each other they can be used for the comparison of the clutter removal schemes. The lowest elevation at 0.3 degree is recorded using KNMI's statistical clutter removal scheme while the second elevation (0.4 degree) is recorded using Doppler-

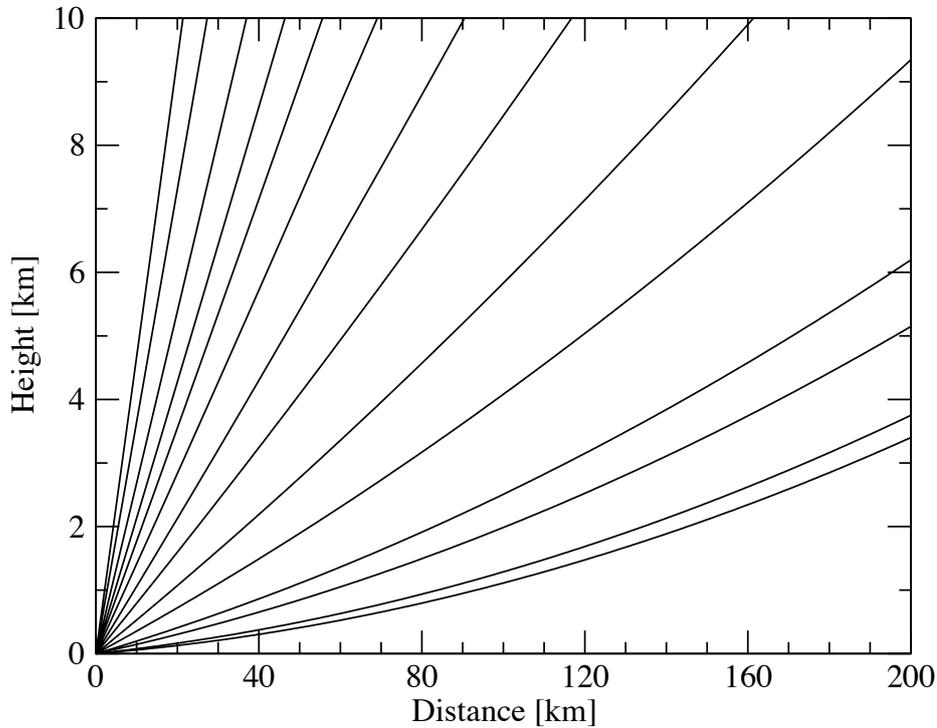


Figure 3.1: Graphical representation of the 14-elevation volume coverage pattern of the operational radar scan of KNMI. This scan is repeated every 5 minutes and thus 288 volume scans are made every day.

based clutter removal. These data are the basis for the comparison presented in the remainder of this report.

Furthermore datasets with raw IQ-samples are available for approximately one week in 2012 and these data can be employed to replay the radar signal processing with different settings of the clutter removal algorithms. Naturally the physical settings of the radar scans, like elevation and PRF, cannot be changed. The radar signal processor (SELEX ES GmbH) has the option to reprocess archived datasets with raw IQ-samples.

For the objective verification of the clutter removal methods Cloud Mask products from the Meteosat satellite from EUMETSAT are used. This product has been developed in the Satellite Application Facility on support to Nowcasting and Very Short-Range Forecasting (SAFNWC). Cloud Mask products over a 10-day period from 30 March till 8 April 2014 are used for this purpose.

Chapter 4

Comparison of clutter removal schemes

The difference between the statistical filter and the time-domain Doppler filter is analyzed by comparing the lowest two elevations in the operational volume scan schedule. These two elevations are performed with a time difference of approximately 25 seconds and an elevation angle difference of 0.1° (relative to a 1° beam width). The main difference between these two scans is that the first one uses the statistical clutter filter and the second one uses a time-domain Doppler filter. For the purpose of evaluating these different clutter filtering methods, we will assume that this is the only difference between these scans.

Both pseudoCAPPI and low-level PPI data will be used in the evaluation of the clutter removal schemes. The pseudoCAPPI data that KNMI produces operationally is used by many parties inside and outside the institute, and hence it is an important product to evaluate. Low-level PPI data are important for hydrological applications, and the corrections that are necessary in order to use radar data for this purpose require the data to be relatively clutter-free (Hazenbergh et al., 2011, 2013).

4.1 Analyses of two cases

There are two aspects to the performance of clutter filters: (1) it should remove most echoes not related to precipitation; and (2) it should not affect echoes of precipitation. This is why two different cases are studied here: an event with severe clutter caused by anomalous propagation ('anaprop') and a rainfall event with limited clutter.

4.1.1 Severe anaprop clutter on 30 March 2014

On 30 March 2014 there were severe anomalous-propagation conditions above the Netherlands, causing very intense land and sea clutter. Figures 4.1 and 4.2 show low-level PPIs and composite images, respectively at 06:00 UTC. This was the time when the anaprop clutter was at its worst.

It is clear from Fig. 4.1 that both the statistical filter and the time-domain Doppler filter remove a large part of the clutter. The clutter that remains is predominantly located over the sea. This is to be expected in case of the statistical filter because anomalous propagation clutter over water tends to be more widespread in nature than clutter over land, which makes the statistical filter (which relies on small-scale spatial inhomogeneities) fail. Another difference between land and sea clutter is that sea clutter generally has a non-zero velocity (i.e., waves move). This typically results in failure of Doppler clutter filters (which rely on clutter targets having a near-zero velocity) over sea.

For both radars it can be seen that the Doppler filter is the most effective in removing clutter for this case. Note that the statistical filter is designed such that it does not remove any data closer than 14 km from the radar (see Holleman and Beekhuis, 2005).

The added value of applying a speckle filter depends on how effective the statistical- or Doppler filters have been in removing speckled clutter. It is apparent from Figure 4.1 that the speckle filter is more effective for the De Bilt radar (top two rows) than for the Den Helder radar (bottom two rows). This is due to the fact that the remaining clutter is more homogeneous for the Den Helder radar because the sea clutter is more severe for this radar (see the left-hand column of Fig. 4.1). The speckle filter is also more effective for the Doppler-filtered data than for statistical-filtered data. This is caused by the nature of the speckle filter which is more similar to a statistical filter than to a Doppler filter. Because the statistical filter has already removed most of the speckle, application of an additional speckle filter will not greatly improve the results.

From the left-hand column of Fig. 4.1 it is also apparent that the assumption that the 0.3° and 0.4° scans are the same apart from the applied clutter filter does not always hold. This is especially true for the Den Helder radar (bottom two rows), see the strong echoes southwest of the radar above the North Sea. The cause of this may be the difference in elevation angles, pulse repetition frequencies (250 Hz and 450/600 Hz, respectively) and pulse lengths ($2.0 \mu\text{s}$ and $0.8 \mu\text{s}$, respectively) employed for these two scans. Another cause could be that the propagation conditions have changed in the 25 seconds between the two scans. It is therefore important to also look

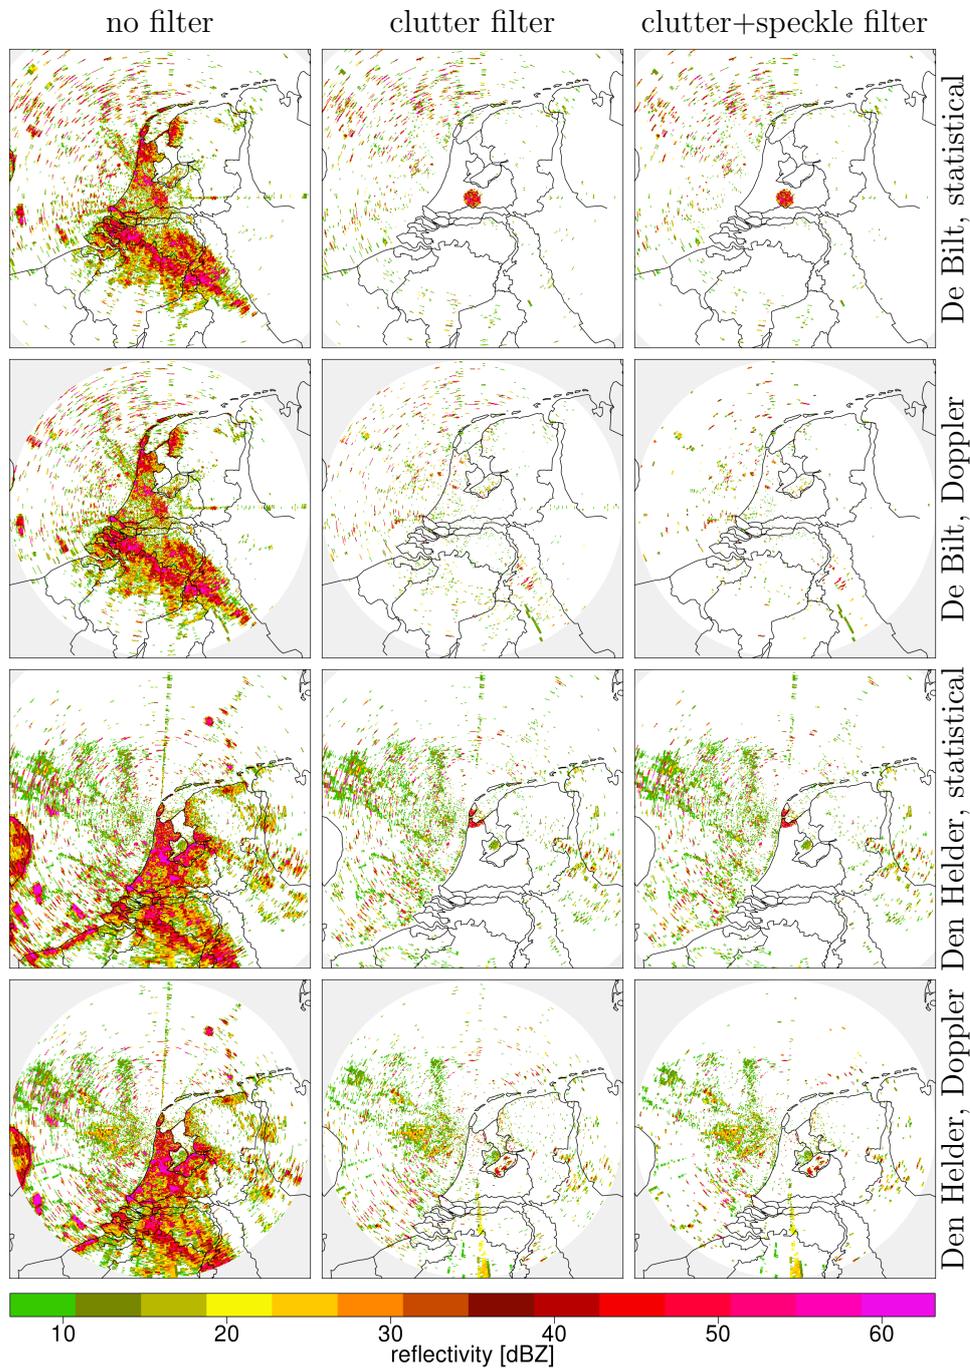


Figure 4.1: Images of low-level PPIs for the severe clutter event on 30 March 2014 at 06:00 UTC. Images in the top row of this figure are based on pseudoCAPPPIs made using the lowest elevation with a statistical filter, whereas a time-domain Doppler filter has been used for the second row. The third and fourth rows are the same as the first and second rows, respectively, but for the Den Helder radar. Images in the left, center, and right columns are made without filters, with a clutter filter, and with a clutter and speckle filter, respectively.

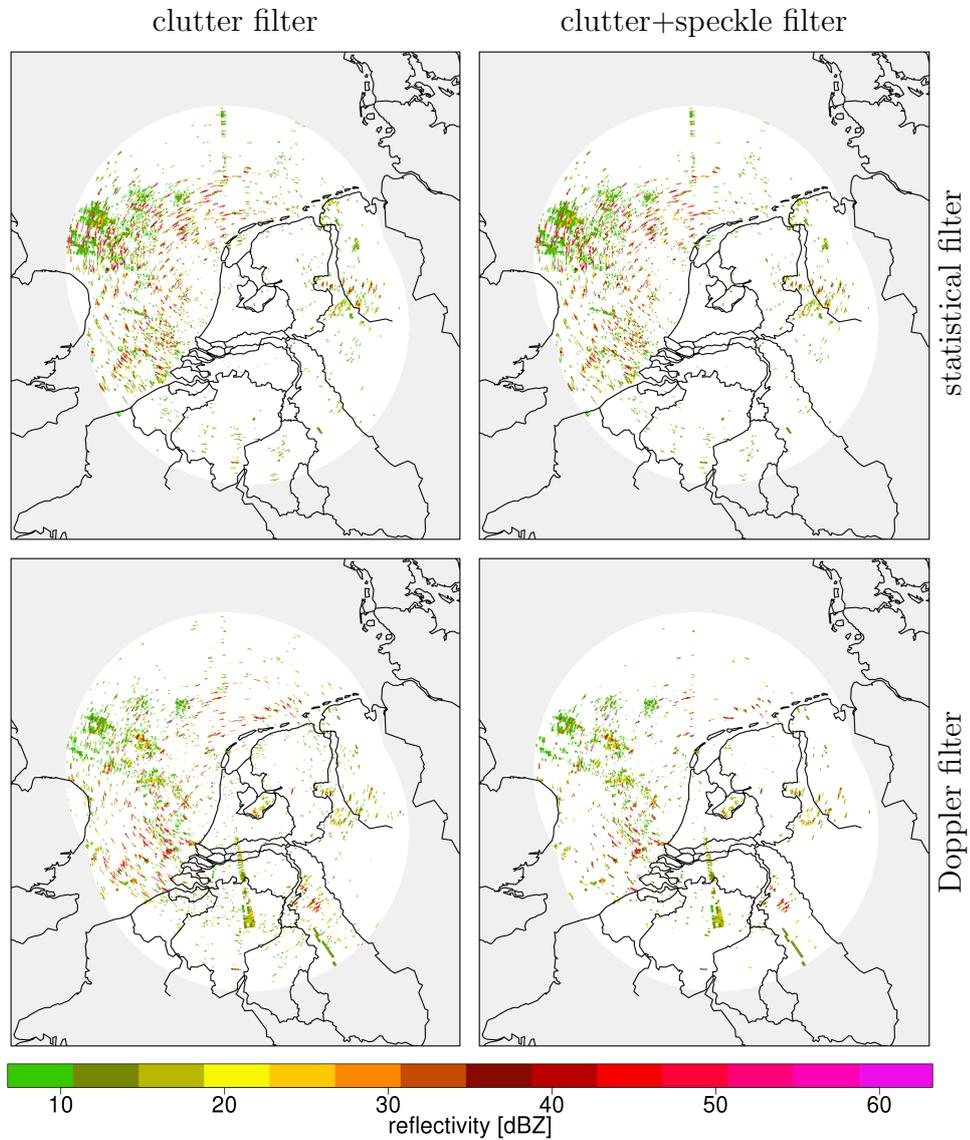


Figure 4.2: Images of composites of pseudoCAPPs for the severe clutter event on 30 March 2014 at 06:00 UTC. Images in the top row of this figure are based on pseudoCAPPs made using the lowest elevation with a statistical filter, whereas a time-domain Doppler filter has been used for the bottom row. Images in the left column are made without a speckle filter, and a speckle filter is applied to the volume data underlying the images in the right column.

at the uncorrected data when comparing these two scans (left column) and there the differences in echo-strengths are also seen. The apparent poorer performance of the Doppler clutter filter on the Den Helder radar is thus explained by the use of different scans.

Looking at KNMI's most widely used radar product, the composite of 1500-m pseudoCAPPIs (see Fig. 4.2), the conclusions drawn from the PPI data (Fig. 4.1) are confirmed: The Doppler filter is more effective in removing clutter than the statistical filter and the speckle filter is most effective when applied after the Doppler filter. At short range, data from higher elevations (on which Doppler filters have been applied) are used to generate the pseudoCAPPI products, and the compositing of these products is done such that the weight is very low at locations close to the radar. Therefore, clutter at these locations is much less severe than in Fig. 4.1.

4.1.2 Rainfall event on 7 April 2014

On 7 April 2014 a line of precipitation moved over the Netherlands from west to east. Figures 4.3 and 4.4 show low-level PPIs and composite images, respectively, at 16:00 UTC when the line of precipitation arrived at the Dutch coast. Because there is very little clutter in these reflectivity data, the discussion of these figures is focused on how much the different clutter filters affect the precipitation signals. However, it is clear from these images that both the statistical and the Doppler filters remove the sidelobe clutter close to the radar (for the statistical filter at ranges further than 14 km, see Holleman and Beekhuis, 2005).

The right-hand column of Fig. 4.3 shows the ratio of the unfiltered and the filtered signals (in dB). The figure only shows values where pixels have not entirely been removed. It is clear that the Doppler filter can remove more than 5 dB from the precipitation signal near the zero isoDop (red pixels in left column images). This corresponds to reducing the estimated rainfall intensity by a factor of 3 or more (when applying the Marshall-Palmer $Z - R$ relation). In areas away from the zero isoDop, the reduction of the precipitation signal is still more than 1 dB (green pixels in left column images), which corresponds to a 15% reduction in precipitation intensity. Because the statistical filter only removes pixels if they are identified as clutter, this filter does not affect the intensity of the rainfall. It is clear from Fig. 4.3 that the statistical filter does not remove many pixels in these rain areas.

The Doppler filter removes the part of the received signal power that has near-zero velocities. When the wind is perpendicular to the line between the radar and the pixel of interest, the radial component of the wind will be close to zero. This causes the Doppler filter to remove a relatively large fraction

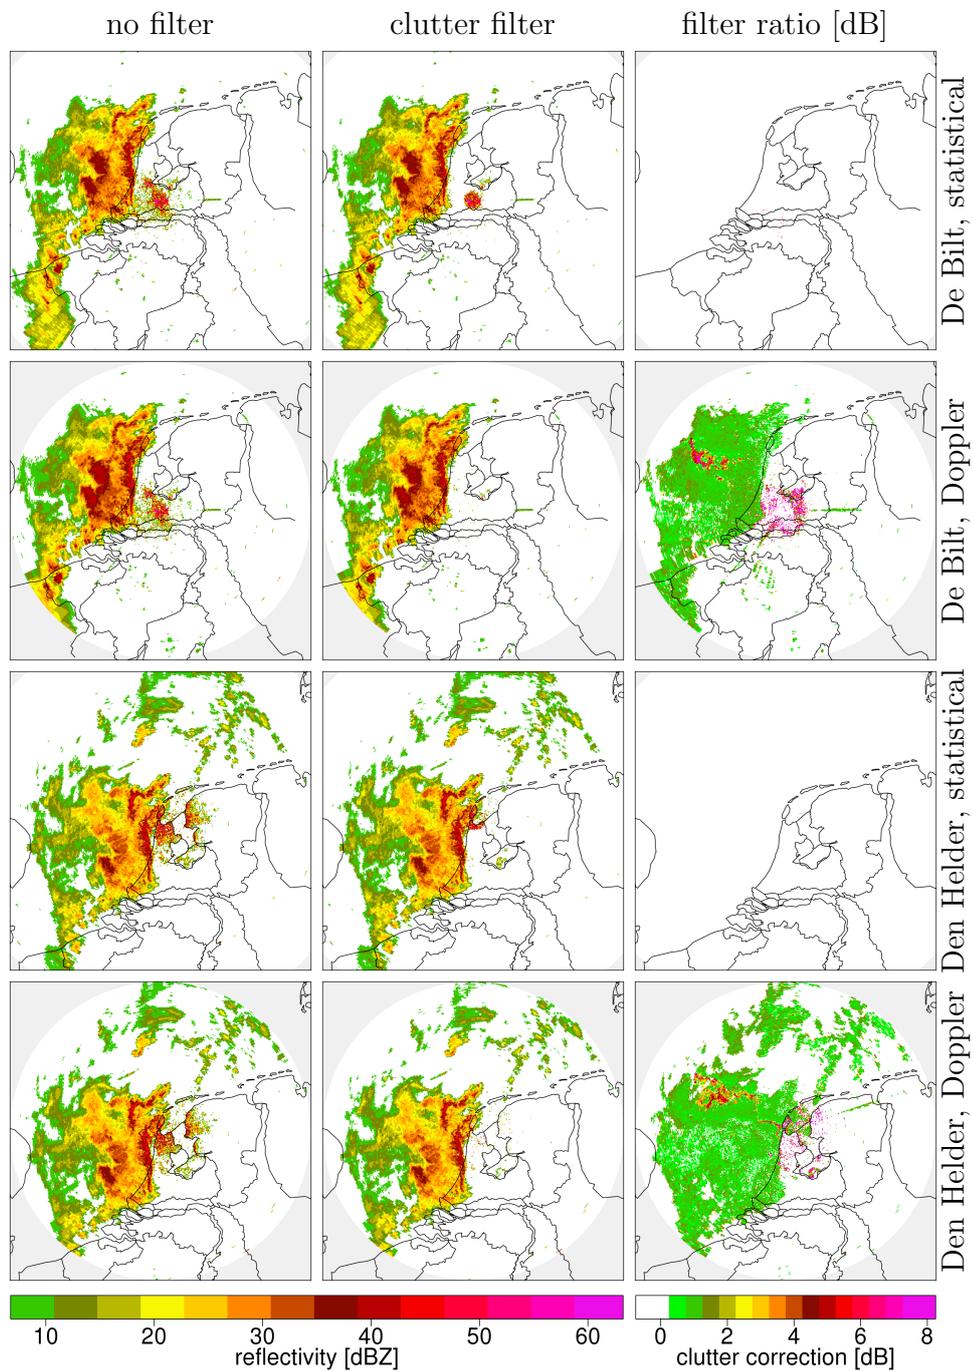


Figure 4.3: Left and center columns as Fig. 4.1, but for the rain event on 7 April 2014 at 16:00 UTC. The right-hand column shows the ratio of the unfiltered and filtered signals in dB. Note that values are only shown if pixels have not been removed through thresholding the clutter-to-signal ratio (in the case of the Doppler filter) or through the filter itself (in the case of the statistical filter).

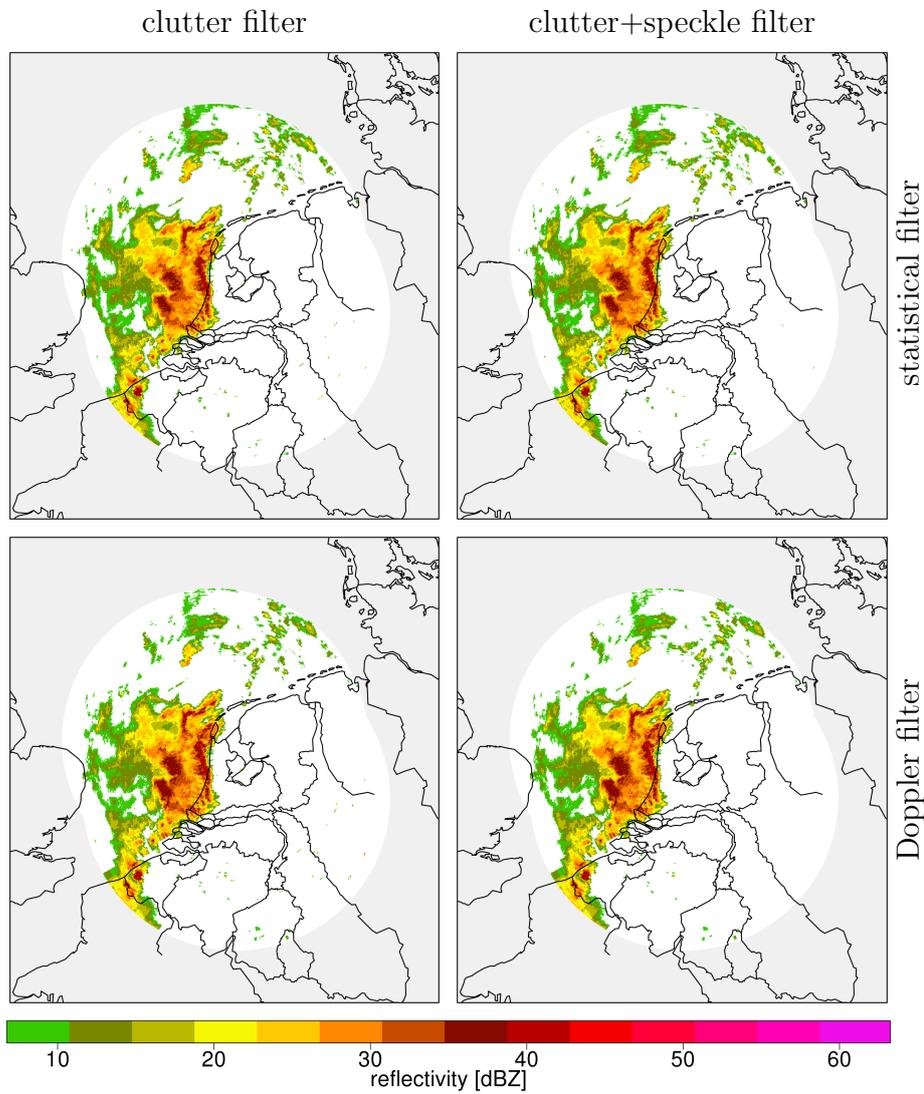


Figure 4.4: As Fig. 4.2, but for the rain event on 7 April 2014 at 16:00 UTC.

of the received echo power. A frequency-domain Doppler filter, which allows for spectral reconstruction, would greatly reduce the severity of this 'zero-isoDop' issue. However, the KNMI radars are currently not configured to apply this kind of Doppler filtering in an effective manner.

The weakening of the echoes by the time-domain clutter filter can also be seen by comparing the top panels of Fig. 4.4 to the bottom panels of that figure, although it is less clear than from Fig. 4.3. The effect of the speckle filter is seen to be minor for this case. This suggests that a speckle filter can

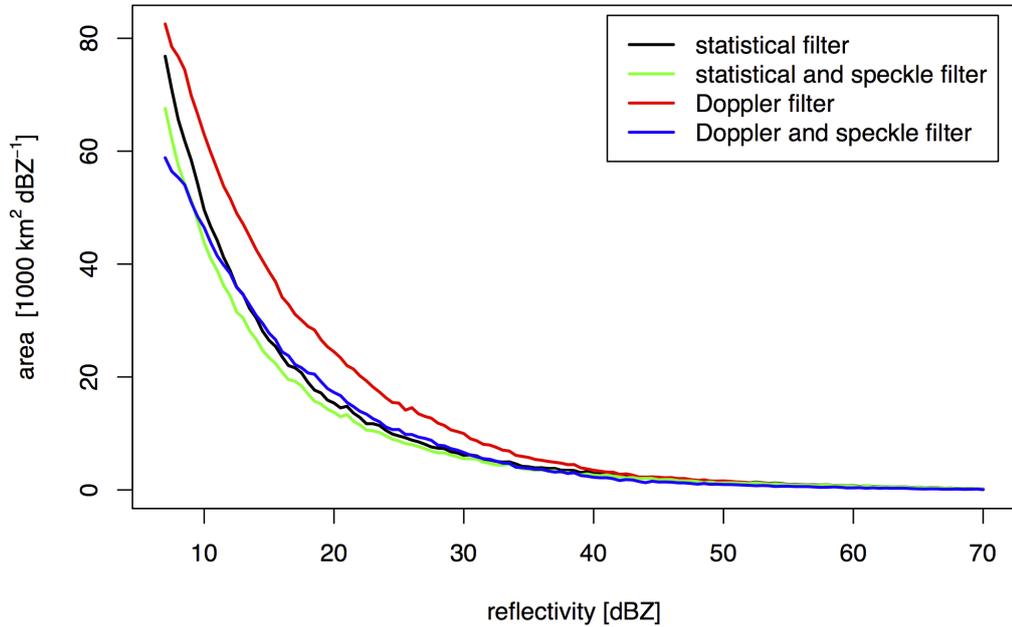


Figure 4.5: Area per dBZ class covered by pixels with a given dBZ value in cloud-free conditions. This histogram is based on composites of 1500-m pseudoCAPPIs.

be used safely as an addition to other clutter filters.

4.2 Objective verification using cloud masks

In order to objectively evaluate the performance of the different clutter filter methods, we use EUMETSAT's SAFNWC cloud mask product (Derrien and Le Gléau, 2005). The rationale behind this approach is that there cannot be any precipitation when there are no clouds. Hence the application of the most effective clutter filter method will result in the lowest number of pixels indicating nonzero rainfall in cloud-free areas (low false alarm rate).

The objective evaluation of the different clutter filter methods is carried out over a 10-day period (30 March till 8 April 2014). This period was characterized by frequent occurrence of anomalous propagation conditions, leading to severe ground and sea clutter. In addition several rainfall events occurred in this period (see also Figures 4.1, 4.2, 4.3, and 4.4) making it an ideal period for this evaluation.

Figures 4.5 and 4.6 show the total area covered by radar pixels having a certain dBZ (non-zero) value in cloud-free zones according to the SAFNWC

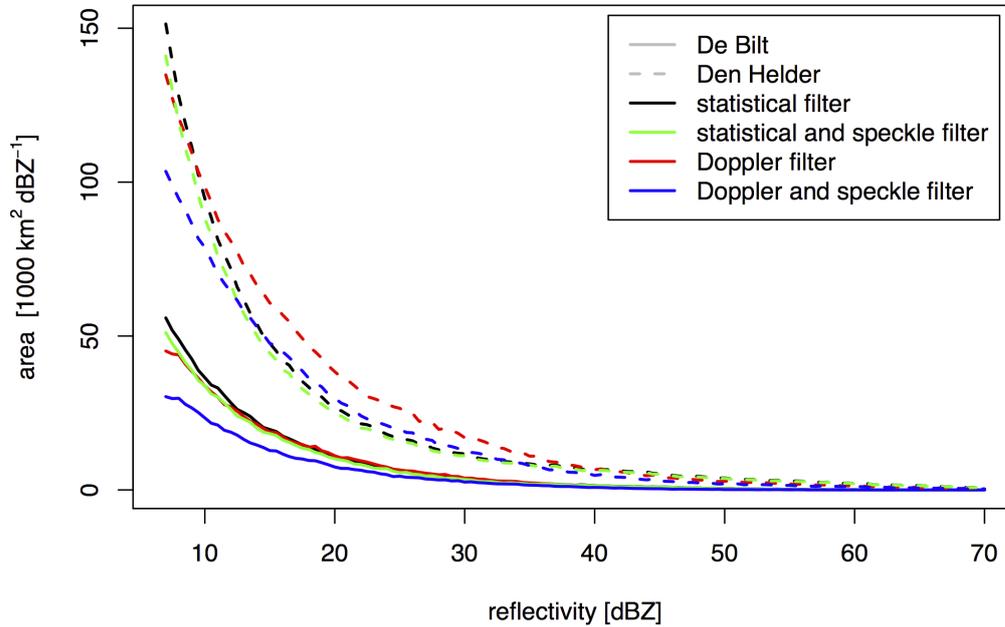


Figure 4.6: Area per dBZ class covered by pixels with a given dBZ value in cloud-free conditions. This figure is based on low-level PPI scans for both the De Bilt (solid lines) and the Den Helder (dashed lines) radars. Only data from ranges of at least 15 km are used.

cloud mask. Ideally this area should be zero because we assume that there cannot be any rain without clouds. Figure 4.5 shows the results for the pseudoCAPPI data. It is clear from this figure that the statistical filter actually outperforms the time-domain Doppler filter when no speckle filter is used. The application of a speckle filter to the time-domain Doppler-filtered data, however, results in a performance that is similar to that of the statistical filter. Adding a speckle filter to the statistically-filtered data only yields a marginal improvement.

Figure 4.6 shows similar results as Fig. 4.5 but for the individual radars in De Bilt and in Den Helder. For the De Bilt radar the Doppler filter is slightly more effective in removing clutter than the statistical filter (even without the speckle filter), but for the Den Helder radar the statistical filter is more effective. The statistical filter outperforms the Doppler filter in Den Helder because a large part of the clutter around this radar is sea clutter. This type of clutter is notoriously difficult for Doppler filters to remove because the waves that reflect the radar signal have non-zero velocities. Note that the

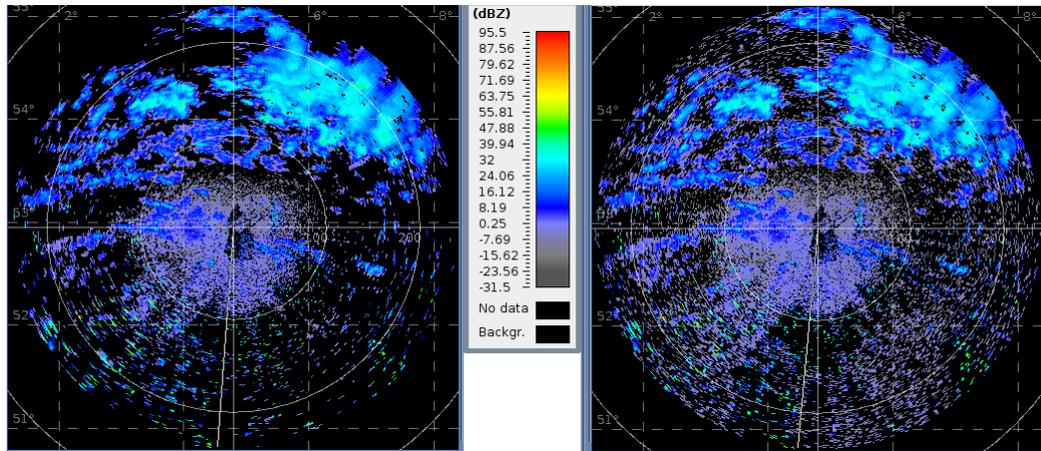


Figure 4.7: Example of using the IQ replay facility to test different Doppler filters on data from the 0.4° elevation scan of the Den Helder radar on 23 October 2012, 18:15 UTC. Left: filter settings as currently used (time-domain filter). Right: same scan but using a frequency-domain Doppler filter with spectral reconstruction.

central idea behind a Doppler clutter filter is that clutter targets have zero velocity. Other clutter filters (e.g. based on polarimetric data) will be better capable to remove sea clutter (Hubbert et al., 2009).

It is clear from Figures 4.5 and 4.6 that a combination of a time-domain Doppler and a speckle filter performs better than the statistical filter with or without speckle filter. This holds true for both pseudoCAPPi's composites (the well-known radar imagery product for surveillance purposes) and for low-level PPIs of the individual radars.

4.3 IQ-replay

In order to test whether a frequency-domain filter with spectral reconstruction can indeed be used effectively to remove clutter without affecting precipitation (as was shown in Fig. 4.3), in-phase and quadrature data (I and Q, or IQ data) have been recorded for approximately 1 week in October 2012. IQ-replay is a facility made by SELEX Gematronik that can be used to re-process these IQ data with different filter settings. Fig. 4.7 shows an example of using this facility. The example shown is from the Den Helder radar on 23 October 2012, 18:15 UTC, with some precipitation to the north of the radar and clutter to the south-west.

The left panel of Fig. 4.7 shows a 0.4° elevation scan, with the current

time-domain Doppler filter applied. The right panel of this figure shows the application of a frequency-domain Doppler filter on the same data. The result is that the precipitation is less affected than when using a time-domain clutter filter, but there seems to be more residual clutter. Note that this additional residual clutter is very low in intensity so that it is likely that this will not appear on any operational imagery. However, ideally a new Doppler filter should be at least as effective in removing clutter as the one that is currently operationally used.

One of the reasons that the frequency-domain filter is not as effective as the time-domain filter in Fig. 4.7 is that these data were measured while switching PRFs halfway through a 1-degree sector. Because of this the frequency-domain Doppler filter can use only 14 pulses to compute the Doppler spectrum (the time-domain filter can use 28 pulses). This is generally not considered to be enough to base Doppler notch filtering and spectral reconstruction on. This means that a thorough evaluation of such a frequency-domain Doppler filter can only be carried out on data measured where PRF switching is done between 1-degree sectors or even better on single-PRF data. Hence it is strongly recommended to carry out experiments with such scan settings in the future.

Chapter 5

Summary and next steps

As part of the Radar Sensor Replacement (RASER) project the performance of KNMI's statistical clutter removal scheme was compared to that of standard Doppler-based algorithms. If the Doppler algorithms would perform at least as well as the statistical scheme, there is no need to port the non-standard statistical clutter removal scheme of KNMI to the new polarimetric weather radars and this would lead to a considerable simplification of the radar sensor replacement project. The quality of the two lowest elevations of the KNMI operational radar scan was evaluated under varying meteorological circumstances. The Cloud Mask product from the Meteosat satellite was used for objective verification of the clutter removal methods. Moreover, datasets with raw IQ-samples from the De Bit radar were reprocessed in order to gain more insight in the optimum settings of the clutter removal algorithms. These studies led to the following conclusions:

- Performances of the statistical clutter and Doppler clutter (without speckle filter) removal schemes are quite similar and depending on the meteorological circumstances and on the radar location the one or the other is slightly better. Overall the Doppler clutter removal method performs somewhat better than the statistical method.
- Application of an one-dimensional speckle filter on the volume data after Doppler-based clutter removal gives a substantial improvement in data quality. For statistical filtered data application of a speckle filter has very little impact.
- At short ranges the Doppler method works much better because the statistical clutter removal is not applied at ranges below 14 km. When both radars are operational this is masked by the radar compositing

algorithm, but when only one radar is active the improvement will be evident.

- Effectiveness of both methods in the removal of (intense) sea clutter is limited due to the variability of the sea surface.
- Performance of the Doppler-based clutter removal can be further improved by the use of frequency-domain filtering instead of time-domain filtering and by optimizing the scan settings for the lowest elevation, most notable application of single-PRF instead of dual-PRF scanning (more pulses and optimal filtering window). Additional evaluations are required to confirm this and to find the optimum settings.

And based on these conclusions the following recommendations are given for the RASER project and the operation of the new polarimetric weather radars:

- Replace the statistical clutter removal scheme on the lowest elevation with the Doppler clutter removal combined with a one-dimensional speckle filter. In the RASER project there is no need to port the KNMI method to the new signal processors.
- Perform the replacement of the clutter removal scheme for the lowest elevation of the operational radar scan on the current radars, i.e. decoupled from the radar sensor replacement.
- As a first step, set up a test elevation in the operational scan of the current Doppler radars using a frequency-domain Doppler filter with spectral reconstruction, and optimize the settings of this scan (single PRF). Evaluate the data quality of this test elevation during a period with frequent occurrence of anomalous propagation and many rain events.
- Implement a pragmatic but effective method to reduce contamination by sea clutter. Modern radar signal processors can be instructed to apply different settings above land and sea surface and thus apply more stringent clutter removal above sea. Alternatively this could be implemented as post-processing of radar volume data. In this way sea clutter can be removed more effectively while not reducing data quality above land (e.g. for quantitative precipitation estimation).
- Investigate the effectiveness of the use of polarimetric clutter detection and removal algorithms.

Follow-up of these recommendations will lead to improved performance of the clutter removal at the lowest elevation, more transparency of the applied methods (for both inside and outside KNMI), easier radar maintenance, and lower costs for the RASER project.

Acknowledgments

SELEX ES GmbH, and more in particularly Peter Gölz, is gratefully acknowledged for support with the IQ-replay module.

References

- Aoyagi, J.: 1983, A study on the MTI weather radar system for rejecting ground clutter. *Papers in Meteor. and Geophys.*, **33**, 187–243.
- Derrien, M. and H. Le Gléau: 2005, MSG/SEVIRI cloud mask and type from SAFNWC. *International Journal of Remote Sensing*, **26**, 4707–4732.
- Doviak, R. J. and D. S. Zrnić: 1993, *Doppler Radar and Weather Observations*. Academic Press, second edition, 562 pp.
- Hazenberg, P., H. Leijnse, and R. Uijlenhoet: 2011, Radar rainfall estimation of stratiform winter precipitation in the Belgian Ardennes. *Water Resources Research*, **47**, W02507.
- Hazenberg, P., P. J. J. F. Torfs, H. Leijnse, G. Delrieu, and R. Uijlenhoet: 2013, Identification and uncertainty estimation of vertical reflectivity profiles using a Lagrangian approach to support quantitative precipitation measurements by weather radar. *Journal of Geophysical Research: Atmospheres*, **118**, 10,243–10,261.
- Holleman, I. and H. Beekhuis: 2005, Review of the KNMI clutter removal scheme. Technical report TR-284, Royal Netherlands Meteorological Institute (KNMI).
- Hubbert, J. C., M. Dixon, and S. M. Ellis: 2009, Weather radar ground clutter. Part II: Real-time identification and filtering. *Journal of Atmospheric and Oceanic Technology*, **26**, 1181–1197.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery: 1992, *Numerical Recipes in C: the Art of Scientific Computing*. Cambridge University Press, second edition, 994 pp.
- Wessels, H. R. A. and J. H. Beekhuis: 1992, Automatic suppression of anomalous propagation clutter for noncoherent weather radars. Scientific report WR-92-06, Royal Netherlands Meteorological Institute (KNMI).

- 1994, Stepwise procedure for suppression of anomalous ground clutter. *COST-75 Seminar on Advanced Radar Systems*, EUR 16013 EN, 270–277.