Performance of the EarthCARE Cloud Profiling Radar in marine stratus clouds

David Burns

Department of Atmospheric and Oceanic Sciences
McGill University, Montreal
June 2016

A thesis submitted in partial fulfilment of the requirements of the degree of Master of Science

Abstract

Marine stratiform clouds are a challenging target for spaceborne radars due to their proximity to Earth's surface, limited vertical extent, and low radar reflectivity. The joint European Space Agency-Japanese Aerospace Exploration Agency (ESA-JAXA) Earth Clouds, Aerosol, and Radiation Explorer (EarthCARE) mission is scheduled for launch in 2018 and features the first spaceborne Cloud Profiling Radar (CPR) with Doppler capability. In this work, the performance of the CPR in detecting these clouds, their boundaries, and their drizzle droplet velocities is evaluated. Observations from the Atmospheric Radiation Measurement (ARM) Mobile Facility at the Azores and the Marine ARM GPCI Investigation of Clouds are used as input to a CPR simulator. The extensive ground-based observations are treated as the truth and compared to the simulated EarthCARE CPR observations. The impact of the surface echo return, radar sensitivity, and range resolution are discussed, and post-processing techniques such as a feature mask algorithm, range resolution inversion, and matched filters for Doppler velocity are presented and tested. The EC-CPR detected cloud fraction is found to range from approximately 70–80 % that of ground-based W-band radars, depending on the along-track integration and configuration of the feature mask algorithm. The range resolution of the EC-CPR introduces an average reflectivity bias of +1 dB and cloud top overestimation bias of 100 m (equal to the range sampling rate of the CPR), but these may be significantly reduced (to 0.1 dB and 35 m respectively) by the application of a range resolution inversion technique. The analysis indicates a CPR velocity uncertainty of approximately $0.5\,\mathrm{m\,s^{-1}}$ is achievable with either a 5 km along-track integration or a combination of

matched filtering and along-track integration to 1 km of the CPR Doppler velocity field.

Résumé

Les nuages stratiformes marins sont une cible difficile pour les radars spatiaux en raison de leur proximité de la surface de la Terre, de leur étendue verticale limitée et de leur faible réflectivité. La mission “Earth Clouds, Aerosol et Radiation Explorer” (EarthCARE), une colaboration entre les agences spatiales européenne et japonaise (ESA-JAXA), est prévue pour 2018 et propose de mettre en orbite le premier profileur de nuages radar (CPR) avec capacité Doppler. On évalue la performance du CPR dans la détection de ces nuages, leurs limites et leurs vitesses de gouttelettes de bruine. Les observations de la Atmospheric Radiation Measurement (ARM) Mobile Facility sur les Açores et la Marine ARM GPCI Investigation of Clouds sont utilisés comme entrée à un simulateur de CPR. Les vastes observations au sol sont traités comme la vérité et comparées aux observations EarthCARE CPR simulées. L’impact du retour de l’écho de surface, la sensibilité du radar, et la résolution en portée sont discutées, et les techniques de post-traitement tels qu’un algorithme de masquage d’entités, l’inversion de la plage de résolution, et des filtres associés à la vitesse Doppler sont présentés et testés. La fraction de nuage détectée par le EC-CPR est d’environ 70–80 % de celui des radars W-band au sol, les résultats dépendent de l’intégration along-track et la configuration de l’algorithme de masquage d’entités. La résolution de portée de la EC-CPR introduit un biais de réflectivité moyenne de +1 dB et un biais dans l’estimation du haut des nuages de +100 m (égale à la fréquence d’échantillonnage de portée du CPR), mais ceux-ci peuvent être considérablement réduits (0.1 dB et 30 m respectivement) par l’application d’une technique d’inversion de résolution de portée. L’analyse indique qu’on peut atteindre un incertitude de environ $0.5\,\mathrm{m\,s^{-1}}$

pour la vitesse CPR avec soit une intégration along-track de 5 km ou une combinaison de filtres associés à la vitesse Doppler et une intégration de 1 km du champ de vitesse CPR Doppler.

Acknowledgements

I would like to first express my sincere gratitude to my supervisor, Pavlos Kollias, for providing the opportunity to work with him and to become involved with the EarthCARE project. His guidance, support and knowledge have been invaluable to me throughout my degree, particularly in the multiple revisions of the manuscript upon which this thesis is based. I also extend my thanks to the manuscript coauthors, Aleksandra Tatarevic, Alessandro Battaglia, and Simone Tanelli, and the many others whose efforts and insight have proven extremely helpful in shaping this research, and to Matthew Henry for providing the French translation of the abstract. Finally, to my friends at McGill, both inside and out of the Department of Atmospheric and Oceanic Sciences, who have made my two years here an incredibly enjoyable experience and for whose company I am immensely grateful, thank you.

Contribution of Authors

The work presented in this thesis was carried out by myself and is based on a manuscript to be submitted to the Journal of Geophysical Research, a journal of the American Geophysical Union. This manuscript is co-authored by my supervisor Pavlos Kollias and colleagues Aleksandra Tatarevic, Alessandro Battaglia, and Simone Tanelli who provided comments on and revisions of the manuscript as well as the initial EC-CPR simulator and feature mask algorithm that were modified for use in this study.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Marine stratiform clouds play a critical role in Earth's radiation budget and hydrological cycle, and therefore in climate and climate change. Their influence on Earth's climate can be attributed to their vast horizontal coverage, their ability to strongly reflect incoming shortwave radiation (Hartmann et al., 1992), and their regulating effect on the marine boundary layer structure through drizzle and turbulence production (Stevens et al., 2003). Evaluation of marine stratus representation in climate models requires large-scale, long-term observational datasets. Such observations are challenging to conduct from ground-based platforms. Therefore, spaceborne instruments, which offer global coverage, are key for monitoring the properties of marine clouds.

Earth Clouds, Aerosol, and Radiation Explorer (EarthCARE) is a joint European Space Agency (ESA) and Japanese Aerospace Exploration Agency (JAXA) satellite due to launch in 2018 that will carry a cloud profiling radar (CPR), a high-spectral-resolution lidar (ATLID), a broadband radiometer (BBR), and a multi-spectral imager (MSI) (Illingworth et al., 2015). Both independently and synergistically, these instruments will provide profiles of clouds, precipitation, and aerosol properties and top-of-atmosphere radiative fluxes and heating rates. EarthCARE may be considered the successor to the National Aeronautics and Space Administration (NASA) Afternoon Train, or A Train, a formation of Earth-observing satellites that carry a mix of active and passive remote sensing instruments. Among these are the Cloud Profiling Radar aboard CloudSat (Stephens et al., 2008) and the Cloud-Aerosol Lidar with Orthogonal Polarization (CALIOP) aboard CALIPSO (Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations; Winker et al., 2009), which have, since their launches in 2006, provided extensive climatologies of marine boundary layer clouds, e.g., Leon et al. (2008).

The 94 GHz EarthCARE CPR (EC-CPR) will feature several improvements over the CloudSat CPR (CS-CPR), including being the first spaceborne atmospheric radar with Doppler capabilities. This will allow it to provide information on not only cloud extent and structure but also vertical motions within clouds and precipitation. In addition to its Doppler capability, the EC-CPR employs a higher along-range sampling rate (100 m) and sensitivity (−35 dBZ) than the CS-CPR (240 m and −30 dBZ respectively; Tanelli et al., 2008), with the enhanced sensitivity of the EC-CPR primarily due to its lower flying altitude and larger antenna. A detailed comparison of the technical parameters of the two CPRs is given in Table 1.1. Despite these improvements, however, observations from spaceborne radars remain challenging. For instance, both the EC-CPR and CS-CPR use pulse lengths of 1000 m, compared to below 100 m for typical ground-based radars (e.g., Widener and Mead, 2004). A longer pulse length allows a greater maximum emitted pulse power, necessary for spaceborne radars to achieve reasonable signal-to-noise ratios (SNRs), but also reduces the range resolution, which is dictated by pulse length and receiver bandwidth. Uttal and Kropfli (2001) investigated the effect long pulse lengths (450 m) in the context of spaceborne radar observations by artificially reducing the vertical resolution of ground-based radar data and found that this introduced an average reflectivity bias of +4 dB. Vertical smoothing of CPR profiles also hampers the accurate definition of cloud boundaries; this is a particular problem for observations of marine stratiform clouds, which are thin (typically less than 500 m, the vertical resolution of the EC-CPR; Wood, 2012) and therefore susceptible to proportionally large errors in cloud boundary estimates.

Parameter	EC-CPR	CS-CPR
Frequency (GHz)	94	94
Altitude (km)	393 (mean)	705-732
Antenna diameter (m)	2.5	1.85
Pulse length (µs)	3.33	3.33
Range resolution (6 dB) (m)	500	480
Vertical sampling rate (m)	100	240
Along-track sampling rate (m)	500	1100
Antenna beamwidth (deg)	0.095	0.12
PRF (kHz)	6.1-7.5	3.7-4.3
Sensitivity (dBZ)	-35	-30

Table 1.1: Technical characteristics of the EarthCARE and CloudSat Cloud Profiling Radars.

A further effect of the EC-CPR range resolution is that surface clutter - strong reflections of radar waves by the ground - extends into range gates up to 1 km above the ground. These surface echo signals can be several orders of magnitude stronger than hydrometeor signals, often causing observations of low-lying and low reflectivity marine strati to be obscured or biased. Furthermore, precipitation-free marine strati typically have reflectivities close to or below the sensitivity of the EC-CPR, such that many of these clouds will go undetected even without the presence of surface clutter.

Spaceborne Doppler velocity measurements are affected by these and other factors. At low SNRs (relative to the single-pulse noise of the EC-CPR, approximately $-21.5\,\mathrm{dBZ}$), random noise-induced errors will dominate velocity estimates. In addition, the satellite's motion of $7\,\mathrm{kms}^{-1}$ combined with the EC-CPR antenna's beamwidth will both broaden the observed Doppler spectra and, in the case of non-uniform beam filling (NUBF), introduce biases that can reach several metres per second (Tanelli et al., 2002; Schutgens, 2008; Kollias et al., 2014). Biases of similar magnitudes are produced by mispointing of the CPR antenna away from the nadir direction by even a few µrad (Battaglia and Kollias, 2015). Together, these effects restrict the quality of the CPR reflectivity and Doppler estimates. However, post-processing methods have the potential to mitigate these errors. For example, increased along-track integration can reduce random noise-induced uncertainties; reference targets with known velocities, such as the surface, allow correction for antenna mispointing; and NUBF errors can be corrected for with knowledge of the along-track reflectivity gradient.

The aims of this thesis are twofold. First, to quantify, in terms of biases and uncertainties, how well the EC-CPR captures i) marine stratiform cloud morphology (i.e., reflectivity, cloud fraction, and cloud boundaries), and ii) drizzle Doppler velocities. Second, to identify and evaluate post-processing techniques that can reduce these errors – specifically, to apply a constrained linear inversion process, similar to those described in Schutgens and Donovan (2004) and Galati et al. (1996), to EC-CPR reflectivity profiles in order to achieve an increased vertical resolution, and to use matched filtering, following Sy et al. (2014), to reduce random noise in the EC-CPR Doppler velocity.

In order to address these issues, ground-based radar observations from two deployments of the US Department of Energy Atmospheric Radiation Measurement (ARM) program (Stokes and Schwartz, 1994) in marine stratus cloud regimes are used as input to an EC-CPR simulator. This outputs synthetic EC-CPR observations of the same scene, which can then be compared with the "true" ground-based data for a quantitative evaluation of the EC-CPR measurements.

The remainder of this thesis is structured as follows: details of the ground-based radar data and the simulator are presented in Chapter Two. This includes a description of the feature mask algorithm used to initially process the EC-CPR data. In Chapter Three, the simulated results are presented and evaluated with respect to the simulator input data. The effect of the surface echo, the horizontal sampling rate, and the configuration of the feature mask algorithm are discussed. The range resolution inversion and Doppler velocity matched filtering techniques are presented and evaluated in Chapter Four. A summary of the thesis is provided in Chapter Five.

Chapter 2

Background

2.1 Ground-based data

Observations from two recent deployments of the ARM Mobile Facility (AMF; Mather and Voyles, 2013) are used as input to the EC-CPR simulator. The first is the AMF-1 deployment at Graciosa Island, Azores (GRW; Rémillard et al., 2012; Wood et al., 2015). The second is the AMF-2 Marine ARM GPCI Investigation of Clouds (MAGIC), in which observations were carried out aboard the cargo ship Horizon Spirit during several trips between Los Angeles, California, and Honolulu, Hawaii (Zhou et al., 2015). In both deployments, a vertically pointing 95 GHz (W-band) ARM Cloud Radar (WACR (Widener and Mead, 2004) at GRW and the marine WACR (M-WACR) in MAGIC) was deployed alongside several other instruments (e.g., ceilometer, microwave radiometer). The W-band radars provide estimates of radar reflectivity factor, mean reflectivity-weighted Doppler velocity, and spectrum width. The WACR and M-WACR use integration times of 2 s and 0.2 s and vertical sampling rates of 43 m and 21 m respectively. Both radars have sensitivities of -40 dBZ at 2 km.

While several days of observations from the two deployments are used to produce statistics of the marine stratus scenes, two example cases, one from each AMF deployment, are presented in detail. The GRW case was observed on 29$^{\text{th}}$ November 2009 and the MAGIC case on 30$^{\text{th}}$ July 2013. Figure 2.1 displays the fields of reflectivity factor and Doppler velocity for the two scenes. Both contain drizzle-free and drizzling regions, with cloud thicknesses varying correspondingly from a few hundred metres (cloud only) to over a kilometre (cloud and drizzle), the latter often extending into the lowest 1 km above the surface. Reflectivity values vary from approximately 10 dBZ within strong drizzle patches to around -40 dBZ near cloud boundaries. Velocities are generally small, ranging from close to zero within clouds to magnitudes of 1–2 ms^{-1} in drizzle.

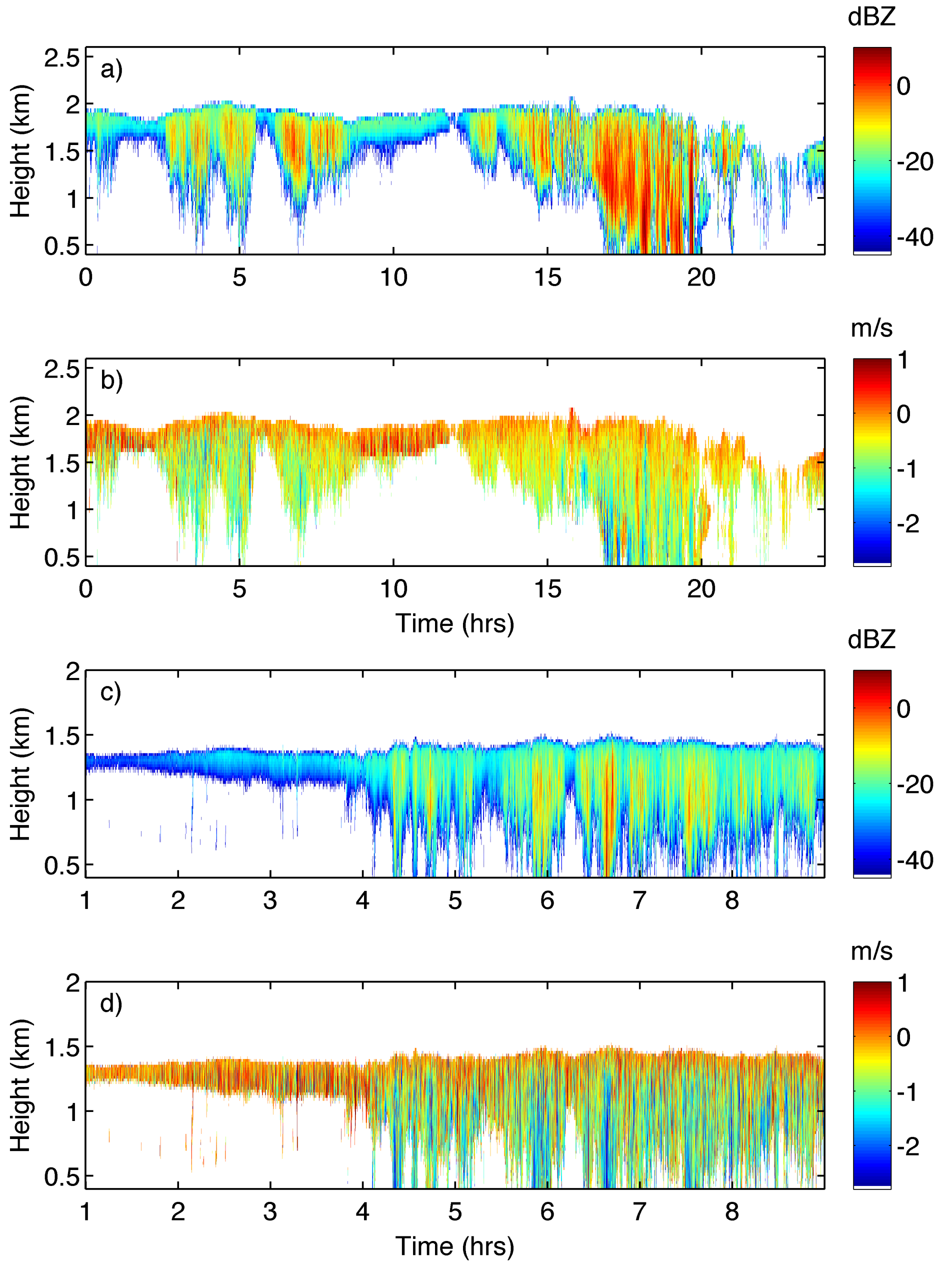

Figure 2.1: Time-height ARM radar observations of reflectivity factor and Doppler velocity for two AMF deployments in marine stratus regimes: (a, b) 29 November 2009 (GRW); and (c, d) 30 July 2013 (MAGIC). Upward (downward) vertical velocities are assumed positive (negative).

2.2 Simulator

The EC-CPR simulator used in this study, described in detail in Kollias et al. (2014), produces EC-CPR observations from ground-based and/or airborne radar data. The cloud scenes are first converted from a time-height to an along-track-height field using local wind speed soundings averaged over the cloud layer. The simulator then applies the effects of measurement by the EC-CPR, such as radar receiver noise, the horizontal and vertical resolutions, spectrum broadening due to satellite motion, along-track integration, and Doppler moment estimation that emulates the real-time on-board processing of EC-CPR signals. In Kollias et al. (2014), the focus was on high-level ice clouds and thus simulation of the surface return was not included. However, the surface echo is an important feature that restricts the detection of low-level clouds by spaceborne radars (Tanelli et al., 2008); therefore, for this study, the surface echo is introduced in the simulator and is described here.

The profile of the surface echo is determined by the normalised backscatter cross section of the Earth's surface and the combined effect of the radar's range-weighting function and receiver bandwidth. In the case of the CS-CPR, the average clear-sky surface return over ocean can be deduced from the CloudSat 2B-GEOPROF reflectivity product. For the EC-CPR, which uses a similar pulse length (and range-weighting function) to the CS-CPR, a best approximation of the clear-sky surface return is produced by convolving the estimated point response function of the EC-CPR with the surface backscatter. Variability of the surface return due to heterogeneous surface conditions, attenuation of the radar signal through the atmosphere, and changes in satellite altitude, which moves the surface's position within the range gate it intersects, have not been included in the EC-CPR simulator. This therefore constitutes a best-case scenario, where the only variation in the surface return is due to finite-sampling errors of the radar. A reduced surface echo EC-CPR profile is also considered, in which the surface clutter has been removed above 500 m, in order to isolate the effect of the surface return. The three profiles are plotted in Figure 2.2. At height, the profiles approach their respective noise floors. In the EC-CPR simulator, the mean Doppler velocity of the surface is assumed zero (valid under the condition that antenna mispointing has been sufficiently corrected for) and a narrow spectrum width (0.5 ms^{-1}) assigned, as the main contributor to Doppler spectral broadening is introduced later in the simulator. When concurrent (within the lowest 1 km), the WACR and surface clutter moments are combined to produce the total observed radar moments.

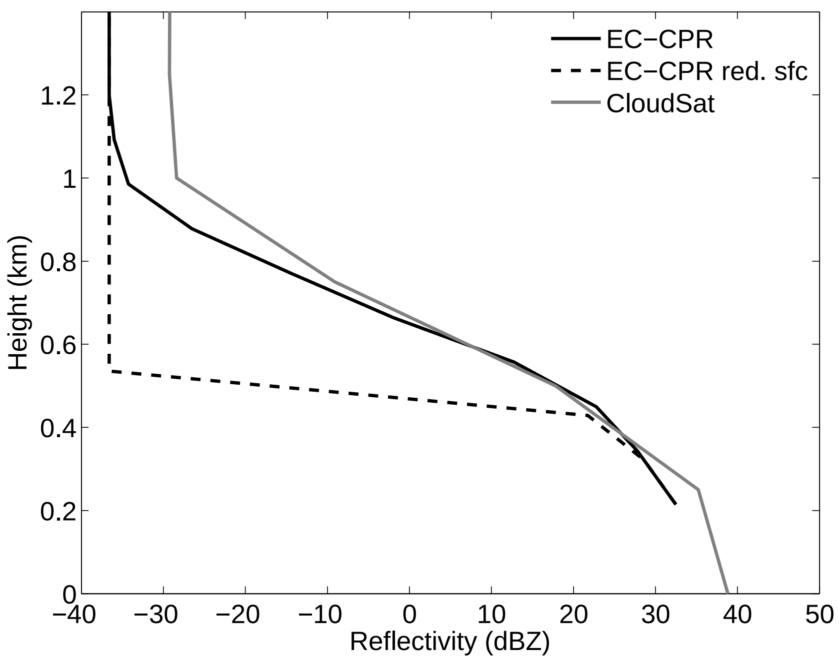

Figure 2.2: Average clear-sky surface return profiles for the CloudSat and EarthCARE CPRs, including a reduced surface case for EarthCARE in which the clutter has been completely suppressed above 500 m.

The simulator outputs estimates of reflectivity factor, mean Doppler velocity, and spectrum width based on pulse pair processing (Doviak and Zrnic, 1993) every 20–25 m of along-track displacement of the EC-CPR. These estimates are subsequently integrated horizontally to 500 m to reduce the impact of the receiver noise and satellite motion. Additional integration up to 10 km is performed in order to evaluate the quality of the measurements. The simulated Doppler velocities are corrected for NUBF according to the along-track reflectivity gradient method (Kollias et al., 2014).

2.3 Processing

Since this study is concerned with the detection of weak targets often embedded in the surface clutter, it is necessary to introduce the feature mask (FM) detection algorithm that is similar to that proposed for the EC-CPR. The FM algorithm takes as input the raw simulated EC-CPR signals described in 2.2 and indicates resolution volumes that contain radar signal returns that are statistically significantly higher than the background signal (receiver noise, atmospheric noise, and surface clutter).

The first step in the FM is the calculation of the mean and standard deviation of the noise in echo-free (i.e., surface clutter- and cloud-free) returns within each EC-CPR profile. The mean noise power in each profile is estimated using a variant of the Hildebrand and Sekhon (1974) method for noise estimation in Doppler spectra. In the original method, a coloured Doppler spectrum has a power threshold applied to it. Any frequency band in which the power is above this threshold is removed from the spectrum and a white noise test is applied. If the spectrum is not white, the threshold is lowered and the process repeated until the conditions for white noise are satisfied. Here, rather than a Doppler spectrum, a "range gate spectrum" is formed for each vertical profile, with each range gate considered equivalent to a frequency band within a Doppler spectrum. Initially considering all range gates, the test, shown in Equation 2.1, is applied.

$$n\sum_{i=1}^{n} P(i)^2 - 2\left(\sum_{i=1}^{n} P(i)\right)^2 \begin{cases} < 0, & \text{White noise} \\ > 0, & \text{Coloured noise} \end{cases} \tag{2.1}$$

$P(i)$ is the power of the i^{th} range gate, with $i = 1$ referring to the range gate with the lowest signal power, and $i = n$ the range gate with the highest. Initially, $n = N$, where N is the total number of range gates in the profile. If the condition for white noise is not met (i.e., Equation 2.1 is greater than zero), then n is set to $N - 1$, removing the range gate with the greatest reflectivity from the spectrum. This process is repeated until the condition for white noise is achieved. From the remaining spectrum, the mean power is computed, which is exactly the mean power of every individual pulse making up the profile. From the assumption of white noise, the standard deviation of the single-pulse noise is equal to this. Horizontal averaging reduces the standard deviation of the noise power to

$$\sigma_M = \frac{\sigma}{\sqrt{M}} \tag{2.2}$$

σ_M is the standard deviation after averaging M pulses, and σ is the standard deviation of the single pulse noise. Therefore, in each vertical profile, the standard deviation of the noise is estimated as the mean noise in that profile divided by the square root of the number of pulses per horizontal estimate. At a 1 km integration, this reduces the standard deviation of the noise to approximately -36 dBZ. The mean noise plus one, two, or three standard deviations (hereafter referred to as 1–3 σ) is used as a dynamic threshold for hydrometeor detection in subsequent processing. Range gates in which the signal power is higher than the threshold are classified as significant and the rest as non-significant returns. The mean noise is then subtracted from the reflectivity profile. Extensive sensitivity tests that illustrate the effects of the choice of the threshold value are shown in Chapter 3.

Surface clutter identification is the next step in the FM algorithm. Given the strength of the surface return relative to the background noise, range gates in approximately the lowest 1 km will be uniformly marked as significant in the previous step. In CloudSat data, surface clutter identification is performed by comparing the observed profile to a reference clear-sky profile. If the observed return power is above the 99^{th} percentile of surface returns at that height, the signal at that range gate is classified as significant and the surface component of the signal is subtracted (Marchand et al., 2008; Tanelli et al., 2008). Here, a similar method is used, with average surface returns produced from clear-sky regions within each case. In each profile, if the signal power in a range gate is greater than the mean surface signal at that level plus three standard deviations, the volume is classified as significant; otherwise, the volume is marked as non-significant. The average surface profile is then subtracted.

Finally, a two-dimensional (along-track and along-range) filter is applied, similar to that used for CloudSat data (Marchand et al., 2008). This runs a box, sized 3 by 3 resolution volumes, over the entire field (e.g., Clothiaux et al., 1995; Uttal et al., 1993). At each point, if more than half-plus-one of the surrounding volumes are classed as significant in the previous mask, then the central point is also marked as significant. If fewer than this are significant, the central point is classed as non-significant. This filter is applied iteratively twice with the aim of reducing both false negatives and false positives. Any volumes with a negative power are marked as non-significant. The mask is produced at horizontal sampling rates of 500 m, 1 km, 5 km, and 10 km.

Chapter 3

Initial EC-CPR results

Simulated EC-CPR reflectivity fields of the two example scenes (Fig. 2.1) are depicted in Figure 3.1. Displayed in each column is the raw EC-CPR reflectivity field (prior to application of the FM algorithm), followed by the reflectivity after noise subtraction and surface clutter reduction, and finally the EC-CPR results after applying the FM algorithm. The results in Figure 3.1 are produced using a 1 km horizontal integration and a $3\,\sigma$ significant detection threshold. In both the GRW and MAGIC input reflectivity fields (Fig. 2.1), there are many drizzle-free regions, characterised by relatively low reflectivities and thicknesses of only a few hundred metres. The signals here are generally close to or below the level of the noise and therefore are often undetected in the masked EC-CPR result. Increasing the along-track integration (e.g., to 5–10 km) reduces the effective noise level, making many of these weak cloud signals detectable. Alternatively, the signal threshold for detection in the FM algorithm may be lowered while maintaining the integration length. The effects of these two methods are discussed in 3.1.

Surface clutter is seen in Figure 3.1 as the horizontal banding in the lowest 1 km of the non-masked reflectivity fields. Even after surface clutter reduction (Fig. 3.1c, 3.1d), below 700–800 m, the surface echo is stronger than the majority of the hydrometeor returns by several dB. Subsequently, at these heights, the extraction of meteorological

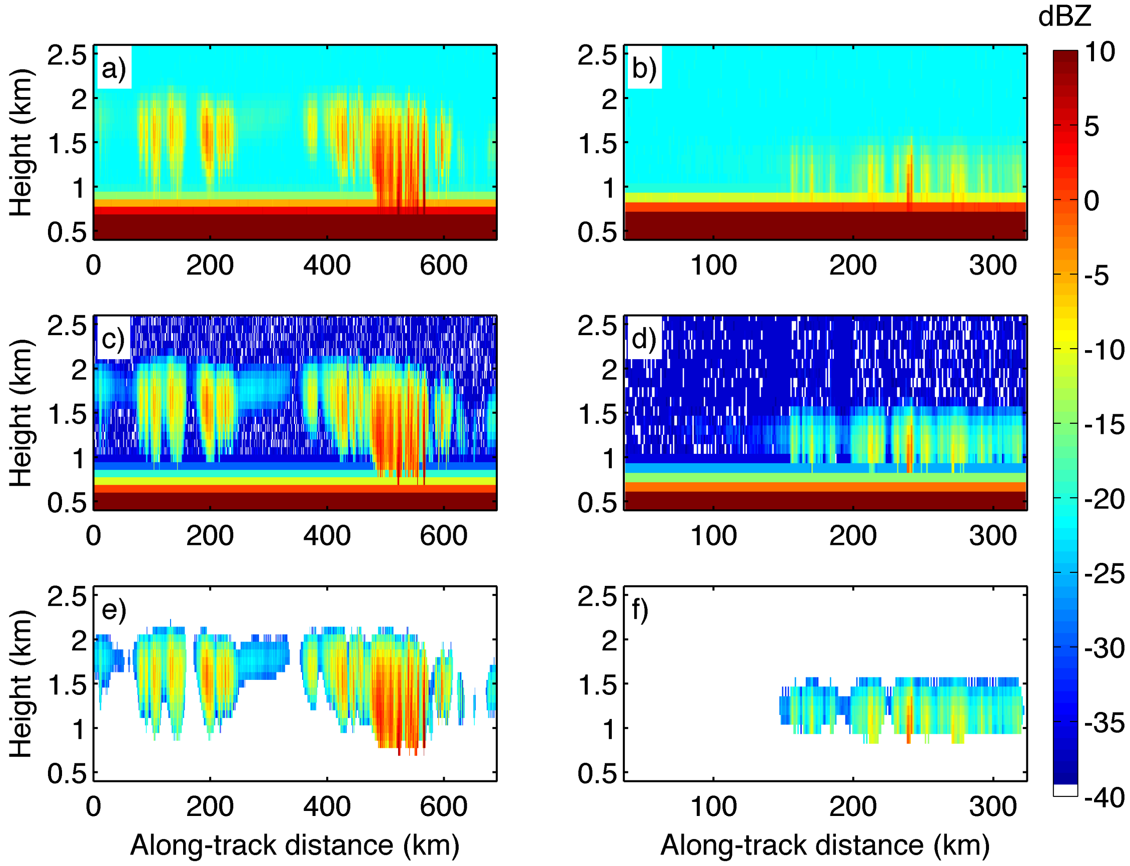

Figure 3.1: (a, b): Simulated EC-CPR reflectivity field of GRW and MAGIC example cases respectively at 1 km along-track integration prior to noise subtraction; (c, d): Simulated EC-CPR reflectivity after receiver noise subtraction and surface clutter reduction; (e, f): EC-CPR results after applying feature mask algorithm (using a $3\,\sigma$ significant detection threshold).

signals is challenging. The impact of the surface echo is explored further in section 3.1.

3.1 Impact of sampling rate, significant detection threshold, and surface echo

From the two AMF deployments, approximately 250 hours of observations of marine stratiform clouds are combined to produce statistics of the marine strati cloud properties. How well the EC-CPR reproduces these, and how this depends on integration length and the threshold for detection in the FM algorithm, is evaluated. Average cloud fraction is calculated using the ground-based WACR and M-WACR and the EC-CPR. The ground-based radars observed an average cloud fraction of 88 %. The EC-CPR value varies significantly by integration length (ranging from 500 m to 10 km) and significant detection threshold. Using a 1 km integration, the detected cloud fraction is 60 % or 70 % with 3 σ and 1 σ thresholds respectively. For a 10 km integration and 3 σ threshold, the value is 73 %. However, increasing the integration length or lowering the significant detection threshold increases the risk of false detections. Employing a 1 km integration and 3 σ detection threshold, the number of false detections (defined as profiles in which cloud is detected by the EC-CPR but not the WACR, with the data of both radars sampled to the EC-CPR initial 500 m along-track rate) is virtually zero, while at a 10 km integration, false detections account for approximately 3 % of all cloud-containing profiles. In simulations where surface clutter is removed above 500 m, the detected cloud fractions are 63 %, 74 %, and 77 % for 1 km, 3 σ; 1 km, 1 σ; and 10 km, 3 σ configurations, with this improvement due to profiles in which clouds are located entirely within the surface clutter.

These results confirm that an increase in integration length or a decrease in detection threshold result in increased detected cloud fraction. This increase comes almost entirely from thin clouds that have reflectivities close to that of the background noise. This is illustrated in Figure 3.2, which shows contoured-frequency-by-altitude-diagrams (CFADs) of WACR reflectivity (Fig.3.2a, 3.2c) and the proportion of true signals that are undetected by the EC-CPR as a function of height and significant detection threshold for the two scenes (Fig. 3.2b, 3.2d). Overlaid on Figure 3.2a and 3.2c are minimum detectable signal thresholds of the EC-CPR for the 1 km, 1 σ and 1 km, 3 σ configurations and the reduced-surface clutter simulation. Cloud and drizzle regimes are distinguishable in the CFAD and it can be seen that cloud base and cloud top are often below the EC-CPR sensitivity while drizzle is consistently above this (excepting range gates affected by the surface echo, where only the reduced surface profile is below the majority of WACR signals). Missed detections are generally higher for the 3 σ threshold than the 1 σ threshold as is expected due to the lower sensitivity of the former, though both show similar features: missed detections approach 100 % in the lowest 700 m where surface clutter dominates and are minimal in range gates below cloud base where only drizzle is present. In the MAGIC scene (Fig 3.2d), a peak in missed detections occurs in range gates where drizzle-free volumes are common at approximately 1.3 km. While these results suggest a lower threshold for detection is preferable, the choice of detection threshold and integration length also impact the accuracy of EC-CPR cloud boundary and reflectivity estimates. This is illustrated in Figures 3.3 and 3.4.

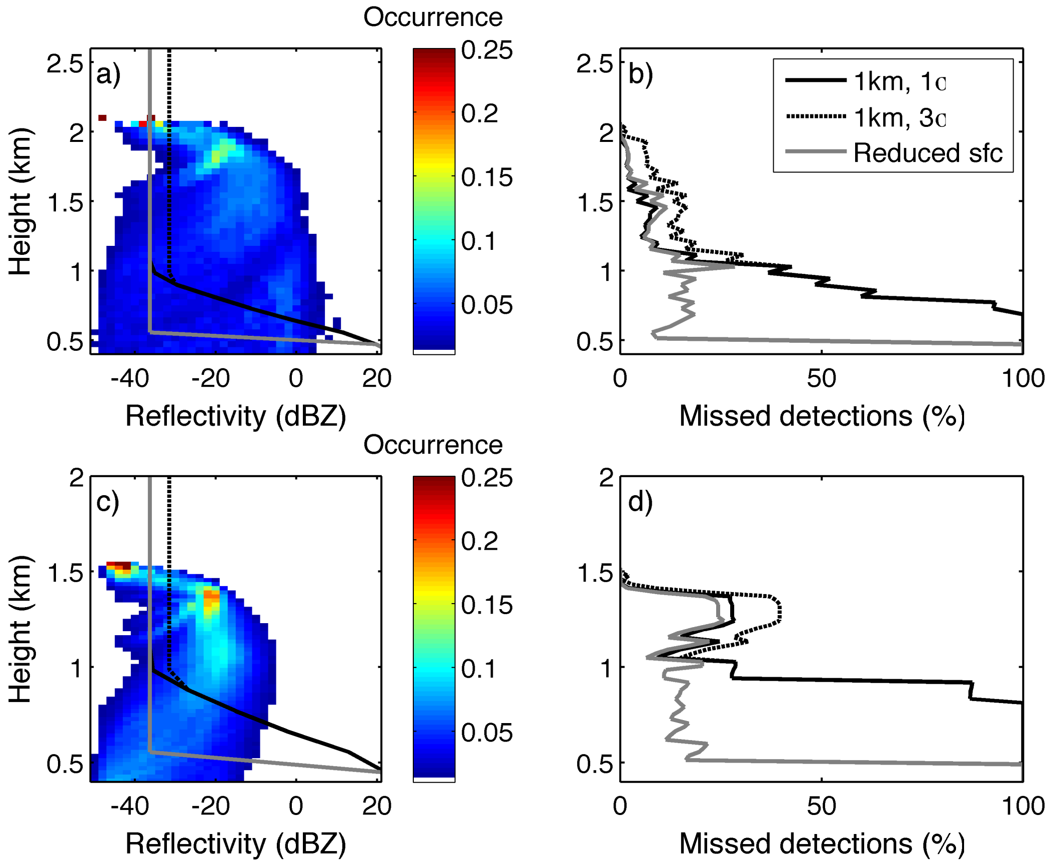

Figure 3.2: Contoured-frequency-by-altitude-diagrams (CFADs) of WACR reflectivity factor for (a) GRW and (c) MAGIC example cases shown in Figure 2.1a and 2.1c. The expected EC-CPR sensitivity for 1 σ and 3 σ detection thresholds at a 1 km integration are shown, including the reduced-surface case sensitivity. The fraction of WACR detections that are not captured by the EC-CPR as a function of height and significant detection threshold for the (b) GRW and (d) MAGIC cases is also shown.

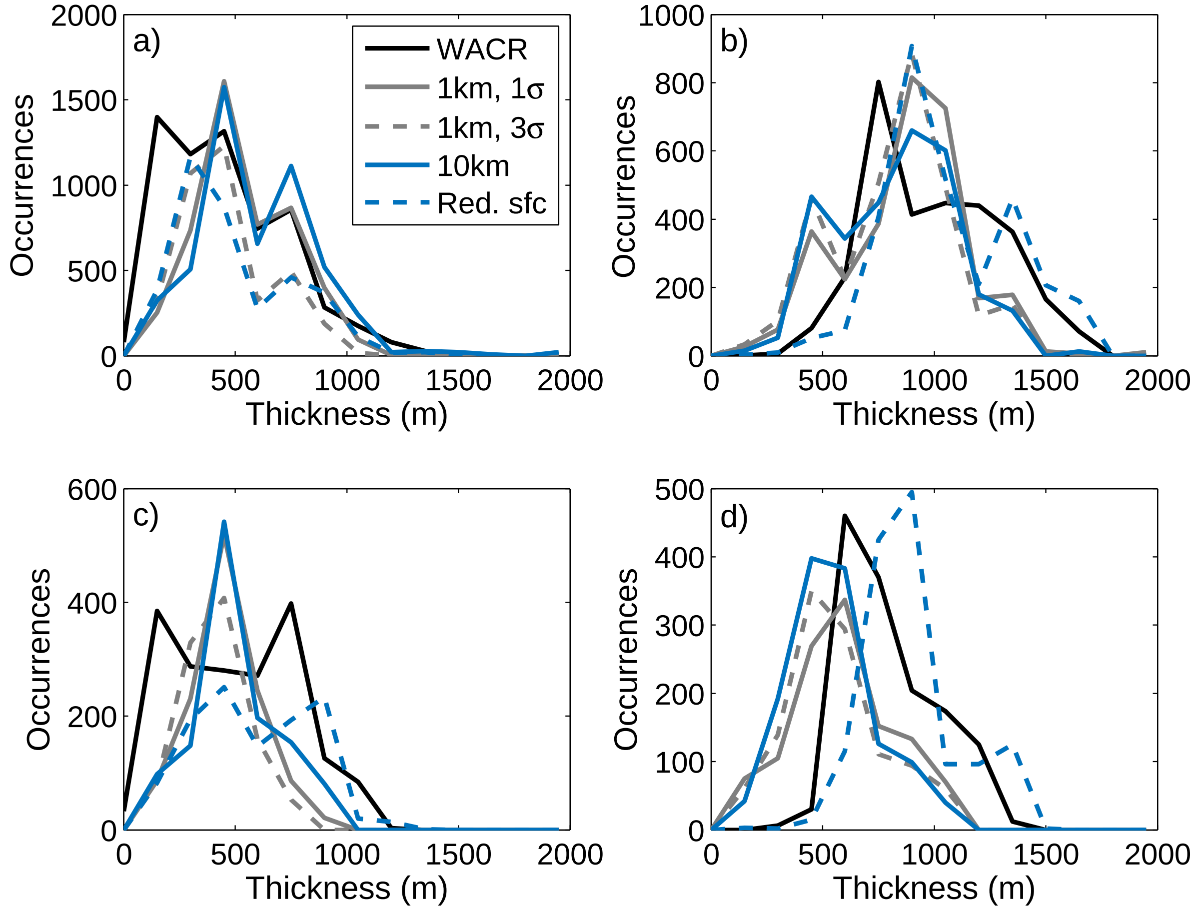

Figure 3.3: Average cloud thickness distributions measured by WACR and EC-CPR for a) GRW non-drizzling; b) GRW drizzling; c) MAGIC non-drizzling; and d) MAGIC drizzling profiles.

Figure 3.3 displays distributions of cloud thickness measured by the ground-based radars and the EC-CPR in the GRW and MAGIC campaigns, separated into drizzling and non-drizzling profiles (the presence of drizzle within a profile is determined with a height-dependent reflectivity threshold following Wang and Geerts, 2003). For comparison purposes, the WACR and EC-CPR data has been sampled to the same 500 m along-track grid. The EC-CPR distributions are produced using $1\,\sigma$ and $3\,\sigma$ thresholds at a 1 km integration and a $3\,\sigma$ threshold at a 10 km integration (the reduced surface clutter distribution uses a 1 km integration and $3\,\sigma$ threshold). The ground-based radar drizzle-free distributions (Fig. 3.3a, 3.3c) are centered around thicknesses of approximately 500 m, with a significant number of profiles having thicknesses less than this. The EC-CPR, in all configurations, under-represents these clouds due to a combination of stretching by the EC-CPR range resolution and low reflectivity, thin clouds not being detected. Average cloud thicknesses are generally lower in the 1 km, $3\,\sigma$ distributions due to the lower sensitivity that often masks the cloud boundaries, especially in the relatively low reflectivities at cloud base.

The drizzling distributions are broader, particularly for the GRW campaign (Fig. 3.3b) where thicknesses reach almost 2 km, compared to below 1.5 km for MAGIC (Fig. 3.3d). Here, the best fit to the WACR data is achieved by the reduced-surface simulation data, with all other distributions biased toward lower thickness values. This indicates that the surface clutter obscures cloud base in drizzling profiles (as is seen in Figures 3.1 and 3.2), leading to an over-representation of relatively thin drizzling clouds. In both drizzling and non-drizzling clouds, the accuracy of the EC-CPR-derived cloud top height is calculated (with the WACR cloud top averaged to 500 m along-track, the initial sampling rate of the EC-CPR, and the EC-CPR data oversampled to this rate). As with cloud fraction, this depends upon the value of the FM detection threshold. Using a 1 km along-track integration and $3\,\sigma$ threshold, the root-mean-square-error (RMSE) is 91 m and the mean bias is +58 m (an overestimation by the EC-CPR). Using the same integration length and a $1\,\sigma$ threshold, the values are 128 m and 106 m respectively, i.e., when a $1\,\sigma$ threshold is in use, the cloud top overestimation bias is comparable in magnitude to the range sampling rate of the EC-CPR (100 m).

Distributions of reflectivity factor are shown in Figure 3.4, separated between GRW and MAGIC campaigns in drizzling and non-drizzling resolution volumes. In drizzling volumes where SNR is large, the choice of significant detection threshold has no effect on the distribution and the 1 km, $1\,\sigma$ and 1 km, $3\,\sigma$ distributions are the same. In volumes containing drizzle (Fig. 3.4b, 3.4d), reflectivities peak at approximately -10 dBZ in the ground-based radar distributions, and this is well matched by the EC-CPR in all config-

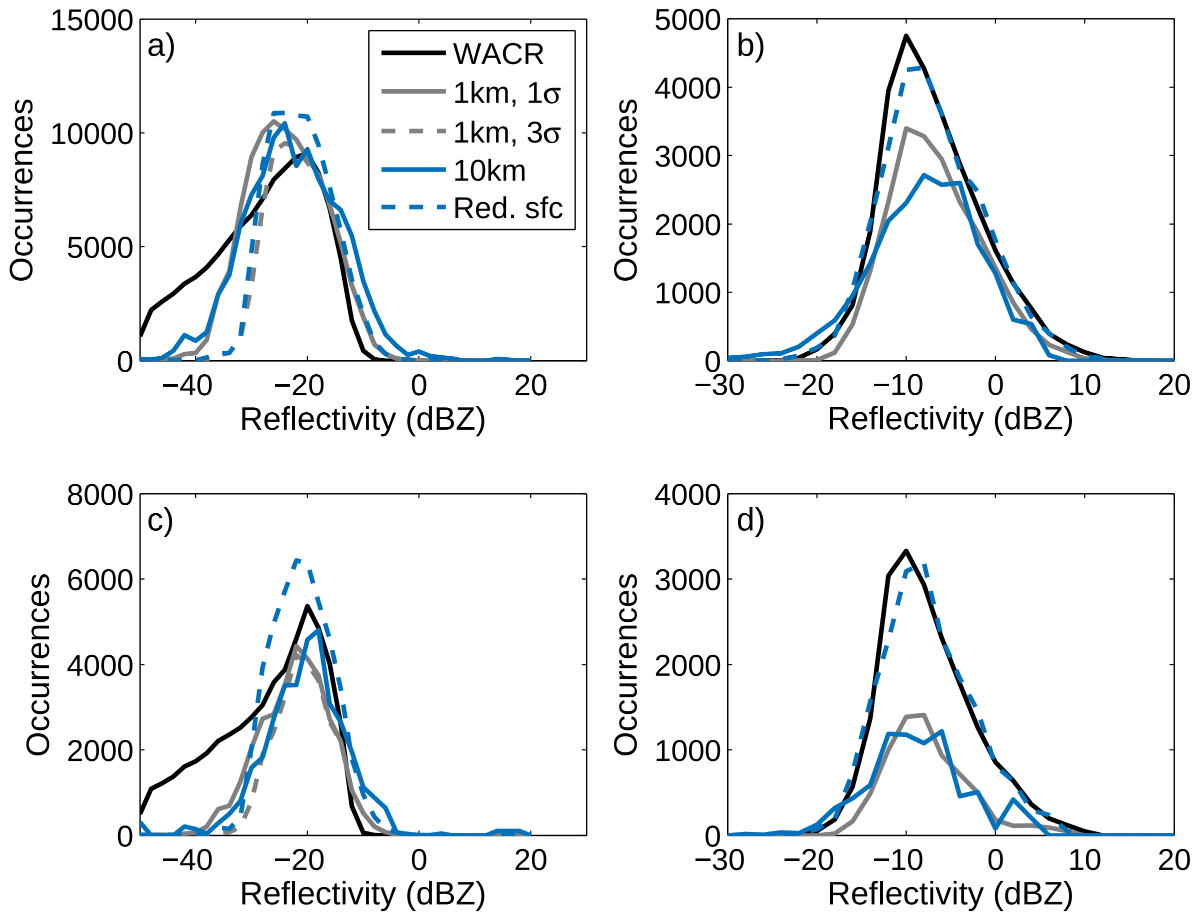

Figure 3.4: Average reflectivity distributions measured by WACR and EC-CPR for a) GRW non-drizzling; b) GRW drizzling; c) MAGIC non-drizzling; and d) MAGIC drizzling resolution volumes.

urations. The total number of volumes shown is lower for the EC-CPR, excepting the reduced-surface cases, due to surface clutter effects on the detection of drizzle. In non-drizzling volumes (Fig. 3.4a, 3.4c), the location of the peak reflectivity (around -20 dBZ) is again reproduced by the EC-CPR, however greater reflectivities (-10 dBZ and above) are over-represented in the EC-CPR distributions. This is an effect of the EC-CPR resolution, both along-track and along-range, which introduces strong drizzle signals into cloud-only volumes. At lower reflectivities, the effect of an increased integration or lower significant detection threshold on the EC-CPR sensitivity can be seen, which allows more low-reflectivity features to be captured relative to the 1 km, $3\,\sigma$ configuration. However, the ground-based radar distributions contain reflectivity values of -40 dBZ and below, which the EC-CPR fails to reproduce in any configuration tested here.

As with cloud top height, the average deviation between the EC-CPR and WACR reflectivity factor values is calculated (in logarithmic units). At a 1 km horizontal integration and $3\,\sigma$ threshold, the bias is +2.1 dB (an overestimate by the EC-CPR) and the RMSE is 5.9 dB. Using a $1\,\sigma$ threshold at the same along-track integration, these values are 2.0 dB, and 5.9 dB. The primary source of these errors, which are calculated by oversampling the EC-CPR data to the WACR grid, are the range and along-track resolution of the EC-CPR, which introduce strong drizzle signals into the surrounding low-reflectivity cloud, with biases of more than 10 dB often present in the uppermost portions of drizzling clouds. After averaging the WACR data along-track to the EC-CPR horizontal resolution (including both along-track integration and antenna pattern effects), the average EC-CPR bias is 1.3 dB ($1\,\sigma$ threshold), which is attributable to the range resolution and some small noise and surface clutter residuals (the introduction of a net bias due to averaging is explained by the fact that the logarithm of a linear average is always greater than the average of the logarithms).

3.2 Drizzle Doppler velocity

Simulated EC-CPR velocities for a region within the GRW example case are shown in Figure 3.5 at 1 km and 5 km integration lengths. The EC-CPR velocities are significantly affected by random noise, particularly at short integration lengths. Because of this, velocities within regions where SNR is low do not reproduce the true velocity field well. Integrating the data to 5 km reduces much of the noise-induced variance, better representing the average velocities, however, in doing so, this also removes much of the variability observed in the WACR velocity field.

Figure 3.6 shows CFADs of the velocity fields of the two example scenes as measured by the WACR and the EC-CPR at 1 km and 5 km integrations. In both cases, the WACR data (Fig. 3.6a and 3.6b) show similar relations of velocity with height. At cloud top, velocities tend toward zero. As altitude decreases, the mean velocity decreases to around -1 $\mathrm{m\,s^{-1}}$ at 500 m where drizzle dominates. The mean WACR velocities as a function of height are overlaid on the EC-CPR CFADs, showing that the EC-CPR tends to overestimate the magnitude of the mean velocity. This is a result of the observed velocity being a reflectivity-weighted average of the velocities within each CPR volume; it is therefore biased towards velocities corresponding to higher reflectivities, which typically are larger in magnitude. In the lowest 1 km, the surface echo, with its stationary velocity and strong reflectivity, biases the reflectivity-weighted velocity towards zero, reducing the usefulness of EC-CPR velocity estimation for weak targets in the lowest 1 km. The 5 km integration distribution is much narrower than that of the 1 km integration, consistent with what is seen in Figure 3.5.

The average EC-CPR velocity error across all cases is calculated by averaging the WACR data to the 500 m along-track initial sampling rate of the EC-CPR and oversampling the EC-CPR data to the same rate. In keeping with the effects observed in Figures 3.5 and 3.6, volumes within the lowest 1 km or where SNR is less than one are ignored for this calculation. At 1 km and 5 km integrations, the EC-CPR RMSE is 0.97 $\mathrm{m\,s^{-1}}$ and 0.49 $\mathrm{m\,s^{-1}}$ respectively (prior to correcting for NUBF, these values are 1.37 $\mathrm{m\,s^{-1}}$ and 0.59 $\mathrm{m\,s^{-1}}$). Given the velocity magnitudes seen in Fig. 3.6 of close to 1 $\mathrm{m\,s^{-1}}$, this suggests a minimum along-track integration of 5 km is necessary to reduce the EC-CPR velocity uncertainty to reasonable levels. These errors are produced from simulations using a PRF of 7.0 kHz, a mid-range value within the EC-CPR 6.1–7.5 kHz operating range. At a PRF of 6.5 kHz, the NUBF-corrected RMSEs are 1.22 $\mathrm{m\,s^{-1}}$ and 0.59 $\mathrm{m\,s^{-1}}$ at 1 km and 5 km integrations. At the maximum PRF of 7.5 kHz, these errors are 0.82 $\mathrm{m\,s^{-1}}$ and 0.48 $\mathrm{m\,s^{-1}}$. The uncertainty produced using a 5 km integration is only slightly improved by increasing the PRF from 7.0 kHz to 7.5 kHz, potentially due to the EC-CPR averaging of the true velocity field, through the range resolution, horizontal resolution, and along-track integration, which introduces a lower limit to the uncertainty in the EC-CPR velocity that depends on the variability within the true cloud scene. Changes in PRF were found to have negligible impact on cloud detection and cloud boundary estimation.

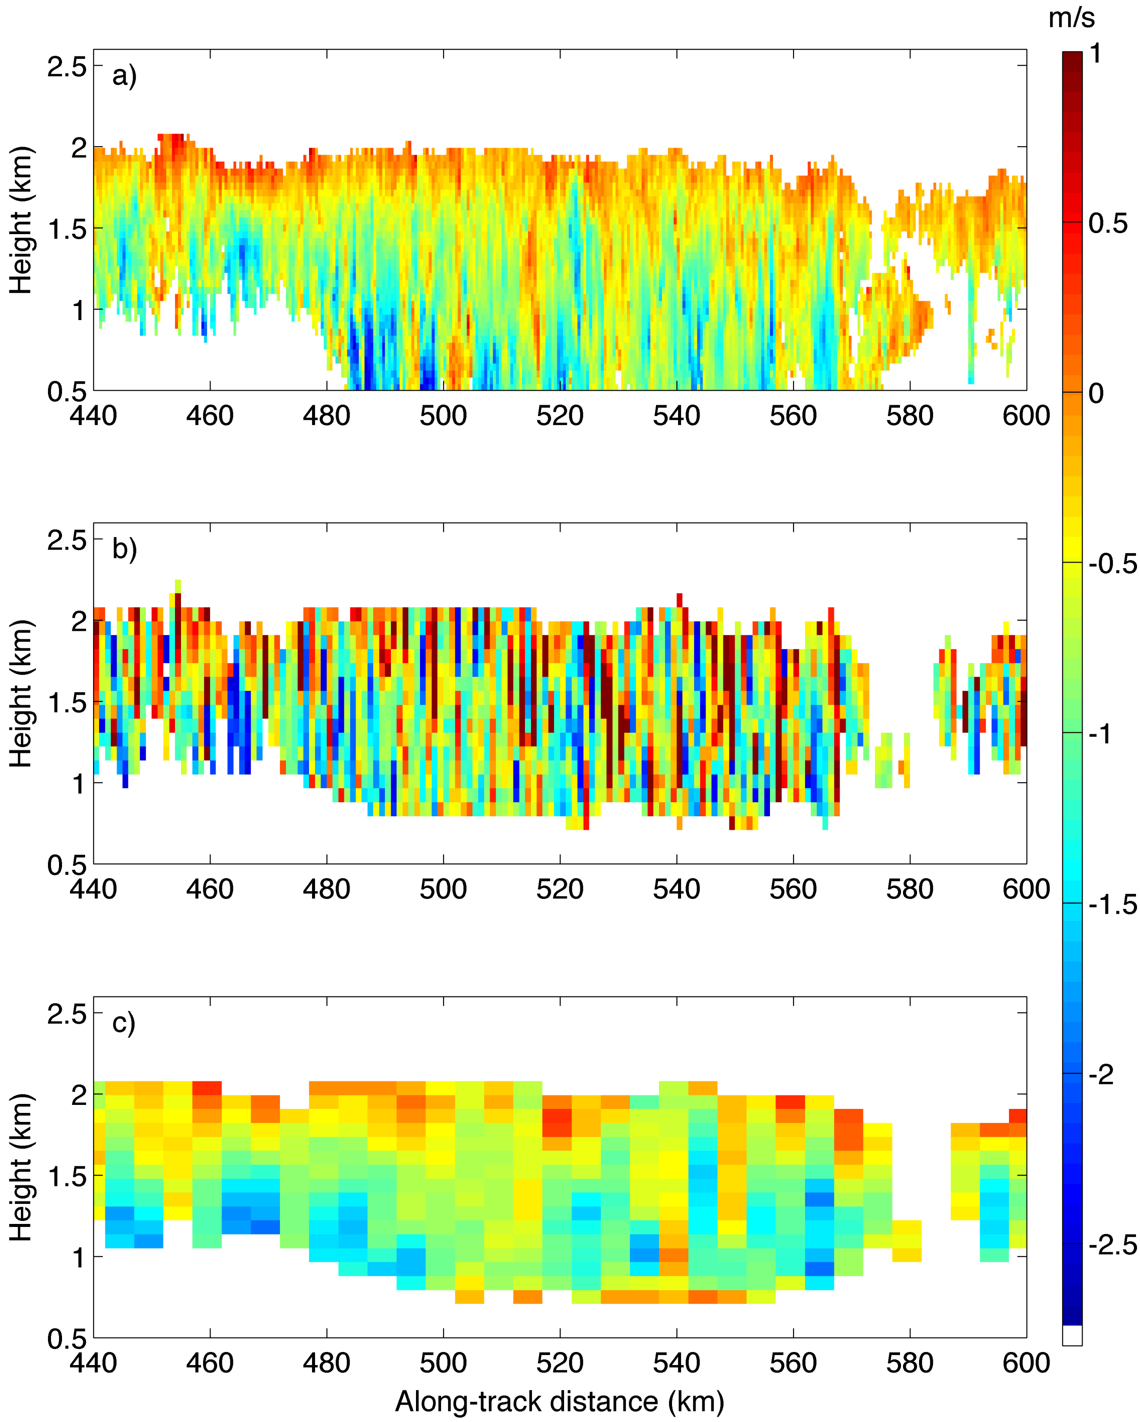

Figure 3.5: Along-track-height velocity fields of GRW case. a) WACR-measured velocity; simulated EC-CPR velocity at b) 1 km integration, and c) 5 km integration.

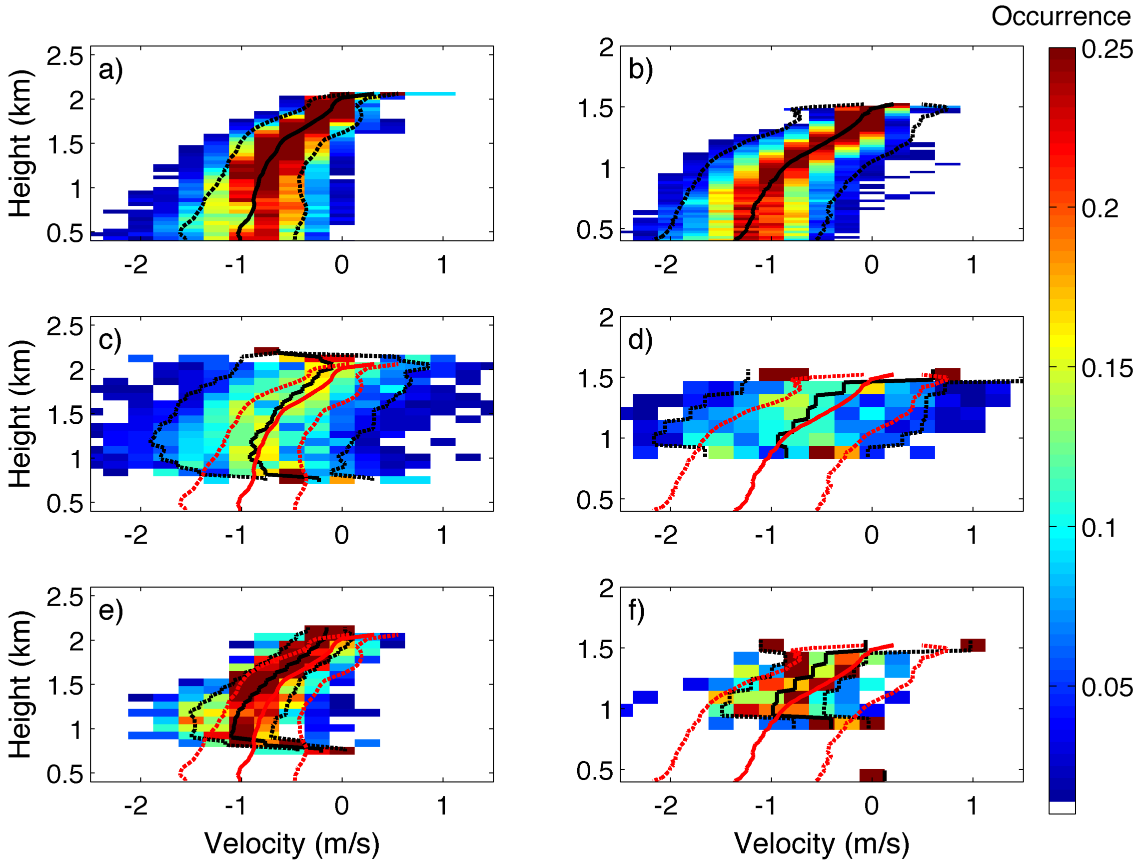

Figure 3.6: (a, b): Contoured-frequency-by-altitude-diagrams (CFADs) of WACR-measured Doppler velocity from (a) GRW and (b) MAGIC example cases shown in Figure 2.1b and 2.1d; (c, d): CFADs of EC-CPR synthetic Doppler velocities from (c) GRW and (d) MAGIC at 1 km integration and (e, d): CFADs of EC-CPR synthetic Doppler velocities from (e) GRW and (f) MAGIC at 5 km integration. Solid and dotted black lines depict mean and standard deviation of the velocity as a function of height. Mean WACR velocities are reproduced as red lines in the EC-CPR CFADs.

Chapter 4

Post-processing algorithms

4.1 Range resolution deconvolution

In Chapter Three, it was seen that the EC-CPR measurements of reflectivity factor (obtained with a 1 km integration and a 1 σ detection threshold) contained an average bias of 2.0 dB and a RMSE of 5.9 dB when compared with the input WACR reflectivities. While these errors can be partially attributed to the residuals of the noise and surface clutter, they are primarily a result of the averaging effects of the EC-CPR resolution, both in range and along-track. The radar's range resolution also causes an overestimation bias in cloud top of approximately 100 m when using a 1 σ significant detection threshold. However, due to the range oversampling in use by the EC-CPR (sampling every 100 m within the 500 m range resolution), it is possible to retrieve an improved-resolution reflectivity profile in which these errors are reduced. Here, a technique to achieve this via a constrained linear inversion, or deconvolution, of the EC-CPR range resolution is presented and tested on the simulated data introduced in Chapter 2 and evaluated in Chapter 3.

The effect of the EC-CPR range resolution on the observed profile of reflectivity may be expressed mathematically as a convolution of the true profile of reflectivity Z_{true} (in linear units) with the range weighting function $G(r)$ of the EC-CPR. For the EC-CPR, the range weighting function is assumed to be a Gaussian with a full-width-half-maximum of 500 m (and is modelled as such in the EC-CPR simulator described in Chapter Two).

$$Z_{CPR}(r) = \int_{r-\Delta r}^{r+\Delta r} G(r'-r) Z_{true}(r') dr' \tag{4.1}$$

In Equation 4.1, $Z_{CPR}(r)$ and Z_{true} are the EC-CPR-observed and true reflectivity at a range r from the radar. $G(r'-r)$ is the normalised range weighting function at a displacement $r'-r$ from r and $2\Delta r$ is the width of G (i.e., 500 m). Negating the effects of the satellite motion, the equivalent equation for the observed EC-CPR (reflectivity-weighted) Doppler velocity may be obtained by including the true profile of Doppler velocity in the integral in Equation 4.1 and dividing the integral by Z_{CPR}.

Equation 4.1 describes the continuous, analytical form of the observed profile of reflectivity as a function of range. In reality, however, the EC-CPR samples the reflectivity at a discrete rate (100 m). The retrieved high-resolution profile of reflectivity will be limited to this same sampling rate. The observed profile of reflectivity is therefore approximated in Equation 4.2 as a discrete convolution of the range weighting function and the true profile of reflectivity, both sampled at the EC-CPR range sampling rate.

$$Z_{CPR}[n] \approx \sum_{m=-dn}^{dn} G[m] Z_{true}[m+n] \tag{4.2}$$

In Equation 4.2, $Z_{CPR}[n]$ and $Z_{true}[n]$ are the observed and true reflectivities at the n^{th} range gate and G is the discrete approximation of the range weighting function with a range-gate-width of $2dn+1$ (for the EC-CPR, $dn = 2$ range gates). Equation 4.2 may be equivalently expressed as a matrix-vector equation, as in Equation 4.3.

$$\mathbf{Z}_{CPR} \approx \mathbf{A}\ \mathbf{Z}_{true} \tag{4.3}$$

Here, $\mathbf{Z}_{CPR}, \mathbf{Z}_{true} \in \mathbb{R}^N$ are vectors representing the profiles of the observed and true reflectivity respectively, each consisting of N range gates, and $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the matrix representation of the discrete approximation of the range weighting function, in which each row is constructed of G from Equation 4.2 centred on the diagonal element. If the two sides of Equation 4.3 were exactly equal, the profile of true reflectivity could be retrieved by left-multiplying both sides by $\mathbf{A}^{-1}$, the inverse of $\mathbf{A}$. In practice, however, $\mathbf{Z}_{CPR}$ contains small errors arising from noise and the finite sampling approximation that may be considered an additional term on the right hand side of Equation 4.3. These errors are amplified when multiplied by the inverse of $\mathbf{A}$, producing large errors in the retrieved reflectivity profile $\mathbf{Z}_{true}$ that typically manifest as high-frequency oscillations about zero – for reflectivity, which is a measurement of backscattered power, these are non-physical. This is a well-documented issue in remote sensing that may be mitigated by applying a constrained linear inversion technique (Phillips, 1962; Twomey, 1977). A brief description of the constrained linear inversion method is now given.

For an equation of the form $\mathbf{b} = \mathbf{A}\ \mathbf{x}$ in which $\mathbf{x}$ is to be solved for and in which there is some noise in $\mathbf{b}$ such that a left-multiplication of $\mathbf{A}^{-1}$ is inappropriate, a level of acceptable error in the residual, $\|\mathbf{b} - \mathbf{A}\ \mathbf{x}'\|$, is defined, where $\mathbf{x}'$ is some approximate solution. Then, from the family of solutions that produce an error less than or equal to this acceptable value, the solution that minimises some other constraint condition is chosen. The constraint condition is chosen according to prior knowledge of the form the true solution should take, e.g., a smooth solution can be obtained by minimising the magnitude of the second differences of $\mathbf{x}'$.

The application of a constrained linear inversion to oversampled radar data in order to improve the resolution is not new to this thesis. Schutgens and Donovan (2004) tested the technique, as here, on simulated EC-CPR data as a means of mitigating the effect of the EC-CPR's long pulse length while in Galati et al. (1996) the method was applied to scanning radar data with a focus on increasing the vertical resolution. However, the form of the constraint and how the level of acceptable error in the solution is set varies between approaches. Here, in constraining the retrieved reflectivity profile to be reasonably physical (i.e., no negative values and relatively smooth), we take advantage of the fact that the original observed profile, $\mathbf{Z}_{CPR}$, already fulfils these criteria. Therefore, from the family of solutions with an acceptable error, we choose the solution that minimises the departure from the observed profile. This is one of two possible constraints proposed in Schutgens and Donovan (2004) (the second being to minimise the deviation of the retrieved profile from a zero reflectivity profile), however, in this work, the difference between the retrieved and observed profile is calculated as the sum of the relative, rather than absolute, differences between the two profiles at each range gate. This distinction is necessary due to the orders of magnitude variability of the reflectivity within a profile; if only the absolute differences between the profiles are considered, the total deviation is dominated by high-reflectivity range gates and the inversion method performs poorly at low-reflectivity range gates, which are under-constrained. Applying this constraint, the high-resolution reflectivity profile, $\mathbf{Z}'_{true}$, is obtained from Equation 4.4 (Twomey, 1977).

$$\mathbf{Z}'_{true} = (\mathbf{A}^T\ \mathbf{A} + \gamma J)^{-1}(\mathbf{A}^T\ \mathbf{Z}_{CPR} + \gamma\ \mathbf{Z}_{CPR}) \tag{4.4}$$

$\mathbf{A}^T$ is the transpose of matrix $\mathbf{A}$ (which, given $\mathbf{A}$ is symmetric, is exactly equal to $\mathbf{A}$), $\mathbf{J} \in \mathbb{R}^{N\times N}$ with $J_{ii} = (\mathbf{Z}_{CPR,i})^{-2}$ and $J_{ij} = 0$, $i \neq j$, and γ is a real scalar that dictates how closely the solution $\mathbf{Z}'_{true}$ is forced toward $\mathbf{Z}_{CPR}$. γ is also directly related to the level of error in the residual, $\|\mathbf{Z}_{CPR} - \mathbf{A}\ \mathbf{Z}'_{true}\|$, and therefore, by choosing an acceptable error level, the appropriate value of γ can be computed. Alternatively, defining an acceptable error threshold may be bypassed by using some other criteria to independently set γ. It is found that suitable values of γ vary from approximately 10^{-7} in low-reflectivity, cloud-only profiles to 10^{-2} or larger in drizzling profiles. Using a value of γ too large or too small results in either over-constraining the retrieved profile, in which case it will almost exactly match the observed profile, or under-constraining the profile, leaving oscillations (and negative values) in the solution; therefore, it is necessary to tailor γ according to the conditions in each profile. The steps used to apply the constrained inversion and set γ are as follows.

Within each profile, $\mathbf{Z}_{CPR}$ is formed of the masked reflectivity profile of a single cloud layer plus dn (i.e., two) hydrometeor-free range gates above and below, with the reflectivity in these range gates set to a very low but non-zero value. The matrices $\mathbf{A}$ and $\mathbf{J}$ are produced, γ is given some initial, small value (e.g., 10^{-7}), and the first estimate of the profile $\mathbf{Z}'_{true}$ is computed. If this contains any negative values, γ is increased (here, by a factor of 1.25) and the solution recalculated. This process is repeated until the retrieved profile contains no negative reflectivities (in practise, very small negative values in $\mathbf{Z}'_{true}$ may occur due to rounding errors and so are allowed; these are set to zero prior to the next step). The magnitude of the retrieved profile is, in general, lower than that of the original observed profile, particularly for large values of γ. This is an effect of the inversion method, rather than the range resolution of the EC-CPR, which is an averaging that should conserve total power. Therefore, the magnitude of the retrieved profile $\mathbf{Z}'_{true}$ is normalised to that of $\mathbf{Z}_{CPR}$ by Equation 4.5. Finally, any range gates with reflectivity below the previously-estimated standard deviation of the noise in that profile are assigned reflectivities of zero.

$$\mathbf{Z}'_{true} = \mathbf{Z}'_{true} \frac{\sum_i \mathbf{Z}'_{true,i}}{\sum_i \mathbf{Z}_{CPR,i}} \tag{4.5}$$

The range resolution deconvolution is not applied in cloud layers in which the cloud base is below 1 km or the thickness of the layer is less than the range resolution (five range gates) as these conditions imply a significant error in Equation 4.3 due to surface clutter or noise respectively. Large errors in Equation 4.3, in contrast with the relatively small errors considered previously, generally cause solutions with physical characteristics to have large residual errors. This is equivalent to a large value of γ being necessary for a reasonably well constrained solution, which implies that the retrieved solution will be essentially equal to the observed solution, i.e., the application of the deconvolution algorithm will, at best, leave the reflectivity profile unchanged and, at worst, produce a non-physical solution.

4.1.1 Results

The deconvolution method is applied to the simulated EC-CPR data set described in Chapters Two and Three. Across all simulated EC-CPR cloud scenes, in cloud layers where the conditions for the application of the method were met (i.e., cloud layer of at least five range gates and cloud base above 1 km), the average reflectivity bias was reduced from 1.5 dB to 0.8 dB and the RMSE was reduced from 5.3 dB to 4.8 dB. The remaining bias and uncertainty are attributable to a combination of the residual range resolution errors and errors due to the horizontal averaging. If the effect of the latter is removed by averaging the WACR data along-track to the EC-CPR horizontal resolution (including both the antenna pattern and along-track integration), the bias and RMSE are reduced from 1.2 dB and 4.8 dB to 0.1 dB and 4.1 dB respectively by applying the deconvolution. The remaining bias and uncertainty are a mix of residuals of the method and errors due to surface clutter and noise. In EC-CPR-derived cloud top, in cloud layers where the deconvolution was applied, the average overestimation bias of 125 m (using a $1\,\sigma$ significant detection threshold) was reduced to 35 m and the RMSE was reduced from 140 m to 60 m. Examples of the results for two profiles, one cloud-only and one drizzling, are shown in Figure 4.1, which compares the input WACR profiles, the observed EC-CPR profiles, and the deconvolved EC-CPR profiles. The improvement achieved in the deconvolved profile can be seen in the location and magnitude of the reflectivity peaks and the derived cloud boundaries which more closely match those of the WACR profiles.

While these results are promising, the overall effectiveness of the method is limited by its applicability. The specified conditions for the application of the deconvolution algorithm are met in only 25 % of the profiles considered here. Many drizzle-free clouds are precluded due to their observed thickness being less than 500 m, while drizzling profiles commonly extend into the lowest 1 km where surface clutter is present. The application of the inversion technique in this data set is therefore typically limited to lightly-drizzling clouds only. However, in higher level, thicker clouds such as cirrus and precipitating systems, the conditions for application should be met far more often.

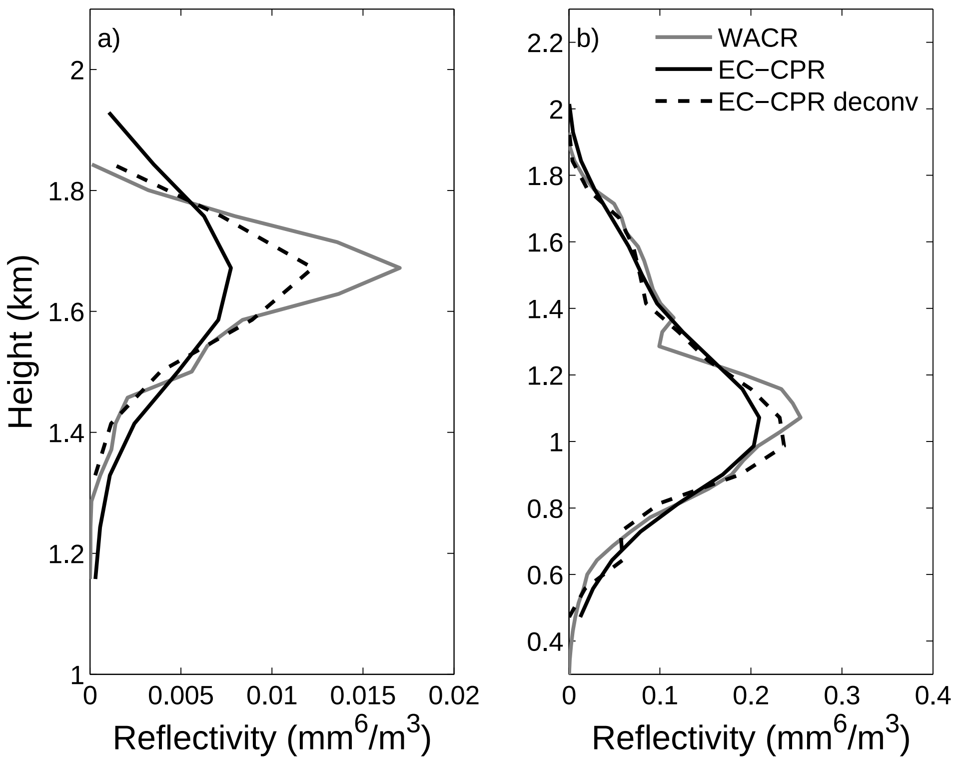

Figure 4.1: Profiles of WACR, EC-CPR, and deconvolved EC-CPR reflectivity in a) cloud and b) drizzle. Note that surface clutter has not been included in these EC-CPR simulations.

4.2 Matched filtering of Doppler velocities

Previously, a 1 km along-track integration of the EC-CPR Doppler velocity was found to produce an uncertainty of approximately $1\,\mathrm{m\,s^{-1}}$, comparable to the magnitudes of the Doppler velocities in drizzle. This error is primarily due to frequently occurring low SNR conditions and signal decorrelation due to the satellite motion. By integrating the data to 5 km along-track, the error is reduced to $0.5\,\mathrm{m\,s^{-1}}$. However, in doing so, the resolution of the data is decreased and small scale variations in the true velocity field are lost; this effect is seen in Figure 3.6, in which the 5 km integration EC-CPR distributions are significantly narrower than those of the WACR data. Therefore, a method of achieving uncertainties similar to those of a 5 km integration while preserving a higher along-track resolution is desirable.

In Sy et al. (2014), the application of matched filters to simulated EC-CPR Doppler velocity measurements was found to achieve uncertainties of less than $0.5\,\mathrm{m\,s^{-1}}$ while maintaining the initial 500 m along-track integration of the EC-CPR. Here, the method is recreated with minor changes and applied to the simulated EC-CPR scenes described in Chapter Two as a verification of its performance in marine stratiform clouds.

In general, a matched filter is one that maximises the SNR for a signal contaminated by noise, typically requiring prior knowledge of the form of the true signal. For radar Doppler velocity measurements, this produces a reduction in random noise-induced errors. In the approach described in Sy et al. (2014), the EC-CPR scene is first split into along-track sections of 100 km in which the filter is matched and applied separately. The discrete spatial Fourier transform of the lag-one pulse pair correlation is produced at each level within the section and multiplied by the matched filter. The filter, a function of spatial frequency and defined in Equation 4.6, increases signal strength at frequencies where true signals are present and suppresses the signal at frequencies where noise dominates, thereby increasing the SNR. The inverse discrete Fourier transform of the altered spectrum is performed to retrieve the lag-one pulse pair correlation and the Doppler velocity is estimated via pulse-pair processing (Doviak and Zrnic, 1993). This process is illustrated for one level in an along-track section in Figure 4.2.

$$L(f) = \frac{1}{1 + |\alpha f|^{\beta}} \tag{4.6}$$

The filter, $L(f)$, is a function of spatial frequency f and defined by two variable parameters, α and β. For a 100 km along-track section sampled every 500 m, the spatial frequency f ranges from -1 km^{-1} to +1 km^{-1} in increments of 0.01 km^{-1}. By choosing appropriate values of α and β, which are varied in the ranges 0.01–100 km and 0.5–3 respectively, the filter is matched to the true signal. However, this requires the form of the true signal, i.e., the true velocity field, to be known, which is generally not the case. Therefore, filter optimisation criteria that depend only on the pre- and post-filter velocity fields and knowledge of the form of the noise are necessary. Several optimisation procedures that adhere to this restriction were proposed in Sy et al. (2014). Here, the Constrained Residue Variance Minimising, or RV, filter, which achieved the lowest velocity errors in Sy et al. (2014), is recreated and tested on the simulated EC-CPR data.

The RV filter uses the difference between the pre- and post-filter velocity fields as a proxy for the difference between the true and post-filter velocity fields. By minimising the former, which is always available, the latter is assumed to be minimised. However, the equivalency between the two differences is only true for certain filter shapes, and therefore the range of possible filter shapes (i.e., combinations of α and β) must be constrained. The filter constraint method, described in detail in Sy et al. (2014), requires comparing the cumulative density function (CDF) of the difference between the pre- and post-filter velocity fields with a reference CDF of the difference between a set simulated EC-CPR velocities, again before and after filtering. For a given filter shape, if the difference between the two normalised CDFs exceeds a threshold value at any point, this indicates that the difference between the pre- and post-filter velocity fields is not directly related to the difference between the post-filter and true velocity fields and this combination of α and β is excluded. Of the filter shapes for which these two quantities are directly related, the filter that minimises the difference between pre- and post-filter velocity fields is selected as the best-matched filter. A limitation to the RV method is that, for the constraint threshold condition to be reliable, a reasonably smooth CDF is required, i.e., a large set of EC-CPR data points (with SNR > 1) are necessary. If there are relatively few appropriate data points in an along-track section, the pre- versus post-filter velocity CDF is coarse and, in general, there are large differences between the two CDFs such that all filter shapes are excluded. In sections where this is the case, an alternative filter optimisation method, new to this work, is used. This method, hereafter referred to as an error optimisation (EO) filter, is also tested as a standalone optimisation approach for comparison purposes.

In the EO method, the filter shape remains unconstrained (i.e., any combination of α and β is allowed), and a minimisation criteria (different to that of the RV filter optimisation) is used. By not applying any constraint threshold, the method can be reliably applied even when relatively few data points are available. In this approach, in each along-track section, the SNR distribution of the EC-CPR data is estimated from the reflectivity field (considering only volumes with SNR greater than 0 dB). The aim is to then estimate the distribution of the random noise errors that should occur for this SNR distribution, i.e., to estimate the probability density function (PDF) of $V_{PP} - V_{TRUE}$ where V_{PP} is the pulse-pair derived EC-CPR velocity field and V_{TRUE} is the true velocity. The filter that produces the PDF of $V_{PP} - V_F$, where V_F is the filtered velocity, that best matches this is then chosen. To estimate the true error PDF, a large number of EC-CPR velocities are simulated from a known velocity ($0\,\mathrm{ms}^{-1}$ is chosen for simplicity), following the method in Kollias et al. (2014). This is done for each value of SNR (binned in 1 dB increments) within the SNR distribution of the along-track section. The velocity error PDFs at each SNR are normalised and summed, weighted by the true SNR distribution of the observed data, producing an estimate of the true error PDF of the EC-CPR velocity field. The filter shape that produces the pre- versus post-filter velocity distribution that most closely matches the true error distribution (in a root-mean-square-difference sense) is chosen as the best-matched filter and used to to produce the best-estimate filtered velocity field.

4.2.1 Results

Matched filters are applied to the simulated EC-CPR cloud scenes described and evaluated in Chapters Two and Three. In addition to the RV and EO filters, best-matched filters are produced by comparing the filtered velocity to the true velocity field (averaged along-track to 500 m) and selecting the filter that produces the lowest RMSE. This provides the upper limit to the uncertainty reduction achievable through these filters and allows evaluation of the performance of the RV and EO methods of filter matching.

Prior to any filtering, the average EC-CPR Doppler velocity RMSE at a 500 m integration and PRF of 7.0 kHz (for SNR>0 dB and range gates above 1 km) is $1.66\,\mathrm{m\,s^{-1}}$ ($1.79\,\mathrm{m\,s^{-1}}$ and $1.54\,\mathrm{m\,s^{-1}}$ at PRFs of 6.5 kHz and 7.5 kHz respectively). Using the true velocity field to match the filters, this is reduced to $0.43\,\mathrm{m\,s^{-1}}$ ($0.46\,\mathrm{m\,s^{-1}}$; $0.42\,\mathrm{m\,s^{-1}}$). The RV and EO filters produce similar average errors of $0.61\,\mathrm{m\,s^{-1}}$ ($0.63\,\mathrm{m\,s^{-1}}$; $0.66\,\mathrm{m\,s^{-1}}$) and $0.62\,\mathrm{m\,s^{-1}}$ ($0.65\,\mathrm{m\,s^{-1}}$; $0.66\,\mathrm{m\,s^{-1}}$) respectively. As the RV filters are only replaced by the EO filters in along-track sections with relatively few data points, these regions to not contribute significantly to the average RV filter velocity error and so the two optimisation methods are comparable in performance.

The best-matched filters perform significantly better in terms of error reduction than the RV and EO methods due to the fact that the EC-CPR velocity field includes not only random noise-induced errors, which the RV and EO filters aim to suppress, but also NUBF residuals and vertical averaging errors caused by the EC-CPR range resolution. These other sources of error, particularly NUBF residuals, may be reduced through extended horizontal averaging. The best-matched filters, determined by all errors in the EC-CPR velocity field, are therefore typically comparable to lengthy along-track integrations, which reduce the total error but also remove much of the true variability of the velocity field. The RV and EO methods, however, are designed only to reduce the random noise-induced errors and are unaffected by the presence of other sources of uncertainty. These optimisation methods therefore do not achieve the error reduction of the best-matched filters but tend to better preserve the horizontal resolution of the EC-CPR velocity field and thus may be preferable to the best-matched filters despite the associated increase in uncertainty. This is illustrated in Figure 4.3, which compares the EO and best-matched filtered velocity fields, as well as the WACR and EC-CPR 500 m integration velocity fields.

The presence of multiple sources of error in the EC-CPR velocity field also provides a possible explanation for the increase in the filtered velocity errors when the PRF is increased from 7.0 kHz to 7.5 kHz. The RV and EO methods of filter matching rely on the simplifying assumption that random noise is the only error source in the EC-CPR velocity field. When other error sources are negligible, this assumption is approximately valid and the filter performs well. However, as other sources of error become proportionally larger (i.e., when PRF is increased and the random noise errors decrease), this assumption breaks down and the filter matching method is less appropriate. This would also explain the relatively small increase in uncertainty when decreasing the PRF from 7.0 kHz to 6.5 kHz (compared to the equivalent change in error of the 5 km integration, non-filtered velocity field - from $0.49\,\mathrm{m\,s^{-1}}$ to $0.59\,\mathrm{m\,s^{-1}}$), as the changes in filter optimisation and random noise partially cancel out.

The uncertainties achieved through matched filtering here are slightly worse than those reported in Sy et al. (2014) ($0.5\,\mathrm{m\,s^{-1}}$ at the EC-CPR 500 m sampling rate). In that work, a range of cloud types were considered, such as cumulus and cirrus clouds, snowstorms, and stratiform rain, as opposed to marine stratocumuli in this study; this impacts the SNR distribution of the scenes and the sources of error and therefore the performance of the filters. By integrating the RV and EO filtered velocity fields along-track to 1 km, the error is reduced to $0.53\,\mathrm{m\,s^{-1}}$ ($0.57\,\mathrm{m\,s^{-1}}$; $0.52\,\mathrm{m\,s^{-1}}$) and $0.53\,\mathrm{m\,s^{-1}}$ ($0.58\,\mathrm{m\,s^{-1}}$; $0.53\,\mathrm{m\,s^{-1}}$) respectively, comparable to that achieved with a standard 5 km along-track integration of the non-filtered velocity field.

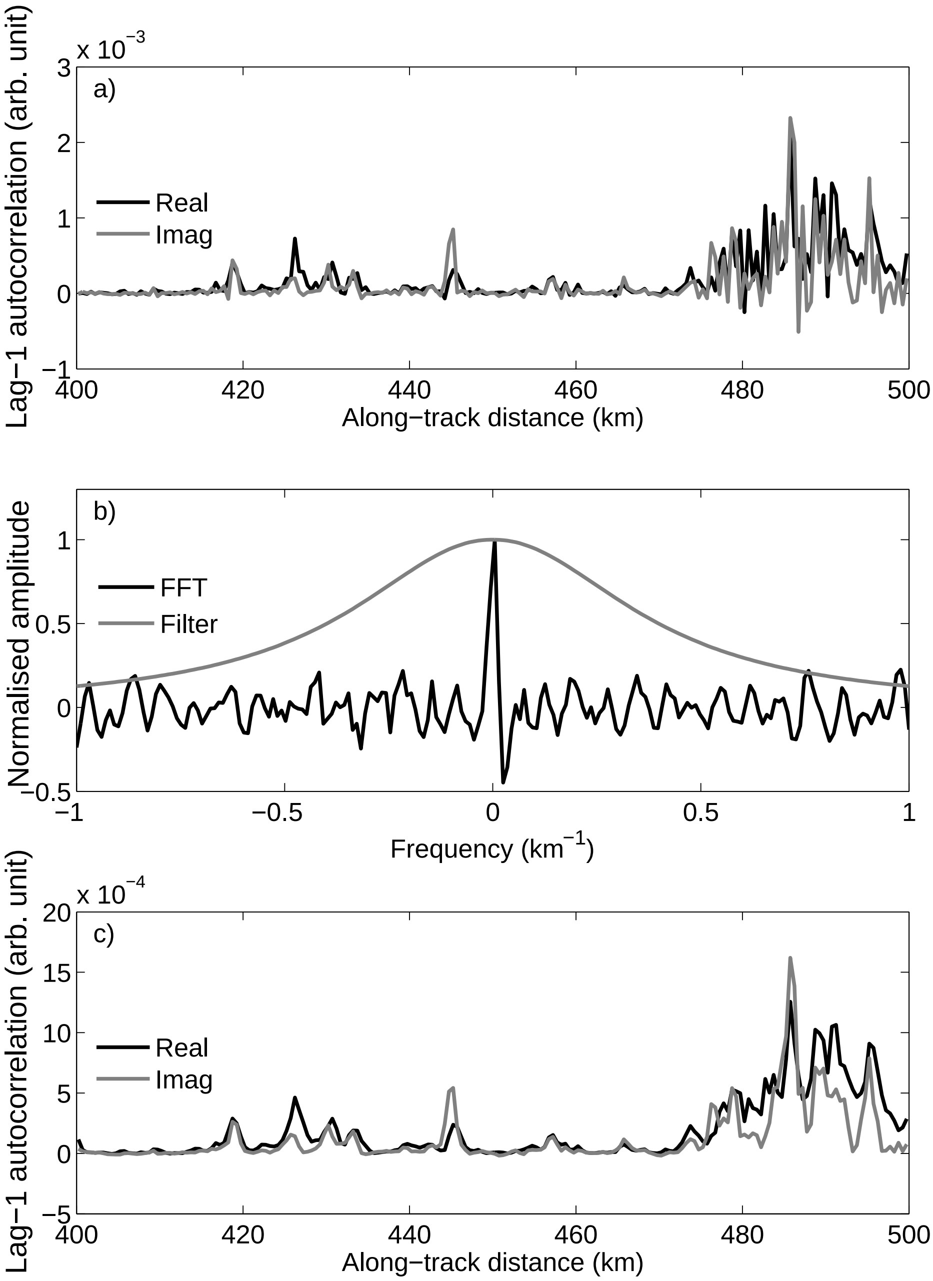

Figure 4.2: Process of matched filtering of EC-CPR Doppler velocities. a) The real and imaginary parts of the initial lag-one pulse pair correlation as a function of along-track position. b) The discrete Fourier transform of the lag-one correlation and the filter as functions of spatial frequency. c) The real and imaginary components of the lag-one pulse pair correlation after filtering.

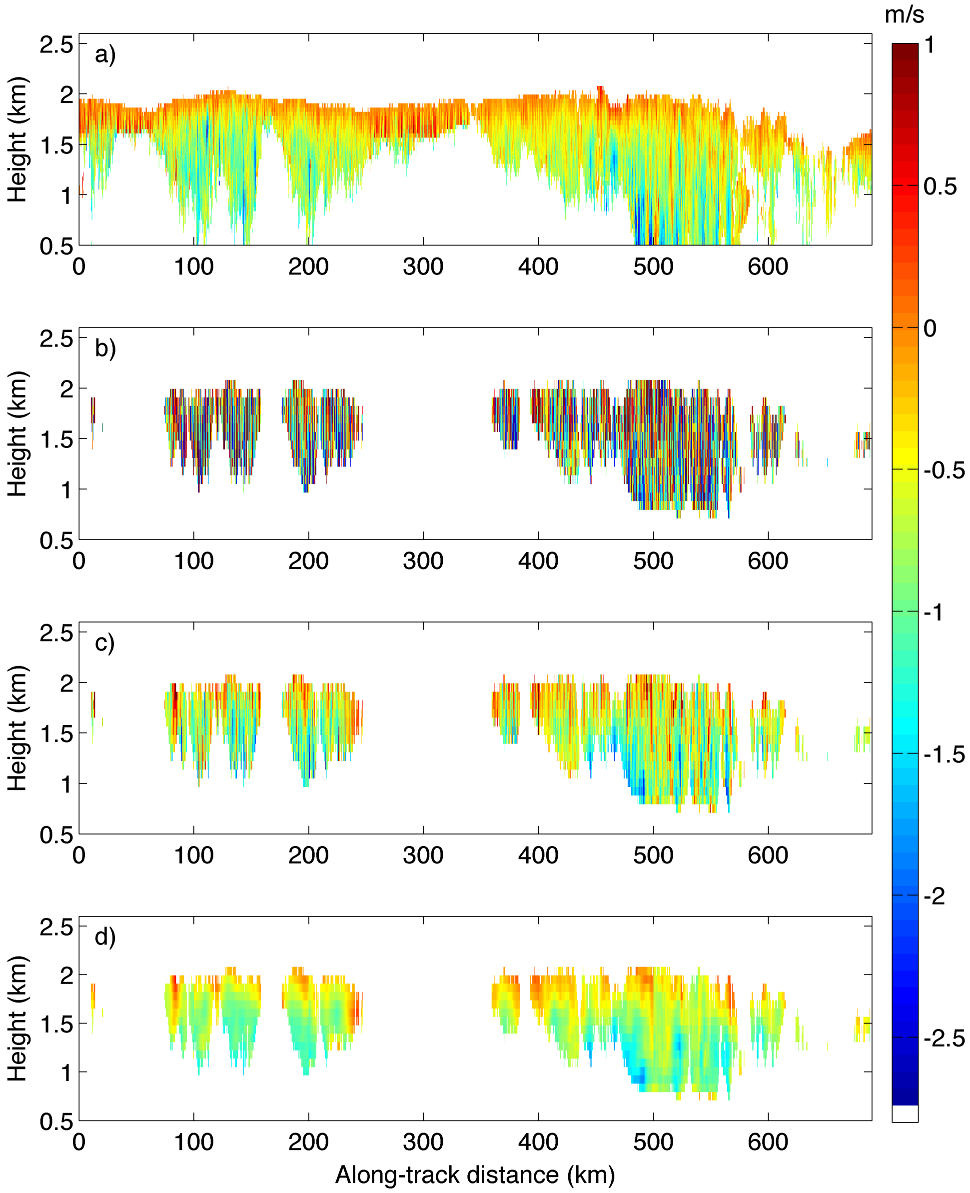

Figure 4.3: Along-track-height velocity fields of a) WACR; b) EC-CPR 500 m integration; c) EC-CPR 500 m integration after EO matched filtering; and d) EC-CPR 500 m integration after best-matched filtering. Only regions of SNR>1 are shown in EC-CPR fields.

Chapter 5

Conclusion

The critical role of marine stratiform clouds on the climate makes them an important target for the EarthCARE Cloud Profiling Radar after its launch in 2018. Accurate measurements of radar reflectivity and cloud boundaries will improve our understanding of the macroscopic structure of these clouds, while high quality Doppler velocity measurements will help us to identify and retrieve drizzle properties. However, radar receiver noise, the proximity of these clouds to the surface, and the coarse radar resolution will limit the quality of the EC-CPR measurements. In this thesis, ground-based radar data of marine stratocumulus cloud scenes are used as input to an EC-CPR simulator and the ability of the EC-CPR to detect marine stratocumulus clouds, retrieve their boundaries and measure their Doppler velocities is evaluated. A range resolution inversion technique and matched filtering of Doppler velocities are tested as possible methodologies of reducing EC-CPR measurement errors.

Along-track integration and the choice of threshold for detection in the FM algorithm are found to significantly impact the EC-CPR results in detected cloud fraction and boundaries. Increasing along-track integration or lowering the detection threshold results in an increase in sensitivity and therefore detected cloud fraction, however the former reduces the EC-CPR's ability to resolve small-scale features and the reflectivity accuracy, and the latter degrades cloud boundary accuracy. For example, a detection threshold of $1\,\sigma$ produces an uncertainty in cloud top height of approximately 130 m and an overestimation bias of 100 m, equal to the EC-CPR range sampling rate, while a threshold of $3\,\sigma$ produces values of 90 m and 60 m respectively. Reflectivity uncertainties are largely unaffected by the choice of FM detection threshold, with an average EC-CPR bias of +2 dB (at a 1 km integration) relative to the ground-based radar data. The bias is caused by the EC-CPR resolution, both along-range and along-track, which introduces strong drizzle signals into low reflectivity cloud regions. By applying a constrained linear inversion, or deconvolution, to the EC-CPR range resolution, the vertical resolution is increased, reducing the cloud top bias to 35 m while maintaining the $1\,\sigma$ detection threshold and almost entirely removing the reflectivity bias caused by the range resolution; if horizontal averaging effects on the bias are accounted for, the deconvolution reduces the average bias from 1.2 dB to 0.1 dB in profiles where it is applied. The conditions for application of the deconvolution are met in approximately one quarter of all marine stratiform cloud profiles in this data set.

The surface clutter masks hydrometeor signals within the lowest 700–800 m, biasing estimates of cloud base height and therefore cloud thickness. Furthermore, the surface's stationary velocity biases velocity estimates within the lowest 1 km towards zero, reducing the usefulness of velocities within these range gates.

Horizontal integration of 5 km is found to be necessary to reduce the average EC-CPR Doppler velocity error to approximately $0.5\,\mathrm{m\,s^{-1}}$, while integration of 1 km produces errors of close to $1\,\mathrm{m\,s^{-1}}$. These errors are not negligible when compared with average velocity magnitudes of $1\,\mathrm{m\,s^{-1}}$ within drizzling clouds. Alternatively, by applying matched filters to the EC-CPR Doppler velocity field, uncertainties of approximately $0.6\,\mathrm{m\,s^{-1}}$ are achievable while preserving the initial EC-CPR along-track integration of 500 m, or $0.5\,\mathrm{m\,ss^{-1}}$ by integrating the filtered velocity to 1 km along-track. The results shown in this thesis will substantially improve the development of the formal ESA L2a EC-CPR products.

Bibliography

A. Battaglia and P. Kollias. Using ice clouds for mitigating the EarthCARE Doppler radar mispointing. *IEEE Trans. Geosci. Remote Sens.*, 53:2079–2085, 2015. doi: 10.1109/TGRS.2014.2353219.

E. E. Clothiaux, M. A. Miller, B. A. Albrecht, T. P. Ackerman, J. Verlinde, D. M. Babb, R. M. Peters, and W. J. Syrett. An evaluation of a 94-GHz radar for remote sensing of cloud properties. *J. Atmos. Oceanic Technol.*, 12(2):201–229, 1995. doi: 10.1175/1520-0426(1995)012¡0201:AEOAGR¿2.0.CO;2.

R. J. Doviak and D. S. Zrnic. *Doppler Radar and Weather Observations.* Academic Press, 1993.

G. Galati, M. Naldi, and M. Ferri. Reconstruction of the spatial distribution of radar reflectivity of precipitation through linear-inversion techniques. *IEE Proc. Radar Sonar Navig.*, 143(6):375–382, 1996. doi: 10.1049/ip-rsn:19960722.

D. L. Hartmann, M. E. Ockert-Bell, and M. L. Michelsen. The effect of cloud type on Earth's energy balance: Global analysis. *J. Climate*, 5:1281–1304, 1992. doi: 10.1175/1520-0442(1992)005¡1281:TEOCTO¿2.0.CO;2.

P. H. Hildebrand and R. S. Sekhon. Objective determination of the noise level in Doppler spectra. *J. Appl. Meteor.*, 13(7), 1974. doi: 10.1175/1520-0450(1974)013¡0808:ODOTNL¿2.0.CO;2.

A. J. Illingworth, H. W. Barker, A. Beljaars, M. Ceccaldi, H. Chepfer, N. Clerbaux, J. Cole, J. Delanoë, C. Domenech, D. P. Donovan, S. Fukuda, M. Hirakata, R. J. Hogan, A. Huenerbein, P. Kollias, T. Kubota, T. Nakajima, T. Y. Nakajima, T. Nishizawa, Y. Ohno, H. Okamoto, R. Oki, K. Sato, M. Satoh, M. W. Shephard, A. Velázquez-Blázquez, U. Wandinger, T. Wehr, and G.-J. van Zadelhoff. The EarthCARE satellite: The next step forward in global measurements of clouds, aerosols, precipitation, and radiation. *Bull. Amer. Meteor. Soc.*, 96:1311–1332, 2015. doi: 10.1175/BAMS-D-12-00227.1.

P. Kollias, S. Tanelli, A. Battaglia, and A. Tatarevic. Evaluation of EarthCARE Cloud Profiling Radar Doppler velocity measurements in particle sedimentation regimes. *J. Atmos. Oceanic Technol.*, 31:366–386, 2014. doi: 10.1175/JTECH-D-11-00202.1.

D. C. Leon, Z. Wang, and D. Liu. Climatology of drizzle in marine boundary layer clouds based on 1 year of data from CloudSat and Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations (CALIPSO). *J. Geophys. Res.*, 113, 2008. doi: 10.1029/2008JD009835.

R. Marchand, G. G. Mace, T. Ackerman, and G. Stephens. Hydrometeor detection using CloudSat—an earth-orbiting 94-GHz cloud radar. *J. Atmos.Oceanic Technol.*, 25(4): 519–533, 2008. doi: 10.1175/2007JTECHA1006.1.

J. H. Mather and J. W. Voyles. The ARM climate research facility: A review of structure and capabilities. *Bull. Amer. Meteor. Soc.*, 94(3):377–392, 2013. doi: 10.1175/BAMS-D-11-00218.1.

D. L. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *J. ACM*, 9(1):84–97, 1962. doi: 10.1145/321105.321114.

J. P. Rémillard, P. Kollias, E. Luke, and R. Wood. Marine boundary layer cloud observations in the Azores. *J. Climate*, 25:7381–7398, 2012. doi: 10.1175/JCLI-D-11-00610.1.

N. A. J. Schutgens. Simulated Doppler radar observations of inhomogeneous clouds: Application to the EarthCARE space mission. *J. Atmos. Oceanic Technol.*, 25:26–42, 2008. doi: 10.1175/2007JTECHA956.1.

N. A. J. Schutgens and D. P. Donovan. Retrieval of atmospheric reflectivity profiles in case of long radar pulses. *Atmos. Res.*, 72:187–196, 2004. doi: 10.1016/j.atmosres.2004.03.014.

G. L. Stephens, D. G. Vane, S. Tanelli, E. Im, S. Durden, M. Rokey, D. Reinke, P. Partain, G. G. Mace, R. Austin, T. L'Ecuyer, J. Haynes, M. Lebsock, K. Suzuki, D. Waliser, D. Wu, J. Kay, A. Gettelman, Z. Wang, and R. Marchand. CloudSat mission: Performance and early science after the first year of operation. *J. Geophys. Res.*, 113, 2008. doi: 10.1029/2008JD009982.

B. Stevens, D. H. Lenschow, G. Vali, H. Gerber, A. Bandy, B. Blomquist, J.-L. Brenguier, C. S. Bretherton, F. Burnet, T. Campos, S. Chai, I. Faloona, D. Friesen, S. Haimov, K. Laursen, D. K. Lilly, S. M. Loehrer, S. P. Malinowski, B. Morley, M. D. Petters, D. C. Rogers, L. Russell, V. Savic-Jovcic, J. R. Snider, D. Straub, M. J. Szumowski, H. Takagi, D. C. Thornton, M. Tschudi, C. Twohy, M. Wetzel, and M. C. van Zanten. Dynamics and chemistry of marine stratocumulus—DYCOMS-II. *Bull. Amer. Meteor. Soc.*, 84:579–593, 2003. doi: 10.1175/BAMS-84-5-579.

G. M. Stokes and S. E. Schwartz. The Atmospheric Radiation Measurement (ARM) program: Programmatic background and design of the cloud and radiation test bed. *Bull. Amer. Meteor. Soc.*, 75(7):1201–1221, 1994. doi: 10.1175/1520-0477(1994)075¡1201:TARMPP¿2.0.CO;2.

O. O. Sy, S. Tanelli, P. Kollias, and Y. Ohno. Application of matched statistical filters for EarthCARE cloud Doppler products. *IEEE Trans. Geosci. Remote Sens.*, 52(11): 7297–7316, 2014. doi: 10.1109/TGRS.2014.2311031.

S. Tanelli, E. Im, S. L. Durden, L. Facheris, and D. Guili. The effects of nonuniform beam filling on vertical rainfall velocity measurements with a spaceborne Doppler radar. *J. Atmos. Oceanic Technol.*, 19:1019–1034, 2002. doi: 10.1175/1520-0426(2002)019¡1019:TEONBF¿2.0.CO;2.

S. Tanelli, S. L. Durden, E. Im, K. S. Pak, D. G. Reinke, P. Partain, J. M. Haynes, and R. T. Marchand. Cloudsat's Cloud Profiling Radar after two years in orbit: Performance, calibration, and processing. *IEEE Trans. Geosci. Remote Sens.*, 46:3560–3573, 2008. doi: 10.1109/TGRS.2008.2002030.

S. Twomey. *Introduction to the mathematics of inversion in remote sensing and indirect measurements*, volume 3. Elsevier, 1977.

T. Uttal and R. A. Kropfli. The effect of radar pulse length on cloud reflectivity statistics. *J. Atmos. Oceanic Technol.*, 18(6):947–961, 2001. doi: 10.1175/1520-0426(2001)018¡0947:TEORPL¿2.0.CO;2.

T. Uttal, L. I. Church, B. E. Martner, and J. S. Gibson. Cldstats: A cloud boundary detection algorithm for vertically pointing radar data. In *NOAA Tech. Memo. ERL WPL-233*, 1993.

J. Wang and B. Geerts. Identifying drizzle within marine stratus with W-band radar reflectivity. *J. Atmos. Res.*, 69:1–27, 2003. doi: 10.1016/j.atmosres.2003.08.001.

K. Widener and J. Mead. W-band ARM Cloud Radar – specifications and design. In *Proc. 14th ARM Science Team Meeting, Albuquerque, New Mexico, ARM*, 2004.

D. M. Winker, M. A. Vaughan, A. Omar, Y. Hu, K. A. Powell, Z. Liu, W. H. Hunt, and S. A. Young. Overview of the CALIPSO mission and CALIOP data processing algorithms. *J. Atmos. Oceanic Technol.*, 96:2310–2323, 2009. doi: 10.1175/2009JTECHA1281.1.

R. Wood. Stratocumulus clouds. *Mon. Weather Rev.*, 140(8):2373–2423, 2012. doi: 10.1175/MWR-D-11-00121.1.

R. Wood, M. Wyant, C. S. Bretherton, J. Rémillard, P. Kollias, J. Fletcher, J. Stemmler, S. de Szoeke, S. Yuter, M. Miller, D. Mechem, G. Tselioudis, J. C. Chiu, J. A. L. Mann, E. J. O'Connor, R. J. Hogan, X. Dong, M. Miller, V. Ghate, A. Jefferson, Q. Min, P. Minnis, R. Palikonda, B. Albrecht, E. Luke, C. Hannay, and Y. Lin. Clouds, aerosols, and precipitation in the marine boundary layer: An ARM mobile facility deployment. *Bull. Amer. Meteor. Soc.*, 96(3):419–440, 2015. doi: 10.1175/BAMS-D-13-00180.1.

X. Zhou, P. Kollias, and E. R. Lewis. Clouds, precipitation, and marine boundary layer structure during the MAGIC field campaign. *J. Climate*, 28(6):2420–2442, 2015. doi: 10.1175/JCLI-D-14-00320.1.