

An Investigation of Adaptive Remote Sensing Methods for Spaceborne Cloud Profiling

Radars

Dissertation

Presented in Partial Fulfillment of the Requirements for the Degree Doctor of Philosophy

in the Graduate School of The Ohio State University

By

Jakob DeLong

Graduate Program in Electrical and Computer Engineering

The Ohio State University

2022

Dissertation Committee

Dr. Joel T. Johnson, Advisor

Dr. Emre Ertin

Dr. Brian Slater

Dr. Fernando Teixeira

Copyrighted by
Jakob Alexander DeLong
2022

Abstract

Cloud profiling radars (CPRs) are designed to illuminate atmospheric clouds and estimate their water content and internal structure. Satellite based CPRs offer several advantages and have already been used to significantly improve scientific understanding of meteorology generally, but the physical nature of clouds makes them a challenging target to observe without large, expensive instruments. Adaptive remote sensing has the potential to enable smaller, cheaper systems to produce useful measurements alongside their larger predecessors.

A Matlab simulation has been developed to investigate the benefits of adaptive remote sensing in the CPR use case. This simulation has the ability to adapt the radar's pulse repetition frequency (PRF) in response to observed cloud column height as well as automatically identify and terminate low precision measurements to manage resources. The simulation has been tested using synthetic data derived from NASA's GEOS5 system and real data derived from CloudSat's operational history. The simulation has shown that adaptive remote sensing is capable of achieving resource savings and quality improvement in the CPR case which would make the use of highly resource constrained platforms for the CPR use case more feasible.

Dedication

To my family, whose support brought the completion of this work within reach.

Acknowledgments

I wish to thank the people whose guidance and advice was invaluable to completing this work. First, I wish to thank my advisors Dr. Graeme Smith and Dr. Joel Johnson. Dr. Smith provided the initial direction to start the project and Dr. Johnson provided guidance for most of my PhD and helped in the conclusion of this work. I would also like to thank my NASA Fellowship Technical Mentor, Dr. Rafael Rincon, for providing me with the opportunity to work with both hardware and software and build experience working with actual radar systems and data.

Vita

Winfield High School 2013
Bachelor of Science in Electrical Engineering, West Virginia University 2017

Publications

J. DeLong, M. A. Shattal, A. O'Brien, C. D. Ball, J. T. Johnson, and G. E. Smith, "Fully Adaptive Cloud Profiling Radar Simulation," in IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium, 2019, pp. 10087–10090.

J. DeLong and J. T. Johnson, "An Algorithm for Adaptive Determination of Radar Coherent Integration Time," Int. Geosci. Remote Sens. Symp., pp. 6579–6581, Sep. 2020.

J. DeLong, C. Ball, J. T. Johnson, and A. O'Brien, "USING The Simulation Toolset for Adaptive Remote Sensing (STARS) TO EXPLORE A TRADESPACE OF New observing strategies," in AGU Fall Meeting, 2020, vol. 2020.

J. DeLong and J. T. Johnson "Simulation Study of a Probabilistic Measurement Terminator for Cloud Radar Systems Using CloudSat Data," Submitted to IGARSS 2022, Accepted May 2022.

Fields of Study

Major Field: Electrical and Computer Engineering

Studies in:

Electromagnetics, Remote Sensing, and Microwaves
Communication and Signal Processing

Table of Contents

Abstract.....	ii
Dedication.....	iii
Acknowledgments.....	iv
Vita.....	v
Table of Contents.....	vii
List of Tables.....	ix
List of Figures.....	x
Chapter 1. Introduction.....	1
1.1 Introduction and Overview.....	1
1.2 Cloud Radar Challenges and Past Systems.....	2
1.2.1 Challenges.....	2
1.2.2 CloudSat.....	4
1.2.3 EarthCARE.....	5
1.2.4 Cloud Radar Products.....	5
1.3 Radar Background.....	7
1.4 CubeSats.....	15
1.5 Adaptive Radar.....	19
1.5 Conclusion.....	22
Chapter 2: Simulation Development.....	23
2.1 Simulation Toolset for Adaptive Remote Sensing (STARS).....	23
2.2 Input Data.....	25
2.3 PRF Adaptation.....	30
2.5 Measurement Terminator.....	34
2.6 Code Analysis.....	39
2.6.1 AdaptiveCPR/CloudSatSimulationScript.....	40

2.6.2 adaptiveCPR/CloudSatSimulator.....	42
2.6.3 dopplerGEOS5MANAGERadar.....	46
2.6.4 GEOS5RadarRangeProfile	48
2.7 Initial Results	49
2.8 Conclusions.....	56
Chapter 3: Simulation Results	58
3.1 Performance Predictions	58
3.2 Simulation Results Without PRF Adaptation	62
3.3 Fine Tuning of the Simulation	69
3.4 PRF Adaptation.....	75
3.5 Summary and Conclusions	88
Chapter 4: Summary and Conclusions.....	89
Appendix A: Derivation of Power Estimate Standard Deviation	92
Appendix B: Main Function Headers	95
B.1: adaptiveCPRSimulator.....	95
B.2 dopplerGEOS5MANAGERadar	96
B.3: GEOS5RadarRangeProfile.....	98
Bibliography	99

List of Tables

Table 1: The operating parameters used for all radar range equation calculations in this dissertation, taken from CloudSat.....	49
---	----

List of Figures

Figure 1: The perception-action cycle forming the basis of the FAR framework [56]	20
Figure 2: An example reflectivity field extracted from GEOS5.....	27
Figure 3: An example reflectivity field extracted from the CloudSat GEOPROF product taken from granule 3607 on the first day of 2007.....	29
Figure 4: The logical flow diagram of the simulation from meteorological data input to data output. Each logical component will be discussed in this section.	40
Figure 5: (top row) The measured and predicted cloud to ground height for the GEOS5 simulation, (second row) The predicted cloud height after taking a measurement, (third row) The PRF used for each measurement, (fourth row) The corresponding unambiguous range calculated from the chosen PRF[57]. The x-axis for all plots is the profile number.	50
Figure 6: (a) Mean fraction of integrations terminated early as a function of the truth SNR for SNR_min=-3 dB for Tterm values indicated in legend (b) Mean fraction of 200 pulses saved as a function of the truth SNR for SNR_min=-3 dB for indicated Tterm, (c) Mean fraction of errors as a function of the truth SNR, (d) SNR-integrated error versus SNR_min threshold.....	52
Figure 7: (a) Mean fraction of errors for T_term=0.87 as a function of the truth SNR and SNR_min (b) Mean fraction of pulses saved for T_term=0.87 as a function of the truth SNR and SNR_min.....	54
Figure 8: Test results applying the measurement terminator algorithm to the GEOS5 data as described in Section 2, representing 300 full measurements each integrated over 1.1 km.	56
Figure 9: A track of the expected reflectivity as a function of SNR for CloudSat.	59
Figure 10: The mean percent of measurements per orbit with an observed SNR below a specified SNR goal choice.	60
Figure 11: A CDF-scaled histogram of the highest mask value in each profile showing the fraction of each measurement type over the entire dataset.	61
Figure 12: The cumulative fraction of all measurements having a given peak reflectivity for the CloudSat data.	62
Figure 13: The error rates averaged per orbit for the CloudSat data	63
Figure 14: The correct decision rates averaged per orbit for the CloudSat data.....	64
Figure 15: The cumulative portion of terminated reflectivity values for the -8 dB SNRmin case.....	65
Figure 16: The cumulative portion of terminated reflectivity values for the -3 dB SNRmin case.....	66

Figure 17: The histogram of the number of pulses integrated in each measurement generated from the CloudSat data set.	67
Figure 18: The histogram of mask values for terminated measurements for the -3 dB SNRmin, 0.87 probability goal case.	68
Figure 19: The first profile from the CloudSat set selected for individual analysis converted to SNR with gaps where no reflectivity data was present in the profile.....	69
Figure 20: The second profile chosen for individual analysis with the strongest signal of -7 dB SNR occurring around 19 km.....	71
Figure 21: The portion of terminated measurements falling into each mask category for the -3 dB SNRmin, 0.99 probability goal case.	73
Figure 22: The portion of measurements terminated for each mask value for the -11 dB SNRmin, 0.5 probability goal case.	74
Figure 23: The portion of measurements with a mask value above 10 allowed to complete and terminated for the -11 dB SNRmin, 0.5 probability goal case.....	75
Figure 24: The error rates for the first PRF adaptation test. Note the 1% reduction relative to the -3 dB test without adaptation.	76
Figure 25: The correct decision rates for the first PRF adaptation test. The 1% error reduction was split approximately evenly between the two decision types.....	77
Figure 26: The mean error rates for the second PRF adaptation test, showing a further .5% decrease relative to the first test.....	78
Figure 27: The mean correct decision rates for the second PRF adaptation test, showing a 1% increase in correct completion rate.	79
Figure 28: The set of terminated mask sentinel values for the second PRF adaptation test, showing a small reduction in the portion of detections terminated.	80
Figure 29: The cumulative fraction of pulses integrated in each measurement for the second PRF adaptation simulation.....	81
Figure 30: The reflectivity field used for additional PRF variation testing, showing a strong central cloud structure with weaker edges.	82
Figure 31: The portion of each reflectivity value in the reflectivity field examined for PRF adaptation.....	83
Figure 32: The portion of each peak mask value observed in the reflectivity field used for additional PRF variation testing.....	83
Figure 33: The set of terminated mask values for the field in Figure 30 without PRF adaptation.....	85
Figure 34: The set of terminated mask values for the reflectivity field in Figure 30 with PRF adaptation and a static SNRmin goal.....	86
Figure 35: The set of terminated mask values for the field in figure with PRF adaptation and a varying SNRmin goal.....	87

Chapter 1. Introduction

1.1 Introduction and Overview

It is known that atmospheric clouds play a significant role in global climate and weather patterns. The most recent NASA decadal survey has indicated that observations of cloud characteristics from space provide significant benefit by providing improvements to short and long term weather prediction models as well as improved understanding of water cycle processes and impacts on related ecosystems [1]. Radar systems provide an ability to profile clouds thoroughly by making observations of their extent in the atmosphere as well as their internal structure, providing insight into the processes governing their development [2]. Satellite cloud profiling radars (CPRs), in particular, enable such profiling on a greater scale and in more challenging environments [3], [4], including enabling storm observation at no risk to the observing platform. Radar observations of clouds also enable better understanding of the convective processes that govern their formation and are in turn influenced by their presence, creating a complex feedback loop that is still not fully understood [5]. These benefits make pursuing satellite cloud profiling radar development highly attractive. Recent proof of concept projects have also shown that highly resource constrained satellite platforms are capable of performing atmospheric remote sensing tasks [6], [7] while adaptive remote sensing provides potential ways of improving performance in resource constrained systems [8]. This dissertation seeks to investigate the potential

benefit of adaptive remote sensing for the CPR problem with an emphasis on resource savings to enable the use of resource constrained platforms.

This chapter discusses background information and literature necessary for understanding the current work and placing it in context. Section 2 discusses spaceborne cloud radar and associated topics, including a discussion of prominent example systems, challenges with spaceborne radar observation of clouds, and techniques used for cloud detection. Section 3 discusses radar system basics and how radar measurements are generated. Section 4 discusses CubeSats in the context of satellite development, remote sensing using CubeSats, and the particular challenges associated with the same. Section 5 discusses adaptive radar and define the concept as well as providing examples and discussing the implementation used as the basis for adaptation in this work.

1.2 Cloud Radar Challenges and Past Systems

1.2.1 Challenges

The first challenge facing satellite cloud radar consists in the long ranges of orbital platforms combined with the relatively low return strength received from the illuminated targets. Platform orbits range from around 700 km [9] down to about 400 km [10], [11], which produce significant power reduction given the R^2 power dependence on range arising from the distributed target form of the radar range equation [12]. In addition to this, the strength of the radar return varies significantly over the possible set of illuminated targets. Values can range over more than 50 dB for strong precipitating clouds to thin clouds [13]–[15], of which all are potentially of interest to a given cloud profiling system. In addition to this, significant attenuation due to atmospheric absorption and even

absorption by the intended cloud targets [14] have an additional degrading effect on the returns from clouds. These attenuation effects necessitate the use of particular operating frequencies due to “windows” in the attenuation profiles; 94 GHz offers one such window where attenuation is minimized while responses from the intended meteorological targets remain detectable [13].

An additional problem for spaceborne cloud profiling radars seeking to measure velocity is scene decorrelation. Clouds are dynamic systems whose constituent hydrometeors are constantly shuffling and rearranging, presenting a substantially different scene to an observing radar over a relatively short period of time [13]. Several so-called “spectral broadening” effects exist and their net effect is to impart a Gaussian correlation function onto estimates of reflectivity derived by illuminating cloud targets [13], [16], [17]. Even in clear weather for satellite borne radars, the observed spectral broadening can be significant because the satellite bearing the cloud profiling radar is moving at around 7 km/s [17]. Because of this effect alone, scene decorrelation can occur on the order of $95\mu\text{s}$, meaning that pulses must be emitted at a rate of 10 kHz for velocity measurements to be accurate and additional weather impacts will only worsen this constraint [17]. For an incoherent system, however, like the one modelled in this project, this presents an advantage as it means that samples can be collected up to a rapid rate and remain statistically independent for the purposes of power estimation.

A final concern for CPRs is the return from the ground. Because atmospheric clouds are relatively weak scatterers, the ground return presents a significantly higher apparent reflectivity than any illuminated cloud and typically dominates the received signal in the

range bins at and near the ground surface, up to potentially 2-3 km [6], [14], [15]. This also contributes to the demand for short range resolutions for cloud profiling, as an excessively long pulse with significant sidelobes will increase the range over which the ground return contaminates useful data. Existing systems approach this problem differently, though a common approach is to incorporate additional data sources or a digital elevation map and mark the ground's position in the return [2], [18].

1.2.2 CloudSat

CloudSat is a 94 GHz incoherent cloud profiling radar flown aboard a satellite orbiting at approximately 705 km [9]. Launched in 2006, CloudSat was the first satellite-borne cloud profiling radar and has since produced a great amount of useful scientific data [19]. It was designed to be able to detect clouds down to about -30 dBZ [20], which should have encompassed between 60% and 90% of all clouds of interest [21]. CloudSat directly measures power received from hydrometeors and produces an estimate of the resulting effective reflectivity, from which many other meteorological data products can be calculated [20]. CloudSat was also originally part of a constellation with other instruments including the CALIPSO satellite, allowing measurements from the two instruments to be co-located and combined to provide more insight into cloud structures [20]. CloudSat has been used for various purposes ranging from characterizing global cloud distribution [22], to cross validating other satellite borne instruments like MODIS [23], to probing cloud structures in various regions across the globe [24]–[28] and validating atmospheric modelling tools [29]. CloudSat has been validated using both ground and aircraft-based measurements and been shown to be acceptably accurate to within about 0.5 dB [30]. Data

from CloudSat is archived on the CloudSat Data Processing Center website [31] and all calculated data products can be found there for use by the community. Owing to its nature as the first CPR on orbit and the wide variety of applications for its produced data, CloudSat was used as the model for this project. Its operating parameters were adopted as the base parameters for the simulation as well as its waveform design and signal processing chain.

1.2.3 EarthCARE

The EarthCARE satellite was initially designed as a follow-on to the A-Train [32]. The EarthCARE satellite is intended to carry multiple instruments but among them is a cloud profiling radar also operating at 94 GHz with the capability to measure both cloud reflectivity and velocity, and it would be the first such radar with velocity measurement capability [10]. Its Doppler capability combined with a lower sensitivity threshold of about -37 dBZ, achieved by orbiting at just 393 km instead of CloudSat's 705 km and having a larger antenna, have the potential to present significant new sources of information on cloud structure that could not be captured using information from CloudSat alone [10]. Due to the difficulties associated with Doppler processing from space, however, this approach is not considered in the current project. Development on the EarthCARE satellite is ongoing and the satellite is expected to launch in 2022 [32]–[34].

1.2.4 Cloud Radar Products

Once cloud reflectivity data is collected it is typically downloaded from the satellite to a ground station and processed offline. The satellite thus measures over some fixed portion of each orbit or over a pre-prescribed number of orbits without attention paid to where clouds are in the field and the determination of where the clouds are is made after the fact.

In the case of CloudSat, this is done using the GEOPROF algorithm first proposed in [35] and modified in [18] that creates a “cloud mask” that assigns a number from 0 to 40 to an individual cell representing a portion of the returned reflectivity field which indicates the likelihood that the cell contains returns from clouds, with increasing numbers corresponding to increasing likelihoods of cloud presence and reduced rates of false detection. The algorithm works by estimating the SNR after subtracting off an estimate of the noise power and taking the ratio to the same in addition to considering continuity of high power (considering that the detectable clouds can be as weak as -30 dBZ) cells both within an atmospheric column and across columns [18].

This approach represents what can be done with just one system, however, and other approaches exist that incorporate data streams from multiple systems. For instance, the co-location of measurements from CloudSat and CALIPSO have enabled the generation of combined cloud masks that take input from both systems into account [22] by marking the presence of a cloud if present on one or both of the masks generated by the two systems. This approach has also been used to validate the MODIS cloud mask and conclude that MODIS does not detect about 9.1% of clouds globally that CloudSat and CALIPSO can [23]. For other instruments that use multispectral (visual, IR, radiometer) measurement systems, cloud detection is often implemented by using a battery of threshold tests on the various spectral channels and determining a hit if one or a subset of the tests agree on the presence of a cloud [36], [37]. Finally, the use of convolutional neural networks has also been proposed as an option for offline processing of existing visual spectrum images of potentially cloudy scenes [38]. These approaches have varying levels of complexity but

most have in common the use of data streams from multiple instruments and all are applied to the entire set of data after it has been collected rather than being applied to the data on a column by column basis as it is generated.

1.3 Radar Background

The information in this section is derived from [12] unless otherwise stated. “Radar” is an acronym expanding to “radio detection and ranging.” This is indicative of both the mode of function and primary use of most radar systems; namely, radar systems emit energy in the form of radio waves that can then be used to measure the range to an object. The radar uses a means to excite an antenna which then radiates energy in a particular direction. The energy can then interact with objects in its environment, experiencing refraction, reflection, and absorption depending upon the nature of said environment.

Radar systems can be divided into two large categories based on the energy they emit: continuous and pulsed. Continuous systems are designed to constantly transmit an unbroken signal. Such systems are often used for velocity measurement or in operational modes where the antenna is mechanically or electronically “scanned” through a scene and returns are similarly captured continuously. Pulsed systems, on the other hand, transmit intermittently producing discrete pulses that travel through the medium and listening for returns when the transmitter is off. The system modelled in this work is of the pulsed variety. Pulsed systems present a more convenient way to measure the range to target by taking note of the time at which echoes are received. Specifically, for a target at a fixed distance R from the radar, the time that an echo will be received by the system can be found by [12]

$$t = \frac{2R}{c} \quad (1.1)$$

where c is the speed of light in a vacuum. The factor of 2 in (1.1) arises because the pulse transmitted from the radar first travels to the target a distance of R away and then is reflected by the target, causing it to travel the same distance back to the radar receiver. In addition to this consideration, the radar pulse itself is of a finite duration and therefore a minimum resolvable distance in range can be calculated. For pulsed radars, this quantity, referred to as the range resolution, can be calculated as [12]

$$\Delta R = \frac{c}{2B} \quad (1.2)$$

Here B is the bandwidth of the transmitted pulse, c is again the speed of light, and ΔR is the resulting range resolution which represents the minimum separation between reflective targets that the radar can resolve. Because it is a function of the pulse bandwidth, range resolution is therefore a tunable parameter depending upon the pulse chosen for the system. For meteorological radar, monochromatic or square pulses are a common choice and the bandwidth of such a pulse is well approximated by the reciprocal of the pulse length.

A final parameter of interest for a pulsed radar is the unambiguous range interval, which places a limit on the greatest distance away from the radar that can be measured. This parameter arises as a consequence of the pulsed nature of the system; since range is measured as a function of the time at which an echo returns from a potential target, distance can only be measured accurately if enough time is permitted between pulses for all echoes

from previous pulses are allowed to return before subsequent echoes are emitted. The relationship between pulse repetition interval (PRI) or the pulse repetition frequency (PRF), which are measures of how frequently pulses are emitted by the radar, and the maximum unambiguous range is given as [12]

$$R_u = \frac{cT_p}{2} = \frac{c}{2f_p} \quad (1.3)$$

In Equation 3, T_p is the PRI and f_p is its inverse the PRF, both of which are measures of the rate at which pulses are emitted by the radar. If targets are at a distance from the radar less than R_u , (1.3) indicates that their range can be accurately extracted using (1.1). If a target is at a distance from the radar greater than R_u , then echoes from one pulse will arrive during a different pulses T_p and all echoes will become ambiguous since it is no longer possible to determine their absolute range from the radar. Equation (1.3) also provides a means by which sensor adaptation can improve signal quality subject to constraints. As an example, both CloudSat [19] and EarthCARE [10] permit changes in their PRFs and the number of pulses integrated per processing interval based on the altitude of their orbits in order to keep the range window within 20-30 km of the ground surface. Since radar systems are typically at least 400 km away from the scene being studied, (1.3) would appear to place a limit on the PRF of just 375 Hz. In practice, however, systems tend to use PRFs on the order of 3 kHz to 7 kHz [9], [10] in order to collect enough pulses during the measurement period on the order of tenths of a second to reduce measurement variances to acceptable levels. This is permissible because the range of interest for most weather radars ranges from the surface of the earth up to about 20-30 km [9], [10]. Since

it is reasonable to assume that no illuminated targets will exist between the radar's orbit path and the illuminated region of the atmosphere, the radar operators can simply perform range calculations using the unambiguous region covering the atmospheric column. The two parameters are not continuously varied, however, and are set to one of a set of pre-determined values based on the constraints of the orbit.

In order to extract the distance to a target using received echoes, the received echoes must be processed. This is commonly done using a "matched filter" since it has been shown that doing so produces the highest possible signal to noise ratio (SNR) at the output of the radar signal processor. In modern systems, this operation is implemented as a convolution operation with a particular waveform, such that the output of the radar detector is given as [12]

$$y(t) = x(t) * h^*(\tau - t) \quad (1.4)$$

The filter is said to be "matched" because the waveform used for $h(t)$ is a copy of the transmitted waveform while $x(t)$ is the signal containing the echoes reflected from the targets. If the transmitted waveform is complex valued, then the conjugate in (1.4) is an identity. The output of the matched filter, $y(t)$ will have a peak when echoes are present at the time echoes are received according to (1.1), thus allowing range to target to be measured. Finally, with pulsed radars, measurements generally consist of the returns from several pulses combined over time to improve the quality of the received signal. This process is referred to as integration and can be either coherent or incoherent. Coherent integration is called such because it considers both magnitude and phase information from

the received signals and therefore permits the measurement of velocity. Incoherent integration is called such because it only takes magnitude information into account from the received signals and thus is not able to directly measure velocity. The system being modelled for this project is an incoherent system, and the integration takes the form of taking the average of the received power from the echoes. This averaging has the effect of reducing the standard deviation of the received power estimate (as discussed in Chapter 2), thus averaging more pulses up to the maximum number allowable by both the time over which the target is illuminated and the resource constraints of the system will improve the quality of the received measurement.

Calculating the power returned by an illuminated scatterer can be done using the radar range equation which has the following form for a point target [12]

$$P_r = \frac{P_t G^2 \lambda^2 \sigma}{(4\pi)^3 R^4} \quad (1.5)$$

where P_t is the peak transmit power, G is the gain of the radar's antenna, λ is the radar's operating frequency, σ is the radar cross section of the illuminated target which governs how much energy the target will return, and R is the range to the target. Atmospheric targets like clouds, however, are not simply point targets but rather consist of many such scatterers existing with each other in a volume. For this case, the illuminated resolution volume can be calculated from the radar's range resolution and antenna parameters according to [12]

$$V = \pi \left(\frac{R\theta_3}{2} \right) \left(\frac{R\phi_3}{2} \right) \left(\frac{c\tau}{2} \right) \quad (1.6)$$

where θ_3 refers to the 3 dB beamwidth of the antenna in azimuth, ϕ_3 refers to the same quantity in elevation, and τ is the pulse length of an assumed square pulse, so that the last term in (1.6) is the range resolution of a radar using a square pulse. The beamwidth parameters refer to the angle over which the gain of the antenna's radiation pattern drops by 3 dB or 50%. In order to determine the illuminated cloud's radar cross section, its "reflectivity" can be used [13], [14]

$$\eta = \frac{10^{-10}\pi^5}{\lambda^4} |K|^2 Z \quad (1.7)$$

where Z is reflectivity in $\frac{mm^6}{m^3}$, λ is the wavelength of the radar in centimeters, $|K|$ is the magnitude of the complex dielectric constant of water, and the use of η instead of σ denotes that (1.7) computes radar cross section per unit volume rather than simple radar cross section. The choice of units for Z arises from the use of the Rayleigh approximation in which the backscattered radar cross section is assumed to be proportional to the sum of the illuminated particle diameters in a resolution volume, with the conversion to millimeters producing a range of values more convenient for use in meteorological studies [14]. In cases where the particle size approaches the wavelength of the radar, scattering can no longer be well approximated by the Rayleigh assumption and must be calculated using Mie theory [14]. In the case of cloud radar, however, it is common to report measured reflectivity using (1.7) as if the target was a Rayleigh target, thus giving rise to an "effective

reflectivity” commonly denoted Z_{eq} [14]. In practice, the Rayleigh assumption for cloud radar has been shown to be acceptable for all cloud targets up to 35 GHz and even for fair weather clouds up to 300 GHz and the use of the Z_{eq} parameter enables the use of Rayleigh reflectivity to aid calculations [14]. Reflectivity calculations in this dissertation are done using the effective reflectivity parameter and the Rayleigh approximation.

Taken together, the resulting received power can be calculated by multiplying (1.6) by (1.7) and substituting the result for σ in (1.5). The result is [12], [13], [15]

$$P_r = \frac{P_t G^2 \lambda^2 \eta \theta_3 \phi_3 c \tau}{512 \pi^2 R^2} \quad (1.8)$$

Thus it can be seen that radar return power for a meteorological target has an R^2 dependence instead of an R^4 dependence due to the distributed nature of the illuminated target. In practice, the power received from the target must compete at the receiver with power generated by random thermal noise. Noise power can be calculated according to

$$P_N = k T_o B F \quad (1.9)$$

where k is Boltzmann’s constant, T_o is the standard temperature of 290 K, B is the receiver’s bandwidth (given by the inverse of the pulse length in the case of a square pulse), and F is the system’s noise figure which captures the degree of additional noise corruption generated by the radar’s internal hardware. The signal to noise ratio (SNR) can then be calculated by taking the ratio of (1.8) and (1.9) and is an important figure of merit for radar systems.

The signal received by a radar can be modelled as a sequence of complex-valued voltage samples with a form given by [12]

$$V = V_R + jV_I \quad (1.10)$$

where V_R is the real component of the voltage and V_I is the imaginary component of the voltage. The complex representation is a calculation aid achieved by oversampling the received signal and splitting the resulting data stream into two streams, with the imaginary component having a 90 degree phase shift relative to the real component. In the case of returns from meteorological targets like clouds, the components of the voltage can be assumed to be independent, zero-mean Gaussian random variables [39] and we can write

$$V = \sqrt{\frac{v^2}{2}} (V'_R + iV'_I) \quad (1.11)$$

where the variance $\frac{v^2}{2}$ has been factored out of the real and imaginary components to reduce them to standard normal random variables. The variance in (1.11) is equal to the expected value of the received power and it can be written [39]

$$v^2 = v_s^2 + v_n^2 = v_n^2(1 + SNR) \quad (1.12)$$

where v_s^2 represents the expected power returned by the illuminated hydrometeors according to (1.8) and v_n^2 represents the thermal noise power generated by the receiver electronics according to (1.9). SNR is given as $\frac{v_s^2}{v_n^2}$ and can be calculated with an estimate of the received signal power and the noise floor, which may either be estimated or known

a-priori. The returned signal power can be estimated using an incoherent radar by taking the average of the magnitude of the received voltage samples, according to

$$P \cong \frac{1}{N} \sum_{k=1}^N |V_k|^2 \quad (1.13)$$

which will produce estimates of the signal power. Because V will still contain a contribution from the noise as well as the intended signal according to (1.11) and (1.12), it is common in practice to subtract an estimate of ν , or the true value if known, from estimates produced by (1.13) before calculating SNR. In what follows, it is assumed that ν_n^2 is known to very high precision so that any associated errors can be neglected. Equation (1.13) effectively describes the process of incoherent integration.

1.4 CubeSats

While the data provided by CloudSat has been unambiguously valuable, it is not without shortcomings. The biggest single problem is that the revisit time for a given region is fully dependent upon the orbit of the satellite, which means that an individual scene in the atmosphere will not be illuminated again until the satellite has passed over other points on the earth and the cycle repeats. In the case of CloudSat, this occurs once every 16 days [19], meaning that the global data stream produced by CloudSat is only providing a single snapshot of atmospheric conditions just over once every two weeks. It is known that meteorological systems evolve over much shorter time periods, however, and that observations with a far higher temporal resolution would be highly valuable [3]. Constellations of satellites would provide increased revisit time within the window

mandated by the satellite orbit, like the A-Train constellation [20], but producing multiple systems like CloudSat and EarthCARE that are capable of cloud profiling measurements is a long and expensive proposition. Recent developments in satellite design have enabled the use of smaller, lower-cost CubeSats for science missions [40] which could provide multiple observations of a given atmospheric system during a single orbit.

CubeSats are small, modular satellite platforms designed around a 10 cm cubic “unit” and consisting of multiples of that unit combined together to fit the various components [41]. This design concept enables a standardized approach and provides constraints for the miniaturization of hardware components, and many existing and proposed CubeSat projects have made extensive use of commercially available hardware in addition to custom-built components, further enabling cost reduction [41]. CubeSats have so far been used in multiple contexts, including testing of RF interference mitigation techniques [42], radiometer observations of both total atmospheric energy radiance [43] and ice cloud physical characteristics [44], and, of greatest importance for this work, radar observations of precipitation [11], [45]. These applications varied in their design, ranging from 3U (indicating the use of 3 of the 10cm cubic units) for the radiometers [43], [44] and 6U for the interference testing and precipitation radar [42], [45].

RainCube currently represents the state of the art for active precipitation measurements on a CubeSat and so merits further discussion. It is optimized for detection of precipitation rather than clouds and so has a minimum detectable reflectivity factor of 12 dBZ and uses pulse compression in order to achieve a 250 meter range resolution with lower peak power than CloudSat (about 10W rather than 1.7kW) while still rejecting ground clutter [45].

Pulses are transmitted at a nominal rate of about 600 Hz and each pulse is $166 \mu\text{s}$ long with a linear chirp bandwidth of 2.5 MHz with windowing applied for sidelobe rejection that results in a natural range resolution of about 120 meters which can also be averaged down to 240 meters to improve sensitivity to targets in the atmospheric column [6]. Returns from RainCube are processed incoherently every .28 seconds, corresponding to a rough along-track footprint of 2 km which is additionally degraded by the antenna's 8 km circular footprint, though deconvolution techniques are applied in offline processing to reduce the effective along-track resolution to the level of the along track integration length [46]. In addition, the radar transmits for the entirety of one orbit before using the next three orbits to downlink data and allow the batteries to recharge [45]. RainCube has proven that the CubeSat form factor can be of use for the atmospheric radar problem [3], [41].

Operating an atmospheric radar in the CubeSat form factor is a significant challenge. Because atmospheric hydrometeor reflectivity is inversely proportional to the wavelength, higher frequencies are better but such high frequencies as those suitable for cloud profiling, namely W-band, exist at the high end of current solid state amplifier technology [7]. Additionally, as seen with RainCube, producing an antenna with sufficient gain and footprint to enable atmospheric sensing is another significant challenge. RainCube's footprint of 8 km had a deleterious effect on along-track resolution as compared to CloudSat's 1.1 km resolution that necessitated purpose-built offline processing to accommodate. Larger antennas would also have to be able to be stored within the confines of the CubeSat platform and successfully deploy once the platform is on orbit [47]. RainCube's 0.5 m deployable antenna was one of the first proposed designs proven to

succeed on orbit and represented a significant step forward [6]. Work has also been done to investigate deployable antennas of various kinds for earth science and deep space applications [48], [49] with deployable mesh antennas standing out for their effectiveness at Ka-band being able to provide around 47 dB gain [50]. Another challenge is presented by the power systems able to fit within the constraints of the CubeSat platform. RainCube operates on a 25% duty cycle which results in it effectively measuring for one out of every four orbits [51] as compared to a system like CloudSat which provides a “continuous” sequence of measurements every orbit [19]. This implies that, even in the best case, a CPR using current CubeSat technology would only be able to capture 25% of the clouds that CloudSat captures. Thus, while RainCube proves that the CubeSat format can be of use for atmospheric remote sensing, additional development is needed to achieve better performance under the resource constraints.

A full Doppler-capable radar architecture has been proposed and tested using a 95 GHz system that would be suitable for cloud profiling, but it was designed around the problem of detecting emission from comets in space and uses a small 15 cm diameter horn-fed reflector owing to reduced range to target relative to its task of observing other spaceborne objects [7]. The system uses frequency-modulated continuous waveform (FMCW) operation as a means to maintain amplifier efficiency rather than pulsed waveform operation, but is otherwise comparable in nature to traditional atmospheric radar systems like CloudSat and RainCube and could even be used as such [7]. Initial testing of the system was done using terrestrial atmospheric targets and indicated that the radar as designed was capable of detecting precipitation and clouds out to a range of about 4 km

[7]. The radar system used an average transmit power of 0.58 W, a pulse repetition interval of 34 μ s to achieve Doppler sensitivity, and had an antenna gain of about 41 dB [7]. Work is still being done to pack the components tightly, though the existing prototype has a mass of 6 kg [7], which is comparable to RainCube's mass of 5.5 kg for a 6U system [6], implying that a 6U W-band cloud radar is feasible given existing technology though a new antenna would be necessary to support such a project.

1.5 Adaptive Radar

Work has been done by Guerci [52] and Haykin [53] to move radar closer to what has been termed “cognitive” or “fully adaptive” operation. For both authors, the model for radar adaptation is the decision making process presumed to be employed by cognitive creatures, be they human beings or other animals [52], [53]. In Haykin's work, the fundamental structure guiding adaptation is the “perception-action cycle” which models animal cognition as a three stage cyclic process [53]. Under this model, a cognitive animal makes an observation of their environment, draws a conclusion about the external world based on that observation, and then takes an action based upon the conclusion before repeating the same cycle [53]. For the purposes of generalized adaptive radar, Haykin's concept was formalized into the Fully Adaptive Radar or FAR framework by researchers at The Ohio State University [54], [55]. To illustrate the concept, a perception-action cycle is shown in Figure 1 [56].

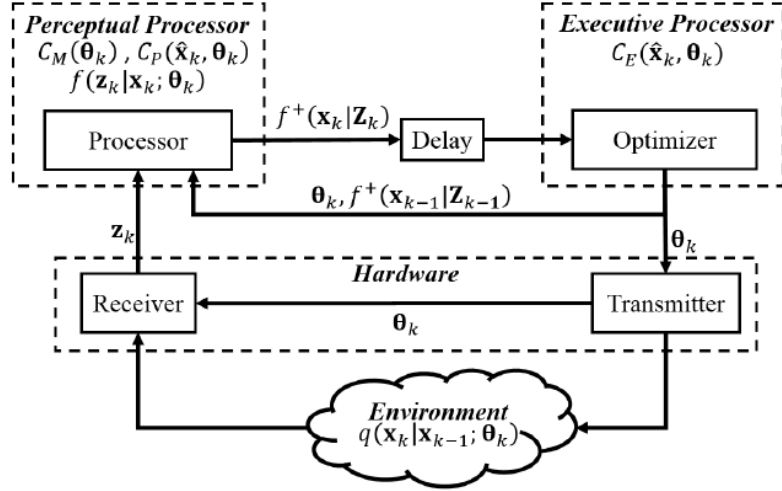


Figure 1: The perception-action cycle forming the basis of the FAR framework [56]

Under the FAR framework, the “perceptual processor” is the logical element responsible for drawing a conclusion based upon measurement data z_k collected by the radar and the “executive processor” is responsible for optimizing system performance by selecting measurement parameters that maximize measurement quality according to a cost function specified by the system designer [56]. The nature of the perception generation and the cost function depend upon the particulars of the radar task being performed. The perception is generated by assuming that the environment evolves according to a prescribed model $q(x_k|x_{k-1}; \theta_k)$, which is a function of the current state x_k , the immediately prior state x_{k-1} , and the parameters used to observe the state in θ_k (like PRF, transmit power, and similar), and that a likelihood function $f(z_k|x_k; \theta_k)$ can be described that can be used to determine the most likely observed state \hat{x} based on a given measurement [57]. Typically, the perception derived from the current set of measurement data is also propagated forward in time according to a prescribed model so that the system can adapt

to the target scene it expects to encounter based on the available data rather than simply reacting to past observations. The forward propagation is often implemented by a Kalman filter [57]. The perceptual processor also defines two cost functions: the measurement cost function $C_m(\theta_k)$ which relates measurement parameters θ_k to resource usage or another user defined “cost”, and the performance cost function $C_p(\hat{x}_k, \theta_k)$ which quantifies the “value” of a given measurement according to user-defined goals using the chosen set of measurement parameters θ_k and the state estimate \hat{x} calculated using the likelihood function $f(\cdot)$ and the state evolution model $q(\cdot)$ [57]. The executive processor chooses the set of measurement parameters given the state estimate \hat{x} using an executive cost function $C_E(\hat{x}_k, \theta_k)$ which may be the sum of the performance and measurement cost functions or may be related in another way to the same with the parameters θ_k being chosen such that the executive cost is minimized [57].

In the case of target tracking, parameters like transmit bandwidth, PRF, and pulse length can be chosen such that track quality is maintained as an illuminated target moves through a scene [54], [55], [58]. Other applications including synthetic aperture radar and low probability of intercept systems [56] have also been investigated, with the common theme being using adaptation to allow the radar to remain effective in challenging environments with possibly adversarial targets [59], [60]. The satellite CPR task presents a challenge for adaptive radar, however, as the set of valid measurement parameters is highly restricted by the relatively weak nature of cloud radar returns when observed from space as mentioned above. Whereas prior work with the FAR was able to consider several different sensing parameters and balance them against each other, the sensor adaptation available to the

satellite CPR is more limited [61], taking into account only the PRF since the other parameters must be set to specific values in order to maintain the ability to detect cloud returns. The FAR framework or a simplified version thereof can nevertheless be used to control the PRF adaptation for the satellite CPR problem in order to improve measurement quality, as will be further described in the next chapter.

1.5 Conclusion

This chapter presented motivation and background information on CPR systems as well as the CubeSat design paradigm. CPRs aid in weather forecasting and climate modelling, and satellite borne CPRs are able to provide global coverage without risk to the observing system. Additionally, sensor parameter adaptation was discussed as a means of improving the quality of measurements obtained by CPRs. Specifically, adapting the PRF to the observed cloud scene can result in improved signal quality assuming that the radar's onboard resources and signal processing chain can accommodate the increase in data. The goals of this dissertation are the following. First, develop a CPR simulator with sufficient fidelity to be able to model current systems like CloudSat. Second, use the developed simulator to investigate the effects of sensor parameter adaptation on the CPR use case with particular interest in effects on measurement quality and resource usage. Third, determine the implications of the previous investigation for the design of CPR systems on CubeSat platforms. The next chapter discusses the design and development of the CPR simulator, and Chapter 3 discusses the results of testing parameter adaptation and implications for CubeSat CPR design.

Chapter 2: Simulation Development

This chapter discusses the development of the adaptive cloud profiling radar simulation used to examine the potential benefits of sensor parameter adaptation in cloud remote sensing. The simulation's origins as a component of the STARS project are first discussed, followed by a description of the input atmospheric data sources used to allow the simulation to run. Next, the first adaptive component of the simulation, PRF adaptation, is discussed. Finally, the second adaptive component, the measurement terminator, is discussed. Following this, an analysis of the code comprising the simulation is provided. This simulation forms the basis for studies presented in the rest of the dissertation.

2.1 Simulation Toolset for Adaptive Remote Sensing (STARS)

The very first version of the adaptive radar simulation was built in the context of the STARS project [61]. STARS sought to provide a reconfigurable, reusable coding tool that could be used to model an arbitrary remote sensing system, be it a radar, radiometer, or other type of system. In order to achieve this, STARS was designed with three primary component libraries: ADAPT, MANAGE, and COLLABORATE. ADAPT is the most relevant of the three components, and it was determined that the core of ADAPT should be structured around a test case to allow the necessary abstract sensor components to be built up. To that end, it was determined that an adaptive satellite-borne cloud profiling radar would serve as an interesting test case, and the simulation was designed accordingly.

Though the goal of the STARS project was to produce open-source libraries using C++ and Python scripting, initial prototyping was done in Matlab owing to its convenient support of matrix math and signal processing techniques like convolution and Fourier transforms that form the basis of much of radar signal processing.

The initial Matlab prototype used an arbitrary user-generated input data set and produced output radar returns by convolving the chosen transmit pulse with the target scene before performing matched filtering on the resulting echo signal. Due to the cloud radar test case, targets in the input scene were considered to be distributed targets rather than point targets, so returns were scaled according to the volume RRE given by (1.9) instead of the point target RRE [12]. This allowed the prototype to better model returns from actual cloud scenes and provided a convenient transition to more realistic input data sets later. It was also initially intended that the prototype would include velocity measurements as well as range measurements, and so the tool had the ability to use either coherent processing via Fourier transform or incoherent processing using simple averaging, as described in Chapter 1. The results in what follows do not consider velocity measurement for several reasons which will be discussed later.

Sensor adaptation in the prototype was achieved using a simplified version of the FAR framework [54] which sought to increase the PRF of the observing radar subject to the unambiguous range constraint provided by the placement of clouds in the observed scene, as described in Chapter 1. Further details of the PRF adaptation will be discussed in Section 2.3. In addition to the basis for the radar simulation coming from ADAPT, the basis of the measurement terminator, to be discussed later, came from the MANAGE

library. Working off of some of the STARS team's experience with CubeRRT [42], a simulation was developed that used a probabilistic model to identify low quality measurements and terminate them. The details of the associated measurement terminator will be discussed in Section 2.4. The contributions of the author consisted of developing the adaptive radar simulation and writing the cloud radar case study used to test the simulation.

2.2 Input Data

In order to make both STARS and the adaptive atmospheric radar project as useful as possible as a design tool, a source of realistic input data was necessary. It was known that cloud reflectivity and turbulence could be modelled as arbitrary Gaussian spectra [13], [39], but having an external description of cloud properties for observation was deemed a priority over arbitrarily creating scenes for the simulation to observe. Based on discussions with STARS team members, it was determined that NASA's GEOS5 simulation could provide a useful data source [62]. NASA's Goddard Earth Observing System, in its fifth version as of 2014, has a long heritage of providing input data for climate models and atmospheric simulations. GEOS5 provided data fields for atmospheric water content per unit air mass and air density, arranged on a regular grid over the entire surface of the earth in its Ganymed release. The Ganymed release was calculated to a 7km horizontal resolution and a variable vertical resolution dependent upon pressure with 72 different levels, though the data was placed onto a .0625 degree latitude longitude grid, meaning that the base level data from the simulation output was already interpolated into a finer grid for data access.

In order to make use of GEOS5, the data were converted into a format more closely resembling radar-observable cloud water content. First, the vertical pressure levels were converted into height above sea level using the hypsometric formula [63] and assuming that the instantaneous temperature field in GEOS5 could serve as a reasonable approximation of the “mean virtual temperature” for the scenes of interest. The formula is given as

$$h = z_2 - z_1 = \frac{R_d \bar{T}_v}{g} \ln \left(\frac{p_1}{p_2} \right) \quad (2.1)$$

where h is the thickness of the calculated layer, R_d is the gas constant for dry air of $287 \frac{J}{Kkg}$, \bar{T}_v is the mean virtual temperature of the layer taken from the GEOS5 temperature field, p_1 is the pressure at the base of the atmospheric column (assumed to be mean sea level pressure), p_2 is the pressure at the altitude levels specified by GEOS5, and g is the acceleration due to gravity. Given that the goal was simply to establish a reference height field for the dataset, this was seen as an acceptable approximation. Next, fields for air density and fractional liquid cloud condensate from GEOS5 were captured and multiplied together to produce a field of airborne liquid water content. This allowed the use of a series of approximations first stated in [64] and restated in [14] to convert directly from airborne liquid water content to volumetric radar cross section.

$$\eta = \frac{1.47}{\lambda_o^4} 10^{-9} |K|^2 m_w \quad (2.2)$$

Here, λ_o is the radar wavelength in centimeters, $|K|^2$ is the squared magnitude of the complex dielectric constant of water assumed to be 0.75 based on figures used for CloudSat

data processing [65], m_w is the cloud liquid water content in $\frac{g}{m^3}$, and η is the resulting radar cross section per unit volume. The resulting data could then be used to generate a simulated radar observation. Figure 2 below shows an example scene representing an atmospheric scene simulated for 05/15/2005 at 2130z over the range from 91.5°E to 98.375°E at -46.875°N, converted to reflectivity factor dBZ as described in Chapter 1, calculated using this method.

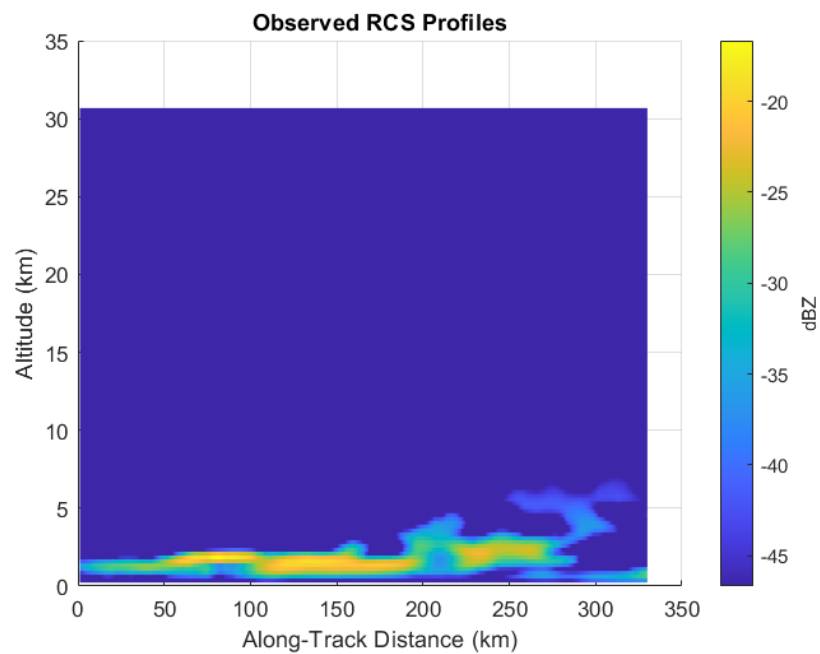


Figure 2: An example reflectivity field extracted from GEOS5.

The data chosen from GEOS5 provided a fairly weak but still detectable target that ranged from about -18 dBZ down to less than -40 dBZ , as can be seen in Figure 2. Additionally, the water content values of interest tend to stay beneath about 10 km above sea level. In order to form a single 1.1 km averaged observation, three longitude steps were used from the full set of GEOS5 data, since the .0625 degree resolution of each step

corresponded to roughly 385 meters, and an interpolation between them was studied to create a full along track data swath.

Due to its regular formatting and ease of access, GEOS5 provided a very convenient source of data early in the project, but it still presented some challenges. First, the approximation formula used to convert water content to radar cross section was just that, an approximation with somewhat loose error bounds in the cited literature. Second, the data was presented at a very coarse horizontal resolution as mentioned previously, on the order of 7 km, which presented a mismatch with the 1.1 km along-orbit distance covered by a single CloudSat processing interval.

Another potential source of atmospheric data exists in the form of data measured by existing atmospheric sensors. For this purpose, data records from the CloudSat mission were very helpful [31]. CloudSat [9], [19], [66] data products include fields for equivalent radar reflectivity correlated to measured height indexes as well as fields for noise floor estimates and standard deviations. An example reflectivity field derived from the GEOPROF product available from the CloudSat website is shown in Figure 3.

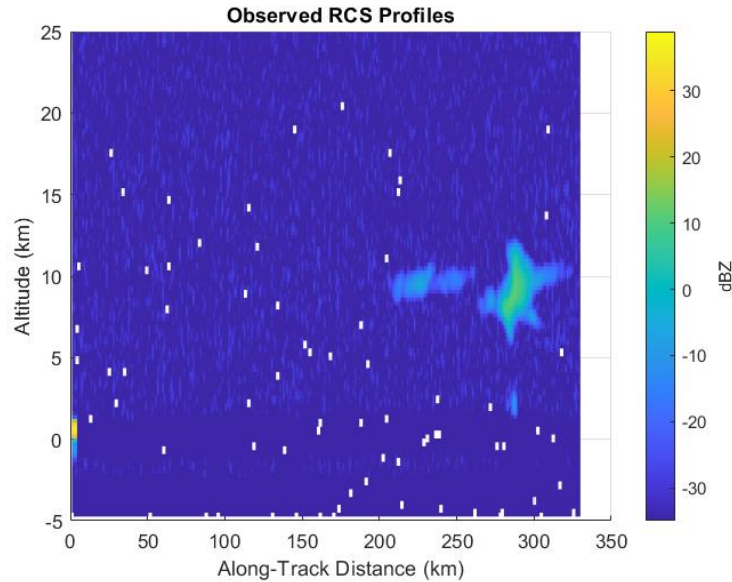


Figure 3: An example reflectivity field extracted from the CloudSat GEOPROF product taken from granule 3607 on the first day of 2007.

Figure 3 contains a portion of the observed reflectivity from the first orbit of 2007 captured by CloudSat. The image contains a typical mid-level cloud around 270-320 km along track, with some holes representing range cells where either no data was recorded or the noise was perfectly canceled out. Compared to the GEOS5 dataset, the available CloudSat data has the advantage of being unarguably realistic and already present on a height grid corresponding to CloudSat’s measurement capabilities. Additionally, the CloudSat GEOPROF product included a cloud mask that identified the location of cloudy bins in the observed reflectivity field as well as possible spurious or ground returns. This represented an additional data point that the GEOS5 set lacked and could be used as a source of truth data for a radar simulation.

Using CloudSat data however also presented its own set of difficulties. As can be seen in Figure 3, the dataset featured holes where bad or missing data was present that had to be

suppressed. In addition to this, the CloudSat data was contaminated by a very strong ground return (suppressed in Figure 3). For the purposes of the simulation, it was assumed that the ground return could be handled by another system through the use of simple thresholding or a similar method in which the cells marked as ground return by CloudSat's mask were suppressed. Another difficulty was the fact that the CloudSat range grid was irregular due to orbital mechanics. Because the range grid was provided in the dataset, however, it was assumed that the same system responsible for generating the range grid would be capable of keeping it in alignment from pulse-to-pulse during operation. Finally, owing to its nature as a radar product, data points in the CloudSat returns were only available for each 1.1 km processing interval rather than for every pulse in each processing interval. In order to provide data for each pulse to illuminate in the radar simulation, it was assumed that the CloudSat reflectivity data represented the ensemble average return from which individual pulse returns would observe individual realizations.

2.3 PRF Adaptation

Owing to the nature of clouds as stochastically varying structures, their radar signature varies considerably over a relatively short period of time. For static targets, integration can be used as a tool to “boost” SNR because the signature of the target remains constant over multiple pulses while the noise signature adds destructively down to a point, resulting in SNR gain whether coherent or incoherent processing is used [12]. For clouds, however, averaging is instead used as a means of improving the quality of the estimate the radar generates of the scene's underlying SNR [13]. Considerations of sample independence notwithstanding [13], it is known that, for incoherent integration, estimate uncertainty

generally scales as $\frac{1}{\sqrt{N}}$, where N is the number of pulses integrated. This would imply that estimate precision can be increased by allowing the number of pulses integrated per measurement to increase. The magnitude of the increase is limited by the aforementioned problem of sample independence and the capabilities of the radar hardware being used. The $\frac{1}{\sqrt{N}}$ bound represents a best case when returns from all pulses are statistically independent, which will only be satisfied if the length of time between pulses is greater than the time the scene takes to decorrelate. If the PRI is shorter than the decorrelation time, the scale factor for estimate uncertainty will decrease according to the correlation still present between samples. Additionally, the number of pulses available for a given measurement will be upper bounded by the radar hardware itself as well as the length of the processing interval afforded to each measurement. If the hardware cannot sustain a PRF above a certain value, then a fixed processing interval will set an upper limit on the number of pulses that can be processed within that time. Finally, onboard resource limits will also set another, possibly more stringent, upper bound on the number of pulses per measurement as more pulses implies more power and data usage. Thus, while an improvement in estimate uncertainty can be realized by allowing the number of pulses used per measurement to increase, the improvement is not unbounded and comes at a cost.

In order to allow N to increase for a given measurement, assuming that the processing interval remains fixed, the PRF must increase. The cloud scene itself sets some limits for the range of allowable PRF values. Considering the example reflectivity fields shown in Figure 2 and Figure 3, it can be seen that the clouds in the scenes observed generally existed between ground level and about 10 km. As discussed in Chapter 1, accurate measurement

of distance using a pulsed radar sets an upper bound on the PRF used based on the unambiguous range viewable with a given PRF, calculated as

$$R_u = \frac{c}{2PRF} \quad (2.3)$$

In order to accurately measure the extent of the clouds scene in Figure 2 and Figure 3, R_u must be kept above about 10 km in the first case and 5 km in the second, placing an upper bound on the PRF of 15 kHz and 30 kHz, respectively. Additionally, due to the nature of range ambiguities and the range folding phenomenon discussed in Chapter 1, it is possible that otherwise useful measurements could be contaminated by false returns or even the ground return if the PRF is allowed to change without bound. A potential workaround for this problem would be to set R_u such that the entire portion of the atmospheric column from the top edge of a detected cloud to the ground surface remains unambiguous. The scene in Figure 2 would not change in this case since the clouds visible in that scene are already at the bottom of the atmospheric column, but the scene in Figure 3 would change in response, increasing the R_u requirement from about 5 km to about 15 km. Though this approach sacrifices some uncertainty improvement, it also prevents new problems from arising as a result of range folding.

Given the considerations mentioned above, how can the PRF actually be allowed to vary under the FAR framework approach? The goal is to simply maximize the PRF for a given measurement subject to the unambiguous range constraints mentioned above. Velocity information could present another constraint on the PRF, but the use of incoherent processing on CloudSat and therefore on the simulation precludes the measurement of velocity. According to the FAR framework instantiation given in Chapter 1, the

measurement vector z_k was defined as the height of the tallest cloud edge detected in a radar profile relative to the ground with the likelihood function being implemented as a simple threshold test on observed SNR. The environment model, $q(x_k|x_{k-1})$, was taken as a constant plus noise model with the noise assumed to be Gaussian distributed with a tunable variance. The external state x_k was therefore the actual cloud top to ground height uncorrupted by observation. Acceptable performance was observed when the standard deviation was estimated by taking the standard deviation of the actual cloud top heights in the CloudSat data set, producing a value of 3610 m. The only parameter permitted to change was the PRF, thus θ_k was simply the PRF value. The measurement cost function was defined as the opposite of the PRF, and the performance cost function was defined according to (2.3) as

$$C_p(\hat{x}_k, PRF) = \begin{cases} 0, & \hat{x} < \frac{c}{2PRF} \\ \infty, & \text{otherwise} \end{cases} \quad (2.4)$$

Given these considerations, the executive cost function for the PRF adaptor becomes

$$C_E = \underset{PRF}{\operatorname{argmin}}(-PRF) \quad (2.5)$$

$$s. t. PRF \leq \frac{c}{2\hat{x}}$$

where \hat{x} is the cloud top to ground height estimated from the measurement z taken by the radar. This has the effect of causing the system to choose the highest PRF permitted for a given measurement based on the observed cloud top to ground height. Absent the consideration of resource constraints, this represents the only driver for PRF behavior.

Based on the above information, it would appear that (2.5) is optimizing for past measurements. While it is generally true that radar measurements of clouds decorrelate from sample to sample when satellite borne measurements are being considered, some properties of cloud structures allow for the assumption of continuity between measurements and the assumption that clouds often have an along track extent greater than that covered by a single CloudSat processing interval [67]. Based on this assumption, the cloud top to ground height of subsequent measurements can be inferred from previous measurements. In order to systematically handle this process, a Kalman filter is used to project measurements forward under the assumption provided by the state evolution model $q(x_k|x_{k-1})$, which will be described in Section 2.6. The smoothing prevents spurious detections from having a drastic impact on the choice of PRF from measurement to measurement and provides a principled way to carry estimates forward in the presence of missed detections with increasing uncertainty with each missed detection. The number of misses, defined here as integrated measurements for which no individual range cell contains cloud returns, permitted before the PRF resets to a default is a tunable parameter, and allowing three missed detections before reset appeared nominal and can be set based upon the particulars of an observed scene and the desired speed with which the radar returns to its default state.

2.5 Measurement Terminator

Considering the reflectivity field in Figure 2 again, it can be seen that significant portions of the field contain open air or weak clouds. The entire field above about 10 km appears devoid of clouds, and even the relatively strong cloud structure has some portions

that fall at or below -30 dBZ (or about -11 dB SNR using CloudSat's parameters), which is the quoted sensitivity figure for CloudSat's radar. Since the mission of a cloud profiling radar is to observe cloud returns, open air represents resource waste and even weak signatures present a problem. Considering the results observed using the cloud mask algorithm developed for CloudSat reveals that even some measurements that are below about -20 dBZ will be marked as not containing clouds, even though they sit above the -13 dB SNR minimum set using the -30 dBZ sensitivity threshold [18]. Section 2.3 stated that radar estimate variance depends upon the number of pulses integrated, but it can also be shown (see Appendix) that the relative standard deviation of reflectivity estimates has the form

$$v_e = \frac{\left(1 + \frac{1}{SNR}\right)}{\sqrt{N}} \quad (2.6)$$

Thus the number of pulses integrated is not the sole driver of estimate uncertainty and low SNR returns, such as those from weak clouds, still present significant reflectivity estimate error. SNRs as low as -3 dB present three times as much uncertainty as those with high SNRs; thus, it would be of potential benefit to a resource constrained cloud radar to reject both clear sky and low-SNR measurements because such measurements will either not provide any information on the targets of interest or provide only inaccurate information.

Having established a motivation for discarding measurements, a principled way of doing so is needed. As discussed in Chapter 1, radar observations of cloud backscattering can be assumed to have the following form

$$V = \sqrt{\frac{v^2}{2}}(V_R + jV_I) \quad (2.7)$$

where the real and imaginary parts are both standard normal random variables and the variance $\frac{v^2}{2}$ has been factored out of each [15]. The radar samples are denoted $V(n)$ where n represents the return from a specific range gate and are assumed to be normalized so that backscattered power is given by $p(n) = |V(n)|^2$ which from (2.7) is a scaled chi-squared random variable with two degrees of freedom. The variance in (2.7) is equal to the expected value of the received power [12] and can be written as

$$v^2 = v_s^2 + v_n^2 = v_n^2(1 + SNR) \quad (2.8)$$

where v_s^2 is the expected signal power, v_n^2 is the expected thermal noise power, and SNR is given by $\frac{v_s^2}{v_n^2}$. Under this formulation, it is assumed that v_n^2 is known a-priori and that the goal of the system is to estimate SNR and therefore reflectivity. Defined in this way, SNR is a property of the scene under observation and is independent of the observing instrument. In addition, it is assumed in the formulation that successive samples of V are independent. While this depends upon the stochastic properties of the cloud scene under observation, most clouds change quickly enough compared to a PRI that the assumption is valid, and the additional spectral broadening that occurs when clouds are observed from space tends to guarantee scene decorrelation within a PRI [13], [15], [17].

Based on (2.8), incoherent integration of multiple power samples results in the following random variable

$$Q'_N = \frac{1}{N} \sum_{n=1}^N p(n) \quad (2.9)$$

where Q'_N is distributed as a scaled chi-squared random variable with $2N$ degrees of freedom, N is again the number of pulses integrated, and $p(n)$ is the observed power calculated according to the squared magnitude of (2.7). The mean of (2.9) is given by (2.8), which means that (2.9) provide estimates of SNR by considering

$$Q_N = \frac{Q'_N}{v_n^2} - 1 \quad (2.10)$$

This new random variable is also a χ_{2N}^2 random variable with a mean equal to the SNR and a standard deviation equal to the mean multiplied by (2.6). Equation (2.10) thus provides a means to estimate both SNR and the resulting estimate uncertainty and indicates that, for a fixed value of N , the estimate uncertainty depends only upon SNR.

The probability density function of Q_N can be written as

$$P(Q_N) = \left(\frac{2N}{1 + SNR} \right)^N \frac{(Q_N + 1)^{N-1} e^{-N \left(\frac{Q_N + 1}{SNR + 1} \right)}}{2^N \Gamma(N)} \quad (2.11)$$

where $\Gamma()$ represents the standard gamma function. Since the goal of the radar is to estimate SNR, (2.11) is not explicitly known during an integration and a Bayesian approach can be used to produce a usable estimator [68]. The probability of observing a particular true SNR given a sample of Q_N can be written as

$$P(SNR|Q_N) = \frac{P(Q_N|SNR)P(SNR)}{\int P(Q_N|SNR)P(SNR)d(SNR)} \quad (2.12)$$

where $P(SNR)$ represents any a-priori information known about the SNR. In the absence of such information, SNR can be assumed to be uniformly distributed over the regions of interest and drops out of (2.12), resulting in

$$P(SNR|Q_N) = \frac{P(Q_N|SNR)}{\int P(Q_N|SNR)d(SNR)} \quad (2.13)$$

Equation (2.11) represents $P(Q_N|SNR)$ since knowing the SNR is required to evaluate it.

Substituting into (2.13) and simplifying results in

$$P(SNR|Q_N) \quad (2.14)$$

$$= \frac{[N(1 + Q_N)]^{N-1}}{[1 + SNR]^N \Gamma(N - 1) \gamma(N(1 + Q_N), N - 1)} \exp\left\{-\frac{N(1 + Q_N)}{1 + SNR}\right\}$$

where $\gamma(x, n)$ represents the incomplete gamma function given by

$$\gamma(x, n) = \frac{1}{\Gamma(n)} \int_0^x q^{n-1} e^{-q} dq \quad (2.15)$$

Equation (13) therefore provides a method of calculating the likelihood of observing a particular true SNR value given a value of the estimator Q_N . In order to use (2.13), first produce a sample of Q_N and then calculate $P(SNR \leq SNR_{min}|Q_N)$ in order to determine the likelihood that the underlying true SNR, given in (2.8), is to fall below a chosen threshold value of SNR_{min} . Performing the integration produces the result

$$P(SNR \leq SNR_{min}|Q_N) = 1 - \frac{\gamma\left(\frac{N(1 + Q_N)}{1 + SNR_{min}}, N - 1\right)}{\gamma(N(1 + Q_N), N - 1)} \quad (2.16)$$

A threshold can then be set on the probability of a measurement falling below the minimum SNR goal, referred to as T_{term} , that will trigger a measurement to be rejected. Thus, (2.16)

provides a means of rejecting measurements that have unacceptable levels of estimate uncertainty based upon their SNR levels in a principled way. The algorithm functions as follows: as pulses are collected, estimates of Q_N are produced. The estimates are used in (2.16) to calculate a likelihood that the resulting scene SNR falls below a minimum SNR_{min} value. If the likelihood of low quality rises above a threshold T_{term} , the measurement is terminated before all pulses are collected.

2.6 Code Analysis

The simulator is a modular, function-based program written in Matlab. The overall logical flow of the simulation is shown in Figure 4.

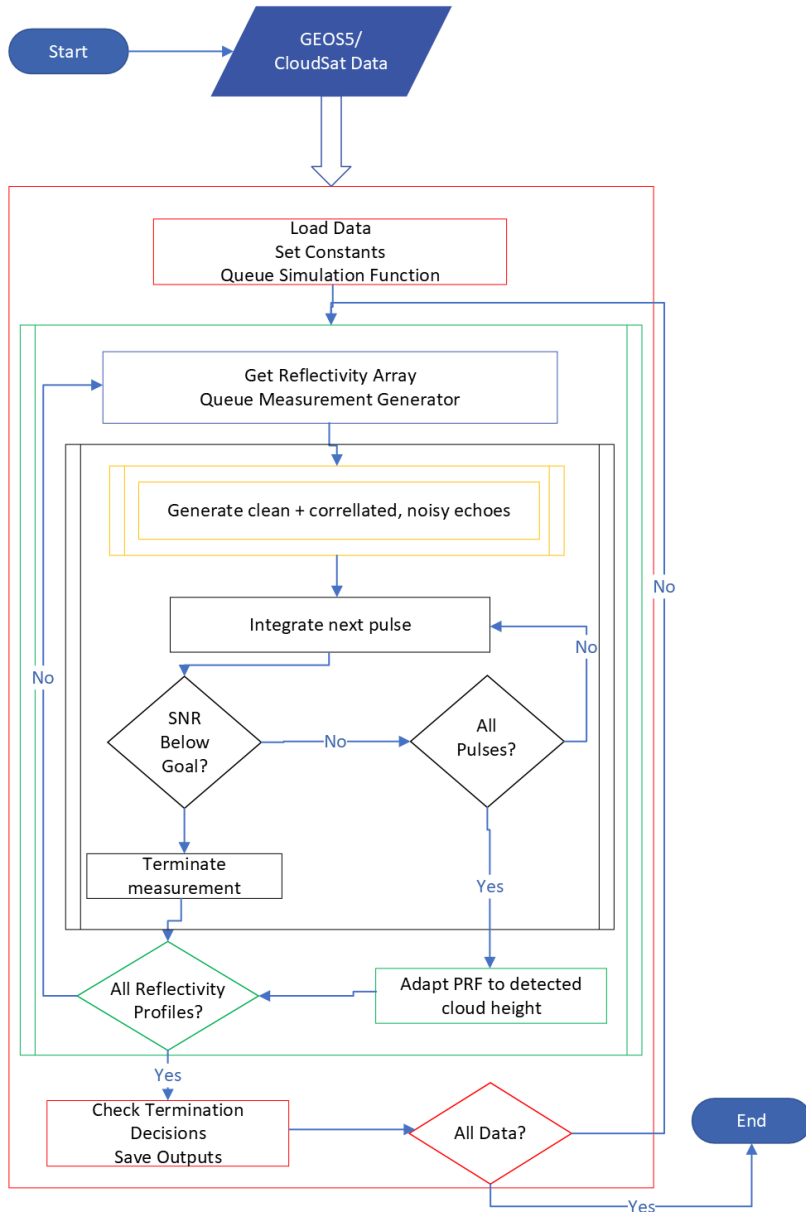


Figure 4: The logical flow diagram of the simulation from meteorological data input to data output. Each logical component will be discussed in this section.

2.6.1 AdaptiveCPR/CloudSatSimulationScript

The highest-level interface to the simulator is a Matlab script that calls the primary simulating function. This represents the red boxes in Figure 4. The script is used to set

initial values for the radar range equation parameters like transmit power, PRF, and noise figure as well as the tunable covariances used for tracking of cloud column height once detected. The script also has logical flags that allow the simulation to be run in incoherent mode, with PRF adaptation, with high-uncertainty measurement rejection, or in any combination thereof. After the initial values are set, the script handles accessing the directory containing the input data coming either from GEOS5 or CloudSat and creating a directory for saving the generated results. Execution then enters a loop where the script calls the appropriate simulator function and passes the initial input parameters to it and handles plotting and saving of data produced afterwards until all the data have been processed.

AdaptiveCPRSsimulationScript was the first version of this script and was intended to handle both GEOS5 data and CloudSat data while AdaptiveCloudSatSimulationScript is more streamlined to handle the CloudSat data specifically. In addition, there is a version of AdaptiveCloudSatSimulationScript that uses a parallel loop instead of a sequential loop to process all the CloudSat data more quickly. When working with CloudSat data, 101 input GEOPROF files are processed. For the sequential loop, each file is processed in order from earliest to latest. For the parallel loop, multiple files are processed at the same time in order to speed computation. This is not problematic as each input file can be treated as a separate entity and continuous uptime is not a parameter of concern for any of the algorithms being tested in this dissertation. When processing a small subset of data, the script calls a function that produces plots showing the PRF adaptation, detections in range, and correctly/incorrectly terminated measurements for each input data file processed with

the option to suppress these plots when many input data files are processed at once. Additionally, the script saves the output data produced by the simulation function, which consists of fields for observed and estimated reflectivity, raw power received by the simulated radar, information on the PRF adaptation, range detections, and measurement termination decisions as well as some derived products.

2.6.2 adaptiveCPR/CloudSatSimulator

The adaptiveCPR/CloudSat simulator functions are the primary simulation functions in the code and are responsible for running the actual radar simulation on the input data. These functions represent the content within the green box in Figure 4. In this case, the adaptiveCPRSimulator function was the first version written and is used to process GEOS5 data while the adaptiveCloudSatSimulator function was written later and is used to process CloudSat data. The two functions are organized similarly and are arranged according to the following.

First, necessary constants and indexing variables are initialized, including variables used to store simulation outputs as they are created. Second, the input data to be read during the current call of the function is loaded. In the case of adaptiveCPRSimulator, this consists of a vector containing the vertical pressure levels used to index the GEOS5 data in terms of height, which is then passed to a different function called pressToHeight which implements the hypsometric formula to convert from pressure to altitude, and a pre-processed array containing atmospheric water content values that can be converted to reflectivity. In the case of adaptiveCloudSatSimulator, this consists of loading the CloudSat 2B-GEOPROF file at the location indicated by

AdaptiveCloudSatSimulationScript and importing the arrays containing the measured cloud reflectivity, the heights at which the reflectivity was measured, and the cloud mask that is used as truth when choosing which measurements to reject for unacceptable precision. In both versions of the function, an array is then populated with a table of values corresponding to the chi-square CDF for small arguments while the built in `chi2cdf` Matlab function is used for large ones to speed computation.

Third, the first measurement is generated using a call to the `dopplerGEOS5MANAGERadar` function, which simulates the measurement operation and will be discussed shortly. Fourth, the initial measurement is processed by passing it to two functions in order. The first function, called `GEOS5_DetectionAnalyzer`, takes the measurement generated by the Radar function and identifies the full set of targets detected in the measurement, identifying the extremes and extracting the highest and lowest detected targets in altitude. This is done using simple intensity thresholding via the `SNRmin` parameter specified at the beginning of the simulation by the user because it is assumed that the simulated radar must use only onboard information to perform calculations. The second function, called `IGARSSDectectionProcessor`, takes the highest and lowest detections identified by the previous function and calculates the resulting observed cloud column height. If a logical flag is checked, this function is modified to take into account the distance between the highest detection and the lowest range cell reported in the data set in order to put a limit on the PRF adaptation in the simulation. With the first measurement in hand, the simulation is ready to enter the loop that manages all subsequent measurements until the full data file is processed.

The first step in the loop is dependent upon whether or not a target was detected in the previous measurement, or in the very first measurement when the code execution first enters the loop. If a cloud was detected, then a measurement of cloud height is available. This measurement is passed to a function called `GEOS5KalmanFilterPredict` which implements the prediction stage of a Kalman filter according to the measurement function and noise covariances set in the `AdaptiveCPR/CloudSatSimulationScript`. The predicted uncertainty is calculated according to [69]

$$P_{pred} = \Phi P_{prev} \Phi + Q \quad (2.17)$$

which allows the current Kalman gain to be calculated as [69]

$$KG = P_{pred} H' (H P_{pred} H' + R) \quad (2.18)$$

and the next state is assumed to be equal to the current state. Here Φ is the state evolution matrix and is assumed to be an identity since the cloud column height is assumed to remain constant between measurements, H is the measurement matrix and is set to be an identity since no modification is made to the radar's height measurement, and Q and R are tunable noise covariances representing noise in the external state and the measurement process, respectively. The external state noise was set by taking the standard deviation of the cloud top heights observed by CloudSat, which will be discussed in Chapter 3, and the measurement noise was treated as a tunable parameter and the value .001 produced acceptable performance. In the absence of a measurement and an excessive number of misses, the immediately prior prediction is instead used to predict the next cloud column height with an according increase in state covariance. The cloud column height prediction

is then passed to the `dopplerOptfun` function which uses Matlab's `fmincon` function to select the next PRF value according to the following simple objective function.

$$PRF_{min} \leq PRF \leq \frac{c}{2.4\hat{h}} \quad (2.19)$$

Here PRF_{min} is the default lower bound of the PRF used by CloudSat and is set to 3.7 kHz, c is the speed of light, and \hat{h} is the current predicted cloud column height. The result of the optimization is that the simulation chooses the greatest possible PRF value while not allowing the next predicted cloud to become ambiguous in range with a buffer factor of 1.2 used to provide an additional limit on the adaptation.

The use of an optimization function for choosing the PRF enables the adaptation behavior to change depending upon the structure of the provided function, even though the default behavior for the simulation is simple. The PRF is maximized in order to improve the measurement uncertainty according to (2.5) by choosing the maximum allowable PRF based upon the current predicted cloud column height. As the loop continues, if three measurements in a row are taken without producing a target detection, then no prediction is made for the next measurement and the PRF is reset to the minimum value.

Once a prediction is calculated and the new PRF is selected, the same sequence of three functions from the preceding paragraph is repeated with a measurement being generated from the next set of available water content or cloud reflectivity and the resulting cloud column height being extracted. If a cloud is detected, then the measurement is passed to a function called `GEOS5KalmanFilterUpdate` which updates the prediction of the current cloud column height with the new measurement information according to [69]

$$\hat{x}_{post} = \hat{x}_{pred} + KG(z - \hat{x}_{pred}) \quad (2.20)$$

where z is the current measurement and \hat{x} is the predicted observed state when the measurement was taken before the loop repeats. The uncertainty associated with the prediction is also updated according to [69]

$$P_{post} = (I - KG)P_{pred} \quad (2.21)$$

where I denotes an identity matrix. The loop thus continues until the entire dataset is processed.

2.6.3 dopplerGEOS5MANAGERadar

The `dopplerGEOS5MANAGERadar` function is responsible for managing measurement generation. This is represented by the black box in Figure 4. It first calls the `GEOS5RadarRangeProfile` function to produce the signal received at the radar via interaction between the transmitted pulses and the illuminated cloud scenes, though without proper correlation statistics or noise contamination. This function will be discussed in detail shortly. Once the clean received echoes are generated, they are multiplied by a random matrix with the covariance specified according to the information in [17] that represents a typical velocity spread observed by a satellite radar with relatively calm conditions. Specifically, it was assumed that all observed cloud scenes would have negligible contributions from sources of decorrelation except for satellite motion since satellite motion was observed in [17] and elsewhere [13], [15] to be the most significant contributor. The value given in [17] for velocity spread due to satellite motion was 15.5

$\left(\frac{m}{s}\right)^2$. Following this, corrupting thermal noise is added to fully model the result of observing the illuminated cloud scene through a real radar receiver.

In order to implement the measurement rejection algorithm, first the raw voltage samples are converted to power and summed along the pulses. The strongest range cell in each pulse is identified and placed into a sequence. The code then enters a loop where the power samples are taken as samples of Q_n as described by (2.9) and the likelihood of the resulting SNR estimate falling below the goal is calculated for each pulse in the integration up to the maximum number of pulses according to (2.14). If the probability calculated by (2.14) ever rises above a set threshold, the integration is stopped and a flag is set indicating that the current measurement will be rejected. Whether or not the integration is terminated prematurely, all the collected data are then incoherently summed over the pulses, the known thermal noise power is subtracted from the result, and a threshold check is done to identify any clouds in the resulting range profile if the measurement was not terminated early. After threshold detection, a clustering algorithm is used to connect detections in adjacent range bins that presumably represent contiguous cloud structures to reduce the total number of detected targets. For each individual detection in the range profile, a tunable number of adjacent cells is searched (currently three) to either side. If any adjacent cells are marked as detections, then the cells between that cell and the current cell under test are also marked as detections and the underlying cloud is assumed to exist across cells. Matlab's `bwconncomp` function is then used to extract the location and boundaries of each detection in the profile, and the range cell corresponding to the highest altitude detection is recorded. The function then passes back the completed, incoherently integrated

measurement, the uncorrupted received echoes, and any detected targets up to the calling function.

2.6.4 GEOS5RadarRangeProfile

This function is responsible for creating the uncorrupted radar echoes resulting from illuminating the cloud scene specified by the input data. This represents the orange box in Figure 4. The function receives the radar range equation parameters, the desired number of pulses, the input vector listing the altitude of each range bin for the current measurement, and the input vector of radar cross section (RCS) data for the current measurement from its calling function as inputs. In order to generate the received echoes, the provided radar cross section vector, which provides a value for radar cross section per unit volume for each range cell in the illuminated profile, is convolved with a copy of the user-specified transmit waveform to create the profile of the received return. The operation is equivalent to (1.4) without taking the complex conjugate of the second input. The convolution is implemented via elementwise multiplication after taking the Fourier transform of the transmit waveform and the RCS vector, essentially treating the cloud profile as a sequence of point targets. The resulting vector is then copied according to the number of pulses requested by the calling function so that the returns from every pulse are calculated at once rather than in sequence. The output of the convolution is also scaled by the distributed target form of the radar range equation taking into account the range to target. Additional attenuation due to atmospheric conditions is ignored in the case of the GEOS5 data and assumed to be already accounted for in the CloudSat data since the effective reflectivity field in the 2B-GEOPROF dataset is not corrected for attenuation. The final received signal

is then calculated by taking the result of the first convolution, representing the received echo, and convolving it again with the complex conjugate of the Fourier transform of the transmit waveform, implementing the matched filter operation according to (1.4). Once this is done, range folding is implemented by calculating the unambiguous extent permitted by the current PRF choice and comparing that to the length of the received signal vector in samples. Any full or partial segments of the received signal vector that extend past the unambiguous interval are then summed with the unambiguous portion and blanked. If the unambiguous interval is the same length as or longer than the received signal vector in terms of samples, then nothing is done and the received signal is simply passed back to the calling function.

2.7 Initial Results

PRF adaptation was initially tested using the first 100 profiles of the GEOS5 data presented in Figure 2. The relevant operating parameters are summarized in Table 1 [9], [19], [66].

Table 1: The operating parameters used for all radar range equation calculations in this dissertation, taken from CloudSat.

Transmit Power	1.7	kW
Gain	63.1	dBi
Azimuth Beamwidth	0.12	°
Elevation Beamwidth	0.12	°
Pulse Length	3.3	μs
Operating Frequency	94	GHz
Min PRF	3700	Hz
Max PRF	10000	Hz
Processing Interval	0.16	s
Radar Altitude	705	km

The default PRF was set to 3700 Hz, within the range of CloudSat’s operational parameters, with the remaining radar parameters set according to the information in [66] and the radar assumed to orbit at a constant altitude of 705 km for simplicity. Initially, the PRF was allowed to increase without bound except for the unambiguous range constraint, but it was found that this case resulted in excessive simulation runtime so the PRF was capped at 10 kHz. At CloudSat’s default PRF of 3700 Hz, 592 pulses are emitted during the 0.16 second processing interval and incoherently integrated. The number of pulses integrated during the processing interval would increase along with the PRF. Initial results are summarized in Figure 5, first discussed in [57].

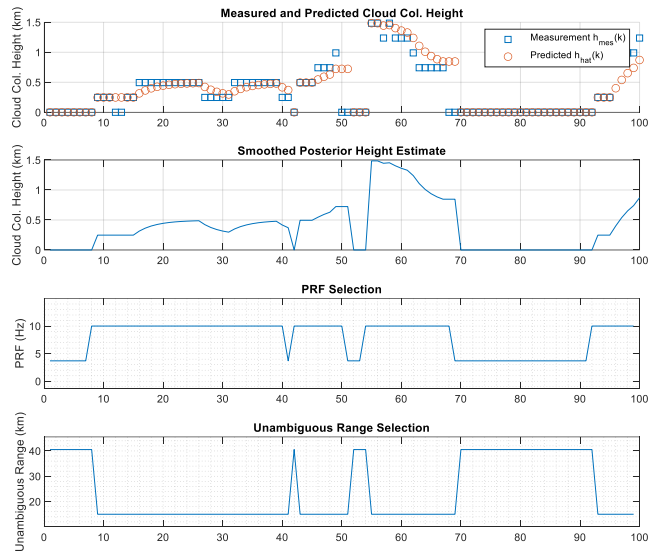


Figure 5: (top row) The measured and predicted cloud to ground height for the GEOS5 simulation, (second row) The predicted cloud height after taking a measurement, (third row) The PRF used for each measurement, (fourth row) The corresponding unambiguous range calculated from the chosen PRF[57]. The x-axis for all plots is the profile number.

Figure 5 shows the simulation detecting the presence of the cloud structure in Figure 2 and the PRF adapting in response. It also shows that the assumption regarding observed cloud column height was generally acceptable, as the majority of the cloud structures in the observation had a constant height for several profiles. As can be seen, profiles 10-40 all consisted of sequences of cloud structures with a constant height for between 5 and 15 profiles, amounting to contiguous cloud structures ranging from 5.5 km to 16.5 km along track assuming a 1.1 km swath length for CloudSat. The structure observed from about profile 45 to profile 70 showed significantly more variation in height along its length as well as a few discontinuities. This represents another potential reason to set a cap on the maximum PRF for measurement as allowing the PRF to adapt all the way down to the cloud top height observed around profile 40 would cause the portion of the cloud structure from profile 55 to profile 65 to either become range wrapped or truncated. The bottom plots show the PRF increasing in the presence of observed clouds and decreasing in their absence. This behavior indicated that the PRF adaptation was behaving as expected without user input and could reasonably be expected to succeed with real data. Results observed when applying the PRF adaptation to CloudSat data will be discussed in Chapter 3. It is important to note that the “predictions” in the plots above are the updated post-measurement predictions generated by the Kalman filter after each measurement is taken which represents the filter’s expectation of the next observed state, rather than an estimate of the measurement just made.

The measurement terminator algorithm was first tested using a single range cell simulation where random complex valued voltage samples were generated according to a

specified SNR and presented to the algorithm. For the tests, 5000 realizations for each combination of true SNR, ranging from -10 to 10 dB, and SNR_{min}, ranging from -3 to 3 dB, were tested. Each realization consisted of 200 samples of $V(n)$ representing 200 radar pulses from which Q_N was computed from N for 1 to 200. Integration was marked as terminated for the first value of N for which $P(SNR \leq SNR_{min} | Q_N) \geq T_{term}$ was satisfied. Because both the true SNR as generated and Q_{200} were known, the portion of correct and incorrect decisions could also be calculated and used to gauge the algorithm's success. Results are presented in Figure 6 for T_{term} values of 0.8, 0.85, and 0.9, and an SNR_{min} value of -3 dB for the first three plots.

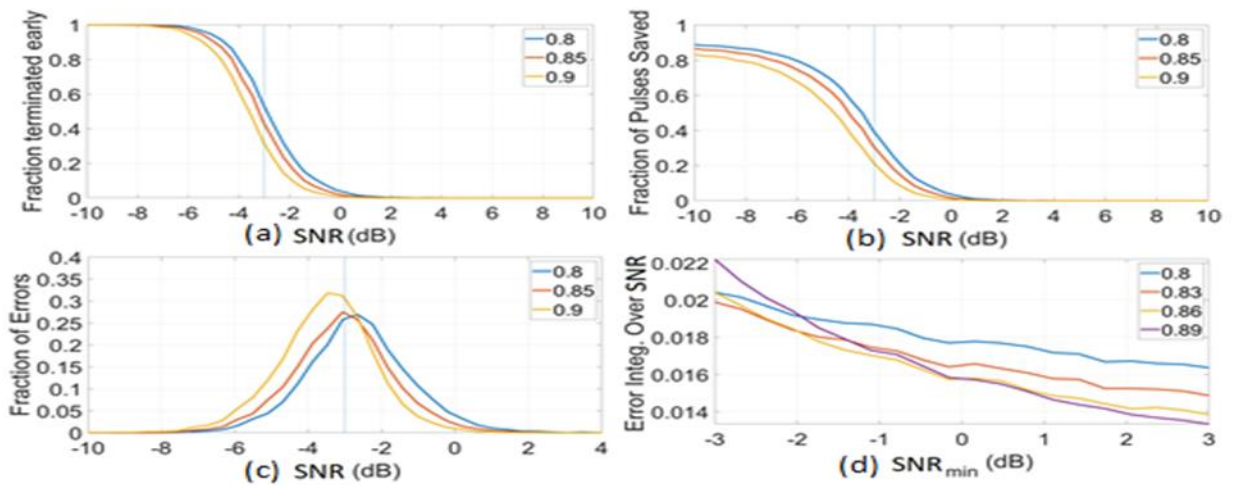


Figure 6: (a) Mean fraction of integrations terminated early as a function of the truth SNR for SNR_{min}=-3 dB for T_{term} values indicated in legend (b) Mean fraction of 200 pulses saved as a function of the truth SNR for SNR_{min}=-3 dB for indicated T_{term} , (c) Mean fraction of errors as a function of the truth SNR, (d) SNR-integrated error versus SNR_{min} threshold

Taken together, the plots in Figure 6 show that the algorithm tends to produce relatively low error rates except for those cases where the underlying true SNR is approximately equal to the SNRmin threshold. Figure 6(a) and (b) show that the algorithm terminated more measurements when the SNR was low and did not when the SNR was high, as expected. More interesting is the fact that the transition from many terminations to few terminations was gradual, meaning that the algorithm would struggle with measurements that were close to the cutoff threshold. This is borne out by Figure 6(c), which shows that the errors were concentrated around the SNRmin threshold of -3 dB. In this test, both false completions and false terminations were weighted equally and the overall error peaked around the SNRmin threshold. Higher probability goals resulted in more false completions (not terminating weak measurements) than false terminations (terminating strong measurements), with the 0.85 case appearing to perform the best overall. Either type of error could clearly be preferred by simply raising or lowering the probability goal. Finally, Figure 6(d) provided a view of errors integrated over all generated SNR values and plotted over each SNRmin value, showing that errors tended to decrease as the quality of desired measurements (as chosen by the SNRmin parameter) increased. This is likely due to the fact that the uncertainty inherent in each measurement would increase directly with SNR, so low SNR measurements were more easily contrasted against a stronger goal. Figure 7 provides another view of the single pulse statistics for a T_{term} value of 0.87, which was observed to produce good results, and further shows that the algorithm handled both very strong and very weak measurements well while struggling somewhat with measurements that were near the cutoff SNR




Figure 7: (a) Mean fraction of errors for $T_{term}=0.87$ as a function of the truth SNR and SNR_min (b) Mean fraction of pulses saved for $T_{term}=0.87$ as a function of the truth SNR and SNR_min

Figure 7(a) recreates Figure 6(c) for the T_{term} value of 0.87 and the same behavior is observed where the algorithm handles extremes well while struggling with measurements around the SNRmin goal. This would indicate that the algorithm could effectively handle the majority of measurement that merit termination and that error performance could be tuned by setting the cutoff thresholds with knowledge of the SNR values that are expected in the scene of interest. Figure 7(b) recreates Figure 6(b) for the T_{term} value of 0.87 and provides a first-principles estimate of resource savings that are achievable with the algorithm. Based on Figure 7(b), the entirety of measurements that are more than 5-7 dB below the SNRmin goal can be saved.

In order to evaluate the algorithm in more realistic circumstances, it was tested against the reflectivity field from GEOS5 presented in Figure 2. The dataset consisted of 300 profiles each producing a 1.1 km along-track integrated measurement. The change from a single range bin represented a theoretical mismatch, but it was determined that the mismatch could be addressed if the terminator was programmed to process only the strongest return in each atmospheric column presented. Two potential issues with this

approach existed. First, it was known that the result of taking the maximum of multiple chi-squared random variables would not necessarily produce a chi-squared random variable, so the assumption that the distribution would not change may not be accurate. Second, typical detectors involve choosing between a target present and target absent state and using a ratio test to calculate a marginal probability rather than calculating a single tail probability as in this case [68]. In this case, however, the decision maker was intended to only determine whether a target was present with a specified likelihood, thus the tail probability calculation was used for simplicity recognizing that this may have an impact on the results. Correctness was determined by the actual scene SNR presented by the GEOS5 reflectivity. If a measurement was terminated and the actual SNR was under the goal or a measurement was not terminated and the SNR was above the goal, the decision assessed as was correct. If a measurement was terminated and the actual SNR measurement was above the goal or a measurement was not terminated and the actual SNR fell below the goal, the decision was assessed as incorrect. It is important to note that the theoretical mismatch between the single pulse formulation of the algorithm and the application to a full range profile means that the SNR_{min} value may not correspond exactly to the value at which measurements are terminated. Without having a cloud mask available for the GEOS5 data, the “errors” are diagnostic only and can be adjusted by simply changing the threshold used. The SNR_{min} value for the test was set to -3 dB and the T_{Term} value was set to 0.87. According to SNR calculation using CloudSat’s parameters, the -3 dB SNR goal would be approximately equal to -24 dBZ reflectivity. The results are presented in Figure 8.




Figure 8: Test results applying the measurement terminator algorithm to the GEOS5 data as described in Section 2, representing 300 full measurements each integrated over 1.1 km.

Figure 8 indicates that the algorithm was largely successful in handling the GEOS5 data, producing only about a 2% error rate with most errors being due to failing to terminate measurements that presented an SNR below -3 dB after full integration and correctly terminating about 45% of the full set of measurements. As stated before, the value of SNR_{min} represents a tradeoff between the importance of observing certain kinds of clouds and the limitations of a resource constrained platform and is tunable depending upon the designer’s goals. Regardless, this test showed that the algorithm is capable of producing significant resource savings for such a system.

2.8 Conclusions

This chapter presented the development history of the simulation used as the investigative tool for this dissertation. The simulation’s heritage from the STARS project was discussed and the nature of the input data sets coming from both GEOS5 and CloudSat

was discussed. A summary of the code and its layout was provided, and initial results when applied to simulated cloud reflectivity data from GEOS5 were presented. The simulation models a version of the CloudSat radar that is capable of changing its PRF in response to observed cloud to ground height and identifying low precision measurements in process and terminating them before completion to save resources. The PRF adaptation routines were tested initially using a full range profile while the measurement terminator was tested initially using a single range cell and then adjusted to account for a full range profile. Chapter 3 will discuss the results observed when applying the simulation to a large set of CloudSat data.

Chapter 3: Simulation Results

This chapter presents results from running the simulation and discusses their implications. The first section discusses diagnostic simulations used to provide an upper bound on potential resource gains achievable with the measurement terminator and details information on the data set taken from CloudSat. The second section presents the result of running the full simulation for three different cases outlined by the diagnostic simulations and provides further estimates for resource savings. The third section presents an analysis of the decision maker applied to two individual range profiles, one containing a peak SNR of about -2 dB and one containing a peak SNR of about -7 dB to investigate the effects of varying the cutoff threshold. The final section presents two cases using both the measurement terminator and PRF adaptation to make use of the resources saved by measurement termination. All simulations run in this chapter using CloudSat data were run using the equivalent reflectivity, Z_e , field from 101 GEOPROF files representing the first week of orbits from 2007. The radar range equation parameters are recorded in Table 1.

3.1 Performance Predictions

In order to provide upper bounds for the portion of the dataset that would be terminated given a choice of SNR_{min} , an additional simulation was produced that would simply process the entire GEOPROF reflectivity data set and calculate the portion of the reflectivity profiles with peaks below specified SNR bounds. Though the choice of whether or not to terminate a measurement is ultimately made based upon estimated measurement uncertainty rather than SNR directly, the SNR observed in the scene remains

the key driver of the algorithm and has already been related to the measurement uncertainty in the previous chapter. A simple plot showing the relationship between observed reflectivity and SNR is shown in Figure 9

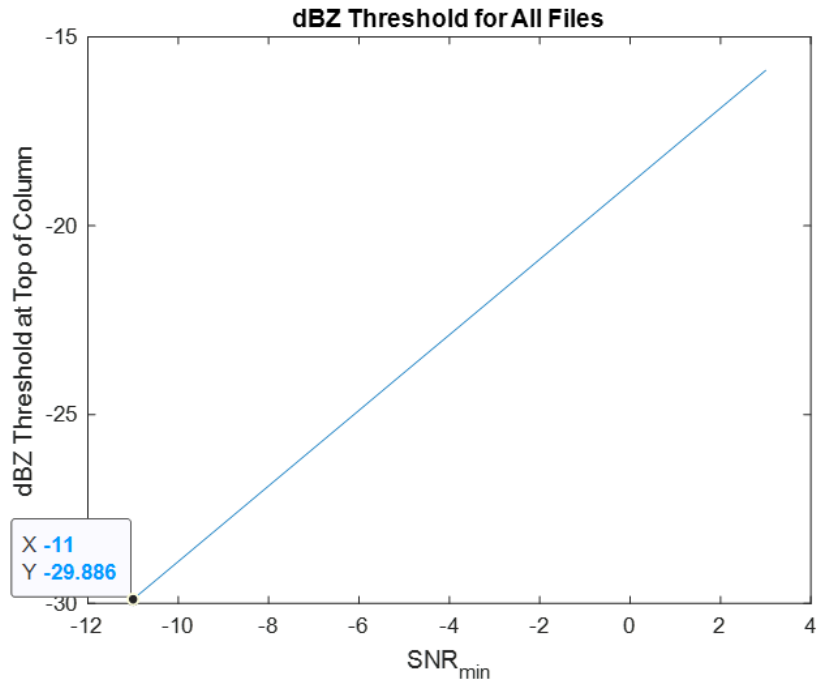


Figure 9: A track of the expected reflectivity as a function of SNR for CloudSat.

Based on Figure 9, the predictor simulation was run using SNRmin values from -11 dB to 3 dB. The results of running this simulation are shown in Figure 10.

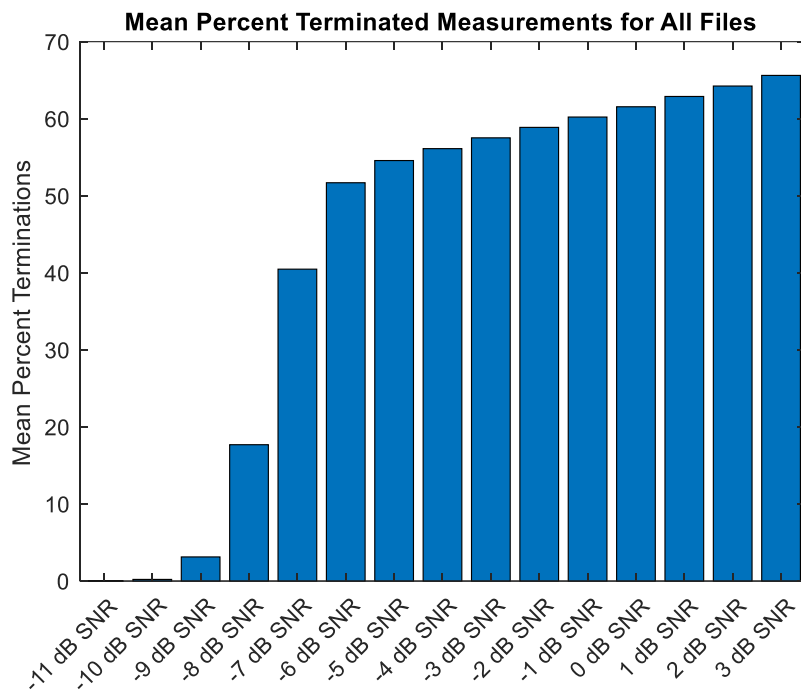


Figure 10: The mean percent of measurements per orbit with an observed SNR below a specified SNR goal choice.

The above plot shows how many measurements on average for the 101 file set would be terminated for the given choices of SNRmin along the x-axis in the absence of noise. The percentages shown above are therefore upper bounds on the possible resource savings achievable with the measurement terminator. As an example, assuming a measurement strategy identical to RainCube’s of measuring one out of every four orbits, choosing an SNRmin value of around -3 dB would allow for measurement of two complete orbits out of the four instead of just one, theoretically increasing the quantity of useful data by a factor of 2. Another measure of interest is the cumulative fraction of profiles with a given peak mask value that exist over the whole set. A plot showing this information is given in Figure 11.

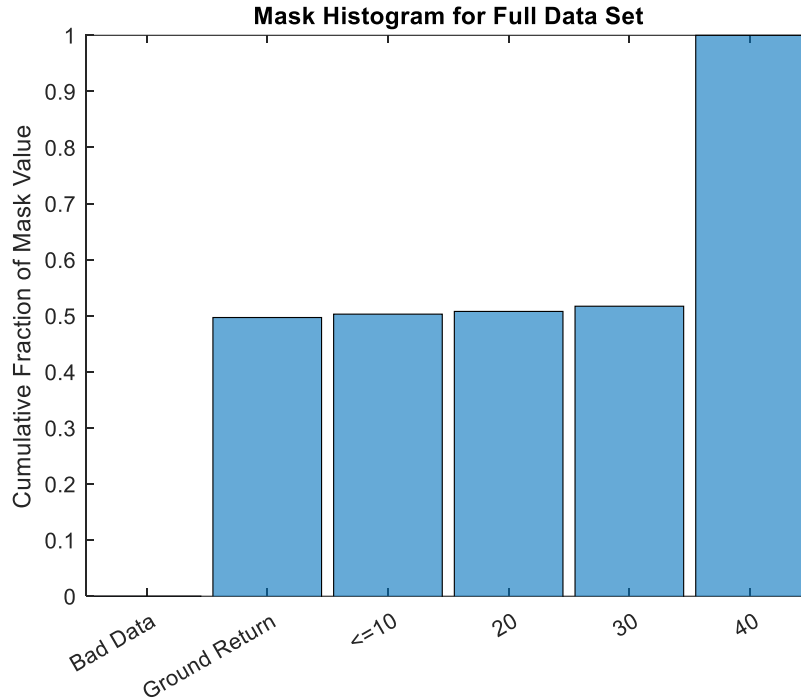


Figure 11: A CDF-scaled histogram of the highest mask value in each profile showing the fraction of each measurement type over the entire dataset.

Figure 11 shows the cumulative fraction of profiles in the data set whose maximum corresponding mask value is in the category shown on the x-axis. Bad data indicates the absence of any discernible return or equipment error, ground return indicates the strongest target in the profile was the ground return, and the corresponding mask sentinel values are as described in [18]. Figure 11 compliments Figure 10 by showing that about half the full dataset contains the highest quality detections, another ~48% contains only ground returns, and only about 2% of the remainder contains lower quality detections between the ground return and the highest quality. Finally, a histogram showing the cumulative fraction of measurements with a given peak reflectivity is shown in Figure 12.

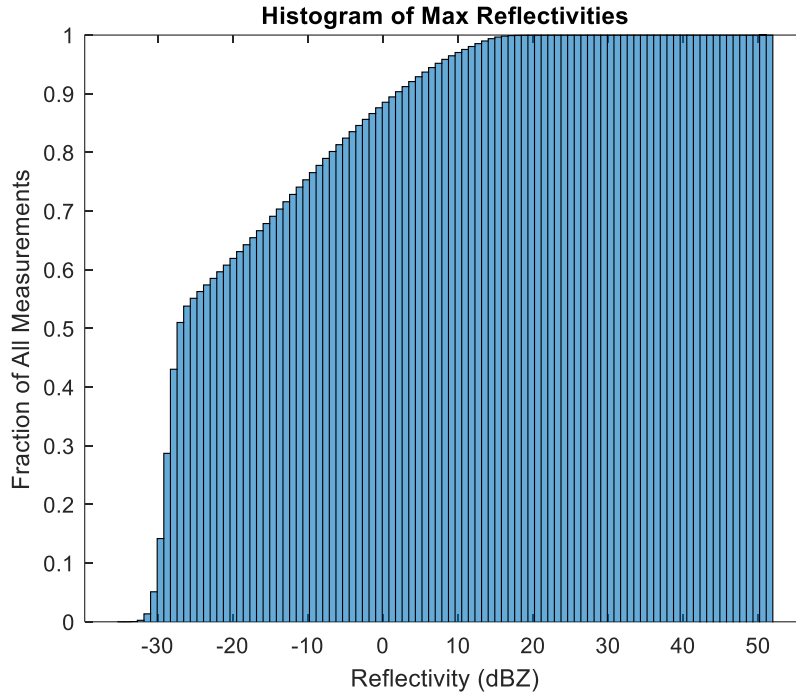


Figure 12: The cumulative fraction of all measurements having a given peak reflectivity for the CloudSat data.

Figure 12 implies that about 50% of all the measurements had a reflectivity peak of at least -20 dBZ and about 60% had a peak of at least -19 dBZ. This reinforces the fact that the resource savings potentially achievable with this method must be balanced in a tradeoff with sensitivity but that significant savings could be achieved depending upon the choice of cutoff. It also implies that the portion of high quality detections given in Figure 11 exist almost exclusively above -20 dBZ reflectivity.

3.2 Simulation Results Without PRF Adaptation

After producing the performance projections, the simulation was run through the 101 file data set for SNRmin values of -11 dB, -8 dB, and -3 dB to examine the algorithm's performance. The probability goal for the algorithm was set at 0.87 based upon the results presented in Chapter 2. Mean error rates for the three cases are shown in Figure 13, with

“completion errors” occurring when a measurement of a profile with mask value below 20 was allowed to complete and “termination errors” occurring when a measurement of a profile with a mask value of at least 20 was terminated. This provides a more rigorous definition of “errors” relative to the GEOS5 test in Chapter 2, but it is important to recall that the thresholds used for this determination could easily be changed based upon the observed terminated reflectivities which are also available for the GEOS5 data.

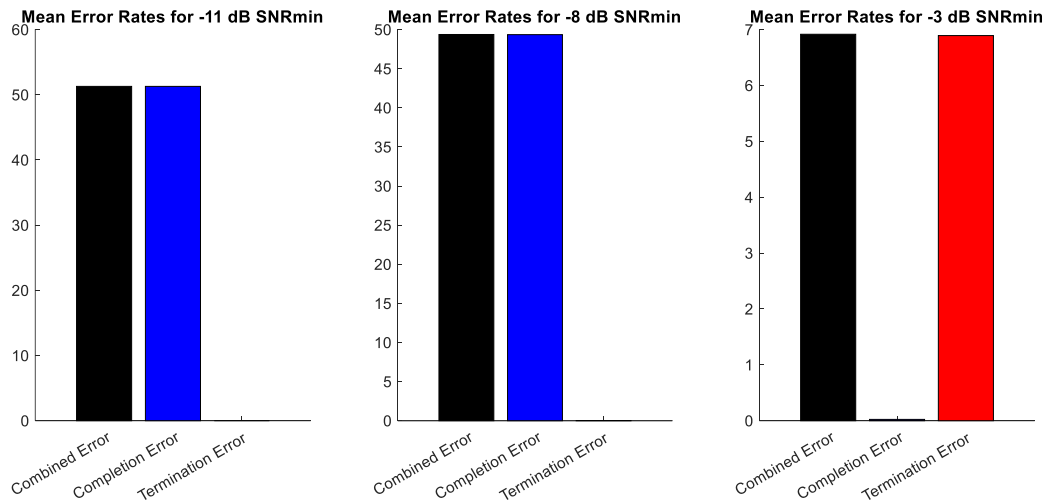


Figure 13: The error rates averaged per orbit for the CloudSat data

The corresponding mean correct decision rates with the definitions following from the error definitions are shown in Figure 14.

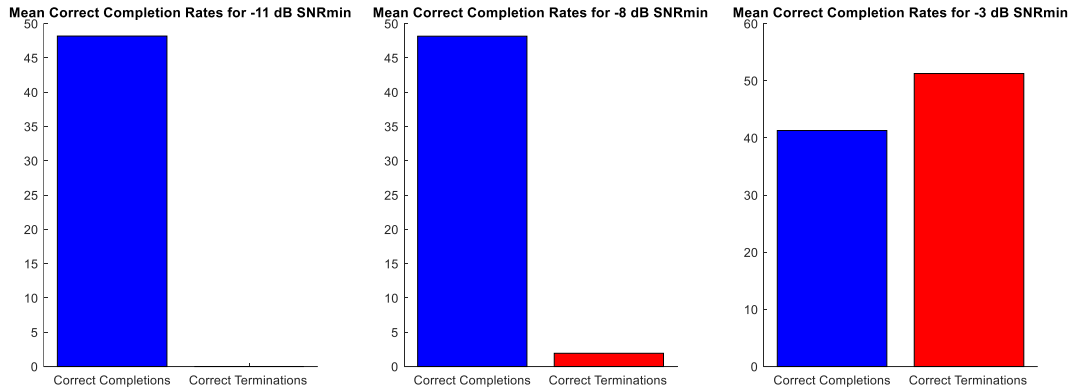


Figure 14: The correct decision rates averaged per orbit for the CloudSat data.

Immediately, it is clear that the simulation did not terminate any measurements for the -11 dB SNR case and only very few in the -8 dB SNR case. For the former, this is exactly in line with the prediction in Figure 10 and makes sense given the fact that -11 dB SNR roughly corresponds to the -30 dBZ sensitivity floor quoted by the CloudSat team. No measurements were terminated because the instrument was incapable of producing measurements below -30 dBZ reflectivity. For the latter, the difference between the actual portion of measurements that were terminated and the prediction given in Figure 10 is likely due to the aforementioned fact that the predictions were made according to perfect information whereas the full simulation had to cope with the presence of noise in the measurements. In order to understand what measurements were terminated by the system and account for the theoretical mismatch, a histogram of the terminated reflectivity values is shown in Figure 15.




Figure 15: The cumulative portion of terminated reflectivity values for the -8 dB SNR_{min} case.

While the previous figure shows that only a few measurements (about 2%) out of the total set were actually terminated, this figure shows that those terminated measurements had reflectivities that were at or below about -25.5 dBZ. Thus a relationship can be derived between the uncertainty requirement related to an SNR goal and the reflectivity of measurements that will be terminated. Now considering the -3 dB case, it was seen earlier that just over 50% of the full set of measurements was correctly terminated with a roughly 6% error rate. This is the expected result of increasing the termination threshold and confirms that improved measurement savings have to come at the cost of some lost good

measurements. A histogram of the terminated reflectivity values for this case is shown in Figure 16.

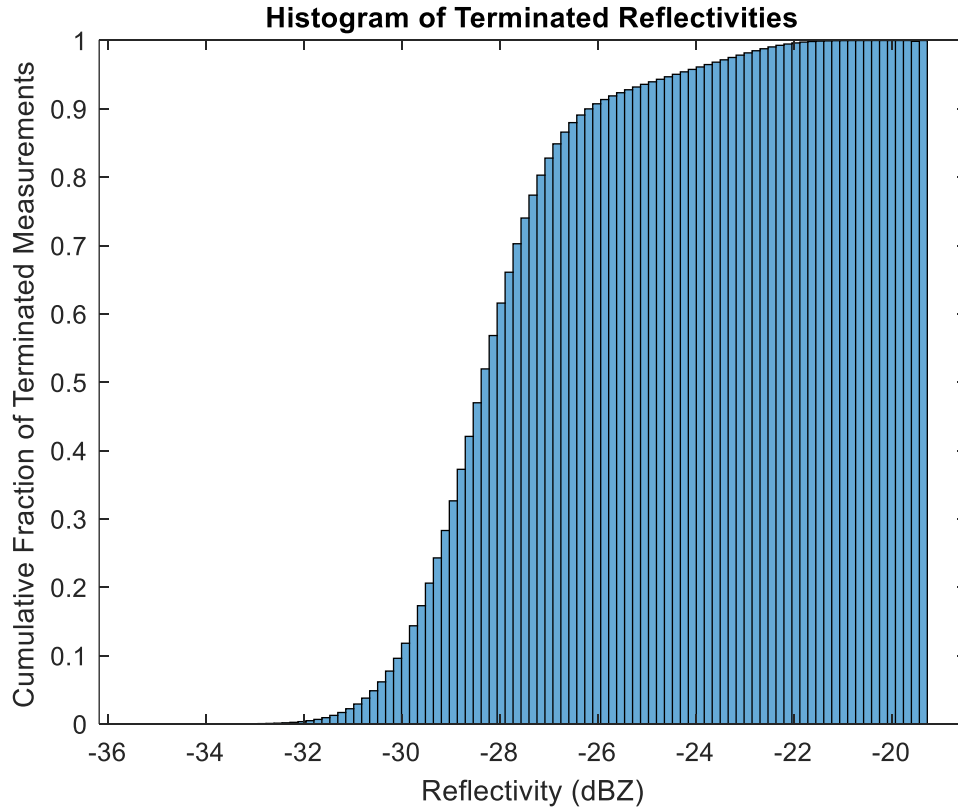


Figure 16: The cumulative portion of terminated reflectivity values for the -3 dB SNR_{min} case.

This plot shows that an SNR goal of -3 dB corresponds to terminating measurements at about -19 dBZ and below, and the majority of terminated measurements are below about -26 dBZ. So, by adopting a cutoff threshold of -3 dB SNR, most of the discarded measurements will present reflectivities below -26 dBZ while only a few measurements presenting higher reflectivity values will be discarded. This implies that a radar designer could balance resource gains against potential loss of science data by allowing the cutoff

threshold to vary based on the plots shown here. To better quantify the potential resource gains, the histogram of pulses integrated for the -3dB case is shown in Figure 17.




Figure 17: The histogram of the number of pulses integrated in each measurement generated from the CloudSat data set.

From the above figure, it can be seen that over 50% of all measurements generated were terminated within the first 115 pulses out of 592, amounting to about 75% resource savings per measurement. It can also be seen that about 58% of all the measurements integrated less than 592 pulses, which corroborates the measurements in the above figures for the -3 dB SNR_{min} test since about 51% of the measurements were terminated correctly and about 7% were terminated in error. Most importantly, this underscores the potential resource savings achievable by establishing a tradeoff between science value and resource use.

Taken together with Figure 11 and the error rates, this proves that using an SNR goal of -3 dB results in up to 50% resource savings over a typical week of CloudSat operation with minimal loss of high quality (mask value 40) cloud detections. The set of terminated mask values for this test is shown in Figure 18.




Figure 18: The histogram of mask values for terminated measurements for the -3 dB SNR_{min}, 0.87 probability goal case.

The resource savings would allow a RainCube-like CPR system to double its useful duty cycle with the same available resource set while Figure 18 confirms that it could do so with minimal loss of science data quality. While it may be desirable to always measure in order to capture clouds down to -30 dBZ, choosing to discard such measurements due to their inherently high uncertainty would enable the deployment of lower-cost constellations with better revisit time and the ability to better study the temporal development of cloud systems with current technology.

3.3 Fine Tuning of the Simulation

CloudSat's cloud mask contains different sentinel values corresponding to different quality levels, with 40 being the highest quality category with the lowest rate of false alarm. The accompanying literature provided with the CloudSat data states that, for most purposes, only the 40 classification should be used in order to ensure the best performance of algorithms performing meteorological parameter estimation using the data. Consequently, two reflectivity profiles were extracted from the full data set in order to provide better tuning for the algorithm, one with a mask value of 40 and one with only a ground return. This would provide a means of determining how the algorithm handles both strong target present and target absent decisions. The profile chosen for individual analysis is shown in Figure 19, which was reflectivity profile 240 of granule 3607.

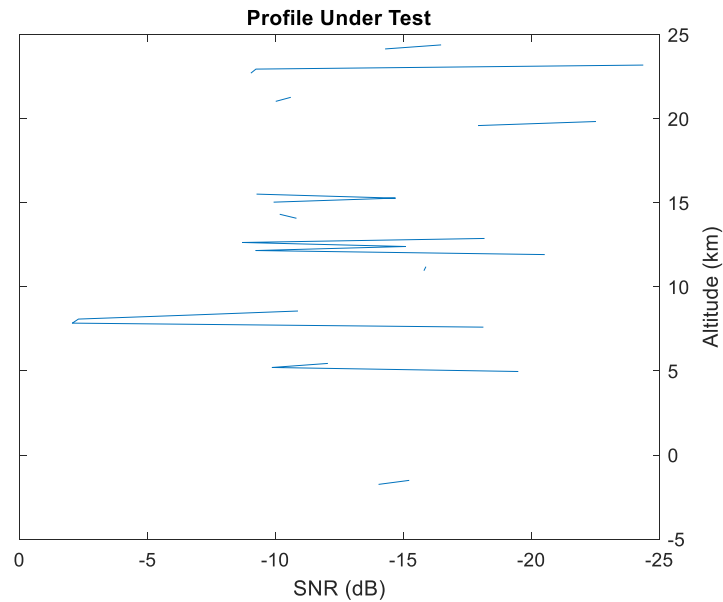


Figure 19: The first profile from the CloudSat set selected for individual analysis converted to SNR with gaps where no reflectivity data was present in the profile

The profile consisted of multiple cloud structures from 5 km above ground to the top of the column, with the strongest return visible around 8 km which was marked as mask value 40. This indicated that the power received from the peak of the profile was at least 3 times the standard deviation of the thermal noise estimate for the file, equating to about -2 dB SNR when related to the mean thermal noise level. The actual reflectivity was around -19 dBZ, which placed it at the transition between termination and non-termination for the -3 dB SNR_{min} case as discussed in Section 3.2. The profile was tested in the simulation over 1000 trials for probability goals of 0.87, 0.90, 0.95, and 0.99. For each of the tested probability goals, the measurement was terminated in 42% of the trials, 32% of the trials, 17% of the trials, and 2% of the trials, respectively. These results revealed two facts. First, the behavior observed in the single range cell tests where the terminator made a significant number of errors when the profile under test had an SNR within 2 dB of the cutoff is evidenced. Second, the problem could be mitigated somewhat by increasing the probability goal for failure, though this would likely have the effect of reducing the number of terminations overall.

The profile selected for ground return analysis was profile 5 of granule 3607, shown in Figure 20.




Figure 20: The second profile chosen for individual analysis with the strongest signal of -7 dB SNR occurring around 19 km

The highest mask value given for the profile was 5, indicating that the strongest detected return in the profile was from the ground even though the peak around 19 km had a reflectivity value of -26 dBZ, which would be 4 dBZ above CloudSat’s quoted minimum detectable signal. This indicates that the return after averaging was still below the standard deviation of the thermal noise for the whole granule. The profile was tested using 1000 trials with a -3 dB SNRmin goal and a 0.87 probability goal, as well as a -8 dB SNRmin goal with probability goals of 0.87 and 0.99. For the first case, the profile was terminated in each of the 1000 trials, indicating that an SNRmin goal of -3 dB is sufficient to reject measurements at this level. Reducing the goal to -8 dB resulted in a termination rate of 72% for the 0.87 probability goal case and 0.1% for the 0.99 probability goal case, indicating consistent behavior from the previous profile regarding the terminator’s performance around the SNRmin cutoff threshold.

The analysis was then extended to the full data set. The base case was the -3 dB SNRmin, 0.87 probability goal test described in Section 3.2. Figure 11 and Figure 13 together indicate that the simulation terminated 57% of the full data set before completion and Figure 16 indicates that terminated measurements had peak reflectivities as high as -19 dBZ with most of the terminated measurements (~80%) having peak reflectivities at or below -26 dBZ. Figure 18 showed that for the given set of parameters, about 12% of the terminated measurements were marked with a mask value of 40 while the majority were marked as ground return. This means that, of the full data set, about 7% of the measurements thrown out were actually of the highest quality discernible by CloudSat. In order to reduce this number, the test was run again with a probability goal of 0.99; the corresponding histogram of terminated mask values is shown in Figure 21.

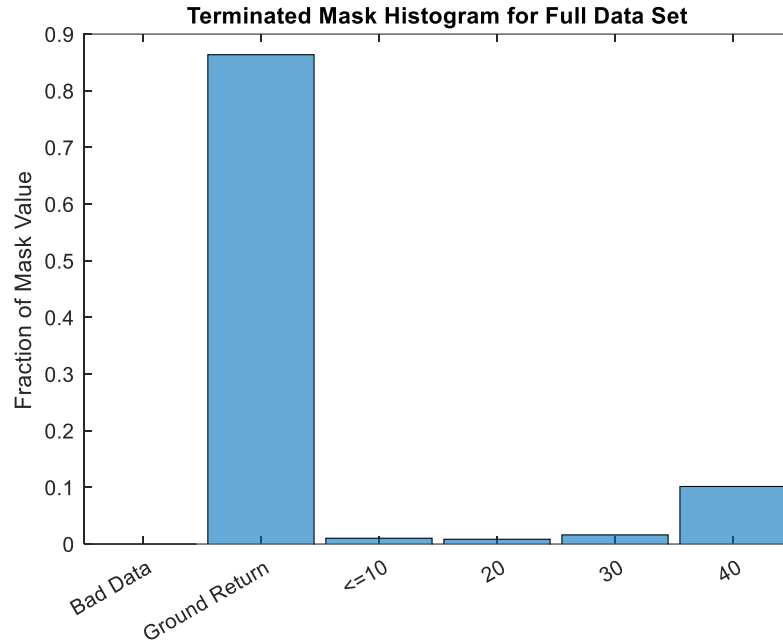


Figure 21: The portion of terminated measurements falling into each mask category for the -3 dB SNR_{min}, 0.99 probability goal case.

Over the full data set, Figure 21 shows that increasing the probability goal to 0.99 had the desired effect but not to the same extent seen for the individual profile 240 of granule 3607, reducing the portion of terminated measurements marked 40 to 10% from the previous 12%. This is likely due to the fact that, over the whole data set, a mask value of 40 corresponds to a range of reflectivity values from about -22 dBZ up, so some profiles marked with a 40 still have weaker SNR values than others.

Lastly, the -11 dB SNR_{min} case was tested again using a 0.5 probability goal over the full data set to determine if exclusively ground returns can be rejected under this regime. The histogram of terminated mask values for this case is shown in Figure 22.

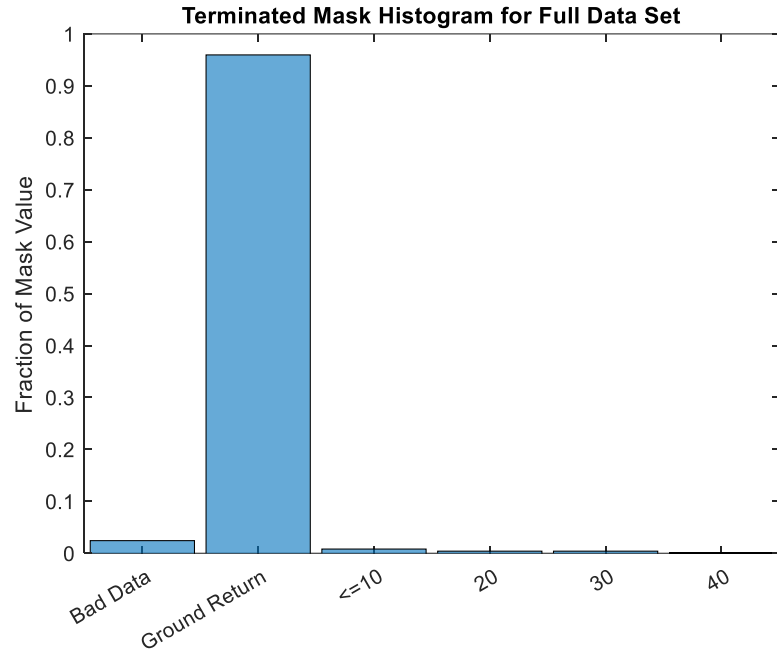


Figure 22: The portion of measurements terminated for each mask value for the -11 dB SNRmin, 0.5 probability goal case.

Thus it seems in this case that the terminator rejected almost exclusively the weak measurements without rejecting any of the categories above mask value 10. Unfortunately, the actual portions of measurements terminated in this case was very small, as shown in Figure 23.




Figure 23: The portion of measurements with a mask value above 10 allowed to complete and terminated for the -11 dB SNRmin, 0.5 probability goal case.

Figure 23 shows that only .006% of the full set of measurements was correctly terminated for this case even with the probability goal being reduced to a coin flip. Given that the terminator is based upon estimating the measured SNR, it appears that the system is unable to generate estimates with sufficient confidence in the absence of a target when the SNRmin goal is at the sensitivity level of the illuminating system. Thus, it can be concluded that the system is able to terminate low quality measurements but not without also discarding a small fraction of the high quality measurements. Given that setting a threshold of -3 dB SNR would have the effect of doubling a RainCube 25% duty cycle, this seems to be an acceptable tradeoff for a resource constrained system.

3.4 PRF Adaptation

PRF adaptation results in similar behavior with most measurements terminated within 200 pulses and an additional step at higher values due to the PRF increase. The simulation

was first tested by allowing the PRF to increase in response to observed cloud top to ground height with no further modifications, meaning the probability goal was held at 0.87, the SNRmin goal was held at -3 dB, and the remaining parameters were as described in Table 1. The initial error rates are shown in Figure 24.

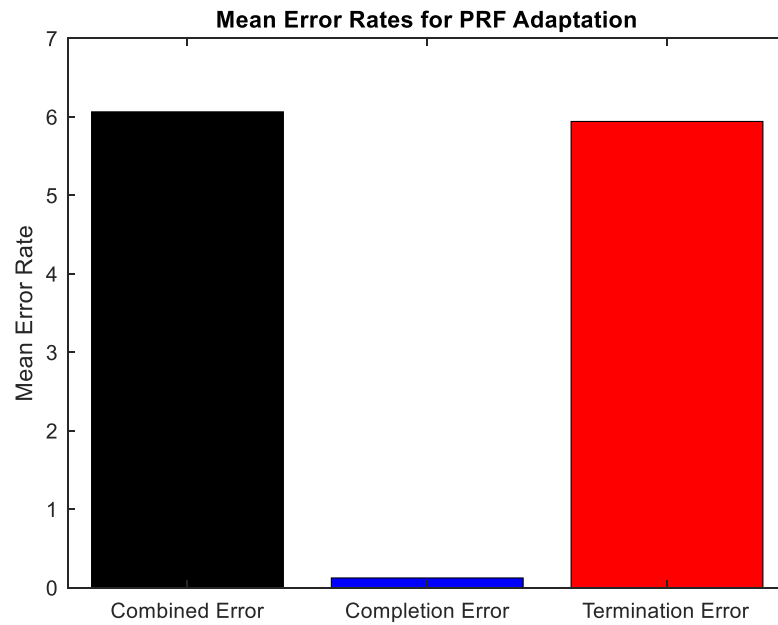


Figure 24: The error rates for the first PRF adaptation test. Note the 1% reduction relative to the -3 dB test without adaptation.

The corresponding correct decision rates are shown in Figure 25.




Figure 25: The correct decision rates for the first PRF adaptation test. The 1% error reduction was split approximately evenly between the two decision types.

As can be observed from the presented data, allowing the PRF to increase resulted in marginal change to the simulation's behavior. Thus it would appear that the PRF increase is of marginal benefit to the CPR case. As an additional test, the simulation was run again with PRF adaptation enabled with the modification that the SNRmin goal would be allowed to decrease as a function of the number of pulses integrated during a measurement according to (2.6) by holding the standard deviation constant with the new value of N and solving for the new value of SNRmin. The error rates for this test are shown in Figure 26.




Figure 26: The mean error rates for the second PRF adaptation test, showing a further .5% decrease relative to the first test.

The corresponding correct decision rates are shown in Figure 27.




Figure 27: The mean correct decision rates for the second PRF adaptation test, showing a 1% increase in correct completion rate.

Allowing the SNR goal to decrease as the PRF increased had the effect of slightly improving the rate of correct measurement completions. This is likely because the SNR goal reduction resulted in the simulation permitting more measurements to complete at the edges of clouds where the reflectivity values begin to fall off rather than judging the trailing edges by the same SNR goal as the stronger central structures of clouds. The set of terminated mask sentinel values is shown in Figure 28.

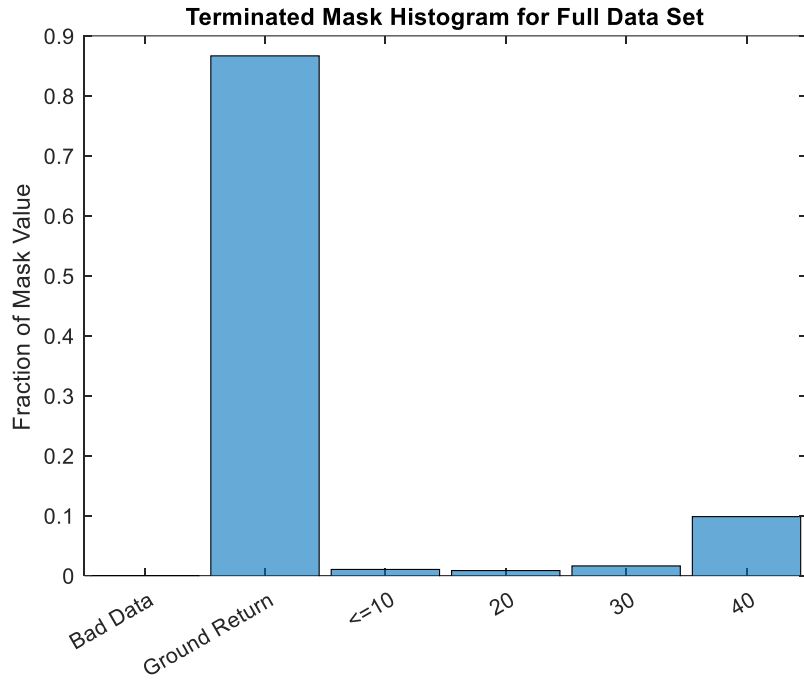


Figure 28: The set of terminated mask sentinel values for the second PRF adaptation test, showing a small reduction in the portion of detections terminated.

The above plot confirms the theory by showing that the portion of terminated sentinel values marking weak detections (20) decreased slightly along with the portion of terminated sentinel values marking strong measurements (40) though again the difference is marginal compared to the previous tests. The corresponding set of terminated reflectivity values was largely unchanged owing to the marginal nature of these differences. The histogram of pulses integrated for each measurement is shown in Figure 29.




Figure 29: The cumulative fraction of pulses integrated in each measurement for the second PRF adaptation simulation.

Figure 29 shows that, even when the PRF is allowed to vary, almost 50% of the measurements were terminated within the first 200 pulses, which is in keeping with the ~50% correct termination rates observed in the rest of the tests. The spread of pulses integrated between 592 and 1600 come about as a result of the simulation detecting clouds of different heights, as well as some clouds initially triggering adaptation then dissipating before the system reached the maximum permitted PRF of 10kHz. The result of allowing the PRF to increase in these cases is reduced uncertainty in the case of strong measurements where clouds are known to be present. The key function of the PRF adaptation in this way would be to take advantage of the resources saved by the measurement terminator by reallocating them to the strong measurements in the current observation pattern rather than by allowing new measurement patterns to be adopted.

An additional test of the PRF adaptation was conducted by isolating a portion of the full CloudSat dataset and running it through the simulation multiple times. The portion under test consisted of profiles 11170 to 12590 of the first data file, which contained results from granule 3607 of the CloudSat data. The reflectivity field is shown in Figure 30.

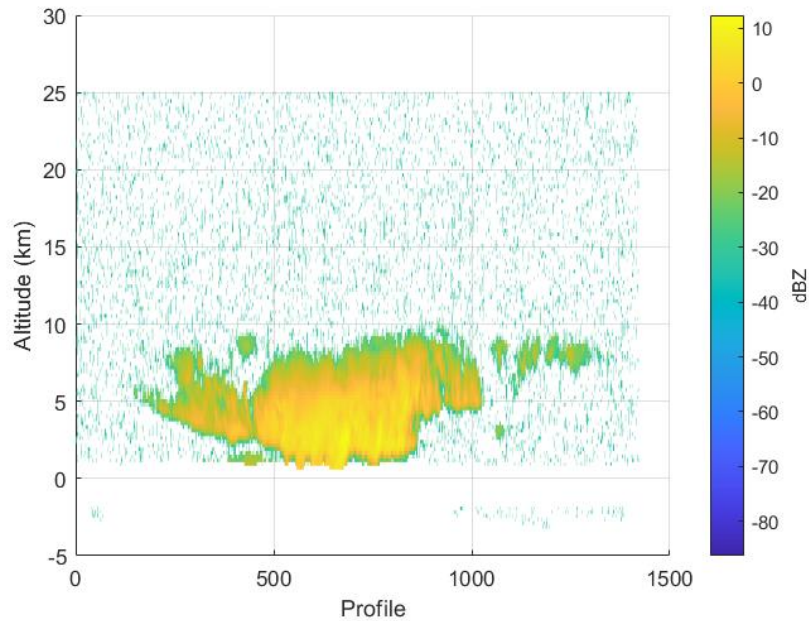


Figure 30: The reflectivity field used for additional PRF variation testing, showing a strong central cloud structure with weaker edges.

Figure 30 was chosen for testing because it presented a strong central cloud structure with relatively weaker edges that would make it more likely for the PRF adaptation to have an impact. The breakdown of reflectivity values in this set is shown in Figure 31.

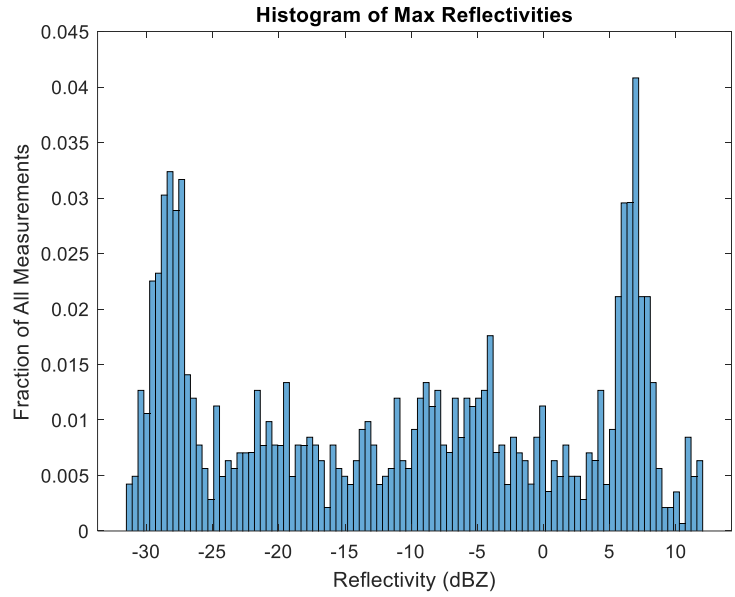


Figure 31: The portion of each reflectivity value in the reflectivity field examined for PRF adaptation.

The corresponding mask value histogram is shown in Figure 32.

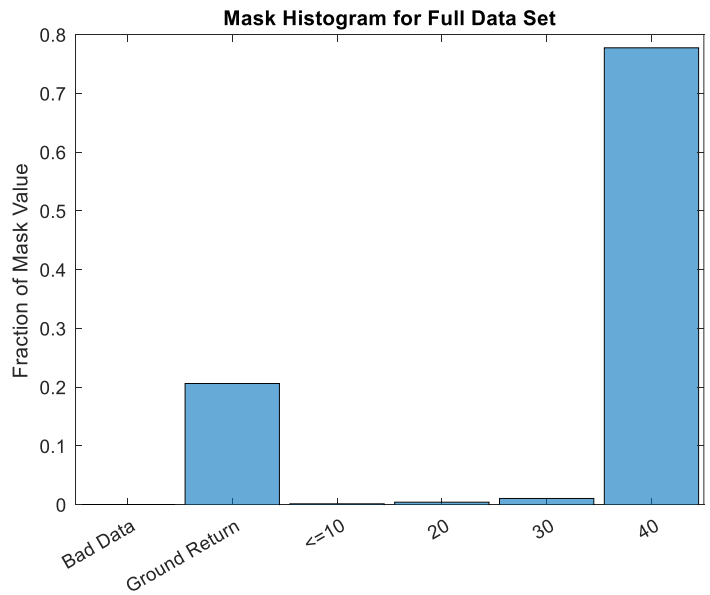


Figure 32: The portion of each peak mask value observed in the reflectivity field used for additional PRF variation testing.

Taken together, the previous two figures show that the majority of this reflectivity field (on the order of 75%) consisted of strong measurements and only about 20% of the field was

marked as containing no cloud. The reflectivity values observed in the field spanned the full range of the CloudSat data set, with the largest individual portions consisting of the strongest and weakest values in the set. To further test the effect of PRF variation, the reflectivity field was run through the simulation 100 times each in three cases. First, the simulation was run without PRF adaptation. Second, the simulation was run with PRF adaptation and a static SNRmin goal as the number of pulses increased. Third, the simulation was run with PRF adaptation and a variable SNRmin goal that decreased as the number of pulses increased according to (2.6) by solving for the uncertainty goal with the default number of pulses and then using that figure to calculate a new SNRmin goal using an increased number of pulses. The set of terminated reflectivity values did not change significantly between the three cases and did not provide a good indicator, so the set of terminated mask values is shown instead as it presented a better indicator of the effects. The set of terminated mask values for the first case is shown in Figure 31.




Figure 33: The set of terminated mask values for the field in Figure 30 without PRF adaptation.

The set of terminated mask values for the second case is shown in Figure 34.




Figure 34: The set of terminated mask values for the reflectivity field in Figure 30 with PRF adaptation and a static SNRmin goal.

The set of terminated mask values for the final case is shown in Figure 35.




Figure 35: The set of terminated mask values for the field in figure with PRF adaptation and a varying SNRmin goal.

Profiles with a peak mask value of 40 accounted for 23%, 24%, and 18% of the terminated measurements in the first, second, and third cases, respectively. As can be seen in the above figures, simply allowing the PRF to vary did not have a significant impact on the portion of mask value 40 measurements that were terminated. This makes intuitive sense given that the probability model does not take the maximum available number of pulses into account and so measurements that would be terminated within the first 200 pulses would still be terminated whether or not an additional 392 or 1400 are used. A 5% reduction in error was observed between the first and third cases, implying that allowing the SNRmin goal to decrease in response to the number of pulses integrated indeed allowed the radar to better observe the measurements with weaker actual reflectivity values that were still marked as high quality detections.

3.5 Summary and Conclusions

This chapter presented a discussion of the results of running the simulation using data collected by CloudSat, essentially simulating a hypothetical version of CloudSat which had PRF adaptation and measurement rejection capabilities. The simulation was tested in various configurations to determine the capabilities of the measurement terminator and it was found that setting a threshold of -3 dB SNR resulted in terminating just over 50% of the collected measurements while allowing 90% of the high-quality measurements to continue and terminating all other measurements. This resulted in measurements with a reflectivity of at least -20 dBZ to be completed. Terminating measurements in this way resulted in halting 50% of all collected measurements, which would have the effect of increasing the effective duty cycle of a resource constrained system. PRF adaptation was found to be of marginal benefit and did not produce a significant change in the overall termination rates during the testing.

Chapter 4: Summary and Conclusions

This dissertation has presented a tool for resource management on a constrained cloud profiling radar system. The tool is formulated around setting an acceptable level of measurement uncertainty and continuously estimating measurement quality during integration, terminating a measurement if it is predicted to fall below the uncertainty goal. The tool also has the option to use an implementation of the FAR framework to increase the PRF of measurements in response to anticipated acceptable levels of measurement precision to take advantage of resources saved via termination.

Chapter 1 discussed background for radar remote sensing, cloud radar, and adaptive radar which presented the motivation for the dissertation. CloudSat and related systems were discussed, with attention paid to CloudSat's various data products and the EarthCARE satellite seeking to continue to advance the CPR problem space. Fundamentals of radar processing were discussed and the implications of the radar range equation when observing clouds from space were presented. RainCube and the CubeSat concept were also addressed as means for enabling more earth-observing platforms to be developed and flown. The challenges associated with the CPR use case and CubeSat design were also presented, along with information about recent developments in satellite and antenna design intended to enable further CubeSat development. Lastly, the concept of adaptive radar was presented and motivated for the CPR task.

Chapter 2 described the development history of the simulation used for the dissertation and detailed its logical structure and mathematical basis. The initial stages of the simulation as a component of the STARS library were discussed, as were the use of GEOS5 as an input source before switching to using CloudSat data. The particulars of the PRF adaptation were then discussed, showing how the FAR framework could be instantiated to tie PRF adaptation to observed cloud column height for the CPR case. The formulation of the measurement terminator was then discussed, which showed how SNR could be used as a proxy for estimating measurement uncertainty and terminating measurements with unacceptable uncertainty. The logical structure and function of the simulation code as then presented, illustrating how the simulation runs.

Finally, Chapter 3 presented the results of testing the simulation in various cases. Ultimately, the measurement terminator was found to be potentially useful for permitting resource savings and extension of measurement duty cycles, while the PRF adaptation was found to be of limited benefit. It was found that an SNR threshold of -3 dB would produce favorable results from a resource conservation perspective while discarding only a small portion of the high quality measurements. It was also shown that the measurement terminator algorithm struggled to effectively terminate measurements when the SNR goal was set well below the typical observed SNR values.

The tool was tested using data derived from both NASA's GEOS5 atmospheric simulation and CloudSat and was shown to be effective in both cases depending upon the chosen input parameters. Future work may include consideration of implementing the algorithm on satellite hardware rather than a general purpose computer as well as a detailed

analysis of the hardware requirements necessary. In addition, investigation into adjustments to the probabilistic model in order to improve performance for low SNRs may also be considered.

Appendix A: Derivation of Power Estimate Standard Deviation

Consider

$$Q_N = \frac{Q'_N}{v_n^2} - 1 \quad (\text{A.15})$$

which forms the basis of the measurement terminator algorithm proposed in this work. The form of Q'_n is given by (2.9)

$$Q'_N = \frac{1}{N} \sum_{n=1}^N p(n) \quad (\text{A.16})$$

where $p(n)$ has the form

$$p(n) = |V(n)|^2 = \left| \sqrt{\frac{v^2}{2}} (V_R(n) + jV_I(n)) \right|^2 \quad (\text{A.17})$$

Here $V(n)$ denotes the complex-valued voltage sample at a particular range bin for pulse n which consists of a real component $V_R(n)$ and an imaginary component $V_I(n)$ which are assumed to be independent, zero-mean Gaussian random variables with variance $\frac{v^2}{2}$ which is factored out in (A.17) [15]. The variance, v^2 is equal to the expected value of the received power and is exponentially distributed.

The mean of (A.16) can be written

$$E[Q'_n] = \frac{1}{N} \sum_{n=1}^N E[p(n)] = \frac{1}{N} * \frac{v^2}{2} \sum_{n=1}^N E[|V(n)|^2] = \frac{1}{N} * \frac{v^2}{2} * 2N = v^2 \quad (\text{A.18})$$

Considering (A.18) and (A.15), it is seen that the mean of the estimator Q_n is equal to the SNR. The variance of Q_n can be written as

$$\text{var}(Q_n) = E[Q_n^2] - E[Q_n]^2 \quad (\text{A.19})$$

Considering the first term of (A.19) we can recognize that starting with $E[Q_n'^2]$ is desirable and write

$$E[Q_n'^2] = \frac{1}{N^2} \left[\sum_{n=1}^N E[p(n)^2] + \sum_{a=1}^N \sum_{b=1, a \neq b}^N E[p(a)p(b)] \right] \quad (\text{A.20})$$

The first sum in (A.20) captures the second moment for an individual pulse while the second sum captures the second moment between different pulses. For the first sum, we can consider

$$\begin{aligned} E[p(n)^2] &= E[|V(n)|^2] = E[(V_R(n)^2 + V_I(n)^2)^2] \quad (\text{A.21}) \\ &= E[V_R(n)^4] + E[V_I(n)^4] + 2E[V_R(n)^2]E[V_I(n)^2] = 8 \end{aligned}$$

Where the last equality follows from the fact that the voltage samples are assumed to be i.i.d. standard normal random variables and the 4th moment of a standard normal random variable is 3. Considering the second sum in (A.20) we can write

$$\begin{aligned} E[p(a)p(b)] &= E[|V(a)||V(b)|] \quad (\text{A.22}) \\ &= E[(V_R(a)^2 + V_I(a)^2)(V_R(b)^2 + V_I(b)^2)] \\ &= E[V_R(a)^2]E[V_R(b)^2] + 2E[V_R(a)^2]E[V_I(b)^2] \\ &\quad + E[V_I(a)^2]E[V_I(b)^2] = 4 \end{aligned}$$

Where the last equality follows because of the assumption that successive pulses are independent, the assumption that the real and imaginary components are independent, and the fact that the real and imaginary components are both standard normal random variables.

Finally, it can be seen that the first sum in (A.20) has N elements going over all N pulses while the second sum has $N(N-1)$ elements since it captures the off-diagonal expectations.

Thus, we can write

$$E[Q'_n] = \frac{1}{N^2} \left[\frac{v^4}{4} * 8N + \frac{v^4}{4} * 4N(N-1) \right] = \frac{v^4(N+1)}{N} \quad (\text{A.24})$$

Moreover, because of (A.15) we can write

$$\begin{aligned} E[Q_n^2] &= E \left[\left(\frac{Q'_n}{v_n^2} - 1 \right)^2 \right] = \frac{1}{v_n^4} E[Q_n'^2] - \frac{2}{v_n^2} E[Q'_n] + 1 \\ &= \text{SNR}^2 + \frac{1}{N} \text{SNR}^2 + \frac{2}{N} \text{SNR} + \frac{1}{N} \end{aligned} \quad (\text{A.25})$$

Where the last equality comes about by substituting (A.24), (A.19), and (2.8) and simplifying. We can thus conclude

$$\text{var}(Q_n) = E[Q_n^2] - E[Q_n]^2 = \frac{1}{N} (\text{SNR}^2 + 2\text{SNR} + 1) = \frac{(1 + \text{SNR})^2}{N} \quad (\text{A.26})$$

which implies that the standard deviation of Q_n is equal to $\frac{1+\text{SNR}}{\sqrt{N}}$ and the relative standard

deviation, which is the ratio of the true standard deviation and the SNR, is given by $\frac{(1+\frac{1}{\text{SNR}})}{\sqrt{N}}$

or equivalently (2.6) after multiplication with the mean of Q_n . This formula therefore provides a means for controlling the uncertainty of measurements via both the single-pulse SNR and the number of pulses being integrated.

Appendix B: Main Function Headers

B.1: adaptiveCPRSimulator

```
%{
%> @brief Simulates a satellite-borne cloud profiling radar with
parameter and measurement adaption capabilities.
@param Ns          number of measurements to take
@param PRF_min     minimum selectable pulse repetition
                  Frequency (PRF)
@param PRF_max     maximum selectable PRF
@param fc          carrier frequency
@param CPI         duration of the (in)coherent processing
                  processing interval
@param swath       the along-track swath length of the radar,
                  in meters
@param pulseLength the time duration of the radar's transmit
                  pulse
@param gain        the radar's antenna gain
@param azimBeamWidth the radar's antenna beamwidth in azimuth
@param elevBeamWidth the radar's antenna beamwidth in elevation
@param Pt         the radar's peak transmit power
@param rangeRes   the radar's range resolution along a
                  profile, in the antenna direction
@param detectThreshdB the intensity threshold above which a target
                  is declared present, in dB (deprecated)
@param RCSBumpdB  the amount by which to artificially
                  strengthen target returns, in dB
@param adaptRange logical flag permitting PRF adaptation for
                  cloud extent in range
@param useCloudToGround logical flag governing whether or not the
                  simulation adapts in range according to the
                  height of the cloud top relative to the
                  ground or the lowest range bin
@param adaptVelocity logical flag permitting PRF adaptation for
                  cloud extent in velocity
@param Q           the state uncertainty for the Kalman filter
@param R           the measurement uncertainty for the Kalman
                  filter
@param H           the measurement matrix for the Kalman filter
@param Phi        the state transition matrix for the Kalman
                  filter
@param P           state uncertainty for the Kalman filter
@param plotIndexes logical array determining which individual
                  profiles to plot
@param personalPC  logical flag set if this sim is run on my
```

```

personal desktop
@param labPC          logical flag set if this sim is run on my
                      office desktop
@param ESL0400        logical flag set if this sim is run on
                      ESL0400
@param useCFAR         logical flag set if the constant false alarm
                      rate detector (CFAR) is used
@param useIncoherent  logical flag set if incoherent integration
                      is to be used
@param goalCNR         intensity threshold below which measurements
                      are declared failures
@param failThresh     probability threshold used with goalCNR to
                      mark failures
@retval h_meas        collected cloud column height measurements
@retval h_hat         collected cloud height estimates from the
                      Kalman filter
@retval prf_record    the PRF chosen for each measurement taken
@retval KG_range      a track of the Kalman filter gain for the
                      cloud column height
@retval KG_vel        a track of the Kalman filter gain for the
                      cloud velocity (deprecated)
@retval P_range       a track of the Kalman filter state
                      Uncertainty for the cloud column height
@retval P_vel         a track of the Kalman filter state
                      for the velocity measurement (deprecated)
@retval ranges        the list of ranges over which the
                      atmospheric profiles exist
@retval RCSHistory    a track of all the illuminated objects seen
                      by the radar
@retval rangeBinRecord a track of (currently) the lowest detected
                      target in the profile
@retval obseruationHistory the collected set of radar measurements
@retval pulsesTrack   a track of the actual number of pulses used
                      and allotted for each measurement
@retval rawProfiles   a record of the uncorrupted power profiles
                      seen by the system

%>
%}
function [h_meas, h_hat, prf_record, RuRecord, KG_range, KG_vel,
P_range, P_vel, ranges, RCSHistory, rangeBinRecord, obseruationHistory,
pulsesTrack, rawProfiles] =
adaptiveCPRSimulator(Ns, PRF_min, PRF_max, noiseFig, fc, CPI, swath,
pulseLength, gain, azimBeamWidth, elevBeamWidth, Pt, rangeRes,
detectThreshdB, RCSBumpdB, adaptRange, useCloudToGround,
adaptVelocity, Q, R, H, Phi, P, plotIndexes, personalPC, labPC, ESL0400,
useCFAR, useIncoherent, goalCNR, failThresh)

```

B.2 dopplerGEOS5MANAGERadar

```

%> @brief Simulates the behaviors of a cloud profiling radar system
%>equipped with a measurement terminator seeking to prevent
%>completion of low-CNR measurements

```

```

%>Operates on provided cloud water volume data to produce radar
observations approximating those produced
%>by the CloudSat radar system. Also implements ambiguous range
folding and forces the output vector to zero for the folded regions.
%> @param PRF          the desired PRF in Hertz
%> @param fc           the desired radar carrier frequency in Hertz
%> @param pulseLength  the desired pulse length in seconds
%> @param azimBeamWidth the desired radar antenna beamwidth in azimuth
                        in degrees
%> @param elevBeamWidth the desired radar antenna beamwidth in
                        Elevation in degrees
%> @param gain         the desired radar antenna gain over the 3dB
                        beamwidth, in dBi
%> @param Pt           the desired radar peak transmit power, in
                        watts
%> @param rangeRes     the desired range resolution of the radar, in
                        meters
%> @param waterVol     an input vector of cloud water content values
                        converted to RCS in m2/m3
%> @param heights     an input vector of heights indicating the
                        altitude in meters from sea level of each
                        cell in the range profile
%> @param sceneStart   a switch instructing the radar system to begin
                        its unambiguous range swath at a certain
                        height
%> @param threshold    a fixed intensity threshold used for simple
                        target detection
%> @param goalCNR      a specified measurement quality level below
                        which termination occurs
%> @param failThresh   a specified probability threshold placing a
                        bound on the likelihood of measurements being
                        low quality
%> @param cdf_arr      a table used to calculate the chi-squared CDF
%> @param ygrid1       a table used to interpolate inputs to the
                        chi-squared CDF
%> @retval surface     a matrix describing a range profile or full
                        range-Doppler surface containing the target
                        returns
%> @retval ranges      a vector containing the range to each cell
%> @retval targets     a vector containing the BoundingBox
                        measurements for all targets in surface
%> @retval fcdfmodTrack a vector containing the running probability of
                        failure as pulses accumulate
%> @retval pulsesTrack a vector containing the number of pulses used
                        in each integration
%> @retval rawProfile  an array containing the observed radar range
%>                    profile uncorrupted by noise
%> @retval Pint1       a vector containing the estimated power
                        calculations used to terminate measurements
%> @retval PintIdx     the location of the peak power values in the
                        column
%}

```

```

function [dopplerSurface, ranges, targets, fcdfmodTrack, pulsesTrack,
rawProfile, Pint1, PintIdx] =
dopplerGEOS5MANAGERadar( noiseFig, PRF, fc, pulseLength, azimBeamWidth,
elevBeamWidth, gain, Pt, rangeRes, waterVol, heights, sceneStart,
rangeFromRadar, threshold, CPI, lambda, useCFAR, useIncoherent,
goalCNR, failThresh, cdf_arr, ygrid1 )

```

B.3: GEOS5RadarRangeProfile

```

%{
%>@brief Simulates a single measurement produced by a CPR by
calculating all the associated range profiles at once.
%>   @param PRF      the desired PRF for the radar, in Hz
%>   @param fc       the desired carrier frequency for the
                    radar, in Hz
%>   @param pulseLength the desired pulse length for the radar, in
                    seconds
%>   @param azimBeamWidth the desired beamwidth in azimuth for the
                    radar, in radians
%>   @param elevBeamWidth the desired beamwidth in elevation for the
                    radar, in radians
%>   @param gain      the desired gain for the radar, in dBi
%>   @param Pt        the desired transmit power for the radar,
                    in watts
%>   @param rangeRes  the desired range resolution for the radar,
                    in meters
%>   @param waterVol  the water content converted to RCS
                    representing the cloud targets to be
                    viewed
%>   @param heights   the vector of altitude from mean sea level
                    Of each cell in waterVol
%>   @param sceneStart the height value at which the radar will
                    begin receiving inputs, to ensure that
                    targets are not missed as the PRF
                    is varied, in meters This should be one of
                    the heights in the actual water profile.
                    If no change in the PRF is required, input
                    -999.
%>   @retval ranges   the final set of ranges used to create the
                    profile
%>   @retval profile  the resulting radar range profile
                    illuminating the provided atmospheric
                    column
%>   @retval selection vector containing the indexes of the bins
                    from the original, full height vector
                    chosen for the final output (deprecated)
%}
function [ranges, profile, selection, fullProfile, fullRanges] =
GEOS5RadarRangeProfile(PRF, fc, pulseLength, azimBeamWidth,
elevBeamWidth, gain, Pt, rangeRes, waterVol, heights, sceneStart)

```

Bibliography

- [1] National Research Council, *Thriving on Our Changing Planet: A Decadal Strategy for Earth Observation from Space*. Washington, DC: The National Academies Press, 2018. doi: 10.17226/24938.
- [2] E. E. Clothiaux *et al.*, “Objective Determination of Cloud Heights and Radar Reflectivities Using a Combination of Active Remote Sensors at the ARM CART Sites,” *Journal of Applied Meteorology*, vol. 39, no. 5, pp. 645–665, May 2000, doi: 10.1175/1520-0450(2000)039<0645:ODOCHA>2.0.CO;2.
- [3] A. Battaglia *et al.*, “Spaceborne Cloud and Precipitation Radars: Status, Challenges, and Ways Forward,” *Reviews of Geophysics*, vol. 58, no. 3, Sep. 2020, doi: 10.1029/2019RG000686.
- [4] S. L. Durden *et al.*, “System Design and Subsystem Technology for a Future Spaceborne Cloud Radar,” *IEEE Geoscience and Remote Sensing Letters*, vol. 13, no. 4, pp. 560–564, Apr. 2016, doi: 10.1109/LGRS.2016.2525718.
- [5] S. Bony *et al.*, “Clouds, circulation and climate sensitivity,” *Nature Geoscience*, vol. 8, no. 4, pp. 261–268, 2015, doi: 10.1038/ngeo2398.
- [6] E. Peral *et al.*, “RainCube: the first ever radar measurements from a CubeSat in space,” *Journal of Applied Remote Sensing*, vol. 13, no. 3, pp. 1–13, Jun. 2019, doi: 10.1117/1.JRS.13.032504.
- [7] K. B. Cooper *et al.*, “A Compact, Low Power Consumption, and Highly Sensitive 95 GHz Doppler Radar,” *IEEE Sensors Journal*, vol. 20, no. 11, pp. 5865–5875, Jun. 2020, doi: 10.1109/JSEN.2020.2972535.
- [8] G. E. Smith, A. E. Mitchell, C. D. Ball, A. O’Brien, and J. T. Johnson, “Fully Adaptive Remote Sensing Observing System Simulation Experiments,” in *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2018, pp. 5839–5842. doi: 10.1109/IGARSS.2018.8518010.
- [9] E. Im, S. L. Durden, and S. Tanelli, “CloudSat: The Cloud Profiling Radar Mission,” in *2006 CIE International Conference on Radar*, Oct. 2006, pp. 1–4. doi: 10.1109/ICR.2006.343540.
- [10] A. J. Illingworth *et al.*, “The EarthCARE Satellite: The Next Step Forward in Global Measurements of Clouds, Aerosols, Precipitation, and Radiation,” *Bull Am Meteorol Soc*, vol. 96, no. 8, pp. 1311–1332, Aug. 2015, doi: 10.1175/BAMS-D-12-00227.1.

- [11] E. Peral, S. Tanelli, Z. Haddad, O. Sy, G. Stephens, and E. Im, “Raincube: A proposed constellation of precipitation profiling radars in CubeSat,” in *2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2015, pp. 1261–1264. doi: 10.1109/IGARSS.2015.7326003.
- [12] M. A. Richards, J. Scheer, and W. A. Holm, *Principles of modern radar. Volume I, Basic principles*. Raleigh NC: SciTech Pub, 2010. Accessed: Aug. 25, 2017. [Online]. Available: <http://osu.worldcat.org.proxy.lib.ohio-state.edu/title/principles-of-modern-radar/oclc/468981923?ht=edition&referer=di>
- [13] R. J. Doviak and D. S. Zrnić, *Doppler radar and weather observations*, 2nd ed., D. Mineola N.Y.: Dover Publications, 2006. Accessed: Oct. 21, 2018. [Online]. Available: http://osu.worldcat.org/title/doppler-radar-and-weather-observations/oclc/67346378&referer=brief_results
- [14] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave remote sensing : active and passive*. Reading Mass.: Addison-Wesley Pub. Co., Advanced Book Program/World Science Division, 1981. Accessed: Nov. 05, 2018. [Online]. Available: <http://library.ohio-state.edu/record=b1998875>
- [15] R. Meneghini and T. Kozu, *Spaceborne Weather Radar*. Boston MA, USA: Artech House, Inc., 1990.
- [16] N. A. J. Schutgens, “Simulated Doppler Radar Observations of Inhomogeneous Clouds: Application to the EarthCARE Space Mission,” *Journal of Atmospheric and Oceanic Technology*, vol. 25, no. 1, pp. 26–42, Jan. 2008, doi: 10.1175/2007JTECHA956.1.
- [17] S. Kobayashi, H. Kumagai, and H. Kuroiwa, “A Proposal of Pulse-Pair Doppler Operation on a Spaceborne Cloud-Profiling Radar in the W Band,” *Journal of Atmospheric and Oceanic Technology*, vol. 19, no. 9, pp. 1294–1306, 2002, doi: 10.1175/1520-0426(2002)019<1294:APOPPD>2.0.CO;2.
- [18] R. Marchand, G. G. Mace, T. Ackerman, and G. Stephens, “Hydrometeor Detection Using *Cloudsat* —An Earth-Orbiting 94-GHz Cloud Radar,” *Journal of Atmospheric and Oceanic Technology*, vol. 25, no. 4, pp. 519–533, Apr. 2008, doi: 10.1175/2007JTECHA1006.1.
- [19] S. Tanelli *et al.*, “CloudSat’s Cloud Profiling Radar After Two Years in Orbit: Performance, Calibration, and Processing,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 11, pp. 3560–3573, Nov. 2008, doi: 10.1109/TGRS.2008.2002030.
- [20] D. Vane and G. L. Stephens, “The CloudSat Mission and the A-Train: A Revolutionary Approach to Observing Earth’s Atmosphere,” in *2008 IEEE Aerospace Conference*, Mar. 2008, pp. 1–5. doi: 10.1109/AERO.2008.4526260.

- [21] T. Uttal and R. A. Kropfli, “The Effect of Radar Pulse Length on Cloud Reflectivity Statistics in: Journal of Atmospheric and Oceanic Technology Volume 18 Issue 6 (2001),” *Journal of Atmospheric and Oceanic Technology*, vol. 18, no. 6, pp. 947–961, 2001, Accessed: Mar. 28, 2022. [Online]. Available: [https://journals.ametsoc.org/configurable/content/journals\\$002fatot\\$002f18\\$002f6\\$002f1520-0426_2001_018_0947_teorpl_2_0_co_2.xml?t:ac=journals%24002fatot%24002f18%24002f6%24002f1520-0426_2001_018_0947_teorpl_2_0_co_2.xml](https://journals.ametsoc.org/configurable/content/journals$002fatot$002f18$002f6$002f1520-0426_2001_018_0947_teorpl_2_0_co_2.xml?t:ac=journals%24002fatot%24002f18%24002f6%24002f1520-0426_2001_018_0947_teorpl_2_0_co_2.xml)
- [22] Y. Hagihara, H. Okamoto, and R. Yoshida, “Development of a combined CloudSat-CALIPSO cloud mask to show global cloud distribution,” *Journal of Geophysical Research: Atmospheres*, vol. 115, no. D4, Feb. 2010, doi: <https://doi.org/10.1029/2009JD012344>.
- [23] T. Wang, E. J. Fetzer, S. Wong, B. H. Kahn, and Q. Yue, “Validation of MODIS cloud mask and multilayer flag using CloudSat-CALIPSO cloud profiles and a cross-reference of their cloud classifications,” *Journal of Geophysical Research: Atmospheres*, vol. 121, no. 19, pp. 11,620–11,625, Oct. 2016, doi: <https://doi.org/10.1002/2016JD025239>.
- [24] M. A. Chan and J. C. Comiso, “Arctic Cloud Characteristics as Derived from MODIS, CALIPSO, and CloudSat,” *Journal of Climate*, vol. 26, no. 10, pp. 3285–3306, 2013, doi: 10.1175/JCLI-D-12-00204.1.
- [25] A. Behrangi, T. Kubar, and B. Lambrigtsen, “Phenomenological Description of Tropical Clouds Using CloudSat Cloud Classification,” *Monthly Weather Review*, vol. 140, no. 10, pp. 3235–3249, 2012, doi: 10.1175/MWR-D-11-00247.1.
- [26] Y. Liu, J. R. Key, S. A. Ackerman, G. G. Mace, and Q. Zhang, “Arctic cloud macrophysical characteristics from CloudSat and CALIPSO,” *Remote Sensing of Environment*, vol. 124, pp. 159–173, 2012, doi: <https://doi.org/10.1016/j.rse.2012.05.006>.
- [27] J. M. Haynes and G. L. Stephens, “Tropical oceanic cloudiness and the incidence of precipitation: Early results from CloudSat,” *Geophysical Research Letters*, vol. 34, no. 9, May 2007, doi: <https://doi.org/10.1029/2007GL029335>.
- [28] Y. Huang, S. T. Siems, M. J. Manton, L. B. Hande, and J. M. Haynes, “The Structure of Low-Altitude Clouds over the Southern Ocean as Seen by CloudSat,” *Journal of Climate*, vol. 25, no. 7, pp. 2535–2546, 2012, doi: 10.1175/JCLI-D-11-00131.1.
- [29] S. di Michele *et al.*, “Interpreting an evaluation of the ECMWF global model with CloudSat observations: ambiguities due to radar reflectivity forward operator uncertainties,” *Quarterly Journal of the*

- Royal Meteorological Society*, vol. 138, no. 669, pp. 2047–2065, Oct. 2012, doi: <https://doi.org/10.1002/qj.1936>.
- [30] A. Protat *et al.*, “Assessment of Cloudsat Reflectivity Measurements and Ice Cloud Properties Using Ground-Based and Airborne Cloud Radar Observations,” *Journal of Atmospheric and Oceanic Technology*, vol. 26, no. 9, pp. 1717–1741, Sep. 2009, doi: 10.1175/2009JTECHA1246.1.
- [31] “Data Access | CloudSat DPC,” 2020. <http://www.cloudsat.cira.colostate.edu/order-data> (accessed Nov. 03, 2020).
- [32] H. Nakatsuka *et al.*, “Design and development status of the EarthCARE Cloud Profiling Radar,” in *2012 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2012, pp. 2415–2418. doi: 10.1109/IGARSS.2012.6351004.
- [33] A. Lefebvre, A. Heliere, K. Wallace, H. Nakatsuka, and E. Tomita, “The ESA EarthCARE mission development status,” in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2017, pp. 4234–4237. doi: 10.1109/IGARSS.2017.8127937.
- [34] A. Lefebvre *et al.*, “Status of ESA’s Earthcare Mission Preparation,” in *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, Sep. 2020, pp. 3459–3462. doi: 10.1109/IGARSS39084.2020.9324497.
- [35] E. E. Clothiaux *et al.*, “An Evaluation of a 94-GHz Radar for Remote Sensing of Cloud Properties,” *Journal of Atmospheric and Oceanic Technology*, vol. 12, no. 2, pp. 201–229, 1995, doi: 10.1175/1520-0426(1995)012<0201:AEOAGR>2.0.CO;2.
- [36] G. B. Gustafson and R. P. D’entremont, “Cloud Detection and Property Retrieval Across the Day/Night Terminator,” 2006. Accessed: Mar. 28, 2022. [Online]. Available: www.aer.com/cloud
- [37] T. Qing, P. Minnis, P. W. Heck, and R. Palikonda, “Improvements in Near-Terminator and Nocturnal Cloud Masks using Satellite Imager Data over the Atmospheric Radiation Measurement Sites,” 2005. Accessed: Mar. 28, 2022. [Online]. Available: https://www.researchgate.net/publication/236495235_Improvements_in_Near-Terminator_and_Nocturnal_Cloud_Masks_using_Satellite_Imager_Data_over_the_Atmospheric_Radiation_Measurement_Sites
- [38] Z. Shao, Y. Pan, C. Diao, and J. Cai, “Cloud Detection in Remote Sensing Images Based on Multiscale Features-Convolutional Neural Network,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 6, pp. 4062–4076, 2019, doi: 10.1109/TGRS.2018.2889677.

- [39] D. S. Zrnić, “Simulation of Weatherlike Doppler Spectra and Signals,” *Journal of Applied Meteorology*, vol. 14, no. 4, pp. 619–620, Jun. 1975, doi: 10.1175/1520-0450(1975)014<0619:SOWDSA>2.0.CO;2.
- [40] S. Tanelli *et al.*, “Radar concepts for the next generation of spaceborne observations of cloud and precipitation processes,” in *2018 IEEE Radar Conference (RadarConf18)*, Apr. 2018, pp. 1245–1249. doi: 10.1109/RADAR.2018.8378741.
- [41] E. Peral *et al.*, “Radar Technologies for Earth Remote Sensing From CubeSat Platforms,” *Proceedings of the IEEE*, vol. 106, no. 3, pp. 404–418, Mar. 2018, doi: 10.1109/JPROC.2018.2793179.
- [42] J. T. Johnson *et al.*, “The CubeSat Radiometer Radio Frequency Interference Technology Validation (CubeRRT) mission,” in *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2016, pp. 299–301. doi: 10.1109/IGARSS.2016.7729070.
- [43] W. Swartz *et al.*, “RAVAN CubeSat Results: Technologies and Science Demonstrated On Orbit,” Aug. 2017. Accessed: Mar. 30, 2022. [Online]. Available: <https://digitalcommons.usu.edu/smallsat/2017/all2017/142>
- [44] D. L. Wu *et al.*, “IceCube: spaceflight demonstration of 883-GHz cloud radiometer for future science,” Aug. 2019. doi: 10.1117/12.2530589.
- [45] E. Peral *et al.*, “Raincube – A New Paradigm to Observe Weather Processes,” in *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2019, pp. 4978–4981. doi: 10.1109/IGARSS.2019.8898687.
- [46] O. O. Sy *et al.*, “Scientific Products From the First Radar in a CubeSat (RainCube): Deconvolution, Cross-Validation, and Retrievals,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–20, 2022, doi: 10.1109/TGRS.2021.3073990.
- [47] S. Gao *et al.*, “Antennas for Modern Small Satellites,” *IEEE Antennas and Propagation Magazine*, vol. 51, no. 4, pp. 40–56, Aug. 2009, doi: 10.1109/MAP.2009.5338683.
- [48] N. Chahat, J. Sauder, M. Mitchell, N. Beidleman, and G. Freebury, “One-Meter Deployable Mesh Reflector for Deep Space Network Telecommunication at X- and Ka-band,” in *2019 13th European Conference on Antennas and Propagation (EuCAP)*, 2019, pp. 1–4.
- [49] Y. Rahmat-Samii, V. Manohar, J. M. Kovitz, R. E. Hodges, G. Freebury, and E. Peral, “Development of Highly Constrained 1 m Ka-Band Mesh Deployable Offset Reflector Antenna for Next Generation CubeSat Radars,” *IEEE Transactions on Antennas and Propagation*, vol. 67, no. 10, pp. 6254–6266, Oct. 2019, doi: 10.1109/TAP.2019.2920223.

- [50] N. Chahat *et al.*, “Advanced CubeSat Antennas for Deep Space and Earth Science Missions: A review,” *IEEE Antennas and Propagation Magazine*, vol. 61, no. 5, pp. 37–46, Oct. 2019, doi: 10.1109/MAP.2019.2932608.
- [51] E. Peral *et al.*, “The Radar-in-a-Cubesat (RAINCUBE) and Measurement Results,” in *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2018, vol. 2018-July, pp. 6297–6300. doi: 10.1109/IGARSS.2018.8519194.
- [52] J. R. Guerci, *Cognitive Radar: The Knowledge-Aided Fully Adaptive Approach*. Norwood, MA: Artech House, Inc., 2010.
- [53] M. Fatemi and S. Haykin, “Cognitive control: Theory and application,” *IEEE Access*, vol. 2, pp. 698–710, 2014, doi: 10.1109/ACCESS.2014.2332333.
- [54] K. L. Bell, C. J. Baker, G. E. Smith, J. T. Johnson, and M. Rangaswamy, “Fully adaptive radar for target tracking part I: Single target tracking,” in *2014 IEEE Radar Conference*, May 2014, pp. 0303–0308. doi: 10.1109/RADAR.2014.6875604.
- [55] K. L. Bell, C. J. Baker, G. E. Smith, J. T. Johnson, and M. Rangaswamy, “Fully adaptive radar for target tracking part II: Target detection and track initiation,” in *2014 IEEE Radar Conference*, May 2014, pp. 0309–0314. doi: 10.1109/RADAR.2014.6875605.
- [56] A. E. Mitchell, “Advancements and Applications of the Fully Adaptive Radar Framework,” Unpublished Doctoral Dissertation, The Ohio State University, 2018.
- [57] J. DeLong, M. A. Shattal, A. O’Brien, C. D. Ball, J. T. Johnson, and G. E. Smith, “Fully Adaptive Cloud Profiling Radar Simulation,” in *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2019, pp. 10087–10090. doi: 10.1109/IGARSS.2019.8898349.
- [58] K. L. Bell, C. J. Baker, G. E. Smith, J. T. Johnson, and M. Rangaswamy, “Cognitive Radar Framework for Target Detection and Tracking,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 8, pp. 1427–1439, Dec. 2015, doi: 10.1109/JSTSP.2015.2465304.
- [59] M. Rangaswamy, A. Jones, and G. Smith, “Recent trends and findings in cognitive radar,” in *2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, Dec. 2015, pp. 1–4. doi: 10.1109/CAMSAP.2015.7465294.
- [60] G. E. Smith *et al.*, “Experiments with cognitive radar,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 12, pp. 34–46, Dec. 2016, doi: 10.1109/MAES.2016.150215.
- [61] J. DeLong, C. Ball, J. T. Johnson, and A. O’Brien, “Using The Simulation Toolset for Adaptive Remote Sensing (STARS) TO

- Explore A Tradespace of New Observing Strategies,” in *AGU Fall Meeting*, Dec. 2020, vol. 2020. [Online]. Available: <https://ui.adsabs.harvard.edu/abs/2020AGUFMIN019..09D/>
- [62] W. Putman, A. M. da Silva, L. E. Ott, and A. Darmanov, “Model Configuration for the 7-km GEOS-5 Nature Run, Ganymed Release (Non-hydrostatic 7km Global Mesoscale Simulation). GMAO Office Note No. 5.” GMAO, p. 18, 2014. [Online]. Available: <https://gmao.gsfc.nasa.gov/pubs/docs/Putman727.pdf>
- [63] A. M. Society, “Hypsometric Formula,” *Glossary of Meteorology*, 2015. https://glossary.ametsoc.org/wiki/Hypsometric_equation
- [64] D. Atlas, “Advances in Radar Meteorology,” *Advances in Geophysics*, vol. 10, pp. 317–478, Jan. 1964, doi: 10.1016/S0065-2687(08)60009-6.
- [65] J. M. Haynes *et al.*, “A Multipurpose Radar Simulation Package: QuickBeam,” *Bull Am Meteorol Soc*, vol. 88, no. 11, pp. 1723–1728, Nov. 2007, doi: 10.1175/BAMS-88-11-1723.
- [66] E. Im, C. Wu, and S. L. Durden, “Cloud profiling radar for the CloudSat mission,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 20, no. 10, pp. 15–18, Oct. 2005, doi: 10.1109/MAES.2005.1581095.
- [67] R. Wood, P. R. Field, R. Wood, and P. R. Field, “The Distribution of Cloud Horizontal Sizes,” *Journal of Climate*, vol. 24, no. 18, pp. 4800–4816, Sep. 2011, doi: 10.1175/2011JCLI4056.1.
- [68] H. V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer-Verlag, 2010.
- [69] S. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Boston, MA: Artech House, Inc., 1999.