

DISSERTATION

APPLICATION OF ADVANCED DATA ASSIMILATION TECHNIQUES
TO THE STUDY OF CLOUD AND PRECIPITATION FEEDBACKS
IN THE TROPICAL CLIMATE SYSTEM

Submitted by

Derek J. Posselt

Department of Atmospheric Science

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2006

UMI Number: 3233360

INFORMATION TO USERS

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

UMI[®]

UMI Microform 3233360

Copyright 2006 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.


ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

COLORADO STATE UNIVERSITY

June 22, 2006

WE HEREBY RECOMMEND THAT THE DISSERTATION PREPARED UNDER OUR SUPERVISION BY DEREK J. POSSELT ENTITLED APPLICATION OF ADVANCED DATA ASSIMILATION TECHNIQUES TO THE STUDY OF CLOUD AND PRECIPITATION FEEDBACKS IN THE TROPICAL CLIMATE SYSTEM BE ACCEPTED AS FULFILLING IN PART REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY.


Committee on Graduate Work




Dr. William R. Cotton




Dr. Richard H. Johnson



Dr. Richard A. Davis



Advisor: Dr. Graeme L. Stephens



Department Head/Director: Dr. Steven A. Rutledge

ABSTRACT OF DISSERTATION

APPLICATION OF ADVANCED DATA ASSIMILATION TECHNIQUES TO THE STUDY OF CLOUD AND PRECIPITATION FEEDBACKS IN THE TROPICAL CLIMATE SYSTEM

The research documented in this study centers around two topics: examination of the response of precipitating cloud systems to changes in the tropical climate system, and assimilation of cloud and precipitation information from remote-sensing platforms on cloud resolving scales. The motivation for this work proceeds from the following outstanding problems:

1. Use of models to study the response of clouds to perturbations in the climate system is hampered by uncertainties in cloud microphysical parameterizations.
2. Though there is an ever-growing set of available observations, cloud and precipitation assimilation remains a difficult problem, particularly in the tropics.
3. Though it is widely acknowledged that cloud and precipitation processes play a key role in regulating the Earth's response to surface warming, the response of the tropical hydrologic cycle to climate perturbations remains largely unknown.

The three above issues are addressed in a two part study. First, uncertainties in the NASA Goddard Cumulus Ensemble (GCE) model's cloud microphysical scheme are assessed and constrained using Bayesian estimation (data assimilation) methods. Because

cloud microphysical parameters are highly nonlinearly related to the cloudy state, traditional data assimilation methods, which are based on a quasi-linear Gaussian-statistics estimation framework are not suitable. Instead, a Markov chain Monte Carlo (MCMC) algorithm, which makes no assumptions on the character of the state space or the linearity of the parameter-state relationship, is used to build a sample of the joint conditional parameter probability space, and to characterize the sensitivity of simulation results to changes in cloud microphysical parameters. It is found that simulations are most sensitive to the parameters that define the rain drop size distribution and the graupel and snow fall velocities, and it is concluded that TRMM retrieved mean precipitation rates, cloud properties, and radiative fluxes and heating rates can be effectively used to constrain these parameters to values consistent with convection in the tropics.

In the second component of this study, the parameter values returned by MCMC are implemented in the GCE model, and the resulting observation-constrained model is used to investigate the response of cloud and precipitation processes to changes in sea surface temperature (SST). Two different experiments are performed--one in which base-state winds are set equal to zero, and one in which base-state winds are set equal to mean flow conditions typical of the active phase of the monsoon over the equatorial western Pacific. In each set of experiments, the SST is fixed at values of 295 K, 300 K, and 305 K, and the characteristics of the resulting equilibrium state are examined for evidence of consistent changes to the tropical hydrologic cycle with increasing SST. It is found that changes to SST alter the mode and scale of the convective organization for both sets of

experiments, and that the manner in which convection organizes on large scales is fundamentally different between sheared and unsheared simulations. In addition, with increasing SST, the spatial and temporal mean cloud fraction decreases, the low cloud amount increases, the outgoing longwave radiation and column integrated latent heat release increase, and the shortwave flux at the top of the atmosphere decreases. At the same time, consistent with an increase in latent heating, the total surface precipitation increases, as does the precipitation intensity and efficiency. In spite of the complex interaction between the large scale circulation, convection, and cloud microphysical processes, increases in convective precipitation efficiency are found to be uniformly associated with decreases in high cloud fraction and decreases in outgoing shortwave radiation.

Derek John Posselt
Atmospheric Science Department
Colorado State University
Fort Collins, CO 80523
Summer 2006

Contents

Abstract	iii
Table of Contents	vi
Acknowledgements	xxiii
1 Introduction	1
1.1 Clouds and the Tropical Climate System	1
1.2 Role of Precipitation Efficiency in Cloud-Climate Feedbacks	3
1.3 CRM Limitations	8
1.4 Assimilation of cloud information from TRMM and A-Train in a cloud resolving model	18
1.5 CRM radiative convective equilibrium climate feedback experiments .	20
1.6 Outline and Key Findings	22
2 MCMC Proof of Concept: Retrieval of Cloud Properties from MODIS	24
2.1 Introduction: MCMC vs Variational Methods in the Retrieval Context	24
2.2 Radiative Transfer Model and Variational Retrieval	28
2.2.1 Description of the Radiant Radiative Transfer Model	28
2.2.2 Variational Retrieval Method	31
2.3 MCMC Sampling Procedure	32
2.4 Results	36

2.4.1	Control Case	38
2.4.2	Double Error Case	42
2.5	Quantitative Assessment and Improvement of Variational Retrieval .	46
2.5.1	Information Content and Assessment of Variational Retrieval .	46
2.5.2	Use of MCMC Results to Modify and Improve the Variational Retrieval	49
2.6	Discussion	53
2.7	Conclusions	55
3	Use of MCMC to Estimate Sensitivity of GCE Simulations to Changes in Microphysical Parameters	57
3.1	Introduction	57
3.2	GCE Model Description and Setup	62
3.3	Observations	65
3.4	MCMC Setup and Convergence Diagnostics	69
3.5	Results	73
3.5.1	Parameter PDFs	73
3.5.2	Assessment of the utility of individual observations for assimilation	80
3.5.3	Assessment of Observation/Parameter Relationships	84
3.6	Discussion	93
3.7	Conclusions	101
4	Assimilation of TRMM Multisensor Retrievals to Constrain GCE Microphysical Parameters	104
4.1	Introduction	104
4.2	3-9 June 1998 TRMM Retrievals	107
4.3	Variational Parameter Estimation	113

4.4	Parameter Estimation Using MCMC	124
4.5	Discussion	135
4.6	Conclusions	138
5	Large-Domain Goddard Cumulus Ensemble Cloud Feedback Experiments	140
5.1	Introduction	140
5.2	Precipitation Efficiency	145
5.3	Experiment Setup	148
5.4	GCE Results	150
5.4.1	Convergence to Radiative Convective Equilibrium	150
5.4.2	Large Scale Convective Organization	154
5.4.3	Response of Clouds and Precipitation to Changes in SST . . .	160
5.5	Discussion	165
5.6	Conclusions	172
6	Summary and Conclusions	174
6.1	Summary of Key Results and General Conclusions	174
6.2	Implications of Current Work and Suggestions for Future Research . .	178
A	GCE Cloud Microphysics Scheme	181
A.1	Introduction	181
A.2	Basic Formulation and Key Assumptions	182
A.2.1	Particle Size Distributions	182
A.2.2	Fall Velocity Formulation	184
A.3	Documentation of GCE Cloud Microphysics Code	186
A.3.1	(PSAUT) Autoconversion of ice to snow	186
A.3.2	(PSACI) Accretion of ice onto snow	187
A.3.3	(PSACW) Accretion of cloud onto snow (riming)	188

A.3.4	(PWACS) Collection of snow by cloud	188
A.3.5	(PRACI) Accretion of ice onto rain	189
A.3.6	(PIACR) Accretion of rain or graupel onto ice	189
A.3.7	(PRAUT) Autoconversion of cloud to rain	190
A.3.8	(PRACW) Accretion of cloud onto rain	191
A.3.9	(PSFW) Bergeron processes for snow with cloud	191
A.3.10	(PSFI) Bergeron processes for snow with ice	192
A.3.11	(PGACS) Accretion of snow onto graupel (dry and wet)	192
A.3.12	(PGACW) Accretion of cloud onto graupel	193
A.3.13	(PGACI) Accretion of ice onto graupel	193
A.3.14	(PGACR) Accretion of rain onto graupel	193
A.3.15	(PRACS) Accretion of snow onto rain	194
A.3.16	(PSACR) Accretion of rain onto snow	194
A.3.17	(PGFR) Freezing of rain to form graupel	195
A.3.18	(PSMLT) Melting of snow	195
A.3.19	(PGMLT) Melting of graupel to form rain	196
A.3.20	(PIHOM) Homogeneous freezing of cloud water into ice	196
A.3.21	(PIDW) Deposition growth of cloud water to ice	196
A.3.22	(PIMLT) Melting of ice to form cloud water	197
A.3.23	(PINT) Initiation of ice from nuclei	197
A.3.24	Cloud water and ice deposition/evaporation	198
A.3.25	(PSDEP) Vapor deposition on snow	200
A.3.26	(PGDEP) Vapor deposition on graupel	201
A.3.27	(ERN) Evaporation of rain	201
A.3.28	(PMLTG) Evaporation of melting graupel	202
A.3.29	(PMLTS) Evaporation of melting snow	202
A.4	Discussion	202

B	Effect of Grid Resolution on 2D CRM Results	206
B.1	Introduction	206
B.2	Results	209
B.3	Discussion	212
B.4	Conclusions	213
C	MCMC: Convergence Diagnostics, Sample Independence, and Adaptive Sampling	215
C.1	Theoretical Foundations of MCMC and Construction of an MCMC Algorithm	215
C.2	MCMC Diagnostics	220
C.2.1	Effective Sample Size	220
C.2.2	Assessment of Convergence to Sampling a Stationary Distribution	221
C.3	Use of Multiple Chains	225
C.4	Gaussian Proposal Distribution	227
C.5	Advanced Topics: Modifications to the MCMC Algorithm	231
D	Application of MCMC to CRM Simulation of a Single Unforced Convective Squall Line	235
D.1	Introduction	235
D.2	GCE Model Setup	237
D.3	Results	239
D.3.1	Parameter PDFs	239
D.3.2	Examination of the conditional PDFs of observations $P(\mathbf{y} \mathbf{x})$.	245
D.3.3	Assessment of Observation/Parameter Relationships	249
D.4	Discussion	253
D.5	Conclusions	254
	References	257

List of Tables

1.1	Set of GCE cloud microphysical parameters that determine the ice crystal habit and liquid and ice particle size distribution. Both control and perturbed values are listed.	10
1.2	GCE output for control and perturbed cloud microphysical parameters.	11
3.1	GCE cloud microphysical parameters used in the MCMC algorithm, along with truth values for the simulated observation experiment and parameter ranges.	60
3.2	Observations used in both the MCMC parameter sensitivity experiment and in later assimilation experiments, along with their units and error estimates.	67
3.3	Evidence of convergence of the MCMC algorithm to sampling a stationary distribution. See text for details.	70
3.4	List of observations vs. mean states for the top 0.1 % of all likelihood values.	77
3.5	List of observation error standard deviations vs. the standard deviation of each simulated state variable.	82
3.6	List of the parameter/observation pairs with the most significant linear relationships, as represented in their linear correlation and signal to noise ratio (see text for description).	86
3.7	List of TRMM retrieved state variables with revised error estimates based on anticipated future sensors.	94

3.8	List of observations vs. mean states for the top 0.1 % of all likelihood values for the experiment with improved error estimates.	100
4.1	List of TRMM observations compared with statistics output from a GCE simulation run with default cloud microphysical parameter values.	110
4.2	List of parameter and state values for MCMC database subsets along with the Jacobian computed from equation 4.9 and from linear regression.	121
4.3	List of TRMM retrieved quantities vs. the GCE state that corresponds to the highest 0.1% of likelihood values. Results are plotted for the case in which both clear-sky and all-sky radiative heating rates are used in the estimation.	123
4.4	Minimum and maximum values of GCE state variables as compared with TRMM retrieved observations. State variables whose range does not include the corresponding TRMM retrieval are marked in red. . .	126
4.5	List of TRMM retrieved quantities vs. the mean GCE state that corresponds to the highest 0.1% of likelihood values. Results are plotted for the case in which clear-sky radiative heating rates are omitted from the estimation.	127
4.6	List of default GCE cloud microphysical parameters alongside the values returned by the MCMC algorithm and their standard deviations.	134
5.1	List of 20-day averaged cloud and radiation fields for the set of simulations with zero base-state wind.	163
5.2	List of 20-day averaged cloud and radiation fields for the set of simulations with base-state winds taken from averages over the TOGA COARE IFA.	166

5.3	Convective vs. stratiform fraction for precipitating systems in the last 20 days of each RCE simulation as computed using the GCE convective-stratiform partitioning method (see text for details). . . .	169
B.1	Comparison of output from MMF-configuration simulations using control vs. perturbed parameter values (Table 1.1). Results of simulations using 4 km, 2 km, and 500 m horizontal grid spacing are shown. . .	212
C.1	Description of the steps involved in computing each iteration in the MCMC algorithm.	219
D.1	List of observations vs. mean states for the top 0.1 % of all likelihood values.	244
D.2	List of the parameter/observation pairs with the most significant linear relationships, as represented in their linear correlation and signal to noise ratio (see text for description).	251

List of Figures

1.1	Conceptual depiction of the operation of the (a) thermostat and (b) infrared iris cloud feedback hypotheses.	4
1.2	Hovmoller plot of ice water path ($\text{g}\cdot\text{m}^{-2}$) for (a) control and (b) perturbed microphysics simulations. Time in hours is plotted on the ordinate axis, while domain length in km is plotted on the abscissa. . .	12
2.1	MODIS observed reflectance for 0.6 micron (a) and 2.13 micron (b) channels. Scene was observed over the northeast Pacific ocean at approximately 2130 UTC 4 July 2001.	29
2.2	Schematic depiction of the MCMC iterative process. (a) A single MCMC iteration and (b) a portion of the MCMC random walk depicting the accept/reject sequence.	35
2.3	Example of convergence of MCMC to sampling a stationary distribution—in this case samples of the joint conditional probability of optical depth and effective radius for pixel number 19,19. Progressively greater numbers of samples are shown in (a) through (e), while the result of brute force integration in increments of 0.1 in optical depth and effective radius is shown in (f) for comparison.	37

2.4	Optical depth and effective radius for the control case from (a and b) variational retrieval, (c and d) mode of the MCMC PDF, and (e and f) mean of the MCMC PDF. Effective radius has been screened out (gray areas) where 0.64 micron reflectance is less than 0.25.	39
2.5	PDFs from MCMC retrieval of (a) optical depth and (d) effective radius for the control case. Shown at right are PDFs for two selected pixels (b) and (e) relatively large optical depth and small effective radius and (c) and (f) relatively small optical depth and large effective radius. Uniform a priori is shown in red, posterior MCMC PDF is shown in black, and Gaussian variational PDF is in dashed blue. The two pixels shown in greater detail are outlined in thick dashed green in the 5x5 matrix. The total integrated absolute difference between MCMC and Gaussian PDFs (see text for explanation) is printed in the upper right hand corner of each pixel in (a).	41
2.6	As in figure 4, but for the double error case.	43
2.7	As in figure 5, but for the double error case.	45
2.8	Comparisons of information content computed from MCMC PDFs (a and c) and variational retrievals (b and d) for optical depth (a and b) and effective radius (c and d). All results are taken from the control case.	48
2.9	Plots of retrieved optical depth from the control case (a) and the control case with log of optical depth estimated (d). Differences between PDFs from variational retrieval and MCMC are plotted for (b) control case and (e) control case with log of optical depth estimated. Optical depth information content from MCMC and variational retrieval of log-optical depth are plotted in (c) and (f), respectively.	51

2.10	Scatterplot of optical depth retrieved from Gaussian variational (red) and log-normal variational (blue) vs. optical depth retrieved from the mode of the MCMC PDF. The one-to-one line is depicted in black for reference.	52
3.1	Time series of five-day running mean precipitation rate during the SC-SMEX field campaign. Figure adapted from Johnson and Ciesielski 2005 (their figure 3).	64
3.2	Hovmoller plot of the evolution of cloud fields in the GCE simulation for “standard” values of microphysical parameters. Color shading is potential temperature and indicates approximate cloud height from blue (low) to red (high). Precipitation rate is contoured in white every 5 mm/hour starting at 5 mm/hour, and the red horizontal line marks the beginning of the “observation” period.	66
3.3	Marginal PDFs for each estimated parameter as represented in parameter histograms. The red horizontal line on each plot indicates the Uniform a priori PDF, while the vertical dashed line depicts the true parameter value.	72
3.4	Comparison of PDFs for subset of the parameters corresponding to secondary mode, full parameter set, and primary mode.	74
3.5	Joint PDFs for each pair of parameters. The color scale is the same for each plot with warmer colors representing higher values of the PDF. Each axis is labelled with its corresponding parameter.	76
3.6	Scatter plots of the largest 0.1% of likelihood values for each parameter.	79

3.7	Marginal PDFs of the set of observations used in the MCMC experiment. In each plot, the range of values of each observed variable is depicted on the abscissa, while the normalized number of counts is displayed on the ordinate. The plots are subdivided into (a) cloud and precipitation observations, (b) long and shortwave cloud radiative forcings, (c) three-layer longwave cooling rates, and (d) three-layer shortwave heating rates.	81
3.8	Relationship between a_g and observations.	87
3.9	Relationship between N_{0r} and observations.	88
3.10	Relationship between b_s and observations.	89
3.11	Relationship between a_s and observations.	90
3.12	Relationship between ρ_g and observations.	91
3.13	Relationship between N_{0g} and observations.	92
3.14	Marginal PDFs for each estimated parameter, as represented in parameter histograms, obtained from the experiment using improved error estimates. As in figure 3.3, the red horizontal line on each plot indicates the Uniform a priori PDF, while the vertical dashed line depicts the true parameter value.	96
3.15	Scatter plots of the largest 0.1% of likelihood values for each parameter for the experiment with improved error estimates.	99
4.1	Depiction of the geographical region involved in the SCSMEX field campaign. The retrievals used in this experiment were derived over the NESR region.	105
4.2	Histograms for each of the TRMM retrieved quantities used in the assimilation. In each plot, the range of values of each observation is plotted on the abscissa and the number of counts in each bin is on the ordinate.	112

4.3	Depiction of isolines of the objective function Φ in a two-parameter state space. The gradient of the cost function is indicated, as is the minimum.	115
4.4	As in figure 3.3, but for marginal PDFs obtained from use of MCMC with TRMM retrievals.	129
4.5	As in figure 3.5, but for joint PDFs of each parameter pair obtained from use of MCMC with TRMM retrievals.	130
4.6	Scatter plots of the value of each estimated parameter for the top 0.1% of likelihood values. Total likelihood values range from 0.0 to 1.0, and were computed excluding the three-layer shortwave all-sky radiative heating rates.	133
5.1	Skew-T/log-P plot of the sounding used to initialize the RCE experiments. Note that the wind profile at right applies only to the cases with shear.	149
5.2	Illustration of convergence of simulations in each experiment to radiative convective equilibrium. Simulations run with zero-mean background flow are depicted at left, while simulations run with TOGA-COARE shear profiles are on the right. Shown are domain means of various cloudy state variables as well as short and longwave radiative fluxes at the top of the atmosphere. In each plot, the blue line corresponds to simulations run with constant SST of 295 K, black to simulations run with SST of 300 K, and red to simulations run with SST of 305 K.	151

5.3	Illustration of convergence of simulations in each experiment to radiative convective equilibrium. Simulations run with zero-mean background flow are depicted at left, while simulations run with TOGA-COARE shear profiles are on the right. Shown are timeseries of domain mean microphysical contributions to the latent heating budget. In each plot, the blue line corresponds to simulations run with constant SST of 295 K, black to simulations run with SST of 300 K, and red to simulations run with SST of 305 K.	153
5.4	Hovmoller plots of perturbation precipitable water vapor for each of the simulations in this study. Experiments with zero mean wind are shown in (a) - (c), while results from simulations with shear are depicted in (d) - (f).	155
5.5	Hovmoller plots of total column integrated condensate (liquid + ice) for each of the simulations in this study. As in Fig. 5.4, simulations run with zero mean wind are depicted in (a)-(c), while experiments run with shear are contained in (d)-(f).	156
5.6	Five-day mean vertical and horizontal wind over the whole simulation domain for each simulation conducted in this study.	158
5.7	Twenty-day time mean and full domain mean profiles of each condensate species in the GCE cloud microphysical scheme. In each plot, the blue line corresponds to simulations run with constant SST of 295 K, black to simulations run with SST of 300 K, and red to simulations run with SST of 305 K.	161
5.8	Vertical profiles of mean vertical velocity and latent heat release for simulations without and with shear. In each plot, the blue line corresponds to SST of 295 K, black to SST of 300 K, and red to SST of 305 K.	168

B.1	Comparison of the time evolution of 2D MMF-configuration GCE simulations with (a) 4 km, (b) 2 km, and (c) 500 m grid spacing. Filled contours represent total cloud condensate mass mixing ratio, while the open red contours depict vertical motion every $1 \text{ m} \cdot \text{s}^{-1}$ starting at $1 \text{ m} \cdot \text{s}^{-1}$. Simulation time is plotted in the upper right hand corner of each plot.	210
C.1	Depiction of the timeseries of each parameter in the first chain of the GCE simulated observations experiment documented in chapter 3. . .	222
C.2	Running computation of the moments of N_{0r} from chain 1 of the GCE simulated observations experiment documented in chapter 3.	224
C.3	Plots of the R-statistic for successively greater numbers of samples (see text for derivation and explanation) for each of the 9 microphysical parameters estimated in chapters 3 and 4. The color scheme is explained in the legend at top right.	228
C.4	Illustration of the departure from Uniform for various values of the (a) Uniform and (b) Normal proposal distribution scale parameters. . .	232
D.1	Figure depicting the evolution of total condensate mixing ratio and upward vertical velocity for the 2D simulation with open boundaries.	240
D.2	Marginal PDFs for each estimated parameter as represented in parameter histograms. The red horizontal line on each plot indicates the Uniform a priori PDF, while the vertical dashed line depicts the true parameter value.	241
D.3	Joint PDFs for each pair of parameters. The color scale is the same for each plot with warmer colors representing higher values of the PDF. Each axis is labeled with its corresponding parameter.	243
D.4	Scatter plots of the largest 0.1% of likelihood values for each parameter.	246

D.5	Marginal PDFs of the set of observations used in the MCMC experiment. In each plot, the range of values of each observed variable is depicted on the abscissa, while the normalized number of counts is displayed on the ordinate. The true value of each observation is depicted in the vertical dashed green line. The plots are subdivided into (a) cloud and precipitation observations, (b) long and shortwave cloud radiative forcings, (c) three-layer longwave cooling rates, and (d) three-layer shortwave heating rates.	247
D.6	As in figure D.5, but for marginal PDFs of the set of observations that corresponds to the primary mode in the joint parameter PDF (see text for details).	250
D.7	PDFs of additional saved GCE state variables for the entire parameter sample. As in the previous PDF plots, the frequency of occurrence of each value is plotted on the ordinate, while the range of values of each variable is plotted on the abscissa. Each PDF is normalized to 1.0. . .	255

Acknowledgments

The past three years of graduate study have left me with an enormous debt of gratitude to family, friends, professors, fellow students, and colleagues. First and foremost is my wife, Julie, who has loved and supported me through many late work nights, and has patiently endured my frequent trips to Washington D.C. I know I can never repay the unconditional love and affection you have shown me, and I will do my best to support you just as thoroughly as you pursue your Ph. D. in the coming years! The arrival of my son, Daniel, in February of 2006 provided me with much of the motivation to finish nearly a year earlier than I had initially anticipated, and is a constant source of joy and wonder to me. My family and in-laws deserve thanks for encouraging me to continue my studies, and for supporting Julie and I as we moved far from home. I am grateful to my parents in particular for always encouraging me to pursue my own interests, for instilling in me an appreciation for higher education, and for teaching me a solid work ethic.

The research performed in the course of my degree could not have been done without the assistance of a host of people. My advisor, Dr. Graeme Stephens, was a constant fount of new ideas, and took a light-handed approach to advising that allowed me the freedom to pursue my own research interests. My Ph. D. committee members included Dr. Arthur Hou at NASA Goddard Space Flight Center, who served as my co-advisor, and provided thoughtful critique of my ideas at each stage of the work. Dr. Tomi Vukicevic at the Cooperative Institute for Research

in the Atmosphere provided invaluable feedback on data assimilation methodologies and applications, and Dr. Richard Davis of the CSU Statistics department helped me to work through the mathematical underpinnings of the assimilation techniques I used. Drs. Richard Johnson and William Cotton served as my Atmospheric Science department committee members, and provided much needed perspective on the workings of the tropical climate system and the details and pitfalls of cloud resolving modeling, respectively. In addition to my committee members, I had the privilege to take classes from several knowledgeable faculty members. Specifically, Drs. Chris Kummerow, Dave Randall, Wayne Schubert, Dave Thompson, Michael Montgomery, and Don Estep provided clear and comprehensive instruction on a wide variety of topics in the Atmospheric Sciences and Mathematics.

In addition to my faculty and professional mentors, several colleagues provided feedback and assistance during the course of the work. In particular, Dr. Tristan L'Ecuyer provided me with data to be assimilated and engaged me in lively discussions on the topics of data assimilation, optimal estimation retrievals, and the field of Atmospheric Science as a whole. Drs. Chung-Lin Shie and Wei-Kuo Tao, as well as Steve Lang at NASA Goddard Space Flight Center provided me with the Goddard Cumulus Ensemble cloud resolving model code, and patiently assisted me as I learned how to configure and run it. Mick Christi, Dr. Philip Gabriel, and Kyle Leesman helpfully fielded questions about the Radiant radiative transfer model, while Dr. Steve Cooper assisted in estimating the errors associated with neglecting absorption in the near-infrared in the MODIS retrieval exercise. My office mates Norm Wood and Dr. Johnny Luo served as sounding boards for many of my research ideas, and patiently listened as I occasionally vented my frustrations about the progress (or lack thereof) in my research. My fellow graduate students and friends served to keep things light, and continually challenged me to perform up to my potential.

Finally, the administrative staff in the Atmospheric Science department deserves

special recognition for keeping track of my progress and helping me to wade through the requirements of the university's bureaucracy. Specifically, Jen Weingardt and Marilyn Hanson in the main department office, Sue Lini, Heather Jensen, and Amy Dykstra in the Stephens research group, and Becky Burke in the Kummerow research group have been extremely helpful to me over the past few years. It is not an exaggeration to say that I could not have finished this degree without their assistance.

This work was funded by through a joint Colorado State University-NASA Goddard Center for Earth-Atmosphere Studies (CEAS) fellowship.

Chapter 1

Introduction

1.1 Clouds and the Tropical Climate System

One of the greatest sources of uncertainty in current estimates of climate system response to external (e.g., anthropogenic) forcing is the influence of clouds. In particular, it is unknown whether feedbacks among clouds, radiation, and the large scale circulation will enhance or moderate the response of the climate system to a given perturbation. Of particular interest is the tropical climate system, in which deep convection plays a key role by linking the radiative balance, the large scale circulation, and the properties of the tropical atmosphere and ocean. Though much effort has been applied toward assessing the role of clouds in the tropical climate system, the response of cloud and precipitation processes to changes in the environment remains largely unknown.

The three key components of the tropical climate system: organized convection, radiation, and the large-scale circulation, are coupled through the microphysics and dynamics of clouds. For example, the strength of the large scale (e.g., Hadley-Walker) circulation is a function of the convective vertical mass flux, which is linked in turn to the latent heat release of condensation and freezing in clouds—a function of cloud

microphysical interactions. In turn, the radiative heating and cooling rates of the tropical atmosphere are a function of cloud optical properties, which depend on cloud mass, particle size distribution (PSD), and ice crystal habit. Though cloud microphysics and dynamics consist of a large and complicated array of processes, it can be argued that the precipitation process exerts a key influence on the links between the three components of the tropical climate system. In particular, the precipitation efficiency (PE) of convective cloud systems, defined as the ratio between the surface precipitation and the total amount of vapor condensed in-cloud, is a useful metric as it indicates the proportion of latent heating in the system relative to the total cloud amount.

The argument for the importance of precipitation efficiency follows from a zero-order consideration of deep convective processes. If we consider two convective systems with different precipitation efficiency developing in an identical environment, the system with smaller PE by definition will contain greater amounts of condensed and suspended liquid and ice condensate relative to the amount of precipitation. It follows then that this system, relative to the higher-PE system, will detrain greater amounts of condensate into the cirrus anvil, and by association, transport more water vapor from lower levels to the upper troposphere. Although it is tempting to draw further conclusions about the role of changes to the PE in affecting climate change, it is important to note that the integrated system PE says nothing about the *radiative* properties of detrained cloud. These are critically dependent on the microphysical characteristics and spatial distribution of the ice clouds themselves, and thus on the cloud microphysical processes operating within the convective system. Nevertheless, PE remains a useful diagnostic, and a discussion of its role in selected cloud-climate feedback hypotheses is presented below.

1.2 Role of Precipitation Efficiency in Cloud-Climate Feedbacks

A close examination of several leading hypothesized cloud feedback mechanisms reveals the role of the precipitation efficiency in coupling the large-scale circulation, radiation, and properties of the large-scale environment. Two key feedback hypotheses are briefly described below:

1. The *thermostat hypothesis* (Ramanathan and Collins 1991) defines a negative feedback in which **increases** in sea surface temperature lead to increased convection and subsequent development of thick stratiform cirrus anvil clouds. The resulting increase in solar albedo then leads to surface cooling and a **reduction** in sea surface temperature. Although it is not explicitly stated, the operation of the thermostat is strongly influenced by cloud microphysical and dynamic processes, as cirrus anvil properties depend on the amount of condensate transported aloft relative to that lost to precipitation (Fig. 1.1a).
2. According to the *infrared Iris hypothesis* (Lindzen et al. 2001), an **increase** in sea surface temperature (SST) leads to increased precipitation efficiency in convective clouds and subsequent decrease in anvil cirrus area as a greater percentage of condensate is precipitated out. The associated increase in area of the cloud-free region would then lead to an increase in outgoing longwave radiation (OLR) from the tropics and corresponding **decrease** in SST (Fig. 1.1b).

Though it can be argued that PE plays a key role in nearly any of the plethora of cloud feedback hypotheses that have been developed over the past few decades, these two are most closely concerned with the partition between condensate and precipitation in tropical convection.¹ It is especially interesting to note that though both thermo-

¹Further discussion of the role of the Earth's hydrologic process in cloud and climate feedbacks can be found in Randall et al. (2001) and Stephens (2005).




Figure 1.1: Conceptual depiction of the operation of the (a) thermostat and (b) infrared iris cloud feedback hypotheses.

stat and iris are negative feedback mechanisms, they operate in a radiatively opposite manner. Specifically, the thermostat assumes cooling from increased shortwave reflection, while the iris assumes cooling from increased longwave emission. Both are tied to the relative area of cloudy vs. clear air in the tropics, and it is assumptions about the precipitation efficiency of convective systems that causes the distinction between them.

Though both the iris and thermostat feedback hypotheses have been heavily debated in the literature, the specifics of the response of cloud and precipitation processes to changes in sea surface temperature have not yet been sufficiently described. To quantify the response of the tropical hydrologic cycle to climate warming, a realistic realization of the dominant microphysical processes and of the simultaneous operation of convection, radiation, and the large-scale circulation is required. Unfortunately, observations of the relevant processes are spatially and temporally limited over the tropical oceans. Because it is difficult to cleanly identify cause and effect from observations even in the case of spatially and temporally continuous observations, cloud-climate feedbacks are generally studied using a numerical model.

Although general circulation models (GCMs) are capable of simulating the interaction among large-scale circulation, radiation, and convection on a global scale, clouds and convection are very crudely parameterized in these models, making it difficult to assess the role of cloud-scale interactions. Since the feedbacks of interest are driven by processes that operate on the scale of individual convective cloud elements—scales of between a few kilometers and a few tens of kilometers—investigation of these processes must employ observations and/or models on corresponding scales. Cloud resolving models (CRMs) explicitly resolve cloud-scale dynamics (to varying degrees), include detailed representations of cloud liquid and ice microphysical interactions, have been used to simulate large-scale organization of convection under differing environmental conditions, and are typically coupled with sophisticated radiative transfer schemes.

As such, CRMs provide a natural framework for the study of microphysical effects on climate feedback mechanisms. Unfortunately, most CRM studies have, by necessity, employed very limited spatial domains and have used specified large scale (and in some cases specified radiative) forcing. This has limited the global applicability of the results.

The combined problems of limited CRM domain and unresolved sub-gridscale cloud processes in global models have recently been addressed through the development of Multiscale Modeling Frameworks (MMFs or “super-parameterization”) or Cloud Resolving Convective Parameterizations (CRCP), which embed a CRM within each grid cell of a GCM. Though the initial results of studies using CRCP are promising (Grabowski 2001, Kharoutdinov and Randall 2001, Randall et al. 2003) questions remain as to the utility of the MMF framework. There are three primary concerns: (1) a CRM embedded within a GCM grid cell with dimensions of a few hundred kilometers cannot sufficiently simulate mesoscale organization (Donner et al. 1999, Tompkins 2000) or the remote effects of subsidence associated with convectively-generated gravity waves, (2) under certain large-scale conditions convection within a small-domain CRM with cyclic boundary conditions can continually recondition the atmosphere for further convection leading to excessive precipitation (Luo and Stephens 2006), and (3) it can be demonstrated (the properties of clouds simulated in the MMF framework are highly sensitive to the orientation of the CRM within the GCM grid cell.

Ideally, climate feedbacks should be studied using a global cloud resolving model with fine-scale topography that is coupled with ocean, land-surface, chemistry, and ecosystem models. Although very short-duration simulations have been produced using a global cloud resolving model with 3 km grid spacing, a fully realistic Earth-system model on cloud resolving scales is not yet computationally feasible. For this reason, and because of the inherent problems with limited CRM domain in the MMF

framework and the necessarily crude treatment of cloud and precipitation processes in GCMs, this study employs a large-domain cloud resolving model to study the response of the tropical hydrologic cycle to changes in sea surface temperature.² Specifically, the null hypothesis is that there is no systematic change in the characteristics of the clouds and precipitation produced by tropical convection with increasing sea surface temperature. This hypothesis will be tested by running large domain cloud resolving simulations to radiative-convective equilibrium for three different fixed sea surface temperatures and examining characteristics of clouds, precipitation, and radiation at equilibrium.

The NASA Goddard Cumulus Ensemble model (GCE, Tao and Simpson 1993, Tao et al. 2003a) is a cloud-resolving model capable of being run on two or three spatial dimensions, and has been used to study both individual convective systems and the tropical climate system. In climate interaction studies, the GCE has been driven with large-scale forcing provided by observations from field programs (Johnson and Ciesielski 2002) and with analyses from the Goddard Global Modeling and Assimilation Office global model (W.-K. Tao, personal communication), and has been shown to realistically reproduce observed surface precipitation and convective heating profiles (Tao et al. 2003b, Shie et al. 2003). To avoid aliasing of large-scale convectively-coupled waves onto mesoscale convective structures (L. Pakula, personal communication), a domain length of $\approx 10,000$ km is necessary. The constraint of a large horizontal domain coupled with long (60 day) integration times restrict this study to use of two-dimensional geometry and grid spacings on the order of 2 km in the horizontal. Drawbacks of 2D geometry and relatively coarse ($\Delta x > 500$ m) grid spacing are discussed in appendix B.

²Since the global climate system tends to respond to external forcing in consistent local patterns, regardless of the pattern of forcing (Boer and Yu 2003a, 2003b), it is assumed that cloud feedbacks can be studied on local temporal and spatial scales and the results applied to an understanding of the global response.

1.3 CRM Limitations

Any study involving a CRM is limited by the fact that CRMs are generally required to parameterize the microscale interaction between different types of cloud and precipitation particles. These cloud microphysical parameterizations make simplifying assumptions about the form of the particle size distribution of ice and liquid condensate, which exerts a large influence on the radiative transfer and a potentially large (and largely unknown) influence on climate. In particular, parameters related to the assumed particle size distribution have a key effect on the details of cloud and precipitation development. These parameters influence particle fall speeds and the details of interspecies interaction within the cloud. As a result, they exert a measure of control over the partition between precipitation and suspended cloud, as well as over the amount and properties of detrained cloud. It has been shown that regardless of the model and form of the microphysical parameterization, the output of CRM simulations are nearly uniformly highly sensitive to small changes in these parameters (Fovell and Ogura 1989, Ferrier 1994, Walko et al. 1995, Ferrier et al. 1996, Meyers et al. 1997, Grabowski 1999, Grabowski 2003). It has also been shown that CRMs tend to exhibit greater internal consistency when simulating clouds whose evolution is dominated by large-scale forcing than when cloud evolution is controlled by microphysical interaction, further illustrating the uncertainty associated with CRM parameterization of microphysics.³

The microphysical parameterization currently used in the GCE consists of a single-moment bulk scheme in which two classes of liquid (non-precipitating liquid and rain) and three classes of ice (non-precipitating pristine crystals, precipitating unrimed pris-

³Another complication stems from the fact that few CRMs even parameterize, much less explicitly resolve, the effects of the distribution of aerosol on cloud microphysical evolution. Those that do (Khain et al. 1999, van den Heever et al. 2005, Khain et al. 2005) report a very large sensitivity of cloud characteristics to changes in aerosol concentrations, however, a thorough investigation of the response of precipitating cloud systems to changes in aerosol concentration and properties is beyond the scope of the current study.

tine crystals and aggregates, and graupel) are assumed to follow an exponential particle size distribution (Lin et al. 1983, Rutledge and Hobbs 1983, 1984, for a detailed description of the GCE cloud microphysical parameterization, the reader is referred to appendix A). Though this single-moment scheme has its drawbacks compared with two-moment and bin schemes,⁴ it represents a significant improvement over earlier parameterizations, which contained only liquid and ice as prognostic variables, and diagnosed the partition between suspended and precipitating condensate.

Since cloud/precipitation processes in the GCE are simulated using parameterized microphysics, prior to use of the GCE to assess the role of precipitation efficiency in cloud-climate feedbacks, it is important to first develop an understanding of the sensitivity of simulated clouds to empirical model parameters, then, if necessary, to constrain the most uncertain parameters using observations. Traditional sensitivity experiments consist of defining a control simulation, perturbing a given parameter or set of parameters, re-running the model, and comparing the results with the control. Several CRM sensitivity experiments have been conducted in this manner (Tao et al. 1995, Walko et al. 1995, Meyers et al. 1997, Grabowski et al. 1999, Grabowski 1999, Wu et al. 1999, Petch and Gray 2001, Grabowski 2003, Saleeby and Cotton 2004), and have shown that CRM simulated cloud fields are particularly sensitive to changes in the parameters that define particle size distributions. Key parameters in the GCE microphysical scheme are summarized in Table 1.2.⁵

Before proceeding with more thorough sensitivity experiments, it is illustrative to first demonstrate the degree of sensitivity of GCE simulations to changes in its

⁴Tropical convective systems are more accurately characterized by a modified gamma PSD (Heymsfield et al. 2002), and because prediction of two moments of the PSD has been shown to more realistically reproduce observed size distributions (Ferrier et al. 1995, Meyers et al. 1997), implementation of the CSU RAMS microphysical parameterization (Walko et al. 1995, Meyers et al. 1997, Cotton et al. 2003, Saleeby and Cotton 2004, van den Heever et al. 2005) is underway.

⁵Though these are certainly not the only uncertain parameters in the GCE cloud microphysical scheme, these nine parameters are those that are most directly associated with assumptions about the liquid and ice particle size distributions, as well as the crystal habit of precipitating ice species (snow and graupel).

Table 1.1: Set of GCE cloud microphysical parameters that determine the ice crystal habit and liquid and ice particle size distribution. Both control and perturbed values are listed.

Parameter	Meaning	Control	Perturbed
a_s	Snow fallspeed coefficient	0.61	0.69
b_s	Snow fallspeed exponent	0.11	0.41
a_g	Graupel fallspeed coefficient	1.5	1.1
b_g	Graupel fallspeed exponent	0.37	0.57
N_{0R}	Slope intercept for rain size distribution	0.08	-
N_{0S}	Slope intercept for snow size distribution	0.16	1.0
N_{0G}	Slope intercept for graupel size distribution	0.08	1.0
ρ_S	Density of snow	0.1	-
ρ_G	Density of graupel	0.4	0.2

cloud microphysical parameters. This is done via a traditional sensitivity experiment, in which control parameters are set equal to values typical of those used in GCE simulations of tropical convection in the published literature. In the perturbation experiment the slope intercept of snow and graupel are increased to simulate the effect of shifting the size distributions of precipitating ice toward smaller mean particle sizes. In addition, the crystal habit of ice crystals within aggregates is changed from unrimed dendrites to columns, side planes, and bullets, and the graupel shape is changed from lump to hexagonal. Fall velocity parameters are changed consistent with these assumptions according to observations (Locatelli and Hobbs 1974, Heymsfield et al. 2002). In order to effectively demonstrate simulation sensitivity, two sets of experiments are performed: one in which a single tropical squall line is simulated, and another in which a large-domain simulation is run to radiative-convective equilibrium. For details of the setup of each simulation, the reader is referred to chapter 3 and appendix D, respectively.

An examination of the results of the convective system simulations (Table 1.2) demonstrates that relatively small changes in cloud parameters have a significant

Table 1.2: GCE output for control and perturbed cloud microphysical parameters.

Output Variable	Control	Perturbed
Total Precipitation	550.4 mm	463.4 mm
Precipitation Rate	10.1 mm/h	8.6 mm/h
Cloud Top Temperature	224.0 K	224.9 K
Mean Liquid Water Path	0.125 kg/m ²	0.106 kg/m ²
Mean Ice Water Path	0.164 kg/m ²	0.201 kg/m ²
TOA Longwave Cloud Forcing	-28.9 W/m ²	-26.0 W/m ²
TOA Shortwave Cloud Forcing	73.9 W/m ²	63.5 W/m ²

effect on key aspects of the simulation. First, the total precipitation amount and precipitation rate decrease, while the mean cloud top temperature slightly increases. In addition, the mean cloud liquid water path (LWP) decreases, while the mean cloud ice water path (IWP) increases, reflecting an increase in the amount of suspended ice. The top of atmosphere longwave (infrared) and shortwave (visible) cloud radiative forcings, defined as the difference between cloudy and clear-air radiative fluxes, both decrease for the perturbed set. In the interest of brevity, a detailed diagnosis of cause and effect is not presented here—instead, we conclude that relatively small changes in cloud microphysical parameters do indeed have a significant effect on the characteristics of individual convective systems.

Hovmoller plots of ice water path for both control and perturbed microphysics runs from large-domain simulations run to radiative-convective equilibrium are depicted in figure 1.2. From this plot, it can be seen that not only does the amount of ice in the simulation increase with changes in microphysical parameters, but the horizontal length scale of convective organization changes as well. This implies a relationship not only between microphysical parameters and the bulk liquid and ice content (and associated radiative properties), but also a significant feedback to the convective and mesoscale circulations.

The sensitivity of precipitation, PE, radiative forcing, and circulation to relatively




Figure 1.2: Hovmoller plot of ice water path (g m^{-2}) for (a) control and (b) perturbed microphysics simulations. Time in hours is plotted on the ordinate axis, while domain length in km is plotted on the abscissa.

small changes in cloud microphysical parameters motivates a more thorough examination of parameter sensitivity. In theory, exhaustive perturbation studies could be used to fully explore the range of sensitivity of model results to changes in parameters, as well as assess which parameters are most important. In this type of exercise, each parameter is systematically varied across its hypothetical range of variability while the others are held constant. The number of model runs required in such an exercise scales as the number of parameter bins raised to the power of the number of parameters, and even a relatively coarse study using only 10 bins per parameter would require 10^9 simulations for the nine parameters of interest. Alternatively, an estimate of parameter sensitivity could be obtained as a linearization about a static set of parameter values $\mathbf{x}_0$. This is mathematically equivalent to computing the Jacobian matrix; the matrix of partial derivatives of model output variables with respect to each unknown parameter.

$$J(x_i, y_j) = \left. \frac{\partial y_j}{\partial x_i} \right|_{\mathbf{x}_0} \quad (1.1)$$

Although it is much more computationally feasible to compute the Jacobian than to perform exhaustive perturbation experiments, J depends critically on the assumed base state $\mathbf{x}_0$. If the parameters and state are linearly (or nearly linearly) related, this is not a problem, however, the parameters and state are *highly nonlinearly* related in the GCE cloud microphysical parameterization. Because of both the issues of nonlinearity and computational expense, past sensitivity studies have perturbed at most a few selected parameters, and have neither been able to span the range of uncertainty for the parameters of interest, nor assess which observations might be used to constrain the uncertainty in any given parameter.

The fact that both individual simulated convective systems and large domain simulations of organized convection exhibit significant sensitivity to changes in microphysical parameters motivates constraint of these parameters using observations prior

to simulation of the response of cloud and precipitation processes to surface warming. This is a two-part problem in which it is first necessary to determine which parameters exert the most control over the climate state, then to assimilate observations that are significantly related to these parameters. Estimation of the parameter values most consistent with reality is not straightforward, primarily because microphysical parameter values cannot be directly measured with current observing systems, and exhaustive perturbation is computationally intractable for even small CRM domains. What is needed is a way to search the space that consists of all possible parameter values, determining which are most important and identifying observations that can be used to constrain them. Statistical estimation methodologies make up the foundation of modern data assimilation techniques, and are well suited to this task.

Given a vector of observations $\mathbf{y}$ that is related to the vector of unknowns $\mathbf{x}$, data assimilation seeks to maximize the probability that $\mathbf{x}$ is the true state. If there is no prior information on $\mathbf{x}$, and if we have some knowledge of the probability distribution of the observation errors, then the solution can be found as the maximum likelihood of the observations conditioned on the state. Observation errors are typically assumed to be distributed Gaussian, leading to the function

$$P(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi)^{\frac{m}{2}} |\mathbf{R}|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] \right\}, \quad (1.2)$$

where $\mathbf{R}$ is the observation error covariance matrix, $\mathbf{y}$ is the vector of m observations, $\mathbf{x}$ is the vector of unknown parameters, and $F(\mathbf{x})$ is the forward model that maps the parameters into observation space. The Gaussian form is convenient in that the probability density function (PDF) is analytically easy to manipulate, and has the added benefit of being fully described by only two parameters; the mean and (co)variance. In the framework of variational estimation, the Gaussian PDF is also unique in that the solution that maximizes the posterior conditional probability is identically that which minimizes the squared differences between state and observations. If reliable

prior information on the distribution of $\mathbf{x}$ is available, or (as is commonly the case), the observations do not provide enough information to constrain the solution, Bayes' theorem can be used to seek the maximum probability $P(\mathbf{x}|\mathbf{y})$

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x}) P(\mathbf{x})}{P(\mathbf{y})} \propto P(\mathbf{y}|\mathbf{x}) P(\mathbf{x}), \quad (1.3)$$

where $P(\mathbf{x})$ is the prior distribution for the state vector, and $P(\mathbf{y})$ describes how a perturbation in the state $\mathbf{x}$ maps into a change in the observations $\mathbf{y}$. Because $P(\mathbf{y})$ is typically either unknown or too computationally expensive to evaluate (see appendix C for further discussion), $P(\mathbf{x}|\mathbf{y})$ is typically computed up to a proportionality constant. If errors in both the observations and the prior estimate are assumed to be distributed Gaussian, then Eqn. 1.2 can be written

$$P(\mathbf{x}|\mathbf{y}) \propto P(\mathbf{y}|\mathbf{x}) P(\mathbf{x}) = \left[\frac{1}{(2\pi)^{\frac{m}{2}} |\mathbf{R}|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] \right\} \right] \left[\frac{1}{(2\pi)^{\frac{n}{2}} |\mathbf{B}|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} [\mathbf{x} - \mathbf{x}_a]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_a] \right\} \right], \quad (1.4)$$

where $\mathbf{B}$ is the a priori error covariance matrix, n is the number of elements in the state vector, and $\mathbf{x}_a$ is the a priori estimate. It is commonly assumed that errors in the prior estimate and the observation errors are independent, leading to the following formulation

$$P(\mathbf{x}|\mathbf{y}) \propto \exp \left\{ -\frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] - \frac{1}{2} [\mathbf{x} - \mathbf{x}_a]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_a] \right\}. \quad (1.5)$$

The solution to 1.5 is often found via variational methods, which by necessity assume a Gaussian distribution for the observation and prior errors. It is worth mentioning that

nearly every modern atmospheric and oceanic data assimilation technique is based on (1.5); methods of solving this equation are discussed in more detail in chapters 2 and 4. The fundamental problem with variational estimation is the fact that the robustness of the solution depends directly on the validity of the Gaussian error assumption. If the true PDF exhibits significant skewness or, worse yet, multiple modes, the solution obtained from variational methods can stray far from the truth. Though information content theory can be used to obtain valuable information on the performance of a variational retrieval, the information content itself is limited by the assumption of Gaussian error statistics. In addition, specification of the error statistics for the prior estimate is extremely tricky, and remains a topic of active research.

In its most general form, data assimilation seeks to characterize the posterior conditional probability $P(\mathbf{x}|\mathbf{y})$, and ideally should not rely on specification of the form of the PDF. In the current application, what is needed is a thorough and efficient search of the probability space defined by the collective uncertainties in the set of empirical microphysics parameters. Given an *a priori* range of realistic values for each parameter as well as observations related to the model state, Markov chain Monte Carlo sampling methods (hereafter MCMC, Metropolis 1953, Tarantola 1987, Mosegaard and Tarantola 1995, Tamminen and Kyrola 2001, Tamminen 2004, Tarantola 2005) can be used to randomly explore this space, efficiently seeking regions of concentrated probability via a predetermined likelihood function that includes a statistical model of the observation errors and quantitatively measures the goodness-of-fit between observations and model state. In the current case, because observations are compared with the model state over the entire length of each GCE simulation, the sampling procedure is analogous to model bias correction methods (Dee and Da Silva 1998, Chepurin et al. 2005). However, unlike variational bias correction methods, MCMC methods are not constrained to any particular choice of a priori error probability distribution, and therefore theoretically allow the freedom to more fully explore the state

space. Perhaps the most flexible choice of prior distribution is a Uniform distribution (equal likelihood for any value of the state within given bounds), where the only constraints on the method are the user-specified range of values for each parameter to be estimated, and definition of a criterion (the *likelihood function*—typically a function of the least-squares distance between model and observations) used in determining whether the algorithm has located a region of concentrated probability.

Because the data assimilation problem is analogous to the problem of retrieving the atmospheric state from remotely-sensed measurements, and because MCMC has not yet achieved widespread use in the atmospheric science community, prior to application of MCMC to the GCE model a proof-of-concept study is performed in which MCMC is used to retrieve effective radius and optical depth from MODIS visible and near infrared reflectances. Variational (optimal) estimation has been applied to this particular retrieval problem, and the results are assumed to be well understood. Because a nonlinear radiative transfer model is used in both variational and MCMC retrievals, this exercise provides a simple, yet robust test of the MCMC algorithm. In this exercise, the usefulness of MCMC and the properties of the algorithm are demonstrated by comparing PDFs retrieved from MCMC with Gaussian PDFs obtained from a variational retrieval. It is found that MCMC not only effectively characterizes the PDF of solutions for optical depth and effective radius in each MODIS pixel, but also can be used to assess and improve a variational estimate. Chapter 2 contains a detailed description of the MODIS proof-of-concept experiment and presentation of its results and implications.

1.4 Assimilation of cloud information from TRMM and A-Train in a cloud resolving model

There are two fundamental difficulties that must be overcome when assimilating cloud information on cloud resolving scales: (1) clouds exhibit significant variability in both space and time, and (2) modern assimilation algorithms are based on linear or weakly-nonlinear Bayesian estimation, and necessarily assume small differences between the model state and the observations. Previous cloud-scale assimilation studies have succeeded to greater or lesser extent, in part because the differences between model state and observations remained small relative to the absolute value of the estimated quantities. On larger (e.g., synoptic) scales and over time periods longer than the evolution of a single cloud system, clouds typically exhibit significant spatial and temporal variation. In this case, assumptions required by discrete-space assimilation techniques are nearly ubiquitously violated, and assimilation of cloud and precipitation properties on cloud resolving scales cannot be performed in a strictly deterministic manner. Because cloud-scale processes evolve on timescales that are short compared to those of the climate system, it is the long-term spatial and temporal statistics of clouds that exert the most important influence on the large-scale radiative budget and circulation. Consequently, assimilation based on the *statistics* of an evolving ensemble of clouds is more appropriate in this context. In addition, when studying the relationship between cloud-scale processes and the climate system, cloud resolving models are typically forced with observations and/or analyses of the large scale. In this case, the model evolution is affected less by errors in the initial conditions than by uncertainties in the physical and dynamical formulation of the model itself. In the current application, this uncertainty is dominated by the choice of semi-empirical parameters in the model's cloud microphysical routine.

The MCMC methods introduced in chapter 2 are extended to estimate the sensi-

tivity and joint conditional probabilities of semi-empirical microphysical parameters within the GCE model in chapter 3. In each sensitivity test, a “true” model state (analogous to the traditional control simulation) is generated from specified values of each parameter, and observations are drawn from the true state. In each MCMC search iteration, a microphysical parameter is randomly perturbed, a new state is generated from a complete GCE model integration, and the value of the likelihood function is computed and compared to the previous maximum likelihood. The new state is accepted if the computed likelihood exceeds the previous maximum likelihood, or if the new likelihood is sufficiently close to the old as determined in a partially probabilistic manner (Tarantola 2005).⁶ The MCMC algorithm is run for many successive iterations until a sufficient sample of each parameter is obtained from the set of accepted states. The resulting *a posteriori* parameter probability density distributions provide valuable information on the stability of the model with respect to changes in each parameter—information which is useful for determining which observations might effectively be used to constrain the evolution of the model state, as well as which parameters might be targeted for tuning. In addition, idealized experiments, in which it is assumed that the observation errors can be reduced consistent with the launch of future sensors, can be used to both determine the impact of future observing systems, and to assess the fidelity of the cloud-precipitation and cloud-radiation relationships in the model.

In chapter 3, the information on parameter error distributions obtained from MCMC is used to design methods with which to assimilate cloud information from sensors on TRMM and the A-Train for the purpose of constraining the model bias in long-term climate feedback experiments. It is found that the GCE results are most sensitive to changes in the snow fallspeed exponent, graupel fallspeed coefficient and density, and the slope intercept of the rain drop size distribution (Table 1.1). In

⁶For a more complete description of how the MCMC algorithm is implemented, the reader is referred to appendix C.

addition, significant and quasi-linear relationships are found to exist between these parameters and several of the TRMM observations, indicating that variational assimilation methods might be employed to constrain them. This is advantageous in that variational methods are far more (temporally and spatially) globally applicable than the much more computationally expensive MCMC algorithm.

In chapter 4, a variational parameter estimation algorithm is first introduced, then tested using synthetic observations. Unfortunately, because of the complexity of the parameter space, variational methods fail to converge to a solution. The MCMC algorithm is instead used to characterize parameter PDFs using TRMM retrievals, and it is found that (1) the probability space contains multiple modes, and (2) that the same four key parameters identified in simulated-observations sensitivity tests are also the most important parameters in real-data experiments. An optimal set of parameters is identified from inspection of the parameter PDFs as well as from an analysis of the state-observation differences via the likelihood function, and these parameter values are used in the large-domain RCE experiments documented in chapter 5.

1.5 CRM radiative convective equilibrium climate feedback experiments

Once the most uncertain cloud microphysical parameters have been identified (chapter 3) and their values constrained through assimilation of cloud and precipitation statistics from TRMM (chapter 4), the GCE is used to assess the response of the precipitation efficiency and key cloud properties to changes in the sea surface temperature in chapter 5. Use of a large horizontal domain allows for the effects of mesoscale organization of convection and accommodates the possibility of studying feedbacks between clouds and radiation and the large scale circulation. In each numerical experiment, the GCE is run for three different constant values of SST (295, 300, and

305 K) until a state of quasi-equilibrium is reached. The thermostat and infrared iris (Fig. 1.1) cloud-climate feedback mechanisms are assessed by studying the characteristics of cloud and precipitation processes in each of the equilibrium states. SST perturbation RCE experiments are performed for two general scenarios. Simulations are first conducted for zero mean wind conditions to assess the idealized response of clouds and precipitation to changes in sea surface temperature. A second set of experiments are then performed, in which moderately-sheared background flow consistent with monsoon conditions over the tropical warm pool is applied.

It is found that simulations run without shear exhibit organization of convection in regular regions approximately 1000 km in scale, which increase in number with increasing SST, while simulations run with vertical shear organize into one large moist region and one dry subsident region, regardless of the value of SST. In simulations with shear, the large-scale overturning circulations associated with regions of convection are located almost entirely above the atmospheric boundary layer, while overturning in unsheared simulations extends from the tropopause to the surface. In both sets of experiments, the precipitation amount, intensity, and efficiency increase with increases in SST, while the high cloud fraction and domain-mean outgoing shortwave radiation decrease. At the same time, the precipitable water vapor and outgoing longwave radiation increase with increasing SST, which is consistent with an overall warmer and more moist equilibrium atmosphere. In general, the results of large-domain CRM experiments appear not to support the operation of a thermostat-type feedback mechanism, as high cloud fraction and outgoing shortwave radiation do not increase with SST. Instead, although the total ice mass increases with SST, precipitation efficiency increases and the high cloud fraction and outgoing shortwave radiation decrease, indicating that an iris-like mechanism may be at work. It should be noted, however, that it is difficult to draw firm conclusions from experiments run in a two-dimensional framework, and that three dimensional simulations run with a coupled

ocean model would be required to properly assess the operation of a cloud-climate feedback mechanism.

1.6 Outline and Key Findings

The remainder of this study is organized as follows: chapter 2 contains the results of the MODIS proof-of-concept study and details implications for use of MCMC to assess and improve variational estimates. In chapter 3, MCMC is applied to the GCE model using simulated observations in an identical twin assimilation study. The resulting parameter histograms and the relationship between observations and state are used to construct a (variational) parameter assimilation system in chapter 4, and this system is tested using simulated observations. It is shown that the variational framework experiences problems due to the complexity of the parameter space, and the MCMC algorithm is instead used to constrain GCE cloud microphysical parameter values using observations taken during the South China Sea Monsoon Experiment. The resulting observation-constrained CRM is used to study the tropical climate system, and results of large-domain 2D CRM experiments are presented in chapter 5. Discussion of the implications of the results contained in this study, general conclusions, and suggestions for future work are offered in chapter 6.

The various components of this study provide unique contributions to the state of knowledge in the fields of remote sensing, data assimilation, and cloud-climate feedback analysis. The MODIS retrieval application (chapter 2) represents one of the first attempts to explicitly assess the effect of Gaussian error assumptions on an optimal estimation retrieval using a MCMC technique. The MCMC-based cloud microphysical parameter sensitivity analysis presented in chapter 3 is unique in that it is a far more thorough CRM sensitivity study than is typically reported in the literature, and is the first documented application of MCMC to a fine-scale nonlinear atmospheric

model. Furthermore, use of TRMM retrievals in the MCMC algorithm (chapter 4) represents one of the first times TRMM observations have been assimilated into a CRM, the first time satellite-based remote sensing observations have been used to obtain PDFs of cloud microphysical parameters. The large-domain RCE experiments documented in chapter 5 represent the first time a CRM with observation-constrained cloud microphysical parameters has been used to assess the relationship between sea surface temperature and the tropical hydrologic cycle, as well as the one of the first times a CRM has been used to explicitly evaluate cloud-climate feedback hypotheses.

Chapter 2

MCMC Proof of Concept: Retrieval of Cloud Properties from MODIS

2.1 Introduction: MCMC vs Variational Methods in the Retrieval Context

In chapter 1, data assimilation methods were proposed as a way to first estimate, then mitigate uncertainty in GCE cloud microphysical parameterization. Most “physical” retrieval approaches and data assimilation techniques are based on statistical estimation (“optimal⁷ estimation”) theory, and assume mean-zero Gaussian error statistics. As a result, a Gaussian probability density distribution is imposed on the estimated state, and description of the error statistics is restricted to mean and (co)variance alone. Unfortunately, there is no way to know whether the probability density function (PDF) of the errors is truly Gaussian in form, and hence whether an optimal estimation method has produced a robust result. In chapter 1, MCMC sampling was proposed as an alternative to variational algorithms as it makes fewer assumptions

⁷The term “optimal” implies that the resulting estimate is the best choice from among several possible estimates, and does not necessarily imply that methods based on minimization of a cost function are superior to other estimation techniques.

about the form of the probability of the estimated state and returns the full PDF of the solution. However, MCMC has not yet been widely used in the atmospheric sciences and prior to applying the method to a highly nonlinear and relatively computationally expensive model, it is first useful to examine the properties of the algorithm in a simpler context.

To this end, the properties of the MCMC algorithm are demonstrated in this chapter using a relatively straightforward two-parameter Moderate Resolution Imaging Spectroradiometer (MODIS) cloud retrieval. In the retrieval context, the optimal estimation approach has been applied to satellite data to obtain a wide variety of atmospheric properties. Engelen and Stephens 1997, Engelen and Stephens 1999, Miller et al. 2000, Stephens et al. 2001, and L'Ecuyer and Stephens 2002, for example, have demonstrated the utility of the optimal estimation technique in the context of ozone, water vapor, cloud property, aerosol, and rainfall retrievals from instruments ranging from radars to lidars to passive measurements at visible, infrared, and microwave wavelengths. These techniques have met with enough success, in fact, that the cloud liquid and ice water content and particle size retrievals central to the CloudSat mission (Stephens et al. 2002) are also formulated in the optimal estimation framework (Austin et al. 2001 and Benedetti et al. 2003). While optimal estimation techniques represent a theoretical advance over empirical retrieval techniques, as with variational assimilation methodologies, the robustness of the solution produced by a variational retrieval depends directly on the validity of the Gaussian error assumption. Though a variety of diagnostics such as the averaging kernel or information content theory can be used to evaluate the performance of a variational retrieval (Rodgers 2000, L'Ecuyer et al. 2005), these properties are themselves limited by the assumption of Gaussian error statistics. To truly assess the efficacy of a variational retrieval, and to quantify the errors associated with Gaussian assumptions and the use of the *a priori*, one must have the full conditional PDF. If the full PDF were available, it would be

possible to directly quantify the degree of departure from the Gaussian form, and this information could then be used to improve the retrieval method. Markov chain Monte Carlo (hereafter MCMC) sampling methods (Metropolis et al. 1953, Hastings 1970, Robert and Casella 1999, Tamminen and Kyrola 2001, Tamminen 2004, Gelman et al. 2004, Tarantola 2005) thoroughly and efficiently search multidimensional parameter spaces for regions of concentrated probability via a guided random walk. The search direction is determined using a measure of error-weighted distance between observations and forward model output, analogous to a variational cost function.

Relative to variational methods, the power of MCMC is that, aside from setting reasonable constraints on the allowable parameter ranges, there are no assumptions made; the algorithm requires no *a priori* state estimate and assumes no particular form of the PDF. Hence, what is output from the MCMC method is the true probability of the state being the true state, given the particular observations and model.⁸ The drawback is that MCMC requires a large number of samples to effectively represent the posterior PDF. These samples are obtained by running the radiative transfer model several thousand times for each pixel retrieved, leading to significant computational cost. While MCMC may be too computationally expensive to be applied in most operational retrieval contexts, because it yields the full conditional PDF, MCMC can be used to assess the accuracy and improve the performance of a variational retrieval.

In this chapter, the usefulness of MCMC is demonstrated by comparing PDFs retrieved from MCMC with Gaussian PDFs obtained from a variational retrieval. To this end, a simplified version of a computationally efficient radiative transfer model is used to retrieve optical thickness (τ) and effective particle radius (r_e) from Moderate Resolution Imaging Spectroradiometer reflectances in 0.64 and 2.13 micron channels

⁸It should be noted that there are assumptions built into any model that relates state to observations. Because of this, the MCMC PDF is always the PDF of the state *conditioned on* the particular observations and model used.

in a manner analogous to Nakajima and King (1990). Observation-dependent error is assigned to each measurement, and two experiments are performed—a control experiment and an experiment in which the measurement error is doubled. Doubling the assumed measurement error allows us to examine the behavior of the MCMC algorithm for the case in which the forward model, the observations, or both are less certain. In each experiment, MCMC is used to generate 20,000 samples of the joint conditional PDF of optical depth and effective radius for each MODIS pixel, and the resulting PDFs are then compared with Gaussian PDFs for the same variables generated from a variational retrieval. It is found that while PDFs of effective radius are generally well-approximated by a Gaussian form, the shape of the PDFs is decidedly log-normal for optical depths greater than approximately 20. The departure of the PDF from Gaussian at higher values of optical depth leads to a loss of sensitivity in the variational retrieval to optical depths above 30. When the variational retrieval is modified such that the natural log of optical depth is estimated in place of the optical depth itself, the sensitivity to larger optical depths is restored and the variational PDFs more closely approximate the true conditional PDFs.

The remainder of this chapter is organized as follows. Section 2.2 presents the details of the radiative transfer model used to compute MODIS reflectances and explains the setup of the variational retrieval. The MCMC algorithm is briefly described in section 2.3, and the specific application of MCMC to the current problem is explained. Results of a comparison between variational and MCMC retrievals are shown in section 2.4; these results are used to modify the variational retrieval, and the resulting improved retrievals are presented in section 2.5. Further uses of MCMC in a retrieval context, as well as general limitations of the MCMC algorithm are discussed in section 2.6. Brief concluding remarks and suggestions for future work, including how MCMC might be applied to the wider data assimilation/parameter estimation problem, are offered in section 2.7.

2.2 Radiative Transfer Model and Variational Retrieval

2.2.1 Description of the Radiant Radiative Transfer Model

In the interest of computational efficiency, since our intent is to compare PDFs produced by a variational retrieval with MCMC PDFs, we employ a simplified radiative transfer model that is physically consistent with the scene of interest. The forward radiative transfer model used in the retrieval is a simplified version of the Radiant model (Gabriel et al. 2005), which is a multi-stream plane-parallel solver applicable to both visible and infrared portions of the electromagnetic spectrum. Radiant employs a multiple scattering solution based on a combination of the well-known interaction principle and the eigenmatrix method and provides an accurate and computationally efficient solution for radiances in both cloudy and clear conditions. For our test case, we retrieve cloud optical depth and effective radius from measured reflectances in 0.64 and 2.13 micron wavelengths in a manner analogous to the method described in Nakajima and King (1990). The scene of interest was observed at approximately 2130 UTC on 4 July 2001 by the MODIS instrument on EOS Terra, and consisted of low broken stratus and stratocumulus clouds over the northeast Pacific Ocean (Fig. 2.1).

Since the scene was viewed near local solar noon and at near-nadir view angles, we neglect the dependence of radiance on azimuth and zenith, as well as the effects of reflection from the ocean surface. We also neglect the effects of absorption by trace gases, which are negligible in the visible and small for the 2.13 micron channel. Absorption by liquid cloud is only allowed to affect the partition between scattering and absorption through the 2.13 micron single scatter albedo, allowing us to assume identical optical depth for both 0.64 and 2.13 micron channels. The atmosphere is assumed to consist of a single layer filled with cloud particles, and because the clouds




Figure 2.1: MODIS observed reflectance for 0.6 micron (a) and 2.13 micron (b) channels. Scene was observed over the northeast Pacific ocean at approximately 2130 UTC 4 July 2001.

are low, we assume them to be entirely ice-free and composed of spherical particles. This in turn allows us to approximate the scattering phase function using a double Henyey-Greenstein phase function (Henyey and Greenstein 1941). We further neglect scattering from the direct beam so that scattering of diffuse radiation is the only form of attenuation considered. The problem thus simplifies to solution of a homogeneous partial differential matrix equation, the solution to which is exponential in form such that the resulting forward model is nonlinear in both optical depth and effective radius.

Having specified the form of the scattering phase function, and assuming a single scatter albedo of approximately 1.0 for the visible channel, the retrieval problem reduces to estimation of only two free parameters: optical depth and 2.13 micron single scatter albedo. Ackerman and Stephens (1987) showed that modified anomalous diffraction theory could be used to relate absorption in clouds to the cloud effective particle radius through the single scatter albedo in near infrared and visible wavelengths. Their results are used to obtain a power relationship between effective radius and single scatter albedo tuned to liquid absorption in the 2.13 μm channel so that the retrieval can be posed in terms of effective radius

$$r_e = \left(\frac{1 - \hat{\omega}_o}{0.00266} \right)^{\frac{1}{0.89}}. \quad (2.1)$$

Sensitivity tests (not shown), in which surface reflectance and gaseous absorption effects were added to the forward model for a range of optical depths and effective radii, yielded small differences in the computed reflectances, and it is our assertion that the simplified radiative transfer model is sufficient for the current study.

2.2.2 Variational Retrieval Method

Given a model $F(\mathbf{x})$ that relates a set of observed reflectances $\mathbf{y}$ to a set of state parameters $\mathbf{x}$ we would like to estimate the true atmospheric state that produces these observations. The solution to this problem can be obtained by maximizing the conditional probability that the state is equal to the true state given the available observations and the radiative transfer model. For our purposes, the state consists of optical thickness and cloud particle effective radius, and the observations are MODIS reflectances at 0.64 and 2.13 microns. Because the state and observation error probability density distributions are often not known, it is common to assume that they take a Gaussian form. This is both useful and appropriate when only the mean and variance are known as the characteristics of the Gaussian PDF are fully determined by only two parameters (Devore 1995). Gaussian distributions are also algebraically easy to manipulate and have the added benefit that the most probable solution is identical to that which minimizes the error-weighted difference between state and observations. The solution is then obtained by minimizing a cost function that consists of the negative log of the Gaussian posterior probability distribution (Eqn. 1.5)

$$\begin{aligned} \Phi(\mathbf{x}) = -2 \ln [P(\mathbf{x}|\mathbf{y})] \propto & [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] + \\ & [\mathbf{x} - \mathbf{x}_a]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_a] + C, \end{aligned} \quad (2.2)$$

where $\mathbf{x}_a$ is the *a priori* value of the state, $\mathbf{R}$ and $\mathbf{B}$ are the observation and *a priori* error covariance matrices, respectively, and C is the normalization constant for the Gaussian probability distribution. The maximum probability (minimum posterior error variance) solution is found by minimizing the gradient of the cost function, $\nabla\Phi$, with respect to the set of unknown $\mathbf{x}$. In this case, we set $\nabla\Phi = 0$ and solve for $\mathbf{x}$

through a Newtonian iterative solution to the resulting equation (Rodgers 2000)

$$\hat{\mathbf{x}}_{i+1} = \hat{\mathbf{x}}_i + \left(\mathbf{B}^{-1} + \mathbf{K}_i^T \mathbf{R}^{-1} \mathbf{K}_i \right) \left[\mathbf{K}_i^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{F}(\hat{\mathbf{x}}_i)) - \mathbf{B}^{-1} (\hat{\mathbf{x}}_i - \mathbf{x}_a) \right]. \quad (2.3)$$

Here, $\hat{\mathbf{x}}_i$ is the current estimate, $\hat{\mathbf{x}}_{i-1}$ is the prior estimate, $\mathbf{K}$ is the matrix of partial derivatives of the forward model with respect to $\mathbf{x}$ (the Jacobian matrix), and $\mathbf{K}^T$ is its transpose. Equation (2.3) makes it clear that the solution depends largely on both the specification of the form and magnitude of error covariances $\mathbf{R}$ and $\mathbf{B}$. Two key assumptions must be made (1) the magnitude of the covariances and (2) the imposition of a Gaussian shape. In the next section, we describe how MCMC can be used to obtain the full unapproximated conditional PDF, and how this PDF can be used to estimate errors that result from the Gaussian constraint.

2.3 MCMC Sampling Procedure

In the retrieval process, what we desire is a search of the probability space for the most likely values of the retrieved state given a set of observations. To be able to evaluate the robustness of the retrieval, as well as the effective information content of the observations, the full joint conditional PDF for the set of retrieved state parameters is required. If we wish to retrieve only one or two parameters such that the state space is one or two dimensional, then it would be feasible to step through the range of possible values for each parameter in small increments, computing an error-weighted distance between state and observations and computing the PDF in the process. However, for state dimensions greater than one or two, this sort of brute force sampling quickly becomes intractable. MCMC is based on the fact that most values of the state are associated with very low probability; the state space is largely empty (Tarantola 2005). In fact, as the dimensions of the estimation problem grow, the state space tends to become increasingly empty. MCMC exploits this fact, searching out areas

of concentrated probability within the state space and sampling them via a guided random walk. In this chapter, we use a form of MCMC similar to the well-known Metropolis-Hastings algorithm (Metropolis et al. 1953, Hastings 1970). For thorough and accessible discussions of the theoretical underpinnings of the MCMC algorithm, the reader is referred to Mosegaard and Tarantola (1995), Gelman et al. (2004), and Tarantola (2005).

The Markov chain is started from a set of state parameters drawn from a bounded Uniform prior, with bounds set to physically realistic values for each parameter. In each MCMC search iteration a randomly-chosen state parameter is perturbed by an amount equal to a uniform random variable multiplied by the allowable parameter range and modified by a tuning factor determined during the adaptive burn-in (see discussion below).⁹ If the newly-perturbed parameter value exceeds its allowable range, the new value is discarded and another random perturbation is applied. This process is repeated until an acceptable parameter value has been obtained, at which point new forward observations are computed and the distance between state and observations is calculated and compared to the previous value. For the current application, we employ a Gaussian (L_2) likelihood function defined identically to the Gaussian cost function in equation (2.2), except that MCMC needs no *a priori* term. For consistency, the observation error covariance is specified as in the optimal estimation technique according to instrument error characteristics and forward radiative transfer model errors. The new state is accepted and stored if the new likelihood exceeds the old, or if the new value is sufficiently close to the previous value as determined in a partially probabilistic manner (Fig. 2.2). This *probabilistic acceptance* is important in that it (1) encourages the algorithm to sample regions close to areas of relatively large probability, and (2) allows the algorithm to leave a local probability

⁹The bounded uniform PDF is appropriate in the retrieval context as it represents the minimum possible prior information on the state, and assumes no other information aside from a reasonable range of values for each retrieved parameter.

maximum. In the current application, a given state is probabilistically accepted if the ratio of the new likelihood to the old is greater than a random variable X between 0 and 1. Because we have computed the log-likelihood, the new state is accepted if

$$X < \exp(\Phi_{new} - \Phi_{old}) \quad (2.4)$$

where X is drawn from the *a priori* probability distribution. As is common practice, we use a burn-in period in which sampling is done, but no states are stored, for the purpose of "forgetting" the initial value in the chain. Similar to the method employed in Tamminen and Kyrola (2001) and later modified in Tamminen (2004), we allow the algorithm to adaptively change the step length during the burn-in period. Tuning the step length during burn-in leads to more efficient sampling of the state space once burn-in has finished. After burn-in, adaptive sampling is turned off to ensure sampling from the correct posterior distribution. It was found that a burn-in length of 2,000 samples was generally sufficient to converge to sampling a stationary PDF.

The MCMC algorithm is run for many successive iterations to allow a thorough sampling of the state space, and a posterior probability density distribution is built for each state parameter from the set of accepted states. Based on initial tests, it was found that a sample size of 1,000 was sufficient to characterize the moments of the posterior PDF. To ensure full sampling of the PDF in each pixel, we chose to run the algorithm until 20,000 samples were accumulated for each pixel. Depending on the error characteristics and the shape of the actual joint conditional PDF, it took between 35,000 and 500,000 iterations after the adaptive burn-in to obtain a 20,000 sample set.

An illustration of the convergence of the MCMC algorithm to sampling a stationary distribution is presented in figure 2.3, which depicts samples obtained from the joint PDF of optical depth and effective radius for pixel (19,19) in the MODIS




Figure 2.2: Schematic depiction of the MCMC iterative process. (a) A single MCMC iteration and (b) a portion of the MCMC random walk depicting the accept/reject sequence.

dataset. It is clear that by as early as 200 samples (Fig. 2.3a), the PDF is starting to take shape, and that by 10,000 samples (Fig. 2.3d) the PDF is well-characterized. For comparison, the PDF obtained from 1,000,000 samples (Fig. 2.3e) is contrasted with the result of brute force sampling across the range of values for optical depth and effective radius (Fig. 2.3f). Brute force sampling was performed by stepping through the range of optical depth and effective radius in increments of 0.1, running the forward model, and computing the likelihood function in each iteration. Though the optical depth PDF has a slightly longer tail in the MCMC result, the PDFs are remarkably similar, demonstrating the effectiveness of MCMC in characterizing the true PDF.

2.4 Results

In this section, we present results from two different comparisons between MCMC and variational retrievals. In the control case (hereafter CTRL), we specify errors in each channel as 5% of the observed reflectance. Prior studies have shown the error to be approximately equal to five percent of the measurement in the visible channel, and approximately equal to two percent of the measured reflectance in the near infrared (L'Ecuyer et al. 2005). Because we are neglecting the effects of trace gas absorption, we inflate the error in the 2.13 micron channel to five percent to account for additional model error in this channel. In the second case (hereafter 2XERR), we double the observation error and examine the effect on the resulting retrieved states and PDFs. In variational retrievals, it is common to combine instrument and forward model error in the observation error covariance; the 2XERR case therefore accounts for either noisier observations, a more highly-parameterized forward model, or a combination of the two. *A priori* optical depth and effective radius for both cases were set equal to




Figure 2.3: Example of convergence of MCMC to sampling a stationary distribution—in this case samples of the joint conditional probability of optical depth and effective radius for pixel number 19,19. Progressively greater numbers of samples are shown in (a) through (e), while the result of brute force integration in increments of 0.1 in optical depth and effective radius is shown in (f) for comparison.

physically reasonable values for liquid stratus clouds; 20 and 10 microns, respectively (Miles et al. 2000). Error variances of 225 and $100 \mu m^2$ were assigned to these *a priori* values for the purpose of spanning ranges of optical depth and effective radius observed in nature; 0 to 100 and 0 to 60, respectively. It was assumed for simplicity that observation errors were uncorrelated between channels, as were *a priori* errors, and convergence to a solution for the scene of interest was uniformly obtained within 10 iterations. In both CTRL and 2XERR cases, we compare both retrieved values and PDFs, as well as the information content, and compare the variational retrieval with the mean, median, and mode of the distribution obtained from MCMC. If the PDF is truly Gaussian, or at least centered and unimodal, the mean, median, and mode retrieved values should be identical.

2.4.1 Control Case

Retrieved optical depth and effective radius are plotted for the variational retrieval in figures 2.4a and 2.4b, for the mode of the MCMC PDF in figures 2.4c and 2.4d, and for the mean of the MCMC PDF in 2.4e and 2.4f. Comparison of mean and median solutions from MCMC (not shown) revealed them to be nearly identical up to optical depths of 75, hence only the mean result is shown. Examination of the retrieved optical depth shows immediately that, though all estimates recover the relatively thin cloud and clear areas well ($\tau < 20$), retrieval of larger values of optical depth varies significantly across the different retrievals. This is a reflection of the loss of sensitivity of cloud albedo with increasing cloud optical depth (Stephens 1978), which translates to a loss of sensitivity in the visible channel to optical depths greater than approximately 50. Retrieved effective radius exhibits almost exact similarity across all retrievals, with the only significant differences occurring in regions of clear air, or where clouds are extremely thin.




Figure 2.4: Optical depth and effective radius for the control case from (a and b) variational retrieval, (c and d) mode of the MCMC PDF, and (e and f) mean of the MCMC PDF. Effective radius has been screened out (gray areas) where 0.64 micron reflectance is less than 0.25.

Retrieved optical depth and effective radius PDFs from both MCMC and variational techniques are compared in figure 2.5 for a 5x5 pixel subregion outlined by the white box in figures 2.4c and 2.4f. In general, it can be seen that all MCMC PDFs are mono-modal, indicating that a quadratic minimization technique can be expected to perform well. In addition, the effective radius PDFs (Fig. 2.5d), as well as PDFs in pixels in which optical depth is less than ≈ 20 , appear to be well-characterized by a Gaussian distribution. However, as optical depths increase beyond 20, the optical depth PDFs increasingly depart from Gaussian. In fact, for optical depths greater than 50 (pixels (15,19), (16,19) and (16,20)), the MCMC PDF exhibits little departure from the *a priori* bounded uniform PDF, indicating that the observations have contributed very little information to the solution. For pixels with optical depths between 20 and 50, the MCMC PDFs appear to be log-normal in shape, with a mode that is increasingly displaced from the variational mode with increasing optical depth. The fact that the PDFs for effective radius are not as skewed as those for optical depth is consistent with the physical limitations imposed on the drop size distribution by the collision-coalescence process. As drop radii increase beyond $20 \mu m$, the drizzle process becomes increasingly active, limiting the presence of large cloud droplets—particularly in liquid clouds over the ocean (Cooper et al. 2003).

To more clearly demonstrate the similarities and differences between variational and MCMC retrieved PDFs, τ and r_e from two selected pixels are shown enlarged at the right side of figure 2.5. Figures 2.5b and 2.5c depict optical depth PDFs for pixels with relatively low (10-15) and relatively high (≈ 40) optical depth respectively, while figures 2.5e and 2.5f depict PDFs of effective radius in these same pixels. In these figures it can be seen that the observations contribute enough information to effectively constrain the retrieval for low optical depths (Fig. 2.5c); the MCMC and variational solutions are nearly identical in this case. However, with increasing optical depth (Fig. 2.5b), the solutions diverge, and the variational PDF provides a poor




Figure 2.5: PDFs from MCMC retrieval of (a) optical depth and (d) effective radius for the control case. Shown at right are PDFs for two selected pixels (b) and (e) relatively large optical depth and small effective radius and (c) and (f) relatively small optical depth and large effective radius. Uniform a priori is shown in red, posterior MCMC PDF is shown in black, and Gaussian variational PDF is in dashed blue. The two pixels shown in greater detail are outlined in thick dashed green in the 5x5 matrix. The total integrated absolute difference between MCMC and Gaussian PDFs (see text for explanation) is printed in the upper right hand corner of each pixel in (a).

approximation to the solution. Though pixels with both relatively small (Fig. 2.5e) and large (Fig. 2.5f) effective radii appear to be well-characterized by a Gaussian PDF, careful examination of all four figures reveals that each of the MCMC PDFs has a tail—none are perfectly Gaussian. This tail is uniformly much smaller for the effective radius than for optical depth, though it can be more clearly seen in the pixel with relatively large effective radius (Fig. 2.5f).

2.4.2 Double Error Case

Plots of retrieved optical depth and effective radius for the 2XERR case are shown in figure 2.6. Comparison between figures 2.6c and 2.6d and 2.4c and 2.4d reveals that the mode of the PDF of both τ and r_e retrieved with MCMC appears to be invariant with changes in observation error. However, τ and r_e retrieved from the mean of the MCMC PDF exhibit a reduced dynamic range compared with figures 2.4e-2.4f, indicating that the PDFs for large values of optical depth have shifted toward lower values of optical depth. The variational retrieval also exhibits this effect; the scene now contains no values of optical depth greater than 40 (Fig. 2.6a). Given the fact that the *a priori* τ was set at 20, this decrease in sensitivity reflects the decrease in influence of the measurements on the retrieval; the optimal estimation solution has shifted toward the *a priori* value. Note that, as in the control case, areas of thin cloud and clear air are largely unaffected by increases in assumed observation error (Figs. 2.6a-2.6c). Comparison of Figures 2.6d-2.6f with 2.4d-2.4f reveals a slight decrease in dynamic range for effective radius, indicating broadening of the PDF in both directions.

Plots of PDFs of τ and r_e for the 2XERR case are presented in figure 2.7. Consistent with the reduced ability of observations with greater “noise” to constrain the solution in both MCMC and variational retrievals, the PDFs in the 2XERR case are uniformly more broad than those in the control case (Fig. 2.5). Plots of PDFs of




Figure 2.6: As in figure 4, but for the double error case.

optical depth (Fig. 2.7a) demonstrate that, almost uniformly, at τ greater than 30, the variance has collapsed to the *a priori* for the variational retrieval. At optical depths less than 20, there is still relatively high sensitivity to the observations, as indicated by the well-defined PDFs in both variational and MCMC cases. As in the control case, many of the MCMC PDFs are decidedly log-normal for larger values of optical depth. Plots of the PDFs of effective radius (Fig. 2.7d) indicate that, while PDFs have broadened, and exhibit a more pronounced tail compared with those in the control case (Fig. 2.5d), the shape and location of the PDF for effective radius appears to be relatively well-approximated by the variational retrieval in all cases except for the extremely thin cloud. Closer examination of the PDFs reveals that for optical depths greater than 50 (Fig. 2.7b), there is almost uniform probability that the solution could lie anywhere between τ of 50 and 100. In the case of thin cloud (Fig. 2.7c), the solution is still well-constrained, but a much larger tail is now evident in the MCMC solution. Examination of effective radius in the same pixels (Figs. 2.7e and 2.7f) reveals that, compared with CTRL, increased instrument noise has led to broader tails for the effective radius PDFs as well. Overall, it is interesting to note that increasing the assumed instrument noise has not changed the fundamental nature of the solution, it has simply enhanced the differences between MCMC and Gaussian PDFs. This is to be expected—for mono-modal PDFs, observations with small errors should provide a reasonable constraint on the solution, even for non-Gaussian PDFs. Having demonstrated the qualitative differences between MCMC and variational solutions, we now proceed to quantify PDF differences and to use the MCMC results to improve the variational solution.




Figure 2.7: As in figure 5, but for the double error case.

2.5 Quantitative Assessment and Improvement of Variational Retrieval

2.5.1 Information Content and Assessment of Variational Retrieval

To be able to effectively retrieve certain state variables from observations, the observations must contain sufficient information on the desired state to be able to constrain the solution. It is, therefore, common to use so-called information content measures to quantify the amount of information contributed to the estimate of the state by an individual observation or sets of observations (Rodgers 2000, L'Ecuyer et al. 2005). One common measure of information is the Shannon Information Content, in which the information content of an observation or set of observations is computed as the reduction in entropy due to the addition of information from observations. In discrete form, the entropy of state P is defined as:

$$S(P) = - \sum_{i=1}^N p_i \log_2 (p_i), \quad (2.5)$$

where p is the discretized PDF, N is the number of discrete bins the PDF is divided into, and log base-2 is used so that the result is defined in bits. Shannon information content is then defined as the difference between the entropy of the *a priori* state and the entropy of the retrieved state

$$H = S(\mathbf{x}_a) - S(\hat{\mathbf{x}}) = \left[- \sum_{i=1}^N p_i(\mathbf{x}_a) \log_2 (p_i(\mathbf{x}_a)) \right] - \left[- \sum_{i=1}^N p_i(\hat{\mathbf{x}}) \log_2 (p_i(\hat{\mathbf{x}})) \right], \quad (2.6)$$

which can be interpreted as the extent to which the number of allowable states is reduced by the addition of information from the measurements. For the variational case, both estimated and *a priori* PDFs are Gaussian, and the information content

reduces to (Rodgers 2000, L’Ecuyer et al. 2005)

$$H = \frac{1}{2} \log_2 |\mathbf{S}_a \mathbf{S}_x^{-1}|. \quad (2.7)$$

Since this expression is predicated on the assumption of Gaussian statistics, a quantitative assessment of the robustness of the Gaussian assumption can be obtained by comparing information content from the variational retrieval with that computed explicitly from the MCMC PDFs. Information content was computed for the variational retrieval according to equation (2.7), while MCMC information content was computed from equation (2.6). For consistency, since information content is a relative measure of retrieval performance, the information content was computed with respect to the variational Gaussian *a priori*.

Comparison of information content between variational and MCMC retrievals (Fig. 2.8) indicates that the variational solution over-estimates the information content of the observations. Mean information content for optical depth was 0.7 bits higher for the variational estimate, while mean information content for effective radius was 0.26 bits higher. Recall that the MCMC-retrieved PDFs of τ and r_e for the control case (Fig. 2.5) were decidedly log-normal for both τ and to a lesser extent r_e . This suggests that larger values of τ will have a higher probability of occurring in the MCMC retrieval. The absence of such a “tail” in the variational solution falsely leads one to believe that the observations do a better job of constraining the solution than they actually do. The fact that there is a greater discrepancy for optical depth as compared to effective radius is due to the fact that MCMC-retrieved optical depth PDFs exhibited a consistently greater departure from Gaussian.

A more quantitative measure of the differences between variational and MCMC solutions can be obtained by computing the integrated absolute difference (IAD) be-




Figure 2.8: Comparisons of information content computed from MCMC PDFs (a and c) and variational retrievals (b and d) for optical depth (a and b) and effective radius (c and d). All results are taken from the control case.

tween the two retrieved PDFs. This measure is computed by calculating the absolute value of the difference between the variational and the MCMC PDF in each discrete bin, then summing the result. Since each PDF uses the same bin intervals, and since the PDFs are normalized so that they each integrate to one, this measure is independent of bin size. Since each PDF integrates to 1.0, a value of IAD of 1.0 indicates a 100% difference between the PDFs. For illustration, values of IAD for optical depth PDFs are printed in the upper right hand corner of each pixel in Figures 2.5a-2.5c. From this figure it can be seen that IAD values of 0.0 to 0.2 represent a negligible difference between variational and MCMC PDFs, while values above 0.7 represent relatively large departures.

2.5.2 Use of MCMC Results to Modify and Improve the Variational Retrieval

The preceding section provides examples of diagnostics that can be used to quantify the quality of a Gaussian-based variational retrieval once the full PDF is known. The greater the integrated differences between PDFs (and consequent differences in information content), the more non-Gaussian the statistics are, and the poorer the retrieval will be. It is logical to ask whether this information can, in turn, be used to improve the variational retrieval, since it is still a desirable method by virtue of its computational speed and its basis in the physics of the radiative transfer problem.

In theory, visible reflectance should be sensitive to optical depths up to 50, however, the variational retrieval demonstrated loss of sensitivity at optical depths above 30 (Figs. 2.4 and 2.6). The fact that the MCMC algorithm returned log-normal PDFs for τ and r_e in both CTRL and 2XERR cases suggests that this problem could be partially mitigated by retrieving the natural log of optical depth and effective radius. A log-normal PDF is thus imposed on τ and r_e , and because the natural log of a log-normally distributed random variable is distributed Gaussian, we can retain the

Gaussian variational framework. The only fundamental change is that the Jacobian is computed with respect to variations in the natural log of τ and r_e , and the *a priori* error variance is specified such that the variance of the log-normal *a priori* distributions for τ and r_e is identical to the variance of the Gaussian *a priori* PDFs.

Because retrieval of the log of effective radius was essentially identical to the original retrieval effective radius (not surprising given the relatively small tails on the PDFs of r_e), we focus on the results of the retrieval of log of τ . Comparison of retrieved optical depth from the original variational retrieval, and the variational retrieval of log- τ for the control case (Figs. 2.9a and 2.9d) demonstrates restored sensitivity to larger values of optical depth for the log- τ retrieval. Retrieved optical depth now much more closely resembles the mode of the MCMC retrieval (Fig. 2.4a), and the information content of the variational retrieval (Fig. 2.9e) now more closely approximates that of the MCMC estimate (Fig. 2.9b). In addition, integrated differences between variational and MCMC PDFs are reduced by approximately 40% in regions where optical depth is greater than 30 (Figs. 2.9c and 2.9f). Though fairly large differences still exist in regions of relatively thick cloud, there are no differences larger than 0.8 left in the domain.

Improved correspondence between MCMC-retrieved τ and the log- τ retrieval is further illustrated in a scatter plot of variational optical depth vs. that from the mode of the MCMC PDF (Fig. 2.10), in which it is clear that the variational estimate loses sensitivity at optical depths greater than 30 for the Gaussian case, while the log- τ retrieval retains a closer match to the MCMC result for larger optical depths. Note, however, that because the *a priori* optical depth is still set at 20, larger values of optical depth still tend to exhibit a low bias in the variational retrievals. This highlights the added benefit of the MCMC techniques; they do not rely on any assumed *a priori* information. Because the information from the visible channel saturates at values of $\tau > 50$, results are compared only for τ in the range of zero to fifty.




Figure 2.9: Plots of retrieved optical depth from the control case (a) and the control case with log of optical depth estimated (d). Differences between PDFs from variational retrieval and MCMC are plotted for (b) control case and (e) control case with log of optical depth estimated. Optical depth information content from MCMC and variational retrieval of log-optical depth are plotted in (c) and (f), respectively.




Figure 2.10: Scatterplot of optical depth retrieved from Gaussian variational (red) and log-normal variational (blue) vs. optical depth retrieved from the mode of the MCMC PDF. The one-to-one line is depicted in black for reference.

2.6 Discussion

One fundamental difficulty that is encountered when applying MCMC to estimation problems is the significant computational cost of the algorithm; typically the forward model needs to be run between 2 and 10 times for each accepted sample leading to a total of on the order of 10^6 forward model runs per pixel and 10^{10} model runs per 100 x 100 pixel scene. If the retrieval is run sequentially, one pixel at a time, the retrieval can take a prohibitively long time to finish. On today's parallel computers, MCMC can be run much faster; the simplest way to parallelize the algorithm is to divide the domain into patches and assign different groups of pixels to different processors. This was the method we adopted in the present study, however, a more interesting and potentially more robust way to parallelize MCMC is to run multiple chains for each pixel and stop after fewer samples have been accepted in each chain. Because a portion of each chain must be discarded in the burn-in phase, this method is less efficient than running a single chain per pixel. However, multiple chains are more likely to thoroughly sample the state space for a given set of observations, especially if the PDF has a complex and/or multi-modal shape (Gelman et al. 2004).

It should be noted that faster convergence to the values of the state that maximize the posterior probability can be obtained through so-called simulated annealing methods (Kirkpatrick et al. 1983, Geman and Geman 1984). While the Metropolis-Hastings algorithm lies at the heart of simulated annealing optimization procedures, simulated annealing does not return information on the probabilities of estimated variables. Since our primary goal is to obtain the conditional PDFs of the set of retrieved variables, we have chosen to employ the traditional Metropolis sampler despite its greater computational cost.

In addition to the computational cost of running MCMC, any application of the MCMC algorithm involves a certain amount of subjective tuning; specifically, it is difficult to anticipate in advance how many samples will be required to adequately

represent the true PDF, and much work has gone into defining appropriate “stopping criteria” for determining when to stop sampling (Robert and Casella 1999). We have chosen to be conservative and accumulate a large number of samples for each pixel. It could be argued that a more efficient setup would involve monitoring the moments of the sampled PDF and stopping when the moments became sufficiently stationary. Another concern is whether MCMC has converged to sampling from a stationary posterior distribution; indeed, many different diagnostics have been developed with which to monitor the MCMC algorithm and ensure it is behaving correctly (Robert and Casella 1999, Gelman et al. 2004). In the retrievals performed here, we tested for convergence by drawing random subsets of 1,000 and 5,000 from the 20,000 sample set for a few of the pixels in the scene. It was found that the moments of the PDF (and hence its overall shape and characteristics) matched very closely between subsets and the overall PDF, indicating sufficient convergence to a stationary distribution.

In addition to convergence and stopping criteria, the correlation of samples within a given Markov chain must also be addressed, as intuitively one would expect adjacent samples in the Markov chain to be correlated to some degree. If the sampler moves efficiently through the state space such that the step sizes between subsequent samples are large relative to the allowable parameter range, then the samples may be assumed to be independent. However, if the PDF is sharply defined, the step sizes will necessarily be small and there can be significant correlations within the chain. To ensure independent samples, we have computed the lagged autocorrelation of optical depth and effective radius within the Markov chain for selected pixels. Results (not shown) indicate that the lagged autocorrelation uniformly decreases to less than 0.5 within 2 steps in the chain, and to less than 0.1 within 10 steps. To ensure independent samples, we used every 10th accepted state in the chain to construct the MCMC PDFs. For a much more in-depth discussion of MCMC diagnostics, the reader is referred to appendix C.

2.7 Conclusions

Given a set of observations and a model that relates observations to a set of desired state variables, Markov chain Monte Carlo methods can be used to generate the full joint conditional posterior PDF for the unknown state. In this chapter MCMC has been used to obtain PDFs for optical depth and effective radius retrieved from visible and near-infrared MODIS reflectances. The results have been used to evaluate a more traditional variational retrieval of τ and r_e that assumes Gaussian error characteristics. It was found that assumption of Gaussian errors leads to a loss of sensitivity in the retrieval to optical depths greater than 30. In addition, assumption of Gaussian errors leads to a consistent overestimate of the information content of the observations. Comparison of MCMC PDFs with those assumed in the variational retrieval revealed that the PDFs for both retrieved variables were log-normal in form, though the effective radius PDFs exhibited a much smaller tail compared with those from the optical depth. When the variational retrieval is modified such that the natural logs of τ and r_e are retrieved, the retrieved values of effective radius change little, but sensitivity to large optical depths is restored and values of optical depth up to 50 can be obtained. In addition, assumption of log-normal errors leads to better fit between the assumed PDFs in the variational retrieval and those obtained from MCMC, and information content is more consistent between the MCMC result and the variational. It be noted that, although MCMC might appear to be too computationally expensive to be applied operationally, MCMC could actually be used to *increase* the efficiency of some retrieval methods. In particular, retrievals that explicitly integrate probabilities over a large database of solutions (e.g., Kummerow et al. 2001) compute the difference between observations and state is computed for every solution in the database, leading to a constant and large number of numerical operations for each retrieved profile. It is possible that MCMC could be used to sample the set of solutions, returning a nearly identical solution for a much smaller number of iterations.

The relatively straightforward application of the MCMC algorithm in this chapter demonstrates the utility of MCMC for evaluating variational estimation methods. In the next chapter, the technique is extended to a much more complex estimation problem, in which it is expected that the PDFs may be highly non-Gaussian, and potentially multi-modal. It will be shown that the PDFs of GCE cloud microphysical parameters contain information on the sensitivity of the model evolution to specified constant parameters, and provide guidance on ways to reduce model error.

Chapter 3

Use of MCMC to Estimate Sensitivity of GCE Simulations to Changes in Microphysical Parameters

3.1 Introduction

As was demonstrated in chapter 1, output of cloud resolving models in general, and the Goddard Cumulus Ensemble model in particular, exhibit considerable sensitivity to changes in the microphysical parameters that determine the particle size distribution and ice crystal habit of condensate species. Because the model's response to this uncertainty affects not only the characteristics of individual convective systems, but also the manner in which convection organizes on large scales, it is imperative that the uncertainty in the model's microphysical parameterization first be quantified, then constrained, prior to performing climate perturbation experiments. Though many authors have demonstrated sensitivity of simulated fields to changes in cloud microphysical parameters, none of these studies have thoroughly tested the response of model output to changes in these parameters. This is largely due to the fact that

such a study would be prohibitively computationally expensive. If we consider the nine parameters defined in Chapter 1, and assume that it is necessary to resolve each parameter into at least 10 bins within its physically realistic range, the model would need to be run 10^9 times to fully explore the parameter space. If each simulation finishes in three minutes,¹⁰ such an experiment would take approximately 5700 *years* to run. If a number of parameters could be eliminated, or their bin sizes coarsened, the experiment might become tractable if run on a parallel computer. However, because of the highly non-linear relationship between parameters and model output, particularly in the case of radiative fluxes, it would be impossible to know whether the effect of changes in each parameter on the given output had been adequately resolved. Such brute force techniques ignore the fact that many parameter combinations are simply not physically realizable—in a probabilistic sense, if the nine parameters are treated as random variables, much of the parameter state space is effectively associated with *very* low probability.

Because the Markov chain Monte Carlo procedure used in the MODIS retrieval proof-of-concept performs a guided random walk that avoids low-probability regions of the state space, it is well-suited to the analysis of microphysical parameter uncertainty. As has been shown, the MCMC algorithm is based on the Bayesian estimation framework, and thoroughly explores the state space of interest, sampling the probability density distribution of a set of unknown parameters given a set of relevant observations. The PDF returned by MCMC satisfies three fundamental requirements prior to assimilating actual observations. First, it returns the relative degree of control each parameter has over the simulation output—the effective sensitivity of model output fields to changes in each parameter. This is represented in the *shape* of the PDF—a progressively greater departure from the a priori Uniform distribution indicates a progressively greater effect of changes in the parameter on the set of observa-

¹⁰Three-minute integration time is based use of current computing resources and the GCE configuration used later in this chapter.

tions of interest. Second, because it returns the full PDF, MCMC provides a baseline estimate of correlations between parameters—information that will be crucial if variational parameter estimation is to be done. Third, MCMC lends insight into which observations should be used to constrain the parameters of interest—this information comes through analysis of the PDFs of the model output variables themselves, and through an examination of how each state variable changes as each parameter is systematically varied. Moreover, because the numerical model is run multiple hundreds of thousands of times, an extensive database of additional model output variables over and above the observations used in MCMC is available. It will be shown that these additional model state variables can be used to identify observations that do not yet exist that would have a positive impact on parameter constraint in the model.

For the parameter sensitivity experiments discussed in this chapter, the MCMC algorithm is set up as a so-called “identical twin” experiment, in which a set of *truth* parameters is specified and used to generate simulated observations. These observations are then used in the MCMC procedure to build the joint PDF of the parameters conditioned on the observations. Such an experiment serves the purpose of returning the desired PDFs, since the observation errors used are identically those that would be used in real-data experiments, but has the added advantage of a known truth state. We can thus assess the degree of accuracy to which MCMC has returned the true parameter values (the maximum likelihood point of the PDF), while also assessing the relative importance of each observation for constraining the parameter set.

The particular set of parameters estimated in this study (Table 3.1) are those most directly linked to liquid and ice particle size distributions and to ice particle type. The current GCE cloud microphysical scheme predicts the mass mixing ratio of two liquid and three ice species, and is based on a combination of Lin et al. (1983) and Rutledge and Hobbs (1983, 1984).¹¹ It should be noted at the outset that all of the parameters

¹¹For a detailed description of the processes contained in the GCE cloud microphysical parameterization, the reader is referred to appendix A.

Table 3.1: GCE cloud microphysical parameters used in the MCMC algorithm, along with truth values for the simulated observation experiment and parameter ranges.

Parameter	Truth	Min	Max
a_s	0.61	0.5	1.0
b_s	0.33	0.1	1.0
a_g	0.88	0.6	1.7
b_g	0.25	0.1	0.7
N_{0R}	1.25	0.04	5.0
N_{0S}	1.25	0.04	5.0
N_{0G}	1.25	0.04	5.0
ρ_S	0.33	0.1	1.0
ρ_G	0.33	0.1	1.0

of interest are highly nonlinearly related to the model predicted condensate mass. For example, a population of precipitating snow or graupel particles is assumed to fall with a speed that is proportional to the characteristic particle diameter according to the power relationship

$$V = a_x D^{b_x}. \quad (3.1)$$

In this case, the fallspeed coefficient (a_x) and exponent (b_x) for species x represent adjustable parameters whose values depend on the crystal habit and degree of riming (Locatelli and Hobbs 1974). In addition, the particle size distribution is assumed to follow an exponential (Marshall-Palmer) distribution, in which the number of particles of diameter D are represented as

$$N(D)_x = N_{0x} \exp \{-\lambda_x D_x\}, \quad (3.2)$$

where N_{0x} is an adjustable parameter that defines the slope intercept of the particle size distribution, and thus the associated proportion of small vs. large particles. λ , the slope parameter, is defined as

$$\lambda_x = \left(\frac{\pi \rho_x N_{0x}}{\rho q_x} \right)^{\frac{1}{4}}, \quad (3.3)$$

where ρ_x is the adjustable (and assumed constant) particle density, and q_x is the predicted mass mixing ratio. An indication of the nonlinear relationship between parameters and predicted condensate mass mixing ratio can already be seen in the above equations, but is especially evident in the model's cloud microphysical formulation. For example, consider the snow-graupel collection equation

$$\frac{dq_s}{dt} = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\rho \pi^2 N_{0s} N_{0g} E_{gs} \rho_s |V_g - V_s| \left(\frac{5.0}{\lambda_s^6 \lambda_g} + \frac{2.0}{\lambda_s^5 \lambda_g^2} + \frac{0.5}{\lambda_s^4 \lambda_g^3} \right) \right]. \quad (3.4)$$

Nearly all of the parameters of interest are present in this equation, and most are nonlinearly related to changes in the snow mass mixing ratio.

The range of values for each parameter is based on a number of observational studies found in the literature. The snow and graupel fallspeed parameter ranges are obtained from the detailed observational study of Locatelli and Hobbs (1974), while slope-intercepts of snow and graupel were obtained from both Locatelli and Hobbs (1974) and from Heymsfield et al. (2002). The range of values of the slope intercept of the rain drop size distribution was taken from a surface-based observational study of TOGA-COARE precipitating systems (Tokay and Short 1996), as well as from a surface-based study of tropical convection over coastal India (Roy et al. 2005).

Definition of a realistic range of values for each parameter effectively sets bounds on the MCMC search space, as the algorithm is expressly forbidden to set parameters equal to values outside the set ranges. Provided these ranges are set to physically realistic limits, MCMC will correctly characterize the parameter PDF. It is always the safer option to set the limits large, as the MCMC algorithm would return the same PDFs for a larger range, though the search would proceed much less efficiently.

It should be noted that setting realistic upper and lower bounds for each parameter can help in determining whether the model’s behavior is consistent with reality; if the bulk of the mass of any of the marginal parameter PDFs is located directly adjacent to one of the preset bounds, this may indicate non-physical performance of the model.

The remainder of this chapter is organized as follows: section 3.2 contains a description of the GCE model, including documentation of the configuration used in the MCMC experiments. Section 3.3 describes the simulated observations, and motivates a new way of assimilating cloud and precipitation on cloud resolving scales. MCMC setup and convergence diagnostics are described in section 3.4, results are presented in section 3.5, and a discussion of the limitations of this experiment, as well as its implications for the future of the cloud and precipitation assimilation problem is contained in section 3.6. Conclusions and suggestions for future work are offered in section 3.7.

3.2 GCE Model Description and Setup

The NASA Goddard Cumulus Ensemble model is a nonhydrostatic cloud system resolving model based on the work of Soong and Ogura (1980), Soong and Tao (1980), and Tao and Soong (1986). It employs either an anelastic or compressible solution to the atmospheric governing equations, and can be run in either two (one vertical and one horizontal) or three spatial dimensions. The model prognostic variables include potential temperature, perturbation pressure, turbulent kinetic energy, and the three Cartesian velocity components, as well as water vapor mass mixing ratio and the mixing ratios of two liquid and three ice condensate species. A second order accurate scheme is used to advect velocity components, while a multidimensional positive definite advection transport algorithm (Smolarkiewicz 1983, 1984, Smolarkiewicz and Grabowski 1990) is used to advect all scalar quantities. Forward time differencing is employed for scalar variables, while a time splitting leapfrog scheme is used for the

velocity.

The subgrid turbulence scheme is based on Deardorff (1975), Klemp and Wilhelmson (1978), and Soong and Ogura (1980), and includes the effect of condensation on the generation of subgrid-scale kinetic energy. Interaction of visible and infrared radiation with atmospheric trace gases, clouds, and the surface is parameterized according to Chou and Kouvaris 1991, Chou et al. 1995, Kratz et al. 1998, Chou et al. 1999, and Chou and Suarez 1999. Infrared fluxes are computed via a k-distribution method, while the solar radiation scheme uses a delta-Eddington approximation to compute two-stream radiative fluxes, and accounts for absorption by water vapor, CO_2 , O_2 , and O_3 , as well as the details of the cloud and ice particle size distribution. Surface fluxes are computed using the TOGA-COARE bulk flux algorithm (Fairall et al. 1996), and a 5 km deep Rayleigh relaxation (absorbing) layer is used at the top of the model to prevent reflection of vertically propagating gravity waves.

Though MCMC experiments are significantly more efficient in producing PDFs of key cloud microphysical parameters than an exhaustive search of the state space, there remain significant computational constraints. The most computationally expensive part of each MCMC iteration is the integration of the cloud resolving model, and though it would be preferable to perform each model run using a domain identical to that used in the climate response experiments, this is simply not feasible as each of these experiments takes approximately 10 days to finish. Given current computational resources, we require each GCE simulation to finish in approximately 3 minutes wall-clock. This allows us to obtain a reasonable sample of the PDF for the set of parameters described earlier within a wall-clock time of approximately three weeks. Though it is not possible to perform a full-domain GCE integration in this time, three minutes is fully sufficient to simulate the evolution of a convective ensemble on a limited domain. To the extent that this system is representative of the convective systems contained in large-domain simulations, it can be used as a surrogate for




Figure 3.1: Time series of five-day running mean precipitation rate during the SC-SMEX field campaign. Figure adapted from Johnson and Ciesielski 2005 (their figure 3).

convection on the larger domain.

The GCE configuration is identical to that employed in the NASA Goddard multi-scale modeling framework (W.-K. Tao, personal communication). In this configuration, the GCE runs on a 256 km horizontal domain with 4 km grid spacing and cyclic boundary conditions in the horizontal. 29 vertical levels are used, which stretch in spacing from 80 meters near the surface to approximately 1000 meters in the upper troposphere and lower stratosphere, and the model top is located at 23 km. The model is forced at every grid point and at every timestep with large-scale advective and thermodynamic tendencies of temperature and water vapor, as well as with large-scale horizontal winds. In the current study, the model is forced with a dataset generated from observations taken during the South China Sea Monsoon Experiment (SCSMEX) field campaign (Johnson and Ciesielski 2002, Johnson et al. 2005, Ciesielski and Johnson 2006). Deep convection during SCSMEX occurred during two distinct phases—a monsoon-onset phase from 16-22 May 1998, and a post-onset phase from 3-9 June 1998 (Fig. 3.1). Since the post-onset phase was more consistent

with conditions observed over the Tropical Western Pacific region (e.g., warmer sea surface temperature and more vigorous deep convection), the model is forced with observations taken during 5-6 June 1998.

A Hovmoller plot of cloud fields output from a control simulation run with default values of GCE microphysical parameters (Table 1.1) are depicted in figure 3.2, in which it can be seen that the model generates deep convection four hours into the simulation. Convection organizes into squall lines that propagate at an average speed of approximately $8 \text{ m} \cdot \text{s}^{-1}$, and it can be seen that the domain contains clouds of multiple types and depths. In each MCMC iteration, the simulation is allowed to spin up for 18 hours, and the convective system that evolves over the next 9 hours is “observed”. This system lies half in simulated sunlight and half in simulated darkness—simulation of daylight conditions is crucial for comparison with shortwave cloud radiative forcing and heating rate observations.

3.3 Observations

Cloud and precipitation assimilation efforts have traditionally been hampered by the fact that the goal of most assimilation projects is specification of a spatial pattern of state variables—either for the purpose of initializing a numerical weather prediction model or for retrospective analysis. The problem is that cloud and precipitation features exhibit extreme variability in both space and time, and this is particularly true in the case of convection. For realistic depiction of the role of convection in the climate system, it is patently unnecessary to simulate each convective element at the exact time and location it occurs in reality. Instead, it is important to correctly represent the spatially and temporally integrated effects of convective systems on the large scale environment. Because cloud microphysical parameters in the GCE




Figure 3.2: Hovmoller plot of the evolution of cloud fields in the GCE simulation for “standard” values of microphysical parameters. Color shading is potential temperature and indicates approximate cloud height from blue (low) to red (high). Precipitation rate is contoured in white every 5 mm/hour starting at 5 mm/hour, and the red horizontal line marks the beginning of the “observation” period.

Table 3.2: Observations used in both the MCMC parameter sensitivity experiment and in later assimilation experiments, along with their units and error estimates.

Observation	Units	Error
Precipitation Rate	$mm \cdot hr^{-1}$	20 %
Cloud Top Temperature	K	8 K
Liquid Water Path	$kg \cdot m^{-2}$	40 %
Ice Water Path	$kg \cdot m^{-2}$	100 %
High Cloud Fraction	%	5 %
BOA LW Cloud Forcing	$W \cdot m^{-2}$	20 $W \cdot m^{-2}$
TOA LW Cloud Forcing	$W \cdot m^{-2}$	20 $W \cdot m^{-2}$
BOA SW Cloud Forcing	$W \cdot m^{-2}$	50 $W \cdot m^{-2}$
TOA SW Cloud Forcing	$W \cdot m^{-2}$	50 $W \cdot m^{-2}$
Three-Layer All-Sky LW and SW Radiative Cooling/Heating	$K \cdot hr^{-1}$	50 %
Three-Layer Clear-Sky LW and SW Radiative Cooling/Heating	$K \cdot hr^{-1}$	15 %

model are assumed invariant with space and time,¹² changes in these parameters have an integrated effect on the simulation results, and it is therefore necessary to use observations that are integrated over space and time to constrain these parameters. In this case, it is the *statistics* of the cloudy state, not the individual cloud elements themselves that will be most useful in constraining the parameters of interest.

Hence, in choosing a suitable set of candidate observations for use in the assimilation experiments, two guiding criteria have been used. First, observations must (in theory) be sensitive to changes in liquid and ice particle size distribution and (in the case of ice) habit. Second, they must be readily available from current observing systems. Since the convective systems of interest are predominantly located over the ocean, observations will necessarily be derived from remote-sensing platforms.

¹²Note, there is no a priori reason to expect that the parameters that define the particle size distributions and crystal habit should remain fixed in space and time. Though examination of the case-to-case variability in the model's cloud microphysical parameters is beyond the scope of this study, an investigation of the degree to which parameters are non-stationary is underway.

Finally, because the feedback mechanisms under investigation in this study depend critically on precipitation efficiency (and associated net latent heat release) and radiative effectiveness of convective systems, observations must also be tied to those characteristics of clouds that have an influence on climate. With the above criteria in mind, the 11 observed fields listed in Table 3.2 have been selected.¹³ These observations are derived from a multi-sensor retrieval algorithm that combines information from instruments aboard the Tropical Rainfall Measuring Mission (TRMM) satellite (L’Ecuyer and Stephens 2003).¹⁴ This particular dataset is well-suited to the task at hand as it encompasses observations of clouds, precipitation, and radiative properties. In addition, rigorous uncertainty estimates are attached to each retrieved field; as was demonstrated in Chapter 2, correct characterization of observation errors is crucial if MCMC is to produce the correct PDF. For the observations used, the error estimates represent the best available knowledge of the uncertainty contained in both the measurements themselves and in the associated retrieval procedure.

As in the MODIS proof-of-concept in Chapter 2, all observation errors were assumed to be zero-mean random variables distributed Gaussian so that the MCMC likelihood function is:

$$L(\mathbf{y}|\mathbf{x}) = \sum_{i=1}^N \exp \left\{ -\frac{1}{2} \frac{[y_i - F(x)]^2}{\sigma_y^2} \right\}, \quad (3.5)$$

¹³Note that this list actually contains 21 total fields because radiative heating rates are separated into three layers for both long and shortwave radiation. Low, mid, and high layers were defined as sfc - 743 mb, 743 - 523 mb, and 523 mb - TOA, respectively. GCE model layers were chosen based on which layers contained the middle-layer boundaries, and layers containing boundary levels were split between the upper and lower radiative heating layers.

¹⁴Modern operational data assimilation rarely uses retrieved products, instead opting to assimilate the instrument radiances themselves and implement a suitable forward radiative transfer model within the assimilation procedure. To the extent that the radiative transfer model used in the assimilation may be less accurate than that used in the retrieval, assimilation of radiances may actually produce a less accurate result. The most important consideration is accurate characterization of the observation errors, and in this case, the retrieval error estimates are robust. In addition, it will be shown in section 3.5 that use of retrieved products facilitates a much more linear relationship between the control parameters and the observations, making it more feasible to use a variational assimilation procedure once real observations are considered.

where N is the number of observations, $\mathbf{y}$ is the vector of observations, $\mathbf{x}$ is the vector of microphysical parameters, $F(\mathbf{x})$ is the model that relates the parameters to the observations (in this case the GCE model itself), and σ_{y_i} is the error standard deviation for observation y_i . It should be noted that the error estimates assigned to some of the most relevant observations (e.g., ice water path) are very large, and hence the value of the observations for constraining the parameters of interest is reduced. This problem will be further addressed in section 3.6, in which an examination of the potential benefit of more accurate observations is presented. In addition, it must be pointed out that if the goal were simply to assess the sensitivity of the simulated output to changes in the microphysical parameters, we would define observations that are identically the fields we are interested in examining and let the MCMC algorithm produce the appropriate PDFs. However, the objective of this experiment is to determine *both* the simulation sensitivity to parameter changes *and* identify which observations should be used in the assimilation. The combination of use of observations that are important in the climate system with realistic error estimates accomplishes this goal.

3.4 MCMC Setup and Convergence Diagnostics

In contrast to the MODIS retrieval (Ch. 2), which utilized a single chain per pixel, the GCE MCMC procedure was run in a set of 60 Markov chains on 60 processors of an in-house Linux cluster. Because each chain theoretically converges to sampling the same overall PDF, this is a more efficient way of running MCMC; after the initial burn-in phase the chains can be combined to yield the full sample of the PDF. However, in contrast to running a single chain, there are additional considerations that must be taken into account when multiple chains are employed. First, it must be demonstrated that the chains have sufficiently *mixed*; each chain must converge to sampling the same

Table 3.3: Evidence of convergence of the MCMC algorithm to sampling a stationary distribution. See text for details.

Parameter	Mean	Variance	Variance between subsamples	R-Statistic
a_s	0.760	0.018	1.69E-06	1.016
b_s	0.489	0.038	2.56E-06	1.014
a_g	1.022	0.065	5.76E-06	1.019
b_g	0.413	0.023	2.25E-06	1.013
N_{0r}	2.403	1.325	1.19E-04	1.015
N_{0s}	2.342	1.726	1.74E-04	1.018
N_{0g}	2.757	1.540	1.77E-04	1.010
ρ_s	0.531	0.057	6.25E-06	1.013
ρ_g	0.548	0.053	5.29E-06	1.023

overall PDF. Second, though each chain can be started at the identical point in state space (identical starting values for the set of uncertain parameters), it is generally more optimal to distribute the chain starting points relatively evenly in the parameter space. In the current case, a multidimensional “cross” pattern has been used, with chains evenly distributed around the mid-point of each parameter’s allowable range, and with each “arm” of the cross consisting of three points—the center point, and two points distributed evenly between the center and the minimum/maximum allowable parameter value. Construction of this pattern accounts for 19 of the chains (9×2 plus the center point)—another 10 are used to construct “diagonals” in the state space, and the remaining chains are initialized with randomly-chosen parameter values drawn from within their acceptable ranges. Initializing the MCMC procedure with widely distributed (*overdispersed*) initial points has the benefit of more quickly identifying whether there is a significant probability maximum well-separated from the primary region of probability—if this is the case, one or more chains that are started nearby will find this region and will thoroughly sample it (Gelman et al 2004, p. 295-296).

After three weeks of wall-clock time, the MCMC algorithm yielded approximately 10^6 samples of the joint conditional parameter PDF. Convergence of the algorithm

was assessed via two methods—by monitoring the value of the *R-statistic*, and by comparing the moments of multiple successive 10,000-element random draws from the full sample. The R-statistic is predicated on the notion that an unbiased estimate of the true variance for each parameter can be obtained from a weighted combination of the between- and within-chain variances. The result is an overestimate, but converges to the true variance as the length of each chain goes to infinity. $\hat{R}$ is an estimate of the factor by which the variance estimate would be reduced if each chain were to be run to an infinite length. $\hat{R}$ close to 1.0 ($\hat{R} \leq 1.1$, Gelman et al. 2004, p. 297) indicates sufficient chain mixing.¹⁵ The mean and variance of each parameter within the set of subsamples is listed in the first two columns of Table 3.3, while the variance between sub-sample means is listed in the third column. The fact that the variance between subsamples is uniformly at least three orders of magnitude smaller than the variance within the samples themselves indicates convergence to a stationary sample. In addition, the set of skewness and kurtosis values (not shown) from the 100 subsamples were nearly identical. The R-statistic provides further evidence that the multiple chains have mixed as each parameter exhibits an R-value that is very close to 1.0. From these diagnostics, it is evident that MCMC has generated a well-mixed sample of the parameter PDF, and we now proceed with analysis of the characteristics of the parameter distribution.

¹⁵See appendix C for a detailed description and derivation of the R-statistic.




Figure 3.3: Marginal PDFs for each estimated parameter as represented in parameter histograms. The red horizontal line on each plot indicates the Uniform a priori PDF, while the vertical dashed line depicts the true parameter value.

3.5 Results

3.5.1 Parameter PDFs

The marginal PDFs for each of the nine parameters of interest obtained from the MCMC algorithm are depicted in figure 3.3. As in the MODIS proof of concept, a first indication of the influence of the entire set of observations on parameter estimates can be estimated via the departure of the PDFs from Uniform. However, in this case, we can also draw the inverse conclusion—the greater the departure from Uniform, the more influence the parameter has on the simulation output. It can be immediately seen that the snow fallspeed exponent (b_s , Fig. 3.3b), graupel fallspeed coefficient (a_g , Fig. 3.3c), and the rain slope intercept (N_{0r} , Fig. 3.3e) are all well-estimated by MCMC. PDFs of these parameters exhibit marked departure from the Uniform distribution with a well-defined mode at or near the true parameter value.¹⁶ By contrast, PDFs of the snow fallspeed coefficient (a_s , Fig. 3.3a), the snow size intercept (N_{0s} , Fig. 3.3e), and the snow and graupel densities (ρ_s and ρ_g , Figs. 3.3h and 3.3i, respectively) all appear relatively flat, with the bulk of the probability concentrated around the center of the range. Although PDFs of N_{0s} and ρ_s are much broader than those of b_s , a_g , and N_{0r} , the fact that the modes of their PDFs are located at the truth value indicate the observations may have some ability to constrain these parameters. It is interesting to note that the graupel fallspeed exponent (Fig. 3.3d) and slope intercept (Fig. 3.3g) appear to be relatively well-estimated, but exhibit PDFs with a mode that is well-displaced from the truth value. While this result is initially rather disturbing, it can be shown that this in fact represents a secondary mode in the joint PDF. Implications of this secondary mode are discussed in more detail below. From the marginal PDFs, we may draw the initial conclusion that the snow fallspeed exponent, graupel fallspeed coefficient and slope intercept of the rain

¹⁶The fact that the PDF of N_{0r} is distributed log-normal is interesting, as observed PDFs of N_{0r} (Tokay and Short 1996) have been found to be log-normal as well.

particle size distribution are the parameters which most significantly affect simulation results.




Figure 3.4: Comparison of PDFs for subset of the parameters corresponding to secondary mode, full parameter set, and primary mode.

It is evident from the marginal PDFs of b_g and N_{0g} that a secondary mode exists in the joint parameter PDF, and that a significant portion of the mass of the b_g and N_{0g} PDFs is concentrated in this mode. To assess the extent to which other parameters may be included in this region of the space, PDFs are generated for two

different parameter subsets. The first subset, which we denote MODE2, consists of all parameter sets for which $b_g \sim [0.33, 0.66]$, and $N_{0g} \sim [2.5, 4.5]$ —this subset restricts the sample to parameter sets that are located in the region of the state space that includes the secondary mode. The second subset, which we denote MODE1, consists of the parameter subspace not included in MODE2. Examination of the marginal PDFs for each subset revealed that the PDFs of b_s , a_g , N_{0r} , and ρ_s were left virtually unchanged regardless of which subset was plotted, which indicates that the secondary mode exists in a region of the parameter space within the primary mode of these four parameters. PDFs of the remaining three parameters (a_s , N_{0s} , and ρ_g) for MODE1, MODE2, and the full sample of parameter sets is plotted in figure 3.4. As the sample is shifted into the subspace occupied by the secondary mode, the mass of each PDF shifts toward higher parameter values, and we may conclude that a_s , N_{0s} , and ρ_g must shift toward larger values along with b_g and N_{0g} to generate the a state that is approximately compatible with the observations. Conversely, when the sample is shifted toward the region of the space containing the true parameter values, the PDFs of a_s , N_{0s} , and ρ_g shift toward their truth values as well. Two conclusions may be inferred: first, that the snow and graupel cloud microphysical parameters are coupled in such a way that changes to one set of parameters necessitates changes to the others, and second, if the secondary mode can be avoided, it may be possible to constrain a_s , N_{0s} , and ρ_g using observations.

A further indication of the relative importance of each parameter, as well as a first indication of the correlations between parameters, can be obtained from the joint PDF of each pair of parameters (figure 3.5). An examination of this figure reveals correlations between snow fallspeed coefficient and exponent, graupel coefficient and exponent, and between the graupel density and slope intercept. These correlations could easily have been anticipated in advance from an examination of the cloud microphysical scheme itself (appendix A), and it is reassuring to see them




Figure 3.5: Joint PDFs for each pair of parameters. The color scale is the same for each plot with warmer colors representing higher values of the PDF. Each axis is labelled with its corresponding parameter.

depicted in the data. Additional correlations can be seen between the graupel fall-speed coefficient and the slope intercept of the rain particle size distribution. This result reflects the strong relationship between graupel and precipitation in this particular cloud microphysical formulation. Figure 3.5 provides further evidence of the fact that the snow fallspeed coefficient, graupel fallspeed exponent, snow slope intercept, and graupel density parameters have little influence over the result of the simulation and that we cannot draw any firm conclusion as to the parameter truth values from either marginal or joint PDFs. By contrast, the three parameters that do contribute significantly to the simulation results can be clearly identified from their joint PDFs.

Table 3.4: List of observations vs. mean states for the top 0.1 % of all likelihood values.

State	Observation	Mean of States
Precipitation Rate	7.211	7.234
Cloud Top Temperature	246.509	246.446
Liquid Water Path	799.175	830.366
Ice Water Path	631.586	624.500
High Cloud Fraction	0.729	0.729
BOA LW Cloud Forcing	-15.499	-14.590
TOA LW Cloud Forcing	-33.665	-33.623
BOA SW Cloud Forcing	121.407	125.078
TOA SW Cloud Forcing	110.409	113.477
Low Level All-Sky LW Cooling	-1.294	-1.319
Mid Level All-Sky LW Cooling	-2.389	-2.275
Upper Level All-Sky LW Cooling	-1.916	-1.939
Low Level Clear-Sky LW Cooling	-2.142	-2.148
Mid Level Clear-Sky LW Cooling	-2.756	-2.752
Upper Level Clear-Sky LW Cooling	-1.621	-1.619
Low Level All-Sky SW Heating	-0.929	-0.915
Mid Level All-Sky SW Heating	-1.194	-1.202
Upper Level All-Sky SW Heating	-1.181	-1.192
Low Level Clear-Sky SW Heating	-0.892	-0.891
Mid Level Clear-Sky SW Heating	-1.167	-1.171
Upper Level Clear-Sky SW Heating	-0.993	-0.990

As a final assessment of the robustness of each parameter PDF, we examine the

set of parameter values that correspond to the highest 0.1% of the likelihood values; these parameters represent those that produced the *best fit* states when compared with simulated observations. If the values of a parameter that generated large likelihood values are concentrated around its true value, we can conclude that this particular parameter provided more significant constraint on the observations. By contrast, if the parameter values are distributed evenly across the allowable range, the implication is that nearly any parameter value could have produced a near-optimal state. Mean state values for the top 0.1% of the likelihoods are listed alongside the observation values in Table 3.4, and it is clear that the parameters associated with the very highest total likelihood values did indeed generate state values that were very close to those observed.

Examination of scatter plots of each parameter vs. the total likelihood (Fig. 3.6) reveals, as anticipated, that b_s , a_g , and N_{0r} exhibit a tight grouping about their truth parameter values, while a_s , b_g , N_{0s} , and N_{0g} , and ρ_s display widely scattered parameter values for even the highest likelihood values. It is interesting to note that ρ_g is nearly as tightly grouped as are a_g and b_s —a first indication that ρ_g exerts greater control over the simulation results than its marginal PDF indicates. It is interesting also to note that the secondary mode is missing from a_s and b_g , and is greatly reduced for N_{0g} , an indication that this particular mode constitutes a local probability maximum in the state space.

It should be noted that care must be taken when interpreting the results of MCMC through the PDFs themselves, as a “significant” change in output vs. change in parameter may be masked by the error assigned to a given observation. For example, a $10 \text{ W} \cdot \text{m}^{-2}$ change in spatially and temporally averaged shortwave cloud radiative forcing would be considered significant, even though it may not affect the parameter histograms because of the $50 \text{ W} \cdot \text{m}^{-2}$ error assigned to this observation. Similar concerns exist with respect to the liquid and ice water paths, as assumed errors in these




Figure 3.6: Scatter plots of the largest 0.1% of likelihood values for each parameter.

quantities are also large. The issue of large observation errors will be addressed further in section 3.6. An examination of the potential ability of observations with currently realistic error characteristics to constrain the parameter set of interest follows.

3.5.2 Assessment of the utility of individual observations for assimilation

The parameter histograms obtained from MCMC provide a first indication of how well the full set of observations constrain each given parameter within the set, as well as hinting at the general influence changes in each parameter have on the set of chosen observations. However, it remains to (1) identify which observations contribute most to constraining the set of parameters (e.g., which should be used as the focus of the assimilation), and (2) quantify exactly how much influence a change in each parameter has on any given observation. Recall that the error model for each observation was assumed to be Gaussian (Eqn 3.5), and that the errors in each observation were assumed to be independent of one another. Hence, just as we examined the marginal PDF of each parameter $P(x_i|\mathbf{y})$ conditioned on the set of observations, we can also examine the PDF of each observation conditioned on the set of parameters $P(y_i|\mathbf{x})$. It follows that a first cut at the influence of each observation on the set of given parameters can be obtained through an examination of the PDFs of the observations themselves. The width of the observation histograms (Fig. 3.7) is directly related to the effective range of variability of each observed variable produced by the parameters sample, which in-turn indicates the level of control the set of cloud microphysical parameters has over each observation. The analysis must be carefully interpreted—if a given observation consistently has a narrow PDF, we might be tempted to assume it had a consistently large impact on the parameter estimates. However, this is only true if the PDF is narrow compared to the assumed observation error. In addition, a consistently narrow PDF means that for nearly *any* set of parameter values, the




Figure 3.7: Marginal PDFs of the set of observations used in the MCMC experiment. In each plot, the range of values of each observed variable is depicted on the abscissa, while the normalized number of counts is displayed on the ordinate. The plots are subdivided into (a) cloud and precipitation observations, (b) long and shortwave cloud radiative forcings, (c) three-layer longwave cooling rates, and (d) three-layer shortwave heating rates.

observation maintained a nearly identical value; changes in the set of uncertain parameters had virtually no influence on this observation. In contrast, if an observation exhibits a wide *range* of values, it means that changes in the set of parameters had an influence on the value of this observation, and hence it may be a desirable candidate for the assimilation.

Table 3.5: List of observation error standard deviations vs. the standard deviation of each simulated state variable.

State	Obs Error Std Dev	Std Dev of States
Precipitation Rate	1.442	0.822
Cloud Top Temperature	8.000	3.374
Liquid Water Path	319.670	85.517
Ice Water Path	631.586	109.208
High Cloud Fraction	0.050	0.005
BOA LW Cloud Forcing	20.000	1.473
TOA LW Cloud Forcing	20.000	3.253
BOA SW Cloud Forcing	50.000	16.534
TOA SW Cloud Forcing	50.000	14.558
Low Level All-Sky LW Cooling	0.647	0.094
Mid Level All-Sky LW Cooling	1.194	0.188
Upper Level All-Sky LW Cooling	0.958	0.071
Low Level All-Sky SW Heating	0.321	0.035
Mid Level All-Sky SW Heating	0.413	0.026
Upper Level All-Sky SW Heating	0.243	0.005
Low Level Clear-Sky LW Cooling	0.464	0.083
Mid Level Clear-Sky LW Cooling	0.597	0.110
Upper Level Clear-Sky LW Cooling	0.591	0.037
Low Level Clear-Sky SW Heating	0.134	0.005
Mid Level Clear-Sky SW Heating	0.175	0.008
Upper Level Clear-Sky SW Heating	0.149	0.004

An examination of figure 3.7 reveals that many of the PDFs exhibit a mode that is shifted compared to the observation truth value. This may be due to the existence of the secondary mode in the parameter PDF, however, when observations are sub-setted according to different regions of the parameter space (not shown), the modes

do not shift. In each case, as demonstrated in Table 3.5, the range of variability of each predicted state variable was well-within its assigned observation error standard deviation. If the observation error had been assumed to be smaller, it is likely that the observation modes would lie much closer to their true values. In addition to the modes being displaced from the true observation values, nearly all of the observation PDFs exhibit some degree of skewness. This is particularly interesting, as the observation errors were uniformly assumed to be distributed Gaussian with zero bias. Skewness in the PDFs is a reflection of the fact that (1) MCMC is not restricted to return Gaussian, or even unimodal PDFs, and (2) parameters and state variables are related in a (sometimes highly) nonlinear fashion. The degree of nonlinearity in the relationship is reflected in the degree of skewness—for example, if we consider the three key parameters identified in the previous section (b_s , a_g , and N_{0r}), the precipitation, BOA SW cloud radiative forcing, and low-level all-sky long and shortwave radiative heating, which are all closely related to the properties of liquid clouds, exhibit very little skewness. By contrast, the liquid water path is a function not only of the rain drop size distribution, but is also tied to ice phase processes—especially to the melting of graupel. Similarly, the mean cloud top temperature is a complicated function of cloud dynamics, the thermodynamic state of the environment, and the distribution of cloud types in the domain, while the high cloud fraction is related much more directly to the fractional coverage of ice clouds. The mid and upper level all-sky LW radiative heating rates likely exhibit skewness because they are strongly affected by the properties of cirrus clouds, which are highly nonlinearly related to the ice particle size distribution and habit. Skewness in the low and mid level all-sky shortwave heating rate PDFs is likely due to the complicated interaction between scattering and the ice habit and particle size distribution. All of the clear-sky radiative heating rates are skewed, likely because the properties of the clear air radiative heating are determined by the interaction between the fractional area of clear air and the strength of

subsidence associated with convection—two properties that are themselves related.

3.5.3 Assessment of Observation/Parameter Relationships

Though the observation histograms provide a preliminary indication of the relative usefulness of each observation, it must be noted that we can draw no inference whatsoever as to the *relationship* between individual parameters and observations from the observation PDFs alone. The fundamental problem is that the assumed error determines the effective information content of each observation; a significant change in the value of a given observation may correspond to little change in its *likelihood* if the assumed error is large. Because the likelihood function is analogous to the observational component of the variational cost function, it contributes valuable information as to which observations might be best to assimilate, however, it does not indicate which parameters are most important to constrain in the assimilation; those that have the greatest effect on the properties of the climate. If we are to assess which observations are most directly affected by changes in particular parameter values, and identify which might be most useful in a variational assimilation context, a more quantitative estimate of the relationship between observations and parameters must be demonstrated. To this end, the most relevant observation/parameter pairs are selected in a multi-part procedure that takes advantage of the large database of simulations generated by MCMC, as well as the fact that for this idealized experiment we know the true values of each parameter. The most significant and linear relationships are culled from the database as follows:

1. First, the standard deviation of each parameter is computed from the MCMC database of parameter values
2. Next, for each observation/parameter pair, observations are subsetting into those that lie within $\pm 1\sigma$ of the truth values for all parameters except the parameter

of interest, for which the entire range is allowed

3. For each observation subset, the autocorrelation between parameters is computed and used to obtain an estimate of the number of independent samples in the data subset. The number of independent samples is used to compute the critical value of the correlation necessary to achieve a statistical significance of 99% via the student's T-test. The critical correlation is computed as

$$\rho = \frac{t}{\sqrt{t^2 + (N - 2)}}, \quad (3.6)$$

where ρ is the correlation coefficient, t is the critical t-statistic value, and N is the number of degrees of freedom in the dataset.

4. Each subset of observations is linearly regressed onto the corresponding parameter values and the correlation coefficient is computed.
5. Finally, the *signal to noise* ratio of the observations is computed by considering the *signal* to be the linear change in observation value over the range of parameter values returned from the linear regression, while the *noise* is the standard deviation of the subset of observations themselves.

The set of observation/parameter pairs chosen for the assimilation are those that meet the criteria of having a 99% significant correlation and signal to noise ratio greater than approximately 1.5. In this way, the 198 possible parameter-observation combinations are reduced to the most significant 23. Of these 23, 8 exhibit both signal to noise greater than 2.0 *and* signal greater than observation error standard deviation. Of the remaining 15 pairs, 6 have signal less than the observation error standard deviation—these are retained because the observations still provide partial information on the parameter values.

Table 3.6: List of the parameter/observation pairs with the most significant linear relationships, as represented in their linear correlation and signal to noise ratio (see text for description).

Parameter	State Variable	ρ	SNR
a_g	Sfc-743 mb SW Clear-Sky Radiative Heating	-0.525	3.123
N_{0r}	Sfc-743 mb LW Clear-Sky Radiative Cooling	0.683	3.039
b_s	Cloud Top Temperature	-0.640	2.930
N_{0r}	BOA LW Cloud Radiative Forcing	-0.649	2.885
N_{0r}	743-523 mb LW Clear-Air Radiative Cooling	-0.620	2.758
ρ_g	523 mb-TOA SW All-Sky Radiative Heating	-0.573	2.335
a_g	743-523 mb LW All-Sky Radiative Cooling	-0.377	2.245
ρ_g	743-523 mb LW All-Sky Radiative Cooling	0.533	2.175
a_g	Ice Water Path	-0.357	2.123
a_g	523 mb-TOA LW All-Sky Radiative Cooling	0.346	2.059
N_{0r}	Sfc-743 mb SW Clear-Sky Radiative Heating	0.461	2.051
ρ_g	TOA LW Cloud Radiative Forcing	-0.483	1.967
N_{0r}	BOA SW Cloud Radiative Forcing	0.442	1.965
N_{0r}	TOA SW Cloud Radiative Forcing	0.423	1.883
N_{0r}	743-523 mb SW Clear-Sky Radiative Heating	-0.414	1.839
a_g	743-523 mb SW Clear-Sky Radiative Heating	0.306	1.820
ρ_g	743-523 mb SW All-Sky Radiative Heating	0.437	1.783
a_g	High Cloud Fraction	-0.296	1.763
ρ_g	523 mb-TOA LW All-Sky Radiative Cooling	-0.430	1.752
a_g	Sfc-743 mb LW Clear-Sky Radiative Cooling	-0.250	1.489
b_s	523 mb-TOA LW All-Sky Radiative Cooling	0.324	1.484
N_{0g}	TOA LW Cloud Radiative Forcing	0.357	1.459
a_s	Cloud Top Temperature	0.346	1.450




Figure 3.8: Relationship between a_g and observations.




Figure 3.9: Relationship between N_{Or} and observations.




Figure 3.10: Relationship between b_s and observations.

An examination of the results of the above procedure (Table 3.6) reveals that there were only three parameters that exhibited consistently significant relationships with the observations. As expected, two of the three parameters with histograms significantly different from Uniform (a_g and N_{0r} , figures 3.8 and 3.9, respectively) were included in this list—the snow fallspeed exponent (Fig. 3.10) was only found to be significantly linearly related to the cloud top temperature and the upper level all-sky longwave radiative cooling. The fact that the snow fallspeed coefficient (Fig. 3.11) exhibited a flat PDF in contrast to the fallspeed exponent reflects the fact that there is greater information content in the longwave cloud radiative forcing, as both parameters were significantly related to the cloud top temperature. However, the TOA LW cloud radiative forcing did not contain enough information by itself to constrain the graupel fallspeed exponent (Fig. 3.13).

It is interesting to note that the graupel density exhibited consistent and significant relationships with four of the all-sky radiative heating observations and with the TOA longwave cloud radiative forcing. The fact that the PDF of the graupel density did not significantly differ from PDFs of parameters that had little discernible relationship with any of the observations was likely due to the fact the graupel density was correlated with observations that had large assumed errors. The implication is




Figure 3.11: Relationship between a_s and observations.




Figure 3.12: Relationship between ρ_g and observations.




Figure 3.13: Relationship between N_{0g} and observations.

that changes to the graupel density have a significant effect on the radiation state variables, and that more accurate cloud radiative forcing and radiative heating observations will be required to constrain it.

The fact that few of the observations of cloud and precipitation variables exhibited significant relationships with any of the cloud microphysical parameters is a bit surprising. There are two possible reasons for this: (1) cloud physical and dynamic properties are strongly nonlinearly coupled; changes to cloud microphysical parameters and to the associated cloud and ice mass feed back to the cloud updraft and downdraft dynamics in a way that causes any relationship to be obscured. (2) The simulations were run over a very short period of time on a limited horizontal domain. Systematic relationships would most likely be revealed for longer simulations on a

much larger domain. In fact, evidence for this can be seen in figure 1.2, in which a systematic change in ice water path is demonstrated for a long-term, large-domain simulation. In contrast to the cloud and precipitation properties, the time and domain integrated radiative properties are directly related to the habit and size distribution of liquid and ice condensate, and are not as directly affected by the transient cloud dynamics. The fact that changes in the *radiative* properties are more directly related to changes in cloud microphysical parameters serves to underscore the motivation to constrain the uncertainty in these parameters in simulations of long-term climate variability.

3.6 Discussion

As was demonstrated in the comparison of control and double error cases in the MODIS retrieval example in chapter 2, assumptions about the magnitude of the observation errors used in the MCMC algorithm have a profound effect on the resulting parameter histograms. In the MODIS study, when the assumed observation error was doubled, the parameter PDFs uniformly broadened, though the maximum likelihood point was unmodified. The fact that many of the GCE parameter histograms are rather flat may simply reflect the fact that relatively large errors were assigned to the observations. Though the chosen error specifications were necessary to conform with realistic and current observations, it is interesting to consider how parameter PDFs might change as a result of improvements to current observing systems and in the corresponding accuracy of retrieved products. To make this assessment, the MCMC procedure is repeated using a set of “optimal” observation errors (Table 3.7). These errors are specified consistent with assumptions about the next generation of remote-sensing instruments and platforms; CloudSat cloud radar retrieval of liquid water path, Cloud-Aerosol Lidar and Infrared Pathfinder Satellite

Observation (CALIPSO) lidar identification of cloud top height, Global Precipitation Measurement (GPM) retrieval of precipitation rate, and Submillimeter Infrared Ice Cloud Experiment (SIRICE) high-frequency microwave measurements of ice water path. Because the cloud radiative forcing and radiative heating rates appeared to contribute most to constraining the GCE cloud microphysical parameters in the control experiment, only observations of cloud top temperature, liquid water path, and ice water path are used in the following MCMC analysis. One of the objectives of this set of experiments is to determine whether cloud microphysical parameters can be constrained using only the integrated cloud observations if the errors are reduced.

Table 3.7: List of TRMM retrieved state variables with revised error estimates based on anticipated future sensors.

Observation	Units	Error
Precipitation Rate	$mm \cdot hr^{-1}$	10 %
Cloud Top Temperature	K	2 K
Liquid Water Path	$kg \cdot m^{-2}$	10 %
Ice Water Path	$kg \cdot m^{-2}$	25 %
High Cloud Fraction	%	2 %
BOA LW Cloud Forcing	$W \cdot m^{-2}$	5 $W \cdot m^{-2}$
TOA LW Cloud Forcing	$W \cdot m^{-2}$	5 $W \cdot m^{-2}$
BOA SW Cloud Forcing	$W \cdot m^{-2}$	10 $W \cdot m^{-2}$
TOA SW Cloud Forcing	$W \cdot m^{-2}$	10 $W \cdot m^{-2}$
Three-Layer All-Sky LW and SW Radiative Cooling/Heating	$K \cdot hr^{-1}$	10 %
Three-Layer Clear-Sky LW and SW Radiative Cooling/Heating	$K \cdot hr^{-1}$	5 %

In the idealized observation error experiment, the MCMC procedure has been run in an identical manner as in the control experiment, with an adaptive burnin period at the beginning and a bounded Uniform proposal distribution. As in the experiments run with current estimates of retrieval errors, the algorithm quickly converges to a 50% acceptance rate, indicating a balance of thorough sampling and appropri-

ate observation information content. It is interesting to note the dramatic increase in observation information content implied in this result—50% acceptance rates are obtained in the control experiment using *twenty-one* observations, while a nearly identical acceptance rate is obtained for improved errors using only *three*. As with the control simulation, after burnin, experiments were run until approximately 10^6 samples were collected. Convergence was assessed through examination of the R statistic, and through inspection of parameter timeseries and running computation of moments of the parameter PDFs (not shown). In contrast to the control experiment, which converged to sampling a stationary distribution after approximately 5,000 iterations, convergence was reached for the optimal observation experiment after only 3,500.

Marginal PDFs of each parameter (Fig. 3.14) indicate that, as in the control experiments, the three most significant parameters are a_g , b_s , and N_{0r} . Comparison with control PDFs (Fig. 3.3) indicates that the PDFs that result from the use of improved observations are more sharply peaked; note in particular that the range of values of N_{0r} on the ordinate in figure 3.14e is double that of figure 3.3e. In other respects, the control and optimal error PDFs are quite similar— ρ_g , ρ_s , and N_{0g} are all relatively flat, while N_{0s} exhibits a slight mode around its true value. Though the secondary mode observed in the control experiment is not as pronounced in the PDF of N_{0g} , it is still evident in the PDFs of a_s , b_g , and ρ_g , indicating that this secondary mode is not an artifact of the particular set of observations used, but is instead a real feature of the state space for this particular experiment. The high degree of similarity between PDFs obtained using optimal errors and those of the control experiment is quite interesting, and appears to indicate that, regardless of the observations used, there are just a few parameters that control the output of this particular cloud microphysical scheme.¹⁷

¹⁷It should be noted that the use of qualitatively different observations can result in very different parameter PDFs. Additional sensitivity tests (not shown), in which observations of the mean mass of each condensate species were used in the MCMC algorithm resulted in far sharper PDFs of N_{0s} and N_{0g} in particular.




Figure 3.14: Marginal PDFs for each estimated parameter, as represented in parameter histograms, obtained from the experiment using improved error estimates. As in figure 3.3, the red horizontal line on each plot indicates the Uniform a priori PDF, while the vertical dashed line depicts the true parameter value.

Scatterplots of parameter values for the top 0.1 percent of likelihood values are plotted in figure 3.15. While marginal PDFs for each parameter in this experiment were quite similar to those of the control simulation, the scatter plots are very different. First, the likelihood values for this experiment are generally higher than in the control; the top 0.1 percent of likelihoods for this case are concentrated above 0.98, while the maximum likelihood in the control experiment was equal to 0.98. The implication is that the algorithm was much more easily able to generate parameter sets that very closely matched the observations for the optimal observations case than for the control. Though the three most significant parameters do indeed exhibit parameter values tightly grouped around the truth values for the highest likelihoods, the scatter does not widen away from the truth with decreasing likelihood. Indeed, the scatter does not significantly increase even down to likelihood of 0.9 (not shown). In addition, whereas b_g and ρ_g exhibit slightly more concentrated parameter grouping for high likelihood in the control experiment, all parameters in the current experiment except for the most significant three are evenly distributed across their respective ranges. This is a clear indication that the three observations chosen *only* contribute information on a_g , b_s , and N_{or} , and contain no information as to the values of the other six parameters.

On a final note, it should be pointed out that use of a subset of observations to estimate parameters can provide useful diagnostic information on the operation of the microphysical scheme that is not available when the full set of observations is used. For example, if the microphysical scheme performs in a realistic manner, then the *precipitation* associated with the best estimate parameters that conform to the selected *cloud* observations should closely match the true precipitation rate. Examination of the set of states that correspond to the top 0.1% of likelihood values (Table 3.8) indicates that, as we would have expected, the best estimate liquid water path, ice water path, and cloud top temperature match the observations very closely,

as does the mean high cloud fraction. The mean precipitation rate does indeed closely match the precipitation “observed” in the model, indicating that the link between production of clouds and precipitation in the model is realistically simulated.

While it is not surprising that the GCE cloud microphysical scheme produced precipitation consistent with observations for estimated cloud content, it is interesting to consider whether the radiative observations are similarly well estimated. Examination of Table 3.8 reveals that the longwave bottom and top of atmosphere cloud radiative forcing are remarkably well-constrained, as are the clear-sky and all-sky longwave cooling rates and the clear-sky shortwave heating. In fact, the only state variables that are not consistent with the non-estimated observation are the shortwave cloud radiative forcing and all-sky shortwave heating rates. The general implication is that observations of cloud mass and cloud top can be effectively used to constrain the precipitation through constraint of cloud microphysical parameters. In addition, it appears that the cloud mass is closely related to the mean thermodynamic state, as the longwave fluxes and heating rates are effectively constrained when the cloud content is estimated. Though it is not straightforward to deconvolve the relationship between cloud observations and shortwave fluxes and heating rates, examination of the GCE shortwave radiation code (not shown) indicates that the ice cloud optical depth depends directly on the properties of the particle size distribution, and is calculated using the snow and graupel density and PSD slope intercept—parameters that are not well estimated from observations of cloud top temperature and cloud liquid and ice water content. It is likely that if better observations of outgoing shortwave radiation, shortwave cloud radiative flux, shortwave heating rates, or some combination of the above were used in the estimation, the PDFs of ice particle densities and slope intercepts might be more effectively characterized, and the shortwave states would conform more closely to observations.




Figure 3.15: Scatter plots of the largest 0.1% of likelihood values for each parameter for the experiment with improved error estimates.

Table 3.8: List of observations vs. mean states for the top 0.1 % of all likelihood values for the experiment with improved error estimates.

State	Observation	Mean of States
Precipitation Rate	21.368	20.066
Cloud Top Temperature	236.413	235.917
Liquid Water Path	537.714	531.241
Ice Water Path	448.355	415.538
High Cloud Fraction	0.693	0.709
BOA LW Cloud Forcing	-12.359	-12.661
TOA LW Cloud Forcing	-49.583	-51.552
BOA SW Cloud Forcing	116.523	144.813
TOA SW Cloud Forcing	107.998	134.272
Low Level All-Sky LW Cooling	-1.527	-1.446
Mid Level All-Sky LW Cooling	-2.029	-2.036
Upper Level All-Sky LW Cooling	-1.807	-1.814
Low Level Clear-Sky LW Cooling	-2.276	-2.253
Mid Level Clear-Sky LW Cooling	-2.705	-2.746
Upper Level Clear-Sky LW Cooling	-1.681	-1.637
Low Level All-Sky SW Heating	-0.769	-0.719
Mid Level All-Sky SW Heating	-1.227	-1.242
Upper Level All-Sky SW Heating	-1.175	-1.216
Low Level Clear-Sky SW Heating	-0.883	-0.857
Mid Level Clear-Sky SW Heating	-1.119	-1.124
Upper Level Clear-Sky SW Heating	-0.926	-0.918

3.7 Conclusions

The results presented in this and the previous chapter clearly demonstrate the robustness of the MCMC algorithm for characterizing the full joint conditional PDF of a set of unknown parameters. In the current application, MCMC was used to characterize the joint PDF of a set of uncertain cloud microphysical parameters given TRMM-based retrievals of cloud properties and radiative forcings and heating rates. Examination of marginal and joint PDFs revealed that GCE simulations were most sensitive to changes in the snow fallspeed exponent, the graupel fallspeed coefficient, and the slope intercept of the rain drop size distribution. The PDF also exhibited a secondary mode centered around higher values of the snow fallspeed coefficient, graupel fallspeed exponent, slope intercept, and density. Analysis of the set of parameters associated with the largest likelihood values showed this mode to be of lesser mass than the primary mode, which was concentrated around the parameter truth values.

Following the control MCMC experiment, a second experiment was performed, in which observation errors were decreased consistent with introduction of near-future sensors, and in which only observations of cloud top temperature, liquid water path, and ice water path were used. These “optimal error” experiments returned PDFs that were quite similar to those in the control, and identified the identical three key parameters. In addition, a secondary mode, consistent with the secondary mode seen in the control experiment, was observed in the parameter PDFs. Use of a subset of cloud observations in the estimation allowed for additional assessment of the robustness of cloud microphysics and radiation parameterizations. It was found that the surface precipitation rate was well constrained by cloud observations, as were the longwave fluxes and heating rates. Shortwave fluxes and (all-sky) heating rates were not as well constrained, likely because these state variables are sensitive to changes in the densities and slope intercepts of precipitating ice—parameters that were not well constrained by cloud and ice content observations.

The results presented here underscore the need for careful interpretation of the data returned from an inversion technique that explicitly involves a highly nonlinear model. Though variational techniques have significant limitations, as was demonstrated in chapter 2, they enable the use of powerful diagnostic tools with which the analysis can be effectively scrutinized. Because diagnostic tools based on linear (or locally linear) assumptions are not appropriate for use on a potentially-multimodal PDF that is nonlinearly related to observations, it is essential that the PDF be examined from a number of different perspectives. The current analysis employed rather crude methods, but other more sophisticated methods that attempt to characterize the PDF in terms of its multi-dimensional statistics are available. In particular, Gelman et al. (2004, p. 298) advocate augmenting the R-statistic with examination of extreme quantiles of the PDF to avoid potential problems with highly non-Gaussian distributions. Tarantola (2005) suggests use of a “movie” strategy in which multiple realizations (or sub-samples) of the PDF are viewed at the same time to let the eye identify common patterns in the data. The author is also currently investigating use of advanced visualization techniques with which a PDF might be visualized in more than six dimensions at once.

After performing MCMC experiments with the GCE model, the following parameters are identified as key control variables for the data assimilation procedure. The graupel fallspeed coefficient and density, snow fallspeed exponent, and the rain slope intercept all exhibited not only relatively sharply-peaked PDFs, but also significant correlations with important model output variables. Examination of the PDFs of the observations for the case in which realistic observation errors were used indicated that most of the observations related to radiative forcings and heating rates will be useful in constraining this subset of microphysical parameters. In addition, plots of the relationship between these variables and the set of control parameters demonstrated a near-linear relationship between many of the candidate observation/parameter pairs.

This is fortunate, as it may allow us to use an assimilation methodology that is predicated on quasi-linear relationships between observations and control variables. The chosen assimilation methodology, observations used in the assimilation, and results of assimilation of the chosen set of observations is presented in the following chapter.

Chapter 4

Assimilation of TRMM Multisensor Retrievals to Constrain GCE Microphysical Parameters

4.1 Introduction

In both traditional sensitivity experiments (chapter 1) and MCMC simulated observation experiments (chapter 3) it was shown that changes in GCE cloud microphysical parameters have a significant effect on the operation of the modeled tropical hydrologic cycle and on its interaction with the planetary radiative balance and large scale circulation. In addition, it was found that mean cloud, precipitation, and radiative properties were most affected by changes in a set of four key GCE cloud microphysical parameters—the graupel fallspeed coefficient, the snow fallspeed exponent, the slope intercept of the rain drop size distribution, and the graupel density. Moreover, it was found that these four parameters exhibited clear linear relationships with a subset of the observations such that changes to each of these parameters produced a consistent change in these observed fields across the range of parameter variability. The




Figure 4.1: Depiction of the geographical region involved in the SCSMEX field campaign. The retrievals used in this experiment were derived over the NESA region.

sensitivity of those characteristics of deep convection that are most closely related to climate response to surface warming, along with the demonstrated linear correlation between changes in parameters and response of the model state, provides motivation to use data assimilation methodologies to constrain GCE microphysical parameters using observations.

The observations to be used consist of TRMM retrievals of precipitation rate, cloud top temperature, liquid and ice water path, high cloud fraction, visible and infrared cloud radiative forcings, and three-layer radiative heating and cooling rates

accumulated over the South China Sea Monsoon Experiment (SCSMEX) Northern Enhanced Sounding Array (NESA) domain (Fig. 4.1) during 3-9 June 1998. As was discussed in chapter 3, propagating tropical squall lines characteristic of the active phase of the monsoon were observed in the SCSMEX domain during this time period, and it has been shown that a 2D version of the GCE model can be successfully used to simulate this convection (Tao et al. 2003). For consistency with the results obtained in the idealized MCMC experiments, the GCE configuration is identical to that used in chapter 3.

Because a robust linear correlation was demonstrated between key GCE microphysical parameters and TRMM retrievals, it is possible that quasi-linear (e.g., variational) optimization methods could be used to perform the assimilation. This is desirable, as MCMC PDF estimation is extremely computationally expensive making it difficult to run multiple different MCMC experiments, much less to regularly (or operationally) estimate parameter PDFs. In addition, if variational assimilation can be shown to effectively estimate cloud microphysical parameters in the current case, it may be possible to apply the results to operational assimilation of cloud, precipitation, and radiation statistics for model bias constraint.

In this chapter, use of variational methods will be explored, and the results compared with PDFs of GCE cloud microphysical parameters obtained from real-data MCMC experiments. It is anticipated that comparisons with a parallel MCMC estimate will allow us to assess the efficacy of the variational estimate in a manner similar to the MODIS retrieval in chapter 2, as well as providing an independent estimate of parameter values should variational assimilation prove unfeasible. In either case, the outcome of the assimilation experiments performed in this chapter will be a set of parameters that is most consistent with statistics of the tropical radiative-convective interaction.

The remainder of this chapter is organized as follows. Section 4.2 contains a de-

scription of the TRMM retrievals to be assimilated, as well as a brief comparison between observations and GCE state variables for a simulation run with default values of cloud microphysical parameters. Derivation of the variational framework from Bayes theorem, specifics of how it is applied to the current problem, and results of variational test experiments are presented in section 4.3. Setup and results of real-data MCMC parameter estimation are documented in section 4.4, along with comparison of resulting parameter marginal and joint PDFs with idealized experiments and variational results. A discussion of the results of both MCMC and variational experiments is presented in section 4.5, while conclusions and suggestions for future work are offered in section 4.6.

4.2 3-9 June 1998 TRMM Retrievals

The observations used in both variational and MCMC assimilation procedures are obtained from a Tropical Rainfall Measuring Mission (TRMM) multisensor retrieval algorithm that uses visible and infrared radiances from the Visible and Infrared Scanner (VIRS) and vertically and horizontally-polarized microwave brightness temperatures from the TRMM Microwave Imager (TMI) to produce estimates of precipitation rate, cloud top temperature, liquid and ice water path, high cloud fraction, visible and infrared cloud radiative forcing, and three-layer shortwave and longwave heating and cooling rates (L'Ecuyer and Stephens 2003, see also chapter 3). Rigorous sensitivity estimates are coupled with the error propagation technique of L'Ecuyer and Stephens (2002) to produce observation error estimates for each retrieved state (Table 3.2). For further discussion of the retrieval algorithm and sensitivity analysis, the reader is referred to L'Ecuyer and Stephens (2003).

For the current application, retrievals were obtained for the period spanning 3-9 June 1998 over a geographic region approximately equal to the SCSMEX NES

domain (Fig. 4.1). This resulted in a total of 20 TRMM overpasses that viewed some portion of the SCSMEX domain during this time period. Before examining the observations themselves, it should first be noted that TRMM orbits between ± 40 degrees N/S latitude, and revisits any given area on the globe at the same local solar time approximately every 46 days.¹⁸ Because the seven-day time period of interest does not span one full TRMM precession period, retrieved shortwave fluxes and heating rates are scaled to remove bias (T. L'Ecuyer, personal communication).

Means and standard deviations of each TRMM retrieved quantity over the SCSMEX NESA domain for the entire 3-9 June 1998 time period are listed in Table 4.1, along with mean values of each state variable output from a GCE simulation run with default cloud microphysical parameters (Table 1.1). Comparison between GCE output and TRMM retrievals prior to assimilation is useful, as it indicates the approximate level of consistency between the non-bias-corrected model and the observations. It can be seen that the mean GCE state is generally consistent with TRMM means, particularly for the precipitation rate, liquid water path, and high cloud fraction. Most of the other GCE state variables fall within one standard deviation of the TRMM observations, with the exception of the cloud top temperature, low and mid-level clear air longwave cooling, and the mid-level all-sky shortwave heating. Part of the discrepancy in the clear-air longwave cooling rates may be due to a slight moist bias in the GCE simulations; whereas the mean precipitable water vapor obtained from the objective analysis dataset used to drive the model (Johnson and Ciesielski 2005) is equal to $64.76 \text{ kg} \cdot \text{m}^{-2}$, the mean GCE precipitable water vapor is equal to $68.10 \text{ kg} \cdot \text{m}^{-2}$. This is approximately $3.3 \text{ kg} \cdot \text{m}^{-2}$ higher than the large-scale forcing data. This moist bias is likely due to use of a small domain and relatively coarse grid spacing; when the domain is expanded to 512 km, the mean GCE precipitable water vapor drops to $67.59 \text{ kg} \cdot \text{m}^{-2}$, and when the grid spacing is reduced to 500

¹⁸See <http://tsdis.gsfc.nasa.gov/overflight/PredictLocalSolar.html> for more details.

m on the 512 km domain, the mean precipitable water vapor decreases to 67.04.¹⁹ Note that the GCE exhibits approximately 25% less mean integrated liquid and ice water mass than the TRMM retrievals; because the model is forced with observed large scale advective tendencies of temperature and water vapor, less suspended condensate indicates greater available water vapor. To examine whether the small moist bias explains the GCE bias toward larger clear-air longwave cooling rates, the GCE radiative transfer scheme was run with a cloud-free profile of temperature and water vapor and with a profile in which the precipitable water vapor was 10% larger. It was found (results not shown) that the longwave cooling rates changed less than $0.15 \text{ K} \cdot \text{day}^{-1}$ while the shortwave heating rates changed less than $0.06 \text{ K} \cdot \text{day}^{-1}$, and it is therefore unlikely that the moist bias is causing the discrepancy in heating rates. The persistent low bias in the shortwave heating rates may be associated with the underprediction of ice mass in the GCE; simulated ice water path is less than half of that observed, though it still lies within the one-standard-deviation range.

Further information on the characteristics of clouds and precipitation within the SCSMEX domain can be obtained through inspection of histograms of the observed variables over the entire seven-day period (Fig. 4.2). Each quantity is retrieved on a 0.25 by 0.25 degree grid, which is then aggregated to 1.0 by 1.0 degree, and therefore histograms consist of the set of values retrieved in 1 degree boxes. Examination of histograms of cloud and precipitation variables (Fig. 4.2a) reveals that most boxes are cloud-covered; high cloud fraction is concentrated near 1.0. At the same time, liquid and ice water paths are primarily very low, as is the precipitation rate, and the cloud top temperature is centered around a relatively cold 222 K. While physical relationships cannot be determined from the retrieval histograms alone, retrieved

¹⁹It has also been pointed out (R. Johnson, personal communication) that observed sea surface temperatures (and corresponding surface sensible and latent heat fluxes) over the NESA during the June portion of the SCSMEX intensive observing period were anomalously low. GCE simulated surface sensible and latent heat fluxes (not shown) were consistent with observed values of approximately $0 \text{ W} \cdot \text{m}^{-2}$ and $30 \text{ W} \cdot \text{m}^{-2}$, respectively.

Table 4.1: List of TRMM observations compared with statistics output from a GCE simulation run with default cloud microphysical parameter values.

Observation	TRMM Mean	TRMM Std Dev	GCE State
Rainfall	21.27	50.24	20.17
Cloud Top Temperature	223.21	5.44	236.11
LWP	0.30	0.64	0.35
IWP	0.53	0.86	0.25
High Cloud Fraction	0.65	0.36	0.59
LW Cloud Forcing	64.20	41.95	50.25
SW Cloud Forcing	-62.18	52.42	-111.84
Low Level All-Sky LW Cooling	-1.36	0.77	-1.79
Mid Level All-Sky LW Cooling	-1.43	0.91	-2.15
Upper Level All-Sky LW Cooling	-1.14	1.06	-1.70
Low Level Clear-Sky LW Cooling	-1.86	0.19	-2.54
Mid Level Clear-Sky LW Cooling	-2.14	0.36	-2.60
Upper Level Clear-Sky LW Cooling	-1.51	0.35	-1.65
Low Level All-Sky SW Heating	1.28	1.07	0.68
Mid Level All-Sky SW Heating	1.82	1.26	0.97
Upper Level All-Sky SW Heating	2.60	2.63	1.27
Low Level Clear-Sky SW Heating	1.34	0.49	0.90
Mid Level Clear-Sky SW Heating	1.90	0.59	1.08
Upper Level Clear-Sky SW Heating	1.26	0.93	0.93

cloud variables provide a strong indication that the domain is largely covered with high (physically) thin clouds with very low water content. This is consistent with the general observation that much of the tropics is covered by thin cirrus at any given time. The shortwave cloud radiative forcing histogram (Fig. 4.2b) appears to support the predominance of thin cloud, as most of the histogram mass is concentrated near zero $W \cdot m^{-2}$, indicating that clouds are primarily optically thin in the visible. By contrast, the longwave cloud radiative forcing exhibits a significant secondary mode near $70 W \cdot m^{-2}$, which suggests clouds are radiatively active in the infrared. Further evidence for the persistent high cloud can be seen in a comparison of all-sky and clear-sky longwave cooling rates (Fig. 4.2c); while counterparts to modes in the clear-sky cooling rates are clearly visible in the all-sky cooling, there is an additional mode centered around zero in all three layers, and a third mode centered around $1.2 K \cdot day^{-1}$ heating in the upper layer. A similar result can be seen in the shortwave all-sky heating rates (Fig. 4.2d), in which only the upper-level all-sky heating exhibits a mode that does not exist in the clear-sky. The heating modes in the upper-level longwave and shortwave heating rates provide a further indication that the domain is consistently covered with thin cirrus.

From the analysis of TRMM retrievals, it is natural to conclude that, while there is some information in the temporal and spatial means of each quantity, significantly more information is contained in the full PDFs. While assimilation of the details of the PDFs is beyond the scope of this study, future work will need to consider constraining the full histogram of the model properties with observed histograms. For the purposes of the current work, which considers the integrated effect of clouds on climate, we contend that it is sufficient that the means of modeled cloud, precipitation, and radiation properties be constrained. Use of first variational, then MCMC methods to do so is presented in the next two sections.




Figure 4.2: Histograms for each of the TRMM retrieved quantities used in the assimilation. In each plot, the range of values of each observation is plotted on the abscissa and the number of counts in each bin is on the ordinate.

4.3 Variational Parameter Estimation

While MCMC is extremely powerful in that it samples the full unapproximated joint conditional PDF of parameters conditioned on a set of observations, it is prohibitively computationally expensive. In addition, analysis of the PDF returned by MCMC is not entirely straightforward, and, as was demonstrated in chapter 3, care must be taken when interpreting the results. By contrast, if it can be reasonably assumed that a single solution to the estimation problem exists (e.g., a single dominant mode in the state space), and if it is known that a consistent and robust relationship exists between observations and the state to be estimated, then the problem can be stated in a much simpler manner. As with MCMC, the variational (optimal estimation) solution proceeds from Bayes theorem in that some a priori information in the form of observations and/or prior knowledge of the state is available, and the objective is to estimate the optimal state conditioned on this information. However, instead of characterizing the PDF, we now assume it takes a particular form, and we search for the maximum conditional probability.

Variational methods uniformly impose a Gaussian form on the state space (Fig. 4.3), in which the conditional likelihood of the solution is computed using a least-squares distance between observations $\mathbf{y}$ and the model solution $F(\mathbf{x})$

$$L(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi)^{\frac{m}{2}} |\mathbf{R}|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] \right\}, \quad (4.1)$$

where $\mathbf{R}$ is the observation error covariance matrix. Because the logarithm is a monotonic function, the state that maximizes $L(\mathbf{y}|\mathbf{x})$ will also maximize the log

$$\ln [L(\mathbf{y}|\mathbf{x})] = \text{const} - \left\{ \frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] \right\}, \quad (4.2)$$

and since the quantity in brackets is quadratic, the best estimate $\mathbf{x}$ also *minimizes* the objective function

$$\Phi(\mathbf{x}) = \frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})]. \quad (4.3)$$

Because the observations do not often provide enough information to fully constrain the solution, additional information is typically added in the form of a priori parameter information. This adds an extra term to 4.3, which becomes

$$\Phi(\mathbf{x}) = \frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] + \frac{1}{2} [\mathbf{x} - \mathbf{x}_a]^T \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_a], \quad (4.4)$$

in which it has been assumed that observation and a priori errors are uncorrelated. The importance of the background error covariance matrix $\mathbf{B}$ cannot be overstated, as it can be shown that this matrix defines, in advance of observational constraint, the extent of the parameter solution space. Similar to the definition of allowable parameter ranges in the MCMC algorithm, if $\mathbf{B}$ is specified incorrectly, it will not be possible for the variational algorithm to locate the correct solution. It can be demonstrated that the above cost function lies at the heart of every current variational operational retrieval and data assimilation algorithm. The primary differences between different algorithms are due to different methods of minimization, and to different assumptions regarding the background error covariance matrix $\mathbf{B}$. Details of the minimization procedure, and of the application of variational methods to GCE cloud microphysical parameter estimation are discussed below.

It was stated above that, given a state space that is Gaussian in form, the best estimate $\hat{\mathbf{x}}$ consists of the set of state variables that minimizes 4.4, the error-weighted least squares distance between observations and model state. Because the state space is multidimensional, the solution is found iteratively, using the *gradient* of Φ ; the partial derivative of Φ with respect to the set of states $\mathbf{x}$ (Fig. 4.3).

$$\nabla_x \Phi(\mathbf{x}) = - [\nabla_x F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] + \mathbf{B}^{-1} [\mathbf{x} - \mathbf{x}_a]. \quad (4.5)$$




Figure 4.3: Depiction of isolines of the objective function Φ in a two-parameter state space. The gradient of the cost function is indicated, as is the minimum.

Note that all of the terms in equation 4.5 are available, with the exception of the gradient of the forward model $\nabla_x F(\mathbf{x})$ —the *Jacobian* matrix of partial derivatives of model output with respect to the set of parameters. The solution to the minimization depends on how this Jacobian matrix is computed, and it is common practice to linearize about either the current estimate or the a priori. Operational data assimilation methods typically assume that the true solution is very close to the a priori estimate $\mathbf{x}_a$ —this is because $\mathbf{x}_a$ most often consists of a short term integration of a numerical weather prediction model initialized from a state that is close to the true state. Hence, it is possible to compute the gradient linearized about the a priori as

$$\nabla_x \Phi(\mathbf{x}) = \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H}(\mathbf{x} - \mathbf{x}_a) - \mathbf{H}^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x}_a)] + \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_a), \quad (4.6)$$

where $\mathbf{H}$ is the Jacobian matrix and $\mathbf{H}^T$ is its transpose (adjoint). This equation is then set equal to zero and the estimate $\hat{\mathbf{x}}$ is obtained via iterative minimization.

Though many studies have reported on the particle size distribution of liquid and ice condensate in tropical convective clouds, there is tremendous variability in the parameters that define the particle number concentration and ice crystal type and habit. Hence, it is difficult to determine in advance which set of parameter values to linearize about, and we must instead linearize the forward model about the current estimate $\hat{\mathbf{x}}$. In doing this, the solution to the gradient becomes

$$\nabla_x \Phi(\mathbf{x}) = -\hat{\mathbf{H}}^T \mathbf{R}^{-1} [\mathbf{y} - F(\hat{\mathbf{x}})] + \mathbf{B}^{-1}(\hat{\mathbf{x}} - \mathbf{x}_a), \quad (4.7)$$

where $\hat{\mathbf{H}}^T$ is the adjoint of the Jacobian linearized about the current estimate $\hat{\mathbf{x}}$. An iterative solution to 4.7 can be found using the inverse-Hessian method (Rodgers 2000, p. 85)

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \left(\mathbf{B}^{-1} + \mathbf{H}_i^T \mathbf{R}^{-1} \mathbf{H}_i \right) \left[\mathbf{H}_i^T \mathbf{R}^{-1} (\mathbf{y} - F(\mathbf{x})) - \mathbf{B}^{-1} (\mathbf{x}_i - \mathbf{x}_a) \right]. \quad (4.8)$$

The problem is that, in contrast to radiative transfer problems in which $F(\mathbf{x})$ can be quickly evaluated, in this case the model that relates state to observations is identically the GCE model. Because creation of each element in the Jacobian matrix requires two full integrations of $F(\mathbf{x})$, evaluation of the full-rank Jacobian requires twice the total number of possible parameter/observation pairs. If we consider the 9 parameters estimated in the current experiment along with the 19 different TRMM retrieved quantities, and including the single integration of the GCE model for the current state $F(\hat{\mathbf{x}})$, a single minimization iteration requires 343 GCE model runs. At the rate of approximately 3 minutes per GCE integration, 17 hours of wall-clock time would be required per iteration, and because multiple iterations are typically required to converge to a solution, in this case the variational solution offers no benefit over MCMC.

The problem can be reduced in scale using the results of the MCMC experiments documented in chapter 3, which demonstrated that it was primarily four of the nine estimated parameters that controlled the output of each GCE simulation, and that these parameters were significantly related to only 13 different observations (Table 3.6). In addition, the rank of the Jacobian can be reduced by considering that there were only 19 observation-parameter pairs with significant linear relationships. Reducing the Jacobian to these 19 elements results in an approximate order of magnitude reduction in computational cost, as only 39 GCE integrations are now required per minimization iteration, and the efficiency might be further enhanced by parallelizing the computation of the Jacobian.

It may also be possible to use the results of the MCMC experiments to bypass

computation of the Jacobian entirely. Recall that MCMC generates a large database of GCE model solutions, and that each set of solutions consists of a realization of the set of GCE cloud microphysical parameters, the associated state values, and the computed likelihood function. Given a suitably large database that spans the allowable range of each parameter,²⁰ the Jacobian can be quickly obtained via one of two different methods. First, the Jacobian could be assumed to be *constant*, and could be computed by linearly regressing the full set of states onto the parameters as was done in section 3.5. The problem is that this would have the effect of linearizing the Jacobian around the pre-set truth state defined for the MCMC experiments. To the extent that the actual parameter set is not equal to the set used in synthetic observation experiments, this method will fail to produce the correct solution. In addition, information content and convergence diagnostics (Rodgers 2000) are predicated on an evolving Jacobian matrix; use of a constant $\mathbf{H}$ denies the assimilation procedure access to this set of powerful tools. The alternative is to read in the entire database of parameter-observation combinations from idealized or real-data MCMC experiments, then compute the Jacobian based on subsets of the larger sample. This has the advantage of truly linearizing the Jacobian about the current estimate, but if MCMC did not venture into a region of the state space for which the minimization needs the Jacobian, it will not be possible to use the database to construct the Jacobian element.

On a final note, the set of realizations from MCMC can also be used to construct the parameter a priori error covariance matrix $\mathbf{B}$. In doing this, the mean of each parameter is computed from the full database, and a covariance matrix is computed using perturbations from the mean. As was discussed earlier, construction of an appropriate a priori error covariance matrix is essential if the correct state is to be retrieved as it is the $\mathbf{B}$ matrix that determines the effective extent of the state space.

²⁰For a discussion of how realistic parameter ranges are defined, the reader is referred to chapter 3, section 1.

Though construction of the correct $\mathbf{B}$ does not solve the problem of unknown a priori parameter values, it does allow us the option of using an a priori estimate, should the observations prove not to contain enough information to sufficiently constrain the parameter set.

With the available information from the MCMC results, we now construct a variational parameter assimilation framework to estimate the snow fall speed exponent, graupel fall speed coefficient, slope intercept of the rain drop size distribution, and the graupel density. Using only those observations that are significantly related to them, the solution is found using equation 4.8 with the Jacobian linearized about the current estimate and obtained from the MCMC database. As a first test, we assume that the observations contain enough information to constrain the minimization so that no a priori is needed, and perform a simulated observation (identical twin) experiment. Parameter truth values are set to values used in chapter 3, and non-estimated parameters are set equal to the middle of the parameter ranges.

Unfortunately, the variational minimization never converges for the above experiment, but instead oscillates around the four-parameter state space indefinitely. To test whether the algorithm might work for single observation-parameter combinations, the most significant 8 parameter-observation combinations identified in chapter 3 were tested individually. These tests returned mixed results—if individual parameters are set equal to the truth values used in MCMC, the estimation finds approximately the correct value for either the Jacobian generated either by lookup or via regression. However, if the true value of the parameter of interest is set to a value other than the MCMC truth, then the minimization either fails to converge or converges to the wrong value entirely. The same result is observed for the case in which one or more other (non-estimated) parameters are set unequal to the idealized MCMC truth value.

To determine why the variational estimation fails, it is useful to compare the characteristics of the Jacobian generated from the MCMC database with the Jacobian

obtained from linear regression. To assess the properties of the Jacobian for each of the four most significant cloud microphysical parameters, the full MCMC database of parameters and states is subset into only those realizations centered around the truth values of the full set of parameters. Partial derivatives of ice water path with respect to a_g , cloud top temperature with respect to b_s , and TOA LW cloud radiative forcing with respect to ρ_g are computed in the following manner. First, the appropriate subset of parameter values and corresponding states is obtained from the database by allowing each non-estimated parameter a range of $\pm 10\%$ of its allowable parameter range around the MCMC truth value, while the parameter of interest is allowed a range of $\pm 20\%$. The maximum and minimum parameter values are then pulled from this subset, and the Jacobian element is computed (e.g., for IWP with respect to a_g) as

$$\frac{\partial IWP}{\partial a_g} = \frac{a_{g_{max}} - a_{g_{min}}}{LWP(a_{g_{max}}) - LWP(a_{g_{min}})}. \quad (4.9)$$

The resulting minimum, maximum, and truth values for each parameter under consideration, along with the corresponding state values and the Jacobian computed via linear regression and via equation 4.9 are presented in Table 4.2. It can be seen from these results that the Jacobian computed for two of the parameters (a_g and ρ_g) is of the same sign as that computed from the regression, though the magnitudes differ. However, the Jacobian computed for b_s and for N_{0r} are of the opposite sign. In addition, for the four parameters examined, there are no database points that lie outside of $\pm 20\%$ around the true values. For example, if we wish to compute the Jacobian element of IWP with respect to $a_g = 1.3$ with all other parameters set equal to the MCMC truth, *no* candidate GCE states can be found in the database, and the Jacobian cannot be computed.

The results presented in Table 4.2 illustrate the fundamental problems with using the MCMC database to compute the Jacobian. First, if a parameter has little control

Table 4.2: List of parameter and state values for MCMC database subsets along with the Jacobian computed from equation 4.9 and from linear regression.

	Parameter			State		Jacobian	
	Min	True	Max	Min	Max	Computed	Regression
a_g with IWP	0.71	0.88	0.93	0.637	0.503	-0.609	-0.245
b_s with T Cloud	0.26	0.33	0.44	244.1	245.1	5.56	-11.9
ρ_g with TOA LW CRF	0.19	0.33	0.51	-26.48	-35.26	-27.4	-7.76
N_{0r} with TOA LW CRF	0.90	1.25	1.80	100.9	122.0	-23.39	5.78

over the modeled state, then the MCMC dataset will contain states generated from a wide range of values of this parameter. However, in this case, we will have no need of the Jacobian, as we will not be assimilating on this parameter. If, by contrast, a given parameter does exert significant control over simulation output, such that we would like to constrain it with assimilation, then its PDF will be sharply defined; there will be values of the parameter associated with very low probability, and MCMC, by design, will not venture into these regions of the state space. If the true value of the parameter is relatively close to that used in simulated-observation MCMC, then the variational assimilation could derive its Jacobian from the MCMC database. If not, then the database will not contain sufficient information.

Because the Jacobian obtained from linear regression of GCE state variables onto cloud microphysical parameters does not directly depend on obtaining samples from the MCMC database, it might be used in regions for which the database has no information. However, in doing this, a linear relationship (for quantities that are often highly nonlinearly related) is extrapolated into regions of the parameter space about which we have no direct knowledge, and there is no guarantee that the properties of this relationship will be maintained in low-probability regions of the idealized MCMC state space. Second, in using the regression Jacobian, the values of all of the other parameters aside from the parameter of interest are fixed to values that may be very different from the current variational estimate. This is analogous to linearizing the state about an a priori estimate, however, as was stated earlier, because

parameter values vary widely in observational studies, it is difficult to determine which set of parameter values to linearize about. The problem with use of either the regression or database Jacobian is the noise in the observation-parameter relationships (see Figs. 3.8, 3.9, 3.10, 3.11, 3.12, and 3.13) caused by use of a limited sample. If exhaustive perturbation could be done, a suitable database could be obtained, but as was discussed in chapter 3, this is thoroughly computationally intractable. Alternatively, if MCMC were run using the real observations that are intended for use in the variational assimilation, then, by construction, the MCMC database would contain the appropriate parameter values, and the Jacobian could be computed using the mode of each parameter as the a priori estimate. However, in this case, we have gained nothing from use of a variational technique as the MCMC algorithm has already returned all of the necessary information on both the parameter estimates *and* their PDFs. In fact, because variational assimilation assumes that the state space is Gaussian with a single mode, while MCMC makes no such assumption, it could be argued that the use of a variational technique in this case might lead to an erroneous conclusion.

On a final note, it could still be argued that if the computation of the Jacobian could be parallelized, then variational methods might yet be applied. However, in doing this, a locally linear relationship between changes in cloud parameters and output state must still be assumed, which may not be valid. An even more compelling concern is the possibility of multiple modes in parameter space. As was demonstrated in chapter 3, even with MCMC we might have difficulty diagnosing these secondary modes—in the variational framework we would have no knowledge they ever existed, or whether the assimilation had found the primary mode. All of this discussion simply underscores the utility of MCMC as a method that is insensitive to non-Gaussian statistics and the presence of multiple modes. In the following section, MCMC will be used to estimate parameter PDFs from the set of TRMM SCSMEX retrievals.

Table 4.3: List of TRMM retrieved quantities vs. the GCE state that corresponds to the highest 0.1% of likelihood values. Results are plotted for the case in which both clear-sky and all-sky radiative heating rates are used in the estimation.

	TRMM Obs	Best GCE State
Rainfall	21.27	19.71
Cloud Top Temperature	223.21	229.72
LWP	0.30	0.41
IWP	0.53	0.33
High Cloud Fraction	0.65	0.68
LW Cloud Forcing	64.20	53.22
SW Cloud Forcing	-62.18	-112.33
Low Level All-Sky LW Cooling	-1.36	-1.61
Mid Level All-Sky LW Cooling	-1.43	-2.08
Upper Level All-Sky LW Cooling	-1.14	-1.66
Low Level Clear-Sky LW Cooling	-1.86	-2.35
Mid Level Clear-Sky LW Cooling	-2.14	-2.71
Upper Level Clear-Sky LW Cooling	-1.51	-1.61
Low Level All-Sky SW Heating	1.28	0.68
Mid Level All-Sky SW Heating	1.82	1.14
Upper Level All-Sky SW Heating	2.60	1.19
Low Level Clear-Sky SW Heating	1.34	0.85
Mid Level Clear-Sky SW Heating	1.90	1.12
Upper Level Clear-Sky SW Heating	1.26	0.91

4.4 Parameter Estimation Using MCMC

Having tested for whether the variational method can be used to constrain GCE cloud microphysical parameters, we now apply the MCMC algorithm to characterize parameter PDFs using actual TRMM retrievals. The observations are identical to those used in the idealized MCMC experiments of chapter 3, with the exception of longwave and shortwave bottom of atmosphere cloud radiative forcing, as these fields are not retrieved from TRMM. In the interest of maintaining consistency with the observations, the GCE observation window is extended to twenty-four hours so that averages are taken over a full diurnal cycle—in doing this, we avoid additional error due to scaling shortwave fields to match TRMM. GCE simulations are also modified for consistency with other assumptions made in the TRMM retrieval—surface albedo is changed from 7% to 12%, surface emissivity is changed from 0.97 to 1.0, and solar constant is changed from 1378 to $1367 \text{ W} \cdot \text{m}^{-2}$. In all other ways, the MCMC procedure is identical to that used in chapter 3 with uniform proposal distribution, the width of which is tuned during burn-in and starting from an identical (overdispersed) first guess distribution.

When the MCMC algorithm is run, accept rates quickly converge to approximately 40%, indicating the observations contain sufficient information on the GCE microphysical parameters (appendix C), however, comparison between observations and those GCE variables that correspond to the top 0.1% of likelihoods (Table 4.3) indicates a problem with the sample; only the GCE high cloud fraction is close to the TRMM retrieval—all other GCE state variables are quite different from the observations. The implication is that there are GCE state variables that cannot match a subset of the TRMM observations for any combination of GCE cloud microphysical parameter values. MCMC, in an attempt to sample the highest likelihoods, is sampling parameter sets that lead to state variables that do not match any observations well. Stated another way, the squared state minus observation differences are large

enough that the state space is relatively flat. In this case, the culprit is likely one or more observations with low assumed error—e.g., high cloud fraction, longwave cloud radiative forcing, and clear sky heating rates. This is because observations with small assumed error will make a relatively greater contribution to the computed likelihood.

Problematic observations can be identified by plotting TRMM retrievals along with the range of values of each GCE state variable (Table 4.4); suspect TRMM observations will fall outside of this range. In this case it can be seen that it is uniformly the radiative heating rate observations that lie outside the attainable GCE range. According to L’Ecuyer and Stephens (2003), TRMM clear-sky long and short-wave radiative heating rates are retrieved by running a radiative transfer model with European Centre for Medium Range Weather Forecasts (ECMWF) analyzed temperature and water vapor profiles as input. Hence, the resulting radiative fluxes depend on a thermodynamic profile that may not be entirely consistent with temperature and water vapor observations used to drive the GCE. In particular, though errors in ECMWF analyzed temperature are relatively low (approximately 2 K), errors in the analyzed specific humidity are much larger (approximately 30%). These errors are large enough in magnitude to produce discrepancies of the magnitude listed in Table 4.4. The intent of the current assimilation exercise is not to force the GCE to reproduce ECMWF analyzed temperature and water vapor profiles through adjustment to its cloud microphysical parameters—instead, the goal is primarily to correctly simulate the mean cloud characteristics. With this in mind, the MCMC procedure is rerun omitting the clear-sky heating rates. All-sky longwave and shortwave heating rates are retained, as assumed errors in these retrieved fields are relatively large, and because the all-sky rates are assumed to contain enough useful information to be included in the estimation.

MCMC states that correspond to the top 0.1% of likelihoods for MCMC run with resulting 13 TRMM retrieval variables are tabulated below (Table 4.5). Comparison

Table 4.4: Minimum and maximum values of GCE state variables as compared with TRMM retrieved observations. State variables whose range does not include the corresponding TRMM retrieval are marked in red.

	Min	Obs	Max
Rainfall	16.56	21.27	24.64
Cloud Top Temperature	216.36	223.21	244.71
LWP	0.24	0.30	0.70
IWP	0.17	0.53	0.85
High Cloud Fraction	0.55	0.65	0.82
LW Cloud Forcing	38.12	64.20	70.47
SW Cloud Forcing	-157.78	-62.18	-58.15
Low Level All-Sky LW Cooling	-2.05	-1.36	-1.21
Mid Level All-Sky LW Cooling	-2.68	-1.43	-1.40
Upper Level All-Sky LW Cooling	-2.06	-1.14	-1.22
Low Level Clear-Sky LW Cooling	-2.59	-1.86	-2.15
Mid Level Clear-Sky LW Cooling	-2.83	-2.14	-2.61
Upper Level Clear-Sky LW Cooling	-1.63	-1.51	-1.58
Low Level All-Sky SW Heating	0.55	1.28	0.90
Mid Level All-Sky SW Heating	0.82	1.82	1.55
Upper Level All-Sky SW Heating	1.06	2.60	1.46
Low Level Clear-Sky SW Heating	0.83	1.34	0.88
Mid Level Clear-Sky SW Heating	1.08	1.90	1.17
Upper Level Clear-Sky SW Heating	0.88	1.26	0.93

Table 4.5: List of TRMM retrieved quantities vs. the mean GCE state that corresponds to the highest 0.1% of likelihood values. Results are plotted for the case in which clear-sky radiative heating rates are omitted from the estimation.

	Obs	GCE
Rainfall	21.27	20.71
Cloud Top Temperature	223.21	225.10
LWP	0.30	0.31
IWP	0.53	0.46
High Cloud Fraction	0.65	0.66
LW Cloud Forcing	64.20	57.46
SW Cloud Forcing	-62.18	-86.97
Low Level All-Sky LW Cooling	-1.36	-1.82
Mid Level All-Sky LW Cooling	-1.43	-1.86
Upper Level All-Sky LW Cooling	-1.14	-1.65
Low Level All-Sky SW Heating	1.28	0.69
Mid Level All-Sky SW Heating	1.82	0.92
Upper Level All-Sky SW Heating	2.60	1.22

with Table 4.3 clearly indicates that removal of clear-sky heating rates from the estimation has improved the ability of the algorithm to find GCE states that fit TRMM observations. The exceptions are the shortwave cloud radiative forcing, the upper level longwave all-sky radiative cooling, and all three levels of the all-sky shortwave radiative heating. As with the previous experiment, which included the clear-sky shortwave and longwave heating/cooling rates, the upper level longwave cooling and all of the shortwave radiative heating observations fell outside the range the GCE was able to produce. Possible reasons for the continued discrepancy between GCE simulated radiative fluxes and heating rates and those retrieved from TRMM are discussed in detail in an examination of MCMC likelihood subsets below.

Examination of the histogram of each parameter estimated in the MCMC algorithm (Fig. 4.4) reveals that, as we may have expected from the idealized results presented in chapter 3, b_s , a_g , and N_{0r} exhibit the most sharply-peaked PDFs. While

a_s , b_g , N_{0s} , and ρ_g also demonstrate a well-defined mode, N_{0g} and ρ_s are relatively flat. Comparison with PDFs generated in idealized experiments (Fig. 3.3) reveals that, while the modes are in different locations, the character of the PDFs is consistent, and that the model is most sensitive to changes in the same subset of parameters in the real-observation experiments as when simulated observations are used. The primary difference is the fact that real-data PDFs (Fig. 4.4) are generally sharper than those in the simulated-observations experiments. This is due to the longer observation window (24 hours for real-data vs. 9 hours for simulated observations) used in the real-data experiments; because assimilated statistics are accumulated in time, a longer observation window necessarily translates to greater observation information content. The correspondingly greater magnitude of the observation minus state differences leads to a more well-defined mode. Aside from the general character of the PDFs, there are a few specific features of selected PDFs worth noting. First, there are secondary modes evident in marginal PDFs of a_g , N_{0r} , and ρ_g . Though the mode in the PDF of a_g centered around $a_g = 0.9$ appears insignificant in the plot of the full marginal PDF, it will be shown to be quite important later in this section. The same is true of the apparently secondary mode in the marginal PDF of ρ_g at $\rho_g \approx 0.5$. By contrast, the clearly-identifiable secondary mode in N_{0r} at $N_{0r} \approx 0.5$ will be seen to be uniformly associated with relatively low likelihood values.

As in the idealized MCMC experiments, further information on the characteristics of the state space can be obtained from examination of the joint PDFs for each pair of parameters (Fig. 4.5). It is clear that the most significant parameter is by far N_{0r} , followed by a_g and b_s . It is also clear that b_g , N_{0s} , N_{0g} , and ρ_s are of relatively less importance, as the mass in their joint histograms is uniformly well-dispersed through the state space. Similar to the identical twin experiments, there appears to be a correlation between a_s and b_s and perhaps between N_{0s} and N_{0g} . It is interesting to note that the secondary mode in N_{0r} appears to be coupled with low values of ρ_g ,




Figure 4.4: As in figure 3.3, but for marginal PDFs obtained from use of MCMC with TRMM retrievals.




Figure 4.5: As in figure 3.5, but for joint PDFs of each parameter pair obtained from use of MCMC with TRMM retrievals.

large values of b_s , and in particular with large values of a_g .

As was pointed out in chapter 3, clear identification of the primary mode for each parameter, as well as compelling evidence for whether changes in that parameter have a significant effect on simulated states can be obtained from scatter plots of parameter values for the top 0.1% of likelihood values. To recap, if nearly any value of a parameter results in identical (and high) likelihood, then changes to that parameter, by definition, have little effect on the GCE state. Furthermore, for the most significant parameters, it will be obvious which *mode* is correct, as this mode

will be associated with the highest likelihoods, while any secondary modes should be uniformly associated with relatively low likelihoods.

When the total likelihood is computed using all of the observations employed in the MCMC procedure, the maximum total likelihood is ≈ 0.1 out of a possible range of $L(\mathbf{x}|\mathbf{y}) \sim [0, 1]$ —by contrast, the highest likelihood for the idealized experiment is nearly equal to 1.0. The implication is that, while idealized MCMC was able to find a parameter combination for which the GCE state nearly exactly matched observations, for the real-data experiment there were *no* parameter combinations that were able to produce a state that was close to the observations. Recall, as was demonstrated in Table 4.5, even when clear-sky heating rates are removed from the estimation, GCE was unable to produce strong enough all-sky shortwave heating rates in any level, and that significant differences were also noted in the shortwave cloud radiative forcing and the upper level all-sky longwave cooling. When the total likelihood is computed without the influence of the shortwave heating rates, the maximum likelihood increases to 0.45. When upper-level longwave cooling and shortwave cloud radiative forcing are removed, the maximum likelihood increases to 0.77, and when low and mid-level all-sky longwave heating rates are also removed, leaving only likelihoods associated with cloud and precipitation variables and the longwave cloud radiative forcing, the maximum likelihood jumps to 0.98. In computing all-sky shortwave cloud forcing and radiative heating rates, the TRMM retrieval procedure makes assumptions about the shape of the particle size distribution of precipitating and non-precipitating liquid and ice. Because the parameters that define the particle size distributions in the retrieval are fixed, they cannot be made consistent with GCE parameter values in each iteration of the assimilation, and the resulting differences in fluxes and heating rates most likely make up the bulk of the discrepancy in the shortwave fields. In the case of the upper level longwave heating rates, it was found after running the MCMC procedure that the TRMM retrieval computes upper level heating rates in a

layer that extends to 17 km above the surface. GCE heating rates were computed in a layer that extends to the model top (approximately 24 km), which resulted in the GCE simulating greater upper level longwave cooling. Given the mismatch between assumed particle size distribution in the retrieval and the GCE, and in the interest of diagnosing values of cloud microphysical parameters for use in large-domain climate interaction experiments, likelihood scatter plots are constructed using total likelihood values computed from a subset of the observations that does not include shortwave heating rates.

When these scatter plots (Fig. 4.6) are compared with the marginal PDFs (Fig. 4.4), it is clear that the mode in the marginal PDFs of a_s , b_s , b_g , and N_{0r} is likely the true primary mode for each of these parameters. By contrast, the highest likelihood values for N_{0s} and ρ_s are evenly distributed across the range of each of these parameters (with the exception of the very lowest parameter values), indicating the apparent mode in N_{0s} may have been secondary. It is interesting to note that the secondary mode in N_{0r} is missing completely from the highest likelihood values, and even more interesting that what appeared to be the *primary* mode in the marginal PDF of a_g has been supplanted with the secondary mode centered around $a_g = 0.9$. The same is true for ρ_g , in which the apparent secondary mode at $\rho_g \approx 0.5$ has become the primary mode in the high-likelihood subset. Since modes in these plots were identical to those in scatterplots constructed from likelihood subsets that did not include longwave cooling or shortwave cloud radiative forcing (not shown), we can safely conclude that it was the shortwave heating rates that were contributing to the secondary modes in a_g , N_{0r} , and ρ_g . It appears, thus, that the differences between the marginal PDFs and the high-likelihood scatter plots are due to MCMC attempting to produce greater all-sky shortwave heating rates using the microphysical parameters as control variables.

If we assume the true modes of parameter PDFs are those that correspond to the




Figure 4.6: Scatter plots of the value of each estimated parameter for the top 0.1% of likelihood values. Total likelihood values range from 0.0 to 1.0, and were computed excluding the three-layer shortwave all-sky radiative heating rates.

Table 4.6: List of default GCE cloud microphysical parameters alongside the values returned by the MCMC algorithm and their standard deviations.

	Default	SCSMEX	Std Dev	Minimum	Maximum
a_s	0.61	0.56	0.13	0.50	1.0
b_s	0.11	0.80	0.20	0.10	1.0
a_g	1.5	0.90	0.26	0.6	1.7
b_g	0.37	0.17	0.15	0.1	0.7
N_{0r}	0.08	0.05	0.45	0.04	5.0
N_{0s}	0.16	1.00	1.25	0.04	5.0
N_{0g}	0.08	1.50	1.26	0.04	5.0
ρ_s	0.10	0.35	0.23	0.1	1.0
ρ_g	0.40	0.55	0.23	0.1	1.0

subset of states with the highest likelihoods (Table 4.6), we can conclude that this set of parameters represents those that are most consistent with the mean atmospheric state over the SCSMEX domain. When estimated parameters are compared with default GCE parameter values (Table 4.6) it can be seen that many are similar. In particular, values of a_g and N_{0r} , the parameters that have the greatest effect on the mean behavior of the GCE model, are quite close to the control values. Though a_s , N_{0s} , and ρ_g differ, the default values fall within ± 1 standard deviation of the MCMC mode. The MCMC estimated parameters with the greatest departure from control values are b_g , b_s , N_{0g} , and ρ_s , all of which exhibit greater than one standard deviation departure from default. It appears that the best match between TRMM observations and the GCE state over the SCSMEX domain was obtained using a relatively fast-falling population of more dense snow crystals that were smaller on average than in the control. In a similar manner, graupel particles were set to smaller sizes with slightly slower fallspeeds and slightly higher densities. Though a detailed comparison between estimated parameter values and observed particle size distributions is beyond the scope of this study, because parameter values were required to conform to the range of values observed in nature in the course of the assimilation, the estimated parameters are by definition physically realistic.

4.5 Discussion

While it is disappointing that the variational methodology could not be successfully applied to the estimation of GCE cloud microphysical parameters, it was demonstrated in chapter 3 that improving the error characteristics of only a few observations dramatically improved the quality of the relationship between the parameters and these observations. It is therefore reasonable to expect that once improved estimates of cloud and precipitation variables are obtained (e.g., from future sensors), the variational assimilation framework may yet be of use. In addition, it remains to be seen whether re-computation of the Jacobian in each variational iteration might produce a better result.

In contrast to variational estimation, the MCMC algorithm proved more effective in returning a PDF of cloud microphysical parameters consistent with TRMM observations. However, it must be stated that there is no way to know whether these results can be extended to other tropical environments, or even to different time periods during the SCSMEX field campaign. While an exhaustive examination of how the PDF might change for different times and regions of the globe is beyond the scope of this study, it is a task that must be undertaken if the qualities of the model are to be fully understood. In addition, as increasing numbers of climate models are incorporating CRMs as replacements for traditional convective parameterizations, it is important to determine how changes in the large-scale environment are related to changes in the assumptions underlying CRM cloud and precipitation parameterizations.

It should also be noted here that the set of nine parameters estimated above do not constitute the full set of uncertain cloud microphysical parameters in the GCE model. They are, instead, those parameters that are most closely related to observable physical properties of the system. There are many other parameters (e.g., collection efficiencies and cloud and ice autoconversion parameters) that are not as directly related to observations, and that may have an important effect on the simulated

results. As we are primarily interested in parameters that define the particle size distribution, in the above procedure, these additional parameters have been set to default values. If knowledge of the uncertainty in the full range of cloud microphysical parameters is desired, additional parameters could easily be inserted into the MCMC procedure in future experiments.

The results presented in section 4.4 above illustrate an important point about the character of an estimate returned from any data assimilation procedure. Though MCMC effectively sampled the PDF of cloud microphysical parameters, and while all diagnostics have indicated this PDF to be a sample of the true joint conditional PDF, it cannot be overstated that data assimilation always returns a joint *conditional* estimate. Results are always a function of which observations are used, assumed errors on these observations, and the particular control parameters estimated. If we were to change which observations are used, or even to change the assumed errors on the observations, the PDFs would necessarily change as well. As was shown in chapter 2, simply scaling the observation errors equally will not change the locations of modes of the PDF in state space, however, if the error on one or more observations *relative* to the others is changed, then the locations of the modes may well change. This has already been illustrated in section 4.4—when the full sample is considered, the primary mode in the graupel fallspeed coefficient is located at $a_g \approx 1.5$, and there is a secondary mode in the slope intercept of the rain drop size distribution at $N_{0r} \approx 0.5$. When the parameter sample is subset according to the values of the likelihood function such that the shortwave radiative heating rate observations are effectively removed from the sample, the formerly primary mode in a_g and secondary mode in N_{0r} disappear.

Sensitivity of the estimation to the presence of shortwave heating observations raises a general concern with the application of data assimilation methodologies to the model bias correction problem. All assimilation methodologies use observations

to constrain a set of uncertain control variables, and though MCMC is more flexible than most, it is still guided by the specific choice of control variables and observations. If there are differences between the model state and observations that are *not* associated with uncertainty in the assimilation control variables, the assimilation algorithm will still attempt to adjust the control variables so that the model state better fits the observations. In the current application, there were differences between TRMM retrieved and GCE simulated long and shortwave heating rates that were likely associated with assumptions used in the retrieval procedure. In the assimilation, MCMC seeks to sample those sets of parameters that provide the best fit to *all* of the observations, including the radiative heating rates, and in the process effectively attempts to change the cloud microphysical parameters so that the GCE state conforms more closely to the assumptions used in the retrieval. In the above example, the estimation failed to find any combination of parameters that could satisfactorily match observations of cloud, precipitation, *and* radiative heating rates. It is a strength of MCMC that the algorithm continued to sample the state space in spite of the difficulty matching the observations.

Because of the inherent perils associated with blind use of data assimilation methods, it must be stressed that estimation and correction of model bias can never replace careful improvement to the model's physical formulation. Data assimilation, when properly applied, should yield appropriate information on *how* the model should be improved. The danger in any model bias estimation/correction technique is that fundamental flaws in the model formulation might be masked simply by tuning poorly-known or poorly-understood constants. In the case of the GCE model, work is underway to significantly upgrade the cloud microphysical formulation. Once this has been done, there will be a new set of unknown cloud microphysical parameters, and we anticipate MCMC will be useful in exploring the sensitivity of model output to changes in this new set of parameters as well.

Finally, while it is tempting to conclude that the MCMC procedure has effectively *retrieved* cloud microphysical characteristics, it must be noted that even if the control variables can be changed to produce a state that is a close match to observations, there is no guarantee that this set of control variables is the *true* set. In the above example, we can draw no conclusions from the MCMC result about the actual mean microphysical characteristics of precipitating cloud systems that occurred during SCSMEX. We only have the assurance that the set of parameter values in Table 4.6 are those that produce cloud, precipitation, and radiation statistics that are most consistent with means observed during SCSMEX. In addition, we have chosen only to assimilate the means of TRMM observations, and it could be argued that it is the full PDF of the state that should be consistent between model and observations. To this end, future studies will explore assimilation of higher order moments of the observed PDFs.

4.6 Conclusions

In this chapter, both variational and MCMC methods have been applied to the estimation of GCE cloud microphysical parameters. Following the results of idealized MCMC experiments in chapter 3, those parameter/observation sets that exhibited the most significant linear relationships were tested in the variational framework. Unfortunately, the computation of the partial derivative of observations relative to changes in microphysical parameters (the Jacobian matrix) proved problematic, and the combination of a poorly-specified Jacobian with noise in the observation-parameter relationships prevented even identical twin variational test experiments from converging.

MCMC was applied to cloud microphysical parameter estimation using time and domain means of TRMM retrieved cloud, precipitation, and radiation properties over the SCSMEX northern enhanced sounding array. It was found that, while MCMC

returned a robust sample of the PDF, the results were influenced by a moist bias in the model. This moist bias, in particular, prevented the model from being able to match observed clear-sky radiative heating and cooling rates. When clear-sky heating rates were removed from the estimation, the algorithm successfully produced GCE states whose range encompassed the observed values. However, shortwave all-sky heating rates were still affected by the moist bias, and led to significant (and secondary) modes in the PDFs of a_g and N_{or} . Subsets of parameters associated with the highest likelihoods computed without shortwave heating rates exhibited no evidence of these secondary modes. It was decided that the most optimal values of the microphysical parameters are those that match most closely the cloud, precipitation, and radiative forcing observations, as the moist bias would likely be mitigated in simulations on a much larger domain.

In general, MCMC has proven to be an effective tool for exploring the behavior of a numerical model in the presence of uncertain cloud microphysical parameters. Having characterized the sensitivity of key climate variables to changes in the GCE cloud microphysical parameters in both idealized and real-data experiments, and having obtained a set of parameters consistent with observations of these variables, we now proceed to use the GCE model to study the response of the tropical hydrologic cycle to surface warming. Details on the configuration of large-domain GCE experiments, as well as results and discussion, are contained in the following chapter.

Chapter 5

Large-Domain Goddard Cumulus

Ensemble Cloud Feedback

Experiments

5.1 Introduction

As was discussed in chapter 1, clouds play a critical role in the tropical climate system as they link the planetary radiation budget with the large scale circulation. Despite decades of research and the development of a host of hypothesized cloud feedback mechanisms, the response of the tropical hydrologic cycle to surface warming is not yet well understood. Specifically, it is unknown whether, given a perturbation to the large-scale environment (e.g., surface warming), clouds will enhance or moderate the system response. This is due, for the most part, to the fact that radiation and the large-scale circulation interact in a highly nonlinear manner through cloud dynamics and microphysical processes. In particular, cloud microphysical processes determine the bulk radiative properties of clouds through the evolving ice and liquid particle size distributions, and affect the large-scale circulation through the vertical distribution

of latent heat release. Over the past few decades, many attempts have been made to characterize and quantify cloud feedbacks using simple conceptual models (Ramanathan and Collins 1991, Lindzen et al. 2001), energy balance models (Stephens and Webster 1981), and General Circulation Models (GCMs), however, observational evidence in support of these feedbacks has been extremely difficult to obtain, and model studies have been plagued by uncertainties associated with simplified representation of cloud and precipitation processes.

Cloud resolving models (CRMs) present a unique system within which to study the operation of cloud feedbacks, as they offer a controlled environment in which the interaction of clouds, radiation, and the large-scale circulation can be realistically represented. A number of studies have demonstrated the utility of CRMs for studying the interaction among convection, radiation, and the large scale circulation in the tropical climate system—in this mode, the model employs periodic boundary conditions and is integrated over a long period of time (20+ days) until a state of quasi-equilibrium is reached (Held et al. 1993, Sui et al. 1994, Grabowski et al. 1996a, Tompkins and Craig 1998, Tao et al. 1999, Shie et al. 2003). CRMs have also been used in attempts to assess the sensitivity of climate to changes in SST (Wu and Moncrieff 1999, Tompkins 2001a, 2001b) and to the properties of the large-scale environment and circulation (Grabowski et al. 1996b, Lucas et al. 2000). Though each of the CRMs used in the above studies includes a representation of cloud and ice microphysical processes, all are required to parameterize the details of the interaction of different cloud species on subgrid scales. As was demonstrated in chapter 1, relatively small changes to the parameters that determine the liquid and ice particle size distributions result in large changes to the characteristics of the simulated state at radiative convective equilibrium. Several studies have demonstrated the sensitivity of CRM results to changes in the cloud microphysical representation (Tao et al. 1995, Grabowski et al. 1999, Grabowski 1999, 2000, Wu et al. 1999, Petch

and Gray 2001, and Grabowski 2003), however, none has yet characterized the full range of uncertainty or used observations to constrain it. In addition, though each of the above studies has investigated aspects of the relationship between clouds and the large scale circulation, none has directly sought to assess the validity of hypothesized cloud feedback mechanisms themselves.

The current study is unique in that it employs an observation-constrained cloud resolving model to study the response of the tropical hydrologic cycle to perturbations in sea surface temperature, and in so doing assesses whether either of two prominent hypothesized feedback mechanisms are feasible. It has been shown above (chapter 3) that observations could theoretically be used to constrain selected cloud microphysical parameters such that the mean climate state is faithfully reproduced, and (chapter 4) that data assimilation methods could be used to estimate an optimal parameter set. Having used observations to estimate values of cloud microphysical parameters associated with the definition of ice and liquid particle size distributions in the Goddard Cumulus Ensemble (GCE) model, large-domain CRM simulations are now used to study the response of cloud and precipitation processes to surface warming over the equatorial ocean.

In this chapter, traditional radiative-convective equilibrium (RCE) experiments are performed, in which both sea-surface temperature and insolation are fixed and the model is run for a long period of time (40-60 days) until the model reaches steady-state. In each experiment, a two-dimensional version of the GCE model is run on a domain that spans more than 10,000 km in the horizontal—use of such a large horizontal domain allows ample room for mesoscale interactions and promotes the development of a realistic large-scale overturning circulation. Three simulations are performed, in which sea surface temperature is fixed at values of 295, 300, and 305 K, and comparison of output from these three simulations is used to assess the response of cloud and precipitation processes to surface warming. In the process,

model output is examined for evidence of the well-known Thermostat (Ramanathan and Collins 1991, Fig. 1.1a) and infrared Iris (Lindzen et al. 2001, Fig. 1.1b) feedback mechanisms. Though a proper examination of the operation of a tropical cloud-climate feedback mechanism necessarily requires inclusion of large-scale dynamical mechanisms associated with the Earth's rotation (e.g., the Madden-Julian oscillation; Madden and Julian 1972, 1994), results of idealized two-dimensional simulations will provide valuable insight into how cloud, precipitation, and radiation processes respond to surface warming.

Both the thermostat (Ramanathan and Collins 1991) and the infrared iris (Lindzen et al. 2001) hypotheses predict that cloud and precipitation processes will act to moderate the climate system response to surface warming in a *negative feedback*, but each operates via a very different mechanism. The thermostat (Fig. 1.1a) assumes that convective intensity will increase with increasing SST. In the process, greater amounts of thick cirrus clouds will be produced, which will consequently reflect greater amounts of incoming shortwave radiation, leading to surface cooling. The key assumption is that the fundamental cloud dynamic and microphysical processes are independent of changes to the surface temperature; greater convective mass flux is directly associated with greater amounts of condensate in the upper troposphere. Output from GCE simulations will be used to investigate whether the thermostat is operating by examining how the outgoing shortwave flux, high cloud fraction, and ice mass change with increasing SST. If the high cloud fraction, ice mass, and outgoing shortwave radiation all *increase* with increasing sea surface temperature, this will be evidence that a thermostat-like effect may be operating.

In contrast to the thermostat, the infrared Iris (Fig. 1.1b) assumes that an increase in *precipitation efficiency* accompanies increases in sea-surface temperature; convective mass flux increases with increasing SST, but the rate of precipitation relative to the rate of condensation increases as well, leading to a net *decrease* in upper-level

cirrus with increasing SST. Decreases in cirrus cloud cover, and consequently larger cloud-free area relative to the area covered by upper-level cloud, would in turn lead to increases in outgoing longwave radiation and regulation of SST. There are two key assumptions that underlie the operation of the Iris hypothesis: (1) the domain average reflected shortwave radiation does not increase with changes in SST; changes in the longwave fluxes dominate the radiative forcing, and (2) cloud microphysical processes (in particular the precipitation production process) fundamentally change character with increasing SST. Similar to the thermostat, the Iris hypothesis can be tested using output from the GCE. If the ice mass decreases or remains constant, while the high cloud fraction decreases and the precipitation efficiency increases, this will be evidence in support of the Iris. However, the results will need to be tested by examining the mean shortwave and longwave components of the radiative budget; cloud and precipitation processes can only produce an Iris-like feedback if they can be shown to affect the radiative energy budget.

The remainder of this chapter is organized as follows. Because of its relevance to the cloud-climate feedback discussion, and its rather ambiguous use in the literature, a general discussion of the definition and relevance of PE as a diagnostic tool is presented in section 5.2. The details of the GCE model configuration and descriptions of the Thermostat and Iris hypotheses are presented in section 5.3. Results of radiative-convective equilibrium experiments with the GCE model are presented in section 5.4, and discussion of the experiment limitations and suggestions for improvement are outlined in section 5.5. Brief concluding remarks and an outline of future work are offered in section 5.6.

5.2 Precipitation Efficiency

Though the infrared iris hypothesis is the only currently-proposed feedback mechanism that explicitly depends on changes to precipitation efficiency, it can be shown that nearly every other cloud-climate feedback hypothesis depends on the convective PE. Though it is beyond the scope of this study to discuss in detail the manner in which PE factors into the operation of every feedback mechanism, it is still useful to discuss the general usefulness of PE as a diagnostic. A search of the literature reveals a range of definitions of convective cloud system PE depending on the desired application or available observing systems, making PE appear to be a relatively ambiguous quantity. In observational studies, PE has either been defined as the ratio of precipitation rate observed at the ground to the rate of water vapor flux at cloud base (Braham 1952, Auer and Maurwitz 1968, Foote and Fankhauser 1973, Heymsfield and Schotz 1985, Fankhauser 1988), or more recently as the ratio of the total amount of surface precipitation to the total mass of water condensed in the storm (Gamache and Houze 1983, Chong and Hauser 1989, Oury et al. 2000). In studies of the water budgets of convective systems using numerical models, PE has been typically been defined as the ratio of total surface precipitation to total condensation (Murray and Koenig 1972, Weisman and Klemp 1982, Lipps and Hemler 1986, Ferrier et al. 1996, Li et al. 2002, Tao et al. 2004), though it should be noted that these studies tend to differ in their interpretation of what constitutes total condensation. In addition, the representation of microphysical processes differs widely across the spectrum of models used in these studies, and the PE of a given convective system has been demonstrated to vary widely when simulated with different microphysical treatments (Murray and Koenig 1972, Ferrier et al. 1996). An alternate definition of PE, suitable for use in large scale averaging, defines PE as the ratio of surface precipitation to large scale moisture sources (e.g., advection of water vapor and surface evaporation, Li et al. 2002). This definition of PE avoids the complication of having to quantify the wa-

ter mass budget of the convective system, while providing useful (though implicit) information on the PE of cloud ensembles in a given region or model domain.

The definition of PE can be demystified by noting that each of the above studies effectively seeks to quantify the ratio between the total accumulated surface precipitation and the total water vapor entering a convective system or geographic region—this partition defines the effective latent heat release of the system, as well as the amount of condensate available for detrainment at cloud top. Because the amount of water vapor that enters a given convective system can be difficult to quantify, the water budget of a convectively active region is often approximated as the total amount of *condensation* over some pre-defined time interval such that

$$q_{v(in)} \approx q_{v(condensed)} = E + S + P_{sfc} + \frac{dq_x}{dt}, \quad (5.1)$$

where $q_v (in)$ is the total vapor ingested into the storm, $q_v (condensed)$ is the total condensed vapor, E and S are the evaporation and sublimation of condensate to vapor in-cloud, P_{sfc} is the surface precipitation, and $\frac{dq_x}{dt}$ is the change in suspended condensate between the beginning and end of the time period. The time interval of interest typically encompasses the entire lifecycle of a convective system from initiation to dissipation so that the precipitation efficiency does not exceed 100%. The PE is then defined as the ratio

$$PE = \frac{P_{sfc}}{E + S + P_{sfc} + \frac{dq_x}{dt}} = \frac{P_{sfc}}{q_{v(condensed)}}. \quad (5.2)$$

It can be shown that each of the aforementioned authors employs a form of the above water mass budget, and applies this definition of PE to individual convective cells, cumulus ensembles, or geographic regions.

In models that seek to realistically parameterize cloud microphysical processes, the amount of total condensation and deposition occurring over the life of a convective

system is often explicitly available, and the PE can then be directly computed from the right hand side of 5.2. It should be noted that while this formulation accounts for all of the vapor condensed in a convective system, it neglects the vertical transport of uncondensed water vapor, a term in the water budget that may be of critical importance for upper tropospheric humidity (Luo and Rossow 2004). In the current application, the precipitation efficiency is defined according to (5.2) as the ratio of total surface precipitation to the total mass of water vapor that is converted to liquid or ice through condensation or deposition, respectively. The time interval used in the computation spans a period of 20 days at the end of each RCE simulation so that the resulting PE is the mean efficiency of many convective systems, and the condensation, deposition, and precipitation are accumulated over the entire horizontal domain. The resulting quantity is the PE of the ensemble of all of the clouds that exist in the domain at equilibrium.

CRMs have been used to study the precipitation efficiency of convective systems for more than three decades, and have been applied to both isolated convective systems (Murray and Koenig 1972, Weisman and Klemp 1982, Lipps and Hemler 1986, Ferrier et al. 1996, Shepherd et al. 2001) and convective ensembles (Li et al. 2002, Tao et al. 2004). Though past studies have demonstrated the usefulness of PE in that it is directly related to the effective latent heating produced by a convective system or ensemble, it is limited in that it indicates *nothing* about the radiative effect of clouds on the system. This is because the radiative properties of clouds, particularly in the shortwave, depend critically on the cloud microphysical characteristics. For example, consider two convective systems with identical surface precipitation and precipitation efficiency such that the amount of suspended condensate is the same. Furthermore, assume that each system produces the same proportion of liquid relative to ice. If one storm consistently produces large ice particles relative to the other, then (assuming all ice particles are large compared with visible wavelengths), the storm with more small

particles will reflect more shortwave radiation, and will have a consequently different effect on the climate. Indeed, it is on the shoals of the cloud radiative effect that both the thermostat and Iris hypotheses founder; both hypotheses assume a simple linear relationship between changes in PE and changes in cloud radiative properties. Hence, in this study, while the PE will be carefully examined, it is only one of a host of diagnostics we will use.

5.3 Experiment Setup

In the experiments performed in this chapter, the 2D version of the GCE is run with 2.5 km horizontal grid spacing on a domain 10,240 km in length and approximately 24 km deep. 43 vertical layers are used, which stretch in spacing from 40 meters in the atmospheric boundary layer to approximately 1 km in the free troposphere. Cyclic boundary conditions are employed such that the model is effectively a closed system; energy is only allowed to enter and exit the domain through the top and bottom boundaries. Sea surface temperature is fixed at values of 295, 300, and 305 K to simulate the response of cloud and precipitation processes to surface warming, and a fixed solar zenith angle of $\cos^{-1}(1/\pi)$ (71.44 degrees) is applied, which leads to constant insolation of 438.6 W m^{-2} (roughly equivalent to the annual mean insolation at the equator). Each simulation is initialized from an identical vertical profile (Fig. 5.1), which is characterized by surface-based convective available potential energy (CAPE) of $1530 \text{ J} \cdot \text{kg}^{-1}$ and a surface temperature of 302 K. This sounding was obtained at 0000 UTC 5 December 1992 during the TOGA COARE field campaign. The fact that the CAPE in this sounding is greater than that typically observed over the equatorial western Pacific ocean is a benefit in RCE experiments, as simulated convection will develop more rapidly and, as a result, the system will more quickly settle into an equilibrium state (L. Pakula, personal communication). So that the




Figure 5.1: Skew-T/log-P plot of the sounding used to initialize the RCE experiments. Note that the wind profile at right applies only to the cases with shear.

model does not experience numerical instability in the boundary layer associated with non-physical temperature lapse rates, at the start of each simulation the surface temperature is gradually relaxed from its initial value to the appropriate fixed surface temperature at a rate of $1 \text{ K} \cdot \text{day}^{-1}$. Initiation of convection is encouraged through introduction of random thermal perturbations between -0.2 and 0.2 K in the lowest 2 km at the initial time.

Two sets of simulations are performed—one in which zero mean wind is applied such that the model is allowed to generate its own overturning circulations (and associated vertical shear of the horizontal wind), and a second, in which moderate background

vertical shear is applied through the base state winds. This shear is obtained from a temporal average of wind profiles over the TOGA COARE intensive flux array (Webster and Lukas 1992) from the period spanning 19-27 December 1992. Shear values are typical of post-monsoon onset conditions, and this particular 6-day period has been extensively studied by the GEWEX cloud systems study (GCSS) working group 4 (Moncrieff et al. 1997, Su et al. 1999). In each case, the model generates cloud at approximately two simulated days, and each experiment is run for 40-60 days.

5.4 GCE Results

5.4.1 Convergence to Radiative Convective Equilibrium

Before examining the results of the RCE experiments, it is first necessary to demonstrate that each simulation has indeed reached a steady state. RCE is defined as a state in which the column radiative cooling and surface latent heat flux are in balance with the column latent heating and surface sensible heat flux. Because this study deals with changes in cloud and precipitation processes with changes in sea surface temperature, we apply the additional constraint of steady-state domain mean precipitation rate, total cloud water, and precipitable water vapor. In addition, we examine the individual cloud microphysical processes for evidence that the individual components of the latent heating budget have also reached equilibrium.

Figure 5.2 depicts timeseries of domain-averaged radiative and moist state properties over the length of each simulation. For both zero-mean wind and TOGA shear cases, radiative fluxes rapidly converge to approximately steady state, reaching equilibrium by ≈ 10 days (240 hours). Though there is greater temporal variability in LHR, precipitation rate, and total water path, these fields converge within less than 5 days. An indication of the convergence of the atmospheric state to equilibrium can be seen




Figure 5.2: Illustration of convergence of simulations in each experiment to radiative convective equilibrium. Simulations run with zero-mean background flow are depicted at left, while simulations run with TOGA-COARE shear profiles are on the right. Shown are domain means of various cloudy state variables as well as short and longwave radiative fluxes at the top of the atmosphere. In each plot, the blue line corresponds to simulations run with constant SST of 295 K, black to simulations run with SST of 300 K, and red to simulations run with SST of 305 K.

in plots of the precipitable water vapor, defined as the pressure-weighted integrated mass of water vapor in the column per unit area. This field, which is more sensitive to subsident drying (and thus to the properties of the large-scale circulation) than the others, exhibits a longer convergence period for both sets of simulations, reaching equilibrium around 30 days for the zero-mean wind case, and 20 days for simulations with shear.

The more rapid convergence of large scale circulations to equilibrium for the simulations with shear evident in the PWV timeseries can also be observed in timeseries of the microphysical components of the cloud latent heating budget (Fig. 5.3). As is typically done for ice and liquid water paths, these quantities have been converted from mass mixing ratio to pressure-weighted column values. While condensation and evaporation settle into steady state as rapidly as precipitation, total water path, and mean latent heat release, the domain mean freezing and melting do not appear to reach equilibrium until after 20 days. Sublimation and deposition converge even more slowly, reaching equilibrium after 40 days for the zero-mean wind case, and after 20 days for simulations with shear.²¹ Rapid convergence of precipitation rate, condensation, and evaporation to steady state indicates the liquid processes themselves adjust more rapidly to equilibrium, and the rapid convergence of LHR reflects the fact that the largest terms in the latent heating budget come from the vapor-liquid transitions (condensation and evaporation). From the slower convergence of freezing and melting, and especially the sublimation and deposition processes, it appears that the ice processes take far longer to reach a steady state. Because the precipitable water vapor also exhibits slow convergence, it is reasonable to imagine that this delay might be associated with the details of convergence of the large-scale organization. Properties of the large scale circulation in each simulation are presented in the next section.

²¹Because simulations with shear converge to equilibrium relatively quickly compared with simulations run without shear, sheared simulations are run for 40 days, while non-sheared simulations are run for 60 days.




Figure 5.3: Illustration of convergence of simulations in each experiment to radiative convective equilibrium. Simulations run with zero-mean background flow are depicted at left, while simulations run with TOGA-COARE shear profiles are on the right. Shown are timeseries of domain mean microphysical contributions to the latent heating budget. In each plot, the blue line corresponds to simulations run with constant SST of 295 K, black to simulations run with SST of 300 K, and red to simulations run with SST of 305 K.

5.4.2 Large Scale Convective Organization

Observational studies have consistently demonstrated the fact that convection in the equatorial tropics rarely occurs in the form of isolated cells, but instead typically organizes into cloud clusters (Nakazawa 1988, Lau et al. 1991, Sui and Lau 1992, Mapes 1993, Mapes et al. 1993, Wilcox 2003). These cloud clusters commonly exhibit cold cloud top area in excess of $5,000 \text{ km}^2$ and maintain coherency for approximately 2 days, though the largest clusters can persist for more than a week (Mapes and Houze 1993). Though the mechanism that causes this large scale organization is a topic of current research, because it is a uniformly observed feature of convection in the tropics, it is important to assess the extent to which it occurs in the RCE simulations.

Hovmoller plots of perturbation precipitable water vapor from each SST and for each set of experiments are contained in figure 5.4. As it will be demonstrated later, the domain mean PWV increases with increasing SST, hence deviations from the domain mean PWV computed over the second half of each simulation are plotted to facilitate the comparison. It is immediately apparent that the mode and scale of convective organization is fundamentally different between sheared and unsheared cases. Simulations with zero mean wind organize into alternating moist and dry bands, while sheared simulations exhibit a single moist and dry region. In addition, while the scale of organization appears to decrease with increasing SST in the unsheared simulations, the scale is roughly equivalent in simulations with shear. The one similarity between the two experiments is the relative magnitude of the PWV perturbations, which uniformly increase with increasing SST.

Examination of the total water path (Fig. 5.5) indicates that, for both sets of experiments, the convection does indeed organize into clusters. In addition, though the water vapor field demonstrates a high degree of temporal coherency, the actual convective systems are far more transient. While the PWV field suggested a different




Figure 5.4: Hovmoller plots of perturbation precipitable water vapor for each of the simulations in this study. Experiments with zero mean wind are shown in (a) - (c), while results from simulations with shear are depicted in (d) - (f).




Figure 5.5: Hovmöller plots of total column integrated condensate (liquid + ice) for each of the simulations in this study. As in Fig. 5.4, simulations run with zero mean wind are depicted in (a)-(c), while experiments run with shear are contained in (d)-(f).

scale of organization with changes in SST for the non-sheared simulations, the scale of the cloud clusters themselves is quite similar. It is interesting to note that the large SST run starts with larger convective systems, but that the convection organizes on a finer scale with time, while the cold-SST simulation exhibits fine-scale organization initially, and more large scale organization later. In the absence of shear, cloud clusters in all SST experiments persist for long periods of time—from approximately 10 days for the smallest clusters ($\Delta x \approx 500$ km) to the entire length of the simulation for the largest clusters ($\Delta x \approx 1000$ km). This is consistent with the findings of Held et al. (1993) and Tompkins (2001a) that, in the absence of shear, convection tends to re-occur in regions moistened by previously-occurring convection. In contrast, while convection eventually organizes into one primary geographic region for each SST in sheared simulations, examination of the total water path reveals significant differences on smaller scales. First, a coherent large scale region of convection develops much more quickly in simulations run with cooler SST, and this large scale cloudy region is much more spatially and temporally coherent. By contrast, convection can be seen to propagate in squall-type systems from east to west across the domain for the first 20 simulated days in the warm-SST experiment. Even after the domain settles into a single moist and dry region, the speed of propagation of the individual convective systems is visibly different with changes in SST, with convection propagating fastest in the warm-SST case. In addition, comparison of Hovmoller plots of LWP and IWP (not shown) indicates that the elements seen to move from west to east with consistent speed for all three SSTs are uniformly (low) liquid clouds moving with the low-level steering current. Evidence for low-level convergence in the vicinity of the large moist regions can be seen in the increase in speed of movement of the low-level elements on the upstream (west) side of the moist region, and in the decreased speed of movement downstream (to the east) (Figs. 5.5d-5.5f).

As a final illustration of the differences in large-scale organization, cross-sections of




Figure 5.6: Five-day mean vertical and horizontal wind over the whole simulation domain for each simulation conducted in this study.

the five-day temporal mean of the horizontal and vertical wind velocity accumulated over the final five days of each simulation are plotted in figure 5.6. The indication from the precipitable water vapor that the overturning circulations in the non-sheared simulations are organized onto larger scales for low SST runs is supported to some extent in the horizontal wind plots. It can be seen that, while there are approximately two primary overturning cells for each value of SST, the number of embedded mesoscale overturning circulations dramatically increases with increasing SST. Consistent with the constraint imposed by mass conservation, the regular regions of convection observed for each SST are reflected in the large-scale dynamics; regions of ascent are co-located with regions of convection, and regions of descent can be observed in the clear regions between. It is interesting to note that the horizontal scale of the largest mode of overturning uniformly decreases with increasing SST for *both* sheared and unsheared simulations. While the horizontal scale differs, with the exception of the stratospheric winds, the magnitude of both horizontal and vertical motions is very consistent between runs with different SST and between sheared and unsheared cases. On a final note, it can be seen that the overturning circulations extend to the surface in the non-sheared simulations, while the overturning is largely confined to the free troposphere in above the atmospheric boundary layer in simulations with shear. As a result, while simulations with shear exhibit low cloud in the non-convective regions, the convection-free regions in the non-shear simulations are entirely devoid of cloud at all levels.

Overall, simulations with greater shear (not surprisingly) appear to be more consistent with observed monsoon conditions over the tropical west Pacific, with propagating convective clusters embedded in larger superclusters. Though clusters in the zero-mean wind case persist far longer than those observed in the real atmosphere, it is nevertheless encouraging to see organization on *scales* consistent with reality and in a manner similar to the *gregarious convection* hypothesis introduced by Mapes

(1993).

5.4.3 Response of Clouds and Precipitation to Changes in SST

Having demonstrated convergence to RCE and examined how convection organizes on the large scales in each simulation, we proceed to analyze how the equilibrium values of cloud, precipitation, and radiation state variables vary with SST. To this end, domain mean radiative fluxes and cloud condensate variables are computed over the last twenty days of each simulation—such a long averaging period is chosen so that the variability associated with transient convective systems in both zero-wind and TOGA shear cases is effectively smoothed. Domain mean vertical profiles of each GCE prognosed condensate variable for all three SST values are plotted in figure 5.7. Upon inspection, the most notable feature in each plot is the general upward vertical displacement in the maximum mass mixing ratio for each condensate species. This is consistent with an increase in the thickness of the troposphere with increasing surface temperature, which translates to a higher tropopause and melting layer.

Though there were significant differences in the large scale organization of convection between sheared and unsheared simulations, the general behavior of each condensate species with changes in SST is consistent for both sets of experiments. The maximum mixing ratio of graupel and pristine ice increases with increasing SST, while snow slightly decreases. The rain water mass mixing ratio increases significantly with increasing SST, while the maximum mixing ratio of non-precipitating cloud remains approximately constant. It is interesting to note that, as SST increases from 295 K to 300 K for both sheared and unsheared simulations, the cloud mixing ratio develops a bimodal structure. The lower mode is located at approximately 2.5 km above the surface, and is coincident with the maximum in rain mass mixing ratio. This mode is also coincident with the maximum in mean cloud condensation rate for each SST (not shown), and is likely associated with the domain and time mean lifting




Figure 5.7: Twenty-day time mean and full domain mean profiles of each condensate species in the GCE cloud microphysical scheme. In each plot, the blue line corresponds to simulations run with constant SST of 295 K, black to simulations run with SST of 300 K, and red to simulations run with SST of 305 K.

condensation level for each SST. The higher mode is located just below the mean freezing level, and is the dominant mode for simulations with shear. A maximum in the domain and time mean rate of condensation is colocated with the higher mode for zero-mean wind simulations, while the mean condensation rate just below the freezing level is much smaller for simulations with shear. The implication is that, while the higher mode is produced by condensation in the no-shear runs, another process is contributing to the accumulation of large amounts of liquid cloud just below the freezing level in simulations with shear. The possibility that the upper mode is associated with more prevalent mesoscale stratiform anvil circulations in sheared simulations is explored later in the discussion (section 5.5).

Twenty-day means and totals of condensate, precipitation, and radiation variables are listed in Tables 5.1 and 5.2 for the zero-mean wind and sheared cases, respectively. Consistent with higher values of saturation vapor mixing ratio with increasing mean atmospheric temperature, the mean precipitable water vapor increases monotonically with SST for both sets of experiments. Examination of liquid condensate variables reveals that, though the maximum mass mixing ratio does not change significantly (Fig. 5.7), the total amount of cloud water increases with SST—though the increase is still not as dramatic as for rain. The zero-mean wind case exhibits no clear trend in total ice, pristine ice, or graupel, though the amount of snow decreases steadily with increasing SST. By contrast, the sheared case exhibits clear increases in pristine ice mass, graupel, and total ice, along with a decrease in snow. Though sheared simulations are generally more cloud-covered than unsheared, each experiment exhibits a clear decrease in total cloud fraction with increasing SST. This decrease in cloudiness is entirely contained in the mid and upper level cloud, while the low-level cloud fraction exhibits a general increase. As was pointed out in the discussion of the large-scale organization, subsidence in unsheared simulations occurs through the depth of the troposphere in convection-free regions, while large-scale overturning circulations are

Table 5.1: List of 20-day averaged cloud and radiation fields for the set of simulations with zero base-state wind.

	No Wind		
	295 K	300 K	305 K
Mean PWV ($kg \cdot m^{-2}$)	21.03	31.03	47.05
Total Cloud ($kg \cdot m^{-2}$)	30331.8	37108.9	48510.3
Total Rain ($kg \cdot m^{-2}$)	26758.6	42053.0	62579.8
Total Ice ($kg \cdot m^{-2}$)	8178.3	9033.7	8643.0
Total Snow ($kg \cdot m^{-2}$)	11527.9	9896.4	8045.7
Total Graupel ($kg \cdot m^{-2}$)	21342.5	23151.7	22460.1
Total All Ice ($kg \cdot m^{-2}$)	41048.7	42081.7	39148.8
Clear Fraction	42.99	47.77	50.29
Total Cloud Fraction	57.01	53.23	49.71
High Cloud Fraction	50.38	46.50	45.38
Mid Cloud Fraction	5.24	4.28	2.16
Low Cloud Fraction	1.59	1.67	2.38
Column Integrated LHR ($K \cdot day^{-1}$)	7.80	10.68	14.67
Total condensation ($kg \cdot m^{-2}$)	281586.0	408132.0	538444.0
Total evaporation ($kg \cdot m^{-2}$)	165364.0	241309.0	304431.0
Total deposition ($kg \cdot m^{-2}$)	108227.0	98738.2	90921.1
Total sublimation ($kg \cdot m^{-2}$)	60115.8	48694.7	41472.8
Total melting ($kg \cdot m^{-2}$)	136852.0	152293.0	147909.0
Total freezing ($kg \cdot m^{-2}$)	5915.7	6817.2	6575.9
Mean Upward Vertical Velocity ($cm \cdot s^{-1}$)	1.37	1.30	1.40
Mean Downward Vertical Velocity ($cm \cdot s^{-1}$)	-44.6	-47.1	-41.9
Total Precipitation (mm)	163179.0	215510.0	282113.0
Mean Precip Intensity ($mm \cdot day^{-1}$)	1.999	2.636	3.467
Total Condensation + Deposition (mm)	389812.0	506870.0	629366.0
Precip Efficiency	41.86	42.52	44.83
OSR at TOA ($W \cdot m^{-2}$)	-99.00	-95.47	-91.62
OLR at TOA ($W \cdot m^{-2}$)	-258.83	-268.91	-281.87

located above the atmospheric boundary layer in simulations with shear. As a result, the low cloud fraction is *very* low in simulations with zero mean wind as compared to those run with shear.

Examination of the latent heat release and the cloud microphysical processes that contribute to LHR reveals a general increase in column integrated LHR with increasing SST. Condensation and evaporation increase with increasing SST for both experiments, while deposition and sublimation decrease. Freezing and melting both increase for the sheared simulation, while there is no clear trend for the nonsheared simulations. It is interesting to note that while both the mean updraft strength and subsidence velocity increase with SST for the sheared simulations, there is no clear trend for the nonsheared simulations. This may be due to the fact that a greater proportion of the precipitation produced in simulations without shear falls through convective updrafts, while more of the precipitation in convection that develops in the presence of shear tends to fall outside of the updraft cores.

Consistent with increases in the cloud and rain mass, the total accumulated surface precipitation increases with increasing SST for both experiments, as does mean precipitation intensity. Though the amount of ice deposition decreases with increasing SST, the amount of condensation is uniformly larger, and when surface precipitation is compared to a combination of the condensation and deposition, it can be seen that the overall precipitation efficiency increases. It is interesting to note that the total amount of condensed liquid and ice and total precipitation are significantly larger for simulations run with shear, while PE in sheared simulations is less. This is consistent with previously-reported findings (Ferrier et al. 1996) that the presence of vertical shear, and consequent advection of condensed liquid and ice away from the primary storm updraft, tends to lead to a decrease in convective system PE.

Not surprisingly, the top of atmosphere longwave flux increases with increasing SST, and consistent with greater cloud fraction, outgoing longwave radiation is lower

for sheared runs than for unsheared. While the domain-averaged top of atmosphere shortwave flux exhibits much greater temporal variability compared with the longwave (Fig. 5.2), there is a slight, but consistent, negative trend in the outgoing shortwave flux with increasing SST. Note that the *decrease* in outgoing shortwave radiation occurs in spite of an *increase* in the low cloud amount; an increase in low cloud fraction should contribute to increased domain mean reflected shortwave flux. It is interesting to note that, in spite of highly nonlinear interaction between convection, the large scale circulation, and cloud microphysical processes, the standard PE argument appears to hold; increases in PE are associated with decreases in mid and upper level cloud and decreases in outgoing shortwave radiation.

5.5 Discussion

The presence of a bimodal structure in the domain and time mean cloud water mass mixing ratio was noted in the previous section, and it was hypothesized that the upper mode, which existed only at SSTs of 300 and 305 K, might be associated with mesoscale circulations in the stratiform region of simulated mesoscale convective systems (e.g., Houze 2004). To assess whether increases in SST are associated with an increase in regions of stratiform precipitation relative to convective cores, an examination is first made of the vertical velocity and latent heating profiles in convective regions. In this case, profiles of vertical velocity and latent heat release are averaged over regions in which the upward vertical velocity exceeds $1.0 \text{ m} \cdot \text{s}^{-1}$ anywhere in the column, and the total condensate (liquid plus ice) mass mixing ratio exceeds $0.1 \text{ g} \cdot \text{kg}^{-1}$ over a depth that extends from 250 meters above the surface to 10 km (Fig. 5.8). Examination of profiles of vertical velocity for both unsheared (Fig. 5.8a) and sheared (Fig. 5.8c) simulations reveals that the mean vertical velocity below the freezing level consistently decreases with increasing SST, while ascent above the freezing

Table 5.2: List of 20-day averaged cloud and radiation fields for the set of simulations with base-state winds taken from averages over the TOGA COARE IFA.

	TOGA Wind		
	295 K	300 K	305 K
Mean PWV ($kg \cdot m^{-2}$)	29.94	43.66	63.37
Total Cloud ($kg \cdot m^{-2}$)	62930.5	81570.0	99110.5
Total Rain ($kg \cdot m^{-2}$)	38383.8	59894.4	89052.5
Total Ice ($kg \cdot m^{-2}$)	8462.1	9944.5	10969.6
Total Snow ($kg \cdot m^{-2}$)	12416.2	10616.1	8954.3
Total Graupel ($kg \cdot m^{-2}$)	21171.5	24395.6	25997.1
Total All Ice ($kg \cdot m^{-2}$)	42049.8	44956.2	45921.0
Clear Fraction	29.89	32.66	40.79
Total Cloud Fraction	70.11	67.34	59.21
High Cloud Fraction	57.86	50.84	43.60
Mid Cloud Fraction	4.12	5.86	3.72
Low Cloud Fraction	8.33	10.84	12.11
Column Integrated LHR ($K \cdot day^{-1}$)	10.57	13.99	18.38
Total condensation ($kg \cdot m^{-2}$)	463026.0	662206.0	856381.0
Total evaporation ($kg \cdot m^{-2}$)	314908.0	454307.0	572268.0
Total deposition ($kg \cdot m^{-2}$)	131266.0	114602.0	107345.0
Total sublimation ($kg \cdot m^{-2}$)	78819.2	56945.2	43195.9
Total melting ($kg \cdot m^{-2}$)	145949.0	160180.0	179918.0
Total freezing ($kg \cdot m^{-2}$)	6231.2	6830.9	7710.3
Mean Upward Vertical Velocity ($cm \cdot s^{-1}$)	0.58	0.72	0.82
Mean Downward Vertical Velocity ($cm \cdot s^{-1}$)	-46.7	-49.6	-50.0
Total Precipitation (mm)	199488.0	264349.0	346677.0
Mean Precip Intensity ($mm \cdot day^{-1}$)	2.441	3.224	4.253
Total Condensation + Deposition (mm)	594292.0	776808.0	963726.0
Precip Efficiency	33.57	34.03	35.97
OSR at TOA ($W \cdot m^{-2}$)	-106.11	-105.37	-102.93
OLR at TOA ($W \cdot m^{-2}$)	-254.17	-263.16	-275.53

level uniformly increases. Interestingly, the latent heat release does not follow suit; though there is an increase in LHR with SST above the freezing level, LHR remains nearly constant below the freezing level in simulations without shear, and increases in simulations with shear. It is also interesting to note that there is a single low-level mode in the latent heating profile for unsheared simulations with SST of 295 K, while a secondary mode develops at upper levels for SSTs of 300 K and 305 K. In contrast, the upper level mode exists for all three SSTs in sheared simulations. In comparison with idealized profiles of convective and stratiform heating in mesoscale convective systems (e.g., Houze 2004, his Fig. 4), the profiles of latent heat release appear consistent with a sum of the primary convective and stratiform heating modes for all three SSTs in the case with shear, and for SSTs of 300 K and 305 K in the case with no mean shear. Increases in domain and time mean vertical velocity above the freezing level are consistent with increased LHR in this region, while decreases in vertical velocity below the freezing level may be due to increased precipitation drag; recall that both the total surface precipitation and the rain mass mixing ratio increased with increasing SST for both sheared and unsheared simulations.

To assess whether the fraction of stratiform relative to convective precipitation increases with increasing SST, a convective-stratiform partitioning scheme is applied to the results of each simulation. The partitioning scheme classifies portions of precipitating convective systems into either convective or stratiform regimes based on a method introduced by Churchill and Houze (1984, hereafter CH). The CH scheme identifies convective precipitating regions as having twice the surface precipitation rate as the domain average. The two grid points on either side of an identified convective core are also classified as convective, as are any grid points in which the rain rate exceeds $20 \text{ mm} \cdot \text{h}^{-1}$. Tao and Simpson (1989) and Tao (1993, 2000) added two additional criteria for the purpose of identifying convective cores aloft that are not associated with copious surface precipitation. The modified CH scheme (which has




Figure 5.8: Vertical profiles of mean vertical velocity and latent heat release for simulations without and with shear. In each plot, the blue line corresponds to SST of 295 K, black to SST of 300 K, and red to SST of 305 K.

Table 5.3: Convective vs. stratiform fraction for precipitating systems in the last 20 days of each RCE simulation as computed using the GCE convective-stratiform partitioning method (see text for details).

Shear	No Wind			TOGA COARE		
SST	295 K	300 K	305 K	295 K	300 K	305 K
Convective	0.48	0.507	0.512	0.19	0.26	0.31
Stratiform	0.52	0.493	0.488	0.81	0.74	0.69

been referred to as the “GCE method”, Lang et al. 2003) also classifies grid columns as convective if the cloud mixing ratio exceeds a threshold (in this case either $0.5 \text{ g} \cdot \text{kg}^{-1}$ or half of the current-time maximum cloud water content) or if the updraft exceeds three $\text{m} \cdot \text{s}^{-1}$ in the region below the melting layer. In addition, if the vertical velocity exceeds five $\text{m} \cdot \text{s}^{-1}$ or one half of the current-time maximum vertical velocity in the stratiform region, the column is re-classified as convective. Finally, if the total liquid plus ice mixing ratio exceeds $1.5 \text{ g} \cdot \text{kg}^{-1}$ or one half of the maximum total mixing ratio, the column is classified as convective. The GCE convective-stratiform classification method is applied to the results of each RCE simulation, and average convective-stratiform fractions from the last 20 days of each simulation are listed in Table 5.3. It is interesting that, while the strength of the mean vertical motion in stratiform regions *increases* with increasing SST, the proportion of stratiform precipitation relative to convective *decreases* with increasing SST. The combination of increased vertical velocity and latent heat release above the melting layer with decreased stratiform fraction indicates that, with increases to SST, the vigor of the circulations in the stratiform region of convection increases with increasing SST.

As was discussed in section 5.1, the thermostat feedback hypothesizes an *increase* in high cloud fraction, ice cloud mass, and outgoing shortwave radiation with increasing SST. The key relationship is between increases in convective mass flux and outgoing shortwave radiation; if the outgoing shortwave flux does not increase with increasing SST, then the thermostat hypothesis fails, as the feedback hypothesizes regulation of SST due to increased shortwave reflection. The results presented above–

particularly those in Tables 5.1 and 5.2 do not support the operation of a thermostat-like feedback. In simulations with and without shear, both the high cloud fraction and the outgoing shortwave radiation *decrease* with increases in SST. Though the total ice mass increases for the case with shear, because it does not lead to increases in outgoing shortwave radiation, it is not consistent with the thermostat.

In contrast to the thermostat, the infrared iris postulates an increase in precipitation efficiency with SST, and more importantly, a decrease in high cloud (and corresponding increase in the cloud-free area relative to that covered with cloud). Indeed, this is found to be the case—increases in SST are uniformly associated with increases in PE, and with a steady decrease in high cloud fraction. Outgoing shortwave radiation exhibits a slight decrease, while the outgoing longwave radiation increases significantly. In addition, though the magnitude of the mean vertical motion in stratiform regions increases with increasing SST, the fraction of stratiform precipitation relative to convective decreases. While these results are compelling in that the same general conclusion can be drawn for both sheared and unsheared simulations, it should be noted that an examination of the characteristics of equilibrium states for fixed SST experiments explains only half of the operation of a feedback. To fully assess the operation of a cloud feedback mechanism, it would be necessary to employ a coupled ocean-atmosphere model with interactive (predicted) sea surface temperature. In this type of experiment, the model would first be run to an equilibrium state for a fixed SST, then either warmer or cooler SST would be applied, and the model allowed to freely run with predicted sea surface temperature until a new coupled atmosphere-ocean equilibrium state was established. The character of the final equilibrium state, and the manner in which the cloud and precipitation processes regulate the transition would give a more realistic indication of the operation of the feedback mechanism.²²

²²It should be noted that, because the ocean temperature adjusts to an equilibrium state at a much slower rate than does the atmosphere, coupled atmosphere-ocean simulations would likely need to be run for much longer time periods than the fixed-SST experiments discussed above.

Although the research presented in this chapter is unique compared with previous studies in that it employs a large-domain observation constrained cloud resolving model to study the response of precipitating cloud systems to surface warming, the applicability of the results to the real atmosphere is limited. Perhaps the most serious drawback of the current set of experiments is the use of the two-dimensional framework. Some studies have declared the differences between the statistics of the state produced by two and three dimensional simulations to be small (e.g., Tao and Soong 1986, Tao et al. 1987, Grabowski et al. 1998), however, it has also demonstrated that 2D simulations artificially constrain convergent/divergent circulations, as well as gravity wave propagation. Specifically, as was pointed out by Bretherton and Smolarkiewicz (1989), subsidence associated with convective heating is artificially spread to more remote locations in 2D, increasing the horizontal distance over which convection affects its environment. In addition, it has been noted (Houze 2004) that mesoscale convective systems tend to develop a balanced horizontal mesoscale circulation at the top of the melting layer in the stratiform precipitation region. Such “mesoscale convective vortices” (Houze 2004) are associated with a positive potential vorticity anomaly that develops in response to the vertical gradient of latent heating in stratiform regions of convection. A slab symmetric two-dimensional model, such as that used in the current study, cannot represent the development of this mesoscale circulation, and hence simulated convective systems lack a key dynamical component. Finally, a two dimensional model cannot realistically simulate the interaction of convective cold pools with the environment—in two dimensions, the constraint of mass conservation requires low-level outflow to maintain its strength to a large distance from the convective updraft, while in 3D, the outflow velocity is allowed to decrease appropriately with distance (Tompkins 2000).

Differences between 2D and 3D simulation geometry notwithstanding, it should be noted that the results of the current experiment do not require exact represen-

tation of convective structure on fine scales; it is sufficient to show how the time and domain mean characteristics of convective ensembles change for different equilibrium states. The key question is whether the manner in which the characteristics of clouds and precipitation change with changes to SST are consistent between two and three-dimensional simulations. Presentation of three-dimensional results is beyond the scope of the current study, but work has begun to implement and test a three-dimensional version of the GCE for use in future experiments.

5.6 Conclusions

In this chapter, idealized two-dimensional cloud resolving model simulations have been performed to assess the response of cloud and precipitation processes to changes in sea-surface temperature. In the process, the model results were examined for evidence of whether an observation constrained CRM supports the operation of either the thermostat or Iris feedback mechanisms. It was found that the character of the convective organization differed significantly between sheared and unsheared simulations—in unsheared simulations, convection organized into multiple bands with greater number of bands for higher values of SST, while sheared simulations uniformly generated a single moist and single dry region. In sheared simulations, the moist region was characterized by convection that developed in clusters and propagated from east to west across the moist region in a matter of a few days, while the dry region contained only low clouds that moved from west to east with the low level steering flow. It was also found that, with increasing SST, the high cloud fraction decreased, while the bulk amount of liquid cloud increased. The precipitation amount, intensity, and efficiency increased for increasing SST, while the outgoing shortwave radiation exhibited a slight decrease. Upward vertical motion below the melting layer decreases with increasing SST, while vertical motion above the melting layer increases, and the

proportion of stratiform precipitation decreases with increasing SST. Overall, the results of this study appear to support the operation of an iris-like response of clouds to surface warming, though the constraint of 2D geometry and fixed SST limit the applicability of the results to the true atmosphere.

Chapter 6

Summary and Conclusions

6.1 Summary of Key Results and General Conclusions

This study has addressed the response of cloud and precipitation processes to changes in SST over the equatorial western Pacific using large domain CRM simulations. Because of the demonstrated variability in CRM results associated with uncertain cloud microphysical parameters, the bulk of the work presented above has dealt with the topic of parameter estimation, and specifically with issues of cloud and precipitation assimilation on cloud-resolving scales. Because cloud microphysical parameters are highly nonlinearly related to the cloudy state, traditional data assimilation methods, which are based on a quasi-linear Gaussian-statistics estimation framework were initially assumed not to be suitable. Instead, Markov chain Monte Carlo (MCMC) methods, which make no assumptions on the character of the state space or the linearity of the parameter-state relationship, were used to first identify key microphysical parameters, then to characterize the nature of the state space. Because MCMC is relatively new to the atmospheric sciences, a proof-of-concept study was performed, which applied the MCMC method to a well-understood cloud property retrieval prob-

lem. It was shown that MCMC can be used to effectively characterize the PDFs of liquid cloud effective radius and optical depth, and in the process, it was demonstrated that MCMC provides a robust method with which to characterize the uncertainty in nearly any inversion procedure.

Following the proof-of-concept retrieval study, MCMC was applied to a two dimensional version of the NASA Goddard Cumulus Ensemble (GCE) model to characterize the uncertainty in its cloud microphysical parameters, and to identify which might serve as control variables for assimilation. So-called “identical twin” experiments were performed, in which “true” parameter values were defined and used to generate simulated observations. These observations were then used in the MCMC algorithm to build parameter PDFs, and to characterize the sensitivity of simulation results to changes in cloud microphysical parameters. In the process, the manner in which changes in parameters related to changes in the resulting GCE state was examined for evidence of monotonic linear relationships—the existence of such relationships would enable the use of more traditional (and much less computationally expensive) data assimilation approaches. It was found that three of the nine parameters selected for estimation exerted the most control over the evolution of cloud, precipitation, and radiation fields, and that these parameters did indeed exhibit a quasi-linear relationship with certain observations.

A second set of idealized MCMC experiments was conducted, in which observation error was reduced consistent with the launch of current and near future sensors. Because of the increase in information content with decreasing observation error, only three observations were used to constrain the state; the cloud top temperature, liquid water path, and ice water path. It was found that, even with a subset of observations that included none of the radiation state variables, the estimation returned PDFs that nearly exactly matched those in the control case. The same three key parameters were identified by the estimation, and in addition, it was found that constraint of cloud

content and height caused the precipitation and longwave fluxes and heating rates to closely adhere to observed values. The shortwave fluxes and heating rates were not as well constrained, likely because these state variables depend directly on the ice density and PSD slope intercept—parameters that were not well-characterized by the estimation. The conclusion was drawn that MCMC could be used to evaluate the robustness of interactions between precipitation, clouds, and radiation in the model, and that relationships between these fields are realistically simulated in the GCE model.

The quasi-linear parameter-state relationships identified in idealized experiments were used to construct a variational assimilation system, however, the variational estimate never converged to a solution—even for idealized simulated observation experiments. In subsequent MCMC experiments, the parameter state space was shown to contain multiple modes—in such cases the variational method nearly always fails to produce a unique solution. In place of variational assimilation, the MCMC algorithm was used to estimate the actual parameter PDFs from TRMM retrieved cloud, precipitation, and radiation statistics. It was shown that the three parameters identified in the idealized experiments were again the most important parameters to constrain, and that the state space again contained multiple modes. The true mode was found by subsetting observations and considering the range of parameter values associated with the most likely set of solutions. This set of parameters represents those for which the model most faithfully simulated the statistics of the climate state.

In the final component of this work, the parameter values returned by MCMC were implemented in the GCE model, and the resulting observation-constrained model was used to investigate the response of cloud and precipitation processes to changes in SST. Two different experiments were performed—one in which base-state winds were set equal to zero, and one in which base state winds were set equal to mean flow conditions typical of the active phase of the monsoon over the equatorial western

Pacific. In each set of experiments, the sea surface temperature was fixed at values of 295 K, 300 K, and 305 K, and the characteristics of the resulting equilibrium state were examined for evidence of consistent changes to the tropical hydrologic cycle with increasing SST. Simulations without shear took longer to converge to RCE (≈ 40 days) than simulations with shear (≈ 20 days), and it was found that changes to SST altered the mode and scale of the convective organization for both sets of experiments, and that the manner in which convection organized on large scales was fundamentally different between sheared and unsheared simulations. While unsheared simulations exhibited multiple overturning circulations that extended through the depth of the troposphere, simulations with shear developed only two circulation cells regardless of the value of SST. In addition, these circulations were confined to the free troposphere above the top of the atmospheric boundary layer. Consistent with what is observed, the ascending branch of the each of the circulations was characterized by higher values of precipitable water vapor and regions of convection, while the descending branch was nearly cloud free for the non-sheared simulations, and contained moderate amounts of low cloud in the simulations run with shear. Sheared simulations exhibited propagating squall-type systems that moved from east to west through the large-scale moist regions; forming on the upshear side and dissipating on the downshear side, and with speed of propagation that was observed to increase with increases in SST.

Statistics on the latent heat release, precipitable water vapor, cloud cover, condensate mass, and radiative fluxes were accumulated over the last 20 days of simulation to assess the response of the hydrologic cycle to changes in SST. It was found that, with increasing SST, the cloud fraction decreased, the low cloud amount increased, the outgoing longwave radiation and column integrated latent heat release increased, and the shortwave flux at the top of the atmosphere decreased. At the same time, consistent with an increase in latent heating, the total surface precipitation increased, as did the precipitation intensity and efficiency. It was found that, in spite of the

complex interaction between the large scale circulation, convection, and cloud microphysical processes, that the anticipated zero-order effect of changes in PE generally holds regardless of the presence or absence of shear. Increases in PE were uniformly associated with decreases in high cloud fraction and decreases in outgoing shortwave radiation, however, the general increase in the amount of cloud condensate led to an increase in total ice mass in the domain. When the results were applied to an examination of the well-known thermostat and iris feedback hypotheses, it was found that the statistics of the state at RCE contradicted the thermostat, which depends on an increase in shortwave flux with increasing SST. Rather, the results appear to lend credence to an iris-like effect, as increases in SST lead to increasing precipitation efficiency and decreasing high cloud.

6.2 Implications of Current Work and Suggestions for Future Research

The above study is unique in a number of significant ways. It represents one of the first times MCMC has been used in the atmospheric sciences to explicitly evaluate the assumptions that underlie an optimal-estimation based inversion. It also contains a thorough examination of the sensitivity of CRM output to changes in those microphysical parameters associated with the assumed particle size distribution, and is the first recorded application of the MCMC algorithm to a cloud resolving model. Use of TRMM retrieved statistics of clouds, precipitation, and radiation represent the first time domain and time mean TRMM observations have been assimilated in a CRM, and the first time statistics of the cloudy state have been used to constrain cloud microphysical parameters. In addition, the current study is unique in that it is the first time an observation-constrained CRM has been expressly used to evaluate the operation of hypothesized cloud-climate feedback mechanisms.

The work presented above opens the door to an entirely new avenue of research focused on the evaluation of uncertainty contained in nearly any numerical representation of a process or physical relationship. Though optimal estimation is being increasingly used to retrieve estimates of the atmospheric state, the robustness of the estimate depends directly on the validity of the Gaussian error assumptions, and on poorly known, empirically specified forward radiative transfer model parameters. As was done in the MODIS retrieval example, MCMC could be used to evaluate the robustness of Gaussian error statistics, and to estimate the sensitivity of the result to changes in unknown forward model parameters for a wide range of retrieval algorithms.

Having applied MCMC to a two-dimensional CRM, it is a natural extension of this work to investigate whether parameter PDFs change with time and geographic location. Since model physics parameters are typically set constant over the length of a model run, and are most often left unchanged with changes in season or location, knowledge of the magnitude and character of this variation should prove to be extremely valuable. It was also noted above that many studies have documented difficulties associated with the use of a two-dimensional framework. In this study, a two dimensional CRM was used for computational reasons—both for the application of MCMC, and for the large-domain experiments, however, future experiments will need to explore use of a three-dimensional model. With parallelized code, it may be possible to re-run the RCE experiments in three dimensions, however, application of the MCMC algorithm to a 3D CRM will likely be challenging. If simulations are run over a very short time period, and the evolution of a single convective system is simulated, then MCMC with a three-dimensional CRM may be barely feasible. If we consider a 64×32 gridpoint simulation with four kilometer grid spacing run for three hours, this would be computationally equivalent to running a two-dimensional (64×1 gridpoint) simulation for 96 hours—three times the length of the simulations used

in the current study. This would translate to a wall-clock time of approximately five minutes per simulation, although the model might be made to run faster by streamlining the code. In any case, this is a task that will need to be looked into in the future.

Finally, RCE experiments were only able to assess the *response* of the tropical hydrologic cycle to perturbations in SST. The more compelling question is whether cloud processes act to *regulate* changes to SST. To address this question, future experiments will need to be run with a coupled ocean model—in such a framework, SST could be fixed until the model reaches RCE, then an SST perturbation could be applied and the SST allowed to adjust in the presence of clouds. The manner in which clouds modulate a return to equilibrium, and the character of the final equilibrium state itself would provide the information needed to more fully characterize the operation of cloud-climate feedback mechanisms.

Appendix A

GCE Cloud Microphysics Scheme

A.1 Introduction

This appendix describes in detail the cloud microphysical parameterization used in the current version of the Goddard Cumulus Ensemble model. The description is based on a 2D version of the GCE obtained from Wei-Kuo Tao by way of Chung-Lin Shie on 9 July 2004. The intent of this appendix is not to re-derive each of the equations that form the foundation of the scheme, but instead to clearly document the code with special attention to the use of specified empirical parameters. To this end, each process equation in the cloud microphysics code is documented and a list of parameters is accumulated for each equation. For each equation process is documented exactly as it is described in the code, and any significant differences from the source theory (primarily Rutledge and Hobbs 1983, 1984 and Lin et al. 1983) are noted. Following the description, a discussion of the GCE cloud microphysical parameterization scheme is offered, which follows from the discussion of the state of the art in cloud microphysics modeling contained in Khain et al. (2000), and from observations of ice and mixed phase particle distribution properties (Heymsfield et al. 2002). The remainder of this chapter is organized as follows: section A.2 describes

the mathematical foundation for the single-moment bulk microphysical scheme and documents the key assumptions used. Section A.3 contains a detailed description of each process as it is represented in the GCE code, along with a list of the tunable parameters. Section A.4 contains a discussion of the strengths and limitations of the current three-class ice scheme as compared with other methods of cloud microphysical parameterization and with observations of cloud ice microphysical properties.

A.2 Basic Formulation and Key Assumptions

A.2.1 Particle Size Distributions

In the GCE, condensate is separated into five discrete species: cloud, ice, rain, snow, and graupel/hail. Cloud droplets and ice crystals are defined as liquid and solid water particles small enough to have negligible fall velocity. They are assumed to be monodisperse, e.g. there is no variation in particle diameter throughout the cloud and hence no need for a particle size distribution function. Once cloud and ice grow through vapor deposition past a certain threshold mass mixing ratio, they are converted to precipitation; rain and snow, respectively. These species (as well as graupel/hail) are assumed to follow an inverse-exponential (Marshall-Palmer) particle size distribution, which can be written for species x as

$$N_x(D_x) = N_{0x} \exp(-\lambda_x D_x). \quad (\text{A.1})$$

Where N_{0x} is the slope intercept, and the slope parameter λ_x is derived by multiplying both sides of (A.1) by the particle mass and integrating over all particle diameters. In the process, the left hand side becomes the air density multiplied by the total mass mixing ratio. To perform the integration on the right hand side, we first assume the

particle mass can be written as

$$M_x(D) = \rho_x \left(\frac{\pi}{6} D^3 \right), \quad (\text{A.2})$$

then integrate the result by parts thrice to obtain the following expression for the slope of the size distribution.

$$\lambda_x = \left(\frac{\pi \rho_x N_{0x}}{\rho q_x} \right)^{0.25} \quad (\text{A.3})$$

Here, q_x and ρ_x are the mass mixing ratio and particle density of species x .

The assumption in the equation for mass is that the particle volume can be represented as a sphere—this is not a bad assumption as long as the portion of the sphere occupied by air is accounted for when specifying the particle density for each condensate species. However, the second significant assumption that must be made when performing the integration is that the particle density is constant with changes in diameter. Even this is not necessarily a bad assumption as long as the particle shape (crystal habit) remains the same with changes in diameter. Unfortunately, this is not typically the case as rain drops deform at larger sizes and snow (falling ice) changes habit with size as dendritic branches grow and riming and aggregation occur. The further assumption is made in the GCE that the particle density remains constant with space and time, as does the intercept parameter. The values of each intercept parameter and particle density currently used in the code, are documented below, as are values from other studies and/or models and parameter values resulting from MCMC experiments documented in chapter 4. An idea of the spread in particle density for graupel and snow can be seen in Locatelli and Hobbs (1974).

	Current	RH84	LFO83	Lord 1983	MCMC Estimate
N_{0R}	0.08	0.08	0.08	0.22	0.05
N_{0S}	0.16	0.04	0.03	0.03	1.00
N_{0G}	0.08	0.04	0.0004	0.0004	1.5
ρ_R	1.0	1.0	1.0	1.0	N/A
ρ_S	0.1	0.1	0.1	0.1	0.35
ρ_G	0.4	0.4	0.917	0.3	0.55

A.2.2 Fall Velocity Formulation

Determination of the fall velocities of precipitating ice and liquid particles follow Rutledge and Hobbs (1983, 1984, hereafter RH83 and RH84). In the code, each of the fall velocities is only used in the accretion processes and is computed once, at the beginning of the cloud microphysics subroutine. The equation for rain fall velocity consists of a polynomial fit to laboratory data from Gunn and Kinzer (1949), and follows RH83 (A3). The form of the expression is

$$V_r = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[a_0 + \frac{a_1}{\lambda_r} + \frac{a_2}{\lambda_r^2} + \frac{a_3}{\lambda_r^3} \right]. \quad (\text{A.4})$$

Here, ρ_o is the (near) surface density and ρ is the density at the current level. The snow fall velocity is parameterized as

$$V_s(D_s) = a_s D_s^{b_s} \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}}, \quad (\text{A.5})$$

where a'' and b'' are derived from field observations (Locatelli and Hobbs 1974). Particle mass is assumed to follow the relationship

$$M(D_s) = \frac{\pi}{6} \rho_s D_s^3 \quad (\text{A.6})$$

Using the above equation and integrating over an exponential size distribution gives the snow fall speed as

$$V_s = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{a_s \Gamma(4 + b_s)}{6 \lambda_s^{b_s}} \right]. \quad (\text{A.7})$$

The expression follows RH83 (their eqn. A5), and is based on the assumption that the snow is composed of aggregate and dendrite habits (Locatelli and Hobbs, 1974). The assumption of constant ρ_s , a_s , and b_s indicates the further assumption that snow maintains a single habit consistent with these constants. To be self-consistent, this habit assumption should propagate to the collection efficiencies in the accretion processes as well as to the capacitance and ventilation assumptions in the deposition/sublimation relationships—these have yet to be checked for consistency. It should be noted that crystals with different habits (and degrees of riming) have been shown to exhibit widely differing fall velocities, evaporation, and collision efficiencies. Other investigators have addressed this problem by allowing more than one species of precipitating pristine ice (Lynn et al. 2004).

As clouds age and cloud water and ice are allowed to interact, riming will occur. With riming, the properties of ice crystals change dramatically, necessitating definition of a third ice phase—graupel—which includes crystals with moderate to heavy amounts of rime. In the case of graupel, in addition to habit uncertainty with associated effect on fall speed and collection efficiency, the degree of riming must now be taken into account—a factor which can have an even greater effect on the condensate interactions. The graupel fall velocity is parameterized in a similar manner to that of snow

$$V_g = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{a_g \Gamma(4 + b_g)}{6 \lambda_g^{b_g}} \right], \quad (\text{A.8})$$

which follows RH84 (their eqn. A3), and is also derived in Locatelli and Hobbs (1974).

Similar to snow, the graupel fallspeed relation is a function of particle diameter

$$V_g(D_g) = a_g D_g^{b_g} \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}}. \quad (\text{A.9})$$

The pertinent constants and parameters are listed below for each equation in this subsection.

Name	Value	Definition
a_s	78.63	Snow fallspeed constant
b_s	0.11	Snow fallspeed exponent
a_g	351.2	Graupel fallspeed constant
b_g	0.37	Graupel fallspeed exponent

A.3 Documentation of GCE Cloud Microphysics Code

Having documented the general basis for the GCE single moment bulk microphysical parameterization, a detailed description of the processes contained in the code is now presented. In each case, (where possible) reference is made to the source from which the process is derived, and empirical parameters are highlighted. Note that no attempt has been made in this study to improve this code, only to estimate the PDFs of key parameters and characterize their associated uncertainty.

A.3.1 (PSAUT) Autoconversion of ice to snow

This process is designed to simulate the transition from growing ice particles to falling snow and requires a threshold above which it will operate. This is specified arbitrarily as the mass of an ice particle with diameter of 500 microns. The autoconversion equation roughly corresponds to RH83 (their eqn. A19)

$$PSAUT = \frac{1}{\Delta t} [q_i - \rho M_{max} n_0 \exp(\beta(T - T_0))], \quad (\text{A.10})$$

where M_{max} is the threshold ice mass for a hexagonal plate with $D_i = 500 \mu m$, n_0 is a specified ice number concentration, and β is the constant used in Fletcher's (1962) equation that relates active ice nuclei to temperature. The second term in the autoconversion equation represents the maximum mass mixing ratio of ice possible in the volume, assuming the number concentration of cloud ice follows Fletcher (1962), which gives the number of active ice nuclei at a given temperature (T) as

$$n_c = n_0 \exp[\beta(T_0 - T)], \quad (\text{A.11})$$

where T_0 is the temperature at freezing. From RH83 (their eqn. A17) the ice mass and diameter are related by

$$D_i = 16.3 M_i^{0.5} = 16.3 \left(\frac{q_i \rho}{n_c} \right)^{0.5}. \quad (\text{A.12})$$

Note that the relationship is only applied at temperatures below freezing.

Name	Value	Definition
β	-0.6	Beta parameter in number concentration of cloud
n_0	$1.0 \times 10^{-8} \text{ cm}^{-3}$	Number concentration of cloud
M_{max}	$9.4 \times 10^{-7} \text{ g}$	Threshold ice mass for autoconversion to snow
D_i	$500 \mu m$	Threshold ice diameter for autoconversion to snow

A.3.2 (PSACI) Accretion of ice onto snow

When the continuous collection equation for accretion of ice onto snow

$$\frac{dM(D_s)}{dt} = \frac{\pi}{4} D_s^2 V_s(D_s) q_i E_{si} \quad (\text{A.13})$$

is integrated over the snow particle size distribution, the result is the rate of accretion of ice onto snow in the volume

$$PSACI = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi}{4} \frac{N_{0s} a_s E_{si} q_i \Gamma(3 + b_s)}{\lambda_s^{3+b_s}} \right]. \quad (\text{A.14})$$

where a_s and b_s are taken from the snow fall speed relation above. The relationship follows RH83 (their eqn. A21), and the value of the ice-snow collection efficiency parameter E_{si} is listed below. The process is only allowed to operate at temperatures less than freezing.

Name	Value	Definition
E_{si}	0.1	Collection efficiency for ice by snow

A.3.3 (PSACW) Accretion of cloud onto snow (riming)

$$PSACW = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi}{4} \frac{N_{0s} a'' E_{sc} q_c \Gamma(3 + b_s)}{\lambda_s^{3+b_s}} \right]. \quad (\text{A.15})$$

This is nearly identical to the ice accretion above, and follows RH83 (their eqn. A22). The primary difference is in the magnitude of the collection efficiency E_{sc} , and that this process is allowed to operate at temperatures above and below freezing. The process produces snow at sub-freezing temperatures and rain above freezing.

Name	Value	Definition
E_{sc}	0.5	Collection efficiency for cloud water by snow

A.3.4 (PWACS) Collection of snow by cloud

This process produces graupel through collisions between cloud water and snow, and operates according to

$$PWACS = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi^2}{24} \frac{N_{0s} n_c a'' E_{sc} q_c \Gamma(6 + b_s)}{\lambda_s^{6+b_s}} \right], \quad (\text{A.16})$$

which follows RH84 (their eqn. A10). Here, n_c is the number concentration of cloud droplets, given by

$$n_c = \frac{\rho q_c}{M_c}, \quad (\text{A.17})$$

and the process only operates at temperatures above freezing.

Name	Value	Definition
n_c	$2.5 \times 10^8 \text{ m}^{-3}$	Number concentration of cloud droplets

A.3.5 (PRACI) Accretion of ice onto rain

This equation produces graupel as its end result, and is only allowed to operate at temperatures less than freezing. It is analogous to RH84 (A5) from the continuous collection equation integrated over the rain drop particle size distribution using the V_R relationship documented above.

$$PRACI = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi}{4} \frac{N_{0r} E_{ri} q_i}{\lambda_r^3} \left(a_0 \Gamma(3) + \frac{a_1 \Gamma(4)}{\lambda_r} + \frac{a_2 \Gamma(5)}{\lambda_r^2} + \frac{a_3 \Gamma(6)}{\lambda_r^3} \right) \right]. \quad (\text{A.18})$$

In this equation, E_{ri} is the collection efficiency for ice-rain collisions, and a_0 , a_1 , a_2 , and a_3 are polynomial fit coefficients for the rain fall speed. Note that the constants are taken from RH83 (their eqn. A2), in which inverse λ_r and gamma functions are replaced with diameter of rain raised to the same power. This process is a sink for cloud ice.

Name	Value	Definition
E_{ri}	1.0	Collection efficiency of ice particles by rain

A.3.6 (PIACR) Accretion of rain or graupel onto ice

This equation arises from the continuous collection equation integrated over the rain drop particle size distribution, uses the V_R relationship, and only operates at temper-

atures below freezing. It follows RH84 (their eqn. A7) and represents contact freezing of rain to produce graupel.

$$PIACR = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi^2 N_{0R} E_{ri} \rho_R q_i n_i}{24 \lambda_R^6} \left(a_0 \Gamma(6) + \frac{a'_1 \Gamma(7)}{\lambda_R} + \frac{a'_2 \Gamma(8)}{\lambda_R^2} + \frac{a'_3 \Gamma(9)}{\lambda_R^3} \right) \right]. \quad (\text{A.19})$$

Here, n_i is the number concentration of ice, which is parameterized as

$$n_i = \frac{\rho q_i}{\overline{M}_i}, \quad (\text{A.20})$$

where $\overline{M}_i$ is the average mass of an ice crystal.

Name	Value	Definition
n_i	$1.67 \times 10^8 \text{ m}^{-3}$	Number concentration of ice particles
E_{ri}	1.0	Collection efficiency of ice particles by rain

A.3.7 (PRAUT) Autoconversion of cloud to rain

This process is designed to simulate the transition from monodisperse cloud growing by various processes to falling rain.

$$PRAUT = \alpha (q_c - q_0) \quad (\text{A.21})$$

This formulation is similar to RH83 (their eqn. A7), where α is a rate constant and q_0 is the cloud concentration threshold for autoconversion. This process is not restricted to any particular temperature range.

Name	Value	Definition
α	1.0×10^{-3}	Rate constant
q_0	$1.5 \times 10^{-3} \text{ kg} \cdot \text{kg}^{-1}$	Threshold cloud mass mixing ratio

A.3.8 (PRACW) Accretion of cloud onto rain

This process produces rain mass at the expense of cloud mass

$$PRACW = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi}{4} \frac{N_{0r} E_{rc} q_c}{\lambda_r^3} \left(a_0 \Gamma(3) + \frac{a_1 \Gamma(4)}{\lambda_r} + \frac{a_2 \Gamma(5)}{\lambda_r^2} + \frac{a_3 \Gamma(6)}{\lambda_r^3} \right) \right], \quad (\text{A.22})$$

where E_{rc} is the cloud-rain collection efficiency. As with PRACI, this equation is derived from the continuous collection equation integrated over the rain drop size distribution and using the rain population fall velocity. The process operates at any temperature, and the equation follows RH83 (their eqn. A9).

Name	Value	Definition
E_{rc}	1.0	Collection efficiency for cloud by rain

A.3.9 (PSFW) Bergeron processes for snow with cloud

This is the production of snow mass through vapor transfer from cloud to snow by the Bergeron process, and is computed following LFO83 (their eqn. 33) as

$$PSFW = \Delta t \left[q_i \left(\alpha_1 m_{i50}^{\alpha_2} + \rho q_c \pi E_{iw} U_{i50} R_{i50}^2 \right) \frac{\alpha_1 (1 - \alpha_2)}{m_{i50} \left(m_{i50}^{(1-\alpha_2)} - m_{i40}^{(1-\alpha_2)} \right)} \right]. \quad (\text{A.23})$$

Here, for a 50 micron radius reference crystal, U_{i50} and m_{i50} are the terminal velocity and mass, respectively, of a 50 μm radius reference crystal. To account for different reference radii, R_{i50} is specified as the actual radius. α_1 and α_2 are temperature-dependent parameters taken from Koenig (1971), and are set equal to 31 different values—the value used is chosen based on a temperature index between -1.0 and -31.0 degrees Celsius.

Name	Value	Definition
U_{i50}	100.0	Terminal velocity of a 50 micron ice crystal
R_{i50}	1.0×10^{-2}	Radius of a 50 micron ice crystal
m_{i50}	6.0×10^{-8}	Mass of a 50 micron ice crystal
m_{i40}	3.075×10^{-8}	Mass of a 40 micron ice crystal

A.3.10 (PSFI) Bergeron processes for snow with ice

This is the production of snow due to Bergeron sublimation transfer from ice particles to snow

$$PSFI = q_i \left[\frac{\alpha_1(1 - \alpha_2)}{m_{i50}^{(1-\alpha_2)} - m_{i40}^{(1-\alpha_2)}} \right], \quad (\text{A.24})$$

where the pertinent constants are identical to those in the PSFW process above.

A.3.11 (PGACS) Accretion of snow onto graupel (dry and wet)

This is a production process for graupel, and is written as

$$PGACS = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\rho \pi^2 N_{0s} N_{0g} E_{gs} \rho_s |V_g - V_s| \left(\frac{5.0}{\lambda_s^6 \lambda_g} + \frac{2.0}{\lambda_s^5 \lambda_g^2} + \frac{0.5}{\lambda_s^4 \lambda_g^3} \right) \right], \quad (\text{A.25})$$

where E_{gs} is the graupel-snow collection efficiency. This equation is, as before, obtained by integrating the continuous collection equation, but now the expression takes into account the differential fall velocities of snow and graupel and integrates over both size distributions. Here, we assume dry and wet accretion operate identically, and the collection efficiency is set low. The equation follows RH84 (their eqn. A14).

Name	Value	Definition
E_{gs}	0.1	Collection efficiency for snow by graupel

A.3.12 (PGACW) Accretion of cloud onto graupel

This process accounts for collection of cloud droplets onto graupel, and is written as

$$PGACW = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi}{4} \frac{E_{gc} N_{0g} q_c a_g \Gamma(3 + b_g)}{\lambda_g^{3+b_g}} \right] \quad (\text{A.26})$$

where E_{gc} is the collection efficiency for accretion of cloud interacting with graupel. Because negligible fall velocity is assumed for cloud droplets, this expression is simpler than the accretion of snow onto graupel. See RH84 (their eqn. A11) for the original expression.

Name	Value	Definition
E_{gc}	1.0	Collection efficiency for cloud by graupel

A.3.13 (PGACI) Accretion of ice onto graupel

This equation, which follows RH84 (their eqn. A12), represents collection of pristine ice by graupel

$$PGACI = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\frac{\pi}{4} \frac{E_{gi} N_{0g} q_i a_g \Gamma(3 + b_g)}{\lambda_g^{3+b_g}} \right], \quad (\text{A.27})$$

where E_{gi} is the collection efficiency for ice by graupel. This process is only allowed to occur at temperatures below freezing.

Name	Value	Definition
E_{gi}	0.1	Collection efficiency for ice by graupel

A.3.14 (PGACR) Accretion of rain onto graupel

$$PGACR = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\rho \pi^2 N_{0R} N_{0G} E_{gr} \rho_r |V_G - V_R| \left(\frac{5.0}{\lambda_R^6 \lambda_G} + \frac{2.0}{\lambda_R^5 \lambda_G^2} + \frac{0.5}{\lambda_R^4 \lambda_G^3} \right) \right] \quad (\text{A.28})$$

This equation follows RH84 (their eqn. A13), and is very similar to that of collection of snow by graupel (PGACS) above, where E_{gr} is the graupel-rain collection efficiency.

Name	Value	Definition
E_{gr}	1.0	Collection efficiency for rain by graupel

A.3.15 (PRACS) Accretion of snow onto rain

This process produces graupel through contact freezing of rain with snow, and produces rain at temperatures greater than freezing. It follows RH84 (their eqn. A9)

$$PRACS = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\rho \pi^2 N_{0r} N_{0s} E_{sr} \rho_s |V_r - V_s| \left(\frac{5.0}{\lambda_s^6 \lambda_r} + \frac{2.0}{\lambda_s^5 \lambda_r^2} + \frac{0.5}{\lambda_s^4 \lambda_r^3} \right) \right] \quad (\text{A.29})$$

Here, E_{sr} is the collection efficiency for snow by rain, and as with accretion of rain onto graupel, this equation uses the differential fall velocities of rain and snow, integrating over both size distributions. Wet accretion is treated differently than dry—under wet accretion ($T > 0$ °C) the term is set to the minimum of $(PRACS * \Delta t)$ and (q_s) , preemptively ensuring that the depletion does not exceed the quantity of snow.

Name	Value	Definition
E_{sr}	0.5	Collection efficiency for snow by rain

A.3.16 (PSACR) Accretion of rain onto snow

This equation is taken from RH84 (their eqn. A8)

$$PSACR = \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{2}} \left[\rho \pi^2 N_{0r} N_{0s} E_{rs} \rho_r |V_r - V_s| \left(\frac{5.0}{\lambda_r^6 \lambda_s} + \frac{2.0}{\lambda_r^5 \lambda_s^2} + \frac{0.5}{\lambda_r^4 \lambda_s^3} \right) \right], \quad (\text{A.30})$$

and is simply the reverse of the above relationship—in this case rain is accreted onto snow producing either snow at low rain water contents or graupel when the rain content is large. In this case, dry and wet accretion are treated the same, though the collection efficiency is set to the same efficiency as in PRACS above.

Name	Value	Definition
E_{rs}	0.5	Collection efficiency for rain by snow

A.3.17 (PGFR) Freezing of rain to form graupel

$$PGFR = 20\rho\pi^2 N_{0r}\rho_r B \left[\frac{\exp\{A(T_0 - T)\} - 1}{\lambda_r^7} \right] \quad (\text{A.31})$$

This equation follows LFO (their eqn. 45), and only operates in the range $273.16 > T > 248.16$. The constants A and B are taken from Bigg (1953), and the formulation originated with Wisner et al. (1972).

Name	Value	Definition
A	0.66	Exponential constant in the rain freezing equation
B	5000.0	Multiplicative constant in the rain freezing equation

A.3.18 (PSMLT) Melting of snow

$$PSMLT = \Delta t \left\{ \frac{2\pi N_{0s}}{L_f} K_a (T - T_0) \left[\frac{0.65\rho}{\lambda_s^2} + 0.44 \left(\frac{\rho_o}{\rho} \right)^{\frac{1}{4}} \left(\frac{\rho a_s}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{b_s}{2} + \frac{5}{2}\right)}{\lambda_s^{\left(\frac{b_s}{2} + \frac{5}{2}\right)}} \right] \right\} \quad (\text{A.32})$$

This equation follows RH83 (their eqn. A25), which follows from snow melt per unit time taken from RH83 (their eqn. A23).

$$\frac{dM_{melt}}{dt} = -\frac{2\pi}{L_f} K_a D_s (T - T_0) F', \quad (\text{A.33})$$

where F' is the ventilation coefficient taken from RH83 (their eqn. A24)

$$F' = 0.65 + 0.44 S_c^{\frac{1}{3}} R_e^{\frac{1}{2}} \quad (\text{A.34})$$

where R_e is

$$R_e = V_s(D_s) \frac{D_s \rho}{\mu} \quad (\text{A.35})$$

The constants in the above equation are defined below. This process is only allowed when the air temperature exceeds freezing.

A.3.19 (PGMLT) Melting of graupel to form rain

$$PGMLT = \Delta t \left\{ \frac{2\pi N_{0G}}{L_f} K_a (T - T_0) \left[\frac{0.78\rho}{\lambda_g^2} + 0.31 \frac{\rho_o}{\rho} \left(\frac{\rho a'}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{b'}{2} + \frac{5}{2}\right)}{\lambda_g^{\left(\frac{b'}{2} + \frac{5}{2}\right)}} \right] \right\} \\ + \Delta t \left\{ \frac{C_w}{L_f} (T - T_0) (QGACW + QGACR) \right\} \quad (\text{A.36})$$

This formulation follows RH84 (their eqn. A18) with the addition of the second term to account for the thermodynamic effects of accretion of cloud water and rain. In this case, the ventilation coefficient is

$$F' = 0.78 + 0.31 S_c^{\frac{1}{3}} R_e^{\frac{1}{2}}, \quad (\text{A.37})$$

and the expression for R_e is identical to that of snow above.

A.3.20 (PIHOM) Homogeneous freezing of cloud water into ice

There is no functional relationship here— q_c is simply all converted to ice if the temperature is lower than -40 °C.

A.3.21 (PIDW) Deposition growth of cloud water to ice

The deposition growth of cloud water onto ice is parameterized as

$$PIDW = \Delta t \left[\rho N_c \alpha_3 \left(\frac{\rho q_i}{N_c \exp\left(\left|\frac{T-T_0}{2}\right|\right)} \right)^{\alpha_4} \exp\left(\left|\frac{T-T_0}{2}\right|\right) \right] \quad (\text{A.38})$$

where α_3 and α_4 are temperature-dependent constants with values valid between $-1.0 > T > -31.0$ °C. Pertinent constants are listed below.

Name	Value	Definition
N_c	$1.0 \times 10^{-8} \text{ cm}^{-3}$	Constant cloud number concentration

A.3.22 (PIMLT) Melting of ice to form cloud water

As with homogeneous freezing of cloud water, this is a simple if-test—if the temperature exceeds freezing, all cloud ice melts within one timestep.

A.3.23 (PINT) Initiation of ice from nuclei

This expression is computed as the minimum of two different mechanisms. The first is

$$PINT = q_v - q_{si}, \quad (\text{A.39})$$

where q_{si} is the water vapor saturation mixing ratio over ice, given by

$$q_{si} = \left[\frac{C_1}{p} \exp \left(C_2 - C_2 \frac{(T_0 - 7.66)}{(T - 7.66)} \right) \right], \quad (\text{A.40})$$

which is essentially a water vapor saturation removal mechanism. The alternative is

$$PINT = 1.0 \times 10^{-9} \left[\frac{r_{nci}}{\rho} - \frac{q_i}{m_{i50}} \right] \quad (\text{A.41})$$

where r_{nci} corresponds to the number of activated ice nuclei at the given air temperature, and is computed as the minimum of 1.0 or the expression

$$r_{nci} = 1 \times 10^{-6} \exp(-0.46(T - T_0)), \quad (\text{A.42})$$

which appears to be a form of Fletcher's equation. The first term in brackets in A.41 corresponds to the number of activated nuclei per unit mass of air, while the second

term is the number of crystals of 50 micron radius per unit mass of air. It should also be noted that, in this case, m_{i50} has been redefined. The altered value, used only here, is noted below

Name	Value	Definition
m_{i50}	$3.76 \times 10^{-8} \text{ g}$	Mass of a 50 micron ice crystal

A.3.24 Cloud water and ice deposition/evaporation

The Tao et al. (1989) saturation technique is implemented to compute condensation or evaporation of cloud water and deposition or evaporation of ice. The procedure operates as follows. First, set initial values of condensation and deposition factor assuming the temperature is between freezing and the threshold for homogeneous freezing of cloud water (typically -40°C).

$$CND = \frac{(T - T_{00})}{(T_0 - T_{00})} \quad (\text{A.43})$$

$$DEP = \frac{(T_0 - T)}{(T_0 - T_{00})}, \quad (\text{A.44})$$

where T_0 is the freezing temperature and T_{00} is set equal to -40°C (233 K). In addition, set two other terms equal to

$$y_1 = q_{sw} \frac{(T - T_{00})}{(T_0 - T_{00})} \quad (\text{A.45})$$

$$y_2 = q_{si} \frac{(T_0 - T)}{(T_0 - T_{00})} \quad (\text{A.46})$$

where q_{si} was defined above in equation A.40, and

$$q_{sw} = \left[\frac{C_1}{p} \exp \left(C_3 - C_3 \frac{(T_0 - 35.86)}{(T - 35.86)} \right) \right]. \quad (\text{A.47})$$

Here, C_1 and C_3 are derived from Tetens's formula for saturation vapor pressure. Next, apply conditions depending on whether the temperature is above freezing or below

the homogeneous freezing threshold (233.16 K). If the temperature is above freezing, set

$$CND = 1.0 \quad (\text{A.48})$$

$$DEP = 0.0 \quad (\text{A.49})$$

$$y_1 = q_{sw} \quad (\text{A.50})$$

$$y_2 = 0.0 \quad (\text{A.51})$$

Conversely, if the temperature is lower than the homogeneous freezing threshold, set

$$CND = 0.0 \quad (\text{A.52})$$

$$DEP = 1.0 \quad (\text{A.53})$$

$$y_1 = 0.0 \quad (\text{A.54})$$

$$y_2 = q_{si} \quad (\text{A.55})$$

Once these terms have been set, we define the saturation water vapor mixing ratio as

$$q_{vs} = y_1 + y_2 \quad (\text{A.56})$$

which corresponds to a temperature-weighted saturation vapor mixing ratio as opposed to the mass-weighted quantity mentioned in Tao et al. (1989). CND and DEP terms are then set equal to

$$CND = CND \left[\frac{q_v - q_{vs}}{1 + \pi (C_3(T_0 - 35.86)y_1^3 + C_2(T_0 - 7.66)y_2^3) \left(\frac{L_v}{\pi C_p} (CND) + \frac{L_s}{\pi C_p} (DEP) \right)} \right] \quad (\text{A.57})$$

$$DEP = DEP \left[\frac{q_v - q_{vs}}{1 + \pi (C_3(T_0 - 35.86)y_1^3 + C_2(T_0 - 7.66)y_2^3) \left(\frac{L_v}{\pi C_p} (CND) + \frac{L_s}{\pi C_p} (DEP) \right)} \right] \quad (\text{A.58})$$

As a final check, set

$$CND = \max(-q_c, CND) \quad (\text{A.59})$$

$$DEP = \max(-q_i, DEP) \quad (\text{A.60})$$

A.3.25 (PSDEP) Vapor deposition on snow

Vapor deposition on snow is computed as

$$PSDEP = 2\Delta t \left\{ \frac{4N_{0s}(S_i - 1) \left[\frac{0.65\rho}{\lambda_s^2} + 0.44 \left(\frac{\rho_0}{\rho} \right)^{\frac{1}{4}} \left(\frac{\rho a''}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{b''}{2} + \frac{5}{2}\right)}{\lambda_s^{\left(\frac{b''}{2} + \frac{5}{2}\right)}} \right]}{\left[\frac{1}{T} \left(\frac{L_v L_s M_w}{K_a R^* T} - \frac{L_v}{K_a} \right) + \frac{TR^*}{(0.226)M_w(6.11 \times 10^3) \exp\left(a2 - a2 \frac{(T_0 - 7.66)}{(T - 7.66)}\right)} \right]} \right\}, \quad (\text{A.61})$$

where the ice supersaturation is computed as

$$S_i - 1 = \delta_1 \left(\frac{q_v}{q_{si}} - 1 \right) \quad (\text{A.62})$$

and the delta term is set equal to zero if the sum of the cloud and ice mass mixing ratios falls below a preset threshold value. This relationship nearly exactly follows RH83 (their eqn. A26), with the saturation vapor pressure over ice (the exponential term) given by a modified version of Teten's formula (Tao et al. 1989).

A.3.26 (PGDEP) Vapor deposition on graupel

Vapor deposition on graupel is very similar to that for snow above and follows RH84 (their eqn. A17).

$$PGDEP = 2\Delta t \left\{ \frac{2\pi N_{0_G}(S_i - 1) \left[\frac{0.78\rho}{\lambda_g^2} + 0.31 \left(\frac{\rho_0}{\rho} \right)^{\frac{1}{4}} \left(\frac{\rho a_g}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{b_g}{2} + \frac{5}{2}\right)}{\lambda_g^{\frac{b_g}{2} + \frac{5}{2}}} \right]}{\left[\frac{1}{T} \left(\frac{L_s^2 M_w}{K_a R^* T} - \frac{L_s}{K_a} \right) + \frac{TR^*}{(0.226)M_w (6.11 \times 10^3) \exp\left(a2 - a2 \frac{(T_0 - 7.66)}{(T - 7.66)}\right)} \right]} \right\}. \quad (\text{A.63})$$

where the ice supersaturation is defined in terms of a delta function identical to the one mentioned above.

A.3.27 (ERN) Evaporation of rain

The evaporation of rain is computed as the minimum of two different equations, the first of which is

$$ERN = - \frac{(q_v - q_{sw})}{q_{sw} \frac{L_v}{C_p} \left[C_3 \frac{(T_0 - 35.86)}{(T - 35.86)^2} \right]}, \quad (\text{A.64})$$

where the saturation water vapor mixing ratio over liquid has been defined above.

The alternative expression is

$$ERN = 2\Delta t \left\{ \frac{2\pi N_{0_R}(S_w - 1) \left(\frac{0.78\rho}{\lambda_r^2} + 0.31 \left(\frac{\rho_0}{\rho} \right)^{\frac{1}{4}} \left(\frac{\rho 3 \times 10^3}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma(3)}{\lambda_r^3} \right)}{\left[\frac{1}{T} \left(\frac{L_v^2 M_w}{K_a R^* T} - \frac{L_v}{K_a} \right) + \frac{TR^*}{(0.226)M_w (6.11 \times 10^3) \exp\left(C_3 - C_3 \frac{(T_0 - 35.86)}{(T - 35.86)}\right)} \right]} \right\}, \quad (\text{A.65})$$

where the supersaturation with respect to water is computed in a similar fashion to that of ice above. The end result is the minimum of either of these expressions or the total amount of rain water. The first formulation is a representation of water vapor supersaturation with respect to liquid, while the second is presented in RH83 (their

eqn. A12). Neither equation is computed if the rain water content is equal to zero.

A.3.28 (PMLTG) Evaporation of melting graupel

$$PMLTG = 2\Delta t \left\{ \frac{2\pi N_{0g}(S_w - 1) \left[\frac{0.78\rho}{\lambda_g^2} + 0.31 \left(\frac{\rho}{\rho_0} \right)^{\frac{1}{4}} \left(\frac{\rho a_g}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{b_g}{2} + \frac{5}{2}\right)}{\lambda_g\left(\frac{b_g}{2} + \frac{5}{2}\right)} \right]}{\left[\frac{1}{T} \left(\frac{L_v^2 M_w}{K_a R^* T} - \frac{L_v}{K_a} \right) + \frac{TR^*}{(0.226)(6.11 \times 10^3) M_w \exp\left(a_1 - a_1 \frac{(T_0 - 35.86)}{(T - 35.86)}\right)} \right]} \right\} \quad (\text{A.66})$$

This computation is nearly identical to the vapor deposition equation for graupel, with changes to latent heat of fusion and supersaturation with respect to water and use of Teten's formula for water saturation vapor pressure instead of ice. It closely follows RH84 (their eqn. A19).

A.3.29 (PMLTS) Evaporation of melting snow

$$PMLTS = 2\Delta t \left\{ \frac{4N_{0s}(S_w - 1) \left[\frac{0.65\rho}{\lambda_s^2} + 0.44 \left(\frac{\rho}{\rho_0} \right)^{\frac{1}{4}} \left(\frac{\rho a_s}{\mu} \right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{b_s}{2} + \frac{5}{2}\right)}{\lambda_s\left(\frac{b_s}{2} + \frac{5}{2}\right)} \right]}{\left[\frac{1}{T} \left(\frac{L_v^2 M_w}{K_a R^* T} - \frac{L_v}{K_a} \right) + \frac{TR^*}{(0.226) M_w (6.11 \times 10^3 \exp\left(C_3 - C_3 \frac{(T_0 - 35.86)}{(T - 35.86)}\right))} \right]} \right\} \quad (\text{A.67})$$

This is very similar to the vapor deposition equation for snow, with changes to latent heat of fusion and supersaturation with respect to water and use of Teten's formula for water saturation vapor pressure instead of ice. It closely follows RH84 (their eqn. A27).

A.4 Discussion

Bulk cloud microphysical schemes, such as that used in the GCE, are located at approximately the midpoint on the scale of cloud microphysical complexity. The simplest schemes are those that diagnose cloud from the relative humidity, and prescribe

the resulting cloud properties. At the current limits of complexity are the (spectral) bin microphysics models, which predict mass, number, or both of cloud liquid and ice particles for a number of condensate species over a range of discrete particle sizes. This formulation, in theory, allows the particle size distribution to assume any number of shapes, and allows explicit simulation of cloud interaction with nucleating aerosol. In contrast, bulk schemes assume a particular form for the particle size distribution, and predict the mass (and in some cases also the number concentration) of each condensate species.

Because they represent a compromise between diagnostic cloud models and computationally expensive bin schemes, bulk schemes have attained wide use in models on scales from large eddy ($\Delta x \approx 10$'s of meters) to GCM ($\Delta x \approx 100$'s of kilometers). The advantage of bulk microphysical schemes is that, while they are relatively computationally inexpensive, they still allow for the representation of a size spectrum of particles. This is beneficial, as it leads to the possibility of realistically representing interactions between populations of cloud particles, as well as the interaction between clouds and radiation. The primary problem with bulk schemes is that they are restricted to a particular form of the particle size distribution.

It should be noted, that although bin microphysical schemes are considered to represent the state of the art in cloud microphysical modeling (Khain et al. 2000), they are plagued by their own unique set of problems. For the liquid processes, the physical mechanism by which the cloud droplet size spectrum broadens is still poorly understood, and hence is difficult to properly represent in bin schemes. The same can be said of rain drop break-up mechanisms and for liquid-liquid collisions in general, in which the collection efficiencies are not well known. Representation of collection efficiencies in ice processes is an even more complex problem, as the representation of ice crystal habit, degree of riming, and wet vs. dry particle surface must be considered. Finally, the computational expense of integrating a bulk scheme

is prohibitive, and becomes even more so with the prediction of both the mass and number of cloud particles within each discrete bin.

As was documented above, the GCE bulk microphysical parameterization assumes non-precipitating cloud liquid and pristine ice particles to be monodisperse; all particles have identical size. Precipitating species are all assumed to follow an inverse exponential size distribution. This is to some extent advantageous, as this distribution is fully characterized by only one of two parameters—either the intercept or the slope. Once either of these is fixed, the other can be predicted in the course of the model’s forward evolution—hence, this scheme is designated a *single-moment* scheme. Other bulk schemes (e.g., Verlinde et al. 1990, Walko et al. 1995, Meyers et al. 1997) assume a gamma particle size distribution such that the number of particles of any given diameter is

$$n(D) = \frac{1}{\Gamma(\nu)} \left(\frac{D}{D_n} \right)^{\nu-1} \frac{N_t}{D_n} \exp \left\{ -\frac{D}{D_n} \right\}, \quad (\text{A.68})$$

where ν is the parameter that defines the width of the distribution, N_t is the total number concentration of particles in the volume, Γ is the gamma function, and D_n is a characteristic particle diameter. D_n is related to the mean particle diameter via the relationship

$$D_{mean} = \frac{\Gamma(\nu + 1)}{\Gamma(\nu)} D_n. \quad (\text{A.69})$$

Note that the gamma distribution width parameter ν contributes an additional degree of freedom compared with the exponential distribution, and hence offers a potentially fit to observations. This conjecture has been borne out in analysis of in-situ cloud microphysical data—in particular, Heymsfield et al. (2002) found that, though the exponential distribution provides a suitable fit to observed ice particle size distributions, the slope intercept can vary through the depth of the cloud by up to two

orders of magnitude. By contrast, the slope parameter λ varies relatively little with height. The implication is that specification of constant N_0 coupled with varying λ may not produce realistic results for tropical deep convection. The gamma particle size distribution does not suffer as greatly from these problems as the width parameter provides more flexibility in specifying the shape of the distribution. In addition, since bulk microphysical schemes based on the gamma distribution typically fix the characteristic diameter D_n , which is technically equivalent to fixing the exponential slope, they offer the potential to conform more closely to observations.

With this in mind, an effort is underway to implement a gamma-based cloud microphysical scheme in the GCE model. Once this is done, the sensitivity tests, evaluation vs. observations, and assimilation can be redone to assess the most appropriate values for fixed constants (e.g., the ν and D_n parameters). The resulting particle size distributions can then be evaluated against the observations of Heymsfield et al. (2002).

Appendix B

Effect of Grid Resolution on 2D CRM Results

B.1 Introduction

Since the very first idealized numerical simulations of convection were performed (e.g., Klemp and Wilhelmson 1978) it has generally been accepted that 1 km horizontal grid spacing is sufficient to resolve cloud scale dynamics within deep convection. This argument follows both from convenience and convention, as well as from a first order consideration of the scale of deep convective updrafts. If approximately 10 grid points are necessary to sufficiently numerically represent a simulated structure, then deep convection, which tends to span 10s of kilometers, should be sufficiently resolved by grid spacings on the order of 1 km. Recently, the notion of a necessary horizontal grid spacing for convective simulation has again been debated in the literature. Weisman et al. (1997, hereafter WSK), studied the sensitivity of convective vertical heat and momentum fluxes and mean profiles of rainwater and perturbation water vapor mixing ratio to changes in horizontal grid spacing, and employed grid spacings of 12, 8, 4, 2, and 1 km. They concluded that horizontal grid spacings as coarse as 4 km are

sufficient to resolve overturning circulations in deep convective squall lines, and to correctly characterize the time and domain mean heat transports. However, WSK did not examine simulations with grid spacings less than 1 km, and note that aspects of their simulations had not yet converged at Δx of 1 km. They contend that, while the details of convective circulations might continue to change with progression to finer grid spacings, the system-wide statistics had satisfactorily converged.

In a subsequent study, Bryan et al. (2003, hereafter BWF) used the vertical shear profiles of WSK to examine whether convective structure had indeed converged to a stable solution at 1 km grid spacing. After examining the fine-scale structure of convective updrafts simulated with grid spacings between 1 km and 125 meters, and taking into account the assumptions made in the models subgrid turbulence closure, BWF concluded that a *maximum* grid spacing of 250 meters was necessary to properly simulate the turbulent overturning characteristic of propagating squall lines. When examining vertical profiles of momentum flux, they also concluded that convergence to a stable solution had *not* been achieved at grid spacings of 1 km, and that it could be debated whether convergence had even been achieved at 125 meters. The conclusion offered in BWF is that grid spacings of 125 meters or less are needed for theoretical studies of convective structure and dynamics, but that grid spacings in excess of 1 km may be sufficient (if not entirely appropriate) for convective-scale numerical weather prediction.

Though BWF demonstrate significant changes in the turbulent circulations with changes in grid spacing, the general storm morphology and organization does not change; the mean horizontal and vertical extent of simulated convection, the convective lifetime, and the vertical profiles of heat flux are approximately constant across grid spacings. In fact, the results set forth in BWF seem to support the conclusion of WSK that the system-averaged properties may have approximately converged at grid spacings between 1-2 km. The only field to exhibit large variation with changes to Δx

in both studies is the vertical momentum flux, which did not converge for any of the grid spacings used. It should also be noted that BWF present results only up to 180 minutes of simulated time, and that WSK demonstrate that coarse-resolution experiments exhibit slower development compared with fine-scale experiments. Whereas WSK demonstrate convergence by six hours of simulation time, BWF neither show the six-hour results, nor the system-wide averages of condensate or moisture, and it could be argued that the results presented in BWF might have been somewhat different had the simulations been run for a longer period of time.

Because this study employs a cloud resolving model to study the response of clouds and precipitation to changes in the sea-surface temperature, it is useful to assess the sensitivity of key *climate* variables to changes in CRM grid spacing. The fundamental question is whether grid spacings of 2 km, which are used out of computational necessity, are sufficient to capture the order-one response of the tropical hydrologic cycle to surface warming. Even though it is not feasible to perform MCMC experiments with grid spacing less than 4 km in the case of MMF-framework simulations, it is still desirable to assess whether the statistics of the cloudy state change with changes to grid spacing. Since a fundamental change in convective turbulent dynamics appears to occur between grid spacings of 1 km and 500 meters (Bryan et al. 2003, S. Lang, personal communication, the author’s unpublished work—not shown), and since 2 km grid spacing is used in large-domain RCE experiments, in this appendix experiments utilizing grid spacings of 4 km, 2 km, and 500 meters will be compared. With the exception of changes to the horizontal grid spacing, the GCE setup is identical to that used in the MCMC experiments presented in chapters 3 and 4, and hence a detailed explanation of the model configuration will not be repeated here. We simply note that the model timestep is reduced consistent with reduction in Δx to 10 seconds for $\Delta x = 2$ km and to 3 seconds for $\Delta x = 500$ meters. In each experiment, state variables that correspond to observations used in the MCMC experiments are extracted

from each simulation, and the results are compared. Similar to the methodology used to demonstrate simulation sensitivity to changes in cloud microphysical parameters in chapter 1, traditional sensitivity experiments are repeated here, and the results are examined to assess whether perturbations in microphysical parameters produce a similar response in simulations with differing grid spacings. The remainder of this chapter is organized as follows: results of each perturbation experiment are presented in section B.2, and discussed in section B.3. Conclusions and implications for the current climate system study are offered in section B.4.

B.2 Results

Snapshots of simulated condensate mass mixing ratio and vertical motion fields at one-hour intervals in simulations run with 4 km, 2 km, and 500 m grid spacing are depicted in figure B.1.²³ It can be seen that apart from the initial development, the general characteristics of deep convective development are similar between the three simulations; each produces an organized propagating system by approximately five hours into the simulation that extends through the depth of the troposphere. Differences are most evident in the strength of the vertical motion field, mode and rapidity of convective development, and in the simulation of shallow convection. While the 2 km and 500 m simulations show distinct multicellular organization at 360 and 420 minutes, multicellular structure is only crudely simulated at $\Delta x = 4$ km, if at all. In addition, while 2 km and 500 m simulations rapidly generate deep convective cells at 240 minutes and 480 minutes, development in the 4 km simulation noticeably lags by approximately one hour. The most striking differences are evident in the simulation of shallow and mid-level convection, particularly at 420 minutes of simulated time. While an approximation of towering cumulus can be seen in the 2 km simulation, and

²³The time period between 240 and 540 minutes of simulated time was chosen as each simulation exhibited a robust propagating squall line during this interval.




Figure B.1: Comparison of the time evolution of 2D MMF-configuration GCE simulations with (a) 4 km, (b) 2 km, and (c) 500 m grid spacing. Filled contours represent total cloud condensate mass mixing ratio, while the open red contours depict vertical motion every $1 \text{ m} \cdot \text{s}^{-1}$ starting at $1 \text{ m} \cdot \text{s}^{-1}$. Simulation time is plotted in the upper right hand corner of each plot.

the 500 m simulation exhibits a full spectrum of clouds, the 4 km simulation (at 480 minutes) can only produce a new deep convective cell. Consistent with simulation of convection on smaller horizontal scales, updrafts are much more intense in the 500 m simulation, and the bulge in condensate mass mixing ratio at the storm top, indicative of penetration of updrafts into the lower stratosphere, is much more well-defined in simulations run with 500 m grid spacing.

A comparison of a selection of mean state variables output from control and perturbed simulations with 4 km, 2 km, and 500 m grid spacing is presented in Table B.1. Though the convective structure and organization vary considerably with changes to grid spacing, it is interesting that the time and domain-mean properties of clouds, precipitation, and radiation do not. The largest change appears to be in the high cloud fraction—the 500 m simulations exhibit less high cloud compared with coarse-resolution simulations. This is consistent with the existence of greater amounts cloud at mid and low levels. The shortwave cloud radiative forcing also varies quite a bit between simulations, which is not surprising given the changes to high cloud fraction.

In addition to determining whether cloud and radiation statistics change with changes to Δx , it is also interesting to assess whether simulations with different horizontal grid spacing exhibit similar *sensitivity* to changes in cloud microphysical parameters. If the sensitivities are consistent, we may conclude that results of the MCMC procedure retain validity across different scales. Further examination of Table B.1 shows that simulations with different grid spacing exhibit similar sensitivities to changes in cloud microphysical parameters for a number of different variables. When ice particles are shifted toward smaller mean diameter and lower density, the precipitation rate and amount increase, as does the ice water path, and the longwave cloud radiative forcing decreases. Although the liquid water path, high cloud fraction, and shortwave cloud radiative forcing do not change consistently between 4 km and

Table B.1: Comparison of output from MMF-configuration simulations using control vs. perturbed parameter values (Table 1.1). Results of simulations using 4 km, 2 km, and 500 m horizontal grid spacing are shown.

Observation	4 km CTRL	4 km PERT	2 km CTRL	2 km PERT	500 m CTRL	500 m PERT
Precip Rate ($mm \cdot h^{-1}$)	0.82	0.85	0.71	0.76	0.75	0.81
Total Precip (mm)	19.15	19.69	16.84	17.95	17.87	19.38
Cloud Top Temp (K)	232.87	233.65	238.32	239.56	236.19	235.26
LWP ($kg \cdot m^{-2}$)	0.296	0.313	0.305	0.297	0.295	0.276
IWP ($kg \cdot m^{-2}$)	0.249	0.306	0.212	0.284	0.227	0.339
High Cloud Fraction	0.60	0.65	0.57	0.53	0.45	0.52
TOA LW CRF ($W \cdot m^{-2}$)	-49.33	-45.68	-45.96	-38.99	-42.99	-36.42
TOA SW CRF ($W \cdot m^{-2}$)	94.50	101.92	111.85	95.39	102.09	77.50

2 km grid spacings, the changes in these fields are small relative to the range of variability in the TRMM observations (see chapter 4 for more details).

B.3 Discussion

Though the results of the above comparison at first appear encouraging, a few cautionary statements must be made. First, the simulations performed in this brief study develop in an environment with relatively large values of shear. Similarities between simulations with coarse and fine grid spacing are due in part to the fact that the environmental shear will force the convection to organize in a similar fashion regardless of grid spacing. A similar result was demonstrated in Bryan et al. (2003), in which simulations with weak shear exhibited larger differences than simulations that developed in the presence of larger values of shear. Since the MCMC experiments were forced with observations from SCSMEX during a period of the monsoon that was characterized by fairly strong shear, differences in simulated convective structure for different grid spacings were likely suppressed. If MCMC experiments are repeated for convection that develops in an environment with weaker shear, the above perturbation

experiments would likely need to be repeated.

It should also be noted that the differences in updraft intensity between coarse and fine-grid simulations have an implication for the use of coarse-grid simulations to study stratospheric/tropospheric exchange. While it is not possible to draw conclusions based on the limited results presented here, there is certainly an indication that the coarse-grid simulations, with weaker updrafts and barely-discernible overshooting top, cannot properly represent the penetration of convective updrafts into the lower stratosphere. In addition, though storm morphology was generally consistent between simulations with coarse and fine grid spacing, and while sensitivity to changes in cloud microphysical parameters was consistent, the conclusions offered here cannot be extended to simulations run with a different cloud microphysical parameterization. This is because there is a highly nonlinear feedback between changes to cloud microphysical properties and cloud dynamics through the precipitation process (water and ice loading and cold pool dynamics) as well as through the details of condensation (latent heat release and feedback to convective vertical motion).

B.4 Conclusions

In this chapter, tests were performed to assess the sensitivity of GCE simulations to changes in horizontal grid spacing. When simulations of propagating oceanic squall lines were produced using relatively coarse (2 and 4 km) vs. relatively fine (500 m) grid spacings, it was found that the differences in mean precipitation rate and amount, ice and liquid water paths, and cloud radiative forcings were small. These results support the conclusions offered by WSK and BWF that, although the details of the convective circulations change with changes to horizontal grid spacing, the general storm morphology does not. Because the statistics of properties important in the climate system were stable to changes in horizontal grid spacing, we may conclude

that RCE simulations that employ 2 km horizontal grid spacing, while not ideal, are not inappropriate for the study of the interaction between radiation and convection.

Appendix C

MCMC: Convergence Diagnostics, Sample Independence, and Adaptive Sampling

C.1 Theoretical Foundations of MCMC and Construction of an MCMC Algorithm

Recall that the fundamental goal of data assimilation is to obtain information about a set of unknown variables (or numerical model parameters) given prior knowledge of the system as represented in a set of observations. Since the observations contain noise (errors), and since observations and numerical model parameters are most often indirectly linked, the set of unknown parameters is treated as a set of random variables, as are the set of unknown observation errors. Because of the complexity in the relationship, it is more appropriate to estimate the probability density function of the set of unknown parameters rather than performing an optimization procedure, and (as has been shown previously in chapters 2 and 3), there are many advantages to obtaining the full PDF.

Fundamentally, the problem reduces to estimation of the probability of the set of parameters $\mathbf{x}$ conditioned on a set of observations $\mathbf{y}$, the solution to which was formally stated by the Reverend Thomas Bayes as (see also equation 1.3)

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x}) P(\mathbf{x})}{P(\mathbf{y})} \propto P(\mathbf{y}|\mathbf{x}) P(\mathbf{x}) \quad (\text{C.1})$$

$P(\mathbf{x}|\mathbf{y})$ can be computed by solving for each quantity on the right hand side of C.1. It is first assumed that a deterministic (numerical) model $F(\mathbf{x})$ exists that relates the parameters $\mathbf{x}$ to the set of observations $\mathbf{y}$.²⁴ Assuming a set of m conditionally independent measurements of $\mathbf{y}$, $P(\mathbf{y}|\mathbf{x})$ is computed as the product

$$P(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^m P(y_i|\mathbf{x}), \quad (\text{C.2})$$

where $P(y_i|\mathbf{x})$ is the discrete probability of obtaining the observation y_i given the set of parameters $\mathbf{x}$. The normalizing (proportionality) constant is defined as

$$P(\mathbf{y}) = \int P(\mathbf{y}|\mathbf{x}) P(\mathbf{dx}), \quad (\text{C.3})$$

which effectively describes how a perturbation $\mathbf{dx}$ to the set of parameters relates to a change in the conditional probability $P(\mathbf{y}|\mathbf{x})$. $P(\mathbf{x})$ can be obtained by assuming no prior information on the distribution of parameters (non-informative prior distribution). This corresponds to assigning identical a priori probability to all values of the state vector across the entire (discrete) state space. Alternatively, each parameter can be assumed to have an a priori probability (e.g., $P(\mathbf{x}) \sim \text{Uniform}[\mathbf{x}_a, \mathbf{x}_b]$), be located only within a certain range, or be restricted to take a certain sign.

Given an assumption of the form of $P(\mathbf{x})$, and a way to evaluate (C.2) and (C.3), all of the components are in place to obtain a solution to $P(\mathbf{x}|\mathbf{y})$. However, com-

²⁴As the set of model parameters is typically denoted as θ in the statistical literature, using the notation $\mathbf{x}$ here is a slight abuse of notation. However, $\mathbf{x}$ is retained for greater consistency with standard atmospheric and oceanic data assimilation notation (Talagrand 1997).

putation of the exact solution is often intractable as it involves numerical evaluation of the integral in (C.3)—this is equivalent to performing an exhaustive perturbation analysis over the entire state space. If this were possible, then the optimal set of parameters could be obtained as the maximum likelihood estimate

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} P(\mathbf{y}|\mathbf{x}), \quad (\text{C.4})$$

however, for the case of estimation of a set of unknown cloud microphysical parameters given temporally and spatially integrated observations, evaluation of $P(\mathbf{y})$ is not feasible (see Chapter 3 for further discussion). Even so, the PDF can still be characterized by sampling the un-scaled posterior PDF $P(\mathbf{x}|\mathbf{y}) \propto P(\mathbf{y}|\mathbf{x}) P(\mathbf{x})$. MCMC methods are used to produce this sample by seeking out regions of the state space with relatively large conditional probability, as determined by a measure of goodness-of-fit between the observations $\mathbf{y}$ and numerical model state $F(\mathbf{x})$.

The fundamental components necessary to construct the Metropolis-Hastings-type MCMC procedure (Metropolis et al. 1953, Hastings 1970) used in chapters 3 and 4 are the following.

1. A deterministic (numerical) model $F(\mathbf{x})$ that relates observations to parameters. In this case, the numerical model is a full nonlinear Goddard Cumulus Ensemble model integration.
2. A probabilistic error model (likelihood function) for the observations

$$\mathbf{y} = \mathbf{F}(\mathbf{x}) + \epsilon, \quad (\text{C.5})$$

assumed to be Gaussian in form with mean $F(\mathbf{x})$ and variance σ_y^2 . In the current

(multivariate) case, this is

$$L(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi)^{\frac{m}{2}} |\mathbf{R}|^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2} [\mathbf{y} - F(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{y} - F(\mathbf{x})] \right\} \quad (\text{C.6})$$

(see chapters 2 and 3 for further explanation). It should be noted that the observation likelihood function is not constrained to be Gaussian—if observations are known to contain outliers, or if the error model is known to be non-Gaussian, an alternative likelihood function can easily be defined and inserted into the MCMC algorithm.

3. A proposal distribution

$$\tilde{\mathbf{x}} \sim q(\tilde{\mathbf{x}}|\mathbf{x}^{(i)}), \quad (\text{C.7})$$

defined as a method of randomly generating a new test realization $\tilde{\mathbf{x}}$ of the parameter set given current parameter values $\mathbf{x}^{(i)}$. For simplicity, an unbiased (symmetric) proposal distribution is used, e.g., bounded Uniform or Gaussian centered on the current estimate. Non-symmetric proposal distributions are discussed further in section C.5, and in Gelman et al. (2004).

Each step in the MCMC algorithm (each “link” in the Markov chain) is evaluated according to the procedure outlined in Table C.1, which is repeated multiple times, and each accepted set of $\mathbf{x}^{(i)}$ is stored as a sample of the PDF. Note that the most expensive part of each iteration is step 2, which involves a full integration of either a radiative transfer or numerical process model.

Table C.1: Description of the steps involved in computing each iteration in the MCMC algorithm.

1. Propose a new set of $\mathbf{x}$ randomly distributed according to $q(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})$. For both the MODIS retrieval and GCE cloud microphysical parameter estimation, this corresponds to drawing a perturbation to the current parameter from a bounded uniform distribution centered on the current parameter value.

2. Generate new observations according to $\mathbf{y} = F(\tilde{\mathbf{x}})$

3. Evaluate $L(\mathbf{y}|\tilde{\mathbf{x}})$

4. Compute the *acceptance ratio*

$$\rho(\tilde{\mathbf{x}}; \mathbf{x}^{(i)}) = \frac{L(\mathbf{y}|\tilde{\mathbf{x}}) P(\tilde{\mathbf{x}})}{L(\mathbf{y}|\mathbf{x}^{(i)}) P(\mathbf{x}^{(i)})}. \quad (\text{C.8})$$

In the case of a uniform proposal distribution, $P(\tilde{\mathbf{x}}) = P(\mathbf{x}^{(i)})$, and the acceptance ratio simply becomes the ratio of the current likelihood to the previous likelihood.

5. Compute the accept/reject criterion: if $\rho(\tilde{\mathbf{x}}; \mathbf{x}^{(i)}) \geq 1$ then accept $\tilde{\mathbf{x}}$ as the new sample $\mathbf{x}^{(i+1)}$. Otherwise, generate $u \sim \text{Uniform}[0, 1]$, and *probabilistically* accept the new set of $\tilde{\mathbf{x}}$ if $\rho(\tilde{\mathbf{x}}; \mathbf{x}^{(i)}) \geq u$, otherwise set $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)}$.

C.2 MCMC Diagnostics

C.2.1 Effective Sample Size

What MCMC is designed to provide is a set of *independent* samples from the posterior PDF, however, because samples are obtained using a Markov chain, there will always be some degree of autocorrelation within the chain. The degree of autocorrelation depends on two factors—the width of the proposal distribution, which determines the effective distance proposed in the state space, and the acceptance rate, defined as the ratio of accepted samples to number of proposal steps. If the accept rate is high and the step size is small (the two tend to be related, though this is not always the case), such that $\tilde{\mathbf{x}}$ is always close to $\mathbf{x}^{(i)}$, then the degree of autocorrelation will be high. Conversely, regardless of the step size, if the acceptance rate is *low*, then the chain effectively remains in the same state for many successive iterations, and again the autocorrelation will be high. In either case, the robustness of the PDF returned by MCMC depends on the number of independent samples. If the autocorrelation is high, this translates to a smaller number of independent samples, and hence a less efficient algorithm overall.

Because it is desirable to strike a balance between low autocorrelation and high accept rate, the step size is typically tuned via changes to the width of the proposal distribution until the accept rate reaches 25% - 50% (Gelman et al. 2004, p. 306), depending on whether the state is perturbed in a single direction or many directions at once in each MCMC iteration. Tuning of the proposal width occurs during the burn-in phase of the chain; the set of MCMC iterations that occur between the beginning of the chain and the point at which the chain has converged to sampling from a stationary posterior distribution. Because these iterations are discarded before the sample is analyzed, it is of critical importance to be able to accurately assess whether the algorithm has converged. Tools with which to assess convergence are documented

in the next section.

C.2.2 Assessment of Convergence to Sampling a Stationary Distribution

The theory that underlies MCMC is robust, and a properly constructed chain is virtually guaranteed to converge to sampling from the correct posterior distribution. However, each chain must be started from a set of initial parameter values, and it is difficult to know in advance how many iterations will be needed before the chain converges to sampling from a stationary distribution. Hence, each MCMC chain contains a *burn-in* period at the beginning, during which the algorithm converges to sampling from the correct stationary posterior distribution. Once convergence has been attained, all samples prior to burn-in must be discarded, as they will not be representative of the proper posterior distribution. Although diagnosing convergence is a topic of current research in the statistics community (Gelman et al. 2004), a few common metrics exist—these are briefly outlined below.

1. Convergence can be approximately visually assessed by plotting a “time series” of each parameter in the chain.
2. Moments of the posterior distribution can be computed from the MCMC sample and plotted for a random subset of the chain. If the chain has sufficiently converged, subsets of the full sample should have nearly identical moments.
3. Similar to (2), a running computation of the moments of the posterior sample can be computed, and a timeseries plotted. Once the algorithm has approximately converged, these moments should not change.
4. If multiple chains are used, comparison of the between- and within-chain statistics can be used to assess convergence. This method is discussed in greater




Figure C.1: Depiction of the timeseries of each parameter in the first chain of the GCE simulated observations experiment documented in chapter 3.

detail in the next section.

For a more in-depth discussion of current methods of assessing convergence, the reader is referred to Robert and Casella (1999). Timeseries of each parameter for the first chain in the simulated-observations MCMC experiment (see chapter 3 for explanation) are depicted in figure C.1. Though it is generally acknowledged to be a weak method of diagnosing convergence, valuable information about the parameter PDFs can still be extracted from visual inspection of parameter timeseries. First, the degree of autocorrelation within the chain can be approximately assessed by examining the number of iterations for which each parameter remains constant. In addition, the character of the state space is to some extent revealed in the rapidity of fluctuation of each parameter value—slow transitions between different regions of

the state space can indicate a complex PDF with many local minima, or alternatively may be the result of overconstraint of the set of parameters by the observations. Finally, the parameter modes and relative parameter importance can be approximated by examining how long and in which regions of the state space the chain remains. For example, the secondary modes in the N_{0g} and b_g parameters are evident in the parameter timeseries, as is the robust influence of a_g on the simulation results. Though they contain valuable information, figure C.1 demonstrates why parameter timeseries are not often used to assess convergence. In the case of broad PDFs, MCMC tends to sample efficiently, and can be difficult to determine whether the algorithm has reached stationarity.

A running calculation of the moments of the PDF for N_{0r} is presented in figure C.2, from which it can be seen that this method is a much more effective tool for assessing convergence than plotting parameter timeseries. Examination of the running mean, variance, skewness, and kurtosis reveals that MCMC reaches convergence between 5,000 and 10,000 samples. Indeed, if we return to figure C.1 with the knowledge that convergence happens some time after 5,000 samples have been obtained, convergence can be seen in the timeseries as well.

The problem with all of the above methods is that it is extremely difficult to use them to assess convergence *while* the algorithm is running. Timeseries inspection, random draws, and running computation of moments are primarily useful in determining how much of each chain to retain in the sample of the PDF. It will be argued in the next section that computation of the between- and within-chain statistics offers a much more effective way of assessing convergence.




Figure C.2: Running computation of the moments of N_{0r} from chain 1 of the GCE simulated observations experiment documented in chapter 3.

C.3 Use of Multiple Chains

The preceding discussion of the MCMC algorithm and its convergence characteristics assumed the use of a single chain. In applications for which many samples are needed, or in which burn-in is anticipated to be short (or ideally, both), the algorithm can be made more efficient by constructing multiple Markov chains. After burn-in, all chains will have converged to sampling the identical posterior distribution, and samples from all chains can be combined to give a larger effective sample. In this case, standard convergence diagnostics can be applied to each chain, and assessment of the number of independent samples within each chain is typically done prior to mixing the samples from each chain. In theory, use of multiple chains provides a robust way of increasing the total sample size, and can lead to a more thorough sample of the state space. In addition, because each chain can be run on a different processor of a distributed computing environment, use of multiple chains provides an effective way to parallelize the MCMC algorithm. Care must be taken, however, to ensure that each chain is sampling the same posterior distribution. This is done by comparing the variance (or other moments) *within* each chain to the variance *between* chains for each estimated parameter in the following manner (Gelman et al. 2004, p.296-297). Consider m chains, each of length n samples. First, the within-chain variance is computed for each parameter x as

$$W = \frac{1}{m} \sum_{j=1}^m \left[\frac{1}{n} \sum_{i=1}^n (x_{ij} - \overline{x_{.j}})^2 \right], \quad (\text{C.9})$$

where $\overline{x_{.j}}$ is the mean of each parameter x within each chain

$$\overline{x_{.j}} = \frac{1}{n} \sum_{i=1}^n x_{ij}. \quad (\text{C.10})$$

The between-chain variances are then computed as

$$B = \frac{n}{m-1} \sum_{j=1}^m (\bar{x}_{\cdot j} - \bar{x})^2, \quad (\text{C.11})$$

where $\bar{x}$ is the mean of the given parameter across all chains

$$\bar{x} = \frac{1}{m} \sum_{j=1}^m \bar{x}_{\cdot j}. \quad (\text{C.12})$$

An unbiased estimate of the marginal posterior variance of each parameter x conditioned on the set of observations $\mathbf{y}$ can be obtained from a weighted combination of B and W as

$$\widehat{var}^+(x|\mathbf{y}) = \frac{n-1}{n}W + \frac{1}{n}B. \quad (\text{C.13})$$

This quantity tends to overestimate the true marginal posterior variance, but converges to the true variance as $n \rightarrow \infty$. Proper chain mixing is assessed by comparing the variance estimate to the within-chain variance, and computing $\hat{R}$, an estimate of the factor by which the dispersion in the current sample would be reduced if each chain were allowed an infinite length

$$\hat{R} = \sqrt{\frac{\widehat{var}^+(x|\mathbf{y})}{W}}. \quad (\text{C.14})$$

It can be readily seen that this estimate will converge to 1 in the limit as $n \rightarrow \infty$. According to the literature (Gelman et al. 2004), there is no specific value of the r-statistic for which chains can be said to have sufficiently mixed, though a value of $\hat{R}$ less than 1.1 for each parameter is generally deemed acceptable. In each of the multiple-chain MCMC experiments performed in this study, $\hat{R}$ values of under 1.05 were uniformly obtained for the sample sizes documented.

An illustration of convergence of within and between-chain variance to increasing numbers of samples for the 60-chain MCMC experiment outlined in chapter 3 is

depicted in figure (C.3). It can be seen that by approximately 3500 iterations, the chains can be assumed to have mixed sufficiently as their $\hat{R}$ value drops below 1.1, and between 5,000 and 10,000 iterations (recall figures C.1 and C.2) $\hat{R}$ drops below 1.05 and levels off.

C.4 Gaussian Proposal Distribution

Several studies have suggested that the use of a Gaussian proposal distribution in place of the Uniform distribution employed in this study can lead to a more efficient MCMC algorithm. In both the MODIS and GCE applications, the proposal distribution was constructed as bounded uniform centered on the current parameter value such that

$$q \sim \text{Uniform} \left[x^{(i)} - \delta x^{(i)}, x^{(i)} + \delta x^{(i)} \right], \quad (\text{C.15})$$

where $\delta x^{(i)}$ was computed as the product of a random variable $u \sim \text{Uniform}[-0.5, 0.5]$ multiplied by the parameter's allowable range and by a tunable scale factor (width parameter) between zero and one. Alternatively, it is possible to draw the new set of parameters from a Normal transition distribution

$$q \sim \text{N}(x^{(i)}, c^2 \Sigma), \quad (\text{C.16})$$

centered on the current parameter estimate $x^{(i)}$ and with scaled covariance $c^2 \Sigma$, where Σ is an estimate of the parameter covariance matrix periodically recomputed from the current sample, and c^2 is a scaling parameter that modifies the width of the distribution

$$c = \frac{2.4}{\sqrt{d}}. \quad (\text{C.17})$$




Figure C.3: Plots of the R-statistic for successively greater numbers of samples (see text for derivation and explanation) for each of the 9 microphysical parameters estimated in chapters 3 and 4. The color scheme is explained in the legend at top right.

In this case, d is the number of parameters perturbed in each MCMC iteration. The accept rate in this case is tunable via the numerator in c , which scales the width of the Normal proposal distribution. The c parameter can be tuned during burn-in, just as the width of the uniform proposal distribution was in both the MODIS and GCE applications. In theory, use of a Gaussian proposal results in a more efficient MCMC algorithm for the same accept rate, as it leads to a greater number of independent draws from the posterior distribution for a given sample size. In addition, use of a Gaussian proposal avoids the problems of limited sampling when using a truncated uniform. Because the Gaussian PDF contains mass into the tails out to a near-infinite length, as the width of the proposal distribution is changed the algorithm is theoretically able to access the entire parameter range in any given proposal step. In the case of a truncated uniform, the maximum step size is absolute, which can affect the speed of mixing in an adverse way, particularly for a highly complicated state space, and can lead to artifacts in the posterior parameter PDFs.

The problems that can arise from use of an absolute step size can be demonstrated by using MCMC to simulate a PDF for a parameter whose true posterior distribution is bounded uniform—e.g., a model parameter that is unrelated to the model output. Because there is assumed to be no relationship between the parameter and the state, no likelihood function is necessary—any parameter value would result in the same state, and thus the likelihood ratio would always be equal to unity. Instead, each new parameter is generated as follows. Using a uniform proposal, the next parameter in the chain is computed as

$$x^{(i+1)} = x^{(i)} + T(U - 0.5)(p_{max} - p_{min}),$$

where T is the distribution width parameter, U is a uniform random variable on $[0,1]$, p_{min} is the parameter minimum value, and p_{max} is the parameter maximum value.

For the Gaussian proposal, new parameters are computed as

$$x^{(i+1)} = x^{(i)} + c^2 \sigma^2 N(0, 1),$$

where c is the scale factor for the Gaussian proposal, σ^2 is the parameter variance (set equal to the variance of the Uniform distribution in this case), and $N(0, 1)$ is a Gaussian random variable with zero mean and unit variance. In each case, only parameter values that lie outside the predefined range are rejected, and the experiment is repeated for a range of Uniform and Gaussian scale factors.

Examination of the results of a simulation run for 10^6 samples reveals that, for any value of the Uniform scale factor less than 2.0 (Fig. C.4a), the algorithm preferentially samples the center of the range. The problem is most pronounced for width parameter values around 1.0, which result in a distinct spurious mode in the PDF located in the center of the parameter range. The cause of the anomalous mode is intuitive—when the chain samples close to one of the extrema of the parameter range, the uniform proposal is free to generate a new parameter value at equal distance on either side of the current chain location. However, *all* of the proposed values outside of the preset parameter range will be rejected causing the algorithm to preferentially move back toward the center. The frequency with which the boundary is a problem depends directly on how close to the boundary the parameter value is and on the scale factor itself. The problem is that, when adaptively tuning the step size to reach an optimal sampling rate, it is preferable for the algorithm to tend *more* toward a Uniform posterior with increasing scale factor. In the case of the absolute-step Uniform proposal, the opposite is true.

In contrast to the Uniform, the Gaussian proposal distribution (Fig. C.4b) is much more well-behaved. Though it still generates an anomalous mode in the center of the parameter range, there is no discontinuity in the simulated PDF. In addition, the

PDF tends toward uniform for increasing values of the scale parameter—this is both consistent with the preferred scale parameter behavior, as well as with the behavior of the algorithm itself. For larger scale factors, the algorithm should sample more widely in the parameter range, while for smaller scales it is preferable for the algorithm to sample more locally.

C.5 Advanced Topics: Modifications to the MCMC Algorithm

While the random walk taken by the MCMC algorithm ensures in theory a thorough search of the state space and a correspondingly unbiased sample of the posterior PDF, it is inherently inefficient, as a large number of proposed parameter samples are rejected (and discarded) in the course of the sampling procedure. In addition, while MCMC (and, in particular, Metropolis-Hastings) algorithms work well for relatively well-behaved state spaces (e.g., spaces in which the probability $P(\mathbf{x}|\mathbf{y})$ varies smoothly with changes in $\mathbf{x}$), it can experience difficulty when confronted with a state space that is highly irregular—e.g., for situations in which the state space consists of sharply defined modes that are widely dispersed. In this case, it will be nearly impossible for the chain to find and sample all of the modes of the distribution. In addition, if the chain is started in a region of the space that is far-removed from one of the modes of the PDF, it may *never* find a region of significant probability. This is especially true if the step size is tuned according to the accept frequency of the algorithm. In these special cases, MCMC must be modified to more effectively search the state space.

One solution to the problem of inefficient sampling is to use a *hybrid* or *Hamiltonian* method, in which a “momentum” parameter is coupled to each estimated state parameter. The momentum parameter borrows from concepts in Newtonian physics





Figure C.4: Illustration of the departure from Uniform for various values of the (a) Uniform and (b) Normal proposal distribution scale parameters.

by allowing the algorithm to maintain some “inertia” in its search of the state space. In each MCMC iteration, both the value of the parameter and the value of the momentum are updated together, and the proposed new parameter value is determined to a large extent by the momentum. In this way, because the momentum parameter defines an a priori expectation of the distance and direction to be traversed in state space, successive moves tend to lie in the same direction. This encourages the algorithm to move more quickly through high-probability regions of the state space and, in theory, allows the algorithm to more efficiently sample the posterior PDF. Once hybrid MCMC encounters a low-probability region of the state space, the accept/reject criterion reduces the momentum until an acceptably high-probability region of the space is found. For a much more thorough discussion of hybrid MCMC, the reader is referred to Neal (1994).

Another solution to the problem of inefficient sampling is to use a *non-symmetric* proposal distribution. In principle, many of the MCMC iterations result in discarded parameter sets because the algorithm has ventured into a low-probability region of the state space. Non-symmetric proposal distributions (sometimes referred to as the *Langevin rule* for Metropolis jumping or alternatively *stochastic gradient* methods), attempt to reduce the number of rejected samples by adapting the direction of sampling according to the local *gradient* of the posterior PDF. Because MCMC is encouraged to move up-gradient (toward higher values of probability), fewer iterations are wasted, and the algorithm samples more efficiently. While non-symmetric proposal distributions have been shown to work well for many problems, additional complexity is involved in the setup of the MCMC algorithm. This complexity enters the process through the accept/reject criterion, which must now take into account the fact that the probability of moving *to* a new location in the state space is not equivalent to the probability of moving *back*. Gelman et al. (2004) point out that the Langevin rule can be thought of as a special case of the hybrid method, in which

the momentum is recomputed for each move independently. For a further discussion of non-symmetric proposal distributions, the reader is referred to Robert and Casella (1999, p. 264-268).

Appendix D

Application of MCMC to CRM Simulation of a Single Unforced Convective Squall Line

D.1 Introduction

In chapters 3 and 4 above, MCMC was used to estimate PDFs of GCE cloud microphysical parameters conditioned on observed cloud, precipitation, and radiation statistics for a 2D version of the model forced with large-scale conditions. In this mode, the GCE predicted temperature, water vapor, and wind fields are adjusted at every grid point and timestep using observed large scale advective tendencies. The objective is to encourage the model to generate cloud and precipitation fields that interact with radiation in a realistic manner in the presence of observed large-scale flow conditions. The results are typically used to assess and improve general circulation model (GCM) cloud and convection parameterizations, though cloud resolving models (CRMs) have recently been run in this mode as a replacement for traditional GCM convective parameterizations.

Application of MCMC in such a framework was appropriate for the current study, as the objective was to obtain cloud parameters that were consistent with large-scale mean cloud and precipitation characteristics. However, it is natural to wonder whether parameter PDFs generated from a forced CRM truly represent the sensitivity of CRM results to changes in cloud microphysical parameters. This question is addressed in this appendix through application of MCMC to a freely-evolving 2D version of the GCE with open boundaries. To be consistent with the type of systems observed during SCSMEX, a similar thermodynamic and wind shear profile is used to initialize the model, and the algorithm is implemented exactly as it was for the forced simulations. As in the idealized experiments of chapter 3, time and domain mean clouds, precipitation, and radiative fluxes and heating rates are used as observations, and the same set of “truth” parameters are employed (Table 3.1), but in this case the model is only run for six hours.

Similar to the results obtained with the forced version of the model, a_g and N_{0r} exhibit sharply-peaked PDFs, indicating that these parameters exert a high degree of control over the simulated result. However, in contrast to forced simulations, it is the graupel fallspeed exponent (not snow) that serves as the third key parameter. As in the experiments presented in chapter 3, there is a secondary mode in the parameter space, but it is not as easy to deconvolve in the present case, as it exists for high likelihoods. Indeed, consistently high values of the likelihood function for a wide range of parameter values, coupled with high accept rates, indicate the observations do not provide enough information to constrain cloud microphysical parameters over the relatively short six hour observation window. Examination of the relationship between parameters and observations shows that most of the significantly related parameter-observation pairs identified in the forced simulations also exhibit robust quasi-linear relationships in freely-evolving simulations, though the relationships are much more well-defined in this case. As in forced simulations, though the graupel

density does not exhibit a well-defined PDF, it clearly shows up as a key parameter in the relationships with observations. It is generally concluded that the specifics of which parameters the model output is most sensitive to, as well as the relationship between changes to parameters and changes in model output statistics, is relatively independent of the model configuration or application of forcing. However, it is also noted that, in cases where MCMC is to be used specifically to characterize model uncertainty, appropriate simulated observations must be carefully chosen so as to maximize the applicability of the result. In this case, sensitivity of convective features structures to changes in cloud microphysical parameters was masked by the choice of relatively noisy temporally and spatially integrated observations. Longer simulations, coupled with observations more closely related to physical features of interest would likely have returned more well-defined parameter PDFs for this case. The remainder of this appendix is arranged as follows. Section D.2 describes the configuration of the GCE simulations used in the MCMC experiments, while section D.3 contains the results of the MCMC experiments. A brief discussion is presented in section D.4, and summary and conclusions are offered in section D.5.

D.2 GCE Model Setup

For the simulations presented in this appendix, the 2D version of the GCE is run on 256 horizontal grid points with horizontal grid spacing of 2 km in the innermost 218 grid points. Open radiative boundary conditions are applied by stretching the grid spacing outside the inner domain by a factor of 1.05 to prevent gravity wave reflection from the boundaries. Combination of fixed 2 km grid spacing on the inner 218 points with the stretched grid leads to a domain that is approximately 1000 km in width. 33 vertical levels are used, which stretch from 50 meter spacing in the boundary layer to approximately 1 km in the upper troposphere and lower stratosphere, and the model

domain is approximately 24 km deep. Analysis of both preliminary large-domain simulations and satellite observations (not shown) indicates that the typical time needed for a convective system to develop to maturity under TWP monsoon conditions is approximately 6 hours. In previous studies, the GCE model has been used to simulate a wide variety of tropical convective systems, and we choose use a case that is representative of the large-scale conditions over the SCSMEX domain—in this case, a simulation of a tropical squall line from 22 February 1993 that occurred during the TOGA-COARE field program (Webster and Lukas 1992). This case has been extensively studied as part of the GEWEX Cloud Systems Study Working Group 4 (Moncrieff et al. 1997), and is representative of the type of convection observed over the equatorial Pacific during monsoon conditions. The simulation domain translates in the horizontal at a speed equal to the storm propagation speed, and convection is initiated by applying cooling in the boundary layer over the first 10 minutes of simulation, gradually decreasing to zero over the following 10 minutes. The total length of simulation is six and a half hours, and a spinup of 30 minutes is allowed before observations are accumulated. Total cloud condensate mass mixing ratio and vertical motion at 30 minute intervals over the first three hours of simulation are depicted in figure D.1. As this figure shows, the model develops a vigorous updraft by 30 minutes into the simulation, and produces an extensive cirrus cloud shield by 2 simulated hours. At two and a half hours of simulation, the model is developing a new cell at the leading edge of the surface cold pool (cold pool position not shown), exhibiting the multicellular behavior typical of simulations of tropical squall lines (e.g., Fovell and Ogura 1988, Bryan et al. 2003, Lang et al. 2003). For completeness, additional simulations were performed using 500 meter grid spacing (not shown)—though these simulations produced a longer-lasting system with greater maximum upward and downward vertical velocities, the response of the climate parameters of interest (e.g., mean bottom and top of atmosphere short and longwave cloud radiative forcing,

precipitation rate, and liquid and ice water path) to changes in cloud microphysical parameters was nearly identical in magnitude to that observed in the microphysical parameter perturbation experiments documented in chapter 1.

D.3 Results

D.3.1 Parameter PDFs

As in chapters 3 and 4, the MCMC algorithm was first allowed an initial burn-in period in which sampling rates were tuned, then was run out to approximately 500,000 samples. It was found that, regardless of the width of the uniform proposal distribution, acceptance rates were consistently high. As a result, the chains converged to sampling a stationary distribution much sooner, and hence fewer samples were needed to characterize the PDF. Examination of the marginal parameter PDFs obtained from the MCMC algorithm (Fig. D.2) reveals that the slope intercept of the rain drop size distribution exhibits the most well-defined PDF, indicating that, as in the forced simulations, this parameter exerts a relatively large measure of control over simulation results. Apart from N_{0r} , and especially compared with simulations generated for forced simulations (Fig. 3.3), most of the PDFs for the unforced case are markedly flat, though there is a hint of a mode around the truth value for b_g and ρ_g . It is interesting to note the apparent tri-modal structure in the marginal PDF of a_g and the concentration of mass in the PDF of N_{0g} away from the true value; as with the PDFs generated for forced simulations, it appears that the solution space contains multiple modes. The fact that the bulk of the mass of the N_{0g} PDF in forced simulations was also located in the upper half of the parameter space indicates that this mode may be a robust feature of the cloud microphysical scheme itself. The lack of a definitive mode in most of the parameter PDFs is likely due to the limited simulation and observation time, as well as the fact that, in contrast to forced simulations, which




Figure D.1: Figure depicting the evolution of total condensate mixing ratio and upward vertical velocity for the 2D simulation with open boundaries.




Figure D.2: Marginal PDFs for each estimated parameter as represented in parameter histograms. The red horizontal line on each plot indicates the Uniform a priori PDF, while the vertical dashed line depicts the true parameter value.

contain an ensemble of cloud types and depths, current simulations are restricted to representation of a single deep convective system.

Additional information on the characteristics of the parameter probability space can be obtained from the joint PDF of each pair of parameters (figure D.3). Examination of this figure confirms, as expected, that N_{0r} is clearly the dominant parameter, and shows a_g and N_{0r} to be strongly anti-correlated. This result provides further evidence of the fact that there is a direct relationship between precipitation and the specifics of the graupel particle size distribution in the GCE cloud microphysical scheme. Note that the mode in b_g is more pronounced than it first appeared to be in the marginal PDFs. In fact, from inspection of the joint PDF of b_g and N_{0g} , it appears that the primary mode in this parameter is masked by the secondary mode in N_{0g} ; only in relationship with N_{0g} is there a significant portion of the mass of b_g located far from the truth value in the upper half of the parameter space. The three modes in a_g are clearly visible in the relationship with N_{0r} , and it is clear from their joint histogram that the mode closest to the truth value is the dominant mode. The fact that a portion of the mass of a_g , b_g , and N_{0g} is contained in the secondary mode is reflected in the fact that a sizable fraction of the mass of ρ_g is located far from its true value in combination with these three parameters. In contrast, the bulk of the mass of ρ_g is concentrated around its truth value in joint PDFs with every other parameter. Figure D.3 provides further evidence of the fact that the snow fallspeed coefficient and exponent, snow slope intercept, and snow density parameters have little influence over the result of the simulation and that we cannot as to the parameter truth values cannot be extracted from either marginal or joint PDFs. By contrast, the three parameters that do contribute significantly to the simulation results can be clearly identified from their joint PDFs.




Figure D.3: Joint PDFs for each pair of parameters. The color scale is the same for each plot with warmer colors representing higher values of the PDF. Each axis is labeled with its corresponding parameter.

Table D.1: List of observations vs. mean states for the top 0.1 % of all likelihood values.

State	Observations	Mean of States
Precipitation Rate	16.074	16.064
Cloud Top Temperature	211.141	210.949
Liquid Water Path	2469.890	2381.210
Ice Water Path	2819.540	2709.390
High Cloud Fraction	0.180	0.186
BOA LW Cloud Forcing	-2.775	-2.966
TOA LW Cloud Forcing	-29.901	-29.972
BOA SW Cloud Forcing	69.913	73.277
TOA SW Cloud Forcing	70.009	73.245
Low Level All-Sky LW Cooling	-2.311	-2.292
Mid Level All-Sky LW Cooling	-2.705	-2.701
Upper Level All-Sky LW Cooling	-1.032	-1.050
Low Level Clear-Sky LW Cooling	-2.482	-2.480
Mid Level Clear-Sky LW Cooling	-2.953	-2.951
Upper Level Clear-Sky LW Cooling	-1.293	-1.293
Low Level All-Sky SW Heating	-1.617	-1.603
Mid Level All-Sky SW Heating	-2.277	-2.285
Upper Level All-Sky SW Heating	-1.771	-1.779
Low Level Clear-Sky SW Heating	-1.757	-1.756
Mid Level Clear-Sky SW Heating	-2.392	-2.388
Upper Level Clear-Sky SW Heating	-1.439	-1.440

As a final assessment of the robustness of each parameter PDF, we examine the set of parameter values that correspond to the highest 0.1% of the likelihood values; in chapters 3 and 4 it was demonstrated that such an analysis can help to differentiate the primary mode in the PDF from secondary modes. Mean state values for the top 0.1% of the likelihoods are listed alongside the observation values in Table D.1, and it is clear that the parameters associated with the very highest total likelihood values did indeed generate state values that were very close to those observed. In fact, the states associated with the highest likelihood values exhibited a much closer match to

observations than in the forced case (Table 3.4). An examination of the histograms of each state variable conditioned on the parameter sample (Fig. D.5, see below), demonstrates that this is not due to a relative lack of variability in simulated states.

Examination of scatter plots of each parameter vs. the total likelihood (Fig. D.4) reveals, as anticipated, that a_g , b_g , and N_{0r} exhibit a tight grouping about their truth parameter values, while a_s , b_s , N_{0s} , and ρ_s display widely scattered parameter values for even the highest likelihood values. At first glance, ρ_g appears to be relatively well-constrained, but exhibits values grouped about the center of the range. Similarly, high-likelihood N_{0g} parameter values are relatively tightly grouped, but are concentrated inside the secondary mode far from the true value. This indicates that, for simulations as short as six hours in length, the secondary mode contains approximately the same order of magnitude of probability mass as the true mode. Examination of the joint parameter histograms (Fig. D.3) indicates that the shift in mass of ρ_g toward the center of its range is likely due to the presence of this secondary mode.

D.3.2 Examination of the conditional PDFs of observations

$$P(\mathbf{y}|\mathbf{x})$$

As was shown in chapter 3, further information on the characteristics of the state space can be obtained from an examination of the PDF of each observation conditioned on the set of parameters $P(y_i|\mathbf{x})$ (Fig. D.5). As with PDFs of states in the unforced case (Fig. 3.7), many of the PDFs in figure D.5 exhibit a mode that is shifted compared to the true observation value. However, though the parameter PDFs vary relatively smoothly across the range of each parameter, most of the state PDFs exhibit multiple modes. Interestingly, many of these modes are widely separated from the truth value, indicating the existence of a separate discrete solution in the state space. Of the single-mode PDFs, only the high cloud fraction appears to be approximately Gaussian in




Figure D.4: Scatter plots of the largest 0.1% of likelihood values for each parameter.




Figure D.5: Marginal PDFs of the set of observations used in the MCMC experiment. In each plot, the range of values of each observed variable is depicted on the abscissa, while the normalized number of counts is displayed on the ordinate. The true value of each observation is depicted in the vertical dashed green line. The plots are subdivided into (a) cloud and precipitation observations, (b) long and shortwave cloud radiative forcings, (c) three-layer longwave cooling rates, and (d) three-layer shortwave heating rates.

form—the rest are skewed to varying degrees. Since the “true” state values are known, it can be reasonably hypothesized that parameter values located in the space occupied by the true parameter mode produce states coincident with the observations, while parameters in the region of the space containing secondary modes are associated with modes in the state space that are farther from the observations.

To test this assertion, the full parameter sample is subset into only those parameter values that lie within $\pm 1\sigma$ of the truth value, and PDFs of associated states are correspondingly recomputed. While subsetting the state space had little effect in the forced case, use of parameter subsets produces a dramatic change in the state PDFs for unforced simulations. Modes that are separate from the true observation value decrease markedly in amplitude, while the primary mode shifts closer to the observations and becomes more sharply defined. Only in the case of the precipitation rate is the primary mode not coincident with the observed value—this is likely due to the fact that the subset range still contains some of the mass of the secondary parameter mode in N_{0g} , as well as the secondary modes in a_g . From the differences in the PDFs produced by subsetting the sample, we can draw inferences about the characteristics of the state associated with secondary modes in the parameter space. Comparison of figures D.5 and D.6 reveals that the secondary modes are associated with higher precipitation rate, larger ice and liquid water path, and greater cloud fraction. Consistent with greater cloud coverage and mass, the secondary mode is also associated with larger magnitude cloud radiative forcing in both the longwave and shortwave spectrum at both the top and bottom of the atmosphere. Radiative heating rates are consistent with more high cloud—the secondary mode is characterized by larger mid and upper level clear and all-sky longwave cooling, and with reduced longwave cooling at low levels. The secondary mode also produces decreased low-level shortwave heating and increased mid and upper level heating. It is sobering, in light of the highly non-Gaussian parameter and state conditional PDFs, that variational

parameter estimation procedures rely on the assumption of a uni modal Gaussian state space. It is clear from the results presented above that a blind application of variational parameter estimation methods could easily result in highly erroneous results.

D.3.3 Assessment of Observation/Parameter Relationships

In chapter 3, a rigorous assessment of each observation-parameter relationship was performed for the purpose of identifying observations that could be used in real-data assimilation experiments. In this appendix, we do not intend to constrain parameters using real data, but instead to characterize the control parameters exert over the simulation results. To this end, we repeat the procedure documented in section 3.5.3 and assess which parameter-state pairs exhibit the most robust linear relationships. In so doing, we hope to complete the characterization of the sensitivity of GCE output to changes in cloud microphysical parameters.

The 30 most significant parameter-state pairs are listed in Table D.2 in descending order according to the signal to noise ratio of the relationship. As expected from the parameter histograms, changes to a_g and N_{or} consistently produce a robust response in many of the observations. However, as with the forced simulations, ρ_g is also consistently related to the observations, and it is likely that the PDF of ρ_g is likely not more sharply peaked because of the relatively large assumed errors in the observations it is significantly related to (longwave and shortwave cloud radiative forcings and all-sky heating rates). Though the likelihood scatter plot (Fig. D.4) appeared to suggest a strong relationship between b_g and GCE output states, b_g only exhibits signal to noise ratio greater than 1.5 (not shown) in relation to high cloud fraction. One possible explanation is that several states were affected in small ways by changes to this parameter, and that no single state contained a large amount of information content. By contrast, though b_s exhibits a robust correlation with the upper level all-




Figure D.6: As in figure D.5, but for marginal PDFs of the set of observations that corresponds to the primary mode in the joint parameter PDF (see text for details).

Table D.2: List of the parameter/observation pairs with the most significant linear relationships, as represented in their linear correlation and signal to noise ratio (see text for description).

Parameter	State Variable	ρ	SNR
a_g	Ice Water Path	-0.89	3.46
N_{0r}	Sfc-743 mb SW Clear-Sky Radiative Heating	0.83	3.24
N_{0r}	BOA LW Cloud Radiative Forcing	-0.81	3.17
N_{0r}	Liquid Water Path	0.80	3.14
ρ_g	523 mb-TOA SW All-Sky Radiative Heating	-0.81	3.10
N_{0r}	Sfc-743 mb LW All-Sky Radiative Cooling	0.78	3.04
N_{0r}	743-523 mb SW All-Sky Radiative Heating	-0.75	2.95
N_{0r}	Cloud Top Temperature	0.74	2.89
N_{0r}	BOA SW Cloud Radiative Forcing	0.74	2.89
ρ_g	TOA LW Cloud Radiative Forcing	-0.74	2.82
N_{0r}	Sfc-743 mb SW All-Sky Radiative Heating	0.71	2.78
ρ_g	743-523 mb LW All-Sky Radiative Cooling	0.72	2.77
a_g	523 mb-TOA LW All-Sky Radiative Cooling	0.69	2.68
N_{0r}	TOA SW Cloud Radiative Forcing	0.68	2.67
N_{0r}	Sfc-743 mb LW Clear-Sky Radiative Cooling	0.66	2.60
b_s	523 mb-TOA LW All-Sky Radiative Cooling	0.68	2.57
N_{0r}	523 mb-TOA LW Clear-Sky Radiative Cooling	-0.64	2.50
N_{0r}	743-523 mb LW Clear-Sky Radiative Cooling	-0.61	2.38
a_g	Sfc-743 mb LW Clear-Sky Radiative Cooling	0.61	2.35
a_g	743-523 mb LW All-Sky Radiative Cooling	-0.59	2.27
ρ_g	TOA SW Cloud Radiative Forcing	0.56	2.16
ρ_g	743-523 mb SW All-Sky Radiative Heating	0.56	2.14
a_g	Rain Rate	0.55	2.12
N_{0r}	743-523 mb SW Clear-Sky Radiative Heating	-0.53	2.08
N_{0r}	523 mb-TOA LW All-Sky Radiative Cooling	-0.53	2.07
N_{0r}	High Cloud Fraction	0.53	2.06
a_g	Liquid Water Path	0.53	2.04
ρ_g	Sfc-743 mb SW All-Sky Radiative Heating	0.52	2.00
a_g	743-523 mb LW Clear-Sky Radiative Cooling	-0.51	2.00

sky longwave radiative cooling, its PDF was absolutely flat. In this case, the single state that was strongly affected by changes in b_s did not contribute enough to the likelihood function to direct the course of the MCMC sample. These two examples underscore the importance of examining the parameter space from many different perspectives.

Comparison of Table D.2 with Table 3.6 reveals consistently higher signal to noise ratio and correlation coefficient for the TOGA-COARE squall line simulations compared with the SCSMEX case. This is a reflection of the fact that changes to parameters produce a different effect for different cloud types; whereas unforced simulations contain a full ensemble of low, mid, and high clouds, unforced simulations contain only a single squall line. It is interesting to note that, although the list is not ordered in quite the same way—the most significant parameter-state pairs in the unforced case are not necessarily those that are the most significant in the forced simulations—both lists are populated with the same general set of related pairs, indicating that the underlying physical relationships are robust. As with forced simulations, parameters exhibit the strongest correlation with both clear and all-sky radiative heating rates, as well as with top and bottom of atmosphere cloud radiative forcings.

It is interesting to note that, as in the forced simulations, few of the observations of cloud and precipitation variables exhibited significant relationships with any of the cloud microphysical parameters. Examination of the range of values for cloud and precipitation properties, as represented in the state PDFs (Fig. D.5) indicates that this is not due to a lack of variability. As was mentioned in chapter 3, lack of a robust linear relationship between changes to parameters and the response of output state variables may be due to the highly nonlinear relationship between cloud microphysical parameters and the cloud mass mixing ratios (and, by extension, cloud water content and precipitation rate). In addition, as with the lack of definition in many of the parameter PDFs, it could be argued that longer simulations might produce a more

robust relationship between cloud parameters and time-integrated observations.

D.4 Discussion

In this appendix, temporal and spatial means of precipitation rate, cloud amount, cover, and height, and radiative fluxes and heating rates were used as surrogate observations in the MCMC procedure. The purpose was to characterize the PDFs of cloud microphysical parameters for a simulation of a single squall line, and the observations were chosen for consistency with the forced simulations used earlier. However, apart from allowing a comparison with earlier MCMC results, it is clear that this particular set of observations are not well-suited to characterize the actual uncertainty in the cloud microphysical scheme for this type of simulation. While temporally-accumulated statistics of the cloudy state and radiation budget are useful for characterizing the role parameters play in the evolution of ensembles of clouds on relatively large domains, it is the specifics of the convective structure that are of interest in simulations of individual systems. In this case, fields such as precipitation efficiency, mean updraft and downdraft velocity, convective vs. stratiform partition, storm propagation speed, and ratio of the masses of condensate variables are of greater interest.

Though such variables were not used in the MCMC estimation above, many were stored as part of the resulting database of model output variables, and it is useful to examine the characteristics of their sample. Domain mean precipitation efficiency, updraft and downdraft velocity, and total column-integrated water path for each condensate variable are plotted in figure D.7. Updraft and downdraft velocities were computed as the mean positive and negative vertical velocity, respectively, where total cloud condensate mass mixing ratio exceeded $10^{-5} \text{ kg} \cdot \text{kg}^{-1}$. As with many of the states used in the MCMC estimation, PDFs of additional states are uniformly

non-Gaussian in character, and most exhibit multiple modes. In contrast to the set of states examined above, when PDFs are re-generated for the subset of parameter values located around the true mode, these PDFs do *not* change their shape—all modes remain. The implication is that these variables contain further information on the sensitivity of GCE output to changes in parameters that could be used to further assess the performance of the model.

In all of the above analysis, temporal and spatial means have been used—even in short-term simulations, it could be argued that averaged quantities are useful, as microphysical parameter values are set fixed over the length of each simulation and have a time and space integrated effect on model output. However, many features of squall lines are *transient*—e.g., maximum vertical velocity, rapidity of convective development and decay, and variability in updrafts and downdrafts. Other important characteristics are *structural*—e.g., the magnitude of the cold pool thermal perturbation, the character of multicellular development, and the strength of the overturning circulation. It could be contended that the sensitivity of model output to changes in parameters could be more effectively characterized if observations related to these features were used in the MCMC algorithm.

D.5 Conclusions

In this appendix, the MCMC procedure introduced in chapters 3 and 4 for GCE simulations forced with SCSMEX advective tendencies was used to characterized PDFs of cloud microphysical parameters for simulations of a single convective squall line. It was found that the same general set of parameters that control the evolution of forced simulations was also important in unforced simulations. In addition, though there were slight differences in the structure of the PDFs themselves, there was again a secondary mode clearly evident in the solution space. This secondary mode was




Figure D.7: PDFs of additional saved GCE state variables for the entire parameter sample. As in the previous PDF plots, the frequency of occurrence of each value is plotted on the ordinate, while the range of values of each variable is plotted on the abscissa. Each PDF is normalized to 1.0.

associated with large graupel fallspeed coefficient and exponent, and in particular with large graupel particle size slope intercept. PDFs of the states conditioned on the parameters ($P(\mathbf{y}|\mathbf{x})$) also contained clearly visible secondary modes, and it was shown that these modes disappeared or were greatly reduced when the parameter sample was subset into the region of the state space surrounding the true parameter mode. The general conclusion is first that MCMC can be used as a useful tool with which to characterize parameter uncertainty in numerical models, even in the absence of a real-data component. In addition, it appears that the characteristics of the GCE cloud microphysical scheme revealed in MCMC experiments with forced simulations are also generally true for unforced simulations, indicating that the parameter sample returned by MCMC characterizes the behavior of the cloud microphysical scheme itself and is only weakly dependent on the model configuration. Define a shear profile and forcing dataset representative of the TWP (or SCSMEX?) Use the same 30 hour CRM simulation as used in MCMC with default microphysics parameter values Rotate the CRM through 90 degrees in 15 degree increments (7 simulations) Note the total precipitation and precipitation rate, IWP, LWP, cloud radiative forcings, and radiative and latent heating rates (as well as precipitable water vapor, updrafts, downdrafts...)

References

Ackerman, S. A., and G. L. Stephens, 1987: The absorption of solar radiation by cloud droplets: An application of anomalous diffraction theory. *J. Atmos. Sci.*, **44**, 1574-1588.

Auer, A. H. Jr, and J. D. Maurwitz, 1968: Estimates of air and moisture flux into hailstorms on the high plains. *J. Appl. Meteor.*, **7**, 196-198.

Austin, R. T., and G. L. Stephens, 2001: Retrieval of stratus cloud microphysical parameters using millimeter-wave radar and visible optical depth in preparation for CloudSat. Part I: Algorithm formulation. *J. Geophys. Res.*, **106**, 28233-28242.

Benedetti, A., G. L. Stephens, and J. M. Haynes, 2003: Ice cloud microphysics retrievals from millimeter radar and visible optical depth using an estimation theory approach, *J. Geophys. Res.*, **108**, DOI: 10.1029/2002JD002693

Bigg, E. K., 1953: The supercooling of water. *Proc. Phys. Soc. London.*, **B66**, 688-694.

Boer, G. J., and B. Yu, 2003a: Climate sensitivity and response. *Climate Dyn.*, **20**, 415-429.

_____, 2003b: Climate sensitivity and climate state. *Climate Dyn.*, **21**, 167-176.

Braham, R. R. Jr., 1952: The water and energy budgets of the thunderstorm and their relation to thunderstorm development. *J. Meteor.*, **9**, 227-242.

Bretherton, C. S., and P. K. Smolarkiewicz, 1989: Gravity waves, compensating subsidence and detrainment around cumulus clouds. *J. Atmos. Sci.*, **46**, 740-759.

Bryan, G. H., J. C. Wyngaard, and J. M. Fritsch, 2003: Resolution requirements for the simulation of deep moist convection. *Mon. Wea. Rev.*, **131**, 2394-2416.

Chepurin, G. A., J. A. Carton, and D. P. Dee, 2005: Forecast model bias correction in ocean data assimilation. *Mon. Wea. Rev.*, **133**, 1328-1342.

Chong, M, and D. Hauser, 1989: A tropical squall line observed during the CORT81 experiment in west Africa. Part II: Water budget. *Mon. Wea. Rev.*, **117**, 728-744.

Chou, M.-D., and L. Kouvaris, 1991: Calculations of transmission functions in the IR CO₂ and O₃ bands. *J. Geophys. Res.*, **96**, 9003-9012.

_____, W. Ridgway, and M.-H. Yan, 1995: Parameterizations for water vapor IR radiative transfer in both the middle and lower atmospheres. *J. Atmos. Sci.*, **52**, 1159-1167.

_____, and M. J. Suarez, 1999: A shortwave radiation parameterization for atmospheric studies. **15**, *NASA/TM-104606*, 40 pp.

_____, K.-T. Lee, S.-C. Tsay, and Q. Fu, 1999: Parameterization for cloud long-wave scattering for use in atmospheric models. *J. Clim.*, **12**, 159-169.

Ciesielski, P. E., and R. H. Johnson, 2006: Contrasting characteristics of convection over the northern and southern South China Sea during SCSMEX. *Mon. Wea. Rev.*, In Press.

Cooper, S. J., T. S. L'Ecuyer, and G. L. Stephens, 2003: The impact of explicit cloud boundary information on ice cloud microphysical property retrievals from infrared radiances. *J. Geophys. Res.*, **108**, DOI: 10.1029/2002JD002611

Cotton, W. R., R. A. Pielke Sr., R. L. Walko, G. E. Liston, C. J. Tremback, H. Jiang, R. L. McAnelly, J. Y. Harrington, M. E. Nicholls, G. G. Carrio, and J. P. McFadden, 2003: RAMS 2001: Current status and future directions. *Meteor. Atmos. Phys.*, **82**, 5-29.

Deardorff, J. W., 1975: The development of boundary-layer turbulence models for use in studying the severe storm environment. *Proceedings of the SESAME Opening Meeting*, Boulder, NOAA-ERL, 251-264.

Dee, D. P., and A. M. Da Silva, 1998: Data assimilation in the presence of forecast bias. *Quart. Jour. Roy. Met. Soc.*, **124**, 269-295.

Devore, J., 1995: *Probability and Statistics for Engineering and the Sciences*, 4th ed., Wadsworth, Belmont, CA.

Donner, L. J., C. J. Seman, and R. S. Hemler, 1999: Three-dimensional cloud-system modeling of GATE convection. *J. Atmos. Sci.*, **56**, 1885-1912.

Engelen, R. J. and G. L. Stephens, 1997: infrared radiative transfer in the 9.6 μm band: Application to TOVS ozone retrieval. *J. Geophys. Res.*, **102**, 6929-6940.

_____, and G.L. Stephens, 1999: Characterization of water vapour retrievals from TOVS/HIRS and SSM/T-2 measurements. *Quart. J. Roy. Meteor. Soc.*, **125**, 331-351.

Fairall, C., E. F. Bradley, D. P. Rogers, J. B. Edson, and G. S. Young, 1996: Bulk Parameterization of air-sea fluxes for TOGA COARE. *J. Geophys. Res.*, **101**, 3747-3764.

Fankhauser, J. C., 1988: Estimates of thunderstorm precipitation efficiency from field measurements in CCOPE. *Mon. Wea. Rev.*, **116**, 663-684.

Ferrier, B. S., 1994: A double-moment multiple-phase four-class bulk ice scheme. Part I: Description. *J. Atmos. Sci.*, **51**, 249-280.

_____, W.-K. Tao, and J. Simpson, 1995: A double-moment multiple-phase four-class bulk ice scheme. Part II: Simulations of convective storms in different large-scale environments and comparisons with other bulk parameterizations. *J. Atmos. Sci.*, **52**, 1001-1033.

_____, J. Simpson, and W.-K. Tao, 1996: Factors responsible for precipitation efficiencies in midlatitude and tropical squall simulations. *Mon. Wea. Rev.*, **124**, 2100-2125.

Fletcher, N. H., 1962: *The Physics of Rainclouds*. Cambridge University Press, 386pp.

Foote, G. G., and J. C. Fankhauser, 1973: Airflow and moisture budget beneath a northeast Colorado hailstorm. *J. Appl. Meteor.*, **12**, 1330-1353.

Fovell, R. G., and Y. Ogura, 1989: Effect of vertical wind shear on numerically simulated multicell storm structure. *J. Atmos. Sci.*, **46**, 3144-3176.

Gabriel, P. M., M. Christi, and G. L. Stephens, 2005: Calculation of Jacobians for Inverse Radiative Transfer: An Efficient Hybrid Method. *J. Quant. Spect. and Rad. Trans.* In Press

Gamache, J. F., and R. A. Houze Jr, 1983: Water budget of a mesoscale convective system in the tropics. *J. Atmos. Sci.*, **40**, 1835-1850.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin 2004: *Bayesian Data Analysis, 2nd ed.*, Chapman and Hall/CRC, New York, NY.

Geman, S., and D. Geman, 1984: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Machine Intelligence*, **6**, 721-741.

Grabowski, M. W. Moncrieff, and J. T. Kiehl, 1996a: Long-term behaviour of precipitating tropical cloud systems: a numerical study. *Quart. J. Roy. Meteor. Soc.*, **122**, 1019-1042.

Grabowski, W. W., X. Wu, and M. Moncrieff, 1996b: Cloud-resolving modeling of tropical cloud systems during phase III of GATE. Part I: Two-dimensional experiments. *J. Atmos. Sci.*, **53**, 3684-3709.

_____, X. Wu, and M. W. Moncrieff, 1999: Cloud-resolving modeling of tropical cloud systems during phase III of GATE. Part III: Effects of cloud microphysics. *J. Atmos. Sci.*, **56**, 2384-2402.

_____, 1999: A parameterization of cloud microphysics for long-term cloud-resolving modeling of tropical convection. *Atmos. Res.*, **52**, 17-41.

_____, 2000: Cloud microphysics and the tropical climate: cloud-resolving model perspective. *J. Climate*, **13**, 2306-2322.

_____, 2001: Coupling cloud processes with the large-scale dynamics using the Cloud-Resolving Convection Parameterization (CRCP). *J. Atmos. Sci.*, **58**, 978-997.

_____, 2003: Impact of cloud microphysics on convective-radiative quasi equilibrium revealed by cloud-resolving convective parameterization. *J. Climate*, **16**, 3463-3475.

Gunn, R. and G. D. Kinzer, 1949: The terminal velocity of fall for water droplets in stagnant air. *J. Meteor.*, **6**, 243-248.

Hastings, W., 1970: Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, **57**, 97-109.

Held, I. M., R. S. Hemler, and V. Ramaswamy, 1993: Radiative-convective equilibrium with explicit two-dimensional moist convection. *J. Atmos. Sci.*, **50**, 3909-3927.

Heney, L., and J. Greenstein, 1941: Diffuse radiation in the galaxy. *Astrophys. Jour.*, **93**, 70-83.

Heymsfield, A. J., A. Bansemer, P. R. field, S. L. Durden, J. L. Smith, J. E. Dye, W. Hall, and C. A. Grainger, 2002: Observations and parameterizations of particle size distributions in deep tropical cirrus and stratiform precipitating clouds: Results from in situ observations in TRMM field campaigns. *J. Atmos. Sci.*, **59**, 3457-3491.

Heymsfield, G. M., and S. Schotz, 1985: Structure and evolution of a severe squall line over Oklahoma. *Mon. Wea. Rev.*, **113**, 1563-1589.

Houze, R. A., 2004: Mesoscale convective systems. *Rev. Geophys.*, **42**, RG4003.

Johnson, R. H., and P. E. Ciesielski, 2002: Characteristics of the 1998 summer monsoon onset over the northern South China Sea. *J. Meteor. Soc. Japan*, **80**, 561-578.

_____, S. L. Aves, P. E. Ciesielski, and T. D. Keenan, 2005: Organization of oceanic convection during the onset of the 1998 East Asian summer monsoon. *Mon. Wea. Rev.*, **133**, 131-148.

Khain, A., A. Pokrovsky, and I. Sednev, 1999: Some effects of cloud-aerosol interaction on cloud microphysics structure and precipitation formation: numerical experiments with a spectral microphysics cloud ensemble model. *Atmos. Res.*, **52**, 195-220.

_____, M. Ovtchinnikov, M. Pinsky, A. Pokrovsky, and H. Krugliak, 2000: Notes on the state-of-the-art numerical modeling of cloud microphysics. *Atmos. Res.*, **55**,

159-224.

_____, D. Rosenfeld, and A. Pokrovsky, 2005: Aerosol impact on the dynamics and microphysics of convective clouds. *Q. J. Roy. Met. Soc.*, In Press.

Khairoutdinov, M. F., and D. A. Randall, 2001: A cloud-resolving model as a cloud parameterization in the NCAR Community Climate System Model: Preliminary results. *Geophys. Res. Lett.*, **28**, 3617-3620.

Kirkpatrick, S., C. D. Gelatt, Jr., and M. P. Vecchi, 1983: Optimization by simulated annealing, *Science*, **220**, 671-680.

Klemp, J. B., and R. B. Wilhelmson, 1978: The simulation of three-dimensional convective storm dynamics. *J. Atmos. Sci.*, **35**, 1070-1096.

Koenig, L. R., 1971: Numerical modeling of ice deposition. *J. Atmos. Sci.*, **28**, 226-237.

Kratz, D. P., M.-D. Chou, M.-H. Yan, and C.-H. Ho, 1998: Minor trace gas radiative forcing calculations using the k-distribution method with one-parameter scaling. *J. Geophys. Res.*, **103**, 31647-31656.

Krueger, S. K., F. Qiang, K.-N. Liou, and H.-N. Chin, 1995: Improvements of an ice-phase microphysics parameterization for use in numerical simulations of tropical convection. *Jour. Appl. Meteor.*, **34**, 281-287.

Kummerow, C., Y. Hong, W. S. Olson, S. Yang, R. F. Adler, J. McCollum, R.

Ferraro, G Petty, D. B. Shin, and T. T. Wilheit, 2001: The evolution of the Goddard Profiling Algorithm (GPROF) for rainfall estimation from passive microwave sensors. *J. Appl. Meteor.*, **40**, 1801-1820.

L'Ecuyer, T. S., and G. L. Stephens, 2002: An estimation-based precipitation retrieval algorithm for attenuating radars. *J. Appl. Meteor.*, **41**, 272-285.

_____, and G. L. Stephens, 2003: The tropical oceanic energy budget from the TRMM perspective. Part I: Algorithm and uncertainties. *J. Climate*, **16**, 1967-1985.

_____, P. M. Gabriel, K. Leesman, S. J. Cooper, and G. L. Stephens, 2005: Objective Assessment of the Information Content of Visible and Infrared Radiance Measurements for Cloud Microphysical Property Retrievals over the Global Oceans. Part I: Liquid Clouds. *J. Appl. Meteor.* In Press.

Lau, K.-M., T. Nakazawa, and C. H. Sui, 1991: Observations of cloud cluster hierarchies over the tropical western Pacific. *J. Geophys Res.*, **96** (Suppl.), 3197-3208

Li, X., C.-H. Sui, and K.-M. Lau, 2002: Precipitation efficiency in the tropical deep convective regime: a 2-d cloud resolving modeling study. *J. Met. Soc. Japan*, **80**, 205-212.

Lin, Y.-L., R. D. Farley, and H. D. Orville, 1983: Bulk parameterization of the snow field in a cloud model. *J. Climate Appl. Meteor.*, **22**, 1065-1092.

Lipps, F. B., and R. S. Hemler, 1986: Numerical simulation of deep tropical con-

vection associated with large scale convergence. *J. Atmos. Sci.*, **43**, 1796-1816.

Lindzen, R. S., M. D. Chou, and A. Y. Hou, 2001: Does the Earth have an adaptive infrared iris? *Bull. Amer. Meteor. Soc.*, **3**, 417-432.

Locatelli, J. D., and P. V. Hobbs, 1974: Fall speeds and masses of solid precipitation particles. *J. Geophys. Res.*, **79**, 2185-2197.

Lucas, C., E. J. Zipser, and B. S. Ferrier, 2000: Sensitivity of tropical west Pacific oceanic squall lines to tropospheric wind and moisture profiles. *J. Atmos. Sci.*, **57**, 2351-2373.

Luo, Z., and W. B. Rossow, 2004: Characterizing tropical cirrus life cycle, evolution, and interaction with upper-tropospheric water vapor using Lagrangian trajectory analysis of satellite observations. *J. Climate*, **17**, 4541-4563.

_____, and G. L. Stephens, 2006: An enhanced convection-wind-evaporation feedback in a superparameterization GCM (SP-GCM) depiction of the Asian summer monsoon. *Geophys. Res. Lett.*, **33**, L06707, doi:10.1029/2005GL025060.

Madden, R.A., and P. R. Julian, 1972: Description of global-scale circulation cells in the tropics with a 40-50 day period. *J. Atmos. Sci.*, **29**, 1109-1123.

Madden, R. A., and P. R. Julian, 1994: Observations of the 40-50-day tropical oscillation – a review. *Mon. Weather Rev.*, **122**, 814-837.

Mapes, B. E., 1993: Gregarious tropical convection. *J. Atmos. Sci.*, **50**, 2026-2037.

_____, and R. A. Houze, 1993: Cloud clusters and superclusters over the oceanic warm pool. *Mon. Wea. Rev.*, **121**, 1398-1415.

Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, 1953: Equations of state on fast computing machines. *J. Chemical Physics*, **21**, 1087-1092.

Meyers, M. P., R. L. Walko, J. Y. Harrington, and W. R. Cotton, 1997: New RAMS microphysics parameterization. Part II: the two-moment scheme. *Atmos. Res.*, **45**, 3-39.

Miles, N. L., J. Verlinde, and E. E. Clothiaux, 2000: Cloud droplet size distributions in low-level stratiform clouds. *J. Atmos. Sci.*, **57**, 295-311.

Miller, S. D., G. L. Stephens, C. K. Drummond, A. K. Heidinger and P. T. Partain, 2000: A multisensor diagnostic satellite cloud property retrieval scheme. *J. Geophys. Res.*, **105**, 19955-19971.

Moncrieff, M. W., S. K. Krueger, D. Gregory, J.-L. Redelsperger, and W.-K. Tao, 1997: GEWEX Cloud System Study (GCSS) Working Group 4: Precipitating Convective Cloud Systems. *Bull. Amer. Meteor. Soc.*, **78**, 831-845.

Mosegaard, K., and A. Tarantola, 1995: Monte Carlo sampling of solutions to inverse problems. *J. Geophys. Res.*, **100**, 12,431-12,447.

Murray, F. W., and L. R. Koenig, 1972: Numerical experiments on the relation

between microphysics and dynamics in cumulus convection. *Mon. Wea. Rev.*, **100**, 217-232.

Nakajima, T., and M. D. King, 1990: Determination of the optical thickness and effective particle radius of clouds from reflected solar radiation measurements. Part I: Theory. *J. Atmos. Sci.*, **47**, 1878-1893.

Nakazawa, T., 1988: Tropical superclusters within intraseasonal variations over the western Pacific. *J. Met. Soc. Japan.*, **66**, 823-839.

Neal, R., M., 1994: An improved acceptance procedure for the hybrid Monte Carlo algorithm. *J. Computational Phys.*, **111**, 194-203.

Oury, S., X. Dou, and J. Testud, 2000: Estimate of precipitation from the dual-beam airborne radars in TOGA-COARE. Part II: Precipitation efficiency in the 9 February 1993 MCS. *J. Appl. Meteor.*, **8**, 249-253.

Petch, J. C., and M. E. B. Gray, 2001: Sensitivity studies using a cloud-resolving model simulation of the tropical west Pacific. *Q. J. Roy. Met. Soc.*, **127**, 2287-2306.

Ramanathan, V. and W. Collins, 1991: Thermodynamic regulation of ocean warming by cirrus clouds deduced from observations of the 1987 El Nino. *Nature*, **351**, 27-32.

Randall, D. A., M. E. Schlesinger, V. Galin, V. Meleshko, J.-J. Morcrette, and R. Wetherald, 2001: *Cloud Feedbacks*. In "Frontiers in Climate Modeling". J. T. Kiehl and V. Ramanathan Ed. In press.

_____, M. Khairoutdinov, A. Arakawa, and W. W. Grabowski, 2003: Breaking the

cloud parameterization deadlock. *Bull. Amer. Met. Soc.*, **84**, 1547-1564.

Robert, C. P., and G. Casella, 1999: *Monte Carlo Statistical Methods*, Springer, New York, NY.

Rodgers, C. D., 2000: *Inverse Methods for Atmospheric Sounding, Theory and Practice*, World Scientific, Singapore

Roy, S. S., R. K. Datta, R. C. Bhatia and A. K. Sharma, 2005: Drop size distributions of tropical rain over south India. *Geofizika*, **22**, 105-130.

Rutledge, S. A., and P. V. Hobbs, 1983: The mesoscale and microscale structure and organization of clouds and precipitation in midlatitude cyclones. VIII: A model for the “seeder-feeder” process in warm-frontal rainbands. *J. Atmos. Sci.*, **40**, 1185-1206.

_____, and P. V. Hobbs, 1984: The mesoscale and microscale structure and organization of clouds and precipitation in midlatitude cyclones. Part XII: A diagnostic modeling study of precipitation development in narrow cold frontal rainbands. *J. Atmos. Sci.*, **41**, 2949-2972.

Saleeby, S. M., and W. R. Cotton, 2004: A large-droplet mode and prognostic number concentration of cloud droplets in the Colorado State University Regional Atmospheric Modeling System (RAMS). Part I: module descriptions and supercell test simulations. *J. Appl. Meteor.*, **43**, 182-195.

Shepherd, J. M., B. S. Ferrier, and P. S. Ray, 2001: Rainfall morphology in Florida

convergence zones: A numerical study. *Mon. Wea. Rev.*, **129**, 177-197.

Shie, C.-L., W.-K. Tao, J. Simpson, and C.-H. Sui, 2003: Quasi-equilibrium states in the Tropics simulated by a cloud-resolving model. Part I: Specific features and budget analysis. *J. Climate*, **5**, 817-833.

Smolarkiewicz, P. K., 1983: A simple positive definite advection scheme with small implicit diffusion. *Mon. Wea. Rev.*, **111**, 479-486.

Smolarkiewicz, P. K., 1984: A fully multidimensional positive definite advection transport algorithm with small implicit diffusion. *J. Comput. Phys.*, **54**, 325-362.

Smolarkiewicz, P. K. and W. W. Grabowski, 1990: The multidimensional positive advection transport algorithm: nonoscillatory option. *J. Comput. Phys.*, **86**, 355-375.

Soong, S.-T., and Y. Ogura, 1980: Response of trade wind cumuli to large-scale processes. *J. Atmos. Sci.*, **37**, 2036-2050.

_____, and W.-K. Tao, 1980: Response of deep tropical clouds to mesoscale processes. *J. Atmos. Sci.*, **37**, 2016-2035.

Stephens, G. L., 1978: Radiation profiles in extended water clouds, Part I: Theory. *J. Atmos. Sci.*, **35**, 2111-2122.

_____, and P. J. Webster, 1981: Clouds and climate: sensitivity of simple systems.

J. Atmos. Sci., **38**, 235-247.

_____, M. Vaughn, R. J. Engelen and T. L. Anderson, 2001: Toward retrieving properties of the tenuous atmosphere using space-based lidar measurements. *J. Geophys. Res.*, **106**, 28143-28157.

_____, D. G. Vane, R. J. Boain, G. G. Mace, K. Sassen, Z. Wang, A. J. Illingworth, E. J. O'Connor, W. B. Rossow, S. L. Durden, S. D. Miller, R. T. Austin, A. Benedetti, C. Mitrescu and the CloudSat Science Team, 2002: The CloudSat Mission and the A-Train: A new dimension of space-based observations of clouds and precipitation. *Bull. Amer. Meteor. Soc.*, **83**, 1771-1790.

_____, 2005: Cloud feedbacks in the climate system: a critical review. *J. Climate*, In press.

Su, H., S. S. Chen, and C. S. Bretherton, 1999: Three-dimensional week-long simulations of TOGA-COARE convective systems using the MM5 mesoscale model. *J. Atmos. Sci.*, **56**, 2326-2344.

Sui, C. H., and K.-M. Lau, 1992: Multiscale phenomena in the tropical atmosphere over the western Pacific. *Mon. Wea. Rev.*, **120**, 407-430.

_____, K.-M. Lau, W.-K. Tao, and J. Simpson, 1994: The tropical water and energy cycles in a cumulus ensemble model. Part I: Equilibrium climate. *J. Atmos. Sci.*, **51**, 711-728.

Talagrand, O., 1997: Assimilation of observations, an introduction. *J. Meteor. Soc. Japan*, 81-99.

Tamminem, J., and E. Kyrola, 2001: Bayesian solution for nonlinear and non-Gaussian inverse problems by Markov chain Monte Carlo method. *J. Geophys. Res.*, **106**, 14,377-14,390.

_____, 2004: Validation of nonlinear inverse algorithms with Markov chain Monte Carlo method. *J. Geophys. Res.*, **109**, D19303, doi:10.1029/2004JD004927.

Tao, W.-K., and Y. Ogura, 1986: A study of the response of deep tropical clouds to mesoscale processes: Three dimensional numerical experiments. *J. Atmos. Sci.*, **43**, 2653-2676.

_____, J. Simpson, and M. McCumber, 1989: Ice-water saturation adjustment. *Mon. Wea. Rev.*, **117**, 231-235.

_____, and J. Simpson, 1993: Goddard Cumulus Ensemble Model. Part I: Model Description. *Terrestrial, Atmospheric, and Oceanic Sciences*, **4**, 35-72.

_____, J. R. Scala, B. S. Ferrier, and J. Simpson, 1995: The effect of melting processes on the development of a tropical and a midlatitude squall line. *J. Atmos. Sci.*, **52**, 1934-1948.

_____, J. Simpson, C.-H. Sui, C.-L. Shie, B. Zhou, K. M. Lau, and M. Moncrieff, 1999: Equilibrium states simulated by cloud-resolving models. *J. Atmos. Sci.*, **56**, 3128-3139.

_____, J. Simpson, D. Baker, S. Braun, M.-D. Chou, B. S. Ferrier, D. Johnson,

A. Khain, S. Lang, G. Lynn, C.-L. Shie, D. Starr, C.-H. Sui, Y. Wang, and P. Wetzel, 2003a: Microphysics, radiation, and surface processes in the Goddard cumulus ensemble (GCE) model. *Meteor. Atmos. Phys.*, **82**, 97-137.

_____, C.-L. Shie, J. Simpson, S. Braun, R. H. Johnson, and P. E. Ciesielski, 2003b: Convective systems over the South China Sea: cloud-resolving model simulations. *J. Atmos. Sci.*, **60**, 2929-2956.

_____, D. Johnson, C.-L. Shie, and J. Simpson, 2004: The atmospheric energy budget and large-scale precipitation efficiency of convective simulations during TOGA COARE, GATE, SCSMEX, and ARM: cloud-resolving model simulations. *J. Atmos. Sci.*, **61**, 2405-2423.

Tarantola, A. 1987: *Inverse Problem Theory*, Elsevier, Amsterdam.

_____, 2005: *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, PA.

Tokay, A., and D. A. Short, 1996: Evidence from tropical raindrop spectra of the origin of rain from stratiform versus convective clouds. *J. Appl. Meteor.*, **35**, 355-371

Tompkins, A. M., and G. C. Craig, 1998: Radiative-convective equilibrium in a three-dimensional cloud ensemble model. *Q. J. Roy. Met. Soc.*, **124**, 2073-2097.

_____, 2000: The impact of dimensionality on long-term cloud-resolving model simulations. *Mon. Wea. Rev.*, **128**, 1521-1535.

_____, 2001a: Organization of tropical convection in low vertical wind shears: the role of water vapor. *J. Atmos. Sci.*, **58**, 529-545.

_____, 2001b: On the relationship between tropical convection and sea surface temperature. *J. Climate*, **14**, 633-637.

van den Heever, S. C., G. G. Carrio, W. R. Cotton, P. J. DeMott, and A. J. Prenni, 2005: Impacts of nucleating aerosol on Florida storms. Part I: Mesoscale simulations. *J. Atmos. Sci.*, In Press.

Verlinde, J., P. J. Flatau, and W. R. Cotton, 1990: Analytical solutions to the collection growth equation: Comparison with approximate methods and application to cloud microphysics parameterization schemes. *J. Atmos. Sci.*, **47**, 2871-2880.

Walko, R. L., W. R. Cotton, M. P. Meyers, and J. Y. Harrington, 1995: New RAMS cloud microphysics parameterization. Part I: the single-moment scheme. *Atmos. Res.*, **38**, 29-62.

Weisman, M. L., and J. B. Klemp, 1982: The dependence of numerically simulated convective storms on vertical wind shear and buoyancy. *Mon. Wea. Rev.*, **110**, 504-520.

_____, W. C. Skamarock, and J. B. Klemp, 1997: The resolution dependence of explicitly modeled convective systems. *Mon. Wea. Rev.*, **125**, 527-548.

Webster, P. J., and R. Lukas, 1992: TOGA COARE: The Coupled Ocean Atmosphere Response Experiment. *Bull. Amer. Meteor. Soc.*, **73**, 1377-1416.

Wilcox, E. M., 2003: Spatial and temporal scales of precipitating tropical cloud systems in satellite imagery and the NCAR CCM3. *J. Climate*, **16**, 3545-3559.

Wisner, C. H., H. D. Orville, and C. Myers, 1972: A numerical model of a hail-bearing cloud. *J. Atmos. Sci.*, **29**, 1160-1181.

Wu, X., and M. W. Moncrieff, 1999: Effects of sea surface temperature and large-scale dynamics on the thermodynamic equilibrium state and convection over the tropical western Pacific. *J. Geophys. Res.*, **104**, 6093-6100.

_____, W. D. Hall, W. W. Grabowski, M. W. Moncrieff, W. D. Collins, and J. T. Kiehl, 1999: Long-term behavior of cloud systems in TOGA COARE and their interactions with radiative and surface processes. Part II: Effects of ice microphysics on cloud-radiation interaction. *J. Atmos. Sci.*, **56**, 3177-3195.

Wyant, M., C., M. Khairoutdinov, and C. S. Bretherton, 2006: Climate sensitivity and cloud response of a GCM with a superparameterization. *Geophys. Res. Lett.*, **33**, L06714, doi:10.1029/2005GL02546.

List of Abbreviations

BOA	Bottom of Atmosphere
CALIPSO	Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observation
CAPE	Convective Available Potential Energy
CRCP	Cloud Resolving Convective Parameterization
CRM	Cloud Resolving Model
ECMWF	European Centre for Medium -Range Weather Forecasts
EOS	Earth Observing System
GCE	Goddard Cumulus Ensemble
GCM	General Circulation Model
GCSS	GEWEX Cloud Systems Study
GEWEX	Global Energy and Water Cycle Experiment
GPM	Global Precipitation Measurement
IWP	Ice Water Path
LHR	Latent Heat Release
LW	Longwave
LWP	Liquid Water Path
MCMC	Markov chain Monte Carlo
MMF	Multiscale Modeling Framework
MODIS	Moderate Resolution Imaging Spectroradiometer

NASA	National Aeronautics and Space Administration
NESA	Northern Enhanced Sounding Array
OLR	Outgoing Longwave Radiation
PDF	Probability Density Function
PE	Precipitation Efficiency
PSD	Particle Size Distribution
PWV	Precipitable Water Vapor
RCE	Radiative Convective Equilibrium
SCSMEX	South China Sea Monsoon Experiment
SIRICE	Submillimeter Infrared Ice Cloud Experiment
SST	Sea Surface Temperature
SW	Shortwave
TMI	TRMM Microwave Imager
TOA	Top of Atmosphere
TOGA	Tropical Ocean Global Atmosphere Coupled Ocean
COARE	Atmosphere Response Experiment
TRMM	Tropical Rainfall Measuring Mission
TWP	Tropical Warm Pool
VIRS	Visible and Infrared Sounder