Michigan Technological University
Digital Commons @ Michigan Tech

Dissertations, Master's Theses and Master's Reports - Open

Dissertations, Master's Theses and Master's Reports

2007

In-situ measurement and characterization of cloud particles using digital in-line holography

Jacob P. Fugal
Michigan Technological University

Copyright 2007 Jacob P. Fugal

Recommended Citation

Fugal, Jacob P., "In-situ measurement and characterization of cloud particles using digital in-line holography", Dissertation, Michigan Technological University, 2007.
http://digitalcommons.mtu.edu/etds/104

Follow this and additional works at: http://digitalcommons.mtu.edu/etds

Part of the Physics Commons

IN-SITU MEASUREMENT AND CHARACTERIZATION OF CLOUD PARTICLES USING DIGITAL IN-LINE HOLOGRAPHY

By

JACOB P. FUGAL

A DISSERTATION

Submitted in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

(Engineering Physics)

MICHIGAN TECHNOLOGICAL UNIVERSITY

2007

© 2007 Jacob P. Fugal

This dissertation, "In-situ Measurement and Characterization of Cloud Particles using Digital In-line Holography", is hereby approved in partial fulfillment of the requirements for the degree of DOCTOR OF PHILOSOPHY in the field of Engineering Physics.

DEPARTMENT or PROGRAM:
Engineering Physics

Signatures:

Dissertation Advisor ______________________
Dr. Raymond A. Shaw

Committee ______________________
Dr. Jacek Borysow

____
Dr. Will H. Cantrell

____
Dr. Timothy J. Schulz

Department Chair ______________________
Dr. Ravindra Pandey

Date ______________________

Abstract

Satellite measurement validations, climate models, atmospheric radiative transfer models and cloud models, all depend on accurate measurements of cloud particle size distributions, number densities, spatial distributions, and other parameters relevant to cloud microphysical processes. And many airborne instruments designed to measure size distributions and concentrations of cloud particles have large uncertainties in measuring number densities and size distributions of small ice crystals. HOLODEC (Holographic Detector for Clouds) is a new instrument that does not have many of these uncertainties and makes possible measurements that other probes have never made. The advantages of HOLODEC are inherent to the holographic method.

In this dissertation, I describe HOLODEC, its in-situ measurements of cloud particles, and the results of its test flights. I present a hologram reconstruction algorithm that has a sample spacing that does not vary with reconstruction distance. This reconstruction algorithm accurately reconstructs the field to all distances inside a typical holographic measurement volume as proven by comparison with analytical solutions to the Huygens-Fresnel diffraction integral. It is fast to compute, and has diffraction limited resolution. Further, described herein is an algorithm that can find the position along the optical axis of small particles as well as large complex-shaped particles. I explain an implementation of these algorithms that is an efficient, robust, automated program that allows us to process holograms on a computer cluster in a reasonable time.

I show size distributions and number densities of cloud particles, and show that they are within the uncertainty of independent measurements made with another measurement method. The feasibility of another cloud particle instrument that has advantages over new standard instruments is proven. These advantages include a unique ability to detect shattered particles using three-dimensional positions, and a sample volume size that does not vary with particle size or airspeed. It also is able to yield two-dimensional particle profiles using the same measurements.

Acknowledgments

Oh, where to start to acknowledge all the people that helped me accomplish this. I have a great debt to my advisor, Dr. Raymond Shaw for professional and personal support above and beyond what any advisor (as I see it) would be expected to give. He's a wise and wonderful friend, advisor, teacher and counselor. He tells me that it is an advisor's job to push their students to excellence. His prodding, advice and leadership have always been in my best interest even when it was a sacrifice to him. No graduate student could ask for a better advisor.

I'm grateful to current and former members of my committee. Dr. Will Cantrell has always had something interesting to talk about. Dr. Alex Kostinski has given me breathtaking insight to fundamental problems that I formerly considered solved. I've seen Dr. Jacek Borysow ask poignant, relevant questions even when he isn't familiar with the method. Both Drs. Kostinski and Borysow encourage me to not be shy to look with a critical eye at scientific work and see if the fundamental assumptions and physics are correct. Dr. Tim Schulz has been my other advisor in the fundamentals of light propagation and hologram reconstruction. This work is founded on what he taught me.

I'm grateful to all who taught my classes: Dr. Don Beck, Dr. Michael Worthier, Dr. Raymond Shaw, Dr. Alex Kostinski, Dr. Will Cantrell, Dr. Tim Schulz, Dr. Song-Lin Yang, Dr. Reza Zekavat, Dr. Jeff Burl, and Dr. Ed Nadgorny. Many of these teachers put up with me being gone for research trips or an extended family emergency. Dr. Ravi Pandey, the Physics Department Chair, advised me to have the attitude to just focus on my research and he'd take care of the finances and he lived up to it. I'm grateful to a slough of grad students who helped me and stood by as friends to listen and come in to the lab when I needed help. They include Mike Larsen, Ewe Wei Saw, Eli Ochshorn, Jiang Liu, and Hansen Nordsiek who practically is a grad student. Gabe Johnson was a big help in the lab, and Austin Kuzmik has made life easier by taking over sysadmin responsibilities in our lab.

This work has been supported by an NSF Graduate Research Fellowship (2003-2007), a NASA Earth System Science Fellowship (2006-2007), two Michigan Space Grant Consortium Graduate Fellowships (2004-2006), a Miles Fellowship (2002-2003), a Thornton Fellowship (2002-2006), and NSF grants ATM99-84294 and ATM-0535488 and NASA grant NAG5-10568.

I'm always grateful for the unfailing and huge support of my wife Jenifer and patience of my kids Josiah and Naomi. Finally, I'm grateful for a loving Heavenly Father and his divine son Jesus Christ who has helped me make it through some really tough times to get to this point. Thank you all!

Dedication

To my dear parents, my beautiful wife Jenifer, and my children (so far) Josiah and Naomi. And also to all those who find clouds and water as beautiful and interesting as I do.

Contents

Abstract **v**

Acknowledgements **vii**

1 Introduction **1**

1.1 Introduction 1

1.2 Organization of Dissertation 2

2 Airborne Digital Holographic System for Cloud Particle Measurements **5**

2.1 Introduction 5

2.2 In-line holography 7

2.2.1 Approximate far-field model 8

2.2.2 Implications of Far-Field Model 9

2.3 Instrument overview and electro-optical design 10

2.3.1 Electro-optical system 10

2.3.2 Data system 12

2.3.3 Temperature control system 13

2.3.4 Mechanical design 13

2.4 Holograms and digital reconstruction 14

2.4.1 Reconstruction Method 14

2.4.2 Calibration 15
2.4.3 Cloud Droplets and Ice Crystals 16
2.5 Summary and future work 17

3 Practical Methods for Reconstruction and Characterization of Particles in Digital Inline Holograms 23

3.1 Introduction 24
3.2 Background division by median of neighboring holograms 25
3.3 Reconstruction method 27
3.3.1 Huygens-Fresnel kernels 28
3.3.2 Distances at which the kernels are undersampled 29
3.3.3 Filtering out undersampled frequencies in the Huygens-Fresnel filtering kernel 31
3.3.4 Testing the Huygens-Fresnel filtering kernel with a low-pass filter 32
3.4 Particle Size Independent Sample Volume 34
3.5 Find the Particle Position on the Optical Axis 36
3.5.1 Focusing Methods in the Literature 36
3.5.2 Detection by Edge Sharpness 37
3.6 Profiles by Local and Global Standard Deviations 38
3.7 Reconstructing on a Computer Cluster 43
3.7.1 Algorithm Implementation within Memory Limits 43
3.7.2 Parallel Algorithms vs. Embarrassingly Parallel Algorithms . . 44
3.7.3 Hologram Reconstruction and Memory-Bound Computation . 45
3.8 Summary and Future Work 45

4 Ice Particle Size Distributions Measured with an Airborne Digital In-line Holographic Instrument 47

4.1 Introduction 48

4.2 Holographic Reconstruction Method and Instrument Parameters . . . 50

4.2.1 Hologram Reconstruction and Particle Finding Method 50

4.2.2 Instrument Parameters 54

4.3 Data Sample 55

4.3.1 Cloud Conditions in Data Sample 55

4.3.2 Corrections to FSSP Sizing of Ice Particles 55

4.4 Results 58

4.5 Discussion and Summary 67

5 Summary and Outlook 69

5.1 Summary 69

A MATLAB Hologram Reconstruction Codes 73

1.1 HFFFTPrep.m 73

1.2 HFFFTPropagate.m 76

Bibliography 81

List of Figures

2.1 (a) A photograph of the HOLODEC instrument mounted on the right wing pod of the NCAR/NSF C-130. This photograph gives a sense of the measurement environment and geometry, and the resulting design constraints. (b) A schematic of the electro-optical system of HOLODEC. Cloud particles flow to the left relative to the instrument, at a true air speed of approximately 100 m s^{-1}. 11

2.2 Frame (a) shows a whole reconstruction slice at 58 mm from the CCD. The droplet generator tip is shown at the top of the frame. A 50 μm diameter droplet is visible about two thirds down the frame below the tip. Frame (b) shows a closeup of the 50 μm diameter drop and (c) shows a 25 μ diameter drop. Note that the reconstructed images of the droplets are within a pixel (4.65 μm) of the known diameter. Help with the calibration droplets was graciously provided by U. Maixner and D. Nagel from GKSS Forschungszentrum. 16

2.3 Images from a small section of a hologram obtained during flight traverse of a cloud. Frames (a) through (d) show digitally reconstructed images in 500 μm steps along the optical axis, as indicated above each image. Each frame shows cloud droplets that come in and out of 'focus' as the frames advance along the optical axis. Units on each axis are in mm. Other instruments on the aircraft indicated the largest drops in this cloud are approximately 10 to 20 μm in diameter. 18

2.4 Various reconstructions of large ice crystals and one of the original holograms. The distance from the CCD is indicated above each frame. The units on each axis are mm. Frame (a) shows the section of the hologram from which the ice crystal in frame (b) was reconstructed. Frames (c), (d), and (e) are other reconstructions of large ice crystals. 19

2.5 Various reconstructions of small ice crystals and one of the original holograms. The distance from the CCD is indicated above each frame and the units on each axis are mm. Frame (a) shows the section of the hologram from which the ice crystal in frame (b) was reconstructed. Frames (c), and (d) are other reconstructions of small ice crystals. . . 20

3.1 Shows a hologram divided by the median of its neighboring 7 holograms (including itself). The top left panel is a subarea of the original hologram, which is reconstructed to 43.5 mm in the top center panel. The top right panel shows a closeup of two bright dots which are 20 μm cloud droplets taken on the 2003-09-05 Research Flight in the IDEAS-3 project (Fugal et al., 2004, see Chapter 2). The center row shows the background image and the bottom row is the background divided hologram. Note the presence of water drops stuck to the windows of HOLODEC in the original and background hologram rows and their absence in the bottom row. Note the presence of the two bright dots in the top right panel, their absence in the center panel, and their presence in a relatively noise free background in the bottom right panels. All axes are in mm, and each image is equally contrast enhanced for visibility in printing. 26

3.2 An example of one of the 1394 tests done for the Huygens-Fresnel Filtering propagation method as discussed in the text. The top left panel shows a 40x48 Δx square and 40 Δx diameter circle that are reconstructed to 10 d_l using analytical solutions in the top row, without the low pass filter in the center row, and with the low pass filter in the bottom row. All axes are in units of Δx and the intensity is in arbitrary units. Note that the color scales are not all the same. Note in the error panels that the low pass filtered error in the bottom left panel has about half the error and a much more reasonable noise pattern than the unfiltered field in the center left panel. 33

3.3 Shows a 2 Δx particle (top row) and a 10 Δx particle (bottom row) reconstructed to distances shown above each panel without a low pass filter. The cross section frames on the right show reconstructed intensity on a line through the center of each particle. Units on the axes are in Δx. Notice that the particle reconstruction quality falls greatly beyond the diffraction limit in the third from the left frames. 35

3.4 Same as Figure 3.3 except with the low pass filter described in the text. 36

3.5 Shown are profiles as described in the text for two small ice crystals. All image axes are in units of mm. In the top panel, the intensity and the Sobel gradient profiles are approximately equal as is usually the case for small particles. In the bottom panel, the intensity and Sobel gradient profiles are not equal. The Sobel gradient is three-peaked, but the highest peak is selected and in this case is the correct peak. Note also that the global profile has its highest peak much higher than the other peaks where the local profile does not have such a contrast. . . 40

3.6 A continuation of Figure 3.5. Note again that the Sobel gradient and Intensity profiles do not peak at the same place for these larger ice crystals. In the top panel, the three ice crystals are focused at a compromise location where the three are most in focus together though not the optimal focus for the 3 individual crystals. The bottom panel is evidence that the global profile method works for the largest of crystals. 41

3.7 A continuation of Figures 3.5 and 3.6. Note in the top panel, the ice crystal is in focus because the global profile peaks higher at the focus position at about 40 mm and not at the other peak at about 37 mm. The inverse is true of the local profile. In the bottom panel is an ice crystal that is not in focus. The global profile peaks at about 46 mm where this image is focused. Visually, the focus is at 54.5 mm as shown in Figure 2.4 in Fugal et al. (2004, see Chapter 2) which corresponds to the peak in the local profile in the Sobel gradient profile plot. The largest reason for error is the particle lies on the edge of the hologram and so the information to reconstruct the particle is incomplete. . . . 42

4.1 The top panel shows a hologram from Research Flight 2003-09-17 reconstructed to 66.5 mm. The reconstructed image has many particles in focus at one depth and many others nearly in focus. The bottom panel is a histogram of the z-positions of all particles in the sample volume. This is suggestive of shards of a single impacted crystal tracing an aerodynamic surface about the instrument housing. 52

4.2 The top panel shows an ice particle in the process of shattering lying about 10 mm from one of HOLODEC's windows. Shown in the center and bottom panels are histograms of optical-axis positions of particles from Research Flight 2003-09-17. The center panel shows all particles and the bottom panel excludes particles in holograms containing shattered particles. Note that rejecting shattered particles yields a more uniform distribution of particles along the optical axis. 53

4.3 Particle number densities from the FSSP and the 2D-C probes from the 2003-09-05 and 2003-09-17 Research Flights. Also shown are the times where HOLODEC recorded clear holograms. The labels in brackets from (a) to (p) correspond to size distribution panels in Figures 4.6 to 4.8. 56

4.4 Total, liquid, and ice (total minus liquid) water content for the 2003-09-05 and 2003-09-17 Research Flights. The labels in brackets from (a) to (p) correspond to size distribution panels in Figures 4.6 to 4.8. 57

4.5 The top panel shows the scattered intensity integrated over the forward scattering angles measured by the FSSP. From this, a correction curve is calculated for the measured FSSP size bins as shown in the bottom panel. The black dashed line shows a one-to-one correction curve for comparison. 59

4.6 Size distributions of ice particles for four time periods shown above each frame. The left panels show the uncorrected FSSP sizes and the right panels show FSSP sizes corrected for ice particles. FSSP size distributions are shown as circles and HOLODEC size distributions are shown as pentagrams with error bars. The squares are HOLODEC size distributions excluding shattered particles. The error bars are calculated assuming Poisson statistics and are one standard deviation. The diamonds are data from the PMS 2-DC probe and are only shown on the left panels. 60

4.7 As in Figure 4.6, for four more time periods illustrating a range of cloud microphysical conditions. 61

4.8 As in Figures 4.6 and 4.7 except for Research Flight 2003-09-05 and only shown with uncorrected FSSP Sizes. No PMS 2-DC probe data are included. 62

4.9 Number densities averaged over 10 second intervals from FSSP and HOLODEC. 64

4.10 A sample of ice particles taken during Research Flight 2003-09-17. All particles are shown as reconstructed by the automated particle finding algorithm and are only scaled to improve contrast for printing. The white scale bar in the upper left is 0.5 mm in length. 65

4.11 As in Figure 10 except ice particles are a sample from Research Flight 2003-09-05. The white scale bar in the upper left is 0.5 mm in length. 66

List of Tables

5.1 Particle size ranges, sample rates and path lengths through clouds to sample 1 L for various instruments including HOLODEC II burst and continuous sample modes. The 1000 particle sample path length is estimated by assuming a typical concentration of 10^4 10 μm particles L^{-1}, 100 50 μm particles L^{-1}, and 1 100 μm particles L^{-1}. The final column gives a rough estimate of how long we might have to sample to get a good size distribution in that size range from that instrument. 70

Chapter 1

Introduction

1.1 Introduction

Satellite measurement validations, global climate models and cloud models each depend on understanding important cloud processes such as aerosol-cloud particle interactions, radiative transfer, precipitation, and surface chemistry which are particularly sensitive to accurate measurements of small ice crystals (< 200 μm extent). These sensitivities partly arise due to the overwhelming populations of small ice crystals relative to large ones (Baum et al., 2005a).

Ironically, the small ice particle size range is the most difficult for standard aircraft instrumentation to accurately measure (Baum et al., 2005b) due to several effects, including ice crystals shattering on leading probe parts and then being counted among the pristine ice crystals (Korolev and Isaac, 2005), and the sensitivity of probe sample volume size to ice crystal size (Baumgardner and Korolev, 1997). The shattering effect has remained an unsolved problem for many years and has left cloud scientists uncertain as to whether the measured large number of small ice crystals from standard aircraft probes (e.g. FSSP, and 2D-C) is real or an instrument artifact (Field et al., 2003a; Gardiner and Hallett, 1985).

There are instruments that have been and are being developed to overcome these and other limitations. For example, the Small Ice Detector (SID) can measure the size and morphology of small ice crystals, and the first Holographic Detector for Clouds (HOLODEC, Fugal et al. (2004), see Chapter 2) can measure the size, 2-D profile, and 3-D position of small ice crystals, detect and reject shattered ice crystals and has no particle-size dependency in its sample volume size.

1.2 Organization of Dissertation

In this dissertation, I describe the HOLODEC instrument and its test flights aboard the NCAR C-130 Hercules Q aircraft. Also, I describe algorithms to reconstruct holograms, find both large complex-shaped particles and small particles, measure their three-dimensional position, size and two-dimensional profiles. I explain a method by which holography can detect and reject shattered particles. Finally, I show calculations of size distributions and number densities and show the two-dimensional profiles of these particles and compare them with measurements of other standard instruments.

This dissertation is organized into 3 separate papers adapted and included as Chapters 2, 3, and 4. Chapter 2 (Fugal et al., 2004) is a paper published in Applied Optics that describes the mechanical, electrical, optical, and data-acquisition systems of the first HOLODEC probe. It describes the test flights and dataset we have and illustrates a few first results. Chapter 3 (Fugal et al., 2007) is a paper to be submitted to Measurement Science and Technology that describes the algorithms to reconstruct holograms, find particles, and make measurements on those particles. It also describes how these algorithms are implemented efficiently on a computer cluster. Chapter 4 (Fugal and Shaw, 2007) is a paper already submitted to the Journal of Atmospheric and Oceanic Technology describing accurate size distributions and number densities of ice crystals of a large dynamic size range along with two-dimensional profiles or images of these particles. Because chapters 2, 3 and 4 are written as separate papers, they each have their own abstract, introduction, and conclusions, and are more separate as documents and therefore they overlap somewhat in their coverage. They refer to each other by paper reference and chapter numbers. Chapter 5 is a summary and outlook on the proposed future of the HOLODEC instrument and its associated software to measure cloud microphysical properties that no one has measured before. Appendix A is a printout of the well commented MATLAB code I use to reconstruct holograms.

As a final word in this introduction, it is fair to note how significant this work is to the advancement of digital particle holography. As stated in Chapter 3, there is no dearth of papers on how to refine some aspect of holographic reconstruction, or particle finding, or precision in calculating particle position (Fugal et al., 2007). Almost all of these papers include some physical example with empirical evidence of their new advance in holographic particle measurement. However, there are few papers that report results of experiments where holography is the tool, not the subject. Malkiel et al.'s work of file holograms, digital scanning, and automated particle characterization to study plankton is one of few examples (Malkiel et al., 2003, 2004, 2006). To my knowledge, this is the first work where digital holography is used as a tool to get some scientific insight and where thousands of holograms were processed to obtain it. In essence, this work takes digital holography into production. This work has publishable advances in hologram reconstruction and holographic particle measurements as well, but this work is also a milestone of holography going from a specialized method under development to a practical tool for solving scientific problems.

Chapter 2

Airborne Digital Holographic System for Cloud Particle Measurements

An in-line holographic system for detection of atmospheric cloud particles (Holographic Detector for Clouds, HOLODEC) has been developed and flown on the NCAR C-130 research aircraft*. Clear holograms are obtained in daylight conditions at typical aircraft speeds of 100 m s^{-1}. The instrument is fully digital and is interfaced to a control and data acquisition system in the aircraft via optical fiber. It is operable at temperatures below -30^o C and at typical cloud humidities. Preliminary data from the experiment show its utility for studies of the three-dimensional spatial distribution of cloud particles and ice crystal shapes.

2.1 Introduction

In models describing radiative transfer through clouds, remote sensing of cloud properties, and precipitation formation, it is usually assumed that the cloud particles are distributed with perfect randomness on small (sub-meter) scales (Shaw, 2003). That is, their statistics are characterized by the Poisson distribution on all scales: the number of particles $N(V)$ in a test volume V is a random variable, and is distributed according to:

$$p(N) = \frac{N^{\bar{N}} e^{-N}}{N!}, \tag{2.1}$$

where $\bar{N} = N(\bar{V})$ is the mean number of particles in V.

---

*Adapted from a paper published in November 2004, Applied Optics, 43, pages 5987-5995 by Jacob P. Fugal, Raymond A. Shaw, Ewe Wei Saw, and Aleksandr V. Sergeyev. ©2004 Optical Society of America

However, it has been observed that cloud particles are spatially correlated on submeter scales (Kostinski and Shaw, 2001; Shaw, 2003). The scale-dependence of the spatial correlation can be quantified as a function of scale by the pair correlation function:

$$\eta(r) = \frac{N(r_o)\bar{N}(r_o + r)}{\bar{N}^2} - 1, \tag{2.2}$$

where r is the scale of interest. While there has been some laboratory and ground-based work on three-dimensional spatial correlations (Kozikowska et al., 1984; Uhlig et al., 1998), still nearly all in-situ measurements of these correlations have been restricted to one- or two-dimensional projections of particle positions (e.g. particles detected along the length of a long, narrow tube swept out by an aircraft instrument, or images of particles illuminated by a thin sheet of light). These measurements yield a one- or two-dimensional pair correlation function, with particle positions in the unmeasured dimensions simply projected. With some assumptions it is possible to generalize a one- or two-dimensional pair correlation function to a three-dimensional (3D) one. However this generalization has restrictions at the smallest scales which are of the greatest interest (Holtzer and Collins, 2002). To directly estimate the 3D pair correlation function at small scales, we need the three spatial coordinates of all particles in the sampled volume. Holography is a robust way to obtain these 3D particle positions.

While the most advanced holographic systems use an off-axis configuration (which has low-noise advantages and higher depth precision) and silver-halide holographic film with resolutions of order 0.1 μm, these systems also require film processing and laborious optical reconstruction or high resolution film scanning and numerical reconstruction (Meng et al., 2004). Digital holographic systems require no film processing and allow numerical reconstruction which is much more flexible than optical reconstruction. Because digital camera pixel widths are around 5 μm they do not have the spatial resolution required for off-axis holography which, in practice, limits digital holography to an in-line configuration (Hinsch, 2002). (Note that hybrid optical arrangements exist that have some advantages of off-axis holography yet can be implemented digitally (Meng et al., 2004).) In-line (Gabor) holography also has the advantage of a simple optical setup which is advantageous in adverse environments like ours: an airborne platform flying through an atmospheric cloud.

In this paper we describe an airborne digital in-line holographic detector, called Holographic Detector for Clouds (HOLODEC), and demonstrate its ability to make in-situ measurements of the three-dimensional coordinates of cloud particles. Furthermore, we show the ability to detect ice crystals and reconstruct their shapes. We first (§2) give an overview of in-line holography and describe how it allows us to measure the three-dimensional coordinates of cloud particles. We then give an overview of the electro-optical, data-acquisition, temperature control system, and the mechanical design of the instrument in §3. Finally, in §4 we discuss the calibration and baseline optical detection limits of the instrument, and show examples of holograms and digitally-reconstructed images from recent flight tests.

2.2 In-line holography

In-line holography experienced some of its earliest development as a result of its application to measurement of atmospheric particles (Thompson, 1974). Since then in-line holography has been applied to various problems of atmospheric (Kozikowska et al., 1984; Uhlig et al., 1998; Borrmann and Jaenicke, 1993; Lawson and Cormack, 1995) and of broader scientific and engineering interest (Meng et al., 1993; Owen and Zozulya, 2000; Pan and Meng, 2003; Xu et al., 2002). The scientific literature covering in-line holography and its applications is far too large to review here, but it suffices to say that the technique is well developed, broadly applied, and that it remains an area of active research. Our application is to measure atmospheric cloud particles from an airborne platform. Of primary interest is the spatial distribution of liquid cloud droplets, but the technique is also attractive for measuring shapes and concentrations of ice crystals and the concentrations of large, rare drizzle drops.

Simply stated, an in-line hologram is an interference pattern resulting from the superposition of an incident plane wave E_R (the reference wave) and light scattered by the dilute suspension of illuminated particles, E_S. Thus, the electric field at the measurement plane is $E_H = E_R + E_S$. In practice we measure intensity or the modulus squared of the superimposed waves yielding:

$$\begin{aligned} \langle I_H(x,y)\rangle &= \langle E_H E_H^*\rangle && (2.3)\\ &= \langle E_R E_R^* + E_R E_S^* + E_S E_R^* + E_S E_S^*\rangle && (2.4)\end{aligned}$$

where the angle brackets $\langle\Box\rangle$ represent a time average over many periods. The first term on the second line is the constant background, which can be filtered out before the numeric reconstruction (Pan and Meng, 2003). The second and third terms are the real and virtual images respectively. The last term is negligible for small cloud particles ($<$ 100 μm) because their scattered waves are much smaller in amplitude than the incident wave.

We are interested in reconstructing either the real or the virtual image of the droplets in the hologram. It turns out that in the Fraunhofer regime, we can reconstruct the real image of a cloud particle and the virtual image is present but weak. The Fraunhofer regime of 100 μm particles and our wavelength of 0.527 μm is $d^2/\lambda = 19$ mm. Since the camera is 35 mm away from the sample volume, we are well within this regime and thus can neglect the virtual image. After the background term is filtered out, the dominant noise is caused by the defocused virtual image. This noise imposes a limit on the number density of particles and depth of the sample volume (Meng et al., 1993). However, number densities of cloud particles with $d > 10$ μm rarely exceed a few hundred droplets per cubic cm (< 1 droplet mm^{-3}) (Pruppacher and Klett, 1997, pp. 15–24). This number density, combined with the sample volume depth of 4.5 cm, yields an area density of order 10^3 droplets cm^{-2} which is a factor of 10 less than typical area densities in other applications of particle field digital holography (Meng et al., 2004). This number density is also sufficiently small and the scattered waves sufficiently weak that multiple scattering is negligible.

2.2.1 Approximate far-field model

For the purposes of providing physical insight and in guiding the instrument design we desire an analytical model for the hologram resulting from a single water droplet. The holograms in our experiment display the interference of a reference beam and a wave scattered by a transparent, order 10 to 100 μm diameter, spherical water droplet. This would suggest a complete solution using Mie theory to describe the electric field due to scattering from a sphere and its interference with the incident plane wave. However, we note that the particle size and scattering geometry allow for several useful approximations. Because size parameters are large ($\pi d/\lambda$) and we observe only forward-scattered light (scattering angle $< 3^o$), to good approximation we may neglect the complexities of Mie theory and treat the scattered wave as diffraction from an opaque disk with the same diameter as the water droplet (Slimani et al., 1984; Bohren and Huffman, 1983, Ch. 4, pp. 107–111). And because we record the hologram in the far field ($z \gg d^2/\lambda$ 2 to 20 mm) we may treat the scattered wave with the Fraunhofer approximation (Born and Wolf, 1999, pp. 412–516). These approximations are invaluable in providing convenient analytical results to guide our understanding of the essential physics underlying the technique and our design of the instrument. The eventual digital reconstruction of the holograms is carried out using a less restrictive model (see §4) because actual conditions do not always satisfy the far-field constraint (for example, ice particles larger than 100 μm).

To develop the analytical model we consider an object at $z = 0$, where the z-coordinate is taken to be the optical axis. Coordinates in the object plane, perpendicular to the optical axis, are (ξ, η) and coordinates in the (far field) diffraction plane, also perpendicular to the optical axis, are (x, y). Beginning with the Huygens-Fresnel principle and making appropriate simplifications, including the neglect of quadratic terms of order $(\xi^2 + \eta^2)/(\lambda z)$, the Fraunhofer approximation is obtained. Taking the aperture function to be a disk of diameter d centered on the optical axis, an analytical expression for the total electric field E_H can be obtained for the far-field limit $z \gg d^2/\lambda$. Defining $r = (x^2 + y^2)^{1/2}$, $C = \pi d^2/4\lambda z$, $Q(r) = 2J_1(\pi r d/\lambda z)/(\pi r d/\lambda z)$, and $\Phi(r) = \pi r^2/\lambda z$, the measured intensity I_H has the form (Meng et al., 1993; Tyler and Thompson, 1976)

$$I_H = 1 - 2CQ\sin\Phi + C^2Q^2 \tag{2.5}$$

The first term is the background intensity and C^2Q^2 is the negligible scattered intensity (diffraction) term. In the Fraunhofer limit, therefore, the hologram obtained from a population of cloud droplets may be approximated as the superposition of the fields leading to Eq. 2.5, one for each particle, with r and d adjusted appropriately for droplet position and size, respectively. In practice the cloud of particles is sufficiently dilute that interference of waves from various particles can be neglected.

2.2.2 Implications of Far-Field Model

Equation 2.5 demonstrates several properties that make in-line holography ideal for studies of the spatial distribution of atmospheric particles. First, the interference term $\Phi(r)$ depends only on the position of the particle along the optical axis, not on its diameter d. Hence, the spatial frequencies in this term alone contain sufficient information to provide the particle's position along the optical axis (the position in the (x, y) plane is easily determined). Also, the spatial frequency increases radially as $r/\lambda z$ so that the desired depth of field of the instrument places a constraint on the spatial resolution of the detector. Note also that the increasing spatial frequency with r suggests that the finite pixel size limits the maximum sharpness attainable in reconstructed images. Both of these conclusions can be obtained by considering in-line holography for a point particle, but the disk aperture model makes it clear that the interference fringe pattern described by the $\sin\Phi(r)$ term contains information on particle position, while the modulation of this pattern by the term $2CQ(r)$, depends on both z and d, as expected from common experience with diffraction by a circular aperture.

Assuming that there is a minimum fringe visibility that can be reconstructed or detected, the Fraunhofer model also illustrates how detection limits vary with particle size and position along the optical axis. Fringe visibility is

$$V = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} \tag{2.6}$$

, where we consider the various maxima and minima in $I_H(r)$. Because of the far field condition $C \ll 1$,and we may neglect the diffraction term C^2Q^2 relative to the interference term. Therefore,

$$V = CQ = \left| \frac{d}{2r} J_1 \left(\frac{\pi r d}{\lambda z} \right) \right| \tag{2.7}$$

The maximum visibility will be that of the first fringe and therefore may be estimated by taking the small r limit of J_1: $V(r \to 0) = \pi d^2/8\lambda z$. Thus the visibility is maximized for small z and large d, with the constraint that the Fraunhofer condition must still be satisfied. We may expect, at least approximately, that the minimum detectable particle size will increase with position along the optical axis. In practice, the fringe visibility of detectable particles can be remarkably low because the interference patterns are coherent, whereas noise sources such as shot and speckle noise are spatially uncorrelated in the hologram.

2.3 Instrument overview and electro-optical design

The instrument consists of three main systems, (1) an electro-optical system, (2) a temperature control system, and (3) a data transmission and storage system. In addition to these three, the mechanical design is a major aspect of any airborne instrument; it will be briefly discussed here, but details of air flow distortion and particle sampling will be treated in more detail elsewhere.

The electro-optical system consisted of a laser and camera optimized for the optical measurement, but still meeting requirements for compactness, power, and stability. The temperature control system consisted of thermocouples and actively-controlled heaters. The data transmission system consisted of a Firewire (IEEE 1394) to optical fiber to Firewire conversion, and associated computer hardware and software for data storage.

Operating conditions for the instrument are severe and place significant constraints on the system design. The airborne holographic instrument has to fit inside a 61-cm long, 15-cm diameter tube attached to the wing of the NCAR C-130 research aircraft, as shown in Figure 2.1 (a). It must be powered by less than 10 A of 120 VAC switched from the cabin and less than 500 W of 28 VDC (anti-ice heater power) switched on when the aircraft is in flight. It has to operate in a temperature range spanning -30^o C at 7900 m altitude to 40^o C on the runway, air speeds of order 100 m s^{-1}, ambient relative humidities above 100%, and pressures down to 0.4 atmospheres (4 x 10^4 Pa). The only communication link between the instrument and the in-cabin controls are 12 twisted pair wires and 4 fiber-optic lines. Finally, the instrument is subject to severe vibration because of its proximity to turboprop engines. Inside the cabin, the instrument controls are in a more forgiving environment (close to room temperature, pressure, and humidity) but still experience significant vibration from the engines and the often turbulent flight conditions.

2.3.1 Electro-optical system

In-line holography does not require large laser intensities or coherence lengths. The HOLODEC instrument uses a Q-switched, frequency-doubled, pulsed Nd:YLF laser (CrystaLaser) which met the optical requirements as well as the critical specifications of no active thermal dissipation (no fans or water cooling), and compactness. The laser operates in TEM00 mode with a wavelength of 527 nm, a pulse width of 20 ns, and a pulse energy of 30 μJ. The pulse width is sufficiently short that a particle moving through the sample volume at 100 m s^{-1} will only blur by approximately 0.2 μm during the pulse. The laser pulse energy is large enough that the background sunlight is negligible within the exposure time of the camera.

One sensitivity of the laser is that the operating temperature of the laser head must

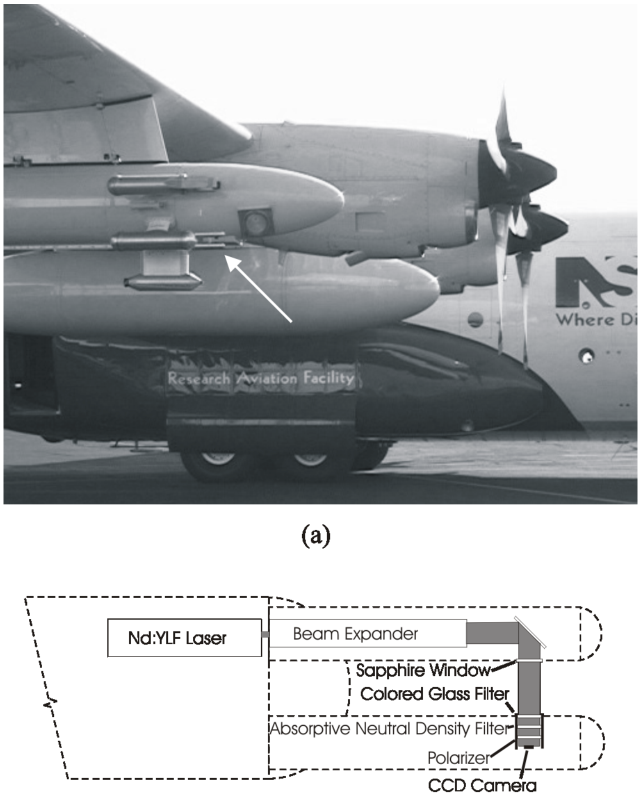

(a)

(b)

Figure 2.1: (a) A photograph of the HOLODEC instrument mounted on the right wing pod of the NCAR/NSF C-130. This photograph gives a sense of the measurement environment and geometry, and the resulting design constraints. (b) A schematic of the electro-optical system of HOLODEC. Cloud particles flow to the left relative to the instrument, at a true air speed of approximately 100 m s^{-1}.

be between 10^o C and 30^o C for stability. This is accomplished by powering down the laser while it is too hot (such as on the runway preparing for takeoff), and heating its aluminum mounting plate while in flight. The heaters are specially positioned and actively controlled in two channels to create a nearly isothermal environment for the laser head.

The detector is a progressive-scan, 1024 x 768, 4.65 μm square pixel, 10-bit (although we solely used the 8-bit mode), monochrome, Firewire CCD camera (Dragonfly, Point Grey Research). It has a maximum frame rate of 15 Hz and a minimum exposure time of approximately 200 μs, which is sufficiently short to keep the background daylight negligible relative to the laser light. Cameras with higher frame-rate, higher resolution, shorter exposure-time were considered, but the unique feature of this camera is the ability to separate the CCD from the camera board by up to 30 cm ("extended head"). Thereby the CCD detector is positioned in an arm of the probe, eliminating the need for relay optics to send the hologram back into the instrument body. This results in great simplification in the optical design, alignment process, and complexity of data reduction.

A significant design challenge was to allow the entire optical system (laser, lenses, mirrors, etc.) to be adjustable, yet remain secure during flight and the associated vibrations. As a result, the optical design, which is shown schematically in Figure 2.1 (b), is necessarily straightforward and minimizes the number of optical elements and the need for precise alignment and adjustment. Between the laser and camera the beam is expanded in a Galilean telescope designed to minimize spherical aberration. The beam expander is composed of a -6 mm, 6 mm dia. focal length planoconcave lens, and a 200 mm focal length, 25.4 mm dia. planoconvex lens which expands the beams about 30 times from less than 1 mm up to about 20 mm. The expanded beam is large enough to cover the diagonal of the CCD with an approximately flat profile. After the steering mirror, the beam passes through a sapphire window, which has a high thermal conductivity (relative to BK7 glass) to aid in the elimination of condensation. The beam then passes through three filters, all of which are non-reflective to minimize unwanted interference fringes. The first is a bandpass (colored glass) filter that cuts out much of the background light at unwanted wavelengths and serves as a window. The second is an absorptive neutral density filter of sufficient optical density so as to bring the laser intensity near the saturation level of the CCD. The third is a polarizer which allows fine tuning of the laser intensity to keep it below the saturation level of the CCD. The polarizer has the added advantage of eliminating roughly half of the unpolarized background light.

2.3.2 Data system

Digitized holograms are transmitted from the camera to the data acquisition computer in the cabin of the aircraft via optical fibers installed in the wing. We used two Firewire-to-optical-fiber converters (OpticLink 1394), one in the instrument on the wing and one in the cabin, to transmit from the Firewire interface of the camera to the Firewire interface on the computer, via the optical fiber. The optical fiber carried both the images transmitted to the computer and control commands to the camera. The maximum data rate of the camera (running in 8-bit mode) is less than 12 Mbytes s^{-1}. This allowed us to stream the data to USB 2.0 external hard drives (Seagate Barracuda 120 GByte) in vibration-damped enclosures.

The data acquisition computer was a laptop (1.6 GHz Pentium M processor, 1 Gbyte RAM). We used custom C++ routines for streaming the holograms with time stamps to disk in an uncompressed format in large contiguous files. The time stamps were synchronized with the other aircraft instruments via an on-board Network Time Protocol server. After a flight the data files are converted into individual compressed images (PNG 8-bit grayscale, lossless compression).

2.3.3 Temperature control system

We used an active temperature control system to keep condensation off the windows, to keep ice off the probe for aircraft safety, and to keep the laser head at a stable operating temperature. Laboratory temperature testing in a cold box (with no wind and below 100% humidity) confirmed that the unheated instrument components (relays, power supplies, etc.) would perform adequately well below (down to -40^o C) their published temperature specifications. The temperature control system consists of four heater channels, two for the laser head, one for the camera-side arm and one for the laser-side arm. The heater channels are switched by relays in the instrument which are remotely controlled by the in-cabin computer via an RS-485 connection over one of the twisted wire pairs. Temperature sensors were read by a second RS-485 connection. The heater power was varied by pulsing the heaters at 1 Hz with variable duty cycle determined by programmable PID feedback control software (Labview 7.0). Finally, the control software also allows individual components (laser, camera, and temperature sensors) to be separately powered up or down via relay. In-flight testing confirmed that the laser head temperature is easily maintained. The instrument arms required all of the available heater power to remain ice and condensation free. In future designs, heaters will be placed closer to the windows to more efficiently eliminate condensation.

2.3.4 Mechanical design

As noted already, the system had to fit into a cylindrical canister mounted to a wing pod on the C-130 research aircraft. The optical detection part of the instrument was allowed to extend from the front of the instrument canister, and the shape of these two arms influences the flow of air around the instrument and, therefore, the flow of particles through the detection volume during flight. This is a long-standing problem in cloud physics instrumentation (MacPherson and Baumgardner, 1998) and will not be treated in detail here. Instead, we provide 'back of the envelope' estimates that guided the mechanical design of the instrument.

First, at air speeds of v = 100 m s^{-1} the flow around the nose of the instrument (where no flow separation has occurred) consists of a viscous boundary layer near the instrument surfaces and the remaining region, which can be considered inviscid and thereby obeying the Euler equations of motion for a continuous fluid (air) (Kundu and Cohen, 2002, pp. 148-192 & 312-377). The thickness δ of the viscous boundary layer can be estimated by considering the diffusion of fluid momentum during the time τ available for boundary layer growth: $\delta \sim (\nu\tau)^{1/2}$, where ν is the kinematic viscosity. That time is related to distance s along the instrument according to $\tau = sv^{-1}$, such that $\delta \sim (\nu s v^{-1})^{1/2}$. Thus, the boundary layer thickness for the length between the instrument arm tips to the observation windows is of order $\delta \sim 100$ μm. This is sufficiently small that it can be ignored, and we then solely deal with the problem of estimating flow distortion in the measurement volume due to pressure gradients (inviscid flow). Theory suggests that flow speeds along the direction of flight are within 10% of the background flow velocity at distances of roughly 2.5 radii upstream of a half sphere, which the front of our instrument can be approximated to be (Kundu and Cohen, 2002). Furthermore, wind tunnel measurements show that the velocity component along the axis of an instrument with nearly identical geometry (an FSSP-100) indeed is within the expected 10% value at the location of the optical measurement volume (Nagel, 2003). These rough arguments give some justification for the geometry of the instrument and provide confidence in our ability to measure droplet spatial distributions that are not strongly distorted from their 'ambient' form.

2.4 Holograms and digital reconstruction

HOLODEC flew on nine flights, for a total of 19 flight hours, on the C-130 Hercules research aircraft (operated by the National Center for Atmospheric Research, NCAR, under sponsorship of the National Science Foundation) in August and September 2003. The flights were made as part of the third Instrument Development and Education in Airborne Science (IDEAS 3) field project IDEAS 3[†]. During the flights we obtained about 225 minutes or 160 GB of recorded holograms. The data demonstrate the ability to detect cloud droplet sizes and three-dimensional positions and, although not our main purpose, to image ice crystals in clouds. We used numerical reconstruction to find the approximate size and position of these particles.

2.4.1 Reconstruction Method

Real images of the particles are obtained by numerical reconstruction of the digital holograms. There are a wide range of possible approaches to the reconstruction problem, each with limitations and advantages. Our purpose in this section is to establish the utility of this new instrument in detecting cloud droplets and ice crystals, and in determining particle size, shape and location. Thus we have chosen a straightforward reconstruction method that is not necessarily optimal for a given automated particle finding algorithm. This is appropriate considering optimal reconstruction methods and hologram filtering might differ for various particle shapes, sizes and experimental conditions. Particle size and location are estimated using relative intensity in reconstructed images stepped along the optical axis.

The reconstruction approach we use incorporates the kernel from Rayleigh-Sommerfeld diffraction theory in the spatial frequency domain. In the frequency domain we can write the propagated field in terms of the diffraction kernel and the measured intensity as

---

[†] IDEAS 3 was held at the National Center for Atmospheric Research, Research Aviation Facility, sponsored by the National Science Foundation. Details of the IDEAS 3 project are available at <http://www.eol.ucar.edu/raf/Projects/IDEAS/>.

$$U\left(\nu_x, \nu_y; z = z_o\right) = U(\nu_x, \nu_y; z = 0)H(\nu_x, \nu_y; z_0) \tag{2.8}$$

where $U\left(\nu_x, \nu_y; z = z_o\right)$ is the Fourier transform of the field at distance z_0 from the CCD surface, $U\left(\nu_x, \nu_y; z = 0\right)$ is the Fourier transform of the field at the CCD camera, and

$$H(\nu_x, \nu_y; z_0) = \exp\left(jkz_0\sqrt{1 - \lambda^2(\nu_x^2 + \nu_y^2)}\right) \tag{2.9}$$

, where $k = 2\pi/\lambda$ (Goodman, 1996, Eq. 3-74).

The hologram processing algorithm takes the following steps:

1. Take the negative of the hologram such that the droplets and ice crystals appear bright (as opposed to shadows).
2. Take the two-dimensional fast Fourier transform of the hologram.
3. Multiply by $H(\nu_x, \nu_y; z_0)$.
4. Take the inverse fast Fourier transform.
5. Take the modulus squared to get the intensity of the image.
6. Rescale the image intensities to emphasize the cloud droplets and ice crystals.

Note that only steps 3 through 6 need to be repeated for each reconstruction distance z_0 from a given hologram.

2.4.2 Calibration

Laboratory tests confirmed that we are able to reconstruct in planes perpendicular to the optical axis with reasonable accuracy both in terms of position and particle size. Figure 2.2 shows reconstructed images of droplets of known size. The droplets were generated with a piezoelectric droplet injector with a set of calibrated glass tips for different radii. Note that the reconstructed diameter is accurate to within one pixel (4.65 μm) of the known diameter. An additional check on the depth accuracy was made by reconstructing dust on the various optical surfaces (e.g., lenses, filters, mirror). We found that we could reconstruct dust particles of characteristic width about 20 μm to within 1 mm at 80 mm from the CCD surface. As is inherent in in-line holographic systems, the depth precision (position along the optical axis) has uncertainty that scales with droplet size. We are able to determine the relative positions of 10 to 20 μm droplets to within about 20 diameters; the position in the two dimensions perpendicular to the optical axis is known to within one pixel. We anticipate that the depth uncertainty can be reduced through more refined numerical techniques (e.g., complex reconstruction (Pan and Meng 2003)).

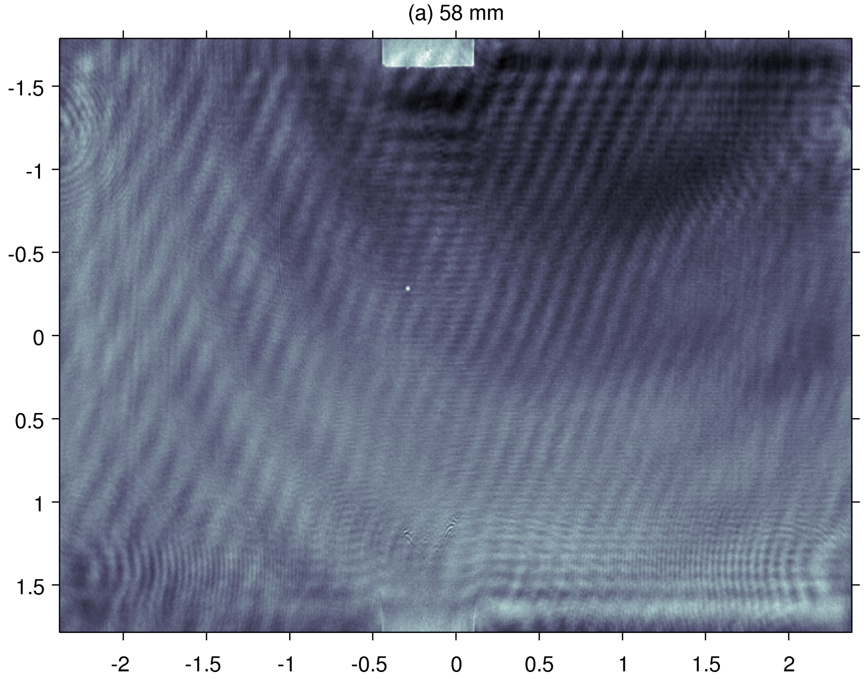
(a) 58 mm
-1.5
-1
-0.5
0
0.5
1
1.5
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2

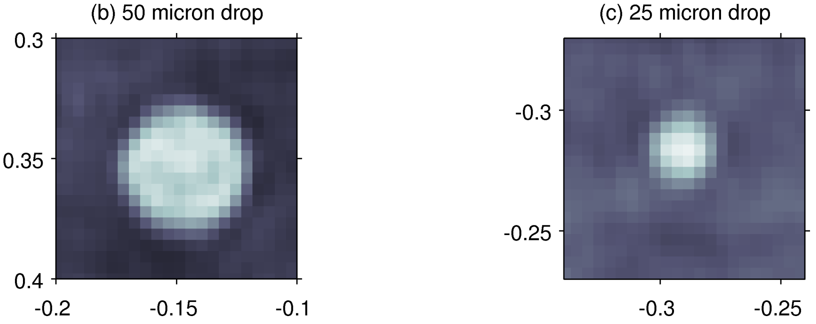
(b) 50 micron drop
0.3
0.35
0.4
-0.2
-0.15
-0.1
(c) 25 micron drop
-0.3
-0.25
-0.3
-0.25

Figure 2.2: Frame (a) shows a whole reconstruction slice at 58 mm from the CCD. The droplet generator tip is shown at the top of the frame. A 50 μm diameter droplet is visible about two thirds down the frame below the tip. Frame (b) shows a closeup of the 50 μm diameter drop and (c) shows a 25 μ diameter drop. Note that the reconstructed images of the droplets are within a pixel (4.65 μm) of the known diameter. Help with the calibration droplets was graciously provided by U. Maixner and D. Nagel from GKSS Forschungszentrum.

2.4.3 Cloud Droplets and Ice Crystals

During the field testing of the instrument we sampled clouds that, although below the freezing point of water, consisted predominantly of supercooled liquid droplets.

Subsequent reconstruction of holograms from these clouds confirms that we are able to image the droplets, thereby estimating their three-dimensional position within the sample volume and estimating their size. Figure 2.3 shows cloud droplets from a small section of a hologram obtained while flying through a cloud where other instruments (e.g., the FSSP-100) indicated the largest droplet sizes were in the range of 10 to 20 μm in diameter. Each frame of Figure 2.3 shows reconstructions at advancing distances along the optical axis, illustrating how droplets come into and go out of 'focus'.

In addition to the liquid cloud particles we obtained holograms of ice crystals. Figure 2.4 shows reconstructions of relatively small ice crystals and Figure 2.5 shows several reconstructions of larger ice crystals. Small ice crystals, although nonspherical, appear to have nearly circular interference fringes and therefore, prior to reconstruction can be mistaken as large, drizzle drops. The interference patterns resulting from large ice crystals are clearly not circular. The reconstructions of large ice crystals are less clear than those for smaller particles, most likely resulting from the non-negligible virtual image (because large particles in the instrument sample volume do not meet the Fraunhofer condition) and the fact that the reference wave is severely occluded by the presence of large particles.

2.5 Summary and future work

Most treatments of cloud processes such as particle collision rates and radiative transfer in clouds, are implicitly based on the assumption that particles are distributed in a perfectly random manner at small scales. Theory and computational work, however, suggest that cloud particles are spatially correlated on small scales if a cloud is sufficiently turbulent. Recent atmospheric measurements tend to confirm this picture, but the experiments are challenging and fraught with instrumental artifacts, and therefore are not yet satisfactory.

We have tackled this problem directly by building an instrument capable of measuring the 3-dimensional spatial distribution of cloud particles. The instrument, called Holographic Detector for Clouds (HOLODEC), digitally records optical holograms of a small volume of cloud. The light source is a Q-switched Nd:YLF laser with pulse energy of 30 μJ and pulse width of 20 ns. This allows clear holograms to be obtained in daylight conditions on an aircraft moving at speeds of approximately 100 m s^{-1}. The detector is a 10-bit extended head CCD camera interfaced with a control and data acquisition system via optical fiber. An actively controlled anti-ice and condensation system is used to allow imaging at temperatures down to -30^o C and in supersaturated water vapor. HOLODEC took its first flight into clouds aboard the NCAR C-130 aircraft in late 2003. Subsequently, the holograms have been reconstructed numerically, resulting in real images of the sampled particles. We therefore obtain all three spatial coordinates and the shape and size of particles in the measurement volume. Initial

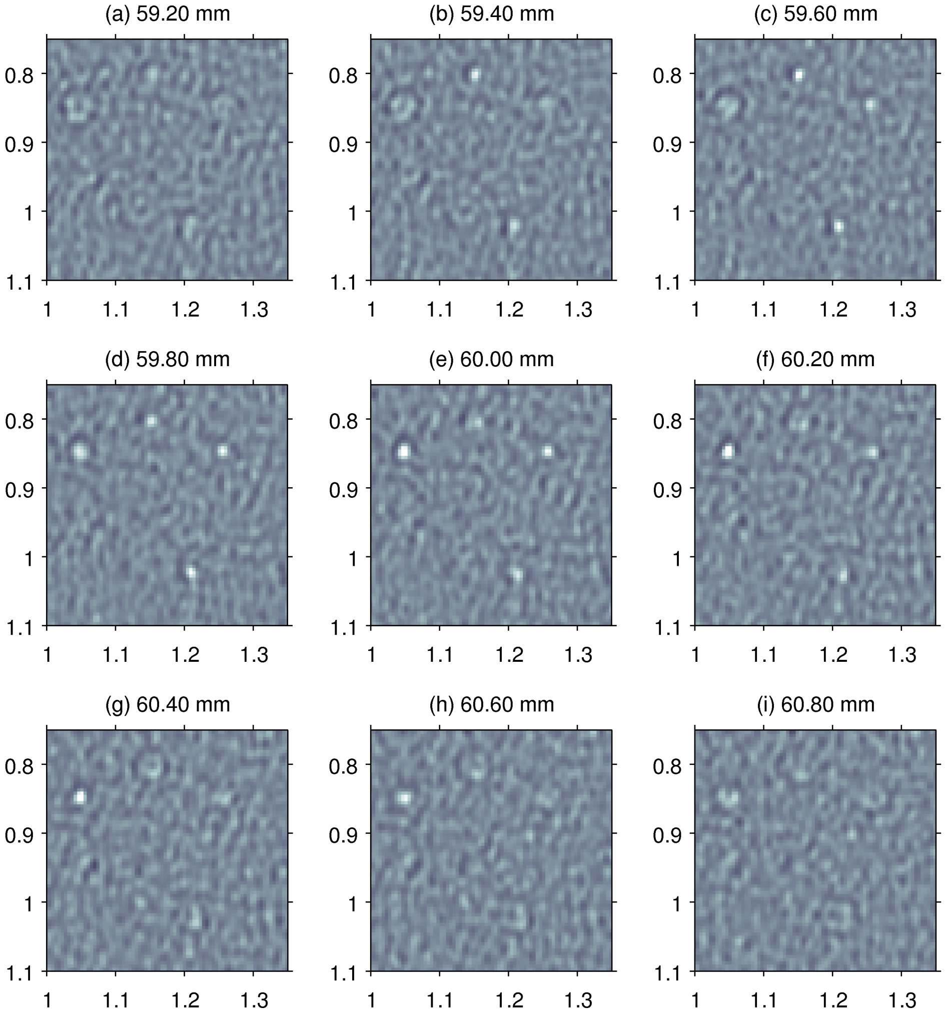

Figure 2.3: Images from a small section of a hologram obtained during flight traverse of a cloud. Frames (a) through (d) show digitally reconstructed images in 500 μm steps along the optical axis, as indicated above each image. Each frame shows cloud droplets that come in and out of 'focus' as the frames advance along the optical axis. Units on each axis are in mm. Other instruments on the aircraft indicated the largest drops in this cloud are approximately 10 to 20 μm in diameter.

data analysis confirms that HOLODEC can detect cloud droplets with diameters of 10 μm and above, and that we are able to reconstruct more complex shapes such as ice crystals.

A primary goal of HOLODEC is to quantify the spatial distribution of cloud droplets through direct measurement of the pair correlation function η. This requires that

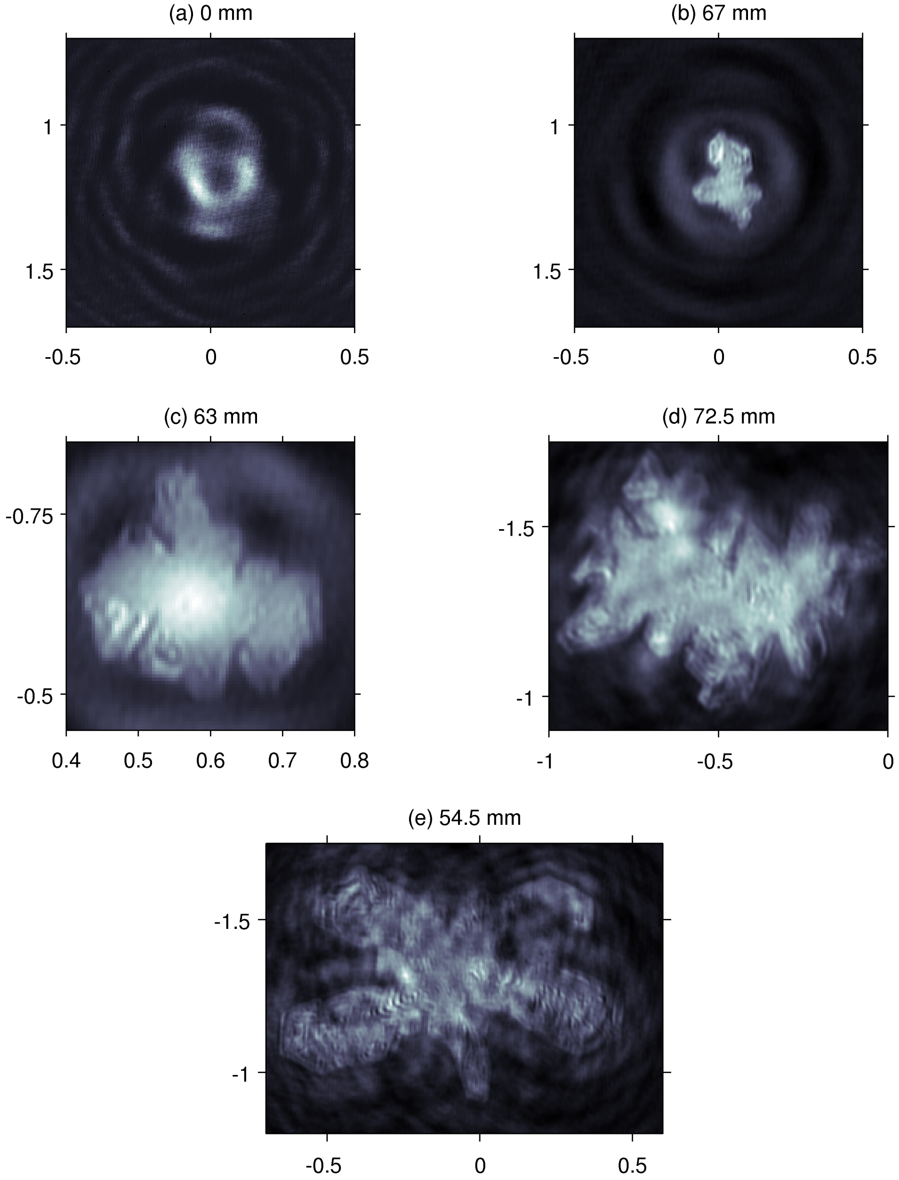
(a) 0 mm
1
1.5
-0.5
0
0.5
(b) 67 mm
1
1.5
-0.5
0
0.5
(c) 63 mm
-0.75
-0.5
0.4
0.5
0.6
0.7
0.8
(d) 72.5 mm
-1.5
-1
-1
-0.5
0
(e) 54.5 mm
-1.5
-1
-0.5
0
0.5

Figure 2.4: Various reconstructions of large ice crystals and one of the original holograms. The distance from the CCD is indicated above each frame. The units on each axis are mm. Frame (a) shows the section of the hologram from which the ice crystal in frame (b) was reconstructed. Frames (c), (d), and (e) are other reconstructions of large ice crystals.

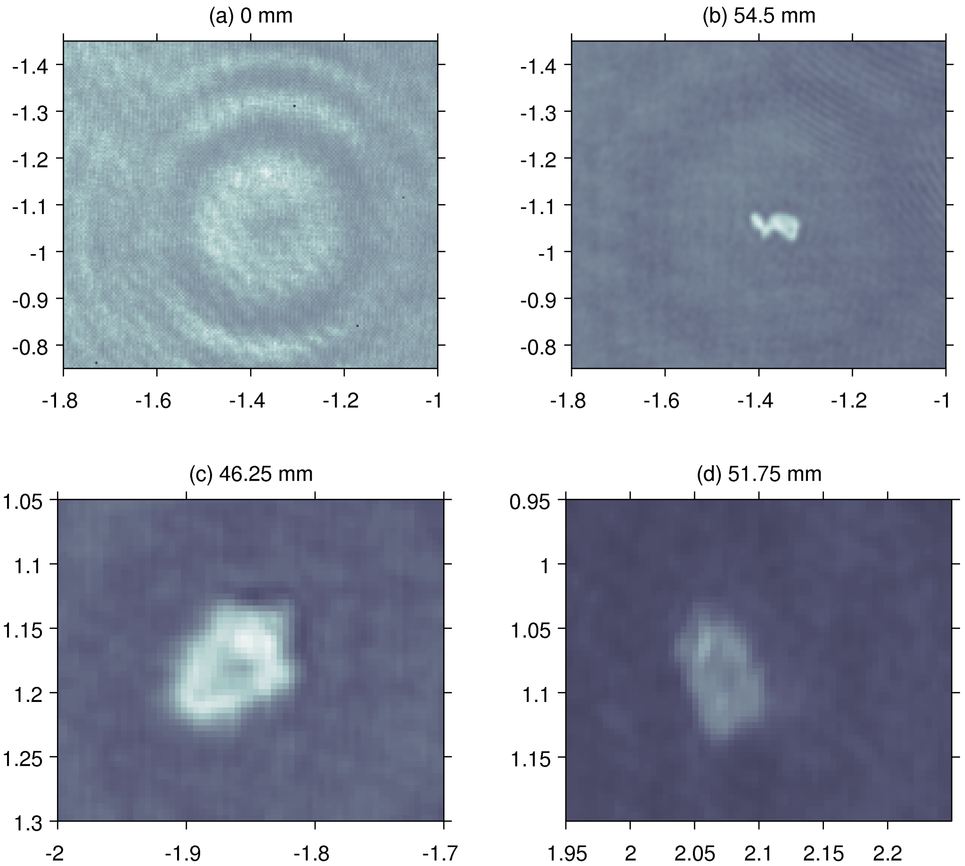

Figure 2.5: Various reconstructions of small ice crystals and one of the original holograms. The distance from the CCD is indicated above each frame and the units on each axis are mm. Frame (a) shows the section of the hologram from which the ice crystal in frame (b) was reconstructed. Frames (c), and (d) are other reconstructions of small ice crystals.

we sample a sufficient number of cloud droplets, so we briefly consider the question of how many holograms will be required for typical cloud conditions. Kostinski and Shaw (2001) made a first estimate of the pair correlation function in clouds using one-dimensional droplet positions measured in a statistically homogeneous cloud. They found empirically that statistical estimates of η converged when droplet numbers of order 10^5 were sampled. The HOLODEC measurement geometry allows us to estimate the pair correlation function at scales of 1 cm and below. What truly matters for the calculation of η is not merely the number of droplets, rather the number of inter-droplet distances of length < 1 cm.

We therefore consider how many inter-particle distances < 1 cm are measured in equal volumes sampled by the one-dimensional measurement (Kostinski and Shaw, 2001) and by the holographic instrument described here. HOLODEC has a measurement volume of order 1 cm^3 so a conservative estimate is that $N = 150$ droplets are sampled. This corresponds to $N(N-1)/2 \approx 10^4$ inter-droplet positions < 1 cm. In the one-dimensional geometry (Kostinski and Shaw, 2001) the instrument has a measurement cross section of approximately 1 mm^2, so a 1-m length of cloud must be sampled to obtain a 1 cm^3 volume with 150 droplets; in this scenario, however, the number of inter-droplet positions < 1 cm is approximately 10^2, a factor of 10^2 smaller than in the three-dimensional measurement. This is a powerful advantage of holography over lower-dimensional measurements. It follows directly that the number of droplets that must be sampled by HOLODEC to estimate η is of order 10^3, which can be achieved by recording approximately 10 holograms. Even with the constraint of statistical homogeneity this is a realistically attainable number.

In the future we will develop automated routines for droplet detection. With sufficient data the droplet pair correlation function can be estimated over scales from approximately 100 μm to 1 cm, thereby quantifying particle spatial correlations in clouds. We also anticipate making improvements to HOLODEC, specifically in the aerodynamic design, the heater control system, and the optical system. Perhaps most importantly, both the quality of individual holograms and the ability to measure greater numbers of droplets in a hologram can be improved by increasing the size of the CCD detector.

Acknowledgments

We thank D. Cook and J. Howe for help with the mechanical design and construction, the NCAR Research Aviation Facility staff for their assistance with integrating the probe with the C-130 instrument system, S. Beaton and J. Borysow for valuable discussions during the instrument design, and T. Jin from CrystaLaser for his flexibility in meeting our specific laser needs. We also thank U. Meixner and D. Nagel of GKSS for their generous assistance with the calibration droplets, and R. Ruth for help accessing data from other instruments flown during IDEAS 3. We are especially grateful to J. Stith, RAF Director, for his invitation to participate in IDEAS 3 and his encouragement on the instrument development. This work was supported by NASA grant NAG5-10568, NSF grant ATM99-84294, and an NSF Graduate Research Fellowship (JPF).

Chapter 3

Practical Methods for Reconstruction and Characterization of Particles in Digital Inline Holograms

Hologram reconstruction algorithms typically undersample in phase for distances of interest in typical optical setups*. Given herein is a algorithm to reconstruct digital particle holograms with a single method over typical distances of interest, with a sample spacing constant in distance, has a diffraction limited resolution, and can be implemented to be as fast as any reconstruction method. This algorithm proves accurate against analytical solutions to the Huygens-Fresnel propagation integral. A low pass filter can be applied to enforce a uniform minimum particle size detection limit throughout a sample volume.

Tens of thousands of holograms of cloud ice particles are digitally reconstructed using the algorithm discussed. Positions of ice particles in the size range of 20 micrometers to 1.5 mm are obtained using a new algorithm that accurately finds the position of large and small particles along the optical axis. The digital reconstruction and particle characterization algorithms are implemented in an automated fashion with no user intervention on a computer cluster. Strategies for more efficient algorithm implementation on a computer cluster are presented.

---

[*] By Jacob P. Fugal, Timothy J. Schulz and Raymond A. Shaw. In preparation to be submitted to Measurement Science and Technology (Fugal et al., 2007).

3.1 Introduction

The problem of reconstructing holograms of large and small ice crystals, finding the particles' focused position so we can calculate their size distributions and number densities is a unique one. While many papers have been published on how to reconstruct small particles in digital holograms, there are few that relate to large particles. Among those that reconstruct large particles are Fugal et al. (2004, see Chapter 2) and Fugal and Shaw (2007, see Chapter 4) where they show holograms of atmospheric ice crystals from an airborne probe. (Raupach et al., 2006) show good reconstructions of ice crystals in their ground-based cross-beam two-camera instrument. Malkiel et al. (2004, 2006) digitally image optically reconstructed holograms and then use an automated algorithm to find 20 μm to 5 mm particles about calanoid copopods (plankton) to study plankton behavior.

For reconstructing small particles, there is no dearth of papers focused on how to refine some holographic measurement of particles or a new and better way to reconstruct holograms or setup the optical system for holography. For example, Soulez et al. (2007); Yang et al. (2005, 2006); Ilchenko and Sattelmayer (2005); Fournier et al. (2004); Pan and Meng (2003); Soontaranon et al. (2002); Coetmellec et al. (2002) all offer ways to reconstruct holograms and find the position along the optical axis of particles in digital inline holography. But few papers talk about processing thousands of holograms en masse in an automated, robust way to gain some scientific insight to a physical problem. However, recently two papers have focused on how to efficiently process holograms: Masuda et al. (2006) describe a specialized computer to reconstruct holograms, and Satake et al. (2007) describe a parallel algorithm to reconstruct holograms on a super computer.

In this paper, we present several methods and refinements to reconstruct holograms of small and large particles (from about 15 to 1500 μm in extent), find and identify the particles, measure their three-dimensional position and record it in a useful form to later compute summary statistics, and do this for thousands of holograms on a computer cluster. In Section 2, we discuss removing noise from holograms via background division. We discuss in some detail our method of reconstructing holograms accurately in Section 3. We relate, in Section 4, how to reconstruct holograms in a way to enforce a particle size independent sample volume for small particles. In Section 5, we describe a robust method for finding the position along the optical axis of particles in a large size range. We explain how we find the peak of edge sharpness in a way that is less sensitive to both local and global noise in section 6. In section 7 we briefly relate how we reconstruct thousands of holograms quickly on a computer cluster. Finally we summarize and give an outlook on future work in section 8.

3.2 Background division by median of neighboring holograms

It is difficult to make an optical setup that allows nearly noise free holograms. Noise comes from sources such as interference fringes from multiple reflections off of optical surfaces especially the coverglass of CCD and CMOS cameras, from particles or contaminants stuck to the windows and optics, and from nonuniformities in the incident laser field. All this is in addition to noise from the neighboring particles' real and virtual images. Several methods have been proposed to remove background noise and not distort the particles of interest. For example, Kreis and Juptner (1997) mention subtracting the local mean to flatten the background. Raupach et al. (2006) removed noise in their holograms by dividing a hologram by a single hologram neighboring in time. The most effective noise removal method we have found so far is much like the latter method.

Our method is to take the pixel by pixel median of several sequential holograms taken before and after (when possible) the hologram we wish to reconstruct and divide this median hologram into the hologram we wish to reconstruct. An example of this is shown in Figure 3.1. The pixel by pixel median of the neighboring holograms yields a background image nearly free of particles of interest which change position from hologram to hologram, but retains stationary fringes that stay from hologram to hologram such as those resulting from particles stuck on the windows. By the same argument, this method eliminates background noise such as interference fringes from the coverglass as long as the laser pulse and optics are stable from hologram to hologram. To the extent these fringes are not stable, the background noise (especially fine interference fringes that change from hologram to hologram) survive the background division.

We find that the median of approximately 7 neighboring holograms is usually sufficient to get a background hologram free of particles of interest. If the hologram we want to reconstruct is missing a particle that appears in the background hologram, then as Raupach et al. point out, the background division will make the missing particle appear in opposite contrast with the particles of interest.

Raupach et al. say that by dividing a hologram by its neighbor and having particles from one hologram appear dark and the particles in the other hologram appear bright, one can save time by reconstructing both sets of particles at once. In this case, the background division process has a contrast negative effect on the particles in hologram in the denominator. However, we have found that for especially large ($\gtrsim 200\mu m$) particles, the division or multiplicative inverse operation makes the particle appear malformed when reconstructed and impossible to focus. Thus, to take a negative of a hologram without distorting particles, especially large particles, it is better to take

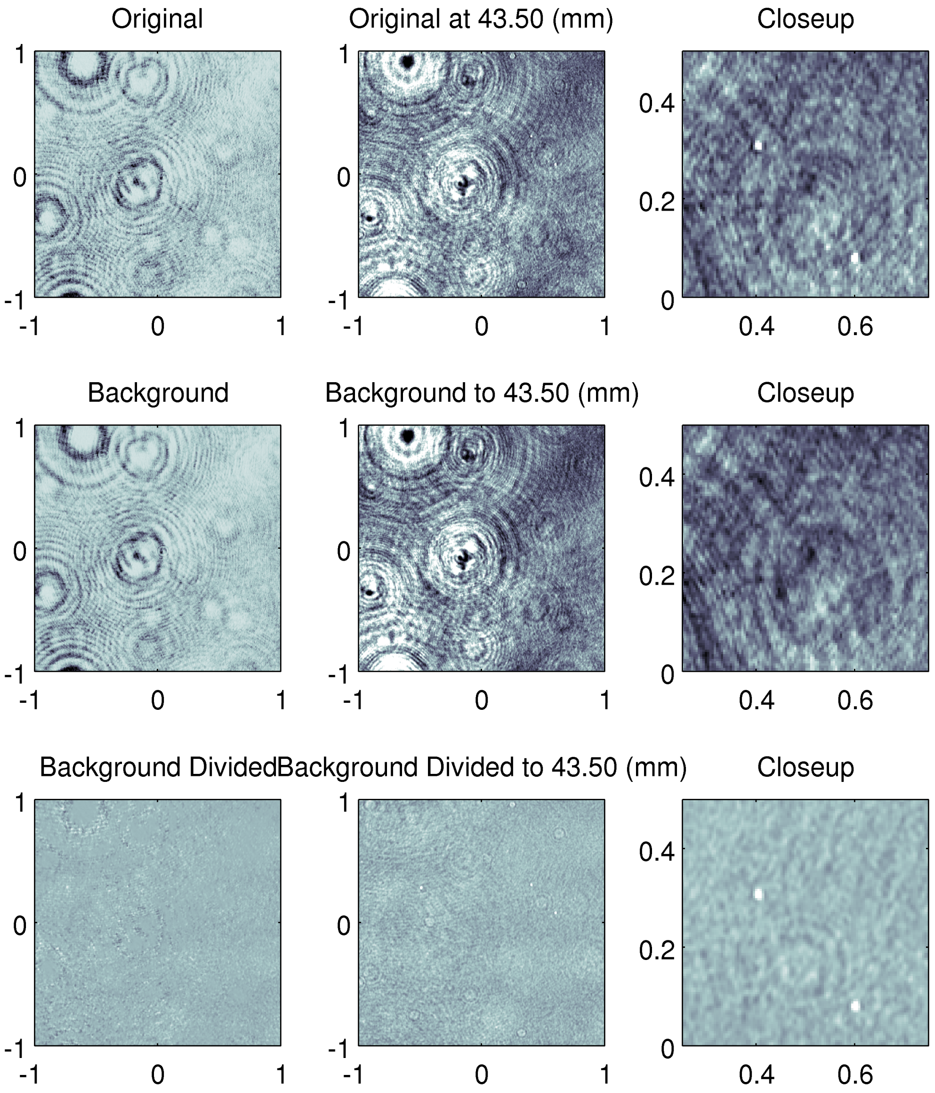
Original
Original at 43.50 (mm)
Closeup
Background
Background to 43.50 (mm)
Closeup
Background Divided
Background Divided to 43.50 (mm)
Closeup
1
0
-1
-1
0
1
0.4
0.2
0
0.4
0.6

Figure 3.1: Shows a hologram divided by the median of its neighboring 7 holograms (including itself). The top left panel is a subarea of the original hologram, which is reconstructed to 43.5 mm in the top center panel. The top right panel shows a closeup of two bright dots which are 20 μm cloud droplets taken on the 2003-09-05 Research Flight in the IDEAS-3 project (Fugal et al., 2004, see Chapter 2). The center row shows the background image and the bottom row is the background divided hologram. Note the presence of water drops stuck to the windows of HOLODEC in the original and background hologram rows and their absence in the bottom row. Note the presence of the two bright dots in the top right panel, their absence in the center panel, and their presence in a relatively noise free background in the bottom right panels. All axes are in mm, and each image is equally contrast enhanced for visibility in printing.

a contrast negative by:

$$\text{neg}(U) = \max(U) + \min(U) - U \tag{3.1}$$

where *neg* is a contrast negative operation, and *max* and *min* take the maximum and minimum pixel value of the whole hologram respectively, and U is the hologram, than to take a negative by multiplicative inverse.

3.3 Reconstruction method

The success to which we can use holography to measure positions and sizes of particles is completely dependent on our ability to accurately reconstruct holograms (assuming that the hologram itself sufficiently samples the interference pattern Kreis (2002a)). Several methods of hologram reconstruction have been used by other researchers. The major possible methods are summarized in Kreis et al. (1997) and for more general applications they are described in textbooks such as Joseph Goodman's Fourier Optics (Goodman, 1996). Pan and Meng (2003) and Fournier et al. (2004) give methods using the Fresnel approximation. Soulez et al. (2007) use the Fraunhofer approximation. In this paper, we use no approximation beyond that of scalar diffraction theory as does Raupach et al. (2006).

One point generally omitted in describing the various methods for numerically solving the propagation of light is the distances we may reconstruct to without undersampling the phase in the propagation kernels. Another fact generally not considered is that, in practice, the distances to which we wish to numerically reconstruct holograms to are typically in the transition region between which two conjugate numerical propagation methods are sufficiently sampled. For example, for the HOLODEC instrument, the distance between propagation methods is 42 mm in a sample volume extending from around 35 to 80 mm (Fugal et al., 2004, see Chapter 2). (The theory underlying this calculation will be discussed subsequently.)

We desire to have a single propagation method so as not to have to concern ourselves with possible biases or errors introduced by trying to combine reconstruction data from two different numerical propagation methods. Further, we would like to have a reconstruction method that has a constant sample spacing for any distance reconstructed to so as eliminate the need for resampling. We desire a method does not lose any spatial resolution as a consequence of computation speed or propagating beyond the undersampled limit. Finally, this method should be proven against known solutions to the propagation of diffracted light. Finally, we need a method that is fast and efficient to compute.

We have developed and implemented a numerical propagation method that has constant sample spacing, has diffraction-limited spatial resolution, is accurate for distances beyond the undersampling limit, is fast and simple to compute and have thoroughly tested it using known propagation solutions to make sure the method is sound in the parameter space in which we wish to use it. The method is well suited to hologram reconstruction in that it accurately reconstructs the hologram to the distances most experimental hologram setups require given typical parameters for hologram size, pixel spacing, and wavelength. A summary of the method is given in Fugal et al. (2004, see Chapter 2), and a more detailed explanation of the method is given below.

3.3.1 Huygens-Fresnel kernels

We use the filtering (as opposed to the convolution form) of the Rayleigh-Sommerfeld or Huygens-Fresnel kernels (Goodman, 1996, Ch. 3). Based solely on the approximations in Scalar Diffraction Theory and assuming monochromatic light, the kernel used in the spatial domain (convolution method) and frequency domain (filtering method) are respectively:

$$u(\overrightarrow{r};d) = \int\int u(\overrightarrow{\rho};0)h(\overrightarrow{r}-\overrightarrow{\rho};d)d\overrightarrow{\rho}, \tag{3.2}$$

and

$$u(\overrightarrow{r};d) = U(\overrightarrow{\nu_r};0)H(\overrightarrow{\nu_r};d), \tag{3.3}$$

where the $\overrightarrow{\square}$ denotes a vector, $\overrightarrow{r}$ and $\overrightarrow{\rho}$ are two-dimensional spatial coordinates perpendicular to the optical axis, u is the field we are propagating, $u \Longleftrightarrow U$ and $h \Longleftrightarrow H$ are two-dimensional Fourier transform pairs, $\overrightarrow{\nu_r}$ is a two-dimensional vector of the conjugate Fourier frequency variables, and d is the distance propagated. The equivalence of these equations can be shown via the Fourier Transform convolution theorem (Goodman, 1996). The kernel h and its conjugate Fourier pair H, without approximation, are respectively:

$$h(\overrightarrow{r}-\overrightarrow{\rho};d) = \frac{de^{ikR}}{R^2}\left(\frac{1}{j\lambda}+\frac{2\pi}{R}\right), \tag{3.4}$$

and

$$H(\overrightarrow{\nu_r};d) = \exp\left(jkd\sqrt{1-\lambda^2\left|\overrightarrow{\nu_r}^2\right|}\right), \tag{3.5}$$

where $R=\sqrt{d^2+\left|\overrightarrow{r}-\overrightarrow{\rho}\right|^2}$ is the distance from any point $\overrightarrow{\rho}$ on the $d=0$ plane to $\overrightarrow{r}$ on the $d=d$ plane where we are propagating to, λ is the wavelength of a monochromatic wave, and $k=2\pi/\lambda$ is the wavenumber. Note that the $2\pi/R$ term in the convolution kernel and the imaginary roots of $\sqrt{1-\lambda^2\left|\overrightarrow{\nu_r}^2\right|}$ that make the exponent partly real in the filtering kernel are related to evanescent waves and can most often be safely ignored and indeed they are often excluded from statements of these kernels in other works (e.g. Goodman 1996).

3.3.2 Distances at which the kernels are undersampled

To propagate correctly, we need to make sure the phase in either kernel is not undersampled after discretizing the variables. That is, with a complex field, we must make sure that the phase does not vary by more than π between samples. We will use $\overrightarrow{\Delta r} = \Delta x\, \hat{\imath} + \Delta y\, \hat{\jmath}$ to represent the diagonal of a rectangular pixel and thus the maximum extent from the center of the CCD is $\overrightarrow{L_r}/2 = [M/2, N/2] \cdot \overrightarrow{\Delta r}$ where N is the number of pixels in the $\hat{\imath}$ direction and M is the number of pixels in the $\hat{\jmath}$ direction, and $\Box \cdot \Box$ represents the vector inner product. Similarly in frequency space we will set $\overrightarrow{\Delta \nu_r} = \Delta \nu_x\, \hat{\imath} + \Delta \nu_y\, \hat{\jmath}$ to be the diagonal of a pixel and the maximum extent in frequency space to be $\overrightarrow{L_{\nu_r}}/2 = [M/2, N/2] \cdot \overrightarrow{\Delta \nu_r}$. We may calculate when the phase in either kernel becomes undersampled using the equation $|\overrightarrow{\Delta r} \cdot \nabla \Phi| < \pi$ and $|\overrightarrow{\Delta \nu_r} \cdot \nabla \Phi| < \pi$ where Φ is the phase in the respective kernel. After some algebra, we get the following distance limits to which we may propagate by convolution with h or filtering with H respectively:

$$\frac{k|\overrightarrow{r} \cdot \overrightarrow{\Delta r}|}{\sqrt{d^2 + |\overrightarrow{r}|^2}} < \pi, \tag{3.6}$$

and

$$\frac{2d\lambda|\overrightarrow{\nu}_r \cdot \overrightarrow{\Delta \nu_r}|}{\sqrt{1 - \lambda^2|\overrightarrow{\nu}_r|^2}} < 1. \tag{3.7}$$

Solving for the most constraining limit by letting $r \to \overrightarrow{L}_r/2$, and $\nu_r \to \overrightarrow{L_{\nu_r}}/2$ we get respectively:

$$\sqrt{\frac{|\overrightarrow{L_r} \cdot \overrightarrow{\Delta r}|^2}{\lambda^2} + \frac{|\overrightarrow{L_r}|^2}{4}} < d, \tag{3.8}$$

and

$$d < \sqrt{\frac{\lambda^{-2} - 2^{-2}|\overrightarrow{L_{\nu_r}}|^2}{|\overrightarrow{L_{\nu_r}} \cdot \overrightarrow{\Delta \nu_r}|}}. \tag{3.9}$$

We may also use the fact that pairs of Fourier conjugate variables are related by $\overrightarrow{L_{\nu_r}} = \left(\overrightarrow{\Delta r}\right)^{-1}$ and $\overrightarrow{\Delta \nu_r} = \left(\overrightarrow{L_r}\right)^{-1}$ where the $(\Box)^{-1}$ operator represents an element by element inverse and so does not commute with some other operators such as the vector inner product. So the latter limit becomes

$$d < \sqrt{\frac{\lambda^{-2} - |2\overrightarrow{\Delta r}|^{-2}}{|(\overrightarrow{L_r})^{-1} \cdot (\overrightarrow{\Delta r})^{-1}|^2}}. \tag{3.10}$$

Making the approximation that the pixels are many wavelengths in extent $|\overrightarrow{\Delta r}| \gg \lambda$ and the assumption that the pixels are square $\Delta x = \Delta y$ and the imaging chip is square, then the convolution and filtering method limits reduce to

$$\frac{L_r \Delta r}{\lambda} < d, \tag{3.11}$$

and

$$d < \frac{L_r \Delta r}{\lambda}, \tag{3.12}$$

respectively where $\Delta r = |\overrightarrow{\Delta r}|$ and $L_r = |\overrightarrow{L_r}|$. Note that this is also an expression for the sample spacing relation for propagation via the Fresnel or Fraunhofer approximation (Goodman, 1996, Section 4.2). Thus the distance at which Huygens-Fresnel convolution and filtering methods' sampling limits approximately meet is the same distance to which Fresnel and Fraunhofer approximations have the same sample spacing in the propagation plane as the source plane.

Now again, we wish to propagate beyond the distance limit using either filtering or convolution so one method can be used for an entire sample volume. We prefer the filtering kernel because it retains the same sample spacing regardless of propagation distance unlike the Fresnel or Fraunhofer approximation of the convolution kernel (Kreis et al., 1997; Kreis, 2002b). It also has practically the same computation speed as the Fresnel approximation for holography applications. To see this, recall that hologram reconstruction by the Fresnel approximation requires a single Fourier transform per reconstruction slice where reconstruction using the filtering kernel requires a single Fourier transform of the hologram and then a Fourier transform per desired propagation distance. When reconstructing holograms with particles, we typically require tens or a hundred propagation distances and thus the single extra Fourier transform that the filtering kernel requires is insignificant when comparing computation speeds of the two reconstruction methods.

What are the consequences of propagating beyond the distance where the phase in the filtering kernel is undersampled? For inline holograms with typical properties of a nearly plane wave background overlaid with relatively weak interference fringes, reconstructing to distances not far beyond the undersampling limit show few visible effects. This is because the first frequencies to be undersampled are the high frequencies which usually have small amplitudes for holograms as described. One workaround of the undersampling problem is to pad the hologram with its mean values or zeros on its edges to extend it without adding information. This would thus increase the distance we can propagate to via the filtering kernel without undersampling. But it also increases the computation time without a great gain in computational accuracy.

Raupach et al. (2006) state that by oversampling the hologram (decreasing $\Delta r \rightarrow \Delta r/2$) allows them to more accurately propagate using the Huygens-Fresnel filtering kernel. This actually shortens the distance to which they can propagate without undersampling. For their experimental setup ($\Delta x = \Delta y = 6.7\mu$m, $\lambda = 532$ nm, L_r 7 mm) d must be less than about 8.8 cm when their desired reconstruction distances go from 8.8 to 12.3 cm. Oversampling the hologram lessens this to 4.4 cm. Note that oversampling the hologram also increases the reconstruction time by nearly an order of magnitude.

Filtering out undersampled frequencies, if done well, removes most of the effects of undersampled frequencies, does not increase the computation time beyond a very small amount to calculate the filter, and allows the Huygens Fresnel filtering method to be used in the entire sample volume even for distances beyond the limit in Eq. 3.10. An appropriate filtering method is discussed in the next section.

3.3.3 Filtering out undersampled frequencies in the Huygens-Fresnel filtering kernel

The limit at which the filtering kernel has undersampled frequencies can be solved numerically using Eq. 3.7. For a given distance d and wavelength λ can we find the ν_r for which:

$$4d^2|\overrightarrow{\nu}_r \cdot \overrightarrow{\Delta\nu_r}|^2 + |\overrightarrow{\nu}_r|^2 = \lambda^{-2} \tag{3.13}$$

or written out in rectangular frequency components:

$$\left(4d^2\Delta\nu_x^2 + 1\right)\nu_x^2 + \left(4d^2\Delta\nu_y^2 + 1\right)\nu_y^2 = \lambda^-2 \tag{3.14}$$

The most basic filter would be to make a rectangular shaped filter by turning the above into a boolean array zeroing out all undersampled frequencies and passing all the correctly sampled ones. But rectangular filters introduce strong aliasing into the result. We implemented the filter using a Super-Gaussian filter given by:

$$\exp\left(-\frac{1}{2}\left(\left(\frac{\nu_x}{\sigma_{\nu_x}}\right)^2 - \left(\frac{\nu_y}{\sigma_{\nu_y}}\right)^2\right)^n\right) \tag{3.15}$$

where we have omitted a normalization factor and we use n = 3. We set the width to be such that the Super-Gaussian reaches has the value of 1/2 at the point the frequencies ν_r become undersampled as calculated numerically by Eq. 3.14. We chose these parameters by comparing propagated fields vs. analytical fields with different filter parameters and looking for parameters that produce lower noise and lower aliasing. What is the minimum resolvable length scale with this low pass filter? Ideally it would be no greater than what other propagation methods offer. In this case ν_x in Eq. 3.14 is approximately the maximum spatial frequency passed through the filter and so $l = \nu_x^{-1}$ is the approximately the minimum resolvable length scale in the x direction. As before, assuming the propagation distance is much greater than the imaging chip size ($d \gg L_x$), and so ($\nu_x d = L_x^{-1} d \gg 1$), we obtain:

$$l = \frac{2d\lambda}{L_x}. \tag{3.16}$$

Or in other words, the length scale l is two sample spacings from the Fraunhofer or Fresnel approximations. Further, it is the diffraction limit for a square aperture. Thus, no resolution is lost with this low pass filter, but it actually neatly enforces a diffraction limited propagation the same as the Fraunhofer and Fresnel approximations do.

3.3.4 Testing the Huygens-Fresnel filtering kernel with a low-pass filter

We verified that propagation via the Huygens-Fresnel filtering kernel with the low-pass filter works by inspecting the similarity between what the described propagation method gives versus what known analytical propagation solutions give. We tested this propagation method for a wide range of particle size, detector size, and propagation distance parameter ranges. We note that most detectors have pixel sizes that are a factor of order 10 larger than the wavelength of visible light, and most detectors have from a few hundred to a few thousand pixels on a side. With these typical values in mind, we chose a pixel size of $\Delta x/\lambda = \Delta y/\lambda = 10$, and a detector size of $L_x = L_y = 1024\Delta x = 1024\Delta y$ to perform these tests. We performed tests using every combination of values for the parameters:

1. For the following object sizes 0.2, 0.5, 1, 2, 4, 10, 20, 40, 100, 200, 400 pixel widths Δx or from $2x10^{-4}$ to 0.4 detector widths L_x;
2. For the following propagation distances in units of $d_l = L_r\Delta_r/\lambda$ from the source plane every .05 d_l from .05 to $1d_l$, every 1 d_l from 1 to 10_l, and 15 and 20 d_l;
3. For an object as a shadow on a unit amplitude background OR a source of unit amplitude on a zero amplitude background; and
4. Propagating with and without applying the described low-pass filter.

The known solutions used to make the test field are the Fresnel approximation for diffraction around a rectangle using the equations given in Goodman (1996, Section 4.5.1), and the equations for diffraction around a circular disk using Sheppard and Hrynevych (1992). We made a test field containing both a rectangle and a circle at off-center locations. Figure 3.2 shows one of the 1364 (11 object sizes × 31 distances × 2 amplitude levels × 2 filter vs. nonfilter) propagation tests. It shows a test field and the intensity and phase of both the analytical field in the top row. It shows the error with respect to the analytical field and the intensity and phase for unfiltered numerical propagation in the center row and the same for filtered propagation in the bottom row. Note that the error range in the unfiltered field is twice that of the filtered propagation (note the colors scales differ), and that the noise pattern in the filtered error panel is far less aliased than in the unfiltered error panel.

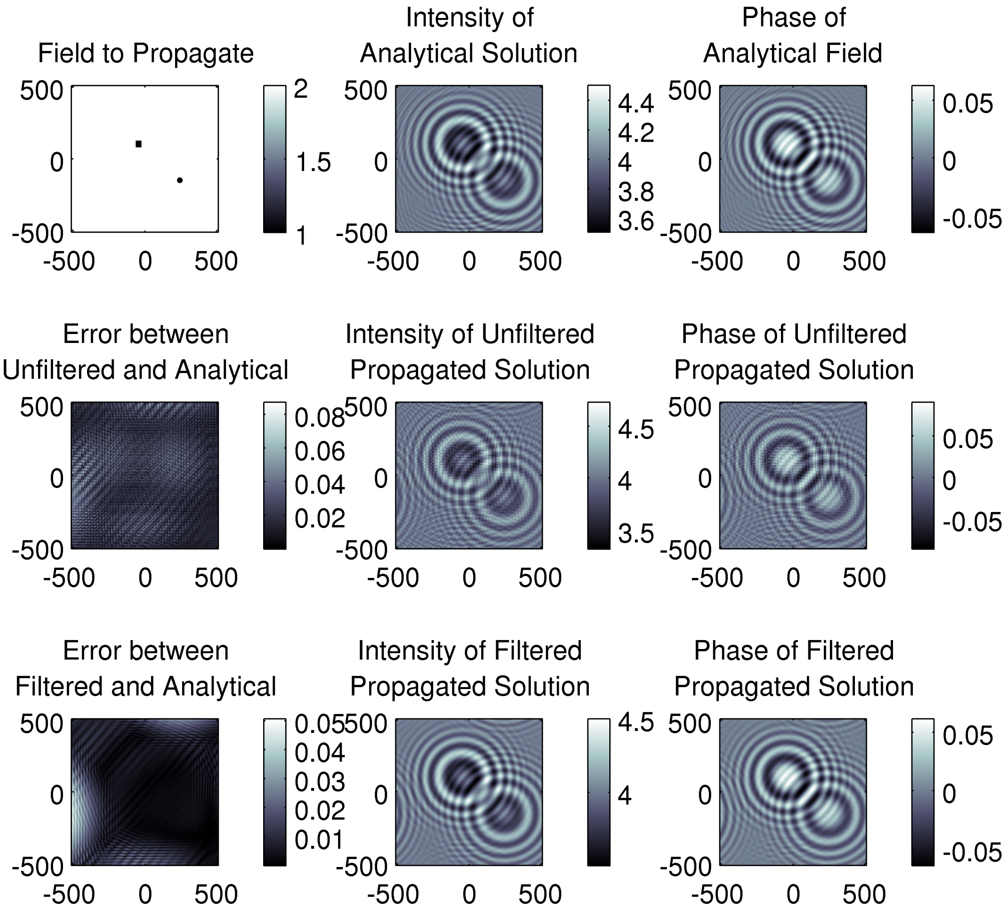

Figure 3.2: An example of one of the 1394 tests done for the Huygens-Fresnel Filtering propagation method as discussed in the text. The top left panel shows a 40x48 Δx square and 40 Δx diameter circle that are reconstructed to 10 d_l using analytical solutions in the top row, without the low pass filter in the center row, and with the low pass filter in the bottom row. All axes are in units of Δx and the intensity is in arbitrary units. Note that the color scales are not all the same. Note in the error panels that the low pass filtered error in the bottom left panel has about half the error and a much more reasonable noise pattern than the unfiltered field in the center left panel.

We found that propagation via Huygens-Fresnel filtering with the low pass filter works well for all resolvable particle sizes. As expected, propagation with and without the low pass filter did not work well for objects that are not well resolved (i.e. that are smaller than the pixel size or sample spacing). The described method worked well for both the fields with the objects as shadows and sources.

We found that applying the low pass filter in reconstructions to distances below the limit given in Eq. 3.10 has little effect on the reconstructions as practically nothing is filtered out. Applying the low pass filter to reconstructions to distances well beyond d_l significantly increases the accuracy of the reconstruction over the unfiltered field. The one exception is the filtered and unfiltered propagation methods reconstruct very large objects (the objects about 0.4 L_x in extent) nearly equally well at distances $d > d_l$.

Because there is no accepted or standard measure to quantify the similarity of two complex fields such as the analytical and numerically propagated fields, it is difficult to quantify the improvement of the low pass filtered propagation over the unfiltered propagation. However, the ratio of RMS deviations of the unfiltered field relative to the analytical field and the filtered field relative to the analytical field was always greater than 2 for resolvable objects propagated beyond d_l or:

$$\frac{\left[\sum_i |U_{unfiltered,i} - U_{analytical,i}|^2\right]^{1/2}}{\left[\sum_i |U_{filtered,i} - U_{analytical,i}|^2\right]^{1/2}} \gtrsim 2, \tag{3.17}$$

where $U_{analytical}$ is the analytical field, $U_{unfiltered}$ is the unfiltered field, $U_{filtered}$ is the filtered field, and the sum over i is over every pixel in the field.

In summary, the Huygens-Fresnel Filtering propagation plus the low-pass filter method has the following advantages in holography applications,

1. It has proven accuracy compared with analytical propagation solutions for a wide range of parameter space.
2. It has a diffraction limited spatial resolution for all propagation distances greater than d_l.
3. It has uniform sample spacing independent of distance propagated to (Kreis et al., 1997; Schnars and Juptner, 2002).
4. It has a computation speed comparable to the fastest of other propagation methods. And,
5. It allows for the use of a single propagation method throughout a sample volume of a typical holographic particle measurement.

3.4 Particle Size Independent Sample Volume

When trying to measure size distributions and number densities (concentrations) of particles, it is desirable to have a sample volume independent of particle size, i.e. to have particle reconstructions appear the same independent of their position inside the sample volume. This is not a problem for particles larger than the diffraction limit at

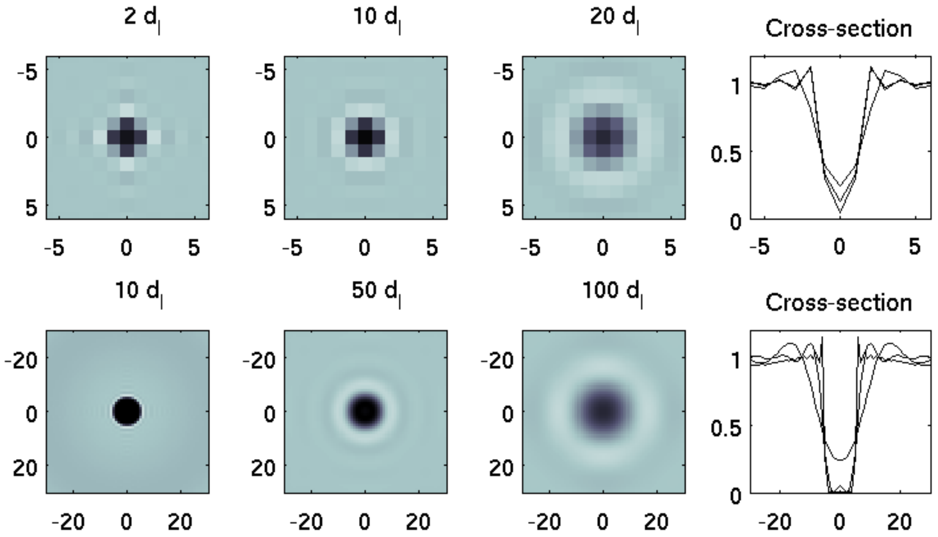

Figure 3.3: Shows a 2 Δx particle (top row) and a 10 Δx particle (bottom row) reconstructed to distances shown above each panel without a low pass filter. The cross section frames on the right show reconstructed intensity on a line through the center of each particle. Units on the axes are in Δx. Notice that the particle reconstruction quality falls greatly beyond the diffraction limit in the third from the left frames.

the furthest reconstruction distance in the sample volume. Particles smaller than this limit appear more visible nearer the camera. The low pass filter described in Section 3 calculated for the furthest distance in the sample volume can be applied to enforce the maximum diffraction limited resolution across all distances in the sample volume, making all particles appear the same throughout the sample volume. This will allow all particles to be detected with the same probability throughout the sample volume. Figure 3.3 shows a 2 and a 10 Δ_x particle reconstructed to 2, 10, and 20 d_l and 10, 50, and 100 d_l respectively without using the low pass filter. The simulation uses the same detector size and pixel size as was done to test the propagation method in Section 3. In these units, it is simple to figure that the diffraction limited resolution of 2 Δx is at 1 d_l away, and 10 Δ_x is at 5 d_l away. You can see in Figure 3.3 that the particles change their reconstructed shape as the distance increases and the shape changes greatly beyond their diffraction limit.

Figure 3.4 shows the same particle reconstructions except with a low pass filter calculated for the furthest distances shown respectively for each size particle in the figure. In this way, there is no variation in the way the intensity levels and edges of particles appear at different locations along the optical axis. Therefore particle detection methods that apply an intensity threshold or edge detection should detect them with equal success independent of depth along the optical axis.

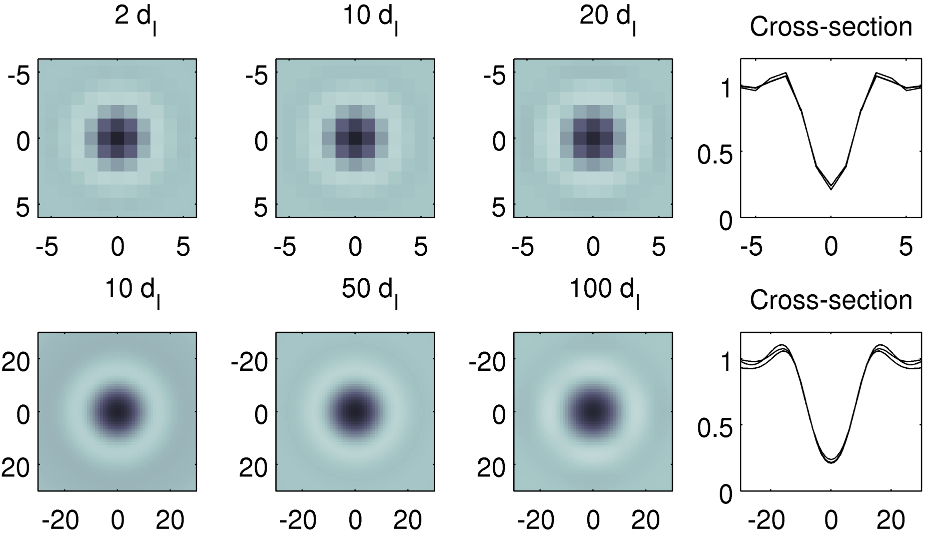

Figure 3.4: Same as Figure 3.3 except with the low pass filter described in the text.

3.5 Find the Particle Position on the Optical Axis

It is likely that there are more papers devoted to obtain a more precise particle position along the optical axis (hereafter referred to as the z-position) than any other problem in the field of particle holography.

3.5.1 Focusing Methods in the Literature

The traditional way is to find the center pixel of a particle and look for the z-position at which the reconstructed intensity peaks (Fournier et al., 2004; Pan and Meng, 2003). This method is the analog method for searching for particles in optical reconstruction (optically reconstruct holograms on film and scanning with a camera on a 3-axis stage, (Pu and Meng, 2000). This is the simplest method and is still commonly used. For example, it is used by Satake et al. (2007) in their parallelized algorithm to find particles in a hologram. Fournier et al. (2004) describe how larger particles have multiple bright peaks along the z-axis through the center pixel of the particle reconstruction, and how the shapes of the peaks and the distance between them can be used to find the z-position of the particle. Pu and Meng (2003) describe how the bright center spot of a particle is aberrated by light scattering of a sphere and not a circular disk (described by Mie theory) as is assumed in reconstruction.

Some methods take advantage of the fact that numerical reconstruction methods that have the complex amplitude of the field. Pan and Meng (2003) show it is possible to more accurately find the z-position of small particles by looking for a dip in the variance of the imaginary part of the field in a small patch about the particle along the z-axis. Their argument is if a particle are purely absorbing, the field must be real at the location of its shadow. Yang et al. (2005, 2006) relate ways to find the z-position of particles using a filter to decouple particle size and position information or reconstructing the phase signature near the particle location respectively. Coetmellec et al. (2002) give a way to find the particle z-position using the fractional Fourier transform and Soontaranon et al. (2002) use wavelets to find the z position. Soulez et al. (2007) present a solution via the inverse problem by iteratively positioning a particle, making a test hologram and comparing to the real hologram, and then refining the particle size and position. Raupach et al. (2006) used an unspecified âĂIJsharpness criterionâĂİ to focus their ice crystals. Ilchenko and Sattelmayer (2005) used the mean of the gradient in a small patch about the particle to focus their particles. Malkiel et al. (2004) use a edge sharpness method involving counting the number pixels with an edge sharpness higher than a threshold. All of the aforementioned papers, save Raupach et al. (2006) and Malkiel et al. (2004), present focusing methods that were designed with small particles in mind. Raupach et al. never presented their “sharpness criterion”. Malkiel et al. searched for particles greater than 50 μm in diameter and then optically reconstructed again with a band pass filter, masked the large particles and searched for particles for smaller than 50 μm in diameter.

3.5.2 Detection by Edge Sharpness

Our experimental conditions require that we have a particle search and finding method appropriate for finding the z-position of very large and small particles. If possible, we would prefer a single method than concern ourselves with biases in combining data from two different methods. Therefore the aforementioned methods will not suit our purpose as they fail when reconstructing larger particles with complex concave shapes like ice particles. We will show the ineptitude of an intensity based method of finding the z-position of large complex particles below.

We have developed a method that consistently is able to focus large and small particles in holograms. There may be more optimal methods, but to our knowledge they do not as yet exist. Yet as shown in Figure 3.5 through 3.7, and in Figures 10 and 11 in Fugal and Shaw (2007, see Ch. 4), we have empirical evidence that the described algorithm is robust and can consistently find the z-position of small circular particles or large complex-shaped particles.

The first element of our particle focusing algorithm is to find a statistic that identifies the focus of a particle regardless of its size (as long as it is resolvable) or its shape, although it need not be the optimal statistic for a particular size or shape. In visual inspection, the eye looks for the sharpness of the edge of a particle to see if the particle is in focus. Further, the statistic needs to be made as insensitive to noise by the virtual image immediately about the particle and other noise throughout the hologram such as background interference fringes. There are multiple ways to measure sharpness including gradient and LaPlacian operators. But in the presence of noise, the higher derivatives quickly lose their utility due to their sensitivity to noise. We chose to use the magnitude of the Sobel gradient filter in both the x and y directions. The numerical derivatives in the z direction are insignificant considering the z-spacing between reconstructed slices is usually much larger than the x and y pixel spacing. The filter has the construction:

$$\nabla_{Sobel} = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} \hat{\imath} + \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} \hat{\jmath}. \tag{3.18}$$

The filter is applied by summing over the element-by-element products of the filter and samples of the field. The edges of the field are padded with values from the adjacent row or column before applying the filter. The output of the overall filter is the root of the sum of the squares of the x and y ($\hat{\imath}$ and $\hat{\jmath}$) components of the Sobel gradient. Note that the extra elements of the Sobel gradient filter have a smoothing effect while the center row or column is emphasized over its neighbors. We apply the Sobel gradient filter to the complex amplitude of the field, not the intensity. In simulations, the peak using the complex amplitude is typically sharper than the same peak using intensity. Examples of the output of the Sobel gradient filter are shown in the top left panels in Figures 5 through 7 where the lower left panels are the intensity of the field over which the Sobel gradient filter was applied.

3.6 Profiles by Local and Global Standard Deviations

Figures 5 through 7 show profiles of the respective particle as a function of z-position. These profiles are based on different summary statistics of each reconstruction slice in z in which we may look for a peak to indicate a particle's z-position. We want to choose a summary statistic that works for very large and small particles. One property that is the same for large and small particles is they have a local neighborhood or patch for which, at the particle's z-position, the edge of the particle is sharpest. This implies that it is also where the Sobel gradient is strongest relative to its neighborhood, including the strong virtual image about large particles as it does not have such sharp edges. For smaller particles, it is also true that the intensities of the particle dominate over all else in its neighborhood. (This is similar to the argument Pan and Meng (2003) use to find the z-position of their small particles.) For larger particles, though, the real and virtual images, and perhaps the interference of the real and virtual images, have constructive fringes at positions other than the focus that are nearly as high as at the focus position. This makes it difficult for any intensity statistic to consistently tell us anything about the z-position of large particles. The intensity profiles in Figures 5 through 7 have a nice peak for the small particle, but do not consistently have distinct peak for the larger particles at the focus position.

The local profiles (red dashed lines) in Figures 5 through 7 are the standard deviation of the top 10% of the magnitudes of the Sobel gradient filter, and the top 25% of the intensities in the local patch, in the upper and lower panels, respectively, for each particle. Our reason for choosing the standard deviation and calculating the statistic only over the highest Sobel gradient magnitudes and brightest intensities is to increase the sensitivity of the profile to the extreme values. The standard deviation statistic is more sensitive to extreme values than the mean statistic. Also by taking the top 10% or 25% of values from the particle in every reconstruction slice in z before calculating the standard deviation for each patch, the profile is focused on the regions of z containing the highest values. This makes it simpler to program a peak search algorithm to find and focus on the correct peak.

The global profile (blue solid line) is calculated for each z position as:

$$\frac{\mathrm{std}(\bar{R})}{\mathrm{std}(\bar{R}-\bar{L})}, \tag{3.19}$$

where $\mathrm{std}(\square)$ is the standard deviation operator, $\bar{R}$ is the set of all pixel values in the reconstruction slice, and $\bar{L}$ is the set of local pixel values used to calculate the local profile and $\bar{R}-\bar{L}$ denotes the set difference between the $\bar{R}$ and $\bar{L}$ sets or $\bar{R}-\bar{L}=\{x : x \in \bar{R} \text{ and } x \notin \bar{L}\}$. Loosely speaking, the ratio of these standard deviations puts the pixel values of the local profile in the larger context of the whole reconstruction slice and has a stabilizing effect on the profile. As a larger complex-shaped particle goes in and out of focus, areas of the particle intensity image have edges and peaks that would otherwise make the Sobel gradient profile peak higher if they are not tempered by the statistical context of the entire reconstruction slice.

As is evident in the bottom ice crystal panel in Figure 3.5 and top ice crystal panel in Figure 3.7, the global profile (blue solid line) is a better indicator of position as long as the particle is away from the edges and other big ice crystals. And in many of those cases the global profile is a better indicator of position than the local profile. Sometimes the converse is the case as is evident in the bottom ice crystal panel in Figure 3.7.

Once the global profile is made, we fit a Gaussian to the highest peak and record the center of that Gaussian as the z-position of the particle. As illustrated in Figures 10 and 11 in Fugal and Shaw (2007, see Ch. 4), we find that this algorithm works very well for particles greater than 100 μm equivalent diameter. For example, reconstructing ice particle holograms from the IDEAS 3 Research Flight 2003-09-17 (Fugal et al., 2004, see Ch. 2), the algorithm found about 10% of the particles so far out of focus that they were not recognizable as particles. These are predominantly particles around 100 to 200 μm equivalent diameter and are near other particles or the edges of the hologram. Some of them are near enough to the camera that interference between the virtual image and the real image confuse the algorithm. Around 8% of the

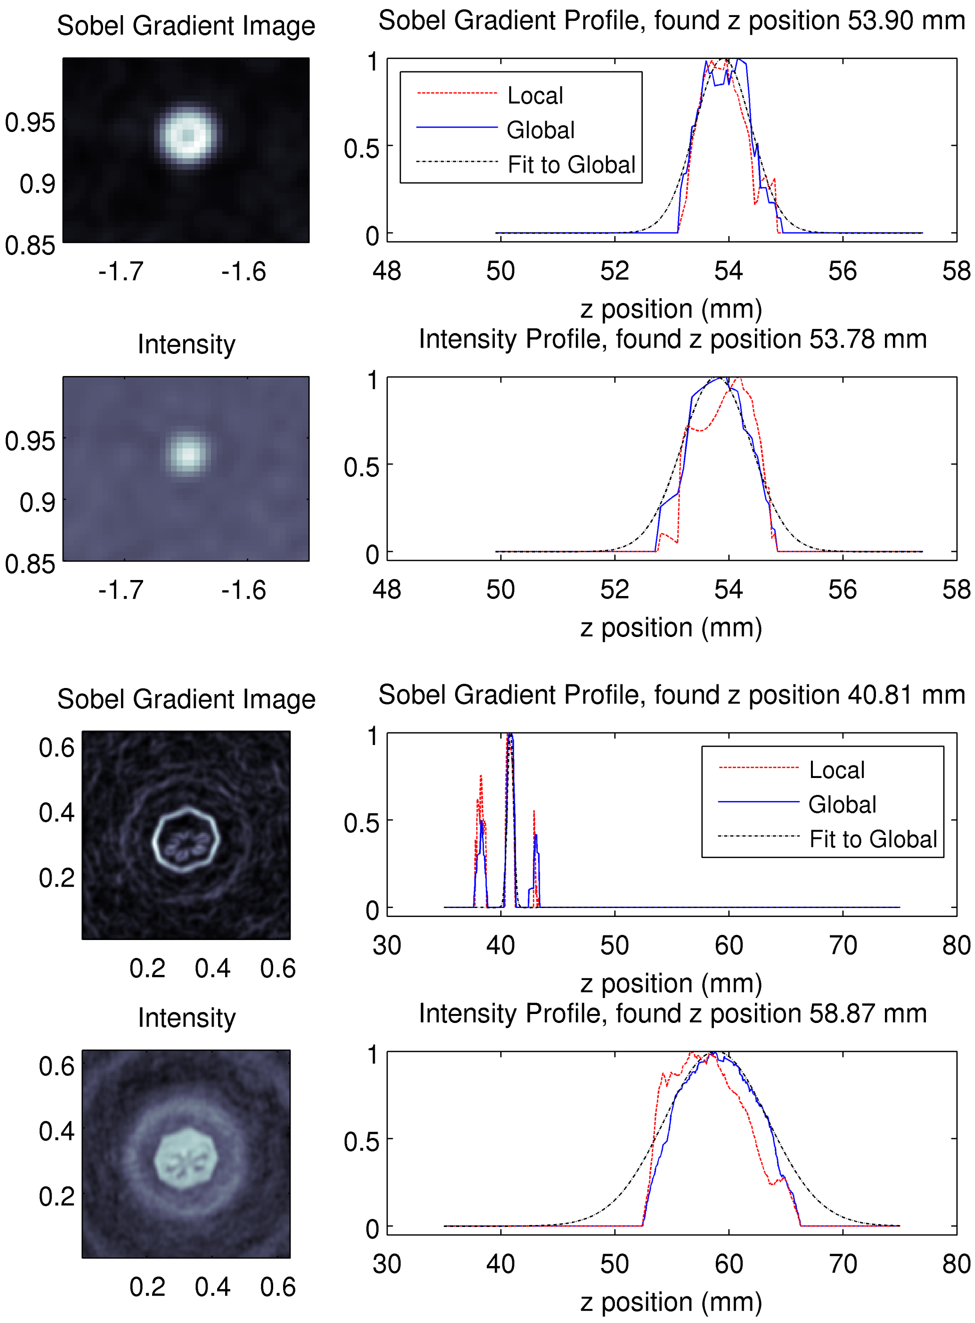

Figure 3.5: Shown are profiles as described in the text for two small ice crystals. All image axes are in units of mm. In the top panel, the intensity and the Sobel gradient profiles are approximately equal as is usually the case for small particles. In the bottom panel, the intensity and Sobel gradient profiles are not equal. The Sobel gradient is three-peaked, but the highest peak is selected and in this case is the correct peak. Note also that the global profile has its highest peak much higher than the other peaks where the local profile does not have such a contrast.

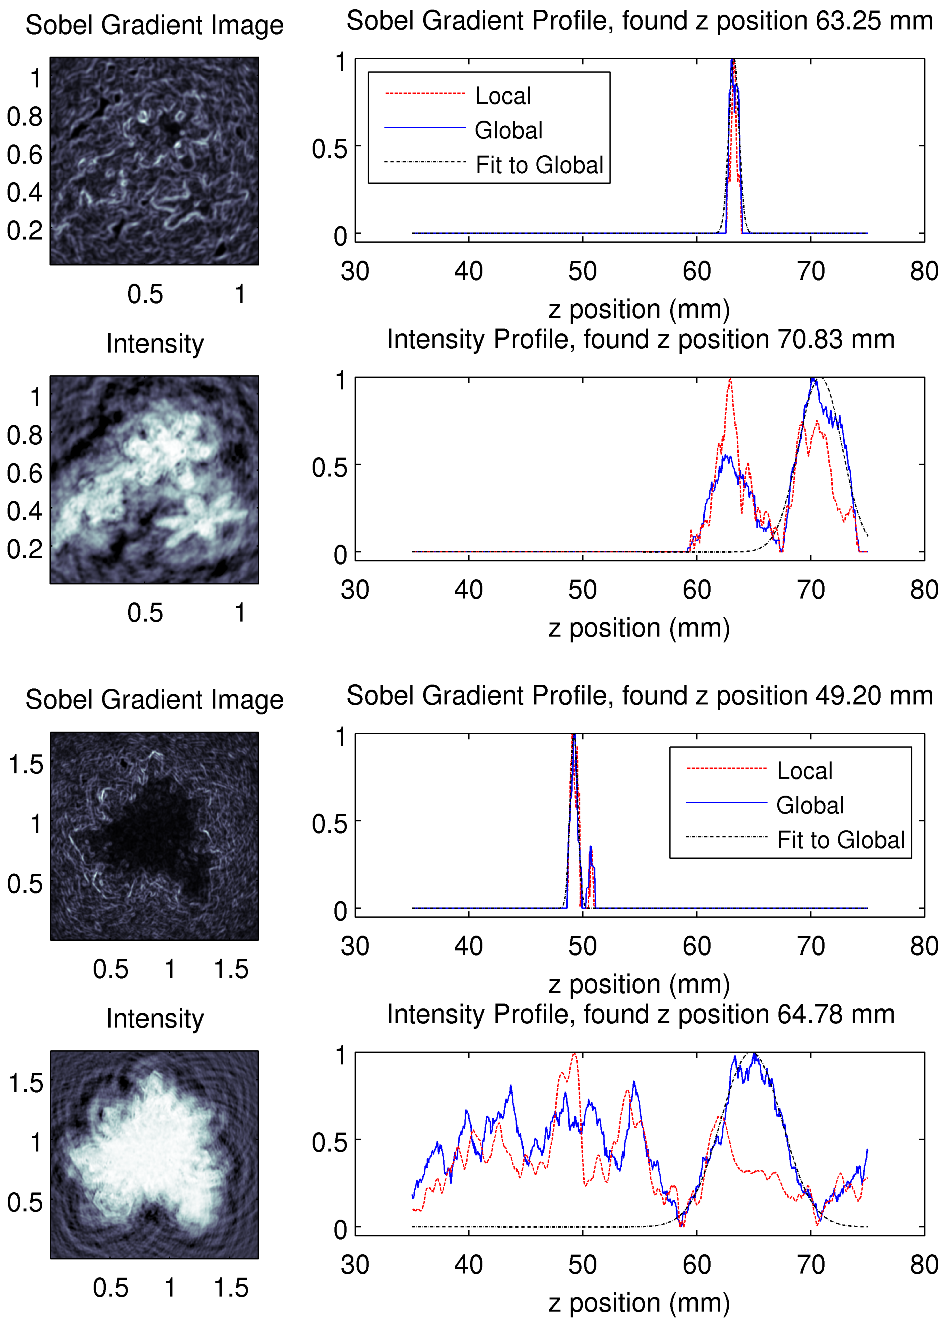

Sobel Gradient Image

Sobel Gradient Profile, found z position 63.25 mm

Intensity

Intensity Profile, found z position 70.83 mm

Sobel Gradient Image

Sobel Gradient Profile, found z position 49.20 mm

Intensity

Intensity Profile, found z position 64.78 mm

Figure 3.6: A continuation of Figure 3.5. Note again that the Sobel gradient and Intensity profiles do not peak at the same place for these larger ice crystals. In the top panel, the three ice crystals are focused at a compromise location where the three are most in focus together though not the optimal focus for the 3 individual crystals. The bottom panel is evidence that the global profile method works for the largest of crystals.

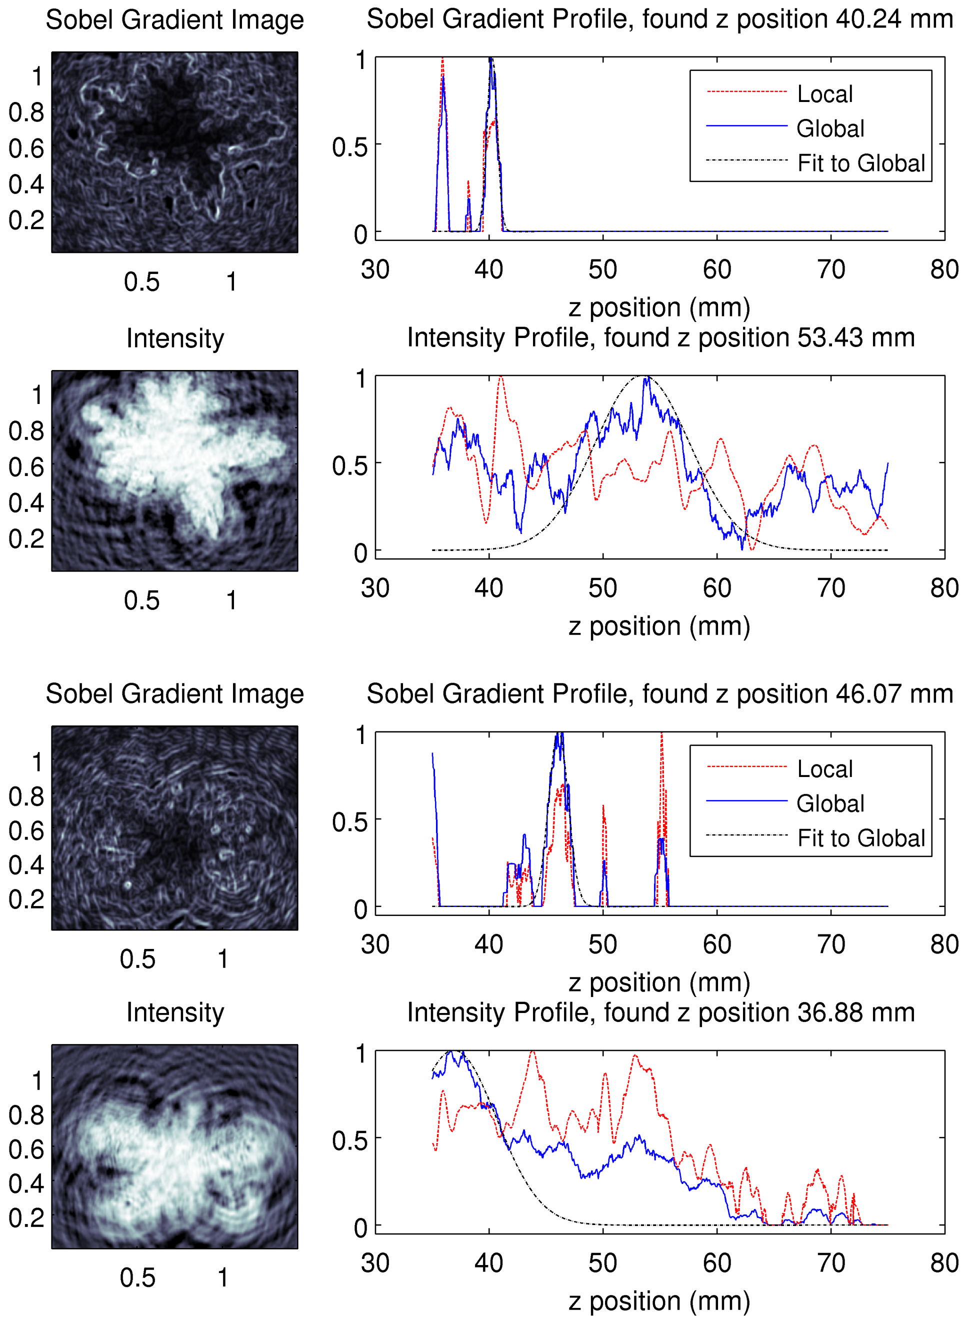

Figure 3.7: A continuation of Figures 3.5 and 3.6. Note in the top panel, the ice crystal is in focus because the global profile peaks higher at the focus position at about 40 mm and not at the other peak at about 37 mm. The inverse is true of the local profile. In the bottom panel is an ice crystal that is not in focus. The global profile peaks at about 46 mm where this image is focused. Visually, the focus is at 54.5 mm as shown in Figure 2.4 in Fugal et al. (2004, see Chapter 2) which corresponds to the peak in the local profile in the Sobel gradient profile plot. The largest reason for error is the particle lies on the edge of the hologram and so the information to reconstruct the particle is incomplete.

particles are located by the algorithm around 1 to 2 mm out of focus.

For small particles (a few pixels or less in diameter), the Sobel gradient filter has broader peaks than does the equivalent intensity profile. So for fields of small particles, such as in holograms of fluid flows seeded with tracer particles, it is computationally more efficient and at least as accurate to calculate and use only profiles of the intensity. We have not compared the accuracy of the aforementioned local and global profiles of intensity to other mentioned methods to find the z-position of small particles, but it is an area for future study.

3.7 Reconstructing on a Computer Cluster

Sections 2 through 6 have described an algorithm for reconstructing a hologram of large and small particles and finding the z-positions of the particles in order to obtain their size distributions and number densities. We have implemented this algorithm on a computer cluster for automated processing of holograms. In order to test the algorithm, we reconstructed thousands of holograms containing ice particles from the HOLODEC instrument and calculated size distributions and number densities. The size distributions and number densities from the holograms were within the uncertainty of measurements of another instrument using a different measurement method (Fugal and Shaw, 2007, see Ch. 4).

We implemented the algorithm in MATLAB (produced by The MathWorks, Inc.) code which allowed for faster algorithm development, shorter test cycle times, and quicker data and analysis visualization than in many other programming languages. Then with only minor code additions we were able to process thousands of holograms on a computer cluster using MATLAB distributed computing products.

3.7.1 Algorithm Implementation within Memory Limits

One point of difficulty in reconstructing holograms is how much memory the entire three-dimensional (3-D) set of reconstruction slices requires. For example, a typical reconstruction of one hologram from HOLODEC has 801, 1024x768, 16-byte complex double-precision numbers, spaced in 50 μm steps from 35 to 75 mm. As a single 3-D array, this would be 10 GB. To make this problem tractable, we implemented the algorithm by reconstructing one slice at a time. In particular, the algorithm implementation (a) reconstructs a single slice, (b) calculates the Sobel gradient values and intensities, (c) applies a threshold to the slice to the predetermined level, (d) dilates the image to include neighboring pixels around the pixels above the threshold, (e) groups neighboring selected pixels together, (f) matches the pixel groups of the current slice with pixel groups in the previously reconstructed slices, (g) merges the new pixel groups in with their respective matched groups, and then reconstructs the next slice and so on. Thus all that MATLAB holds resident in memory is a record of the selected voxels (three-dimensional pixels or volume elements) recorded as groups of neighboring voxels, and one reconstruction slice at a time. This implementation works well to minimize the amount of memory needed. Still, some large particles are so large that their recorded voxels nearly exceed the memory limits of the computer.

To allow the calculation of the standard deviation of the entire reconstruction slice except for values from the local patch of each particle, the algorithm saves a histogram of the intensity and of the magnitude of the Sobel gradient filter for each reconstruction slice. Then approximate standard deviations can be calculated for the entire reconstruction slice from the saved histograms. The advantage of a histogram is it uses so little memory relative to an entire slice. The standard deviation of the set difference of the entire slice and selected local pixels can be calculated by subtracting a histogram of the selected local pixels from the histogram of the entire slice and then calculating the standard deviation from the histogram difference.

3.7.2 Parallel Algorithms vs. Embarrassingly Parallel Algorithms

Satake et al. (2007) used an SGI Altix supercomputer with 128 processors and shared memory. Their parallel algorithm reconstructs a single hologram using multiple processors and holding all the reconstruction slices in memory at once, and then search for particles in the reconstructed 3-D volume. They say, “The algorithm can easily be used for a PC cluster machine by use of the MPI library. The program can be carried out for a high performance computing on-site experiment at an individual laboratory without having to access large-scale computers or supercomputer systems.”

While it is possible, it would be very inefficient to use this algorithm on a computer cluster because typical computer cluster interconnects (e.g. Gigabit ethernet) are orders of magnitude slower in shared memory access than in Satake et al.’s SGI Altix supercomputer. This is especially true for their “axis rotation” step which involves shuffling large amounts of data around in shared memory and would take a very long time (longer than the total computation time) over a Gigabit network. We have implemented our algorithm in an embarrassingly parallel fashion where one core of a processor reconstructs one hologram at a time independent and asynchronous of every other processor core. The speed up advantage gained by a cluster of computers is therefore strictly linear with the number of processor cores as long as the shared filesystem is not overburdened with loading code and saving results.

3.7.3 Hologram Reconstruction and Memory-Bound Computation

One other consideration is that hologram reconstruction, particle searching, and calculating measurements of particles are "memory-bound" computations in that a large amount of the processor time is spent waiting for data to be transferred from RAM to the memory cache on the CPU. This is because the CPU core can compute the necessary computations on data in its onboard memory cache faster than the memory bus can put the next block of new data in the onboard cache. Therefore to speed up the computation, the memory bandwidth or the speed in which the memory bus can transfer data must be increased. In contrast, this is not a problem for "computation-bound" problems like matrix multiplication or matrix inversion. Therefore, little advantage is gained for hologram reconstruction by having more CPU cores on a single chip without an equivalently higher capacity memory bus.

In the production run on our computer cluster for the data presented in Fugal and Shaw (2007, see Ch. 4), we reconstructed about 18000 holograms on a computer cluster with 20 core Intel Xeon 5060, 40 GB of RAM, running Ubuntu 6.06 Linux and MATLAB 2007a. The average time to reconstruct a hologram, find the particles and calculate their statistics was 25 minutes. Holograms with large ice crystals sometimes took more than an hour and the holograms void of particles sometimes took as little as 20 minutes to reconstruct. Both data sets took in all a little less than 15 days for the cluster to process.

3.8 Summary and Future Work

In summary, we have developed an algorithm for reconstructing holograms of large and small particles that has been proven against analytical solutions of light propagation. The reconstruction algorithm has a sample spacing that does not vary with distance propagated to. It can reconstruct beyond the distance at which the phase in its kernel undersamples the highest frequencies by a low-pass filter that filters out the undersampled frequencies. This allows us to reconstruct the entire sample volume using one reconstruction algorithm.

We use the magnitude of the Sobel gradient filter to find the z-position when each particle is in focus. We use both global and local noise statistics to distinguish the correct peak between the multiple peaks large complex-shaped particles produce as they go in and out of focus. And we have implemented these algorithms on a computer cluster to process thousands of holograms in a reasonable time.

We are currently using the reconstruction algorithm to reconstruct holograms of small water droplets in a turbulence chamber. We are using the global profile of these particles intensities to find each particles z-positions. In the future we expect to further compare this method of finding the z-position of small and large ice crystals with other published and unpublished methods. We also expect to apply these algorithms to other data sets from future HOLODEC instruments.

Acknowledgments

This work is supported by an NSF Graduate Research Fellowship, a NASA Earth System Science Fellowship, a Michigan Space Grant Consortium Graduate Fellowship, and NSF Grant ATM-0535488.

Chapter 4

Ice Particle Size Distributions Measured with an Airborne Digital In-line Holographic Instrument

Holographic data from the prototype airborne digital holographic probe, HOLODEC (Holographic Detector for Clouds), taken during test flights are digitally reconstructed to obtain the size, three-dimensional position, and two-dimensional profile of ice particles and then ice particle size distributions and number densities are calculated using an automated algorithm with minimal user intervention*. The holographic method offers the advantages of a well-defined sample volume size that is not dependent on particle size or airspeed, and offers a unique method of detecting shattered particles. The holographic method also allows the volume sample rate to be increased beyond that of the prototype HOLODEC instrument, limited solely by camera technology.

HOLODEC size distributions taken in mixed-phase regions of cloud compare well to size distributions from a PMS FSSP probe also onboard the aircraft during the test flights. In regions of cloud with pure ice, the HOLODEC size distributions compare better to FSSP size distributions corrected for non-spherical ice particles. A conservative algorithm for detecting shattered particles utilizing the particles' three-dimensional positions eliminates the obvious ice particle shattering events from the data set. Resulting size distributions are reduced by approximately a factor of two for particles 15 to 70 μm in equivalent diameter, compared to size distributions of all particles. The comparison with the FSSP under cloud conditions with no large ice particles provides an estimate of the magnitude of shattering biases in the FSSP and a correction for the FSSP's

---

* By Jacob P. Fugal and R. A. Shaw. In preparation to be submitted to the Journal of Atmospheric and Oceanic Technology (Fugal and Shaw, 2007).

calibration for particle sizing.

4.1 Introduction

Accurate ice particle size distributions and number densities are necessary for understanding and modeling cloud processes such as precipitation formation and radiative transfer, and for validation of remote sensing and satellite measurements. Many methods and instruments have been devised to measure ice particles, yet there is still considerable uncertainty in measuring small (less than about 100 μm) ice particles (Baum et al., 2005a; McFarquhar et al., 2007). Beyond the inherent uncertainty in counting statistics, the uncertainty in small ice size distributions results primarily from poorly defined sample volumes, instrument resolution limits, and instrument-induced ice particle breakup. Digital holography is one of several approaches that allows for improvements in the measurement of ice size distributions. In this paper, we present results from the Holographic Detector for Clouds (HOLODEC), which is a prototype airborne digital holographic instrument. In relation to the existing uncertainties, holography has the benefit of providing a well-defined sample volume, a uniform and well-defined resolution, and three-dimensional spatial information that can assist in identifying shattered crystals. The difficulty of using digital holography is the added complexity in data processing, which includes digital reconstruction and particle detection. Therefore, the essential question we address here is, can digital holography be used to successfully measure ice size distributions in an automated fashion?

The holographic method is not new to atmospheric science and, in fact, some of the earliest applications of optical holography were to measure atmospheric particles (e.g. Thompson (1974)). Subsequently, analog holography (i.e., recording with photographic emulsions) was applied by various groups, including in an airborne setting (Kozikowska et al., 1984; Brown, 1989). More recent holographic instruments for atmospheric particle studies include the large-sample-volume HODAR (Borrmann and Jaenicke, 1993) which uses photographic plates, a prototype digital holographic instrument (Lawson and Cormack, 1995), and a new ground-based instrument utilizing two digital cameras in a cross-beam configuration (Raupach et al., 2006). A challenge that remains, however, is to take holography to the point where large numbers of ice crystals can be reconstructed and ice size distributions determined in an automated fashion suitable for regular field use.

The HOLODEC probe measures ice particle size, three-dimensional (3-D) position, and two-dimensional (2-D) profile via digital in-line holography (Fugal et al., 2004, see Chapter 2). It has the advantage of a sample volume that is not particle-size dependent, nor air-speed dependent. It also has a unique ability to detect probable crystal-shattering events utilizing 3-D positions of cloud particles and the fact that shattered particles tend to appear as highly localized clusters in certain parts of the sample volume. The HOLODEC instrument flew aboard the NCAR C-130 aircraft in the IDEAS 3 project. There it established proof of concept that a digital in-line holographic instrument could detect cloud particles, both ice and water, in the size range of about 15 μm up to around 1 mm (Fugal et al., 2004, see Chapter 2). Since then, software routines for automated hologram reconstruction and size distribution calculation have been developed.

To provide a context for the measurements we briefly describe some of the limitations of commonly used optical particle detectors for measuring ice crystal size distributions. First, there is a particle-size-dependent depth of focus and sensitivity for small particles in optical array probes such as the PMS 2D-C (Strapp et al., 2001; Korolev et al., 1998; Baumgardner and Korolev, 1997). Second, light-scattering instruments, e.g. the Forward Scattering Spectrometer Probe (FSSP), are calibrated only for spherical particles. Third, some instruments, e.g. the Cloud Particle Imager (CPI), yield uncertain ice particle size distributions (Baum et al., 2005b) because severe air stream distortion and triggering mechanisms make it difficult to estimate the effective sample volume. Fourth, nearly all probes suffer in some degree from ice particles shattering on the leading parts of the instrument housing. When the shards are swept into the sample volume, the result is an over counting of small particles and an under counting of large particles (Field et al., 2006, 2003b; Korolev and Isaac, 2005; Gayet et al., 1996; Gardiner and Hallett, 1985; McFarquhar et al., 2007). Other errors such as air stream distortion caused by the aircraft, and calibration errors have also been discussed in the literature (Nagel et al., 2007; Korolev, 2007; Twohy and Rogers, 1993).

Remedies for some of these problems exist: for example, given certain assumptions FSSP size measurements can be corrected for ice, effective instrument resolution can be improved via post-processing (Korolev, 2007), and instrument housings can be modified to reduce shattering (Field et al., 2003b). Furthermore, new instruments are being developed to measure small ice particles without some of these problems, such as the Small Ice Detector (SID) (Field et al., 2003a) and the SPEC 2D-S (Stereo) probe (Lawson et al., 2006). The SID probe measures light scattered by ice particles at many angles, and can yield ice particle size and crystal habit within the size range of approximately 1 to 50 μm. The SPEC 2D-S (Stereo) optical array probe can measure cloud particle size and a two-dimensional (sometimes three-dimensional) profile in the size range of about 10 to 1000 um.

In this paper we analyze holograms of ice particles to show that digital inline holography and hologram processing algorithms can measure accurate ice particle size distributions after rejecting shattered ice particles, and that the results are comparable to standard aircraft probes under certain conditions. This is accomplished with minimal user intervention, and therefore is suitable for field measurements. In Section 2 we briefly summarize the automated hologram reconstruction, particle finding and measuring algorithms, and discuss the sampling characteristics of the HOLODEC probe relevant to measuring ice. In Section 3 we discuss the conditions under which we took the ice holograms and we present the results in Section 4. In Section 5, we discuss implications of our results and summarize.

4.2 Holographic Reconstruction Method and Instrument Parameters

In this section, we briefly describe how we reconstruct holograms, detect particles, and determine their three-dimensional location and two-dimensional cross section (and size). We also explain the particle-size detection limits of HOLODEC and their causes, as well as limits on its ability to measure ice particle size distributions and number densities.

4.2.1 Hologram Reconstruction and Particle Finding Method

Holograms from the data sets presented in this paper are reconstructed using the method outlined by Fugal et al. (2007, see Chapter 3). and we only briefly summarize it here. As a preliminary, we review the instrument parameters described in more detail by Fugal et al. (2004, see Chapter 2). The holograms are taken with a 527 nm wavelength pulsed laser, and a 1024x768, 4.65 μm pixel, 8-bit grayscale camera with particles confined to approximately 30 to 80 mm from the CCD chip of the camera. The laser pulse width of 20 ns and flight speed of about 100 m s-1 limits the particle motion in the hologram to about 2 μm or less than half a pixel width. The holograms are recorded in a raw format in real time and later compressed to a lossless image format (.png) and stored for subsequent processing.

Each hologram is divided by a background image made by taking the median of seven of its neighboring holograms in time. This step removes most of the large-scale background intensity variations (e.g., due to non-uniform laser beam profile) and interference fringes from dust and other particles contaminating the optics and windows. No other filtering is done to remove large-scale background intensity variations (i.e. such as dividing by or subtracting a local mean or median pixel value) because such filtering tends to destroy the reconstructed images of ice particles larger than the neighborhood over which a local mean or median is calculated. Finally, we take the contrast negative of the image so the ice particles appear as bright shapes on a dark background.

Holograms are reconstructed using the Huygens-Fresnel filtering method described in (Fugal et al., 2007, see Chapter 3). We use the low-pass filter detailed by Fugal et al. (2007, see Chapter 3), to enforce a minimum particle-size detection limit uniformly throughout the sample volume. We apply a threshold to reconstructed intensities in order to select bright voxels (volume elements or volume pixels) that likely compose particles, and then group neighboring voxels together into individual particles. The position of each particle along the optical axis is then found via the edge sharpness algorithm detailed by Fugal et al. (2007, see Chapter 3), and at that axial position, the centroid in the other two dimensions is calculated as well as an equivalent area diameter. The edge sharpness algorithm is robust in finding the particle's axial position over a range of particle sizes extending from 15 μm to 1 mm (Fugal et al., 2007, see Chapter 3).

Particles are selected from noise by examining each particle's position, peak intensity, and size. These criteria are effective because false particles (bright spots resulting from constructively interfering fringes such as from out-of-focus particles, window contaminants, specular reflection off the surface of the optics, etc.) occur with greatest frequency near the edges of the sample volume and in highly localized regions ($<$ 1 mm-3) inside the sample volume and tend to be fainter and smaller than real particles. As with any other optical detector, it is ultimately the noise level that that determines how small an ice particle HOLODEC can reliably detect and count.

We detect ice particle shattering by taking advantage of HOLODEC's ability to measure the 3-D particle positions inside the sample volume. Shattered particles tend to congregate in highly localized clusters, often near the windows of the instrument, whereas natural ice particles are detected throughout the sample volume. Figure 4.1 shows a hologram reconstructed at a plane (perpendicular to the optical axis) where a disproportionate number of ice particles appear in or nearly in focus as can be seen in the histogram of the location of ice particles along the optical axis. Note the high concentration of particles around 66.5 mm which is the distance at which the hologram is reconstructed. To detect holograms containing shattered ice crystals, we search for high concentrations of particles along the optical axis. We acknowledge there are likely more optimal (and necessarily more complex) methods to detect shattering, but in this paper we have chosen a simple method with a conservative threshold to detect and remove only the most obvious shattering events. In fact, only 6% of the holograms (or sample volumes) in the entire data set are marked as shattered, but they contain 60% all of the particles in the data set.

The top panel of Figure 4.2 shows an example of an ice particle in the process of breaking up. Again, the localized cluster of particles allows them to be identified as resulting from a shattering event. The lower two panels show the effectiveness of this simple shattering detection scheme. The center panel shows a histogram of the position along the optical axis in the sample volume of all detected ice particles and the lower panel shows a histogram of all the particles not in holograms flagged as containing shattered particles. Note that most of these shattered ice particles are near the edges of the sample volume, or near the windows and tips of the probe where particles impact, shatter are then their shards are swept into the sample volume near the probe arms. Further, the bottom panel shows a nearly uniform detection rate for particles of all sizes across the optical axis after rejecting shattered ice particles as we expect to see from a nearly uniform distribution of particles.

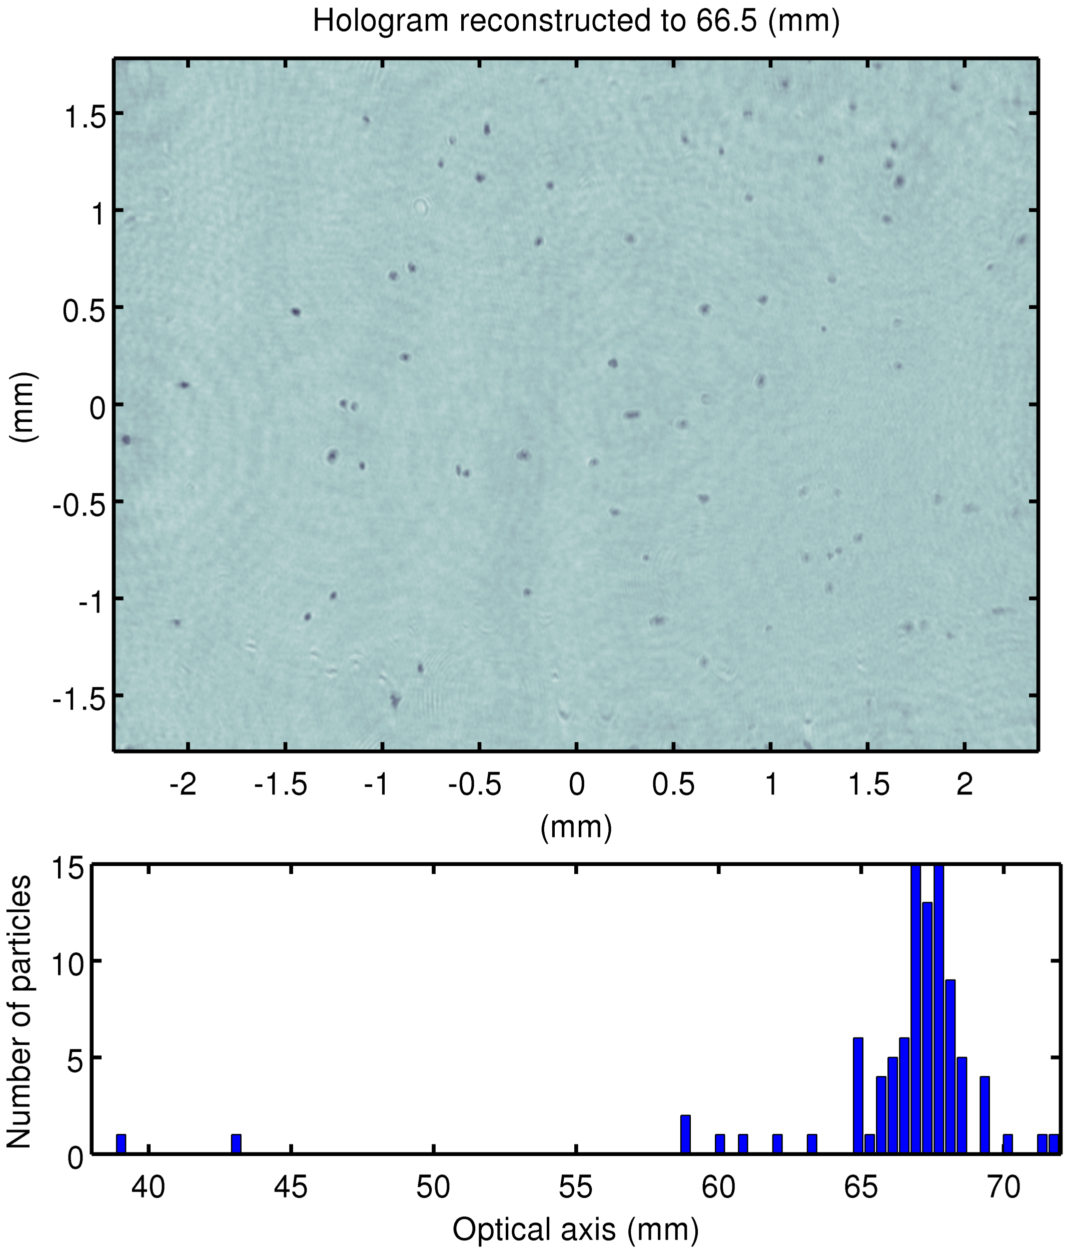

Figure 4.1: The top panel shows a hologram from Research Flight 2003-09-17 reconstructed to 66.5 mm. The reconstructed image has many particles in focus at one depth and many others nearly in focus. The bottom panel is a histogram of the z-positions of all particles in the sample volume. This is suggestive of shards of a single impacted crystal tracing an aerodynamic surface about the instrument housing.

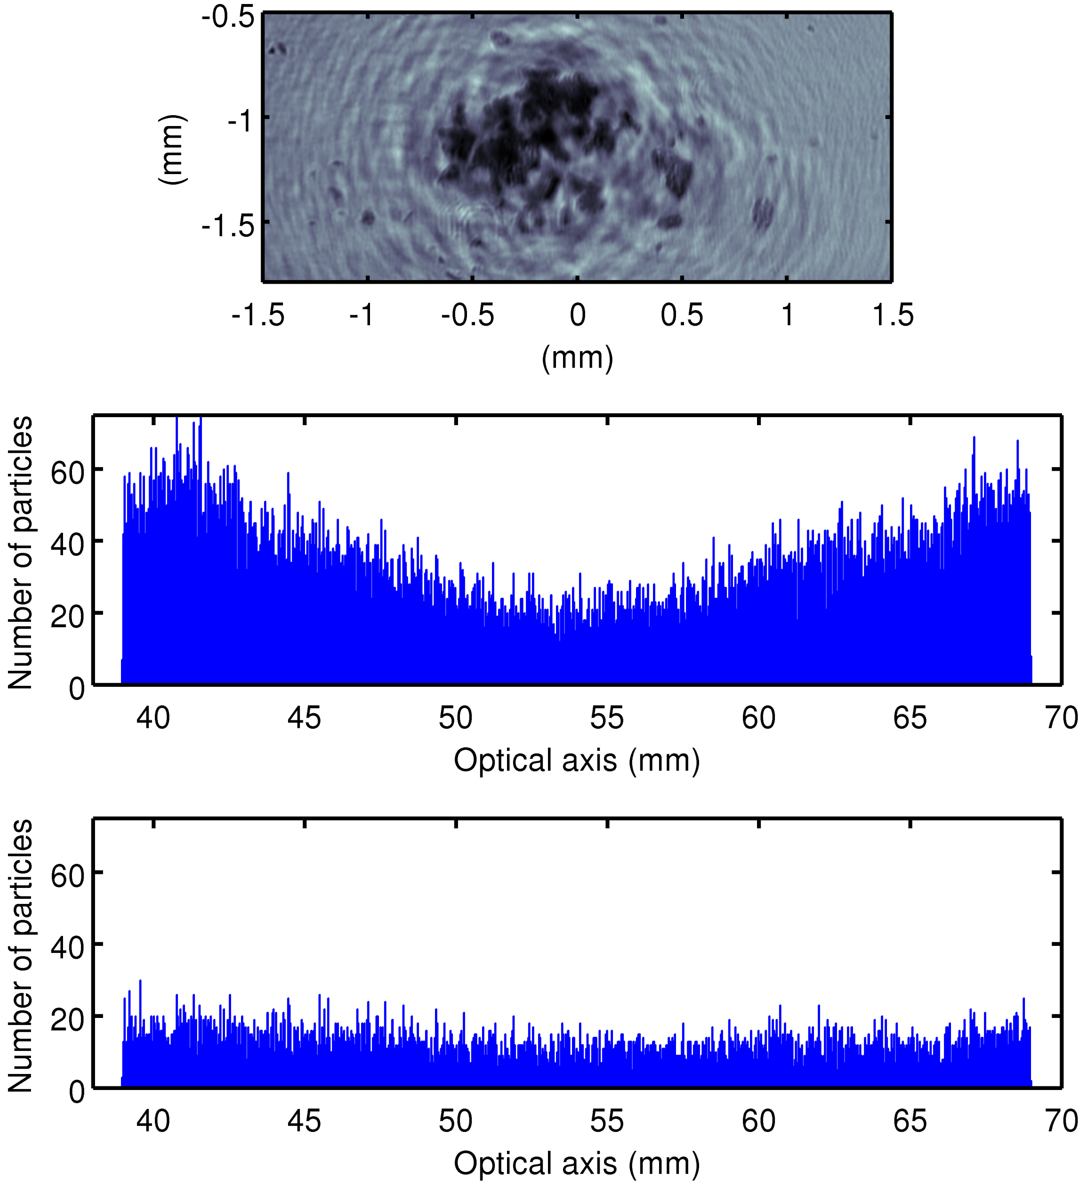

Figure 4.2: The top panel shows an ice particle in the process of shattering lying about 10 mm from one of HOLODEC's windows. Shown in the center and bottom panels are histograms of optical-axis positions of particles from Research Flight 2003-09-17. The center panel shows all particles and the bottom panel excludes particles in holograms containing shattered particles. Note that rejecting shattered particles yields a more uniform distribution of particles along the optical axis.

4.2.2 Instrument Parameters

We now explain what determines HOLODEC's sampling characteristics, including sample volume, volume sample rate, and detectable particle size range. The sample volume is determined by the camera CCD area that is not affected by edge effects (e.g. particle images overlapping the edge of the hologram) and the space between the windows of the probe housing that is not severely affected by airstream distortion or the boundary layer around the instrument housing. For HOLODEC this is about 0.4 cm^3. In estimating the sample volume, we have applied no correction for the distortion of the airstream from the leading probe parts. Simulations of the flow (using Fluent software) based on a simplified probe body indicate that the sample volume is compressed along the optical axis, and stretched in the other two dimensions. The compression and stretching do not completely cancel each other out, but the effect is on the order of 10% of the sample volume size. We do restrict the sample volume to be away from the probe housing so as to minimize the effect of the boundary layer near the probe. HOLODEC's volume sample rate is not dependent on air speed, but on the sample volume size and the frame rate of the camera, provided the frame rate of the camera is slow enough that the sample volumes do not overlap. For a 100 m s^{-1} flight speed, HOLODEC records one sample volume every 7 m and its sample rate is about 5 ml s^{-1}. For comparison, the FSSP probe sample rate is around 60 ml s^{-1}. We note that HOLODEC is a prototype instrument and a production version of the instrument would have a faster or larger camera with a much greater volume sample rate.

The minimum particle detection size is approximately 15 μm or about 3 pixels wide. This limit is determined by the greater of the diffraction limit or a few pixels wide (Fugal et al., 2007, see Chapter 3). It is a larger size when there is more noise from contaminated windows, or a high density of particles. The maximum detection size is around 1.5 mm (or some fraction of the detector size) as determined by automated hologram processing code's ability to reliably bring such a large particle in focus (Fugal et al., 2007, see Chapter 3).

Finally, as with every instrument that counts particles, HOLODEC's ability to measure ice particle size distributions and number densities is limited by the uncertainty inherent in counting statistics. As larger particles tend to have much smaller number densities, this means that while HOLODEC can detect and measure the size of larger ice particles, its sample rate limits its ability to reliably measure the number density of very large ice particles. Again, a faster or larger camera in HOLODEC would yield a much higher volume sample rate and more accurate counting statistics. The number density error bars shown in subsequent plots are one standard deviation assuming Poisson counting statistics.

4.3 Data Sample

HOLODEC flew onboard the National Center for Atmospheric Research C-130 Q Hercules research aircraft over Colorado on 9 test flights in August and September 2003. This was an instrument-testing campaign (Instrument Development and Education in the Atmospheric Sciences Phase 3 âĂŞ IDEAS 3) and therefore much of the flight time was devoted to debugging and evaluation. From the subset of times when HOLODEC was collecting clear holograms, we select times during which there is a likelihood of finding ice particles as indicated from data taken with the FSSP and 2D-C instruments. We analyzed about 10 minutes (8100 holograms) of data taken on the 2003-09-17 Research Flight, and 18 minutes (9900 holograms) of data taken on the 2003-09-05 Research Flight. Together the data consist of approximately 18,000 holograms reconstructed to find around 75,000 particles in the size range 15 μm to 980 μm equivalent area diameter.

We reconstructed the 18000 holograms on a computer cluster (20 core Intel Xeon 5060, 40 GB of RAM, with the Ubuntu 6.06 Linux operating system, and MATLAB Distributed Computing Engine to manage jobs on the cluster) with the algorithm coded in MATLAB R2007a (Fugal et al., 2007, see Chapter 3). There was no parallel computation involved; the reconstructions are done with one hologram on one CPU core independent and asynchronous of every other CPU core. On average, it took 25 minutes to reconstruct each hologram, find the particles and calculate their statistics. Holograms with large ice particles sometimes took more than an hour and the holograms void of particles took as little as 20 minutes to reconstruct. Both data sets in all took a little less than 14 days to process on this cluster.

4.3.1 Cloud Conditions in Data Sample

Figure 4.3 shows 2D-C and FSSP number densities for both data sets as well as times at which HOLODEC recorded good holograms. These time periods are of interest as the low number densities of the 2D-C instrument indicate that there are few large ice particles to shatter on the leading probe parts, and the high number densities of the FSSP indicate we should have measurable number densities of ice particles to compare between the FSSP and HOLODEC probes. Figure 4.4 shows total and liquid water content as measured by the Nevzorov probe showing when we expect nearly all ice or mixed ice and liquid water cloud particles. (On both Figures 4.3 and 4.4, the bars identify the times for which panels in Figures 6, 7 and 8 show size distributions.)

4.3.2 Corrections to FSSP Sizing of Ice Particles

The particle sizes from HOLODEC are calculated by counting the pixels the particle contains at its focus and calculating an equivalent diameter. The FSSP measures par-

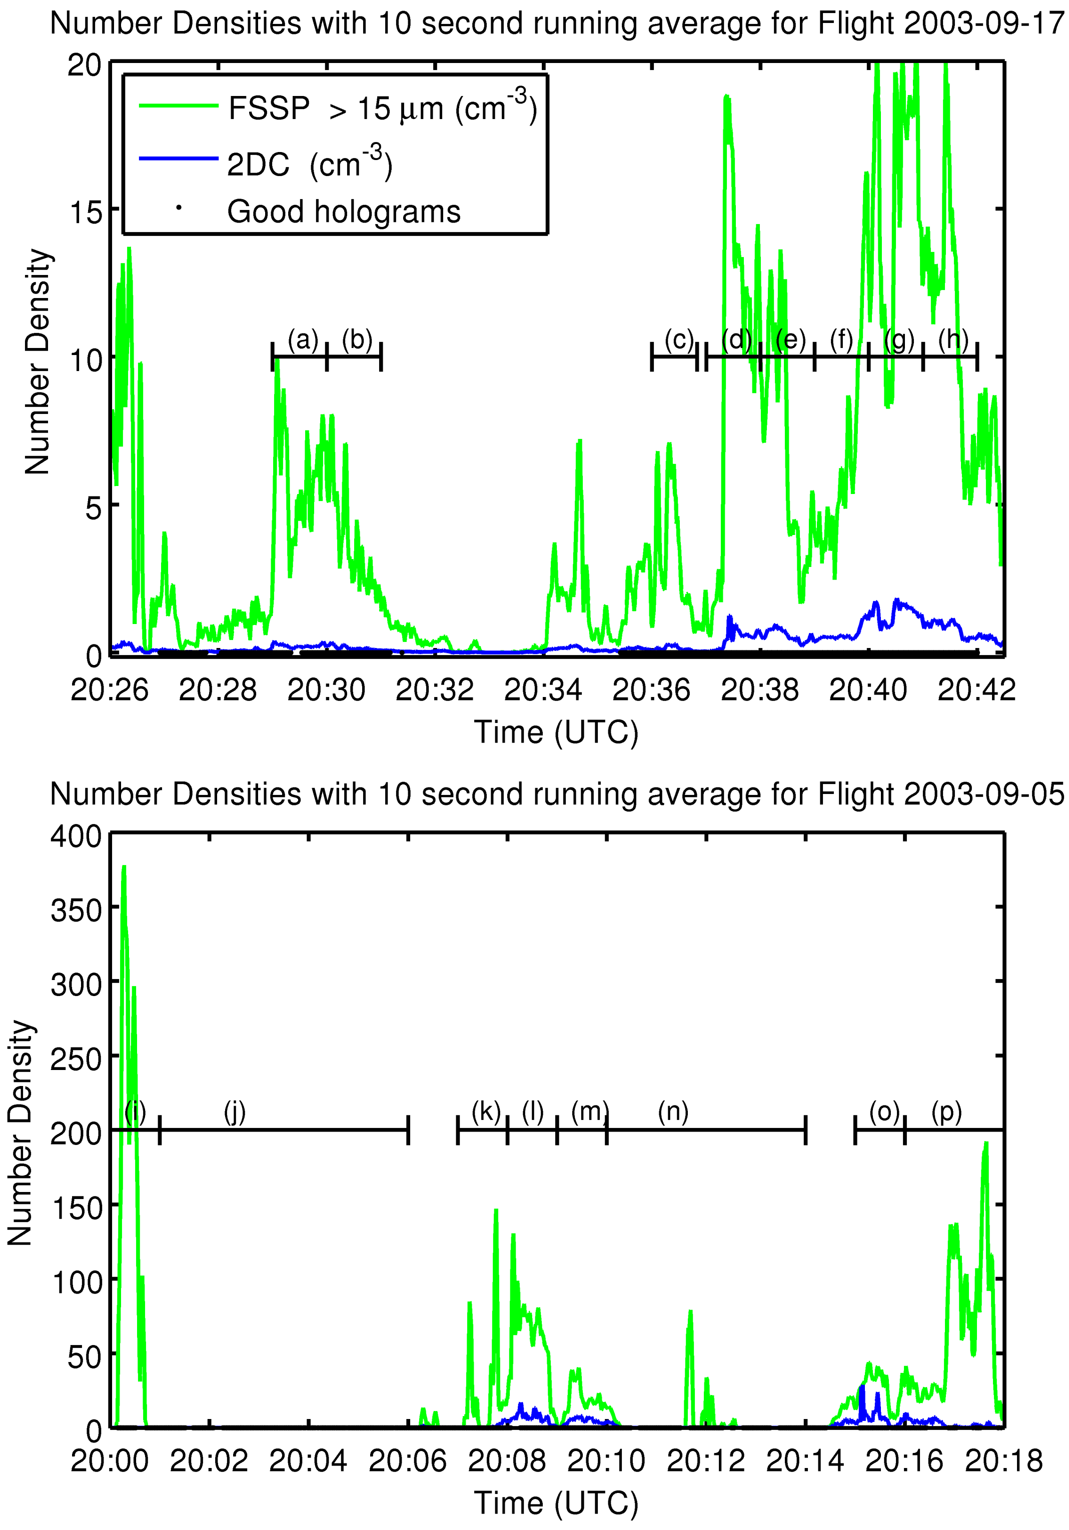

Figure 4.3: Particle number densities from the FSSP and the 2D-C probes from the 2003-09-05 and 2003-09-17 Research Flights. Also shown are the times where HOLODEC recorded clear holograms. The labels in brackets from (a) to (p) correspond to size distribution panels in Figures 4.6 to 4.8.

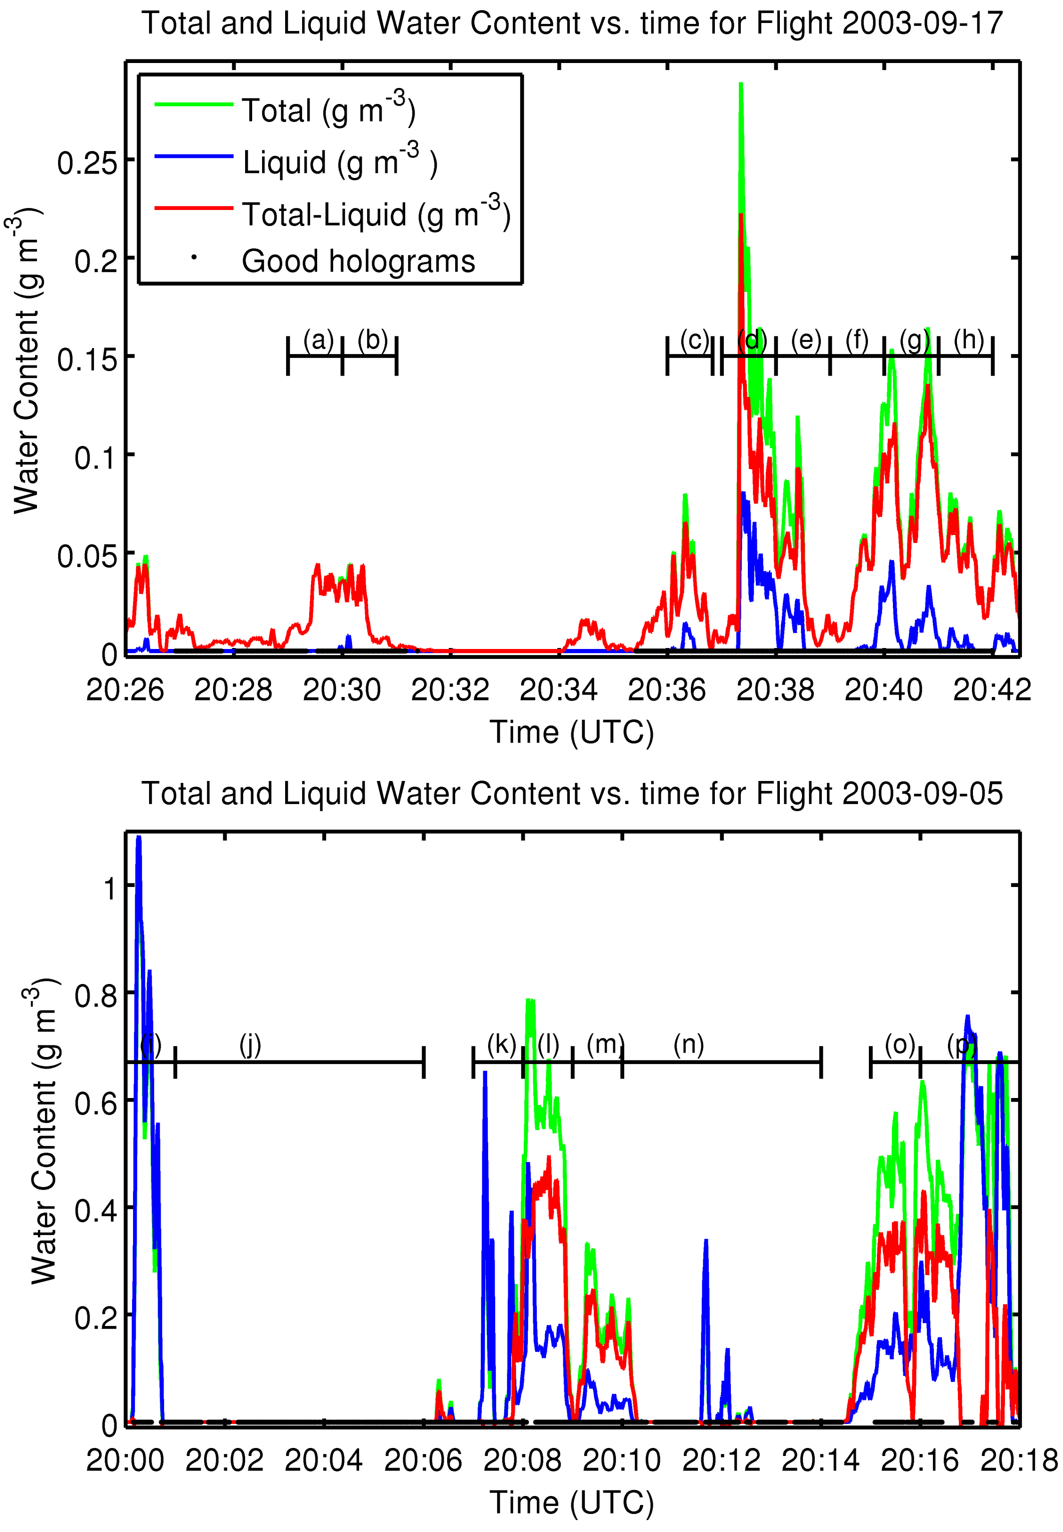

Figure 4.4: Total, liquid, and ice (total minus liquid) water content for the 2003-09-05 and 2003-09-17 Research Flights. The labels in brackets from (a) to (p) correspond to size distribution panels in Figures 4.6 to 4.8.

ticle size by measuring the intensity of forward scattered laser light collected between 5O and 14O. The size of the particle is inferred by calculating, for a given particle size, the intensity of the integrated forward scattered light assuming a spherical shape and an index of refraction for liquid water. Ice particles, however, are not spherical even in the 1 to 10 μm diameter range. A reasonable approximation for the shape of small ice particles is a droxtal which resembles a faceted sphere (Field et al., 2003a). Using a library of light scattering calculations for small droxtals provided by Dr. Ping Yang (Yang et al., 2003; Zhang et al., 2004), we calculated a correction of forward scattering from spherical water drops to droxtals of the same size. The scattering curves as a function of particle size are shown in the top panel of Figure 4.5. The bottom panel shows the size correction between spherical droplets and droxtals. As a heuristic, note that non-spherical particles scatter less light in the forward direction than spherical particles do. Thus the FSSP detects less light for a non-spherical particle and underestimates its size. A correction for this effect will shift the FSSP size bins calibrated for spherical particles to larger sizes for non-spherical particles (see lower panel of Figure 4.5).

4.4 Results

Figures 4.6 and 4.7 show size distributions from 8 different time periods in Research Flight 2003-09-17 comparing HOLODEC, FSSP, and 2D-C data. The panels on the left show the FSSP probe with its uncorrected sizes while the right side panels show the FSSP with sizes corrected for ice particles as described in Section 3. The time periods have been chosen to illustrate the performance of HOLODEC under a variety of conditions ranging from nearly all liquid to all ice. Notice that the 2D-C size distributions in (a) through (d) are always lower than HOLODECs but are usually within the error bars. The 2D-C distributions don't appear in (e) through (h) because they are below the lower number density limit on the plots. Note that for some time periods such as (a) through (e) the ice corrected sizes match best and for others such as (f) through (h) the uncorrected FSSP sizes match better. Also note in the data from the Nevzorov probe from the top panel in Figure 4.4, that panels (a) through (c) are nearly pure ice and so we would expect the corrected size distributions on the right to match better. Panels (d) through (h) are mixed phase but with more ice than water. In mixed-phase clouds, one ideally should partition ice and liquid water as a function of particle size. But without some assumption to justify a particular partition, there is no rationale for determining what weighting of uncorrected and corrected FSSP sizes to use.

Figure 4.8 shows ice particle size distributions from 8 selected times in Research Flight 2003-09-05 and only shows uncorrected FSSP sizes. Nearly the entire data set is mixed phase and in this case the corrected FSSP distributions are not a better match than the uncorrected FSSP sizes for any panel. Note that there is fairly good agreement for each size distribution except in panels (i) and (j) of Figure 4.8. Panel (i) shows

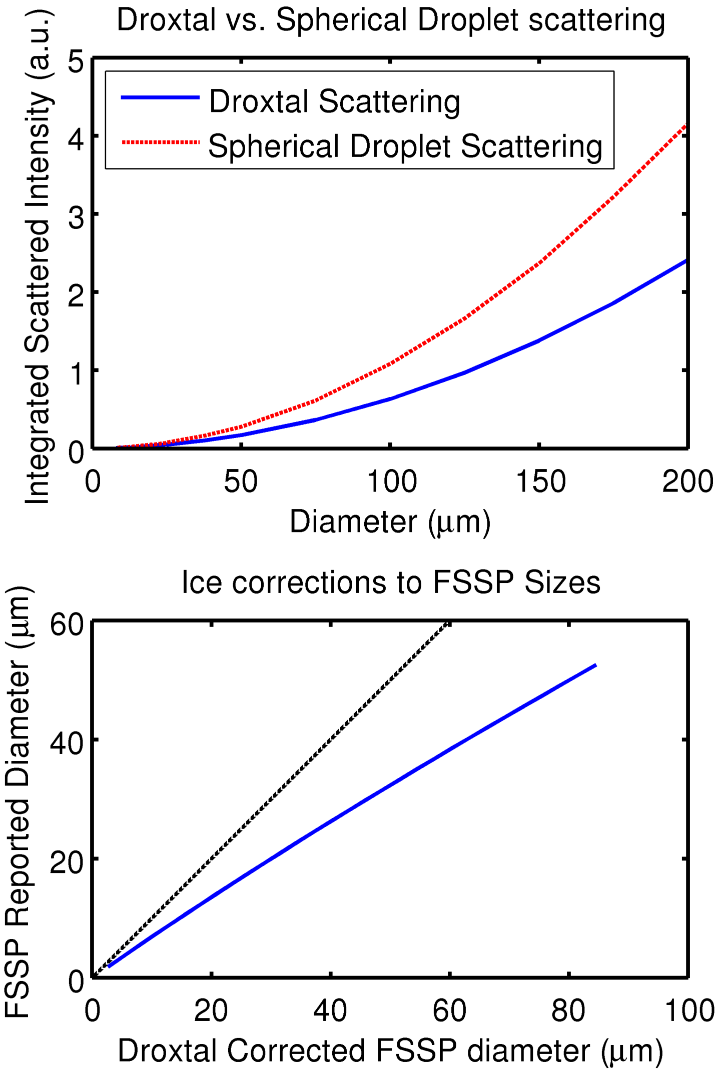

Figure 4.5: The top panel shows the scattered intensity integrated over the forward scattering angles measured by the FSSP. From this, a correction curve is calculated for the measured FSSP size bins as shown in the bottom panel. The black dashed line shows a one-to-one correction curve for comparison.

a time when the FSSP detected a high number density of small liquid water droplets that HOLODEC was not able to detect. This is partly because these holograms had a higher background noise level due to window contamination. Panel (j) shows a time when HOLODEC should not detect any particles of any size according to the FSSP, but the particle detection algorithm falsely detects a small number of particles. In both instances, higher quality holograms obtained with clean windows (e.g. no condensation or impacted cloud particles) would allow for more reliable detection of small particles and fewer false detections.

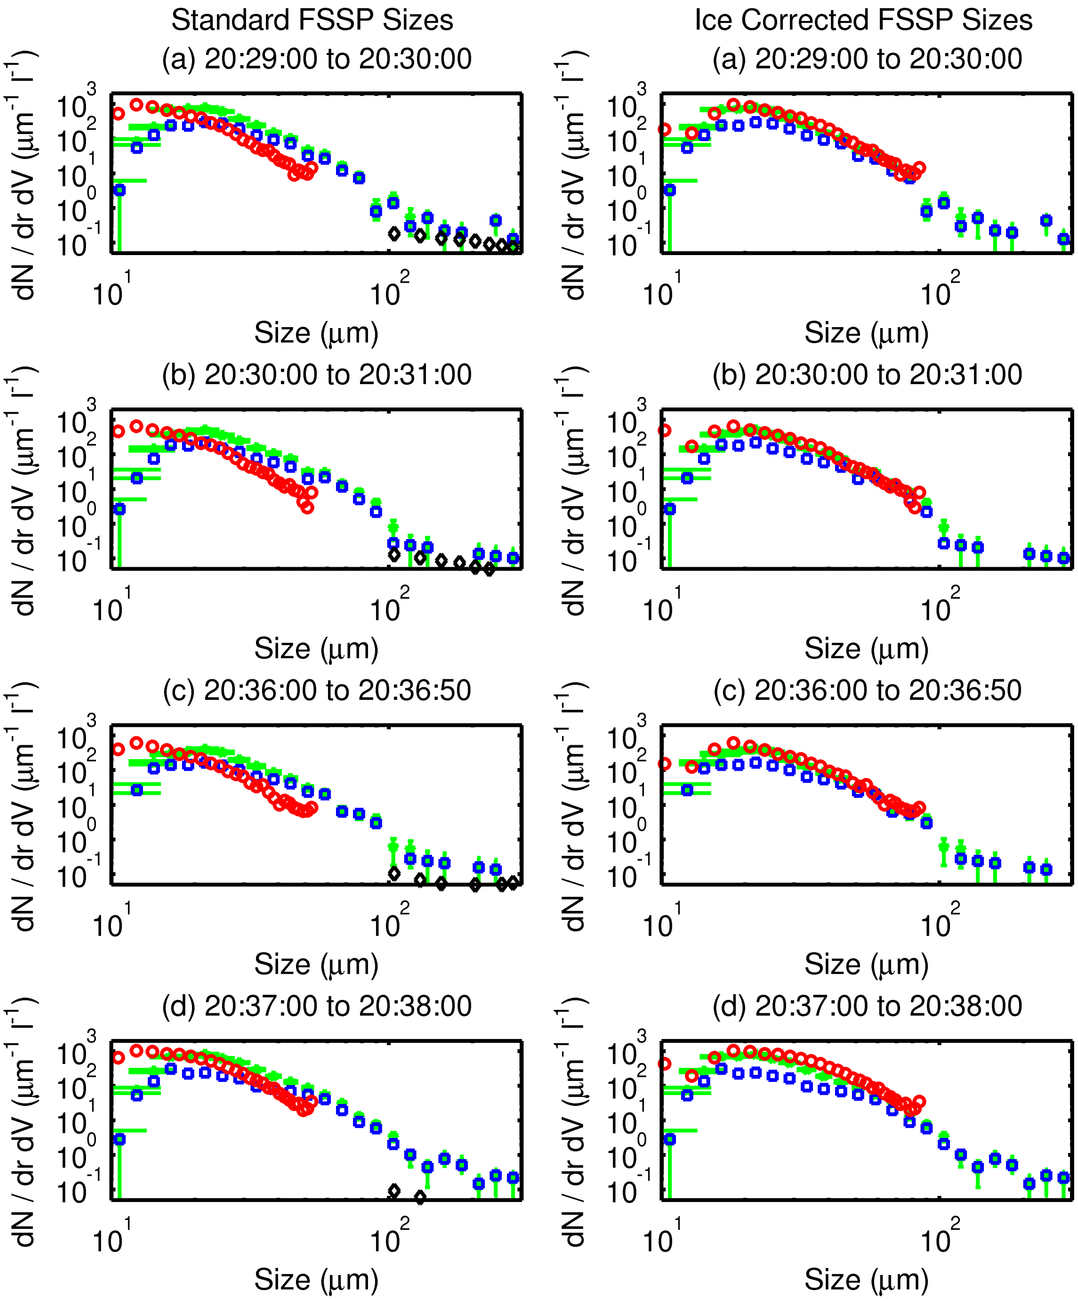

Figure 4.6: Size distributions of ice particles for four time periods shown above each frame. The left panels show the uncorrected FSSP sizes and the right panels show FSSP sizes corrected for ice particles. FSSP size distributions are shown as circles and HOLODEC size distributions are shown as pentagrams with error bars. The squares are HOLODEC size distributions excluding shattered particles. The error bars are calculated assuming Poisson statistics and are one standard deviation. The diamonds are data from the PMS 2-DC probe and are only shown on the left panels.

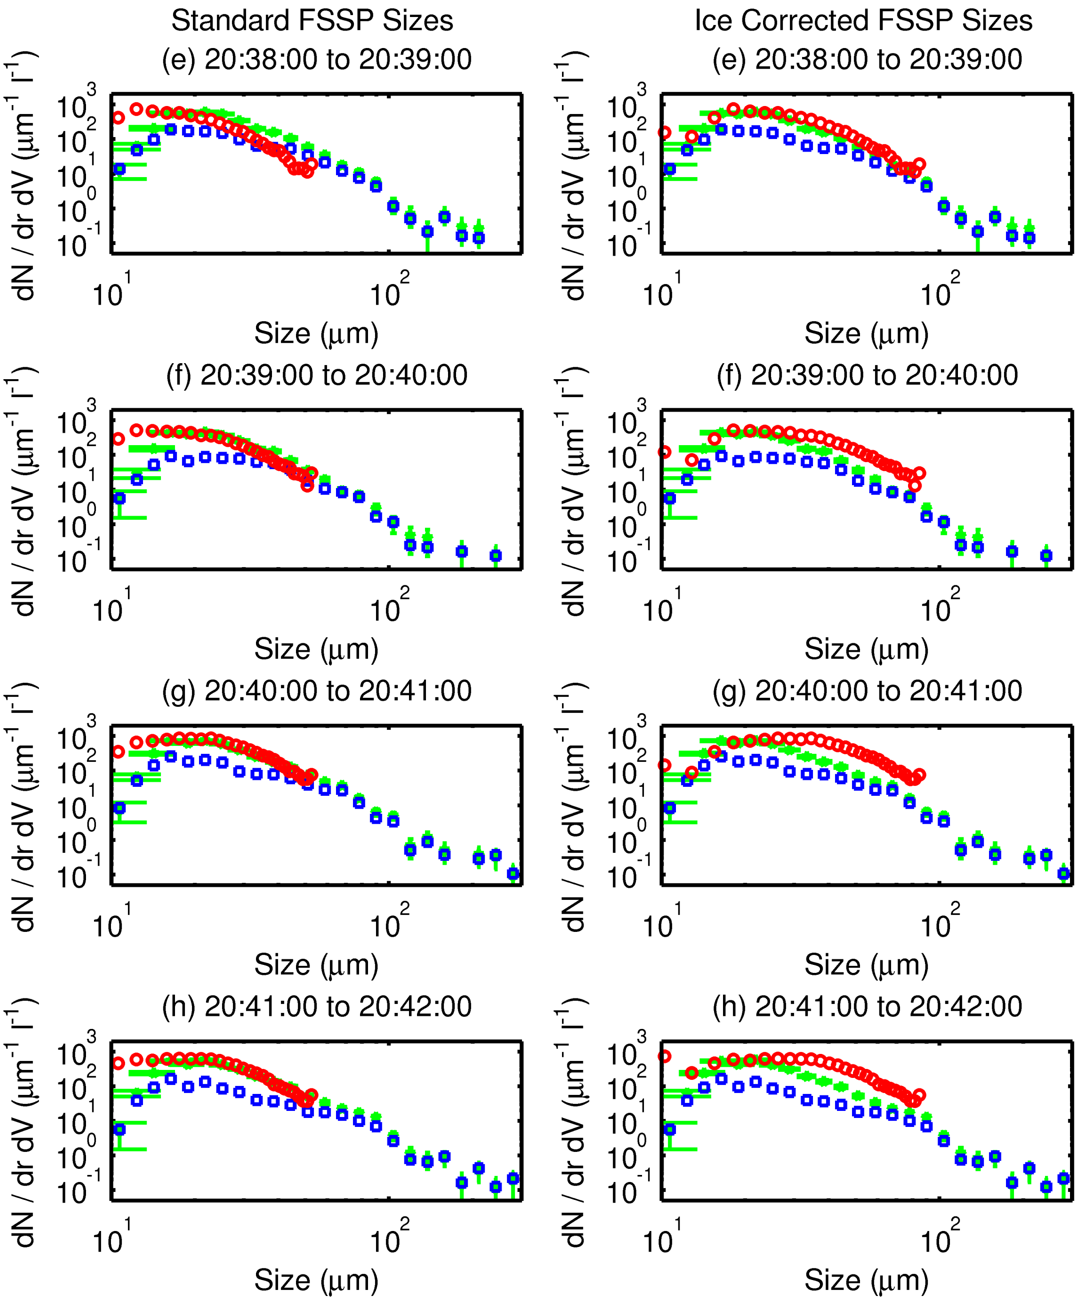

Figure 4.7: As in Figure 4.6, for four more time periods illustrating a range of cloud microphysical conditions.

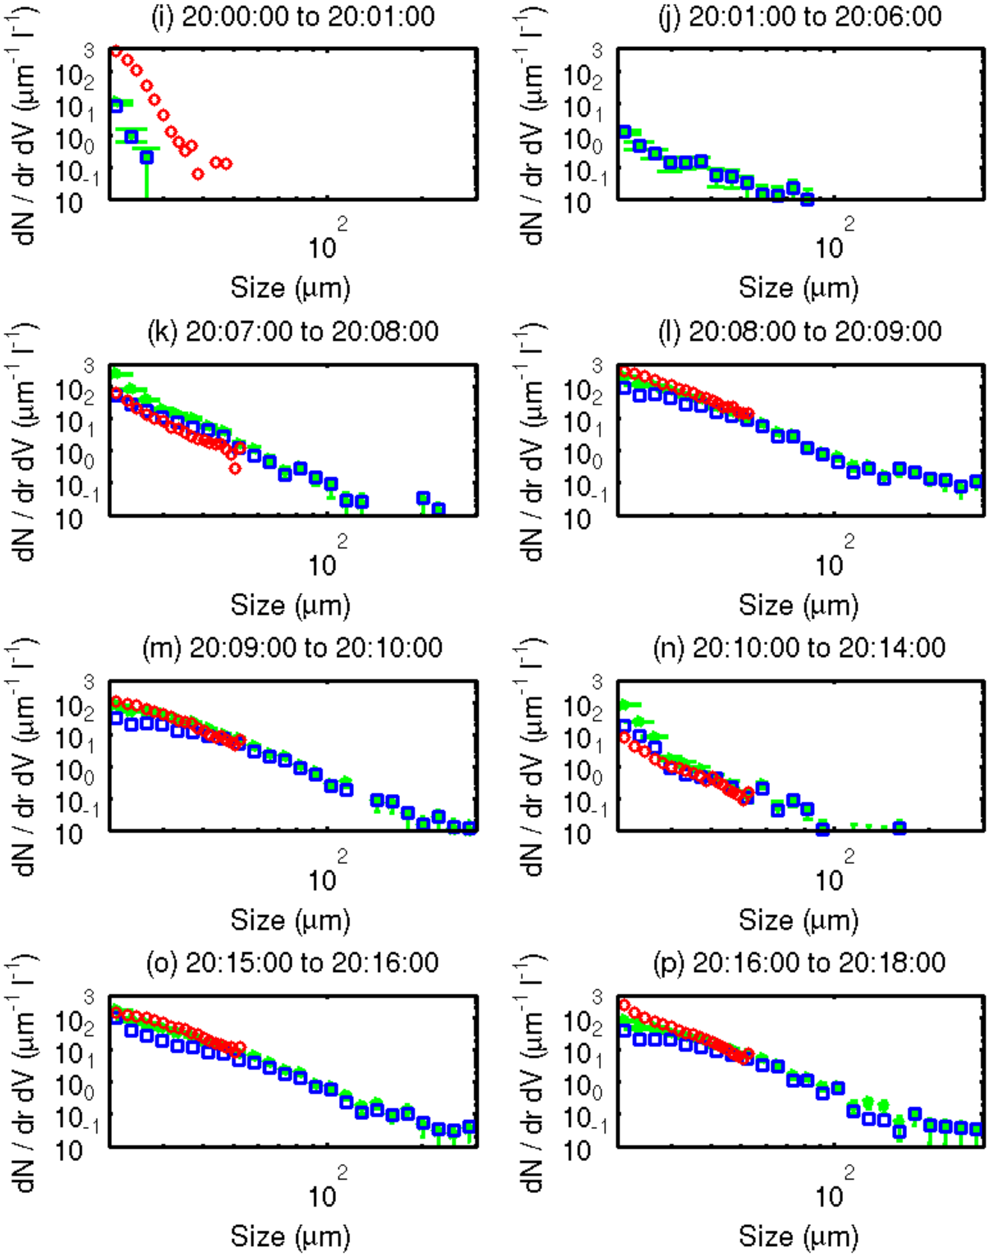

Figure 4.8: As in Figures 4.6 and 4.7 except for Research Flight 2003-09-05 and only shown with uncorrected FSSP Sizes. No PMS 2-DC probe data are included.

The variety of cloud conditions, range of number densities, and particle size range for which the HOLODEC instrument can measure cloud particle size distributions as shown in Figures 6, 7, and 8 is evidence of its reliability for measuring size distributions. We can also see from the HOLODEC size distributions in Figures 6, 7, and 8 that the shattered particles (the difference between the blue squares and green pentagrams) are typically smaller ice particles between about 10 and 70 μm. Further, the shattered particles are typically not more than about 50% of the total particles. This is a useful approximate upper bound on the effect of shattered particles in size distributions in the absence of large ice particles and perhaps can be used to better interpret FSSP size distributions of ice particles. Note that this observation about interpreting the effect of shattering on FSSP size distributions is only based on the assumption that the FSSP and HOLODEC probes measure the same ice size distributions and the fact that the uncorrected size distributions from the FSSP and HOLODEC are approximately similar. It is not dependent on similarities between instrument measuring methods or instrument housing geometries near the sample volume.

The shattering effect is further illustrated in Figure 4.9, which shows a number density time series plot for both the FSSP and HOLODEC. The HOLODEC data are plotted to show the total number density of ice particles and the number density of ice particles in holograms not flagged as containing shattered particles. There is good agreement in the HOLODEC total number density curve for all ice particles and the FSSP. Again, we see that shattering for ice particles amounts for up to about 50% of the total number density for these cloud conditions.

Figures 10 and 11 show a collection of ice particles 100 μm and larger from Research Flights 2003-09-17 and 2003-09-05 respectively, as they are reconstructed by the automated reconstruction and particle finding algorithm (Fugal et al., 2007, see Chapter 3). These figures illustrate the size range and different morphologies or crystal habits of ice particles detected in these two data sets. It also shows the automated algorithm's ability to find the focus and obtain a 2-D profile of these ice particles. The wavy non-uniform background about these particles is the virtual image of each ice particle (Fugal et al., 2004, see Chapter 2). As with optical array probes, the particles smaller than 100 μm reconstruct to appear more like circles with fewer discernible features such as facets. As a curiosity, we note the 8 sided crystal with a 6 pointed star inside it at the lower right of Figure 10. This is likely a thick plate tilted to make it appear 8 sided. Notice also in Figure 11 that several ice particles appear as perspective views resulting from ice particles tilted with respect to the camera of HOLODEC.

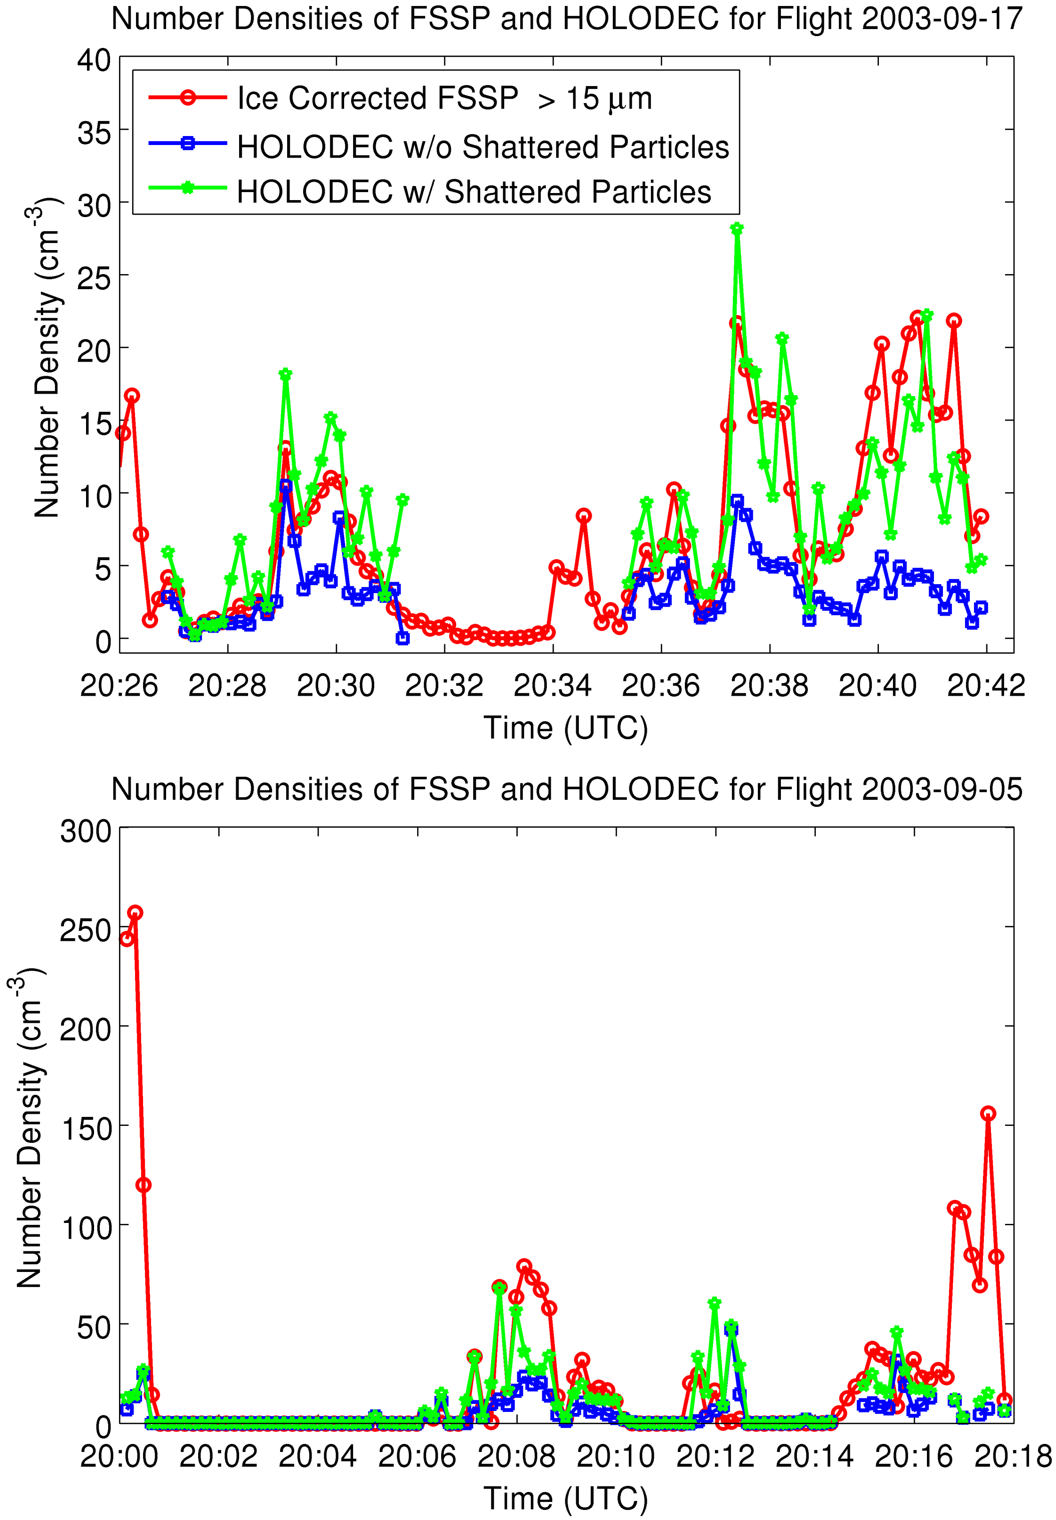

Figure 4.9: Number densities averaged over 10 second intervals from FSSP and HOLODEC.

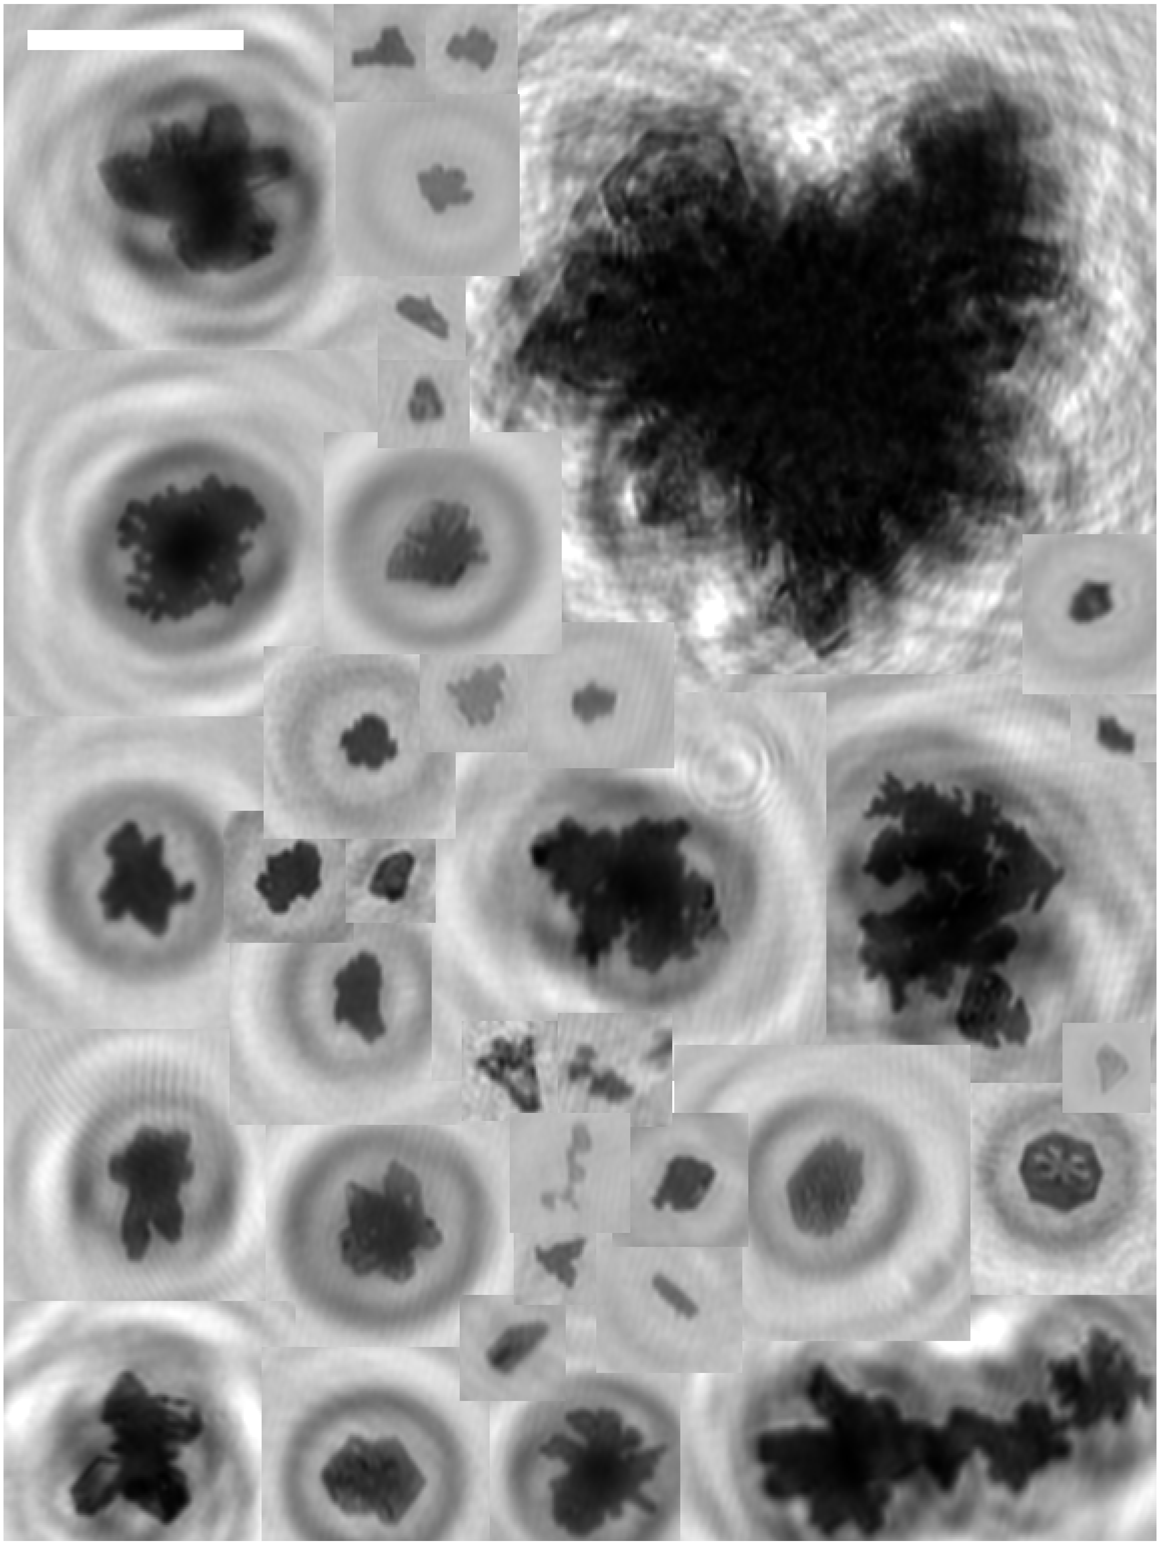

Figure 4.10: A sample of ice particles taken during Research Flight 2003-09-17. All particles are shown as reconstructed by the automated particle finding algorithm and are only scaled to improve contrast for printing. The white scale bar in the upper left is 0.5 mm in length.

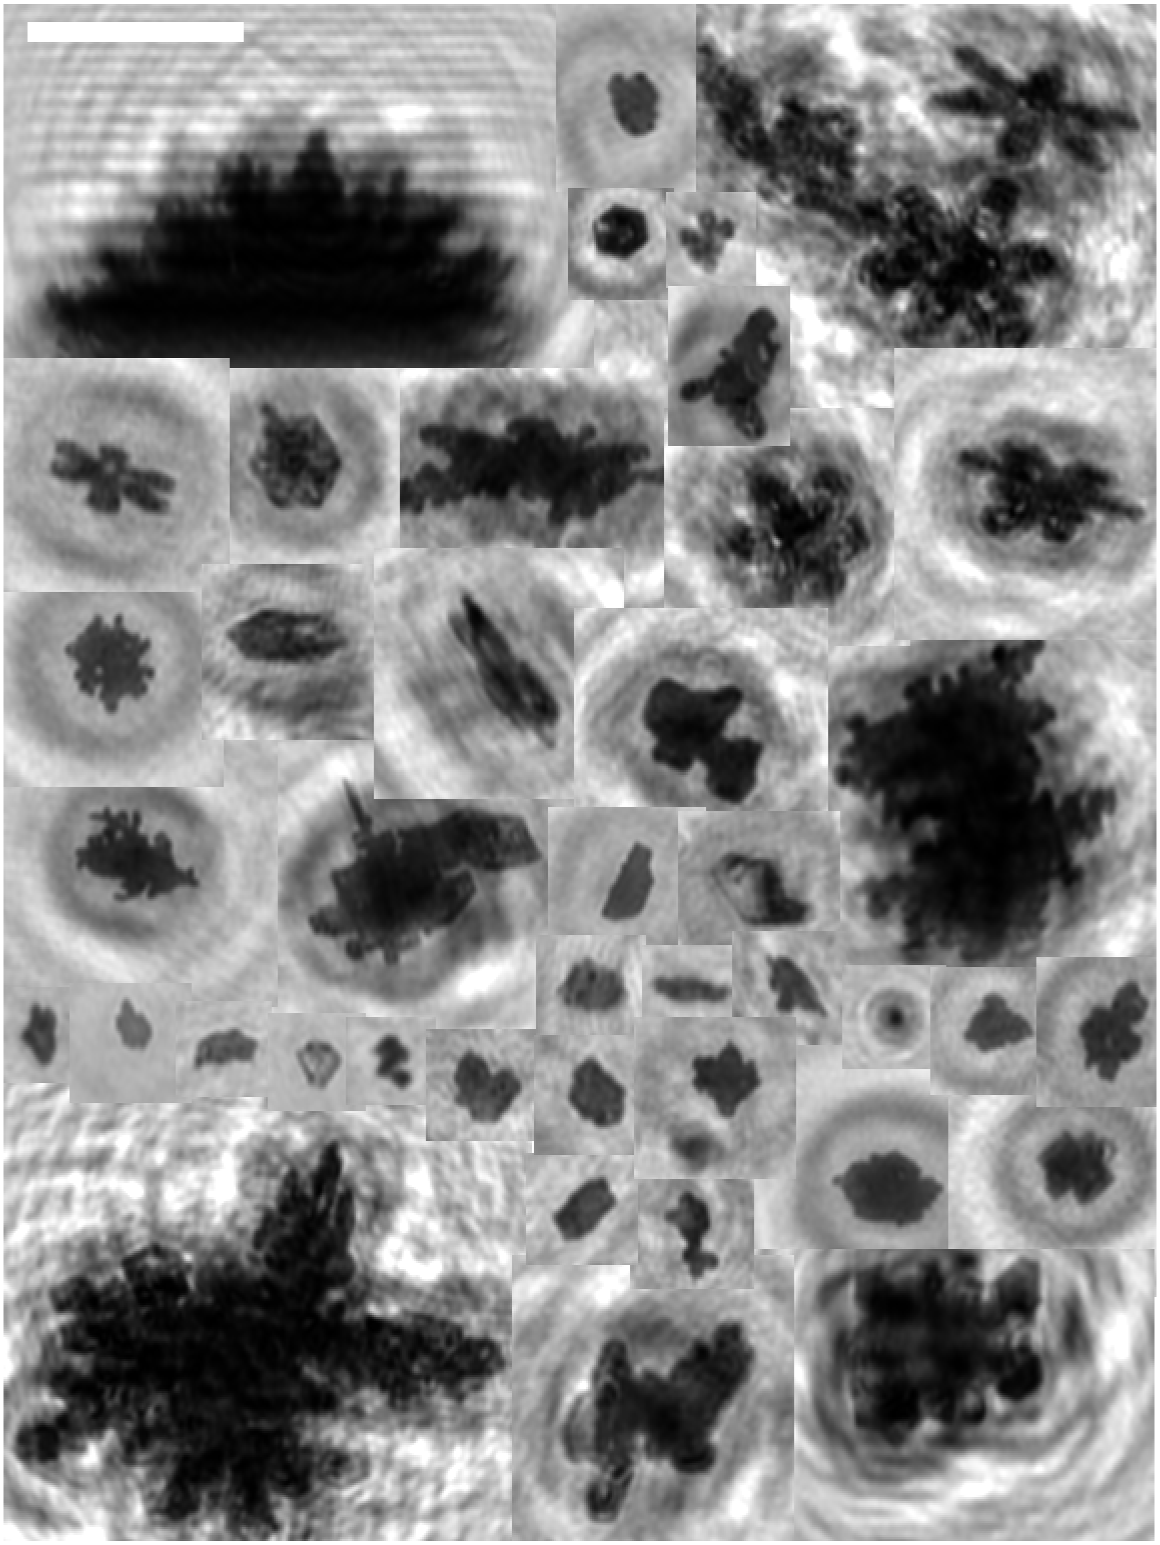

Figure 4.11: As in Figure 10 except ice particles are a sample from Research Flight 2003-09-05. The white scale bar in the upper left is 0.5 mm in length.

4.5 Discussion and Summary

We have shown that the HOLODEC instrument, an instrument using digital inline holography, and the reconstruction algorithm described by Fugal et al. (2007, see Chapter 3) can detect ice particles, find their 3-D positions, calculate their sizes, and give 2-D profiles. From these results, accurate size distributions and number densities can be calculated. We have shown that size distributions from HOLODEC and an FSSP agree given the uncertainty in correcting FSSP data for non-spherical ice particles. We note that these results support the possibility that in clouds of nearly pure ice, such a correction as shown in Figure 5 can be applied to correct for the FSSP calibration for spherical water droplets to faceted non-spherical ice particles (Field et al., 2003a).

We also note the unique ability to detect holograms with shattered ice particles by searching for highly localized clusters along the optical axis and thus obtaining more accurate size distributions. The data presented support the conclusion that in the absence of large ice particles, size distributions from the FSSP are likely overestimated by no more than a factor of two due to particle shattering. This is important because there are many large data sets of FSSP ice particle distributions that are used in satellite measurement validations, and studies of ice processes in clouds. Hence, any evidence that tells us under what conditions we might expect FSSP ice particle data to be valid and any corrections we might apply to make the data more accurate is valuable.

The digital holographic method allows a volume sample rate that is independent of particle size and air speed. And with its associated hologram reconstruction and particle finding algorithm, we are able to find particles in a large size range limited only by a few pixels wide for small particles, and some fraction of the size of the camera for large particles. While holograms require a much longer processing time than instruments that offer real time results, holograms provide more information than current typical optical cloud particle instrumentation. Furthermore, as with any computational problem, obstacles related to hologram construction time will erode and eventually vanish as computation speeds and algorithm development progress. HOLODEC is a prototype instrument designed to give proof of concept of the ability and utility of digital holography for cloud particle measurements. Future versions of this instrument would be improved by including a camera with a much higher and/or larger image size and therefore have a larger volume sample rate. This combined with more refined algorithms will allow researchers to obtain accurate cloud particle size distributions, absent shattered particle contamination along with other cloud particle measurements of interest.

Acknowledgments

This work is supported by an NSF Graduate Research Fellowship, a NASA Earth System Science Fellowship, a Michigan Space Grant Consortium Graduate Fellowship, and NSF Grant ATM-0535488. We thank the staff of NCAR's Research Aviation Facility for aiding us in analyzing the data from other aircraft probes taken during the IDEAS 3 campaign. We also thank Ping Yang for providing the library of scattering calculations for droxtals. We thank Paul Field for helpful comments and discussion.

Chapter 5

Summary and Outlook

5.1 Summary

In this dissertation, I have shown a successful airborne digital holographic instrument that takes holograms of cloud particles. I have presented an algorithm that has properties that make it desirable for reconstructing holograms for typical experiments. Further, I described an algorithm that can find the position along the optical axis (z-position) of small particles as well as large complex-shaped particles. I explained an implementation of these algorithms that is an efficient, robust, automated program that allows us to process holograms on a computer cluster. Finally, I have shown size distributions and number densities of cloud particles, and have shown that they are within the uncertainty of independent measurements made with another measurement method. Thus we have proven the feasibility of another cloud particle instrument that has advantages over new standard instruments. These advantages include a unique ability to detect shattered particles using three-dimensional positions, and a sample volume size that does not vary with particle size or airspeed. It also is able to yield two-dimensional particle profiles using the same method. There are also some cloud particle measurements we can make with an advanced version of HOLODEC that no one has made before that I explain below.

Cloud processes are sensitive to events that occur at interfaces such as mixing of clear and cloudy air at cloud edge and cloud top, and boundaries of liquid water and ice particle layers. These events include the rapid evolution of particle size distributions from mixing (Andrejczuk et al., 2006; Korolev et al., 2003) and the interaction between particles in liquid, ice, and mixed phase regions in clouds. Standard instruments typically have modest sample rates for the populations of the particles they detect as shown in the last column of Table 5.1. Many standard instrument probes do not sample fast enough to resolve size distributions on a cloud length scale (of order 100 m) let alone at cloud edge and cloud top (of order 1 m). Therefore a high sample rate instrument that can measure size distributions in path lengths of order 1 or 10

Instrument	Particle Size Range (μm)	Sample Rate at 100 m s^{-1} (L s^{-1})	Path Length to sample 1 L at 100 m s^{-1} (m)	Path Length to sample 1000 particles at 100 m s^{-1} (m)
HOLODEC I	15 to 200	0.007	14000	1400
FSSP and SID	1 to 50	0.05	2000	200
2D-C	100 to 2000	3	30	33000
2D-S	10 to 100	0.16	600	60
2D-S	100 to 1000	16	6	6000
HOLODEC II burst (6688 fps)	50 to 1000	160	0.6	6
HOLODEC II 5x burst (6688 fps)	10 to 200	2	50	5

Table 5.1: Particle size ranges, sample rates and path lengths through clouds to sample 1 L for various instruments including HOLODEC II burst and continuous sample modes. The 1000 particle sample path length is estimated by assuming a typical concentration of 10^4 10 μm particles L^{-1}, 100 50 μm particles L^{-1}, and 1 100 μm particles L^{-1}. The final column gives a rough estimate of how long we might have to sample to get a good size distribution in that size range from that instrument.

m has potential to greatly increase our understanding of the internal structure of clouds. There are new instruments that have higher sample rates, for example, the SPEC 2D-S has a much higher sample rate than previous instruments (e.g. FSSP and 2D-C, see Table 5.1). Still, with the improvement of fast cameras and computers, and advanced algorithms to find particles in holograms, there is great potential for a new high sample rate instrument and associated data analysis hardware and software to find size distributions of cloud particles at these boundaries.

The new generation HOLODEC instrument could be greatly improved over the first one using a new camera capable of high sample rates and other refinements to make it more autonomous and reliable. A likely best candidate for the next generation HOLODEC camera is the Phantom v7.3 camera that has 800x600 22 μm square pixels and has up to 16 GB of RAM onboard. It has both a burst mode of 6688 fps (which writes to onboard RAM and downloads to hard disk later) to intensely sample a cloud and a continuous sampling mode of 80 fps (streams to hard disk) for long passes through clouds. The frame rate is continuously variable between these two extremes. The burst mode allows for a high sample rate 160 L s^{-1} (see Table 5.1) or up to 5.1 seconds and then takes 7 minutes to download to disk. Assuming particle densities of order 100 L^{-1} in the 50 to 500 μm size range, this mode allows for size distribution determined from 1000 particles every 0.6 seconds or 6 m if flying at 100 m s^{-1}. HOLODEC II could measure smaller particles using magnification (say 5x magnification to see 10 μm particles) and a smaller burst sample rate of around 2 L

s^{-1}. Assuming particle densities of 10 cm^{-3} in the 10 to 200 μm size range, only 5 m of flight path yields about 1000 particles and a size distribution. Both the 5 and 6 m sample lengths to get size distributions allows us to study the internal structure of clouds especially at the aforementioned boundaries that may be key to understanding important cloud events.

Digital holograms require significant computational power to reconstruct and only recently has it been feasible to routinely reconstruct thousands of holograms (Fugal and Shaw, 2007; Fugal et al., 2007). There is great potential for faster processing of holograms that does not sacrifice accuracy. For example, one could employ an adaptive scheme that reconstructs a hologram in coarse steps until it detects a particle nearly in focus and at that point finely resolve the particle.

Consider a concrete example of processing time for the next generation HOLODEC using a research flight with 60 minutes of in-cloud time which would yield around 300,000 holograms or 150 GB of raw images. Applying the simplest adaptive scheme, these holograms would take about 15 minutes per hologram on a Xeon 3.0 GHz 2GB RAM machine and the flight data would take around 75000 cpu hours or 24 days on a 128 processor computer cluster. Current algorithms are already ported to computer clusters and could be adapted to supercomputers with shared memory (e.g. Satake et al. 2007).

A new high sample-rate HOLODEC instrument would be a powerful tool to accurately measure small ice crystal number densities and size distributions with subcloud resolution providing measurements no other instrument has ever provided.

Appendix A

MATLAB Hologram Reconstruction Codes

1.1 HFFFTPrep.m

This code prepares field 'F' to be propagated up to distance 'maxz'. and is called by:

```
 FPrepped = HFFFTPrep(F, dx, dy, wavelength, maxz);
```

where 'dx' is the pixel x width, and 'dy' is the pixel y width, 'wavelength' is the wavelength of the light, and 'maxz' is the distance to which you want the undersampled frequencies filtered out. If you want to enforce a uniform particle detection limit up to 'maxz', the that's the distance you put in. If 'maxz' is equal to 0, then the low pass filter is calculated for every distance 'z' called in HFFFTPropagate.

```
function FPrepped = HFFFTPrep(Field, dx, dy, lambda, maxz)
%This function prepares a field for faster propagation for instances
%when propagating the same field to more than one distances. The
%savings is in eliminating the first FFT for subsequent propagations.
%It is also possible to apply a low pass filter to obtain uniform
%resolution out to some distance maxz. The low pass filter is set to
%cut out frequencies so high they'd be undersampled using this method
%of propagation.
%
%FPrepped = HFFFTPrep(Field, dx, dy, lambda, maxz)
%
%Field, is the field to propagate,
%dx is the x sample spacing
```

```
%dy is the y sample spacing
%lambda is the wavelength of the propagated light
%maxz is the maximum z you want to propagate to with uniform resolution.
%
%If you want to use adaptive low pass filtering (filter out frequencies
%corresponding to undersampled phases per propagation distance) set
%maxz = 0;
%

% Version 1.0
% last edited April 9, 2007


%For convenience
FPrepped.Ny     = size(Field,1);    %Rows of y
FPrepped.Nx     = size(Field,2);    %Columns in x
FPrepped.k      = 2*pi/lambda;
FPrepped.lambda = lambda;
FPrepped.dnux   = 1/(dx*FPrepped.Nx);
FPrepped.dnuy   = 1/(dy*FPrepped.Ny);
FPrepped.dx     = dx;
FPrepped.dy     = dy;
FPrepped.Lx     = FPrepped.Nx * FPrepped.dx;
FPrepped.Ly     = FPrepped.Ny * FPrepped.dy;
FPrepped.Lnux   = FPrepped.Nx * FPrepped.dnux;
FPrepped.Lnuy   = FPrepped.Ny * FPrepped.dnuy;
FPrepped.field  = Field;
FPrepped.resolution = 2.44*FPrepped.lambda*maxz...
        /sqrt(FPrepped.Nx*FPrepped.Ny*FPrepped.dx*FPrepped.dy);

%Take the FFT
FPrepped.FieldFFT = fft2(double(Field));
clear Field;

%Make the phase multiplier
FPrepped.x = -FPrepped.Nx/2:FPrepped.Nx/2-1;
FPrepped.y = -FPrepped.Ny/2:FPrepped.Ny/2-1;
[xx,yy] = meshgrid(FPrepped.x,FPrepped.y);
rho2 = lambda^2*(xx.^2*FPrepped.dnux.^2 + yy.^2*FPrepped.dnuy.^2);

%Now we must select only those rho^2s that keep the square root real or
%the rho^2s where (1-rho^2*lambda^2 > 0). This effectively excludes
%evanescent waves.

if( max(rho2(:)) >=1)
```

```
    warning(['HFFFT:ApTooBig','Aperture exceeds spatial '...
        'resolution allowed by lambda');
    FPrepped.zeroOut = rho2 <=1;
else
    FPrepped.zeroOut = logical(ones(FPrepped.Ny,FPrepped.Nx));
end
FPrepped.root = sqrt(1-rho2);
clear rho2;
% Find the maximum possible z we can go to wihtout undersampling the
% phase. A little work will yield the following formula from:
% | ( dx * d/dx ihat + dy * d/dy jhat ) k d sqrt( 1 - lambda^2 ( nux^2
% + nuy^2)) | < pi %
%
% Or:
%
% d < 1/( 2*lambda ) * sqrt( nux^2*dnux^2 + nuy^2*dnuy^2 )/sqrt( 1 -
% lambda^2 * ( nux^2 + nuy^2) %
%
% where we should plug in nux = Lnux/2 and nuy=Lnuy/2 to get the maximum
% phase difference between samples at the corners or the phase filter.

%maxZ With Out UnderSampling
FPrepped.maxZwous = 1 / (2*FPrepped.lambda) *...
        sqrt( 1 - FPrepped.lambda^2 * ((FPrepped.Lnux/2)^2...
        + (FPrepped.Lnuy/2)^2))...
        ./ sqrt( (FPrepped.Lnux/2*FPrepped.dnux)^2...
        + (FPrepped.Lnuy/2*FPrepped.dnuy)^2);

% minZ with out undersampling via convolution:
FPrepped.minZwous = sqrt( 1/FPrepped.lambda^2 *...
        ( FPrepped.Lx^2 * FPrepped.dx^2...
        + FPrepped.Ly^2 * FPrepped.dy^2 )...
        - (FPrepped.Lx^2 + FPrepped.Ly^2)/4) ;


%If the phase will vary by more than pi between samples then give a
%warning and set all the undersampled pixels to 0.
if(maxz == 0)
    FPrepped.adaptiveFiltering = 1;
else
    FPrepped.adaptiveFiltering = 0;
    if(abs(maxz) > FPrepped.maxZwous);
        ellipse = (4*maxz^2*FPrepped.dnux^2+1)*(xx*FPrepped.dnux).^2 ...
              + (4*maxz^2*FPrepped.dnuy^2+1)*(yy*FPrepped.dnuy).^2;
        zeroOut = ellipse < FPrepped.lambda^-2;
        clear ellipse;
        % Now make a SuperGaussian window around this such that some
        % fraction 'f' of the Gaussian lies on the edge of zeroOut. 'n'
        % is the order of the SuperGaussian such that (ignoring
        % normalization):
        %
        % SG(x,y) = exp(-1/2*((x/sigmax)^2 + (y/sigmay)^2)^n )

        f         = .5;
        n         = 3;
        nuxwidth = sum(zeroOut(end/2+1,:));
        nuywidth = sum(zeroOut(:,end/2+1));
        clear zeroOut;
        sigmax    = nuxwidth/2 * log(1/f^2)^(-1/(2*n));
        sigmay    = nuywidth/2 * log(1/f^2)^(-1/(2*n));
        filter    = exp(-1/2*((xx/sigmax).^2+(yy/sigmay).^2).^n);
        FPrepped.zeroOut = double(FPrepped.zeroOut) .* filter;
        fprintf(['Incidentally, the maxz you gave is in the '...
 'undersampled phase regime.\n']);
    end
end
```

1.2 HFFFTPropagate.m

This code propagates field 'F' distance 'z'. and is called by:

```
FieldOut = HFFFTPropagate(F, z, applyLowpassFilter, dx, dy,...
        wavelength, maxz);
```

where 'dx' is the pixel x width, and 'dy' is the pixel y width, 'wavelength' is the wavelength of the light, 'applyLowPassFilter' is if you want to filter out undersampled frequencies (if any), and 'maxz' is the distance to which you want the undersampled frequencies filtered out. If you want to enforce a uniform particle detection limit up to 'maxz', the that's the distance you put in. If 'maxz' is equal to 0, then the low pass filter is calculated for distance 'z'.

If you've previously have called HFFFTPrep, then you call HFFFTPropagate as:

```
FPrepped = HFFFTPrep(F, dx, dy, wavelength, maxz);
Fieldout = HFFFTPropagate(FPrepped, z, applyLowpassFilter);
```

```
function FieldOut = HFFFTPropagate(F, z, applyLowPassFilter, varargin)
```

```
%
%FieldOut = HFFFTPropagate(FPrepped, z, applyLowPassFilter)
%Propagates a field prepped with HFFTPrep. This method can save time if
%propagating the same field (same wavelength, same sample spacing, same
%amplitude and phase pattern) more than one distance.
%
%FieldOut = HFFFTPropagate(Field, z, applyLowPassFilter, dx, dy,
%        lambda, maxz)
%First calls HFFFTPrep then propagates the field.
%
%This function takes the FFT of Field (actually done in HFFFTPrep),
%multiplies by the Huygens Fresnel kernel, then takes the inverse FFT.
%The result is mathematically identical
%to convolving with the Huygens Fresnel kernel in the space domain.

%If propagated too far the phase begins to jump by more than pi between
%samples at the edges. If applyLowPassFilter = 0 this function ignores
%any undersampling and but provides a warning if undersampling is
%present. If applyLowPassFilter = 1 this function checks for those
%jumps and sets the amplitude of undersampled phases to 0 effectively
%acting as a low-pass filter.
%
%Alernately you can call HFFFTPrep with a maxz which will set
%FPrepped.zeroout to filter out all undersampled frequencies at maxz
%throughout every propagation distance yielding a uniform resolution for
%the propagated fields < maxz.
%
%The function HFFFTPrep is intended to cut down on the processing time
%for multiple propagations of the same field. Calling HFFTPrep
%beforehand and passing FPrepped to HFFFTPropagate function will save
%you an initial FFT, and grid calculation.

% Version 1.0
% last edited April 9, 2007

if(length(varargin) == 4)
    dx = varargin{1};
    dy = varargin{2};
    lambda = varargin{3};
    maxz = varargin{4};
    FPrepped = HFFFTPrep(F, dx, dy, lambda, maxz);
    clear F;
elseif isempty(varargin)
    OK = isfield(F,'root') & isfield(F,'k') & isfield(F,'FieldFFT') ...
& isfield(F,'zeroOut');
```

```
    if(OK)
        FPrepped = F;
        clear F;
    else
        error(’HFFFTPropagate:MissingFields’,[’FPrepped must have’ ...
’ fields ’’k’’, ’’FieldFFT’’, ’’root’’, ’’zeroOut’’. Use’ ...
  ’ HFFFTPrep.’);
    end
else
    error(’Wrong number of arguments’);
end


%Now make the phase:

%The matlab exp(i*phase) function takes about 2 seconds to calculate
%if the phase matrix is 2000x2000 and is uniformly less than
%1.7 million (and greater than 0). If it is greater than 1.7 million
%it takes 8 seconds to calculate. It actually is about the same to do
%modulo 2*pi then take the exponential as to just take the
%exponential if the phase is around 1,000. Since our phase is
%often greater than 1000, I’ll always take the moduli.

phase = FPrepped.k*z*FPrepped.root;
phase = phase - floor(phase/(2*pi))*2*pi;

%If the user doesn’t ask for a lowpass filter applied, then check if the
%phase will vary by more than pi between samples (z > maxZwous (maxZ
%With Out Under Sampling)) then give a warning.
if(~applyLowPassFilter && z >= FPrepped.maxZwous)
    warning(’HFFFT:PropagationTooFar’,[’Propagation is too far to ’...
           ’sufficiently resolve the phase multiplier.’]);
%If the user asks for the lowpass filter to be applied and we are into
%the z > maxzwous where undersampling is present and the user asked for
%adaptive filtering, make a filter:
elseif(applyLowPassFilter && FPrepped.adaptiveFiltering)
    if(abs(z) > FPrepped.maxZwous);
        [xx,yy] = meshgrid(FPrepped.x,FPrepped.y);
        ellipse = (4*z^2*FPrepped.dnux^2+1)*(xx*FPrepped.dnux).^2 ...
                    + (4*z^2*FPrepped.dnuy^2+1)*(yy*FPrepped.dnuy).^2;
        zeroOut = ellipse < FPrepped.lambda^-2;
        clear ellipse;
        %Now make a SuperGaussian window around this such that some
        %fraction ’f’ of the Gaussian lies on the edge of zeroOut. ’n’
  %is the order of the SuperGaussian such that (ignoring
  %normalization):
        %
        % SG(x,y) = exp(-1/2*((x/sigmax)^2 + (y/sigmay)^2)^n )

        f        = .5;
        n        = 3;
        nuxwidth = sum(zeroOut(end/2+1,:));
        nuywidth = sum(zeroOut(:,end/2+1));
        clear zeroOut;
        sigmax   = nuxwidth/2 * log(1/f^2)^(-1/(2*n));
        sigmay   = nuywidth/2 * log(1/f^2)^(-1/(2*n));
        filter   = exp(-1/2*((xx/sigmax).^2+(yy/sigmay).^2).^n);
        FPrepped.zeroOut = double(FPrepped.zeroOut) .* filter;
        fprintf(['Incidentally, the z you gave is in the ' ...
'undersampled  phase regime.\n']);
    end
end

%Make the phase filter
phaseFilter = exp(j*phase);
clear phase;
%multiply in any low pass filter either from undersampled frequencies or
%evanescent waves
if( applyLowPassFilter )
    filter = phaseFilter.*FPrepped.zeroOut;
else
    filter = phaseFilter;
end
clear phaseFilter;

%Apply the overall filter, take the IFFT and return
FieldOutFFT = FPrepped.FieldFFT.*fftshift(filter);
clear filter
FieldOut = ifft2(FieldOutFFT);
```

References

Andrejczuk, M., W. W. Grabowski, S. P. Malinowski, and P. K. Smolarkiewicz, 2006: Numerical simulation of cloud-clear air interfacial mixing: effects on cloud microphysics. *jas*, **63**, 3204–3225.

Baum, A. B., P. Yang, A. J. Heymsfield, S. Platnick, M. D. King, Y.-X. Hu, and S. Bedka, 2005a: Bulk scattering properties for the remote sensing of ice clouds. part ii: Narrowband models. *jam*, **44**, 1896–1911.

Baum, B. A., A. J. Heymsfield, P. Yang, and S. Bedka, 2005b: Bulk scattering properties for the remote sensing of ice clouds. part i: Microphysical data and models. *jam*, **44**, 1885–1895.

Baumgardner, D. and A. J. Korolev, 1997: Airspeed corrections for optical array probe sample volumes. *joat*, **14**, 1224–1229.

Bohren, C. F. and D. R. Huffman, 1983: *Absorption and Scattering of Light by Small Particles*. Wiley-Interscience, New York.

Born, M. and E. Wolf, 1999: *Principles of Optics*. Cambridge Univ. Press, Cambridge, 7 edition.

Borrmann, S. and R. Jaenicke, 1993: Application of microholography for ground-based in situ measurements in stratus cloud layers: a case study. *joat*, **10**, 277–293.

Brown, P. R. A., 1989: Use of holography for airborne cloud physics measurements. *joat*, **6**, 293–306.

Coetmellec, S. D., D. Lebrun, and C. Ozkul, 2002: Application of the two-dimensional fractional-order fourier transformation to particle field digital holography. *J. of Opt. Soc. Am. A*, **19**, 1537–1546.

Field, P. R., A. J. Baran, P. H. Kaye, E. Hirst, and R. Greenaway, 2003a: A test of cirrus ice crystal scattering phase functions. *Geophys. Res. Lett.*, **30**, 1752.

Field, P. R., A. J. Heymsfield, and A. Bansemer, 2006: Shattering and particle inter-arrival times measured by optical array probes in ice clouds. *joat*, **23**, 1357–1371.

Field, P. R., R. Wood, P. R. A. Brown, P. H. Kaye, E. Hirst, R. Greenway, and J. A. Smith, 2003b: Ice particle interarrival times measured with a fast fssp. *joat*, **20**, 249–261.

Fournier, C., C. Ducottet, and T. Fournel, 2004: Digital in-line holography: influence of the reconstruction function on the axial profile of a reconstructed particle image. *Meas. Sci. Technol.*, **5**, 686–693.

Fugal, J. P., T. J. Schulz, and R. A. Shaw, 2007: Practical methods for reconstruction and characterization of particles in digital inline holograms, to be submitted to Measurement Science and Technology, (Adapted and Included as Chapter 3 in this dissertation.).

Fugal, J. P. and R. A. Shaw, 2007: Ice particle size distributions from holodec, an airborne digital in-line holographic instrument, submitted to J. Atmos. Ocean. Tech. (Adapted and Included as Chapter 4 in this dissertation.).

Fugal, J. P., R. A. Shaw, E. W. Saw, and A. V. Sergeyev, 2004: Airborne digital holographic system for cloud particle measurements. *ao*, **43**, 5987–5995, (Adapted and Included as Chapter 2 in this dissertation.).

Gardiner, B. A. and J. Hallett, 1985: Degradation of in-cloud forward scattering spectrometer probe measurements in the presence of ice particles. *joat*, **2**, 171–180.

Gayet, J.-F., G. Febvre, and H. Larsen, 1996: The reliability of the pms fssp in the presence of small ice crystals. *joat*, **13**, 1300–1310.

Goodman, J., 1996: *Introduction to Fourier Optics*. McGraw Hill, Boston, 2 edition.

Hinsch, K. D., 2002: Holographic particle image velocimetry. *Meas. Sci. Technol.*, **13**, R61–R72.

Holtzer, G. L. and L. R. Collins, 2002: Relationship between the intrinsic radial distribution function for an isotropic field of particles and lower-dimensional measurements. *J. Fluid Mech.*, **459**, 93–102.

Ilchenko, V. and T. Sattelmayer, 2005: Potential of the digital holographic particle image velocimetry (dhpiv) for flow investigations, 6th World Conference on Experimental Heat Transfer, Fluid Mechanics, and Thermodynamics, April 17-21, 2005, Matsushima, Miyagi, Japan, 9-a-3.

Korolev, A., 2007: Reconstruction of the sizes of spherical particles from their shadow images. part i: Theoretical considerations. *joat*, **24**, 376–389.

Korolev, A. and G. Isaac, 2005: Shattering during sampling by oaps and hvps. part i: snow particles. *joat*, **22**, 528–542.

Korolev, A. V., G. A. Isaac, S. G. Cober, J. W. Strapp, and J. Hallett, 2003: Microphysical characterization of mixed-phase clouds. *Q. J. R. Meteorol. Soc.*, **129**, 39–65.

Korolev, A. V., J. W. Strapp, and G. A. Isaac, 1998: Evaluation of the accuracy of pms optical array probes. *joat*, **15**, 708–720.

Kostinski, A. B. and R. A. Shaw, 2001: Scale-dependent droplet clustering in turbulent clouds. *J. Fluid Mech.*, **434**, 389–398.

Kozikowska, A., K. Haman, and J. Supronowicz, 1984: Preliminary results of an investigation of the spatial distribution of fog droplets by a holographic method. *Q. J. Roy. Meteorol. Soc.*, **110**, 65–73.

Kreis, T. M., 2002a: Frequency analysis of digital holography. *Opt. Eng.*, **41**.

— 2002b: Frequency analysis of digital holography with reconstruction by convolution. *Opt. Eng.*, **41**.

Kreis, T. M., M. Adams, and W. P. O. Juptner, 1997: Methods of digital holography: a comparison. *Proc. of SPIE*, **3098**.

Kreis, T. M. and W. P. O. Juptner, 1997: Suppression of the dc term in digital holography. *Opt. Eng.*, **36**, 2357–2360.

Kundu, P. K. and I. M. Cohen, 2002: *Fluid Mechanics*. Academic Press, San Diego.

Lawson, R. P. and R. H. Cormack, 1995: Theoretical design and preliminary tests of 2 new particle spectrometers for cloud microphysics research. *Atmos. Res.*, **35**, 315–348.

Lawson, R. P., D. O'Connor, P. Zmarzly, K. Weaver, B. Baker, Q. Mo, and H. Jonsson, 2006: The 2d-s (stereo) probe: design and preliminary tests of a new airborne, high-speed, high-resolution particle imaging probe. *joat*, **23**, 1462–1477.

MacPherson, J. I. and D. Baumgardner, 1998: Airflow about king air wingtip-mounted cloud particle measurement probes. *joat*, **5**, 259–273.

Malkiel, E., J. N. Abras, and J. Katz, 2004: Automated scanning and measurements of particle distributions within a holographic reconstructed volume. *Meas. Sci. Tech.*, **15**, 601–612.

Malkiel, E., J. N. Abras, E. A. Widder, and J. Katz, 2006: On the spatial distribution and nearest neighbor distance between particles in the water column determined from in situ holographic measurements. *J. Plankton Res.*, **28**, 149–170.

Malkiel, E., I. Sheng, J. Katz, and J. R. Strickler, 2003: The three-dimensional flow field generated by a feeding calanoid copepod measured using digital holography. *J. of Exp. Biology*, **206**, 3657–3666.

Masuda, N., T. Ito, K. Kayama, H. Kono, S. Satake, T. Kunugi, and K. Sato, 2006: Special purpose computer for digital holographic particle tracking velocimetry. *Optics Express*, **14**, 587–592.

McFarquhar, G. M., J. Um, M. Freer, D. Baumgardner, G. L. Kok, and G. Mace, 2007: Importance of small ice crystals to cirrus properties: Observations from the tropical warm pool international cloud experiment (twp-ice). *Geophys. Res. Lett.*, **34**, L13803.

Meng, H., W. L. Anderson, F. Hussain, and D. D. Liu, 1993: Intrinsic speckle noise in in-line particle holography. *J. Opt. Soc. Am. A*, **10**, 2046–2058.

Meng, H., G. Pan, Y. Pu, and S. H. Woodward, 2004: Holographic particle image velocimetry: from film to digital recording. *Meas. Sci. Technol.*, **15**, 673–685.

Nagel, D., 2003: personal communication, gKSS-Research Centre, Institute for Coastal Research, Max-Planck-Strasse 1, D-21502 Geesthacht, Germany.

Nagel, D., U. Maixner, W. Strapp, and M. Wasey, 2007: Advancements in techniques for calibration and characterization of in situ optical particle measuring probes, and applications to the fssp-100 probe. *joat*, **24**, 745–760.

Owen, R. B. and A. A. Zozulya, 2000: In-line digital holographic sensor for monitoring and characterizing marine particulates. *Opt. Eng.*, **39**, 2187–2197.

Pan, G. and H. Meng, 2003: Digital holography of particle fields: reconstruction by use of complex amplitude. *ao*, **42**, 827–833.

Pruppacher, H. R. and J. Klett, 1997: *Microphysics of clouds and precipitation*. Kluwer, Boston.

Pu, Y. and H. Meng, 2000: An advanced off-axis holographic particle image velocimetry (hpiv) system. *Exp. Fluids*, 184–197.

— 2003: Intrinsic aberrations due to mie scattering in particle holography. *J. Opt. Soc. of. Am. A*, **20**, 1920–1932.

Raupach, S. M. F., H. J. Vossing, J. Curtius, and S. Borrman, 2006: Digital crossed-beam holography for in situ imaging of atmospheric particles. *J Optic. Pure Appl. Optic.*, **8**, 796–806.

Satake, S., H. Kanamori, T. Kunugi, K. Sato, T. Ito, and K. Yamamoto, 2007: Parallel computing of a digital hologram and particle searching for microdigital-holographic particle-tracking velocimetry. *ao*, **46**, 538–543.

Schnars, U. and W. P. O. Juptner, 2002: Digital recording and numerical reconstruction of holograms. *Meas. Sci. Technol.*, **13**, R85âĂŞ–R101.

Shaw, R. A., 2003: Particle-turbulence interactions in atmospheric clouds. *Ann. Rev. Fluid Mech.*, **35**, 183–227.

Sheppard, C. J. R. and M. Hrynevych, 1992: Diffraction by a circular aperture: a generalization of fresnel diffraction theory. *J. Opt. Soc. Am. A*, **9**, 274–281.

Slimani, F., G. Grehan, G. Gouesbet, and D. Allano, 1984: Near-field lorenz-mie theory and its application to microholography. *ao*, **23**, 4140–4148.

Soontaranon, S., J. Widjaja, and T. Asakura, 2002: Direct analysis of in-line particle holograms by using wavelet transform and envelope reconstruction method. *Optik*, **113**, 489–494.

Soulez, F., L. Denis, C. Fournier, E. Thiebaut, and C. Goepfert, 2007: Inverse-problem approach for particle digital holography: accurate location based on local optimization. *J. of Opt. Soc. of America: A*, **42**, 1164–1171.

Strapp, J. W., F. Albers, A. Reuter, A. V. Korolev, U. Maixner, E. Rashke, and Z. Vukovic, 2001: Laboratory measurements of the response of a pms oap-2dc. *joat*, **18**, 1150–1170.

Thompson, B. J., 1974: Holographic particle sizing techniques. *J. Phys. E Sci. Instrum.*, **7**, 781–788.

Twohy, C. H. and D. Rogers, 1993: Air-flow and water-drop trajectories at instrument sampling points around the beechcraft king air and lockheed electra. *joat*, **10**, 566–578.

Tyler, G. A. and B. J. Thompson, 1976: Fraunhofer holography applied to particle size analysis: A reassessment. *Opt. Acta*, **23**, 685–700.

Uhlig, E., S. Borrmann, and R. Jaenicke, 1998: Holographic in-situ measurements of the spatial droplet distribution in stratiform clouds. *Tellus*, **50B**, 377–387.

Xu, W., M. H. Jericho, I. A. Meinertzhagen, and H. J. Kreuzer, 2002: Digital in-line holography of microspheres. *ao*, **41**, 5367–5375.

Yang, P., B. A. Baum, A. J. Heymsfield, Y.-X. Hu, H.-L. Huang, S.-C. Tsay, and S. A. Ackerman, 2003: Single scattering properties of droxtals. *J. Quant. Spectrosc. Radiat. Transfer*, **79**, 1159–1169.

Yang, W. D., A. B. Kostinski, and R. A. Shaw, 2005: Depth-of-focus reduction for digital in-line holography of particle fields. *Opt. Lett.*, **30**, 1303–1305.

— 2006: Phase signature for particle detection in digital holography. *Opt. Lett.*, **31**, 1399–1401.

Zhang, Z. B., P. Yang, G. W. Kattawar, S.-C. Tsay, B. A. Baum, Y. X. Hu, A. J. Heymsfield, and J. Reichardt, 2004: Geometrical-optics solution to light scattering by droxtal ice crystals. *ao*, **43**, 2490–2499.