

Retrieval of polar mesospheric cloud properties from CIPS: algorithm description, error analysis and cloud detection sensitivity.

J.D. Lumpe¹, S.M. Bailey², J.N. Carstens², C.E. Randall^{3,4}, D.W. Rusch³, G.E. Thomas³, K. Nielsen¹, C. Jeppesen³, W.E. McClintock³, A.W. Merkel³, L. Riesberg³, B. Templeman³, G. Baumgarten⁵, and J.M. Russell, III⁶

¹ Computational Physics, Inc., Boulder, CO

² Bradley Department of Electrical and Computer Engineering, Virginia Tech,
Blacksburg, VA

³ Laboratory for Atmospheric and Space Physics, University of Colorado, Boulder, CO

⁴ Also Department of Atmospheric and Oceanic Sciences, University of Colorado,
Boulder, CO

⁵ Leibniz-Institute of Atmospheric Physics at the Rostock University, Kuhlungsborn
Germany

⁶ Center for Atmospheric Sciences, Hampton University, Hampton, VA

Abstract

The Cloud Imaging and Particle Size (CIPS) instrument has been in operation on the NASA Aeronomy of Ice in the Mesosphere (AIM) satellite since May 2007. CIPS is a multi-camera UV imager that makes unprecedented hemispheric-scale measurements of polar mesospheric clouds (PMC). The primary CIPS data products are cloud frequency, albedo, mean particle radius, ice water content and vertical column particle density. These quantities are retrieved at 25 km^2 resolution at latitudes between ~ 55 and 84 degrees over a range of local times in the summer hemisphere. CIPS has obtained data for six Northern Hemisphere and five Southern Hemisphere PMC seasons to date and is still in operation and performing flawlessly. The CIPS data are made available to the scientific community in a variety of formats and spatial and temporal resolution, including full-resolution single-orbit level 2 data files and images, daily (hemispheric) albedo maps and images, and full-season latitude-binned summary files. In this paper we describe the CIPS measurement strategy and sampling characteristics, calibration and the Version 4.20 processing algorithms and retrievals. We also provide a quantitative evaluation of the CIPS cloud detection sensitivity and estimated random and systematic errors of the V4.20 cloud data products.

1. Introduction

The Cloud Imaging and Particle Size (CIPS) instrument was launched onboard the Aeronomy of Ice in the Mesosphere (AIM) satellite into a circular, sun-synchronous, 600-km orbit on April 25, 2007. AIM is the first satellite mission dedicated solely to the study of polar mesospheric clouds (PMC), and has provided unique measurements that directly address the issues of PMC formation, morphology and variability, and their connection to the meteorology of the polar mesosphere (Russell *et al.* 2009). CIPS is one of three instruments on the AIM satellite, along with the Solar Occultation for Ice Experiment (SOFIE) (Gordley *et al.*, 2009) and the Cosmic Dust Experiment (CDE) in-situ dust detector (Poppe *et al.*, 2011).

CIPS is a nadir-imaging instrument that utilizes a unique four-camera design to discriminate the scattering of ultraviolet (UV) solar photons from PMC ice particles against the background sunlit atmosphere. Cloud single-scatter albedo is measured at high spatial resolution over a range of scattering angles, thus providing a direct measurement of the PMC ice scattering phase function. Through analysis of the measured phase function it is possible to retrieve fundamental microphysical properties of the cloud particles. CIPS began routine measurements on May 24, 2007 and has operated flawlessly since that time. To date it has observed six full PMC seasons in the Northern hemisphere (NH) and five in the Southern hemisphere (SH). Approximately 97% of all possible data has been successfully downloaded and processed.

PMCs have been studied from numerous satellite-borne instruments over the past 30 years, using a variety of remote sensing techniques. These include both nadir- and limb-

viewing geometries, measurements of extinction (via solar and stellar occultation) as well as scattered solar radiance, and utilizing spectral bands from the UV through the infrared (IR). Two primary methods are available for deriving information on the PMC ice particle size parameters. The majority of measurements employ a spectral approach, where the wavelength dependence of the PMC scattering or extinction signal is used to constrain the particle size via comparison to optical model calculations. The alternative approach, employed by CIPS, is to measure the angular dependence of PMC scattering, e.g., the scattering phase function. A small number of direct measurements of the PMC phase function have been previously obtained from rocket experiments (Gumbel et al., 2001) while the satellite-borne Student Nitric Oxide Explorer (SNOE) (Bailey et al., 2005) and Solar Mesosphere Explorer (SME) (Thomas 1984) instruments both made limited (bidirectional) measurements of the phase function, from which particle size information was inferred (Thomas and McKay, 1985; Rusch et al., 2008). CIPS represents the first comprehensive space-based measurements of the PMC phase function over the entire polar summer hemisphere.

The scientific validity of the CIPS data has already been established through its use in a variety of scientific and validation analyses. Comparison of CIPS cloud frequency and albedo to measurements from the solar backscatter ultraviolet (SBUV/2) instrument on NOAA-17 show consistent agreement between these two nadir-viewing datasets when the CIPS data is degraded to the SBUV resolution (Benze et al., 2009; 2011). Baumgarten et al. (2012) have analyzed CIPS data obtained in close coincidence with ground-based lidar measurements and found good agreement in the cloud brightness observed by these two very different methods. The detailed spatial structures observed by CIPS have been used to study mesospheric gravity waves (Chandran et al., 2009; 2010, 2012) and planetary waves (Merkel et al., 2009), while the CIPS ice water content has been used by Stevens et al. (2010) to analyze the effect of tidal signatures on PMC. CIPS cloud frequencies were used in Karlsson et al. (2011) to connect SH PMC variability with the breakdown of the wintertime SH stratospheric polar vortex. The SH intra-seasonal PMC variability observed by CIPS was also used to investigate inter-hemispheric coupling in Karlsson et al. (2009). Stevens et al. (2012) used the CIPS observations of PMC frequency and albedo in July 2011 to help demonstrate a causal link between the occurrence of very bright clouds and the main engine exhaust from the space shuttle's final flight. A more extensive list of publications based on CIPS data can be found at <http://lasp.colorado.edu/aim/publications.html>. Two closely related companion papers in this issue complement the current paper in addressing CIPS data quality and validation from different perspectives. Bailey et al. (2012, this issue) presents a detailed analysis of coincident CIPS and SOFIE measurements in the AIM common volume while Carstens et al. (2012, this issue) provides a theoretical analysis of the information content and error characteristics of the CIPS measurements.

Bailey et al. (2009) provided a brief description of an early CIPS data version and associated algorithms. This paper addresses the current CIPS V4.20 data version. The goals of the paper are: 1) to describe the key features of the CIPS instrument, measurement technique, and sampling characteristics, 2) to provide a thorough description of the V4.20 data processing and retrieval algorithms, 3) to summarize the

available CIPS data products, and 4) to provide quantitative error estimates for those products. The paper is organized as follows. Section 2 summarizes the basic instrument design, measurement technique, and sampling characteristics of the CIPS instrument. Section 3 describes the level 1 data processing, summarizing the key steps used to calibrate the level 1a images and then merge data from the four cameras to produce the critical level 1b data. In Section 4 we present a comprehensive description of the CIPS level 2 algorithms used to separate the cloud signal from the Rayleigh background and retrieve PMC data products. In section 5, we present an analysis of the CIPS out-of-season data to quantify the false cloud detection rate in the data set. Section 6 presents results from a detailed retrieval simulation study, yielding quantitative estimates of the CIPS cloud detection sensitivity and retrieval errors. In Section 7, we briefly describe the publicly released CIPS data products and give examples of each, followed by a summary and conclusions in Section 8.

2. CIPS measurement technique and sampling

McClintock et al. (2009) describes the CIPS instrument in detail. CIPS consists of four nadir-viewing UV sensitive cameras that together provide an instantaneous field of view of 120° (along-track) by 80° (cross-track). The spectral passband is centered at 265 nm with a width of ~ 15 nm (full-width half-maximum). Figure 1 illustrates the orientation of the cameras, which are arranged in a cross pattern with overlapping pixels along the inner edges. We define a spacecraft-centered coordinate system with X-Y-Z axes corresponding to the roll (along-track), pitch (cross-track) axis and yaw (nadir) directions, respectively. The PX and MX cameras are pointed in the +X and -X direction, with similar convention for the Y cameras. The Y cameras are offset $\pm 19^\circ$ along the pitch axis and thus are primarily nadir-pointed. The X cameras are offset $\pm 39^\circ$ along the roll axis and therefore project out to larger off-nadir view angles along the orbit track. Figure 2 illustrates the distribution of offset angles across the four cameras, relative to the spacecraft nadir (Z) axis. The resulting geometric footprint, when projected to the nominal 83-km cloud deck altitude, is characterized by the along-track and cross-track distance scales shown in Figure 1. During routine science data operations the AIM satellite is always oriented such that the PX camera is facing the sunward direction. This corresponds to the forward (ram) direction for NH measurements and the backward (anti-ram) direction in the SH. The yaw (Z) axis is nominally aligned with satellite nadir during the science image sequence, with the notable exception of the CIPS common volume measurements, which are described later in this section.

CIPS operates only in the summer hemisphere, switching between NH and SH measurement modes at the equinoxes. The data sequences executed in the two hemispheres are essentially mirror images of one another. On each orbit, measurements are obtained over a solar zenith angle (SZA) range from approximately 20 degrees to 105 degrees (well past the terminator). The NH measurement sequence starts with three “first-light” images taken by the forward-looking PX camera as it leads the transition from darkness into the sunlit portion of the orbit. The sequence begins approximately six minutes after completion of the SOFIE sunrise occultation, with the first image taken at a solar zenith angle of ~ 105 degrees. 27 full CIPS scenes follow the first-light images. A

scene is comprised of simultaneous images from all four cameras and spans an area approximately 2000 km along-track by 1000 km cross-track (see Figure 1). The image integration times are 0.75 and 1.3 seconds for the X and Y cameras, respectively, and the measurement cadence (time between successive images) is 43 seconds. A typical orbit of science data therefore consists of 30 PX images and 27 images from the other three cameras, and takes roughly 20 minutes to execute. The SH sequence is essentially the time-reverse of the NH, with a series of 27 scenes beginning at low solar zenith angle followed by three PX “last-light” images beyond the terminator. The sequence ends approximately six minutes before the SOFIE sunset occultation scan. The CIPS level 1a data consist of these calibrated and geolocated science images for each orbit (see Section 3.1 for a more complete discussion of level 1a processing).

The CIPS data sequences are illustrated in Figure 3, which shows the footprint of every 6th scene on a polar projection map for typical orbits at summer solstice in each hemisphere. These plots illustrate the CIPS camera orientation, where the PX camera leads in the NH while the MX camera leads in the SH but in both cases PX is facing the sunward direction. Figure 4 provides more detail on the data sampling for the same two solstice orbits shown in Figure 3, with the top and bottom rows corresponding to NH and SH orbits, respectively. The three columns display the latitude, solar zenith angle and local time as a function of elapsed time in the orbit sequence, with one data point per science image. The solid vertical line in each panel separates the ascending and descending portions of the orbit (AIM is ascending on the night side and descending on the dayside of the orbit; see Figure 3). The dashed vertical line indicates the first (NH) or last (SH) time-stamp for complete scenes, and thus serves to separate the PX first- and last-light images.

The solar zenith angle and local time samplings are essentially constant throughout the season (not shown). The solar zenith angle ranges from a minimum of ~20 degrees on the descending node to well past the terminator (~105 degrees) on the ascending portion of the orbit. The corresponding local times range from 12-24 hours in the NH and 0-12 hours in the SH, with “noon” and “midnight” measurements being made on the descending and ascending portions of the orbit, respectively, in both hemispheres. The latitude sampling obtained from this range of solar zenith angles varies as the terminator latitude changes throughout the season. At solstice the latitude is roughly symmetric on the ascending and descending nodes of the orbit (see Figure 3), but this symmetry becomes skewed moving towards the equinoxes as the terminator latitude approaches the pole. To illustrate this point the latitude sampling for a typical end-of-season orbit (70 days after solstice) is shown by the dotted lines. At this time, when the terminator is close to the pole, the ascending node samples a much smaller range of high latitudes while the descending node extends to much lower latitudes. All of these sampling patterns are highly repeatable from year-to-year in both hemispheres.

Figure 5 illustrates the basic CIPS viewing geometry and defines some fundamental quantities that will appear in the following discussion. This figure shows a ray-trace for a single CIPS camera pixel through the atmosphere intersecting a scattering volume at geometric altitude z . The spacecraft view angle ψ , defined as the angle between the line-

of-sight (LOS) vector and the spacecraft nadir, is equivalent to the pixel offset angle plotted in Figure 2 when the satellite Z axis is pointed to true nadir. The most important quantities for describing the scattering geometry, and hence the relevant CIPS radiative transfer problem, are the solar zenith angle, ϕ , view angle, θ , and scattering angle, Φ . Note that “view angle”, as defined here and used in the remainder of this document, is the angle between the LOS vector and the local zenith, at the scattering point. All of these quantities are dependent on the altitude of the scattering point, z . Two physically distinct altitudes are important to the CIPS measurements – the cloud altitude at ~ 83 km and the peak altitude of the Rayleigh scattering contribution function, which is typically ~ 55 km but rapidly shifts to higher altitudes at solar zenith angles greater than ~ 85 deg.

Figure 6 illustrates the typical range of view and scattering angles sampled by each camera over an orbit, and their systematic dependence on solar zenith angle. This sampling is essentially identical in the two hemispheres. From this figure it is apparent that the smallest scattering angles (most forward-scattering measurements) are always obtained by the PX camera. This is a consequence of the fact that PX faces in the sunward direction as described previously. By comparison, the MX camera, pointing in the opposite direction, always measures back-scattered photons ($\Phi > 90$ degrees). The nadir (Y) cameras fill in the gap in the middle and are largely overlapping in $[\theta, \Phi]$ space (although of course not physically in terms of geometric footprint). The view angle sampling is what one would expect from Figure 2, with the Y cameras predominately pointing near-nadir and the X cameras sampling larger off-nadir angles. Note that the scattering point view angle θ is systematically larger than the spacecraft view angle, ψ , as illustrated in Figure 5.

One additional critical feature of the CIPS measurement sampling apparent from Figure 6 is that the smallest scattering angles (most forward-scattering geometry) are obtained at high SZA on each orbit. As the satellite position varies from the terminator (high-SZA) to the noon (low-SZA) portion of the orbit, the minimum scattering angle increases steadily, but is always smaller than 90 degrees. Therefore, CIPS always sees some forward scattering, with $\Phi_{\min} \sim 60$ -70 degrees in the 40-50 degree SZA bin (compared to 20-30 degrees in the 90-95 degree SZA bin). This systematic trend in scattering angle sampling over an orbit has important implications for the CIPS cloud detection sensitivity and retrieval errors, as discussed in Sections 5 and 6.

As mentioned previously the CIPS common volume (CV) measurements represent a departure from the nominal nadir orientation of the CIPS cameras during the science data sequence. The CV is the region of overlapping measurements between CIPS and SOFIE, which occurs near the terminator where SOFIE makes its solar occultation scan. When the solar beta angle (the angle between the satellite orbit plane and the sun) is less than 3 degrees the CV will naturally fall within the normal CIPS measurement swath. However, for larger beta angles the satellite must be rolled to point CIPS at the SOFIE occultation point in order to insure the common volume is sampled. In the NH this maneuver occurs at the beginning of the orbit sequence and the first seven science images are obtained in this rolled orientation, after which the satellite is immediately reoriented to normal nadir viewing for the remainder of the orbit. The reverse sequence occurs in the SH, where the

satellite is rolled for the final seven science images. The solar beta angle has a seasonal variation as well as a long-term increasing trend due to the decay of the AIM orbit, so that the roll maneuver was activated and de-activated as needed during the first four years of AIM operations. It was permanently disabled in February 2010 when the beta angle exceeded its maximum value and therefore no more CV data is obtained after that time.

3. Level 1 data processing

In this section we summarize the primary steps in the CIPS level 1 data processing algorithms. Level 1a data contains the calibrated and geolocated science images and associated metadata for a single camera per orbit. Level 1b data consists of “stacks” of coincident data from all four cameras and is the input data stream for the primary level 2 cloud retrieval algorithm. The CIPS camera alignment, pixel view angle distribution and measurement cadence described in Section 2 were chosen so that, by combining data from all four cameras, CIPS can observe a given air parcel nearly simultaneously at multiple scattering angles. This is the essence of the CIPS measurement technique, which provides a direct measurement of the PMC ice scattering phase function.

3.1 Level 1a data

McClintock et al. (2009) provide a comprehensive discussion of the CIPS instrument design and characterization, including a systematic description of the level 1a calibration and on-board processing procedures. Here we will simply summarize the highlights. The four CIPS cameras each utilize a square, 4-megapixel CCD detector with 2048 pixels on each side. During normal operations, only the central 1360 x 1360-pixel region of the CCD is used for imaging. To satisfy data transmission rate constraints the raw data from this central detector array is binned in three sequential steps – 2x2 on chip, 1x2 on the Field Programmable Gate Array (FPGA) processor and 2x2 in software - to produce the final “science pixels” that are telemetered to the ground. The cumulative binning is therefore 4x8, for a combination of 32 CCD pixels per science pixel. This yields a final level 1a image format of 170x340 pixels (cross-track x along-track).

The spatial resolution of the level 1a data is determined by the geometric footprint of each square pixel when projected onto a given altitude surface in the atmosphere. The size and shape of this footprint is determined by the spacecraft view angle, ψ , which varies across each camera image and is different for the X and Y cameras (see Figure 2). At nadir viewing geometry ($\psi=0$) the individual CCD pixel footprints are square and have their minimum size (highest spatial resolution) whereas for larger off-nadir view angles the footprint increases in size and becomes trapezoidal in shape. This effect is greatest at the outer edges of all cameras, with the leading and trailing edge pixels of the X cameras being the most extreme case. After combining the effects of finite camera point-spread-functions, exposure smearing and binning, the effective 1a pixel resolution for the Y cameras is ~ 1.1 km x 2.1 km (along-track x cross-track) for the inner edge pixels, and ~ 1.5 km x 3.8 km at the outer edges. For the X cameras the values range from ~ 1.2 km x 2.4 km along the inner edge to ~ 6.0 km x 4.7 km at the center of the forward and trailing edges, with a maximum of ~ 6.7 km x 5.3 km in the extreme corner pixels.

The level 1a processing includes basic packet checking and noise spike filtering, time tagging, calibration and geolocation. The net result is radiometrically calibrated data in physical units (albedo). Each pixel in the 170x340 science image array is geolocated using the known pixel point spread functions, camera alignment offsets and satellite ephemeris (position and orientation) at each image time stamp. The geolocation algorithm traces the LOS vector for each pixel through the atmosphere and determines its intersection with a desired reference altitude layer. The location of the pierce point is determined (in latitude and longitude) and the associated view, scattering, and solar zenith angles are calculated for the scattering volume based on the known satellite and sun positions (see Figure 5). Geolocation is performed independently for two reference altitude surfaces – 83 km and 55 km, corresponding to the nominal cloud-deck altitude and the typical Rayleigh scattering peak altitude.

3.1.1 Level 1a image calibration

For a detailed description of the CIPS level 1a image calibration see McClintock et al. (2009). The calibration process includes the following primary steps: subtraction of electrical offsets and dark counts, non-linearity and gain correction, and camera flat fielding and normalization. Images from each camera are calibrated individually, and the resulting calibrated images and associated geolocation and metadata are stored in the level 1a data files (one file per camera per orbit). We note that while the fundamental steps of the calibration process never change, the specific methods used to implement some steps have changed over the CIPS mission in response to changes in measurement sequences and satellite commanding constraints. The most notable is the flat fielding and camera normalization analysis, which will be described in some detail here.

3.1.2 Camera flat-fielding

One critical element of the CIPS level 1a calibration sequence that has evolved considerably since launch is the camera flat-field correction. The flat field is the pixel-to-pixel variation across the camera due to the lens system and the photocathode. Each raw camera image contains a flat field variation that was mapped out in pre-flight laboratory testing. The variation, normalized to unity in the center of the image, provides a pixel-by-pixel correction factor that is applied to each science image in the calibration process.

The laboratory flat field correction was designed to remove all instrument-induced pixel-to-pixel variation from the cameras. However, on-orbit science images were found to contain residual variability across the detector. We refer to this residual as the Delta Flat Field (or Δ -flat for short), as it is a secondary correction to the laboratory flat fielding. This residual non-uniformity must be corrected in the calibration procedure to avoid systematic biases in the CIPS cloud retrievals. Additional calibration procedures were developed that make use of special operational datasets obtained on-orbit to characterize and remove this residual variation.

Ideally, obtaining an accurate estimate of the Δ -flat field requires a uniformly illuminated camera image, in order to minimize natural geophysical gradients across the image. On-orbit, this condition is best approximated at the subsolar point and in nadir viewing geometry. This scenario has the advantage of minimizing the solar zenith and view angles (thus minimizing scattering angle variation) as well as atmospheric variation across the image (e.g., ozone, which has minimal spatial variability at the near-equatorial subsolar point). New satellite and instrument commands were generated early in the AIM mission to obtain CIPS images from each camera at the subsolar point a number of times during the year. These “special 1a” images are taken either before (SH) or after (NH) the normal science data sequence. The AIM satellite is rotated so that one camera is pointed directly nadir for a series of images, and the sequence is rotated through the four cameras on sequential orbits. This special image data is calibrated by the standard method for normal 1a science images up through, but excluding, the final Δ -flat field correction.

For each measured subsolar image, a simulated image is calculated from a Rayleigh scattering forward model using the identical viewing geometry (for more details on the so-called C/σ Rayleigh model see Section 4). The measured image is then divided by the model image and the resulting ratio is normalized to unity at the image center, to isolate the pixel-to-pixel variation and eliminate any absolute offset between model and data. The resulting quantity, a 170x340 array, is the Δ -flat field correction. This process is performed for all subsolar images separately, and results for the entire season are averaged to minimize random noise on the corrections. The result is a final Δ -flat image for each camera for each season, which is then applied to each science image in the next-to-last step of the calibration process. The residual variation observed in the Δ -flat corrections is on the order of 4% for PY, MY and MX and up to 11% for PX, and shows significant structure across the camera. A sample set of Δ -flat images is shown in Figure 7. The top panel of this figure shows the Δ -flats for all four cameras in the NH 2010 season, illustrating the camera-to-camera morphology of the flat-field corrections. The lower panel isolates a single camera (PY), showing both the season-to-season trends in the Δ -flat field as well as small residual differences in the correction derived from NH and SH data (note the different color scale for the bottom panel). For all four cameras the year-to-year variability and long-term stability of the Δ -flat field arrays are generally well within 1-1.5%.

In calculating the model images for this analysis, the ozone column density parameter (see Section 4) is assumed constant across the image footprint, which necessarily introduces some error into the analysis. Comparison with results from an independent technique based on statistical analysis of overlapping pixels obtained from special fast-cadence subsolar images indicates that the operational V4.20 calibration could still have systematic flat field errors up to 1.5% across the cameras (along-track direction). While this error seems small, it does affect the retrieval of the dimmest clouds. A typical atmospheric background albedo measured by CIPS is approximately $200 \times 10^{-6} \text{ sr}^{-1} = 200 \text{ G}$ (in the remainder of this paper the symbol G is used to represent the fundamental albedo unit of 10^{-6} sr^{-1}). Because 1.5% of 200 G is 3 G, the threshold for cloud detection can vary by this amount across the detector. This is significant compared to the brightness of the dimmest clouds CIPS is able to detect, which are less than 5 G. New

calibration methods are being explored with the goal of reducing the systematic flat fielding errors in the next CIPS data version.

3.1.3 Modified CIPS measurement sequence

Beginning September 2011, due to constraints imposed by loss of satellite bitlock, CIPS was no longer able to make the “special_1a” subsolar measurements described above. However, from the beginning of the mission, CIPS has measured one extra scene every sixth orbit at low latitudes, outside the normal range of science images where PMCs occur. These special “low latitude” (LL) images are similar to the subsolar images in that they are obtained in conditions of relatively uniform illumination and low atmospheric variability. They therefore provide a suitable substitute for the subsolar images for the purposes of calibration and are available throughout the mission. While the viewing geometry is not as ideal (they are not at the subsolar point, nor are the cameras rotated to point nadir), this data set has the advantage that there are a large number of images, and they are available continuously throughout the year. The use of this data allows more Δ -flat images to be averaged and thus reduces random error, and in principle opens up the possibility of doing time-dependent calibration during a season. In addition, beginning in September 2011 the CIPS measurement sequence was modified to obtain two LL scenes every orbit (rather than 1 every 6 orbits), thereby increasing the data density significantly. All CIPS seasons starting with SH 2011/2012 use Δ -flat corrections obtained from this data source. A detailed regression analysis was performed to compare results from these two alternative calibration sources during overlapping time periods and the results confirmed minimal changes to the level 2 data products.

3.1.4 Camera normalization

The instantaneous field of view of each camera overlaps at least two other cameras as illustrated in Figure 1. Since the basic CIPS measurement technique involves combining spatially coincident measurements from different cameras made at different scattering angles (level 1b data, see below), it is critical that the calibration provide consistent normalization between the cameras. The Δ -flat field correction is concerned with fixing the correct pixel-to-pixel variation across each camera, and thus leaves undetermined an overall calibration constant. The final step of the calibration procedure is therefore to normalize the relative sensitivity of each camera to the others by forcing the ratios of observed albedo in overlapping camera pixels to be equal. One camera must be used as a standard against which the others are normalized and the MY camera is chosen because 1) it overlaps all other cameras and 2) it exhibits the most stable long-term trends.

The LL scene data is ideal for this purpose as it provides the overlapping pixel measurements required to calculate the inter-camera normalization. In addition, the LL data provides a self-consistent measurement of both the Δ -flat field for all four cameras, and the normalization factors, from every scene. Prior to the operational change in September 2011, this data set was used to derive the camera normalization factors, however it is now used for both the Δ -flat field and normalization analysis. A normalization factor is obtained for each camera (PX/MX/PY) from each LL scene by

forcing the mean albedo ratio in all overlapping MY pixels to be unity. These factors are averaged to calculate a final normalization factor per season, and this single number is multiplied into every science image as the final step in the calibration.

Figure 8 shows the normalization factors over the entire CIPS mission to date. Each vertical line represents a separation of PMC seasons (NH to SH measurements or vice versa). The red circles represent the median normalization in each PMC season and are the factors used in the final calibration process. For seasons prior to the measurement sequence change in September 2011, there are approximately 400 LL scenes available per season. After that point, the number increases to ~5000 per season. This sampling difference is apparent in the much higher density of points in the final two seasons.

The full season average, while always available for reprocessing of past data, is never available for routine operational processing of the current season. For all seasons prior to SH 2011/12 each new season was started using Δ -flat and normalization factors obtained from pre-season “special 1a” subsolar data, if available. Sometimes the new calibration data had not yet been obtained before the start of a cloud season, since commanding for these measurements required satellite bitlock to uplink commands, and AIM often went through significant periods without bitlock. In this situation operational processing of a season began using the previous year’s calibration data. This does not generally present a problem as the Δ -flats are consistent year-to-year. (The normalization factors, as Figure 8 shows, do exhibit trends at the 1-2%/year level, so this is more problematic). At the end of each season, a final calibration is calculated using all available calibration data obtained during the season and the season is then re-processed for consistency. The use of a different set of Δ -flat fields for each season can cause small discontinuities in the normalization factors between seasons, as seen in Figure 8.

3.2 Level 1b data

Level 1b processing is the critical step where data from all four cameras are merged to create the fundamental CIPS measurement quantity – total directional albedo as a function of scattering angle. The CIPS measurement technique is illustrated schematically in Figure 9. By combining overlapping images from multiple cameras along the orbit track we can identify spatially coincident measurements of the same scattering volume obtained from different viewing geometries, and hence over a range of scattering angles. The first step in this process is to map-project the level 1a data onto a common grid. The CIPS processing uses the Lambert Azimuthal Equal Area Polar Aspect Projection (hereafter Lambert Projection). This projection was chosen because it has minimal distortion near the center of the projection and is equal-area, meaning each 1b pixel represents the same area at the cloud deck. The map pixel size used in the V4.20 level 1b processing is 25 km^2 (~5x5 km), so a single 1b pixel represents a 25 km^2 area anywhere on the globe, and counting pixels is equivalent to summing in area.

Using the geolocation referenced to the 83-km cloud deck altitude, the 1a pixel latitudes and longitudes are first mapped into the Lambert projection X/Y coordinates for all images in the orbit. Each grid point value is obtained by interpolation, using only the 1a

pixels within the immediately surrounding X/Y grid box (extrapolation is not allowed). Individual images (i.e., cameras) are mapped separately since they are identified by a unique time stamp and characterized by a well-defined solar zenith angle (and hence scattering angle) range. In general the level 1b map grid encompasses a significant number of 1a pixels, but the 1a spatial sampling varies by camera and pixel position within each camera, as discussed in Section 3.1. The Y cameras and the inner (nadir) pixels of the X cameras have the smallest footprint (highest spatial density) while the X camera edge pixels have the coarsest sampling, which is comparable to the 1b grid resolution. By overlaying 1a images onto the Lambert projection grid in this way we form a level 1b data “stack”, where each grid point contains multiple spatially coincident measurements from different camera images, each with a different scattering angle. We refer to this 1b data quantity as a scattering profile, as it consists of total albedo vs. scattering angle.

Figure 10 provides a visual representation of part of the image “stack” for orbit 22600 on June 19 2011. Seven image footprints are shown, including PX images from scenes 16-18, MX images from scenes 18-20 and a single PY image from scene 19. All seven images contribute measurements to the subset of level 1b pixels represented by the black symbols in the overlapping region and therefore all of these pixels have a minimum of seven data points in the scattering profile. This figure illustrates how the 1a pixels from subsequent scenes come from different regions of each camera, moving from front to back (along-track). As discussed in Section 2 the PX camera samples the most forward scattering angles, while MX provides only back-scattered measurements and the nadir PY/MY cameras fill the gap in between. Since PY and MY only overlap over a very small range of edge pixels (see Figure 1) the vast majority of CIPS 1b scattering profiles contain data from one nadir camera or the other, but not both.

The data arrays contained in the 1b files (which are not a standard data product and are not publicly released) take the form of 3-dimensional arrays, where the first two dimensions correspond to spatial location on the map projection (essentially latitude/longitude) and the third dimension is the “stack” of coincident measurements obtained at that grid point. We refer to each such stack, which has a 25-km² footprint, as level 1b “pixel”. The scattering profile contained within each 1b pixel will contain data from at most seven different 1a scenes and therefore, while these measurements are spatially coincident, they can be separated by as much as 5 minutes in time (7 scenes at 43-second cadence). The total number of level 1b pixels in a CIPS orbit is typically around 350,000 and therefore the arrays can be quite large.

The number of data points per pixel, denoted NLAYERS, is an important diagnostic quantity in the CIPS data and retrieval characterization. Figure 11 illustrates the distribution of NLAYERS in the CIPS 1b data. The map on top shows the spatial distribution over a typical orbit. The densest sampling, with 6 to 10 points per pixel, occurs down the middle of the orbit strip where the cameras have the highest density of overlapping pixels. The number of points decreases towards the orbit edges where the level 1a spatial resolution degrades and the overlapping camera pixel pattern becomes less dense. This also occurs at both ends of the orbit strip, where the stacks contain only

pixels from a single X camera. The plot in the bottom panel summarizes the fractional distribution of the NLAYERS parameter over a CIPS orbit. This result, which is highly repeatable in both hemispheres and all seasons, shows that the most common sampling is 7 measurements per scattering profile, which occurs for approximately 30% of all pixels. All values from 1 to 6 occur with roughly equal frequency, at 10-13% of all pixels, while a small number (~ 7%) of pixels contain more than 7 data points. This occurs when the grid point coincides with the overlapping pixel regions on the camera inner edges, in which case more than one camera will contribute per scene.

4. Level 2 algorithm description

The fundamental CIPS data analysis problem is to separate the total albedo measured by the instrument into components due to Rayleigh scattering from the background atmosphere and scattering from PMC ice particles. When a cloud signal is detected, subtraction of the Rayleigh background provides a direct measurement of the PMC ice phase function, from which cloud microphysical parameters are derived. The cloud scattering signal is distinguished from the background Rayleigh signal by its scattering angle dependence. Figure 12 compares the normalized scattering phase function for Rayleigh scattering and PMC ice particles for a range of particle sizes. The Rayleigh scattering limit is characterized by a phase function that is symmetric about 90 degrees scattering angle, so that forward- and back-scattering look identical. The ice scattering phase function, on the other hand, is strongly forward-peaked, with the forward/back-scattered asymmetry increasing sharply for larger particle sizes.

Figure 13 illustrates the relative contribution of the cloud signal to the total measured albedo for a range of solar zenith angles and realistic cloud particle sizes and brightness. These results were generated using the CIPS forward model (described below) and realistic CIPS scattering geometry for each solar zenith angle. The Rayleigh background signal is modulated by the local ozone concentration, but is driven primarily by the solar zenith angle. It decreases monotonically with increasing solar zenith angle, with a rapidly increasing gradient above 85 degrees, approaching the terminator. Thus a given cloud will appear relatively brighter when viewed against the dark atmosphere near the terminator, as compared to the much brighter overhead sun conditions. This largely accounts for the increasing cloud fraction of the total albedo at higher solar zenith angles. Another contributing factor is the CIPS measurement geometry (see Figure 6), which provides more forward scattering at high solar zenith angles and thus increases discrimination between clouds and background, particularly for larger particles. This sampling bias has important implications for the CIPS cloud detection sensitivity, as the combination of increased background atmospheric brightness and loss of the peak forward-scattered cloud signature means that CIPS has more difficulty detecting dim clouds at low solar zenith angles than near the terminator. The cloud detection sensitivity therefore decreases towards the noon portion of the orbit in both hemispheres. This issue is discussed in more detail below and in Section 6.

4.1 CIPS forward model

The albedo measured by CIPS at a given location is a sum of contributions due to the background atmosphere plus cloud scattering, if a cloud is present:

$$A_{meas}(\theta, \varphi, \Phi) = A_{Ray}(\theta, \varphi, \Phi) + A_{cloud}(\theta, \Phi, r_o) \quad (1)$$

where θ , φ , and Φ are the view angle, solar zenith angle and scattering angle, respectively (see Figure 5) and r_o is some measure of cloud particle radius. The atmosphere is optically thick at 265 nm due to the large ozone absorption cross section, and multiple scattering is negligible in these optically thin clouds; therefore the single-scattering approximation can be made with no loss of accuracy (Thomas and McKay 1985). This makes the CIPS radiative transfer problem particularly simple. The essential task of the retrieval algorithm is to separate the two terms in Equation (1), remove the Rayleigh background from the measured signal, and retrieve effective cloud parameters from the resulting residual PMC scattering phase function. We therefore need a parameterization of these two terms that allows for a direct retrieval of both background and cloud parameters.

The Rayleigh background term may be parameterized in a concise analytical form by the so-called C/σ model. This model was first applied to the analysis of nadir Solar Backscatter Ultraviolet (SBUV) ozone data by McPeters (1980) and is generalized for the CIPS viewing geometry in Bailey et al. (2009). It describes the Rayleigh albedo for a given measurement geometry configuration (θ, φ, Φ) as a simple function of two variables: the ozone vertical column density above a reference altitude (C) and the ratio of the ozone to atmospheric scale height (σ , assumed constant). Specifically:

$$A_{Ray} = \frac{P_{Ray}(\Phi) \Gamma(\sigma + 1) \beta_{Ray} N_{air}}{\mu \left(\mu^{-1} + \text{ch}(\varphi, H_{O_3}) \right)^\sigma \left(\beta_{O_3} C \right)^\sigma} \quad (2)$$

Here P_{Ray} is the Rayleigh phase function, Γ is the Gamma function, β_{Ray} and β_{O_3} are the Rayleigh scattering and O_3 absorption cross-sections, respectively, H_{O_3} is the ozone scale height, and N_{air} is the total atmospheric column density. The ozone absorption cross section at 265 nm is taken to be $9.261 \times 10^{-18} \text{ cm}^2$, consistent with several currently available UV ozone cross section compilations, while the Rayleigh scattering cross section is $9.708 \times 10^{-26} \text{ cm}^2$ (Bates 1984). Both C and N_{air} are evaluated at the peak altitude of the Rayleigh scattering contribution function, which changes with solar zenith angle (for a discussion of how this dependence is handled in the CIPS algorithms see Appendix B). The quantity $\mu = \cos(\varphi)$ is the geometrical path-length factor corresponding to the photon path from the scattering point to the satellite (see Figure 5). The more general Chapman function, $\text{ch}(\varphi)$, is used to describe the photon path from the sun to the scattering point. For solar zenith angles less than 70 degrees it reduces to $\text{ch}(\varphi) \rightarrow \mu^{-1}(\varphi) = \sec(\varphi)$.

The cloud scattering contribution to the albedo is written in the following form:

$$A_{\text{cld}}(\theta, \Phi) = A_{\text{PMC}} P_{\text{ice}}(\Phi, r_o) \mu^{-1}(\theta) \quad (3)$$

where P_{ice} is the ice scattering phase function (normalized to unity at 90 degrees), r_o is the particle mode radius and A_{PMC} is the cloud single-scatter albedo referenced to nadir viewing geometry and 90° scattering angle ($\theta=0; \Phi=90$). Note that the use of any consistent mathematical representation of the CIPS ice phase function implicitly assumes that the cloud layer thickness is uniform over the 1b spatial grid scale, and remains constant over the ~5 minutes required to assemble a phase function measurement. Also the use of a constant slant path factor, μ , will cause a random retrieval error if there is a large tilt in the horizontal cloud surface, since this observations from different view angles will see different effective path lengths.

4.2 Cloud retrieval strategy

In previous CIPS data versions the Rayleigh background was retrieved pixel-by-pixel, simultaneously with the cloud parameters. This approach resulted in increased retrieval noise in the presence of clouds, requiring additional spatial binning of the 1b data to increase signal-to-noise at the expense of retrieval resolution. The Version 4.20 algorithm takes a fundamentally different approach, assuming that the underlying ozone parameters that modulate the Rayleigh albedo vary smoothly over an area much larger than the CIPS pixel resolution. This provides a smoother and more constrained background removal, allowing the cloud parameter retrieval to be implemented at significantly higher resolution than previous retrieval versions. The Version 4.20 algorithm operates directly on the level 1b data, so that the level 2 cloud data products are retrieved and archived on this same 25 km²-resolution grid. The level 1b and level 2 grids are therefore identical and can be referred to interchangeably. By comparison, previous Version 3.X data versions binned the 1b data 3x3 before input to the cloud retrievals, resulting in a much coarser retrieval resolution of 225 km².

The essential steps in the level 2 algorithm are to 1) remove the Rayleigh background from the measured total albedo in each CIPS pixel, 2) apply a quantitative detection algorithm to the resulting residual signal to determine if a cloud is present, and 3) if a cloud is detected, analyze the PMC scattering phase function to derive cloud microphysical parameters. The retrieval process is iterative and designed to incrementally improve the separation of background and cloud components in the measured signal. Cloud detections, and the characterization of cloud and background components, are updated and improved with each iteration by selectively fitting the data using the analytical forward model parameterizations in Equations (2) and (3).

The logical flow of the V4.20 algorithm is summarized in Figure 14. The objective of the primary iterative loop is cloud detection and separation/characterization of the Rayleigh background and cloud albedo for all pixels in the orbit. This is accomplished by retrieving the parameters $[C, \sigma]$ and $[A_{\text{PMC}}, r_o]$ in Equations (2) and (3), respectively. The final step uses the retrieved cloud albedo and particle size information to calculate ice water content (IWC) and ice vertical column density (ICD) for all cloud pixels.

Each iteration starts by isolating the background Rayleigh albedo over the orbit, which is estimated from the data by calculating the residual $A_{Ray}^{resid} = A_{meas} - A_{cloud}^{ret}$. Here A_{meas} is the measured level 1b total albedo and A_{cloud}^{ret} is the cloud signal retrieved from previously detected cloud pixels (described below). The first iteration necessarily assumes $A_{cloud}^{ret} = 0$, since no cloud detections have yet occurred. This residual signal is then fit with the analytic Rayleigh model to derive $[C/\sigma]$ parameters over the orbit, after which Equation (2) is evaluated to generate an updated Rayleigh albedo retrieval, A_{Ray}^{ret} , for all pixels.

These steps are described in more detail in Section 4.3 and Appendix A. The retrieved background is then subtracted from the data to isolate the residual cloud signal

$A_{cloud}^{resid} = A_{meas} - A_{Ray}^{ret}$, which represents a direct measurement of the PMC phase function when ice is present. A quantitative cloud detection algorithm, described in Section 4.4, is applied pixel-by-pixel to this residual to determine cloud presence or absence and the results are used to update the detected cloud pixel list. For cloud pixels the residual PMC phase function is fit using the cloud forward model expression in Equation (3) to derive the cloud albedo and particle radius values, using an analysis described in Section 4.5. Equation (3) is then evaluated for all cloud pixels to update the cloud albedo retrieval, A_{cloud}^{ret} (for non-cloud pixels $A_{cloud}^{ret} = 0$). At this point the iterative loop begins again with an updated estimate of the Rayleigh residual albedo.

By iterating these steps the algorithm achieves sequential improvement in the background retrieval and subtraction, as more cloud pixels are detected and the fits become more accurate. The current V4.20 algorithm does not implement any quantitative convergence criteria on this iterative process. Experience has shown that three iterations are sufficient to provide convergence and further iterations do not generally improve the results. Therefore the number of iterations has been fixed at three. The following sections provide more detail on the implementation of each step in the algorithm.

4.3 Rayleigh background characterization

The objective of this step is to characterize and remove the Rayleigh background over the entire orbit. To accomplish this the background parameters $[C, \sigma]$ must be retrieved so that Equation (2) may be applied to calculate A_{Ray}^{ret} . We make the assumption that these parameters can be represented solely by their solar zenith angle dependence: $C=C(\varphi)$ and $\sigma=\sigma(\varphi)$. The algorithm works by dividing the orbit into 0.25-degree SZA bins and retrieving a single $[C, \sigma]$ value in each bin. These bins are predominantly cross-track slices, spanning an area of approximately $25 \times 800 \text{ km}^2$ and typically containing ~ 1000 level 2 pixels. In the V4.20 algorithm retrievals are performed over a SZA range from 40 to 95 degrees, and thus the working grid has 221 elements. By constraining the ozone parameters to be constant across a bin, we ignore smaller-scale ozone gradients that can be predominantly zonal or meridional depending on the location in the orbit. Retrieval of the background in each bin is accomplished by fitting the Rayleigh residual albedo, A_{Ray}^{resid} , with a linear version of the C/σ model using a process described in Appendix A.

The complete profile of $C(\varphi)$ and $\sigma(\varphi)$ over the orbit is constructed using multiple sources of information. The initial fit is done bin-by-bin, deriving a single $[C, \sigma]$ pair for each bin. Two independent fits are done, one using all the data points in the bin and the other using only the back-scattered measurements at scattering angles ≥ 110 degrees. Because the primary signature of cloud scattering is an asymmetry between forward- and back-scattered albedo, in the absence of clouds the ozone column density derived from the two fits, C_{all} and C_{back} , should be the same within experimental noise and known systematic biases. Larger discrepancies between C_{all} and C_{back} indicate the presence of clouds, which will bias the fit. As an illustration, Figure 15 shows sample data in three different SZA bins for both a cloud-free post-season orbit (top row) and an in-season orbit with clouds (bottom row). The quantities plotted are the linear X/Y parameters used to derive $[C, \sigma]$ as described in Appendix A. The cloud-free data in the top panels conforms closely to the linear Rayleigh model and there is good agreement between the background parameters derived from fitting either all data or back-scattered data only. The cloud signature is readily apparent in the forward-scattering measurements from orbit 1240 on the bottom row, destroying the linearity of the background Rayleigh model. The $[C, \sigma]$ retrievals for this orbit are not meaningful and therefore are not shown.

A relative difference array $\delta(\varphi) = |C_{\text{all}}(\varphi) - C_{\text{back}}(\varphi)| / C_{\text{back}}(\varphi)$ is constructed and used to screen for cloud contamination. Based on trial and error a threshold difference of $\delta \sim 0.1$ has been found to be a reliable indicator of cloud presence. Therefore, all bins which satisfy $\delta < 0.1$ are assumed to have a good background determination and we take $C = C_{\text{back}}$ and $\sigma = \sigma_{\text{back}}$ for these bins. The fit derived from back-scattered data is assumed to be the most accurate because it minimizes the contribution due to any residual undetected cloud signal. This is especially critical in the first iteration where cloud removal has not yet occurred in the calculation of $A_{\text{Ray}}^{\text{resid}}$. The bottom panel of Figure 15 illustrates how the linear behavior of the X/Y relationship is generally preserved for back-scattered measurements even in the presence of significant cloud signal.

We fill in those bins where significant cloud presence is detected (as defined by $\delta > 0.1$) using climatological profiles $C_{\text{clim}}(\varphi)$ and $\sigma_{\text{clim}}(\varphi)$ obtained from fitting out-of-season data. These climatologies are updated each season by averaging fits over two-week time periods before and after the cloud season in each hemisphere. The climatological profiles are scaled to match the retrieved back-scattered values in the 40-70 degree SZA range and then used to fill in all bins missing due to significant cloud contamination. In practice the climatology is primarily used in the first iteration of the retrieval, where no cloud signals have yet been removed and the difference array δ is at its maximum. On subsequent iterations, as cloud pixels are detected and the Rayleigh residual converges the dependence on the climatology vanishes.

As a final step, the complete $C(\varphi)$ and $\sigma(\varphi)$ arrays are fit using a 4th order polynomial over the SZA range 40-85 deg. This smoothes out any bin-to-bin scatter from the fitting procedure and ensures a continuous background albedo over the orbit. For solar zenith angle bins above 85 degrees the retrieval of σ becomes noisier and less constrained by the data, and therefore the operational code holds σ fixed to the mean value retrieved

over the 80-85 degree range. The retrieved value of $C(\varphi)$ in these high-SZA bins responds to this constraint, adjusting as needed to fit the data. We have found no degradation in the ability of the C/σ model to fit the data at high solar zenith angle due to this adjustment, and the retrievals are more stable when this constraint is implemented. This result is consistent with the findings of Carstens et al. (2012, this issue).

With $C(\varphi)$ and $\sigma(\varphi)$ thus determined over the entire 40-95 degree SZA range, these arrays are then linearly interpolated to obtain $[C, \sigma]$ values for every level 2 pixel in the orbit. The retrieved Rayleigh albedo A_{Ray}^{ret} is then calculated for all data points using Equation (2) and the known view and scattering angles for each data point. The final step is to remove the Rayleigh background from the data to obtain the cloud residual

$$A_{cloud}^{resid} = A_{meas} - A_{Ray}^{ret}.$$

4.4 Cloud detection algorithm

The CIPS cloud detection algorithm is applied pixel-by-pixel to the residual cloud scattering profile A_{cloud}^{resid} . If no ice is present this residual simply represents the error in the background Rayleigh characterization, which can have both random and systematic components. The objective of the detection algorithm is to determine whether the residual signal represents a statistically significant enhancement relative to the estimated error in the Rayleigh background subtraction. The key to this problem is therefore to accurately quantify the error in the background removal, and use that error estimate to derive a detection threshold that optimizes the tradeoff between maximizing sensitivity to dim clouds while minimizing the number of false cloud detections. The Rayleigh subtraction error is quantified by analyzing statistics on the mean error (bias) and standard deviation of C/σ model fits to ensembles of non-cloud data. This analysis is updated every PMC season and the results are used to construct look-up tables (LUT) that are incorporated into the operational cloud detection algorithm. Details of the LUT construction are described in Appendix B, including some sample plots from the LUT used to analyze the NH 2010 season.

As described in Appendix B the LUT are constructed for each camera separately and contain the mean error (Δ_{mean}) and standard deviation (Δ_{std}) in uniform 1-degree grids in both solar zenith angle and view angle. In addition, there are separate LUT for both forward- and back-scattering measurements. These input files are utilized in the level 2 cloud detection algorithm in the following way. Mean error and standard deviation values can be assigned to each data point in the level 1b scattering profiles directly from the LUT files before the retrieval process begins. This is completely determined by information available in the 1b data: camera ID, solar zenith and view angles (which are rounded to the nearest degree and picked directly off the LUT grid, rather than interpolated), and scattering angle (to discriminate forward- vs. back-scattering). At this point these values are still “normalized errors”, relative to the (as yet undetermined) Rayleigh background albedo, as per Equation (B.2). During each iteration of the retrieval, as the Rayleigh background retrieval A_{Ray}^{ret} is updated, the mean and random errors are

converted to absolute albedo units by multiplying by the value of A_{Ray}^{ret} for each data point.

The mean error is added back to the retrieved Rayleigh albedo to correct for systematic bias in the Rayleigh subtraction. Note that this is equivalent to subtracting the mean error from the cloud residual signal, since

$$A_{Ray}^{ret} \rightarrow A_{Ray}^{ret} + \Delta_{mean} \Rightarrow A_{cloud}^{resid} \equiv A_{meas} - A_{Ray}^{ret} \rightarrow A_{cloud}^{resid} - \Delta_{mean} \quad (4)$$

Thus the residual cloud signal is “corrected” by the estimated bias in the Rayleigh background subtraction. The standard deviation term Δ_{std} is then used to define the cloud detection threshold. Each data point in the residual scattering profile is compared to a threshold value, and at least two points are required to exceed their threshold for a cloud detection to occur. This reduces the chance of a single noisy or bad data point triggering detection. The cloud threshold is taken to be a multiplicative factor of the Rayleigh standard deviation error for each data point, $A_{thresh} = \gamma \Delta_{std}$. The CIPS Version 4.20 algorithms use $\gamma = 2.4$ (i.e. 2.4 standard deviations above the mean) for this threshold. This factor was determined by trial and error to optimize the trade-off between false cloud detection and sensitivity to dim clouds (see the discussion in Sections 5 and 6). An additional constraint imposed on the detection criteria is to set a minimum (floor) value of $\Delta_{std} \geq 1$ G. This has no effect except at very high solar zenith angles where the absolute Rayleigh background albedo drops rapidly to zero and the relative errors compiled in the LUT become meaningless (see the discussion in Appendix B). Experience has shown the 1 G value to be a good estimate of the lower limit on the absolute albedo error obtained from fitting non-cloud data with the C/σ model.

The detection algorithm, and the subsequent step of fitting the cloud phase function described in Section 4.5, are performed on a pixel-by-pixel basis and thus operate only on the NLAYERS data points contained in a given pixel. For the orbit edge pixels, which have $NLAYERS \leq 3$ (see Figure 11) there are too few data points to reliably detect and analyze clouds. Therefore, for these pixels the data are first binned 3x3, co-adding data points from the 8 immediately adjacent pixels. This generally provides enough data sampling to perform the cloud detection and phase function fitting analysis, at the cost of degrading the effective resolution at the edges of the orbit. This technique tends to break down, however, for pixels at the extreme orbit edge where $NLAYERS = 1$. In this situation the adjacent pixels generally originate from the same camera and image and therefore have essentially identical view and scattering angles. No useful new information is added by binning in this case and the data quality for these pixels is low.

4.5 Cloud phase function fitting

For each pixel in which a cloud is detected the residual cloud signal is analyzed to retrieve the cloud albedo and particle size parameters archived in the CIPS level 2 data file. From Equation (3), the cloud albedo is represented analytically as

$$A_{cloud}^{resid} \equiv A_{cloud}(\theta, \Phi) = A_{PMC} P_{ice}(\Phi, r_o) \mu^{-1}(\theta) \quad (5)$$

Here we have written the explicit dependence on view angle and scattering angle (the SZA dependence is implicit). The parameters A_{PMC} and r_o are determined directly by fitting a theoretical ice phase function, P_{ice} , to the measured phase function A_{cloud} . P_{ice} is calculated using the T-matrix code described in Mishchenko and Travis (1998). The ice refractive index is taken from Warren (1984), and has the values 1.357090 and 1×10^{-8} for the real and imaginary parts, respectively.

The CIPS measurements contain no direct information about the shape of the PMC ice particles or the specific form of the underlying particle size distribution, and therefore these fundamental properties have to be assumed in calculating P_{ice} . The physical assumptions made about ice particle shape and size distribution are constrained to be consistent with PMC measurements made by lidar (Baumgarten et al., 2010), and also with the ice retrievals from SOFIE (Hervig et al., 2009). The particles are assumed to be oblate spheroids with an axial ratio of 2 and the size distribution is assumed to be Gaussian. The CIPS retrieved particle size parameter r_o in Equations (3) and (5) is then interpreted as the mean radius of the Gaussian distribution, where particle size is defined by the volume-equivalent sphere radii. Another critical parameter is the Gaussian distribution width, s . The lidar data retrievals show that the distribution width is a function of particle size, and that this relationship depends on altitude within the cloud layer. We have adopted a parameterization of the $s(r_o)$ relationship derived from the vertically integrated lidar data, as appropriate for the CIPS nadir viewing geometry (Baumgarten et al. 2010). This dependence is represented by the solid red curve in Figure 22 (the rest of this figure is discussed in Section 6.3). The distribution width varies as approximately $0.39 \cdot r_o$ for radii up to 40 nm and then stays essentially fixed at ~ 15.8 nm for larger particles. To eliminate unphysical artifacts in the CIPS retrievals the relationship derived in Baumgarten et al. (2010), which is a piece-wise linear fit with discontinuous derivative at $r_o = 40$ nm, has been smoothed at the inflection point to give the dashed red curve in Figure 22 which is used operationally.

With these assumptions built into the T-matrix calculations, the normalized theoretical phase function P_{ice} depends on a single free parameter – the Gaussian mean radius r_o . The CIPS residual cloud phase function, normalized to the slant path factor, is fit to the theoretical ice phase function for a range of particle radii from 1 to 100 nm. For any radius value the albedo is determined by minimizing the squared error between data and model as follows:

$$E^2 = [\mathbf{d} - A_{PMC} \mathbf{P}_{ice}]^2$$

$$\frac{\partial E}{\partial A_{PMC}} = 0 \Rightarrow A_{PMC} = \frac{\mathbf{d} \cdot \mathbf{P}_{ice}}{\mathbf{P}_{ice}^2} \quad (6)$$

where

$$\mathbf{d} = \mathbf{A}_{cloud}^{resid} \boldsymbol{\mu} \quad (7)$$

and the vector notation represents the dependence on scattering/view angle (NLAYERS points per pixel). A χ^2 metric is calculated for every r_o value as follows:

$$\chi^2(r_o) = \sum_{i=1}^{NLAYERS} \frac{[d^i - A_{PMC} P_{ice}^i(r_o)]^2}{2|A_{meas}^i|} \quad (8)$$

The retrieved particle radius is chosen to be that value for which $\chi^2(r_o)$ is minimized, and the corresponding value of A_{PMC} is given by Equation (6).

As mentioned in Section 3.2 a typical CIPS orbit contains ~350,000 level 1b (hence level 2) pixels. In the middle of the PMC season it is common for the CIPS orbit-averaged cloud fraction to exceed 50%, meaning that the CIPS level 2 data file contains ~175,000 individual cloud phase functions. Figure 16 illustrates the results of this fitting procedure for a very small random sample of pixels from a single orbit. The measured (residual) PMC phase function is shown by the black symbols and the corresponding best-fit theoretical phase functions are plotted in red. The derived cloud albedo and particle radius for the pixel are written inside each panel, along with the SZA of the measurement. We have determined that the fitting procedure works quite well, especially for bright clouds in the CIPS data (albedos greater than ~10-15 G). Two general characteristics illustrated by the results in Figure 16 are that the residual cloud phase functions, and the corresponding fit and cloud parameter retrieval, tend to improve for brighter clouds and higher SZA. This is due to the fact that, for dim clouds and/or low solar zenith angles, the cloud albedo represents a smaller fraction of the total measurement and thus the retrieval is more sensitive to small errors in the background Rayleigh subtraction. That is why all of the phase functions at 55 degrees SZA (left column), no matter how bright the cloud, appear relatively noisier than those at higher SZA in the right columns. Also the brighter clouds at the top of each column generally yield better fits.

Finally, a word about errors in this procedure. The assumptions made about the ice particle shape and size distribution directly impact the particle radius and, to a lesser extent, the cloud albedo retrieved from CIPS data. While these assumptions are physically reasonable and strongly constrained by complementary PMC measurements, they are still assumptions and are unconstrained by the CIPS measurements themselves. Thus they represent a source of systematic error in the CIPS retrievals and these errors are examined in more detail in Section 6.3.

4.6 Calculation of ice water content and ice column density

After the last iteration is complete the final retrieved values of cloud albedo and particle radius are used to calculate IWC and ICD for all level 2 pixels in which a cloud was detected (step 8 in Figure 14). Previous CIPS retrieval versions (V3.X) utilized an empirical analysis in which the IWC was obtained directly from the retrieved cloud

albedo. This relationship was derived from a linear regression analysis of IWC vs. A_{PMC} generated from a detailed PMC microphysical model. In the V4.20 algorithms IWC and ICD are calculated directly from first principles using the retrieved albedo and particle size, and the assumed underlying particle size distribution used in the T-matrix phase function calculations. The calculation is straightforward, with the following definitions:

$$\begin{aligned}
 r_o &= \text{retrieved particle radius [nm]} \\
 A_{PMC} &= \text{retrieved cloud albedo at } 90^\circ \text{ scattering angle [} 10^{-6} \text{sr}^{-1}\text{]} \\
 V(r_o) &= \text{volume per ice particle [cm}^3\text{]} \\
 \sigma(r_o) &= \text{scattering coefficient at } 90^\circ \text{ scattering angle [cm}^2 \text{sr}^{-1}\text{]} \\
 \rho &= \text{ice density [g cm}^{-3}\text{]}
 \end{aligned}$$

Note that V and σ (not to be confused with the parameter σ in the C/σ model) come directly from the assumed particle size distribution and T-matrix calculations described in Section 4.5 and are both functions of the particle radius. Two microphysical quantities that can be derived from the combination of retrieved albedo and particle radius are then the ice column density

$$\text{ICD} = A_{PMC} / \sigma(r_o) \quad [\text{cm}^{-2}] \tag{9}$$

and the ice water content

$$\text{IWC} = \text{ICD} \cdot V \cdot \rho = \rho A_{PMC} \frac{V(r_o)}{\sigma(r_o)} \quad [\text{g km}^{-2}] \tag{10}$$

The dependence on the particle size, r_o , has been made explicit in the right-hand side of both expressions. For each cloud pixel in the CIPS level 2 orbit file the retrieved $[A_{PMC}, r_o]$ values are plugged into these equations to produce IWC and ICD.

5. CIPS false cloud detection rate

Understanding the false cloud detection rate in the CIPS data is important primarily because it provides a quantitative constraint on the accuracy of the CIPS level 2 cloud frequency data. Analyzing the frequency and morphology of false detections can also help to identify and diagnose retrieval artifacts and errors. Some number of false detections will always be present, due to both random measurement noise and systematic artifacts in the data that are not fully accounted for in the detection and retrieval algorithms. To a large extent, however it is possible to tune the detection criteria to achieve any desired false detection level. The primary tool available for this is the cloud detection threshold $A_{thresh} = \gamma \Delta_{std}$ discussed in Section 4.4. Specifically, increasing the factor γ will decrease the false detection rate, which can be effectively made as low as desired. The obvious trade-off is that this will also increase the number of real clouds that are missed by the detection algorithm on the low (dim) end of the albedo spectrum. The

reality is that there is a continuum of cloud brightness extending all the way down to the Rayleigh limit, and certainly below the effective noise floor of the measurements. Therefore it is critical to have the best understanding possible of what that noise floor is for the CIPS data and set the detection threshold accordingly. This information is quantified in the standard deviations Δ_{std} obtained from the error LUT described in Appendix B, and the operational factor γ has been set based on extensive testing of the algorithms on both cloudy and cloud-free data.

The CIPS false cloud detection rate can be quantified directly by considering cloud occurrence statistics in out-of-season data, where there are no clouds present. For this purpose we analyze the same two-week pre- and post-season time periods used for generating the Rayleigh error LUT, as described in Appendix B (specifically, 40-55 days before, and 70-85 days after solstice, respectively). Special diagnostics codes are routinely run on the level 2 data to calculate cloud fraction, median cloud parameters and many other quantities used for routine diagnostic analyses. Data from this analysis were averaged over these two-week periods to produce the results shown in Figure 17. The top panels of this plot show the mean orbit-average cloud fraction for each CIPS PMC season to date. Blue (red) curves correspond to NH (SH) data, and the solid and dashed lines represent pre- and post-season averages, respectively. The cloud fraction data has been further screened by the quality flag (QF) parameter in the level 2 data. In the Version 4.20 data set this flag is a proxy for NLAYERS, the number of independent measurements per pixel (see Figure 11). A QF value of 0 indicates 6 or more data points per pixel while QF=1 indicates 4 or 5 data points and QF=2 is used to flag edge pixels that have 3 or fewer data points. Pixels with QF=0 or 1, which occupy the center of the CIPS orbit swath and account for ~75% of the level 2 data, are considered the best quality whereas users are urged to use caution with the QF=2 data, since data sampling is sparser for these pixels.

The top left panel of Figure 17 shows the false detection rates for the best quality data, $QF \leq 1$. In the NH the rates are generally at or below the 1% level, whereas the SH rates are somewhat higher but still below 2%. Although the post-season rates do show some decrease over time in both hemispheres (particularly in the SH), in general there are no significant long-term trends in the false detection rates. The small hemispheric difference in false detection rates is not completely understood, but is likely related to differences in the background ozone parameters in the two hemispheres. A closer inspection of the SZA dependence of the false detections (not shown) reveals that they are not uniformly distributed over the orbit but tend to be localized at either ~40 degrees or close to 80 degrees. These features are thought to be artifacts of the Rayleigh retrieval algorithm, most likely due to small systematic errors from the final smoothing step on the $[C, \sigma]$ parameters. The 40-degree bin is the lower end point of the retrieval SZA grid and this artifact is likely due to small residual end-point errors in the fit. The 80-degree region, on the other hand, corresponds to the maximum curvature of the ozone column density, C , as it drops rapidly to zero at higher solar zenith angles. Thus it is the most difficult part of the profile to fit with the polynomial smoothing. This step will be improved in the next retrieval version, most likely by introducing a sliding spline smoothing function that can fit all regions of the SZA range. As the top right panel of Figure 17 shows, there are more

false cloud detections in the edge (QF=2) pixels, with rates approximately 1.5 to 2 times higher than the center pixels. This is not surprising as the sparser data sampling in these pixels generally correlates with more noise and small systematic artifacts in the data. Close inspection of the level 2 albedo image strips released with the V4.20 data often shows these small artifacts along the orbit edges.

The bottom panel of Figure 17 shows the median albedo retrieved from these false clouds averaged over the same 2-week periods. This gives a quantitative estimate of the typical magnitude, in terms of brightness, of the false cloud detections that appear in the CIPS data. The median albedo of these features is generally on the order of ~ 1 G or less, with slightly higher values in the SH than the NH, consistent with the higher false detection rate in the south. Thus an albedo level of 1-2 G is our best estimate of the CIPS retrieval noise floor. It is recommended in the level 2 data documentation that screening out clouds with albedo below this level is an effective way for users to eliminate false detections from the CIPS data.

6. Retrieval simulations, cloud detection sensitivity and error analysis

Detailed retrieval simulation studies were performed to quantify the CIPS cloud detection sensitivity and retrieval algorithm performance. The objective of the simulations is to generate realistic level 1b data containing simulated cloud signals from the CIPS forward model. Processing this data with the level 2 algorithms allows a direct comparison of the retrieved cloud parameters with the “true” cloud fields used in the simulation. To be as realistic as possible the simulated data is overlaid on actual level 1b data files, thus using the existing substrate of pixel sampling, measurement geometry (solar zenith angle, view and scattering angles) and the number of data points per pixel. In principle the data could be entirely simulated, by generating both the Rayleigh background and the cloud component for every pixel. However, this requires simulation of realistic ozone and atmospheric density/pressure fields over the entire orbit. More importantly, while realistic random noise could be added to the resulting total albedo field, it would not be possible to simulate all known or unknown systematic errors existing in the actual 1b data (e.g., calibration errors, camera mis-alignments and normalization errors, etc.). The alternative approach, adopted here, is to allow the Rayleigh background component to be defined by actual CIPS measurements and simulate only the cloud component. We therefore take as our starting point level 1b data for out-of-season orbits and simply add a simulated cloud component to the measured albedo. In this scenario the Rayleigh background, while unknown to us, is realistic in the sense that it contains real measurement noise as well as any biases and systematic errors contained in the actual CIPS 1b data. This provides the most realistic test of the algorithm in terms of the Rayleigh background subtraction and its interplay with the cloud detection and retrievals.

The simulated cloud component added to the 1b data is random in the sense that the cloud pixels are both randomly distributed in space and the albedo and particle radius for each cloud pixel are also generated randomly. The only constraint imposed on the cloud field is that a predetermined cloud fraction profile as a function of solar zenith angle is built into the simulation (see below). This allows one to study the convergence and accuracy

of the retrievals over a range of cloud fraction scenarios. The data simulator steps through the standard 0.25-deg SZA grid used in the level 2 retrievals, adding clouds one bin at a time. In each bin the number of cloud pixels is simply given by the total number of pixels in the bin times the assumed cloud fraction in that bin. The algorithm distributes these pixels randomly throughout the bin so that they are equally sampled in the cross-track dimension. This is important because we want to ensure a uniform sampling of clouds in both the middle and edge pixels, in order to test for retrieval artifacts sensitive to the NLAYERS parameter (e.g., see Figure 17).

For each assigned cloud pixel the PMC albedo and particle size (A_{PMC} and r_0) are chosen at random from a Gaussian distribution, the mean and standard deviation of which were tuned to mimic the distributions of retrieved cloud parameters observed in actual CIPS data. The albedo distribution had a mean value of 10 G and a width of 30 G while the mean and width of the radius distribution are 40 nm and 15 nm, respectively. For each cloud pixel the selected albedo and radius values, along with the level 1b geometry parameters $[\theta, \Phi]$, are then used in Equation (3) to calculate the cloud component, which is subsequently added to the measured 1b scattering profile. The forward model simulation uses the same theoretical ice phase function, P_{ice} , used in the retrievals and therefore the simulations do not evaluate the component of forward model error due to uncertainties in the phase function. This issue is addressed in a separate study presented in Section 6.3.

After adding the cloud signal the simulated level 1b data file is input to the level 2 retrieval codes, which are run in standard operational mode exactly as for real CIPS data. The resulting level 2 data files are then used to compare the retrieved cloud fraction and cloud parameter values against the known “true” values that went into the simulation. Ten orbits of simulated data were used in the analysis, corresponding to pre- and post-season orbits from both hemispheres and several CIPS PMC seasons. This provides a good random sampling of 1b data and the final results were derived by averaging over all ten orbits.

6.1 Cloud detection sensitivity

One important quantity that can be derived from the simulations is the CIPS cloud detection sensitivity. This is defined as the fractional probability that a cloud of given brightness will be detected by the CIPS algorithm. Due to the nature of the CIPS measurement sampling and the geophysical variation of the background Rayleigh brightness, we expect the detection sensitivity to depend on solar zenith angle (see Figures 6 and 13). This quantity is calculated from the simulations as follows: the data are binned in both solar zenith angle and albedo, using a discrete grid of cloud albedo values of interest. In each solar zenith angle bin we first identify all pixels with “true” albedo equal to the grid point value ± 0.5 G. We then isolate these same pixels in the level 2 retrievals and count how many of those clouds were actually detected. The detection rate is simply the ratio of these two numbers, expressed as a percent.

The results are plotted in Figure 18, in terms of detection rate as a function of solar zenith angle for representative cloud albedo values ranging from 2 to 10 G. For example, consider the black curve, corresponding to a true cloud albedo of 2 G. Recall from the discussion in Section 5 that this value is just above the suggested CIPS noise floor for screening false detections. According to Figure 18 clouds of this brightness are only detected ~30% of the time at the lowest solar zenith angles, whereas by 70 degrees SZA the detection rate rises to 60% and at the terminator (90 degrees SZA) even clouds this dim are detected almost 100% of the time. The detection rate increases rapidly with cloud brightness so that clouds just a factor of 2.5 brighter (5 G, the gold curve) are already detected with ~90% efficiency at the lowest solar zenith angles. Clouds with an albedo of 10 G or brighter are detected essentially 100% of the time everywhere. These curves provide guidance to CIPS data users about the confidence level to assign to cloud detections and retrievals at various brightness levels. For example, combined with the results of Figure 17, they suggest that clouds at the 5 G level and above are good detections to the 90% confidence level or greater anywhere in the CIPS orbit. Likewise, 3-G clouds are also most likely good detections based on Figure 17, but could be undercounted by as much as 40% at the lowest solar zenith angles.

6.2 Characterization of retrieval errors

6.2.1 Cloud fraction retrieval

We first consider the retrieved cloud fraction, which provides a direct measure of both missed cloud detections as well as the false detection rate in the simulations. For each orbit the data were binned in 2.5-deg SZA bins and the true and retrieved cloud fractions were calculated in each bin for various albedo threshold levels. The true cloud fraction is simply the fraction of pixels in a given bin for which clouds at or above the threshold albedo were added to the simulation. The retrieved cloud fraction is similarly given by the fraction of pixels where cloud detection occurred and the retrieved albedo is above the threshold value. The simulations were constructed using a constant total cloud fraction (all clouds) of 50% for solar zenith angles from 50 to 95 degrees. Low-latitude clouds are turned off by linearly forcing the cloud fraction to 0 between 50 and 40 degrees SZA.

Results are summarized in Figure 19 in terms of the cloud fraction error, defined as the difference between the retrieved and true cloud fraction (in %) for a given albedo threshold. Five different thresholds were considered, ranging from 0 (all clouds) to 10 G, which corresponds to a fairly bright cloud as observed by CIPS. These results are generated by considering only pixels with QF=0 or 1, and thus excluded the edge pixels (QF=2) with three or fewer data points. Note that positive errors as defined here indicate net false detections, in the sense that the retrieval detected more clouds above a given brightness than were actually present in the simulated data. Similarly, negative values indicate that some fraction of “true” clouds present in the data was missed (i.e., not detected) by the retrieval.

The results highlight two systematic features: 1) the error in retrieved cloud fraction decreases systematically (in absolute terms) as the albedo threshold increases, and 2) there is a transition from negative errors, indicating missed clouds, at low SZA to positive errors, indicative of false detections, at SZA above ~ 80 -85 degrees. This SZA dependence occurs for all cloud brightness levels. Both conclusions are easily understood from the results presented previously. The predicted cloud detection sensitivity curves in Figure 18 show that brighter clouds are easier to detect at all solar zenith angles and thus the algorithm misses fewer bright clouds than dim ones. Similarly, we know that false detections are generally characterized by very low retrieved albedo (see Figure 17), so statistically we expect to see false detections enhanced for low albedo thresholds. When considering all possible clouds (0 G threshold) the maximum errors range from -3% (missed clouds) at 55 degrees SZA to almost +40% (enhanced false detections) at 90 degrees. The errors decrease rapidly when the dimmest clouds are eliminated. Using a 2 G threshold, which from previous discussions is the recommended screening threshold for CIPS data, the fraction of missed clouds at low SZA drops to 1-1.5% with less than 10% false detections at 90 degrees. Finally, for 5 G and brighter clouds the cloud fraction error is within ± 1 % for the entire solar zenith angle range spanned by the CIPS data.

The cloud fraction errors can be better understood by looking at the error in the retrieved Rayleigh background albedo shown in Figure 20. This quantity, the Rayleigh albedo normalized to 90-degree scattering angle and nadir viewing, is not released in the level 2 data products but is routinely archived in a separate “ozone” data file. As discussed previously the true Rayleigh background is unknown in these simulations, since it comes from the actual CIPS level 1b measurements. The “error” plotted in Figure 20 is defined as the difference between the Rayleigh albedo retrieved from the simulation (i.e., in the presence of clouds) and the operational CIPS retrieval for each cloud-free orbit. Defined this way, systematic errors in the actual Rayleigh retrieval cancel and we can directly quantify the effect of clouds on the background retrieval. The mean absolute error for the ten orbits, indicated by the red curve, shows that the retrieved background is biased high by as much as ~ 1 G at lower SZA, transitioning to a low bias of ~ 0.5 G at 90 degrees.

The bias in the retrieved cloud albedo is inversely correlated with this error, since the cloud component is obtained by simply subtracting the retrieved background Rayleigh signal from the measured data (step 4 in Figure 14). Viewed in this way the results in Figures 19 and 20 make sense and are self-consistent. At low SZA the cloud/background partitioning is biased slightly towards the background, with a corresponding deficiency in the cloud component, whereas at high SZA the bias is in the opposite direction, with too much cloud and too little background. The more difficult question is how to attribute cause and effect to these competing scenarios. The most plausible scenario is the following. At low SZA the decreased detection sensitivity (see Figure 18) causes a significant fraction of dim clouds to be missed and this “missed” cloud albedo component is attributed directly to the Rayleigh background. At high SZA it is more likely that the Rayleigh background retrieval itself is biased slightly low, probably by small positive errors in the polynomial fitting of the $C(\varphi)$ profile in this region of maximum gradient. The decreased Rayleigh background subtraction then directly enhances the cloud

residual, resulting in false detections. The absolute magnitude of the Rayleigh background errors is consistent with previous results showing that the vast majority of missed detections, as well as the false detections produced by the algorithm, correspond to very dim clouds of 2 G or less.

6.2.2 Cloud parameter retrieval

The dataset of simulated retrievals has been used to estimate both systematic and random retrievals errors on the CIPS level 2 cloud parameters. The process is straightforward; we consider only those pixels that contain a “true” cloud in the simulation and are correctly detected by the algorithm. For each such pixel i we calculate errors in the cloud parameters by simply taking the difference between the retrieved and true values:

$$\delta X_i \equiv X_i^{ret} - X_i^{true} \quad ; \quad X = [A_{PMC}, r_o, IWC] \quad (11)$$

Note that we are dealing with absolute, rather than relative, errors. The analysis presented here only addresses the cloud albedo, particle radius and ice water content. The ice column density, while archived as a V4.20 level 2 data product, has not yet been sufficiently evaluated and therefore is not included. For each orbit of simulated data, the retrieval errors are binned on a fixed three-dimensional grid of solar zenith angle, albedo and particle size. This allows us to evaluate the error dependence on all three of these key parameters. Solar zenith angle averaging is done using two 22.5-degree bins between 40 and 85 degrees and a single 10-degree bin from 85 to 95 degrees. The albedo binning uses 3-G wide bins centered at 2, 5, 10, 25 and 50 G while radius is binned in three 20-nm wide bins centered at 30, 50 and 70 nm. The bins were constructed to be fine enough to resolve the relevant retrieval dependencies yet large enough to provide good statistics in each 3-D grid box. Figure 21 shows the resulting error estimates for all three parameters, obtained by averaging over all ten orbits of simulated data. Each column corresponds to a different SZA bin and the solid and dashed curves represent the mean error (bias) and standard deviation, respectively, within each bin. For each retrieved quantity, the errors are plotted as a function the parameter itself using the grid values on the abscissa. A family of curves is plotted, corresponding to a range of independent-variable values, which are color-coded using the legend in the first column.

The albedo errors are plotted on the top row for a range of particle sizes. Both the mean error and standard deviation of the retrievals are less than 2 G everywhere, and more commonly in the range of 1-1.5 G. However, even though the errors are small several systematic trends are apparent. While the standard deviation is relatively constant with SZA, the mean error has a systematic dependence on SZA, with negative errors of 1-1.5 G at low SZA transitioning to a positive bias of ~0.3-0.5 G at the terminator. This result implies that on average the retrieved cloud albedo is too low at lower SZA and too high at high SZA. This behavior is consistent with the observed biases in retrieved Rayleigh albedo in Figure 20. The small positive Rayleigh bias at lower SZA depresses the retrieved albedo for those clouds that are detected, and the opposite occurs at high SZA where the cloud albedo is enhanced by the small negative Rayleigh bias. Both the mean error and standard deviation also generally increase in magnitude for larger particles. The

size dependence of the mean error is largest at low SZA and decreases with increasing SZA, becoming negligible at 90 degrees. Finally, we note that both the mean error and standard deviation are essentially constant for 10 G and brighter clouds, and decrease only slightly for dimmer clouds. This indicates a true systematic (i.e., absolute) error, since the relative errors are much larger for dim clouds. For example, consider the results for the 73-degree SZA bin and 50-nm particles (the solid blue curve in the top middle panel). The mean error on a 2 G cloud detection is 25% (-0.5 G), but only 2% (-1.0 G) for a 50 G cloud.

The particle radius errors are plotted on the middle row, where the color-coding now corresponds to a range of albedo values. Both the mean error and standard deviation have a strong dependence on cloud albedo, decreasing rapidly in magnitude for brighter clouds. For 25 G and brighter clouds both the mean errors and standard deviation are less than 2 or 3 nm for all SZA and particle sizes. There is a more pronounced SZA dependence in the mean radius errors than for the albedo. At 90-degrees the mean errors are essentially zero for all cloud brightness levels except the dimmest clouds (2 G, black curve), for which there is a very small negative bias for the largest particles. The mean error becomes positive in sign and increases in magnitude at lower SZA. In the lowest 40-62.5 degree SZA bin errors for the dimmest (2 G) clouds can be as high as 15 to 25 nm, whereas they are less than 15 nm for 5 G clouds and 10 nm for 10 G clouds. We also note that the radius and albedo mean errors are anti-correlated, with negative (positive) albedo biases coinciding with positive (negative) radius biases. This anti-correlation is consistent with the results derived in Carstens et al. (2012, this issue), which also finds particle size errors that are generally consistent with the results presented here. Generally the errors decrease with increasing particle size (except for the dimmest 2 G clouds, which show essentially flat behavior with radius). The standard deviation of the particle size retrievals is approximately constant with SZA, but like the mean error shows a strong inverse dependence on cloud brightness. For all but the dimmest clouds the random error is less than ~5 nm everywhere, but increases to 10-20 nm for the dimmest clouds (≤ 5 G) in the lowest SZA bin.

Finally, the IWC errors are plotted on the bottom row. As discussed in Section 4.6, IWC is completely determined by the cloud albedo and particle radius. From Equation (10) it is directly proportional to A_{PMC} with the proportionality factor $\gamma \equiv V(r_o) / \sigma(r_o)$. The function γ depends only on particle size and is determined by the underlying size distribution assumed in the T-matrix calculations. It is relatively flat for radius values above 30 nm, but increases rapidly for smaller particles. Thus, while IWC scales linearly with albedo it is strongly dependent on particle size for small particles, and this dependence is captured by color-coded curves in Figure 21. The IWC grid used on the abscissa was calculated from the albedo and radius binning grids using Equation (10), which explains why the grids differ for different radius values. The IWC mean error and standard deviation both increase with decreasing particle size, particularly for the smallest particle (30-nm) bin. They also show a strong dependence on IWC (or alternatively, albedo) for these small particles. For larger particles the errors show little dependence on particle size and a smaller correlation with IWC/albedo. The SZA dependence of the mean error is very similar to albedo. At low SZA there is a negative

bias for all particle sizes and IWC values, which decreases with increasing SZA. At 90 degrees the errors are very small but can change sign. They are slightly positive for the dimmest clouds and smallest particles, but still tend towards negative values as IWC increases, particularly for the smallest r_o values. The standard deviation shows little or no dependence on SZA.

6.3 Phase function sensitivity

One potential source of systematic error in the CIPS retrievals is the choice of the theoretical ice phase function P_{ice} used in the algorithm. This directly affects the cloud albedo and particle size retrieved in the phase function fitting routine (see Section 4.5) and through the iterative process of the level 2 algorithm, can potentially feed back into the Rayleigh background subtraction and cloud detection. While the underlying assumptions made about the ice particle shape and size distribution have a strong geophysical motivation based on consistency with the lidar and SOFIE measurements, other choices are possible. In fact, Baumgarten et al. (2010) discuss several variations of the derived Gaussian width vs. particle radius curve, $s(r_o)$, obtained from alternative analyses of the lidar data. Additionally, there is some ambiguity concerning the true particle shape (axial ratio, AR), which also impacts the retrievals.

A sensitivity analysis was performed to estimate potential errors in CIPS cloud parameter retrievals due to these uncertainties. A sample of CIPS high cloud-frequency orbits were processed using a set of alternative theoretical phase functions and the resulting retrievals were compared statistically with the standard V4.20 products. A total of nine different alternative phase functions were generated with the T-matrix code. All of these alternative cases retain the assumption of a Gaussian particle size distribution. Four of them, denoted Cases 2-5, differ from the operational V4.20 P_{ice} (Case 1) by assuming slightly different shapes for the $s(r_o)$ relationship. The $s(r_o)$ curves for these cases are shown in the top panel of Figure 22 along with the original and smoothed versions of Case 1 (red curves) for reference. Cases 2-5 are discussed in Baumgarten et al. (2010) and are derived using different assumptions in fitting the lidar data. Finally, five additional phase functions are included in the study that retain the same $s(r_o)$ relationship as Case 1 but make different assumptions on the particle shape. Recall that the operational P_{ice} assumes the particles are oblate spheroids with AR=2. The alternatives explored here include AR=1 (spheres), AR=0.5 (needle-like), AR=0.25 (extreme needle-like), AR=5 (extreme plate-like), and AR=2 for cylinders rather than oblate spheroids.

Six CIPS orbits were used for this study to provide a sample of data from different years and both hemispheres. Each orbit was processed with the operational level 2 algorithm using all 10 model phase functions as input. As a technical note, Case 1 was re-run for this study using the piecewise linear fit as shown in Figure 22 (solid red curve), for consistency with the other cases rather than using the operational V4.20 results, which use the smoothed version of this function. Thus any systematics introduced by the discontinuity in the derivative of $s(r_o)$ will cancel out of the comparisons. For each orbit and each test retrieval case the common cloud pixels are identified and used to generate errors relative to the reference Case 1:

$$\delta X_N^i = X_N^i - X_1^i \quad ; \quad X = [A_{PMC}, r_o, IWC] \quad (12)$$

where $N (= 2-10)$ is the phase function case and the superscript i denotes a single pixel in the orbit. The errors for each orbit were binned on a uniform grid for each retrieval parameter and the results from all six orbits were then averaged together.

The results of this analysis are shown in Figure 23, where the curves represent the mean difference over all orbits. Panels (a)-(c) summarize the sensitivity to the assumed $s(r_o)$ relationship (Cases 2-5). The largest differences in all three parameters occur for Cases 3 and 5, which is expected since these $s(r_o)$ curves have the largest deviation from the reference Case 1. The albedo and particle size mean retrieval differences are shown in panels (a) and (b), respectively. These results show the same systematic dependence on the shape of the $s(r_o)$ function: as the width of the Gaussian distribution is increased the retrieved albedo and particle size both decrease. The albedo dependence on P_{ice} is a small, second-order effect and comes directly from Equation (6). Mean albedo differences are less than 1 G for all clouds, which corresponds to relative errors of less than 2-3%. The radius retrievals show a greater sensitivity, which can be understood as a direct consequence of the shape of the ice phase functions, P_{ice} . These are shown in the bottom panel of Figure 23 for the case of 50-nm particles (results are qualitatively similar for all particle sizes). Consider Case 5, given by the blue curve. This phase function is more strongly peaked in the forward-scattering direction than the reference case. Therefore, using this P_{ice} a measured phase function will be better fit by a smaller radius curve, which is consistent with the results in panel (b). Case 3 shows the opposite behavior, as P_{ice} is less strongly forward-peaked than the reference case resulting in a larger retrieved particle size. For all cases other than Case 2 the sign of the radius difference is the same for all particle sizes. Case 2 is unique in that the distribution width is larger than Case 1 for radii less than ~ 55 nm and smaller for larger particles. This is the reason that the retrieved radius error changes sign depending on particle size. The change in shape of the differences for all cases at the 40-nm point is due to the piecewise dependence of the $s(r_o)$ curves and the discontinuity in slope at this point. The magnitude of the particle size errors are within 5 nm ($\sim 10\%$) for all cases and particle sizes, with the exception of Case 5, where the errors reach a maximum of 7 nm for the largest particles.

The mean IWC errors, plotted in panel (c), are qualitatively similar to the albedo errors but are opposite in sign to both the albedo and radius error. As discussed previously, while IWC is directly proportional to albedo it is primarily driven by the radius dependence through the term $\gamma \equiv V(r_o) / \sigma(r_o)$. This factor varies inversely, and non-linearly, with radius over most of the range of PMC particle sizes (i.e., γ increases for decreasing r_o). This dependence explains why the IWC mean errors are opposite in sign to both albedo and radius. Overall the IWC sensitivity to P_{ice} is small, with mean differences less than 6 g km^{-2} ($\sim 3\%$) everywhere.

The retrieval sensitivity to assumptions about particle shape (Cases 6-10) is shown in the same format in panels (d)-(f) of Figure 23. For the most part, the differences are similar in magnitude for all cases except the $AR = 0.25$ and 5.0 cases. These are expected to produce the largest differences because they represent the most extreme changes to the

assumed particle shape (extreme needle-like and plate-like particles, respectively). If these particle shapes were indeed sampled by the CIPS observations then mean albedo errors could be as much as 1G (15%) for the dimmest clouds for the AR=5 case. Likewise the needle-like particle shape assumption (AR=0.25) produces the largest differences in retrieved radius, with a maximum of ~10 nm (20%). The radius differences for all other particle shape assumptions are within 2-3 nm (5-10%). The IWC retrievals are the most sensitive to these extreme particle shapes, as they have a much larger impact on the radius-dependent factor γ discussed above. For all but the AR=5 case the errors are still within 15 g km⁻², or 10-15%. The AR=5 case results in mean errors of up to 80 g km⁻² (40%), which are off the scale in this plot.

The conclusions we draw from this study are that error bounds on the retrieved cloud parameters due to reasonable uncertainties in the underlying ice phase function are within 1 G (~2%) for albedo, 5 nm (~10%) for radius and 10 g km⁻² (5-10%) for IWC. It was also mentioned previously that changes in the assumed P_{ice} could potentially alter the cloud detections through the iterative retrieval process. Analysis of the cloud fractions derived from these test cases confirmed that this is a very minor effect, affecting a small fraction of the dimmest cloud pixels near the detection threshold. Total retrieved cloud fractions for all test cases for a given orbit were within a few tenths of a percent of the reference case.

7. CIPS data products

The CIPS publically released data products include the primary single-orbit level 2 cloud retrievals and multiple value-added level 3 products (levels 1a and 1b are not released). The level 3 products are all generated by combining level 2 data from multiple orbits, either for a single day or an entire season. In this section we briefly summarize the content of these products and show some examples. All CIPS data is available online through two separate sources: the University of Colorado/Laboratory for Atmospheric and Space Physics (LASP) CIPS web site at <http://lasp.colorado.edu/aim/> and the primary Hampton University AIM web site at <http://aim.hamptonu.edu/>. The LASP site contains extensive documentation on all data products as well as software for reading the files.

7.1 Level 2 data

Level 2 data is the primary CIPS data product. It consists of cloud parameter retrievals, derived cloud phase functions and associated geolocation and meta-data at 25 km² spatial resolution for each orbit. The data are provided in three separate NetCDF files to make it easier for the user to select only those quantities of interest. The “CAT”, or catalog, file contains geolocation and meta-data including latitude, longitude, UT time, date, solar zenith angle, and data processing information. The “CLD”, or cloud, file contains the retrieved cloud parameters – A_{PMC} , Γ_0 , IWC, ICD – and associated error bars (currently not populated in V4.20). Also contained in this file are key data diagnostic quantities such as the NLAYERS parameter (# of data points per pixel), quality flag (QF) index, cloud presence map (cloud/no-cloud indicator), and a χ^2 metric from the phase function fit (cloud pixels only). Finally a “PSF”, or phase function, file contains the cloud phase

function for each pixel and associated errors (currently not populated in V4.20), as well as view angle and scattering angle arrays associated with each data point in a pixel. Also provided with each orbit of the level 2 data are image (png format) files of retrieved albedo, particle radius and IWC. The image products are purely visual but very useful for identifying PMC extent, morphology and dynamical structures (Thurairajah et al. 2012).

Figure 24 shows a sample of level 2 data from a high cloud frequency NH orbit, number 22789 on July 1 2011. Included at the bottom is the albedo image, giving an idea of the extent of the cloud cover and the complexity of the spatial structure on this orbit. The left(right) side of the image corresponds to high(low) solar zenith angles. Panels (a)-(c) show the retrieved values of albedo, radius, and IWC as a function of SZA for all cloud pixels in the orbit. There are 199,258 detected cloud pixels in this orbit out of a total of 349,464 good level 2 pixels, yielding a total-orbit cloud fraction of 57%. This is at the upper range of what CIPS typically observes. We see that clouds are present at some level over almost the entire 40-95 degree SZA retrieval range and there are many very bright clouds (> 20 G). The pixels highlighted in red have a retrieved particle radius less than 20 nm (see panel (b)). These pixels are typically flagged and removed from higher-level 3X data products. Forward-model studies suggest that CIPS should not be sensitive to particles this small and hence these retrievals, while they might represent real clouds, are generally assumed to be an indicator of an incorrect background subtraction. Note that they are predominately associated with dim clouds and high IWC values (panels (a) and (c)).

Panels (d)-(f) result from some simple analysis on the level 2 data and are included to point out some larger-scale features in the cloud field. Plotted in panel (d) is the cloud fraction obtained by averaging in 2.5-deg SZA bins. Four different albedo thresholds are used ranging from 1 G, which essentially includes all clouds, to 10 G, which includes only the brightest clouds. Obviously the cloud fraction decreases as the albedo threshold is increased but for all, but the brightest clouds the fraction equals or approaches 100% over a large portion of the orbit. Panel (e) shows the median albedo (red), radius (green) and IWC (blue) on the same 2.5-deg SZA grid, with different scale factors applied to each parameter in order to plot them together. Note how the albedo and IWC are in general highly correlated, except for pixels with abnormally low retrieved radius as discussed above. Finally, plotted in panel (f) are the frequency distributions for the three cloud parameters over the orbit. The albedo, radius and IWC retrieval distributions peak at 11 G, 34 nm and 77 g km^{-2} , respectively. This is typical of the particle size distributions CIPS generally sees, but the albedo and IWC distributions are shifted higher than the norm because of the large number of bright clouds in the orbit.

7.2 Level 3a data

Level 3a data contains cloud albedo for a single calendar day. It is produced by combining level 2 data from all individual orbits on that day, using the common 25 km^2 Lambert Projection grid. When pixels from different orbits overlap, the brightest pixel (not the average) is used, with the objective of providing maximum enhancement of the cloud features observed during the day. One NetCDF and one image (png) file are

produced per day. Each day's image uses a color scale dynamically scaled to the observed albedo range on that day. Figure 25 is an example of the 3a image (called a "daily daisy") for June 29 2012. Note that discontinuities at orbit boundaries tend to be enhanced by the logic of choosing the maximum albedo in overlapping pixels. This emphasizes the rapid time scale of PMC formation and evolution. These images are probably the most viewed feature on the LASP/CIPS web site as they provide users with quick and direct feedback on the cloud activity on a daily basis.

7.3 Level 3b data

Level 3b data consists of movies of daily cloud albedo maps for an entire PMC season. One mpeg4 file is created per season, by combining all level 3a files for the season and using a fixed color scale dynamically scaled for the season.

7.4 Level 3c data

The level 3c product consists of season-long files of latitude-binned level 2 data. Retrieved cloud albedo, particle radius, and ice water content from each orbit are binned in one-degree latitude bins and the mean and standard deviations are output for an entire PMC season. The ascending and descending node data are binned separately in order to keep local time effects separate. Also output are the mean solar zenith angle, UT time and local time in each bin, as well as the total number of pixels and the number of cloud pixels. The latter two quantities can be used to calculate latitude-binned cloud frequencies over the season.

The data are provided in both ascii text and IDL binary save file formats. Nine files are generated per season, corresponding to three cloud brightness thresholds (1, 2 and 5 G) and screening by cloud-only, non-cloud or all pixels. Figure 26 shows sample results from the NH 2011 season, generated from the 2 G albedo threshold files. Each panel shows a time vs. latitude contour for the entire season with cloud frequency, albedo and radius in the left, middle and right columns, respectively. Ascending node data are plotted in the top row and descending in the bottom. This data product is very useful for identifying both temporal and latitudinal cloud morphology during a single season, or for studying inter-annual variability.

7.5 Level 3d data

CIPS level 3d data files contain a season-long summary of level 2 data from the AIM common volume (CV). From the CIPS perspective, the CV is defined as the level 2 pixels that are co-located with the SOFIE line-of-sight (LOS) on each orbit. As described in Section 2, the AIM satellite was intentionally rolled each orbit to ensure sampling in the CV region. This roll maneuver was discontinued in February 2010 and therefore the 3d data files are only available for the first three seasons in each hemisphere. This dataset is a valuable resource for co-validation of the CIPS and SOFIE retrievals, and for studying the effects of horizontal inhomogeneity on the SOFIE limb retrievals (see Bailey et al. 2012, this issue).

The level 3d data are pulled directly from the level 2 data files, where the subset of CV pixels is already defined. No additional processing is done on the data. The level 3d files contain the CIPS level 2 albedo, radius and IWC retrievals (no phase functions) and associated auxiliary data, in the CV, over an entire season in ascii text format. By construction, the CV extends ~400 km on either side of the 83-km occultation tangent point. In the SOFIE limb-viewing geometry this corresponds to a ray path length through a vertical range approximately seven times the SOFIE vertical resolution element of 1.6 km, or ~11 km total. The CV width of ~5 km is defined by the SOFIE field-of-view perpendicular to the LOS.

Figure 27 shows an example of the CV data for a single orbit on July 15 2007. The top left panel shows a map of the CV region, which spans approximately 6 degrees in longitude by 7 degrees in latitude on this particular orbit. The dense set of small blue crosses denote all the CIPS level 2 pixels in the region while the larger circles are the CV pixels, color-coded brown or black for cloud and non-cloud pixels, respectively. The single red circle indicates the location of the SOFIE 83-km tangent point and the shaded region identifies the day-night terminator boundary. The number of CV pixels is 305 (a typical number) and 249 of those pixels contained cloud as detected by CIPS. The cloud fraction in the CV is therefore 81.6%, which we note is higher than average. The CV region as seen in the CIPS data is highly variable, and mean cloud fractions over a season are generally in the 40-50% range.

The remaining three panels of Figure 27 show the CIPS albedo, radius and IWC as a function of SZA for the CV cloud pixels. For reference the LOS distance from the SOFIE occultation point is also shown along the top axis. By convention the dayside (SZA < 90) pixels, which correspond to the SOFIE far field, are assigned negative distance. For each parameter the mean CIPS value in the CV is indicated by the horizontal dashed line and written in black in the top right corner (rounded to the nearest integer). In the radius and IWC panels the corresponding SOFIE retrievals for this orbit are represented by the single red circle at 90-degree SZA, and written in red in the top right corner.

CIPS observes a wide range of PMC brightness and particle sizes in the CV on this orbit (the IWC tracks well with the observed albedo). The cloud field seen by CIPS is mostly confined to the dayside of the terminator, with the brightest clouds and the largest particles occurring at the lowest SZA. There is also significant and coherent spatial structure in the cloud field. For this particular orbit the SOFIE radius retrieval agrees well with the mean CIPS value, and more importantly is almost identical to the CIPS retrievals from pixels closest to the occultation location. SOFIE provides several particle size products; we have used the Gaussian mean radius obtained from the vertically integrated optical depth, which is consistent with the CIPS nadir-viewing geometry. The SOFIE IWC retrieval is significantly higher (~ 60%) than the CIPS mean value in the CV, although it is consistent with the brightest far field clouds seen by CIPS. Overall, a careful analysis of the CV data over the entire AIM mission reveals systematic biases in particle size and IWC retrieved by the two instruments. SOFIE generally reports smaller particle sizes and higher IWC than CIPS. The accompanying paper by Bailey et al. (2012,

this issue) describes this analysis in detail and attempts to reconcile the observed biases to within a level of consistency with the stated errors in both data sets, by considering differences in instrument sensitivity, viewing geometry, and retrieval methods.

7.6 Level 3e data

The level 3e data product is very similar in design and construction to level 3d, but instead of containing the CV pixel subset these files contain subsets of CIPS data that are coincident with specific ground stations. The motivation for this product is to provide the ground-based observing community with easy access to CIPS data in a format that allows for rapid identification and analysis of coincident measurements within the ground station field of view. Currently this product is being produced for 22 different stations (18 NH and 4 SH locations), with new ones added by request. The coincidence criteria are based purely on spatial proximity, using an observation field of view that is defined for each station. There is no attempt to identify CIPS measurements that are temporally coincident with known ground-based measurements – we provide all CIPS data that are obtained within the defined station coincidence range.

Two distinct types of coincidence criteria are used in constructing the 3e product. The most common coincidence definition (Type 0) uses a simple fixed geographical distance from the station location. The maximum distance can in principle vary from station to station but currently all are set to 100 km. The second coincidence definition (Type 1) uses a fixed reference coincidence region, defined by set latitude/longitude ranges. These boundaries may or may not contain the actual ground station location, but are constructed to encompass the known ground-based observation space. This type of criterion is used for ground observations that are not looking directly in the zenith, the most common example being ground-based photography of noctilucent clouds where the camera is typically pointed towards the horizon.

The structure, format and content of the level 3e files is very similar to the 3d files. A single file is produced for each station for each PMC season in ascii text format. The data are pulled directly from the level 2 files with no additional analysis or processing. One notable difference between the two products is that the level 3d files contain CV pixels for every single AIM orbit, whereas only a subset of orbits will be coincident with any given ground station. The 3e files only contain entries for AIM orbits where actual coincident pixels exist. For every orbit the CIPS pixel locations are checked against each ground station; if any pixels fall within the coincidence region the corresponding data are written to the file, otherwise that orbit is skipped. Stations at higher latitudes tend to have a much higher density of CIPS overlaps than lower-latitude stations.

There are essentially two types of information included for each orbit. A single summary line contains definitive information about the orbit, including AIM orbit number, date, time, number of coincident pixels, cloud fraction in the coincidence space, and so forth. Also provided are mean cloud parameters (albedo, radius, ice water content and cloud fraction) averaged over all pixels within a 500-km radius of the station. This is meant to give the user additional information about the morphology of the cloud field in the larger

general vicinity of the station (“loose” coincidence). This is followed by detailed information on each coincident pixel including geolocation, cloud/no-cloud discrimination and the level 2 cloud parameter retrievals for cloud pixels.

Figure 28 shows a sample image illustrating this product for a single orbit coincident with the Alomar lidar station at ~69 N latitude on July 8 2010. The station location is represented by the solid black dot. The green area represents the 100-km radius coincidence criteria used for this station and the blue area is the larger 500-km “loose” coincidence region. The small grey symbols represent all the CIPS level 2 pixels within that portion of the orbit swath falling inside the plotted region. CIPS cloud detections are indicated by red symbols. We see that CIPS observed significant cloud features in the Alomar region on this day. For this orbit there are 1263 coincident CIPS pixels (green symbols) and the CIPS cloud fraction is 40.2% in the coincidence region and 36.4% in the larger “loose” region. Baumgarten et al. (2012) have analyzed CIPS and Alomar lidar data on a different day in 2008 when the measurements were closely coincident both spatially and temporally. They found very good agreement in the cloud brightness observed by the two instruments, as well as a high degree of correlation between the temporal variations observed by the lidar and the horizontal structure observed by CIPS when compared using an advection model derived from nearby radar measurements of horizontal winds.

8. Summary and conclusions

We have presented a thorough description of the CIPS version 4.20 retrieval algorithms and data products. Cloud microphysical parameters are retrieved from the CIPS data via a three-step process. The level 0 data stream is first time-tagged, geolocated and calibrated to produce level 1a images. This process is implemented separately for each of the four CIPS cameras. Some critical level 1a calibration steps require special on-orbit image sequences, which have been modified over the AIM mission in response to the loss of satellite bitlock and subsequent changes in orbit data sequences. In particular, the camera flat-field corrections and cross-camera normalization factors have been generated each seasonal using a variety of different low-latitude, subsolar or near-subsolar calibration images. This process was described in detail and results were shown demonstrating that the available data sets provide a consistent long-term calibration of the CIPS data over the entire mission.

The second step is level 1b processing, which merges data from all four cameras onto a common 25-km² resolution Lambert Projection grid to generate the unique CIPS data stacks. This data level consists of “scattering profiles” of total albedo vs. scattering angle for each level 1b pixel, which is the essence of the phase function approach used by CIPS. Level 1b provides the input data stream for the level 2 processing, which is the primary CIPS cloud retrieval process. The level 2 algorithm utilizes an iterative approach to the fundamental steps of Rayleigh background removal, cloud detection, and fitting of

the residual PMC scattering phase function to retrieve cloud parameters. The cloud parameters derived include albedo (normalized to 90-degree scattering angle), the Gaussian mean particle radius, ice water content, and ice column density. The V4.20 level 2 retrievals are performed operationally over a SZA range from 40 to 95 degrees. Near the peak of the PMC season, at solstice, this yields retrievals down to approximately 50-55 degrees latitude on the descending node of the orbit. Somewhat lower latitudes are sampled earlier and later in the season.

A quantitative analysis of the CIPS cloud detection sensitivity and retrieval errors has been performed based on a dataset of simulated retrievals. We have shown how fundamental characteristics of the CIPS measurement sampling result in a solar zenith angle dependence in both the cloud detection sensitivity and cloud parameter retrieval errors. CIPS loses its optimal forward-scattering measurements, which provide the greatest cloud detection sensitivity and particle size discrimination, at low SZA where the sunlit Rayleigh atmosphere is brightest. This combination of factors makes it more difficult for CIPS to detect a cloud of a given brightness at low SZA than at high SZA.

Results presented here indicate that CIPS should detect dim, 2 G clouds approximately 40% of the time at 50 degrees SZA, increasing to 60% at 70 degrees and essentially 100% at the terminator. For brighter 4 G clouds the corresponding detection confidence rates are 85, 95, and 100%, respectively, and any cloud brighter than 4 G is detected with greater than 90% confidence at any SZA sampled by the data. The simulations also show that CIPS retrieves the total cloud fraction (all clouds considered) to within 2-3 percent for SZA below 75 degrees. The small fraction of clouds that are missed is due to the above-mentioned inability to detect dim clouds at lower SZA. If only clouds brighter than 2 G are considered the retrieved cloud fraction is within ~1% of the “truth”. At high SZA, on the other hand, some false cloud detections occur leading to cloud fractions biased high by 5-10% (again, screening for clouds > 2 G).

Both the mean error (bias) and standard deviation (random error) of the CIPS cloud parameter retrievals have been estimated by comparing retrieved quantities to “truth” values used in the simulated data. For all cloud parameters, the mean retrieval errors are largest at low SZA and decrease with increasing SZA, becoming quite small at the terminator. Random errors, on the other hand, are roughly independent of SZA. The albedo and IWC errors are generally small in absolute terms. Maximum estimated albedo errors are less than 2 G at all SZA, while the IWC errors are less than 10 g km⁻² for particle sizes ≥ 40 nm, but increase sharply (to 20–40 g km⁻²) for smaller particles at low SZA. Both quantities are biased low at low SZA and have essentially zero or slightly positive bias at the terminator. Albedo errors increase with increasing particle size while IWC errors are largest for small particles. The particle size retrieval errors have the

strongest SZA-dependence, with a high bias at low SZA and essentially no bias, or a slight negative bias, at the terminator. These biases are therefore anti-correlated with both albedo and IWC. Both the random and mean errors in the radius retrieval decrease with increasing cloud brightness.

The estimated radius errors at high SZA are relevant to the CIPS/SOFIE validation study in the AIM common volume region. The findings presented here suggest that any systematic bias in the CIPS radius retrievals are small relative to the particle size differences observed in direct CIPS/SOFIE comparisons (Bailey et al. 2012, this issue). If anything the dimmer clouds (5 G or less), which are typical of CIPS observations in the common volume, are predicted to have a slight low bias in particle size. This does not explain the observed CIPS/SOFIE differences, where CIPS particle size values are systematically higher than SOFIE.

The observed SZA dependence in the CIPS mean retrieval errors must be considered when interpreting ascending/descending node (i.e., local time) differences in the data. For example, analysis of the CIPS level 3c data routinely shows small but systematic ascending/descending node differences in particle size, with descending node values generally larger than the ascending node at a given latitude. The error analysis results presented here show the same trend, indicating that at least part of this dependence could be explained by systematic retrieval error rather than local time dependence in cloud properties. Estimated errors in the other CIPS cloud parameters, however, have a much smaller SZA-dependence and therefore are not as affected as the retrieved particle size. Understanding and eliminating the source of these dependencies continues to be one of the highest priorities for future algorithm improvements.

The results presented here generally support the recommendation made in the CIPS data documentation that users employ a 2-G albedo screen for conservative applications. The estimated retrieval errors all increase significantly for clouds dimmer than this threshold, while the predicted and observed CIPS false cloud detections are largely eliminated if these dimmest clouds are excluded. The observed false detection rate in CIPS data varies from 1 to 2 % per orbit, with somewhat higher rates in the SH compared to the NH and no significant long-term trends.

Work is currently underway on a number of algorithm improvements, with the goal of producing a version 5 dataset within the next year. This work will focus on two distinct areas: improvements to the level 1a calibration and to the level 2 cloud retrieval algorithms. A primary goal is to extend the CIPS cloud retrievals down to lower SZA (hence lower latitudes), to detect dimmer clouds with greater confidence, and to remove some known retrieval artifacts that exist in the V4.20 dataset. There are known

systematic cross-camera biases in the CIPS level 1a data at the $\sim 1\%$ signal level that are not well understood or effectively removed by the current calibration analysis. While these residual errors are small, as discussed in Section 3.1 they can affect our ability to detect the dimmest clouds, particularly at low SZA. A related issue is to understand and correct the observed small hemispheric discontinuities in the Δ -flat field corrections and camera normalization factors. The result will be a more consistent NH/SH level 1a calibration and increased confidence that observed inter-hemispheric differences in the CIPS data are purely geophysical and thus scientifically relevant.

We continue to explore new ways to improve the Rayleigh background characterization and removal in the level 2 algorithm. The current method of retrieving one effective $[C/\sigma]$ pair in a narrow SZA bin can still leave significant cross-track ozone gradients within a bin, which produces artifacts near the orbit edges. Future versions will explore using a denser sampling for the background retrieval binning, perhaps using a 2-dimensional background grid to better capture cross-track ozone gradients. We will also improve the fitting and smoothing routine used to constrain the final $[C/\sigma]$ values as a function of SZA over the orbit. The emphasis will be on getting a better fit at low SZA (below the current 40-degree limit) and in the high SZA range where the background gradient is steepest. This should provide improved retrievals at lower latitudes and eliminate the known false-detection artifacts now in the CIPS data at ~ 40 and 85 degrees SZA.

Finally, the version 5 algorithms will attempt to incorporate some of the techniques and results derived from the analysis of Carstens *et al.* (2012, this issue). This analysis reproduces many of the systematic features observed in the error analysis study presented here. It describes in a fundamental way the limitations on retrieving particle size due to the overlapping cloud and background components of the measured albedo, and helps to quantify the dependence of retrieval errors on the measurement sampling and noise characteristics of the CIPS data. It also derives key diagnostic quantities, such as the leverage to noise ratio (LNR), which can be used to improve the retrievals and to provide additional quantitative information in the definition of the level 2 Quality Flag metrics.

Appendix A. Derivation of Rayleigh parameters from non-cloud data

The C/σ model described in Section 4.1 (see Equation (2)) provides a simple analytic parameterization that can be used to characterize and remove the background Rayleigh albedo in the CIPS data processing. This model is easily cast into a form that makes it straightforward to retrieve effective C/σ parameters from fitting CIPS non-cloud data. This practical problem is encountered in two critical areas of the CIPS processing: the separation of background and cloud albedo components in the level 2 retrieval algorithms (see section 4.3) and generation of error lookup tables used to define cloud detection thresholds (see Appendix B).

By re-arranging terms in Equation (2) and taking the logarithm we obtain:

$$\ln\left(\frac{\mu A}{P_{Ray}(\Phi)}\right) = -\sigma \ln(\mu^{-1} + \text{ch}(\varphi)) + \ln\left(\frac{\Gamma(\sigma + 1)\beta_{Ray} N_{air}}{(\beta_{O_3} C)^\sigma}\right) \quad (\text{A.1})$$

This may be cast as a simple linear equation $Y = ax + b$ using the following definitions:

$$Y \equiv \ln\left(\frac{\mu A}{P_{Ray}(\Phi)}\right) \quad ; \quad X \equiv \ln(\mu^{-1} + \text{ch}(\varphi)) \quad (\text{A.2})$$

$$a = -\sigma \quad (\text{A.3})$$

$$b = \ln\left(\frac{\Gamma(\sigma + 1)\beta_{Ray} N_{air}}{(\beta_{O_3} C)^\sigma}\right) \quad (\text{A.4})$$

For any set of CIPS Level 1b data points the quantities X and Y are known – they are determined purely from the measurement geometry and the measured albedo. A simple linear regression to these X/Y points determines the slope and offset values [a,b]. The scale height ratio σ is obtained directly from the slope via equation (A3.) after which equation (A.4) may be solved for C. This approach can be used to derive effective background [C/σ] parameters from any subset of CIPS data, ranging from a single-pixel scattering profile to the ~5000 data points typically contained in the 0.25-deg solar zenith angle bins used in the level 2 background removal.

We note that in this final step the atmospheric column density N_{air} must be assumed. In the operational CIPS algorithms this constant is generated as follows. A radiative transfer code is run offline using climatological high-latitude summer atmospheric and ozone density profiles to generate a lookup table of peak Rayleigh scattering altitude as a function of solar zenith angle, $z_o(\varphi)$. The same model generates the vertical column density profile $N_{air}(z)$. To evaluate Equation (2) or (A.4) for a given measurement solar zenith angle φ this profile is interpolated to $N_{air}[z_o(\varphi)]$. Since this approach uses a static

reference atmosphere, an overall systematic error is introduced that is proportional to $N_{\text{air}}(z_0)/N_{\text{true}}(z_0)$, where N_{true} is the true atmospheric column density at the time and location of the CIPS measurement. In practice this simply means that there is an overall unknown bias in the retrieved O_3 column density factor C . This error is irrelevant to the primary CIPS science since the retrieved $[C/\sigma]$ parameters are not released publicly and are not strictly interpreted as valid geophysical parameters. Rather, they are viewed as convenient fitting parameters and an overall scaling error in C has no effect on the ability of the C/σ model to fit the CIPS cloud-free data, and hence to remove the background signal from the measurements. Thus it has no effect on the derived residual cloud phase function, or the cloud parameters retrieved from it.

Appendix B. Construction of Residual Error Look up Tables.

In order to define meaningful cloud detection thresholds it is critical to have a quantitative characterization of the errors in the CIPS Rayleigh background subtraction. A straightforward way to do this is to analyze the residual error obtained when applying the level 2 Rayleigh albedo retrieval to cloud-free data. Because these data are a pure measurement of the Rayleigh albedo this residual directly quantifies the Rayleigh subtraction error in the operational algorithm. Using a large ensemble of data, we compile statistics on both the mean error and standard deviation of these retrievals. The CIPS V4.20 algorithm implements a look-up table (LUT) approach in which these errors are characterized separately as a function of solar zenith angle and view angle.

For each season we analyze two-week time periods (~ 200 orbits) immediately before and after the nominal cloud season (~ 36 -50 days before solstice and 80-94 days after solstice). For each orbit the Rayleigh albedo is retrieved using the exact analysis described in Section 4.3, consisting of steps 2 through 4 of the algorithm flowchart in Figure 14. That is, the level 1b data are fit in 0.25-degree solar zenith angle bins over the entire orbit, including the final polynomial smoothing in solar zenith angle. The retrieved Rayleigh albedo $A_{\text{Ray}}^{\text{ret}}$ calculated from Equation (2) is then subtracted from the measurements to generate the residual error signal

$$A_{\text{resid}}(\theta, \varphi) \equiv A_{\text{meas}}(\theta, \varphi) - A_{\text{Ray}}^{\text{ret}}(\theta, \varphi) \quad (\text{B.1})$$

The dependence on solar zenith angle and view angle has been made explicit in Equation (B.1). There is no separate dependence on scattering angle, as this angle is determined by the solar zenith angle and view angle at each location. Note that this residual error is identical to the quantity called $A_{\text{cloud}}^{\text{resid}}$ in Section 4.3, however in this case the ‘‘cloud’’ subscript is dropped because the data is known to be cloud-free. The residual error is converted to a relative error, referenced to the Rayleigh albedo:

$$\Delta(\theta, \varphi) \equiv A_{\text{resid}}(\theta, \varphi) / A_{\text{Ray}}^{\text{ret}}(\theta, \varphi) \quad (\text{B.2})$$

This quantity, representing the total ‘‘error’’ in the retrieved Rayleigh background relative to the data, can arise from a number of sources, only some of which originate with the

A_{Ray}^{ret} term in Equation (B.1). Examples of errors in the retrieval include algorithm errors such as ozone gradients not resolved by the 0.25-deg SZA binning, and limitations in the C/σ forward model itself. However errors in the measurement term, A_{meas} , also contribute to the differences. Random measurement noise contributes to the standard deviation of the difference, whereas systematic measurement errors (e.g., due to calibration errors) will contribute to non-zero mean errors (biases) in this difference even for a perfect retrieval algorithm. For this reason data from each camera are fit separately in order to independently characterize any systematic measurement errors such as camera-to-camera biases.

The level 1b “stack” data retains information on the camera origin for every data point in the entire orbit. Once the relative error Δ has been calculated, the data are separated by camera and the error for each camera is binned on a uniform view angle (θ) and SZA (ϕ) grid, using 1-degree bins over the ranges $\theta=[0,90]$ degrees and $\phi=[40,95]$ degrees. Both the mean and standard deviation of Δ are calculated in each bin and separate LUT are generated for both quantities. As an additional level of complexity, forward-scattered and back-scattered measurements are tracked separately in the analysis. This allows a more nuanced characterization of camera non-uniformities in the error matrix, as different pixel sets on a given camera tend to sample only forward- or back-scattering at a given solar zenith angle (see Figure 6). Thus the final result of the analysis is four LUTs for each camera on the standard $[\theta,\phi]$ grid – mean and standard deviation for both forward- and back-scattered measurements. The data from all ~ 200 orbits are then averaged to obtain the final LUT for each time period. In constructing the final tables the grid is padded out by interpolation to fill in any portions of the $[\theta,\phi]$ space that are not sampled by a given camera in the four weeks of test data used in the analysis.

Figure B.1 shows sample results from the operational error LUT for the NH 2010 season (results are generally very consistent for all seasons). We see that over most of the solar zenith angle range, from 40 to 85 or 90 degrees, the standard deviation is on the order of 1%, and no higher than 2%. Thus over the vast majority of the CIPS measurement space the model fits the background Rayleigh measurements to within a random error of 1-1.5%. Only near the terminator, at solar zenith angles above 90 degrees, do the random errors increase sharply. However, recall from Equation (B.2) that this is a relative error. It is referenced to the background Rayleigh albedo, which is dropping rapidly above 90 degrees SZA and essentially goes to zero by 95 degrees. The absolute fitting error A_{resid} , in albedo units, reaches a noise floor and stays roughly constant at high solar zenith angle, causing the relative error to rapidly increase to $\sim 100\%$. The mean error (lower row), which can be either positive or negative, is generally positive and roughly the same size as the random error. It also shows the same qualitative dependence on solar zenith angle. For the most part, there is neither a significant view angle dependence or a forward- vs. back-scattering asymmetry evident.

In normal CIPS data processing, the pre-season error maps are used operationally throughout the season. After the season ends and the post-season analysis is completed, the two time periods are averaged to produce a final set of LUTs and the season is re-

processed (this coincides with the routine post-season reprocessing done with final level 1a calibration data). Using the average helps to correct for any small temporal trends in sensitivity or other instrument characteristics over the season. While no long-term trends or systematic year-to-year variations have ever been detected in these error maps, the calculation of both pre- and post-season LUT every season, and the averaging and re-processing procedure are done for the sake of caution and consistency.

Acknowledgments

This work was supported by the NASA Small Explorer Program through contract # NAS5-03132. We acknowledge the tremendous effort and ingenuity of the AIM Mission Operations team in keeping the satellite operational and returning almost all the CIPS data to the ground for processing.

References

Bailey, S. M., C.E. Randall, J.D. Lumpe, M.E. Hervig, G.E. Thomas and J.M. Russell, Nadir and limb viewing observations of Polar Mesospheric Clouds from the Aeronomy of Ice in the Mesosphere (AIM) Explorer, 2012 (this issue).

Bailey, M. Scott, Gary E. Thomas, David W. Rusch, Aimee W. Merkel, Chris Jeppesen, Justin N. Carstens, Cora E. Randall, William E. McClintock, and James M. Russell, III, Phase Functions of Polar Mesospheric Cloud Ice as Observed by the CIPS Instrument on the AIM Satellite, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.09.039, 2009.

Bailey, S.M., A. W. Merkel, G. E. Thomas, and J. N. Carstens, Observations of polar mesospheric clouds from the SNOE satellite. *J. Geophys. Res.* 110 (D13), D13203, 2005.

Bates, D.R., Rayleigh scattering by air, *Planet. Space Sci.* 32, 785–790, 1984.

Baumgarten, G., A. Chandran, J. Fiedler, P. Hoffman, N. Kaifler, J. Lumpe, A. Merkel, C. E. Randall, D. Rusch, and G. Thomas, On the horizontal and temporal structure of noctilucent clouds as observed by satellite and lidar at ALOMAR (69N), *Geophys. Res. Lett.*, 39, L01803, doi:10.1029/2011GL049935, 2012.

Baumgarten, G., Fiedler, J., and Rapp, M., On microphysical processes of noctilucent clouds (NLC): observations and modeling of mean and width of the particle size-distribution, *Atmos. Chem. Phys.*, 10, 6661-6668, doi:10.5194/acpd-10-6661-2010, 2010.

Benze, S., C.E. Randall, M.T. DeLand, G.E. Thomas, S.M. Bailey, J.M. Russell, III, and A.W. Merkel, Evaluation of AIM CIPS measurements of polar mesospheric clouds by comparison with SBUV data, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2011.02003, 2011.

Benze, S., C. E. Randall, M. T. DeLand, G. E. Thomas, D. W. Rusch, S. M. Bailey, J. M. Russell, III, W. McClintock, A. W. Merkel, and C. Jeppesen, Comparison of polar mesospheric cloud measurements from the Cloud Imaging and Particle Size Experiment and the Solar Backscatter Ultraviolet instrument in 2007, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.07.014, 2009.

Carstens, J. N., S.M. Bailey, J.D. Lumpe and C.E. Randall, Understanding uncertainties in the retrieval of polar mesospheric clouds from the cloud imaging and particle size instrument, 2012 (this issue).

Chandran, A., D.W. Rusch, G.E. Thomas, S.E. Palo, G. Baumgarten, E.J. Jensen, and A.W. Merkel, Atmospheric gravity wave effects on polar mesospheric clouds: A comparison of numerical simulations from CARMA 2D with AIM observations, *J. Geophys. Res.*, doi:10/1029/2012JD017794, 2012.

Chandran, A., D. W. Rusch, A. W. Merkel, S. E. Palo, G. E. Thomas, M. J. Taylor, S. M. Bailey, and J. M. Russell III, Polar Mesospheric Cloud structures observed from the CIPS experiment on the AIM spacecraft: Atmospheric gravity waves as drivers for longitudinal variability in PMC occurrence, *J. Geophys. Res.*, doi:10.1029/2009JD013185, 2010.

Chandran, A., D. W. Rusch, S. E. Palo, G. E. Thomas, and M. Taylor, Gravity wave observation from the Cloud Imaging and Particle Size (CIPS) Experiment on the AIM Spacecraft, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.09.041, 2009.

Gordley, L.L., M. Hervig, C. Fish, J. Russell III, J. Cook, S. Hanson, A. Shumway, S.M. Bailey, G. Paxton, L. Deaver, T. Marshall, J. Burton, B. Magill, C. Brown, E. Thompson, and J. Kemp, The Solar Occultation For Ice Experiment (SOFIE), *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.07.012, 2009.

Gumbel, J., J. Stegman, D. P. Murtagh, and G. Witt, Scattering phase functions and particle sizes in noctilucent clouds. *Geophys. Res. Lett.* 28 (8), 1415–1418, 2001.

Hervig, M.E., L.L. Gordley, M. Stevens, J.M. Russell, and S. Bailey, Interpretation of SOFIE PMC measurements: Cloud identification and derivation of mass density, particle shape, and particle size, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.07.009, 2009.

Karlsson, B. M., C. E. Randall, T. G. Shepherd, V. L. L. Harvey, J. D. Lumpe, K. Nielsen, S. M. Bailey, M. E. Hervig, and J. M. Russell III, On the seasonal onset of polar mesospheric clouds and the breakdown of the stratospheric polar vortex in the southern hemisphere (2011), *J. Geophys. Res.*, 116, D18107, doi:10.1029/2011JD015989, 2011.

Karlsson, B., C.E. Randall, S. Benze, M. Mills, V.L. Harvey, S.M. Bailey, J.M. Russell III, Intra-seasonal variability of polar mesospheric clouds due to inter-hemispheric coupling, *Geophys. Res. Lett.*, VOL. 36, L20802, doi:10.1029/2009GL040348, 2009.

McClintock, W., D. W. Rusch, G. E. Thomas, A. W. Merkel, M. R. Lankton, V. A. Drake, S. M. Bailey, and J. M. Russell III, The Cloud Imaging and Particle Size experiment on the Aeronomy of Ice in the Mesosphere mission: Instrument concept, design, calibration, and on-orbit performance, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.10.011, 2009.

McPeters, R. D., The behavior of ozone near the stratopause from two years of BUUV observations, *J. Geophys. Res.*, 85, 4545-4550, 1980.

Merkel, A. W., D. W. Rusch, S. E. Palo, J. M. Russell III, and S. M. Bailey, Mesospheric planetary wave activity inferred from AIM-CIPS and TIMED-SABER for the northern summer 2007 PMC season. *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2006.05.01, 2009.

Mischenko, M. I., and L. D. Travis, Capabilities and limitations of a current FORTRAN implementation of the T-matrix method for randomly oriented, rotationally symmetric scatters, *J. Quant. Spectrosc. Radiat. Transfer*, 60, 309-324, 1998.

Poppe, A., D. James, and M. Horanyi, Measurements of terrestrial dust influx variability by the Cosmic Dust Experiment, *Planet. Space Sci.* 59, 319-326, doi:10.1016/j.pss.2010.12.002, 2011.

Rusch, D.W., S. M. Bailey, G. E. Thomas, A. W. Merkel, Seasonal-latitude variations of PMC particle size from SME measurements for the Northern 1983 season and SNOE measurements for the Northern 2000 and Southern 2000/2001 seasons, *J. Atmos. Sol. Terr. Phys.*, doi:10.1016/j.jastp.2007.10.006, 2008.

Russell III, J. M., S. M. Bailey, M. Horanyi, L. L. Gordley, D. W. Rusch, M. E. Hervig, G. E. Thomas, C. E. Randall, D. E. Siskind, M. H. Stevens, M. E. Summers, M. I. Taylor, C. R. Englert, P. J. Espy, W. E. McClintock and A. W. Merkel, Aeronomy of Ice in the Mesosphere (AIM): Overview and early science results, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2008.08.011, 2009.

Stevens, M. H., S. Lossow, J. Fiedler, G. Baumgarten, F.J. Lübken, K. Hallgren, P. Hartogh, C.E. Randall, J. Lumpe, S.M. Bailey, R. Niciejewski, R.R. Meier, J.M.C. Plane, A.J. Kochenash, D.P. Murtagh, C.R. Englert, Bright polar mesospheric clouds formed by main engine exhaust from the space shuttle's final launch, *J. Geophys. Res.*, 117, D19206, doi:10.1029/2012JD017638, 2012.

Stevens, M. H., D. E. Siskind, S. D. Eckermann, L. Coy, J. P. McCormack, C. R. Englert, K. W. Hoppel, K. Nielsen, A. J. Kochenash, M. E. Hervig, C. E. Randall, J. Lumpe, S. M. Bailey, M. Rapp & P. Hoffmann, Tidally induced variations of polar mesospheric cloud altitudes and ice water content using a data assimilation system, *J. Geophys. Res.*, 115, D18209, doi:10.1029/2009D013225, 2010.

Thomas, G.E., and C. P. McKay, On the mean particle size and water content of polar mesospheric clouds, *Planet. Space Sci.*, 33 (10), 1209–1224, 1985.

Thomas, G.E., Solar mesosphere explorer measurements of polar mesospheric clouds (noctilucent clouds). *J. Atmos. Terr. Phys.* 46, 819, 1984.

Thurairajah, B., S. M. Bailey, K. Nielsen, C. E. Randall, J. Lumpe, M. J. Taylor, J. M. Russell III, Morphology of polar mesospheric clouds as seen from space, *J. Atmos. Solar-Terr. Phys.*, doi:10.1016/j.jastp.2012.09.009, 2012.

Warren, S. G., Optical constants of ice from the ultraviolet to the microwave, *Appl. Optics* 23, 11906-11,926, 1984.

Figure Captions

Figure 1. This diagram illustrates the relative alignment of the four CIPS cameras. The two “X” cameras point in the forward and aft direction, while the “Y” cameras have a primarily nadir orientation. The distance scales indicate the geometric footprint of a CIPS scene at the nominal PMC cloud deck altitude of 83 km. The PX camera always points in the sunward direction, which corresponds to the satellite ram (anti-ram) for Northern (Southern) Hemisphere measurements. Note that the color-coding used here for the four cameras is preserved in all figures to follow.

Figure 2. Distribution of offset angles from the satellite Z (yaw) axis for the four CIPS cameras. Pointing offsets range from 0 to 44 degrees for the Y cameras and 18 to 62 degrees for the X cameras.

Figure 3. Polar maps illustrating typical CIPS image sequences at solstice in the NH (left panel) and SH (right panel). Image footprints are drawn to scale and for clarity only every 6th scene is plotted. The legend in the top left hand corner identifies the cameras by color. The arrow indicates the direction of satellite motion and the black meridional line indicates the transition between ascending and descending nodes of the orbit.

Figure 4. CIPS data sampling in latitude, solar zenith angle and local time for sample orbits on solstice in the NH (top panel) and SH (bottom panel). The independent variable is elapsed time since the first image, with the transition between ascending (ASC) and descending (DSC) nodes denoted by the solid vertical line. Each data point corresponds to an average over a single image, and cameras are color-coded according to the legend at the top. Dashed vertical lines denote transition between the PX first- and last-light images and normal full-scene images. The dotted curves in the latitude panel show sampling for a typical end of season orbit (70 days past solstice).

Figure 5. Schematic depiction of CIPS viewing geometry, showing the line-of-sight (LOS) of a single CIPS camera pixel projected to a scattering volume at a reference height z . This picture is simplified in that it shows the sun as coplanar with the satellite, scattering point and earth center (i.e., zero solar azimuth angle), which is not generally the case. The offset angle ψ is the angle between the LOS vector and the satellite nadir. The solar zenith angle ϕ and view angle θ are the angles between the local zenith (at the scattering point) and the sun and satellite directions, respectively. The symbol Φ is used to denote the solar scattering angle.

Figure 6. This figure illustrates the typical range of view angle and scattering angle sampled by each CIPS camera, and their dependence on solar zenith angle. Each panel corresponds to a different solar zenith angle bin, ranging from 40 to 95 degrees. This sampling pattern is the same in both hemispheres.

Figure 7. Top panel shows the Δ -flat field correction for all four cameras for the NH 2010 season. Bottom panel shows the time series of the PY camera Δ -flat field correction

over all seasons in sequential order. Temporal trends and season-to-season changes are similar for all four cameras.

Figure 8. Time series of the camera normalization factors for the PX, MX and PY cameras (relative to MY) through the NH 2012 season. Data points for the different cameras are color-coded according to the legend at the lower left. There is a single data point for each LL scene used in the calibration analysis. Vertical dashed lines denote the season (NH/SH) boundaries and the red circles indicate the median value for each season, used operationally in the V4.20 level 1a calibration algorithms.

Figure 9. Schematic illustration of the basic CIPS measurement technique. By combining sequential images spaced 43 seconds apart, CIPS makes multiple coincident measurements of the same air parcel at different scattering angles. In this case seven measurements contribute to the scattering profile, three each from the PX and MX cameras and one from a nadir camera (PY or MY). The arrows indicate the direction of satellite motion in either hemisphere. In the NH the PX images are executed first, followed by PY/MY and then MX; the sequence is reversed in the SH.

Figure 10. Spatial footprints of seven level 1a images from the PX, MX and PY cameras over northern Europe on orbit 22600 (June 19, 2011). Cameras are color-coded according to the legend at the top. Outlines of consecutive images for the PX and MX cameras alternate between solid and dashed lines for clarity. Time is moving forward as images move from northeast to southwest on the map. The black symbols in the center represent the level 1b grid points with contributions from all seven images shown.

Figure 11. The polar map on top illustrates the spatial distribution of the number of data points per pixel (NLAYERS) in the CIPS 1b data for a random orbit. Blue pixels have the densest sampling, with 6 or more points, whereas green pixels have 4 or 5 points and red pixels have 3 or less. This bottom plot shows the fraction of pixels in a CIPS orbit with a given NLAYERS value. The grey shaded area represents the maximum range and the solid black symbols are the mean over a complete season. This pattern is very repeatable for both hemispheres and all CIPS seasons.

Figure 12. A comparison of the relative scattering phase functions (normalized to 90 degrees scattering angle) for Rayleigh scattering (black) and PMC ice particles (grey). Ice phase functions are shown for particle sizes ranging from 20 to 80 nm, with larger particles producing the higher forward-scattering signals.

Figure 13. Relative contribution of cloud signal to total albedo, in %, for a range of scenarios. Each panel corresponds to a different solar zenith angle bin, with scattering and view angle sampling for each bin consistent with actual CIPS sampling patterns (see Figure 6). The Rayleigh background is derived from CIPS climatology and the cloud component is calculated from Equation (3) for a range of PMC albedos and particle sizes. Solid and dashed curves correspond to two broad cloud types – small particle/dim and large particle/bright, respectively, as defined by the legend.

Figure 14. Flowchart of the CIPS Version 4.20 level 2 algorithm.

Figure 15. Scatterplots of the X and Y parameters from Equation (A.2) for three representative solar zenith angle bins (columns) on two CIPS orbits (rows). The top row shows data from a non-cloud (out of season) orbit. The black and red data points are the back- and forward-scattered measurements, respectively. The $[C, \sigma]$ parameters derived from two fits to this data are printed in each panel; top and bottom numbers correspond to the back-scattered only fit and all-points fit, respectively. Also shown is the calculated Rayleigh albedo for nadir viewing angle and 90 degrees scattering angle. The bottom panel shows data from an in-season orbit with clouds. No fits are done on this data due to corruption by cloud signals.

Figure 16. A random sample of CIPS PMC ice phase functions from orbit 1240 on July 18 2007. The CIPS data, corresponding to the residual of Equation (5) are represented by the black symbols. The best-fit theoretical phase function is plotted in red, and the cloud albedo and particle radius values derived from the fit are shown in each panel, along with the SZA. The vertical and horizontal lines are the reference 90-degree scattering angle and the albedo at that angle (defined to be the cloud albedo). These pixels were sampled at random but are sorted on the page in order of decreasing albedo (top to bottom) and increasing SZA (left to right).

Figure 17. CIPS false cloud detection statistics for all NH (blue) and SH (red) seasons. Each point represents the mean over a two-week time period outside the PMC season. Solid and dashed curves correspond to pre- and post-season time periods, respectively. Orbit-averaged cloud fraction is plotted on the top row and is separated by the CIPS quality flag (QF) value, with QF = [0,1] data on the left and QF=2 on the right. The bottom row shows the corresponding orbit-averaged median albedo retrieved from these false detections.

Figure 18. CIPS cloud detection sensitivity derived from simulated retrievals. The five curves correspond to different cloud albedo values, as denoted by the color legend (in standard albedo units, $G = 10^{-6} \text{ sr}^{-1}$). The cloud detection rate (in %) is defined as the fraction of clouds of that brightness actually detected at a given solar zenith angle.

Figure 19. Error in the retrieved CIPS cloud fraction, defined as the difference between the retrieved and true cloud fraction from the retrieval simulations. Cloud fraction is determined as a function of solar zenith angle for all pixels with quality flag (QF) values of 0 and 1. The five curves represent different cloud albedo thresholds from 0 (all clouds) to 10 G, and are color-coded according to the legend.

Figure 20. Error in the retrieved nadir Rayleigh albedo, defined as the albedo observed at 0-degree view angle, as a function of solar zenith angle. The error is given in absolute albedo units (G) and is defined as the difference between the Rayleigh albedo retrieved in the presence of simulated clouds and the reference no-cloud case. The light grey curves represent the error for each of the ten orbits used in the simulations while the red curve is the median value for all orbits.

Figure 21. Estimated errors in retrieved cloud parameters obtained from averaging over the 10 sample orbits of simulated retrievals. Albedo, particle radius and ice water content errors are plotted in the top, middle and bottom rows, respectively, with each column representing a different solar zenith angle bin. Solid curves are the mean absolute error (retrieved-true), while dashed curves are the standard deviation. Errors are plotted over the range of the corresponding grid used to bin the data for each parameter. The IWC grid was generated from the combined albedo and radius grids using the defining relationship in Equation (11). Multiple, color-coded curves are plotted for each parameter, corresponding to the primary dependence of interest (particle size for albedo and IWC errors, and albedo for the particle size errors). Color keys for all parameters are in the left panel of each column.

Figure 22. Five alternative parameterizations of the Gaussian width vs. radius relationship derived in Baumgarten et al. (2010) from fits to lidar data. These size distribution assumptions were used in the T-matrix calculations to generate alternative ice phase functions used in the sensitivity study. The top panel shows the underlying width vs. radius curves. Case 1 (smoothed version, dashed red curve) is the form used operationally in the CIPS V4.20 retrievals. The bottom panel shows the resulting ice phase functions for 50-nm particles (normalized to unity at 90-degree scattering angle).

Figure 23. Summary of the absolute differences in albedo (left), particle radius (middle) and ice water content (right), each binned over its typical range of values. Differences are shown for all alternative phase function cases, relative to the reference Case 1, and are averaged over all solar zenith angles and all six test orbits. Solid curves are the mean differences while dotted curves are the standard deviations of the differences. The top row shows the results from Cases 2-5, which use different assumptions for the Gaussian $s(r_0)$ relationship (see Figure 22). The bottom row shows results for the five cases that assume different particle shapes. Results are color-coded according to the legend in panels (a) and (d).

Figure 24. CIPS level 2 data for orbit 22789 on July 1 2011 (NH). The image on the bottom is the standard albedo strip for the orbit. Plotted in panels (a)-(c) are the retrieved albedo, radius and IWC as a function of solar zenith angle for all cloud pixels in the orbit. Pixels plotted in red have a retrieved particle radius less than 20 nm. Panel (d) shows the cloud fraction binned in 2.5-deg SZA bins for four different albedo threshold values. In panel (e) the albedo, radius and IWC are binned as a function of SZA, color-coded according to the legend (note the different scale factors applied to each parameter). Finally, panel (f) shows the retrieval histograms for each cloud parameter, using the same color-coding as before.

Figure 25. A sample CIPS level 3a “daily daisy” for June 29 2012. The image shows a polar projection plot of albedo for all 15 orbits in the day. The orbit number labels around the outside edge identify individual orbits. The color bar in the top left corner sets the albedo scale in standard CIPS albedo units (G).

Figure 26. Contour plots of PMC frequency, albedo and particle radius for the NH 2011 season from the level 3c 2 G threshold files. Ascending and descending node data are plotted on the top and bottom rows, respectively.

Figure 27. Sample CIPS level 3d data for orbit 1200 on July 15 2007. The top left panel shows the spatial distribution of CIPS common volume (CV) pixels. The red circle denotes the location of the SOFIE solar occultation point while small blue cross symbols represent all CIPS level 2 pixels within this range. The circles are the 305 CV pixels for this orbit (brown and black denote cloud and cloud-free pixels, respectively). The shaded lower region indicates location of the day-night terminator. The remaining three panels show the CIPS CV albedo, radius and IWC as a function of SZA. The scale at the top indicates distance of the CV pixels from the SOFIE occultation point. Dashed horizontal lines represent the CV-mean value of each parameter, also written in black in the top right-hand corner. For radius and IWC the corresponding SOFIE retrieval for this orbit is indicated by the red circle, with the numerical value also printed in red.

Figure 28. Sample results from a level 3e data file. This map shows a single coincidence with the Alomar ground station (69.28 N, 16.01 E) on July 8 2010. The black circle indicates the ground station location. The small grey symbols denote all CIPS level 2 pixels, with cloud pixels highlighted in red. The green and blue areas represent the 100-km coincidence radius and the 500-km “loose” coincidence region, respectively, for this station.

Figure B.1 Sample results from the Rayleigh error look-up-tables for the NH 2010 season. Each column corresponds to a different camera. The top and bottom rows show the standard deviation (random error) and the mean error (bias), respectively, expressed as a percent of the Rayleigh albedo. The solar zenith angle dependence is shown explicitly for three representative view angles, which are color-coded according to the legend. Solid and dashed curves correspond to errors for forward- and back-scattered measurements, respectively.

CIPS Scene Footprint

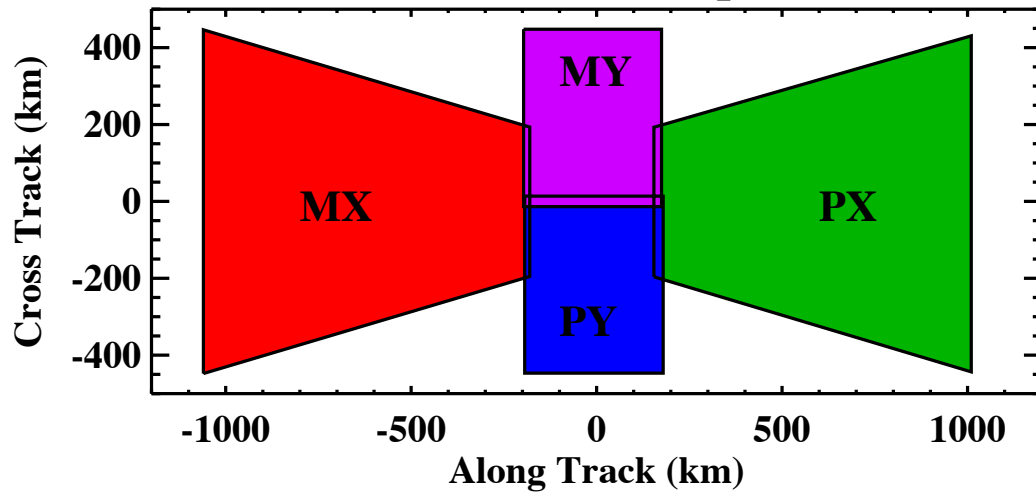


Figure 1

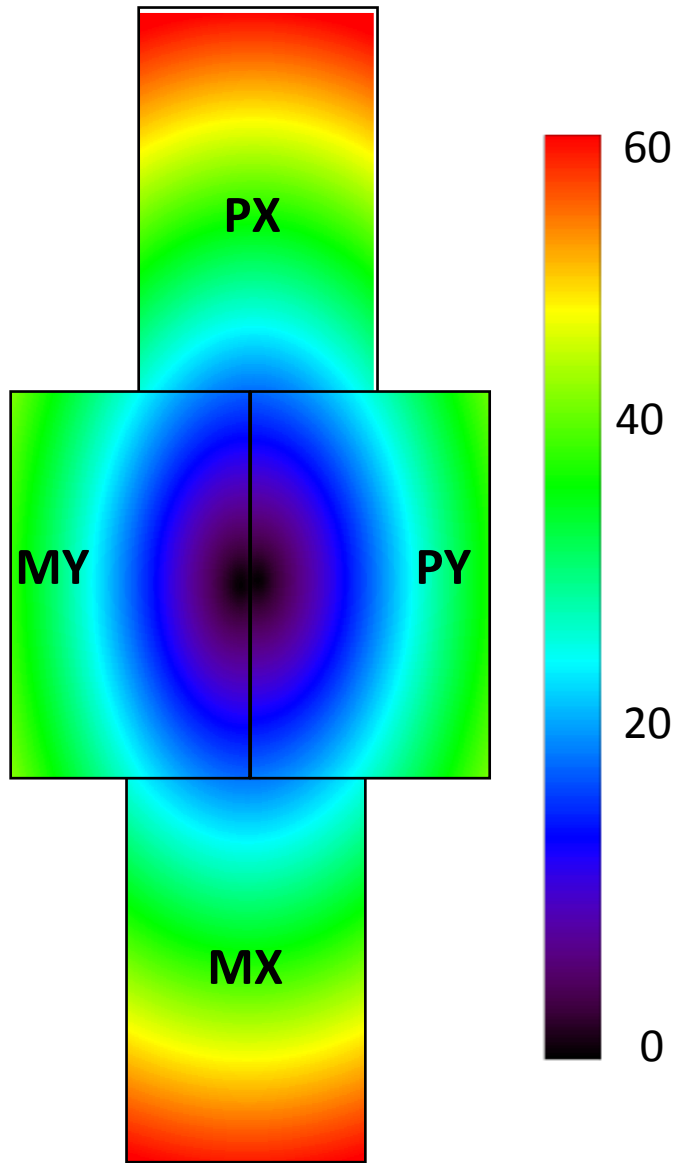



Figure 2

PX
MX
PY
MY

CIPS Level 1A Sampling - NH Solstice



PX
MX
PY
MY

CIPS Level 1A Sampling - SH Solstice




Figure 3

PX MX PY MY

CIPS Data Sampling - Solstice




Figure 4

PX PY MY MX

CIPS Data Sampling




Figure 6

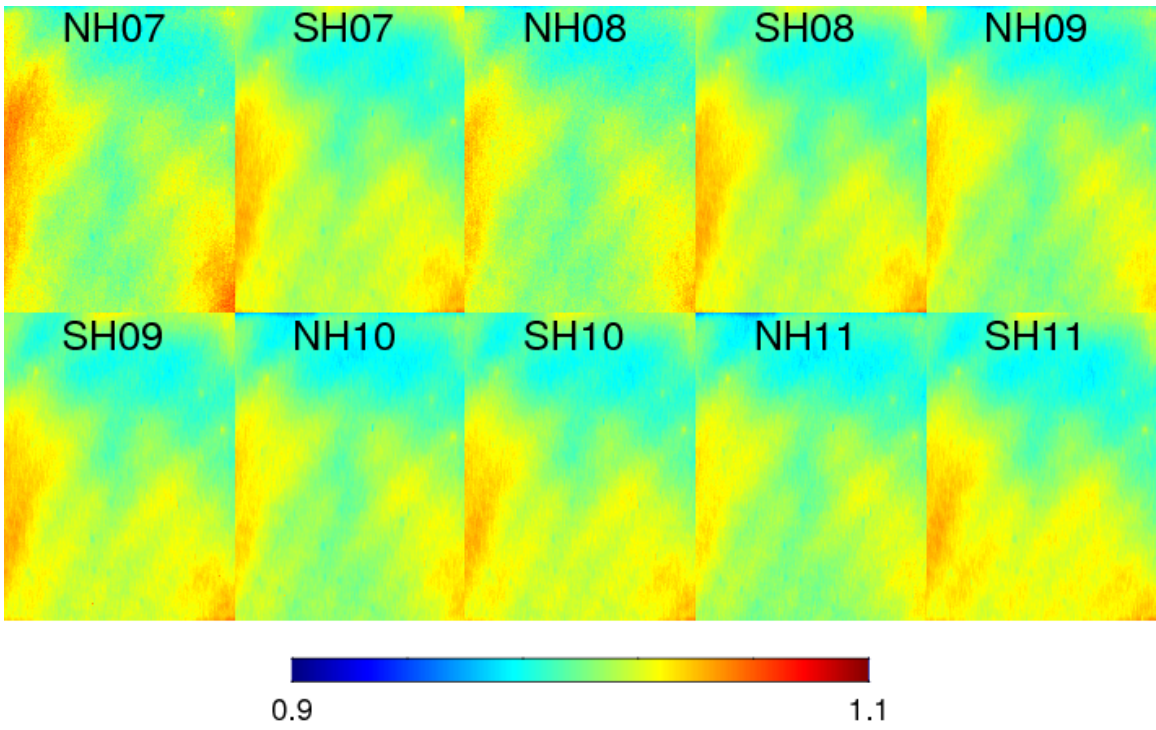
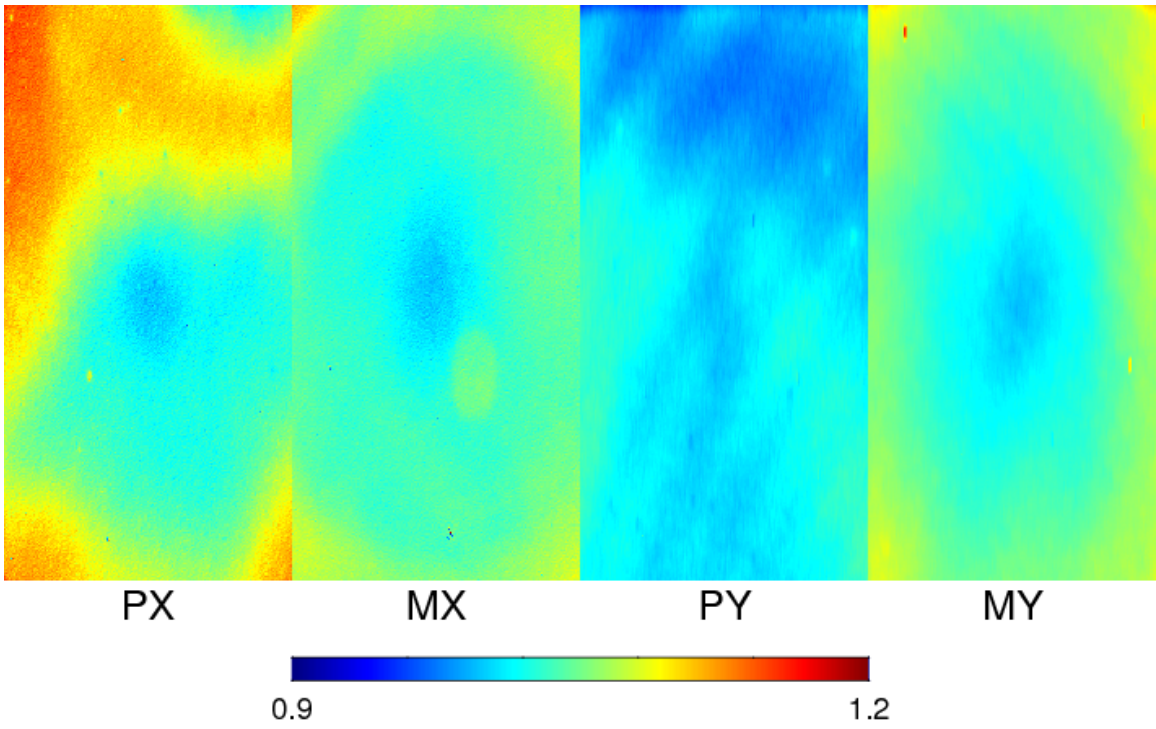


Figure 7

CIPS Camera Normalization Factors

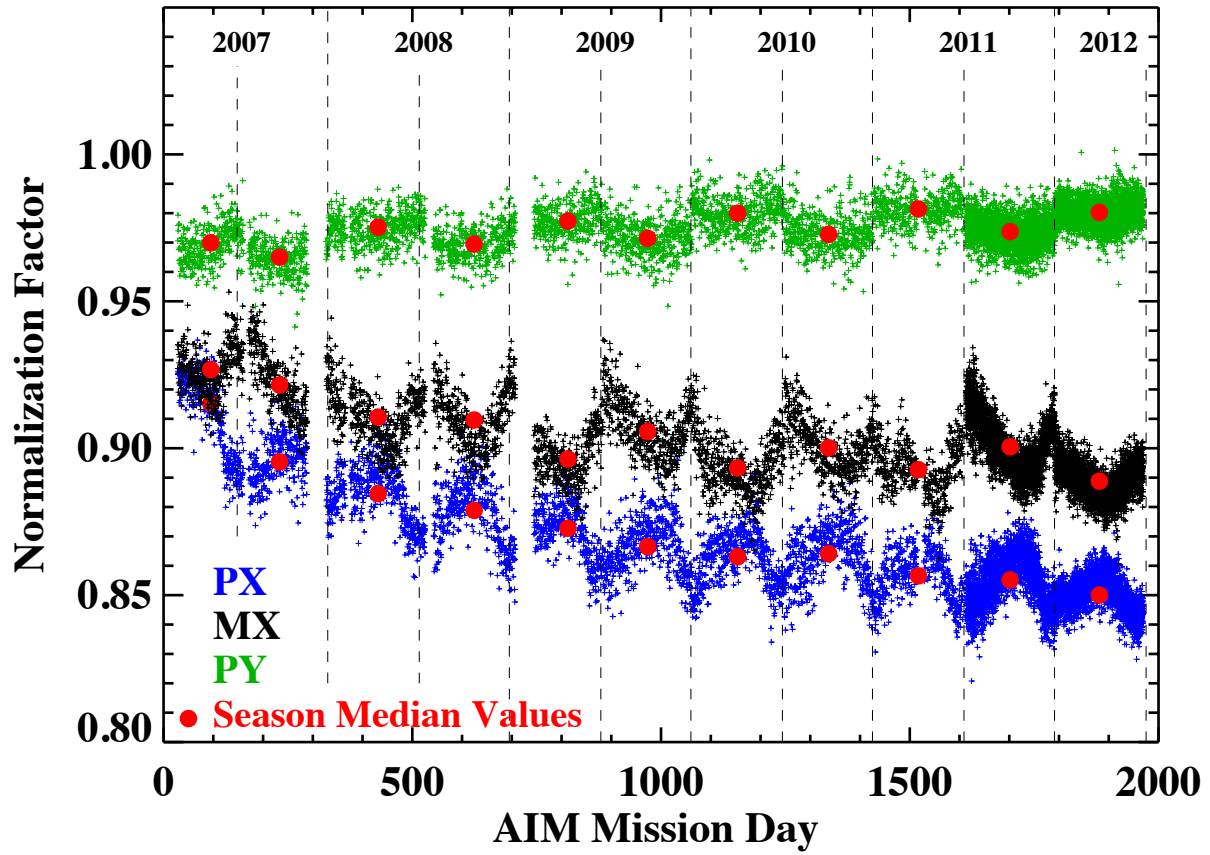


Figure 8




Figure 9

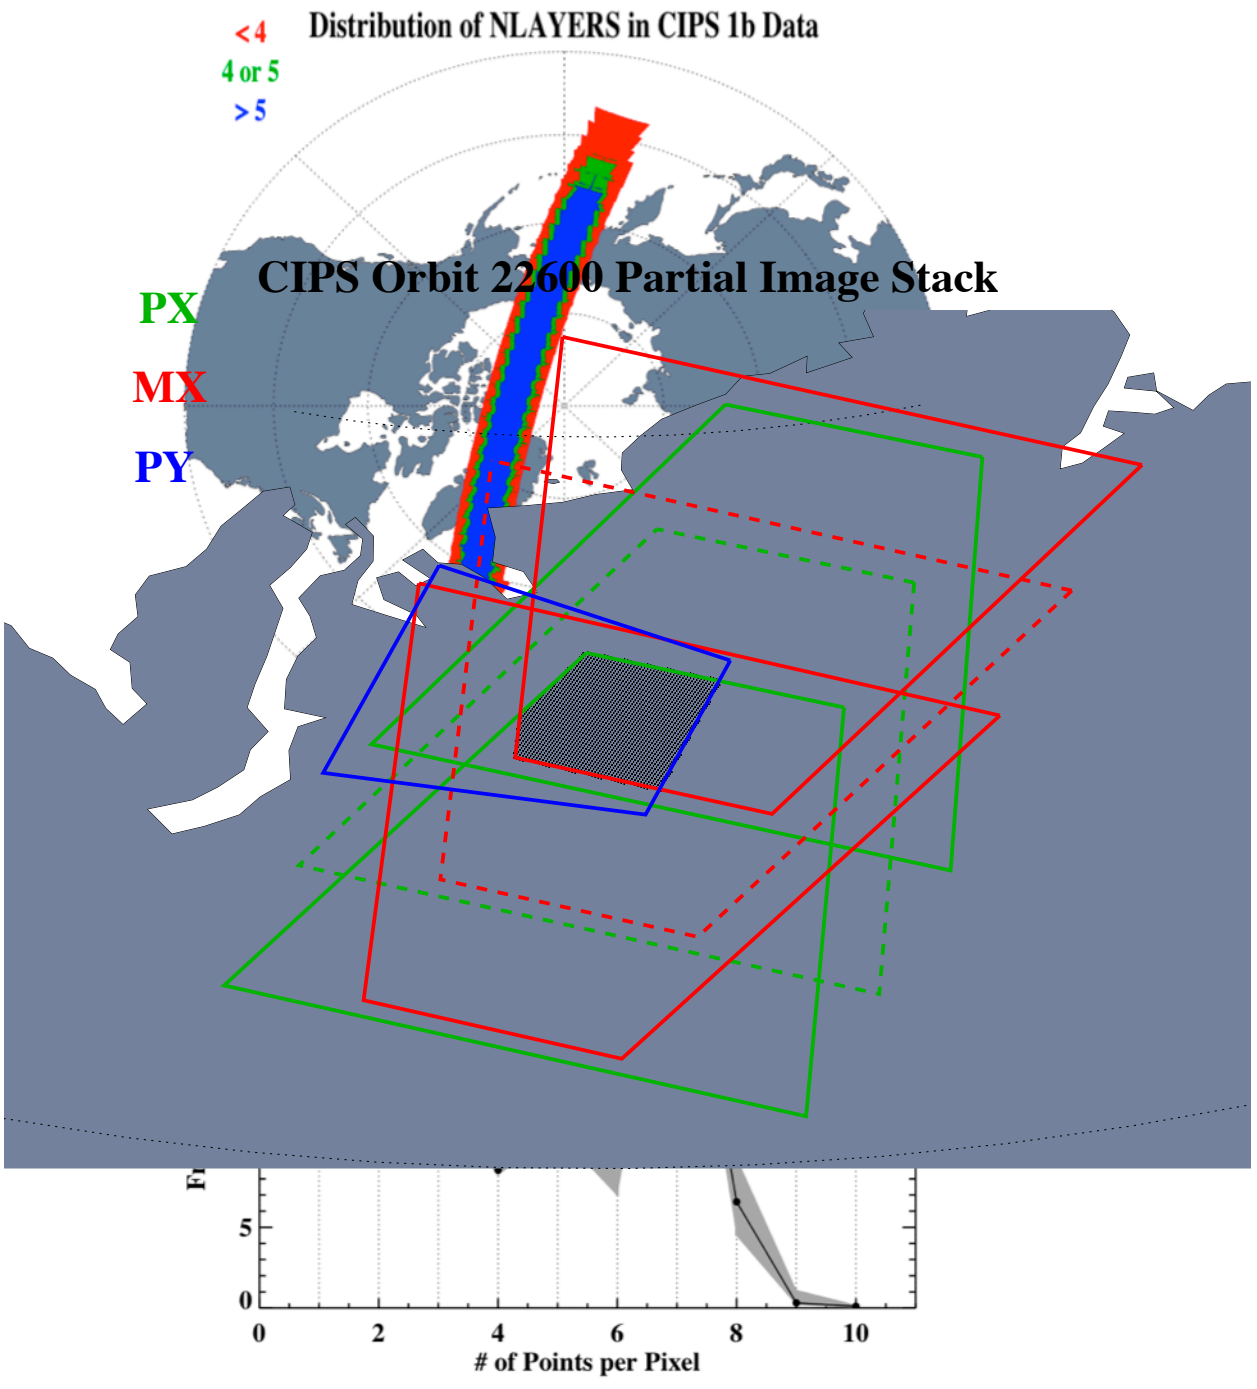


Figure 11

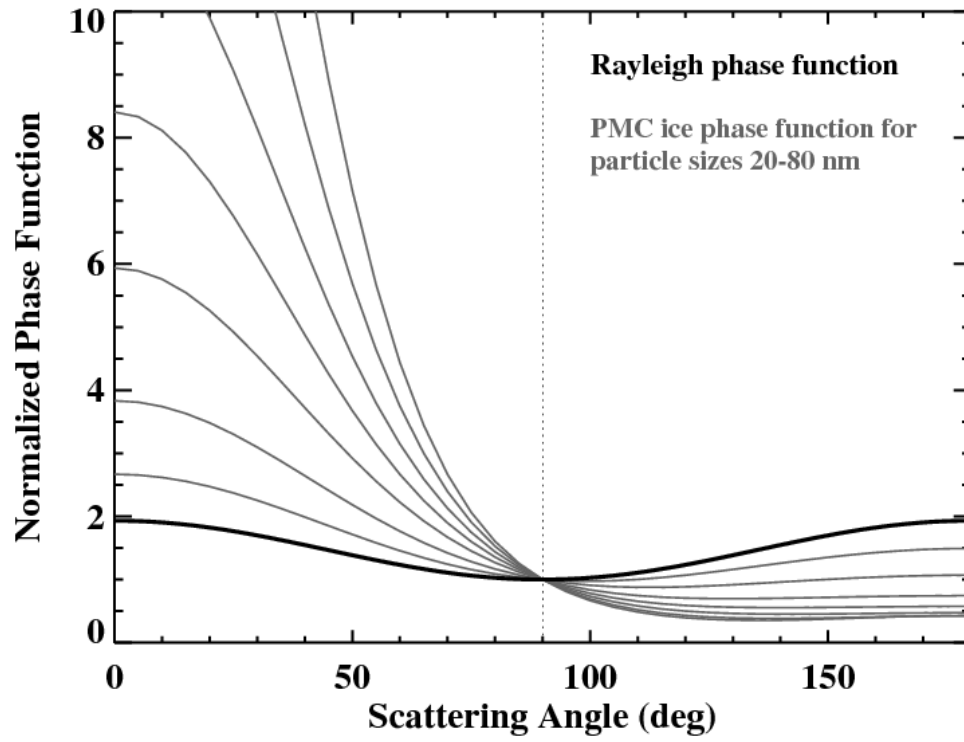


Figure 12

Cloud Albedo Contribution

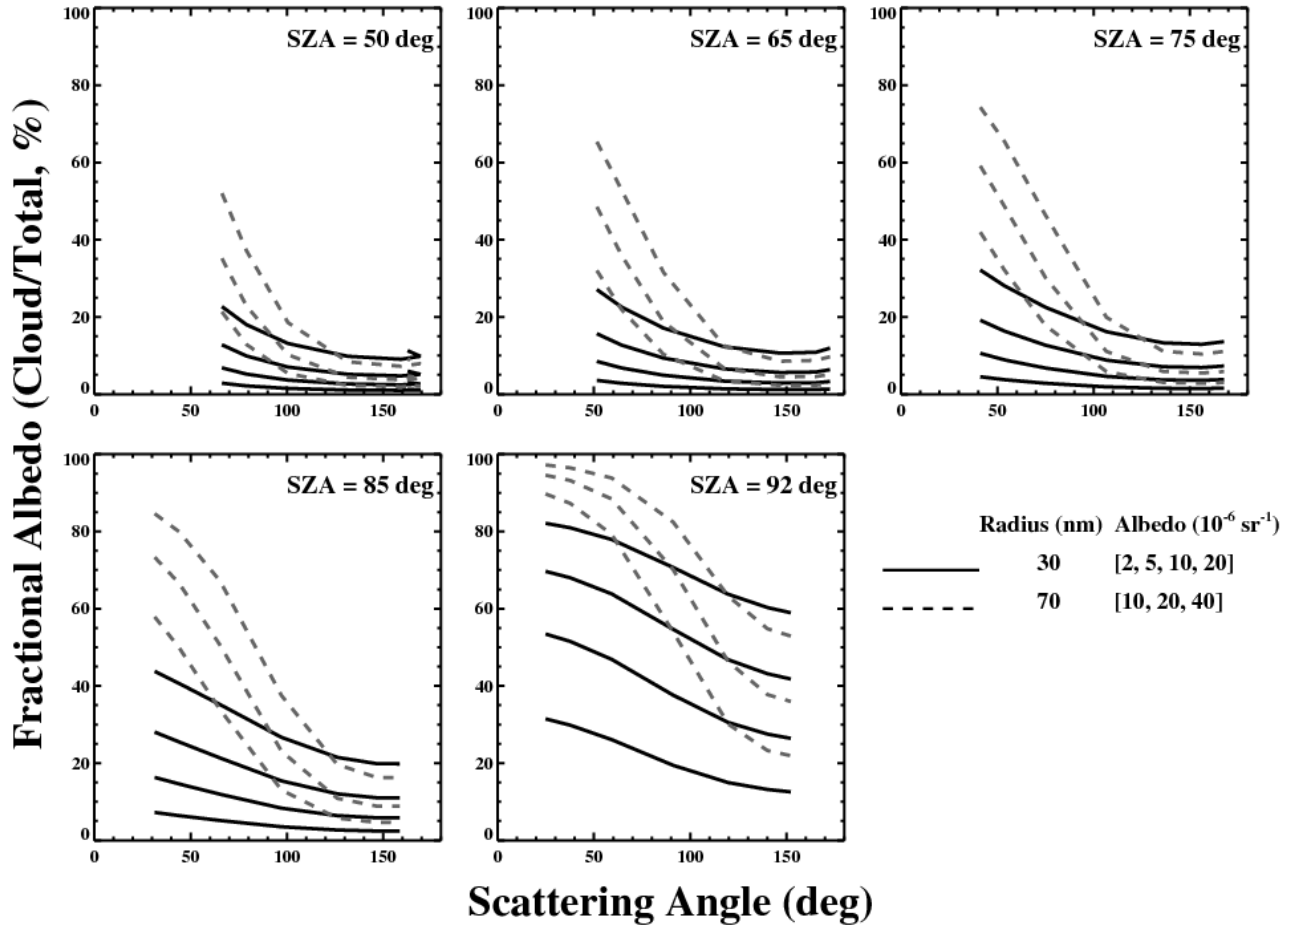


Figure 13

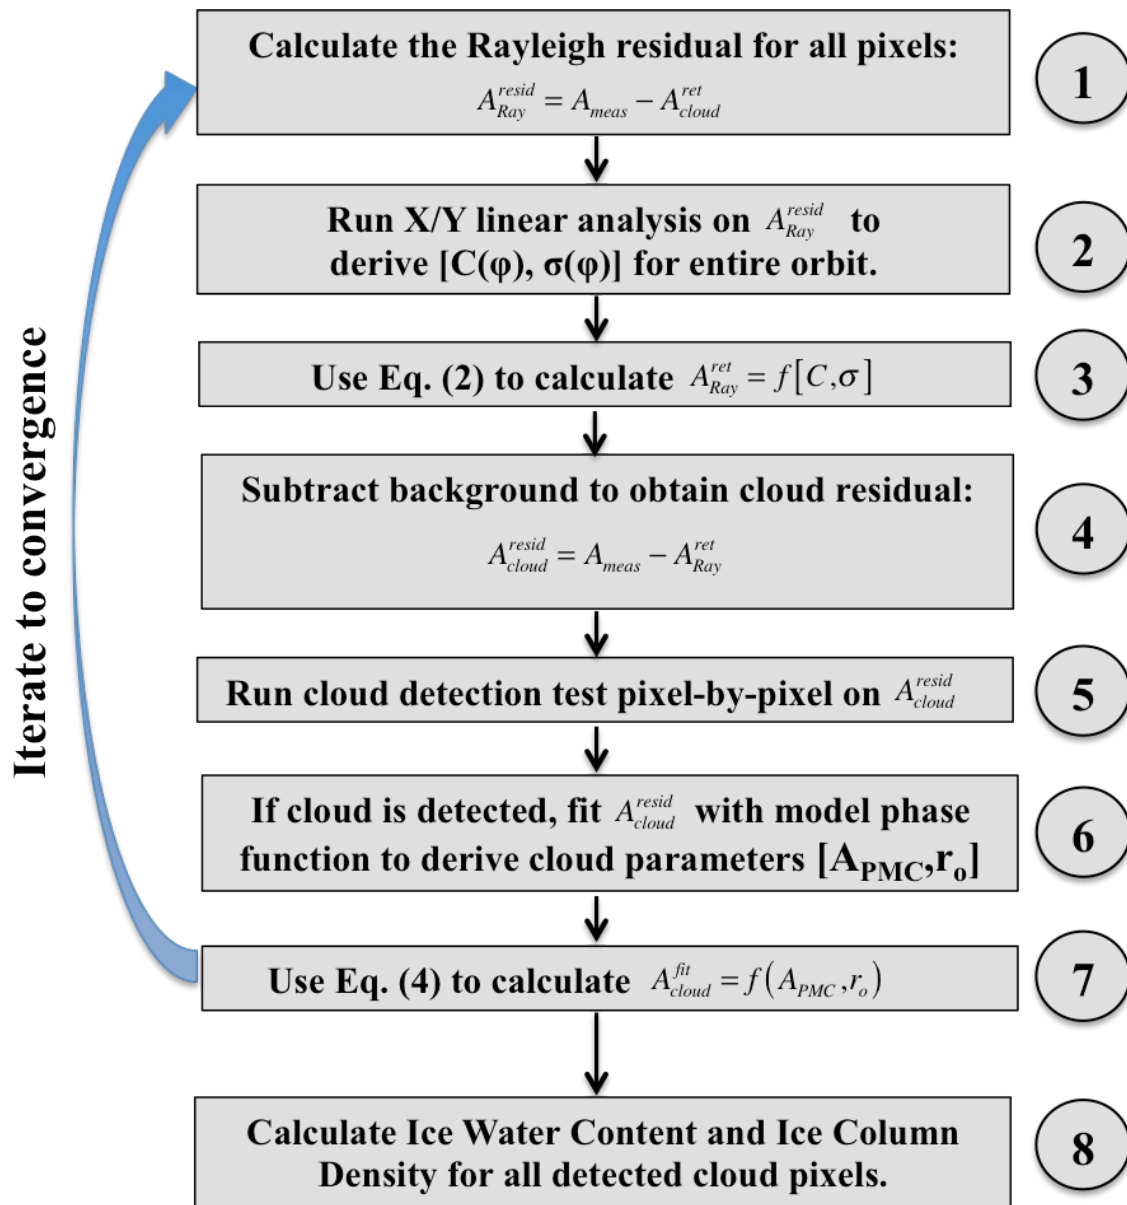


Figure 14

X/Y Fit Analysis - 0.25 deg SZA bins




Figure 15

Sample Phase Function Fits

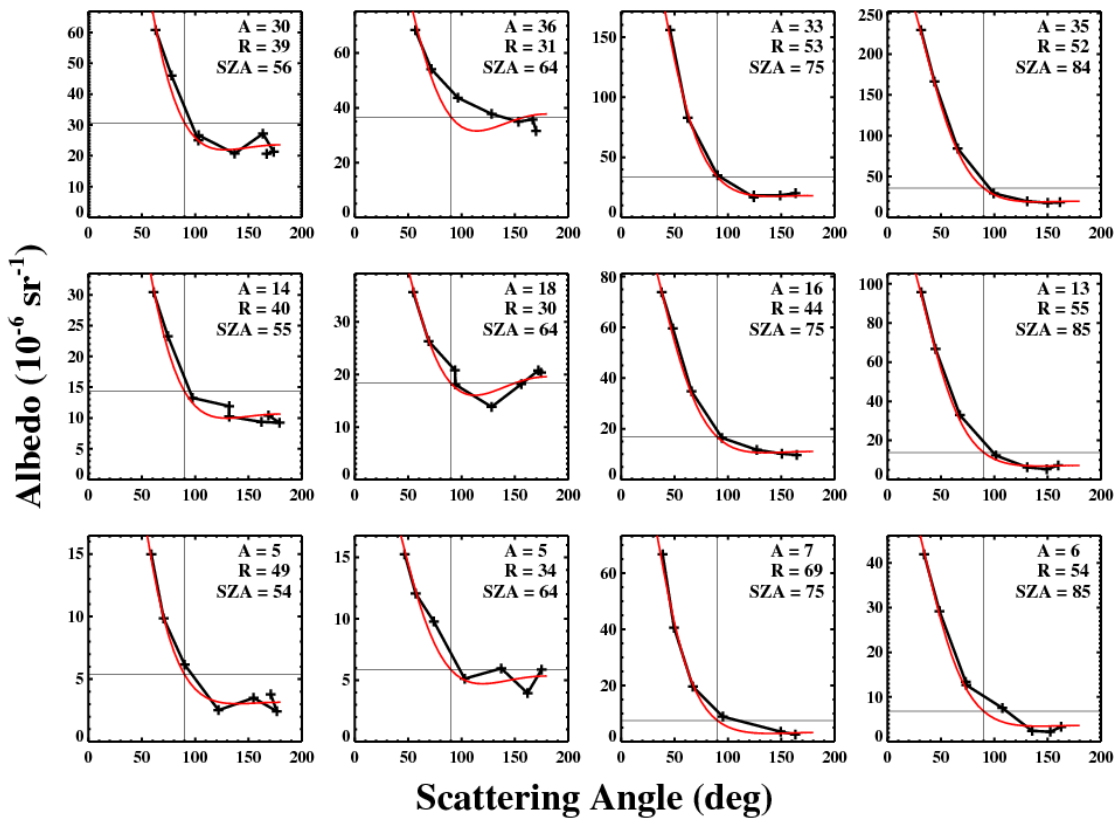


Figure 16

CIPS False Cloud Detections

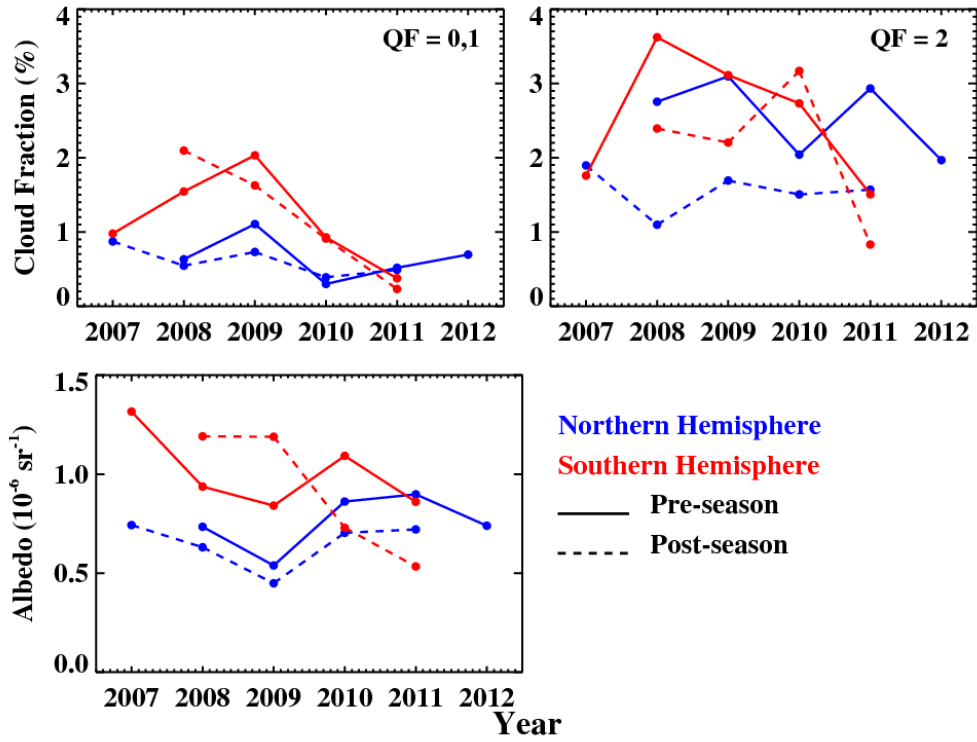


Figure 17




Figure 18

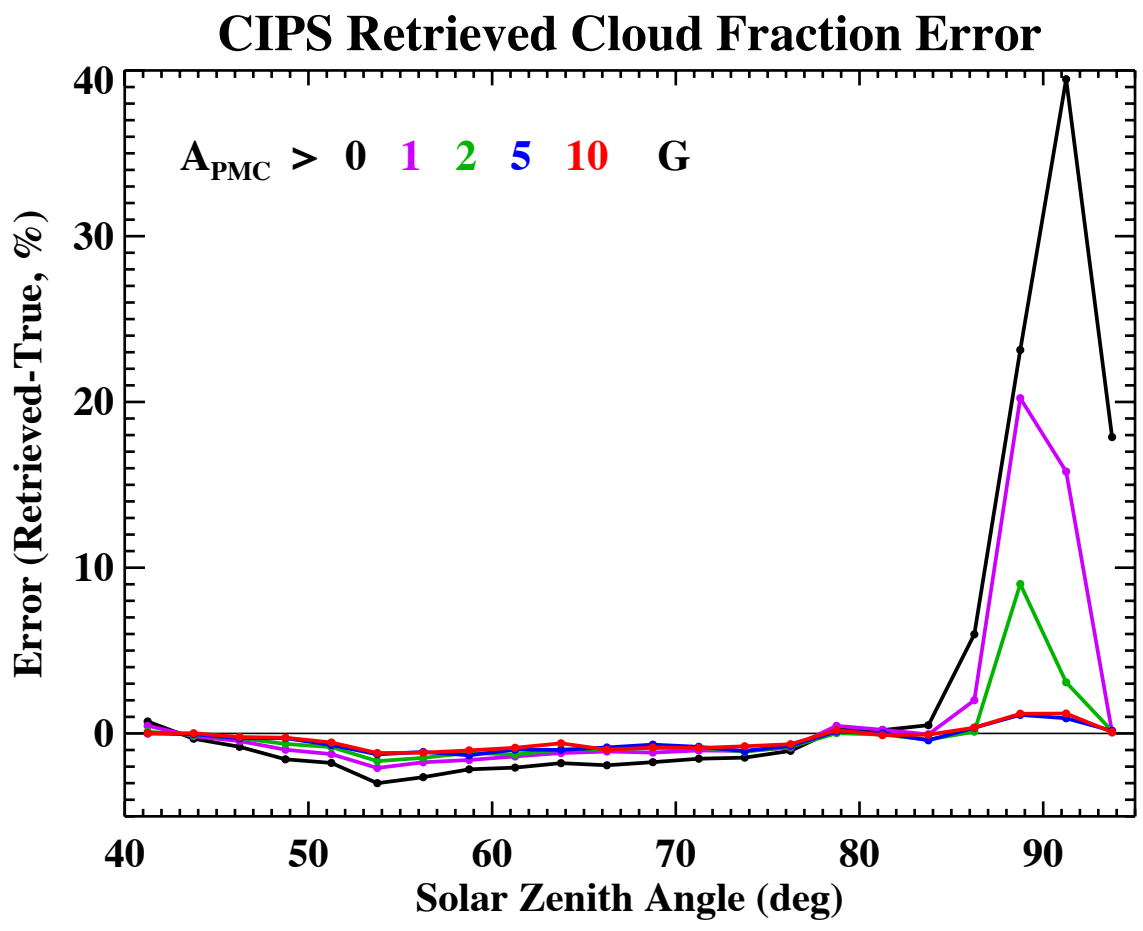


Figure 19

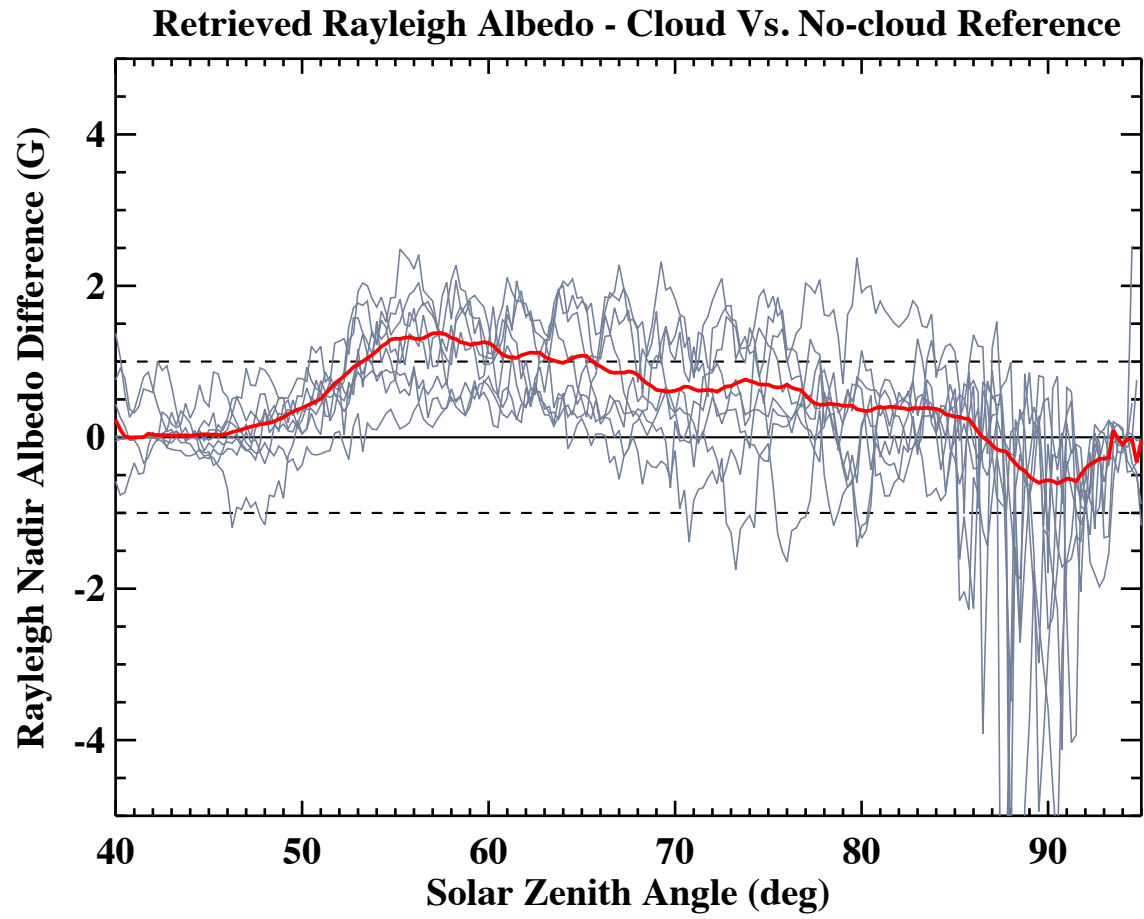


Figure 20

Estimated CIPS Retrieval Errors

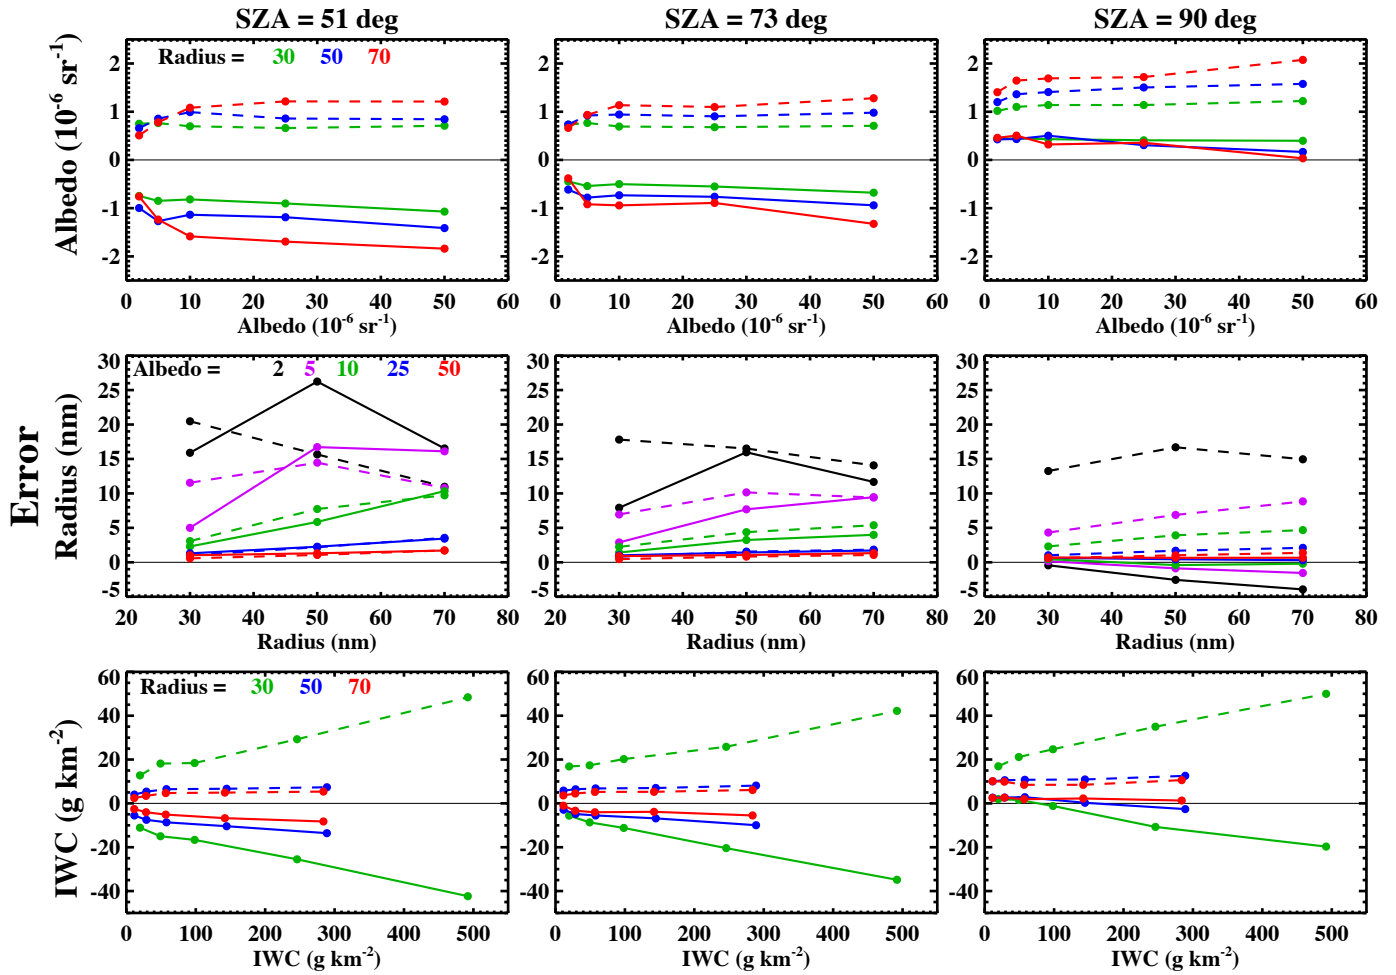


Figure 21

Alternative Ice Phase Functions

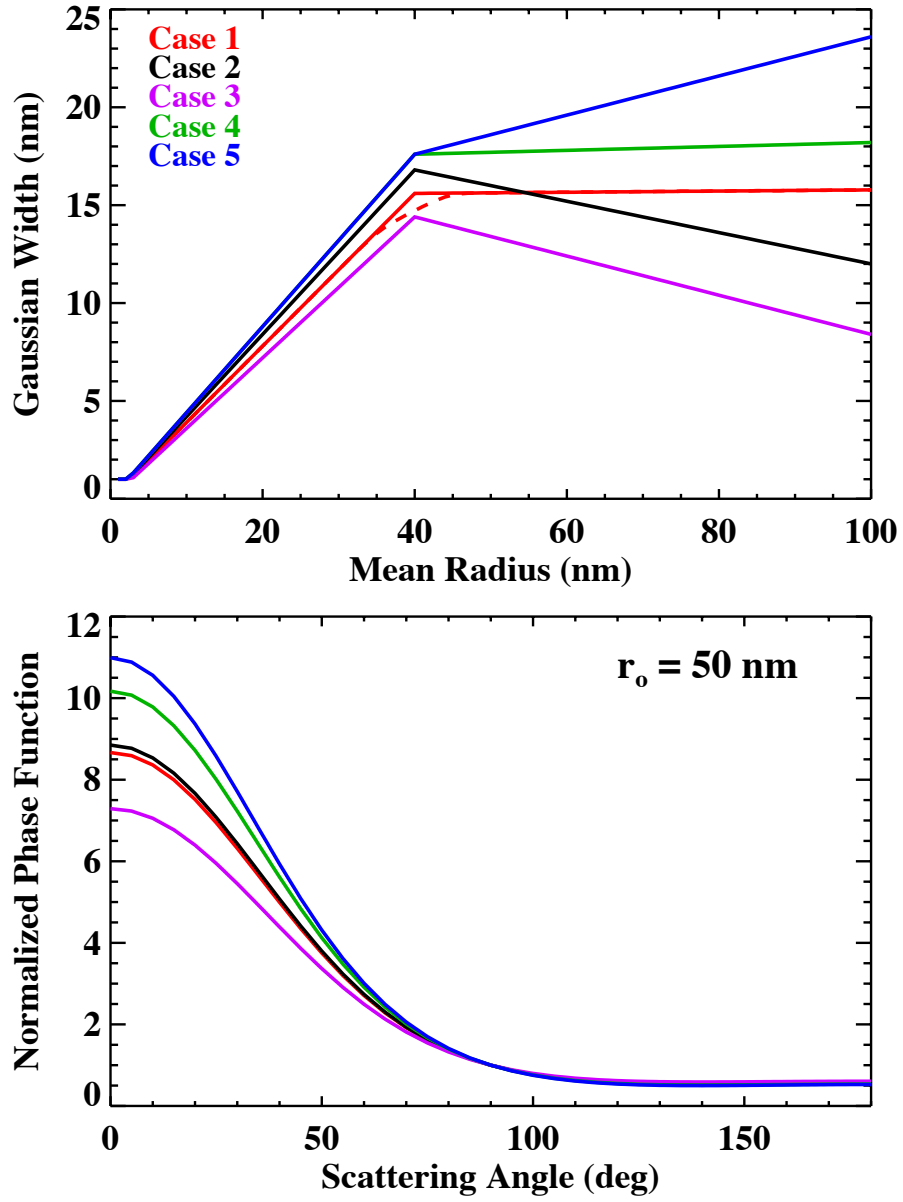


Figure 22

Retrieval Sensitivity to Phase Function

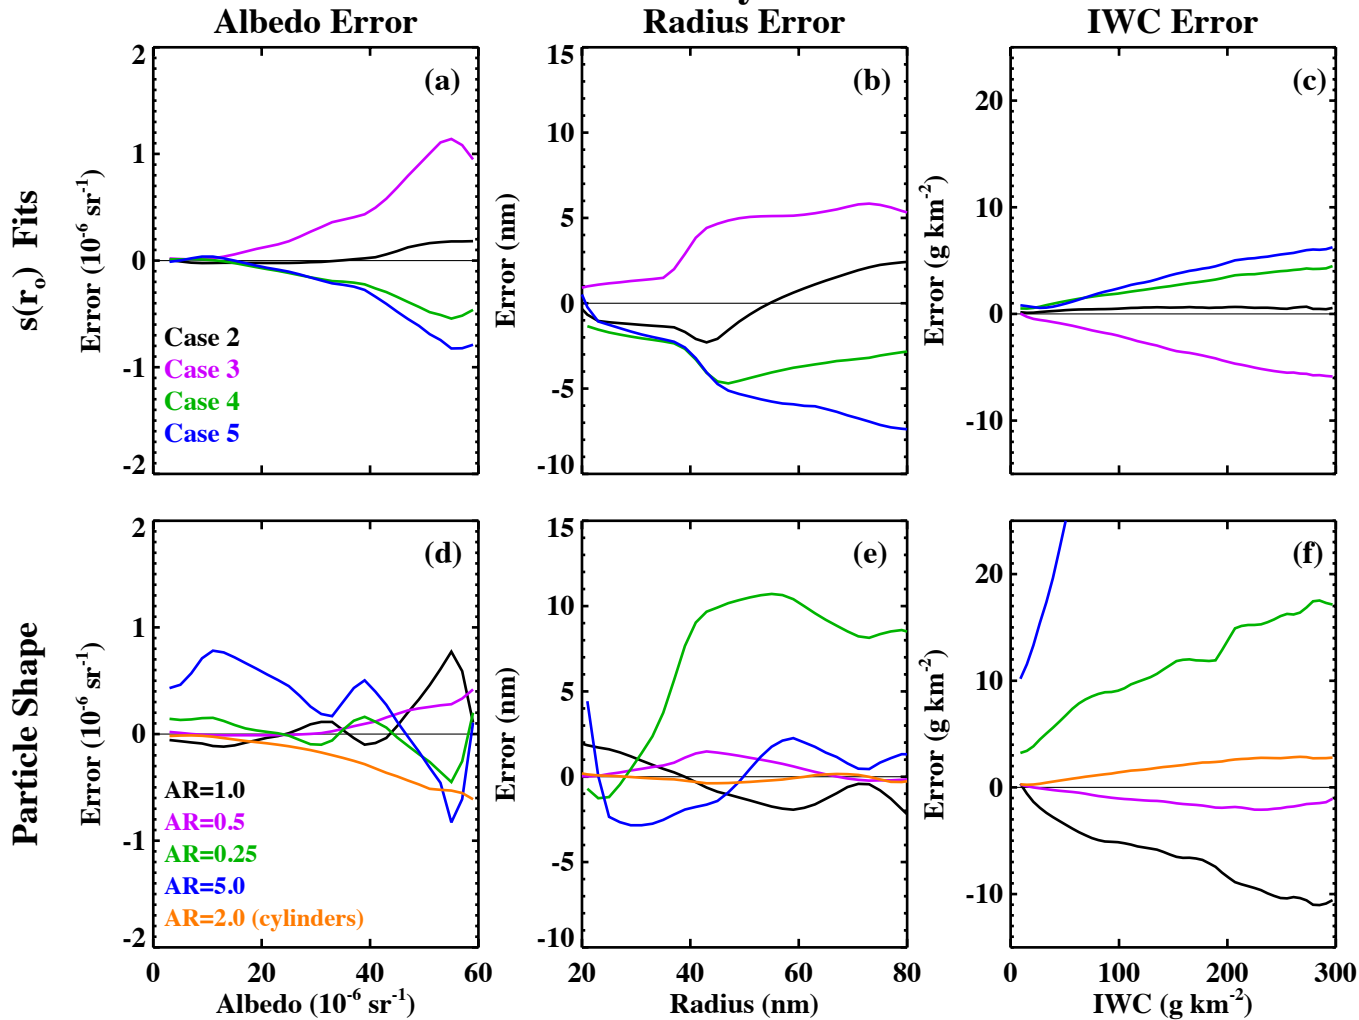


Figure 23

CIPS Level 2 Data - Orbit 22789

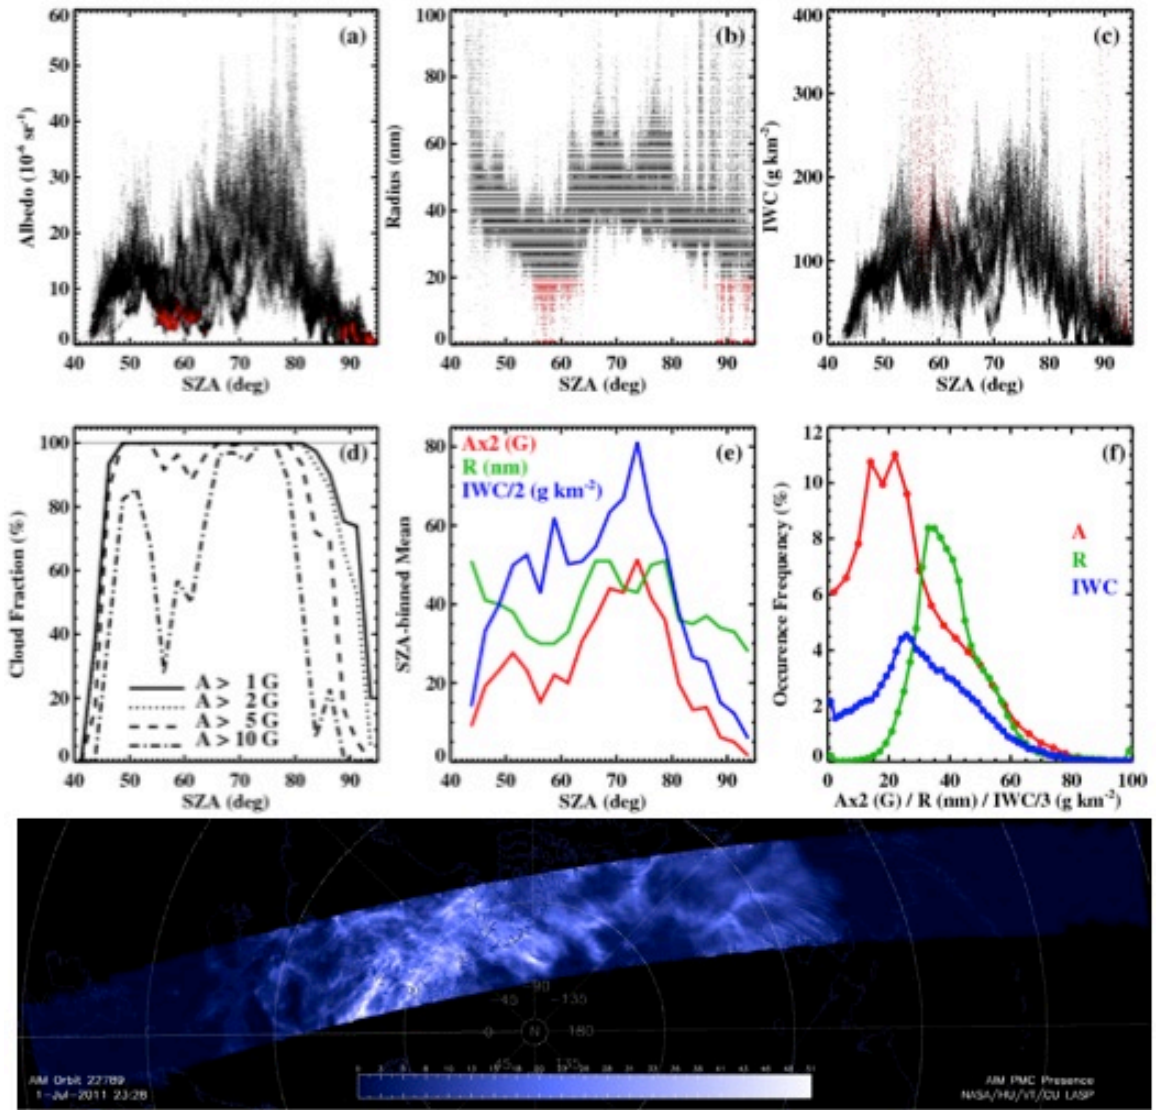


Figure 24




Figure 25

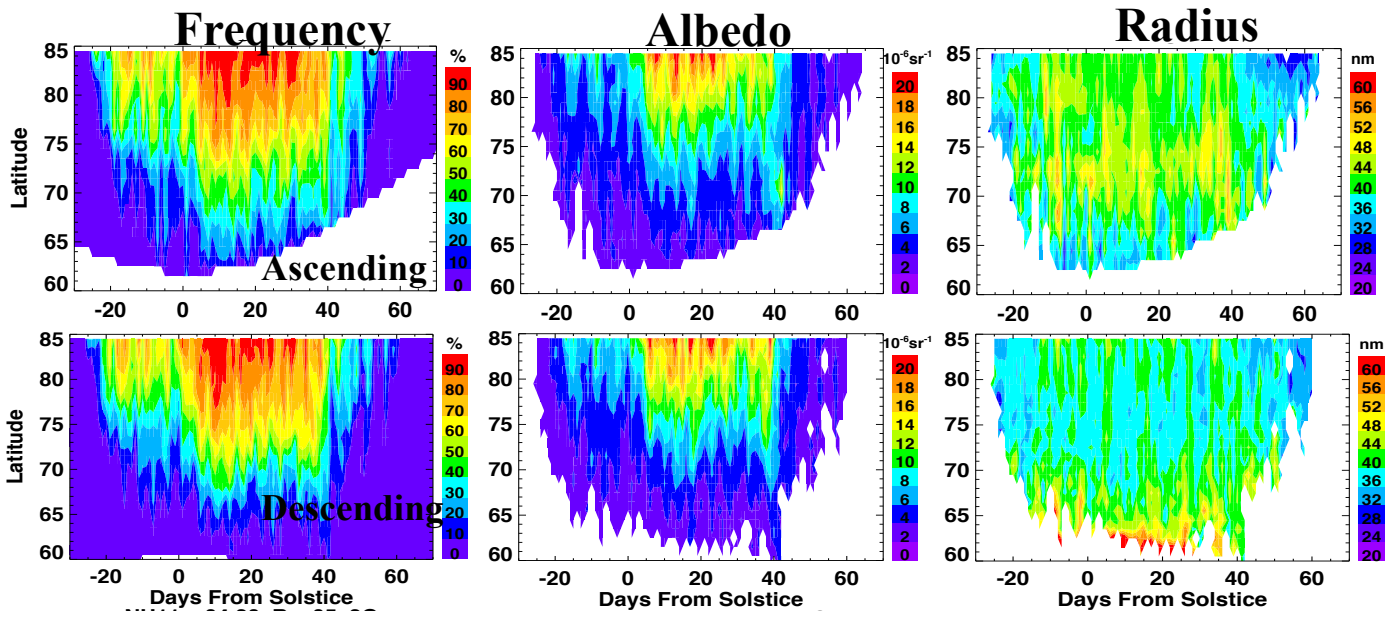


Figure 26

CIPS Common Volume Pixels - Orbit 1200

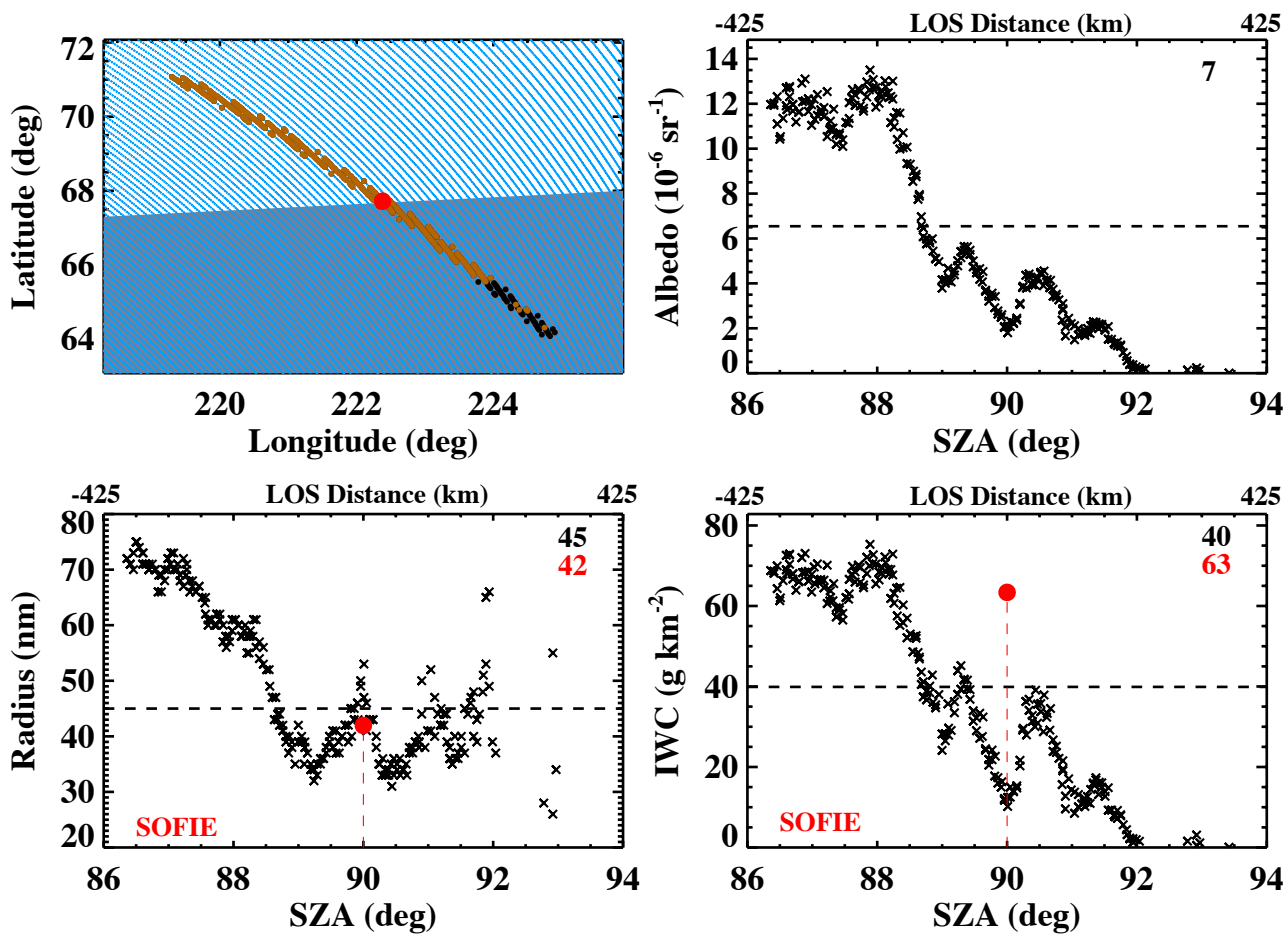


Figure 27

CIPS Coincidence with ALOMAR on 7/8/2010




Figure 28

NH 2010 Error Maps

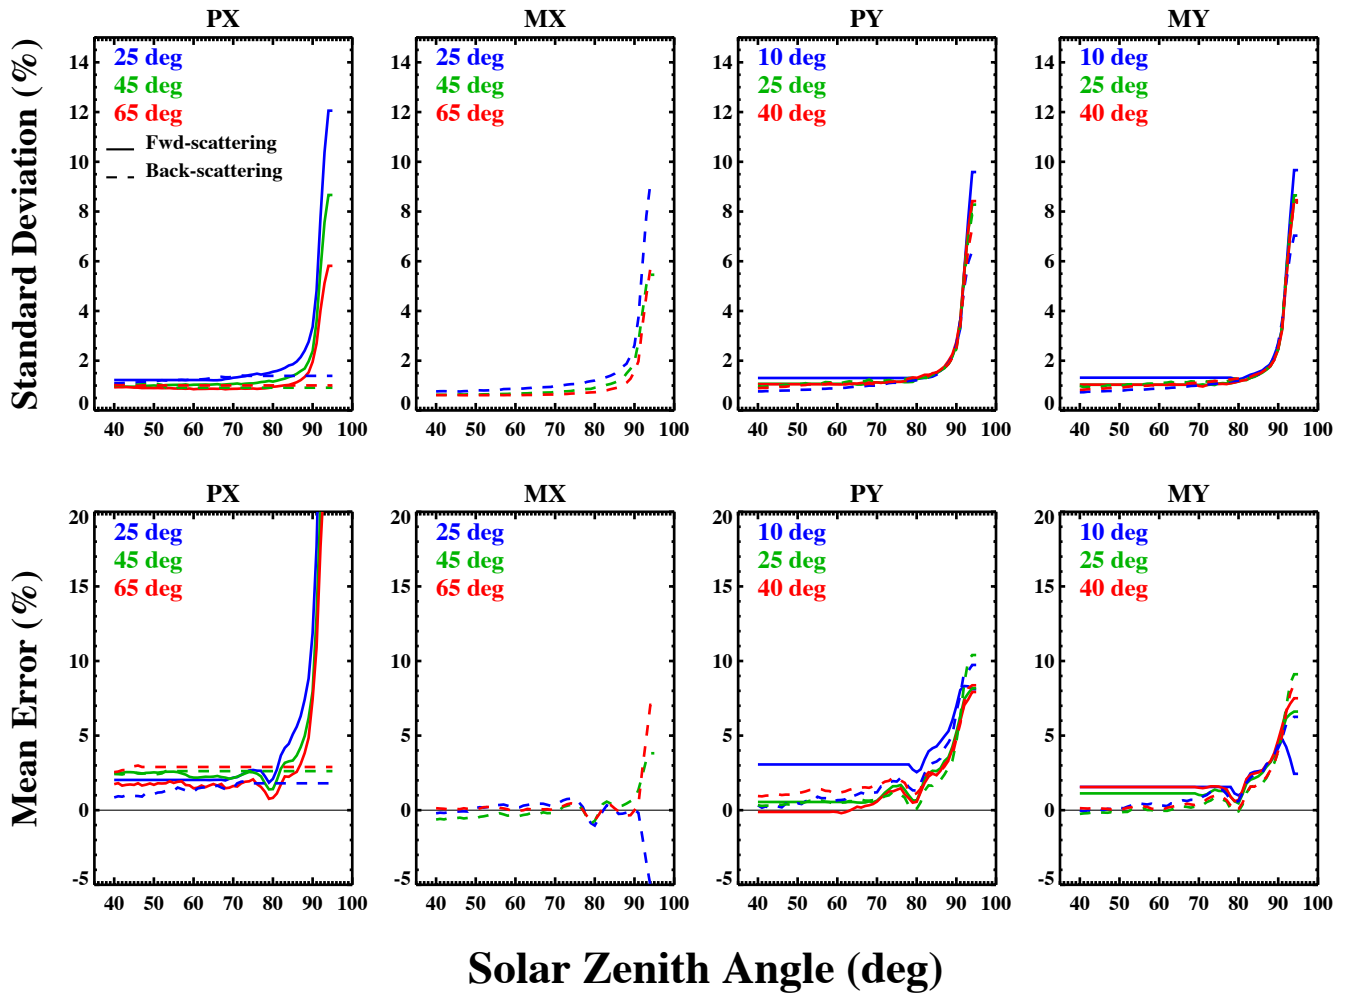


Figure B.1