JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 112, D06204, doi:10.1029/2006JD007091, 2007

An assessment of Multiangle Imaging Spectroradiometer (MISR) stereo-derived cloud top heights and cloud top winds using ground-based radar, lidar, and microwave radiometers

Roger T. Marchand,[1] Thomas P. Ackerman,[1] and Catherine Moroney[2]

Received 18 January 2006; revised 4 November 2006; accepted 6 November 2006; published 17 March 2007.

[1] In this article stereoscopically derived cloud top heights and cloud winds estimated from the Multiangle Imaging Spectroradiometer (MISR) are assessed. MISR is one of five instruments on board the NASA Terra satellite. The cloud top height assessment is based on a comparison of more than 4 years of MISR retrievals with that derived from ground-based radar and lidar systems operated by the U.S. Department of Energy Atmospheric Radiation Measurement program. The assessment includes a comparison of the MISR cloud top heights and ground-based data sets as a function of cloud optical depth and a simple cloud classification. Overall, we find that the MISR retrieval is working well with little bias for most cloud types, when the cloud is sufficiently optically thick to be detected. The detection limit is found to be around optical depth 0.3 to 0.5, except over snow and ice surfaces where it is larger. The standard deviation across all clouds is less than about 1000 m for the MISR best winds retrievals at all ARM sites, and the standard deviation for the MISR without winds retrieval varied between about 1000 to 1300 m, depending on the site. The performance for various cloud types is explored.

Citation: Marchand, R. T., T. P. Ackerman, and C. Moroney (2007), An assessment of Multiangle Imaging Spectroradiometer (MISR) stereo-derived cloud top heights and cloud top winds using ground-based radar, lidar, and microwave radiometers, *J. Geophys. Res.*, *112*, D06204, doi:10.1029/2006JD007091.

1. Introduction

[2] Satellites are an immensely valuable tool for monitoring many cloud properties on a global basis. One cloud property that satellites are in a particularly advantageous position to monitor is cloud top height. In this article we provide an assessment of stereoscopically derived cloud top heights and cloud winds produced from observations by the Multiangle Imaging Spectroradiometer (MISR). MISR is one of five instruments on board the NASA Terra satellite, which was launched in December of 1999 [*Diner et al.*, 2002, 2005].

[3] The cloud top height assessment is based on a comparison of MISR-retrieved cloud top heights, which we will refer to as stereo-heights, with those derived from ground-based cloud radar and lidar systems operated by the U. S. Department of Energy Atmospheric Radiation Measurement (ARM) program. We use ARM data from three of its primary sites located in the U.S. Southern Great Plains (SGP) at Lamont Oklahoma, the North Slope of Alaska (NSA) at Barrow Alaska, and in the Tropical Western Pacific (TWP) on Nauru Island. The cloud top height assessment includes a comparison of the MISR and ground-based data sets as a function of cloud optical depth and a simple cloud classification. The cloud optical depth is determined from ARM ground-based observations using several retrieval techniques including a radar-reflectivity-velocity and a radar-reflectivity-microwave-radiometer technique, as well as simple lidar and radar only techniques. The assessment of MISR cloud top winds is based on a comparison of MISR-retrieved cloud velocity (commonly referred to as cloud winds) with wind speeds derived from a 404 MHz radar wind profiler (RWP).

[4] Section 2 of the paper briefly describes the MISR and ground-based data sets, as well as the optical depth retrievals and the approach used to combine the retrievals into a single best estimate (along with an estimated uncertainty). A comparison of the MISR cloud winds and stereo-heights with ground-based data is given in sections 3 and 4, respectively, along with a few examples that illustrate limitations of the MISR approach in section 5.

[5] A number of articles have been published over the last few years that examine MISR stereo-heights, including *Horvath and Davies* [2001a], *Marchand et al.* [2001], *Moroney et al.* [2002], *Muller et al.* [2002], *Naud et al.* [2002, 2004, 2005a, 2005b], and *Seiz et al.* [2006]. These articles are narrower in focus than this study and examine individual cases or small sets of a particular cloud type. The results of these previous studies are largely consistent with those presented here. We summarize our results and comment on the similarities and difference with these early studies in section 6 and make a few final remarks on the

[1]Pacific Northwest National Laboratory, Richland, Washington, USA.
[2]NASA Jet Propulsion Laboratory, Pasadena, California, USA.

0148-0227/07/2006JD007091$09.00

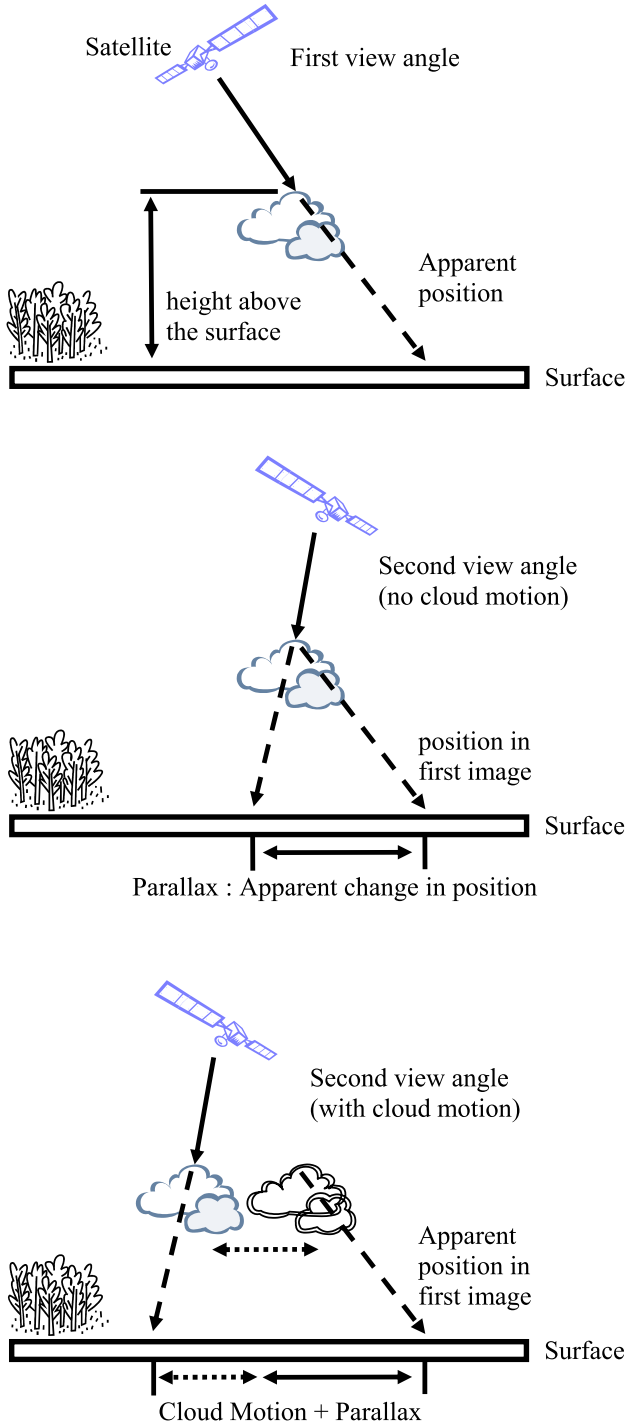

Figure 1. Depiction of parallax effect, with and without cloud motion/clouds winds.

implications of the results for future satellite missions in section 7.

2. Description of Data

2.1. MISR Cloud Top Height and Cloud Winds

[6] The Multiangle Imaging Spectroradiometer (MISR) was launched on board the EOS Terra spacecraft in December 1999. The orbit is sun-synchronous at a mean height of about 705 km, with an inclination of 98.5°, an equatorial crossing time at about 10:30 am, and an orbit repeat cycle of 16 days. The MISR instrument consists of nine pushbroom (or line imaging) cameras, each of which makes high-resolution images (with approximately 275 m sampling) in four narrow spectral bands located at 443, 555, 670, and 865 nm. These cameras collect data at nine view angles (nadir plus 26.1, 45.6, 60.0, and 70.5° forward and aft of the direction of flight). The time delay between adjacent camera views is 45–60 s with a total acquisition time between the 70.5° aft and 70.5° forward images of about 7 min. In normal operation (called Global Mode) all the data in the red band from all nine cameras and all the data in the blue, green and NIR bands in the nadir camera are saved at the full 275 m sampling. The data of the blue, green and NIR bands of the remaining eight cameras are averaged to 1.1 km [*Diner et al.*, 2002].

[7] The MISR stereo-height retrieval is conceptually simple, as depicted in Figure 1. An object located above the surface (such as a cloud) will appear in two different positions when viewed from two different angles. From the apparent change in position (or parallax) one can calculate the height of the object relative to the surface. Using multiple views, one can estimate both the height and velocity of the object (assuming the object moves sufficiently far in the time between when images are acquired).

[8] While the approach is conceptually simple, there are many practical difficulties. For example, one must identify not just the same object, but the same point on the object to get an accurate result. In the case of MISR (where the height retrieval is based on the nadir and 26.1° view angles), a one pixel error in the estimate of the parallax results in about a 560 m error in the height retrieval [*Moroney et al.*, 2002; *Muller et al.*, 2002]. If the apparent shape of the object changes in between the multiple views (either because the cloud structure evolves in time or simply because one sees different "sides" or parts of a cloud), it may be impossible to find a matching point in two or more views. Also, many clouds do not have distinct boundaries, but rather are diffuse with photons being scattered toward the satellite from a range of altitudes (with different view angles "seeing" photons scattered from different depths with in the cloud). Multilayer scenes with a thin layer on top of lower and more reflective clouds, in particular, cause problems when trying to match oblique cameras view with near nadir cameras (as performed in the MISR cloud winds retrieval).

[9] The MISR operational algorithm is described by *Moroney et al.* [2002] and *Horvath and Davies* [2001b]. In brief, the operational algorithm first uses images from three forward viewing cameras to estimate the height and velocity for a subset of the most distinctive features (i.e., objects in the scene) on a 70.4 × 70.4 km scale. The velocity is assumed to be constant on this scale for a maximum of two cloud populations, designated as either the "low-cloud" or "high-cloud" wind speed depending on their relative height. A second set of stereo-matching algorithms is then run using two cameras (e.g., nadir and +26 forward) to obtain the stereo heights at high resolution [*Moroney et al.*, 2002; *Muller et al.*, 2002]. The full resolution (275 m) MISR red band data are used by the stereo-matchers, but in order to save on processing time only 1 out of every group of 4 pixels is matched yielding a

stereo height field at a reduced 1.1 km sampling. (We stress the data is not averaged to ¼ the original resolution. The matching is applied on the original 275 m images, the retrieval is just not run on every possible pixel.) The conversion of image disparity (horizontal displacement) to height is performed in two ways, first by assuming that there is no cloud/object motion (i.e., a velocity of 0 m/s), which yields the without winds heights, and once using the existing 70.4 km domain cloud motion vector, if any. (This second retrieval actually yields two separate height flavors: best winds and raw winds. The best wind retrieval only uses those winds that pass a wind quality test. In the version of the algorithm analyze here this is based on how strongly peaked a histogram of the retrieved disparities is [*Moroney et al.*, 2002]. The raw winds are calculated using all available wind vectors regardless of their quality. The raw winds are not recommended for general use and we concentrate only on the best winds retrievals in this article). The 1.1 km cloud top height retrieval is performed using the MISR nadir and +26 forward view and again using the nadir plus −26 aftward view. If both retrievals are successful, only those areas which yield the same height (to within a given threshold, defined as twice the standard deviation of all fwd-aft height differences in the 70.4 m domain) are kept, otherwise if only one height retrieval is successful it is automatically retained.

2.2. Ground-Based Data Sets

[10] Among the best available independent data sources for evaluating MISR stereo-heights are ground-based cloud radar and lidar systems. In this paper we use cloud top boundaries determined form a combination of these two types of sensors at three U.S. DOE ARM sites. The particular data set we used is known as the active remote sensing of cloud layers (ARSCL) data set, which is described in detail and evaluated by *Clothiaux et al.* [2000]. The cloud top boundaries contained in the ARSCL product are essentially taken from the lidar observations when the lidar penetrates the cloud and from the cloud radar otherwise. The results shown in sections 3 and 4 are based on a 2 min median values centered on the time of the MISR overpasses. We investigating using a single closest-in-time sample and a 5 min median value and found no significant differences in the results.

[11] While the ARM cloud radars are among the most sensitive cloud radars available, there are high thin ice clouds which are not sufficiently reflective (at millimeter wavelengths) for theses radars to detect. As a result, there is more uncertainty in the location of cloud top in this data set when the cloud is sufficient optically thick to attenuate the lidar. In situations where the lidar is able to penetrate a cloud, *Clothiaux et al.* [2000] found that the bias between the radar and lidar cloud top heights to be small, that is less than the 300 m resolution of the lidar system used in that study. Also, as we will see in section 4, the radar is routinely able to detect clouds which are too optically thin to be retrieved by the MISR stereo algorithm. Therefore the sensitivity limits of the radar do not appear to limit the statistics reported in this article.

[12] We will also compare MISR-retrieved cloud velocities, which we will refer to as clouds winds, with estimates of the wind velocity obtained from a 404 MHz radar wind profiler (RWP) at the ARM SGP site. This is the only ARM site with a 404 MHz profiler. Radar wind profilers are long wavelength (74 cm in this case) Doppler systems that measure backscattering from index of refraction variations of the air. The need to combine radial velocity observations from several directions (in order to obtain the horizontal wind speed and direction) in combination with interference from clutter and frequent low signal-to-noise ratios require careful data processing. In this analysis we use 1 hour consensus wind averaging (of wind profiles gathered once every 6 min) centered on the MISR overpass times [*Barth et al.*, 1994]. Comparisons of RWP winds with radiosondes typically show no significant biases and standard deviations between the two measurements of 3 to 5 m/s [*Martner et al.*, 1993; *Cohn and Goodrich*, 2002]. Much of this difference can be attributed to spatial and temporal mismatch between these observing systems. Comparisons of RWP measurements against tower measurements and Doppler lidar systems generally show better agreement with standard deviations of 1 to 2 m/s [*Angevine et al.*, 1998; *Cohn and Goodrich*, 2002].

2.3. Ground-Based Retrievals of Cloud Optical Depth

[13] Many satellite retrievals have problems with thin clouds and (as will be demonstrated in section 4) the MISR stereo-height retrieval is no exception. An important consideration in using the MISR stereo-heights is to know how optically thin a cloud can be before the stereo-algorithm fails to detect the cloud (that is, the retrieval produces no result or retrieves the height of the local surface). To address the thin cloud limit, we will examine the MISR stereo-height retrieval in the context of the optical depth as determined for a combination of ground-based data.

[14] Cloud optical depths and water paths vary over several orders of magnitude and there is no single ground-based remote sensing algorithm which covers the full range of clouds types and optical thicknesses observed in the atmosphere. We have therefore created an algorithm which combines the results of several retrieval techniques into one combined algorithm to obtain a best estimate along with an estimate for the uncertainty. Our combined algorithm incorporates a radar-reflectivity-only retrieval, a radar-reflectivity-and-Doppler-velocity retrieval, a radar-reflectivity-microwave-radiometer retrieval, and a lidar-only based retrieval. In Appendix A, we describe each of the retrievals and the rules used to combine them.

3. Comparison of MISR Cloud Winds With Radar Wind Profiler Data

[15] The first step in the MISR stereo height retrieval is estimating the velocity of clouds and so we begin by comparing MISR-retrieved cloud winds with estimates of the wind velocity obtained from a 404 MHz radar wind profiler (RWP). While the RWP can detect some clouds, most clouds (and the location of cloud top, in general) are not easily inferred from radar observations at this frequency. Therefore we use two methods to determine at which altitude to compare the RWP data with MISR cloud velocities: (1) We use ARM colocated cloud radar and lidar observations and (2) we use cloud top as estimated by the MISR stereo-height algorithm.

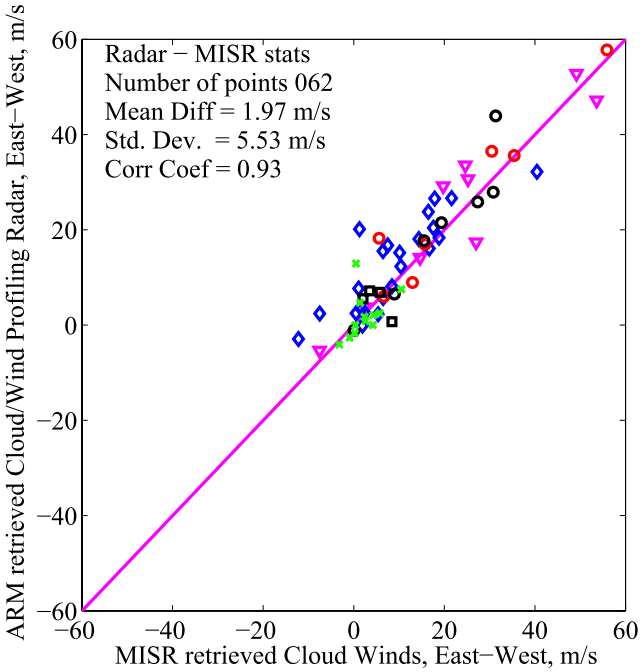

Figure 2. Scatterplot of MISR retrieved eastward cloud velocity against 404 MHz RWP-derived 1 hour consensus wind velocity at cloud top. The cloud top is determined from collocated millimeter-wavelength cloud radar or cloud lidar. Each symbol represents one overpass of the ARM SGP site, and the different symbols represent different types of clouds (see Table 3).

[16] Figure 2 shows a scatterplot of the MISR retrieved eastward cloud velocity against the 404 MHz RWP-derived eastward wind velocity at cloud top at the ARM SGP site. In Figure 2, the cloud top is determined from collocated millimeter-wavelength cloud radar and cloud lidar. As discussed in section 2, MISR retrieves up to two velocities in each 70.4 × 70.4 km domain. The value plotted here is the value chosen by the stereo-height retrieval algorithm for the observed cloud. Each symbol represents one overpass of the ARM site. The different symbols represent different types of clouds, as will be explained later. Figure 2 indicates that there is a strong correlation between the MISR-retrieved and the RWP values for the eastward component of the wind with a mean bias of less than 2 m/s and a standard deviation of about 6 m/s.

[17] Figure 3 shows a similar plot to Figure 1, except for the northerly component of the velocity. Figure 3 shows only a weak correlation between the MISR-retrieved and RWP values for the northward component. While the bias remains small, the standard deviation is nearly double. It is likely that the difference in performance between the two velocity components is a consequence of parallax and the MISR orbit. MISR is in a sun-synchronous orbit and moves overhead from north to south (on the daylight side of the Earth). Therefore parallax (see Figure 1) induces little effect on the observed change in the east to west position of cloud features, but a large effect on the north to south position. Figures 2 and 3 demonstrate that the much larger parallax in the north to south direction substantially reduces the quality of the retrieved cloud winds in this direction.

[18] All the data in Figures 2 and 3 are from version 11, 12 or 13 (file format F07) of the MISR cloud stereo retrieval algorithm. These versions contain a wind quality assurance (QA) flag which separates good wind retrievals (called best winds in the context of the MISR stereo-height retrieval as described in section 2) from the set of all wind retrievals (sometimes called raw winds). There are a total of 62 cases in Figures 2 and 3 which are all of the available cases between March 2000 and April 2004 where there is (1) a MISR overpass with a cloud over the ARM site, (2) any MISR cloud wind retrieval (raw winds), (3) a good RWP consensus wind profile, and (4) good cloud radar/lidar boundaries. The requirement of a good consensus wind speed from the RWP reduces the possible data set by about 50%. Restricting the comparison to include only the MISR best winds cases removes about 20 points from the scatterplots (not shown). This does not significantly change the comparison of the eastward velocity component, but does improve the correlation coefficient and reduces the standard deviation of the northward velocity component from about 13 to 9 m/s, as underscored in Table 1.

[19] Finally, the RWP data used in Figures 2 and 3 is based on winds at cloud top, where cloud top was obtained from collocated ARM cloud radar and lidar. We also compared wind velocities using the MISR stereo-height. The scatterplots (not shown) are characteristically similar to Figures 2 and 3. Table 1 lists the mean differences, standard deviations and correlation coefficients. The standard deviations are a little larger and correlation a little weaker when using the MISR retrieved heights, but not much. These results support expanding the comparison of MISR cloud

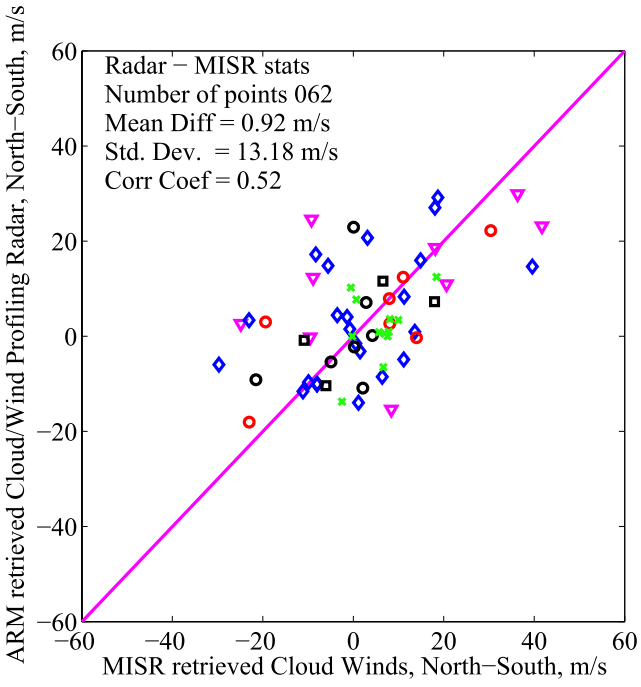

Figure 3. As Figure 1 except for the northward cloud velocity.

Table 1. Difference Between MISR Retrieved Cloud Winds and RWP Retrieved Wind Speed at Cloud Top

	Number of Points	Mean Difference, m/s	Standard Deviation, m/s	Correlation Coefficient
Radar Winds at Cloud Top From Colocated Millimeter-Wavelength Cloud Radar				
MISR wind QA flag indicates good retrieval				
East-west	042	1.68	5.85	0.91
North-south	042	1.04	10.12	0.66
MISR wind QA flag ignored				
East-west	062	1.97	5.53	0.93
North-south	062	0.92	13.18	0.52
Radar Winds at Cloud Top as Retrieved by MISR Stereo Height Algorithms				
MISR wind QA flag indicates good retrieval (calculated using best winds stereo-height)				
East-west	043	0.27	7.64	0.84
North-south	043	1.69	9.37	0.68
MISR wind QA flag ignored (at without winds stereo-height)				
East-west	066	1.78	6.64	0.89
North-south	066	0.65	13.01	0.51

winds with RWP data sets to those locations with RWP but without collocated cloud radar and lidar data.

4. Comparison of MISR Stereo-Height With ARM Cloud Radar/Lidar Data

[20] In this section we examine four aspects of the MISR stereo-height without winds and best winds retrievals: (1) coverage, (2) differences between MISR stereo-heights and cloud top height from ARM cloud radar and lidar, (3) the effect of the MISR patch size used to match up the MISR spatial data with ARM temporal data, and (4) the optical depth limit at which the MISR stereo-height algorithm fails to find the highest cloud top.

4.1. Coverage

[21] A summary of the available number of MISR overpasses for the ARM SGP, Barrow, and Nauru sites is listed in Table 2. Only those overpasses where the cosine of the solar zenith angle is 0.2 or greater are considered. Table 2 shows that there are many more overpasses of the ARM site in Barrow, Alaska (71.3°N) than the ARM site at Nauru Island (0.53°S). This is a consequence of the satellite polar orbit. MISR views the Barrow site 6 times over each orbital repeat cycle of 16 days, while it views the Nauru site 2 times each 16 days and the southern great plans (SGP) site 3 times each 16 days.

[22] Table 2 also list the number (and percentage) of these overpasses where the MISR stereo height retrieval was successfully applied. Each of these table entries specifies a patch size (e.g., 11 × 11 pixels) indicating the number of pixels (centered on the ARM site) used to obtain a median stereo−height. We require at least half the pixels in the patch to have a retrieved value for the retrieval to be consider successful over the patch. We will examine the impact of the patch size in section 4.3. Table 2 shows that the MISR without winds retrieval has good coverage, varying from about 78% at Barrow to 96% at SGP. The difference between the coverage at these two sites is not a result of random sampling. The stereo retrieval works by identifying common features at several view angles (see description in section 2). At SGP, the surface has many

Table 2. Summary of Available MISR Overpasses for Three ARM Sites From March 2000 Through (at Least) April 2004[a]

Total Overpasses	Nauru	SGP	Barrow
With MISR stereo data (versions 8 thru 13)	180	277	375
With best winds 11 × 11 pixels	71	164	155
Percentage of total	39.7%	58.0%	35.8%
Without winds 11 × 11 pixels	152	272	337
Percentage of total	84.9%	96.1%	77.8%
Without winds 5 × 5 pixels	170	265	353
Percentage of total	95.0%	93.6%	81.5%
Without winds 3 × 3 pixels	171	260	361
Percentage of total	95.5%	91.9%	83.4%
Without winds 1 pixel	160	255	360
Percentage of total	89.4%	90.1%	83.1%
Good ARM cloud boundaries	65	227	273
Percentage of total	36.1%	81.9%	72.8%
Clear-sky cases	10	93	57
Percentage of cases with good ARM boundaries	15.4%	41.0%	20.9%

[a]The patch size (e.g., 11 × 11 pixels) indicates the number of pixels (centered on the ARM site) used to obtain a median stereo-height. We also require at least half the pixels in the patch have a successful retrieval value.

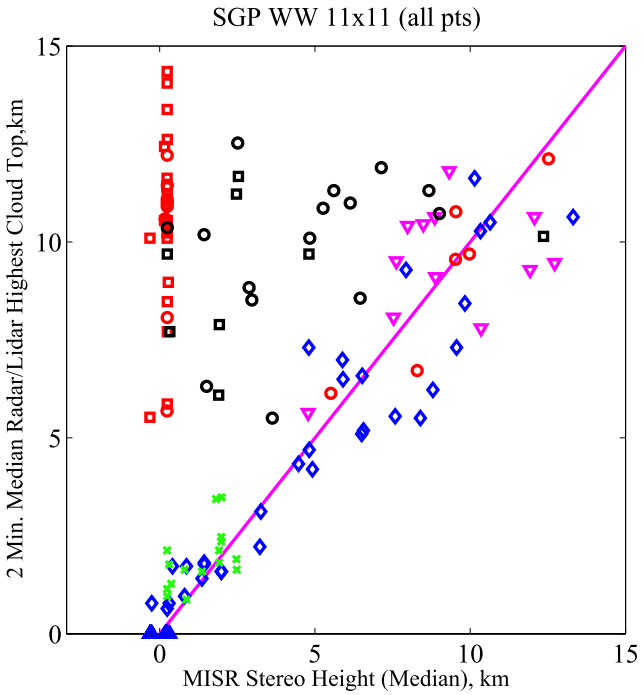

Figure 4. Scatterplot of ARM Radar/Lidar derived cloud top (median value for 2 min surrounding satellite overpass) against MISR stereo-derived without winds cloud top (median value for 11x11 pixel patch centered on ground site).

visible features such that the only time the retrieval tends to fail to return a result (that is no stereo-match) is when there is an exceptionally bright and homogeneous cloud over the site or to a lesser extent around the edges of broken clouds, especially those with extended shadows. At Barrow, however, the surface is frequently bright and appears nearly featureless (at 275 m), especially after new snow falls. This results in a larger proportion of no-retrievals under clear-sky and thin cloud conditions. At Nauru, we similarly find a lower percentage of successful retrievals, because the ocean surface often appears featureless (at 275 m). We provide examples that demonstrate these effects in section 5. While the MISR without winds retrieval has reasonably good coverage, the same is not true of the best winds retrieval. The best winds coverage varied from about 36% coverage at Barrow to 58% at SGP.

[23] Table 2 also lists the number (and percentage) of overpasses where ARM radar or lidar data was collected and judged to be of sufficient quality to make a good estimate of cloud top. While the ARM program strives to operate their cloud radar and lidar instruments continuously, they have had difficult in achieving this goal especially for their lidar systems and especially at the remote Nauru island site.

4.2. Height Difference Statistics

[24] Figure 4 shows a comparison of MISR and ARM (radar and lidar) retrieved cloud top heights for the SGP ARM site. In Figure 4 each symbol represents one overpass of the ARM site where we have divided the overpasses into a number of categories, as listed in Table 3. These categories are based on a visual inspection of the MISR nadir imagery and ARM radar and lidar measurements. Figure 4 shows that the MISR retrieval works well for many of the overpasses, and in particular for those represented by the blue, magenta and green symbols. These symbols represent visually opaque clouds. There are, however, a large number of overpasses where the ARM cloud tops are much higher than the MISR cloud tops. These overpasses are often denoted by red and black symbols indicating that these clouds appear thin, by visual inspection. In the case of the red symbols, the pattern matcher used by the MISR stereo algorithm tends to identify features at the surface (rather than the location of the thin cloud) while in the case of the black symbols the algorithm identifies lower cloud decks. Many satellite retrievals have problems with thin clouds, and an important consideration in using the MISR stereo retrieval is to know how thick a cloud must be in order to be retrieved. We will return to this topic in section 4.4.

[25] Information on the number of cases, as well as the mean, median, and standard deviation of the differences between the ARM and MISR retrievals is given in Table 4. Table 4 shows that we found a total of 216 overpasses where MISR retrieved a without winds stereo-height and we had sufficiently good cloud radar and lidar data at the ARM site to make a determination of cloud top (if any). Of these 216 cases, MISR detected a cloud in 93 cases or about 43% of the time and of the remaining 123 cases there was a cloud too thin for MISR to detect on 30 occasions or about 14% of the time (most of the red symbols in Figure 4).

Table 3. Simple Classification of MISR Overpasses Based on a Visual Inspection of Composite Color MISR Nadir Imagery (i.e., True Color Image Using the MISR Red, Green, and Blue Channels)

Symbol	Description
Blue triangle	Clear-sky (retrieved height should be at ground level).
Green crosses	Broken boundary layer clouds.
Blue diamonds	Opaque clouds with well defined structure visible at cloud top. One cannot see surface or lower clouds. If a boundary layer cloud, the cloud must be stratiform.
Red symbols	Thin cloud over surface. Squares mean NO hint of a cloud is visible in nadir imagery. Circles mean one can see some cloud "whitening" in nadir imagery, but also clearly make out the surface (or other lower cloud decks).
Black symbols	Thin cloud over other clouds.
Magenta symbols	Opaque to mostly opaque clouds with diffuse clouds top. One can at best marginally identify underlying surface features or lower clouds and there are no well defined structures visible at cloud top in nadir imagery.

Table 4. Difference Statistics (ARM Cloud Top − MISR Stereo Height) for the ARM SGP Site

Height Difference Statistics (ARM − MISR)	Number of Cases	Percent	Mean, m	Median, m	Standard Deviation, m
SGP without winds 11 × 11					
Distinct cloud/blue[a] and green symbols	051 of 216	23.61%	216	92	1224
Diffuse cloud/magenta symbols	013 of 216	6.02%	557	220	2031
Thin cloud/red circles[b]	006 of 216	2.78%	−126	−61	958
Lower cloud/black symbols	023 of 216	10.65%	302	189	1154
All MISR detected clouds	093 of 216	43.06%	229	124	1314
Clear/blue triangles	093 of 216	43.06%	−240[c]	−204[c]	130
Surface under cloud/other red symbols	030 of 216	13.89%	−242[c]	−199[c]	142
SGP best winds 11 × 11 (BW pts)					
Distinct cloud/blue[a] and green symbols	028 of 134	20.90%	709[c]	519[c]	893
Diffuse cloud/magenta symbols	006 of 134	4.48%	1571	1274[c]	1330
Thin cloud/red circles[b]	003 of 134	2.24%	−172	42	442
Lower cloud/black symbols	006 of 134	4.48%	352	448	508
All MISR detected clouds	043 of 134	32.09%	551[c]	581[c]	926
Clear/blue triangles	073 of 134	54.48%	−244[c]	−126[c]	380
Surface under cloud/other red symbols	018 of 134	13.43%	184[c]	413[c]	867
SGP without winds 11 × 11 (subset including only cases with best winds retrieval) (BW pts)					
Distinct cloud/blue[a] and green symbols	028 of 134	20.90%	195	73	1154
Diffuse cloud/magenta symbols	006 of 134	4.48%	548	−341	2086
Thin cloud/red circles[b]	003 of 134	2.24%	626	96	1474
Lower cloud/black symbols	006 of 134	4.48%	263	22	748
All MISR detected clouds	043 of 134	32.09%	302	10	1251
Clear/blue triangles	073 of 134	54.48%	−240[c]	−202[c]	133
Surface under cloud/other red symbols	018 of 134	13.43%	−242[c]	−200[c]	133

[a]Blue diamonds.
[b]Only those clouds where MISR does locate cloud top.
[c]Mean or median bias that is statistically different from 0 at the 95% confidence limit.

[26] Considering all 93 MISR detected clouds as a group, we find that the without winds retrieval shows a small positive bias with respect to the ARM data and a standard deviation of about 1300 m, at SGP. Treating each overpass as an independent sample, we used the well known statistical t-test or sign-test to determine if the mean or median difference, respectively, is significantly different from zero (that is, likely represent a statistically significant bias). In Table 4 (as well as in Tables 5 and 6 for the other ARM sites) we mark those means or medians which indicate a statistically significant bias at the 95% confidence level using a footnote. At the 95% confidence level, none of the without winds cloud top heights show a statistically significant bias.

[27] Under clear skies, the mean and median difference between the MISR stereo-height and the true-surface elevation, however, are found to be statistically significant. This may seem odd at first glance, but the MISR stereo-height algorithm has an expected resolution of 1 pixel, which is equivalent to a height offset of about 560 m [*Moroney et al.*, 2002]. A careful inspection of the MISR stereo-heights in Figure 4 shows this quantization, particularly for the clear sky overpasses (blue diamonds). As a result of the quantization and fixed MISR data grid, the retrieved surface height will tend to have a fixed offset from the true surface height at any point on the globe. This offset should be no more than half the height quantization or about 280 m (as first shown by *Muller et al.* [2002]).

[28] Of the 93 cloudy without winds cases at SGP, 43 of them also have a best winds retrieval. Table 4 shows that the best winds retrieval has a standard deviation of less than 1000 m. Application of the well-known statistical *F* test for the difference in standard deviations shows that the reduction from 1300 m from the without winds retrieval to less than 1000 m is statistically significant at the 95% confidence level.

[29] Examining the various cloud categories listed in Table 4 shows that those clouds with visually diffuse cloud tops (magenta symbols) have both a large observed bias and large standard deviation relative to other cloud types. It appears that contrast for many clouds with diffuse tops is not generated at cloud top, but within the cloud (that is below millimeter-wavelength radar-derived cloud top). Later in this section we will investigate the optical depth level at which the MISR stereo-height tends to locate cloud top for these clouds. We speculate that the large standard deviation is due to both variations in the actual in-cloud contrast level and poor wind retrievals for these cases, few of which tend to have best winds retrievals.

[30] The Table 4 best winds result also indicates that the best winds retrieval shows what appears to be a large (greater than 500 m) statistically significant bias for opaque clouds with distinct cloud top features. The bias at the other two ARM sites (Tables 5 and 6) are both estimated to be less than 200 m for this same cloud type. The large bias at SGP does not appear to be a problem with the MISR retrieval, rather it appears to be a problem with the ARM radar data as most of the bias can be attributed to cases of boundary layer clouds that occur during the summer time, a period when the ARM radar data is known to be contaminated by clutter from insects [*Clothiaux et al.*, 2000].

[31] Finally, the bottom of Table 4 gives the height difference statistics for the without winds retrieval, but limiting the cases considered to those points that also have a best winds result. These statistics clearly show that the improvement between the without winds and best winds

Table 5. Same as Table 4 Except for Barrow, Alaska, Site

Height Difference Statistics (ARM − MISR)	Number of Cases	Percent	Mean, m	Median, m	Standard Deviation, m
Barrow without winds 11 × 11					
Distinct cloud/blue[a] and green symbols	085 of 217	39.17%	168	50	643
Diffuse cloud/magenta symbols	021 of 217	9.68%	−203	−18	1201
Thin cloud/red circles[b]	004 of 217	1.84%	−1674	−1425	511
Lower cloud/black symbols	040 of 217	18.43%	196	112	867
All MISR detected clouds	150 of 217	69.12%	158[c]	68	802
Clear/blue triangles	033 of 217	15.21%	8	−21	282
Surface under cloud/other red symbols	034 of 217	15.67%	8	−255[c]	483
Barrow best winds 11 × 11					
Distinct cloud/blue[a] and green symbols	056 of 102	54.90%	139	69	699
Diffuse cloud/magenta symbols	008 of 102	7.84%	990	966	1478
Thin cloud/red circles[b]	000 of 102	0.00%	0	0	0
Lower cloud/black symbols	017 of 102	16.67%	242	−106	1041
All MISR detected clouds	081 of 102	79.41%	187[c]	121	910
Clear/blue triangles	015 of 102	14.71%	−12	−80	330
Surface under cloud/other red symbols	006 of 102	5.88%	61	−124	597
Barrow without winds 11 × 11 (subset including only cases with best winds retrieval)					
Distinct cloud/blue[a] and green symbols	056 of 102	54.90%	147	44	656
Diffuse cloud/magenta symbols	008 of 102	7.84%	427	151	945
Thin cloud/red circles[b]	001 of 102	0.98%	−1695	−1695	0
Lower cloud/black symbols	017 of 102	16.67%	−334	−392	753
All MISR detected clouds	082 of 102	80.39%	136	−24	727
Clear/blue triangles	015 of 102	14.71%	8	−19	156
Surface under cloud/other red symbols	005 of 102	4.90%	8	−109	242

[a]Blue diamonds.
[b]Only those clouds where MISR does locate cloud top.
[c]Mean or median bias that is statistically different from 0 at the 95% confidence limit.

results is not an artifact of the reduced coverage of the best winds retrieval.

[32] Tables 5 and 6 present the same difference statistics for the ARM Barrow and Nauru sites as given in Table 4 for the SGP site. The performance of the MISR stereo retrieval is similar at the Barrow and Nauru sites to the SGP site with respect to cloud type. Overall best winds results are also similar at Barrow and Nauru to SGP with a standard deviation in ARM to MISR cloud height of less than 1000 m. The most noteworthy differences in the perfor-

Table 6. Same as Table 4 Except for Nauru Island Site

Height Difference Statistics (ARM − MISR)	Number of Cases	Percent	Mean, m	Median, m	Standard Deviation, m
Nauru without winds 11 × 11					
Distinct cloud/blue[a]	025 of 058	43.10%	44	30	751
Diffuse cloud/magenta symbols	002 of 058	3.45%	1582	1582	137
Thin cloud/red circles[b]	004 of 058	6.90%	−153	45	1312
Lower cloud/black symbols	013 of 058	22.41%	−85	6	416
All MISR detected clouds	044 of 058	75.86%	23	95	769
Clear/blue triangles	007 of 058	12.07%	−1036[c]	−862[c]	316
Surface under cloud/other red symbols	007 of 058	12.07%	−997[c]	−892[c]	559
Nauru best winds 11 × 11					
Distinct cloud/blue[a] and green symbols	008 of 028	28.57%	212[c]	317	499
Diffuse cloud/magenta symbols	001 of 028	3.57%	1043	1043	0
Thin cloud/red circles[b]	003 of 028	10.71%	201	−117	669
Lower cloud/black symbols	006 of 028	21.43%	156	63	1496
All MISR detected clouds	018 of 028	64.29%	216	200	941
Clear/blue triangles	005 of 028	17.86%	−50	−209	304
Surface under cloud/other red symbols	005 of 028	17.86%	−600	−872	1126
Nauru without winds 11 × 11 (subset including only cases with best winds retrieval)					
Distinct cloud/blue[a] and green symbols	008 of 028	28.57%	−72	69	429
Diffuse cloud/magenta symbols	001 of 028	3.57%	1679	1679	0
Thin cloud/red circles[b]	003 of 028	10.71%	12	166	1579
Lower cloud/black symbols	006 of 028	21.43%	−51	43	377
All MISR detected clouds	018 of 028	64.29%	15	166	745
Clear/blue triangles	005 of 028	17.86%	−554	−746	301
Surface under cloud/other red symbols	005 of 028	17.86%	−554	−772	637

[a]Blue diamonds.
[b]Only those clouds where MISR does locate cloud top.
[c]Mean or median bias that is statistically different from 0 at the 95% confidence limit.

Table 7. Same as Table 4 Except for Differing Patch Sizes

Height Difference Statistics (ARM − MISR)	Number of Cases	Percent	Mean, m	Median, m	Standard Deviation, m
SGP without winds 31 × 31 pixels					
Distinct cloud/blue[a] and green symbols	048 of 210	22.86%	111	−26	1268
Diffuse cloud/magenta symbols	012 of 210	5.71%	699	401	2120
Thin cloud/red circles[b]	006 of 210	2.86%	−2	−11	1210
Lower cloud/black symbols	023 of 210	10.95%	222	22	1202
All MISR detected clouds	089 of 210	42.38%	160	45	1373
SGP without winds 21 × 21 pixels					
Distinct cloud/blue[a] and green symbols	048 of 210	22.86%	129	−33	1234
Diffuse cloud/magenta symbols	012 of 210	5.71%	797	536	2230
Thin cloud/red circles[b]	007 of 210	3.33%	248	279	1282
Lower cloud/black symbols	023 of 210	10.95%	−163	180	1396
All MISR detected clouds	090 of 210	42.86%	219	122	1432
SGP without winds 5 × 5 pixels					
Distinct cloud/blue[a] and green symbols	051 of 209	24.40%	144	91	1337
Diffuse cloud/magenta symbols	009 of 209	4.31%	479	−113	1493
Thin cloud/red circles[b]	004 of 209	1.91%	−136	145	916
Lower cloud/black symbols	023 of 209	11.00%	273	−246	1937
All MISR detected clouds	087 of 209	41.63%	160	−17	1503
SGP without winds 3 × 3 pixels					
Distinct cloud/blue[a] and green symbols	050 of 208	24.04%	152	95	1351
Diffuse cloud/magenta symbols	011 of 208	5.29%	896	286	1653
Thin cloud/red circles[b]	005 of 208	2.40%	30	590	1274
Lower cloud/black symbols	022 of 208	10.58%	284	−235	1982
All MISR detected clouds	088 of 208	42.31%	252	64	1552
SGP without winds 1 pixel					
Distinct cloud/blue[a] and green symbols	048 of 201	23.88%	270	200	1666
Diffuse cloud/magenta symbols	011 of 201	5.47%	421	−408	3041
Thin cloud/red circles[b]	005 of 201	2.49%	−104	−149	381
Lower cloud/black symbols	021 of 201	10.45%	−390	−552	2922
All MISR detected clouds	085 of 201	42.29%	170	−85	2193

[a]Blue diamonds.
[b]Only those clouds where MISR does locate cloud top.

mance of the MISR stereo height retrieval at the Barrow and Nauru sites relative to the SGP site are (1) the coverage is higher at SGP (as discussed previously) and (2) the without winds results are on par with the best winds results at both Barrow and Nauru. It is not immediately apparent why the without winds results are better at Barrow and Nauru than at SGP, but we speculate that clouds are less likely to have a strong northward velocity component at either the Barrow and Nauru sites as at SGP.

4.3. Effect of Patch Size on Difference Statistics

[33] The height difference statistics presented in section 4.2 are based on the median value of the MISR stereo-height retrievals for an 11 × 11 pixel patch (or equivalently about a 12 × 12 km area) centered on the ARM sites. We also require that at least half the pixels in the patch have a retrieved height. This median filter removes the effect of many speckle-like retrieval blunders (some examples of which can be found later in section 5). These blunders are caused by the stereo pattern matcher misidentifying the same feature in two images. While the occurrence of blunders has some spatially correlation (e.g., they tend to be concentrated around the edges of clouds, regions with low contrast, or regions with sun glint contamination), their effect seems to be easily minimized using a simple median filter. Table 7 shows the effect of using a patch size of 31 × 31, 21 × 21, 5 × 5, 3 × 3 and 1 × 1 pixels in the without winds retrieval at the SGP site. As one increases the patch size, the standard deviation (in the difference between the MISR and ARM heights) decreases from more than 2000 m for a single pixel to about 1300 m for the 11 × 11 patch. Increasing the patch size above 11 × 11 pixels does not make any significant improvement. At the 11 × 11 pixel scale (or larger), the error in the MISR stereo-height retrievals are dominated by errors in estimating the cloud velocity. Results of the best winds retrieval (not shown) follow this same trend, as do retrievals at Barrow and Nauru (also not shown). This analysis suggests that for many applications, the MISR stereo height retrieval should be filtered to reduce the speckle-like blunder errors and that a median filter using 11 × 11 pixels is reasonably effective.

4.4. Optical Depth Detection Limit/Detecting Cloud Top With Stereo

[34] In section 4.2, we found that the MISR stereo height algorithm was able to make an estimate of cloud top height much of the time, but that some clouds appeared too thin (based on visual inspection of the MISR nadir imagery) for the stereo-matcher to find. The "thin cloud over surface" category (where the MISR stereo-height retrieval returns the height of the surface rather than that of the overlying clouds) accounts for 12 to 16% of the total overpasses at each of the ARM sites.

[35] Figure 5 (top left) shows a plot of the total cloud optical depth (retrieved via ARM observations) for those clouds where MISR does retrieve a cloud top height. (In the case of the black symbols, the cloud top is not that of the highest layer but of a lower cloud deck, though not necessarily the second layer if there is more than one high thin cloud layer). Figure 5 (top right) shows a similar plot

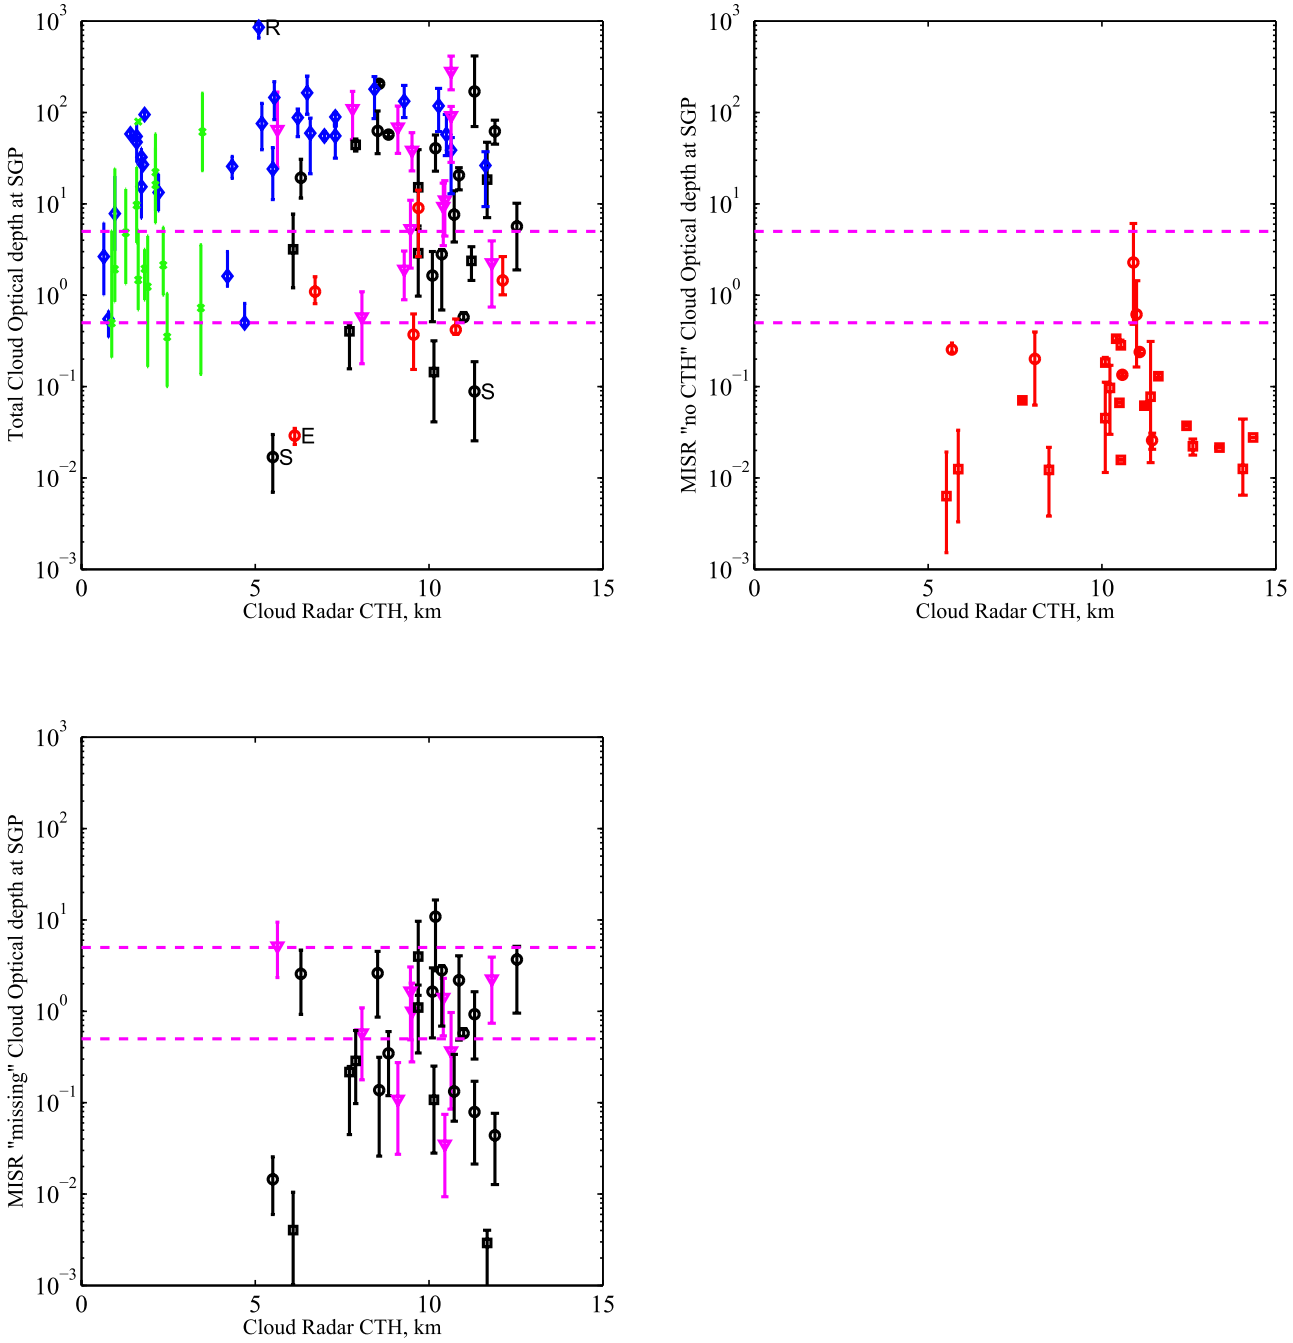

Figure 5. Scatterplot showing ground-based optical depth (vertical axis) against retrieved cloud top height (horizontal axis). Error bars denote the estimated uncertainty in each retrieval. The letter "R" denotes rain contaminated, "E" denotes a cloud edge and "S" denotes a case with supercooled liquid water, for which the optical depth retrieval may be poor as explained in the text. The dashed lines indicate optical depth thresholds of 0.5 and 5.0. (top left) Total optical depth for clouds with a successful MISR stereo retrieval. (top right) Total optical depth of those clouds NOT detected by the MISR stereo retrieval. (bottom left) Optical depth between MISR retrieved cloud top and cloud top derived from ARM radar/lidar.

but for those clouds where MISR does not find any cloud top, and instead retrieved the surface height through the clouds. Error bars show the estimated uncertainty in the retrieved cloud optical depth (Appendix A). The size of the error bars differ from case to case depending on the type of retrieval. For example, lidar-based estimates where the cloud is penetrated by the lidar have relatively little uncertainty whereas those in broken cloud fields or those based solely on radar-reflectivity observations tend to have large uncertainties. A few of the cases are labeled by the letters

"R," "E," or "S." Those labeled "R" are cases where the optical depth retrieval is contaminated by rain, such that the retrieved optical depth is likely biased much too high. Those labeled "E" are retrievals that occur at the edge of optically thicker cloud where the thick part of the cloud does not appear to have passed directly over the ARM vertically pointing radar and lidar systems. In these edge cases, the MISR stereo algorithm is assigning the correct height on the basis of the contrast generated by the cloud edge and the optical depth retrievals, while correct, do not characterize the optical depth of the thicker portion of the cloud. We will discuss the MISR stereo-height retrieval spatial resolution further in section 5. Finally, those points labeled "S," appear to be cases where the cloud contains supercooled water. These clouds often contain only a small amount of liquid water such that the microwave radiometer used in the retrieval processes cannot accurately determine the amount of water present [*Marchand et al.*, 2003] and yet the contribution of the liquid water to the total cloud optical depth is substantial or even dominates the total optical depth. The optical depth estimated by our ground-based retrievals of these "S" clouds is generally biased low by much more than our automated algorithm uncertainty interval indicates.

[36] The results in Figure 5 (top) show that the MISR stereo-height retrieval detects most clouds with an optical depth greater than about 0.3 to 0.5 (as indicated by the lower horizontal dashed line). For the SGP, we see that all of the clouds with cloud top below 5 km and most of the clouds with cloud tops below 7 km are observed to have distinct features at cloud top (blue and green symbols). The MISR stereo-retrieval does well with these clouds. Most of the clouds with cloud top above 7 km are observed to be either multilayered where MISR stereo frequently detects a lower cloud deck (black symbols) or have diffuse cloud tops (magenta symbols).

[37] From the ARM retrievals we can estimate the optical depth of the upper clouds through which the MISR stereo-height retrieval was able to detect a lower cloud. For these cases (black symbols) MISR is returning a true cloud top but missed the highest cloud deck. This occurred in about 10% to 20% percent of the cases studies, depending on the site (see Tables 4–6). The thickness of the missed upper clouds is shown in Figure 5 (bottom) (black symbols). In a similar fashion we can estimate the amount of missing optical depth for diffuse clouds, magenta symbols. (Recall that in section 4.2 we found that the MISR stereo-height was biased below ARM highest cloud top for optically opaque clouds with diffuse cloud tops.) For both the cloud types represented by the black and magenta symbols, we estimate that the missing optical depth can be as large as 10 (worst case, possibly a poor optical depth retrieval) and usually in the range from 0.5 to 5.

[38] Figures 6 and 7 show a similar set of plots for the ARM Nauru and Barrow sites, respectively. As discussed in section 4.1, there are considerably fewer overpasses of the Nauru site than the other sites. Overall, Figure 6 shows that the performance of the MISR stereo-height at Nauru appears similar to that of SGP. However, Figure 7 shows that the Barrow site is different. We find many more cases where clouds with optical depths greater than 0.5 are being seen through by the MISR stereo-height algorithm (red symbols, Figure 7, top right). Also unlike the SGP and Nauru sites where these thin clouds only occurred above 5 km, the thin cloud at Barrow have cloud top altitudes ranging from 1 to 11 km, throughout the entire troposphere. A detailed examination of individual cases shows that the reduced cloud detection at Barrow is caused by cases with high contrast at the surface generated by a combination of snow or sea ice and open water or dark (snow free) land. One of the more extreme examples of this is shown in section 5. The optical depth threshold at which MISR fails to detect clouds over high contrast surfaces (i.e., mixed snow or ice with dark water or land features) is, not surprisingly, similar to the missing cloud optical depth for multilayer clouds at all three sites.

[39] In Figure 8, we summarize the MISR stereo-height detection limit by plotting the percentage of cases where MISR does detect cloud top as a function of the cloud optical depth. In this plot, the cloud optical depth bins are logarithmic with bin centers at log10(optical depth) = -3, -2.5, -2, . . . 1, 1.5, 2. Figure 8 demonstrates that there is a step transition at the SGP and Nauru sites with most clouds with a total of optical depth 1 or greater detected and few clouds with optical depth 0.1 or less detected. Because of the bright surface a Barrow (for part of the year), the overall detection of clouds with optical depths greater than 0.5 is not as good as it is at SGP or Nauru.

5. Case Studies

[40] In this section we discuss a few examples of the MISR retrieved stereo-height to demonstrate in more detail limitations in the MISR approach. We stress that these cases were selected to represent limitations. Examples and analysis of less problematic cases that highlight the strengths of the MISR approach are given by *Moroney et al.* [2002], *Muller et al.* [2002] and *Diner et al.* [2005].

5.1. Broken Clouds and Clear Skies Over Water

[41] Figure 9 illustrates how the MISR without winds stereo-height retrieval works over two broken cumulus cloud fields. Figure 9 (top) shows a scene composed entirely of fair weather cumulus, which are often observed at the ARM SGP site in the summer. The MISR stereo-height does an excellent job of retrieving the height of the clouds and the surface height in adjacent clear-skies areas. Notice however that the stereo-height retrieval does not separate most of the individual cloud cells. The MISR stereo-height is obtained using a pattern-matcher with a patch size of 6 by 10 pixels (cross-track by along-track). Thus while the MISR stereo-height data set is stored at 1.1 km sampling, the true resolution of the stereo-height field is several times larger than this sampling.

[42] Cloud fraction as measured by satellites is notoriously sensitive to resolution for broken clouds and the true area covered by the clouds in this scene (Figure 9, left) is much smaller than one would estimate by taking the ratio of the 1.1 km pixels with a retrieved stereo-height above the surface to the total number of pixels.

[43] Figure 9 (bottom) shows an example of trade cumulus near the ARM Nauru island site. Trade cumulus are a persistent feature of this oceanic region throughout the year, often showing linear cloud streaks as in this case. In

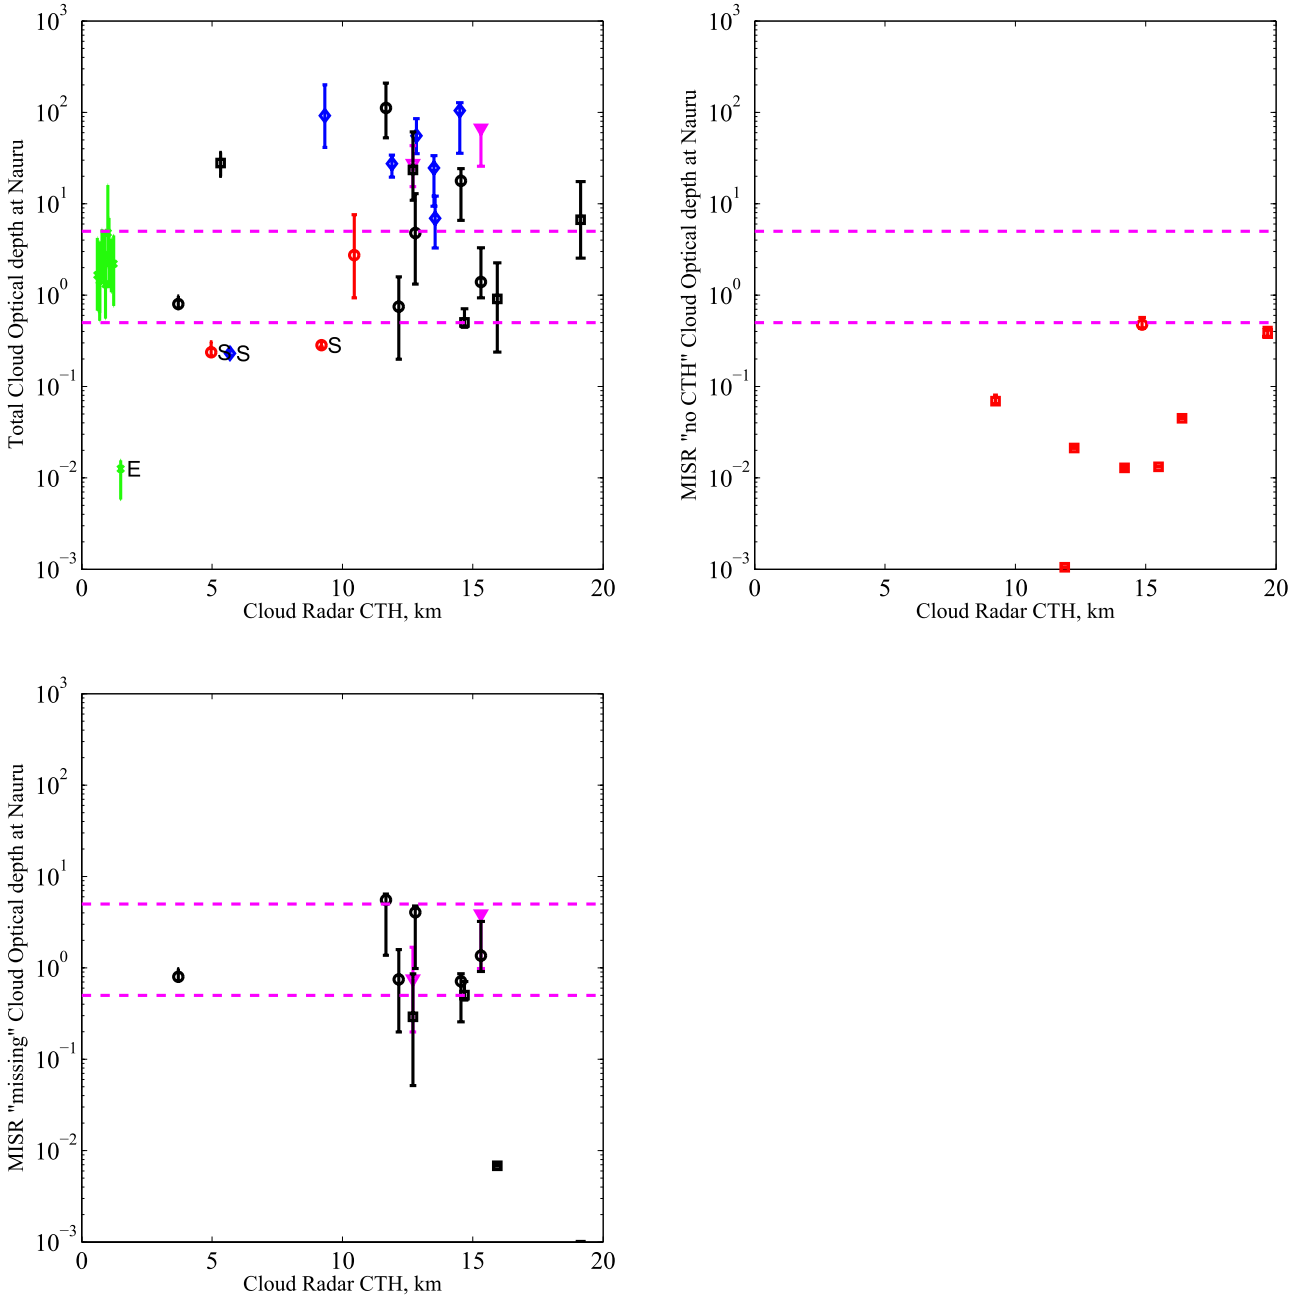

Figure 6. Same as Figure 5 except for Nauru.

particular, we have observed a cloud streak forming off Nauru Island itself for almost every satellite overpass, and this is helpful in gauging the wind directions (from the northeast in this example). The stereo-height retrieval is doing an excellent job of identifying clouds in this scene, but areas devoid of clouds have no features for the pattern-matcher to locate (from one MISR view angle to the next) and so clear sky regions tend to have no retrievals (denoted by white pixels) or blunders. The blunders become more common as sun glint increases (not shown). As with the SGP site, the size of many of the cloud elements is less than the resolution of the stereo-height algorithm, such that the ratio of the pixels with a stereo-height retrieval to the total pixels (or worse to the total number of pixels with a retrieval) will be much larger than the true area covered by clouds.

5.2. Blockiness in Best Winds Result

[44] Several publications [e.g., *Moroney et al.*, 2002; *Muller et al.*, 2002] have noted the MISR stereo-height field sometimes displays a blocky or mottled appearance. This blocky appearance is a direct result of discontinuities in the wind retrieval across adjacent domains (70.4 × 70.4 km regions) where the wind is assumed constant in the retrieval process (see description section 2). Figure 10, for example, shows the MISR best winds stereo-height for the same two

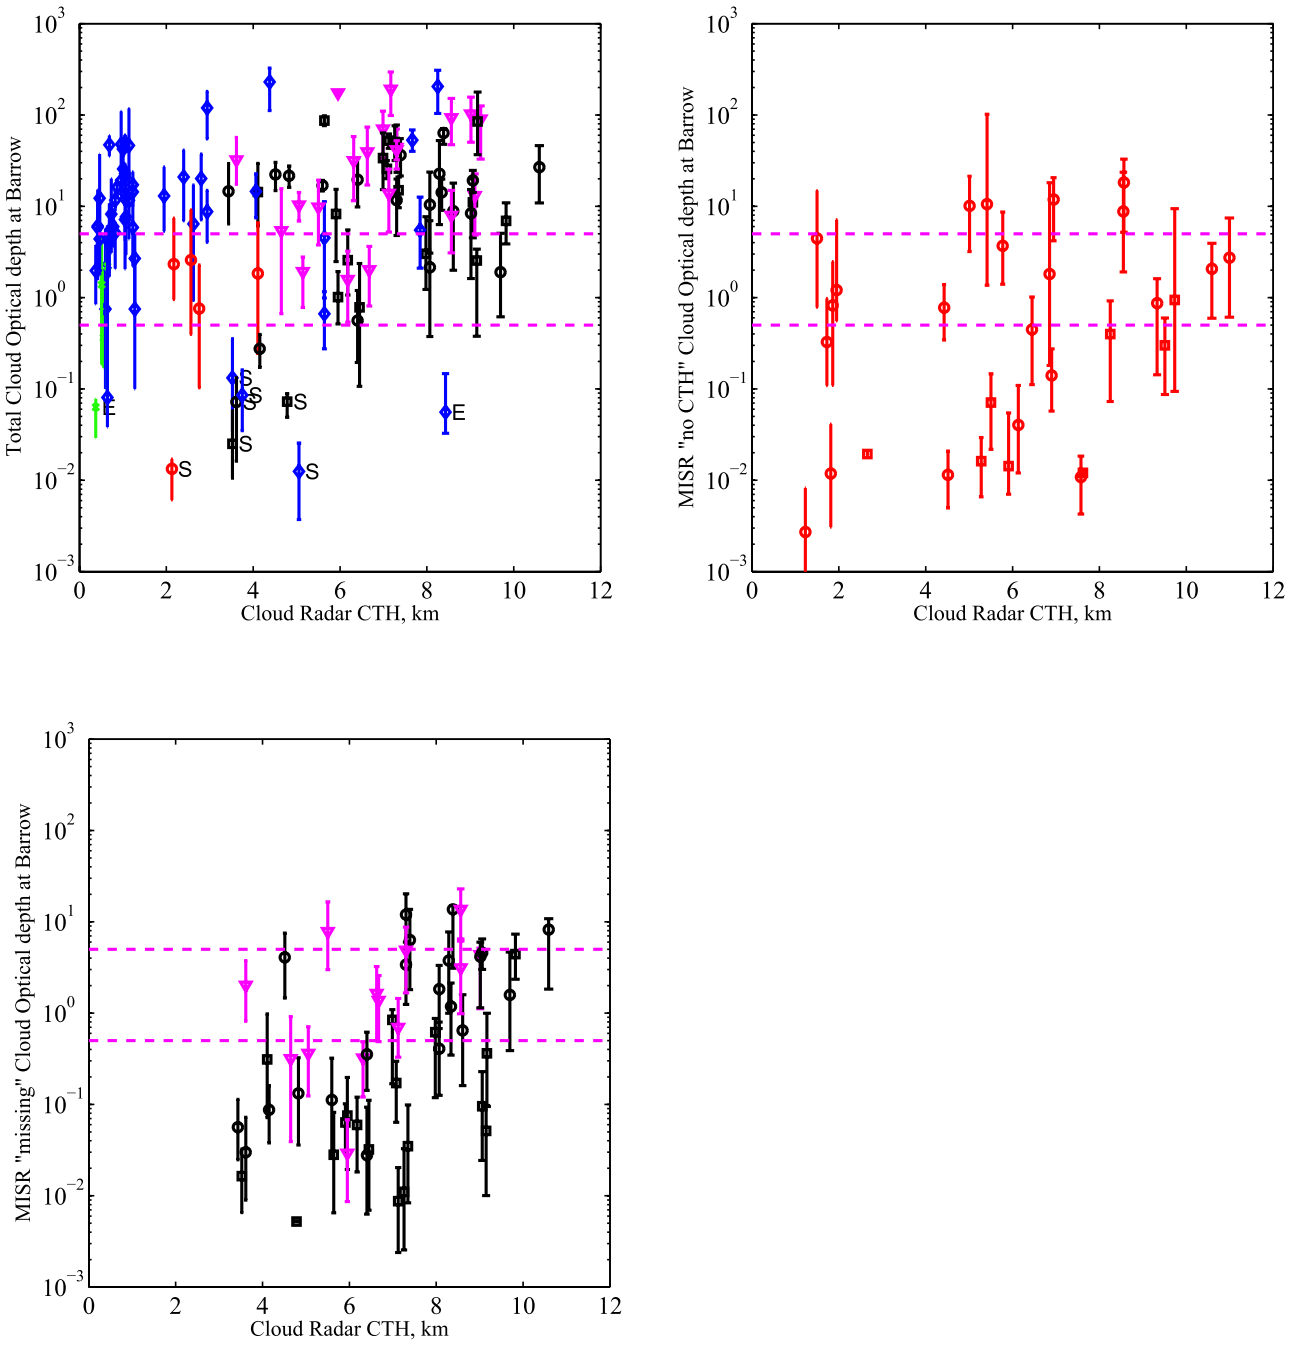

Figure 7. Same as Figure 5 except for Barrow.

scenes as shown in Figure 9. The fair-weather cumulus case at SGP, Figure 10 (top) shows that MISR did not successfully retrieve a wind speed for much of the eastern portion this scene (with the result that there is no best winds retrieval over this region). For the western portion of the scene, including the effect of cloud velocity lowers the retrieved stereo-height of the clouds by approximately 500 to 1000 m. This is equivalent to a shift in the cloud parallax of one or two pixels. For the Nauru case (Figure 10, bottom), the best winds results displays even greater differences between neighboring 70.4 km domains. We note that the without winds stereo-heights for the Nauru case shows cloud top from about 0 to 500 m, which is too low. In each of the four 70.5 km domains surrounding Nauru, including the cloud velocity increases the height from about 500 to 1500 m. Including the cloud motion is correcting the heights, as it should. The difficulty is that the uncertainty in the along track cloud wind retrieval is about 10 m/s (section 3 and Table 2) which in turns makes the correction uncertain to about 1000 m in a discontinuous fashion across domain boundaries.

5.3. High Contrast at the Surface

[45] In section 4.4, we found that the optical depth limit at which the surface height was retrieved by MISR through an overlying cirrus cloud was much higher at Barrow than at

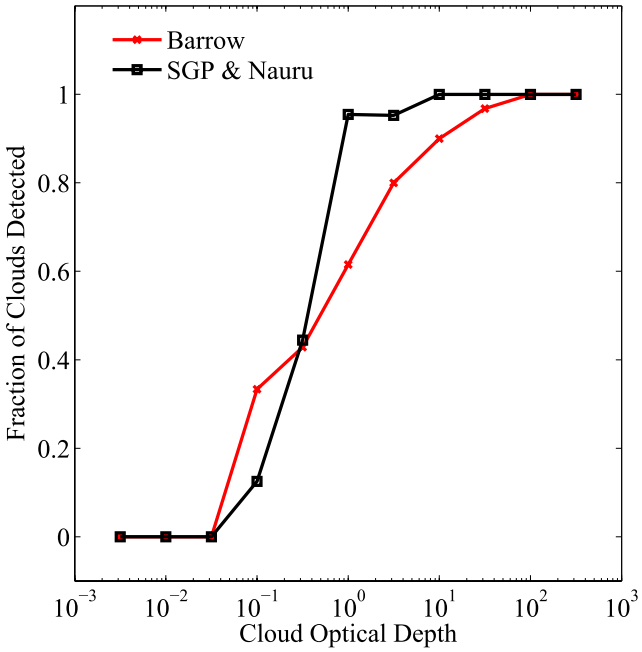

Figure 8. Relative frequency of MISR stereo-height cloud retrieval by cloud optical depth. A value of 1.0 means all clouds (in the optical depth range) have good MISR cloud top heights, and a value 0 means the MISR stereo-height algorithm identified the surface rather than any cloud. These curves are based on the best estimate values shown in Figures 5–7 (top). Rain contaminated, edge and supercooled water cases (previously labeled "R," "E" and "S") are not included because the ground-based optical depth retrievals are not considered reliable for these cases.

the other two ARM sites. An examination of the Barrow cases shows that the problem occurs when there is very high contrast at the surface generated by a combination of snow or sea ice and open water or dark (snow free) land. Figure 11 shows such an example. The overlying cloud is physically thick and is observed to have large values of radar reflectivity (indicative of large particles or large amounts of condensate). Despite the physical size and radar reflectivity, this cloud has an estimated optical depth (from a radar-reflectivity-and-Doppler-velocity retrieval) ranging from only 2 to at most 11 (meaning the large reflectivity is a reflection of particle size not amount of condensate in this case). Our experience with this optical depth retrieval suggests to us that the actual optical depth is probably closer to the low end of the uncertainty range. In any event, the MISR stereo-retrieval returns the surface height or is unable to retrieval any height most of the time.

6. Summary and Discussion

[46] Coverage of the MISR stereo-height retrievals (that is, the percentage of cases in which the MISR retrieval returns what the retrieval algorithm believes is a good quality result) varies with nature of the surface type. The without winds retrieval has very high coverage at the ARM SGP site with a no-retrieval condition (at the 12 km scale) only about 4% of the time, while the best winds has a no retrieval condition almost 40% of the time. We expect similar coverage over most land surfaces. Coverage of the stereo-height retrieval is lower at the Nauru and Barrow ARM sties, primarily because the retrieval does not retrieve the surface height when clear skies over open water or clear skies over homogeneous snow cover are present. The lack of stereo-retrievals in clear skies over open water and over some snow surfaces, in addition to an effective resolution of several kilometers (see section 5.1), suggests that one should be judicious in using the stereo-height field as a "cloud mask" or to calculate cloud fractions. The use of MISR multiangle radiances in combination with stereo-heights over snow and ice to augment tradition multispectral satellite approaches for detecting clouds and retrieving cloud properties is a promising and ongoing area of research.

[47] The MISR stereo-height algorithm has two major sources of error. One is a result of blunders in matching each pixel in MISR nadir view with a matching pixel in one of the 26° views. While the blunders that pass the retrieval algorithm quality tests are not random (for example, they tend to cluster around the edges of clouds and other difficult areas), the blunders are reasonably well filtered by taking the median value on scales of 11 × 11 pixels or larger.

[48] A second major source of error occurs in retrieving the cloud motion, which is a difficult problem. For the MISR geometry, an error of 1 m/s in the estimated along-track cloud velocity introduces an error of approximately 100 m in the cloud stereo-height [*Seiz et al.*, 2006]. Ignoring cloud velocities (i.e., the MISR without winds retrieval) or errors in estimated velocity (in the MISR best winds retrieval) account for most of the difference between MISR retrieved cloud stereo-heights and cloud top height obtained from ARM ground-based radars and lidars. In section 3, we found that the standard deviation in the difference between the MISR cloud wind retrieval and radar wind profiler estimates of cloud top winds at SGP to be about 10 m/s (see Table 1). Not surprisingly in section 4, we then found the standard deviation between the MISR best winds retrieval and ARM cloud tops to about 1000 m (see Table 4).

[49] Comparison of MISR cloud winds with other data sets has been sparse to date. *Horvath and Davies* [2001a] compare MISR clouds winds with GOES track winds for a single case. These authors found that, when GOES and MISR identified the same cloud, the winds were in reasonable agreement but made no quantitative assessment. *Horvath and Davies* [2001b] determined from simulated clouds fields that the error for a single cloud element due to the 275 m MISR pixel size to be about 10 m/s, but also estimated from these simulated cloud fields that for a mesoscale (70.4 km) domain the average along-track and cross-track horizontal wind components should be obtained with an accuracy of 3 to 4 m/s and 1 to 2 m/s, respectively. The comparisons in section 3 between MISR retrieved clouds winds and radar wind profiler data suggest that the MISR wind retrieval error in the along track direction is close to the Horvath and Davies' estimate for a single cloud.

[50] Having examined more than 4 years of MISR overpasses at 3 ARM sites we find that the MISR retrieval is working reasonably well with little observed bias between

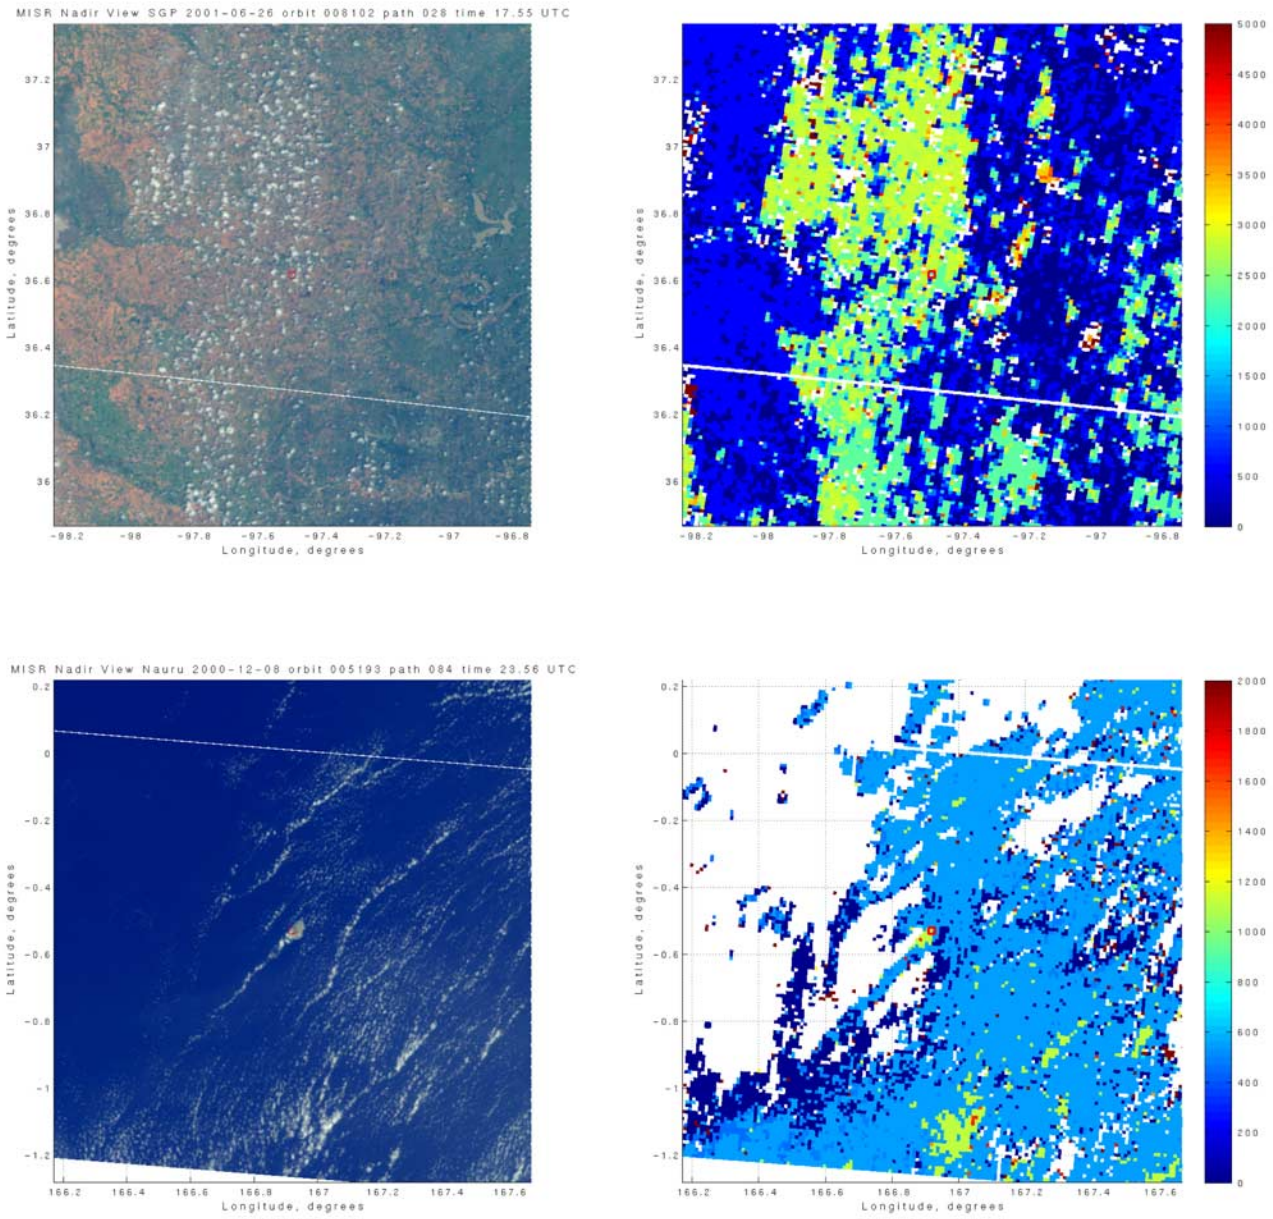

Figure 9. Stereo-height retrieval for two cumulus clouds scenes. (top) Centered on the ARM SGP site and (bottom) set on the ARM site at Nauru Island in the tropical Western Pacific. (left) MISR nadir camera view (color composites formed from MISR red, green, and blue wavelength observations). (right) MISR without winds stereo-height retrieval, in meters. Pixels where the stereo-height algorithm failed to find a stereo-match are shown in white. Clear-sky ocean areas have no features that the stereo-matcher can identify and generally have no retrieval.

the MISR retrievals and ARM cloud radar and lidar boundaries for most cloud types. The standard deviation (across all cloud types) is less than about 1000 m for the MISR best winds retrievals and the standard deviation for the MISR without winds retrieval is less than 1000 m at Barrow and Nauru and about 1300 m at SGP.

[51] The performance is not the same across all cloud types. It is in the nature of the stereo-imaging technique at the heart of the MISR stereo-height retrieval to identify the location of points where contrast in the observed radiance imagery is generated and so the performance of the stereo-heights likewise depends on the scene contrast.

6.1. Optically Thick Clouds (Optical Depth > ~5) With Well-Defined Contrast at Cloud Top

[52] The MISR retrieval works particularly well for optically thick clouds (optical depth greater than about 5) with well-defined contrast at cloud top, as is usually the case for water clouds. In particular, we found no cases of optically thick clouds at SGP with cloud tops below 5 km

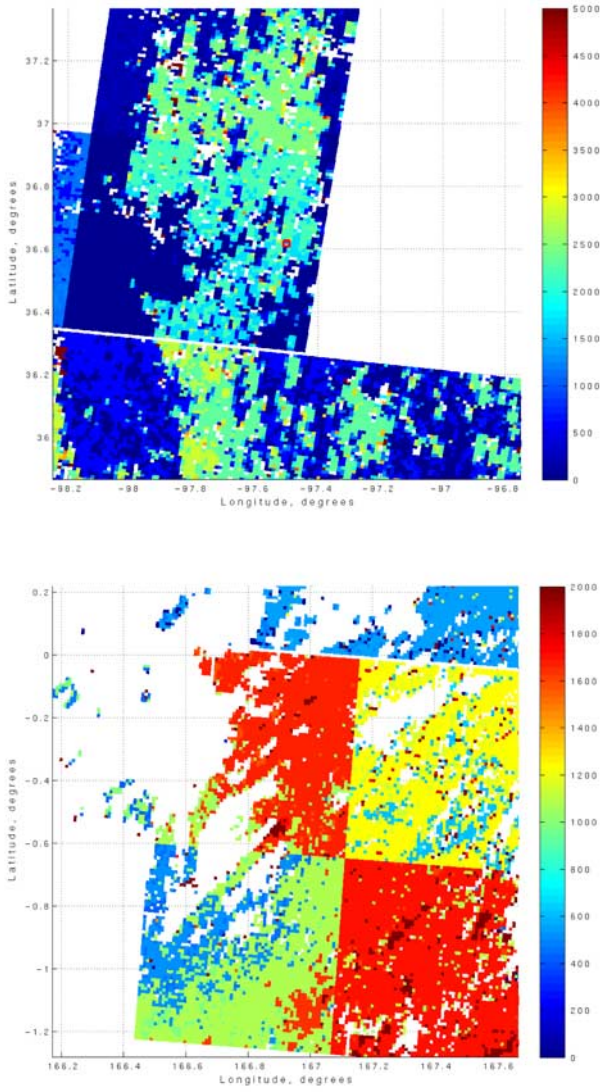

Figure 10. MISR best winds stereo-height retrieval for the same cases shown in Figure 9.

where MISR failed to obtain a good cloud top height retrieval. The good performance of the MISR stereo-height retrieval for water clouds is in agreement with results published by *Marchand et al.* [2001], who compared airborne lidar with a MISR-like retrieval (from aircraft data) for a single arctic stratus case. *Naud et al.* [2005a] also found good agreement between MISR and radar observations for 9 cases of water clouds with cloud tops below 5 km using a combination of 3 and 95 GHz radar near Chilbolton, UK. Similarly, *Naud et al.* [2005b] found good results when examining 30 cases of single-layered nonbroken clouds over Chilbolton and the ARM SGP site.

6.2. Optically Thick Clouds (Optical Depth > ~5) Without Contrast at Cloud Top

[53] For optically thick clouds (optical depth greater than about 5) without well-defined contrast at cloud top (i.e., visually diffuse cloud tops, primarily ice clouds), the MISR-stereo height tends to identify structures within the cloud, typically 500 m to 1.5 km deep, rather than the true cloud top. The apparent optical depth penetration (estimated from ground-based radar retrievals of cloud particle size and condensate amount) varied considerably from case to case with typical values between 0.5 and 5. While the variability in the penetration optical depth is likely due in part to the uncertainty in the height retrieval and the optical depth retrieval, it is also likely that the actual contrast level also varies greatly from case to case.

6.3. Optically Thin Single-Layer Clouds Over Land or Water Surfaces

[54] For single layer clouds at the SGP and Nauru sites (as well as at the Barrow site in the snow-free portion of year), we find that the MISR stereo-height algorithm works well for thin clouds that have an optical depth greater than about 0.3 to 0.5. Clouds with an optical depth less than this are usually not detected (see Figure 8). This results is consistent with a previously published study by *Naud et al.* [2004] who examined 18 cases over a site in France (SIRTA; 48.7°N, −2.2°E), where a lidar system was able to penetrate thin clouds. *Naud et al.* [2004] were only able to make optical depth estimates from the lidar measurements for six of their 18 cases and according these authors did not attempt a quantitative estimate of the cloud optical depth detection limit for the MISR stereo-height algorithm. Be that as it may, for these six cases MISR did not retrieve cloud top for the four cases when the optical depth was 0.2 (or lower) but did retrieve the cloud top for the two cases with an estimated optical depth of 0.5 and 0.6, which is consistent with the results presented here.

6.4. Optically Thin Clouds Over High Contrast Surfaces or Other Clouds

[55] The MISR stereo-height retrieval routinely identifies features below optically thin (cirrus) clouds. The optical depth threshold at which the algorithm identifies the upper thin cloud rather than the height of the underlying object depends on the contrast of the underlying feature. This effect was also noted by *Naud et al.* [2002, 2004] in the case of multilayer clouds. We find that for cases with dark features (such as open water) near bright features (such as most boundary layer clouds, fresh snow, or sea ice), the MISR stereo will preferentially identify the surface or lower cloud layer through all clouds with optical depths less than 0.5 and frequently for clouds with optical depths between 0.5 to about 5 (with a median value of about 2). While the data set examined here contains many such cases, the uncertainty in the radar-based retrievals makes it difficult to set a firm upper boundary on the optical depth. A few extreme cases were found where the best estimate for the optical depth of the missed cloud was greater than 5, but the lower-bound uncertainty for all of these cases is much less than 5 (see example, in section 5.3).

7. Final Remarks

[56] In this study we found that most of the uncertainty in the stereo-height retrieval was associated with uncertainty in the cloud winds. Accordingly, efforts to improve the MISR

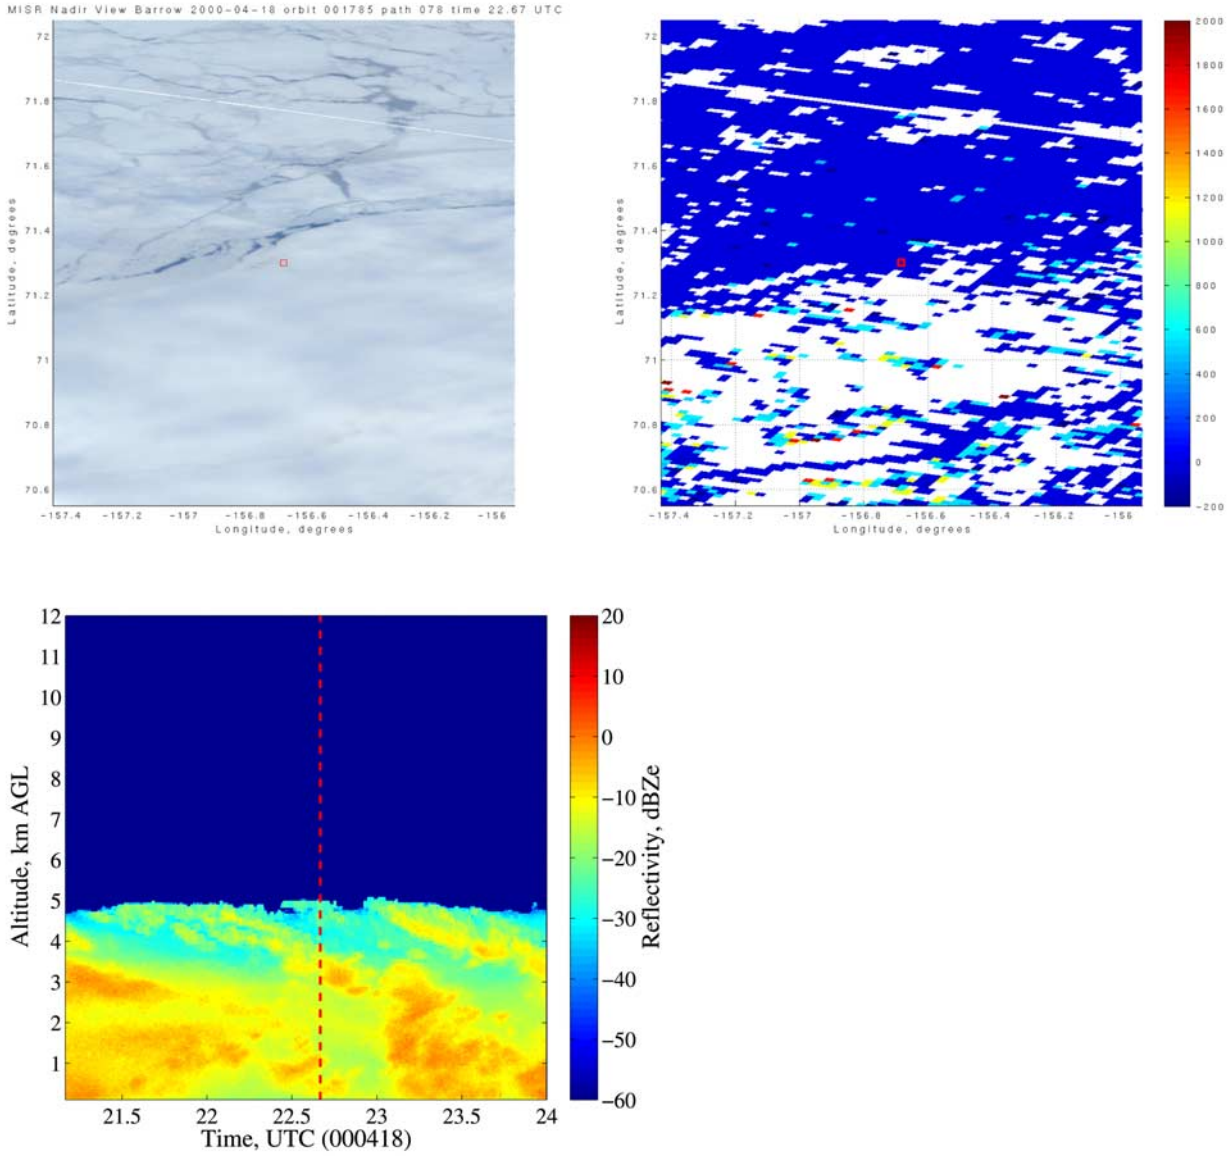

Figure 11. Example of large surface contrast at ARM Barrow site resulting in a stereo-height retrieval at the surface through a physically thick cloud deck. This cloud is composed entirely of ice and a combined radar-reflectivity-Doppler-velocity retrieval indicates an optical depth from 2 to 11. (top right) MISR nadir view. (top left) MISR without winds stereo-height retrieval. (bottom) ARM program millimeter-wavelength cloud radar data, 1 hour centered on time of MISR overpasses.

stereo-height retrieval are currently focused on improving the cloud winds retrieval or determining with greater confidence when the wind retrieval is accurate. In particular, the MISR project is planning to do the wind retrieval twice in future versions of the stereo-height algorithm (versions 015 and higher), once using three forward viewing cameras and once using three aftward viewing cameras. The two sets of velocities will be required to agree (to within a fixed threshold of 10 m/s for the north-south component of the wind vector) for the wind retrieval to be considered good. Preliminary results using this new quality flag indicate that it reduces the observed blockiness. The project is also moving toward using a more time-demanding pattern-matcher with subpixel accuracy [*Davies et al.*, 2007].

[57] Retrieving cloud motion is difficult for MISR because there is only a short time interval to observe the cloud motion and this requires using the more oblique viewing cameras, which in turn makes matching common features difficult. One must also simultaneously solve for the height and motion requiring matching the same feature in at least three views. Increasing the time interval would make for a better wind retrieval, but would also make the height retrieval more sensitive to the value of the retrieved wind. Perhaps the best way to deal with this obstacle in future satellite missions would be to have two observing platforms viewing the same scene from different angles, simultaneously. For example, one small platform with one sensor looking only downward (nadir) and one platform

flying several minutes behind with a forward overlapping view angle and a nadir camera. The first-nadir and forward overlap camera can provide a high-precision cloud top, with no dependence on the cloud motion, and the two nadir cameras can be used to retrieve motion, the accuracy of which will depend on the sensor resolution (and georegistration) with very little dependence on the height. Finally we note for future satellite missions that the accuracy of the stereo-height and winds retrievals will also benefit from increasing the image resolution and need not be restricted to visible wavelengths. For example, one could use a channel similar to the MODIS 1.6 um or 1.38 um band which would improve stereo-height over snow and ice surfaces or enable determination of stereo-height for very optically thin high-altitude cirrus clouds, respectively.

Appendix A: Description of a Combined Retrieval Algorithm for Cloud Water Content, Effective Radius, and Optical Depth

[58] The retrieval algorithm described below combines the output from two radar-reflectivity-only retrievals, a radar-reflectivity-and-Doppler-velocity (ZV) retrieval, a radar-reflectivity-microwave-radiometer (Z-LWP) retrieval, a lidar-only retrieval and a very limited radar-reflectivity-and-Doppler-velocity-and-lidar retrieval. All of these retrievals (except for the limited radar-lidar retrieval) have either been compared in case studies against aircraft data or compared against observed surface shortwave fluxes (after application of 1-D radiative transfer) or both.

[59] While each of these retrievals has thus been established for a particular set of conditions, we are in several respects pushing the retrievals into arenas that they have not been thoroughly tested. For example, none of these retrievals are believed to work well for thick mixed phase or deep convective clouds. We will discuss particular problem areas at the end of this section, after we present details on the combined algorithm. As a guiding principle, we have therefore tried to designed the combined algorithm to return not only a best estimate of the cloud microphysical properties (effective radius and water content), but also to return a conservative set of uncertainty bounds. That is, we generally believe that the combined retrieval uncertainty boundaries will error toward being too broad rather than too narrow at any instant, with a couple of notable exception which we will discuss momentarily.

[60] The input to the algorithms, in addition to obvious millimeter radar reflectivity and Doppler velocity, include liquid water path (estimated) from observation by a passive microwave radiometer), cloud extinction (estimated from micropulse-lidar backscattering profiles) and the closest-in-time available radiosonde observations of the atmospheric state (pressure, temperature, and relative humidity as function of altitude). When no radiosonde data was available within 1 day, a typical atmospheric profile is assumed. All of this data is freely available through the Department of Energy Atmospheric Radiation Measurement (ARM) data archive (http://www.arm.gov).

[61] Using the retrieval algorithm of *Comstock and Sassen* [2001], the ARM micropulse lidar measured backscattering profile is first converted into an extinction profile (on the basis of the assumption of only molecular Rayleigh scattering above the cloud). A test on the lidar signal to noise level and a comparison against radar reflectivity observations are used to determine if the lidar penetrated the cloud. If a good extinction profile is obtained, the optical depth implied by this profile is taken as the best estimate value. This technique is the most direct and most accurate of the retrievals in terms of optical depth, but the ARM micropulse lidar typically only penetrates clouds with an optical depth ranging from about 1 to 3 (depend on the amount of background daylight/noise). J. M. Comstock (personal communication, 2006) places the accuracy of the derived optical at ±20%, which we take as our uncertainty bounds. If the radar also detects the cloud, we use a radar-reflectivity-and-velocity retrieval to estimate the particle size. In our radar-reflectivity-and-velocity retrieval (described in more detail below), we estimate the cloud ice water content and effective radius using several possible crystal habits/shapes. We then compare the radar derived extinction profiles for each habit against the lidar profile and choose the crystal habit (and associated effective radii) that give the best extinction match to the lidar profile. When the radar does not detect the cloud, we know the particle size must be small and therefore assume a particle effective radius of 5 to 10 microns with a best estimate value of 8.

[62] If we fail to obtain a good lidar-based estimate (either because the lidar was not functioning or did not penetrate the cloud), we then examine the radiosonde profile to determine the freezing level and the location of any inversions (of more than 0.5°C). We assume that all radar detections below the freezing level (or below the highest inversion with a temperature of no less than −20°C) are entirely liquid clouds. We assume only ice above this level.

[63] For the liquid water cloud region, we apply a radar-reflectivity-microwave radiometer (Z-LWP) retrieval following the work of *Frisch et al.* [2002]. The microwave radiometer is used to derive the cloud liquid water path, with an uncertainty of approximately 25 g/m^2 [*Marchand et al.*, 2003]. The derived total column liquid water path and observed radar reflectivity (at each height bin) are used to determine the total number of cloud particles and the effective radius assuming a single-mode lognormal size distribution with a fixed lognormal width of 0.35. Our experience with this technique (gained primarily through applying it to field experiment case studies and comparing with aircraft data) indicates that the retrieval works well for nondrizzling water clouds (which are likely well modeled by a single lognormal size distribution). We have found that when the clouds are vertically thin and broken, or drizzling, or contain ice or mixed phased hydrometers the Frisch retrieval scheme tends to fail in such a way that the algorithm returns a layer mean effective radius less than 4 um or greater than 25 um. The retrieved effective radius will be too small if either the LWP is too large (e.g., because of water on the MWR ray dome) or if the observed radar reflectivity is too small (which can occur because the cloud does not fill the entire radar range bin or in a more general sense because of beam width and sampling differences between the MWR and radar instruments). Conversely, when the observed radar reflectivity is large because of the presence of ice or drizzle particles (which invalidate the single-mode-lognormal assumption) the retrieval tends to return effective radius which is too large. Therefore if we

retrieve a layer mean effective radius less than 4um or greater than 25 um, we reject the Frisch retrieval and turn to radar-reflectivity-only and an LWP-only based solution. First we try to apply the reflectivity-only based approach by applying the droplet concentration to water content relationship suggested by *Sassen and Liao* [1996], with an assumed total number concentration ranging from 25 to 500 cm^{-3} and with a value of 100 cm^{-3} for the best estimate. This radar-only solution is then checked against the microwave radiometer liquid water path (when the liquid water path is less than 500 g/m^2 as values larger than this are generally associated with rain or water on the instrument the results) to ensure the liquid water path returned by the radar-only solution is not excessive (that is not more than LWP+2*LWP_uncertainty larger than the MWR derived value). Unfortunately, we often find that the when Frisch radar-reflectivity-MWR approach fails, the Sassen and Liao approach also fails because it is usually the radar-reflectivity that does not meet the single-mode approximation rather the retrieved LWP that is the source of the problem. In such cases, we default to an even simpler scheme that does not depend on the absolute value of radar-reflectivity. Rather we assume a fixed effective radius ranging from 5 to 15 um, with a value of 7.5 um for the best estimate, typical of continental status [*Miles et al.*, 2000] and simply distributed the microwave-radiometer-derived liquid water path in the vertical by weighting it according to the square root of the observed reflectivity. We also use LWP+2*LWP_uncertainty and LWP-2*LWP_uncertainty to obtain uncertainty bounds. It is not unusual for the value of LWP-2*LWP_uncertainty to be less than 0. Whenever the value is less than 5 g/m^2, we set the value to this minimum.

[64] For the ice cloud region, we apply a radar-reflectivity (Z) and mean-Doppler-velocity (V) retrieval technique. Two such ZV algorithms have been published: one by *Mace et al.* [2002] and one by *Matrosov et al.* [2002]. These ZV retrievals are attractive in that they do not require single layer or optically thin clouds and can therefore be applied during many periods when most other retrievals are inapplicable. In our combined retrieval, we follow the approach given by *Mace et al.* [2002] but have modified the Mace algorithm in several ways. Most significant is that our version of the algorithm requires an explicit choice/model of the ice crystal habit (or an explicit mixture) rather than relying on a more generic power law relationship that relates particle fall velocity and particle mass to particle maximum dimension. Significant limitations in the radar-reflectivity-and-Doppler-velocity approach include (1) the relationship between the measurements (both reflectivity and fall velocity) and the particle microphysics depends on the crystal habit, which is unknown; (2) one must estimate the particle fall velocity from the measured Doppler velocity, which includes effects due to updrafts and other air motions; and (3) the algorithm assumes a single-mode particle size distribution. With regard to the uncertainty due to crystal habit, in the combined retrieval algorithm we run the retrieval using clouds composed of 100% aggregates, 100% rosettes, 100% solid columns, 100% hollow columns, and 100% plates. We take the solution with the minimum and maximum optical depth to form our uncertainty limits. We choose rosettes to form our best estimate, partially on the basis of the fact that this solution tends to fall in the middle of the range of values and partially on the basis of the frequent observations of rosettes in aircraft data sets over the ARM SGP site. With regard to the fall velocity, as discussed by *Mace et al.* [2002], a vertically pointing millimeter-wavelength Doppler radar (such as those operated by the DOE ARM program) measures the absolute velocity of the cloud particles in the vertical direction (along the line of sight of the radar) not only the particle fall velocity. Vertical air motions occur on a wide range of spatial scales and are generally of sufficient magnitude that they cannot be neglected in the retrieval process. Mace et al. adopt an approximate model where the air motions are separated into those that are external to the cloud system being observed (by which Mace means they cannot be estimated or reduced substantially in magnitude by averaging the radar Doppler velocity measurements) and those that are internal (and can be averaged out). We adopt this same conceptual model for the air motion and like Mace we assume that if we take all the measured Doppler velocities (over a period of 1–3 hours) associated with a particular value of the reflectivity (at a particular altitude) and average these velocities we obtain a measure of only the mean particle fall speed plus the external (or large-scale) atmospheric motion. Another way to look at this approach is that we are effectively constructing a tuned regression, under the assumption that similar reflectivities (in a particular altitude region) have similar particle size distributions and should therefore have similar particle fall velocities and water contents. Clearly this assumption is not valid during periods of strong convective activity. We have noticed that such periods often reveal themselves as an apparent decrease in the average velocity with increasing reflectivity (whereas one normally observes the opposite to occur). With regard to the remaining external (or large-scale) vertical velocity, on the basis of results from *Gultepe and Starr* [1995], *Mace et al.* [2002] suggest that the large-scale vertical velocity will typically be in the range of ±10 to 15 cm/s. In our combined retrieval, we run the retrieval three times assuming a large-scale velocity of +15, 0, and −15 cm/s (for each crystal habit) and use the minimum and maximum values to define the uncertainty limits (and the rosettes-only solution with 0 cm/s large-scale velocity as the best estimate).

[65] The combined algorithm, described above, is robust in the sense that it returns a result all of the time (at least when millimeter-wavelength radar data is available). However, there are a number of significant limitations that need to be considered in using the results. One is that the decision on how to combine the results is made on a time-sample-by-time-sample basis. As a result, there can be large discontinuities in time series of retrieved values. For example, as a cumulus cloud passes over the radar the combined algorithm will often accept the Frisch retrieval result over the bulk of the cloud. However, the edge of the cumulus cloud as observed by the radar does not generally correspond to the edge of the cloud as observed by the microwave-radiometer/LWP data set. Thus the Frisch scheme tends to fail at the edge of water clouds and the combined algorithm defaults to a reflectivity-only solution. One approach to fix this problem might be to design a ''on-the-fly'' tuned-radar-reflectivity-regression which used the values from the Frisch scheme obtained in the middle of the cloud and applied the tuned parameters to the entire cloud as observed by the radar.

Of course, it is entirely possible and maybe even probable that the cloud properties are not the same in the middle as at the cloud boundaries. We may eventually create such an algorithm, but for the present we merely note the presence of unphysical discontinuities in our retrieval process.

[66] Another problem with the combined algorithm is that none of the retrievals are explicitly designed to work during rain, deep convection, or mixed phase conditions. Rain is problematic in that even a light rain will generally corrupt the MWR estimates of LWP (making them too large) and the radar reflectivity can also be attenuated during heavier rains. At present, we try to flag the presence of rain by monitoring the MWR LWP values and recognize that the retrievals are of low quality in this circumstance. Cases of mixed phase clouds (which are not uncommon) are treated in ad hoc fashion, with the warm region (below the inversion or 0° line) treated as a water cloud and above as an ice clouds. For thick/high-reflectivity mixed phase or deep convective clouds or in rain, the combined algorithm generally defaults to a reflectivity-only approach with an effective uncertainty that is estimated to be an order of magnitude. The reflectivity-only retrievals were not designed for mixed phase conditions and even an order of magnitude uncertainty is little more than a guess. Thin mixed phase clouds (including this altostratus composed of supercooled water) often have very little liquid water such that the microwave radiometer retrievals does not unambiguously indicate that the cloud does have water, and yet this small amount of water can dominate the cloud optical thickness. The combined retrieval will generally badly underestimate the optical thickness of these clouds treating them as pure ice clouds with relatively large particles. In general, for clouds with little liquid water (less than 25 to 50 g/m^2) that are in the "warm region" we assume a minimum value of 5 g/m^2 in the combined retrieval, but this is just a fill value that gives an optical depth greater than about 1, which we know to be true because the lidar does not penetrate them. While the ARM retrieval community has known about this low liquid water problem for a number of years, little headway has been made toward a broad solution [*Turner et al.*, 2007].

[67] **Acknowledgments.** This research would not have been possible without the hard work of many individuals on the MISR science team and at the EOS DAAC. Thank you all! Ground-based data were obtained from the Atmospheric Radiation Measurement Program sponsored by the U.S. Department of Energy Office of Science, Office of Biological and Environmental Research, Environmental Science Division. We would like to express extra thanks to the many members of the ARM infrastructure who keep the data flowing. Primary funding for this work was provided by the NASA MISR project under the direction of David J. Diner at the NASA Jet Propulsion Laboratory.

References

Angevine, W. M., P. S. Bakwin, and K. J. Davis (1998), Wind profiler and RASS measurements compared with measurements from a 450-m-tall tower, *J. Atmos. Oceanic Technol.*, *15*(3), 818–825.

Barth, M. F., R. B. Chadwick, and D. W. van de Kamp (1994), Data processing algorithms used by NOAA's Wind Profiler Demonstration Network, *Ann. Geophys.*, *12*, 512–528.

Clothiaux, E. E., T. P. Ackerman, G. G. Mace, K. P. Moran, R. T. Marchand, M. A. Miller, and B. E. Martner (2000), Objective determination of cloud heights and radar reflectivities using a combination of active remote sensors at the ARM CART sites, *J. Appl. Meteorol.*, *39*(5), 645–665.

Cohn, S. A., and R. K. Goodrich (2002), Radar wind profiler radial velocity: A comparison with Doppler lidar, *J. Appl. Meteorol.*, *41*(12), 1277–1282.

Comstock, J. M., and K. Sassen (2001), Retrieval of cirrus cloud radiative and backscattering properties using combined lidar and infrared radiometer (LIRAD) measurements, *J. Atmos. Oceanic Technol.*, *18*(10), 1658–1673.

Davies, R., A. Horvath, C. Moroney, B. Zhang, and Y. Zhu (2007), Cloud motion vectors from MISR using sub-pixel enhancements, *Remote Sens. Environ.*, in press.

Diner, D. J., J. C. Beckert, G. W. Bothwell, and J. I. Rodriguez (2002), Performance of the MISR instrument during its first 20 months in earth orbit, *IEEE Trans. Geosci. Remote Sens.*, *40*(7), 1449–1466.

Diner, D., et al. (2005), The value of multiangle measurements for retrieving structurally and radiatively consistent properties of clouds, aerosols, and surfaces, *Remote Sens. Environ.*, *97*(4), 495–518.

Frisch, S., M. Shupe, I. Djalalova, G. Feingold, and M. Poellot (2002), The retrieval of stratus cloud droplet effective radius with cloud radars, *J. Atmos. Oceanic Technol.*, *19*(6), 835–842.

Gultepe, I., and D. O. Starr (1995), Dynamical structure and turbulence in cirrus clouds: Aircraft observations during FIRE, *J. Atmos. Sci.*, *52*(23), 4159–4182.

Horvath, A., and R. Davies (2001a), Simultaneous retrieval of cloud motion and height from polar-orbiter multiangle measurements, *Geophys. Res. Lett.*, *28*(15), 2915–2918.

Horvath, A., and R. Davies (2001b), Feasibility and error analysis of cloud motion wind extraction from near-simultaneous MISR measurements, *J. Atmos. Oceanic Technol.*, *18*, 591–680.

Mace, G. G., A. J. Heymsfield, and M. R. Poellot (2002), On retrieving the microphysical properties of cirrus clouds using the moments of the millimeter-wavelength Doppler spectrum, *J. Geophys. Res.*, *107*(D24), 4815, doi:10.1029/2001JD001308.

Marchand, R. T., T. P. Ackerman, M. D. King, C. Moroney, R. Davies, J. A. L. Muller, and H. Gerber (2001), Multiangle observations of Arctic clouds from FIRE ACE: June 3, 1998, case study, *J. Geophys. Res.*, *106*(D14), 15,201–15,214.

Marchand, R., T. Ackerman, E. R. Westwater, S. A. Clough, K. Cady-Pereira, and J. C. Liljegren (2003), An assessment of microwave absorption models and retrievals of cloud liquid water using clear-sky data, *J. Geophys. Res.*, *108*(D24), 4773, doi:10.1029/2003JD003843.

Martner, B. E., D. B. Wuertz, B. B. Stankov, R. G. Strauch, E. R. Westwater, K. S. Gage, W. L. Ecklund, C. L. Martin, and W. F. Dabbert (1993), An evaluation of wind profiler, RASS, and microwave radiometer performance, *Bull. Am. Meteorol. Soc.*, *74*, 599–613.

Matrosov, S. Y., A. V. Korolev, and A. J. Heymsfield (2002), Profiling cloud ice mass and particle characteristic size from Doppler radar measurements, *J. Atmos. Oceanic. Technol.*, *19*(7), 1003–1018.

Miles, N. L., J. Verlinde, and E. E. Clothiaux (2000), Cloud droplet size distributions in low-level stratiform clouds, *J. Atmos. Sci.*, *57*(2), 295–311.

Moroney, C., R. Davies, and J.-P. Muller (2002), Operational retrieval of cloud-top heights using MISR data, *IEEE Trans. Geosci. Remote Sens.*, *40*, 1532–1540.

Muller, J.-P., A. Mandanayake, C. Moroney, R. Davies, D. J. Diner, and S. Paradise (2002), Operational retrieval of cloud-top heights using MISR data, *IEEE Trans. Geosci. Remote Sens.*, *40*, 1547–1559.

Naud, C., J. Muller, and E. E. Clothiaux (2002), Comparison of cloud top heights derived from MISR stereo and MODIS CO_2-slicing, *Geophys. Res. Lett.*, *29*(16), 1795, doi:10.1029/2002GL015460.

Naud, C., J. Muller, M. Haeffelin, Y. Morille, and A. Delaval (2004), Assessment of MISR and MODIS cloud top heights through intercomparison with a back-scattering lidar at SIRTA, *Geophys. Res. Lett.*, *31*, L04114, doi:10.1029/2003GL018976.

Naud, C. M., J. P. Muller, E. C. Slack, C. L. Wrench, and E. E. Clothiaux (2005a), Assessment of the performance of the Chilbolton 3-GHz advanced meteorological radar for cloud-top-height retrieval, *J. Appl. Meteoro.*, *44*, 866–877.

Naud, C. M., J.-P. Muller, E. E. Clothiaux, B. A. Baum, and W. P. Menzel (2005b), Intercomparison of multiple years of MODIS, MISR and radar cloud-top heights, *Ann. Geophys.*, *23*, 1–10.

Sassen, K., and L. Liao (1996), Estimation of cloud content by W-band radar, *J. Appl. Meteorol.*, *35*(6), 932–938.

Seiz, G., R. Davies, and A. Gruen (2006), Stereo cloud-top height retrieval with ASTER and MISR, *Int. J. Remote Sens.*, *27*(9), 1839–1853.

Turner, D. D., et al. (2007), Thin liquid water clouds: Their importance and our challenge, *Bull. Am. Meteorol. Soc.*, in press.

T. P. Ackerman and R. T. Marchand, Pacific Northwest National Laboratory, P.O. Box 999, Richland, WA 99352, USA. (roj@pnl.gov)

C. Moroney, NASA Jet Propulsion Laboratory, Pasadena, CA 91109-8099, USA.