SPECIAL CALIPSO COLLECTION

Fully Automated Detection of Cloud and Aerosol Layers in the CALIPSO Lidar Measurements

MARK A. VAUGHAN,* KATHLEEN A. POWELL,* RALPH E. KUEHN,+ STUART A. YOUNG,# DAVID M. WINKER,* CHRIS A. HOSTETLER,* WILLIAM H. HUNT,+ ZHAOYAN LIU,@ MATTHEW J. MCGILL,& AND BRIAN J. GETZEWICH+

* *National Aeronautics and Space Administration, Hampton, Virginia*
\+ *Science Systems and Applications, Inc., Hampton, Virginia*
\# *CSIRO Marine and Atmospheric Research, Aspendale, Victoria, Australia*
@ *National Institute of Aerospace, Hampton, Virginia*
& *National Aeronautics and Space Administration, Greenbelt, Maryland*

(Manuscript received 19 September 2008, in final form 16 April 2009)

ABSTRACT

Accurate knowledge of the vertical and horizontal extent of clouds and aerosols in the earth's atmosphere is critical in assessing the planet's radiation budget and for advancing human understanding of climate change issues. To retrieve this fundamental information from the elastic backscatter lidar data acquired during the Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations (CALIPSO) mission, a selective, iterated boundary location (SIBYL) algorithm has been developed and deployed. SIBYL accomplishes its goals by integrating an adaptive context-sensitive profile scanner into an iterated multiresolution spatial averaging scheme. This paper provides an in-depth overview of the architecture and performance of the SIBYL algorithm. It begins with a brief review of the theory of target detection in noise-contaminated signals, and an enumeration of the practical constraints levied on the retrieval scheme by the design of the lidar hardware, the geometry of a space-based remote sensing platform, and the spatial variability of the measurement targets. Detailed descriptions are then provided for both the adaptive threshold algorithm used to detect features of interest within individual lidar profiles and the fully automated multiresolution averaging engine within which this profile scanner functions. The resulting fusion of profile scanner and averaging engine is specifically designed to optimize the trade-offs between the widely varying signal-to-noise ratio of the measurements and the disparate spatial resolutions of the detection targets. Throughout the paper, specific algorithm performance details are illustrated using examples drawn from the existing CALIPSO dataset. Overall performance is established by comparisons to existing layer height distributions obtained by other airborne and space-based lidars.

1. Introduction

On 28 April 2006, eight years of close collaboration between the National Aeronautics and Space Administration (NASA) and the Centre National d'Etudes Spatiales (CNES) came to fruition with the launch of the Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations (CALIPSO) mission from Vandenberg Air Force Base in California (Winker et al. 2007). Launched simultaneously with the Cloudsat satellite aboard a single Delta-II rocket, CALIPSO is now an integral part of NASA's A-Train of Earth-observing remote sensing satellites (Stephens et al. 2002). The primary payload aboard CALIPSO is the Cloud-Aerosol Lidar with Orthogonal Polarization (CALIOP). CALIOP is an elastic backscatter lidar that transmits linearly polarized laser light at 532 and 1064 nm and measures range-resolved backscatter intensities at both wavelengths using a three-channel receiver. For the 532-nm signal, a polarizing beam splitter separates the backscattered light into components polarized parallel and perpendicular to the polarization plane of the laser output. These are then measured

Corresponding author address: Mark Vaughan, Mail Stop 475, NASA Langley Research Center, Hampton, VA 23681-2199.
E-mail: mark.a.vaughan@nasa.gov

DOI: 10.1175/2009JTECHA1228.1

© 2009 American Meteorological Society

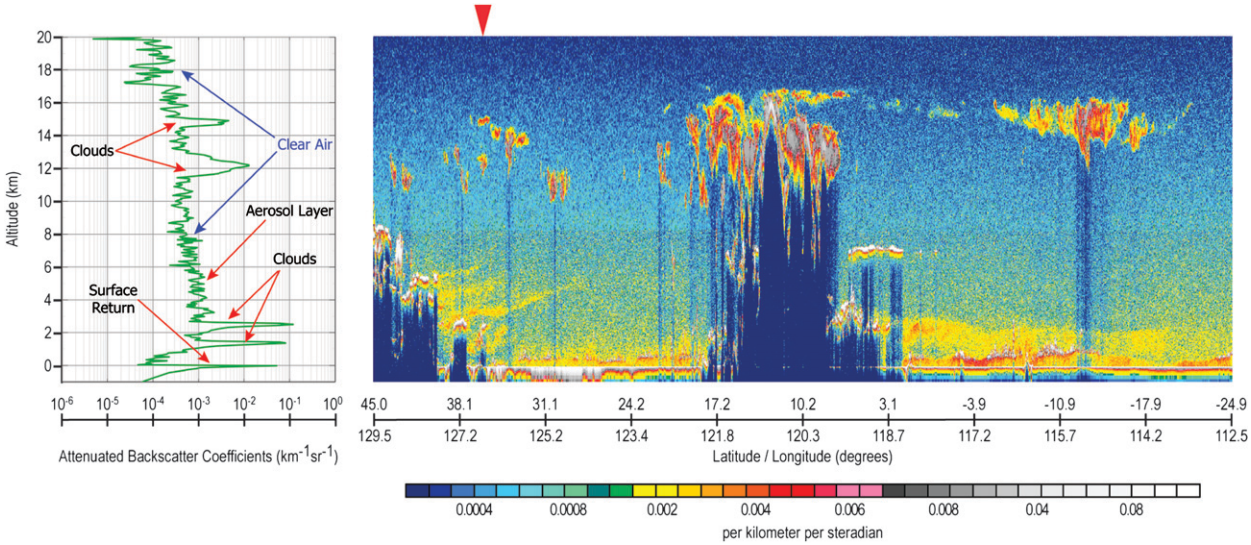

FIG. 1. (left) CALIOP 5-km horizontal average profile (15 laser pulses) of 532-nm attenuated backscatter coefficients measured over the Yellow Sea (36.1°N, 126.6°E) on 13 Aug 2006 at ~1741:18 UTC. Clouds, aerosols, and the ocean surface return all appear as signal enhancements rising above the ambient molecular backscatter intensity. (right) Time history of CALIOP backscatter measurements acquired on 13 Aug 2006 from ~1738:49 to ~1758:15 UTC, illustrating the wide range of feature locations and backscatter intensities that can be encountered during a single orbit segment. The red arrow at the top of the image denotes the location of the profile data shown in the left panel.

independently using a pair of photomultipliers. The total backscatter at 1064 nm (i.e., both parallel and perpendicular polarizations) is measured using a single avalanche photodiode. An in-depth description of the CALIOP instrument can be found in Hunt et al. (2009). Winker et al. (2009) provide a summary of the CALIPSO mission goals and a top-level overview of the end-to-end CALIOP data analysis architecture.

The fundamental measurements made by CALIOP are calibrated range-resolved profiles of backscatter intensity (Powell et al. 2009). Embedded within these profiles are the signals from a variety of geophysical entities, including clouds, aerosol layers, regions of clear air, and the returns from the earth's surface. As shown in the left panel of Fig. 1, all of these can be present within a single profile. Examination of the right panel of Fig. 1 shows that, even over a short time, CALIOP can encounter a large number of dissimilar scenarios. In the span of ~20 min, CALIOP observes instances of multiple cloud layers (e.g., at ~43°N and ~6°N); faint, possibly subvisible, cirrus (~20°S, at ~15 km MSL); lofted aerosol layers (~32°N, ~4 km vertically); aerosol layers beneath overlying cirrus (at the equator and at ~16°S); cumulus embedded in boundary layer aerosols (~26°N); and aerosol extending above broken cloud decks (~10°S).

The fundamental data products derived from the CALIOP profile measurements are the spatial locations of these many different types of geophysical entities. The function of the CALIOP layer detection algorithm is thus to untangle scenes such as that shown in Fig. 1, to identify those portions of the profiles backscattered from clouds, aerosols, and/or the earth's surface, and to clearly separate those backscattered portions from the ambient "clear air" scattering (i.e., from regions of purely molecular atmosphere). To refer in general to any of the several types of CALIOP detection targets, we adopt the term "feature." Features are composed of a generic substance called particulates and are defined as any extended, vertically contiguous region of enhanced backscatter that rises significantly above the signal magnitude expected from a purely molecular atmosphere. A successful detection scheme must be able to identify naturally occurring features characterized by backscatter intensities that vary over many orders of magnitude. To accomplish this goal, we have constructed a selective, iterated boundary location (SIBYL) algorithm, in which a nested multiresolution spatial averaging scheme is driven by a robust profile scanning engine that incorporates an adaptive context-sensitive threshold algorithm. After averaging the profile data to an initial horizontal resolution, SIBYL invokes the profile scanner to search for the presence of features. Where features are detected, the backscatter data between feature top and feature base are removed from the profile, and the data below the feature base are corrected for the estimated signal attenuation that occurs as the lidar beam propagates through the layer. These feature-cleared profiles are then averaged to a new, coarser spatial resolution, thereby increasing

the signal-to-noise ratio (SNR) of the remaining data so that subsequent profile scans can identify progressively fainter features.

The complexity of the SIBYL scheme arises as a direct consequence of the harsh retrieval environment in which the algorithm must function. Because of the mass and power restrictions imposed on space-based platforms and the very large distance between the satellite and the measurement targets (~705 km), the SNR of the CALIPSO measurements is substantially lower than is typical for ground-based and/or airborne lidars. Some degree of spatial averaging is thus required to detect all but the strongest cloud and surface signals. Figure 2 illustrates the necessity of averaging. The left panel shows a full-resolution CALIOP profile taken from the data shown in Fig. 1 at 19.2°S, 113.9°E. Although a strong echo from the earth's surface is plainly visible at 0 km, the remainder of the data appears to be feature-free noise. When these same data are averaged to a 20-km horizontal resolution, as shown in the right panel of Fig. 2, an aerosol layer clearly emerges at ~2.2 km and the molecular signal profile begins to take its expected shape. However, SIBYL's capacity to improve profile SNR via averaging is limited. Taken together, the relatively low pulse rate of the laser (20.16 Hz), the high speed at which the satellite traverses the ground track (~7 km s^{-1}), and the constantly changing atmospheric content impose stringent limitations on the amount of horizontal averaging that can reasonably be done. These restrictions on spatial averaging are further constrained by the essential science objectives of the CALIPSO mission, which require accurate separation of clouds and aerosols at the highest possible spatial resolution (Winker et al. 2009).

The retrieval difficulties cited thus far are exacerbated by the solar background signals present during daytime observations. Unlike passive sensors that rely heavily, or even entirely, on reflected sunlight for their measurements, CALIOP is an active sensor that is equipped with its own light source and thus can acquire profile data continuously during both the daytime and nighttime portions of every orbit. Although the magnitude of the additional background signal introduced during daytime operations is relatively easy to measure and remove, the concomitant noise remains and exerts a pronounced and deleterious effect on the backscatter SNR.

Although the task of detecting layer boundaries within the CALIOP data is handled exclusively by the SIBYL algorithm, the equally important task of classifying layers as clouds or aerosols is, with one exception, accomplished externally by separate data processing modules. The algorithms used to discriminate between clouds and aerosols are described in Liu et al. (2009). The analyses subsequently applied to identify different aerosol types are outlined by Omar et al. (2009). Similarly, Hu et al. (2009) describe the methods used to determine cloud ice-water phase. SIBYL's sole contribution to the layer classification task is the high-resolution boundary layer cloud-clearing process described in section 3b.

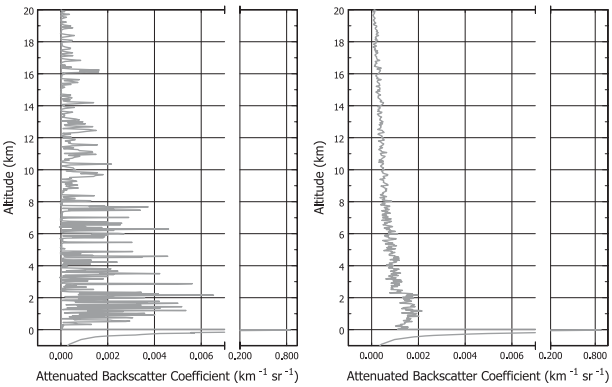

FIG. 2. (left) An example of the native SNR of the CALIOP data shown by a single, full-resolution 532-nm attenuated backscatter profile acquired 13 Aug 2006 at 19.2°S, 113.9°E; (right) the same data incorporated into a 20-km (60 profile) horizontal average. The aerosol layer at ~2.2 km is present in both profiles, but its presence only becomes apparent in the averaged data.

2. The CALIOP profile scanning engine

The profile scanning technique implemented in SIBYL relies on several basic assumptions. First, because CALIOP is an accurately controlled, near-nadir-viewing instrument, all nominal backscatter profiles will contain at least one feature (i.e., either a totally attenuating cloud or aerosol layer or the earth's surface). We further assume that the global structure of the molecular components of the atmosphere (i.e., clear air) is well understood and can be reliably modeled, and that reasonably accurate digital elevation maps of the surface of the planet are readily available. Rayleigh scattering theory is well developed and easily applied to scattering from molecules at the CALIOP wavelengths so that, based on our knowledge of the molecular density profile and surface elevation at any point along the CALIPSO orbit track, we can compute an accurate model that describes the expected backscatter return from the purely molecular components of the column being observed. Yet despite being well calibrated (Powell et al. 2009), the backscatter profiles measured by CALIOP can and will deviate from these theoretical expectations. The two most prominent perturbations are the presence of clouds and aerosols in the atmosphere and the existence of random noise and occasional systematic uncertainties in the lidar measurements. The sole function of the CALIOP profile scanner is to positively identify all instances of the former while summarily rejecting all instances of the latter.

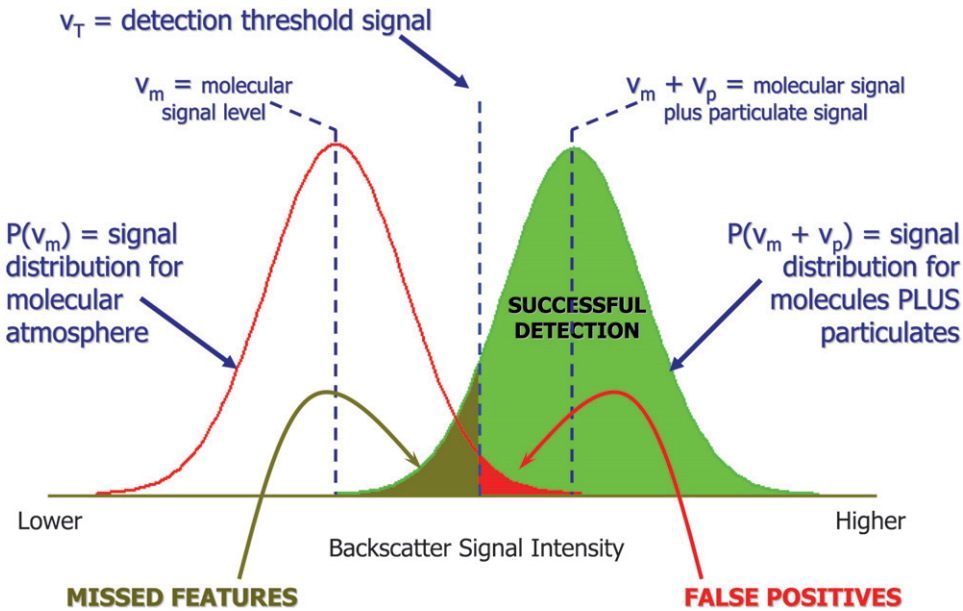

FIG. 3. Schematic comparing (left) the expected signal distribution due to molecules only to (right) the distribution expected from a combination of molecules and particulates. A detection threshold V_T partitions the feature histograms into three distinct regions: the fraction of successful detections is represented by the green area, the probability for false positive detections is represented by the red area, and the probability for false negatives (missed features) is shown in brown.

In section 3a, we briefly review the fundamentals of threshold detection schemes for noisy signals. The signal regime (attenuated scattering ratios) in which the scanning algorithm is applied is described in section 3b. The remaining subsections explain the rationale and the procedures for building a range-dependent threshold array and the justification and mechanics for the context-sensitive adjustments to the initial array that are required to compensate for signal attenuation.

a. Threshold detection fundamentals

In the atmosphere, the backscatter intensity measured at any range can come either from molecules alone or from the combination of molecules and particulates. For well-calibrated and noise-free measurements, differentiating between a signal from molecules only, V_m, and one from molecules plus some particulate contribution, V_{m+p}, is relatively straightforward, because the expected molecular contribution can generally be well characterized. The two classes of signals can easily be separated by establishing a threshold value $V_T = V_m + \Delta V_m$ such that only those signals exceeding V_T are identified as features. When noise is not a consideration, ΔV_m does no more than place an upper bound on the expected fidelity of the molecular model used to estimate V_m. In practice, however, the lidar signal is always contaminated with some amount of noise. Uncertainties are introduced from a variety of sources, including the stochastic processes governing photoelectron multiplication in the detectors, natural variations in the solar background signals, and the Poisson-distributed photon arrival rates of the backscattered laser light (Liu et al. 2006). As a result, the measured magnitudes of both V_m and V_{m+p} are not fixed values but are instead characterized by probability distributions. Estimating V_T now becomes somewhat more complex in that, for any range bin, separation of V_{m+p} from V_m requires some knowledge of the first two moments of the V_m distribution. Furthermore, successful detection of the complete vertical extent of the initial feature and of any secondary features that may also be present requires updated assessments of the signal attenuation resulting from overlying layers.

Figure 3 illustrates the detection problem. In the presence of noise, a threshold detection scheme is susceptible to two kinds of errors: the measured backscatter intensity of legitimate features may fall below the detection threshold, which results in missed features (false negatives), or the noise excursions from a molecular signal may exceed the threshold, which results in the identification of phantom features (false positives). Given a sufficiently stable atmospheric scene, the probability of successful detection can be increased considerably by applying additional signal averaging. In such cases, averaging will reduce the standard deviations about V_m and V_{m+p}, which in turn will decrease the overlap region between the two distributions and diminish the fraction

of false positives and false negatives reported. However, as discussed earlier, the combination of the CALIPSO science requirements and CALIOP design constraints frequently compromises our ability to confidently average sufficient amounts of data while simultaneously avoiding the spatial smearing of separate, dissimilar features. The threshold selection problem therefore remains the same, irrespective of the underlying distributions and the amount of averaging applied. In all cases, an optimal value of V_T must be chosen to maximize the likelihood of successfully detecting a feature while simultaneously minimizing the occurrence of spurious identifications due to noise.

The existing literature on mathematical detection theory is well developed (e.g., Kay 1998). As several specific derivations of the theoretical expectations for detection limits and detection efficiencies for space-based lidar measurements are given elsewhere (Liu et al. 2002; Vaughan et al. 2005), further discussion of these topics is postponed and will appear in future publications. The focus of this work is on the practical considerations required to establish the threshold values used in the CALIOP retrieval processes.

b. Attenuated scattering ratios

The CALIOP level 1 data products report range-dependent profiles of attenuated backscatter coefficients $\beta'_\lambda(r)$, defined as

$$\beta'_\lambda(r) = [\beta_{\lambda,m}(r) + \beta_{\lambda,p}(r)]\, T^2_{\lambda,m}(r)\, T^2_{\lambda,\mathrm{O}_3}(r)\, T^2_{\lambda,p}(r), \qquad (1)$$

where $\beta'_{\lambda,m}(r)$ and $\beta'_{\lambda,p}(r)$ are, respectively, the volume backscatter coefficients for molecules (m) and particulates (p) at wavelength λ (either 532 or 1064 nm); and $T^2_{\lambda,m}(r)$, $T^2_{\lambda,\mathrm{O}_3}(r)$, and $T^2_{\lambda,p}(r)$ represent signal attenuation terms (two-way transmittances) due to, respectively, molecules, ozone (O_3), and particulates (Powell et al. 2009). SIBYL searches are conducted using the 532-nm total attenuated backscatter signal, as small aerosol particles are more efficient scatterers at 532 nm than at 1064 nm.

For a space-based, near-nadir pointing lidar operating at the CALIOP wavelengths, $\beta'_\lambda(r)$ measured in clear air is an increasing function of range (i.e., a decreasing function of altitude above the surface). To transform this range-dependent function into one that is constant with range, we use molecular and ozone number density profiles supplied by the NASA Global Modeling and Assimilation Office (GMAO; Bloom et al. 2005) to convert the 532-nm attenuated backscatter coefficients into attenuated scattering ratios, $R'(r)$, such that

$$R'(r) = \frac{\beta'_{532}(r)}{\beta'_{\mathrm{GMAO}}(r)} = \left[1 + \frac{\beta_{532,p}(r)}{\beta_{532,m}(r)}\right] T^2_{532,p}(r), \qquad (2)$$

where

$$\beta'_{\mathrm{GMAO}}(r) = \beta_{532,m}(r)\, T^2_{532,m}(r)\, T^2_{532,\mathrm{O}_3}(r) \qquad (3)$$

represents the molecular attenuated backscatter coefficients derived from the GMAO model data. In completely clear air, $\beta'_{532,p}(r) = 0$ and $T^2_{532,p}(r) = 1$, so that, absent any excursions due to noise, $R'(r) = 1$ for the entire atmospheric portion of the profile.

c. Establishing an initial threshold level

Conceptually, the CALIOP profile scanner is very similar to the algorithms used to detect clouds in uplooking ground-based radar data (Uttal et al. 1993) and lidar data (Winker and Vaughan 1994). Once a threshold level has been established, the profile data are examined sequentially, beginning immediately below the 532-nm calibration region (~30 km) and moving downward toward the surface. Feature boundaries are determined by locating the first (top) and last (base) points of those regions where the profile data exceeds the threshold value for all points within some predetermined altitude range (i.e., over some minimum feature thickness). Because the backscatter signal is attenuated by passing through the feature, this first estimate of feature base must be further refined by searching downward to identify that point at which $R'(r)$ is no longer a decreasing function of range. Where the CALIOP scheme deviates from previous methods is in defining an initial threshold, which, rather than being a constant, is instead a range-dependent array that accounts for expected variations in profile magnitude and SNR due to both continuous changes in molecular density as a function of altitude and step changes in the CALIPSO onboard averaging scheme (Hunt et al. 2009). The initial magnitude of the threshold array is taken to be identical to a molecular attenuated backscatter model derived from GMAO-provided profiles of molecular and ozone number densities indexed to the latitude and longitude of the profile footprint. To estimate the measurement uncertainties necessary to construct the final threshold array, two categories of noise are considered. The first category consists of those contributions that, for a single-shot profile, remain constant with respect to range from the lidar. Included in this portion of the noise budget are the detector dark current and the noise resulting from solar background light. This quantity is measured on board the satellite for each laser pulse by computing the standard deviation of the background-subtracted backscatter profiles in a region of the profile where molecular scattering is essentially negligible (~80 to ~65 km). Paradoxically, however, because the CALIOP profile data are averaged both vertically and horizontally on board the satellite prior to being downlinked (see

TABLE 1. CALIOP data averaging scheme applied to all backscatter profile measurements prior to downlink. Onboard data acquisition resolution is 20.16 Hz horizontally (~0.33 km) and 10 MHz vertically (~15 m; Hunt et al. 2009). The standard deviation of the background signal is measured and downlinked single-shot (~1/3 km) resolution. The rightmost column provides the scale factors required to compute β'_{BKG} as a function of onboard averaging region [i.e., as in Eq. (4)] for application to a single-shot 532-nm total attenuated backscatter profile obtained from the CALIOP level 1 data products.

Altitude region		Horizontal resolution	Vertical resolution	Threshold correction
Top	Base	(km/shots)	(meters/bins)	(background)
40.0	30.1	5.00/15	300/20	1
30.1	20.2	0.67/5	180/12	$\sqrt{5}$
20.2	8.2	1.00/3	60/4	5
8.2	−0.5	0.33/1	30/2	$5\sqrt{6}$
−0.5	−2.0	0.33/1	300/20	$\sqrt{15}$

Table 1), these range-invariant background noise measurements must by transformed into a range-dependent array before they can be properly incorporated into the threshold. For feature detection using averaged profiles, a single composite value P_{BKG} is computed as the root sum square of the background noise measurements made on board the satellite for each of the single-shot profiles used to construct the average. Then, P_{BKG} is converted to units of attenuated backscatter coefficients, and an array of $\beta'_{\mathrm{BKG}}(r)$ values is generated by scaling the converted P_{BKG} value according to the spatial averaging strategies employed in various altitude regimes. Doing so produces step discontinuities in the final threshold array that mirror the step changes in SNR seen in the CALIOP attenuated backscatter profiles.

The second noise category incorporated into the CALIOP threshold scheme encompasses all sources for which the magnitudes of the contributions are range dependent within a single-shot measurement. For the purposes of constructing the threshold, this range-dependent component is modeled as Poisson-distributed detector noise, proportional to the square root of the magnitude of the expected backscatter intensity. No attempt is made to incorporate an excess noise factor or to model the more complex Neyman type-A distribution (Teich 1981) that is appropriate for the photomultipliers used for the 532-nm channels.

In the attenuated backscatter coefficient regime, the range-dependent threshold is defined as

$$\beta'_T(r) = \beta'_{\mathrm{GMAO}}(r) + C_0\beta'_{\mathrm{BKG}}(r) + C_1\sqrt{\beta'_{\mathrm{GMAO}}(r_0)\beta'_{\mathrm{GMAO}}(r)}, \quad (4)$$

where, as above, $\beta'_{\mathrm{BKG}}(r)$ represents contributions from range-invariant noise that have been scaled to account for CALIOP variable onboard signal averaging. The range-dependent noise sources are quantified by the $\sqrt{\beta'_{\mathrm{GMAO}}(r_0)\beta'_{\mathrm{GMAO}}(r)}$ term, which expresses the relative change in the standard deviation of the molecular signal that can be expected with respect to the top of the profile (at r_0) and any other point ($r \geq r_0$). A complete derivation of this range-dependent noise term is given in Vaughan et al. (2005). For averaged profiles, the molecular model (β'_{GMAO}) used to compute the range-dependent components is constructed by simple averaging of the molecular models associated with each laser pulse used to create the averaged profile. Here, C_0 and C_1 are independent, empirically determined scaling constants used to balance the contributions of the various noise components. Higher values are appropriate for more conservative search strategies that minimize the occurrence of false positives but simultaneously risk increasing the number of false negatives. In general, both C_0 and C_1 are set to the same value, which is usually between 1.5 and 2.0, depending on the detection sensitivity desired and the amount of horizontal profile averaging done. (The exact values used are reported in the metadata included with all CALIOP data products.) For use within the SIBYL profile scanner, $\beta'_T(r)$ is transformed into the attenuated scattering ratio regime by dividing by $\beta'_{\mathrm{GMAO}}(r)$, so that

$$R'_T(r) = \frac{\beta'_T(r)}{\beta'_{\mathrm{GMAO}}(r)} = 1 + C_0\frac{\beta'_{\mathrm{BKG}}(r)}{\beta'_{\mathrm{GMAO}}(r)} + C_1\sqrt{\frac{\beta'_{\mathrm{GMAO}}(r_0)}{\beta'_{\mathrm{GMAO}}(r)}}. \quad (5)$$

Figure 4 shows $R'_T(r)$ applied to an averaged CALIOP attenuated scattering ratio profile drawn from the scene shown in Fig. 1. The step discontinuities introduced into $\beta'_T(r)$ by the variable onboard averaging scheme are clearly seen at altitudes of 20.2, 8.3, and −0.5 km.

d. Context-sensitive threshold adjustments

Within any profile, the uppermost feature top is identified at that altitude where the profile data first exceeds the threshold values over some minimum feature thickness. Base determination is less straightforward, and the correct identification will depend on the (as yet undetermined) effects of feature attenuation. In any clear-air region above the first feature detected, the expected value of $R'(r)$ is always 1. However, as is evident by inspection of Eq. (2), in the clear air immediately beneath the first feature, where $\beta_P(r)$ once again equals 0, the expected value of $R'(r)$ is a new, lower constant equal to the two-way transmittance of the feature just detected. The data below cloud base in Fig. 4 illustrate this change. Base identification is thus a two-step

process. While $R'(r)$ remains above $R'_T(r)$, the scanner increments the estimate of base altitude. When $R'(r)$ first drops below $R'_T(r)$, the scanner will continue to increment the base altitude estimate as long as $R'(r)$ is a decreasing function of range. Once the base altitude has been ascertained, the feature two-way transmittance T^2_{feature} is estimated by computing the mean attenuated scattering ratio over some fixed distance extending downward from the feature's lower boundary. Assuming $0 < T^2_{\text{feature}} < 1$, which cannot be guaranteed for optically thin layers in noisy profiles, to maintain the detection efficiency of the threshold in the search for additional features, all threshold values below feature base are multiplied by T^2_{feature} so that below the base $R'_{\text{updated}}(r) = T^2_{\text{feature}} R'_{\text{initial}}(r)$. An example of this context-sensitive threshold rescaling scheme is shown in Fig. 4, where the solid gray line represents the initial threshold computed according to Eq. (5), and the dashed blue line represents the updated threshold computed following the detection of the cirrus cloud between 15.6 and 12.4 km. We note too that although the text above references only the uppermost feature in a profile, the rescaling procedure described is applied in an identical fashion for every feature detected within a profile, so that, for example, below the base of the third feature found in a profile, $R'_{\text{updated}}(r) = T^2_{\text{feature3}} T^2_{\text{feature2}} T^2_{\text{feature1}} R'_{\text{initial}}(r)$.

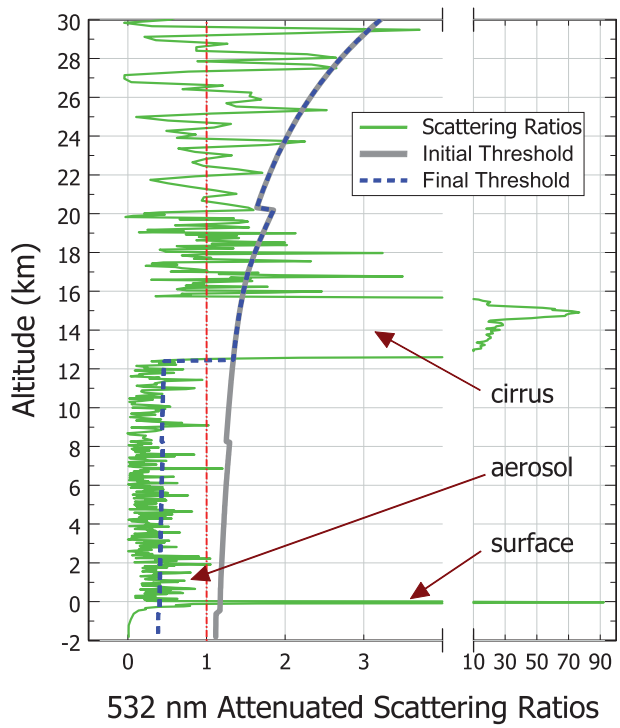

FIG. 4. Attenuating cirrus over aerosol in a 5-km averaged profile of attenuated scattering ratios measured on 13 Aug 2006 at 13.2°N, 115.2°E. The solid gray line shows the initial threshold array computed using Eq. (5). The dashed blue line shows the modified threshold recomputed after detecting the base of the cirrus cloud located at ~12.4 km. In the clear-air regions both above and below the cirrus, the value of the attenuated scattering ratios is seen to remain constant (albeit noisy). At the 5-km averaging level, the aerosol layer between ~2 km and the surface appears to be little more than a region of enhanced noise.

3. A nested multigrid averaging scheme

In the spatial analysis of the CALIOP backscatter data, we seek to detect features at the highest possible spatial resolution while simultaneously ensuring that the SNR within each feature is sufficient to meet the requirements of the extinction retrieval algorithm (Young and Vaughan 2009). Strong features, such as stratocumulus clouds, can be easily identified at the highest possible horizontal resolution (single shot), as, even in the presence of strong solar background noise, their backscatter intensities contrast sharply with the much weaker scattering from the ambient molecular atmosphere. At the other end of the intensity spectrum, very faint features, such as disperse aerosols and subvisible cirrus, cannot be reliably identified at high spatial resolutions, simply because the magnitude of the scattering from these features is often indistinguishable from the local molecular background and its associated noise. To locate the boundaries of these weaker features requires enhancing the contrast between the scattering from the feature itself and the scattering from contiguous regions of clear air; that is, we seek to reduce the degree of overlap between the histograms shown in Fig. 3. Traditionally, this enhancement is accomplished by averaging profile data both horizontally and vertically prior to searching for weaker features (e.g., as shown in Fig. 2). However, any such averaging scheme must be applied judiciously; otherwise, optically and/or meteorologically dissimilar features could be irretrievably commingled. Even within a single horizontally extensive feature (e.g., the cirrus deck centered at ~14°S in Fig. 1), capturing the natural spatial variability within the layer requires that averaging be limited to the minimum amount necessary. To detect spatial boundaries over the full range of feature backscatter intensities measured by CALIOP, SIBYL relies primarily on an iterated multiresolution averaging scheme. The function of the profile scanner described in section 3 is to identify feature boundaries within individual profiles. The function of the averaging engine described in this section is to enhance the contrast between features and clear air and then to repeatedly feed the newly constructed profiles back into the profile scanner.

Figure 5 shows a schematic of the SIBYL averaging engine. The engine consists of two loops or cycles. The

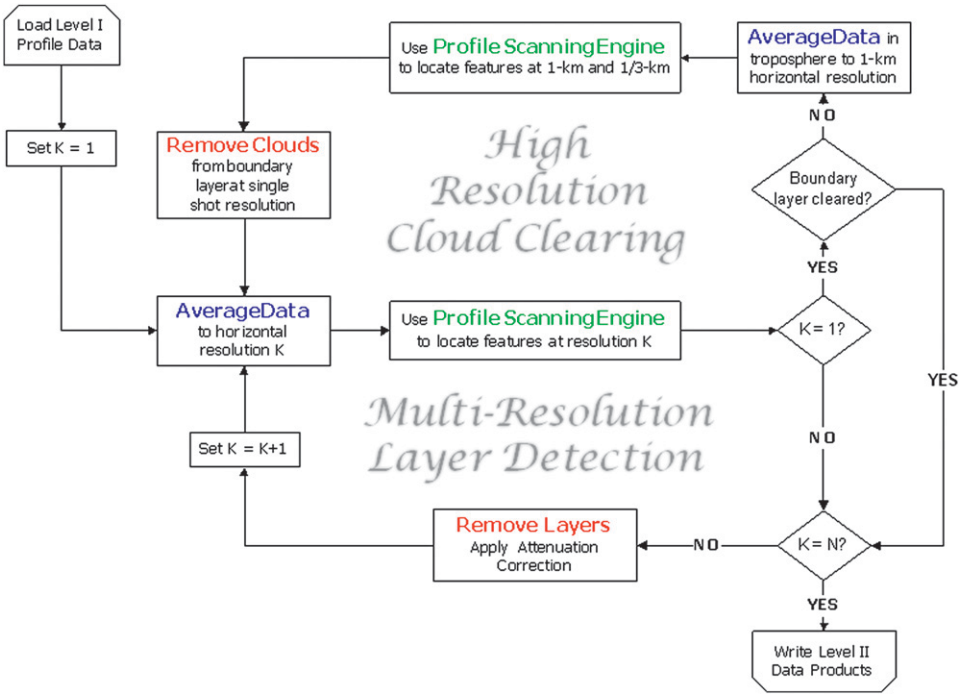

FIG. 5. SIBYL flow diagram showing the relationship between the lower MRLD loop (see section 4a) and the upper HRCCL (see section 4b). The current implementation uses three horizontal averaging resolutions (i.e., $N = 3$): 5, 20, and 80 km. Only the 5-km resolution ($K = 1$) is processed in the upper loop.

lower, multiresolution layer detection (MRLD) loop scans a sequence of profiles constructed using successively greater horizontal averaging, corresponding to successively coarser spatial resolutions. Features detected at higher spatial resolutions (less averaging) are removed from consideration in successive scans. Doing so allows subsequent additional averaging to increase the visibility (i.e., contrast) of weaker features while simultaneously reducing the risk of spatial smearing. The process is initiated by averaging the profile data to a 5-km horizontal resolution. As currently configured, the lower loop scans data averaged to three different horizontal resolutions: 5, 20, and 80 km. This progression of resolutions, which increases the backscatter SNR by a factor of 2 at each step, was chosen after consultations with members of the cloud modeling and passive sensor measurements communities. The initial 5-km resolution is determined by the maximum horizontal averaging distance used by the CALIOP onboard data averaging scheme. The upper high-resolution cloud-clearing loop (HRCCL) in the averaging engine decomposes the 5-km profiles into their constituent high-resolution profiles and rescans only those regions where a feature was initially detected at 5 km. The following subsections provide details on the operation of both cycles of the averaging engine.

a. Multiresolution layer detection

Because SIBYL is configured for a maximum averaging distance of 80 km, level 1 data are analyzed over 80-km intervals that contain an uninterrupted sequence of 240 profiles. The intermediate steps of the MRLD process (i.e., the lower loop in Fig. 5) are illustrated in Fig. 6, using the 20-km data segment that begins at 17.2°S, 114.3°E in Fig. 1. Figure 6a shows the attenuated scattering ratios and initial threshold array for the first of four consecutive profiles that have been averaged to a 5-km horizontal resolution. Contained within this profile are three features: a transparent cirrus layer between 15.4 and 12.4 km, a strong surface return at 0 km, and an aerosol layer between 2.1 km and the surface. As shown in Fig. 6b, the cirrus and the surface spike both rise well above the threshold and thus are detected and removed in the initial scan of the profile. Despite the adaptive threshold adjustments to account for the initial estimate of cirrus two-way transmittance (see the step change at 12.4 km), the aerosol layer goes undetected because it never exceeds the threshold for the required minimum distance. Immediately following the 5-km scan, those profile segments lying between any two features are assessed by using a sliding window technique that identifies the most likely regions of clear air therein (Vaughan et al.

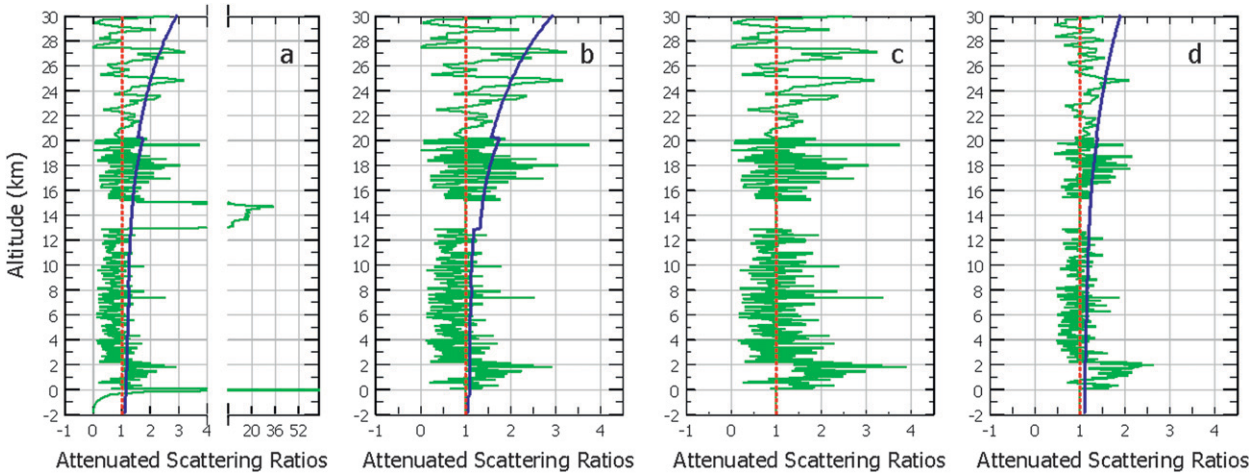

FIG. 6. The illustrated version of the SIBYL multiresolution layer detection process: (a) representative profile from Fig. 1 at 17.2°S, 114.3°E, which shows cirrus over moderate aerosol (green line) and the initial threshold array (blue line); (b) detection and removal of cirrus, in which in-line adaptations to the threshold magnitude at 12.4 km are insufficient to detect the aerosol layer; (c) correcting for cirrus attenuation enhances the aerosol backscatter intensity; and (d) average of four consecutive cloud-cleared attenuation-corrected profiles, in which aerosol now easily exceeds the revised threshold. All profiles shown have been averaged to a nominal horizontal resolution of 5 km (15 laser pulses). The vertical resolution of the profiles varies as a function of altitude, as described in Hunt et al. (2009).

2005). The attenuation resulting from each feature detected is estimated by using the mean attenuated scattering ratios from these "most likely to be clear air" regions, and the $R'(r)$ profile below feature base is renormalized to remove these attenuation effects. This step is illustrated in Fig. 6c, where $R'(r)$ between 12.4 and 2.1 km is now seen to be centered around a mean value of 1, which indicates clear air. The goal of this remove-and-renormalize procedure is to produce a profile representing the backscatter that would have been measured had the detected feature(s) not been present. However, full realization of this ideal is prevented by the loss of the photons backscattered and extinguished within each feature, and this loss of signal degrades the SNR in the data below. At this point, we also assume that the attenuation of all data below the lowest feature detected at the 5-km resolution is so severe that these data cannot be reliably used in subsequent coarser spatial averages. These data are therefore excluded when constructing subsequent coarser-resolution profiles.

After all of the sixteen 5-km averaged profiles within the 80-km horizontal data segment have been processed in a manner identical to the first, the resulting feature-cleared, attenuation-corrected profiles are averaged together to form a sequence of four 20-km averages. The first 20-km profile in this sequence is shown in Fig. 6d. Also shown is the profile scanner's updated threshold array, which automatically accounts for the improved SNR resulting from averaging. The aerosol layer between ~2.1 km and the surface now clearly exceeds the threshold and thus is detected at the 20-km averaging resolution. Following the invocation of the profile scanner, the analysis of the 20-km profile proceeds exactly as at 5 km: all detected layers are removed, their two-way transmittances are estimated, and the profile data below are renormalized to account for attenuation losses. When all four 20-km profiles have been analyzed, the resulting feature-cleared attenuation-corrected profiles are averaged together to form a single 80-km averaged profile. Applying the profile scanner to this 80-km average completes the processing cycle for the lower loop of the averaging engine.

b. *High-resolution cloud clearing*

For the ensuing layer typing and extinction retrievals, it is important that SIBYL identify homogeneous features that consist solely of a single class of scattering species (i.e., either cloud or aerosol). In the free troposphere and stratosphere, this task is made relatively easy by the larger horizontal spatial scales of the features that typically occupy these regions. However, the spatial distribution of targets in the planetary boundary layer (PBL) in particular, and the surface-attached aerosol layer in general, is significantly different. For example, the horizontal extents of fair weather cumulus can be on the order of 200 m or less (Lane et al. 2002), whereas the aerosol layers in which these clouds are embedded can span hundreds of kilometers (Anderson et al. 2003). To separate these strongly scattering fine-scale clouds from the more extensive, fainter aerosol layers, SIBYL integrates an HRCCL into its two-cycle averaging engine.

The operation of the HRCCL is similar in concept to the multiresolution search routine used to detect layers in the backscatter data acquired by the Geoscience Laser Altimetry System (GLAS; Palm et al. 2002)[1]. Rather than beginning the search for features at an intermediate spatial resolution, as does SIBYL, the GLAS algorithm initiates its detection scheme using the coarsest spatial resolution (0.25 Hz, which is equivalent to ~28 km horizontally). Subsequent searches at finer resolutions are conducted only in those regions where features were first identified in the coarse-resolution scan. Similarly, those CALIOP profiles for which a feature is identified at the 5-km averaging level are decomposed into finer horizontal averages of 1 km and, if possible, 1/3 km, and additional searches are conducted in those altitude regimes where features were first detected in the 5-km data. The CALIOP onboard averaging scheme dictates the resolutions at which these additional searches are conducted. From ~8.3 to ~20.2 km, the backscatter data are averaged to a 1-km horizontal resolution prior to being downlinked, hence the final 1/3-km search cannot be conducted above ~8.3 km.

The search for high-resolution features occurs on the initial pass through the MRLD loop, after the profile scanner has identified feature boundaries at 5 km but prior to the feature-removal and attenuation-correction steps. For all layers with top altitudes less than ~20.2 km, a secondary search is conducted at a 1-km resolution. The spatial and optical properties for features detected at 1 km are recorded in the data products. However, the feature-removal and attenuation-correction steps so critical to the performance of the MRLD scheme are not implemented for the 1-km search, as they are not required within the context of the HRCCL analyses.

If layers with tops below ~8.3 km are detected at 1 km, a final scan is conducted at the highest horizontal resolution (1/3 km) of the downlinked data. Once again, the spatial and optical properties for features detected at the 1/3-km resolution are separately recorded in the CALIOP data products. Unlike the searches at all other SIBYL resolutions, the 1/3-km scan is no longer looking for features of all types but is instead devoted solely to the detection of clouds and surface returns. To guard against the inclusion of aerosol layers in the 1/3-km results, an additional term is added to the search threshold. In the attenuated scattering ratio regime, the augmented expression for the threshold array becomes

$$R'_T(r) = \frac{\beta'_T(r)}{\beta'_{\mathrm{GMAO}}(r)} = 1 + C_0 \frac{\beta'_{\mathrm{BKG}}(r)}{\beta'_{\mathrm{GMAO}}(r)} + C_1 \sqrt{\frac{\beta'_{\mathrm{GMAO}}(r_0)}{\beta'_{\mathrm{GMAO}}(r)}} + C_2 \frac{\max(\beta'_{\mathrm{aerosol}})}{\beta'_{\mathrm{GMAO}}(r)}. \quad (6)$$

The $\max(\beta'_{\mathrm{aerosol}})$ term represents the largest aerosol backscatter coefficient that one could reasonably expect to occur in nature. Here, C_2, like C_0 and C_1, is an empirically determined scale factor and thus far C_2 has been set uniformly to 1. Consistent with a global distribution of aerosol backscatter coefficients derived by the authors from measurements made during the Lidar In-space Technology Experiment (LITE; Winker et al. 1996), $\max(\beta'_{\mathrm{aerosol}})$ is set to 0.0075 $\mathrm{km}^{-1}\ \mathrm{sr}^{-1}$ for nighttime data acquisition. The daytime value is adjusted upward to account for additional noise and possible diurnal increases in aerosol backscatter intensities. In retrospect, these values appear quite conservative, because CALIPSO has measured aerosol backscatter intensities far in excess of anything encountered during LITE.[2]

If, within any 5-km data segment, the HRCCL detects cloud tops below an altitude of 4 km at the 1/3-km resolution, a special high-resolution cloud-clearing routine is triggered. This fixed height of 4 km was chosen because it lies above the PBL, the most heterogeneous region of the atmosphere, for most CALIOP measurements and thus it represents a nominal, not actual, demarcation of the PBL. In those 1/3-km profiles for which clouds were detected within the nominal PBL, all backscatter data from the top of the highest cloud detected in the PBL downward to the end of the profile are removed. A revised 5-km averaged profile is then constructed using these newly cloud-cleared 1/3-km profile segments. This revised profile contains all data contained in the 5-km original profile except for the data at and below the tops of the clouds detected in the PBL. A second 5-km profile scan is then conducted using this newly created cloud-cleared profile, with the express purpose of attempting to detect the presence of any aerosols that may be surrounding the just removed clouds. Following this scan, the HRCCL terminates and the MRLD process resumes. Whenever this PBL cloud-clearing process is required, the results obtained in the initial 5-km scan are discarded and only those features (if any) detected during the

[1] The multiresolution layer detection loop described in section 4a is, on the other hand, radically different from the GLAS approach. Users who seek to compare results from the two datasets are urged to be cognizant of the differences inherent in the two schemes.

[2] The difference in the maximum aerosol backscatter intensities measured by the two systems is largely due to the limited dynamic range of the LITE receiver electronics (12 bits versus ~22 bits for CALIPSO).

second 5-km scan are reported in the CALIOP 5-km layer products. The cloud-cleared and re-averaged 5-km backscatter profile is used in all subsequent steps of the MRLD, and any layers detected during the second 5-km scan are treated exactly as described in section 3a. Operating together, the HRCCL and the MRLD enable the separate detection of small-scale clouds lying within faint aerosol layers that can only be detected after extensive horizontal averaging (up to 80 km).

c. *Eliminating false positives*

Each stage of the MRLD implicitly focuses on a specific, though not precisely defined, class of targets; that is, the 5-km scans are designed to identify the most robust features, whereas the 80-km scan seeks out the most tenuous layers. Furthermore, given the SNR requirements imposed by the CALIOP extinction retrieval, we specifically do not want to identify those weaker features that may be detectable at the 5-km resolution, because the poor SNR within these features will generate suboptimal retrievals of extinction and optical depth. SIBYL therefore imposes a lower bound on the layer-integrated attenuated backscatter, γ', on all features detected. Candidate features that do not exceed this minimum value are rejected, and thus these backscatter data are not removed in the feature clearing process. Weak, albeit genuine, features rejected at higher resolutions will presumably be retrieved on subsequent lower-resolution scans, where the γ' thresholds are necessarily less stringent. An additional and equally important benefit of the γ' threshold test is the reduction of the false positive rate. In effect, weaker features are required to exhibit some degree of horizontal persistence before they are positively identified by SIBYL.

d. *Systematic errors*

Given an atmosphere where the backscatter coefficients remain essentially unchanged in the CALIPSO along-track (i.e., horizontal) direction, the unwanted signal contributions from random noise can be reduced to arbitrarily small levels by increased signal averaging. Systematic uncertainties, however, impart biases to the signal that cannot be reduced by the SIBYL multigrid averaging scheme. There are two primary sources of bias errors in the feature detection process: input errors from the level 1 calibration coefficients used to convert raw backscatter profiles into attenuated backscatter coefficients, and self-inflicted errors resulting from erroneous estimates of the feature two-way transmittances made in the MRLD loop (the second of these is essentially a localized version of the first). Small systematic errors (e.g., $<5\%$) are generally benign at higher spatial resolutions because the random components typically dominate the profile noise budget. However, because successive averages will not reduce their contributions, bias errors will exert relatively more influence at coarser spatial resolutions.

4. Performance examples

During the development phase, SIBYL was rigorously and repeatedly tested using synthetic data generated by high-fidelity simulation software that models all components of the CALIOP transmitter and receiver and generates both the signal and the noise characteristic of analog detection of backscattered laser light (Powell 2005). Detailed results for many of these tests are given in the CALIPSO algorithm theoretical basis document (ATBD; Vaughan et al. 2005). When using synthetic data where all layer boundaries are known exactly, quantitative performance metrics can be developed to describe the SIBYL detection efficiency with respect to some underlying truth. When assessing the algorithm's performance on real-world measurements, however, access to this ''underlying truth'' is no longer available; the degree to which SIBYL succeeds (or fails) can only be determined by internal consistency checks (i.e., did the algorithm detect all the features that a trained observer would identify in an image of the data?) and by comparisons to other reliable measurements of the same set(s) of spatial properties. Brief explorations of both tactics are given in the following subsections.

a. *Qualitative assessment of algorithm performance*

Figure 7 shows all layer boundaries detected by SIBYL for the backscatter data shown in Fig. 1. The features are color coded according to the amount of horizontal averaging required for their detection. Consistent with the operation of the HRCCL, all $1/3$-km features are embedded in 1-km features, which in turn are embedded in features detected at some coarser resolution. Similarly, spatial inhomogeneities and/or attenuation effects can result in layers detected at the 5-, 20-, and 80-km resolutions lying in direct contact with one another. Examples of such ''vertically adjacent'' layers occur frequently in the lower altitudes between $\sim 30°$ and $\sim 20°$N, wherever an 80-km (black) feature lies immediately above layers detected at some higher resolution. Layers can also be wholly embedded within other layers that are detected at coarser resolutions. An example is seen at $\sim 18.1°$S, between 14 and 15 km, where a 20-km feature lies embedded within a second feature detected at 80 km.

Following the execution of SIBYL, a cloud-aerosol discrimination (CAD) algorithm is applied to all features detected. As explained in Liu et al. (2004, 2009),

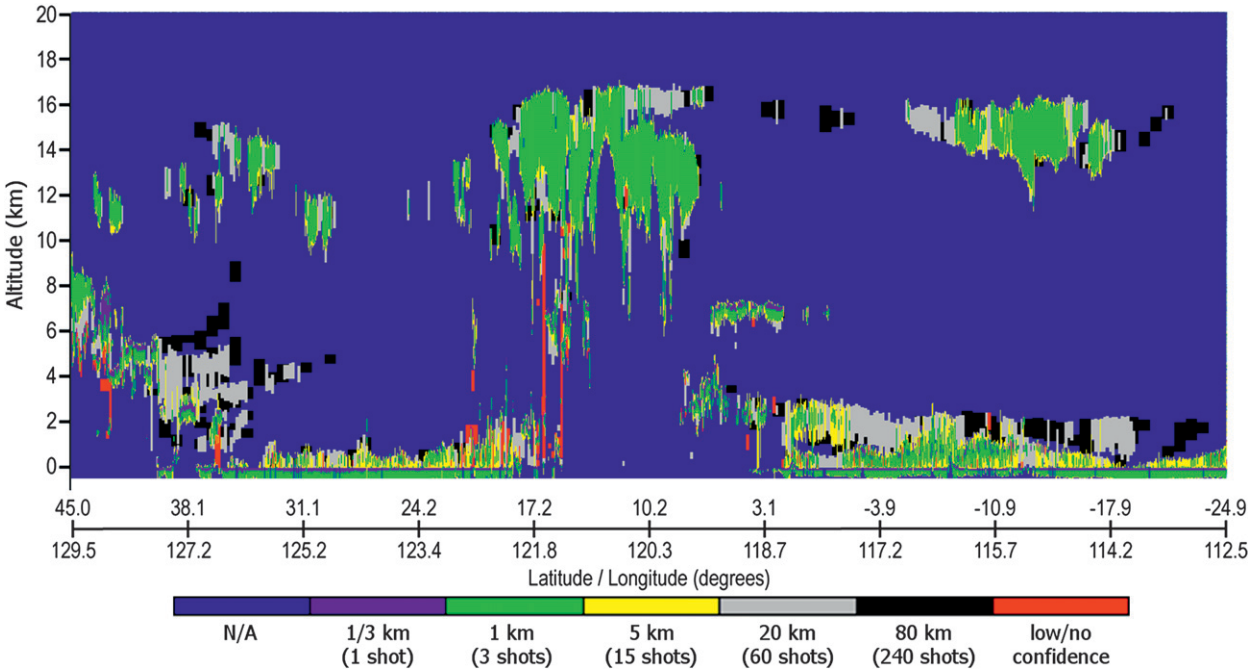

FIG. 7. Vertical feature mask image showing the location of all layers detected in the 13 Aug 2006 data shown originally in Fig. 1. Layers detected at 1/3, 1, 5, 20, and 80 km are shown in plum, light green, yellow, silver, and black, respectively. Low confidence layers (CAD scores less than 20) are shown in red, irrespective of the amount of averaging required for detection.

the CAD algorithm consists of a confidence function applied to the spatial and optical properties measured for each feature. The return value is a signed floating point number between −1.0 and 1.0. The sign of the result indicates feature type, whereas the magnitude quantifies the confidence with which the classification is made. In addition to the features rendered using the standard color palette, a small number of features in Fig. 7 are displayed in red. The red color indicates those features that, irrespective of feature type and the averaging required for detection, have very low confidence values (<0.2) or were otherwise flagged by the CAD algorithm as being suspicious. Though the CAD values were not originally designed to provide an assessment of overall feature viability, by happy coincidence we have found that uncertainties in the layer classification process (i.e., low CAD values) are strongly correlated with false positives in the detection process.

As demonstrated in Fig. 7, the detection performance of SIBYL is generally quite good. All major features of the scene are captured, including tentacles of lofted aerosol between 3 and 9 km centered at ~35°N, broken cumulus embedded in deep aerosol layers (2°–15°S), aerosol beneath clouds (e.g., 6°–18°S), multilayer cloud decks (e.g., 42° and 19°N), and subvisible cirrus (3°N). The number of false negatives is acceptably small and confined to weakly scattering layers with small spatial scales (e.g., 21°N at 16 km). The few false positives are readily identified by the CAD score assigned to the feature.

As can be seen by comparing Figs. 7 and 1, the data products that result from CALIOP's iterated multiresolution layer detection scheme are best understood as a sequence of coarse contour plots of backscatter intensity reported at regular intervals of 80 km horizontally. As would be expected from the "detect, remove, and reaverage" sequence of the MRLD, layers detected at finer resolutions are seen to be embedded within other layers detected at coarser resolutions. The histograms of backscatter intensity shown in Fig. 8 further illustrate this notion. These distributions were compiled using all data acquired during August 2006. To reduce the statistical side effects imparted by estimating (and sometimes misestimating) attenuation corrections for overlying layers, only the uppermost feature in any column was included in this analysis. The resulting 5-km histogram is seen to be asymmetric and multimodal, befitting the very broad distribution of scattering targets that can be detected at that resolution. At 20 and 80 km, the histograms appear roughly lognormal, albeit with a pronounced tail in the upper range of values. The enhancement in this region is thought to arise partly from those instances when intermittent, finer-resolution features differentially attenuate portions of lower lying, coarser-resolution features (e.g., a single 5-km feature

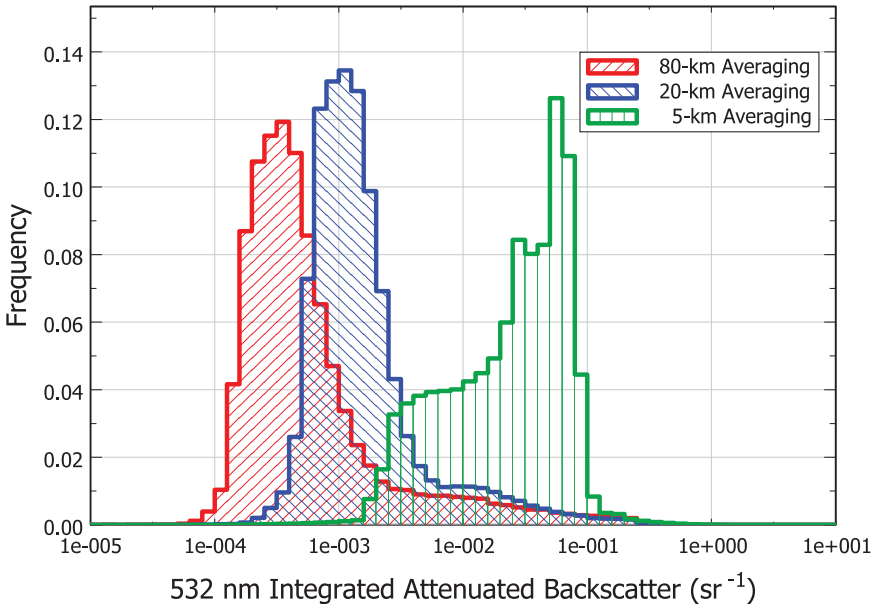

FIG. 8. Distributions of integrated attenuated backscatter measured at 532 nm for features detected at 5 (green), 20 (blue), and 80 km (red) during August 2006. Data were extracted from the 5-km cloud layer products and the 5-km aerosol layer products distributed in CALIOP data release 2.01. To reduce overestimates of γ'_{532} at 20 and 80 km, the histograms were compiled by using only the uppermost feature in each 5-km column.

at 12 km overlying a 20-km feature at 9 km). This partial obscuration leads to underestimates of backscatter intensity in the detection phase of the analysis. These underestimates are subsequently corrected (and, on occasion, overcorrected) during the extinction retrieval process (Young and Vaughan 2009).

b. *Interinstrument comparisons*

Comparing CALIPSO cloud-detection accuracies with those reported by other instruments is an inherently difficult task. Spatial and temporal mismatches, vertical and horizontal sampling rates, and differing SNRs, viewing geometries, and retrieval schemes, all complicate the picture. Ideally, CALIOP validation would be done using aircraft instruments such as the NASA airborne high spectral resolution lidar (HSRL; Hair et al. 2008) and the Cloud Physics Lidar (CPL; McGill et al. 2002) that flies on board the NASA high-altitude ER2 aircraft. Numerous flights of both validation platforms have already been conducted, and the initial comparisons are highly encouraging (e.g., McGill et al. 2007; Liu et al. 2008).

Figure 9 compares the layer boundaries retrieved by CALIPSO and CPL for data acquired 12 August 2006 as part of the CALIPSO-Cloudsat Validation Experiment (CC-VEX; McCubbin et al. 2006). Figure 9 shows the 532-nm attenuated backscatter coefficients measured by CALIPSO (top panel) and CPL (bottom panel) in the region of the exact coincidence of the two instruments (32.1°N, 75.4°W). The boundaries of the uppermost layer detected by each instrument are shown using white (tops) and maroon (bases) lines. At the time of coincidence, the ER2 was 37 m from the CALIPSO subsatellite point and thus, considering pointing uncertainties, was likely making measurements within several hundred meters or less of the CALIOP footprint. The temporal matching of the datasets is less exact; although both lidars flew identical flight tracks for the data segment shown, the ER2 required over 36 min to span the same distance that CALIPSO covered in 66 s.

Examining the uppermost layer at and immediately to the south of the coincidence point, the top and bottom panels of Fig. 9 show that both CALIPSO and CPL layer detection algorithms correctly identify the full depth of the feature. However, although the top altitudes are in excellent agreement, as seen in the middle panel, the base altitudes reported by the two instruments differ considerably. The largest disparities occur within the opaque regions of the layer, where the backscatter signal is totally extinguished, so that neither instrument can reliably detect the true base. In such cases, only an apparent base can be reported. The additional penetration (i.e., lower apparent base heights) reported by CALIPSO is the result of multiple scattering. In those areas where the uppermost layer is transparent and additional features are visible below, both instruments can make estimates of the true base. For

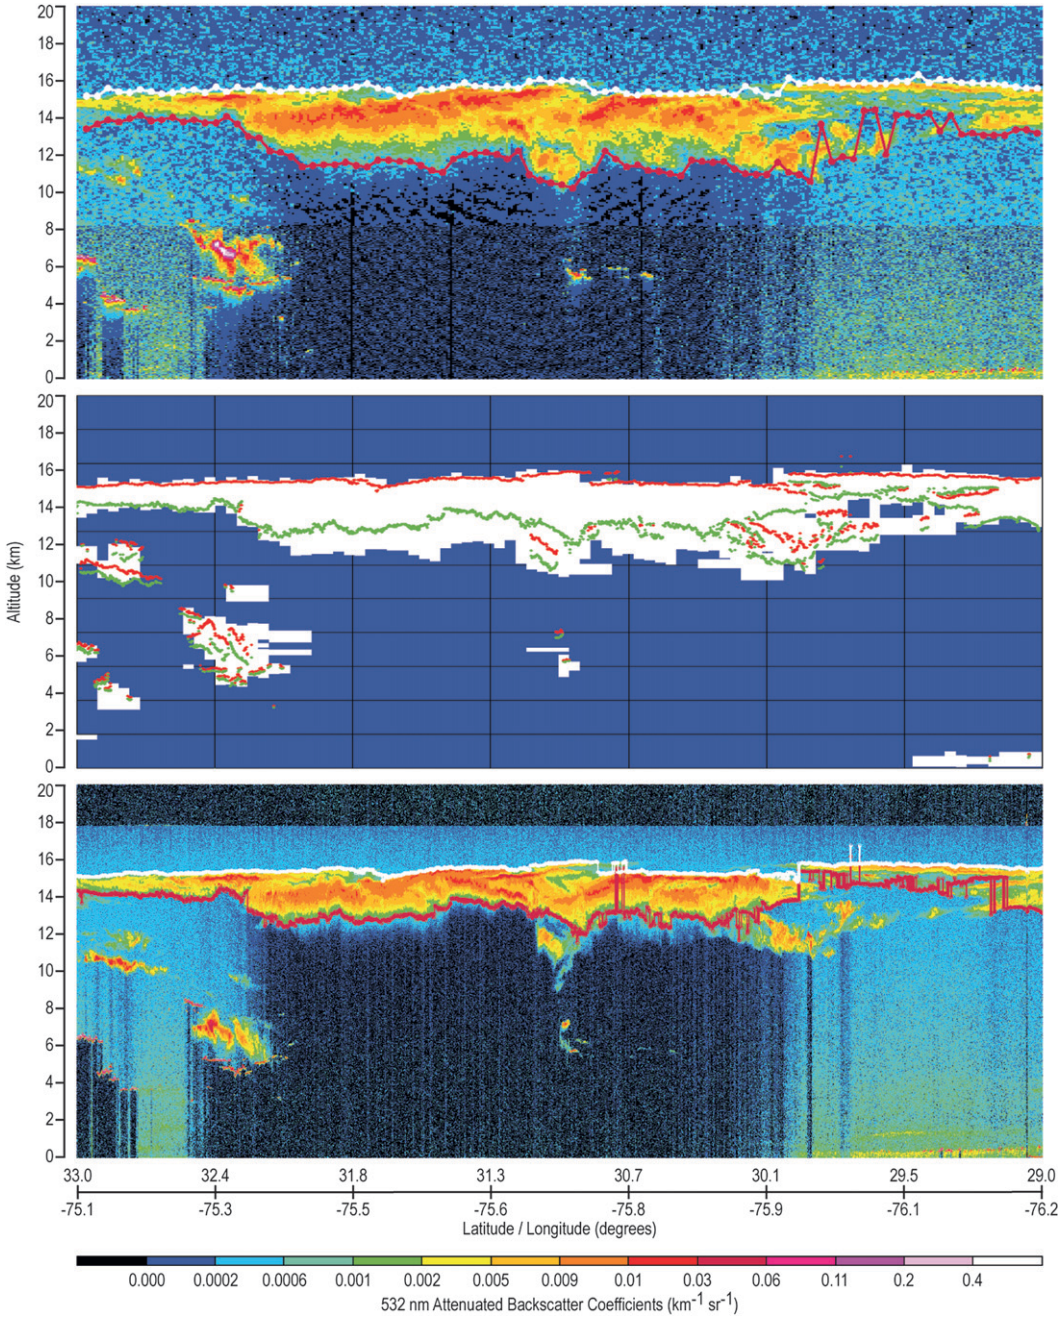

FIG. 9. (top) CALIPSO measurements from 12 Aug 2006 and (bottom) coincident CPL measurements. The top altitude of the uppermost layer detected is indicated with a white line and the base altitude is shown in maroon. (middle) The spatial extent of layers detected by CALIPSO is shown in white, while CPL feature altitudes are overplotted using red (tops) and green (bases) lines. Exact coincidence was at ~0705:50 UTC (32.1°N, 75.4°W).

these transparent regions, the agreement between the CALIPSO and CPL base altitudes is greatly improved.

c. *Statistical validation*

In addition to point-to-point comparisons of independent measurements, we are also investigating ways to validate the CALIOP retrievals via comparisons to the spatial and/or temporal statistics amassed from other space-based sensors. Several different approaches can be taken. First we can compare distributions of cloud heights derived solely from space-based lidars (i.e., from CALIOP, GLAS, and LITE). The drawbacks to this

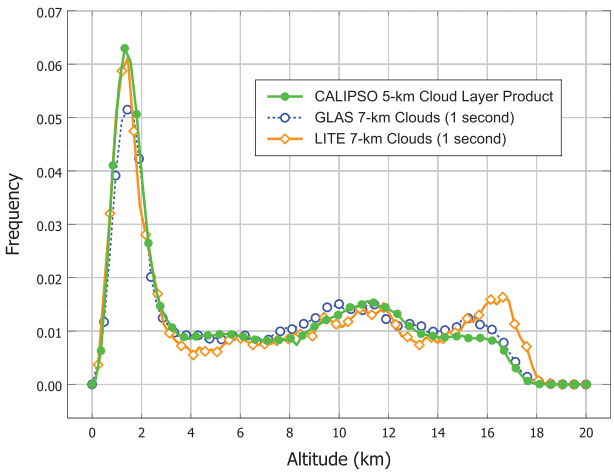

FIG. 10. Globally compiled (60°N–60°S) distributions of uppermost cloud-top height for CALIOP (green filled circles), GLAS (blue open circles), and LITE (orange diamonds). The LITE measurements were acquired over 50 h during September 1994. The GLAS dataset was acquired between 25 Sep and 18 Nov 2003. The CALIOP data are from 25 Sep through 18 Nov 2006. To ensure a more representative comparison to the LITE data, the data used for GLAS and CALIOP are restricted to latitudes between 60°N and 60°S. CALIOP and GLAS data were extracted from the publicly distributed data products. The LITE distribution was compiled by using a semiautomated, single-resolution forerunner to the fully automated CALIOP detection scheme described in this paper.

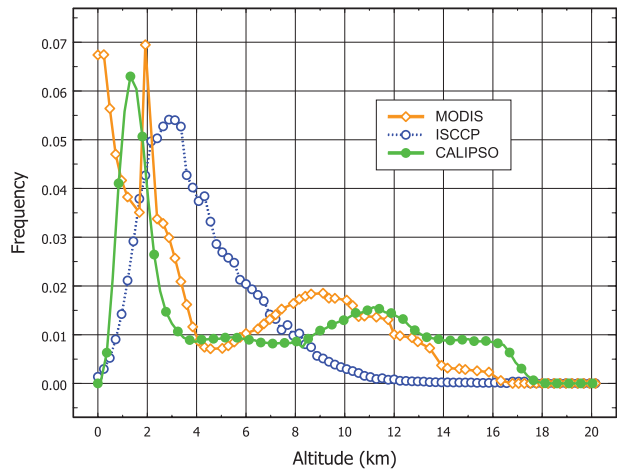

FIG. 11. Globally compiled (60°N–60°S) distributions of uppermost cloud-top height obtained from three different sources: CALIOP (green filled circles), MODIS (orange diamonds), and ISCCP (blue open circles). All data were acquired between 25 Sep and 18 Nov 2006.

approach are immediately obvious; for example, the three systems were never in orbit simultaneously, their orbit tracks and ranges of spatial coverage are not identical, and they have widely differing SNR. Nevertheless, because all three instruments acquired large volumes of data during similar periods of the calendar year (although at a mere 50 h, LITE acquired substantially less than the others) and because CALIOP and GLAS both fully encompass the more limited latitudinal range spanned by LITE, the data can be compiled and compared as seasonal statistics. An examination of the maximum cloud-top heights detected by each system is revealing. Figure 10 shows the globally compiled (60°N–60°S) probability density functions (PDFs) for the highest cloud top detected in each profile for all three space-based lidars. The similarities, especially between 2 and 14 km, are extremely strong. The differences are likewise predictable and easily understood. Because the LITE SNR is ~5 times higher than either GLAS or CALIOP, LITE can be expected to detect a greater frequency of high, faint cirrus. As seen in Fig. 10 between 16 and 18 km, this is indeed what happens (although LITE's targeted sampling strategy and much lower data volume may also contribute somewhat to the disparity).

A second statistical approach to validating CALIOP layer boundaries is to compare CALIOP retrievals to the cloud and/or aerosol heights derived from other space-based instruments. However, as is evident from the distributions shown in Fig. 11, it is not at all clear that CALIOP layer boundaries can be sensibly compared to traditional passive sensor retrievals. Recent comparisons of CALIOP data to collocated cloud-top heights derived from measurements made by the Moderate Resolution Imaging Spectroradiometer (MODIS; Platnick et al. 2003) show that MODIS underestimates cloud-top height by ~2.5 km (Holz et al. 2008). Comparisons to cloud-top heights estimated by the International Satellite Cloud Climatology Project (ISCCP; Rossow and Shiffer 1999) show similar discrepancies (see Fig. 11). Comparisons of this sort are confounded by the fact that both active and passive sensors use the same terminology to report fundamentally different quantities. Lidars are sensitive to abrupt changes in particulate concentrations and are especially well suited to the detection of small particles. Passive sensors, on the other hand, report an infrared effective radiating height that, even for optically dense clouds ($\tau > 8$), lies typically ~1.6 km below the cloud tops detected by lidar (Minnis et al. 2008). Differences in the active versus passive cloud-top comparisons also arise because of detection sensitivities. For example, recent validation studies conclude that MODIS is relatively insensitive to clouds with optical depths less than 0.4 (Ackerman et al. 2008), whereas CALIPSO is capable of detecting high, thin cirrus with optical depths of 0.01 or less, even during daylight operations (McGill et al. 2007). Such detection sensitivity issues are in no way confined solely to comparisons between active and passive sensors. Even when simultaneously

viewing the same region of space, radars and lidars frequently detect substantially different portions of the cloud and/or aerosol layers therein (McGill et al. 2004; Mace et al. 2008). In fact, the highly complementary nature of the CALIPSO and Cloudsat profiling abilities is a major contributor to the success of the A-Train observational strategy.

5. Conclusions and summary

In this work, we have presented an architectural overview and performance summary of the selective, iterated boundary location (SIBYL) algorithm used to detect cloud and aerosol layers in the CALIOP backscatter signals. The SIBYL scheme embeds a generic profile scanning engine within an iterated, multiresolution spatial averaging scheme. Each iteration of the profile scanning engine builds a range-varying detection threshold that scales automatically according to the magnitudes of the background noise and the expected molecular backscatter signal in the profile being examined. During execution of the scan, the threshold is further modified to account for the estimated attenuation of each feature encountered. By applying the multiresolution averaging scheme, SIBYL reliably culls increasingly fainter features from increasingly coarser spatial averages of the same 80-km horizontal data segments.

SIBYL has been deliberately designed and constructed as a one-size-fits-all solution to the (phenomenally complex) problem of detecting features of arbitrary backscatter intensity and arbitrary vertical and horizontal extent, irrespective of layer type. The same algorithm used to detect dense stratus clouds off the coast of California during nighttime operations is also used to detect subvisible cirrus in and around the intertropical convergence zone during daylight operations and Arctic haze and polar stratospheric clouds in all lighting conditions. There is a price to be paid for this kind of very general applicability: sometimes SIBYL misses weaker layers that might otherwise be detected by an algorithm more focused on identifying specific feature types (e.g., Pitts et al. 2007). Furthermore, although SIBYL is quite adept at detecting multiple layers, the laser backscatter signal becomes totally attenuated at particulate column optical depths of ~ 3, so that there are occasions where CALIOP cannot measure the full extent of the vertical column. Despite these few caveats, we find the implementation of SIBYL employed for the CALIOP measurements to be a robust, highly effective layer detection scheme. Comparisons of cloud and aerosol boundaries with near-simultaneous measurements made by CPL show excellent agreement for transparent layers. For totally attenuating features, the SIBYL base altitudes are consistently lower than those reported by CPL. This is expected, because the larger contributions from multiple scattering allow the CALIOP signal to penetrate deeper into optically thick layers.

In addition to the evidence derived from one-to-one comparisons with airborne validation measurements, the fidelity of the SIBYL retrievals is also demonstrated via statistical comparisons to datasets acquired by other space-based lidars. The distribution of highest cloud-top heights measured by CALIOP during the fall of 2006 is essentially identical to the distribution measured by GLAS during the fall of 2003. Both of these distributions are in turn quite similar to the LITE distribution acquired during September 1994, with differences resulting from LITE's superior SNR appearing, as expected, in the upper reaches of the troposphere. The results obtained by SIBYL from the fall 2006 CALIOP measurements are also compared to similar quantities derived from MODIS and ISCCP. Here, we again we see the expected, albeit stark, differences. Consistent with differing retrieval targets and detection sensitivities and with the results reported in recent studies, the SIBYL cloud tops are notably higher than those reported by passive sensors. Those readers seeking additional details about the inner workings of SIBYL will find a complete description of all facets of the algorithm and its prelaunch implementation in Vaughan et al. (2005).

Acknowledgments. The authors wish to acknowledge the many and substantial contributions of the members of the CALIPSO Lidar Science Working Group. We also extend special votes of thanks and appreciation to Sharon Rodier and Kam-Pui Lee, chief architects of the CALIOP level 1 analysis codes; to Steve Palm, Bill Hart, and Dennis Hlavka for many valuable discussions about space-based lidar data processing in general and layer detection techniques in particular; and to John Hair, Ray Rogers, Mike Obland, and Chip Trepte for their phenomenal and ongoing validation support. The ISCCP D1 data were obtained from the NASA LaRC ASDC. The GLAS GLA08 and GLA09 data were obtained from the National Snow and Ice Data Center (NSIDC) at the University of Colorado. The MODIS level 2 cloud data were obtained from the NASA Goddard Space Flight Center Level 1 and Atmosphere Archive and Distribution System.

REFERENCES

Ackerman, S. A., R. E. Holz, R. Frey, E. W. Eloranta, B. C. Maddux, and M. McGill, 2008: Cloud detection with MODIS. Part II: Validation. *J. Atmos. Oceanic Technol.,* **25,** 1073–1086.

Anderson, T. L., R. J. Charlson, D. M. Winker, J. A. Ogren, and K. Holmén, 2003: Mesoscale variations of tropospheric aerosols. *J. Atmos. Sci.,* **60,** 119–136.

Bloom, S., and Coauthors, 2005: Documentation and validation of the Goddard Earth Observing System (GEOS) data assimilation system—Version 4. NASA Global Modeling and Data Assimilation Tech. Rep. NASA/TM-2005-104606, Vol. 26, 165 pp. [Available online at http://gmao.gsfc.nasa.gov/pubs/tm/docs/Bloom168.pdf.]

Hair, J. W., and Coauthors, 2008: Airborne high spectral resolution lidar for profiling aerosol optical properties. *Appl. Opt.,* **47,** 6734–6752.

Holz, R., S. Ackerman, F. Nagle, R. Frey, S. Dutcher, R. Kuehn, M. Vaughan, and B. Baum, 2008: Global MODIS cloud detection and height evaluation using CALIOP. *J. Geophys. Res.,* **113,** D00A19, doi:10.1029/2008JD009837.

Hu, Y., and Coauthors, 2009: CALIPSO/CALIOP cloud phase discrimination algorithm. *J. Atmos. Oceanic Technol.,* in press.

Hunt, W. H., D. M. Winker, M. A. Vaughan, K. A. Powell, P. L. Lucker, and C. Weimer, 2009: CALIPSO lidar description and performance assessment. *J. Atmos. Oceanic Technol.,* **26,** 1214–1228.

Kay, S. M., 1998: *Fundamentals of Statistical Signal Processing: Detection Theory*. Prentice-Hall, 560 pp.

Lane, D. E., K. Goris, and R. C. J. Somerville, 2002: Radiative transfer through broken clouds: Observations and model validation. *J. Climate,* **15,** 2921–2933.

Liu, Z., and N. Sugimoto, 2002: Simulation study for cloud detection with space lidars using analog detection photomultiplier tubes. *Appl. Opt.,* **41,** 1750–1759.

——, M. A. Vaughan, D. M. Winker, C. A. Hostetler, L. R. Poole, D. Hlavka, W. Hart, and M. McGill, 2004: Use of probability distribution functions for discriminating between cloud and aerosol in lidar backscatter data. *J. Geophys. Res.,* **109,** D15202, doi:10.1029/2004JD004732.

——, W. Hunt, M. Vaughan, C. Hostetler, M. McGill, K. Powell, D. Winker, and Y. Hu, 2006: Estimating random errors due to shot noise in backscatter lidar observations. *Appl. Opt.,* **45,** 4437–4447.

——, and Coauthors, 2008: CALIPSO lidar observations of the optical properties of Saharan dust: A case study of long-range transport. *J. Geophys. Res.,* **113,** D07207, doi:10.1029/2007JD008878.

——, and Coauthors, 2009: The CALIPSO lidar cloud and aerosol discrimination: Version 2 algorithm and initial assessment of performance. *J. Atmos. Oceanic Technol.,* **26,** 1198–1213.

Mace, G. G., Q. Zhang, M. Vaughan, R. Marchand, G. Stephens, C. Trepte, and D. Winker, 2008: A description of hydrometeor layer occurrence statistics derived from the first year of merged Cloudsat and CALIPSO data. *J. Geophys. Res.,* **114,** D00A26, doi:10.1029/2007JD009755.

McCubbin, I. B., and Coauthors, 2006: Overview of the Multi-Aircraft CALIPSO-CloudSAT Validation Experiment. *Eos, Trans. Amer. Geophys. Union,* **87** (Fall Meeting Suppl.), Abstract A51E-0146.

McGill, M. J., D. L. Hlavka, W. D. Hart, V. S. Scott, J. D. Spinhirne, and B. Schmid, 2002: The Cloud Physics Lidar: Instrument description and initial measurement results. *Appl. Opt.,* **41,** 3725–3734.

——, and Coauthors, 2004: Combined lidar-radar remote sensing: Initial results from CRYSTAL-FACE. *J. Geophys. Res.,* **109,** D07203, doi:10.1029/2003JD004030.

——, M. A. Vaughan, C. R. Trepte, W. D. Hart, D. L. Hlavka, D. M. Winker, and R. Kuehn, 2007: Airborne validation of spatial properties measured by the CALIPSO lidar. *J. Geophys. Res.,* **112,** D20201, doi:10.1029/2007JD008768.

Minnis, P., C. R. Yost, S. Sun-Mack, and Y. Chen, 2008: Estimating the top altitude of optically thick ice clouds from thermal infrared satellite observations using CALIPSO data. *Geophys. Res. Lett.,* **35,** L12801, doi:10.1029/2008GL033947.

Omar, A. H., and Coauthors, 2009: The CALIPSO automated aerosol classification and lidar ratio selection algorithm. *J. Atmos. Oceanic Technol.,* **26,** 2048–2068.

Palm, S., W. Hart, D. Hlavka, E. J. Welton, A. Mahesh, and J. Spinhirne, 2002: GLAS atmospheric data products. NASA Goddard Space Flight Center Geoscience Laser Altimeter System Algorithm Theoretical Basis Document Version 4.2, 141 pp. [Available online at http://www.csr.utexas.edu/glas/pdf/glasatmos.atbdv4.2.pdf.]

Pitts, M. C., L. W. Thomason, L. R. Poole, and D. M. Winker, 2007: Characterization of polar stratospheric clouds with spaceborne lidar: CALIPSO and the 2006 Antarctic season. *Atmos. Chem. Phys.,* **7,** 5207–5228.

Platnick, S., M. D. King, S. A. Ackerman, W. P. Menzel, B. A. Baum, J. C. Riedi, and R. A. Frey, 2003: The MODIS cloud products: Algorithms and examples from Terra. *IEEE Trans. Geosci. Remote Sens.,* **41,** 459–473.

Powell, K. A., 2005: Development of the CALIPSO lidar simulator. M.S. thesis, Dept. of Applied Science, College of William and Mary, 228 pp.

——, and Coauthors, 2009: CALIPSO lidar calibration algorithms. Part I: Nighttime 532-nm parallel channel and 532-nm perpendicular channel. *J. Atmos. Oceanic Technol.,* **26,** 2015–2033.

Rossow, W. B., and R. A. Schiffer, 1999: Advances in understanding clouds from ISCCP. *Bull. Amer. Meteor. Soc.,* **80,** 2261–2288.

Stephens, G. L., and Coauthors, 2002: The CloudSat mission and the A-Train: A new dimension of space-based observations of clouds and precipitation. *Bull. Amer. Meteor. Soc.,* **83,** 1771–1790.

Teich, M. C., 1981: Role of the doubly stochastic Neyman type-A and Thomas counting distributions in photon detection. *Appl. Opt.,* **20,** 2457–2467.

Uttal, T., L. I. Church, B. E. Martner, and J. S. Gibson, 1993: CLDSTATS: A cloud boundary detection algorithm for vertically pointing radar data. NOAA Tech. Memo. ERL WPL-233, 28 pp.

Vaughan, M. A., D. M. Winker, and K. A. Powell, 2005: Part 2: Feature Detection and Layer Properties Algorithms. CALIOP Algorithm Theoretical Basis Document PC-SCI-202.01, 87 pp. [Available online at http://www-calipso.larc.nasa.gov/resources/pdfs/PC-SCI-202_Part2_rev1x01.pdf].

Winker, D. M., and M. A. Vaughan, 1994: Vertical distribution of clouds over Hampton, Virginia observed by lidar under the ECLIPS and FIRE ETO programs. *Atmos. Res.,* **34,** 117–133.

——, R. H. Couch, and P. McCormick, 1996: An overview of LITE: NASA's Lidar In-space Technology Experiment. *Proc. IEEE,* **84,** 164–180.

——, W. H. Hunt, and M. J. McGill, 2007: Initial performance assessment of CALIOP. *Geophys. Res. Lett.,* **34,** L19803, doi:10.1029/2007GL030135.

——, M. A. Vaughan, A. Omar, Y. Hu, K. A. Powell, Z. Liu, W. H. Hunt, and S. A. Young, 2009: Overview of the CALIPSO mission and CALIOP data processing algorithms. *J. Atmos. Oceanic Technol.,* in press.

Young, S. A., and M. A. Vaughan, 2009: The retrieval of profiles of particulate extinction from Cloud-Aerosol Lidar Infrared Pathfinder Satellite Observations (CALIPSO) data: Algorithm description. *J. Atmos. Oceanic Technol.,* **26,** 1105–1119.