

THOR – Cloud Thickness from Offbeam Lidar Returns

ROBERT F. CAHALAN, MATTHEW MCGILL, AND JOHN KOLASINSKI

Laboratory for Atmospheres, NASA/Goddard Space Flight Center, Greenbelt, Maryland

TAMÁS VÁRNAI

*Joint Center for Earth Systems Technology, University of Maryland Baltimore County,
Baltimore, Maryland*

KEN YETZER

Raytheon, SSAI, Lanham, Maryland

Prepared for *Journal of Atmospheric and Oceanic Technology*

June 2004

Corresponding author address: Robert F. Cahalan, Laboratory for Atmospheres,
NASA Goddard Space Flight Center, Greenbelt, MD 20771.

E-mail: Robert.F.Cahalan@nasa.gov

Abstract

Conventional wisdom is that lidar pulses do not significantly penetrate clouds having optical thickness exceeding about $\tau = 2$, and that no returns are detectable from more than a shallow skin depth. Yet optically thicker clouds of $\tau \gg 2$ reflect a larger fraction of visible photons, and account for much of Earth's global average albedo. As cloud layer thickness grows, an increasing fraction of reflected photons are scattered multiple times within the cloud, and return from a diffuse concentric halo that grows around the incident pulse, increasing in horizontal area with layer physical thickness. The reflected halo is largely undetected by narrow field-of-view (FoV) receivers commonly used in lidar applications. THOR – Thickness from Offbeam Returns – is an airborne wide-angle detection system with multiple FoVs, capable of observing the diffuse halo, detecting wide-angle signal from which physical thickness of optically thick clouds can be retrieved. In this paper we describe the THOR system, demonstrate that the halo signal is stronger for thicker clouds, and validate physical thickness retrievals for clouds having $\tau > 20$, from NASA P-3B flights over the Department of Energy/Atmospheric Radiation Measurement/Southern Great Plains site, using the lidar, radar and other ancillary ground-based data.

1. Introduction

Conventional wisdom is that lidar pulses quickly fade away after penetrating clouds to an optical thickness of about 2. Beyond this limit, multiple scattering by cloud particles increasingly spreads the distinct pulse into a diffuse halo. Since this halo lies outside the narrow field-of-view of most lidars, they are able to probe only thin clouds and the edges of thick clouds. As a result, much of the Earth's cloud cover remains outside the reach of lidar remote sensing.

However, empirical studies reveal a "radiative smoothing scale" in satellite-derived cloud reflectances (Cahalan and Snider, 1989), subsequently found to depend on cloud physical thickness (Davis et al., 1997). Further theoretical studies reveal that a lidar pulse entering a cloud spreads by multiple scattering, creating a bright halo that can be mined for cloud information (Davis et al. 1997; 1999). The studies show that in homogeneous clouds the size of the bright halo is proportional to $\sqrt{\frac{\Delta z}{\sigma(1-g)}}$, where Δz is the geometrical cloud thickness, σ is the extinction coefficient, and g is the cloud particle scattering asymmetry parameter. This relationship arises because photons, as they undergo random walks by scattering from cloud droplets, create wider halos in thicker clouds, because they travel farther without escaping through cloud base. In contrast, photons travel less far—and create a smaller halo—if σ and $(1-g)$ are larger, i.e. when the cloud is more dense and when the scattering tends to change photon directions more abruptly. Such results raise the possibility of using halo observations for retrieval of cloud geometrical thickness, as well as cloud internal properties.

Empirical, theoretical, and experimental advances spurred the simultaneous development of three wide field-of-view lidar systems that exploit the information content of bright halos. The THOR (THickness from Offbeam Returns) system described in this paper was first operated in ground-based up-looking configuration for mid to high clouds, then in an aircraft in down-looking configuration, well above cloud top. The up-looking WAIL (Wide-Angle Imaging Lidar) was developed at Los Alamos National Laboratory for ground-based observations, and having a wider field-of-view can image a lower cloud halo than a ground-based THOR (Love et al. 2001; Polonsky et al. 2004). The third system, an *in situ* cloud lidar, was built in Colorado for aircraft flights inside clouds (Evans et al. 2003). Initial results from these new systems confirmed the value of halo observations in cloud remote sensing. In addition, observations of diffuse halos have also proven fruitful in a wide range of other scientific disciplines including medicine, dentistry, glaciology, and astronomy (e.g., Groenhius et al. 1983; Sparling and Weiss 1993; Haines et al. 1997; Maffione et al. 1998; Predehl et al. 2000).

The THOR project started with a demonstration of the concept and feasibility of halo observations. In September 1996, the transmitter of a ground-based up-looking lidar at the NASA Goddard Space Flight Center (GSFC) was tilted slightly, so that instead of measuring the direct back scattering signal observed during normal operations, the receiver measured parts of the diffuse halo. The instrument detected a clear signal from the halo, thus demonstrating that current technology allows the development of halo-observing lidars (Davis et al. 1999; 2002).

Experiments with “clouds” of water suspensions of polystyrene microspheres as well as micro-fat globules in Liposyne confirmed the halo theory in a situation for which

cloud thickness and internal scattering properties were precisely known (Figure 1a). These suspensions provide approximately 1:1000 scale models of homogeneous clouds consisting of spherical particles having known size distributions. The observations confirmed theoretical predictions of halo properties, such as the spatial distribution of halo signals (Figure 1b) for a variety of physical thicknesses and opacities, thus demonstrating the potential for cloud property retrievals based on halo measurements.

Following the successes of the demonstration and the experiments, the THOR instrument was designed in 1998-1999 and built in 1999-2000. THOR's first ground-based measurements of mid- and high-level clouds took place at NASA GSFC in April 2001, and THOR's first airborne measurements of boundary layer clouds were collected over Oklahoma in March 2002.

THOR reveals the structure of diffuse halos by collecting time-dependent return signals not only from the immediate vicinity of the spot illuminated by its laser, as most lidars do, but also from seven additional rings around this spot (Figure 2). These observations are then used for retrieving the geometrical and optical thickness of optically thick stratiform clouds, as well as the vertical profile of cloud volume extinction coefficient. While the current study focuses on retrievals of geometrical cloud thickness, in future studies THOR halo observations may also be used for retrieval of microphysical cloud properties (see Roy et al. 1997).

This paper describes the THOR project as follows. First, Section 2 describes the instrument itself, and then Section 3 presents the methodology and sensitivity of cloud property retrievals. Next, Section 4 discusses the results from the March 2002 THOR

validation campaign conducted over central Oklahoma. Finally, Section 5 offers a brief summary and a few concluding remarks.

2. The instrument

a. System description

The basic structure of the THOR system is illustrated in Figure 3a, and its main parameters are listed in Table 1. Photos of the actual instrument inside the NASA P-3B aircraft are shown in Figures 4a (above deck control system) and 4b (below deck). This section discusses the workings of each component by following the path of a single pulse through the system.

When the control unit determines that it is time to emit the next laser pulse, it sends out two simultaneous signals. One signal goes to the data system computer, which then records the pulse's timing and starts archiving the photon counts reported by the detectors. The other signal goes to the laser unit, and causes the solid-state, fiber-coupled Nd:YALO laser to emit a single pulse of green light at 540 nm wavelength. The pulse then passes through an energy monitor, which determines the pulse energy by measuring the tiny fraction of light that gets scattered back into a sensor from the surface of a lens. The exact energy of the outgoing pulse is then reported back to the data system. Next, the pulse is collimated by a 4X beam expander, reducing its divergence to 325 μ rad (full angle). The pulse then leaves the system and passes through a window at the bottom of the aircraft, and travels until scattered by atmospheric particles.

The photons returning from the atmosphere pass through the same window, and are collected by the THOR telescope, a 5-lens custom telecentric design with 19.05 cm diameter primary lens, and 28.3 cm effective focal length, giving an f -ratio of 1.25. This telescope is precisely aligned with the outgoing laser beam, and the telescope edge is separated from the laser beam by just under 6 cm. This 6 cm separation means that the laser beam must travel about 1 km before it fully enters the central channel's field-of-view. Fortunately this 1 km minimum range does not pose serious problems, as THOR typically flies several kilometers above cloud top.

The telescope creates a 2.5 cm diameter image at its focal plane, about 37 cm behind the primary lens. The light forming this image is then collected by the most unique component of THOR, a custom-made optical fiber bundle that transports collected photons to an array of 10 photon-counting detectors.

The bundle consists of approximately 250,000 optical fibers, each 66 cm long, that guide photons to the appropriate detectors. Each fiber has a diameter of about 50 μm , except for a single 200 μm fiber that originates at the center of the focal plane. This central fiber captures photons coming from the central field of view—that is, the direct backscatter signal. The remaining fibers are organized into seven concentric annular rings, each collecting photons from a corresponding ring in the focal plane image. Each fiber in a given ring then leads to an associated detector, except for the outermost ring, whose ~150,000 fibers are approximately equally divided among three detectors, that each “see” one of three 120° azimuthal sectors of the outermost annular ring.

The reason for this arrangement is that in order to keep the signal level high even for the more distant, fainter parts of cloud halos, the outer rings (or channels) must collect light from larger fields of view. In particular, THOR compensates for the outward weakening of the halo signal by doubling the width of each successive fiber ring. This results in each ring covering an area four times larger than its inner neighbor. One consequence of this arrangement is that the outermost annulus, which covers 75% of the focal plane, and which can have up to 80% packing efficiency, contains up to 60% of all the fibers, or as many as 150,000 fibers. This many fibers cannot be connected to a single detector. Therefore the outermost ring is divided into three segments of up to about 50,000 fibers each, and fibers from each segment lead to one of 3 detectors. This yields a total of 10 THOR channels: the central spot, 6 annular rings, and 3 sectors of the outermost annulus (Figure 3c). We note, however, that most applications use data from the three outermost sectors combined together (Channels 8+9+10).

Upon leaving the optical fiber bundle, photons pass through spectral filters, each with about 7 nm bandwidth. These pass nearly all the returning lidar signal, but block most background illumination. This is crucial, because strong background illumination implies poor signal-to-noise ratio, especially in the outer channels, as discussed in more detail in Appendix A. The 7 nm effective bandwidth of the filters is sufficient for nighttime observations, but would pass too much sunlight for effective daytime measurements. As a result, THOR is currently operated only at night—but is expected to be extended to daytime operations as suitable ultra-narrow wide-angle filters become available.

Photons that pass through the filters are counted by single-photon-counting photomultiplier-tubes. The efficiency of these counters is about 10%. To keep the data volume manageable, the data system is set for a range-resolution of 30.8 m. In addition, the data system further reduces the data volume by accumulating the photon counts of each 30.8 m range bin for 500 subsequent laser pulses. Since the laser emits 1000 pulses per second, this makes the temporal sampling equal to 0.5 s. Considering the P-3B aircraft's cruising speed of about 550 km/h, this time-resolution corresponds to a spatial resolution of about 77 m.

Finally, the data system merges the accumulated photon counts with the pulse energy values coming from the energy monitor and with the navigational data coming from the P-3B aircraft (e.g., position, speed, pitch and roll angles). The results are then stored on a removable hard disc, and also displayed on a monitor for quick looks in real time.

b. Calibration

The THOR system performance was characterized through numerous component-level and system-level testing procedures. For brevity, we describe here only the tests that are most unique to THOR: the radiative calibration of its multiple fields-of-view. This calibration involves three stages.

First, laboratory experiments establish the relative calibration of the 10 THOR channels. A calibration sphere is attached to the THOR telescope to illuminate the front lens uniformly and isotropically (Figure 5a). Since increases of the fields-of-view from

one channel to the next (successively doubling in radius) are known, we can predict that with two exceptions, the photon counts of successive uniformly illuminated channels would increase by a factor of 4 in an ideal instrument. One exception to this ideal rule is the outermost ring, where the ratio of each 120° sector, Channels 8, 9, and 10, to Channel 7 is 4/3. The second exception is that the ideal ratio of Channel 2 to Channel 1 is 3, because Channel 1 covers a full circle, not a ring.

Figure 5b reveals that the ratios observed for the actual instrument are somewhat different from the above ideal values. The differences are especially large at high illumination levels, when the detectors become saturated (as they cannot keep up with the large number of incoming photons). Because of saturation, calibration data were collected for Channels 6-10 only at the lower illumination levels. The figure also reveals that Channels 2 and 5 are relatively less sensitive and register lower-than-expected photon counts.

THOR's deviations from ideal behavior are compensated by a relative calibration that anchors all channels to Channel 1. The calibration converts the observed raw counts, c_i^{obs} for $i = 2, 3, \dots, 10$, into calibrated counts c_i^{calib} using

$$c_i^{calib} = c_i^{obs} \cdot \prod_{j=2}^i \frac{r_j^{ideal}}{r_j^{obs}}, \quad \text{where } r_i^{obs} = \begin{cases} \frac{c_i^{obs}}{c_{i-1}^{obs}}, & \text{if } 2 \leq i \leq 7 \\ \frac{c_i^{obs}}{c_7^{obs}}, & \text{if } 8 \leq i \end{cases}, \quad (1)$$

where r_i^{ideal} are the ideal ratios and r_i^{obs} are the observed ratios from Figure 5b.

The second stage of calibration is an in-flight relative calibration of THOR channels. The idea is the same as in the laboratory, but this time the uniform illumination

is provided not by a spherical calibration lamp, but by the moonlight reflected from extended cloud fields. (This light is measured during the brief intervals between successive laser pulses.) Since this method allows calibration only at low illumination levels, and since it makes the uncertain assumption that all channels see statistically similar clouds when data are averaged over long flight segments, in-flight calibration is used only for checking whether or not the instrument's behavior changed substantially since the last laboratory experiment. The left side of Figure 5b shows that THOR performed quite consistently throughout the science flight discussed in Section 4.

The final, third stage of calibration provides absolute calibration for Channel 1, using the Rayleigh scattering signal returning from clear air. Since the calibration of all other channels is anchored to Channel 1 through Eq. (1), Rayleigh scattering can provide an absolute calibration for all THOR channels. Unfortunately, this absolute calibration can be performed only when THOR flies over cloud-free and aerosol-free air. As a result, THOR's cloud property retrievals, discussed below in Section 3b, are designed to work even if only relative calibration or no calibration at all is available, although retrieval uncertainties may increase in such situations.

3. Data processing methodology

The analysis of THOR data starts with an initial processing that prepares the data for scientific interpretation. This initial processing starts with radiometric calibration (see Section 2b), followed by the removal of the constant background illumination. Next, the navigational data coming from the P-3B aircraft is interpolated to 0.5 s time-resolution and is merged with the THOR photon counts. Finally, the data is flagged unsuitable for

cloud retrievals whenever the aircraft pitch and roll caused THOR to deviate from the nominal vertical pointing angle by more than 3°.

The second step of THOR data analysis estimates cloud properties by comparing the observations to look-up tables that contain simulated THOR data for a wide variety of clouds—and by selecting the case whose simulated data are most similar to the observations.

This section discusses the retrieval method as follows. First, Section 3*a* presents the methodology for generation of look-up tables. Next, Section 3*b* describes how the look-up tables are compared to observations in order to determine the best matches. Finally, Section 3*c* examines expected retrieval accuracy using theoretical calculations.

a. Look-up table generation

The look-up tables of simulated THOR data were generated using a suitably modified version of the Monte Carlo model that was thoroughly tested as part of the International Intercomparison of 3-dimensional Radiative Codes (I3RC) under the name UMBC5 (see <http://i3rc.gsfc.nasa.gov/>; Cahalan and Oreopoulos et al, 2004; Cahalan and Davies, 1999). The model obtains reflectance values using the method of local estimates (e.g., Marchuk et al. 1980). The simulated reflectances are then converted into photon counts detected by THOR using the method described in Appendix A.

The main challenge in creating look-up tables of simulated THOR data is to keep the computational time manageable while simulating enough photons to bring the random Monte Carlo noise low. (We typically simulate 10^8 - 10^9 photons per cloud.) To speed up

the simulations, the model uses a noise reduction technique similar to the one described in Section 5b2 of Barker et al. (2003). The technique smoothes out the forward peak of the scattering phase function at the last scattering event, for those photons that had already undergone several scattering events, including some wide-angle scatterings. This smoothing allows us to reduce the number of simulated photons by a factor of three without increasing the random noise or introducing significant biases.

The Monte Carlo model further reduces the simulation time by another factor of ten through using the same photon paths in simulations of many clouds. The scattering angles are identical for all clouds, while the path lengths between subsequent scattering events are resized according to each cloud's profile of extinction coefficient.

The Monte Carlo simulations and the equations in Appendix A are used to calculate THOR photon counts ($F(\rho, t)$) coming through a unit area at the cloud top surface—measured in units of *photon count* · m^{-2} · s^{-1} —as a function of radial distance (ρ) from the illuminated spot and photon travel time (t). (The travel time is obtained by dividing the total photon path length by the speed of light.) These values are then used to calculate photon counts for THOR's i^{th} channel ($F_i(t)$) through

$$F_i(t) = \int_{\rho_i^{\min}}^{\rho_i^{\max}} 2\pi\rho \cdot F(\rho, t) \cdot d\rho, \quad (2)$$

where ρ_i^{\min} and ρ_i^{\max} are the inner and outer limits of radial distance for the area viewed by channel i . Their values can be calculated as

$$\begin{aligned} \rho_i^{\min} &= z_{THOR} \cdot \tan \theta_i^{\min} \\ \rho_i^{\max} &= z_{THOR} \cdot \tan \theta_i^{\max}, \end{aligned} \quad (3)$$

with z_{THOR} being THOR's altitude above the cloud top, and θ_i^{\min} and θ_i^{\max} being the minimum and maximum viewing zenith angles of Channel i , respectively (see Table 1). We note that Eq.s (2) and (3) can accommodate changes in z_{THOR} , and so we did not need to redo the Monte Carlo simulations for each flight altitude discussed in Section 4.

To reduce the Monte Carlo simulation time, radiative transfer is calculated only for clouds that have a single geometrical cloud thickness of $\Delta z_0 = 2$ km. Once these simulations are done, we obtain results for other geometrical thicknesses (Δz) by rescaling the radial distance and time-delay values according to:

$$F_{\Delta z}(\rho, t) = F_{\Delta z_0}\left(\rho \frac{\Delta z_0}{\Delta z}, t \frac{\Delta z_0}{\Delta z}\right). \quad (4)$$

The idea behind this rescaling is that if Δz is reduced from 2 km to, for example, 500 m (while τ and the shape of the extinction profile remain unchanged), the extinction coefficient must increase four-fold everywhere inside the cloud. This means that the path lengths between subsequent scatterings get reduced by a factor of 4, which prompts a four-fold decrease in the radial distances the photons travel and in the time they spend inside the cloud. Because the rescaling goes much faster than the Monte Carlo simulations, the reduction in computational time is nearly proportional to the number of required Δz values. For example, if we needed results for cloud thicknesses ranging from 200 m to 1200 m at 50 m resolution, the rescaling method would allow a 20-fold reduction in simulation time.

Unfortunately, Monte Carlo simulations are much too slow to fill the required look-up-tables even if we use the efficiency enhancement methods described above. Therefore we perform Monte Carlo simulations for only three values of each independent

cloud parameter other than Δz (such as τ) and we use multidimensional cubic interpolation to fill the entire high-resolution look-up-tables. Unlike the techniques discussed above, the interpolation is not fully accurate and it can introduce errors into the retrievals of interpolated parameters (e.g., τ). Eventually, we plan to reduce the interpolation errors by performing direct Monte Carlo simulations for more than three values of each cloud parameter. We will also explore replacing the interpolation by fitting suitable analytical functions to match the way Monte Carlo results depend on each parameter.

Since we currently do simulations for three values of each cloud parameter, the total number of required simulations is 3^{N_p} , with N_p being the number of independent cloud parameters. Given our current computational resources, we can simulate up to about 2000 cloud cases (all with $\Delta z = 2$ km), which limits the description of clouds to no more than seven independent parameters (in addition to Δz). Our strategy has been to start by generating look-up tables for simple idealized clouds that can be described by only a few parameters, and then to proceed to more and more complex cloud structures. While there are no inherent limitations regarding the cloud structures one can use, our current cloud models have vertical extinction coefficient profiles consisting of linear segments (Figure 6). The results for each cloud model are stored in a separate look-up table, and the cloud retrievals discussed in Section 3*b* consider all available look-up tables.

We note that the simulations did not consider horizontal cloud variability, variations in cloud particle size, Rayleigh scattering, and surface reflection. These effects

are believed to be not significant for the relatively uniform cloud cases discussed in this paper, but will have to be considered in future applications.

In contrast, the calculations did need to consider Cirrus (or aerosol) layers occurring between THOR and its targeted boundary layer cloud, because such a Cirrus layer occurred during the validation campaign discussed in Section 4. Ignoring such Cirrus or aerosol layers cause THOR retrievals to overestimate the thickness of the targeted low-level cloud, because their particles scatter the narrow laser beam into a wide disc of downwelling photons, and this widening makes the observed halo of reflected photons wider as well. In addition, Cirrus or aerosol layers also scatter photons returning from low-level clouds, and this scattering further complicates retrievals.

Since the scattering properties of Cirrus and aerosol particles vary substantially, Monte Carlo simulations would need very long times to generate enough results for creating comprehensive look-up tables even if interpolations were used extensively. Instead, we use the Monte Carlo model to generate results only for the case of clear sky between THOR and the boundary layer cloud, and we are developing a separate scheme to incorporate Cirrus and aerosol effects into the clear sky results. Because this scheme has not yet been completed, Section 4 only discusses retrievals for flight segments with thin Cirrus cover.

b. Estimation of cloud properties

A sample THOR observation used for cloud retrievals is illustrated in Figure 7. The figure shows that the central channel observes intense direct backscatter from near

the cloud top, while the outer channels observe a fainter halo formed by multiple scattering deep inside the cloud. The signal in the outer channels is delayed because photons need time to reach the halo's outer portions, and it is stretched, because some photons meander more, while others follow more straight paths.

As it will be described below, the retrieval algorithm must adapt to wide variations in observational and environmental conditions. (For example, THOR's actual field of view varies with the aircraft altitude above the cloud.) To maintain maximum flexibility, the retrievals treat the spatial aspect and time-dependence of THOR observations separately.

The spatial aspect is characterized through each channel's contribution (C) to the overall detected photon count:

$$C^j = \begin{cases} \int_{t=0}^{\infty} c^j(t) \cdot dt, & \text{if absolute calibration is available} \\ \frac{\int_{t=0}^{\infty} c^j(t) \cdot dt}{\sum_{k=1}^{10} \left(\int_{t=0}^{\infty} c^k(t) \cdot dt \right)}, & \text{if absolute calibration is not available} \end{cases}, \quad (5)$$

where j and k are the THOR channels ($j, k \leq 10$), c is the calibrated photon count, and t is time (or equivalent vertical range).

The time-dependence of THOR observations is characterized through the width of time-intervals (Δt) that contain certain percentiles of the time-integrated return signal:

$$\Delta t_{a_i}^j = t_{a_i}^j - t_{a_{i-1}}^j. \quad (6)$$

Here, the a_i values ($i = 1, 2, \dots, N_i$) identify the boundaries of percentile intervals ($0 \leq a_i \leq 1$), and $t_{a_i}^j$ is determined by the equation

$$\int_0^{t_{a_i}^j} c^j(t) dt = a_i \int_0^\infty c^j(t) dt. \quad (7)$$

As mentioned earlier, the retrievals work by finding in the look-up tables the cloud whose signal is most similar to the observations. The degree of dissimilarity (D) is measured as

$$D = B \sum_{j=1}^{10} \left(W_j \frac{|C_{obs}^j - C_{sim}^j|}{C_{obs}^j} \right) + (1-B) \frac{\sum_{j=1}^{10} \left(W_j \sum_{i=2}^{N_i} \left(w_i \frac{|\Delta t_{a_i,obs}^j - \Delta t_{a_i,sim}^j|}{\Delta t_{a_i,obs}^j} \right) \right)}{\sum_{j=1}^{10} \left(W_j \sum_{i=1}^{N_i} w_i \right)}. \quad (8)$$

In this equation, B ($0 \leq B \leq 1$) determines whether the retrievals focus more on the spatial or the time dependence of THOR observations (i.e., on the first or the second term). The subscripts *obs* and *sim* identify observed and simulated photon counts, respectively. Finally, W_j and w_i are the weights given to each channel and percentile bin, respectively.

In practice, we set B to focus on the time-dependence of signals, because the time dependence does not require accurate calibration, and set the values of W_j and w_i to focus the retrievals on the outer channels and on the tails of time distributions, where cloud thickness has the largest influence (Figure 8). This has the added benefit of reducing the errors arising from uncertainties in cloud droplet size, as multiple scattering washes out the influence of details of the scattering phase function by the time photons reach the outer rings. On the other hand, focusing on the outer channels and on the tails can create difficulties if the observational noise is large or if the surface reflection can be mixed up

with reflection from the lower portions of the cloud. Thus the B , W_j and w_i values must be selected by considering several factors such as the spot size observed by each channel (determined by aircraft altitude above the cloud top), the calibration accuracy, the noise level (governed by the intensity of background illumination), and surface reflection (influenced by cloud altitude and surface albedo). Typical values are: $N_i = 6$, $B = 0$, $a_i = (0.40, 0.60, 0.80, 0.90, 0.95, 0.97)$, $w_i = (0, 1, 1, 1, 1, 1)$, and $W_j = (0, 0, 0, 0, 0, 1, 1, 1/3, 1/3, 1/3)$.

c. Retrieval accuracy

This section addresses two aspects of the accuracy of THOR cloud thickness retrievals. First, it briefly discusses the main sources of retrieval errors, and then it describes a simple method that was used to estimate retrieval uncertainties for the 2002 THOR validation campaign described in Section 4.

The most important sources of retrieval errors are as follows:

Vertical cloud variability causes retrieval uncertainties because the look-up tables consider idealized models of vertical variability (see Fig. 6), which do not fully represent the complexity of real clouds. Theoretical simulations and results from the THOR validation campaign show that using inappropriate extinction coefficient profiles can potentially lead to errors exceeding a few hundred meters. To reduce such uncertainties, the look-up tables are expanded to include more cloud structures (see Figure 6). During the validation campaign discussed in Section 4, vertical variability did not appear to cause uncertainties greater than 50 m.

Horizontal cloud variability causes retrieval uncertainties because our look-up tables were built for horizontally homogeneous, plane-parallel clouds. One source of

error is that the observed signal depends on cloud thickness in a non-linear way. Similarly to the plane-parallel bias arising in cloud optical thickness retrievals (e.g., Cahalan et al. 1994a; Oreopoulos and Davies 1998), this nonlinearity can cause systematic biases when observations used for a single retrieval include areas of different thicknesses. Fortunately, in the cases chosen here, these biases tend to be small: Theoretical simulations show them to be less than 10 m even if the thicknesses at the two halves of an observed area differ by 100 m. This implies that the biases were negligible ($\approx 2\text{-}3$ m) at the 2002 validation campaign, because the standard deviation of cloud thickness variations inside any retrieval segments was less than 20 m. 3D radiative effects associated with cloud inhomogeneity may cause additional retrieval uncertainties. These uncertainties will be examined in a separate study that will use simulated THOR data for 3D cloud fields generated by Large Eddy Simulations. We expect 3D uncertainties to be smaller than uncertainties caused by nonlinearity effects because they are smaller in optical thickness retrievals and in flux calculations (e.g., Cahalan et al. 1994b; Zuidema and Evans 1998). Though inhomogeneities were small in the cases used here, in general they are likely to be important in more convective clouds.

Variations in cloud droplet size cause only moderate uncertainties because the effects of any changes in droplet scattering properties are largely washed out by multiple scattering that creates the observed halos. For example, when test retrievals applied look-up tables generated for 10 μm effective radius to artificial THOR data that was simulated assuming a 4 μm effective radius, retrieval errors were around 30 m. Fortunately, we can easily reduce these uncertainties by generating look-up tables for several droplet sizes.

Cirrus and aerosol layers between THOR and the targeted cloud can cause retrieval errors exceeding several hundreds of meters, if their effects are not considered properly. We are currently developing a scheme for this purpose, but since the scheme is not yet functional, Section 4 discusses retrievals only for areas with very thin Cirrus. Since Cirrus makes the retrievals overestimate cloud thickness, and since the retrieved cloud thicknesses are, on average, 10 m higher than the thicknesses obtained from combining THOR Channel 1 with ground-based data, we estimate Cirrus effects to be small in the flight segments selected for retrievals.

Monte Carlo simulation noise unavoidably introduces some errors into the generated look-up tables, but these errors can be reduced to minimal levels by simulating sufficiently large numbers of photons. A comparison of retrieval results from simulations that used different random number seeds indicates that uncertainties caused by Monte Carlo noise are well below 10 m in our retrievals.

Observational noise also causes retrieval errors, especially for intense background illumination. In fact, this noise prohibits daytime THOR retrievals. At night, spatial averaging can reduce the noise to acceptable levels. We examined these errors by first estimating the observations' signal-to-noise ratios (see Appendix A), and adding a corresponding amount of Gaussian white noise to otherwise perfect simulated observations. The results indicate that if the laser works near full power and if observations are averaged over about 800 m (\approx field-of-view of THOR's outermost ring), retrieval uncertainties remain under 10 m even for full Moon. We note, however, that thermal control issues prevented the laser from operating at full power during the 2002

validation campaign, and so we estimate that observational noise caused retrieval uncertainties around 20–40 m even for retrievals based on 3–5 km long flight segments.

Table 2 provides a brief summary of the sources of retrieval errors discussed above. We note that, as mentioned in Section 3*b*, retrieval uncertainties can change significantly with cloud, environmental, and observational conditions. Thus the rest of this section describes a simple approach that was used for obtaining first-order uncertainty estimates for the 2002 THOR validation campaign.

The method estimates retrieval uncertainties based on how closely the observations fit the best-matching simulated cloud in our look-up tables: Retrievals are deemed more reliable if the fit between observed and simulated signals is good, and they are deemed less reliable if the fit is not so good. To make this decision, the method selects all simulated clouds in our look-up tables that satisfy two criteria:

- The deviation D (see Eq. 8) between a cloud and the retrieval's best guess cloud must be similar (to within 0.25%) to the deviation D between the retrieval's best guess cloud and the actual observations.
- The deviation D between a selected cloud and the retrieval's best-guess cloud cannot be lowered by changing only the thickness of the retrieval's best-guess cloud (while keeping its other parameters fixed).

Hypothetical retrievals for a cloud satisfying these two conditions and using look-up tables limited to the extinction coefficient profile of the actual retrieval's best-guess cloud would give results similar to the actual retrievals in two respects: Both would yield the same physical thickness, and the deviation D between input signal and the retrieved cloud's simulated signal would also be similar ($\Delta z_{\text{actual retr.}} = \Delta z_{\text{hypo. retr.}}$; $D_{\text{actual retr.}} \approx D_{\text{hypo. retr.}}$).

The main difference between the actual and hypothetical retrievals is that in the hypothetical retrievals, all retrieval errors come from differences in the vertical extinction profile, whereas in actual retrievals, a variety of sources can contribute to the uncertainties. Therefore the method must assume that the ratio of retrieval error to D is similar if the retrieval errors arise from vertical cloud variability or from other sources of uncertainty. Using this assumption, the method calculates the mean of the absolute value of hypothetical retrieval errors for all simulated clouds that satisfy the two criteria mentioned above, and report this mean error as our estimate for the actual retrieval's uncertainty. We plan to test the validity of the method's assumption using 3D cloud fields generated by Large Eddy Simulations.

4. THOR validation campaign

The first THOR airborne cloud observations took place during the March 2002 THOR validation campaign. During this campaign the NASA P-3B aircraft was dedicated solely to THOR measurements, and based at McConnell Air Force Base near Wichita, Kansas, from which it made repeated passes over the Department of Energy (DOE) Atmospheric Radiation Measurement (ARM) program's Southern Great Plain (SGP) site in central Oklahoma. This site was chosen because of its collection of ground-based instruments that provided a wealth of information for validating THOR cloud thickness retrievals. (See <http://www.arm.gov/>) The idea was to compare halo-based cloud thickness estimates retrieved from THOR with "true" thickness values that were obtained as the difference between the cloud top altitude measured by the central channel and the cloud base altitude obtained by use of ground based instruments.

The campaign's first two flights took place on the nights of March 19 and 20, respectively. These engineering flights were dedicated to testing the instrument behavior during THOR's first airborne operations. The flight data indicated that most of the THOR system performed well, but also revealed some problems in system software and in laser thermal control. After these problems were fixed, a third flight took place on March 24. Clear-sky conditions during this flight allowed for thorough system testing. THOR performed well, and it was clear that the software and thermal problems seen during the first flight had been adequately addressed.

The campaign's main science flight took place on the next day, March 25, 2002. The flight departed base around 04 UTC (10 PM local time) and lasted for about 5 hours. The synoptic pattern during this flight was dominated by a strong high pressure system containing a very cold and dry air mass dropping southward out of Southern Manitoba. The southern extent of this air mass was represented by a strong frontal boundary that stretched from Southern New Mexico to Northeastern Oklahoma and then eastward toward the Ohio Valley. Across this boundary, from the panhandle of Oklahoma to southeastern Oklahoma, temperature ranged from 0° C to 20° C, and specific humidity ranged from 3.0 g kg⁻¹ to 9.5 g kg⁻¹. Between 00 UTC and 04 UTC several waves of low pressure had developed along this boundary, and by 06 UTC one of these lows was centered over northeastern Oklahoma. Over the next few hours this low slowly moved toward the northeast while the trailing cold front dropped southeastward across Oklahoma.

During the flight period, the vicinity of the ARM SGP site was covered by two distinctive cloud layers (Figure 9). First, a low-level stratus cloud covered the sky

completely, thus providing an excellent target for the THOR halo-based retrievals. The cloud base varied from 200 to 500 m above the ground (see below), and the cloud thickness ranged from 500 to 1000 m according to the cloud radar. For a brief period around 06 UTC, the cloud produced a light drizzle around the ARM site. The sounding data from 23.5 UTC, on March 24, and that from 11.5 UTC, on March 25 (Figure 9b) (the 5.5 UTC sounding was unavailable) showed that the cloud layer mean temperature had dropped from 10° C to about -3° C. Since this latter sounding was obtained about 3 hours after THOR data collection had ended, and since no observed melting layer appears in the ground based cloud radar measurements, it is reasonable to assume that this layer consisted primarily of liquid water droplets. Figure 9c shows that the ARM Micropulse Lidar (MPL) clearly detected the cloud base, though it could not provide information on the clouds' inside and top, as cloud droplets quickly scattered the MPL laser pulses outside its field-of-view. In contrast, the ARM Millimeter-Wave Cloud Radar (MMCR) provided good information on cloud top and internal structure (note the presence of two joining sub-layers in Figure 9d), but it could not clearly identify the cloud base, because it is not sensitive to small droplets that dominate near cloud base, and because it is not able to distinguish signals reflected by numerous small cloud droplets or by a few large drizzle drops.

The second cloud layer was a layer of highly variable Cirrus clouds approximately 5.5 km above the ground. Figure 9 shows that for the most part, this cloud was optically too thin to produce a significant halo, and that the THOR laser pulses easily passed through it. Still, the Cirrus cloud complicated the retrievals for the low-level stratus cloud: It scattered parts of the well-focused downwelling laser pulses into a wider

cone, and this widened the halo observed at the low cloud. This means that to correctly interpret the halo observations, cloud retrievals need to properly account for the Cirrus spreading effect. Although flying below the Cirrus cloud can certainly eliminate the retrieval uncertainties caused by the Cirrus, low-altitude observations pose different challenges. Most importantly, the fixed viewing angles of THOR (see Table 1) imply that if it flies at lower altitudes, THOR observes smaller areas of a given cloud top, so in that case even its outermost channels may see only the relatively inner portions of the bright halo. (Methods for obtaining close-range halo observations using even wider fields-of-views are presented by Polonsky et al. (2004).)

To provide observations from a variety of conditions, THOR made passes over the ARM site at several altitudes. The first flight segments were flown at 8 km (well above the Cirrus), then the aircraft dipped to 5.0 km (just below the Cirrus). Subsequently, the aircraft rose to 6.2 km (inside the Cirrus), and to 7.3 km and 8.5 km (above the Cirrus). We note that the Cirrus cloud thinned out several times during the flight, and this allowed relatively more simple retrievals for a few segments even when THOR flew at high altitudes.

Figure 9 also shows that the central channel could clearly detect the stratus cloud's cloud top altitude. To validate the THOR halo-based cloud thickness retrievals, we obtained an independent estimate of cloud thickness as the altitude difference between the cloud top observed by THOR and the cloud base observed by surface instruments at the ARM SGP site. Of the instruments available, including the cloud radar, micropulse lidar and ceilometer, it was determined that data from the Micropulse Lidar (MPL) gave the most reliable cloud base. This conclusion is based on Figure 10 which shows the time

series of cloud base height from several sources, including the estimated value from the MPL and ceilometer instruments and from the calculation of the lifting condensation level (*LCL*) that was derived using tower data of temperature (*T*) and dew point temperature (*T_d*) at 25 and 60 m above ground level.

The *LCL* is calculated as the altitude at which an air parcel rising dry adiabatically from its current level would become saturated:

$$LCL = \frac{T - T_d}{\gamma - \gamma_d}, \quad \text{where} \quad \gamma_d = \frac{gT_d^2}{\epsilon l_v T}. \quad (9)$$

$\gamma = 9.8 \text{ K km}^{-1}$ is the dry adiabatic lapse rate, γ_d is the lapse rate of dew point temperature, $g = 9.807 \text{ m s}^{-2}$ is the gravitational acceleration, $\epsilon = 0.622$ is the gas constant ratio of dry air to water vapor and l_v is the latent heat of vaporization. The main assumption concerning the *LCL* calculation is that the mixed layer is "well mixed" (i.e., the potential temperature and humidity are approximately constant with height). During our flight, the vertical profile of wind speed obtained from both soundings and wind profilers near the central facility indicate that a large magnitude of shear existed throughout the subcloud layer, suggesting that the well-mixed assumption is valid for this particular event. In addition, temperature sounding data (Figure 9b) also indicate a well-mixed sub-cloud boundary layer, as *T* follows the dry adiabat and *T_d* remains nearly constant with altitude. Figure 10 clearly shows that the MPL cloud base height is very close to the derived *LCL*'s from both the 25 and 60 m levels whereas the ceilometer is on average about 100 m too high.

We also made use of temperature and dew point temperature surface data from five nearby Oklahoma Mesonet sites to determine the cloud base height (i.e., *LCL*) time

series in these areas (Figure 11). The figure reveals that as the cold front passed through the area, the cloud base rose significantly. As indicated by grey bars in the figure, Cirrus-free THOR retrievals could be performed at various stages of the cold front's passage. Because the cloud base rose most sharply at the central facility and at the two mesonet stations West from it (BREC and MEDF), the cloud base height exhibited a great deal of spatial variability during the flight period, with a maximum change of up to 350 m from East to West. The predominantly NorthEast-SouthWest orientation of flight tracks allowed THOR to observe clouds having a wide range of physical thicknesses (Figure 12). As mentioned earlier, the aircraft passed near the ARM central facility numerous times at various altitudes. Most flight paths passed well within 4 km from the ARM site (this distance is indicated by a dashed ellipsoid in the map's projection), although direct overflight was avoided to prevent interference with ARM observations. Cirrus-free THOR retrievals (indicated by grey bars in Figure 11) were performed at the flight segments shown as white in Figure 12.

We estimated the cloud base at THOR's constantly changing location by using a 2D interpolation scheme that ingested cloud base values at the central facility and at the five mesonet stations. Figure 13 displays the interpolated cloud base values, along with cloud top altitudes inferred from THOR Channel 1 direct backscatter data. The high-frequency variability in cloud base (and to some degree, cloud top) is caused by the aircraft flying back and forth between the thicker clouds in the East and the thinner clouds in the West. The grey bars indicating Cirrus-free retrieval segments (also appearing in Figure 11) show that retrievals were performed at a variety of cloud base

and top altitudes. Finally, the bold symbols at the top of the figure show that THOR approached the central facility to less than 4 km at all stages of the flight.

Both estimating the cloud base at the extended facilities and interpolating these values to THOR's location involve considerable uncertainties. As a result, we regard the cloud thickness estimates that combined THOR cloud top measurements with ARM cloud base measurements as most accurate when THOR flew near the central facility. The other main sources of errors in estimating cloud thickness by combining THOR Channel 1 with ground-based data fall into three categories: (1) Errors in cloud top height. (2) Errors in cloud base height at the ARM central facility. (3) Errors in cloud base height at the five Oklahoma mesonet stations. Table 3 shows a brief summary of these errors.

The first source of error in cloud top height is uncertainty in aircraft flight altitude. Because of the failure of other aircraft instruments, flight altitude could be estimated only from the aircraft's pressure altimeter, which can be off by as much as 200 m. To remove any such biases, we compared cloud top heights calculated from THOR's channel 1 data to those calculated from cloud radar (MMCR) data at times when THOR flew within 4 km of the central facility. Although in some situations the cloud radar's insensitivity to small droplets makes it difficult to detect the cloud top, the MMCR cloud top heights appear quite accurate for this particular night. (A comparison of the central facility's sounding data at 11:30, 17:30, and 23:30 UTC shows good agreement with MMCR cloud tops.) The average bias, calculated as the mean difference between the cloud top heights from THOR and MMCR, was determined to be -135 m. We removed this bias by lifting the aircraft altitude and cloud top altitudes by 135 m.

After removing the systematic bias, the main uncertainty in cloud top height is caused by the range resolution of instruments. THOR has a bin resolution of 30.8 m while the cloud radar has a bin resolution of 45 m. Since the maximum error for either instrument is one half the respective range resolution, the total combined uncertainty is around 40 m.

Cloud base height uncertainty at the central facility was estimated to be half of the range resolution of the micropulse lidar (~15 m), except for times between 5.5 and 6.7 UTC. In this time period both the micropulse lidar and the cloud radar show the possible existence of drizzle (Figure 9), which can cause these instruments to underestimate cloud base height. In fact, Figure 10 clearly shows the lidar observing a cloud base that is much lower than a cloud base calculated from either the ceilometer or from the tower and ground data. This, however, may not mean an error in MPL cloud base estimates, because drizzle can saturate part of the subcloud layer through evaporation, thereby extending the cloud base downward. Therefore if one assumes that the 60-meter tower observation represents a maximum value for cloud base height then the uncertainty can be estimated as the mean difference between this value and that derived from the lidar (~100 m).

Cloud base height uncertainty at the surrounding mesonet sites comes from two main sources: instrument uncertainty and the assumption of a well-mixed boundary layer. Instrument uncertainty (~ 0.5°C in temperature and ~ 2% in relative humidity) leads to errors of up to 50 m in the lifting condensation level (LCL), with the mean error being around 35 m. To estimate the uncertainty caused by the assumption of a well-mixed subcloud layer, we refer back to Figure 10. This figure shows that the LCL calculated from EBBR data is too low when compared with the tower and MPL results (except

when drizzle occurs). The difference is usually less than 50 m, except during a sharp rise in cloud base between 7 and 7.2 UTC, when the difference approaches 100 m. It is possible that the drizzle created this increase in cloud base and that there is a lag in surface response to this forcing. Thus the uncertainty is estimated to be 100 m for periods of drizzle and sharp rise in cloud base, and it is estimated to be 35 m for all other times. We note that only two mesonet sites west of the central facility experienced drizzle with the associated sharp rise in cloud base (Figure 11).

These cloud thicknesses (obtained by subtracting ARM estimates of cloud base from THOR Channel 1 estimates of cloud top) were then used to validate the THOR halo-based retrievals. Figure 14 shows this comparison for 16 approximately 5 km long flight segments where the Cirrus cloud was very thin. (Our current look-up tables do not consider Cirrus effects, so that thick Cirrus makes THOR retrievals overestimate cloud thickness by hundreds of meters. Improved tables with Cirrus effects included are expected to add many points to Figure 14.) These 16 segments were among 21 segments selected *a priori*, by visually examining Channel 1 data for information on Cirrus thickness. For the remaining 5 of the 21 segments the algorithm returned a result of “no valid retrieval”, as the deviation D (see Eq. 8) exceeded a threshold of 3%. It appears that despite our guess based on a visual inspection of THOR Channel 1 data, these 5 segments were contaminated significantly by Cirrus. We note that all retrievals used the retrieval parameters listed in Section 3b, except that the weight of Channel 6 (W_6 in Eq. 8) is set to zero for 5 km flight altitude. This was done because the low flight altitude made Channel 6 observe only inner parts of the diffuse halo, which are not sensitive to cloud thickness (see Figure 8). For most retrieval segments, the agreement between THOR observations

and simulated look-up tables proved best for 3-segment extinction coefficient profiles (Figure 6e); The only exceptions were segments #12-14, for which the profile shown in Figure 6f worked better. The mean of the absolute value of differences between thicknesses retrieved from THOR data and thicknesses obtained by combining THOR Channel 1 with ground data is 30 m.

5. Summary

This paper describes the new airborne lidar instrument called THOR, its data processing algorithms and retrieval methodology, and illustrates and validates these using data from the March 2002 THOR Validation Campaign. THOR was built to probe the multitude of clouds that have optical thicknesses larger than about 2. The inside of these clouds is beyond the reach of conventional lidar, because cloud particles scatter the distinct laser pulses into a diffuse halo that lies outside the narrow field of view of conventional lidar. Multiple wide fields-of-view allow detailed observations of the reflected halo from aircraft flying several kilometers above clouds. Currently, the primary use of halo observations is retrieving the geometrical thickness of optically thick stratiform cloud layers, although the retrievals simultaneously estimate cloud optical thickness and the vertical profile of cloud extinction as well. THOR also has potential application to ice, snow, and other multiple scattering media.

The instrument itself was first described, focusing on the components most unique to its multiple field-of-view design. The single most unique component is the custom-built optical fiber bundle, which funnels photons collected by the THOR telescope to the appropriate detector. The fibers in this bundle are organized into concentric rings, each of

which collects the signal that returns from a different ring surrounding the spot illuminated by the THOR laser. This arrangement allows THOR to measure the number of photons returning from the clouds as a function of photon travel time and radial distance from the illuminated spot. The signal level is kept high even at the outer, fainter parts of the halo by making the outer rings increasingly wide.

Next, data processing algorithms developed for THOR were described. These algorithms estimate cloud properties—the geometrical and optical thicknesses and the vertical profile of extinction coefficient—by comparing halo observations to look-up tables that contain simulated THOR data for a wide variety of clouds. Because of the time-consuming nature of three-dimensional Monte Carlo radiative transfer simulations, creating sufficiently comprehensive look-up tables has been a major challenge and required the development of several efficiency-enhancement methods.

The main sources of retrieval uncertainties (such as cloud variability or observational noise) were discussed, as were approaches to reducing the uncertainties caused by each source. A method of estimating the uncertainty of actual retrievals was described. While retrieval uncertainty depends on the exact situation, the uncertainties estimated for typical boundary layer cloud retrievals are comparable to the 30 m vertical resolution of the THOR system.

Finally, we described the THOR validation campaign that took place in central Oklahoma in March 2002. Data from the first flights aboard the NASA P-3B aircraft helped test and improve system performance, and demonstrated that THOR provides high-quality observations in optically thick clouds. The observations over an extensive layer of thick stratus cloud were used for validating THOR cloud thickness retrievals.

This validation compared THOR halo-based cloud thickness estimates to cloud thicknesses obtained by combining cloud top altitude from a central narrow angle channel with ground-based estimates of cloud base altitude at the DOE ARM SGP site. The comparisons confirmed that the halo retrievals work well, and the results lie well within the range of uncertainty of ground-based observations of cloud base. The average discrepancy between ground-based observations and THOR retrievals is 30 m (see Figure 14), for a range of thicknesses between 500 m and 1000 m.

In addition to estimating physical cloud thickness, THOR retrievals simultaneously estimate vertical extinction coefficient profile and optical thickness as well. The retrievals often recognize that the stratus clouds observed during the validation campaign contain two separate sub-layers and, accordingly, choose a 3-segment model of vertical extinction coefficient profile (Figure 6e) in most cases. The retrievals agree with ARM cloud radar data (Figure 9d) that the two sub-layers were most distinct around 6–7 UTC. THOR estimates of optical thickness are higher (≈ 30 -35) than the ARM Microwave Radiometer (MWR) estimates in Polonsky et al. (2004). Unfortunately it is difficult to evaluate the accuracy of these estimates, because the discrepancies may come from errors in either THOR or MWR retrievals or both. The higher THOR estimates are consistent with the possibility that cloud droplets were smaller than the 10 μm effective radius used in converting MWR liquid water paths into optical thicknesses. To fully understand the potential and limitations of THOR extinction profile and optical thickness retrievals, detailed theoretical studies are planned.

Other future developments in the THOR project are expected in several areas. One area is the further testing and development of cloud property retrieval algorithms. In

addition to performing more extensive sensitivity studies, we will explore the possibility of THOR retrievals of microphysical cloud parameters. The microphysical retrievals would complement the current macrophysical retrievals that focus on the outer portions of the halo and on the tail of time-dependent signals (where multiple scattering minimizes the sensitivity to cloud microphysics and maximizes the sensitivity to cloud thickness). By emphasizing more the inner channels and earlier returns, the retrievals may estimate particle scattering asymmetry parameter, and perhaps cloud phase and droplet size or ice crystal behavior (Roy et al. 1997). Examining the details of early returns may provide useful information also on the overlying Cirrus cloud, if the partial spread of downwelling laser pulses can reveal the shape of the forward peak of the Cirrus scattering phase function.

Another area of future development is to expand THOR capabilities to measure the thickness of snow cover and sea ice. In-situ experiments have demonstrated the concept of such retrievals using halo observations (Haines et al. 1997), and THOR sea ice observations will be collected during the Antarctic AMSR-E Sea Ice (AASI) experiment planned for Fall 2004. Planned low flight altitudes will allow the inner channels to observe the halo that in sea ice extends only to a few meters. Subsequently, THOR may be upgraded by increasing its time resolution and by adding several more tightly focused channels that can better resolve the details of halos in snow and ice.

Finally, we intend to adapt THOR to perform daytime cloud observations as soon as suitable spectral filters (now under development) become available to sufficiently reduce the effects of background illumination.

Acknowledgments: THOR has been funded by the Goddard Director's Discretionary Fund (G. Soffen) and the NASA Radiation Sciences Program (D. Anderson and H. Maring). This research was also supported by the Office of Science (BER), U.S. Department of Energy, Interagency Agreement No. DE-AI02-00ER62939. We are grateful to Anthony Davis and Steve Love for fruitful discussions, and to Robert Kolasinski and Mindy Schwasinger for help with laboratory experiments of halos in artificial clouds. We also thank Luis Ramos-Izquierdo for the design of THOR's telescope, Ray Disilvestre and Ken Cory for manufacturing some parts of THOR, and Stan Scott for testing and aligning THOR's laser and telescope. In addition, we are also grateful to all who contributed to the success of the March 2002 THOR validation campaign, including Dave Easmunt who managed the THOR mission at Wallops Flight Facility, Steve Spears who helped integrating THOR into the NASA P-3B aircraft, Willie Dykes and Mike Singer who piloted the aircraft, Scott Dunn who provided ground support at McConnell Air Force Base, Jim Liljegren, Peter Lamb and Jim Teske who supported the mission at the ARM SGP site, and Chad Bahrmann and John Schatz who provided timely, detailed, and accurate weather forecasts for the mission. Administrative support by Cathy Newman, Nadja Hardy, and Marcy August is also gratefully acknowledged.

References

- Barker, H. W., R. K. Goldstein, and D. E. Stevens, 2003: Monte Carlo simulation of solar reflectances for cloudy atmospheres. *J. Atmos. Sci.*, **60**, 1881-1894.
- Cahalan, R. F., and J. B. Snider, 1989: Marine stratocumulus structure. *Remote Sens. Environ.*, **28**, 95-107.
- Cahalan, R. F., and R. Davies, eds., 1999: *Abstracts of the first and second international workshops on the international workshop on Intercomparison of 3-dimensional Radiation Codes*. ISBN 0-9709609-0-5.
- Cahalan, R. F., W. Ridgway, W. J. Wiscombe, T. L. Bell, and J. B. Snider, 1994a: The albedo of fractal stratocumulus clouds. *J. Atmos. Sci.*, **51**, 2434-2455.
- Cahalan, R. F., W. Ridgway, W. J. Wiscombe, S. Gollmer, and Harshvardhan, 1994b: Independent pixel and Monte Carlo estimates of stratocumulus albedo. *J. Atmos. Sci.*, **51**, 3776-3790.
- Cahalan, R. F., L. Oreopoulos, and the I3RC Team, 2004: International Intercomparison of 3-dimensional Radiative Codes (I3RC). *Bull. Amer. Meteor. Soc.*, in preparation.
- Davis, A., A. Marshak, R. Cahalan, and W. Wiscombe, 1997: The Landsat scale break in stratocumulus as a three-dimensional radiative transfer effect: Implications for cloud remote sensing. *J. Atmos. Sci.*, **54**, 241-260.
- Davis, A. B., R. F. Cahalan, J. D. Spinhirne, M. J. McGill, and S. P. Love, 1999: Off-beam lidar: An emerging technique in cloud remote sensing based on radiative Green-function theory in the diffusion domain. *Phys. Chem. Earth (B)*, **24**, 757-765.

- Davis, A. B., S. P. Love, R. F. Cahalan, and M. J. Mc Gill, 2002: Off-beam Lidar senses cloud thickness and density. *Laser Focus World*, October Issue, **38**, 101-103.
- Evans, K. F., R. P. Lawson, P. Zmarzly, and D. O'Connor, 2003: In situ cloud sensing with multiple scattering cloud lidar: Simulations and demonstration. *J. Atmos. Ocean. Tech.*, **20**, 1505-1522.
- Groenhius, R. A. J., J. J. Ten Bosch, and H. A. Ferwerda, 1983: Scattering and absorption of turbid materials determined from reflection measurements. 2: Measuring method and calibration. *Applied Optics*, **22**, 2463-2467.
- Haines, E. M., R. G. Buckley, and H. J. Trodahl, 1997: Determination of the depth dependent scattering coefficient in sea ice. *J. Geophys. Res.* **102**, 1141-1151.
- Love, S. P., A. B. Davis, C. Ho, and C. A. Rohde, 2001: Remote sensing of cloud thickness and liquid water content with Wide-Angle Imaging Lidar. *Atmos. Res.*, **59-60**, 295-312.
- Maffione, R. A., J. M. Voss, and C. D. Mobley, 1998: Theory and measurements of the complete beam spread function of sea ice. *Limnol. Oceanogr.*, **43**, 34-43.
- Marchuk, G. I., G. A. Mikhailov, M. A. Nazaraliev, R. A. Darbinjan, B. A. Kargin, and B. S. Elepov, 1980: *The Monte Carlo methods in atmospheric optics*, 208 pp., Springer-Verlag, New York.
- Oreopoulos, L., and R. Davies, 1998: Plane-parallel albedo biases from satellite observations. Part I: Dependence on resolution and other factors. *J. Climate*, **11**, 919-932.

- Polonsky, I. N., S. P. Love, and A. B. Davis, 2004: Wide-Angle Imaging Lidar (WAIL) deployment at the ARM Southern Great Plains site: Intercomparison of cloud property retrievals. *J. Atmos. Ocean. Tech.*, Submitted to this issue.
- Predehl, P., V. Burwitz, F. Paerels, and J. Trümper, 2000: Chandra measurement of the geometrical distance to Cyg X-3 using its X-ray scattering halo. *Astron. Astrophys.* **357**, L25-L28.
- Roy, G., L. R. Bissonnette, C. Bastille, and G. Vallée, 1997: Estimation of cloud droplet size density distributions from multiple-field-of-view lidar returns. *Opt. Eng.*, **36**, 3404-3415.
- Sassen, K., and K.-N. Liou, 1979: Scattering of polarized laser light by water droplet, mixed-phase and ice crystal clouds. Part I: Angular scattering patterns. *J. Atmos. Sci.*, **36**, 838-851.
- Sparling, L. C., and G. H. Weiss, 1993: Some effects of beam thickness on photon migration in a turbid medium. *J. Modern Optics*, **40**, 841-859.
- Takano, Y., and K.-N. Liou, 1989: Radiative transfer in cirrus clouds. Part III: Light scattering by irregular ice crystals. *J. Atmos. Sci.*, **52**, 818-837.
- Welch, R. M., S. K. Cox, and J. M. Davis, 1980: *Solar radiation and clouds*. Meteorological monographs, **17**, #39, American Meteorological Society, 96 pp.
- Zuidema, P., and K. F. Evans, 1998: On the validity of the independent pixel approximation for boundary layer clouds observed during ASTEX. *J. Geophys. Res.* **103**, 6059-6074.

Appendix A. Calculation of THOR signal-to-noise ratios

If several minutes of THOR observations are available under unchanging conditions, the signal-to-noise ratio (SNR) of THOR observations can be deduced directly from the statistical spread of $c_i^{obs}(t)$ photon counts observed by channel i at time-delay (or range) bin t . Unfortunately, suitable steady conditions seldom occur outside laboratory experiments. As a result, it is more practical to estimate the SNR using a theoretical approach, rather than trying to separate noise from variability in the observed clouds. Noise levels estimated by this theoretical approach have been confirmed by laboratory experiments in which a steady lamp was used to illuminate the THOR telescope. The theoretical procedure consists of three steps: estimating the signal strength, the background illumination, and the SNR itself.

Step 1. Signal strength

First, the expected reflectance values $R_i(t)$ can be obtained through Monte Carlo radiative transfer simulations. These reflectances are related to radiances through the equation

$$R_i(t) = \frac{\pi \cdot I_i(t)}{\mu_0 \cdot F_0}, \quad (\text{A1})$$

where μ_0 , the cosine of illumination zenith angle is equal to 1, because THOR points vertically downward. For our purposes it is more practical to express the radiance $I_i(t)$ and irradiance F_0 in units of $photon\ count \cdot s^{-1} \cdot m^{-2} \cdot sr^{-1}$, and $photon\ count \cdot s^{-1} \cdot m^{-2}$,

respectively, instead of the usual units involving Watts. In these units, F_0 represents the number of photons emitted by the laser, and is given by

$$F_0^{laser} = \frac{t_{acc} \cdot N_{pulse} \cdot E_{pulse} \cdot \lambda}{v_{light} \cdot h}. \quad (A2)$$

In this equation t_{acc} is the time interval through which observations are accumulated (separately for each time-delay or range bin t ; $t_{acc} \geq 0.5$ s), N_{pulse} is the number of laser pulses per second (1000 s^{-1}), E_{pulse} is the energy of a single laser pulse (225 μ J), λ is the wavelength (540 nm), v_{light} is the speed of light ($3 \cdot 10^8$ m s^{-1}), and h is the Planck constant ($6.63 \cdot 10^{-34}$ s J^{-1}).

Next, we can calculate $\Delta\Omega$, the solid angle THOR's telescope lens occupy in the field-of-view of an imaginary observer standing at the cloud top as

$$\Delta\Omega = 2\pi \left(1 - \cos \left(\arctan \frac{\rho_{tel}}{z_{THOR}} \right) \right) \approx \pi \left(\frac{\rho_{tel}}{z_{THOR}} \right)^2. \quad (A3)$$

In this equation, ρ_{tel} is the radius of THOR's telescope lens, and z_{THOR} is THOR's altitude above the cloud top.

Finally, we can combine Equations (A1-A3) to estimate the desired signal strength $F_i^{laser}(t)$ as the number of laser photons reflected from the clouds and intercepted by THOR's telescope:

$$F_i^{laser}(t) = \mu_i \cdot I_i(t) \cdot \Delta\Omega = \mu_i \frac{R_i(t) \cdot \mu_0 \cdot F_0}{\pi} \Delta\Omega. \quad (A4)$$

Here, μ_i is the cosine of channel i 's effective viewing zenith angle (≈ 1).

Step 2. Background illumination (solar and lunar)

To calculate the number of background photons reaching THOR's telescope, we first need to calculate $F_{0,i}^{bg}$, the number of incoming lunar or solar photons that reach the cloud inside the field-of-view of THOR channel i :

$$F_{0,i}^{bg} = \frac{t_{acc} \cdot t_{bin}}{t_{repeat}} \cdot \frac{F^{bg} \cdot \mu_0 \cdot A^{bg} \cdot A_i \cdot \lambda}{v_{light} \cdot h}, \quad (A5)$$

where t_{bin} is the time corresponding to THOR's 30.8 m range-resolution ($2 \cdot 10^{-7}$ s), t_{repeat} is the time between subsequent laser pulses (0.001 s), F^{bg} is the lunar or solar irradiance at 540 nm (about $3.6 \cdot 10^{-6}$ W m⁻² nm⁻¹ or 1.8 W m⁻² nm⁻¹, respectively), and μ_0 is the lunar or solar zenith angle. For solar illumination $A^{bg} = 1$, whereas for lunar illumination A^{bg} is the fractional area of the Moon's bright surface, a value that depends on the lunar phase ($0 \leq A^{bg} \leq 1$). Finally, A_i is the cloud top surface area viewed by channel i . Given channel i 's minimum and maximum viewing zenith angles θ_i^{\min} and θ_i^{\max} (see Table 1), A_i can be calculated as

$$A_i = \pi \cdot \left(z_{THOR} \cdot \left(\tan \theta_i^{\max} - \tan \theta_i^{\min} \right) \right)^2. \quad (A6)$$

Knowing $F_{0,i}^{bg}$ from (A5) and $\Delta\Omega$ from (A3), we can calculate the number of background illumination photons reflected by the cloud into THOR's telescope (F_i^{bg}) as

$$F_i^{bg} = \mu_i \cdot I_i \cdot \Delta\Omega = \mu_i \cdot \frac{R_i(\mu_0) \cdot \mu_0 \cdot F_{0,i}^{bg}}{\pi} \Delta\Omega. \quad (A7)$$

In this equation, $R_i(\mu_0)$ is the clouds nadir reflectance (obtained from Monte Carlo or other radiative transfer calculations).

Step 3. Signal-to-noise ratio

Based on the number of photons reaching the telescope, we can estimate the number of photons actually counted by the THOR detectors as

$$c_i^{laser}(t) = F_i^{laser}(t) \cdot E_i \quad (\text{A8a})$$

$$c_i^{bg} = F_i^{bg} \cdot E_i. \quad (\text{A8b})$$

Here, $E_i = 0.04$ is the overall system efficiency for channel i . The value of 0.04 was estimated by considering the following photon losses: 15% loss in spectral filters, 50% loss in focal plane (many lost photons arrive to the focal plane at the unavoidable gaps between the fibers, which have circular cross-sections), 4% loss inside the fibers, and 90% loss at the detector (which detect 1 out of 10 arriving photons). The approximate magnitude of E_i was confirmed by the number of lunar photons that were detected during the field campaign described in Section 4. If needed, the E_i values can be refined using calibration experiments.

The c_i values calculated in (A8) are sufficient for estimating signal-to-noise ratio, because virtually all of the observational noise arises from the random nature of photon arrivals to the detector. This so-called photon shot noise can be estimated using the theory of Poisson distributions, which reveals that the uncertainty, or noise, is equal to the square root of observed photon counts. Therefore the signal-to-noise ratio (*SNR*) can be estimated as

$$SNR_i(t) = \frac{c_i^{laser}(t)}{\sqrt{c_i^{laser}(t) + c_i^{bg}}}. \quad (\text{A9})$$

The *SNR* values estimated for typical observation conditions are discussed in Section 3c.

Tables

Table 1. Main parameters of the THOR system

Laser	
Pulse rate	1 kHz
Wavelength	540 nm
Pulse duration	8 ns (pulse length \approx 2.4 m)
Pulse energy	Adjustable, max. 225 μ J
Beam divergence	325 μ rad at $1/e^2$ level
Receiver	
Telescope lens diameter	19.05 cm
Telescope effective focal length	23.80 cm (f/1.25)
Spectral filter bandwidth	7 nm
Fields of view (full angle, mrad)	
Channel 1	0.000-0.840
Channel 2	1.029-1.681
Channel 3	1.681-3.361
Channel 4	3.361-6.723
Channel 5	6.723-13.40
Channel 6	13.40-26.72
Channel 7	26.72-53.40
Channel 8+9+10	53.40-106.7

Table 2. Summary of major sources of uncertainty in THOR cloud thickness retrievals.

Source of uncertainty		Estimated magnitude at the 2002 validation campaign	Approach to reducing uncertainties
Vertical cloud variability		< 50 m	Generate look-up tables for more cloud structures
Horizontal cloud variability	Nonlinearity effect	< 10 m	Perform retrievals at high spatial resolution
	3D radiative effects	< 10 m	Perform theoretical study
Drop size variations		< 30 m	Create look-up tables for several drop sizes
Cirrus effects		≈ 10 m	Currently, retrievals only performed for very thin Cirrus. Working on a scheme to consider Cirrus effects.
Monte Carlo noise		< 10 m	No need for improvement.
Observational noise		20-40 m	Install better thermal control for full laser power.

Table 3. Summary of major sources of uncertainty in cloud thicknesses estimated by combining THOR Channel 1 data with ground-based measurements.

Source of uncertainty		Estimated magnitude at the 2002 validation campaign
Cloud top height	THOR and MMCR range resolution	~ 40 m
Cloud base height at ARM site	MPL range-resolution	~ 15 m
	Drizzle	~ 100 m
Cloud base height at Oklahoma mesonet sites	Instrument resolution	< 50 m
	Assumption of well-mixed sub-cloud layer	~ 35 m
	Drizzle	~ 100 m

Figure captions

Figure 1. Laboratory observations of diffuse halos. (a) Experimental setup. (b) Signal strength measured by the photo-multiplier tube (PMT) detector, as a function of the fiber optic probe's horizontal position. Successive jumps at discrete positions indicate insertion of a series of 10X neutral density filters, used to maintain the detector's linear response. The decrease to noise level closely follows theoretical predictions.

Figure 2. A schematic view of THOR observations.

Figure 3. (a) A schematic of the THOR system. The thin black arrays indicate the flow of electric signals, whereas the gray arrows indicate the path of photons between THOR components. (b) A drawing of THOR. (c) A schematic view of the fiber bundle at the focal plane.

Figure 4. Photos of the THOR instrument inside the NASA P-3B aircraft. (a) View of THOR above the flight deck. (b) A downward-looking view of the part of THOR that lies below the flight deck. Starting from the lower left corner and moving clockwise through the image, the components are as follows. In the left foreground, we see the data system computer. The little box right behind the computer (and slightly toward the center) is the base plate on which the laser unit is mounted. Further behind we can see the shiny metal box of the laser unit, and behind the laser, the black frame of the beam expander. The camera flash reflects from the aircraft's down-looking 18 inch port window. To the right

of the beam expander lies THOR's telescope. The black-coated optical fiber bundle couples to the telescope as a single bundle, but then divides into 10 branches, each coupled to one of the 10 channel's spectral filters. In the right foreground we see the cylindrical housings of the filters.

Figure 5. THOR calibration experiments. (a) A view of THOR oriented horizontally in the laboratory, with the calibration sphere attached to the telescope. (b) Ratio of photon counts in neighboring THOR channels. The data on the right side was obtained in laboratory experiments that used 10 illumination levels. The data on the left side is from the last flight of the THOR validation campaign described in Section 4. The illumination level of in-flight data corresponds to about a third of the lowest illumination level used in the laboratory.

Figure 6. A schematic illustration of cloud models used in creating THOR look-up tables. τ is the total cloud optical thickness and $\bar{\sigma}$ is the cloud's mean extinction coefficient. z and σ are altitude above cloud base and extinction coefficient, respectively, at points identified in the figure. As mentioned in the text, each parameter listed as "Variable" can take three possible values in Monte Carlo simulations. Results for additional parameter values are then obtained using cubic interpolation.

Figure 7. A sample of THOR observations for a typical one of the 16 retrieval segments collected during the THOR validation campaign, as described in Section 4. See also Figure 9a, which shows a continuous time sequence for channels 1, 5 and 8+9+10.

Figure 8. Influence of cloud thickness in simulated THOR data. The calculations assume homogeneous clouds with an extinction coefficient of 25 km^{-1} . The scattering phase function is obtained from Mie calculations assuming the Sc_{top} drop size distribution of Welch et al. (1980). The simulations assume that, similarly to some validation flights discussed in Section 4, THOR is about 7.3 km above the cloud top. (a) Influence on the time-integrated signal of each channel; (b) Influence on the time-dependence of the outermost ring (Channels 8+9+10).

Figure 9. Observations of clouds encountered during THOR's flight on March 25, 2002. (a) An RGB color composite image of THOR observations taken from a flight altitude of 8540 m. Red is the direct backscatter photon count observed by THOR's central channel (Channel 1), while green and blue are photon counts observed at two of THOR's outer rings (Channel 5 and Channel 8+9+10, respectively). Therefore red means single or low order scattering, while blue indicates multiple high order scattering. The Cirrus cloud and the very top of the Stratus cloud appear bright red due to an intense direct backscatter (though some thick parts of cirrus appear yellow because of some multiple scattering). The signal received at longer time-delays (displayed as if it had returned from lower altitudes) comes from outer parts of the halo, and so the red signal turns white at first, then green, and eventually blue. The 5 colored segments along the x -axis represent data records that were used in the retrievals. Although the RGB coloring scheme does not show this clearly, the Cirrus was much thinner at the retrieval segments than at other times. (b) Skew T -log P diagram of balloon sounding data from 11.5 UTC at the ARM

SGP central facility. Vertical scaling is height in km (labeled by the black numbers). The cyan, magenta and green colors represent dry adiabats, mixing ratios and saturated adiabats, respectively. (c) ARM MPL normalized relative backscatter data. The colored dots on this figure show both the times and the flight altitudes of the data records that were used in the retrievals. (d) ARM MMCR boundary layer mode data during the flight period.

Figure 10. Cloud base height determined from the following data sources at the ARM Central Facility on March 25, 2002: Ceilometer (CEIL), 60 meter tower temperature and humidity (TWR60 - at the 60 meter level, TWR25 - at the 25 meter level), Energy balance Bowen ratio temperature and humidity (EBBR - at the 2 meter level) and the Micropulse Lidar (MPL).

Figure 11. Cloud base heights calculated from temperature and humidity data at five Oklahoma Mesonet sites (BLAC, REDR, BREC, MEDF, NEWK) and the cloud base height determined from the micropulse lidar at the ARM central facility (ARM CF) on March 25, 2002. The locations of mesonet sites are indicated in Figure 12 (the same symbols are used in both figures). The grey bars indicate records used in THOR retrievals, and are numbered sequentially.

Figure 12. Flight track on March 25, 2002 over the ARM site. The two arrows on the track show the heading of the aircraft. The sections of the track colored in white represent data records that were used in the retrievals. The central facility is indicated by the letter "c" while the other five symbols represent the locations of Oklahoma mesonet sites. The

thin black lines represent county boundaries. The Kansas-Oklahoma border lies exactly on the 37° latitude line. The dashed ellipse around "c" indicates a 4 km distance from the central facility in the map's projection.

Figure 13. Cloud top and base heights along the flight track. Upper line segments represent cloud top height for valid records (where the aircraft tilt angle was less than 3°). Bottom curve represents cloud base height interpolated from the values in Figure 12. Bold marks at the top of the plot indicate periods where the aircraft was within 4 km of the central facility (these records were used to determine the flight altitude bias). The thickness of each mark reveals the time the aircraft spent within the 4 km range. The grey bars indicate records that were used in THOR cloud retrievals, and are numbered sequentially. The inset illustrates the 3-segment model of vertical extinction coefficient profile (also shown in Figure 6e) that the retrievals chose for all retrievals except #12-14, for which the profile shown in Figure 6f worked better.

Figure 14. Comparison of cloud thicknesses obtained from THOR halo retrievals (vertical axis) with those obtained from combining THOR cloud tops with ARM cloud base values (horizontal axis). The numbers inside symbols indicate the time-sequence of retrievals, as shown in Figures 11 and 13.

Figures

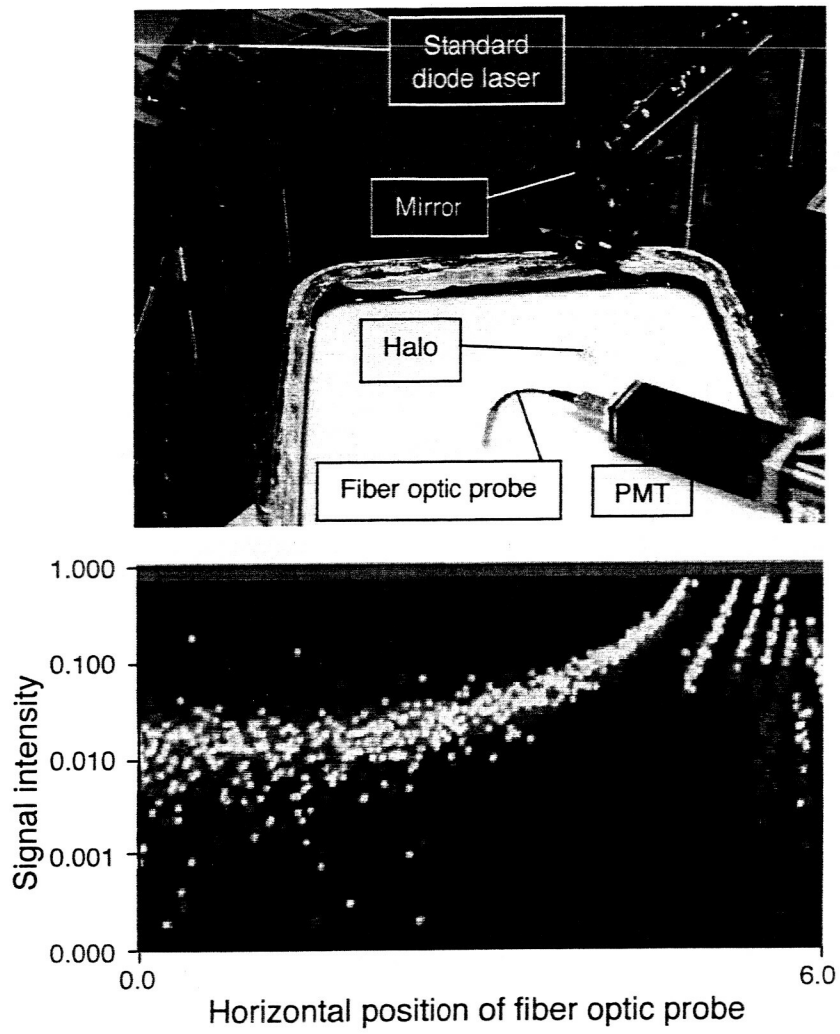


Figure 1.




Figure 2.

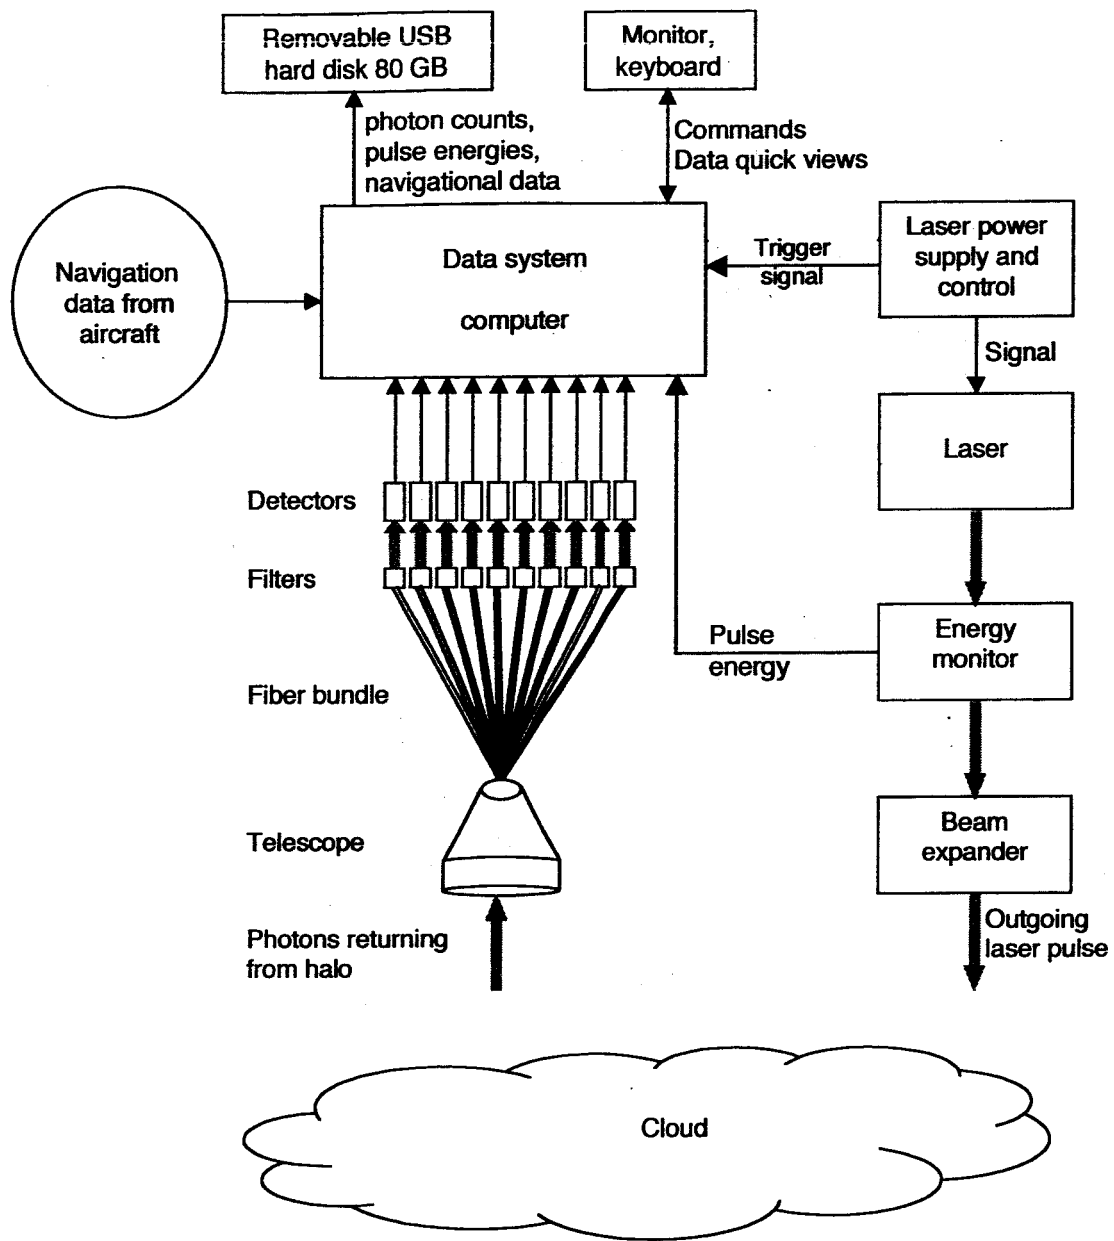


Figure 3a.

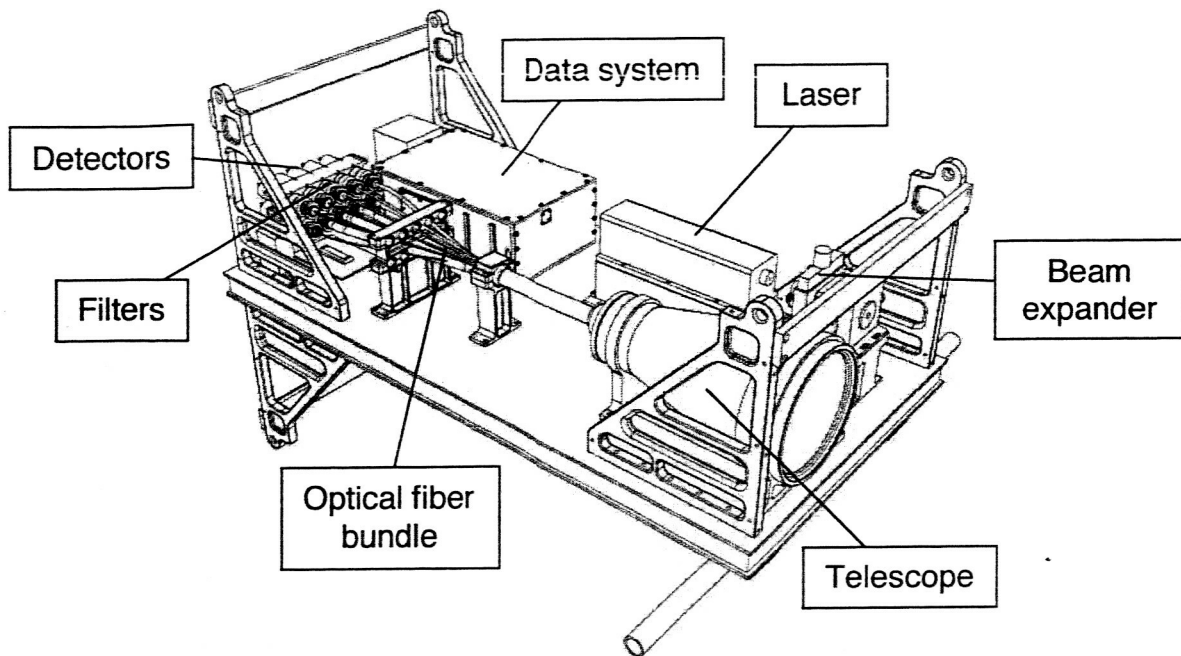


Figure 3b.





Figure 3c.

a.



b.




Figure 4.




Figure 5a.




Figure 5b.

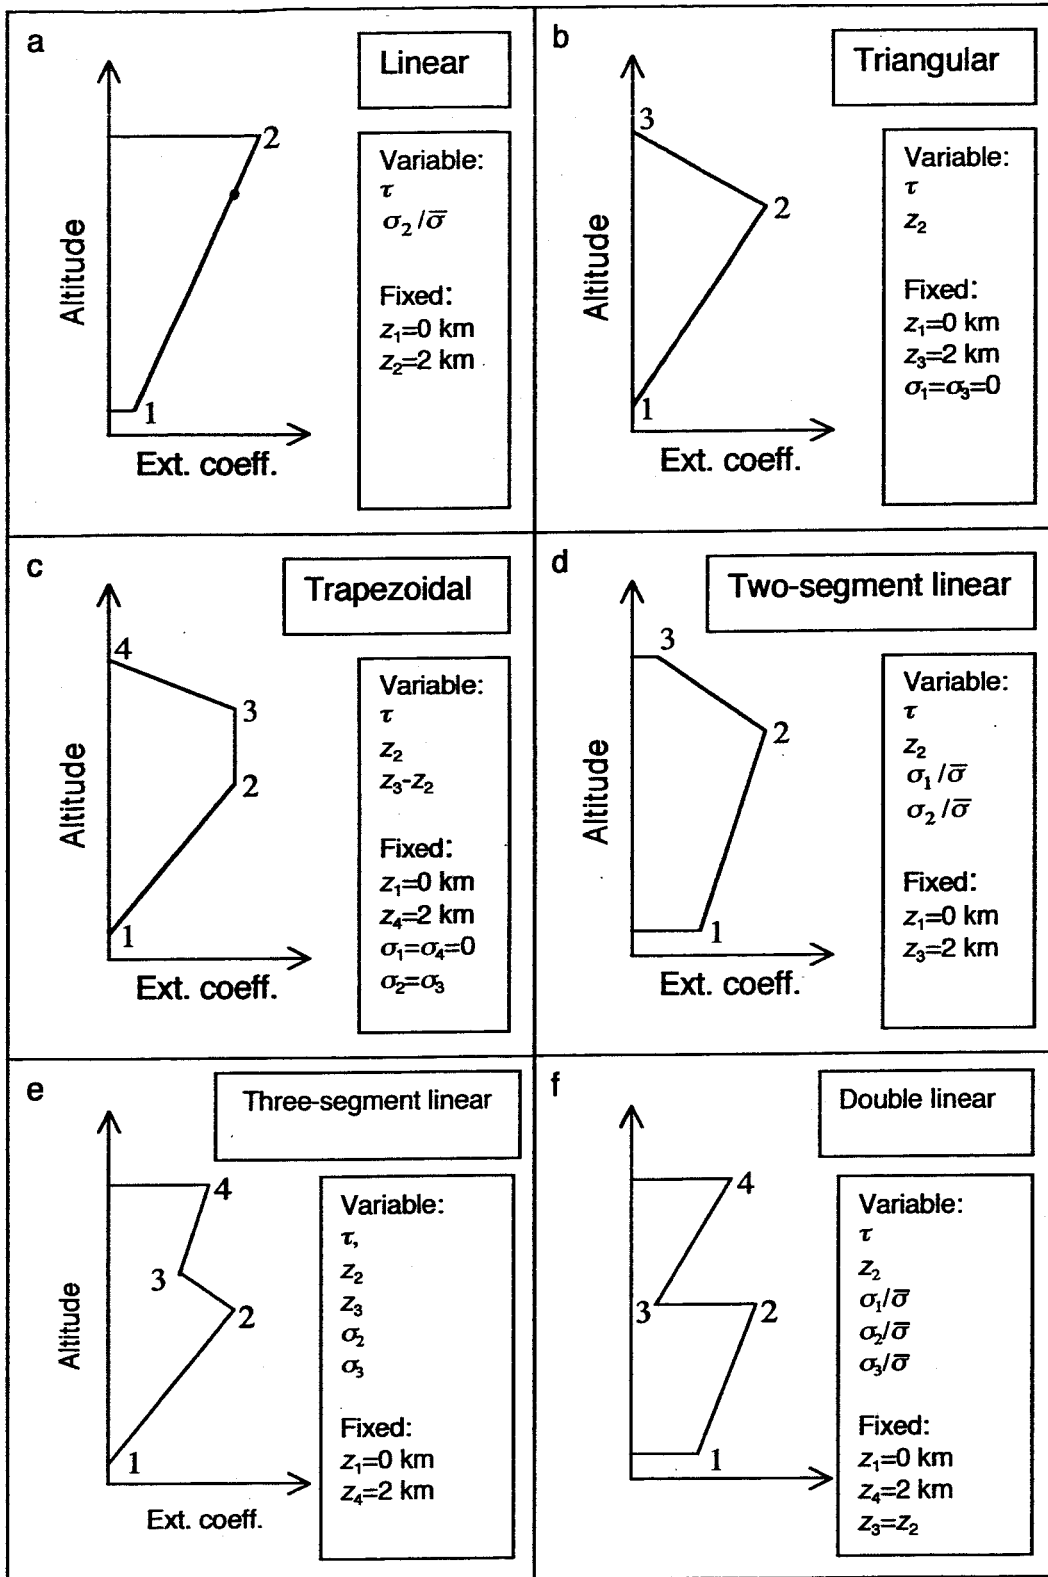


Figure 6.

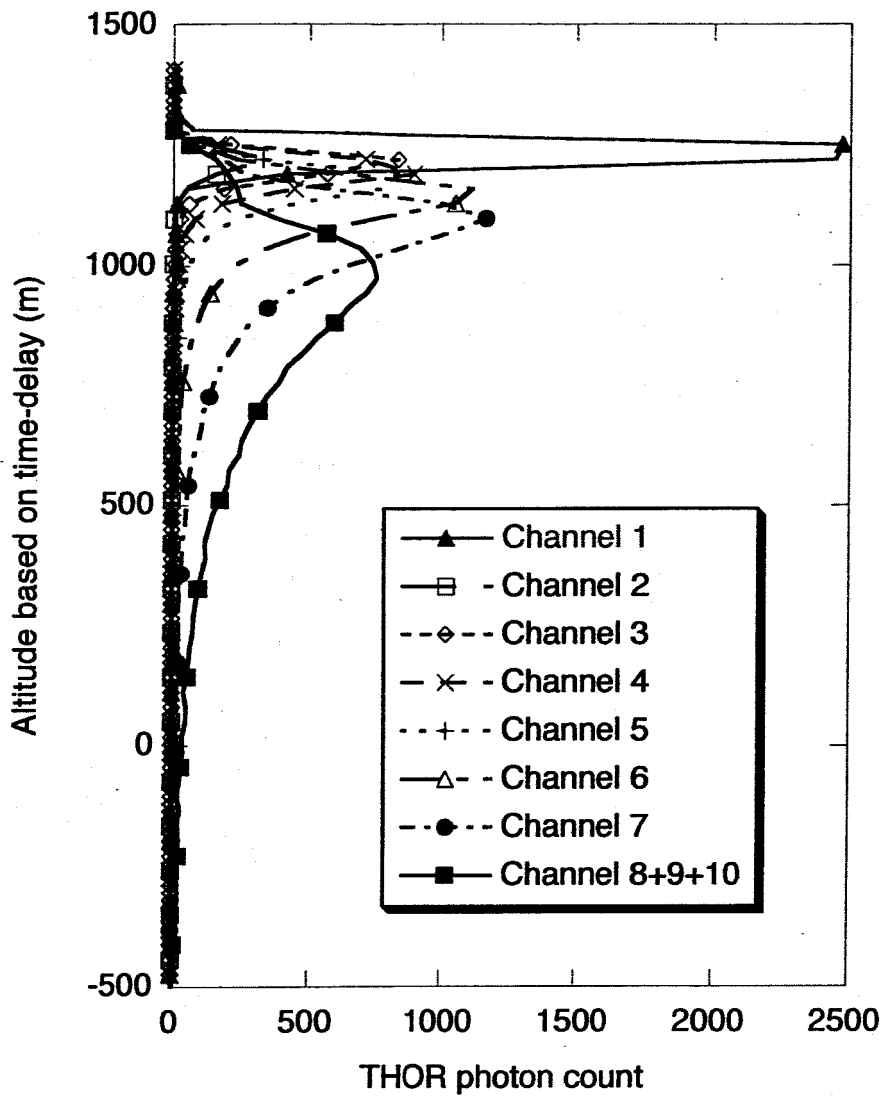


Figure 7.

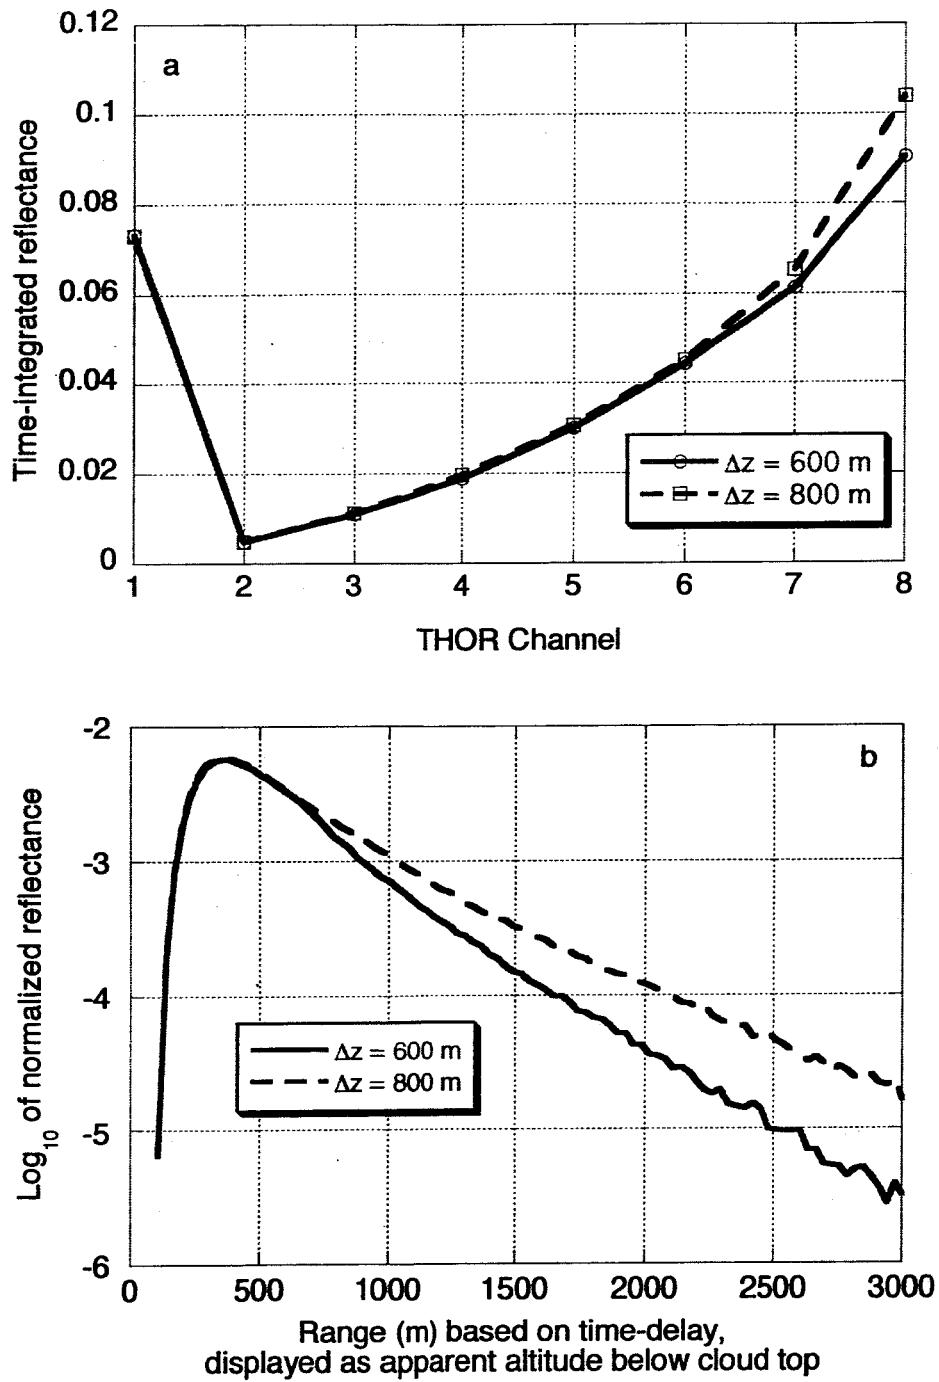


Figure 8.

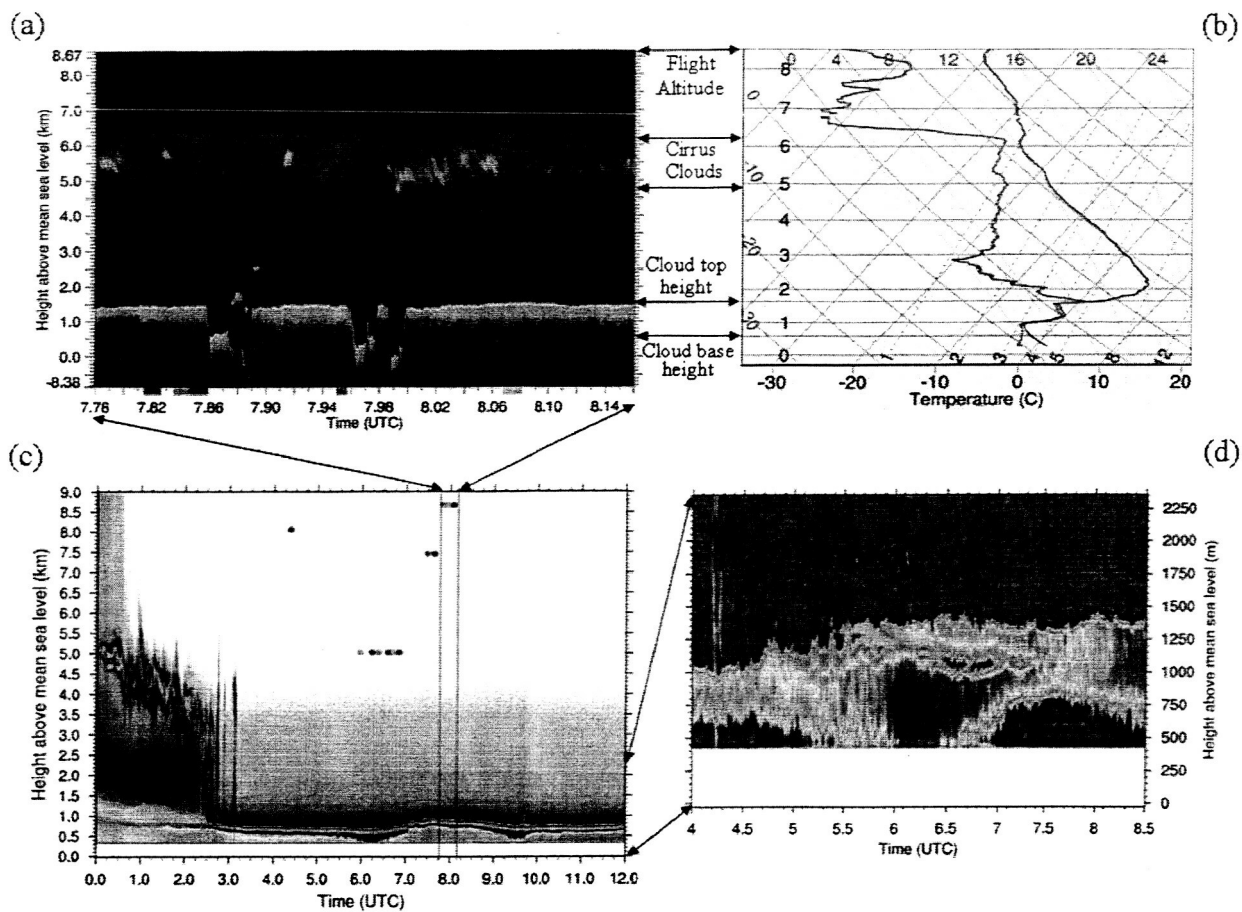


Figure 9.




Figure 10.




Figure 11.




Figure 12.




Figure 13.

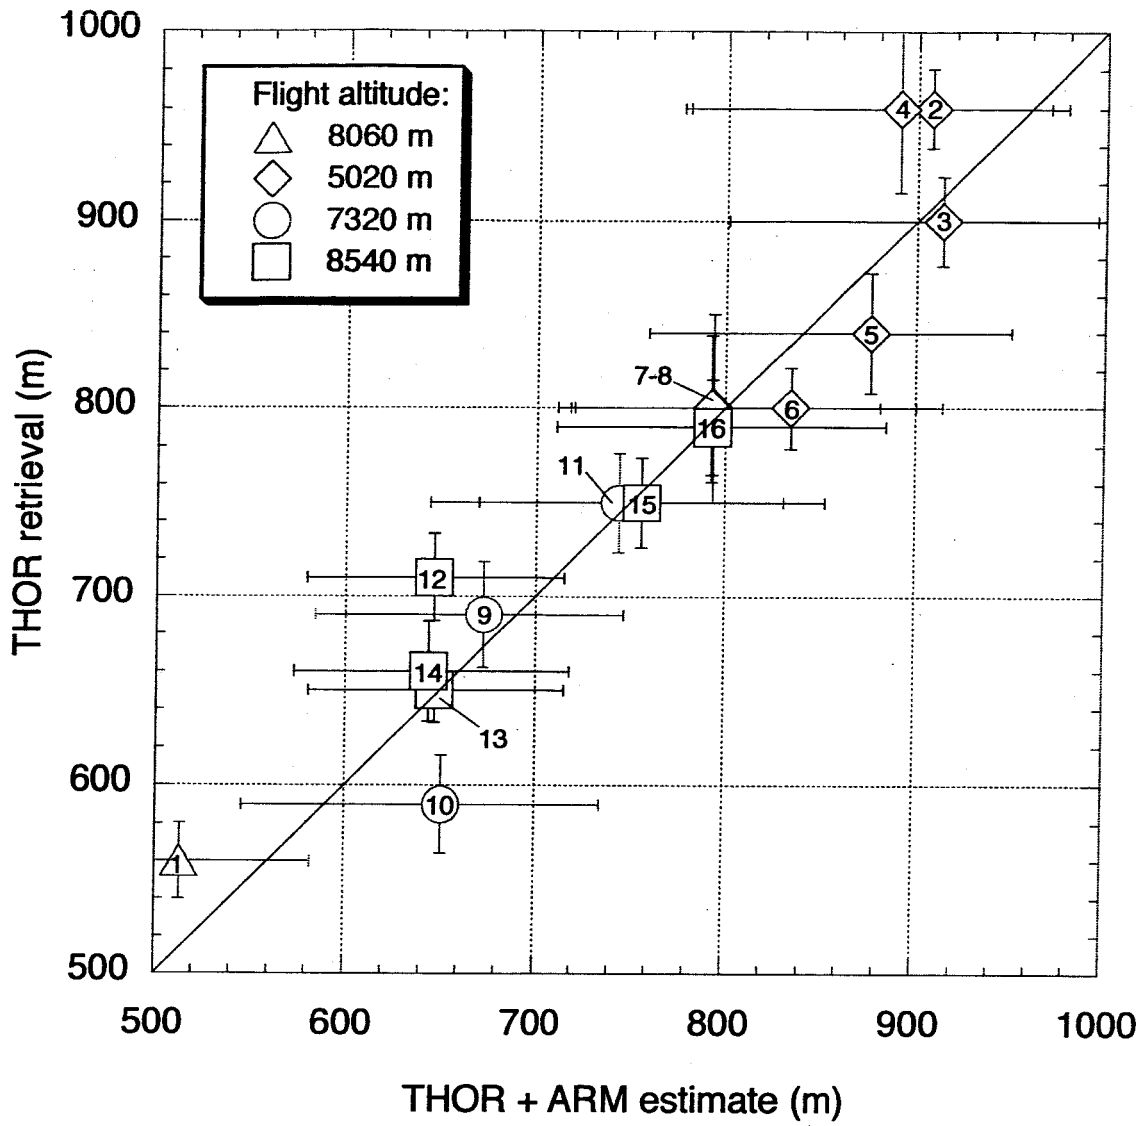


Figure 14.

THOR – Cloud Thickness from Offbeam Lidar Returns

Robert F. Cahalan, Matthew McGill, John Kolasinski, Tamás Várnai, and Ken Yetzer

Prepared for the Journal of Atmospheric and Oceanic Technology

June 2004

Popular summary

This paper describes the new airborne lidar instrument called THOR, its data processing algorithms, and the March 2002 THOR validation campaign. THOR was built to probe optically thick clouds. The inside of these clouds is beyond the reach of conventional lidars, because cloud particles scatter the distinct laser pulses into a diffuse halo that lies outside the narrow field of view of conventional lidars. Multiple wide fields-of-view allow THOR to observe the details of a reflected halo from aircraft flying several kilometers above clouds. Currently, the primary use of halo observations is obtaining the geometrical thickness of thick stratiform cloud layers, although the retrievals simultaneously estimate the internal structure of clouds as well.

The THOR validation campaign took place in March 2002 over the Department of Energy Atmospheric Radiation Measurement program Southern Great Plains site in central Oklahoma. This validation compared THOR halo-based cloud thickness estimates to cloud thicknesses obtained by combining cloud top altitude from THOR with ground-based estimates of cloud base altitude. The comparisons confirmed that the halo retrievals work well, and the results lie well within the range of uncertainty of ground-based observations. The average discrepancy between ground-based observations and THOR retrievals is 30 m, for a range of thicknesses between 500 m and 1000 m.