Lidar measurements taken with a large-aperture liquid mirror. 1. Rayleigh-scatter system

R. J. Sica, S. Sargoytchev, P. S. Argall, E. F. Borra, L. Girard, C. T. Sparrow, and S. Flatt

A lidar system has been built to measure atmospheric-density fluctuations and the temperature in the upper stratosphere, the mesosphere, and the lower thermosphere, measurements that are important for an understanding of climate and weather phenomena. This lidar system, the Purple Crow Lidar, uses two transmitter beams to obtain atmospheric returns resulting from Rayleigh scattering and sodium-resonance fluorescence. The Rayleigh-scatter transmitter is a Nd:YAG laser that generates 600 mJ/pulse at the second-harmonic frequency, with a 20-Hz pulse-repetition rate. The sodium-resonance–fluorescence transmitter is a Nd:YAG-pumped ring dye laser with a sufficiently narrow bandwidth to measure the line shape of the sodium D_2 line. The receiver is a 2.65-m-diameter liquid-mercury mirror. A container holding the mercury is spun at 10 rpm to produce a parabolic surface of high quality and reflectivity. Test results are presented which demonstrate that the mirror behaves like a conventional glass mirror of the same size. With this mirror, the lidar system's performance is within 10% of theoretical expectations. Furthermore, the liquid mirror has proved itself reliable over a wide range of environmental conditions. The use of such a large mirror presented several engineering challenges involving the passage of light through the system and detector linearity, both of which are critical for accurate retrieval of atmospheric temperatures. These issues and their associated uncertainties are documented in detail. It is shown that the Rayleigh-scatter lidar system can reliably and routinely measure atmospheric-density fluctuations and temperatures at high temporal and spatial resolutions.

1. Introduction

Ground-based determinations of temperature profiles through the use of active optical remote sensing in the middle atmosphere, that is, the stratosphere, the mesosphere, and the lower thermosphere, were first performed in the early 1950's by Elterman.[1] The development of the laser in the early 1960's provided an ideal light source for lidar (*li*ght *d*etection *a*nd *r*anging) systems. The use of lasers in lidar systems improved on the initial experiments by Elterman, who used a search light for a transmitter. The first Rayleigh-scatter temperature determinations made with a laser radar were performed by Fiocco and Smullin[2] in 1963. Rayleigh-scatter lidars have now been developed that can routinely determine the temperature, wind, and density in the middle atmosphere for climatological studies. Rayleigh-scatter lidar measurements of density fluctuations are beginning to have a significant impact on the study of gravity waves.[3,4] However the power–aperture product of most systems limits their ability to measure the gravity-wave spectrum at sufficient temporal and spatial scales to challenge current theoretical progress.

For lidars operating at identical transmitter frequencies, the power–aperture product is a relative measure of the system's performance and thus of its temporal and spatial resolutions. The laser of choice for most Rayleigh-scatter applications is the Nd:YAG laser, which is typically used at the second-harmonic frequency (532 nm). Excimer lasers enjoy higher peak powers, in addition to offering the advantage of the larger Rayleigh-scatter cross section in the blue region of the visible spectrum. However, the near "plug-and-play" ease of operation of modern Nd:YAG systems is attractive for routine lidar measurements when compared with the extensive maintenance time

R. J. Sica, P. S. Argall, and C. T. Sparrow are with the Department of Physics, The University of Western Ontario, London, Ontario N6A 3K7, Canada; S. Sargoytchev and S. Flatt were with the same institution when this work was performed. E. F. Borra and L. Girard are with the Department of Physics and Astronomy, Université Laval, Québec, Québec G1K 7P4, Canada. S. Sargoytchev is now with the Institute for Space and Terrestrial Science, New York, Ontario M3J 3K1, Canada; S. Flatt is now with Scitech, Inc., Saskatoon, Saskatchewan S7N 3L7, Canada.

Received 9 November 1994; revised manuscript received 2 June 1995.

0003-6935/95/306925-12$06.00/0.

and operating costs of a high-power excimer laser. Typical Nd:YAG lasers employed in Rayleigh-scatter systems have pulse powers of 100–1000 mJ, with repetition rates of 10–50 Hz. Systems that have significantly higher power are not presently available, and large increases in Nd:YAG power output is not on the immediate horizon for commercial systems.

System performance can also be improved by an increase in the receiver aperture. Most systems use conventional telescopes made from glass substrates that have metallic coatings. Diameters of typical primary mirrors used in lidar systems are between approximately 0.25 to 1 m. The size is typically limited by the quality of mirror surface needed for the experiment, cost, logistics, and maintenance. Depending on the application, the need for access to a recoating facility may exist, which would further increase costs. In the past decade it has been demonstrated that large mirrors of high quality can be made inexpensively if liquid mercury is rotated in a container. The idea has been known for decades, but many practical issues needed to be addressed to make the operation of large-diameter liquid mirrors robust enough for routine geophysical and astronomical applications. The effort to make the liquid-mirror technology practical has been lead by E. F. Borra's group at Université Laval. A recent review by Borra[5] documents the history of liquid mirrors and the developments that have led to their practical use for scientific applications.

The basic idea behind liquid mirrors is to spin a container containing a liquid metal, say mercury, on a turntable. The liquid will assume a near-parabolic shape with a focal length L, given by the equation

$$L = \frac{g}{2\omega^2}, \tag{1}$$

where g is the magnitude of the acceleration due to gravity and ω is the angular velocity of the turntable. Liquid mirrors offer many advantages for lidar work, including a large area for a low cost, high reliability, and ease of cleaning. Mercury vapor from the surface is not a problem except when the mercury surface has been disturbed, for instance, after being cleaned. At these times protective masks are worn in the telescope room, but within a day the mercury levels return to approximately an order of magnitude below the current occupational-safety guidelines. The large liquid mirror assembled for the lidar system described here did create some engineering challenges, associated with its large aperture, because of the large photocount signal from low altitudes.

2. Description of the Rayleigh-Scatter System

The Purple Crow Lidar (PCL) is located at the Delaware Observatory (at 42°52′ N, 81°23′ W, and 225 m in elevation), which is approximately 30 km southwest of The University of Western Ontario in London, Ontario, Canada. The PCL laser system employs a seeded Nd:YAG laser with two output beams. One beam is used for Rayleigh-scatter measurements while the other beam simultaneously pumps a pulse dye amplifier that is seeded with a narrow-band tunable ring–dye laser for sodium-resonance–fluorescence measurements. The sodium system is used to determine the temperature in the mesosphere and the lower thermosphere through the measurement of the line shape of the sodium D_2 line and will be the subject of the second paper in a series started here.

The configuration of the transmitter and the receiver for the Rayleigh-scatter lidar (henceforth called simply the lidar) is shown in Fig. 1. The laser system is in a separate room, with the temperature controled to within a degree to help maintain the frequency stability of the system. On the laser table are the four lasers that compose the transmitter system: (1) an injection-seeded Nd:YAG laser, (2) a 6-W argon-ion laser used to pump the lasing medium in (3) a ring–dye laser and in (4) a custom-pulse dye amplifier. The Nd:YAG and sodium beams are each expanded; they then exit the laser room and are directed to the top of the tripod in the telescope room, which is at ambient temperature. The transmitted beams are then directed upward along the telescope's axis. The coaxial configuration is important to ensure that the field of view of the receiver system fully includes the area illuminated by the laser beam at all altitudes of interest.

The backscattered light incident on the liquid mirror is focused on an optical fiber. Just below the focus is a mechanical chopper, which blocks the intense low-altitude returns from the detection system. The optical fiber guides the light to the detector table, where it is collimated and split into the Rayleigh and sodium components by a dichroic mirror. The light then passes through the appropriate narrow-band interference filter and is directed to the photocathode of a photomultiplier, whose output is amplified and sent to a Stanford Research Systems multichannel scaler–averager, which includes a 100-MHz discriminator. The range-gated signal is then co-added for the desired number of laser shots and

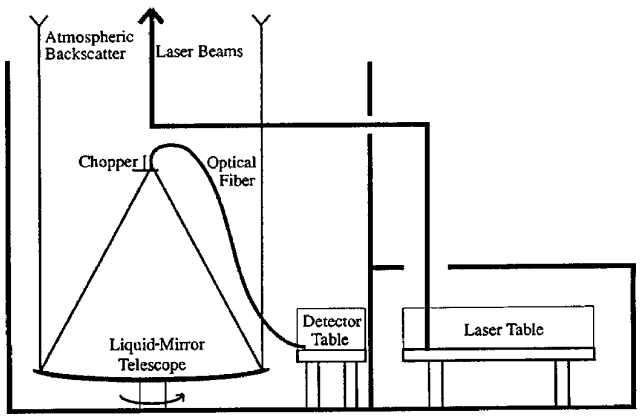

Fig. 1. Schematic diagram of the Purple Crow Lidar system. The laser table is 1.2 m × 3 m. The liquid-mirror telescope is 2.65 m in diameter and has a focal length of 5.175 m.

stored on a personal computer. Important system parameters are shown in Table 1, and specifics of the subsystems of the lidar are described below.

A. Receiver

1. Liquid Mirror

The receiving telescope is a 2.65-m-diameter liquid mirror, shown schematically in Fig. 2. The liquid-mirror technology offers the advantages of a large mirror at a fraction of the cost of a conventional glass telescope. The equipment and materials needed to build the 2.65-m-diameter mirror cost approximately $30,000, albeit with a considerable amount of labor costs extra! A comprehensive technical review of the liquid-mirror technology, including design, safety issues, and optical shop testing is given by Borra *et al.*[6] The container is made from foam covered with epoxy-laminated Kevlar and must be of a sufficient rigidity to avoid its bending, which could lead to instability and possible damage to the air bearing. The container is based on a design developed for the University of British Columbia–Université Laval 2.7-m telescope.[7] It was fabricated at the University of British Columbia by P. Hickson, L. Girard, and F. Arrien. The surface of the Kevlar container is approximately parabolic and is improved when a polyurethane resin is spin cast at the period appropriate for the desired focal length (approximately 10 rpm for our mirror).

A thin layer of mercury is then poured into the container. A minimum amount of mercury must be used in order to avoid overloading the air bearing, in addition to damping any disturbances that are caused by wind or building vibrations and whose magnitude increases with an increasing depth of mercury. Without the use of any special techniques a mercury surface thickness of approximately 4 mm could be achieved, but the weight of that amount of mercury would exceed the load limit of the air bearing. Several techniques have been developed to decrease the surface layer to less than half this amount. For the PCL mirror, a channel that is approximately 12 mm in width and in depth is formed around the outside of the mirror during the spin-casting process. The channel ensures that there is a sufficient depth of mercury at the liquid–surface interface, where the disruptive effects of surface tension are largest.

Table 1. PCL Rayleigh-Scatter System Parameters

Parameter	Value
Transmitter	
Wavelength	532 nm
Energy	600 mJ/pulse
Repetition rate	20 Hz
Line width	7 ns FWHM
Beam divergence at e^{-2}	0.2 mrad, full width
Beam diameter	27 mm
Receiver	
Aperture	2.65 m
Focal length	5.175 m
Obscuration of the liquid-mercury mirror	0.11
Bandwidth	1.0 nm
Height resolution[a]	24 m
System efficiency[b] (with an optical fiber)	1.2%
Field of view (full width)	0.39 mrad

[a]This value represents the nominal height resolution; it is adjustable from 6 m to >7.5 km in discrete steps.

[b]This value includes the obscuration ratio and the reflectivity of the liquid-mercury mirror.

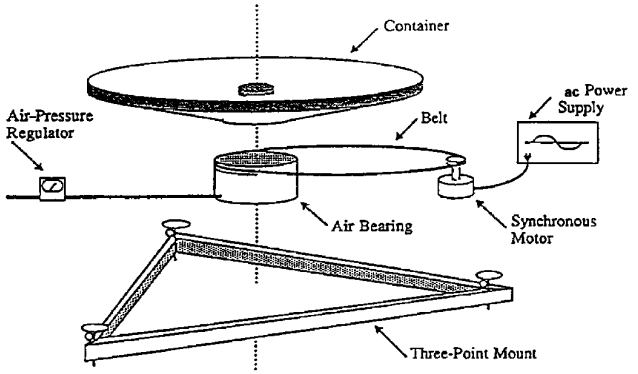

Fig. 2. Basic components of a liquid mirror.

2. Telescope Optics

An Optech, Inc., mechanical chopper is positioned just below the focal point of the mirror (Fig. 1). The mechanical chopper is important, as it eliminates signal-induced noise from the return signal. The chopper motor is mechanically isolated from the mirror tripod assembly to avoid vibrations that could affect the optical alignment. As a contingency, the system is designed to operate either with or without the chopper system. The photomultipliers can be electronically gated if the focusing dynode voltage is switched to protect them from the low-altitude returns.

The use of a mechanical chopper requires the laser system to be triggered externally from a signal derived from the chopper-position sensor. Several other functions depend on the external trigger, so a control unit was designed and built for this purpose. The chopper-trigger signal is used by the control unit to fire the Nd:YAG flashlamps and Q switch externally, to gate the phototubes, to synchronize the multichannel scaler–averagers, and to interface to a radar system that can shut the laser off in the event of an aircraft overflight.

Two optical systems have been used to redirect the light collected by the mirror to the detectors. Initially the light was collimated and directed down to an optical bench by means of lenses and mirrors. Because of the large distance between the focus of the telescope and the detector table, this simple design was inefficient. The efficiency of this system could have been improved by the addition of field lenses, but a simpler and less-expensive alternative was to install a 2-mm-diameter, single-strand plastic optical fiber. The light exiting the fiber is collimated by a condenser lens on the detector table and then encoun-

ters a dichroic mirror angle of 45° relative to the incoming beam. The dichroic mirror reflects the 589-nm sodium light and transmits the 532-nm Rayleigh-scatter light. The two light beams then pass through the appropriate interference filters and into the photomultiplier tubes.

3. *Detectors*

The lidar's large power–aperture product imposed very stringent requirements on the detector system. Of particular importance is to have a large dynamic range with an acceptable linearity. An EMI-9817B multialkaline photomultiplier tube was selected because of its large dynamic range. This tube can operate at a high anode current, which permits the use of a low-resistance voltage divider with an increased bias voltage in the final three stages. This design yields a lower dependence of the anode current pulse on the input signal.

A wideband pulse preamplifier allows the commercial multichannel-scaler averages to be located in the control room, away from the optical bench. The multichannel-scaler averages, which are used to range-gate the incoming signal, have a measured dead time of 5 ns. The incoming measurements are transferred from the averagers by means of a general-purpose interface bus (GPIB) bus (IEEE 488) and recorded on a personal computer, along with relevant header information.

3. System Characterization

A. Laser Performance

The laser system has performed consistently and reliably thus far. The output power of the laser is generally within specifications, and the overall power drift over a night is less than 5%. Injection seeding of the Nd:YAG is not critical for the Rayleigh-scatter measurements, but does provide a small increase in the power of the laser, in addition to improving the pulse shape.

B. Liquid-Mirror Performance

1. *Surface Quality*

Detailed measurements of a 1.5-m- and a 2.5-m-diameter liquid mirror through standard optical shop tests, such as the knife-edge and Ronchi tests, in addition to more sophisticated interferometric tests, have shown the liquid-mirror surface to be near-diffraction limited.[8] Presently three 2.5-m-class liquid mirrors are operating in Canada: one at the University of British Columbia, one at Université Laval, and one at The University of Western Ontario. The University of British Columbia's liquid mirror was designed for astronomical observations and has produced stellar images limited only by the local seeing conditions. The best of these images has an angular full width at half-maximum that is less than 2″.[9] The mirror at Université Laval is operated in a laboratory designed for optical shop tests on large mirrors. Interferometric measurements on this mirror have shown the quality of the mirror to be $\lambda/20$, which results in a diffraction-limited resolution of the order of 0.1″.[7] A 3-m liquid mirror used to study space debris has also been built by the Johnson Space Flight Center. The University of Western Ontario's liquid mirror was crudely tested for image quality through the observation of stars with a commercial-grade video system. Seeing of approximately 7″ was estimated from these images, which is roughly 20% higher than expected for rural southern Ontario. Part of this discrepancy may be due to the temperature of the telescope room, which was not at the ambient temperature for this test. The analysis of the images was further complicated by a lack of a proper alignment mount for the camera, which caused errors in the position of the detector at the prime focus and added coma to the star images. The star measurements were terminated in the early morning as a result of ground fog, so a high relative humidity may also have contributed to poor seeing. A more thorough investigation of the image quality for this mirror is not warranted, as the mirror is not intended for imaging work and is of sufficient quality on the basis of this and other tests for the PCL requirements.

2. *Quality of the Mirror as a Lidar Receiver*

An important test of the mirror is its performance as part of the lidar system. Does the full aperture of the mirror contribute to the Rayleigh-scatter signal or do surface defects limit the effective aperture? The uniformity of the mirror was evaluated through the use of a series of opaque masks placed over the mirror. First, the lidar returns were obtained with the full available mirror surface over a period of 10 min in the normal configuration. An integration time of 10 min was used to provide an adequate signal-to-noise ratio as well as to minimize the likelihood of significant changes in the atmospheric transmission. A 1-m-diameter mask was then put over the central part of the telescope, which decreased the effective area by an additional 11%. Rayleigh-backscatter returns were again acquired over a 10-min period with the 1-m mask, and then a 2-m-diameter mask was put over the central part of the telescope, decreasing the effective area by 57% as compared with the normal configuration. Measurements were then obtained over a 10-min period, which was followed by another series of measurements over the 10-min period in the normal configuration as a check on the stability of the optical alignment. The results of this test, shown in Table 2, demonstrate that the full effective area of the mirror is contributing to the return photocounts, as in both cases the total received photocounts over the 10-min integration decreased in the proportion expected for the different mirror areas.

Table 2. Uniformity Test Results for the Liquid Mirror

Operating Conditions	Area (m^2)	Ratios	
		Predicted	Measured
Normal (no mask)	4.66	1.00	—
Masked			
1-m Mask	4.14	0.89	0.88 ± 0.06
2-m Mask	2.02	0.42	0.37 ± 0.05

3. *Robustness of the Mirror*

The mirror has been in near-continuous operation since August 1992. It has continued to operate through the heat of summer and the cold of winter at ambient temperatures. A major difference between liquid mirrors and conventional glass mirrors is the definition of their "breaking." When a conventional glass mirror is struck with a heavy object, it breaks in a catastrophic and irreversible manner. A liquid mirror breaks much more easily, but reversibly. Breaking a liquid mirror means that the surface tension of the mercury is disrupted (hence broken), causing the mercury to flow into the center of the container. The mirror can then be quickly restarted ("closed"), often in less than 15 min. A pictorial sequence of the closing procedure for the mirror is shown in Fig. 3. The top-left photograph illustrates the broken condition. A stable liquid mirror usually breaks only when the rotational velocity is seriously perturbed.

The ease of reformation of the mirror is a significant advantage of liquid mirrors. A conventional glass mirror must be recoated in an expensive, often

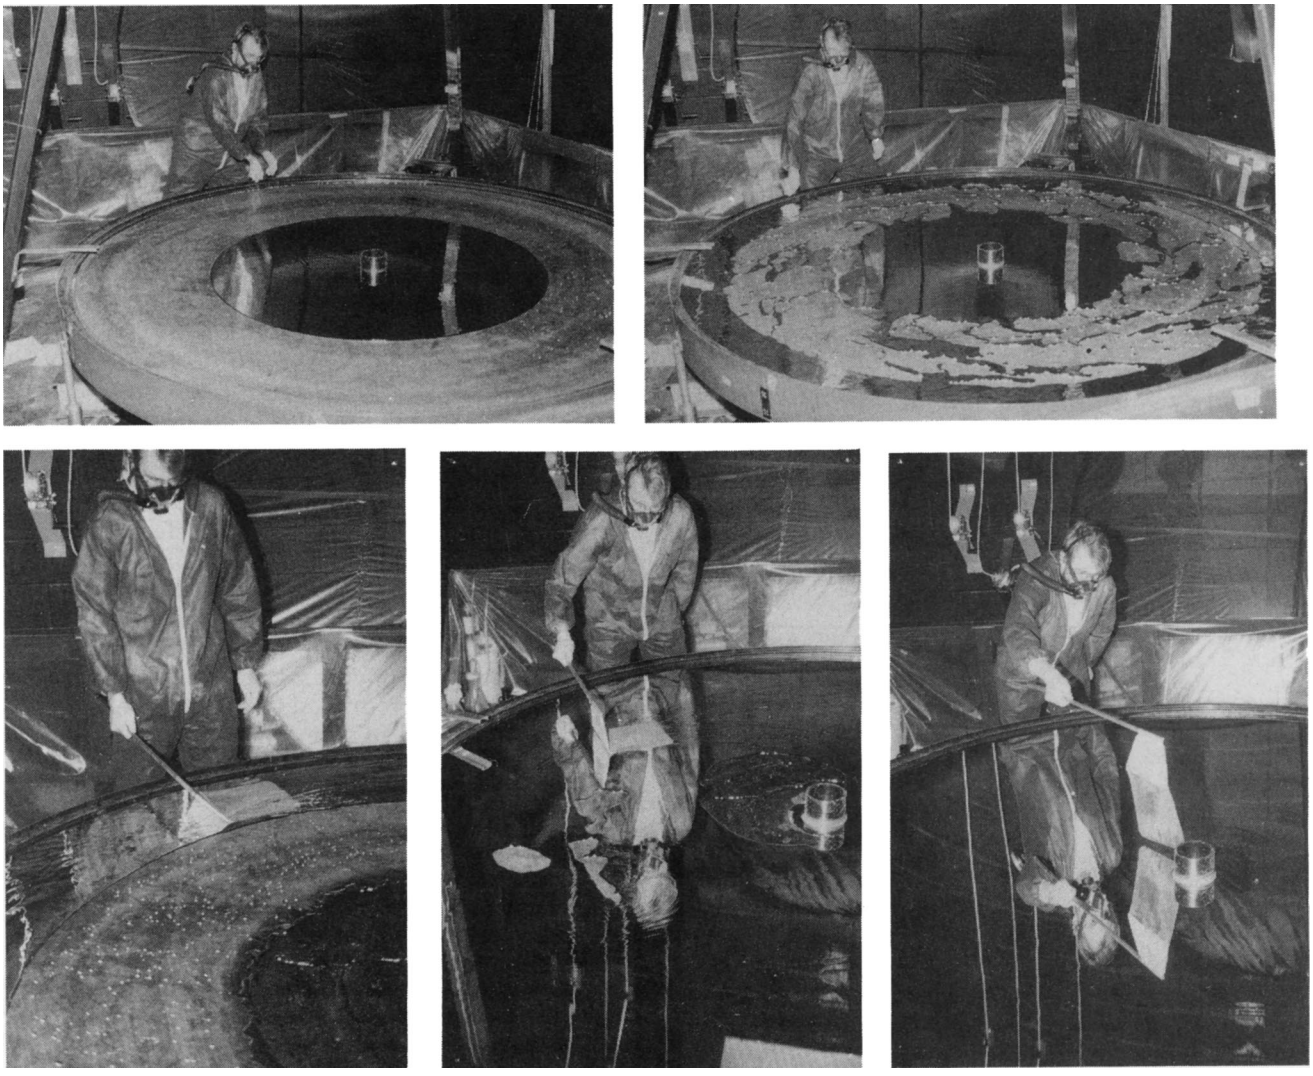

Fig. 3. Closing procedure for the mercury surface of the 2.65-m liquid mirror (read the sequence from the top-left photograph to the lower-right photograph). In the top-left photograph, the operator has just cleaned the mercury in the center of the container by skimming the surface with a weighted piece of tubing and then vacuuming off the dirty mercury, which is filtered and reused. The operator begins to spin the container sufficiently fast to move the mercury out to the edge of the container (top-right and lower-left photographs). The middle-bottom photograph shows a Mylar flag as it is used to drag the mercury toward the center of the container. After a few minutes, when the surface is closed, the flag is carefully removed and cleaned.

dedicated, facility as its reflective coating degrades. This procedure is complicated and expensive, particularly for large telescopes, which often need to have recoating facilities built into the observatory. The liquid mirror can be quickly and simply cleaned when a weighted piece of plastic tubing is dragged across the surface, thus skimming the dirty mercury from the surface when the mirror is stopped and the mercury has pooled to the center of the container.

Figure 3 shows the closing procedure during which the mirror is spun by hand at approximately twice the nominal period to drive all the liquid to the edge of the container. The surface is then closed by the operator's dragging a sheet of Mylar across the surface to spread the mercury toward the center of the container. The mirror does not instantly regain its equilibrium focal point once a closed surface is established. Measurements of the focal point on a 1.5-m-diameter liquid mirror show a decrease of approximately 2 mm/h over a 4-h period, which is correlated with a decrease in the amount of mercury vapor above the mirror surface.[6]

It was expected that the focus-stabilization time would be significantly longer for the 2.65-m-diameter mirror. The stabilization time of our mirror was measured with Rayleigh-scatter returns after the surface was cleaned and closed. Figure 4 shows the sum of the photocounts between 30 and 70 km, as a function of time, commencing 1 h after the mirror was closed. The step-like increases in the photocount totals at 85 and 115 min correspond to adjustments of the laser beam–telescope alignment. Approximately 130 min after startup the total number of counts were at nominal values. The rapid stabilization time is probably related to the modest imaging requirements of the system.

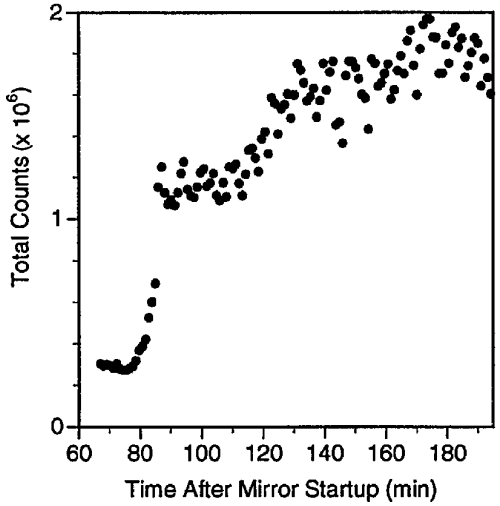

Fig. 4. Total Rayleigh-scatter photocounts between 30 and 70 km measured during the stabilization of the liquid mirror. This test shows that the liquid mirror's focal length is sufficiently stable for lidar observations approximately 130 min after the mirror surface has been closed.

C. Corrections to the Raw Backscatter Profiles

Implicit in the determination of Rayleigh-scatter density fluctuations and temperatures is that any signal corrections, such as the removal of background resulting from the dark count, scattered light, signal-induced noise, detector linearity, switching delays, or mismatches in the transmitter and the receiver fields of view, are properly accounted for. The last point is not a concern in our system, because the laser beam is coaxial with the telescope axis. The other issues mentioned above must be addressed. Normally the system is run with the mechanical chopper and thus does not require the photomultipliers to be gated. However, some of the earlier measurements were made without the chopper and with dynode gating. An additional correction to the photocounts is required when electronic gating is used. Background removal is important at the upper altitudes, particularly when the signal-to-noise ratio of the measurement becomes less than ≈10. Detector characterization is important to ensure that the relative shape of the profile is determined correctly at the high-photocount rates incident on the photocathode from low-altitude returns.

1. Background Removal

The lidar system normally collects data from the ground up to 147 km. Above the level at which Rayleigh returns are no longer detected, nominally 110 km, the photocounts are used to estimate the background. The background level depends on the operating mode of the system. When the chopper is used, the background is small and is due mainly to sky background with a small contribution from the photomultiplier dark counts. This type of background is constant with time over the nominal 1-min integration performed and is therefore easily removed from the signal through the simple determination of the average of the background signal between 120 and 147 km. When photomultiplier dynode gating is used a large amount of signal-induced noise is present because of phototube after pulses due to the large low-altitude returns, despite the excellent hold-off characteristics of the phototube. This background is more difficult to remove because it changes with height (i.e., time). The signal-induced noise varies exponentially with time with a large time constant, so in practice it can be reasonably approximated by a straight line and removed from the signal.

2. Detector Linearity

The gain of the photomultiplier tube varies as a function of time when dynode switching is employed. It is therefore necessary to correct the data for this effect. Variations in gain were measured when light from a stable source was directed into the photomultiplier and the detector system was run (Fig. 5). The results show that the gain initially increases until an equivalent altitude of 24 km is reached, then decreases with height. The increase in gain is below the nominal height of 30 km, above which the photo-

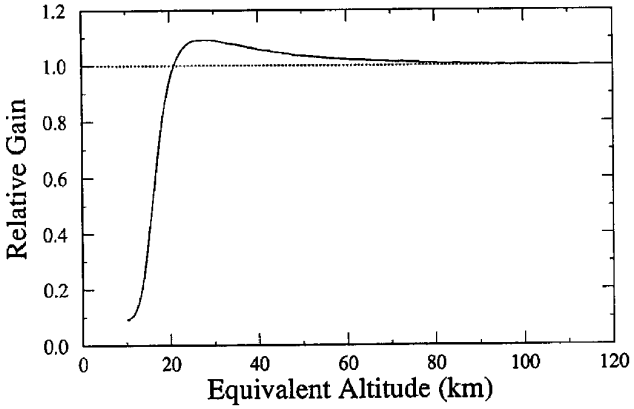

Fig. 5. Measured gain variation resulting from dynode-gating the photomultiplier with a constant, stable light source. The gain of the photomultiplier decreases 8% over the height range of interest for the Rayleigh-scatter measurement in this mode of operation. The solid curve represents the gain-variation measurement, and the dotted line indicates unity gain.

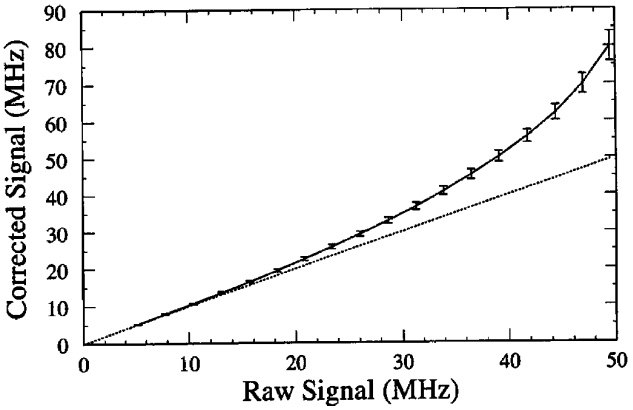

Fig. 6. Correction curves from five nights and seven LED calibration runs conducted during the past year through the use of the ratio of counts with and without a sufficient neutral-density filter to keep the count rates linear. This correction introduces an uncertainty in the accuracy of the measurements at the lower heights that is proportional to the vertical bars shown at the $\pm 1\sigma$ level of variation.

count returns are considered uncontaminated by aerosols and the scattering is fully attributed to molecules. The gain decreases in the region of interest for temperature measurements taken above 30 km. The measured gain decrease in this region is used to correct the count rates when electronic gating is necessary.

The returns from the stratosphere are large enough that the signal recorded is less than the true signal because of photomultiplier paralyzation and pulse pile-up effects. It is necessary to be able to correct for these effects because the shape of the photocount profile is critical for the proper retrieval of density and temperature. For the calculation of a correction factor for these effects, measurements have been taken both with and without various neutral-density filters in front of the detector. The linear response of the detector with the attenuated signal can be compared with the nonattenuated signal to calculate the correction as a function of the count rate. This correction has been measured with both the Rayleigh-scatter signal and measurements of the photocounts from a light-emitting diode (LED). The LED is cycled with a sawtooth pulse that is similar in shape and intensity to the sky returns. These two methods for the determination of this correction give consistent results. Measurements of the correction on five nights (6.2 h total) and during seven different days with the LED source (16 hours total) have been combined into a correction curve that gives the corrected counts in terms of the measured raw counts (Fig. 6). The error bars in the curve are $\pm 1\sigma$ standard deviation of the variation of the 12 calibration runs used to determine the correction. The correction becomes significant above a count rate of 2.6 MHz, when the measurements begin to deviate from the dotted line that indicates the linear relation between output and input. Because of the variations in the raw versus the corrected counts evident from Fig. 6, the correction introduces an accuracy uncertainty in the temperature retrieval, which is discussed in Subsection 3.C.3. When simultaneous sodium-resonance–fluorescence measurements are available they provide Rayleigh-scatter returns at the lower altitudes with much lower count rates, and these can be used to obtain independent temperature measurements in the stratosphere.

3. *Overlap Function*

Because the temperature is found from the relative shape of the density profile, it is important that the return beam be completely within the receiver field of view or that variations in the overlap are well characterized. In practice it is extremely difficult to correct for variations in the overlap function, and so it is preferable to ensure that the laser beam is contained within the field of view of the receiver throughout the altitude range over which the measurements are obtained. The PCL was designed to have a unity overlap function over the height range of interest.

The laser-beam divergence is specified by the manufacturer as 0.45 mrad full width for 86% energy, which, when passed through the $3\times$ beam expander, yields a divergence of 0.15 mrad full width for the transmitted beam. The spot diameter at the focus of the telescope is then the focal length (5.175 m) times the beam divergence, that is, 0.78 mm. The location of this spot varies its position along the optic axis as a function of altitude. For the PCL the range of interest is nominally 30–150 km of altitude, over which the separation of the minimum spot-size position varies 0.71 mm. The lidar is set up so that the optical fiber is located at the position for the minimum spot size for light backscattered from 30 km, where the spot size is 0.78 mm. Thus, the spot size for light backscattered from 150 km will be larger by 0.71 mm times the focal length of the telescope, or 1.1 mm. The laser-beam-pointing stability causes the focal spot of the backscattered light to wander

across the fiber with variations of approximately 0.5–0.75 mm over a few seconds. This alignment results in a focal spot that varies in diameter from 0.78 mm for returns from 30 km to 1.1 mm for returns from 150 km, with a center position that varies approximately 0.75 mm. This spot is input on to an optical fiber with a active collection diameter of 2 mm. Hence an adequate safety factor exists to ensure that the light from all heights is incident on the optical fiber. In fact, the variations in the position of the focal spot on the fiber that are due to laser jitter helps to smooth out any variations that may exist in the transmission of the fiber over its entrance area. Visual confirmation of the spot size and variation from the 30-km returns support these simple calculations.

D. Overall System Performance

The overall performance of a lidar system can be estimated from the lidar equation, which permits the calculation of the number of return Rayleigh-scatter-signal photocounts as a function of height, $N(z)$, to be found from the relation[10]

$$N(z) = \frac{\beta[\sigma_R \rho(z)]}{z^2}, \tag{2}$$

where z is the altitude, β is proportional to the product of the following: system efficiency, the square of the atmospheric transmission, the laser power, the wavelength, the integration time, and the telescope area, σ_R is the Rayleigh-scatter cross section (5.15×10^{-31} m^2), and $\rho(z)$ is the atmospheric density (Table 1). The calculated Rayleigh-scatter signal can then be compared with a measurement. The atmospheric transmission, which enters as a square in β as a result of the passage of the incident beam up and then back through the lower atmosphere, cannot be independently determined by a single-frequency sounding, thus requiring an independent normalization of the absolute-density measurements. Furthermore, the β term is a function of time, since the atmospheric transmission and laser power can fluctuate over the integration time of the measurement.

A comparison between the photocount calculated from the lidar equation and a measurement with the background removed is shown in Fig. 7. The measurement is obtained over a period of 364 min, but as a general comment photocount profiles can be scaled by the measurement period, whereas density and temperature statistical errors can be scaled by the square root of the measurement period. The agreement between measurement and calculation is excellent, within approximately 10%. An atmospheric transmission of 0.58 was used for the calculation. The atmospheric transmission was calculated from the LOTRAN 7 program run with a rural aerosol model.[11] The actual atmospheric transmission varies for a variety of reasons, including the humidity, the aerosol concentration and type, and variations in the composition of different air masses. The system efficiency may also be less than optimum because of the presence of dust and insects on the optical surfaces, which is inevitable in the telescope room because it has a 3-m hole in the ceiling during operation. Nevertheless, this comparison supports our contention that the liquid mirror's performance is consistent with its reflectivity and area.

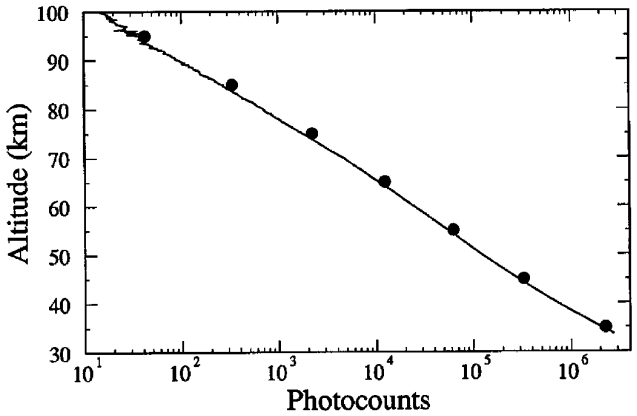

Fig. 7. Comparison of measured and calculated photocounts: The solid curve represents 364 min of Rayleigh-scatter returns at a height resolution of 24 m obtained on 31 August 1994, with the background removed as described in Section 3. The measurements have been smoothed with a 15-point running average. The circles are the anticipated returns calculated from the lidar equation (2) with the system parameters listed in Table 1 and the appropriate **MSIS-90** model-atmosphere density profile. The agreement between the measurement and the calculation values are within approximately 10% and suggest that the liquid-mirror receiver's performance is equivalent to that of a conventional glass mirror with the same reflectivity.

4. Retrieval of Density-Fluctuations and Temperatures

A. Methodology

The technique for the determination of density fluctuations and temperature from the Rayleigh-scatter returns is straightforward in concept but subtle in practice. The useful range of measurements is typically between 30 and 100 km. Below 30 km the Rayleigh-scatter signal increases, but aerosols cause additional scattering that cannot be independently estimated to permit separation of the Rayleigh-scatter component. Above 100 km or so the atmospheric density becomes sufficiently low that an unacceptable signal-to-noise ratio is available to determine temperatures. Temperatures are found from the density profiles with the assumption that hydrostatic equilibrium exists between adjacent atmospheric layers, by application of the ideal gas law, integration of the resulting equation, and solution for temperature. The limits of the integration are from the height of the measurement to the top of the atmosphere, which in practice is the height of the highest available measurement. The integration of the ideal gas law results in a constant of integration, which requires that either the pressure or the temperature at the greatest height be known, either from an indepen-

dent measurement or an atmospheric model. A particularly clear exposition of this technique is given by Chanin and Hauchecorne.[12] Their formalism is particularly well suited for temperature retrievals, as the density integration is performed over pressure surfaces rather than height surfaces. An example of the retrieved temperature for a 6-h measurement period is shown in Fig. 8. Most of the variation seen in the profile is above the noise limit and is due to wave activity in the atmosphere. This particular night shows excellent agreement with the MSIS-90 model[14] (the dotted curve). Other nights show large variations with interesting temporal and spatial structures.

Many of the studies planned with the PCL involve the measurement of rapid changes in the atmosphere that are due to the transport of energy and momentum by waves traveling from the lower to the middle atmosphere. Signal-processing algorithms to determine density fluctuations have been given by Gardner *et al.*[13] In their scheme, rather than direct integration of the relative-density profile, the profile is detrended and the residuals corrected for bias that is due to atmospheric transmission variations and laser power fluctuations. The resulting spatial–temporal time series can then be spectrally analyzed. The fluctuations are detrended by the removal of a fifth-order polynomial fit to the average-density profile. A fifth-order polynomial is used to achieve a fit to the average profile, which is sufficiently (i.e., at least 10 times) better than the size of the atmospheric fluctuations, which typically range from less than 0.1% to 15% or more. The median is then removed from the residuals. An example of the temperatures obtained from these density fluctuations is shown in Fig. 9. The average-temperature profile shown in Fig. 8 was calculated from the total counts obtained over the 6-h measurement time. Figure 9 shows 10-min-interval measurements at three selected heights to emphasize both the high signal-to-noise ratio of the measurements as well as to demonstrate the rapid fluctuations in temperature on this night,

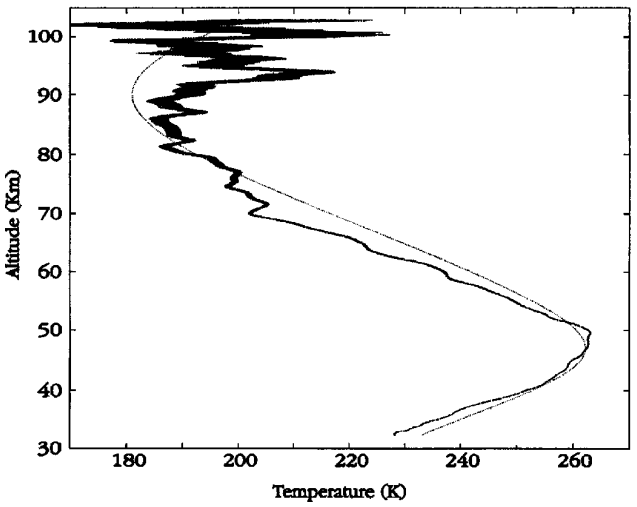

Fig. 8. Rayleigh-scatter temperature measurements for the same measurement period described for Fig. 7. The measurements were obtained at a 48-m height resolution then smoothed with a digital filter with a 1-km cutoff. The dotted curve is the MSIS-90 model-atmosphere temperature profile for the geophysical conditions appropriate for that night. The shading is the $\pm 1\sigma$ statistical error of the temperature. The temperature-profile integration commenced at a signal-to-noise ratio of 2 at a height of 103 km.

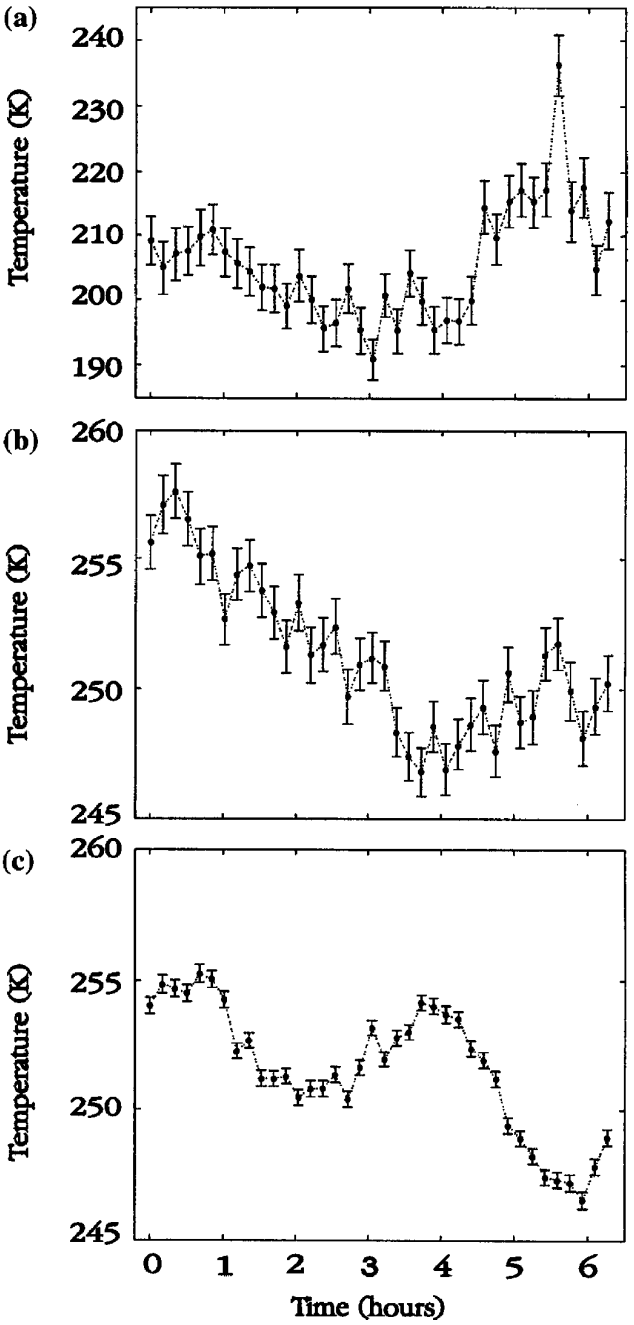

Fig. 9. Temperature variations at a height resolution of 48 m at three selected altitudes, (a) 72 km, (b) 54 km, and (c) 40 km, at 10-min intervals over the same period described for Fig. 8. The 10-min profiles have been processed similarly to the average profile shown in Fig. 8. The ordinate value of 0 h is 0332 UT (midnight Eastern Daylight Time is 0400 UT). Temperature variations of the order of 10 K, significantly above the measurement error, are apparent.

despite the relatively smooth appearance of the average-temperature profile. These time series show large and statistically significant variations in the temperatures in the stratosphere and mesosphere that are connected with lifting and subsiding of the pressure surface by atmospheric waves. A detailed spectral analysis of the density-fluctuation measurements on this and other nights is in progress to help improve our understanding of the behavior of waves in the middle atmosphere.

B. Density Normalization

Rayleigh-scatter lidars, like falling spheres, grenade sounding, or other similar techniques, measure the relative density as a function of height. The absolute density is found from either ancillary measurements, such as those provided by radiosondes, or by normalization of the relative densities to an atmospheric model. Often density profiles are normalized over a small range of height intervals or at single locations. This procedure may have the undesired effect of weighting a profile too heavily to one location, which could be strongly affected at any instant by atmospheric variations. The Rayleigh-scatter measurements are normalized with the total column density between 45 and 60 km found from the MSIS-90 model atmosphere for the appropriate geophysical conditions.[14] The altitude range between 45 and 60 km is used because the total error of the lidar measurement is smallest in this region, the count rates are reasonably high but linear, and no aerosols are present.

C. Temperature Uncertainties

1. *Initial Temperature Guess*

Uncertainties in the relative densities are primarily due to the Poisson noise of the received photocounts. Because the temperature-retrieval method involves an integration, a boundary condition must be supplied to define the constant of integration. This constant requires the temperature or pressure at the highest altitude to be supplied, as the temperature at the greatest height is multiplied by the ratio of the density at the desired height to the density at the greatest height. In the absence of an independent measurement of the temperature at the top of the measurement region, a temperature from a model atmosphere must be assumed. The uncertainty of this assumption, σ^{top}, decreases with decreasing height as

$$\sigma^{\mathrm{top}}(z) = |(T^{\mathrm{top}} - T_{\mathrm{model}}{}^{\mathrm{top}})| \frac{\rho(z^{\mathrm{top}})}{\rho(z)}, \qquad (3)$$

where T^{top} is the environmental temperature and $T_{\mathrm{model}}{}^{\mathrm{top}}$ is the temperature obtained from a model at the greatest (i.e., the top) height available. In addition to the above uncertainty, this term contributes a statistical error to the temperature because of the statistical error in the density ratio. Measurements of temperature from various sources, primarily rocket experiments, have been statistically analyzed by Gringorten *et al.*[15] The measurements from midlatitude differ by approximately 20 K from the median (80 km) 10% of the time. Because the atmospheric density decreases by nearly an order of magnitude every 16 km, this 20-K uncertainty in the initial temperature would be approximately ± 2 K at 64 km and approximately ± 0.2 K at the stratopause. On the other hand, if the model temperature is accurate this error is ± 0 K! Hence, the uncertainty in the assumed initial temperature may be significant for temperature studies in the mesosphere, particularly for atmospheric-change studies. For a system with a high power–aperture product, such as the PCL, this uncertainty may be of the order of the statistical error. Future PCL measurements will address this problem through the use of the simultaneous absolute-temperature measurements from the sodium-resonance–fluorescence temperature experiment to seed the integral. A temperature accuracy of less than 1 K is expected because of instrumental effects for the sodium line-shape measurements. At the present time the uncertainty of the temperatures that is attributable to the initial temperature estimate is not included in the error budget because there is no convenient way to independently determine the precision of the model.

2. *Temperature Accuracy and Precision*

The temperature-error sources are statistical errors that decrease as the square root of the photocounts and accuracy errors that are due to detector nonlinearities at high count rates, as discussed in Section 3. The variation between the correction curves shown in Fig. 6, when converted to an uncertainty in temperature, is shown as a function of height in Fig. 10 (solid

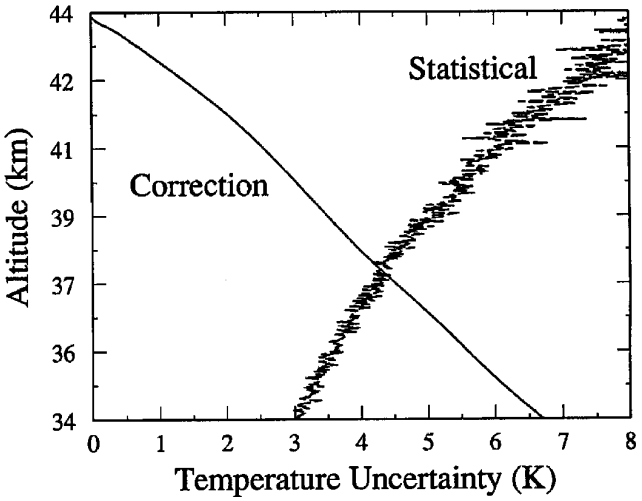

Fig. 10. Uncertainty resulting from photomultiplier corrections (solid curve) compared with the statistical error (unsmoothed curve) for a 1-min interval on 31 August 1994. The correction uncertainty is due to variations at high count rates for the photomultiplier-correction curve shown in Fig. 6. The correction uncertainty is independent of the integration time, whereas the statistical error scales as the square root of the photocounts.

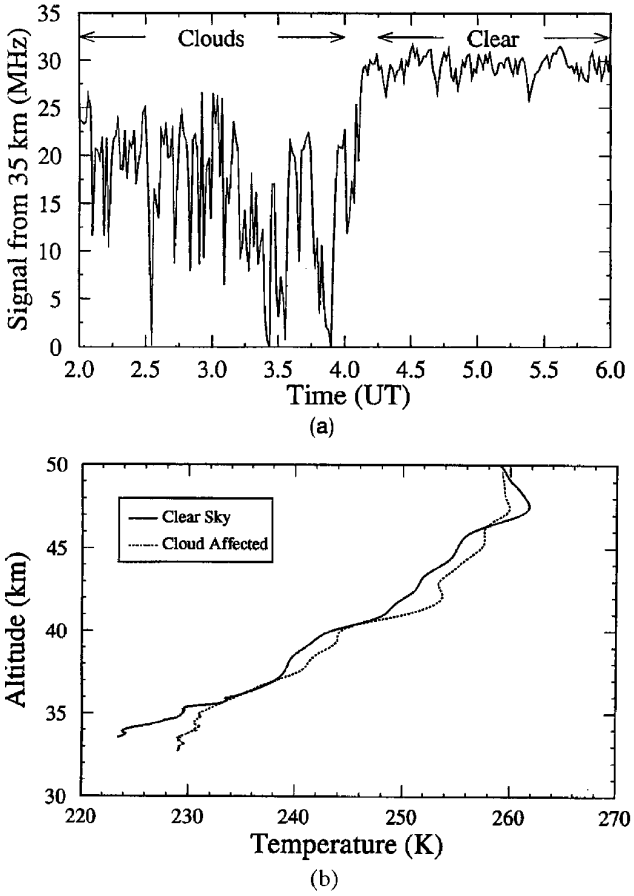

Fig. 11. Estimate of the temperature error through a comparison of two temperature-measurement intervals: (a) Rayleigh-scatter signal at 35 km altitude on 30 September 1994 from 2035–0034 solar local time (0200–0600 UT) showing a period of varying cloud cover followed by clear skies. (b) Stratospheric temperature measurements averaged over the clear period (solid curve) and the cloudy period (dotted curve). The difference between the two profiles below 35 km, where the correction of the photomultiplier counts is the largest, is consistent with the underestimation of the count rate during the cloudy period because of rapid atmospheric-transmission variations over the 1-min time interval in which the profile was acquired. The clear-sky measurements are in close agreement with the predictions of the MSIS-90 model.

curve) and compared with the statistical uncertainty in temperature without smoothing (i.e., unsmoothed solid curve) for a 1-min integration time. The contribution of the correction curve decreases rapidly below a count rate of approximately 5 MHz, which for nominal operations is a height of roughly 40 km. Thus, the correction uncertainty is the dominant uncertainty in the stratosphere because, unlike the statistical error, the correction uncertainty is independent of the integration time. The temperature error is entirely statistical above 50 km. This uncertainty scales as the square root of the photocounts. At altitudes below 50 km, for which the linearity correction is applied, the statistical error carried throughout the calculations is the fractional error of the raw measurement.

3. *Effects of Clouds on Temperatures*

When photocount profiles are corrected for nonlinearities, caution must be exercised if the profiles have been obtained during periods of large atmospheric-transmission variations occurring on shorter time scales than that over which the individual laser shots are co-added. If the photocount profiles are co-added over a series of laser shots during which large atmospheric-transmission variations occur, the photocount correction at low altitudes will underestimate the true signal. The correction shown in Fig. 6 was made with the assumption that all scans within a record sample the same signal. If measurements are collected over 1200 laser shots (1 min) and a cloud passes overhead obscuring the beam for part of the record, the integrated signal will be undercorrected at the lower heights, causing an overestimate of the temperature. For an estimate of the error in temperature caused by this effect under actual operating conditions, the temperature has been calculated for two periods on 29 September 1994. The first period was during an interval of patchy cloud cover; the second period was after the sky had cleared [Fig. 11(a)]. Stratospheric temperatures measurements from these two periods show a reasonable agreement above 35 km [Fig. 11(b)]. At the very lowest heights, for which the nonlinear correction is largest, the temperatures taken during the clear period, which are very close to those anticipated from the MSIS-90 model, are approximately 6 K cooler than those taken during the cloudy period. This result is consistent with the underestimation of the nonlinearity correction that is due to rapid changes in the atmospheric transmission over the summation of the 1200 laser shots/1-min record.

5. Summary

The Rayleigh-scatter Purple Crow Lidar (PCL) has a power–aperture product system that is sufficient for high temporal–spatial resolution measurements of density and temperature over a large height range, from the stratosphere to the lower thermosphere. The large power–aperture product is obtained, in part, through the use of a 2.65-m-diameter liquid-mercury mirror. The high light-gathering power of this system created some interesting engineering problems to solve. Our experiences in the building, testing, and retrieval of the primary geophysical parameters from the Rayleigh-scatter system is summarized in the list below:

(1) A 2.65-m-diameter liquid mirror has been built and maintained in near-continuous operation since August 1992.

(2) Masking the mirror aperture has shown the liquid mirror to perform uniformly across its diameter.

(3) The detector system has been characterized to permit high photocount rates from the stratosphere to be corrected for detector nonlinearities. High-altitude returns have been characterized for their

relative contributions of signal-induced noise and background sky light for various instrumental configurations.

(4) The correction curve for detector nonlinearities is provisional at this time. The errors that are due to this correction (Fig. 10) will be reduced as more calibration measurements are performed and included in the correction average; this work is in progress. Simultaneous measurements of Rayleigh scatter from the sodium-resonance–fluorescence experiment at 589.1 nm will also provide an independent determination of the stratospheric temperature.

(5) The Rayleigh-scatter lidar returns are in excellent agreement with predictions from the lidar equation, suggesting that systems are functioning as anticipated.

(6) The methodology and some caveats on the retrieval of density fluctuations and temperatures are discussed, including the estimation of the temperature at the top of the profile and the temperature errors from count-rate corrections.

Presently more than 60 nights of useful measurements are being analyzed for scientific studies. The characterization of the sodium-resonance–fluorescence lidar, which is the second beam of the PCL transmitter, for the determination of kinetic temperature in the upper mesosphere and lower thermosphere is in progress. Initial temperature measurements obtained with this system are encouraging. The combination of these temperature measurements should provide comprehensive and consistent temperature- and density-fluctuation profiles over a wide range of heights at high temporal–spatial resolutions.

We would like to thank the National Research and Engineering Council of Canada for their support of this project through the Canadian Network for Space Research, a Collaborative Special Project, and the Research Grants Program. Special appreciation is extended to P. Hickson of the University of British Columbia for his help in the design and fabrication of the mirror container and the detector-tripod assembly. The design and building of the PCL was also made easier by the moral, technical, and scientific support of many individuals, including P. Sica, D. Moorcroft, W. Hocking, C. Gardner, D. Black, C. Castle, A. Carswell, J. Whiteway, S. Pal, L. Cogger, L. Freyer-Davis, C. Tepley, G. Shepherd, J. McKaig, A. French, J.-P. St.-Maurice, B. Smith, J. DuBois, S. Baverstock, P. Johnson, L. Caldwell, M. St.-Maurice, M. Williams, M. Roseman, A. Wienerman, R. De Serranno, and the technical-support group at Continuum and Coherent Lasers.

References

1. L. Elterman, "A series of stratospheric temperature profiles obtained with the searchlight technique," J. Geophys. Res. **58,** 519–530 (1953).
2. G. Fiocco and L. D. Smullin, "Detection of scattering layers in the upper atmosphere (60–140 km) by optical radar," Nature Phys. Sci. **199,** 1275–1276 (1963).
3. C. S. Gardner, M. S. Miller, and C. H. Liu, "Rayleigh lidar observations of gravity wave activity in the upper stratosphere at Urbana, Illinois," J. Atmos. Sci. **46,** 1838–1854 (1989).
4. R. Wilson, M. L. Chanin, and A. Hauchecorne, "Gravity waves in the middle atmosphere observed by Rayleigh lidar, 1. Case studies," J. Geophys. Res. **96,** 5153–5167 (1991); "Gravity waves in the middle atmosphere observed by Rayleigh lidar, 2. Climatology," J. Geophys. Res. **96,** 5169–5183 (1991).
5. E. F. Borra, "Liquid mirrors: a review," Can. J. Phys. (to be published).
6. E. F. Borra, R. Content, L. Girard, S. Szapiel, L. M. Tremblay, and E. Boily, "Liquid mirrors: optical shop tests and contributions to the technology," Astrophys. J. **393,** 829–847 (1992).
7. P. Hickson, B. K. Gibson, and D. Hogg, "Large astronomical liquid mirrors," Publ. Astron. Soc. Pac. **105,** 501 (1993); P. Hickson, E. F. Borra, R. Cabanac, R. Content, B. K. Gibson, and G. A. H. Walker, "UBC/LAVAL 2.7-meter liquid mirror telescope," Astrophys. J. Lett. **436,** L201–L203 (1994).
8. E. F. Borra, R. Content, and L. Girard, "Optical shop tests of a $f/1.2$ 2.5-meter diameter liquid mirror," Astrophys. J. **418,** 943–946 (1993).
9. P. Hickson, E. F. Borra, R. Cabanac, R. Content, B. K. Gibson, and G. A. H. Walker, "UBC/Laval 2.7-meter liquid mirror telescope," Astrophys. J. (to be published).
10. R. M. Measures, *Lidar Remote Sensing: Fundamentals and Applications* (Wiley, Ontario, 1984), Chap. 7.
11. F. X. Kneizys, E. P. Shettle, L. W. Abreu, J. H. Chetwynd, G. P. Anderson, W. O. Gallery, J. E. A. Selby, and S. A. Clough, "Users guide to LOWTRAN 7," AFGL Env. Res. Pap. 1010 (Air Force Geophysics Laboratory, Hanscom Air Force Base, Mass., 1988), Chap. 1.
12. M. L. Chanin and A. Hauchecorne, "Lidar studies of temperature and density using Rayleigh scattering," in *Handbook for MAP: Ground-Based Techniques,* Vol. 13 of the Middle Atmosphere Program Series (Scientific Committee on Solar Terrestrial Physics, International Council of Scientific Unions, Urbana, Ill., 1984), paper 7.
13. C. S. Gardner, D. C. Senft, T. J. Beatty, R. E. Bills, and C. A. Hostetler, "Rayleigh and sodium lidar techniques for measuring middle atmosphere density, temperature, and wind perturbations and their spectra," in *World Ionosphere/Thermosphere Study Handbook,* (Scientific Committee on Solar Terrestrial Physics, International Council of Scientific Unions, Urbana, Ill., 1989), Vol. 2, Chap. 6.
14. A. E. Hedin, "Extension of the MSIS thermospheric model into the middle and lower atmosphere," J. Geophys. Res. **96,** 1159–1172 (1991).
15. I. I. Gringorten, A. J. Kantor, Y. Izumi, and P. I. Tattelman, "Atmospheric temperatures, density, and pressure," in *Handbook of Geophysics and the Space Environment,* Doc. ADA 167000 (National Technical Information Service, Springfield, Va., 1985), Section 15.1.