CONF-960163--15

LA-UR- 96-333

Title: LASER SPECKLE EFFECTS ON HARD TARGET DIFFERENTIAL ABSORPTION LIDAR

Author(s): E. P. MacKerrow, J. J. Tiee, C. B. Fite, M. J. Schmitt, M. C. Whitehead, R. J. Nemzek, G. E. Busch, C. R. Quick, R. Remelius, P. J. Schafstall, and D. C. Thompson

Submitted to: Gas & Chemical Lasers Conf./SPIE PHOTONICS WEST '96, San Jose, CA , February 1996

Los Alamos
NATIONAL LABORATORY

Laser speckle effects on hard target differential absorption lidar

E.P. MacKerrow, J.J. Tiee', C.B. Fite, M.J. Schmitt, M.C. Whitehead, R.J. Nemzek,

G.E. Busch, C.R. Quick, D. Remelius, P.J. Schafstall, and D.C. Thompson

Los Alamos National Laboratory, Los Alamos, New Mexico 87545

ABSTRACT

Reflection of laser light from a diffuse surface exhibits a complex interference pattern known as laser speckle. Measurement of the reflected intensity from remote targets, common to "hard-target" differential absorption lidar, requires consideration of the statistical properties of the reflected light. We have explored the effects of laser speckle on the noise statistics for CO_2 DIAL. For an ensemble of independent speckle patterns it is predicted that the variance for the measured intensity is inversely proportional to the number of speckle measured. We have used a rotating drum target to obtain a large number of independent speckle and have measured the predicted decrease in the variance after correlations due to system drifts were accounted for. Measurements have been made using both circular and linear polarized light. These measurements show a slight improvement in return signal statistics when circular polarization is used. We have conducted experiments at close range to isolate speckle phenomena from other phenomena, such as atmospheric turbulence and platform motion thus allowing us to gain a full understanding of speckle. We have also studied how to remove correlation in the data due to albedo inhomogenuties producing a more statistically independent ensemble of speckle patterns. We find that some types of correlation are difficult to remove from the data.

Keywords: laser speckle, autocorrelation, remote sensing, lidar, CO_2 DIAL, averaging statistics, segmented averages, moving averages, polarization, and smoothing.

1. INTRODUCTION

One method of differential absorption lidar (DIAL) is to use the reflection from natural targets to enhance the return laser power. In this method the laser beam(s) are directed through the region of atmosphere which is being interrogated for chemical effluents. The difference in the return laser power between absorbed and non-absorbed wavelengths can be used to measure the concentration of the absorber.[2] In effect a DIAL measurement of this type is really a measure of the reflected laser light (accounting for absorption in the path of the laser beam).

Since natural or man-made targets are rough on the scale of a laser wavelength, path length differences will result across the footprint of the laser beam on the target. These path length differences lead to a complicated interference pattern at the receiver. This interference pattern,"speckle pattern", gives a "granularity" to the return laser light. The accuracy of an intensity measurement, made with one speckle pattern, increases with the number of speckle (correlation areas) which comprise the pattern. The individual speckle , inside a pattern, are statistically independent and follow random Gaussian statistics where the standard deviation obeys

$$\sigma_{N_s} \propto \frac{1}{\sqrt{N_s}}, \tag{1}$$

where N_s is the number of speckle lobes (correlation areas) in the speckle pattern. The number of speckle, N_s, is a

function of the size of each speckle and the entrance pupil of the receiver telescope. The speckle size is determined by the wavelength of the laser light, λ, the target illumination size, D_T, and the distance to the target, z_T. The speckle diameter, D_S, can be expressed analytically as[1]

$$D_S = \frac{4\lambda z_T}{\pi D_T}, \tag{2}$$

or more usefully in terms of the laser divergence, θ, as

$$D_S = \frac{4\lambda}{\pi\theta}. \tag{3}$$

For remote sensing applications it may be difficult to reduce the speckle size with the intent of improving the signal-to-noise ratio (S/N $\equiv \langle I\rangle/\sigma$). For example, a CO_2 laser operating at 10.6 μm, with a 1 mrad divergence, will have a speckle size of 1.35 cm in diameter. If the entrance pupil of the lidar receiver is 40 cm, as it is our case, the S/N would be 29.6 ($\sigma_1 = 3.37$ %); this can be found by calculating the number of speckle, the ratio of the receiver telescope area to the correlation area, and then using Eq. (1). Note that σ_1 is the standard deviation that one would measure for an ensemble of laser returns, each return being totally uncorrelated to the other returns. In general, σ_1 decreases with increasing divergence, but laser divergence requirements of remote sensing can set a lower limit for the single-shot standard deviation. Therefore, averaging over many independent speckle patterns is typically used to improve S/N.

If we have collection of speckle patterns, one pattern for each return laser pulse, then it may be possible to improve the standard deviation on an intensity measurement by averaging over the speckle patterns. What is meant here by averaging is a segmented average. For example, if we have a total number of laser pulses, Γ, we average between pulses $1 \rightarrow N$, (N+1) $\rightarrow 2N$, $(2N+1) \rightarrow 3N$, ..., $(\Gamma - N) \rightarrow \Gamma$; the result is a set of Γ/N values which make up the segmented N-pulse average. The improvement gained by averaging pulses depends on the correlation between the individual pulses, known as the temporal autocorrelation. In general the standard deviation of a segmented N-pulse average[4] is

$$\sigma_N = \frac{\sigma_1}{\sqrt{N}}\left[1 + 2\sum_{j=1}^{N-1}(1 - j/\Gamma)(1 - j/N)\,\rho_j\right]^{1/2}, \tag{4}$$

where the temporal autocorrelation, ρ_j, is given by

$$\rho_j = \frac{1}{\sigma^2}\langle I(t_k)\, I(t_k + j\tau)\rangle, \tag{5}$$

or

$$\rho_j = \frac{1}{\sigma^2(\Gamma - j)}\sum_{k=1}^{\Gamma-j}\left[\frac{P_k - \langle P\rangle}{\langle P\rangle}\right]\left[\frac{P_{k+j} - \langle P\rangle}{\langle P\rangle}\right]. \tag{6}$$

Here $\langle P\rangle$ is the average return laser power, P_k is the return laser power of the k^{th} pulse, τ is the time interval between laser pulses, and σ^2 is the "normalized variance" defined as

$$\sigma^2 \equiv \left\langle\left[\frac{P_k - \langle P\rangle}{\langle P\rangle}\right]^2\right\rangle = \frac{1}{\Gamma}\sum_{k=1}^{\Gamma}\left[\frac{P_k - \langle P\rangle}{\langle P\rangle}\right]^2. \tag{7}$$

The standard deviation given in Eq. (4), sometimes referred to as the "standard deviation of the mean", reduces to the result for a Gaussian distributed random variable, Eq. (1), when the temporal autocorrelation is zero. The temporal autocorrelation should not be confused with the spatial autocorrelation used in statistical optics used to describe the correlation area or speckle size.[3]

In an ideal situation we would like to have minimal temporal autocorrelation between the return lidar pulses; allowing for increased sensitivity by averaging over large data sets. In "real-life" remote sensing applications it

may be difficult to obtain data completely free of temporal autocorrelation. In order for each speckle pattern to be unique, the illumination spot (or angle of incidence) on the target must be completely different for each laser pulse. Two pulses are completely independent of each other if there is no overlap between their beam footprints on the target. Depending on the beam divergence, and distance to target, it may be difficult to obtain many independent speckle patterns. A good analogy (made by G. Busch) is that of a cookie-cutter; for a given size sheet of cookie dough there are only so many "full-size" cookies that can be cut out of the dough.

In order to circumvent this difficulty we have used angular sampling, along with spatial sampling, to obtain a large number of independent speckle patterns. This is achieved by using a rotating right circular cylinder as a target. The laser illuminates the cylindrical side of a 2 m diameter x 2 m height drum, which rotates at 2 rpm. We sample a small angular section of the total speckle pattern; the angle that the telescope aperture subtends at the target. This allows for a much larger number of independent speckle patterns than could be generated by spatial sampling only.

In sampling different areas of a target, differences in albedo may introduce correlation in the data. This will result in some return pulses being much stronger than the "average" return. If there is a well defined frequency to the albedo variations (as with our rotating drum), then we may be able to "filter" the anomalous data in post-processing. In the "standard" DIAL configuration, two-lasers toggle between two different wavelengths, it is hoped that albedo fluctuations are removed when the ratio between adjacent pulses is taken. In order for the effect of albedo sampling to be completely removed, the returns from the two lasers must have the same temporal variations. This is not always the case.

There are many other possible sources of systematic errors which may also introduce temporal autocorrelation into lidar returns. Slight drifts in detector power supplies, drifts in the overall optical alignment, and electromagnetic interference are a few examples. A baseline drift in the data, which is not completely removed when taking pulse-pair ratios, results in a very distinct lower limit on the standard deviation.

We will discuss data that display some of these effects. Our main emphasis will be on the averaging statistics and the relationship between the temporal autocorrelation and the standard deviation of the mean.

2. EXPERIMENTAL CONFIGURATION

Our data were obtained from a lidar platform consisting of two CO_2 lasers, (Laser Science Inc. model PRF-150 WAL). These two lasers can operate, in a wavelength-agile mode, from 0 $\rightarrow$ 150 Hz p.r.f. (10 $\rightarrow$ 20 Hz for optimum stability). The laser pulse has a narrow ($\approx$ 200 ns) peak containing about 30% of the pulse energy with a long trailing edge ($\approx 2\mu s$). Typical output pulse energies range from 30 mJ/pulse to 60 mJ/pulse. The two laser beams are combined before they enter a 10X (Space Optics Research Labs, Model COAR 65X02-VF) beam expander. The combined beam is folded onto the optic axis of the receiver telescope (OCA Applied Optics, 406 mm, 1.37 m focal length, f/3.4 Dall-Kirkham configuration). The light-collecting area of the telescope, accounting for obscuration by the secondary mirror, is 1056 cm^2. The field of view of the telescope is approximately 1° off axis. The beam is then steered outside the platform trailer by a 24" diameter, R = ∞, copper mirror. This mirror also serves as the collection mirror for the return lidar light. A relay optic is used between the telescope and the detector in order to match the focal ratio of the detector and the telescope. This prevents the detector from seeing thermal radiation from any part of the optical system. Different detectors are used (0.5 mm $\rightarrow$ 2.0 mm in diameter), depending on the exact DIAL phenomena that is being measured; in general these are photovoltaic Mercury Manganese Telluride (MMT) or Mercury Cadmium Telluride (MCT) detectors , LN_2 cooled to 77K.

The outgoing energy, termed the "reference signal", is monitored for each laser pulse so that the return signal from the target can be normalized. Pyroelectric detectors with an internal battery-powered preamplifier are used to measure the outgoing energy. Optical triggers are provided for each laser by a room-temperature MCT

detector.

The return and reference signals are routed to an oscilloscope, transient digitizer, and boxcar integrator. The entire system (laser controls, scanning mirror position, wavelength, and range gate) is computer controlled. The most important data parameters are displayed in real-time so the operator can observe chemical absorptions as they take place.

3. ACQUISITION OF SPECKLE PATTERNS

We have made measurements of lidar return statistics from short range targets in order to minimize atmospheric effects. In general we have used fixed targets and a rotating drum target (described above).

3.1. Fixed targets, stationary laser beam.

If both the laser beam and target remain *perfectly* motionless during acquisition of lidar returns, then the standard deviation for the ensemble of returns would be zero. This is because the speckle pattern for each return pulse would be identical. In this case our measurement would be very *precise* but it may be inaccurate. The *accuracy* of our intensity measurement depends on how many speckles we have in the speckle pattern associated with a return pulse. An extreme example of an inaccurate reflectivity measurement is the case where the detector is looking at a "dark" region of the speckle interference pattern. In this case the measured reflectivity would be zero!

Of course, the case of a perfectly motionless system is not very realistic. In reality there will be some slight motion of the laser beam, receiver, and target. These slight motions are not large enough to result in an unique speckle pattern for each return pulse. Instead, this type of motion results in a *slightly* different speckle pattern for each return pulse. If we average over many of these returns then we will approach the mean intensity that we would have measured with one speckle pattern in the perfectly motionless situation; however, the accuracy of this intensity measurement (no matter how many returns are averaged over) depends on the number of speckles in the pattern. This would not be true if each speckle pattern was completely independent of the others.

A sample set lidar return data taken at 100 meters from a stationary brick wall are shown in Fig.1. These data are an example of the situation described in the previous paragraph. Two lasers were used to obtain these data (referred to as "laser 0" and "laser 1"); each laser operated at 10 Hz p.r.f. with the same wavelength (10P20). The separation between the laser pulses from the two lasers was $50\mu s$. Long-term drifts that appear in this data set limit the effect that averaging over pulses has in reducing the standard deviation. The source of these drifts is unknown at this time. It has been our experience that in obtaining large data sets, with the intent on averaging to reduce noise, slow baseline drifts or "jumps" can set a limit on noise reduction obtained by averaging techniques.

When a smaller sample of this data is averaged over, we find two important facts. The first is that the single-shot standard deviation, σ_1, is much lower than that for the entire data set. Since the smaller sample data is taken over a short time, the difference between the speckle patterns that make up the small sample is minimal. As we discussed earlier, if the speckle pattern at our detector were perfectly motionless, then $\sigma_1 = 0$. The second fact is that the averaging statistics, for the very small sample, approach random Gaussian statistics, $\sigma_N = \sigma_1/\sqrt{N}$. During the time duration of the small sample the long-term drifts do not add much autocorrelation to the data set.

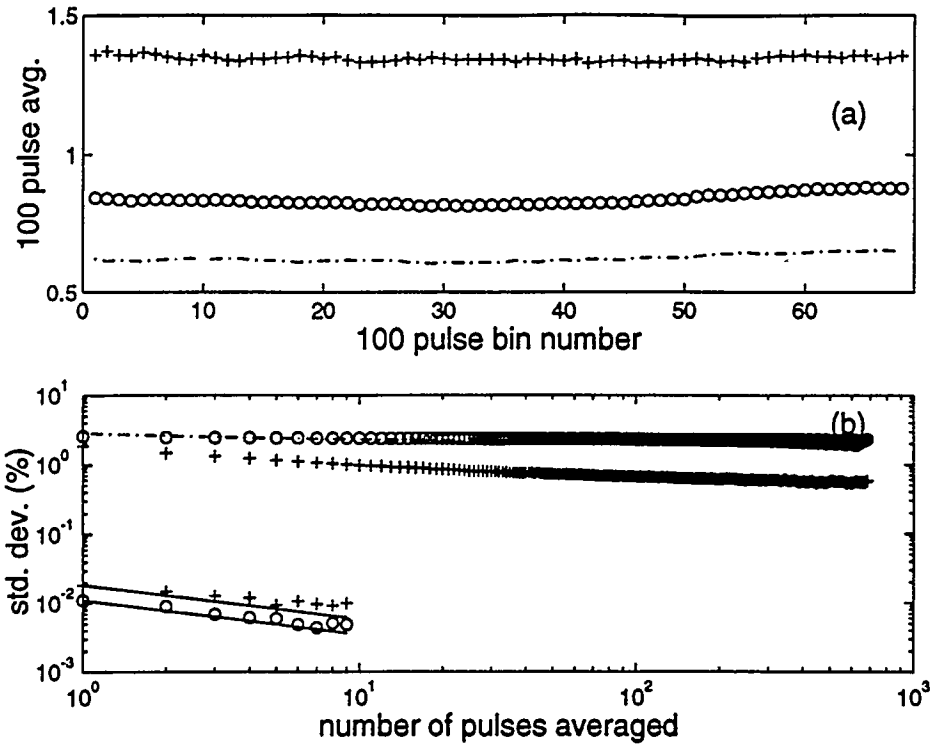

Figure 1: Return signal statistics from a brick wall at short range (100 m). In (a) the 100 pulse averages are shown for "laser 0" by circles, "laser 1" by pluses, and the ratio of the laser 0 to laser 1 returns by the dash-dotted line. In (b) the averaging statistics are shown for the entire data set and for a very small sample of the data set. For the small sample we have also shown the $\sigma_1/\sqrt{N}$ curve (solid lines) predicted for random noise. The long-term drifts in the data, shown in (a), add correlation to the data which diminishes improvement by averaging. The small sample data is not affected as much by these long term drifts and thus averages better. Note that the long term drifts also increase the "single-shot" standard deviation, σ_1; which is much lower for the small sample data.

We will see later how important long-term drifts are in averaging statistics. In the example discussed here we do *not* have an ensemble of statistically independent speckle patterns. For this reason σ_1 for the ensemble of speckle patterns is relatively small and is greatly influenced by small drifts in the data. If each speckle pattern of the ensemble were unique, then σ_1 would be determined by the average number of speckles per pattern, Eq. (1). More importantly, as we averaged over this ensemble of independent speckle patterns our measurement would approach the true mean reflectance of the target.

3.2. Lidar returns from a rotating target.

We have acquired large data sets of independent speckle patterns from the rotating target described in Sec.2. The cylindrical side of the drum was illuminated by the laser while the drum rotates around the axis of the cylinder. The total number of independent speckles available using this drum depends on the range to the target, the size of the receiver aperture, and the number of speckle per pattern. For measurements taken at a 100 meter range, the 40 cm diameter receiver subtends an angle of 4 mrad at the drum. In this case the number of independent speckle patterns is $2\pi/0.004$, or 1571 independent speckle patterns. Each pattern contains 878 speckle (for 1 mrad divergence). The total number of independent speckle (linear polarization) is 1.37 x 10^6.

This suggests that the lower limit on the standard deviation is approximately 8 x10^{-4}, if no other noise source is present. Of course this number can be lowered by moving the laser beam on the rotating target.

The rotating drum was coated with an aluminum roofing compound (used for trailer roofs) in hopes of obtaining a Lambertian reflectance for the drum. There are 72 equal sized facets on the drum. The facets and variations in the roofing compound result in a non-Lambertian surface. The result is that a periodic correlation exists in the data, the period is given by the drum rotation period. The effect can be seen in the 100 pulse averages shown in Fig.2. The drum period is also seen in the standard deviation of the mean curve. When the size of the segments in the segmented average equal the period of the drum rotation the periodic components will average out. The albedo variation of the drum is larger than the variation in intensity between independent speckle patterns – supported by the fact that we can see the albedo variations in the data!

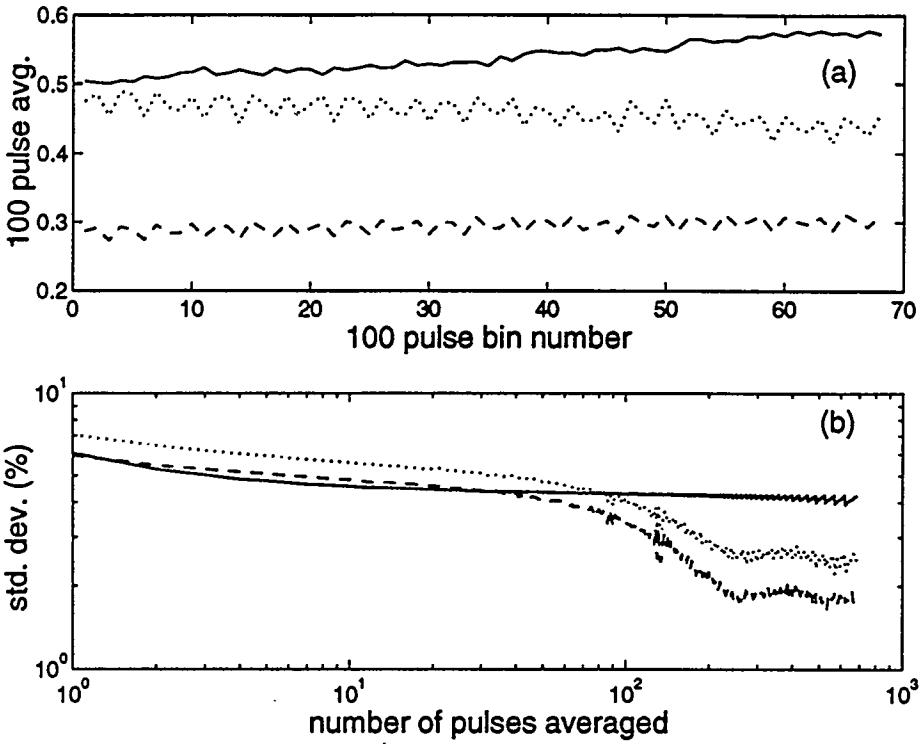

Figure 2: Lidar returns from a rotating drum target at 115 meters. The 100 shot averages, (a), show that a periodic correlation exists in the data. This is due to variation in the drum albedo. The return from laser 0 (dotted line), laser 1 (dashed line), and the ratio, laser 0 : laser 1, (solid line) are shown. Note that taking the ratio between adjacent pulses, seperated by 6 μs, removes most of the periodic correlation, but un-correlated baseline drifts for each laser results in poorer overall averaging statistics for the ratio data, shown in (b).

The data shown in Fig.2 were obtained using the "standard" DIAL configuration: two lasers toggling between wavelengths. The ratio of the adjacent lidar return pulses should cancel out correlation which is common to both lasers. If both lasers "track" each other *perfectly*, the ratio of the two returns will be a flat baseline with a random noise component. The random noise component exists only if the speckle patterns between adjacent pulses are uncorrelated. The adjacent pulses will have uncorrelated speckle patterns if their time separation is long compared with time it takes the drum to sweep through the angular extent of the receiver aperture. The

adjacent speckle patterns will also be uncorrelated if the phase fronts for each laser beam are different, and if there is a large wavelength separation between the lasers. If a chemical absorber exists, which only absorbs at one of the two wavelengths, the ratio between adjacent laser pulses will show the chemical absorption.[5]

Taking the ratio of adjacent laser pulses only helps when there is a high degree of cross-correlation between the two lasers. If the laser returns for each laser are *highly* correlated then autocorrelation in the return data set, for each vector, should be removed by taking the ratio. The averaging statistics for this data will then follow random Gaussian statistics.

The correlation, associated with the data shown in Fig.2, is shown in Fig.3. In Fig.3 we see the periodicity of the drum in the temporal autocorrelation of the data for both lasers and also in the temporal cross-correlation[6] between the lasers. The temporal cross-correlation is defined as

$$\rho_{jxy} = \frac{1}{\sigma_x \sigma_y (\Gamma - j)} \sum_{k=1}^{\Gamma - j} \frac{(P_{kx} - \langle P_x \rangle)}{\langle P_x \rangle} \frac{(P_{(k+j)y} - \langle P_y \rangle)}{\langle P_y \rangle}, \tag{8}$$

where the subscripts x and y identify each of the two lasers, and j is the lag (or lead) index.

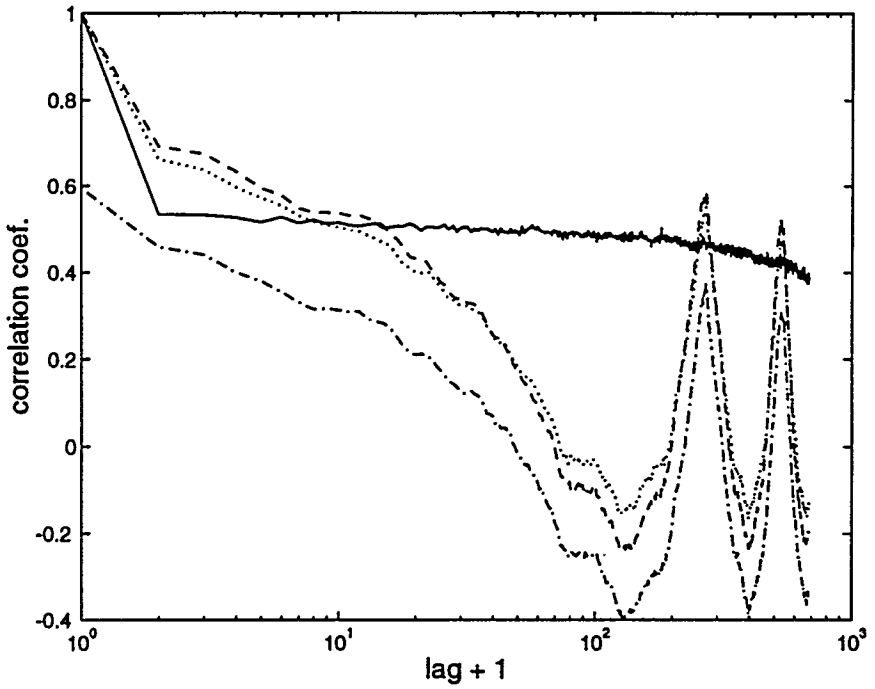

Figure 3: The correlation associated with the data of Fig. 2. The temporal autocorrelation is shown for laser 0 (dashed curve), laser 1 (dotted curve), the ratio of laser 0 to laser 1 (solid curve). The temporal cross-correlation between the two lasers (dash - dotted curve) is also shown. Note how taking the ratio of adjacent pulses removes the periodic autocorrelation at the drum rotation frequency, but adds a "drift" correlation. The un-correlated drift between the two lasers results in the sustained temporal autocorrelation for the ratio. In this case taking pulse-pair ratios results in poorer overall averaging statistics.

Taking the ratio of adjacent lidar returns from the drum does a good job of removing the periodic autocorrelation for the data set associated with each laser. However, each laser has a baseline drift which is "anti-correlated"

with the drift from the other laser. In this case taking the ratio of adjacent pulses actually hurts the overall statistics. If each laser had the *same* drift, then it would be removed by the ratio. For this situation it would be best to remove the correlations using other methods, such as digital filtering or baseline renormalization.

4. REMOVAL OF CORRELATION IN DATA PROCESSING

We have just shown how temporal autocorrelation in a data set reduces the benefit that averaging has on noise reduction. It is therefore beneficial to minimize or remove temporal autocorrelation in the data. This is the main reason for using a two-laser DIAL system and taking the ratio of adjacent pulses.

If the two lasers behave *exactly* during the data acquisition interval, then the ratio of the adjacent returns will have minimal autocorrelation. In reality two lasers never give *exactly* the same returns – they are not perfectly correlated. Slight differences in the return signals from the two lasers results in temporal autocorrelation remaining in the ratio data set. Sometimes data processing techniques may be used to remove autocorrelation from the individual laser returns before taking pulse-pair ratios and then averaging.

4.1. Correction of baseline drifts.

Perhaps the easiest autocorrelation component to remove from a data set is a linear baseline drift. In this case we simply fit a straight line to the data and then ratio the data to this straight line. When measuring atmospheric or chemical absorption we must be very careful when "re-normalizing" data in this manner. Suppose that a chemical absorber was present during our data collection and the concentration of this chemical changed in a linear fashion during this time; "re-normalizing" would then remove all information about the absorption.

Averaging statistics for "baseline re-normalization" of the data in Fig.2 is shown in Fig.4. A linear curve fit was made to the data vector from each laser and then each data vector was divided by the linear fit. This effectively removes the baseline drift and allows for more correlation to be removed by taking the ratio of adjacent pulses. The autocorrelation for the data vectors, and the ratio, are shown in Fig.5.

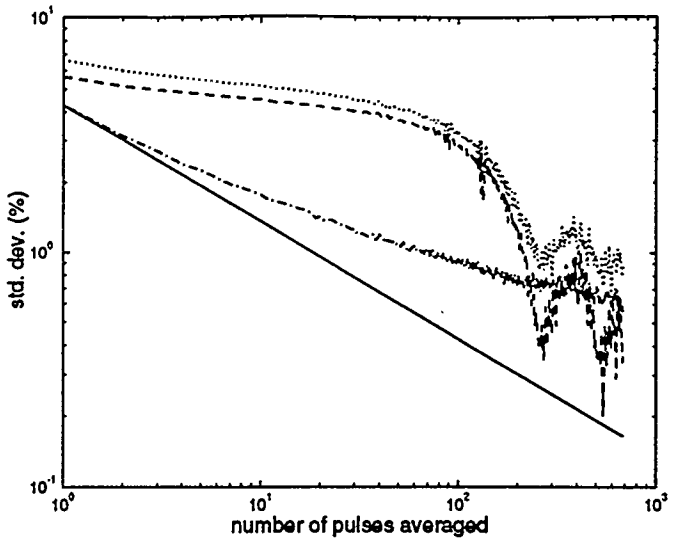

Figure 4: The averaging statistics of the data shown in Fig. 2 after each laser has been "re-normalized" to remove the baseline drifts associated with each laser. The averaging statistics for laser 0 (dashed line), laser 1 (dotted line), and their ratio (dash-dotted line) are shown. Taking pulse-pair ratios (after baseline renormalization) significantly improves the averaging statistics. The ratio data follows closer to the averaging statistics for random Gaussian noise (solid line).

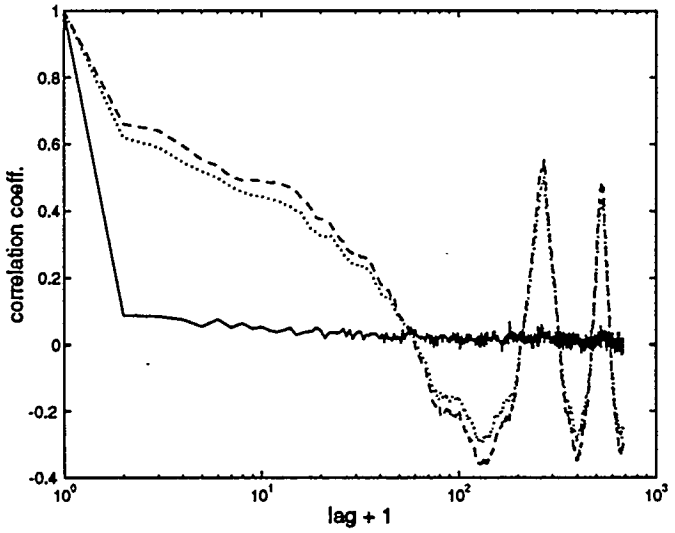

Figure 5: The temporal autocorrelation for the "baseline re-normalized" data shown in Fig. 4. The autocorrelation of laser 0 (dashed line), laser 1 (dotted line), and their ratio (solid line) are shown. Note that the ratio data now has a much smaller autocorrelation than in Fig. 3.

4.2. Correction for periodic correlation.

We have shown how taking the ratio of adjacent return pulses from the rotating drum removes most of the periodic correlation in the data (due to the albedo variation over the surface of the drum). If the ratio removed all of the correlation the averaging statistics would follow $1/\sqrt{N}$ exactly. We can see in Fig. 4 that the averaging has been improved by taking the ratio of pulses between the lasers, but not to the "random" noise limit. It is more difficult to remove the periodic autocorrelation in the rotating drum data if pulse pair ratios are not used.

An obvious method of removing the periodic noise components is to digitally filter, in the frequency domain, the rotation harmonics of the drum. The power-spectral density of the drum data is shown in Fig. 6. The rotation frequency ($\approx$ 0.033 Hz), and higher harmonics, are obvious in the power-spectral density function. The power-spectral density of the ratio data is also shown in this figure. Note that the ratio data has some "residue" signal at the frequency of the drum rotation.

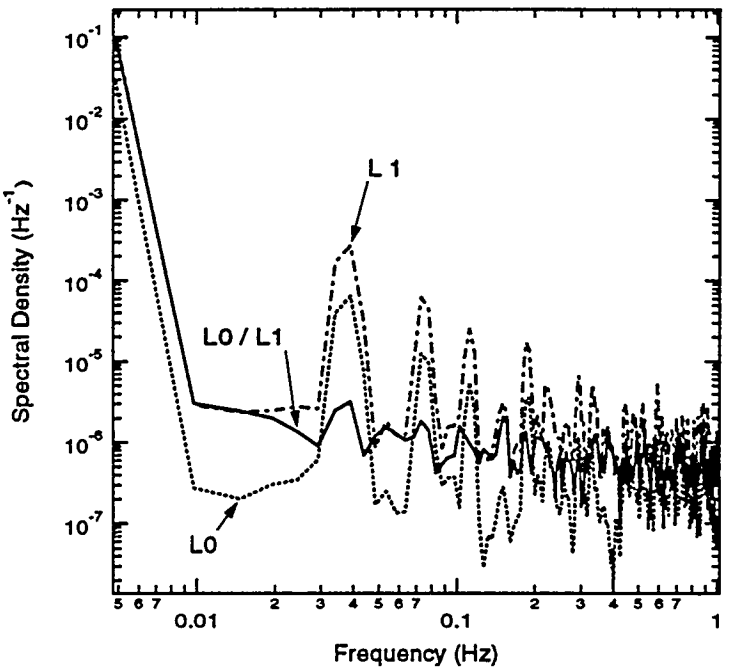

Figure 6: The power-spectral density for the raw lidar returns. The returns from laser 0 (L0), and laser 1 (L1) show the drum rotation period. The data was obtained at 10 Hz, the drum rotation frequency is approximately 0.033 Hz. The power-spectral density of the adjacent pulse pair ratios, between the lasers, is shown by the solid line. Note how taking pulse-pair ratio removes the periodic correlation. The fundamental and second harmonic were digitally filtered from this data set; the averaging statistics of the filtered data are shown in Fig. 7.

A digital filter was used on the laser 0 data; the fundamental and second harmonic components were spectrally filtered using a simple finite-impulse filter. A segmented average was then performed on the filtered data to see how much correlation was removed. The results are shown in Fig. 7. Also shown in this figure are the averaging statistics for the raw data, the adjacent pulse-pair ratio data, the ratio data after correcting for baseline drift, and the digitally filtered data corrected for the baseline drift. Note that the digital filtered data is single laser data (laser 0); no ratio was taken between adjacent pulses. The results indicate that the "standard" DIAL technique

of taking pulse-pair ratios works quite well (assuming baseline drifts are accounted for) as compared with this simple digital filtering method.

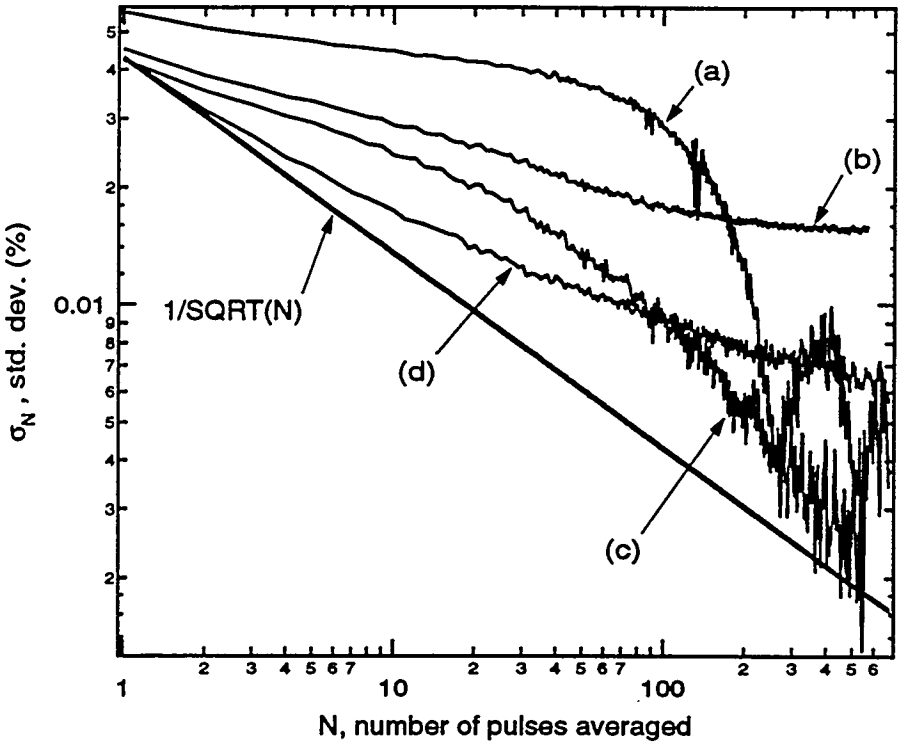

Figure 7: Averaging statistics for the raw data from the drum (laser 0): (a), the digitally filtered drum data from laser 0 (b), the digitally filtered and then drift corrected data (c), and the adjacent pulse pair ratio and then drift corrected data (d). Removing higher harmonic components (only the fundamental and second harmonic were filtered) in the spectrum should result in better averaging for curve (c).

Digital filtering may remove more of the periodic correlation if higher-order harmonics are also filtered. It is interesting to see the change in the averaging statistics when the first two harmonics are filtered as compared to the raw periodic data. One must remember that "random" data would result from filtering *all* significant frequency components from the data, but information on an absorption signal may also be removed. Care must be taken to "surgically" remove only those frequency components which are expected to appear in the power spectrum from *known* noise sources. The drum rotation frequency is an example of a known periodic noise source, 60 Hz AC line noise is another example of a known source of periodic correlation.

4.3. Running averages can introduce autocorrelation!

A "standard" practice in averaging over signals is to use a "moving average" to smooth out high frequency noise. When a moving average is used, individual data points are "counted" many times during the averaging process. This is not true of a segmented average. For example, if a ten-pulse moving average is used, the average of pulses 1-10, 2-11, 3-12, etc. are obtained. It is obvious that these average values, obtained with a moving average , will be correlated with a correlation length roughly equal to the size of the averaging window. Smaller

sized averaging windows will introduce more correlation to the smoothed data set than a larger window. In fact, the correlation *introduced* by moving averages is used a tool for measuring relative trends in financial markets.[7] In order to obtain true "random" statistics, each lidar return should only be used once in averaging processes.

Smoothing of data, usually done by a moving average, is a useful tool for *visualizing* general trends in a noisy data sample. Accurate concentration measurements should *not* be based on results obtained from smoothed data. In particular, one must be careful of any further data processing used on a data set which has been smoothed. An example of how things can go wrong with smoothed data follows.

We have tried to remove the periodic correlation from the rotating drum returns by smoothing the data using a moving average and then "renormalizing" the raw data to the smoothed data. The hope is that this will remove the albedo variations in the data and leave us with a random set of data, which can then be averaged to give better statistics. The standard deviation of the mean, σ_N, of data which has been renormalized using this method is shown in Fig. 8. The averaging statistics show a spurious averaging improvement – greater than that for random data! The error in using this method is based upon the fact that individual data points are included more than once in the averaging process. The temporal autocorrelation for this smoothed data is shown in Fig.9.

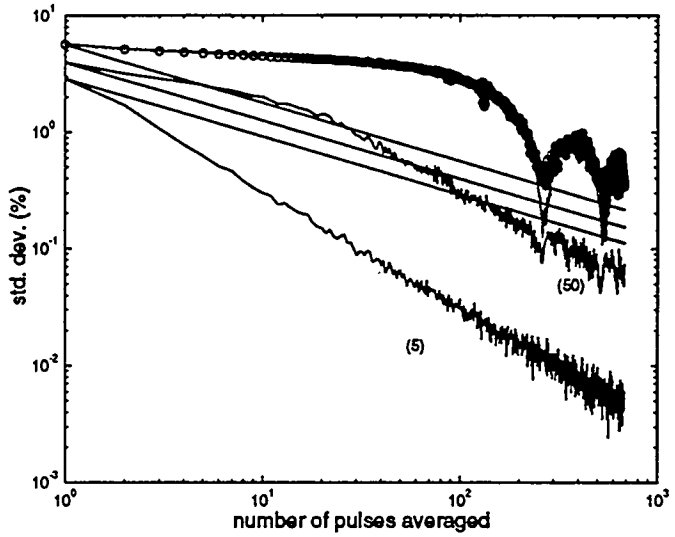

Figure 8: The averaging statistics for single-laser lidar returns from the rotating target after the raw data has been normalized to "smoothed" data (moving average). The averaging statistics for the raw data are shown by circles. The straight lines represent $\sigma_1/\sqrt{N}$ for random noise. The raw data was normalized to data smoothed by a 5 pulse, 50 pulse, and 275 pulse moving average. The 275 pulse moving average follows the statistics of the raw data closely. The drum rotation period is equivalent to 275 laser pulses. Note the spurious averaging obtained using this method – the apparent averaging is better than random data. The fault in this technique is that the same data is averaged over more than once.

In order to quantify the amount of correlation that is introduced to a data set by smoothing we have generated a 6000 point random set of "fake" data (using a random noise generator). We have then smoothed this random data, using various smoothing algorithms, and then calculated the temporal autocorrelation of the smoothed data. The results are shown in Fig. 10. The Savitzky-Golay[8] algorithm is a type of least-squares polynomial smoothing. A least-squares fit, of a fixed order polynomial, is made to the data inside the smoothing window; the box is then shifted one point and the process is repeated throughout the length of the data vector. The binomial smooth[9] is basically a Gaussian filter, a box with a Gaussian weighting factor. The other smoothing routines used were simple box averages.

Smoothing of data should really only be used for *visualization* of trends in data and *not* used for any concentration measurements.

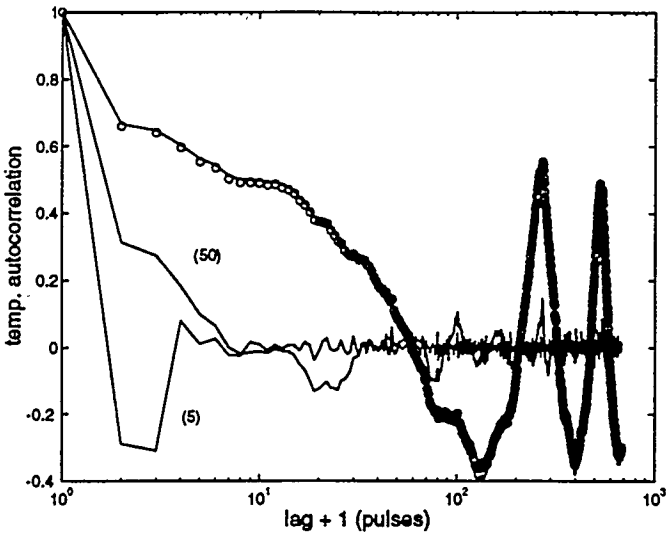

Figure 9: The temporal autocorrelation for the raw data normalized to "smoothed" data, whose averaging statistics are shown in Fig. 8. When the size of the moving average window, used to smooth the data, is small compared to the drum rotation period (275 pulses), the smoothed data will be highly correlated with the raw data. When the raw data is normalized to the smoothed data, "too much" correlation is removed – resulting in spurious averaging statistics. The temporal autocorrelation is shown for data averaged by a 5 pulse, 50 pulse, and 275 pulse moving average. The autocorrelation for the raw data is shown by circles, the autocorrelation for the 275 pulse average follows that for the raw data.

4.4. Overall goal of correlation removal

One must think very carefully before attempting to remove correlation in the data to improve averaging. Consider for example the case of a baseline drift type of correlation. This type of correlation can be removed very easily by renormalizing the data to give a "flat" baseline. There are certain instances when the "real" lidar absorption signal will have a drift component, for instance a changing concentration for an effluent. If we renormalize the drift out of the data do we really know the mean of the return intensity any better? In fact, averaging over the now randomized data, obtained by baseline drift corrections, may give us a higher precision (smaller σ_N), but we may be averaging to an inaccurate absorption value.

In general, other "external" knowledge must be used to correctly decide on which correlations are not related to real lidar absorptions and which should be removed from the data set before any averaging. An example of using prior knowledge is the method we have used above for the rotating drum; we knew in *advance* that a periodic component would be present in the data.

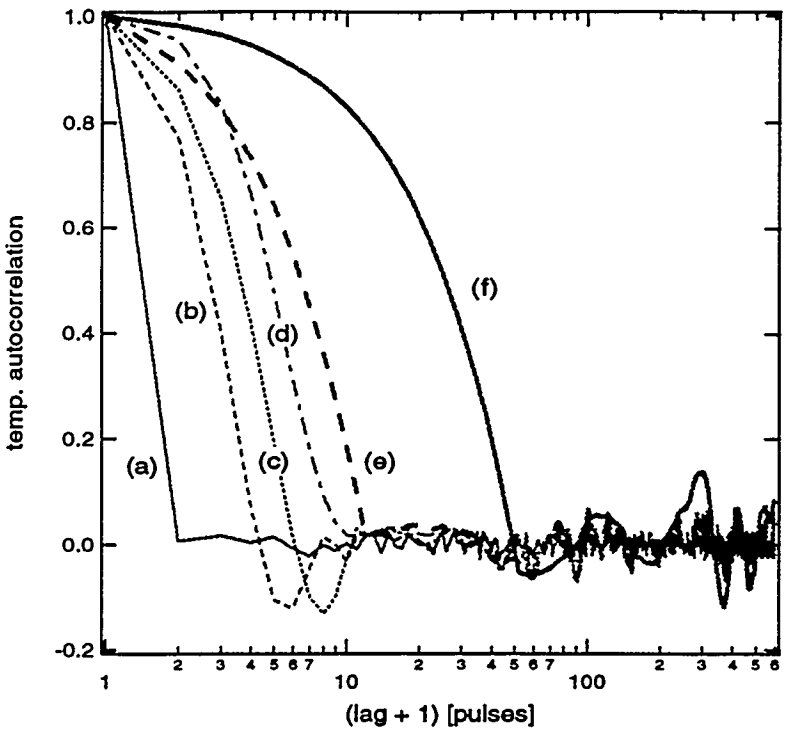

Figure 10: The temporal autocorrelation of random noise which has been smoothed using a variety of smoothing routines: (a) raw data, no smoothing, (b) 11 point – 4^{th} order Savitzky-Golay, (c) 11 point – 2^{nd} order Savitzky-Golay, (d) 11 point – binomial smooth, (e) simple 11 point moving box, and (f) simple 51 – point box smooth. All of the smoothing routines introduce some correlation to the original un-correlated data.

5. POLARIZATION EFFECTS

The speckle pattern generated from circular polarized light can be decomposed into two speckle patterns; one pattern for each orthogonal linear polarization component of the total field. The number of speckle generated by circular polarization is twice as large as the number generated from a single linear polarization component. This suggests that the standard deviation for circular polarization will be smaller, by a factor of $1/\sqrt{2}$, than for linear polarization. In practice, however, linear polarized light, scattered from a rough surface, may become depolarized upon reflection. The depolarization results in larger *total* number of speckle than will be collected if the reflected light remained linear polarized. The overall statistics of the received intensity is dependent on the incident polarization and the amount of depolarization from reflection.

The total light field of a speckle pattern can be written in terms of two orthogonal linearly polarized field components,

$$\vec{A}(x,y) = A_x(x,y)\hat{x} + A_y(x,y)\hat{y}, \tag{9}$$

where the total intensity, $I(x,y)$ is given by

$$I(x,y) = \vec{A}\cdot\vec{A}^* = |A_x(x,y)|^2 + |A_y(x,y)|^2 \tag{10}$$

$$= I_x(x,y) + I_y(x,y). \tag{11}$$

Correlation existing between the two orthogonal field components is given by the off-diagonal elements of the coherency matrix,

$$[J] = \begin{bmatrix} \langle |A_x|^2 \rangle & \langle A_x A_y^* \rangle \\ \langle A_x^* A_y \rangle & \langle |A_y|^2 \rangle \end{bmatrix}. \tag{12}$$

The coherency matrix can be diagonalized, by a coordinate rotation and a relative retardation of the two components,[3] to give a new coherency matrix

$$[J\prime] = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} = \begin{bmatrix} \lambda_2 & 0 \\ 0 & \lambda_2 \end{bmatrix} + \begin{bmatrix} \lambda_1 - \lambda_2 & 0 \\ 0 & 0 \end{bmatrix}, \tag{13}$$

where λ_1 and λ_2 are the eigenvalues of the coherency matrix $[J]$. The first of these matrices represents totally unpolarized light with a mean intensity equal to $2\lambda_2$, and the second matrix represents linear polarized light with a mean intensity of $\lambda_1 - \lambda_2$.[10] The degree of polarization, P, is the ratio of the average polarized intensity to the total average intensity,

$$P = \left| \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \right|. \tag{14}$$

Diagonalizing the coherency matrix allows us to regard the total speckle pattern as being the sum of two uncorrelated speckle patterns, one with an "effective" average intensity λ_1, and one with an "effective" average intensity λ_2. The "effective" average intensities of the two speckle patterns depends on the degree of polarization through,[3]

$$\begin{aligned} \lambda_1 &= \frac{1}{2}\langle I \rangle (1+P) \\ \lambda_2 &= \frac{1}{2}\langle I \rangle (1-P). \end{aligned} \tag{15}$$

The statistical properties of the total speckle pattern depend on the degree of polarization. This can be seen from the probability density function, $p(I)$:

$$p(I) = \frac{1}{2}(1+P)\exp\left[\frac{-2I}{\langle I \rangle (1+P)}\right] - \frac{1}{2}(1-P)\exp\left[\frac{-2I}{\langle I \rangle (1-P)}\right] \tag{16}$$

We see that if the degree of polarization is zero, then pure negative exponential statistics follow. In this case the standard deviation of the intensity, for one speckle pattern, is reduced by a factor of $\sqrt{2}$ as compared with a speckle pattern due to one polarization component alone. In other words, if we use circular polarized light, and the polarization state is *not* scrambled on reflection, then the resultant speckle pattern will be the sum of two uncorrelated speckle patterns, whose mean intensities are equal. We then effectively coolect twice as many *independent* speckle and therefore standard deviation of the total intensity improves by a factor of $1/\sqrt{2}$. If the polarization is scrambled upon reflection, then we are still better off than we would be using pure linear polarization.

The amount of depolarization in the return laser light depends on how much multiple scattering occurs, and the target reflectivity. A target which introduces a large amount of multiple scattering must have a high enough reflectivity to return the fraction of the light which has gone through the multiple reflections needed for depolarization. The rotating drum used in these experiments seems to fit this criteria.

We have made measurements of return lidar statistics from the drum using a circular polarized laser beam and a linear polarized beam. The experiment was configured to give a large speckle noise component by masking off the aperture of the receiver telescope to the size of one correlation area. In this configuration speckle noise dominates the correlated albedo noise associated with albedo variations on the surface of the drum. The resulting averaging statistics are shown in Fig. 11. The data for linear polarization is only slightly noisier than the data for circular polarized light. This is because the drum target does a good job of depolarizing the linear polarized beam and the number of speckle for linear and circular polarization is approximately the same.

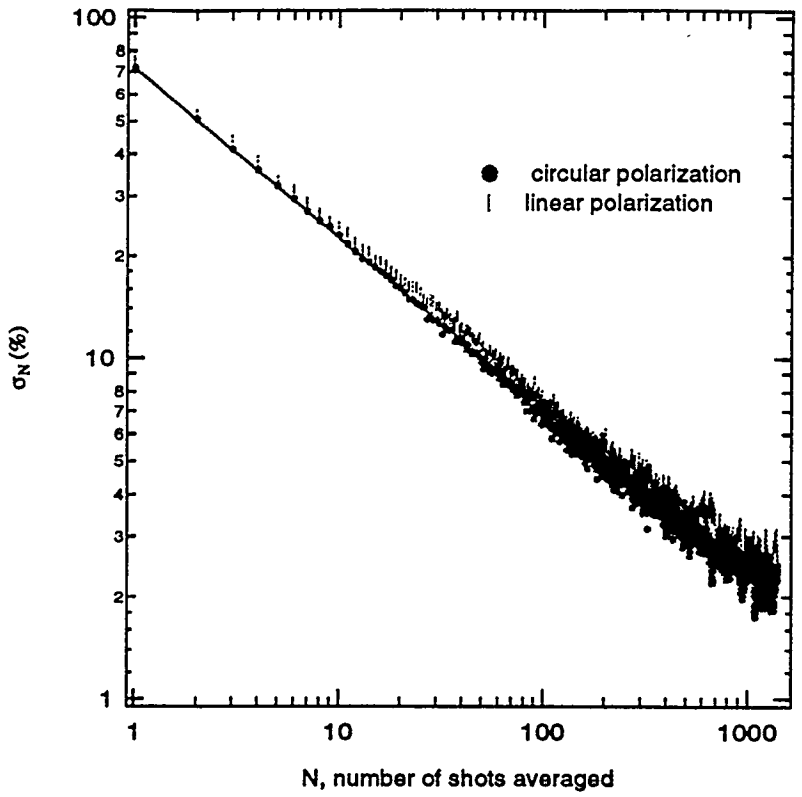

Figure 11: The averaging statistics from the rotating drum at 115 meters range using a linear polarized laser beam and a circular polarized beam. The noise level is high since the receiver aperture was completely masked off except for a 6.3 mm square hole, the speckle size is approximately 13.5 mm. If the target did *not* depolarize the linear polarized beam upon reflection, then the difference in averaging statistics, between the two polarization states, would be larger.

We have made measurements on the amount of depolarization from various lidar targets. The overall results show much less depolarization than from our rotating drum target. Most natural targets have such a low reflectance that the intensity from multiple reflections is very weak compared to the primary reflected intensity. The degree of polarization measured for a sample of different target types is shown in Table (1). The degree of polarization can be used in Eq. (16) to obtain the probability density function – which can then be used to find statistical moments such as the variance etc.

Of the various targets that we have measured depolarization from, the radio telescope reflection retained its polarization the most. The surface of the radio telescope target was cylindrical and therefore the reflection was highly specular and little to no multiple scattering occurred. The flame-sprayed target aluminum target showed the most depolarization due to the high reflectivity and rough surface; the return intensity from multiple reflections is large – resulting in large depolarization. The other natural targets did not depolarize the lidar returns significantly–probably due to their overall low reflectivity.

Based upon these depolarization measurements one can see that it is advantageous to use circular polarized light to increase the total number of received speckle, thereby increasing the signal-to-noise ratio. If the target depolarizes a linear polarized beam then the signal-to-noise ratio will also improve. Using circular polarization guarantees better noise statistics overall. We hope to measure the improvement on noise statistics from other

natural targets to gain a better understanding of how polarization relates to speckle noise.

Target	P, Deg. of Polarization
metal structure (radio telescope)	0.94
plants (juniper trees)	0.85
rocks (welded tuff)	0.72
flame-sprayed aluminum target	0.53

Table 1: Measured degree of polarization, Eq. (14), from various lidar targets at 1 – 1.5 km range using the 10P20 CO_2 laser line. In general, the largest depolarization results from multiple reflections from the higher reflectance targets.

6. SUB-SPECKLE SIZE RECEIVER DATA

We stated earlier that the noise statistics will improve as more speckle are received by the lidar telescope. In order to quantify this we have measured the standard deviation in the intensity for an ensemble of speckle patterns as a function of the receiver diameter. This was accomplished by masking off our receiver telescope, with a "light-tight" mask, and measuring the return the signal statistics from the rotating drum. Lidar returns were collected for various mask sizes. The outgoing laser beam divergence was held constant (1.0 mrad) for all of these measurements so that the speckle size (13.5 mm, from Eq.(3)) would remain constant. The smallest mask sizes were smaller than the predicted speckle size. The results are shown in Fig. (12). This data still contain the correlated noise associated with the albedo variation across the surface of the rotating drum. As the size of the mask aperture increases we can see the speckle noise decrease until it reaches the noise level associated with the albedo variation. For the smallest size apertures a differential amplifier was used to amplify the very weak signal (due to the limited optical throughput). Other noise was introduced by amplifying these very weak signals and the data point showing a percent standard deviation greater than 100 percent is most likely the result of amplifier noise.

Earlier we had stated that the noise due to speckle is inversely proportional to the square root of the number of speckle received in the aperture. This is only approximately true. A more accurate calculation must consider the distribution of the laser intensity on the target and the correlation between the individual fields which comprise the total laser intensity. A more accurate calculation has been made by Goodman.[3]

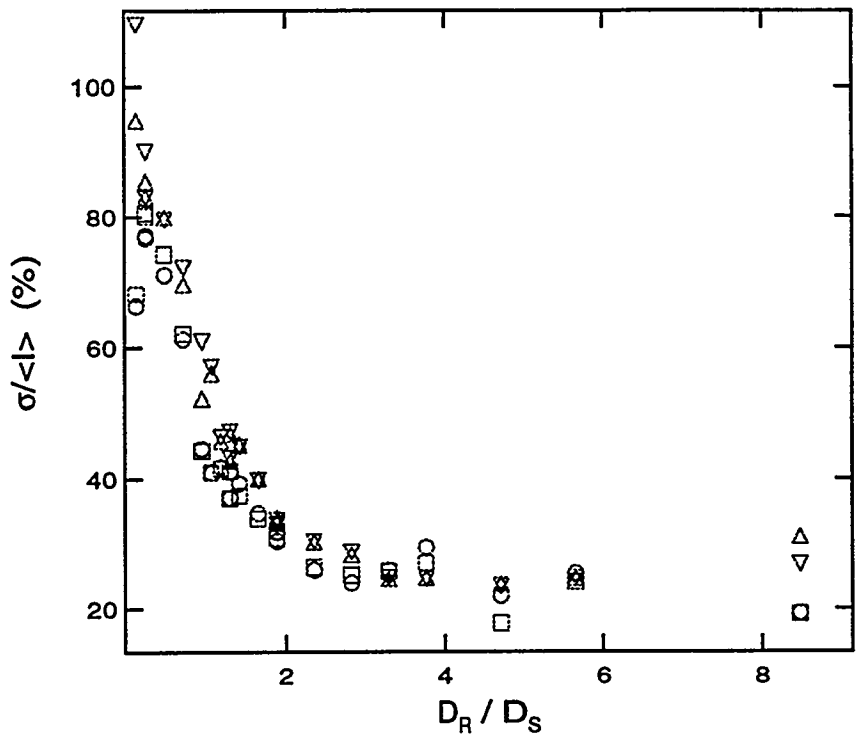

Figure 12: The standard deviation vs. the ratio of the receiver aperture, D_R, to the speckle diameter, D_S. The data was obtained from the rotating drum target. As the aperture size increases, the effects of the correlated noise associated with the variations in the drum become dominant over the speckle noise.

7. CONCLUSION

We have found that speckle is usually the dominant noise source in hard-target CO_2 DIAL; however, speckle noise is not the only noise source that we have measured. Other sources of noise, such as albedo variations, which are not entirely removed by pulse-pair ratios, may also dominate. Most importantly we have found that averaging over pulses can improve noise statistics only to the level where temporal autocorrelation in the data being averaged becomes significant. Systematic noise, which is not random, will not average away and will limit the improvements made by averaging.

Our results show that some types of correlated noise, such as drifts, may be removed in data analysis. Other correlated noise components can also be removed by filtering methods. The pulse-pair ratio method appears to work well if both lasers have a high cross-correlation, as expected. In practice you must either collect data over a long time interval at low p.r.f., or use a very high p.r.f. laser, to obtain large enough data sets to average over. We have made some preliminary measurements of lidar returns with 5 kHz p.r.f. CO_2 lasers to investigate this phenomena. For either configuration we find that stability in the laser intensity is crucial for obtaining good averaging statistics.

Our results on noise, versus the degree of polarization of the laser, indicate that circular polarization should always be used. Using circular polarization is a "win – win" proposition. If the target depolarizes any incident linear polarized components, then the noise statistics will approach that of pure circular polarization. In any case, the best signal-to-noise will obtained with circular polarized light. We are currently examining the effect

polarization has on "glint" type reflections. Our hypothesis is that circular polarization will be less sensitive to the direction of the normal of the reflecting surface.

In summary, we can say that independent speckle does average as $1/\sqrt{N}$ but other correlated noise sources associated with DIAL may prevent one from measuring this type of averaging directly. Care must be taken to prevent all sources of correlated noise in the lidar measurements. Some types of correlated noise, that remain in the lidar measurements, can be adjusted (or removed) by data analysis techniques. The best averaging results will be obtained for the least correlated data. Using two-laser DIAL, and taking pulse-pair ratios, appears to be the best experimental method of correlation removal in data. appear to be the best method

8. ACKNOWLEDGEMENTS

We would like to thank Dr. John F. Schultz, Daniel Beatty, and Dr. Stephen Czuchlewski for their valuable help and support. We also thank Beverely Dickinson for her valuable assistance. This work was supported by the United States Department of Energy through the Los Alamos National Laboratory, Chemical Science and Technology Division.

9. REFERENCES

[1] M.J. Scmitt personal communication January 3, 1995.

[2] R.M Measures (1992), *Laser Remote Sensing, Fundamentals and Applications*, Krieger, Florida.

[3] J.W. Goodman, in *Laser Speckle and Related Phenomena*, J.C. Dainty Ed., (Springer-Verlag, Berlin, 1984).

[4] N. Menyuk, D.K. Killinger, and C.R. Menyuk, Appl. Opt. **21**, 3377 (1982).

[5] M.J.T. Milton and P.T. Woods, Appl. Opt. **26**, 2598 (1987).

[6] N. Menyuk, D.K. Killinger, and C.R. Menyuk, Appl. Opt. **24**, 118 (1985).

[7] R.W. Colby and T.A. Meyers (1988), *The Encyclopedia of Technical Market Indicators*, Irwin, New York. See in particular the discussion of MACD, pg. 281.

[8] A. Savitzky and M.J.E. Golay, Anal. Chem. **36**, 1627 (1964).
J. Steiner, Y. Termonia, and J. Deltour, Anal. Chem. **44**, 1906 (1972).
H. Madden, Anal. Chem. **50**, 1383 (1978).
W.H. Press, *et al.(1988), Numerical Recipes in C*, 2nd. ed., Cambridge University Press.

[9] P. Marchand and L. Marmet, *Revues of Scientific Instrumentation*, **54(8)**, 1034 (1983).

[10] E. Wolf, Nuovo Cimento **13**, 1165 (1959).

DISCLAIMER