Wide Area Wind Field Monitoring

Status & Results

Alan Marchant & Jed Simmons

USU Energy Dynamics Laboratory

Report Number: DOE/EE0002734-1

DOE Award #: DE-EE0002734

Contents

Executive Summary.

Volume-scanning elastic has been investigated as a means to derive 3D dynamic wind fields for characterization and monitoring of wind energy sites. An eye-safe volume-scanning lidar system was adapted for volume imaging of aerosol concentrations out to a range of 300m. Reformatting of the lidar data as dynamic volume images was successfully demonstrated. A practical method for deriving 3D wind fields from dynamic volume imagery was identified and demonstrated. However, the natural phenomenology was found to provide insufficient aerosol features for reliable wind sensing. The results of this study may be applicable to wind field measurement using injected aerosol tracers.

I. Introduction.

The report presents the activities and results of a study funded by the Department of Energy under award DE-EE0002734 with additional analyses and recommendations for further

investigation. The general project goal is to advance commercialization of wind energy by evaluating a new laser-based method for mapping and monitoring the local area wind field.

In this context, the "area wind field" refers to the 3-dimensional air flow vector treated as a function of the 3 spatial dimensions over a scale comparable to an industrial wind turbine or wind farm, i.e. up to 300 m high and 300 - 3,000 km wide with a spatial resolution much smaller than 100m. Monitoring the local wind field also entails identification of temporal changes in the wind flow which can be expected to occur as fast as 1 Hz (i.e. downstream fluctuations due to turbine rotation).

Two classes of available instruments are capable of remote wind sensing: SODAR and Doppler Lidar. Both of these sensors types measure wind speed by detecting the Doppler shift from backscattered waves: SODAR projects sound waves that are back-scattered by local fluctuations in air density; Doppler lidar projects a narrow-bandwidth laser beam that is backscattered by dust (elastic lidar) or molecular fluctuations (molecular lidar). An individual Doppler reading detects only the line-of-sight (LOS) component of the wind velocity. Full 3D velocity vectors can be derived from multiple readings for which the source is pointed in at least 3 different directions. SODAR and Doppler Lidar systems usually obtain their multi-directional readings using a conical scan. An important essential drawback is that the 3D Doppler velocities from a single transceiver are averages over a relatively large area and this volume increases with range. At a range of 300m with a cone angle of 30°, the spatial resolution exceeds 150 m. For SODAR and pulsed Doppler Lidar, the line-of-sight spatial resolution can be shortened by employing short pulses. The sample distance is then derived from the wave velocity (sound or light) times half the echo delay. But the pulse length is limited by a tradeoff with wind speed resolution: the resolvable speed is not much smaller than the wave velocity divided by the number of cycles per pulse. For example, a Doppler Lidar with a pulse length of 200 ns is limited to a (longitudinal) spatial resolution of 60 m and a velocity resolution approaching 1 m/s. Extensive averaging is required to obtain finer velocity resolution.

An important goal of this study is to evaluate the utility of a new method for lidar wind profiling that is not subject to the spatial resolution limitations of Doppler sensing. We find that scanning elastic lidar can generate volume images of aerosol concentrations in the atmosphere with relatively high spatial resolution. Dynamic 3D imagery (repeated volume images) can be analyzed to find the wind velocity field that drives the motions of ambient aerosols.

It is hoped that the study of this new method for lidar wind measurement can lead to development of a tool for sensing wind patterns on the scale of wind turbines or wind energy sites for purposes of micrositing (optimal placement of turbines), operational look-ahead (advance warning of sudden wind changes), and turbine wake evaluation (to optimize turbine design and the siting of multiple rows of turbines).

II. Volume-scanning Lidar System.

The current study makes use of a volume-scanning lidar system developed at the Utah State University Research Foundation. This system is known as the Wind Lidar Demonstrator (WiLD).

WiLD Optics

A Multiwave Photonics MOPA-L-1550-nm-5kW fiber laser generates the output pulses for WiLD. It is capable of 1.19 W average output power and 150 kHz pulse repetition frequency (PRF), but is typically operated at a PRF of 50 kHz where the average output power is 1 W. As reported by Multiwave, the center wavelength is 1549.55 nm and the spectral width is 0.23 nm. The pulse width is approximately 3 ns.

The laser pulse exits the fiber, which has a collimating lens attached, then passes through a hole at the edge of the mount of the receiver objective lens before reflecting off two scan mirrors and into the target space. The beam diameter from the laser is 1.4 mm at the 1/e points and the corresponding divergence is 0.7 mrad.

A Zemax ray trace simulation was performed to design the collection optics. Aspheric lenses were selected because the 200 μm diameter APD active element required a small spot size at the focal plane. A 1 mrad field of view (FOV) was used in the simulation to account for the spot size of the laser beam in the target volume; four off-axis fields were included in the model. The objective lens is an AL100100 aspheric lens from Asphericon, antireflection coated for the near infrared, with an effective focal length of 100 mm, a numerical aperture (NA) of 0.478, a diameter of 100 mm, and a clear aperture of 92 mm. In order to maximize signal collection, we chose the largest aperture compatible with the scanner optics.

Because the narrow band-pass interference filter requires collimated input, a focal relay was introduced into the optical system. Both filters were placed between two AL1210 lenses from Asphericon. The AL1210 lenses have an effective focal length of 10 mm, a NA of 0.545 (greater than the objective lens), and a diameter of 12.5 mm. The filters were modeled as BK7 glass windows. After the last lens in the system is a 1 mm thick window which is positioned 0.5 mm in front of the APD. As shown in Figure 1, the optimal distance from the objective lens to the first lens in the focal relay is 86.8 mm. The distance from the last lens in the focal relay to the detector window is 6.7 mm. As illustrated by the spot diagram inset in Figure 3, the full field falls well within the 200 μm diameter of the detector.

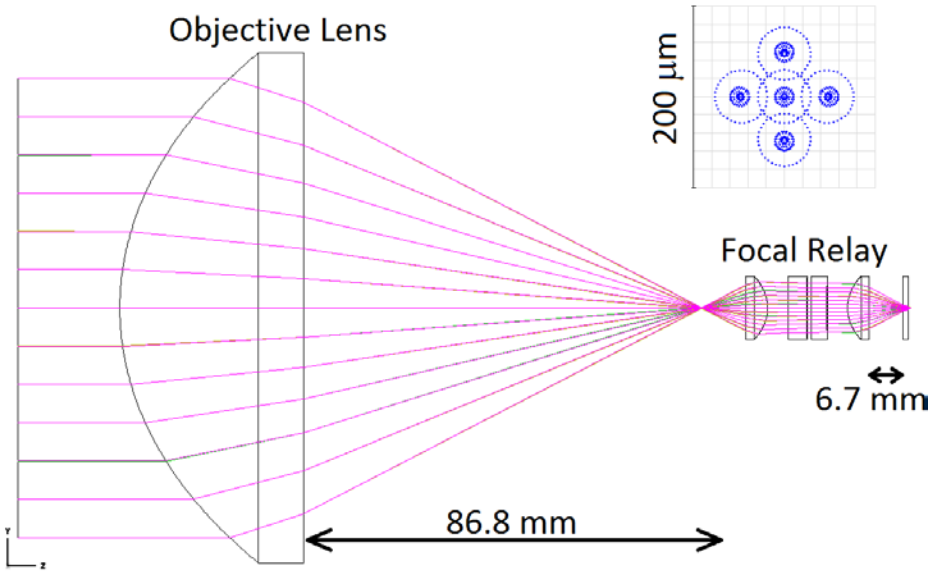

Figure 1. Ray trace illustrating the function of the collection optics in WiLD. The inset image is the spot diagram at the image plane.

An iris was placed in the system just prior to the focal relay at the prime focus of the objective lens. In the initial stages of system alignment, the iris can be opened to allow a signal to pass to the detector despite the misalignment. The focal relay is secured to the APD detector housing and alignment is achieved by translating this assembly in the x, y, and z directions. When the system is fully aligned, the iris can be reduced almost to its closed limit without a noticeable decrease in signal at the detector.

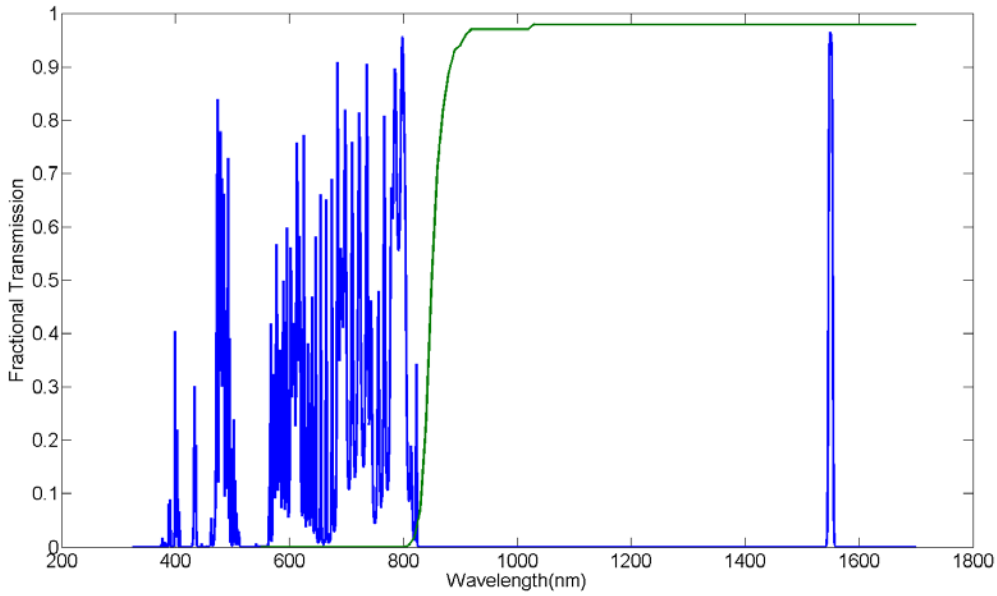

Figure 2. Fractional transmission of the narrow band-pass filter from Semrock (blue) and the long pass Schott filter from Edmund (green).

It is necessary to spectrally filter the return in order to reduce the interference of ambient light with the laser return signal. For this purpose a Semrock NIR01-1550/3-25 narrow band-pass filter with an effective bandwidth of 7.2 nm is used. In addition, a long pass filter, Schott RG850, is used to block visible light that can pass through the narrow band filter. Together these filters block most of the solar spectrum in the sensitive range of the avalanche photodiode (APD) detector. The transmission spectra of the two filters are shown in Figure 2. The sensitivity of the detector is shown in Figure 3. Using plotted radiance data from [1] the values for radiance from the clear sky and sunlit cloud are found to be $2*10^2$ $\mu W\ cm^{-2}\ sr^{-1}\ \mu m^{-1}$ and $2*10^3$ $\mu W\ cm^{-2}\ sr^{-1}\ \mu m^{-1}$ respectively. Multiplying by the area of the APD detector and the solid angle of the focusing light, we estimate that the background signal level coming out of the APD module is approximately 1.3 mV from clear sky and 13 mV from a sunlit cloud. These values are net of an estimated optical system throughput of 78.7% at 1550 nm.

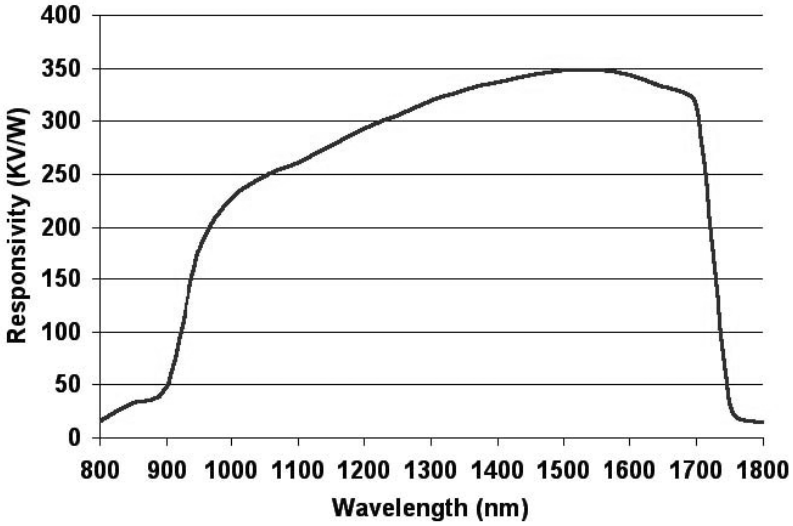

Figure 3. Responsivity of the APD used in WiLD.

A Tasco 4x32 rifle scope is mounted on WiLD above and to one side of the objective lens (see Fig. 4). A periscope is mounted to the end of the rifle scope such that the first scan mirror is in the field of view of the rifle scope. The rifle scope then provides a continuous view of the target. With a reticle in the finder scope aligned to a target, the scope can be used to precisely identify the lidar target direction including deflection by the scan mirrors. The rifle scope is particularly useful for alignment and measurements made while the scan mirrors are stationary. (See discussions of the 1D collection mode below.)

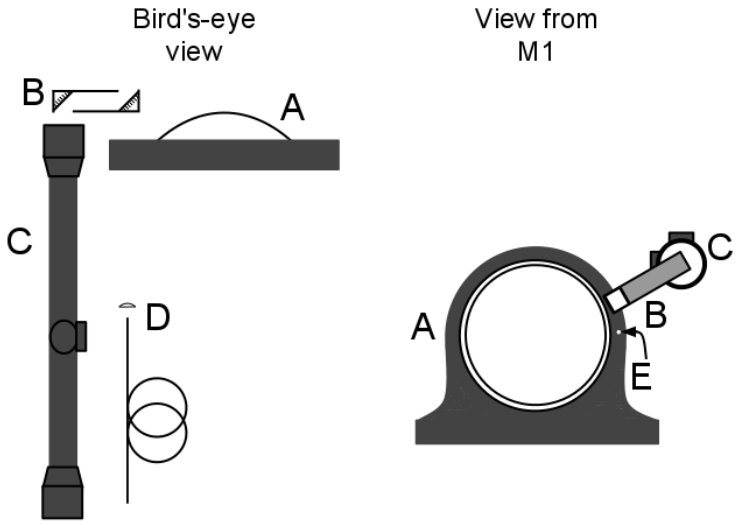

Figure 4. Relative positions of the objective lens (A), periscope (B), rifle scope(C), fiber and collimating lens (D), and aperture for the laser beam (E) illustrated from two perspectives.

WiLD Scanner

The WiLD transceiver optics - including the laser beam, the receiver aperture, and the periscope field of view (FOV) point into a two-axis scanner. The scanner is illustrated in Figure 5.

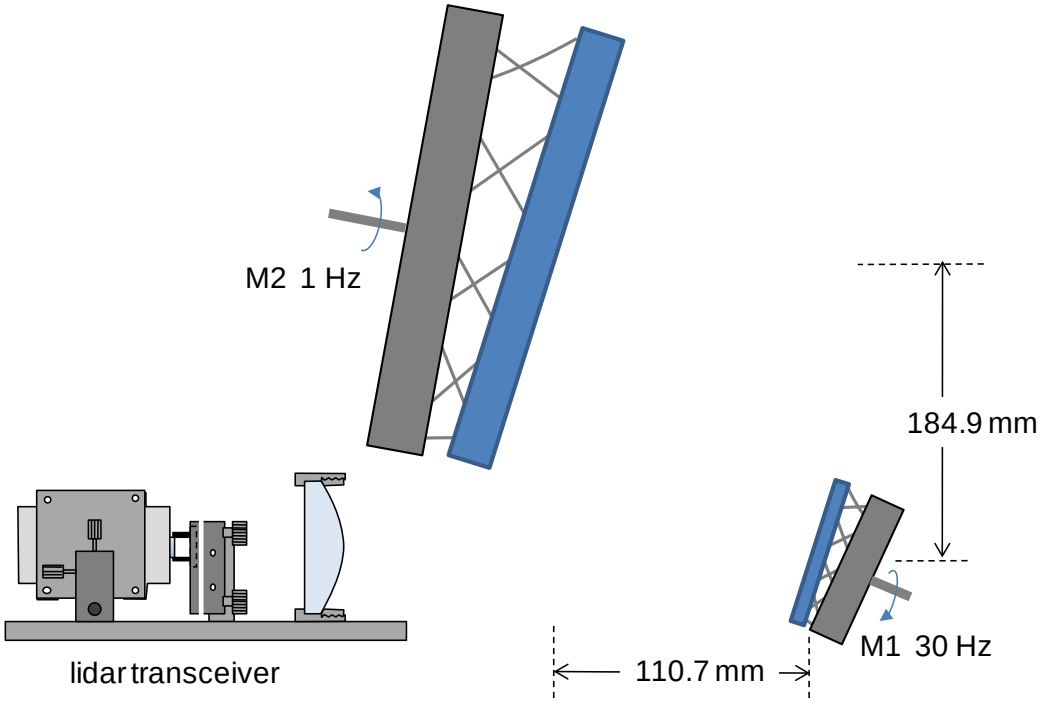

Figure 5. WiLD two-axis scanner configuration.

The horizontal receiver axis is centered on the first scan mirror (M1, diameter 121 mm). The scan motor axis for M1 points upward by 25.0° from the horizontal and M1 is tilted 7.0° with respect to the motor axis. As M1 spins (continuously at 30 Hz), the reflected elevation of the lidar beam between M1 and M2 oscillates rapidly in the range 36.0° - 64.0° above the horizontal. The mounts for M1 and M2 include counterbalances to balance the moment of inertia and control vibration. Both mirrors rotate in a counter-clockwise sense as viewed from their motors.

The scan motor axis for the second scan mirror (M2, diameter 288 mm) points downward by 9.6° from the horizontal and M2 is tilted 7.5° with respect to its motor axis. As M2 spins (continuously at 1 Hz), the doubly reflected lidar beam emerging from M2 oscillates through the range 1.8° - 58.8° above the horizontal forming the closed hypotrochoid-like scan pattern shown in Figure 6. The scan rate (LOS motion) is 23 mrad/ms, modulated slightly by the secondary rotation of M2. The mirror and motor tilts were designed so that the scan pattern has an angular diameter of 1 radian and a solid angle coverage of 0.8 steradian, with the bottom of the scan pattern a few degrees above horizontal. Subsequent field experience suggests that utility of the scanner could be improved by tilting the scan pattern upward a few more degrees because the bottom of the scan frequently encounters foreground objects that interfere with the lidar channel. In practice, this adjustment was made by tilting the entire WiLD scanner platform.

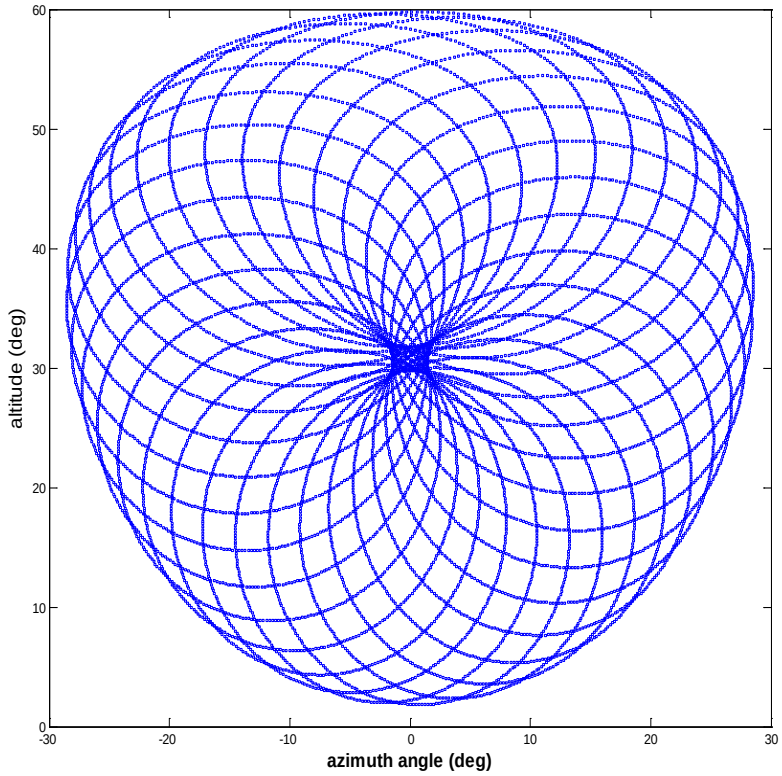

Figure 6. WiLD scan pattern.

The size of M2 and the spacings between the scan mirrors is constrained by the requirement that the full lidar aperture – including the offset laser beam and periscope FOV – not be vignetted for all rotational positions of the mirrors. Thus the lidar system can operate continuously without any interruption associated with the scan. Possible obstructions that must be considered in the scanner design include 1) the lidar beam leaving the transceiver and passing M2 from behind, 2) the lidar beam incident on M1, 3) the lidar beam incident on M2, and 4) the lidar beam from M2 passing M1. The minimum size and volume of the 2-axis scanner scales with the required angular extent of the scan pattern and the diameter of the lidar aperture. The WiLD scanner packaging is somewhat clumsy (35 x 29 x 41 cm including motors); it would be even more so if the receiver aperture were increased.

Both M1 and M2 are driven by rotary servo motors (QCI-M23L and QCI-A34LC, respectively) with closed-loop stepper controllers. Although in principle either M1 or M2 could be driven at the higher rate (30 Hz vs. 1 Hz) it is advantageous with respect to power, accuracy, and safety to drive the smaller motor at the faster rate.

The scan mirror motors include controllers (40,000 counts per rev for M1 and 16,000 counts per rev for M2) that provide motion feedback for a pair of QCI-D2-IG8 Quicksilver Controllers. High-level firmware resident in the two controllers automatically initializes the mirror starting orientations and slaves M2 to the actual rotation of M1 at an exact ratio of 1:30. Volume scanning is initiated simply by turning on the two controllers in sequence (M2, then M1).

Measurements of the scanner stability is are shown in Figure 7. The upper plot shows that the small mirror, after the first few revolutions, maintains a constant rotation rate within ±0.015 rps, with a dominant once-around residual. This contributes an open-loop scanning error of only ±0.16 mrad in the WiLD target space. The lower plot shows that M2 settles down to 0.4% flutter as it follows the rotation of M1 with a dominant residual frequency at about 6 Hz. The M2 motor contributes ±0.17 mrad to the open-loop scanning error, for a total rms error < 0.13 mrad for each axis. Thus the WiLD scanner easily supports scanning applications with resolution and co-registration requirements of order 1 mrad.

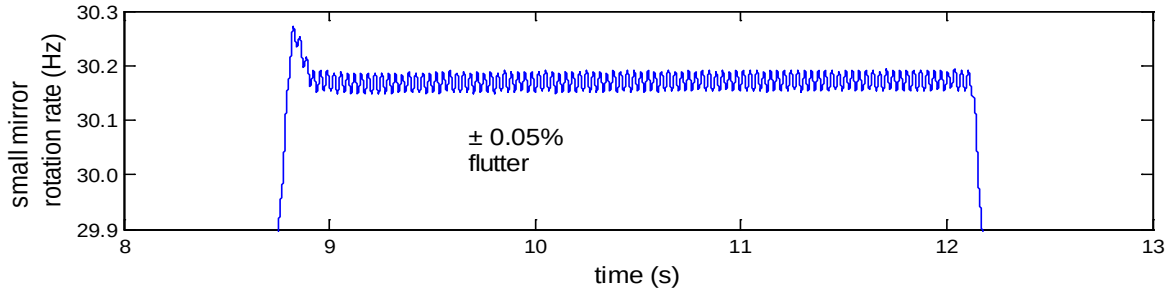

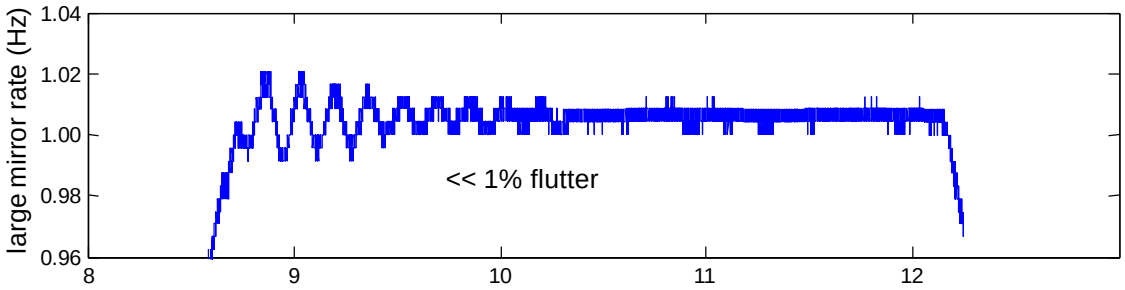

Figure 7. Scan motor stability.

As a cost-saving measure, and recognizing the likelihood of damage to unprotected mirrors in the field, we designed first-surface float-glass mirrors for the WiLD scanner. With a unit price < $60 each, these mirrors cut from bulk sheets are much less expensive than standard optical mirrors that must be custom ground and polished. M1 and M2 have thickness 3 mm and 6 mm, respectively, and mass 168 g and 954 g, respectively. Figure 8 shows inspection data for the scanning runout caused by thickness non-uniformity for two thicknesses of float-glass mirror (both from GrayGlass). Because the thicker float glass is significantly flatter, we used it for both M1 and M2. Together, the mirrors contribute an open-loop scanning error less than 0.2 mrad rms per axis, comparable to the motor errors.

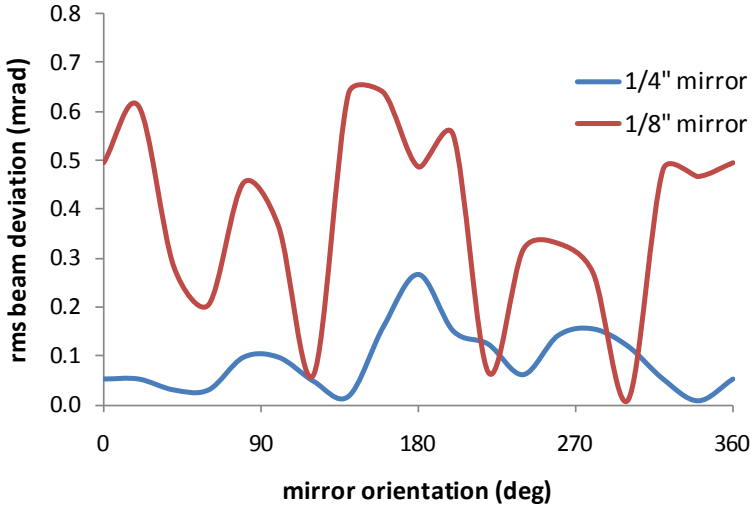

Figure 8. Scanning errors associated with the float-glass mirrors.

In addition to contributing to pointing error, mirror non-flatness can also cause blurring of the lidar receiver FOV and periodic signal fluctuation due to runout of the laser beam relative to the receiver FOV. However, in the WiLD system these effects are much smaller than the receiver FOV and the laser beam divergence as described above.

WiLD Analog Channel

Signal detection in the WiLD transceiver is performed by the avalanche photodiode (APD) module, model C30659-1550-R2A, illustrated in Figure 9a. This module combines a 200 μm diameter InGaAs APD with a maximum spectral response near 1550 nm. The detector is integrated with a trans-impedance amplifier (TIA); the circuit diagram is shown in Figure 9b.

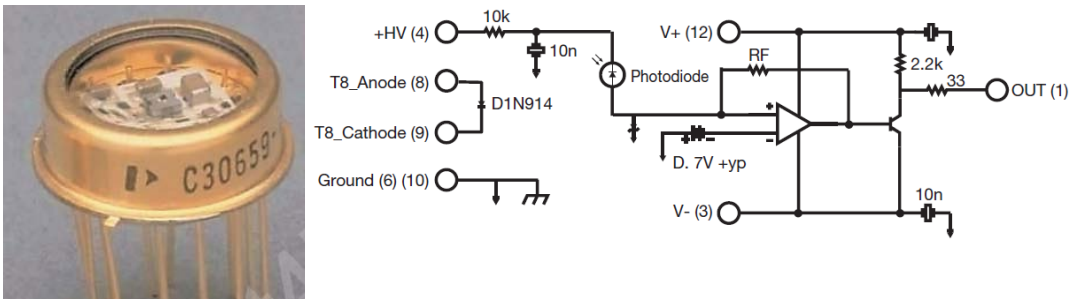

Figure 9. a) The APD module. b) Integrated TIA circuit.

The TIA is tuned for a bandwidth of 50 MHz and the APD module has a rise-time of 7 ns for short laser pulses, equivalent to a range resolution of 7ns·c/2 = 1m. Note however that in the WiLD system subsequent channel components filter the signal to a bandwidth of 25 MHz. This increases the resolvable range increment (of the analog channel) to 2 m.

The APD module is rated for a nominal responsivity of 340 kV/W and a typical noise-equivalent power (NEP) of 150 fW/√Hz at 1550 nm. With a 25 MHz bandwidth, the noise on a 100 megasamples per second (MSPS) data sample is equivalent to the signal from a 10 ns pulse of 56 photons incident on the detector surface. The nominal noise amplitude under these conditions is 0.26 mV.

To achieve avalanche amplification within the APD, a high voltage bias must be applied across the APD. Optimal performance is achieved when the bias is a maximum consistent with stable operation, approximately 43V for this APD. The APD module for WiLD is housed in a metal enclosure - on the left in Figure 10a. A separate APD control box - on the right in Figure 10a - generates the bias voltage, monitors the bias, and powers the APD circuit (via battery or A/C adapter).

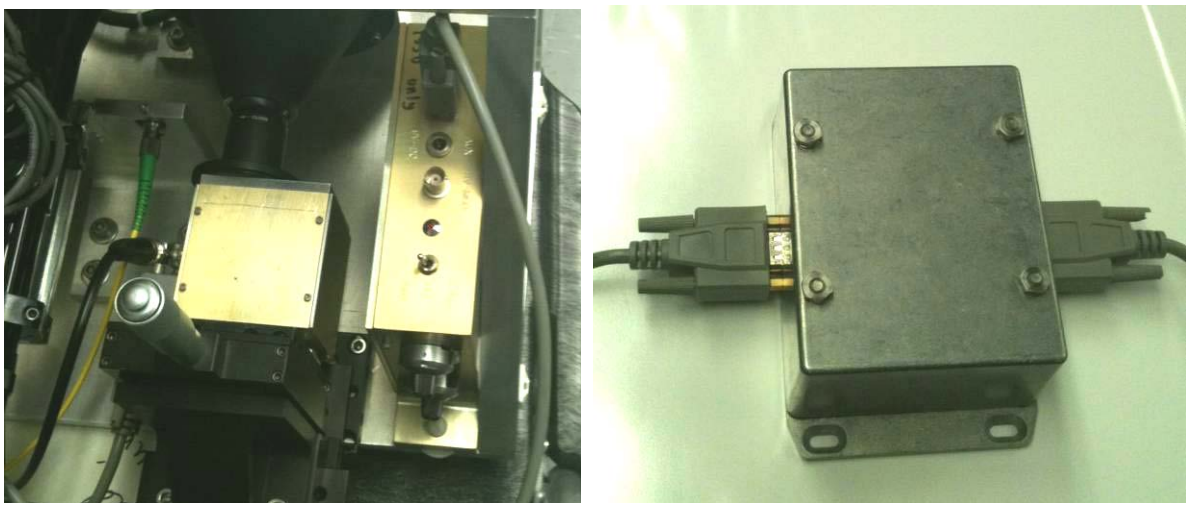

Figure 10 a) APD housing (left) and APD control box (right). b) Bias voltage low-pass filter.

Stabilization of the APD bias voltage is critical for volume imaging applications. Any slight low-frequency modulation of the bias is directly superimposed on the signal waveforms, and can create spurious spatial patterns in derived aerosol images. Initial testing of the WiLD receiver showed the presence of noise around 7 kHz from the high-voltage supply. These oscillations were mitigated by adding a passive low-pass filter (Figure 10b). Even with the addition of the low-pass filter, low-frequency fluctuation of APD output continues to be noticeable. Battery operation of the APD control box and electrical isolation of the low-pass filter help. At least two remaining kinds of low-frequency fluctuation are apparently due to EMI.

- Whenever the laser is triggered, a waveform about 1 μs long with 50-100 ns oscillations is observed. This waveform precedes any optical output from the laser. The amplitude is several times the noise level.
- While the scanner is operating, very low frequency fluctuations (< 20 Hz) are observed at the APD output. (See section V.)

In order to realize the full capabilities of a high rep-rate lidar system, it must be possible to enhance SNR by averaging multiple data samples and waveforms. Digitization must not add noise that is comparable to the intrinsic analog noise. This observation motivated us to add a low-noise amplifier to the analog data channel. An Analog Modules model 353A amplifier (Figure 11) provides 30 dB gain with input-referenced noise performance of 2nV/√Hz. For our 25 MHz channel, the amplifier input noise level is 10 μV - much lower than the APD readout noise. The amplified signal can be sampled at an A/D resolution of 1 mV without loss of any potential SNR.

Figure 11. Low-noise signal amplifier.

A drawback to the Analog Modules amplifier is that the amplified output has a DC offset that cannot be shifted below 30 mV. This makes it difficult to make high precision readings using a typical oscilloscope.

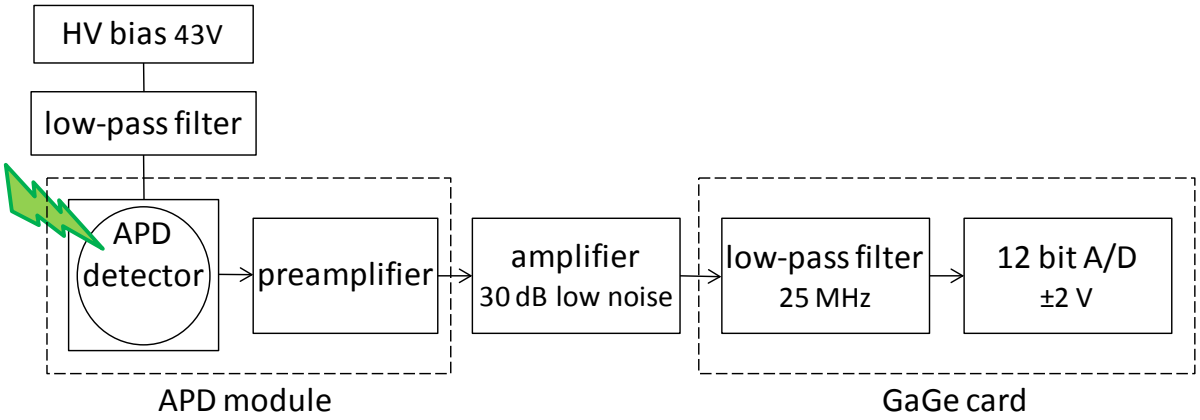

Figure 12. The WiLD channel.

The structure of the WiLD data channel from the photodetector to digitization is illustrated in Figure 12. The APD preamplifier has a nominal output impedence of 40 Ω and the amplifier input and output impedences are 50 Ω. 50 Ω termination is also selected (via software) for the data acquisition card and the channel components are connected directly (without additional terminators). The analog input to the data acquisition card includes a 25 MHz selectable low-pass filter that sets the effective bandwidth of the analog channel. We experimented with additional low-pass filtering (20 MHz and 12 MHz) and with the 25 MHz filter selected or deselected. The 25 MHz filter without additional attenuation yielded the minimum input-referenced noise.

Digital Data Collection

The WiLD control system is a hybrid of manual and automatic elements. Following is the control sequence for a scanned, 3D WiLD data collection. See Figure 13 for an illustration of the control and data acquisition (DAQ) components and their interconnections.

a. The user manually turns on the laser system in triggered mode; the laser warms up and awaits the presence of trigger signals.

b. The user starts a LabView program on the laptop control computer that initiates communication with the encoder DAQ card (National Instruments USB-6212). An input parameter is set equal to the number of laser pulses planned for the collection.

c. The user starts a C++ routine on the control computer that sets the collection parameters of the waveform DAQ card (GaGe model RAZOR 1621) including number of pulses, waveform length, and digitization range.

d. The user manually turns on first the M2 motor controller, then the M1 motor controller. This sequence initializes the mirror orientations and sets them spinning in tandem. The mirrors spin for approximately 1 minute, long enough to start and complete a 10-second data collection.

e. The user clicks Run in the LabView program GUI; the encoder DAQ is now waiting for a trigger to start encoder data collection. At each trigger event, the encoder DAQ records the values of free-running counters for the two motor encoders. These values are reset to zero once per revolution by the encoder index signal.

f. The user enters a run command in the C++ routine to initiate waveform collection by the waveform DAQ (pending a trigger). The waveform DAQ collects one waveform (programmable-length sequence of data values at 100 MSPS) starting at each trigger. Generation of the laser pulse is delayed by approximately 500ns, so the actual length of the lidar signal stream is about 50 samples shorter than the collected waveform.

g. The user manually turns on a portable pulse generator (BK Precision model 3003). The 50 kHz output is distributed three ways to simultaneously trigger the laser, the trigger the waveform DAQ, and trigger the encoder DAQ.

h. When the programmed number of pulses have occurred, the user sees the LabView collection sequence shut off. The encoder data is automatically saved to a predefined directory on the control computer.

i. To maintain laser safety, the user turn off the trigger before the motors stop spinning. (See Section VII.)

j. The C++ routine automatically starts downloading data from the waveform DAQ to a predefined directory on the control computer. This data transfer can take several minutes. Progress is illustrated by repeated characters printed to the screen.

k. After all data tranfers are complete, the user renames data files (so they won't later be over-written). The WiLD system is now ready for a new collection (returning to step d).

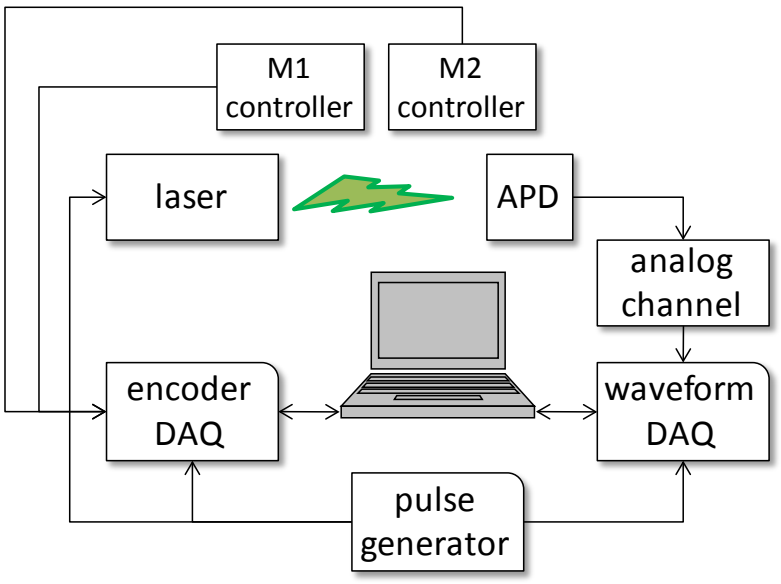

Figure 13. Data acquisition and system control architecture.

Following are some suggested operational improvements for WiLD based on experience to date.

- Identify robust settings (amplitude and duty cycle) for the trigger source for simultaneous triggering of the endoder DAQ, waveform DAQ, and laser. Provide a fourth trigger output for diagnostic purposes (including verification of the trigger settings in the field).
- Combine the functionalities of the encoder DAQ and waveform DAQ software into a single program.
- Provide a GUI that displays the DAQ settings and allows them to be modified without manually editing configuration files.
- Create time-tagged file names and save the data in a partition that has sufficient memory for many collections.

The waveform DAQ card has 2 GBytes (1 GSample) of onboard memory (Figure 14). The WiLD system utilizes only one channel (of this dual-channel card). In the preferred operating mode it collects data at 100 MSPS with 16 bit resolution over a signal range of ±2V. The "effective number of bits" (ENOB) from the digitizer is 12 bits, for a full signal resolution of 1mV. Thus digitization does not add significantly to the noise from the amplified analog channel (approximately 7 mV for each sample).

Figure 14. GaGe data acquisition card.

The GaGe card is linked to the WiLD control computer by a PCI bus. The card and bus support data transfer at an instantaneous rate of 200 MB/s. However, discontinuous data flow and the short record length of lidar waveforms slow the channel significantly (to about 10 Mbyte/s, net), so that WiLD cannot be operated continuously. Instead, we operate WiLD for discrete intervals consistent with the GaGe-card memory and then transfer that data to the computer before repeating the operation.

1D Collection Mode

Although the WiLD system is designed for volume lidar imaging, for purposes of debugging, channel characterization, and atmospheric phenomenology it is useful to operate the system without scanning. In this 1D collection mode, the waveform DAQ collects a long sequence of lidar waveforms that sample the atmosphere along a single, unmoving line of sight (LOS). In order to extend the duration of the 1D collection, the sampling rate of the waveform DAQ is reduced to 800 Hz using a second pulse generator. The laser is still triggered at the nominal 50 kHz rate.

The laser pulses and waveform acquisition triggers are unsynchronized in this mode. Therefore, to ensure that complete lidar return is included with each capture, it is necessary to capture waveforms that are longer than the laser pulse period. We set the captured waveform to 2500 samples = 2.5 μs (at 100 MSPS) which is 25% longer than the laser period. If a laser pulse happens to occur just prior to the beginning of waveform collection, the total capture will include another laser pulse late in the sequence with at least 500 lidar data samples following. Note that approximately 80% of the collections will include two laser pulses.

A 1D collection consists of 80,000 waveforms, collected over a 100 s interval. The extended interval allows for identification of transient phenomena (e.g. dust clouds) that may drift slowly through the LOS. The GaGe card stores the waveforms internally using a format that takes 10 bytes per value. The memory required for the full collection is 800 x 100 x 2500 x 10 = 2 GB, which fills the on-board memory.

Lidar data from the 1D collection mode must be synchronized after the fact by analysis of the waveforms. A small fraction of the light from each emitted laser pulse is doubly scattered by the mirror surfaces such that it returns directly into the receiver FOV. Because this scattered light source is so close to the receiver, it creates a strong pre-pulse from the APD that is uniquely identifiable in the collected waveform (see Figure 15). The risetime of the pre-pulse is consistent with the bandwidth of the APD module (50 MHz) but the trailing edge is typically slower because of detector saturation effects due to the large pre-pulse amplitude. The pre-pulse is delayed by only 1-3 ns relative to the actual laser pulse (much less than the 10 ns sampling period). Preprocessing of the 1D WiLD waveform starts with identification of the (first) laser pre-pulse and selection of a fixed set of samples following the pre-pulse rising threshold.

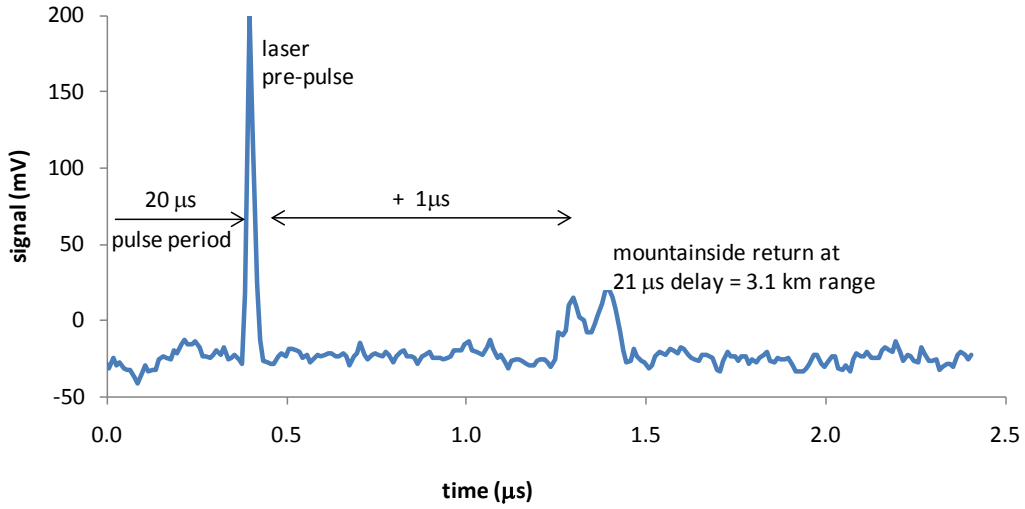

Figure 15. WiLD lidar waveform showing the laser pre-pulse and a return from a surface.

3D Collection Mode

For lidar volume imaging, data collection includes both lidar waveforms and encoder counts. These data, although they are stored in different files, are synchronized by virtue of the common trigger. Since waveform acquisition is synchronized to the laser pulse, a waveform length of 400 samples (4 µs at 100 MSPS) covers a lidar range of 500 m after trimming the delay from trigger to pre-pulse.

The WiLD laser exhibits jitter of order 10 ns rms in the delay from trigger to outgoing laser pulse. To minimize the effects of this jitter, the 3D waveforms are preprocessed like the 1D waveforms to identify the laser pre-pulse and optimize the relative waveform timing.

A full 3D collection consists of 50,000 waveforms per second, collected over a 10 s interval. The memory required for the full 3D collection is 50,000 x 10 x 400 x 10 = 2 GB, using up the card capacity.

III. 1D Lidar Observations.

WiLD Range Demonstration

Although WiLD is intended to acquire signal returns from aerosols, hard target signal returns aid in alignment and provide unambiguous information about system performance. A typical aerosol data collection is preceded by a field alignment. In a field alignment WiLD is pointed at an object such as a building; non-specular surfaces are preferred. Pointing WiLD is achieved by moving the cart on which the scanner and laser sit and by manually rotating the scan mirrors. Manual rotation of the mirrors is performed when the power to the motors is disengaged. Weak detent forces hold the motors fixed when they are released.

A building within 50 m is a good initial target. After the signal return has been maximized via micrometer adjustment of the detector position, a more distant object can be chosen. For aerosol measurement, a target at 100 m is desired for the final alignment.

To test the capability of the instrument, targets at progressively longer distances can be chosen and the signal return maximized at each of the positions. As illustrated in Figure 15, a hard target signal return was detected from a mountainside at 3.1 km. The figure also shows that if the range exceeds $0.5 \cdot c/PRF = 3$ km the next laser prepulse will occurs before the return signal is detected. (This effect can create confusion in interpreting the lidar signal if a short range aerosol signal from the second pulse is coincident with the long-range, hard target return from the first pulse.) Note that the signal from hard targets at large distances is weak but sharp. Interference with aerosol returns is only likely when the hard-target range is close 3 km. To date, the largest demonstrated hard-target range is approximately 5 km, coming from the face of a nearby mountain.

1D Waveform Processing

Because the 1D collection mode maintains a constant LOS, it supports a high degree of waveform aggregation. The typical processing sequence is as follows.

a) Identify the laser pre-pulse in each waveform (as discussed in section II) and trim the data to include 100 samples before the pre-pulse and at least 200 samples after the pre-pulse (300 m lidar range).

b) Aggregate groups of 100 consecutive waveforms (mean waveforms). The purpose of aggregation is to reduce temporal noise. Aggregation also filters out any fast temporal phenomena, but little change in aerosol scattering is expected over the 1/8 s required to collect 100 waveforms.

c) For each of the 800 aggregate waveforms, subtract an offset equal to the mean of the first 75 samples, a region that is unaffected by EMI from the laser. This step reduces the impact of low frequency signal offsets.

d) Create a "typical" waveform from the median values at each sample time. This step eliminates to first order the disturbance due to ring-down of the laser pre-pulse and EMI from the laser.

e) Zero-out any outliers that are more than 5-sigma from the mean, where sigma is the standard deviation at each sample time. Outliers (e.g. due to bugs, EMI, or noise glitches) may be very large and tend to visually obscure underlying signals.

f) Apply a box-car filter that is 2 samples long (3-meter range) and 16 waveforms wide (2 s duration). This further heightens the lidar sensitivity without compromising spatial or temporal resolution for the ambient aerosol features that we have observed.

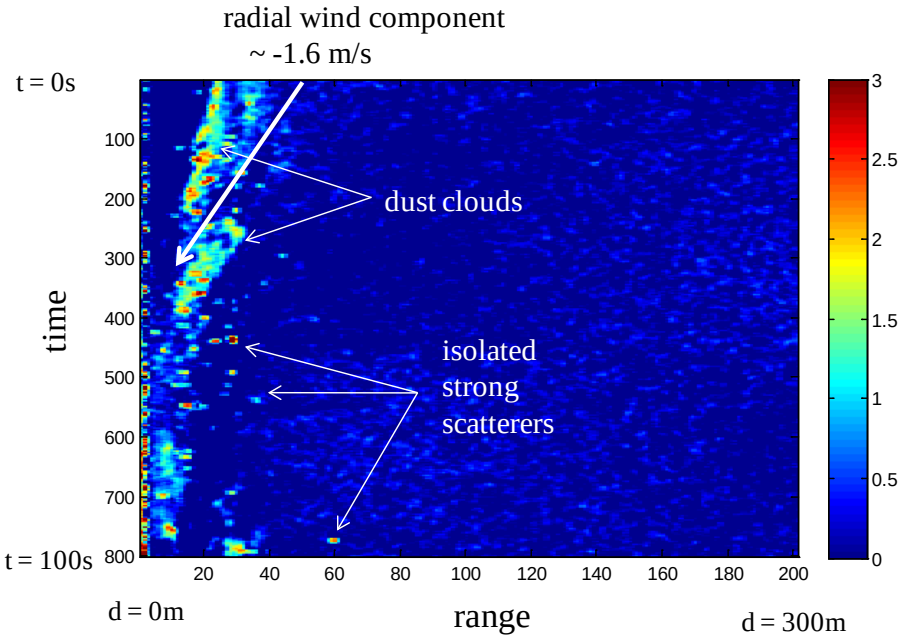

Figure 16. Processed 1D collection showing ambient aerosol motions.

An early example of processed 1D data is shown in Figure 16, collected at 12:40 pm on 14 September, 2010. The view angle is ~5° above the horizon. Aerosol clouds are evident at short range (<60 m). The elongated shape of the features suggests that the aerosols clouds are moving generally toward the WiLD system with a radial wind component of -1.6 m/s. In this example, isolated strong scatterers (speculated to be insects or wind-blown seeds) are also seen at short range.

Stack Emissions

Data collections on 25 March 2011 included two 1D collections with the WiLD LOS pointed toward the south southeast above the roof of the USU BioInnovations Center near the a small exhaust vent that shows a steady flow of visible condensation (steam) during cold weather. These 1D collections are illustrated in Figure 17, with logarithmic color coding to display the

large dynamic range. Both collections show strong returns from the exhaust aerosols, nearly continuous on the right and intermittent on the left. Slight range variations are seen on the left-hand image, indicating wind fluctuations driving the plume north and south. Intermittency (especially in the right-hand image) indicates that the plume also fluctuates east and west. The plume signatures shows up at a delay of 2.5 ms, for a range of 375 m.

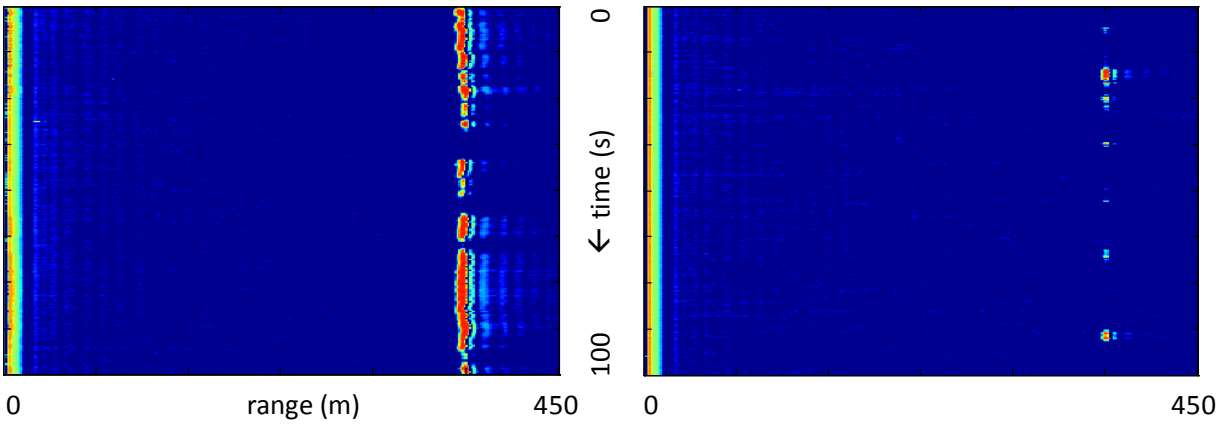

Figure 17. Waveforms from exhaust "steam".

The signature from the plume is approximately 5 samples wide, consistent with the 25 MHz bandwidth of the WiLD receiver. The amplitude of the plume waveform exceeds 0.8 V - much higher even than the laser pre-pulse. Thus exhaust plumes at ranges less than 400 m are detectable with high signal to noise ratio (SNR) without waveform aggregation. Note that the left-hand image shows ringing behind the plume signature. This may be the same oscillation phenomenon discussed below in connection with validation of the overlap function.

WiLD Overlap Function

The lidar equation relates the signal amplitude (S) to the pulse energy (E), aerosol concentration (ρ), and range distance (r):

$$S(r) = \frac{E \cdot \rho \cdot \sigma_b}{2c} \cdot R \cdot \frac{\pi \cdot d_A^2}{4r^2} \cdot e^{-2D(r)} \cdot f_{overlap}(r) \qquad \text{(III.1)}$$

where R is the receiver responsivity (including optical losses), σ_b is the backscattering cross-section per particle, d_A is the receiving aperture diameter, and D(r) is the cumulative optical density out to range r (a small number under conditions of good visibility). The overlap function, $f_{overlap}$, describes the fraction of the aperture that is illuminated within the field of view from the laser beam position at range r. This function describes deviations from the asymptotic square-law fall-off of lidar response versus range. Figure 18 plots the expected overlap function for several alignment conditions (pointing of the laser beam with respect to the receiver optical axis). The model assumes that the laser beam divergence is much smaller than the receiver FOV

(approximately 0.7 mrad vs. 2 mrad). The x and y coordinates correspond to horizontal and vertical offsets, respectively. The laser beam offset is illustrated in Figure 4.

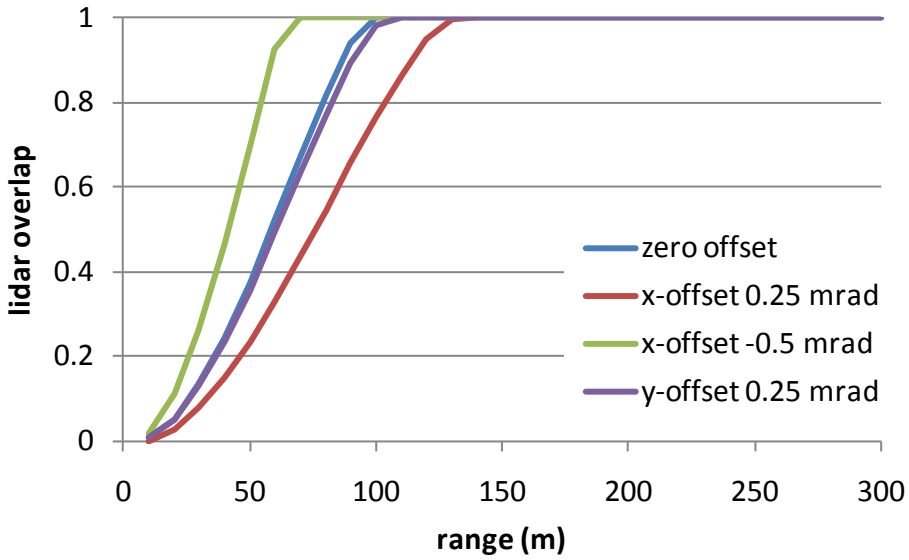

Figure 18. Predicted overlap function WiLD.

Comparison of the first and fourth curves shows that the overlap function is relatively insensitive to vertical alignment offsets. (An offset of 0.25 mrad corresponds to 1 mil at the APD alignment micrometers.) However, the overlap at short range is sensitive to the horizontal alignment. Short range sensitivity can be optimized without loss of longer range sensitivity (out to 300 m) by tipping the receiver FOV toward the offset laser beam by 0.5 mrad. This alignment condition is accomplished by pointing the system toward a target at 100 m range and centering the alignment with respect to the signal maxima for horizontal and vertical adjustments.

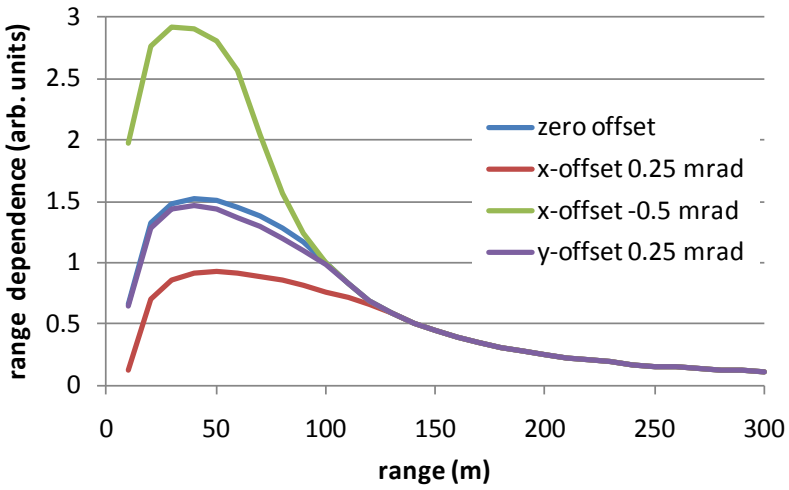

Figure 19. Combined range dependence of WiLD.

The combined range dependence ($f_{overlap}$ / range^2) is plotted in Figure 19. WiLD is quite insensitive to alignment beyond 100m. At the recommended alignment (-0.5 mrad horizontal offset), the response follows the inverse-square law down to 70 m and the system shows a strong response down to a range of 10m. Unlike many lidar systems, WiLD has virtually no dead zone at short range.

On 25 March 2011, a set of 1D collections was obtained with the intent of verifying the WiLD overlap function. The weather was overcast, 40°F, and not very hazy. The wind was blowing at < 3 m/s from the southwest while WiLD was pointed toward the southeast. Four collections were obtained, two pointed near the horizon, one pointed 30° upward, and one pointed 60° upward.

After scaling the time-averaged waveforms for each collection by range^2 (in meters), we find a strong high-frequency component (spatial frequency > 0.1 m^{-1} with respect to range) that is similar in all four collections (see Figure 20). This may be the tail of detector ring-down from the laser pre-pulse or the EMI signal from the pulsing laser. Further study is needed to establish the cause and optimize the post-processing treatment of this phenomenon.

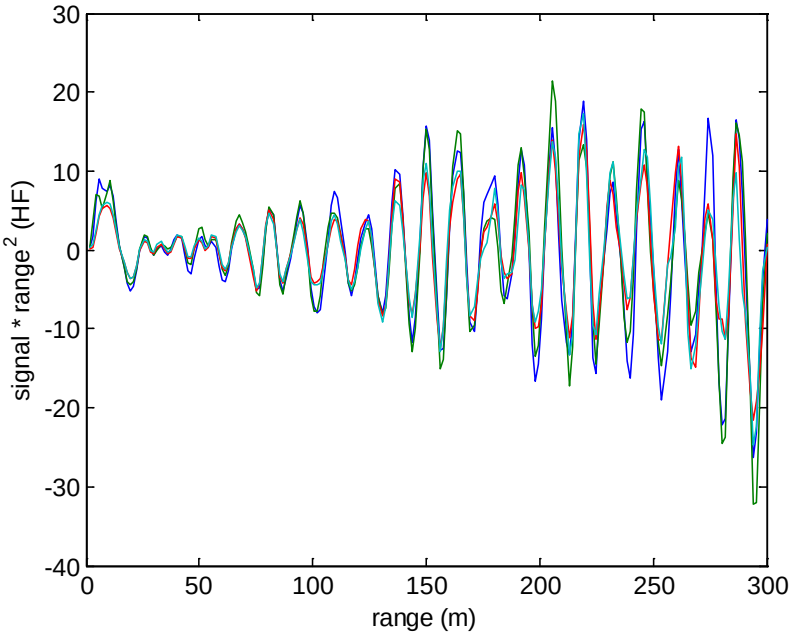

Figure 20. Residual signal disturbance after high-pass filtering of range-scaled 1D waveforms.

The range-scaled 1D waveforms are shown in Figure 21 with the mean high frequency disturbance subtracted. Adjustments have also been made for offset and slope of the detector reference voltage (fit to minimize range dependence around 300m). The results are qualitatively similar for all four collections. But none of the traces rise as rapidly as the predicted overlap

function (c.f. Figure 18). Possibly, all of the waveforms are affected by increasing aerosol concentration with range, although this is especially unlikely for different view elevations.

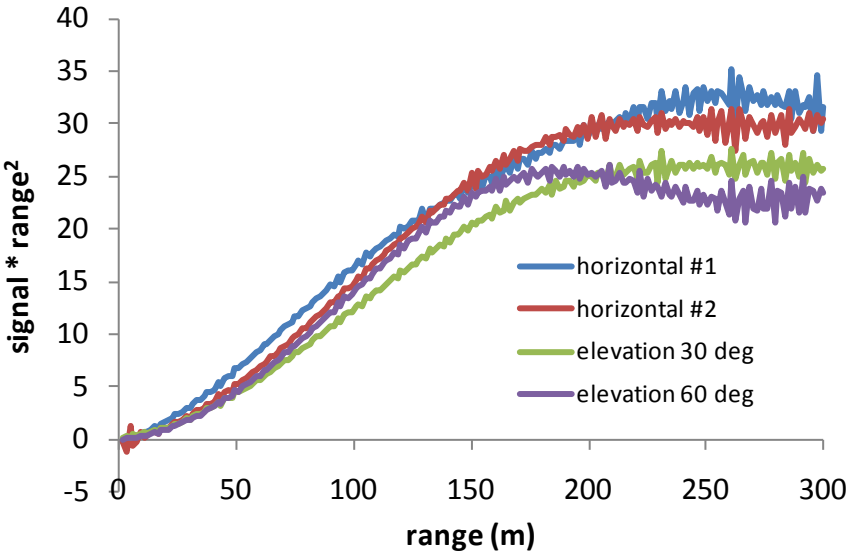

Figure 21. Range-scaled 1D waveforms with common disturbance removed.

Validation of the overlap interval will be easier to establish by taking collections under hazy or foggy conditions where the return signal is greater and the aerosol uniformity can be qualitatively confirmed.

Altitude Dependence.

A prerequisite for wind-field sensing using ambient aerosols is that aerosol non-homogeneity be present in all regions of the atmospheric volume and, in particular, at all altitudes. 1D data collections over a range of elevation angles can be used to assess relative aerosol concentrations vs. altitude.

The first attempt at vertical aerosol sounding using WiLD was performed on the morning of 4 November 2010. The sensor was pointed east over the 600 East roadway. The weather was fair with very light haze. The 1D data shown in Figure 22 shows that aerosols were observed at short range. Rows in each image show the lidar waveforms aggregated over 1/8 second. The laser pre-pulse shows up in each row at sample 100. The right-most edge of each image corresponds to a range of 300 m. The sequence of 1D collections included two vertical "scans" with three positions each. In each vertical scan, the overall aerosol concentration decreased steadily with elevation. Untilted vertical streaks in these images are due to residual signal disturbances (see Fig. 20) that were not backed out of this data. This result suggests that under ambient conditions the aerosol concentration decreases strongly with altitude. In this data set there is no positive indication of aerosols above 50 m.

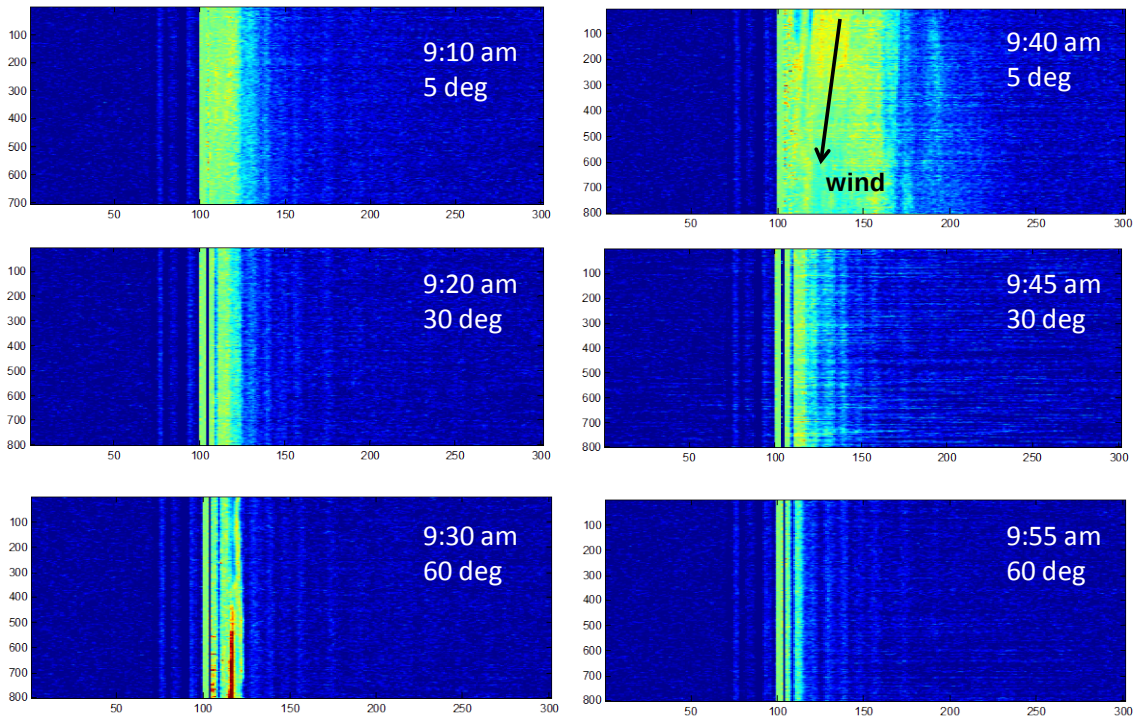

Figure 22. 1D observations at varying elevation angles.

Spatially varying aerosol patterns are evident in two of these observations. During the second half of the 9:30 am collection, an aerosol spike shows up at short range (20-25 m). This is likely an exhaust plume rising off of the nearby roadway. During the 9:40 am collection, aerosol patterns show up over the range 0 - 150m. These patterns appear to shift inward with time. The slope of the patterns is indicative of the radial velocity component of the wind, in this case approximately 1 m/s toward the WiLD receiver.

Ambient Aerosol Observations

An early example of aerosol cloud detection by WiLD was presented in Figure 16. A series of additional 1D collections were performed to investigate the presence and detectability of ambient aerosol clouds. Figure 23 shows ambient aerosol observations from 9/17/2010 (left) and 10/19/2010 (right). The lidar system is pointed upward at approximately 5°. The horizontal axes count the waveform samples (at 100 MSPS so that e.g. the 200th sample corresponds to a range of 300 m). The vertical axes count the 100-waveform aggregates (at 800 waveforms per second so that 800 rows correspond to 100 s). The collection from 9/17/2010 shows a compact aerosol cloud with a depth of about 8 m crossing the lidar beam transversely (with very little radial velocity) over the interval t = 80 - 90 s. If this is a roughly spherical cloud, then its velocity is about 1 m/s. The collection from 10/19/2010 shows several clearly defined aerosol clouds traveling away from the WiLD receiver with a radial velocity component of +1.2 m/s (16 range

bins in 20 s). The cloud thicknesses are 10 - 20 m. If the clouds are roughly spherical, then the transverse velocity is also of order 1 m/s.

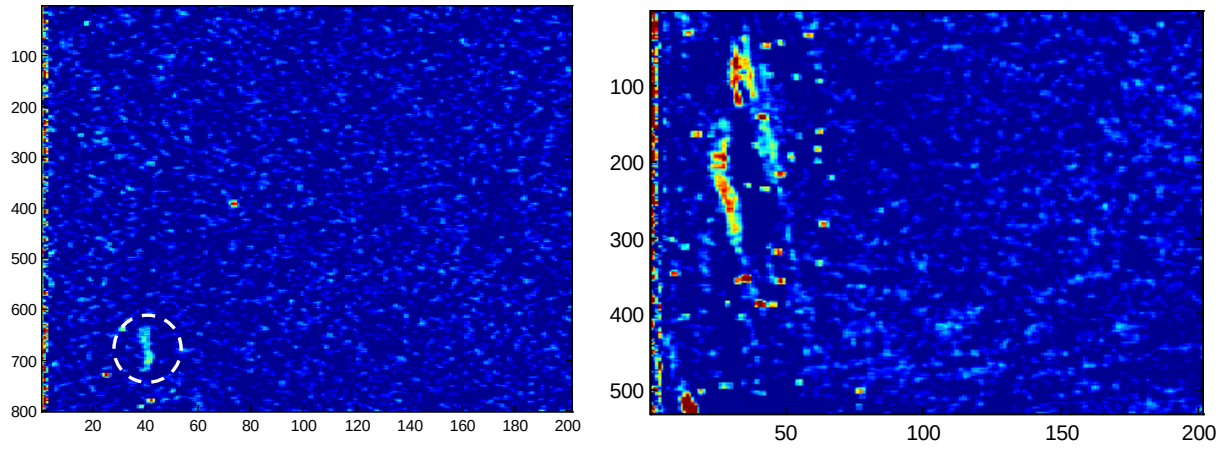

Figure 23. Two 1D observations of ambient aerosol clouds.

The aerosol features in Figure 23 are all nearer than 90 m (60 range bins). An important question suggested by these observations is whether this range limitation is due to sensitivity limitations (cloud concentration x $1/r^2$ falloff) or low-altitude localization of the aerosol clouds (i.e. below about 10m). Figure 24 shows a similar 1D collection obtained on 10/19/2010 with the WiLD LOS elevated more steeply at 15°. This image shows an aerosol cloud with a radial thickness of 20 m at range of 110 m and thus an altitude of 30m. To date, this is the highest altitude ambient aerosol cloud observed by WiLD.

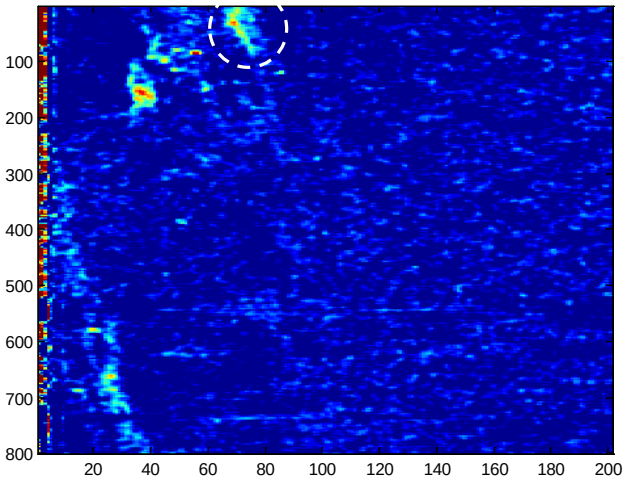

Figure 24. Aerosol clouds detected at a 15° elevation.

The composition of the ambient aerosols observed in Figures 16 and 22-24 is unknown. Strong lidar returns were noted from condensed water droplets (Figure 17). On 20 September 2010 we

had an opportunity to observe lidar returns from mineral dust. Figure 25 shows a 1D collection obtained while a tractor was plowing a nearby field. The tractor initially plowed toward the south while WiLD was pointed eastward. A lightly visible dust plume from the tractor drifted across the lidar LOS about 15 s into the collection. The tractor then turned toward the west, and the plume began to drift toward the WiLD system. The range to the plume started at 400 m and ended at 270 m. The lidar returns from the dust plume are very strong, evident at high SNR without aggregation.

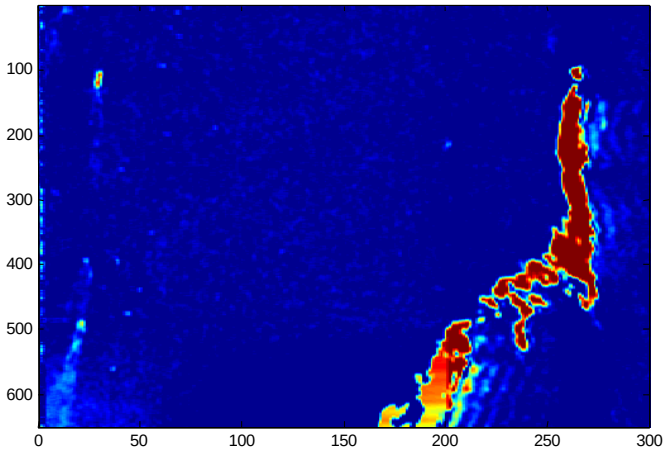

Figure 25. WiLD observation of an agricultural dust plume.

Note starting about 50 s into the collection of Figure 25 (row 400) that the WiLD collection begins to show a high degree of waveform structure. This demonstrates that the dust plume was essentially transparent and further demonstrates that WiLD detects true range profiles (not just initial threshold detections).

IV. Constructing Volume Images.

Scan Angle Conversions

The rotational orientations of the scan mirrors as described in section II determine the hypotrochoid-like scan pattern. The pattern can be calculated relative to the reference coordinate system (i.e. associated with the baseplate) by first finding the unit normal vectors from the mirrors in the reference coordinate system. This is accomplished by multiplying the mirror's unit normal vector in the corresponding motor's reference frame by a transformation matrix representing the rotation of the motor axis. A subsequent transformation brings the unit normal vector into the reference frame. Equations (IV.1) and (IV.2) illustrate this calculation. Angle values in these equations are have units of degrees.

$$\hat{n}_s = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(25) & -\sin(25) \\ 0 & \sin(25) & \cos(25) \end{bmatrix} \begin{bmatrix} \cos(\theta_s) & -\sin(\theta_s) & 0 \\ \sin(\theta_s) & \cos(\theta_s) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ -\sin(7) \\ -\cos(7) \end{bmatrix}$$

$$= \begin{bmatrix} \sin(\theta_s)\sin(7) \\ -\cos(25)\sin(7)\cos(\theta_s) + \sin(25)\cos(7) \\ -\sin(25)\sin(7)\cos(\theta_s) - \cos(25)\cos(7) \end{bmatrix} \quad \text{(IV.1)}$$

$$\hat{n}_L = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(9.6) & -\sin(9.6) \\ 0 & \sin(9.6) & \cos(9.6) \end{bmatrix} \begin{bmatrix} \cos(\theta_L) & \sin(\theta_L) & 0 \\ -\sin(\theta_L) & \cos(\theta_L) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ -\sin(7.5) \\ \cos(7.5) \end{bmatrix}$$

$$= \begin{bmatrix} -\sin(\theta_L)\sin(7.5) \\ -\cos(9.6)\sin(7.5)\cos(\theta_L) - \sin(9.6)\cos(7.5) \\ -\sin(9.6)\sin(7.5)\cos(\theta_L) + \cos(9.6)\cos(7.5) \end{bmatrix} \quad \text{(IV.2)}$$

In equations (1) and (2), $\hat{n}_s$ and $\hat{n}_L$ are the unit normal vectors pointing into the small and large mirrors, respectively, in the reference frame. The angles θ_s and θ_L are the rotation angles of the small and large mirrors, respectively. Both mirrors undergo left-handed rotation as indicated in Figure 26.

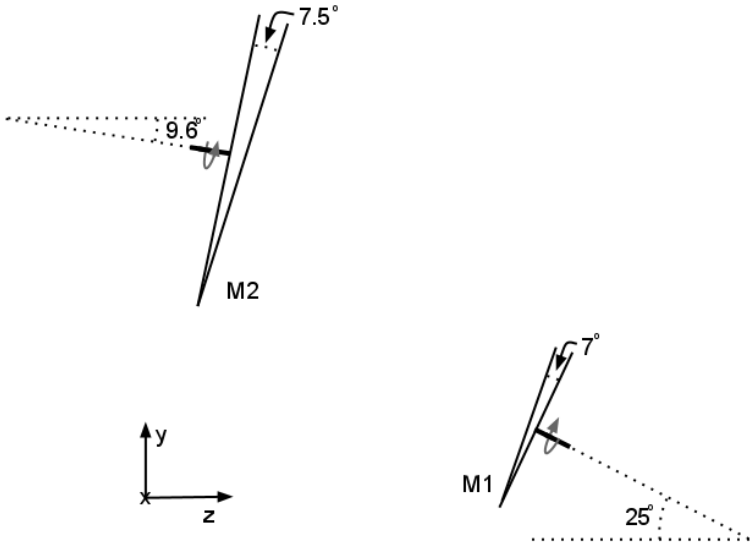

Figure 26. Schematic of scan mirrors M1 and M2 showing relative angular orientations and laboratory reference frame coordinate system. The mirror spin is also depicted.

In the reference coordinate system, the laser beam's initial direction is given by

$$\hat{b} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \quad \text{(IV.3)}$$

After reflection from mirror M1 the beam direction is given by

$$\hat{b}' = \hat{b} - 2(\hat{b} \cdot \hat{n}_s)\hat{n}_s \quad \text{(IV.4)}$$

Similarly, after reflection from mirror M2 the beam direction becomes

$$\hat{b}'' = \hat{b}' - 2(\hat{b}' \cdot \hat{n}_L)\hat{n}_L \quad \text{(IV.5)}$$

The beam direction after the second mirror can be represented as an elevation and azimuth. Elevation, represented by θ and taken to increase from zero at the horizon, and azimuth, represented by φ and taken to increase to the right from directly ahead of the scanner, are related to $\hat{b}''$ according to

$$\hat{b}'' = \begin{bmatrix} -\cos(\theta) \cdot \sin(\varphi) \\ \sin(\theta) \\ \cos(\theta) \cdot \cos(\varphi) \end{bmatrix} = \begin{bmatrix} \hat{b}_x'' \\ \hat{b}_y'' \\ \hat{b}_z'' \end{bmatrix} \quad \text{(IV.6)}$$

Therefore,

$$\theta = \text{asin}\left(\hat{b}_y''\right) \text{ and}$$

$$\varphi = atan2\left(-\hat{b}_x'', \hat{b}_z''\right) \quad \text{(IV.7)}$$

A plot of the hypotrochoid-like scan pattern based on the above formulation is shown in Figure 4. Note that the size of spots in Figure 4 is scaled for visibility and does not represent the 0.7 mrad divergence angle of the beam. The actual beam size would be less than 1 pixel in extent on a square 1 megapixel figure representing the full extent of the scan pattern.

Scan-angle Phasing

The information extracted from the encoders includes the count number of the small motor. With 40,000 counts per revolution, the angular orientation of M1 can be determined according to

$$\delta = \frac{n_S \cdot 2\pi}{40{,}000} \quad \text{(IV.8)}$$

where n_s is the encoder count. If the initial mirror position is not as indicated in Figure 26, then a motor phase offset must be included. The true angular orientation of M1 is then given by

$$\theta_s = \Delta\theta_s + \delta \quad \text{(IV.9)}$$

where $\Delta\theta_s$ is the initial phase offset. We also extract information from the M2 encoder indicating the large motor index when the large motor completes a revolution. The M2 index and the M1 count (n_s) maintain a constant phase relationship after the scan motions have stabilized. Therefore, in post processing the number of small mirror rotations can be tracked and used to calculate the position of the large mirror. That relationship is given by

$$\theta_L = \Delta\theta_L + \frac{(\delta + 2\pi \cdot N_s)}{30} \quad \text{(IV.10)}$$

where $\Delta\theta_L$ is the initial phase offset of the large mirror, and N_s is the number of small mirror rotations. Once $\Delta\theta_s$ and $\Delta\theta_L$ are determined, the values for θ_s and θ_L are known and the beam directions can be calculated according to the formulation in equations (IV.1-7).

To determine the values of $\Delta\theta_s$ and $\Delta\theta_L$ for a particular data collection, we need to start with some angular calibration information. Angular calibration is performed using a hard-target

collection that includes a sharp edge with known orientation. A frame of this data is processed with initial assumptions for $\Delta\theta_s$ and $\Delta\theta_L$(e.g. zero). Those values are adjusted until the sharp, straight-line edge is seen with the correct orientation. The value for $\Delta\theta_s$ is then known for all subsequent data collections.

The value for $\Delta\theta_L$may vary for each data collection because synchronization of M2 is not achieved immediately upon motor startup. But the M2 index signal provides the needed information. When the M2 index occurs, the large motor has an fixed phase. This phase can be calculated using equation (IV.10) at the moment of the M2 index (i.e. from the angular calibration collection). For subsequent collections, the value for $\Delta\theta_s$ can be used to determine θ_s according to eqn. (IV.9) and $\Delta\theta_L$can be determined according to

$$\Delta\theta_L = \Theta_L - \frac{(\delta_0 + 2\pi \cdot N_{s0})}{30} \qquad \text{(IV.11)}$$

where Θ_L is the fixed phase offset at the M2 index, δ_0 is the value of δ at the large motor index, and N_{s0} is N_s at the M2 index.

Hybrid Cartesian Mapping

The scan-angle transformation described above and the range/time-delay proportionality map each sample point in each lidar waveform into object space. However, the angular mappings do not represent conventional images. Consider the initial WiLD image shown in Section V on the left of side of Figure 30. The image sampling is inhomogeneous and discontinuous, with samples concentrated on segments of the arcuate scanning path.

For purposes of visualization and analysis it would be ideal to further convert the WiLD data to a conventional Cartesian form. However, the scanned volume has a conical shape that does not fit well into a rectangular prism as would be expected for a conventional Cartesian volume image. We choose instead to construct a hybrid image format with coordinates of azimuth angle, elevation angle, and range distance. Because the coordinates are locally orthogonal, many kinds of image analysis can be performed directly on these volume images. Note that the system FOV is a circle for azimuth/elevation image slices (perpendicular to the LOS) and a rectangle for slices parallel to the LOS.

Transformation of a WiLD frame into a hybrid Cartesian volume image proceeds in two stages: 1) preprocessing to associate waveforms with the grid, a one-time operation; and 2) accumulation of waveforms to construct each image frame.

In the preprocessing step, the angular grid is defined by altitude and azimuth with grid spacings of 1° in both directions. Using the techniques described above, altitude and azimuth are computed for each of the laser pulses in the first 1-second scan. The number of pulses per frame is nominally N=50,000. However, the motor encoder clock may not exactly match the trigger pulse generator. Therefore N is measured by counting samples between M2 index pulses. As discussed in Section II, the scan angles repeat to excellent precision from frame to frame. In

order to minimize global distortion associated with the angular coordinates, the coordinate system is tilted back 30.8° so that the center of the FOV is near 0° elevation.

Each lidar waveform is associated with up to 8 surrounding pixels as illustrated in Figure 27. (Throughout this discussion, "pixel" refers to an angular interval in azimuth and elevation, as opposed to a voxel that is also extended over a range interval.) A heuristic maximum distance function, δ_{max}, is used to decide whether each of these pixels is a "near neighbor" with respect to the distance δ from the scan angle to center of each pixel. It is desirable to associate approximately the same number of waveforms with each pixel, but the density of samples is somewhat concentrated at the center and edges of the hypertrochoid scan (c.f. Figure 6). Therefore the distance function δ_{max} is made smaller at the center and edges of the scan pattern and larger in the intermediate annulus.

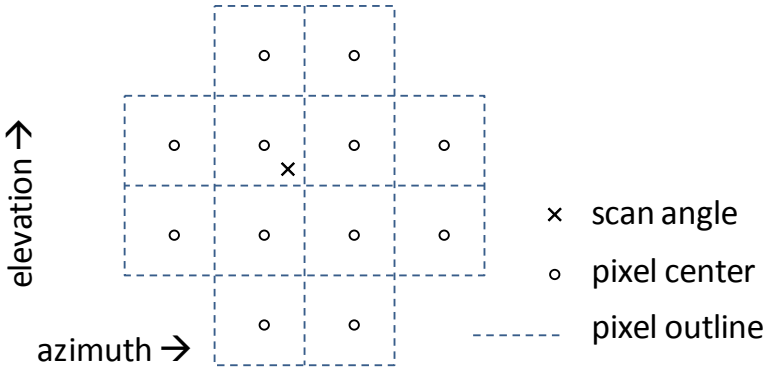

Figure 27. Association of a waveform scan angle with surrounding pixels.

On average, each waveform is associated with 6 pixels. Conversely, most pixels are associated with 100 - 200 waveforms. Because each voxel value combines information from many lidar scans, the SNR of the image is much higher than the individual waveforms. Figure 28 shows the degree of uniformity in the count of waveforms associated as near neighbors with each pixel. The outline of the scan in this figure is much more symmetric than Figure 6 because of the 30.8° coordinate rotation.

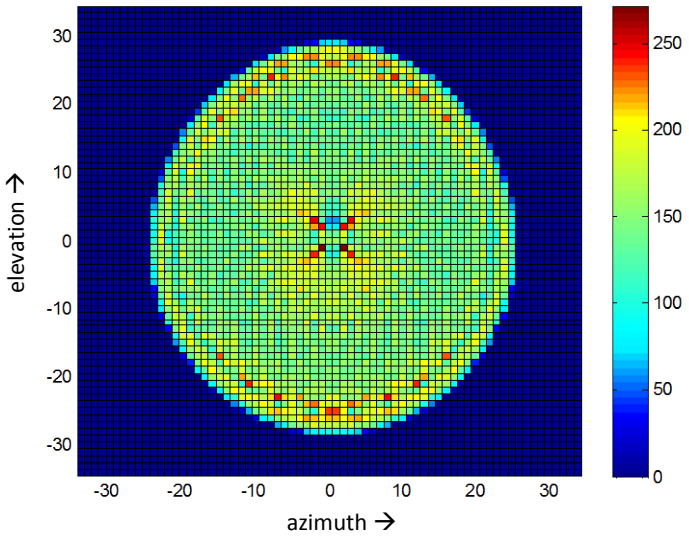

Figure 28. Number of waveforms contributing to each pixel.

By combining multiple waveforms into a single pixel, the image construction process reduces the intrinsic spatial resolution. Figure 29 plots the rms spread of the waveform angles contributing to each pixel. The typical spread is approximately 1° in both azimuth and elevation. Thus the aggregated pixel construction does not substantially degrade the spatial resolution of the grid.

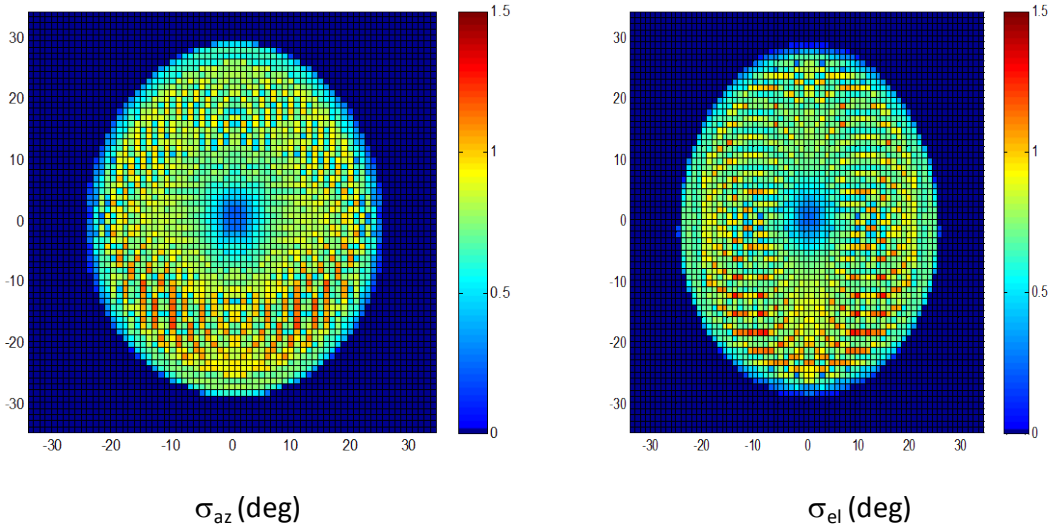

Figure 29. The rms spread of waveforms contributing to each pixel.

It is desirable to weight the contributing waveforms for each pixel to eliminate spatial biases, i.e. due to an off-center distribution of the contributing waveforms with respect to the pixel. The average offset for the set of waveforms contributing to a pixel is typically ±0.25°, but up to 1° at the edges of the FOV. The weighting that provides a least-squares estimate for the pixel-centered image value is given by the first row of

$$\begin{bmatrix} count & \sum \delta_{xi} & \sum \delta_{yi} \\ \sum \delta_{xi} & \sum \delta_{xi}^2 & \sum \delta_{xi}\delta_{yi} \\ \sum \delta_{yi} & \sum \delta_{xi}\delta_{yi} & \sum \delta_{xi}^2 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 1 & \dots \\ \delta_{x1} & \delta_{x2} & \dots \\ \delta_{y1} & \delta_{y2} & \dots \end{bmatrix} \quad \text{(IV.12)}$$

where δ_{xi} and δ_{yi} are the azimuth and elevation offsets, respectively, for the i^{th} contributing waveform. Like the waveform/pixel assignments, these weight evaluations need only be performed once for a sequence of WiLD image frames.

Preprocessing of waveform/pixel assignments and weights takes approximately 1 s in a MatLab implementation. Once preprocessing is complete, each frame of the volume image is created in real-time (<< 1 s per frame) as follows.

a) Trigger the image framing process based on each 30^{th} reset of the M1 encoder or the reset of the M2 encoder.

b) Clear the volume image. The volume image consists of a range vector, S(r), for each pixels.

c) Clear the waveform count, n_{count}, for all pixels. The waveform count indicates how many waveforms have been added into the corresponding range vector.

d) Receive a lidar waveform s(r) with a corresponding pulse count n (i.e. 1:50,000).

e) Precondition the lidar waveform, applying range offsets, range scaling, background subtraction, spatial filtering, etc. (See the discussion in Section III.) All of these operations should be structured as linear operations if possible.

f) For each pixel in the n^{th} mapping set, increment n_{count} and accumulate the range vector.

$$S_n(r) \leftarrow S_n(r) + W_n(n_{count}) \cdot s(r) \quad \text{(IV.13)}$$

g) Iterate to step (d) until all waveforms in the frame have been accumulated.

h) Store the completed frame and iterate to step (a) until all frames of the volume image have been processed.

For waveforms extending 200 samples beyond the laser pre-pulse (300 m range at 100 MSPS), the completed volume image includes 70x70x200 = 980,000 voxels. Because the FOV is circular, the actual uncompressed storage requirements are 470 kBytes at 1 byte per voxel.

V. 3D Lidar Observations.

The 3D collection mode of the WiLD sensor was described at the end of section II and construction of volume images from scanned waveforms was discussed in section IV. Here we describe the experimental implementation of the 3D collection mode and present some resulting volume images.

Solar Interference

Initial attempts at volume imaging resulted in burn damage to the conical barrel situated between the objective lens and the APD housing. The original barrel was fabricated of black flocked paper to minimize stray light while accommodating alignment of the APD. The objective lens concentrates sunlight onto this barrel surface when the sun is present within about 45° of the sensor LOS. But in 3D collection mode, the LOS scans a 60° FOV. If the sun is located within 75° of the scan center, there is a risk of thermal damage from sunlight. Such damage is not limited to burning of the barrel – any burning creates smoke that can contaminate the optical surfaces. Subsequently the paper barrel was replaced by a black-anodized aluminum part that is not damaged by concentrated sunlight. The new barrel has an oversized opening that fits around the barrel of the relay optics to allow for fine alignment of the receiver.

The large FOV of WiLD also makes it desirable to ensure that the system can tolerate a direct view of the sun. The solar radiance within the filter bandwidth is $2.8 \cdot 10^4$ W/m^2/sr. For an optical throughput of 79%, the detector flux for unobscured solar viewing is 400 μW for a scan duration of 0.4 ms. This flux is less than the maximum survival rating for the APD (2 mW). But it is much higher than the saturation level (approx. 3 μW). The recovery time for nominal overloads is rated at 150 ns. But we have seen that secondary ringing effects last much longer and a severe overload may take much longer to recover, especially if it causes the HV bias supply to sag. This effect should be investigated experimentally to establish the nature of a solar transient and to identify how it can be handled in processing of 3D collections.

Ladar images

The first scanned data from WiLD was collected on 3/15/10. The system was set up immediately outside the front door of B620 (620 East 1600 North, Logan, Utah) directed generally eastward toward the USU BioInnovations building (then under construction).

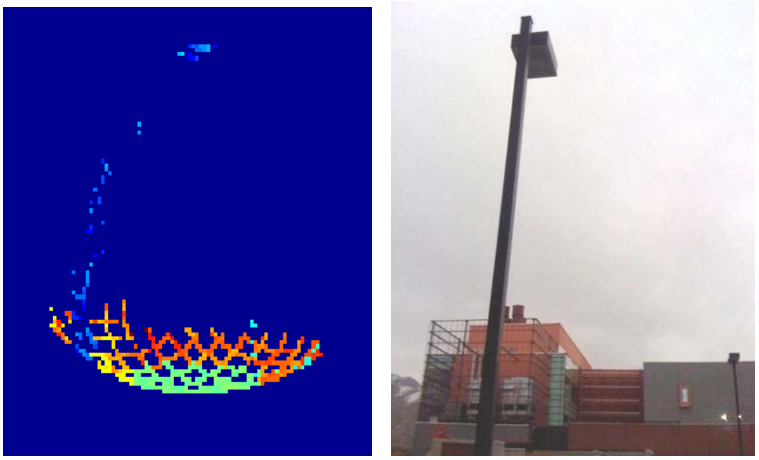

Figure 30. First scanned data collection from WiLD.

The lidar data was processed for scan angle using manual derivations of the mirror angle offsets and each waveform was evaluated for a strong (hard-target) return following the laser pre-pulse.

The left-hand image in Figure 30 is a scatter plot showing the color-coded range for these returns. (No attempt was made at this time to aggregate the waveforms or smooth the image.) This sort of hard-target detection is a form of Ladar (Laser Radar), a term that usually refers to the application of lidar techniques for surface ranging and profiling. The WiLD lidar image clearly identifies a light pole in the foreground and sections of the BioInnovations Building that are visible in the photograph on the right.

Another set of scanned data was collected on 29 July 2010, with WiLD located outside at the south end of B620, directed toward the southwest. Ten frames (1 s each) were collected and processed for scan angle and target range. Figure 31 is a scatter plot from one of the frames, showing color coded elevation vs the (projected) horizontal location of the target. Low elevation returns (i.e. from the ground) are not shown. The (empty) 600 East roadway is clearly visible, as is the outline of a building across the street. Other concentrations of laser returns represent small trees. The sequence of frames from this collection show the movement of people (isolated returns) and the passage of a vehicle in the roadway.

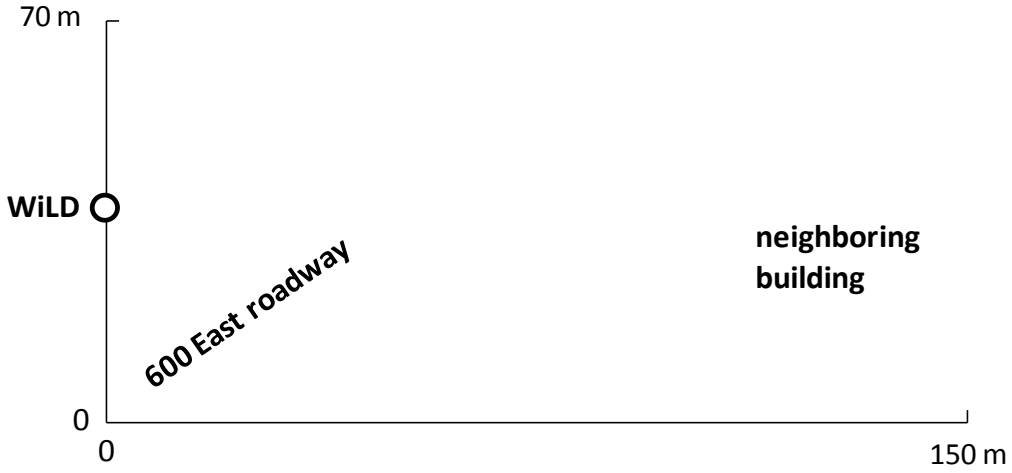

Figure 31. Object ranging from a WiLD scan.

Aerosol Plumes

The first volume images of an aerosol plume were captured by WiLD on 15 March 2010. The sensor was situated near the south driveway of 620 East 1600 North, Logan, UT, pointed toward the west. The lidar waveforms were processed as described in section IV to create consecutive volume images (4 frames). Images slices at close range from the first two frames (Δt = 1 s) are illustrated in Figure 32. An aerosol plume with a strong central concentration at 27 m range appears in frame #1. In frame #2, the plume has dispersed a bit and drifted westward approximately 1m. The plume likely represent exhaust from a passing vehicle, although this was not noted at the time. The strong, localized return at the bottom of both frames at 30 m range is a building sign situated near the road.

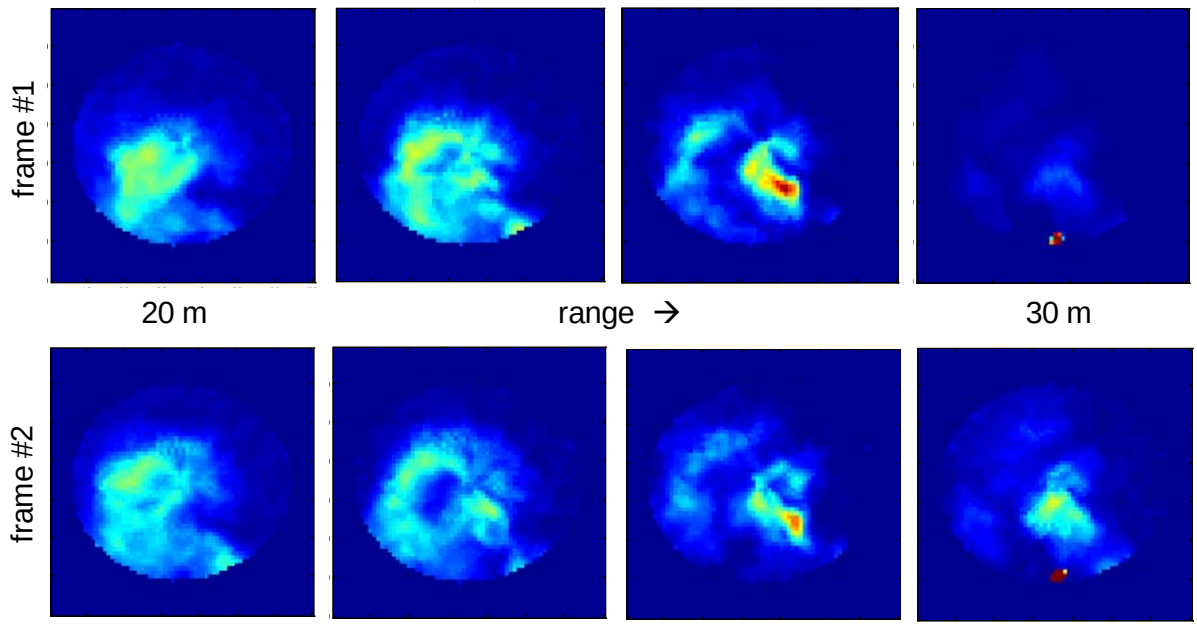

Figure 32. Close-range aerosol plume.

In order to demonstrate the dynamic 3D aerosol imaging capabilities of WiLD, we set up idealized conditions using artificially generated smoke plumes. Smoke plumes were created using smoke signal flares. The smoke flares produce a stream of white, magnesium-oxide smoke for a duration of about 30 s. The flare was lit at a location approximately 200 m upwind and 200 m down-range from the WiLD system. A 3D scanned lidar collection was initiated after smoke from the flare had begun to cross the sensor FOV.

The first smoke plume collection occurred on 10/20/2010. The WiLD system was set up immediately north of B1750 on the USU Innovation Campus, with WiLD directed over the 600 East roadway toward the east. The smoke plume was generated in the large, empty agricultural field across the road. The lidar data was processed as described in section IV. Figure 33 shows several images that illustrate the motion and evolution of the smoke plume. The rows in Figure 33 represent three image slices at ranges of 150, 160, and 175 m. The columns represent four consecutive volume-image frames (at 1 s intervals). A sharp cutoff at the lower left resulted from obstruction by a tree trunk near the WiLD system. The elevation scales in Figure 33 (70 is near horizontal while 10 represents an elevation of 60°). At 160 m range, an angular spread of 10° corresponds to a width of 30 m, which is also comparable to the depth of the smoke-plume. The evolving detail in these images illustrates the spatial resolution of WiLD - approximately 7m range and 2° direction for situations like this where the target signal is strong.

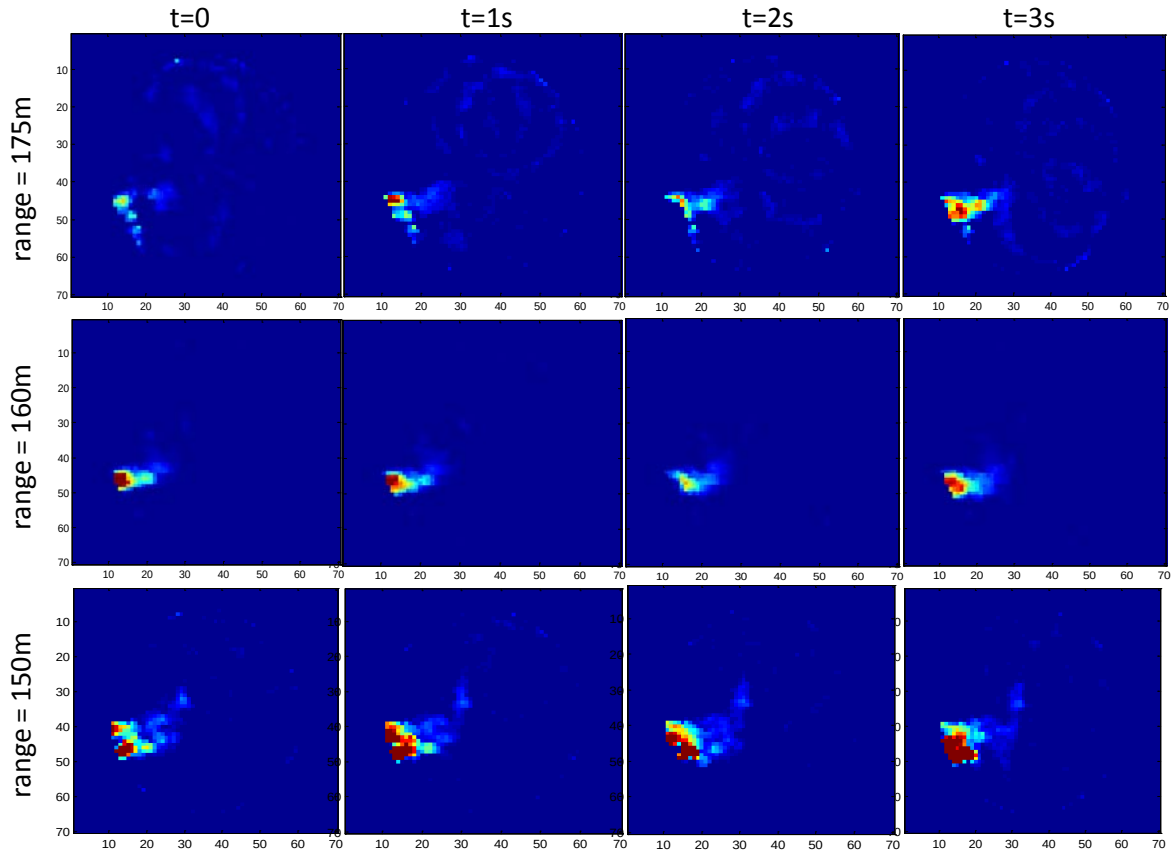

Figure 33. Motion and evolving shape of a smoke plume on 10/20/2010.

Another set of smoke plume collections were performed on 5 April 2011 at the same location and under similar conditions. No plume was captured in the first collection, but the second collection was successful. Figure 34 is an embedded 10-frame movie of the smoke plume dynamics. The plume passed through the lower section of the FOV, so the movie shows four horizontal slices through the FOV, 7° to 10° in 1° increments. Range in these movie slices extends from left to right (0 to 300 m). The vertical axis in these slices corresponds to azimuth (up is left / north). The lower slice is narrow because it lies near the bottom of the (approximately circular) FOV.

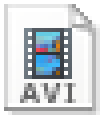

plumeMovie.avi

Figure 34. Smoke plume movie with 4 horizontal image slices (click to activate).

The plume movie of Figure 34 shows a dominant smoke cloud and several weaker segments of the smoke plume blowing across the FOV from the south at a range of 150 - 180 m.

Low Frequency Fluctuation

Close inspection of the 175 m range images in Figure 33 shows faint circular features. The circle diameter is 14°, corresponding to the M1 scan angle. The circles are caused by small fluctuations in the APD output at very low frequency, < 20 Hz. The fluctuation amplitude is large enough to interfere with imaging of faint aerosol features. And because the pattern changes from frame to frame it would interfere with motion detection.

Further study is needed to characterize these low frequency fluctuations and identify the source. If the fluctuation is an ultra-low-frequency bias, it can be filtered out before processing of the volume images. But if it is a range-dependent transient modulated at ultra-low-frequency, then it must be eliminated at the source.

We have noted the possibility that contamination of the scan mirrors (by dust and/or condensation) may modulate the amplitude of the laser pre-pulse signal and the subsequent ring-down phenomenon. This is probably not the cause of the observed low-frequency fluctuation, because contamination-induced artifacts would repeat each second. Another possible cause to be considered is EMI from the scan motors, controllers, or other WiLD components.

Failure of the APD

Following the 3D collection shown in Figure 34, several additional smoke plume collections were attempted. But that subsequent data did not show any lidar signatures. Initially it was thought that the laser trigger had failed. But subsequent testing showed that the APD module was no longer responsive. This put an end to the final data collections planned for DOE contract DE-EE0002734.

The APD module in the WiLD detector box has since been replaced and at this writing the system is fully functional. However the cause of failure is not known. From the solar flux analysis at the beginning of this section, we see that the safety factor for direct solar viewing is approximately 5x. Therefore accidental sunlight viewing is unlikely to have caused the failure. However, this failure mode cannot be completely ruled out.

VI. Velocity Field Analysis.

Velocity Analysis Methods

Early in the project, the team investigated and developed a variety algorithms for analysis of dynamic volume imagery to derive 3D wind fields.

a) Cross-Correlation Method. This was method was the initial concept considered for wind velocity analysis using dynamic volume imagery. Consider a vector offset applied to an initial image. The offset vector that maximizes the correlation between the shifted initial image and the next (unshifted) image is a regional fit to the velocity field (multiplied by the time interval).

b) Semblance Method. This method is similar to method (a), but with a semblance function taking the place of a correlation function. Semblance is a generalization of cross-correlation that weighs the correspondence between absolute image amplitude, not just scaled similarity.

c) Translational Phase Shift. A pure translation leaves the Fourier Transform (FT) of an image unchanged except for a constant phase shift (scaled by $1/\lambda$). The Translational Phase Shift method compares the FT of consecutive images and finds the phase-shift that creates a least-squares best fit.

Methods (a) through (c) depend on an assumption of localization. That is, they require many separate evaluations to evaluate a non-homogeneous velocity field. This can be accomplished by appropriate segmentation of the source imagery. But this will not be effective unless the images have sufficient spatial resolution and significant image detail in each segmented region. Phenomenological data collected in this study suggests that volume images of aerosol concentrations are most likely to be relatively sparse, so these methods will not be effective.

d) Translational Mapping Method. Christian Marchant developed this global adaptation of the Translational Phase Shift method for an inhomogeneous velocity field. The Translational Mapping method constructs a (non-affine) transformation function that when applied to a first image matches its FT most closely to the FT of the next image. The algorithm starts with the lowest frequency transformation component and adds successively higher frequencies.

Two experimental results using the Translational Mapping method are illustrated in Figure 35. Both examples (a and b) show consecutive photographs of clouds with the derived 2D velocity fields superimposed. This is an distinct improvement over method (c) with respect to efficiency and the required sampling frequency, but it still has limitations with respect to sparse imagery.

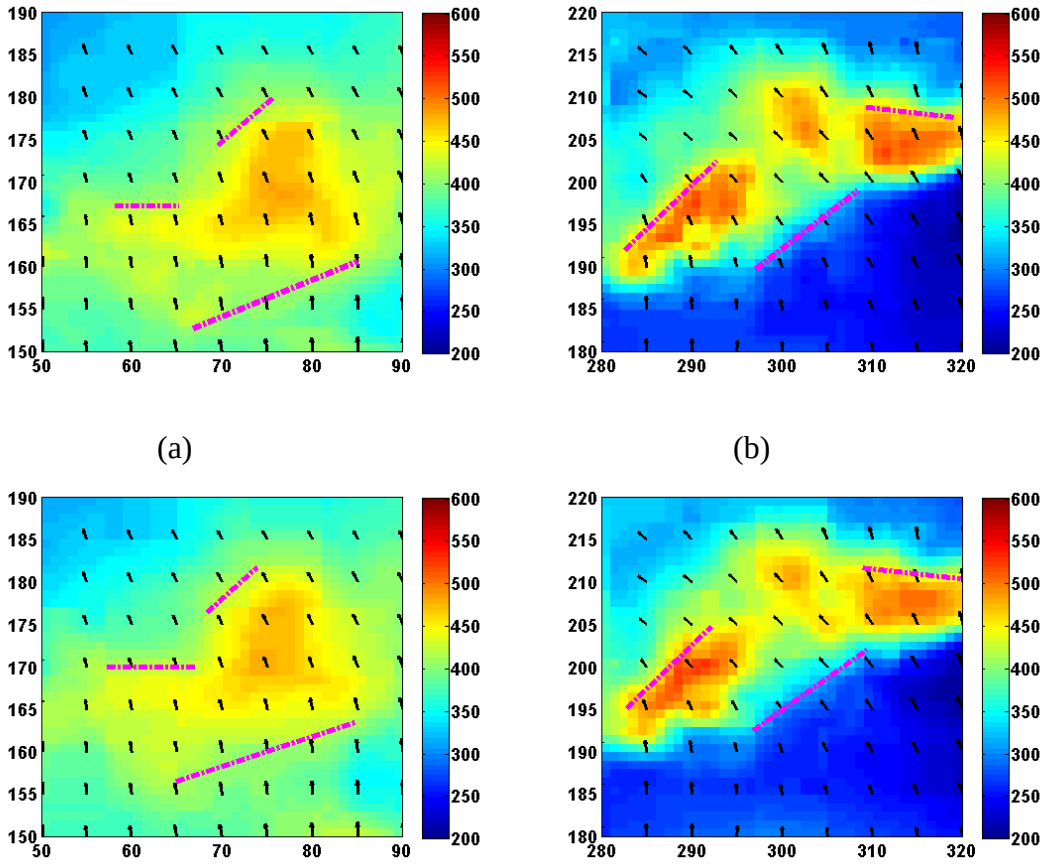

Figure 35. 2D wind field results from the Translational Mapping Method.

e) Lucas-Kanade Optical Flow. For a density function ρ tied to an incompressible fluid (a good approximation for air), the velocity field ū relates the spatial gradient of ρ to its temporal derivative through

$$\overline{\nabla}\rho \bullet \overline{u} + \rho_t = 0. \qquad \text{(VI.1)}$$

In a neighborhood over which the velocity is well behaved, a least-squares velocity estimate can be found by inverting the equation

$$\mathbf{M} \cdot \overline{u}^T = -\mathbf{b} \qquad \text{(VI.2)}$$

where **M** is a cross-correlation matrix of spatial derivatives and **b** is a vector of spatio-temporal cross-correlations:

$$\mathbf{M} = \begin{bmatrix} \sum g \cdot \rho_x^2 & \sum g \cdot \rho_x \cdot \rho_y & \sum g \cdot \rho_x \cdot \rho_z \\ \sum g \cdot \rho_x \cdot \rho_y & \sum g \cdot \rho_y^2 & \sum g \cdot \rho_y \cdot \rho_z \\ \sum g \cdot \rho_x \cdot \rho_z & \sum g \cdot \rho_y \cdot \rho_z & \sum g \cdot \rho_z^2 \end{bmatrix} \text{ and} \qquad \text{(VI.3)}$$

$$\mathbf{b} = \begin{bmatrix} \sum g \cdot \rho_x \cdot \rho_t \\ \sum g \cdot \rho_y \cdot \rho_t \\ \sum g \cdot \rho_z \cdot \rho_t \end{bmatrix} \qquad \text{(VI.4)}$$

where subscripts represent partial deriviatives with respect to the indicated coordinates, sums are taken over the neighborhood points, and g are (optional) weights at each point. This is the basis of the Lucas-Kanade optical flow method. The spatial coordinates can be any locally orthogonal set, i.e. azimuth, elevation, and range in the context of the WiLD scanner. Note that the image gradients should be estimated such that they are consistently localized, e.g. best estimates for the center of each voxel, and midway between the two consecutive images.

Like methods (a) through (c), the Lucas-Kanade method depends on an assumption of localization. It requires relatively dense sampling and non-sparse image detail. Regional estimate of ū using the Lucas-Kanade method is biased if the frame-to-frame motion is not small with respect to the scale of spatial features in ρ. But the bias problem is readily addressed by iteration: given an initial estimate of the velocity, form a "warped" image by backing out the approximate motion from the second image; then re-apply the Lucas-Kanade method to update the velocity estimate.

f) Hierarchical Lucas-Kanade Method. The Lucas-Kanade method can be extended to sparse imagery through hierarchical analysis. The following steps are suggested.

- Estimate the velocity vector over a coarse grid, i.e. with the volume imagery low-pass filtered to a sufficient degree that it sparsity is not a problem.
- Subdivide the image and apply the Lucas-Kanade method to each sub-region.
- Evaluate the rank and condition of the **M** matrix on each sub-region. If rank(M) is less than 3, equation VI.2 cannot be inverted. If cond(M) is large, numerical errors in b are magnified. (We suggest a threshold of cond(M) < 20.) In regions with insufficient rank or excessive condition number, the velocity field is interpolated rather than evaluated using the Lucas-Kanade method.
- For regions where the Lucas-Kanade method is successful, iterate until the desired spatial resolution is achieved.

g) Horn-Schunck Optical Flow. The Horn-Schunck method [2] is an iterative, non-localized approach to Optical Flow analysis. The velocity field is iterated globally as follows:

$$\bar{u}' = \bar{u} - \bar{\nabla}\rho \cdot \frac{\left(\nabla\rho \bullet \bar{u} + \rho_t\right)}{\left(\alpha^2 + |\nabla\rho|^2\right)} = 0 \,. \qquad \text{(VI.5)}$$

The parameter α^2 is a smoothing parameter that enforces minimization of velocity gradients in regions where there is no image structure ($|\nabla\rho|^2 \sim 0$). Iteration proceeds until convergence is observed. For a volume image with 20x20x20 voxels, 40 or more iterations are required initially. For dynamic flow evaluations, the velocity field from the previous time step may be

used as an initial guess, greatly speeding up convergence. As discussed below, we were successful in applying the Horn-Schunck method for 3D velocity retrieval from volume images. Additional analysis steps included:

- Apply a low-pass spatial filter to each volume image frame.
- Apply offset and normalization corrections so that each frame has zero mean and unit variance.
- Estimate spatial image gradients on the spatial grid (the MatLab grad function does this by default) and the temporal gradients at the midpoint between consecutive frames.
- Set the smoothing parameter to α^2=0.1 (for unit spacing on the spatial grid) when image features are much larger than the image grid scale. Larger values are needed when the features are relatively small.
- Iterate until the rms change is less than 0.1% of the spatial interval divided by the time step.

Temporal aliasing is a common problem in motion estimation. It is a necessary consideration for all of the analysis methods that we have considered. For reliable convergence, the dominant spatial scales in the aerosol image must be large compared to the motion between frames. That is, the feature scale must be much larger than $v_{wind} \cdot \Delta t$. For WiLD observations (Δt = 1 s) at wind energy sites (v ~ 10 m/s), the aerosol features of interest must be larger than 20 m. This is consistent with the resolution of the WiLD sensor with substantial low-pass filtering. However, observations to date have not shown ambient aerosol features that are clearly large enough.

Simulated Wind Fields

Since field experiments with WiLD have not yet yielded usable, wide-scale volume images, we turned to simulation for evaluation and validation of the wind analysis methods. The first simulation step is creation of realistic 3D wind fields.

We model the atmosphere as an incompressible fluid, i.e. with a zero-divergence constraint $\nabla \bullet \bar{u}=0$. The wind field is constructed by summing a set of wave components

$$\bar{v} = \sum_n \bar{v}_n \exp\left(i\left[\bar{k}_n \bullet \bar{x} - \varpi_n t\right]\right). \qquad \text{(VI.6)}$$

The incompressibility condition must apply to each wave component individually. Consequently, each wave vector must be perpendicular to the corresponding peak velocity vector

$$\bar{k}_n \bullet \bar{v}_n = 0. \qquad \text{(VI.7)}$$

The simulated wind field is constructed by randomly choosing five wave components with uniformly distributed amplitudes $|v_n|$, frequencies ω_n, and directions for $\mathbf{v}_n$ and $\mathbf{k}_n$, subject to constraint VI.7. The vertical wind component v_z was scaled down relative to the horizontal components to imitate the usual dominance of horizontal wind. Figure 36 illustrates the synthetic wind field that will be used throughout the rest of this section. The images in this figure show the magnitude of the wind vector components on a plane at z = 10 and t = 15.

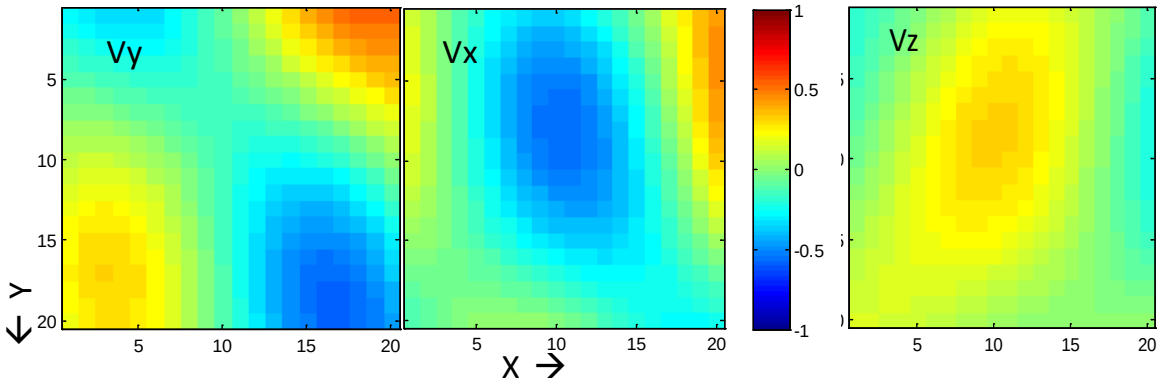

Figure 36. Slice through a simulated velocity field comprised of five wave components.

Simulated Volume Images

Simulated dynamic volume images (21 frames) were created from the wind field of Figure 36 by placing tracer "particles" in four columns. These tracers represent aerosols that are visible to the sensor and drift with the wind. The tracer particles at time t=0 are shown in Figure 37, with a vertical view on the left and a view down the Y axis on the right. Particle trajectories are calculated by Runge-Kutta integration of the wind field. For enhanced visualization, the particles are color-coded with respect to depth position (z on the left and y on the right). Embedded movies show how the plumes are displaced, rotated, and distorted by the dynamic wind field.

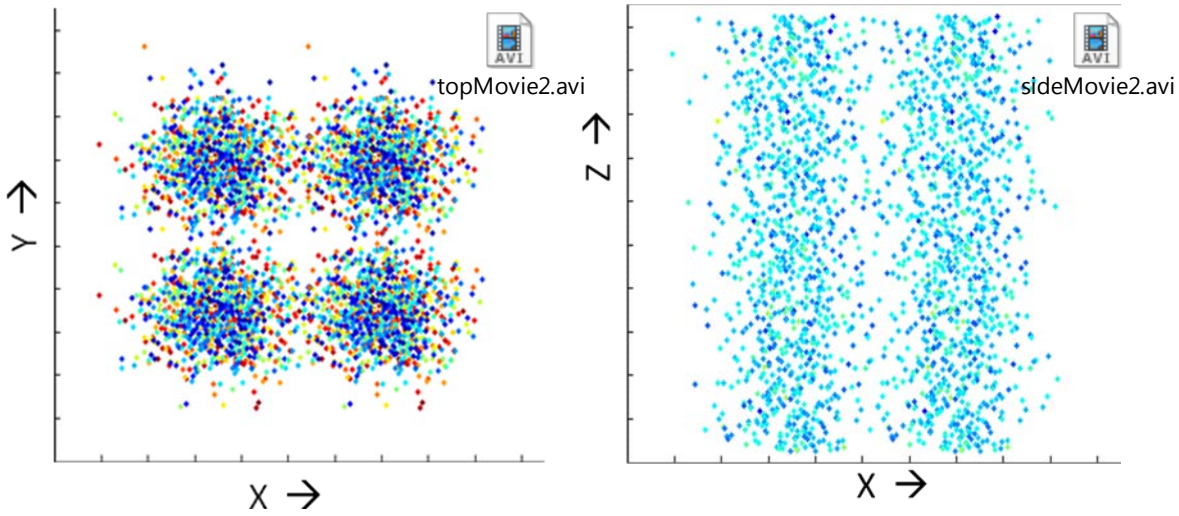

Figure 37. Tracer particles permit visualization of the velocity field. Click on embedded file icons to activate movies.

Aerosol volume images represent not individual tracer particles, but 3D concentrations of aerosol particles. The tracer patterns shown in Figure 37 were converted into dynamic volume imagery by counting the number of particles on a coarse, 3D spatial grid at each time increment. Slices

through the raw simulated imagery at z = 19 with t=14 and t=16 are shown in Figure 38. Corresponding smoothed images (pre-processing in preparation for optical flow analysis) are shown in Figure 39. Plume motions are qualitatively evident in Figures 38 and 39, although quantitative wind speeds would be hard to estimate through visual assessment.

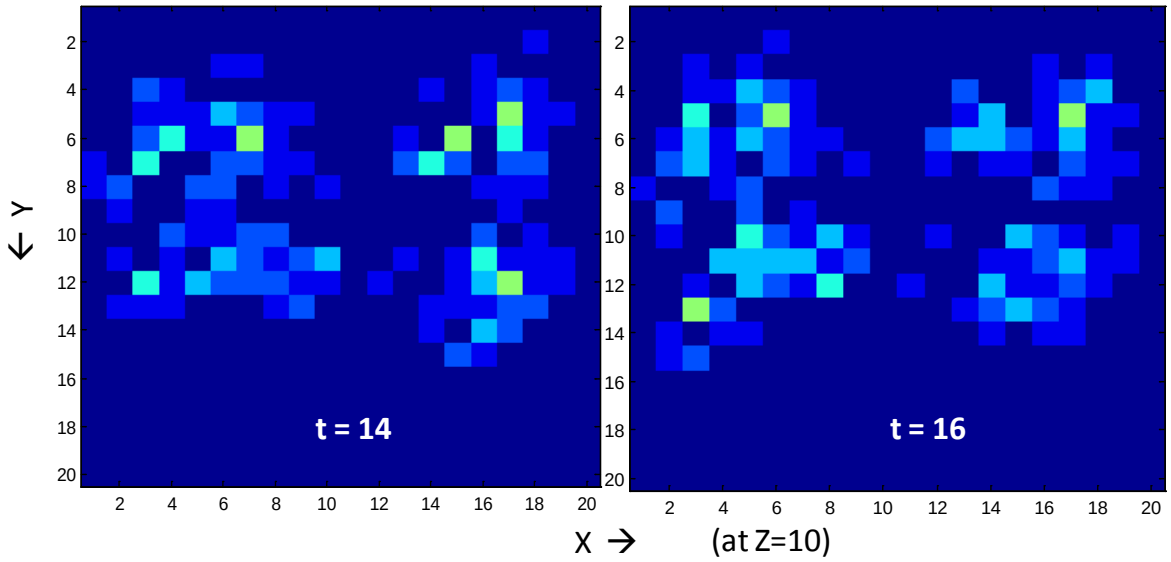

Figure 38. Raw volume imagery derived from the tracer paths (c.f. Figure 37).

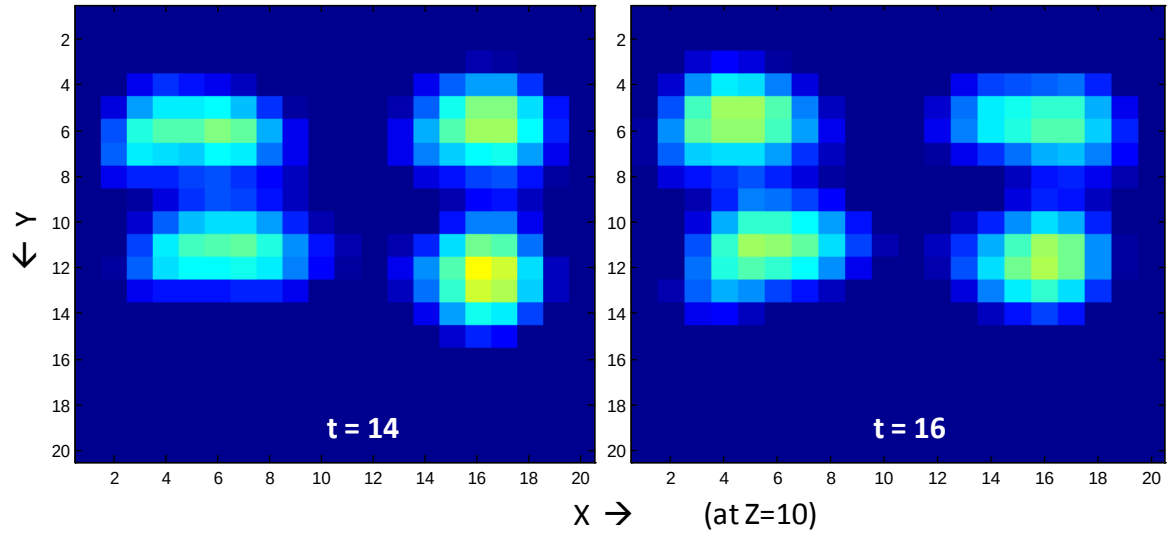

Figure 39. Volume imagery smoothed in preparation for optical flow analysis (c.f. Figure 38).

Horn-Schunk Method Results

Attempts to analyze the underlying wind field from the simulated imagery using the Lucas-Kanade optical flow method were not successful. Small analysis subregions, for which the method is efficient, do not respond to the relatively large moving image features. If the analysis

subregions are increased sufficiently in size, then the matrix inversion that is required for the Lucas-Kanade method is too slow.

The Horn-Schunk optical flow method proved to be effective. For a 20x20x20 image size, the calculation time was ~0.4 s per frame - faster than the collection rate for WiLD. Examples of accurate velocity field reconstruction are shown in Figures 40 and 41.

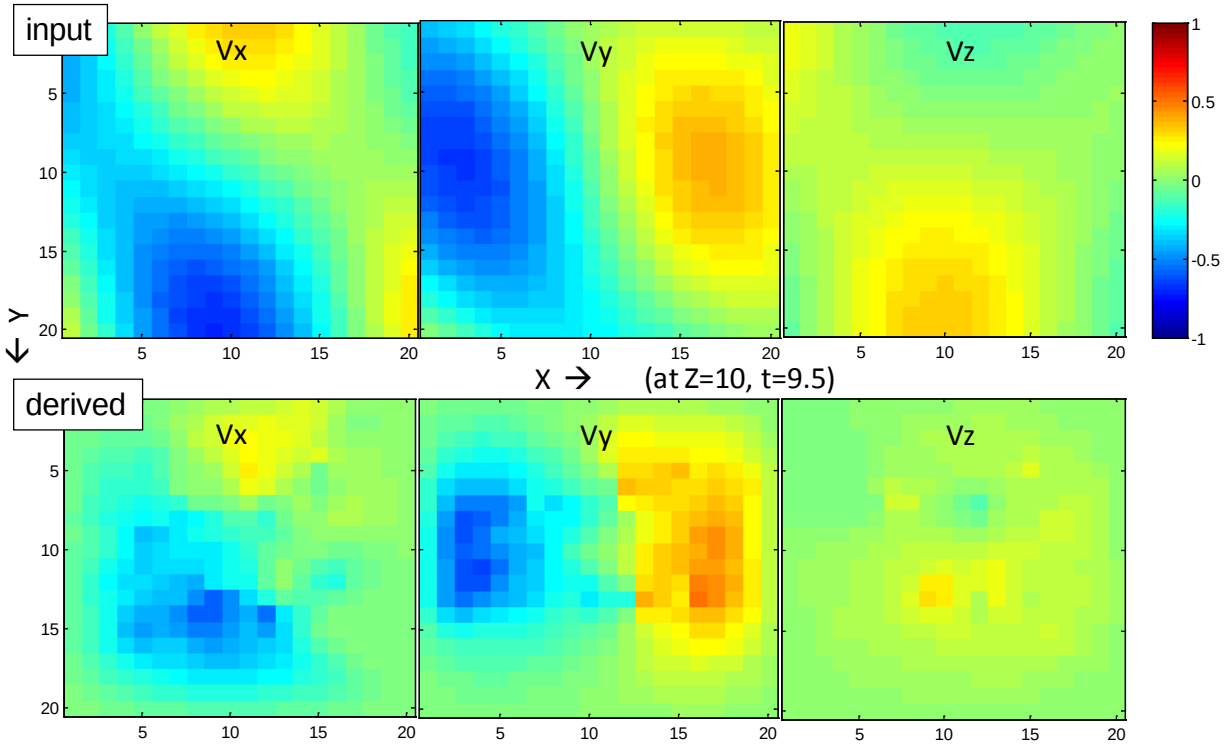

Figure 40. Velocity field reconstruction at z = 10 and t = 9.5 (using frames 9 and 10).

The examples of Figures 40 and 41 demonstrate that the Horn-Schunk method is effective for 3D velocity fields and results in quantitatively accurate velocities. Some of the input features are not reproduced in the derived velocity fields; these are features that do not intersect the tracer plumes. That is, a non-homogeneous vector field cannot be completely derived from sparse image data. Even so, the Horn-Schunck method preserves smoothness in the velocity field and faithfully reproduces the dominant characteristics of the velocity field.

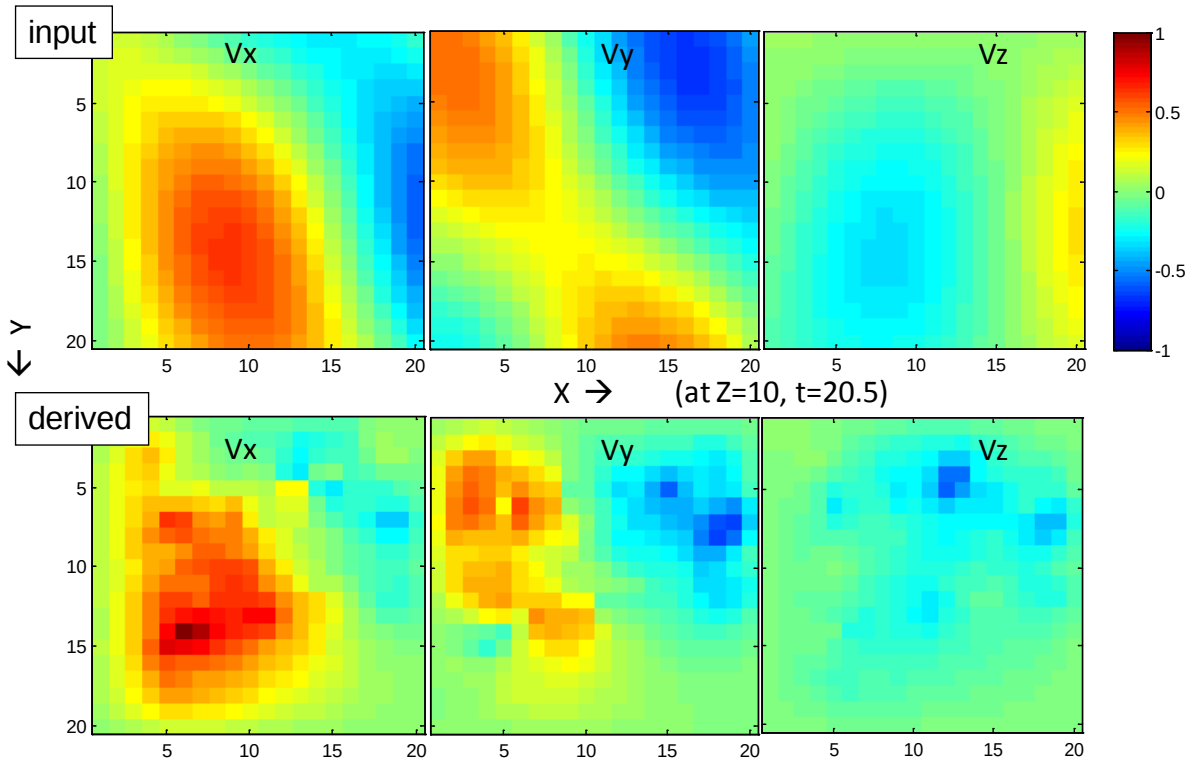

Figure 41. Velocity field reconstruction at z = 10 and t = 20.5 (using frames 20 and 21).

VII. Laser Safety Considerations.

A principle design criterion for the WiLD system is lidar operation in the "retina-safe" wavelength range between about 1.4 and 2 μm. But just because a laser beam cannot penetrate the eye does not make it "safe." Safety analysis and corresponding operational constraints are needed before a lidar system can be operated without risk to operators and the public.

Safe operation of the pulsed retina-safe laser for WiLD is achieved by reducing human laser radiation exposure to a Class 1 hazard level. By reducing the dwell time of the beam, the Class 1 hazard level can be reached without diminishing instantaneous irradiance or pulse energy. Dwell time reduction can be accomplished by either reducing the duty cycle or scanning the laser. Reducing the duty cycle is an effective means of reducing the laser hazard class for laboratory tests, while scanning the beam is useful for reducing the hazard class for field measurements. Furthermore, beam scanning is implemented in WiLD for collecting atmospheric data over a large area. However, it can be a challenge to establish the hazard class of a scanned laser.

The first step of our safety analysis is the determination of the accessible emission limit, AEL, which is a product of the maximum permissible exposure (MPE) and the area of the limiting aperture. The AEL can be used to determine the number of pulses allowed during the exposure duration for Class 1 operation.

From Table 5a of [3], we find the maximum permissible exposure for our laser, MPE = 1 J/cm^2. The limiting aperture diameter D_M (defined as, "the diameter of a circle over which irradiance or radiant exposure is averaged for purposes of hazard evaluation and classification") is given [ref. 3, table 8a] as D_M=3.5 mm in the wavelength range 1.4 - 100 μm. For a retina-safe laser, corneal damage is a concern, and for hazard calculations the limiting aperture for skin is the same as that for the eye.

A critical parameter to consider in the safety hazard assessment is the limiting exposure duration, T_{max} which can be determined by the laser safety officer. It represents the maximum total direct exposure time to the laser. If the value for T_{max} for a given system is unclear, an upper limit for T_{max} for some spectral ranges can be found in [3]. For the spectral region which includes 1550 nm, we can consult example 25 in ref. [3]. In determining the MPE, we find from table 5a of [3] that for exposures 10^{-9} s to 10 s a value of 1 J/cm^2 can be used, and for exposures from 10 s to 30000 s an MPE of 0.1 W/cm^2 can be used. For an exposure time of 10 s and an MPE value of 1 J/cm^2, the AEL will be

$$AEL = MPE \times \pi \left(\frac{D_M}{2}\right)^2 = 96.2\ mJ. \qquad \text{(VII.1)}$$

Here we have used a limiting aperture diameter, D_M = 3.5 mm from table 8a of [3] as required by Rule 2, the average power limit rule. Rule 2 yields the lowest AEL of the 3 Rules presented in ANSI. For any exposure duration from 10 s to 30000 s, the Class 1 AEL will be 9.62 mW. We do not anticipate any exposure longer than 30000 s, and since the result will be the same, we use the most conservative T_{max} for an MPE of 1 J/cm^2, 10 s.

The Class 1 hazard level is therefore 96.2 mJ. To solve for the number of pulses permitted through the limiting aperture in the interval T_{max}, we apply the following equation from Example 25 of [3].

$$\phi_d = n\phi_o \left[1 - e^{-\left(\frac{D_M}{D_L}\right)^2}\right] \qquad \text{(VII.2)}$$

where ϕ_d is the energy through the limiting aperture, n is the number of pulses, ϕ_o is the total energy in a pulse, D_M is the diameter of the limiting aperture, and D_L is the diameter of the beam at the 1/e points. Solving for n, we obtain the maximum number of pulses per interval, T_{max}.

$$n = \frac{\phi_d}{\phi_o\left[1-e^{-\left(\frac{D_M}{D_L}\right)^2}\right]} \qquad \text{(VII.3)}$$

From eqn. (VII.3) we calculate that a maximum of 4772 pulses are allowed over a 10 second interval when the pulse energy is 20.2 μJ and ϕ_d = MPE = 96.2 mJ. We will refer to this pulse-count limit as $N_{max.}$

Gating the repetitive pulsed laser

In order to run our laser in the laboratory as a Class 1 level laser without reducing the per-pulse energy, we can gate the laser at 1 Hz with a 9.5 ms gate. Pulse packets recurring at 1 Hz with approximately 475 pulses in each packet will be the result. Our laser has a buildup time such that for the first few milliseconds of the gate, the pulse energy is increasing to a steady-state energy of 20.2 μJ; this further reduces the integrated exposure.

A disadvantage to this approach is that there are fewer pulses to be measured, analyzed, and averaged. Also, the 1 Hz refresh rate slows down alignment of the lidar optical system. The advantages are that in this operating condition safety goggles and an interlock aren't required.

Scanning the repetitive pulsed laser

The worst case scenario occurs when 1) the eye is looking directly at the laser, 2) the laser scans over a diameter of the pupil, 3) the product of laser beam divergence and propagation distance is small compared to the initial beam diameter, and 4) the laser trace passes across the pupil of an eye 30 times per second (This could occur if the large mirror accidentally stopped rotating, or perhaps at the center of the scan pattern in figure 4.). For pulsed laser scanning, our objective is to find the "nominal ocular hazard distance" defined as "the distance along the axis of the unobstructed beam from a laser, fiber end, or connector to the human eye beyond which the irradiance or radiant exposure is not expected to exceed the applicable MPE." This is the distance from the first scan mirror someone must be in order to limit the number of possible laser pulse exposures to a value less than N_{max}.

For conservative evaluation, we find the maximum possible number of scans across the eye and the maximum possible number of pulse exposures into the eye during each scan. Then we compare N_{max} to the product of these two upper limits.

The maximum number of scans across the eye is,

$$n_s = floor(T_{max} \times f_s) + 1, \tag{VII.4}$$

where T_{max} is the limiting exposure duration and f_s is the scanning mirror rotation frequency. The maximum number of exposures in each scan is given by

$$n_e = floor\left(\frac{\theta_L}{2\pi} \times \frac{f_L}{f_s}\right) + 1, \tag{VII.5}$$

where θ_L is the radial displacement of the scan required to pass over the pupil, and f_L is the frequency with which the laser is pulsing. The number of pulse exposures during the limiting exposure duration therefore satisfies the inequality

$$n_{exp} \leq (T_{max} \times f_s + 1)\left(\frac{\theta_L}{2\pi} \times \frac{f_L}{f_s} + 1\right). \tag{VII.6}$$

Solving for θ_L we obtain

$$\theta_L \geq \frac{2\pi f_s}{f_L}\left(\frac{n_{exp}}{T_{max}\times f_s+1} - 1\right). \quad \text{(VII.7)}$$

In order to obtain a conservative hazard distance as a function of the rotation frequency of the scan mirror we use

$$\theta_L = 2sin^{-1}\left(\frac{r}{R}\right) \text{ and} \quad \text{(VII.8)}$$

$$R = L \times tan(\theta_d) \quad \text{(VII.9)}$$

where r is the radius of the limiting aperture, R is the radius of the circular scan trace, θ_d is half the cone angle of the circular scan, and L is the distance from the scanner to the screen. Solving equations (VII.7 & 9) for L, we obtain

$$L \leq \frac{r}{\sin\left(\frac{\pi f_s}{f_L}\left(\frac{n_{exp}}{T_{max}\times f_s+1}-1\right)\right)tan(\theta_d)} \quad \text{(VII.10)}$$

If we now substitute N_{max} for n_{exp} we have a conservative hazard boundary as a function of rotation frequency, f_s. Hazard boundary estimates can be derived similarly for other types of scanning, such as an oscillating scanner or polygonal scanner.

Plots of $L(f_s)$ for the conically-scanned lidar system described in sections II and III can be seen in Figures 42a and 42b. In this case, the hazard distance has a minimum value of approximately 104 mm at a scan rotation frequency of approximately 6.8 Hz; it is less than 110 mm for rates up to 30 Hz (the nominal rate for the small scan mirror). This result locates the hazard boundary conveniently inside the instrument, between the two scanning mirrors.

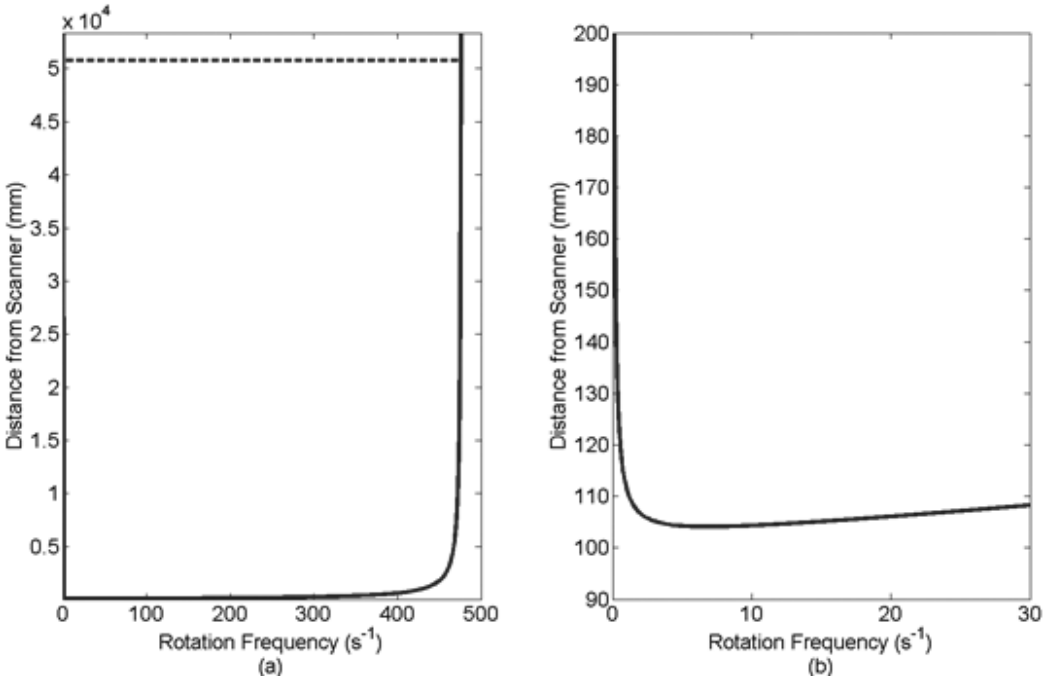

Figure 42: (a) Solid line - hazard boundary as a function of the rotation frequency of the scan mirror. Dashed line - hazard boundary for the staring beam. (b) Zoomed-in view.

The scanning-based hazard boundary (eqn. VII.10) does not account for beam divergence. A hazard distance that does account for the beam divergence can be calculated for the case of the staring beam. This divergence-limited hazard distance, calculated via eqn. B54 from [3] is shown as a horizontal dotted line in Figure 41a. At extremely low and extremely high scan frequencies, eqn. 10 sets a hazard boundary that is larger than the hazard distance for the staring case. Under such scan conditions, the staring limit (~50 m in this case) should be applied.

VIII. Summary

Demonstrated Results:

- a volume-scanning elastic lidar system based on an Erbium fiber laser operating at 1.55 μm wavelength and 50 kHz pulse repetition;
- construction of the volume-scanning using commercial-off-the-shelf components including laser, detector module, refractive optics, spectral filters, analog and digital DAQ instruments, scan mirrors, and scan motors;
- Class I eyesafe operation of a 1W pulsed laser, eliminating the need for special safety restrictions in the field;

- real-time construction of aerosol volume images, including scan angle calculations and conversion from a hypertrochoid scan pattern to a hybrid Cartesian format;
- 1D phenomenological studies showing that under typical ambient conditions aerosol inhomogeneity is insufficient for lidar volume imaging at ranges beyond 100m; and
- extension of the Horn-Schunk optical flow method to three spatial dimensions and successful application of this method for accurate real-time analysis of wind fields from volumetric images.

Lessons Learned:

- The lidar receiver design must include re-imaging optics that collimate the beam internally at the position of the spectral filter(s).
- The receiver design must block off-axis solar images formed by the objective lens and tolerate the resultant heat.
- Float-glass mirrors are flat enough for a precision lidar scanner.
- For a volume-scanning lidar, the bottom edge of the scan pattern should be set 5° or more above horizontal.
- Higher pulse energy and/or a larger receiver aperture are needed for imaging ambient aerosol features.
- High performance APD modules require an secondary low-noise preamp in addition to the built-in TIA. Ideally the preamp output should be referenced near 0V.
- Stability of the APD bias voltage and elimination of other low-frequency interference is critical for atmospheric lidar applications.
- The data transfer channel and protocol can severely limit the effective data from the waveform DAQ instrument to the host computer.
- The compound trigger that synchronizes laser pulses, waveform acquisition, and encoder data acquisition should include a monitor capability to ensure nominal amplitude and duty cycle.

Project Status and Recommendations

- USU filed for a US patent on the concept of wind field derivation from dynamic aerosol imagery [4].
- The APD module failed during 3D collections on 5 April 2011. It was subsequently replaced and the system is currently operational.
- We plan to use WiLD for further phenomenological observations associated with ambient aerosols under conditions of higher wind speeds and inversion conditions

- 1D observations in the presence of widespread haze could help us further validate the WiLD overlap function.
- The WiLD control software will be consolidated and enhanced.
- We need to identify the cause(s) of low-frequency fluctuations in the waveforms from 3D collections.
- We plan consider the applicability of WiLD for high-resolution field observations using injected aerosols to increase image contrast and coverage.
- We plan to investigate the utility of WiLD for spatially-resolved ceilometry including remote observations of inversion boundaries.
- If the risk of permanent APD damage can be mitigated, we would like to characterize the transient response of WiLD to solar transits.

References

1. Bell, E. E., L. Eisner, J. Young, and R. A. Oetjen, Spectral Radiance of Sky and Terrain at Wavelengths between 1 and 20 Microns. II. Sky Measurements, JOSA, vol. 50, no. 12, pp. 1313–1317 (1960).

2. Horn, B. K. P. and B. G. Schunck, Determining Optical Flow, Artificial Intelligence, vol. 17, pp. 185-203 (1981).

3. American National Standards Institute (ANSI), *American Nation Standard for Safe Use of Lasers, ANSI Z136.1-2007*, Laser Institute of America, Orlando, FL, 2007.

4. Marchant, A. B., R. D. Barson, M. Wojcik, and A. Q. Howard, Jr., Dynamic 3D wind mapping system and method, Patent Application US 2011/0149268 A1 (2010).