

Two-component winds from an eye-safe scanning elastic backscatter lidar

Shane D. Mayor

Department of Physics, California State University Chico

The following chapter is based on a manuscript (Mayor et al., 2012) recently accepted for publication in the American Meteorological Society’s *Journal of Atmospheric and Oceanic Technology*.

1 Introduction

The use of Doppler lidars are likely to pervade the discussion when it comes to wind measurement by lidar. The size and cost of Doppler systems has decreased while performance and reliability have improved. However, Doppler lidars are designed with one primary measurement goal: the line-of-sight (radial) velocity component through the detection of the Doppler frequency shift of the backscattered radiation. The frequency shift generated by the wind (relative to the laser frequency) is so small and challenging to measure that heterodyne (i.e. “coherent”) detection has emerged as the leading technique to obtain the Doppler frequency shift. While it is very accurate, heterodyne detection imposes strict requirements of optical coherence on the transmitted and detected laser radiation and limitations on the resulting data due to interferometric issues.

Doppler lidars (Grund et al., 2001, Henderson et al., 1991, Post and Cupp, 1990, Pearson et al., 2009) provide high quality *direct* measurements of only the radial component of motion (Mayor et al., 1997, Mann et al., 2009, Sathe et al., 2011, Sathe and Mann, 2012). In order to obtain two or more components with a Doppler lidar, additional techniques have been devised. They include: (1) The use of two Doppler lidars separated by several kilometers distance while simultaneously scanning over the same area (known as dual-Doppler technique) (Newsom et al., 2005, 2008, Hill et al., 2010); (2) the use of full or wide sector scans in azimuth, assumption of homogeneous flow, and curve fitting to obtain vertical profiles of horizontal velocities; or (3) the use of flow retrieval models (Newsom and Banta, 2004a,b).

This paper presents two-component wind data derived from sector scans of a single eye-safe elastic backscatter lidar. The system employs direct-detection and it is *not* a Doppler lidar. The flow vectors are derived using a motion estimation algorithm that has been applied to consecutive pairs of backscatter images. While the technique may not result in line-of-sight vectors with the precision and spatial and temporal resolution and independence as those from Doppler lidars, the approach is not encumbered by the optical coherence requirements that dictate the primary design characteristics of Doppler systems using heterodyne detection. These heterodyne design aspects include velocity precision that is directly related to laser pulse length (therefore discouraging the use of short laser pulses), diminishing benefits of larger telescope diameters (due to turbulence induced speckle), and (likely most significant) speckle-limited backscatter intensity data (Hardesty, 1984, Rye and Hardesty, 1993a,b, Frehlich et al., 1997). Without these restrictions, direct-detection backscatter lidars with high pulse energy, short pulse length, and large telescope diameter can produce richly textured backscatter images that appear to hold useful multi-component wind information.

Deriving velocity information from the movement of features in image sequences is a field known as *optical flow*. The cross-correlation method (also called phase correlation) is perhaps the most widely known and understood form of optical flow. The cross-correlation technique has been applied




Figure 1: (Color required.) Mean offshore wind speed and direction derived from aerosol backscatter scans of the University of Wisconsin Volume Imaging Lidar during the 1997-1998 Lake-Induced Convection Experiment (Lake-ICE) using the cross-correlation technique. 41-minutes of cross-correlation functions were averaged. Vectors are spaced every 250 m. (From Mayor and Eloranta (2001)).

to aerosol lidar data several times previously (Eloranta et al., 1975, Sroga et al., 1980, Kunkel et al., 1980, Sasano et al., 1982, Hooper and Eloranta, 1986, Kolev et al., 1988, Schols and Eloranta, 1992, Piironen and Eloranta, 1995, Mayor and Eloranta, 2001). Mayor and Eloranta (2001) showed two-component vector fields with 250 m horizontal resolution over areas as large as 60 km². These



Figure 2: Photo of the NSF/NCAR Raman-shifted Eye-safe Aerosol Lidar (REAL) located at the California State University Chico farm. The lidar resides only in the shipping container on the left. Although all atmospheric lidars must detect backscatter to perform their intended function, the REAL provides the required combination of (1) eye-safety, (2) invisible beam, (3) rapid scanning, (4) continuous operation, (5) high sensitivity to small changes in aerosol properties, and (6) high temporal and spatial resolution data. It does this by transmitting laser pulses with sufficient energy to generate strong SNR backscatter for several kilometers range without any temporal or spatial averaging. In 2012, we expect to complete software that enables the REAL to produce vector flow fields such as those shown in Figs. 19 and 20 in real-time. A discussion of the software to perform the calculation in real-time is given in section 8.5.

vectors however were the result of averaging many cross-correlation functions over time—up to 41 minutes in one case (see Fig. 1). Furthermore, the lidar system used in Mayor and Eloranta (2001) was not eye-safe and the vectors were not verified with an independent wind measurement. In this paper, new results are presented from the 2007 CHATS experiment. The new data were collected with an eye-safe system and the presence of an instrumented tower enabled validation. This paper also briefly discusses more advanced forms of optical flow in section 8.4 and the ability to perform the cross-correlation computation in real-time in section 8.5.

2 Instrument

The following work is based on data collected with the NSF/NCAR Raman-shifted Eye-safe Aerosol Lidar (REAL, see Fig. 2). The instrument is described by Mayor and Spuler (2004), Spuler and Mayor (2005), Mayor et al. (2007). It is a ground-based, scanning, elastic backscatter lidar operating at 1.54 microns wavelength. Table 1 lists the specifications of the system as configured for the experiment.

The REAL is often criticized as being too large and complex. This is likely to result from the impression one gets when looking at Fig. 2. However, the actual instrument only occupies a portion of the shipping container on the left. No attempt was made to reduce the size or weight of this experimental prototype. (The container on the right is an office area for use during field

deployments.) Duplicates of the REAL have been created by ITT Corp. for the Department of Defense and they are approximately the size of half of one shipping container (Ponsardin et al., 2006, Mayor et al., 2008). The NSF/NCAR REAL operates at 10 Hz pulse repetition frequency. DOD versions operate at 20 Hz. Spuler and Mayor (2007b) demonstrate a transmitter in the laboratory with twice the pulse energy of the NSF/NCAR REAL at a repetition rate of 50 Hz. This is significant because faster pulse rates enable faster scans (with equivalent angular resolution). Faster scans provide imagery with smaller changes in aerosol features which reduces the challenge for motion estimation algorithms to retrieve the flow field.

The REAL uses a commercially-available, flash-lamp pumped Nd:YAG laser and a high-pressure gas cell to convert 1.064 μm radiation to 1.543 μm radiation. The 10 Hz Nd:YAG laser requires changing flash-lamps approximately every 20 million pulses or every 3 weeks of continuous operation. A careful undergraduate student can be trained to do this. The Raman cell requires no maintenance if pumped correctly.

Table 1: Specifications of the NSF/NCAR Raman-shifted Eye-safe Aerosol Lidar (REAL).

Wavelength	1.543 microns
Pulse energy	170 mJ
Pulse rate	10 Hz
Pulse duration	6 ns
Beam diameter at lidar	66 mm ($1/e^2$ points)
Beam divergence	0.24 mrad (full-angle)
Telescope dia.	40 cm
Receiver FOV	0.54 mrad (full-angle)
Digitizer speed	100 MSPS
Digitizer range	14 bits
Detector type	200- μm InGaAs APD

3 Experiment

The data presented in the upcoming pages were collected between 19 March and 11 June 2007 near Dixon, California, during the Canopy Horizontal Array Turbulence Study (CHATS) (Patton et al., 2011). The REAL was located 1.61 km directly north of the NCAR Integrated Surface Flux Facility (ISFF) 30-m vertical tower (VT) (see Fig. 3). The VT was located inside an 800 m \times 800 m orchard of 10 m tall walnut trees. The VT was located 100 m from the northern edge of the orchard in order to maximize the fetch over the canopy during the prevailing southerly flow. The VT supported 13 Campbell Scientific CSAT3 3-D sonic anemometers of which 5 were located above the tree tops at 12.5, 14, 18, 23, and 29 m above ground level (AGL). The flat terrain in the vicinity of the experiment, relatively short height and uniformity of the orchard, and absence of additional obstructions between the lidar and the orchard enabled nearly horizontal atmospheric cross-sections (hereafter referred to as plan position indicator scans or PPI scans) that were collected over a wide area surrounding the orchard. The laser beam, which was projected from the REAL at a height of 4.2 m AGL, passed above the tops of the trees in the orchard and intersected the VT at a height of approximately 18 to 20 m AGL. This corresponds to a slope of 8.6 m per kilometer. Therefore, the PPI scans were within the atmospheric surface layer and the resulting cross-sections were essentially planar and horizontal. The laser beam at the point where it is emitted from the lidar is 6.6 cm in

diameter and has a half-angle divergence of 0.12 mrad. This results in a beam diameter of 45 cm at 1.61 km range and 1 m at 4 km range. The laser pulse duration is 6 ns which corresponds to 1.8 m length.




Figure 3: Plan-view of the 2007 CHATS experimental area. The shaded regions from 150° – 210° and 175° – 185° azimuth represent the areas covered by “wide” and “narrow” PPI scans, respectively. The vertical tower (VT) was located 1.61 km directly south of the REAL. The squares centered on the VT represent the image blocks extracted from the gridded PPI scan data that were used to compute motion vectors via the cross-correlation technique.

We estimate that the majority of PPI cross-sections intersected the VT between 18 and 20 m height (see Fig. 4). This estimate is based on observations of hard-target reflections from the horizontal span of VT guy wires. No attempt was made during the experiment to observe the beam on the tower with an infrared viewer. Furthermore, the height of the lidar beam at the VT is likely to have changed over time resulting from drifts and adjustments in the attitude of the lidar trailer.¹

¹Due to the unexpected trenching and flooding of a nearby irrigation ditch by land-owners, one side of the lidar trailer often stood in mud while the other side on dry soil. Staff attempted to compensate for the sinking of one side by periodically adjusting the leveling system. However, measurements and records were not kept. A 1 mm change in




Figure 4: Diagram showing the approximate altitude, spacing, and size of REAL laser pulses with respect to the tree tops and the vertical tower at CHATS. This diagram is an east-west cross section looking either towards or away from the lidar. Red circles represent laser pulses from one scan at an azimuthal scan rate of 4° s^{-1} . At this scan rate, the pulses are spaced 11 m apart at 1.61 km range from the lidar.

At an azimuthal scan rate of 4° s^{-1} , and a pulse repetition frequency (PRF) of 10 Hz, the distance between laser pulses in the east-west direction at the range of the VT was 11 m. The relatively narrow profile of the tower², small diameter of the laser beam, and distance between pulses at the range of the tower, resulted in intermittent hard target returns. The beam steering unit operated independently of the 10 Hz transmitter and, as a result, laser pulses were usually not projected at the same azimuth angles in consecutive scans.

The aerosol backscatter performance of the lidar has an impact on the vectors derived from the backscatter data. For example, a low performance system may not detect variations in backscatter intensity and could result in an absence of features to track. For this study, a two-channel polarization-sensitive receiver was used Mayor et al. (2007). This arrangement of hardware splits the total backscatter power into two detection channels thereby increasing the amount of electronic noise by $\sqrt{2}$ compared to that of an equivalent single channel detection system. However, dividing the signal into two channels enabled better quantization of the signal. This configuration likely resulted in higher signal-to-noise ratio (SNR) in the first few km of range but lower SNR at far distances when compared with a single channel system. This issue, and other system variables such as transmitted laser pulse energy and noise equivalent power of the photodetectors have a significant impact on the SNR of the backscatter data and therefore affect the quality of the retrieved velocities. Therefore, we document the SNR performance.

Signal-to-noise ratio (SNR) of the backscatter data was calculated by first computing the standard deviation (σ_b) and mean ($\overline{S_b}$) of 375 digitizer samples recorded during the $3.75 \mu\text{s}$ prior to the discharge of each laser pulse. σ_b is a measure of the intensity of the electronic noise in the

the height of one side of the trailer may have resulted in a 0.5 m change in the height of the beam at the tower.

²The central column of the tower was 32 cm wide. Guy wires 0.47 cm in diameter were attached to the tower at 4 heights on the tower. Guy wires from 7.9 and 18.8 m AGL were anchored into the ground 13.4 m from the tower base. Guy wires 23.8 and 32.0 m AGL were anchored into the ground 26.8 m from the base.




Figure 5: Frequency distribution of single pulse SNR of the aerosol backscatter data as a function of range resulting from a subset of 208 905 laser pulses that were directed almost horizontally and toward 179° azimuth throughout the 3-month long CHATS experiment. The vertical dotted lines at 1.1 and 2.1 km range correspond to the northern and southern edges of the largest (1 km^2) block used to compute horizontal velocity vectors. The ISFF VT was located at 1.6 km range.

detection system and \overline{S}_b is proportional to the background intensity. Then, for each element in the waveform, the background is subtracted and the result divided by the noise value as shown in equation 1, below.

$$SNR(r) = \frac{S(r) - \overline{S}_b}{\sigma_b} \quad (1)$$

This procedure is applied to each pulse individually and the resulting arrays are not averaged over time. Figure 5 shows the frequency distribution of single shot SNR as a function of range for a subset of 208 905 laser pulses (1 pulse per scan) that were directed nearly horizontally over the center of the CHATS experimental area throughout the approximately 3 month period. It shows that individual pulses in the range of the experimental area between 1.1 and 2.1 km range typically result in SNRs of approximately 100 at 1.1 km range and decrease to approximately 20 at 2.1 km range. The broad range of SNRs results from variability in both the atmosphere and the instrument performance. For example, the transmitted laser pulse energy decreased slowly (2.5 – 5.0 mJ per day) during operation until operators either increased flash lamp voltage (every 5 to 10 days) or replaced the flash lamps (approximately every 21 days). Figure 7 of Spuler and Mayor (2007a) shows the transmitter performance as a function of time for the entire CHATS experiment. Also contributing to the range of SNR is the variability of the atmosphere. The concentration and microphysical characteristics of the aerosol fluctuate over time. In addition to changes in air mass, aerosol plumes resulting from local agricultural activities pass through the scan area and result in substantial perturbations in SNR. Therefore, the distribution shown in Fig. 5 represents the broad range of SNRs for the scans used to extract aerosol motion vectors reported herein.

4 Algorithm

In the present work, no temporal averaging of the cross-correlation functions was performed in order to evaluate the ability of the algorithm to measure the wind from only pairs of consecutive scans. The data were collected with an eye-safe elastic lidar operating at 1.54-microns wavelength (Mayor et al., 2007) and an instrumented tower intersected the nearly-horizontal lidar scan surface. The remote and in situ data were collected nearly continuously over three months in the atmospheric surface layer over an agricultural area. Comparison of mean remote and in situ velocity components are avoided due to uncertainty in the precise altitude of the lidar beam at the tower and frequent strong vertical wind shear profiles. Instead, this paper focuses on documenting the variability of velocity components resulting only from pairs of consecutive scans. Two examples of two-dimensional flow fields that result from the application of the algorithm to the entire scan area are presented.

The algorithm applied to the backscatter data to derive two-component aerosol motion vectors follows that described by Schols and Eloranta (1992). Lidar scans are composed of multiple “beams” with each beam made of uniformly-spaced backscatter samples; the beams form a polar grid of samples with the lidar at the origin. Each beam is produced by calculating and subtracting the mean background from its waveform, multiplying the waveform by the square of the range to remove the approximately one-over-range-squared dependence of the raw signal, and converting the waveform to decibels. Each beam then undergoes low-pass and high-pass median filtering to remove single point outliers and large scale features, such as the effects of attenuation and the inability to normalize for shot-to-shot laser pulse energy variations, respectively. For the results shown here, the low-pass filter length was set to 7 points corresponding to 10.5 m and the high-pass filter length was set to 333 points corresponding to 500 m.³

After processing the beams of a scan as described above, the scan data are projected onto a uniform rectangular grid using bilinear interpolation. The grid is oriented with the abscissa directed positive to the east and the ordinate positive to the north. In this study, grid spacing of 4, 6, 8, 10, and 20 m were tested. In addition to the processed backscatter data, the time of each pulse at millisecond resolution and the range dependent SNR are projected onto grids of the same dimensions and resolution, enabling this information to be easily and precisely extracted from the same regions of the backscatter data that are analyzed by the cross-correlation technique. A temporal median image is computed based on all of the filtered PPI scans with elevation angles greater than an amount sufficiently large to prevent hard target reflections from the tree tops and within the time span of each raw data file. A single raw data file typically contains dozens to hundreds of scans and includes data spanning periods of minutes to more than an hour in most cases.

Next, subsets of the gridded data are extracted in square regions corresponding to the 4 blocks shown in Fig. 3. For this work, vectors were calculated from $250\text{ m} \times 250\text{ m}$, $500\text{ m} \times 500\text{ m}$, $750\text{ m} \times 750\text{ m}$, and $1\text{ km} \times 1\text{ km}$ blocks to investigate the effect of block size on the motion vectors. Blocks of the same size and position and from pairs of consecutive scans are used to calculate each motion vector. Histogram equalization is applied to each block (Schols and Eloranta, 1992). Two-dimensional cross-correlation functions (CCF) are computed using fast-Fourier transforms and the Wiener-Khinchin theorem.

The resulting CCFs have resolutions equal to the grid spacing. Therefore, at this stage of the process, the measurement of aerosol feature displacement is limited to the grid spacing and the quantization of velocity estimates is limited to the grid spacing divided by the time between scans. In order to improve the velocity estimate, a two-dimensional polynomial based on the 5 by 5 set of points centered on the peak of the CCF was calculated as described by Piironen and Eloranta (1995). The location of the peak of the numerically fit function (with respect to the

³The transmitted beam falls entirely within the receiver’s field-of-view by approximately 500 m range.

coordinate system of the CCF) corresponds to the displacement caused by the predominant aerosol feature movement within the block area over two scans. While the peak refinement does reduce the quantization of the displacements, we found that it does not eliminate it entirely.

The velocity is determined by dividing the displacement by the time between scans. In addition to the above, the average lidar SNR of the block regions is computed for data analysis. These regions often contain a small number of pixels resulting from hard target reflections from the vertical tower and a grove of nearby trees that were substantially higher than the orchard canopy.⁴ However, the number of pixels influenced by these hard target reflections is very small compared to the total number of pixels in the blocks and we found that they do not significantly effect the mean SNR.⁵ Temporal median filtering reduces the biasing effects of these stationary features in the images.

The horizontal wind components from the sonic anemometer data at 18 m AGL on the VT were used as a standard reference. The nature of the sonic anemometer measurement is very different from the lidar measurement. Sonic anemometer velocity measurements are essentially of the air flow in a small volume in space and 60 Hz in time. The lidar vectors are based on the drift of macroscopic aerosol features over relatively large areas and are relatively sparse in time (once every 10 to 30 s) in comparison to the sonic anemometer data. Therefore, substantial time averaging of the sonic anemometer wind components are required to make a comparison with each aerosol motion vector.

As a first guess for a suitable time interval to average the sonic anemometer data over, we chose to begin the averaging interval when the lidar beam enters the block area on the first scan and end the averaging interval when the lidar beam exits the block area on the second scan. Therefore, each block size has a slightly different averaging duration (larger blocks result in longer averaging times). Figure 6 depicts how the sonic anemometer time series data are averaged relative to the lidar scans. To investigate the effect of varying the interval over which to average the sonic anemometer data, we expanded and contracted the duration of the intervals. No significant improvement was found.

The REAL operated at CHATS with a constant pulse repetition frequency (PRF) of 10 Hz. However, all other parameters controlling the scans were variables that could be altered in order to optimize the scans to achieve a variety of experimental objectives, some that were related to the direction of the flow. The variables include the angular scan rate, the angular width of a scan, and the sequence of scan types to be performed. In summary, a large number (approximately 200 000) of the PPI scans collected were intended to reveal a broader view of the experimental area. These are called “wide” and ranged from approximately 150° to 210° azimuth and were performed at a rate of 4° s⁻¹. Another large percentage of the PPI scans (approximately 75 000) were termed “narrow” and were intended to probe the fine-scale structure and motion of the atmosphere in the immediate vicinity of the tower. The narrow scans ranged from 175° to 185° azimuth. A majority of the narrow scans were collected during gentle southerly winds in an effort to observe canopy-scale turbulent coherent structures over the orchard.

Changing scan strategies during the experiment resulted in a variety of time periods between scans. The most frequent time between scans, hereafter referred to as dt, was 30 s. The second most frequent was 17 s. Scans with update periods of 45 s and 10 s were also collected occasionally. In summary, the population of PPI scans available for analysis have a wide variety of associated parameters (dt, elevation angle, angular width, angular resolution) as well as atmospheric conditions in which they were collected. This complicates the analysis. But, it also provides an opportunity to explore the effects of adjustable parameters so that they can be optimized for best results in the

⁴The associated hard-target returns from these trees were located in a region near 1.50–1.65 km south and 0.35 km west of the lidar.

⁵Given a pulse rate of 10 Hz, an angular scan rate of 4° s⁻¹, and a digitizer speed of 100 mega samples per second, the number of data points in the native spherical coordinate system of the lidar falling into blocks that are 250 m × 250 m, 500 m × 500 m, 750 m × 750 m, and 1 km × 1 km, and centered on the VT are 3692, 14 885, 33 650, and 60 222, respectively. After interpolation to a Cartesian grid with spacing of 10 m in both east-west and north-south dimensions, the number of points (or pixels) in the blocks are 625, 2500, 5625, and 10 000, respectively.




Figure 6: Visual depiction of how anemometer time series data were averaged for comparison with the vectors derived from lidar scans. This example considers a series of four consecutive PPI scans that result in three vectors. The anemometer data are averaged over the periods corresponding to when the lidar’s beam enters the block on the first scan (positions 1, 5, and 9 in time) and exits the block on the subsequent scan (positions 6, 10 and 14 in time).

future experiments.

5 Time Series Comparisons

Of the approximately 275 000 PPI scans collected during CHATS, our analysis used only the vectors resulting from a subset of approximately 180 000 PPI scans. A large number of PPI scans were contaminated by intermittent reflections from top foliage of the orchard trees as a result of inadvertently setting the elevation angle too low. Our analysis excludes vectors that resulted from partially filled blocks or blocks with substantial contamination from hard target returns.

Time-series plots of the velocity components, like those shown in Figs. 7, 8, and 9 were created for the entire data set. We have examined them all and concluded that the agreement between lidar and anemometer velocity components varies substantially and depends on a number of factors. In this section, we present examples of time-series of the velocity components with good agreement but from different meteorological and wind conditions. There are also, of course, many examples of poor agreement. However, poor agreement can be easily described as noisy results and there is little value in showing it. We conclude this section by showing the results of a statistical analysis that includes all 180 245 vectors.




Figure 7: (Color required.) Time series of lidar-derived (black points) and averaged sonic anemometer (color traces) velocity components for a 4-hour period during light wind conditions. A $250 \text{ m} \times 250 \text{ m}$ block size was used for the lidar-derived velocity estimates.

5.1 Light and variable wind example

Figure 7 shows the velocity components derived from the lidar data (black dots) and the averaged sonic anemometer data from 5 heights on the VT during a weakly stable 4-hour period (10 - 14 UTC on 26 March 2007) with light and variable winds. For this case we present the results from application of the $250 \text{ m} \times 250 \text{ m}$ block size to show the skill of the cross-correlation algorithm with the smallest of the 4 block sizes tested. Larger block sizes result in smoother time-series. The scan update rate was 17 s during this period.

The mean wind speed during this 4-hour period, according to the sonic anemometers on the VT, ranged from 1.3 m s^{-1} at 12.5 m AGL to 2.25 m s^{-1} at 29 m. The mean turbulent kinetic energy ranged from 0.04 to $0.08 \text{ m}^2 \text{ s}^{-2}$ among the 5 sonic anemometers. The mean gradient Richardson number was 0.02 indicating weak static stability. The lidar SNR over the $250 \text{ m} \times 250 \text{ m}$ block ranged from 25 to 100 with a mean of 78. CCF maxima ranged from 0.1 to almost 0.9 with a mean of 0.4. Particular attention however should be taken to notice the very good correlations between perturbations in the lidar velocity components and those from the anemometers. Linear correlation coefficients from this period range from 0.70 to 0.79 with the strongest correlation corresponding to the anemometer at 23 m AGL.

5.2 Strong wind example

Figure 8 shows the velocity components derived from the lidar data using $1000 \text{ m} \times 1000 \text{ m}$ blocks (black dots) and the averaged sonic anemometer data from 5 heights on the VT (colored lines) during a 4 hour period with strong NNW winds (20 - 24 UTC on 21 March 2007). Lidar velocity estimates for blocks that were $500 \text{ m} \times 500 \text{ m}$ and smaller produced substantially noisier results with little to no skill in matching the anemometer data. According to the tower anemometer data, the mean wind speed above the tree tops ranged from 8 m s^{-1} at 12.5 m AGL to 12.2 m s^{-1} at 29 m AGL for the 4 hour period. The mean TKE ranged from $4.93 \text{ m}^2 \text{ s}^{-2}$ at 12.5 m AGL to $1.58 \text{ m}^2 \text{ s}^{-2}$ at 29 m AGL. The mean gradient Richardson number was -0.005 indicating slightly unstable




Figure 8: (Color required.) Time series of lidar-derived (black points) and averaged sonic anemometer (color traces) velocity components for a 4-hour period during strong wind conditions. A 1 km^2 block size was used for the lidar-derived velocity estimates.

to nearly-neutral conditions.

The lidar collected wide PPI scans every 17 s during this period. The mean SNR over the 1 km^2 block ranged from 50 to 75 with only few very brief excursions exceeding 100. The mean CCF maximum ranged from 0.1 to 0.45 and averaged 0.29. The fluctuations in the time-series of lidar u-components show slightly increasing correlation with the averaged anemometer u-components as a function of height on the tower. At 12.5 m, the linear correlation coefficient, R , is 0.31 and at 29 m AGL, R is 0.39. The correlation coefficients for the v-components are essentially constant ranging from 0.19 at 12.5 m to 0.22 at 29 m. Overall, the time series (Fig. 8) suggest the algorithm has captured the strong mean flow for this period and the larger timescale u-component fluctuations on the order of 15 minutes or more. However, that is not to say that the technique is not capable of resolving sudden changes of the mean flow during turbulent conditions. An abrupt change in v-component of the flow can be seen at 13:30 UTC in Fig. 7. In the next section we show an example of a near reversal in wind direction that occurred over a 10-minute period with the passage of a density current front.

5.3 Changing wind during the passage of a density current front

Figure 9 shows the velocity component time-series comparisons from the afternoon of 26 April 2007 between 21:30 UTC on 26 April and 01:30 UTC on 27 April 2007. The lidar was programmed to collect alternating RHI and PPI scans resulting in one PPI scan (or one RHI scan) every 30 s. The PPI scans were directed between 151° and 211° azimuth at a scan rate of 4° s^{-1} . During this period, a density current front passed over the experimental site at 23:25 UTC on 26 April (Mayor, 2011). The z/L stability parameter at 12.5 m height ranged from -2.0 to -0.6 (strongly to moderately unstable) before the arrival of the front to -0.5 to -0.2 (moderately to weakly unstable) after the passage of the front. Wind speeds at the beginning of the period ranged from 3 to 6 m s^{-1} and decreased until the front passed. The wind direction veered dramatically from 350° (N) before the front to 221° (SSW) after the front. The scatter in both forms of wind measurement decrease with time, as the atmosphere becomes more stable and the turbulence intensity decreases. Figure 10




Figure 9: (Color required.) Time series of lidar-derived (black points) and averaged sonic anemometer (color traces) velocity components for a 4-hour period during moderate wind conditions with a frontal passage. A 1 km^2 block size was used for the lidar-derived velocity estimates.

presents the same data shown in Fig. 9 except as speed (top panel) and direction (bottom panel).

5.4 Velocity component difference distributions

Figures 11 through 14 show the distributions of velocity differences as functions of the corresponding mean block SNRs, CCF maxima, observed horizontal wind speeds, and TKEs. The left panels in Figs. 11 through 14 correspond to the east-west velocity component differences ($U_{lidar} - \bar{U}_{sonic}$) and the right panels in Figs. 11 through 14 correspond to the north-south velocity component differences ($V_{lidar} - \bar{V}_{sonic}$). Figures 11 and 12 show that the velocity component differences tend to decrease as the mean SNR and CCF maxima increase. This is an encouraging result since the SNR of data from future experiments may be improved by increasing laser pulse energy, using larger collection optics, and lowering the noise intensity in the detection electronics. CCF maxima may also be improved by decreasing the time between scans (i.e., faster scanning). Figures 13 and 14 confirm that the velocity component differences tend to increase as the observed wind speed and TKE increase.

5.5 Lidar velocity components versus sonic anemometer velocity components

Figure 15 shows the distributions of 180 245 pairs of velocity components. The lidar-derived components resulted from 1 km^2 blocks and 17 s and 30 s scan update rates. The data points were accumulated into bins of 0.2 m s^{-1} by 0.2 m s^{-1} resolution. The numbers above the gray scale on the top edge are the accumulation limits for the shades below. The total number of comparisons for a given shade is printed in that block of the gray scale. We chose to shade bins containing as few as 1 pair to reveal the behavior of the algorithm over all conditions including infrequent high wind speed events. Doing so reveals a broad distribution of lidar-derived u-components when the sonic anemometer measurements are between -1 and 3 m s^{-1} . Similarly, a broad distribution of lidar-derived v-components exist when the sonic anemometer v-components exceed $\pm 2 \text{ m s}^{-1}$. These distributions are attributed to the prevailing north-south flow at the site during the field




Figure 10: (Color required.) Time series of lidar-derived (black points) and averaged sonic anemometer (color traces) wind speed and direction for the same 4-hour period shown in Fig. 9.

experiment. Figure 16 shows the distribution of observed wind speeds and directions that occurred while PPI scans were being collected during CHATS. The most frequently occurring air flow was from 180° to 200° at less than 4 m s^{-1} . The second most frequently occurring flow regime was from 330° to 360° between 2 and 6 m s^{-1} . Weak westerly flow was less common and easterly flow was rare. As a result of this non-uniform distribution of wind velocities, we randomly sampled the population of velocity differences to achieve a uniform distribution in wind direction. Figure 17 shows this distribution. The linear correlation coefficients are 0.74 and 0.89 for the u-components (left panel) and v-component (right panel), respectively.

6 First attempt at Quality Control

As shown in previous sections, the correlation between the lidar-derived motion components and sonic anemometer wind components varies according to a number of factors, including some that can be derived solely from the blocks of lidar data. In this section, we describe a quality control (QC) model that enables the exclusion of vectors that are likely to be in poor agreement with the anemometer data. This functionality is desired when applying the cross-correlation method to blocks of image data in the scan area that do not have anemometer data for validation. This allows us to create two-dimensional flow fields with the low quality vectors removed.

6.1 Creating the Quality Control Model

The data used to create the QC model were sampled from the lidar-derived velocity components using 1 km^2 blocks from scans with 17 s and 30 s update rates. In order to eliminate the biasing effects of the uneven distribution of wind directions at the CHATS site (as shown in Fig. 16), a random sample of 300 points per 10° interval was taken to form a training set. Therefore, a training set for the predictive model contains 10 800 points out of a total population of 180 245 data points. The dependent, or response, variable for the multiple linear regression is the magnitude of the difference between the lidar-derived velocity component and the sonic anemometer velocity

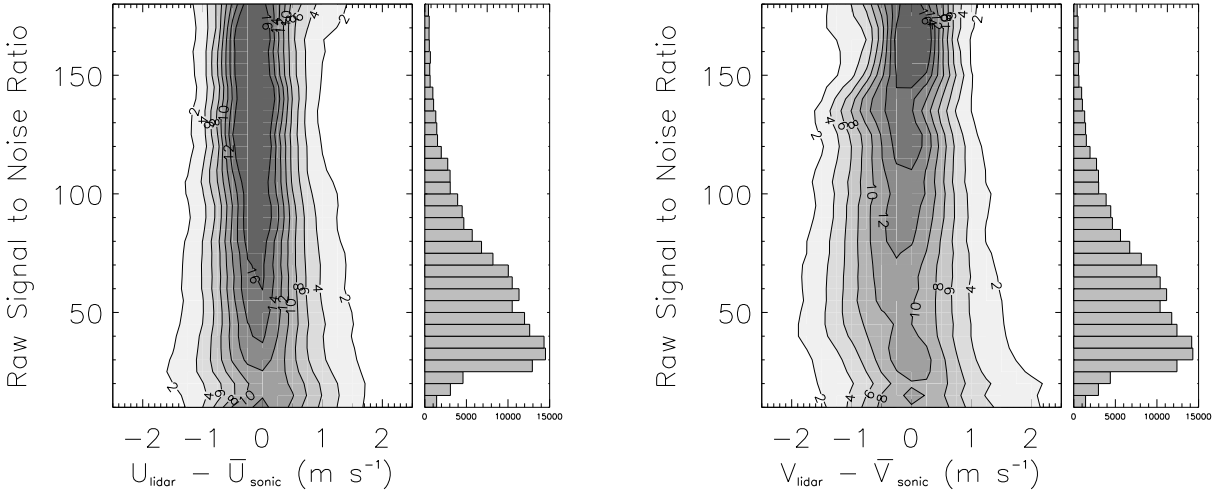


Figure 11: Distributions of velocity component differences as a function of average SNR in the blocks used to derive the lidar vectors. Contour interval labels are in percent. The bar chart to the right of each contour plot shows the number of data points in each SNR bin interval. The width of the distribution of velocity component differences tends to decrease as the SNR increase.

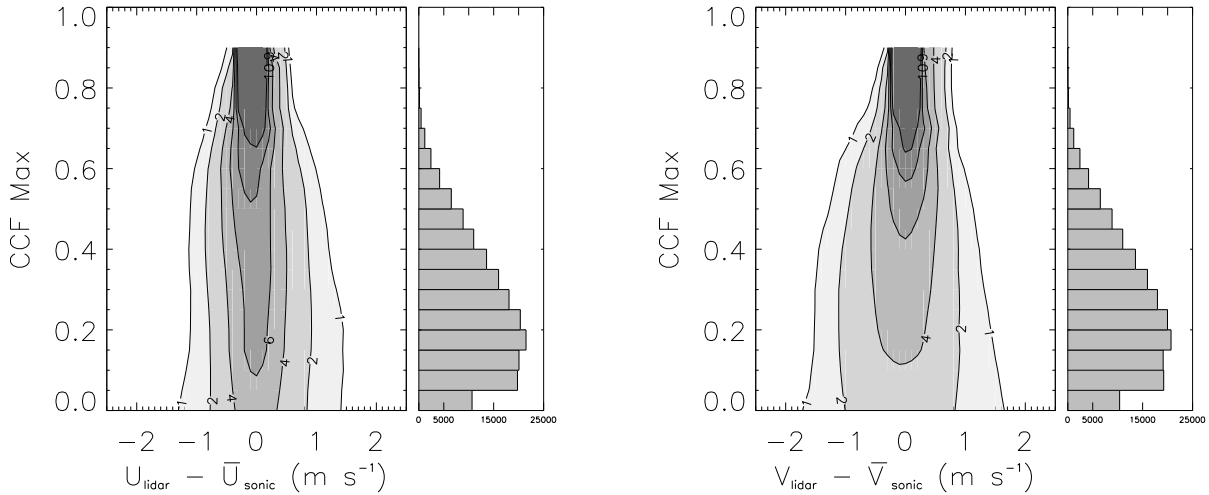


Figure 12: Distributions of velocity component differences as a function of the maximum of CCF corresponding to each lidar derived vector. The bar chart to the right of each contour plot shows the number of comparison points in each CCF maximum bin interval. The width of the distribution of velocity component differences tends to decrease as the CCF maxima increase.

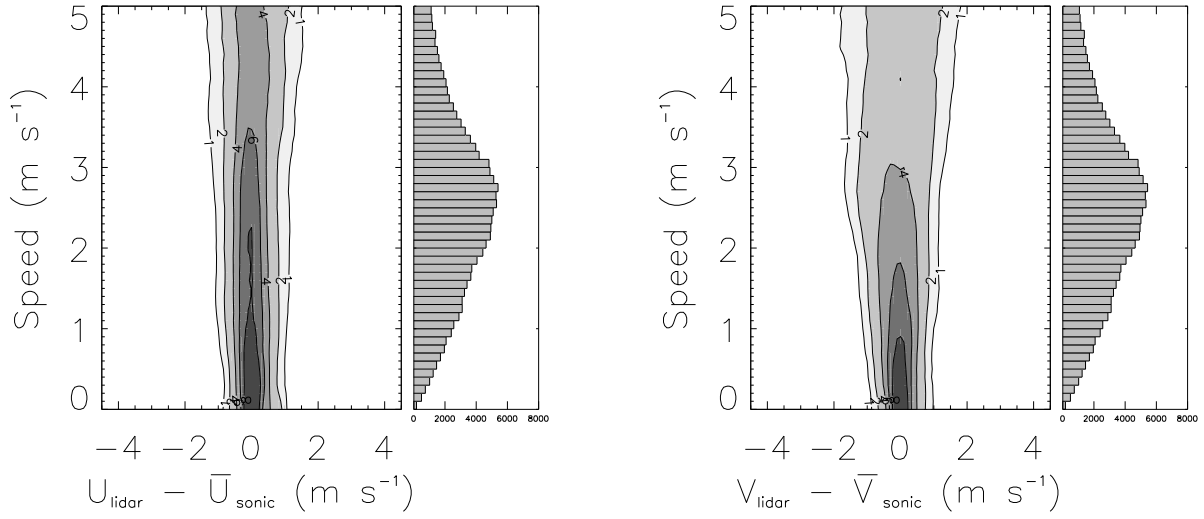


Figure 13: Distributions of velocity component differences as a function of the corresponding observed wind speed as measured by the sonic anemometer at 12.5 m on the ISFF VT. The bar chart to the right of each contour plot shows the number of comparison points in each wind speed interval. The width of the distribution of velocity component differences tends to increase as the wind speed increases.

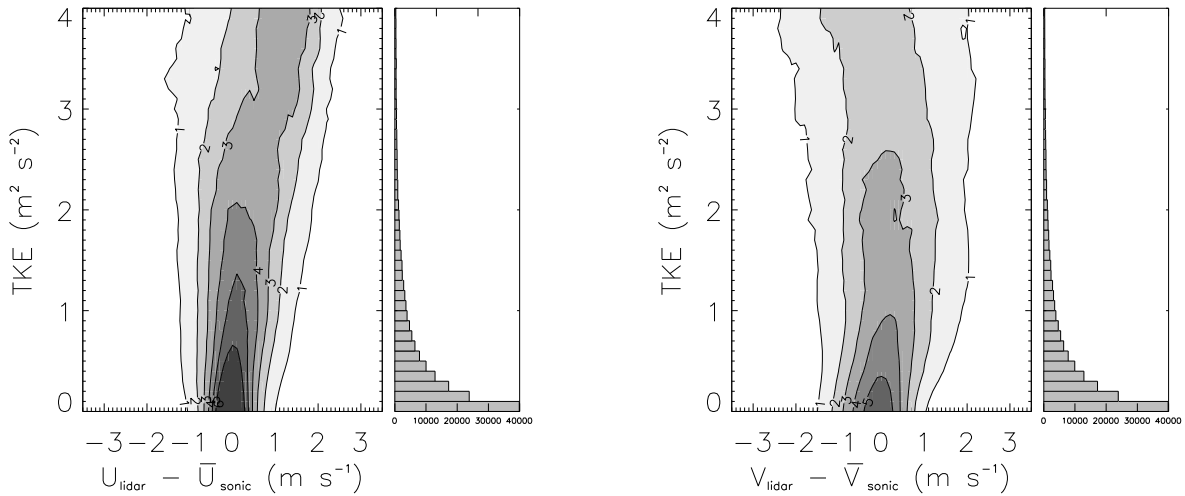


Figure 14: Distributions of velocity component differences as a function of TKE calculated from sonic anemometer data at 12.5 m on the ISFF VT. The bar chart to the right of each contour plot shows the number of comparison points in each TKE interval. The width of the distribution of velocity component differences tends to increase as the TKE increases.




Figure 15: Distributions of 18 m AGL sonic anemometer wind components versus aerosol motion components derived from the lidar data for 180 245 points using 1 km² blocks and 17 and 30 s scan update rates. The vast majority of points fall in the darker bins close to the dashed identity line. This distribution is the result of the non-uniform distribution of wind speeds and directions during the experiment shown in Fig. 16.

component, $|\widehat{U}_l - \widehat{U}_s|$ and $|\widehat{V}_l - \widehat{V}_s|$. Predictor variables are the corresponding wind speed (S), raw SNR (R), and the maximum of the CCF (C).

Multiple linear regression was applied on (1) first degree predictor variables, S , R , and C ; (2) first degree and interaction predictor variables (e.g., $S*R$, $S*R*C$); and (3) first degree, interaction, and second degree predictor variables (e.g., S^2 , R^2). Random subsampling was repeated 20 times for cross-validation of the models. Table 2 shows the mean adjusted R^2 (coefficient of determination) values of the regression models for u and v components in each of these cases. Including interaction and second-degree terms increases the ability of the model to account for variation in the data, so these terms were included in all subsequent regressions. Terms with significance less than 0.05 were not included in the model. The adjusted R^2 results for the u component ranged from 0.39 to 0.47, with a mean R^2 of 0.43, and adjusted R^2 results for the v component ranged from 0.24 to 0.33, with a mean adjusted R^2 of 0.30. Thus the model for the u component explains more variability in the data than does the v component model. Predictive models for u and v differences created from one stratified random sample are shown in equations 2 and 3.

$$\begin{aligned}
 |\widehat{U}_l - \widehat{U}_s| = & 1.066 - 0.5026 * C + \\
 & 0.05059 * S^2 + 0.00003382 * R^2 + 3.649 * C^2 - \\
 & 0.002577 * S * R - 0.09743 * S * C - 0.02463 * R * C + \\
 & 0.008761 * S * R * C
 \end{aligned} \tag{2}$$

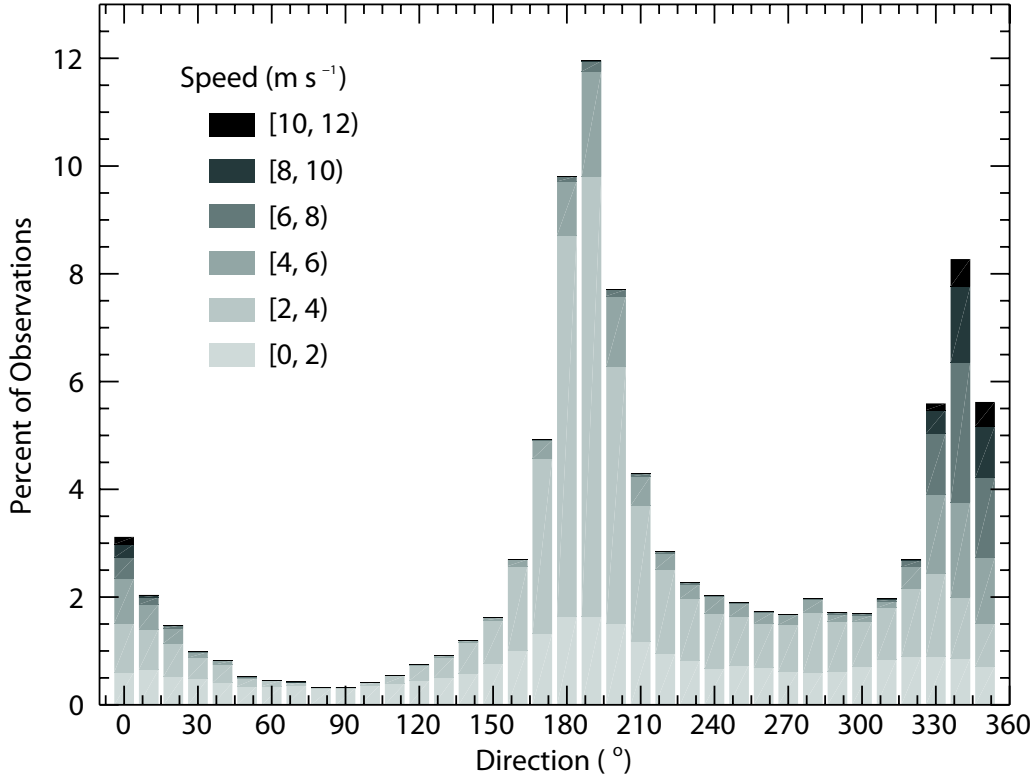


Figure 16: Distribution of 18 m AGL sonic anemometer wind measurements for the set of PPI scans collected during CHATS. The ideal environment for testing would provide a uniform frequency distribution of wind speeds and directions.

$$\begin{aligned}
 |\widehat{V}_l - \widehat{V}_s| = & 1.156 + 0.003156 * R - 2.875 * C + \\
 & 0.02572 * S^2 + 0.00002885 * R^2 + 5.576 * C^2 - \\
 & 0.00139 * S * R - 0.4965 * S * C - 0.0249 * R * C + \\
 & 0.004234 * S * R * C
 \end{aligned} \tag{3}$$

6.2 Testing the Quality Control Model

The testing set for each of the 20 random samples was the complete data set with the training set removed. Model testing was performed by removing model-predicted differences greater than 2 m s^{-1} from the testing set. The mean results show that the u model removes 9.6% of the data, of which 42.1% is correctly identified for removal. The u model fails to identify 49.8% of differences greater than 2 m s^{-1} . The v model removes 5.4% of the data, of which 58.2% is correctly identified for removal. The v model fails to identify 75.5% of differences greater than 2 m s^{-1} . Figure 18 shows the distributions of lidar-derived velocity components as a function of corresponding sonic anemometer velocity components after removal of model-predicted differences greater than 2 m s^{-1} . Correlation between the sonic and lidar u measurements for the complete data set is 0.52;




Figure 17: Distribution of 18 m AGL sonic anemometer wind components versus motion components derived from the lidar data after random sampling to account for the wind distribution shown in Fig. 16.

correlation between the u measurements after applying the QC model is 0.75. Correlation between the sonic and lidar v measurements for the complete data set is 0.84; correlation between the v measurements after applying the QC model is 0.90.

Table 2: Mean Adjusted R^2 Values for Regression Models

	1 st deg.	1 st deg. w/ interactions	1 st and 2 nd deg. w/ int.
u	0.22	0.36	0.43
v	0.18	0.26	0.30

7 Flow Fields

In addition to computing vectors for blocks centered on the tower location for time-series comparisons as described in section 5, we systematically applied the block and cross-correlation algorithm to all possible locations in the scan area to compute two-dimensional flow fields. Although the block size is large compared to the grid spacing, it can be moved laterally (in x and y) in increments as small as the grid spacing. The resulting vector flow fields may be spatially dense, with vectors at the grid spacing. Neighboring vectors in such cases have a large degree of common heritage. For example, applying a 1 km \times 1 km block on data with grid spacing of 10 m \times 10 m, a shift in one direction by one row or column results in 100 different points out of a total of 10 000 points in the block area, a 1 percent change in input data. Therefore, neighboring vectors are far from being independent measurements. However, as can be seen in the flow fields, modest changes in the box location can have significant impacts on the resulting vectors and reveal significant microscale circulations. We show two cases to make this point. We also apply the QC method described in section 6 to eliminate vectors that are likely to be in significant error.




Figure 18: Distribution of 18 m AGL sonic anemometer wind components versus motion components derived from the lidar data after application of the quality control algorithm. The linear correlation coefficient for the u-component comparisons (left panel) is 0.75. The linear correlation coefficient for the v-component comparisons (right panel) is 0.90.

7.1 Cellular surface layer convection

Time-lapse animations⁶ of high-pass median filtered aerosol backscatter data on the afternoon of 25 March 2007 reveal a roiling surface layer with broad regions of divergence, narrow bands of convergence, and numerous vortices. We note the remarkable resemblance of this flow with large-eddy simulation results in Fig. 14 of Sullivan and Patton (2011). The flow field shown in Fig. 19, our estimate of the two-component horizontal wind field at 00:19:47 UTC, was calculated using 1 km \times 1 km blocks and one pair of scans separated by 17 s. Vectors were calculated every 10 m in the horizontal Cartesian dimensions and streamlines were drawn by a particle trace procedure in Interactive Data Language with 2nd order Runge-Kutta integration. At the time, a vortex is located 2.7 km south and 0.2 km west of the lidar. The flow field also reveals a saddle point 3.5 km south of the lidar and 0.3 km east of the lidar. In situ data show light ($< 3 \text{ m s}^{-1}$) and variable winds until approximately 01:30 UTC when a uniform WSW flow swept across the region.

7.2 Density current front

Mayor 2011 describes seven density current fronts that passed over the experimental area during CHATS. Here we apply the cross-correlation method to a pair of scans from one of those cases and ending at 23:08:58 UTC on 26 April 2007.⁷ A block size of 1 km by 1 km was used and the scans were separated by 30 s. Figure 20 reveals the flow field when the front approximately bisected the scans area. Flow north of the front was northerly (indicated by blue streamlines) and flow south of the front was southerly (indicated by red streamlines). However, in addition to the narrow band of convergence at the front, the lidar-derived flow fields reveal eastward transport of air that flows

⁶The following time-lapse animations are available for viewing:

25 March 2007, 20:00 – 23:59 UTC: <http://www.youtube.com/watch?v=DojxQVSkIGc>

26 March 2007, 00:00 – 03:59 UTC: <http://www.youtube.com/watch?v=rN4BqRpjBlg>

⁷A time-lapse animation of this event covering the 3-hour period between 22:00 UTC on 26 April 2007 and 01:00 UTC on 27 April 2007 is available for viewing at: <http://www.youtube.com/watch?v=DFsLaUTEATw>

Start: 00:19:47 UTC
PPI Scan 49

Elevation: 0.20°
26 Mar 2007




Figure 19: Streamlines from the application of the cross-correlation algorithm to a pair of PPI scans through a convective afternoon atmospheric surface layer when winds were light and variable. Scans were separated in time by 17 s. A block size of 1 km \times 1 km was applied at the grid interval of 10 m. Streamlines were launched every 100 m.

into a vortex centered 3.7 km south of the lidar and 1.5 km east of the lidar. These observations show that flow may not rise over the front uniformly and rather may be transported significant horizontal distances before being swept up into narrow and rapidly rising currents.

8 Discussion

8.1 Hypothesis

A hypothesis prior to beginning this research project was that the velocity component agreement would be best during periods of sufficient turbulence intensity and poorest during periods of stability. The justification of the hypothesis was that turbulence mixes particulate matter and would result in aerosol features that were good tracers of the mean air motion in the block area. Conversely, stability at night was expected to result in worse agreement due to stratification and gravity waves with propagation vectors that may be significantly different from the local wind vector. The results

Start: 23:08:58 UTC
PPI Scan 301

Elevation: 0.20°
26 Apr 2007




Figure 20: (Color required.) Streamlines from the application of the cross-correlation algorithm to a pair of PPI scans when a density current front was advancing from the south. Streamlines are colored according to the sign of the v component with blue indicating northerly flow north of the front and red indicating southerly flow south of the front. The scans were separated by 30 s and a $1 \text{ km} \times 1 \text{ km}$ block was applied at the grid spacing of 10 m. Streamlines were launched every 100 m.

from CHATS indicate that the aerosol motion vectors are most correlated with sonic anemometer wind measurements during periods of stability and light winds. In general, as long as aerosol features are present, better results will be found at night with light winds than during more turbulent daytime conditions.

8.2 Resolution

The two-component vectors derived from the cross-correlation technique are not representative of the average motion over the block area. They are representative of the motion of the strongest coherent feature recognized by the algorithm. For example, a block covering an area that is composed of equally significant features moving in opposite directions will result in a CCF with two peaks. The algorithm we applied uses the strongest CCF peak (the largest maximum) to determine

a single velocity vector to represent the block. It is for this reason that the algorithm does not smooth the actual velocity field (within the area of a block) and that the derived velocity fields reveal variations at scales smaller than the block size. The response of the algorithm is not without comparison: a similar distribution of velocities may occur within the pulse volume of a Doppler lidar and signal processing algorithms are programmed to provide a single representative velocity (Rye and Hardesty, 1993a,b).

8.3 Unique environment

The unique environment, including the altitude of the measurements, spatially varying land uses below the blocks, and seasonally dependent agricultural activities, are likely to have influenced the data set and statistics of velocity differences. For example, as shown in Fig. 1, the four blocks were over different amounts of two dramatically different vegetation types: (1) walnut orchard and (2) either bare field or short crops such as tomato plants (depending on the season). The 250 m \times 250 m block was entirely over walnut orchard. However, for all of the larger blocks, only the southern portions were over orchard while northern portions were over bare field or short crops.⁸ As a result, only portions of the blocks sampled the atmospheric roughness sublayer where the presence of the canopy strongly influences the character of the turbulence (Kaimal and Finnigan, 1994, Adrian, 2007, Finnigan et al., 2009, Shaw et al., 1995, Su et al., 2000). (The depth of the roughness sublayer is about 3 times the height of the canopy.) We also note that the orchard was bare of leaves at the beginning of the CHATS experiment and fully leaved by the end (Patton et al., 2011). In addition to these two general surface types, an east-west oriented road passed beneath the three largest blocks on the northern edge of the orchard and a north-south oriented road passed beneath the largest block on the eastern side. Inspection of time-lapse animations of backscatter data reveal moving point sources of particulate matter from this roadway. Agricultural activities (e.g., plowing, spraying, harvesting) also resulted in significant moving point sources of dust.

8.4 Moving beyond cross-correlation

The cross-correlation method is perhaps the simplest form of optical flow. However, more advanced techniques exist. The other techniques have heritage in the fields of artificial intelligence and machine vision. Early work on optical flow focused on deriving the motion of solid objects (Horn and Schunck, 1981). These successes led to significant advancements in robotic navigation (visual odometry), video compression, and flight control (Ruffier and Franceschini, 2005). However, most optical flow methods are designed for application to images of moving solid bodies. Such problems have the advantage of more consistent shape and brightness over very short periods of time. Image sequences of fluid motions add significant complexities: changing brightness and shape deformation. In addition, images from remote sensors such as the lidar have more spatially and temporally varying amounts of SNR.

Through a collaborative effort with the group *Fluminance* led by Dr. Etienne M emin at The National Institute for Research in Computer Science and Control (Institut national de recherche en informatique et en automatique, INRIA) in Rennes, France, Ph.D. student Pierre Derian applied advanced optical flow as described in a paper by Corpetti et al. (2002) to the REAL data from CHATS. The results of the test are described in D erian et al. (2010) and Mayor et al. (2010).

To summarize the results, the French optical flow method applied has the advantage of providing a vector at every grid point regardless of the availability of aerosol features in the immediate vicinity of the grid point. (The application cross-correlation to entire scans often results in areas with

⁸Seventy-eight percent of the 500 m \times 500 m block was over orchard while the remaining northern 22% was over bare fields or relatively short-crops such as tomato plants. Sixty-eight percent of the 750 m \times 750 m block was over orchard and 32% over non-orchard. The 1 km \times 1 km block was 63% over orchard and 37% over non-orchard.

obviously bogus vectors in patches were the algorithm was not able to find aerosol features. This is especially true at far ranges.) The French optical flow method is a global (i.e. spanning the entire image area) numerical solution and therefore able to estimate velocities at all grid points. However, results presented in Dérian et al. (2010) and Mayor et al. (2010) suggest that the Corpetti et al. (2002) method provides a solution that is too smooth. In the future, we plan to test a newer method based on wavelets and described in the recent papers by Kadri-Harouna et al. (2012), Dérian et al. (2011), and Dérian et al. (2012).

8.5 Use of Multi-core and Graphical Processing Units (GPUs)

Without optimization for execution speed, the original program to compute aerosol feature velocity required approximately 4 days to calculate all of the velocities (4 blocks centered on tower) used in our study. In practice, we often applied the program to different periods of the data set concurrently to hasten the process. Doing so reduced production of the results to 24 hours. However, the approach did not take advantage of the parallel processing capabilities of CPUs and GPUs, respectively. Therefore, we explored and implemented methods to dramatically accelerate the execution of the program by using software that employs modern processors more efficiently. Since these efforts were helpful in producing results more quickly, and will likely be of use in real-time calculation of wind velocity fields in the future, we provide a brief overview of the steps taken to accelerate the execution.

The algorithms used to calculate wind velocity from lidar data were originally implemented in Interactive Data Language (IDL). IDL provides many array operations that streamline application development, and provides an easy means of processing many data files by running multiple instances of the application. However, some routines did not fully utilize multiple CPU cores for maximum performance. We decided to write our own routines that utilized multiple threads and the CPU cores' single instruction, multiple data (SIMD) operations. However, running multiple instances of the application to process multiple files simultaneously requires substantial processor resources and reduces the efficiency of multithreading on the CPU. We decided to investigate GPU computing as a viable method of accelerating application performance by offloading computationally-intensive algorithms to a graphics processor. Our goal was to accelerate the execution speed of a single instance of processing using one GPU, and to run multiple instances with multiple GPUs.

We profiled parts of our application to gauge which routines required the most execution time. Median filtering, polar-to-rectangular projection, and calculation of the CCF were the most time-consuming parts of our application. Fortunately, these routines map easily to data parallelism and thus are ideal for GPU computing. For this application we used a library developed by Tech-X Corp called GPULib. The library provides various array routines found in IDL that have been programmed to run on nVidia GPUs using Compute Unified Device Architecture (CUDA).⁹ This made it easy to rewrite parts of the original application to run on the GPU by keeping most of the programming in IDL. The version of the library used in this study, GPULib 1.4.4, did not have all of IDL's routines such as the median filter, which had to be created through CUDA and executed using IDL's CALL_EXTERNAL utility. We ran this application on a workstation with a six-core 3.33 GHz Intel Xeon 5680 and one nVidia Tesla C2070 card.

8.5.1 Median Filter

Median filtering is the most time-consuming part of the algorithm. Each beam of a scan must pass through a low-pass and a high-pass median filter. The computationally intensive part of the median filters are the sorting that must be performed for each new window position. As the

⁹<http://developer.nvidia.com/what-cuda>

width of the filter window increases, the execution time of the filter also increases. We initially leveraged the parallelism of the CPU by implementing the branchless vectorized median (BVM) algorithm developed by Marc Kachelrieß (Kachelrieß, 2009). The BVM was implemented in an external program that was coded in C and used Streaming SIMD Extensions (SSE) and OpenMP for vectorization and multithreading respectively. For the GPU version of BVM, we implemented the algorithm described in Chen et al. (2009) using CUDA. Filtering scans containing 150 radial beams of 7500 backscatter samples each, the CPU version of BVM implemented with OpenMP and SSE required 160-180 ms to complete. The GPU version of BVM using CUDA C completed filtering in ~ 90 milliseconds, making it almost twice as fast as the CPU version.

8.5.2 Polar-to-Rectangular Projection

The projection of the scan data onto a rectangular grid begins by calculating the polar coordinates of all the grid points. The polar coordinates are then used for the interpolants of the polar scan data, which are obtained using bilinear interpolation. The backscatter samples in a scan are uniformly spaced along the range of a beam, but the azimuthal angle of each beam is slightly different from the last. To accurately measure the interpolation values along the azimuth of the scan, each position of the rectangular grid must be checked to determine which pair of beams they lie between and to determine the distances to the beams. Finding which two beams an interpolant lies between is the most time-consuming part of the projection process. Initially, IDL was used but later we used GPULib to speed the process. The execution times for grid resolutions of 20 m, 10 m, and 8 m are lower on the CPU than on the GPU. At 6 m and 4 m, however, the GPU becomes much faster than the CPU. We notice that the execution times for the CPU start small but increase exponentially as grid spacing decreases. The GPU’s execution time is much longer than that of the CPU at 20 m, but does not increase as sharply as the CPU as the grid spacing decreases.

8.5.3 Cross Correlation

The fast Fourier transforms (FFT) are the most time-consuming part of the calculation of the CCF. Two-dimensional (2-D) FFTs of n -by- n matrices have a time complexity of $O(N^2 \log_2 N)$ which causes their execution time to increase exponentially as the matrix size increases. The CCF calculation involves two forward transforms and one inverse transform. This runs quickly with $250 \text{ m} \times 250 \text{ m}$ blocks, but becomes significantly slower using $1 \text{ km} \times 1 \text{ km}$ blocks. At a grid spacing of $4 \text{ m} \times 4 \text{ m}$ the CCF of a pair of $250 \text{ m} \times 250 \text{ m}$ blocks can be computed in less than 4 milliseconds using one CPU core, whereas the $1 \text{ km} \times 1 \text{ km}$ blocks require more than 250 milliseconds to be computed on one CPU core. Due to the inherent parallelism of FFT algorithms, libraries have been made that perform FFTs on GPUs such as nVidia’s CUFFT.¹⁰ GPULib provides a GPU-accelerated FFT routine with the same functionality as the one in IDL, making this the easiest optimization in our application. At block sizes of $250 \text{ m} \times 250 \text{ m}$ and $500 \text{ m} \times 500 \text{ m}$ the CPU is slightly faster than the GPU, especially at low resolutions. However, at larger blocks sizes and higher resolutions, the GPU can become several times faster than the CPU.

8.5.4 Discussion on use of GPSs versus multi-core CPUs

By using the optimizations described, we were able to process the entire CHATS data set consisting of 4 blocks centered over the tower in approximately 12 hours. Holding all else constant, except optimizations, the process requires almost 24 hours. The use of GPU computing in the processing of lidar data can accelerate execution speed over CPU based methods but, as described above, may result in slower execution for some cases. The larger the amount of data in a computation, the

¹⁰<http://developer.nvidia.com/cufft>

more efficient the GPU becomes at processing compared to the CPU. This is especially beneficial for the calculation of vector flow fields such as those shown in Figs. 19 and 20. For a pair of 60° wide PPI scans, we are able to compute approximately 5000 vectors using a $1 \text{ km} \times 1 \text{ km}$ block every 50 m to a range of 5 km within 10 s. This amount of time is smaller than the time required to collect the two scans thereby demonstrating the feasibility of providing two-dimensional horizontal vector flow fields in real-time.

9 Summary

The results presented show that the aerosol motion components derived from pairs of scans are in best agreement with the averaged sonic anemometer data when the mean wind speed and TKE are low; maximum of the CCFs are large; mean SNRs over the block areas are large; and the atmosphere tends toward stability. Use of larger blocks improved the agreement. Conversely, the agreement tends to become worse when the wind speed is strong and TKE are high; the maximum of the CCFs are small; the SNR is low; and the atmosphere is unstable. The study was limited to vectors resulting only from pairs of scans and single CCFs. Doing so resulted in vectors with the highest possible temporal resolution given the scan update rates used during CHATS. Had the ground beneath the lidar trailer been constantly firm and the tower been avoided by scanning just above it, cross-correlation functions could be averaged together as was done in Mayor and Eloranta (2001) to investigate mean velocities.

In the near future, we plan to implement an improved lidar control system to encourage long-term and unattended operation of the REAL. We also plan to compute vector flow fields in real-time and make them available via the web in near-real-time from our Department's server. We also hope to implement and test more advanced forms of optical flow at altitudes above the surface layer. Ideally, comparisons will be made with winds from a Doppler lidar at altitudes of 50 - 150 m AGL. More care will be taken to maintain constant platform attitude and avoidance of hard target reflections in the intercomparison area. Testing in a coastal environment may have significant merit given the current importance of developing improved observational methods for offshore wind resource assessments.

10 Acknowledgments

The author thanks Drs. Scott Spuler, Sara Tucker, Ron Schwiesow, and Alan Brewer for very helpful discussions about coherent Doppler lidar. The CHATS experiment was funded by the NCAR Director's Opportunity Fund. Analysis of the data and this paper was funded by Grant 0924407 from the National Science Foundation's Physical and Dynamic Meteorology Program. NCAR EOL In-situ Sensing Facility staff provided the tower data. NVIDIA Corporation provided a C2070 GPU card to perform the calculations. The author thanks Jennifer P. Lowe and Christopher F. Mauzey for their vital assistance in data processing and analysis.

References

- Adrian, R. J., 2007: Hairpin vortex organization in wall turbulence. *Phys. Fluids*, **19**, 041 301–1–041 301–16.
- Chen, W., M. Beister, Y. Kyriakou, and M. Kachelrieß, 2009: High performance median filtering using commodity graphics hardware. *IEEE Nuclear Science Symposium Conference Program*.
- Corpetti, T., E. Mémin, and P. Perez, 2002: Dense estimation of fluid flows. *IEEE Trans. Pattern Analysis Machine Intelligence*, **24**, 365–380.

- Dérian, P., P. Héas, C. Herzet, and E. Mémin, 2011: Wavelets to reconstruct turbulence multifractals from experimental image sequences. *7th Int. Symp. on Turbulence and Shear Flow Phenomena*, Vol. TSFP-7.
- Dérian, P., P. Héas, C. Herzet, and E. Mémin, 2012: Wavelets and optical flow motion estimation. *Numerical Mathematics Theory, Methods, and Applications*.
- Dérian, P., P. Héas, E. Mémin, and S. D. Mayor, 2010: Dense motion estimation from eye-safe aerosol lidar data. *Proceedings of the 25th International Laser Radar Conference*, Vol. 1, 377–380.
- Eloranta, E. W., J. M. King, and J. A. Weinman, 1975: The determination of wind speeds in the boundary layer by monostatic lidar. *J. Appl. Meteor.*, **14**, 1485–1489.
- Finnigan, J. J., R. H. Shaw, and E. G. Patton, 2009: Turbulence structure above a vegetation canopy. *J. Fluid Mech.*, **637**, 387–424.
- Frehlich, R., S. M. Hannon, and S. W. Henderson, 1997: Coherent doppler lidar measurements of winds in the weak signal regime. *Appl. Optics*, **36**, 3491–3499.
- Grund, C. J., R. M. Banta, J. L. George, J. N. Howell, M. J. Post, R. A. Richter, and A. M. Weickman, 2001: High-resolution Doppler lidar for boundary layer and cloud research. *J. Atmos. Ocean. Technol.*, **18**, 376–393.
- Hardesty, R. M., 1984: Coherent DIAL measurement of range-resolved water vapor concentration. *Appl. Optics*, **23**, 2545–2553.
- Henderson, S. W., C. P. Hale, J. R. Magee, M. J. Kavaya, and A. V. Huffaker, 1991: Eye-safe coherent laser radar system at 2.1 micron using tm,ho:yag lasers. *Opt. Lett.*, **16**, 773–775.
- Hill, M., R. Calhoun, H. J. S. Fernando, A. Wieser, A. Dörnbrack, M. Wiessmann, G. Mayr, and R. K. Newsom, 2010: Coplanar Doppler lidar retrieval of rotors from T – REX. *J. Atmos. Sci.*, **67**, 713–729.
- Hooper, W. P. and E. W. Eloranta, 1986: Lidar measurements of wind in the planetary boundary layer: the method, accuracy and results from joint measurements with radiosonde and kytoon. *J. Clim. Appl. Meteor.*, **25**, 990–1001.
- Horn, B. K. P. and B. G. Schunck, 1981: Determining optical flow. *Artificial Intelligence*, **17**, 185–203.
- Kachelrieß, M., 2009: Branchless vectorized median filtering. *IEEE Medical Imaging Conference Program*.
- Kadri-Harouna, S., P. Drian, P. Has, and E. Mmin, 2012: Divergence-free wavelets and high order regularization. *Int. J. Comput. Vision*.
- Kaimal, J. C. and J. J. Finnigan, 1994: *Atmospheric Boundary Layer Flows*. Oxford University Press, 200 Madison Ave., New York, NY 10016, 289 pp.
- Kolev, I., O. Parvanov, and B. Kaprielov, 1988: Lidar determination of winds by aerosol inhomogeneities: motion velocity in the planetary boundary layer. *Appl. Optics*, **27**, 2524–2531.
- Kunkel, K. E., E. W. Eloranta, and J. Weinman, 1980: Remote determination of winds, turbulence spectra and energy dissipation rates in the boundary layer from lidar measurements. *J. Atmos. Sci.*, **37**, 978–985.

- Mann, J., et al., 2009: Comparison of 3D turbulence measurements using three staring wind lidars and a sonic anemometer. *Meteorol. Z.*, **18**, 135–140.
- Mayor, S. D., 2011: Observations of seven density current fronts in Dixon, California. *Mon. Wea. Rev.*, **139**, 1338–1351.
- Mayor, S. D., P. Benda, C. E. Murata, and R. J. Danzig, 2008: Lidars: A key component of urban biodefense. *Biosecur. Bioterror.*, **6**, 45–56, DOI: 10.1089 bsp.2007.0053.
- Mayor, S. D., P. Dérian, P. Héas, and E. Mémin, 2010: Two-component horizontal motion vectors from scanning eye-safe aerosol lidar. *19th Symp. on Boundary Layers and Turbulence*, AMS.
- Mayor, S. D. and E. W. Eloranta, 2001: Two-dimensional vector wind fields from volume imaging lidar data. *J. Appl. Meteor.*, **40**, 1331–1346.
- Mayor, S. D., D. H. Lenschow, R. L. Schwiesow, J. Mann, C. L. Frush, and M. K. Simon, 1997: Validation of near 10.6 micron Doppler lidar radial velocity measurements and comparison with a 915-mhz profiler. *J. Atmos. Ocean. Technol.*, **14**, 1110–1126.
- Mayor, S. D., J. P. Lowe, and C. F. Mauzey, 2012: Two-component horizontal aerosol motion vectors in the atmospheric surface layer from a cross-correlation algorithm applied to elastic backscatter lidar data. *J. Atmos. Ocean. Technol.*, **29**, In Press.
- Mayor, S. D. and S. M. Spuler, 2004: Raman-shifted Eye-safe Aerosol Lidar. *Appl. Optics*, **43**, 3915–3924.
- Mayor, S. D., S. M. Spuler, B. M. Morley, and E. Loew, 2007: Polarization lidar at 1.54-microns and observations of plumes from aerosol generators. *Opt. Eng.*, **46**, DOI: 10.1117/12.781902.
- Newsom, R., R. Calhoun, D. Ligon, and K. J. Allwine, 2008: Linearly organized turbulence structures observed over a suburban area by dual-doppler lidar. *Bound. Layer Meteorol.*, **127**, 111–130.
- Newsom, R. K. and R. M. Banta, 2004a: Assimilating coherent Doppler lidar measurements into a model of the atmospheric boundary layer. Part I: Algorithm development and sensitivity to measurement error. *J. Atmos. Ocean. Technol.*, **21**, 1328–1345.
- Newsom, R. K. and R. M. Banta, 2004b: Assimilating coherent Doppler lidar measurements into a model of the atmospheric boundary layer. Part II: Sensitivity analyses. *J. Atmos. Ocean. Technol.*, **21**, 1809–1824.
- Newsom, R. K., D. Ligon, R. Calhoun, R. Heap, E. Cregan, and M. Princevac, 2005: Retrieval of microscale wind and temperature fields from single- and dual-Doppler lidar data. *J. Appl. Meteorol.*, **44**, 1324–1345.
- Patton, E. G., et al., 2011: The Canopy Horizontal Array Turbulence Study (CHATS). *Bull. Amer. Meteor. Soc.*, **92**, 593–611.
- Pearson, G., F. Davies, and C. Collier, 2009: An analysis of the performance of the UFAM pulsed Doppler lidar for observing the boundary layer. *J. Atmos. Ocean. Technol.*, **26**, 240–250.
- Piironen, A. K. and E. W. Eloranta, 1995: Accuracy analysis of wind profiles calculated from volume imaging lidar data. *J. Geophys. Res.*, **100**, 25 559–25 567.
- Ponsardin, P. L., C. S. Kletecka, R. D. Babnick, K. Krubsack, S. D. Mayor, and S. M. Spuler, 2006: Autonomous eye-safe lidar for continuous monitoring of atmospheric aerosols. *International Symposium on Spectral Sensing Research*.

- Post, M. J. and R. E. Cupp, 1990: Optimizing a pulsed Doppler lidar. *Appl. Optics*, **29**, 4115–4158.
- Ruffier, F. and N. Franceschini, 2005: Optic flow regulation: the key to aircraft automatic guidance. *Robot. Auton. Syst.*, **50**, 177–194.
- Rye, B. J. and R. M. Hardesty, 1993a: Discrete spectral peak estimation in incoherent backscatter heterodyne lidar. I. Spectral accumulation and the Cramer-Rao lower bound. *IEEE Trans. Geosci. Remote Sensing*, **31**, 16–27.
- Rye, B. J. and R. M. Hardesty, 1993b: Discrete spectral peak estimation in incoherent backscatter heterodyne lidar. II. Correlogram accumulation. *IEEE Trans. Geosci. Remote Sensing*, **31**, 28–35.
- Sasano, Y., H. Hirohara, T. Yamasaki, H. Shimizu, N. Takeuchi, and T. Kawamura, 1982: Horizontal wind vector determination from the displacement of aerosol distribution patterns observed by a scanning lidar. *J. Appl. Meteor.*, **21**, 1516–1523.
- Sathe, A. and J. Mann, 2012: Measurement of turbulence spectra using scanning pulsed wind lidars. *J. Geophys. Res.*, **117**, D01 201, DOI:10.1029/2011JD016786.
- Sathe, A., J. Mann, J. Gottschall, and M. S. Courtney, 2011: Can wind lidars measure turbulence? *J. Atmos. Ocean. Technol.*, **28**, 853–868.
- Schols, J. L. and E. W. Eloranta, 1992: The calculation of area-averaged vertical profiles of the horizontal wind velocity from volume imaging lidar data. *J. Geophys. Res.*, **97**, 18 395–18 407.
- Shaw, R. H., Y. Brunet, J. J. Finnigan, and M. R. Raupach, 1995: A wind tunnel study of air flow in waving wheat: two-point velocity statistics. *Bound. Layer Meteor.*, **76**, 349–376.
- Spuler, S. M. and S. D. Mayor, 2005: Scanning eye-safe elastic backscatter lidar at 1.54 microns. *J. Atmos. Ocean. Technol.*, **22**, 696–703.
- Spuler, S. M. and S. D. Mayor, 2007a: Eye-safe aerosol lidar at 1.5 microns: progress towards a scanning lidar network. *SPIE Lidar Remote Sensing for Environmental Monitoring VIII*, San Diego, CA, DOI: 10.1117/12.739519.
- Spuler, S. M. and S. D. Mayor, 2007b: Raman shifter optimized for lidar at 1.5-micron wavelength. *Appl. Optics*, **46**, 2990–2995.
- Sroga, J. T., E. W. Eloranta, and T. Barber, 1980: Lidar measurements of wind velocity profiles in the boundary layer. *J. Appl. Meteor.*, **19**, 598–605.
- Su, H.-B., R. H. Shaw, and K. T. Paw-U, 2000: Two-point correlation analysis of neutrally stratified flow within and above a forest from large-eddy simulation. *Bound. Layer Meteor.*, **94**, 423–460.
- Sullivan, P. P. and E. G. Patton, 2011: The effect of mesh resolution on convective boundary layer statistics and structures generated by large-eddy simulation. *J. Atmos. Sci.*, **68**, 2395–2415.