

## Two-Component Horizontal Aerosol Motion Vectors in the Atmospheric Surface Layer from a Cross-Correlation Algorithm Applied to Scanning Elastic Backscatter Lidar Data\*

SHANE D. MAYOR, JENNIFER P. LOWE, AND CHRISTOPHER F. MAUZEY

*Department of Physics, California State University Chico, Chico, California*

(Manuscript received 26 December 2011, in final form 11 April 2012)

### ABSTRACT

Two-component horizontal motion vectors of aerosol features were calculated by applying a cross-correlation algorithm to square image blocks extracted from consecutive pairs of elastic backscatter lidar scans. The resulting vector components were compared with corresponding horizontal wind components from tower-mounted sonic anemometers located at the center of the image blocks. In the analysis 180 245 pairs of vectors derived from 75 days of field data collected between 19 March and 11 June 2007 were used. Examples of time series comparisons from 4-h periods during light, strong, and changing wind conditions are presented. Mean signal-to-noise ratios (SNRs) of the block backscatter data, maxima of the cross-correlation functions (CCFs), observed wind speed, and turbulent kinetic energy (TKE) were also calculated for each velocity component comparison. The correlation between the lidar-derived motion components and sonic anemometer wind components tends to be highest during light wind conditions with low TKE. An empirical relationship is presented that enables the elimination of vectors that are likely to be significantly different than the anemometer measurement. When applied to the entire set of scans available, this quality control (QC) method increases the correlation between the two forms of measurements. Finally, the cross-correlation algorithm and QC method are applied to a mesh of locations over pairs of scans. Two examples of two-dimensional and two-component vector flow fields are shown. In one case, the flow field reveals a rotational circulation associated with a vortex and, in the other case, convergence and transport near the leading edge of a density current front.

### 1. Introduction

Remote measurements of the vector wind field in the lower atmospheric boundary layer from ground-based platforms in the range of 1–20 km away from areas of interest are likely to be of value in several applications. Examples include wind resource assessments (especially offshore), very short-term wind predictions near established wind farms, initial transport and hazardous materials dispersion observations (i.e., over nuclear power and industrial chemical sites), wind shear and wake

vortex detection near airports, and meso- and micro-scale meteorological research. Doppler lidars provide high-quality direct measurements of only the radial component of motion (Mann et al. 2009). However, in many of the above applications it may not be practical to 1) collect 360° azimuth scans or assume horizontal homogeneity of the atmospheric boundary layer (Browning and Wexler 1968), or 2) combine radial measurements with numerical flow retrieval models to derive two or more wind components from a single Doppler lidar (Newsom and Banta 2004a,b). The use of two Doppler lidars separated by distances similar to the range of interest can provide multiple wind components over a common area (Newsom et al. 2005), but using two instruments increases the cost. Instead, a direct observation of two or more components of the vector wind field from scanning over a sector with a single instrument is desired. In this case, the application of the cross-correlation technique to aerosol backscatter lidar images may hold significant value.

---

\* Supplemental information related to this paper is available at the Journals Online website: <http://dx.doi.org/10.1175/JTECH-D-11-00225.1>.

---

*Corresponding author address:* Shane D. Mayor, Department of Physics, California State University Chico, 400 West First St., Chico, CA 95929.  
E-mail: [sdmayor@csuchico.edu](mailto:sdmayor@csuchico.edu)

The cross-correlation technique has been applied 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). The most recent of these papers (Mayor and Eloranta 2001) showed two-component vector fields with 250-m horizontal resolution over areas as large as 60 km<sup>2</sup>. These vectors, however, were the result of averaging many cross-correlation functions over time, up to 41 min in one case. 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 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 pairs of consecutive scans. The data were collected with an eye-safe elastic lidar operating at 1.54- $\mu\text{m}$  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 3 months in the atmospheric surface layer over an agricultural area. We avoid comparison of mean remote and in situ velocity components because of 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 from pairs of consecutive scans. We conclude by showing two examples of two-dimensional flow fields that result from the application of the algorithm to the entire scan area.

## 2. Experiment

The lidar system used in this experiment, known as Raman-shifted Eye-safe Aerosol Lidar (REAL), is described by Mayor and Spuler (2004), Spuler and Mayor (2005), and Mayor et al. (2007). It is a ground-based, scanning, elastic backscatter lidar operating at 1.54- $\mu\text{m}$  wavelength. Table 1 lists the specifications of the system as configured for the experiment.

The data used in this study 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 National Center for Atmospheric Research (NCAR) Integrated Surface Flux Facility (ISFF) 30-m vertical tower (VT; see Fig. 1). 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

TABLE 1. Specifications of the lidar system.

Wavelength	1.543 $\mu\text{m}$
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 diameter	40 cm
Receiver field of view	0.54 mrad (full angle)
Digitizer speed	100 MSPS*
Digitizer range	14 bits
Detector type	200- $\mu\text{m}$ InGaAs APD**

\* Megasamples per second

\*\* 200- $\mu\text{m}$ -diameter indium gallium arsenide avalanche photodiode

flow. The VT supported 13 Campbell Scientific CSAT3 3D 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 (PPI) scans] that were collected over a wide area surrounding the orchard. The laser pulses, which were projected from 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 18–20 m AGL. This corresponds to a slope of 8.6 m km<sup>-1</sup>. 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.

We estimate that the majority of PPI cross sections intersected the VT between 18- and 20-m height (see Fig. 2). 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.<sup>1</sup>

<sup>1</sup> Because of the trenching and flooding of a nearby irrigation ditch after the installation of the lidar, one side of the lidar trailer occasionally stood in mud while the other side was on dry soil. Staff attempted to compensate for the sinking of the one side by periodically adjusting the leveling system. However, measurements and records were not kept. Note: a 1-mm change in 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.




FIG. 1. Plan view of the experimental area. The yellow shaded regions from  $150^{\circ}$ – $210^{\circ}$  and  $175^{\circ}$ – $185^{\circ}$  azimuth represent the areas covered by wide and narrow PPI scans, respectively. The VT was located 1.61 km directly south of REAL. The white squares centered on the VT represent the outer edges of the image blocks extracted from the gridded PPI scan data that were used to compute motion vectors via the cross-correlation technique.

At an azimuthal scan rate of  $4^{\circ} \text{ s}^{-1}$  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 tower,<sup>2</sup> 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 not always 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

<sup>2</sup> The central column of the tower was 32 cm wide. Guy wires 0.47 cm in diameter were attached to the tower at four 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.




FIG. 2. Diagram showing the approximate altitude and diameter of REAL laser pulses with respect to the treetops and the vertical tower at CHATS. This diagram is an east–west cross section looking either toward or away from the lidar. Laser pulses from one scan at an azimuthal scan rate of  $4^{\circ} \text{ s}^{-1}$  are shown (shaded circles). At this scan rate, the pulses were spaced 11 m apart at 1.61-km range from the lidar.

channels enabled better quantization of the signal. This configuration likely resulted in a higher signal-to-noise ratio (SNR) in the first few kilometers of range, but a 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.

The SNR of the backscatter data was calculated by first computing the standard deviation ( $\sigma_b$ ) 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;  $\sigma_b$  is a measure of the intensity of the electronic noise in the 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 is divided by the noise value, as shown in Eq. (1), below:

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

This procedure is applied to the return of each pulse individually and the resulting arrays are not averaged over time. Figure 3 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 the 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




FIG. 3. 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<sup>2</sup>) block used to compute horizontal velocity vectors. The ISFF VT was located at 1.6-km range.

slowly (2.5–5.0 mJ day<sup>-1</sup>) during operation until operators either increased flash lamp voltage (every 5–10 days) or replaced the flash lamps (approximately every 21 days). Figure 7 of Spuler and Mayor (2007) 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. 3 represents the broad range of SNRs for the scans used to extract aerosol motion vectors reported herein.

### 3. Algorithm

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.<sup>3</sup>

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 treetops and within the time span of each raw data file. A single raw data file typically contains

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

dozens to hundreds of scans and includes data spanning periods from minutes to more than an hour in most cases.

Next, subsets of the gridded data are extracted in square regions corresponding to the four blocks shown in Fig. 1. 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 those 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 (CCFs) are computed using fast Fourier transforms (FFTs) 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. To improve the velocity resolution, a two-dimensional polynomial based on the  $5 \times 5$  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 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.<sup>4</sup> 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 affect the mean SNR.<sup>5</sup> 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 representative of the airflow 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 10–30 s) in comparison to the sonic anemometer data. Therefore, substantial time averaging of the sonic anemometer wind components is required to make a comparison with each aerosol motion vector.

As a first guess for a suitable time interval over which to average the sonic anemometer data, 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 4 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.

REAL operated at CHATS with a constant 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 of which 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 collected PPI scans were intended to reveal a broader view of the experimental area. These are called “wide” and ranged from approximately  $150^\circ$  to  $210^\circ$  azimuth and were performed at a rate of  $4^\circ\text{ s}^{-1}$ . Another large percentage of the PPI scans (approximately 75 000) were termed “narrow” and were intended to probe the finescale structure and motion of the atmosphere in the immediate vicinity of the tower. The narrow scans ranged from  $175^\circ$  to  $185^\circ$  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 ( $dt$ ) was 30 s. The second most frequent was 17 s. Scans with update periods of 45 and 10 s were also collected occasionally. In summary,

<sup>4</sup> The 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.

<sup>5</sup> Given a pulse rate of 10 Hz, an angular scan rate of  $4^\circ\text{ s}^{-1}$ , and a digitizer speed of 100 megasamples per second, the number of data points in the native spherical coordinate system of the lidar fall into blocks that are  $250\text{ m} \times 250\text{ m}$ ,  $500\text{ m} \times 500\text{ m}$ ,  $750\text{ m} \times 750\text{ m}$ , and  $1\text{ km} \times 1\text{ km}$ , 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.




FIG. 4. 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).

the population of PPI scans available for analysis has 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. However, it also provides an opportunity to explore the effects of adjustable parameters so that they can be optimized for best results in future experiments.

#### 4. 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, such as those shown in Figs. 5–7 were created for the entire dataset. We have examined them all and concluded that the correlation 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 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.

##### a. Light and variable wind example

Figure 5 shows the velocity components derived from the lidar data (black dots) and the averaged sonic anemometer data from five heights on the VT during a weakly stable 4-h period (1000–1400 UTC 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 four 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-h period, according to the sonic anemometers on the VT, ranged from  $1.3 \text{ m s}^{-1}$  at 12.5 m AGL to  $2.25 \text{ m s}^{-1}$  at 29 m. The mean turbulent kinetic energy ranged from 0.04 to  $0.08 \text{ m}^2 \text{ s}^{-2}$  among the five sonic anemometers. The mean gradient Richardson number was 0.02, indicating weak static stability. The mean lidar SNR over the  $250 \text{ m} \times 250 \text{ m}$  block ranged from 25 to 100, with a 4-h 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




FIG. 5. Time series of lidar-derived (black points) and averaged sonic anemometer (color traces) velocity components for a 4-h period during light wind conditions. A  $250 \text{ m} \times 250 \text{ m}$  block size was used for the lidar-derived velocity estimates.

this period range from 0.70 to 0.79, with the strongest correlation corresponding to the anemometer at 23 m AGL.

*b. Strong wind example*

Figure 6 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 five heights on the VT (colored lines) during a 4-h period

with strong north-northwest winds (2000–2400 UTC 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  $8 \text{ m s}^{-1}$  at 12.5 m AGL to  $12.2 \text{ m s}^{-1}$  at 29 m AGL for the 4-h 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




FIG. 6. As in Fig. 5, but during strong wind conditions. A  $1\text{-km}^2$  block size was used for the lidar-derived velocity estimates.




FIG. 7. As in Fig. 5, but during moderate wind conditions with a frontal passage. A  $1\text{-km}^2$  block size was used for the lidar-derived velocity estimates.

AGL. The mean gradient Richardson number was  $-0.005$ , indicating slightly unstable to nearly neutral conditions.

The lidar collected wide PPI scans every 17 s during this period. The mean SNR over the  $1\text{-km}^2$  block ranged from 50 to 75 with only a 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. 6) suggest the algorithm has captured the strong mean flow for this period and the larger time-scale  $u$ -component fluctuations on the order of 15 min or more. However, that is not to say that the technique is not capable of resolving sudden changes of the mean flow




FIG. 8. As in Fig. 7, but for wind speed and direction for the same 4-h period.



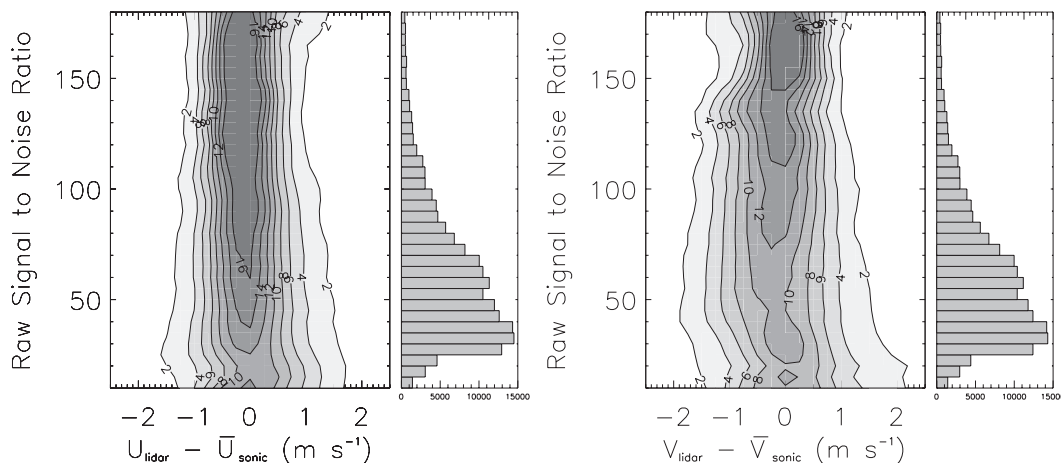


FIG. 9. Distributions of velocity component differences as a function of average SNR in the blocks used to derive the lidar vectors. Contour interval labels are in percent. The bar chart to the right of each contour plot shows the number of data points in each SNR bin interval. The width of the distribution of velocity component differences tends to decrease as the SNR increase.

during turbulent conditions. An abrupt change in the  $v$  component of the flow can be seen at 1330 UTC in Fig. 5. In the next section we show an example of a near reversal in wind direction that occurred over a 10-min period with the passage of a density current front.

*c. Changing wind during the passage of a density current front*

Figure 7 shows the velocity component time series comparisons from the afternoon of 26 April 2007 between 2130 UTC 26 April and 0130 UTC 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  $151^\circ$  and  $211^\circ$  azimuth at a scan rate of  $4^\circ \text{ s}^{-1}$ . During this period, a density current front passed over the experimental site at 2325 UTC 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 from  $-0.5$  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  $6 \text{ m s}^{-1}$  and decreased until the front passed. The wind direction veered dramatically from  $350^\circ$  (north) before the front to  $221^\circ$  (south-southwest) after the front. The scatter in both forms of wind measurement decreases

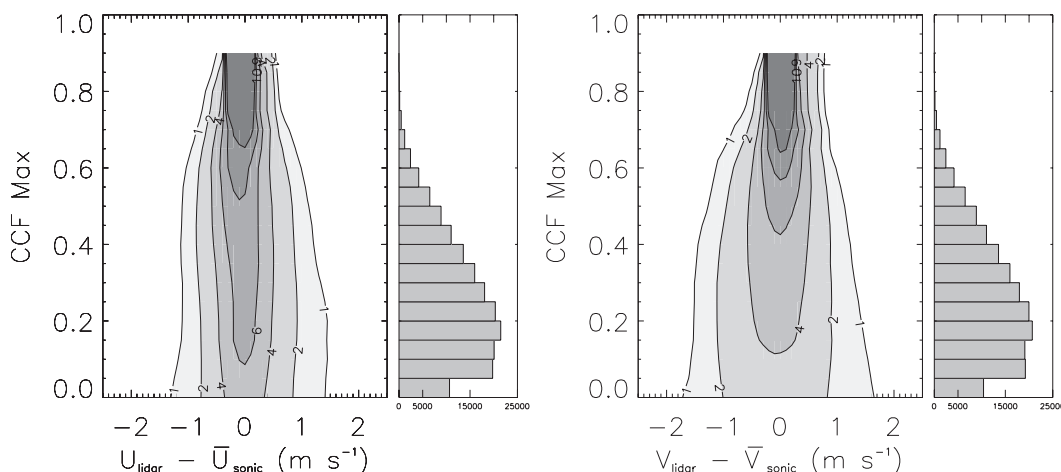


FIG. 10. Distributions of velocity component differences as a function of the maximum of CCF corresponding to each lidar-derived vector. The bar chart to the right of each contour plot shows the number of comparison points in each CCF maximum bin interval. The width of the distribution of velocity component differences tends to decrease as the CCF maxima increase.

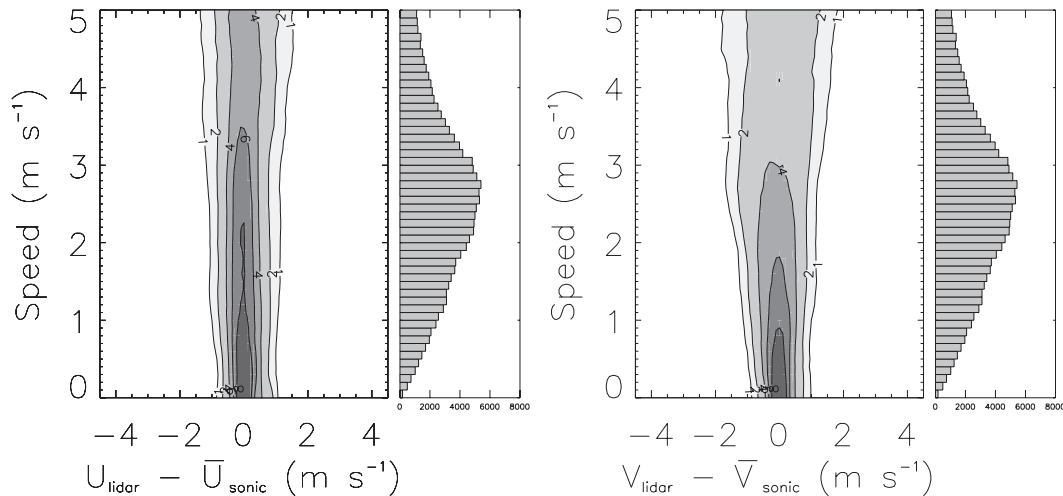


FIG. 11. Distributions of velocity component differences as a function of the corresponding observed wind speed as measured by the sonic anemometer at 12.5 m on the ISFF VT. The bar chart to the right of each contour plot shows the number of comparison points in each wind speed interval. The width of the distribution of velocity component differences tends to increase as the wind speed increases.

with time, because the atmosphere becomes more stable and the turbulence intensity decreases. Figure 8 presents the same data shown in Fig. 7 except as speed (top panel) and direction (bottom panel).

#### d. Velocity component difference distributions

Figures 9–12 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. 9–12 correspond to the east–west velocity component differences ( $U_{\text{lidar}} - \bar{U}_{\text{sonic}}$ ), and the right panels in Figs. 9–12

correspond to the north–south velocity component differences ( $V_{\text{lidar}} - \bar{V}_{\text{sonic}}$ ). Figures 9 and 10 show that the velocity component differences tend to decrease as the mean SNR and CCF maxima increase. This is an encouraging result because 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 11 and 12 confirm that the velocity component differences tend to increase as the observed wind speed and TKE increase.

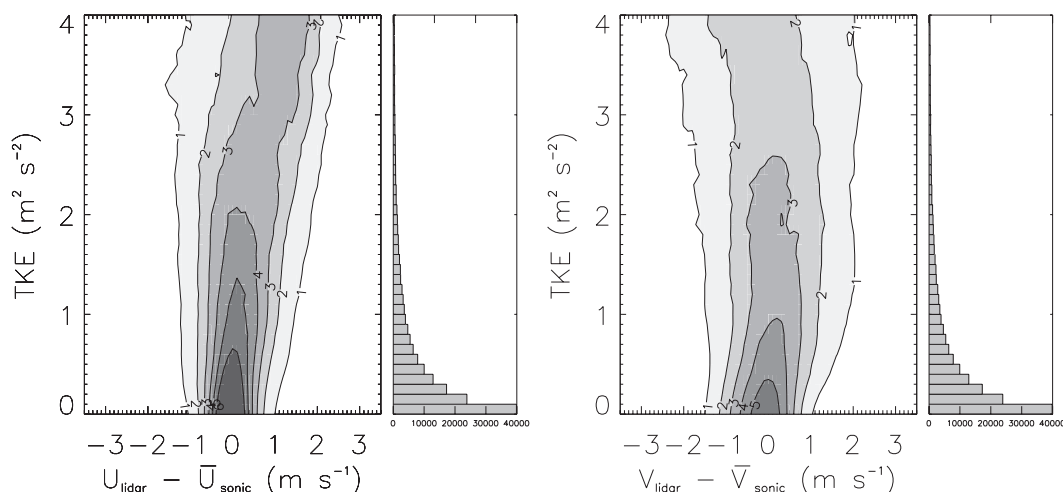


FIG. 12. Distributions of velocity component differences as a function of TKE calculated from sonic anemometer data at 12.5 m on the ISFF VT. The bar chart to the right of each contour plot shows the number of comparison points in each TKE interval. The width of the distribution of velocity component differences tends to increase as the TKE increases.




FIG. 13. Distributions of 18 m AGL sonic anemometer wind components vs aerosol motion components derived from the lidar data for 180 245 points using 1-km<sup>2</sup> 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 nonuniform distribution of wind speeds and directions during the experiment shown in Fig. 14.

*e. Lidar velocity components versus sonic anemometer velocity components*

Figure 13 shows the distributions of 180 245 pairs of velocity components. The lidar-derived components resulted from 1-km<sup>2</sup> blocks and 17- and 30-s scan update rates. The data points were accumulated into bins of 0.2 m s<sup>-1</sup> × 0.2 m s<sup>-1</sup> resolution. The numbers above the grayscale 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 grayscale. We chose to shade bins containing as few as one 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<sup>-1</sup>. Similarly, a broad distribution of lidar-derived *v* components exists when the sonic anemometer *v* components exceed ±2 m s<sup>-1</sup>. These distributions are attributed to the prevailing north-south flow at the site during the field experiment. Figure 14 shows the distribution of observed wind speeds and directions that occurred while PPI scans were being collected during CHATS. The most frequently occurring airflow was from 180° to 200° at less than 4 m s<sup>-1</sup>. The second most frequently occurring flow regime was from 330° to 360° between 2 and 6 m s<sup>-1</sup>. Weak westerly flow was less common and easterly flow was rare. As a result of this nonuniform distribution of wind velocities, we randomly sampled the population of velocity differences to achieve a uniform distribution in wind direction. Figure 15 shows this distribution. The linear correlation coefficients are 0.74 and 0.89 for the *u* (left panel) and *v* (right panel) components, respectively.

**5. 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




FIG. 14. 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.

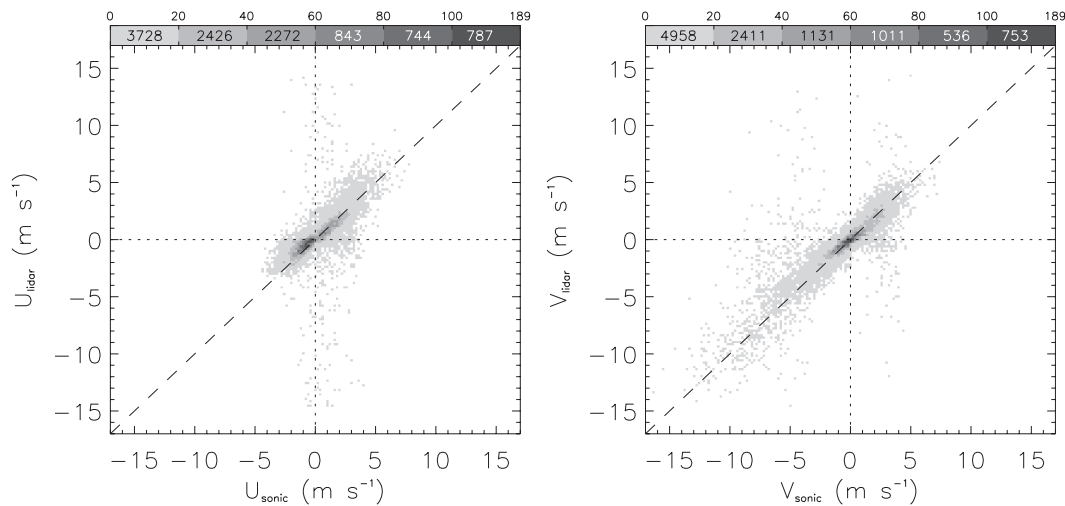


FIG. 15. Distribution of 18 m AGL sonic anemometer wind components vs motion components derived from the lidar data after random sampling to account for the wind distribution shown in Fig. 14.

not have anemometer data for validation. This allows us to create two-dimensional flow fields with the low-quality vectors removed.

#### a. Creating the quality control model

The data used to create the QC model were sampled from the lidar-derived velocity components using 1-km<sup>2</sup> blocks from scans with 17- and 30-s update rates. To eliminate the biasing effects of the uneven distribution of wind directions at the CHATS site (as shown in Fig. 14), 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 component,  $|U_l - U_s|$  and  $|V_l - V_s|$ , respectively. Predictor variables are the corresponding wind speed ( $W$ ), SNR ( $S$ ), and the maximum of the CCF ( $C$ ).

Multiple linear regression was applied on 1) first-degree predictor variables,  $W$ ,  $S$ , and  $C$ ; 2) first-degree and interaction predictor variables (e.g.,  $WS$ ,  $WSC$ ); and 3) first-degree, interaction, and second-degree predictor variables (e.g.,  $W^2$ ,  $S^2$ ). 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 captures 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 Eqs. (2) and (3):

$$\begin{aligned} |U_l - U_s| &= 1.066 - 0.5026C + 0.05059W^2 + 0.00003382S^2 + 3.649C^2 - 0.002577WS \\ &\quad - 0.09743WC - 0.02463SC + 0.008761WSC, \end{aligned} \quad (2)$$

$$\begin{aligned} |V_l - V_s| &= 1.156 + 0.003156S - 2.875C + 0.02572W^2 + 0.00002885S^2 + 5.576C^2 - 0.00139WS \\ &\quad - 0.4965WC - 0.0249SC + 0.004234WSC. \end{aligned} \quad (3)$$

#### b. Testing the quality control model

The testing set for each of the 20 random samples was the complete dataset with the training set removed.

Model testing was performed by removing model-predicted differences greater than 2 m s<sup>-1</sup> 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

TABLE 2. Mean adjusted  $R^2$  values for regression models.

	First degree	First degree with interactions	First and second degrees with interactions
$u$	0.22	0.36	0.43
$v$	0.18	0.26	0.30

removal. The  $u$  model fails to identify 49.8% of differences greater than  $2 \text{ m s}^{-1}$ . 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  $2 \text{ m s}^{-1}$ . Figure 16 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  $2 \text{ m s}^{-1}$ . Correlation between the sonic and lidar  $u$  measurements for the complete dataset is 0.52; correlation between the  $u$  measurements after applying the QC model is 0.75. Correlation between the sonic and lidar  $v$  measurements for the complete dataset is 0.84; correlation between the  $v$  measurements after applying the QC model is 0.90.

**6. Flow fields**

In addition to computing vectors for blocks centered on the tower location for time series comparisons as described in section 4, 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 are spatially dense, with vectors at the grid spacing. Neighboring vectors in such cases have a large degree of common heritage. For example, applying a  $1 \text{ km} \times 1 \text{ km}$  block on data with grid spacing of  $10 \text{ m} \times 10 \text{ m}$ , 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% 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 microscale circulations. We show two cases to demonstrate this. We also apply the QC method described in section 5 to eliminate vectors that are likely to be in significant error.

*a. Cellular surface layer convection*

Time-lapse animations (see the following: 2000–2359 UTC 25 March 2007, <http://dx.doi.org/10.1175/JTECHD1100225.s1> and 0000–0359 UTC 26 March 2007, <http://dx.doi.org/10.1175/JTECHD1100225.s1>) 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. 17, our estimate of the two-component horizontal wind field at 0019:47 UTC, was calculated using  $1 \text{ km} \times 1 \text{ km}$  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




FIG. 16. Distribution of 18 m AGL sonic anemometer wind components vs motion components derived from the lidar data after application of the quality control algorithm. (left) The linear correlation coefficient for the  $u$ -component comparisons is 0.75. (right) The linear correlation coefficient for the  $v$ -component comparisons is 0.90.

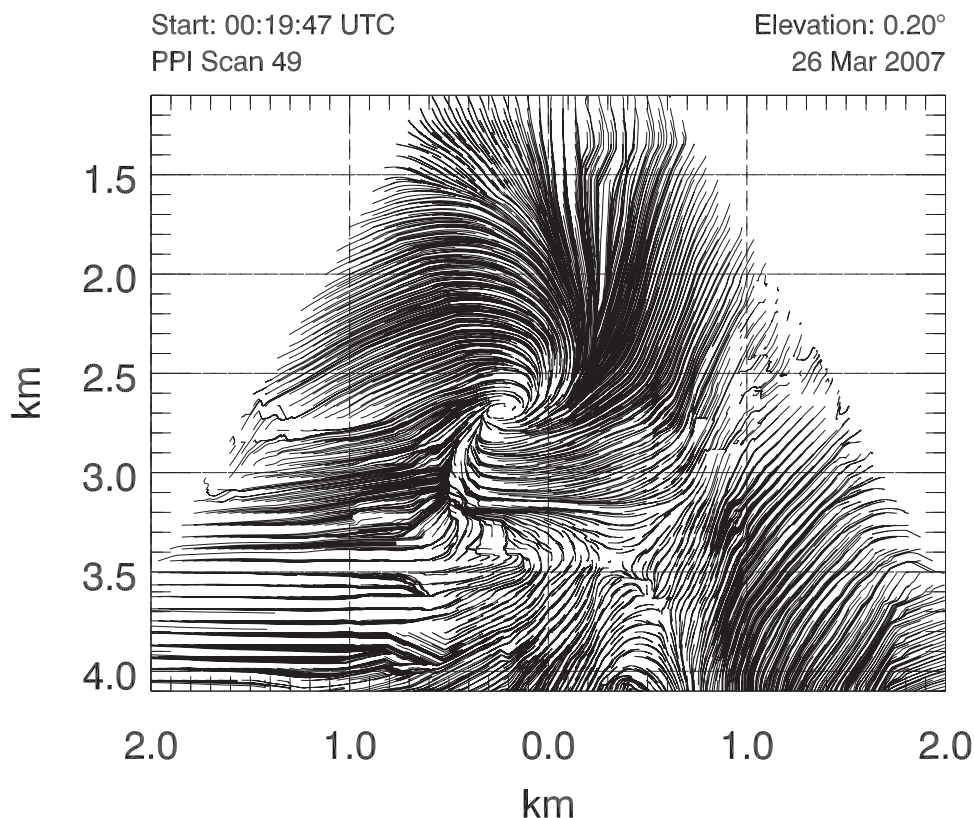


FIG. 17. Streamlined flow field resulting from the application of the cross-correlation algorithm to a pair of PPI scans through a convective afternoon atmospheric surface layer when winds were light and variable. Scans were separated in time by 17 s. A block size of  $1 \text{ km} \times 1 \text{ km}$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

by a particle trace procedure in Interactive Data Language (IDL) with second-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 0130 UTC, when a uniform west-southwest flow swept across the region.

#### *b. 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, ending at 2308:58 UTC 26 April 2007 (a time-lapse animation of this event covering the 3-h period between 2200 UTC 26 April 2007 and 0100 UTC 27 April 2007 is available online at <http://dx.doi.org/10.1175/JTECHD1100225.s1>). A block size of  $1 \text{ km} \times 1 \text{ km}$  was used and the scans were separated by 30 s. Figure 18 reveals the flow field when the front approximately bisected the scan 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 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.

## 7. Discussion

### *a. 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 because of stratification and gravity




FIG. 18. Streamlined flow field resulting 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 northerly flow north of the front (blue) and southerly flow south of the front (red) indicated. 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.

waves with propagation vectors that may be significantly different from the local wind vector. The results 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.

#### *b. 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 predominant coherent features. 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 may 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).

#### *c. 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 dataset 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) a walnut orchard and 2) either bare field or short crops, such as tomato plants (depending on the season). The  $250 \text{ m} \times 250 \text{ m}$  block was entirely over a walnut orchard. However, for all of the larger blocks, only the southern portions were over the orchard while northern portions were over

either bare field or short crops.<sup>6</sup> 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 was 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 reveals moving point sources of particulate matter from this roadway. Agricultural activities (e.g., plowing, spraying, and harvesting) also resulted in significant moving point sources of dust.

## 8. Summary

The results presented show that the aerosol motion components derived from pairs of consecutive scans are in best correlation with the averaged sonic anemometer data when the mean wind speed and TKE are low, maximum of the CCFs are large, means of the SNR over the block areas are large, and the atmosphere tends toward stability. The use of larger blocks also improved the agreement. Conversely, the correlation tends to become worse when the wind speed is strong and TKE is high, the maximum of the CCFs is small, the SNR is low, and the atmosphere is unstable. We chose to limit our study to vectors only resulting from pairs of consecutive scans and single CCFs. Doing so results 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, as was done in Mayor and Eloranta (2001) to investigate mean velocities.

In the future, we recommend tests of the technique with more care taken to maintain constant platform attitude, verification of the laser beam location near the reference wind measurement, and avoidance of hard-target reflections in the block areas. We also recommend evaluation of the technique at altitudes above the

surface layer. Testing in a coastal environment may have significant merit given the current importance of developing improved observational methods for offshore wind resource assessments.

*Acknowledgments.* This project 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.

## APPENDIX

### Use of Multicore and Graphical Processing Units

#### a. Introduction

Without optimization for execution speed, the original program to compute aerosol feature velocity required approximately 4 days to calculate all of the velocities (four blocks centered on the tower) used in our study. In practice, we often applied the program to different periods of the dataset concurrently to hasten the process. Doing so reduced production of the results to 24 h. However, the approach did not take advantage of the parallel processing capabilities of CPUs and graphics processing units (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. Because 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 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. Therefore, 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. Therefore, 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

---

<sup>6</sup> Of the 500 m × 500 m block 78% was over orchard while the remaining northern 22% was over bare fields or relatively short crops, such as tomato plants. Of the 750 m × 750 m block 68% was over orchard and 32% was over nonorchard. The 1 km × 1 km block was 63% over orchard and 37% over nonorchard.



processing using one GPU, and to run multiple instances with multiple GPUs.

### b. Methodology

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. 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). 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 version 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.

#### 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 is the sorting that must be performed for each new window position. As the 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 Kachelriess (2009). The BVM was implemented in an external program that was coded in C and used Streaming SIMD Extensions (SSE) and Open Multiprocessing (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 ms, making it almost twice as fast as the CPU version.

#### 2) POLAR-TO-RECTANGULAR PROJECTION

The projection of the scan data onto a rectangular grid begins by converting the grid's coordinates from a Cartesian system to a polar system with its origin at the location of the lidar. The polar coordinates are then used to compute interpolants of the scan data 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, 10, and 8 m are lower on the CPU than on the GPU. At 6 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.

#### 3) CROSS CORRELATION

FFTs are the most time-consuming part of the calculation of the CCF. Two-dimensional (2D) FFTs of  $n \times n$  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 ms using one CPU core, whereas the  $1 \text{ km} \times 1 \text{ km}$  blocks require more than 250 ms to be computed on one CPU core. Because of the inherent parallelism of FFT algorithms, libraries have been made that perform FFTs on GPUs such as nVidia's CUDA FFT library (CUFFT). 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.

### c. Discussion

By using the optimizations described, we were able to process the entire CHATS dataset consisting of four blocks centered over the tower in approximately 12 h. Holding all else constant, except optimizations, the process requires almost 24 h. 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. As the size and resolution of the blocks increases, processing them on the GPU becomes more efficient than on the CPU. This is especially beneficial for the calculation of vector flow fields, such as those shown in Figs. 17 and 18. For a pair of 60°-wide PPI scans, we are able to compute approximately 5000 vectors using a 1 km × 1 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.

## REFERENCES

- Adrian, R. J., 2007: Hairpin vortex organization in wall turbulence. *Phys. Fluids*, **19**, 041301, doi:10.1063/1.2717527.
- Browning, K. A., and R. Wexler, 1968: The determination of kinematic properties of a wind field using a Doppler radar. *J. Appl. Meteor.*, **7**, 105–113.
- Chen, W., M. Beister, Y. Kyriakou, and M. Kachelrieß, 2009: High performance median filtering using commodity graphics hardware. *IEEE Nuclear Science Symp.*, Orlando, FL, IEEE, 4142–4147.
- 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.
- 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. Climate Appl. Meteor.*, **25**, 990–1001.
- Kachelrieß, M., 2009: Branchless vectorized median filtering. *IEEE Nuclear Science Symp.*, Orlando, FL, IEEE, 4099–4105.
- Kaimal, J. C., and J. J. Finnigan, 1994: *Atmospheric Boundary Layer Flows*. Oxford University Press, 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. Opt.*, **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., and Coauthors, 2009: Comparison of 3D turbulence measurements using three staring wind lidars and a sonic anemometer. *Meteor. Z.*, **18**, 135–140.
- Mayor, S. D., 2011: Observations of seven density current fronts in Dixon, California. *Mon. Wea. Rev.*, **139**, 1338–1351.
- , and E. W. Eloranta, 2001: Two-dimensional vector wind fields from volume imaging lidar data. *J. Appl. Meteor.*, **40**, 1331–1346.
- , and S. M. Spuler, 2004: Raman-shifted eye-safe aerosol lidar. *Appl. Opt.*, **43**, 3915–3924.
- , —, B. M. Morley, and E. Loew, 2007: Polarization lidar at 1.54- $\mu\text{m}$  and observations of plumes from aerosol generators. *Opt. Eng.*, **46**, 096201, doi:10.1117/1.2786406.
- 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. Oceanic Technol.*, **21**, 1328–1345.
- , and —, 2004b: Assimilating coherent Doppler lidar measurements into a model of the atmospheric boundary layer. Part II: Sensitivity analyses. *J. Atmos. Oceanic Technol.*, **21**, 1809–1824.
- , 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. Meteor.*, **44**, 1324–1345.
- Patton, E. G., and Coauthors, 2011: The Canopy Horizontal Array Turbulence Study (CHATS). *Bull. Amer. Meteor. Soc.*, **92**, 593–611.
- 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.
- 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 Sens.*, **31**, 16–27.
- , and —, 1993b: Discrete spectral peak estimation in incoherent backscatter heterodyne lidar. II. Correlogram accumulation. *IEEE Trans. Geosci. Remote Sens.*, **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.
- 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** (D17), 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  $\mu\text{m}$ . *J. Atmos. Oceanic Technol.*, **22**, 696–703.
- , and —, 2007: Eye-safe aerosol lidar at 1.5 microns: Progress towards a scanning lidar network. *Lidar Remote Sensing for Environmental Monitoring VIII*, U. N. Singh, Ed., International Society for Optical Engineering (SPIE Proceedings, Vol. 6681), 668102, doi:10.1117/12.739519.
- 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.