

## ANALYSIS OF THE VELOCITY FIELD OF CMEs USING OPTICAL FLOW METHODS

ROBIN C. COLANINNO

George Mason University, Fairfax, VA 22030; robin.colaninno@nrl.navy.mil

AND

ANGELOS VOURLIDAS

Code 7663, Naval Research Laboratory, Washington, DC 20375; vourlidas@nrl.navy.mil

Received 2006 April 14; accepted 2006 July 24

### ABSTRACT

Optical flow is a powerful image processing tool for measuring motion in digital images. The optical flow algorithm provides an estimate of the velocity vector at every pixel from a pair of successive images. Here we present an application of this method to images of coronal mass ejections (CMEs). The technique is first tested and validated on a simulated CME. It is then applied to several CMEs observed with the LASCO C2 coronagraph to derive their velocity fields. The resulting velocity measurements allow us to visualize the evolution of the CME plasma and to separate the “bulk” velocity from the expansion velocity of a given CME. To our knowledge, this is the first time that such information has been extracted from CME observations. We discuss the limitations and accuracy of our optical flow method and propose further improvements.

*Subject headings:* methods: data analysis — Sun: coronal mass ejections (CMEs) — techniques: image processing

*Online material:* mpeg animations

### 1. INTRODUCTION

Virtually all analyses of coronal mass ejections (CMEs) start with the measurement of their physical properties, such as their width, position angle, and velocity. The most common is the measurement of the height of the CME front as a function of time. These so-called height-time (HT) plots are generated interactively by a user selecting the height of the CME front in successive frames containing the event. The user-defined heights are then plotted against the time of their observation, and a fit is applied to the data to derive the average CME velocity (from a linear fit) and acceleration (from a second-order fit).

This simple analysis method has several drawbacks. It provides the velocity along a single position angle only. It relies on the judgment of the user to determine the CME height. It is also time consuming and labor intensive, when measurements of the velocities of several features are needed. These costs are why multipoint velocity measurements of CMEs are rare. To avoid these problems, it is necessary to develop procedures that calculate velocities throughout a CME with minimal user intervention. Low & Hundhausen (1987) calculated the velocity field of a CME observed during the *Solar Maximum Mission*. The velocity field was interpolated from multiple position measurements of the complex, evolving structure of the CME. More recently, Tappin et al. (1999) applied an automated cross-correlation method to Large Angle Spectrometric Coronagraph (LASCO) images in order to measure the outflow speeds of the solar wind, but they did not apply their method to CMEs. Other automated velocity measurement methods have been used on coronal EUV and solar granulations observations (Gissot et al. 2003; November & Simon 1988). Finally, the automated CME detection software CACTus calculates a velocity at each radial angle along the front of the CME (Robbrecht & Berghmans 2004). Although CACTus provides an automated method of measuring the height of CMEs over multiple angles, it does not generate a velocity field over the whole CME.

In this paper, we introduce an image-processing algorithm from the field of computer vision and demonstrate that the velocity field of a CME can be obtained reliably over most of the area, with certain assumptions and restrictions. In § 2, we outline the algorithm and the expected limitations of the method. In § 3, we describe the application and validation of the algorithm on an artificial CME and several real CMEs, and in § 4 we discuss the results from our initial analysis. We conclude in § 5.

### 2. OPTICAL FLOW METHOD

#### 2.1. Optical Flow Constraint Equation

Optical flow can be defined as the vector field that warps one image into another (Horn & Schunck 1981). If the two images are the same scene separated by some time  $dt$ , then these vectors can be interpreted as velocities. These velocities will be two-dimensional (2D) projections of the *true* three-dimensional (3D) motion of the observed scene. We can express this definition of optical flow as an equation. If  $f(s, t)$  is the intensity at each pixel  $s$  on the image grid  $S$  at time  $t$ , we can define the optical flow,  $\mathbf{w}_s$ , as the vector at every pixel location that warps the image at time  $t$  into the image at time  $t + dt$ :

$$f(s + dt \mathbf{w}_s, t + dt) - f(s, t) \approx 0. \quad (1)$$

However, equation (1) is nonlinear with respect to the unknown velocity vector. In order to solve for  $\mathbf{w}_s$ , a first-order expansion is usually performed around  $(s, t)$ . This linearization leads to the optical flow constraint equation (OFCE):

$$\nabla f(s, t) \cdot \mathbf{w}_s + f_t(s, t) \approx 0, \quad (2)$$

where  $f_t$  is the partial temporal derivative calculated as the frame-to-frame difference, and  $\nabla f = [f_x, f_y]^T$  is the spatial gradient at  $t$ .

The OFCE is a very simple interpretation of motion in a scene. It does not use objects or patterns in the scene, but relies solely on

changes in intensity to estimate the motion. The reliance on changes in intensity requires important assumptions. We discuss these assumptions and their effects on the velocity estimation in § 2.3.

## 2.2. Regularization and Energy Minimization

The OFCE is *ill posed* and does not allow us to solve for the two component velocity vector from the scalar intensity. To constrain the possible solutions of the OFCE, a regularization term that favors smooth solutions can be added. The smoothness regularization term arises from the a priori assumption that the velocity vectors vary smoothly. Horn & Schunck (1981) proposed a global criterion where the Laplacian of the vector field is approximated by taking the difference of  $\mathbf{w}_s$  and its local average. This deterministic approximation leads naturally to a statistical framework in which the regularization is a Markov random field (MRF). The estimation of the optical flow can then be expressed as a global optimization problem (Mémin & Pérez 1998). The computation reduces to a search for a global minimum of an energy function. An energy function typically describes the interaction between variables and generally involves two components. The first component expresses the interaction between the observed and unknown variables, while the other component expresses some a priori knowledge about the unknown variables.

For our problem, the first component is the OFCE, which expresses the interaction between the image intensity and the vector field. The second component is the smoothing regularization MRF, where  $\mathbf{w}_r$  are the neighboring site pairs. Thus the energy function that we minimize is

$$H(\mathbf{w}; f) = \sum_{s \in S} [\nabla f(s, t) \cdot \mathbf{w}_s + f_i(s, t)]^2 + \alpha \sum_{\langle s, r \rangle \in C} \|\mathbf{w}_s - \mathbf{w}_r\|^2. \quad (3)$$

To calculate  $\mathbf{w}_r$ , we use a four-neighborhood system for which any pixel has a set of four neighbors;  $C$  is the set of neighboring pixels. The parameter  $\alpha$  controls the balance between the two components and is dependent on the signal-to-noise ratio of the input images. It is a user-defined quantity that determines the smoothness of  $\mathbf{w}_s$ .

The energy function in equation (3) is quadratic with respect to the unknown vector field and can be solved through an iterative scheme. However, this minimization is nonlinear and contains numerous local minima. To find a global minimum, we adopt the multigrid relaxation method (Heitz et al. 1994; Mémin & Pérez 1998). The multigrid relaxation method should not be confused with the multiresolution approach used in most optical flow methods for tracking large motions. The velocity field is minimized on a reduced grid, whereas observations are still viewed at the natural image scale. The energy function is minimized through a hierarchy of nested subspaces of the whole configuration space. These subspaces correspond to configurations constrained to be piecewise constant over smaller and smaller pixel subsets. The multigrid relaxation method leads to faster and more accurate minimization of the global energy function (Heitz et al. 1994).

## 2.3. Assumptions and Limitations

To formulate our optical flow algorithm, we had to make a few assumptions that affect the accuracy of the velocity estimation. In this section, we discuss some of these assumptions and their inherent limitations. We refer the reader to Horn (1986) for an in-depth discussion of optical flow. Here, we limit ourselves to those problems that most effect the application to CMEs.

The first assumption arises from the definition of optical flow, which associates a change in image intensity with motion in the observed scene. In other words, we have to assume that an observed change in brightness, at a given location, is caused solely by motion and not by a change in the intrinsic brightness of the object. This is a reasonable assumption for solid opaque objects. However, it is not such a good assumption for coronal images in which the features have intrinsic brightness variability. CMEs are often optically thin. The CME intensity naturally decreases as a CME expands and its density decreases. The intensity decrease is interpreted by the algorithm as slower velocities. Another situation where this assumption is violated is when there is motion with no observable change in intensity. This effect is particularly noticeable along the front of the CME where the intensity profiles of two successive images overlap.

Our second assumption comes from the linearization of equation (2). We assumed that intensity and hence velocity varies slowly over  $dt$ . This assumption imposes an upper limit to the intensity evolution between images that can be accurately estimated. Also, the smoothness regularization assumes that the velocity is varying smoothly, which limits the time steps  $dt$ . These two assumptions require that the sequence of images be closely spaced in time and that the observed motion between frames does not cause extreme intensity changes.

Another limitation is that the smoothness assumption can be easily violated at sharp edges. It is a good assumption here, since CMEs do not usually have sharp features. However, it can lead to oversmoothing in the velocity field. Oversmoothing reduces the velocity intensity of the features by smoothing them into the background levels. Small-scale structures can also be completely oversmoothed and get lost in the background. The combined effect of all these limitations is a reduction of the estimated velocities.

## 3. APPLICATION OF OPTICAL FLOW METHOD TO CMES

### 3.1. Application to a Simulated CME

To develop the optical flow algorithm and understand how it operates on CME images, we applied it first to a simulated CME. The simulated CME is a toroidal 3D object with plasma located only on its outer surface; the plasma density is constant ( $1 \times 10^5 \text{ cm}^{-3}$ ). The resulting brightness is calculated using the standard Thomson scattering equations. The CME expands self-similarly. This particular simulation was created for another application without the projected velocity parameterized. We assigned a cadence of 10 minutes to calculate the velocity. The resulting sequence, like the observed data, is a series of 2D projections of a 3D structure moving in space. The simulated CME is well behaved, with a smooth front expanding at a constant rate. It is an excellent case for testing the optical flow method for application to CMEs. A full description of the simulated CME can be found in Thernisien et al. (2006).

We used a sequence of 15 images in which the CME appears from behind the occulter and expands to the edge of the field of view. The field of view of the simulated images is equivalent to that of LASCO C2 ( $2.5\text{--}6 R_\odot$ ). Figure 1 shows three frames from the optical flow estimation. The results from the full sequence of images can be viewed as an mpeg file in the online version of the paper. The images on the left are the first of the two images used to estimate the optical flow. The images on the right are the optical flow results. The shading in the images represents the magnitude of the velocity, on a scale ranging from 0 to  $300 \text{ km s}^{-1}$ . The arrows show the direction of the velocity for an  $8 \times 8$  grid on the image. We actually have a velocity vector for every pixel in the image.




FIG. 1.—Simulated CME, showing the calculated mass images (*left*) and optical flow results (*right*). The shading represents the velocity magnitude. The arrows show the direction of the velocities sampled over an  $8 \times 8$  grid in the image. The scale bars are in units of kilometers per second. [This figure is available as an mpeg file in the electronic edition of the Journal.]

At the front of the toroid, where the intensity is changing, the velocity vectors reveal a smooth radial expansion of the CME. In the region between the two leading edges of the toroid, the optical flow algorithm reports zero velocity. In this region the intensity does not change, so the program does not estimate a velocity. There is also a band inside the front where there is no velocity, and where the intensity profiles of the two images overlap. Thus there is no intensity change in this narrow region. In

Table 1, we compare the average front velocity derived from the optical flow field and the velocity derived from the standard HT plot. We used the LASCO software to measure the height of the CME front. Figure 2 (*top*) shows the standard user-generated HT plot for the front of the simulated CME. To find the average front velocity from the optical flow results, we extracted the velocities along the position angle (PA). Then we averaged the maximum velocity from the profile of each image. We see that

TABLE 1  
CME FRONT VELOCITY

CME	Time	HT (km s <sup>-1</sup> )	HT PA (deg)	AOF (km s <sup>-1</sup> )	Bulk PA (deg)	Bulk Velocity (km s <sup>-1</sup> )	Expansion Velocity (km s <sup>-1</sup> )
Simulated .....	...	299	269	218	264	109	204
2000 Jun 18 .....	18:50	233	262	166	264	43	177
2001 Feb 20 .....	17:00	202	301	236	288	94	100
2001 May 13 .....	12:20	273	234	191	228	123	120
2002 Jan 13 .....	07:35	215	55	184	53	152	76
2002 Apr 2 .....	15:35	152	88	132	84	...	...
2002 Sep 1 .....	04:35	233	261	204	257	120	120
2002 Sep 1 .....	19:40	215	30	191	30	72	106
2003 Sep 28 .....	13:00	220	138	167	142	...	...
2004 Jun 26 .....	07:48	211	85	216	83	155	18
2004 Jun 27 .....	09:12	220	299	185	298	148	93

NOTES.—CME time is the time of the first appearance of the CME in the LASCO C2 field of view. HT is the front speed derived by a standard height time plot. The AOF (average optical flow) is describe in § 3.2. HT PA, and bulk PA are the position angles of the CME front and center of motion, respectively. The bulk and expansion velocities are described in § 4. The structures of the CMEs at 2002 Apr 2 15 : 35 and 2003 Sep 28 13 : 00 do not conform with the model used to measure the bulk and expansion velocities.

<sup>a</sup> HT velocity calculated using only LASCO C2 data.

the average optical flow (AOF) velocity is lower than the HT velocity by about 30%.

### 3.2. Application to Observed CME (LASCO C2)

We applied our optical flow estimation to images taken with LASCO C2. We found that the LASCO C2 synoptic cadence of 24 minutes was too low for successful optical flow estimation, so we used data from special observing sequences between 1999 and 2004. These observing sequences have an alternating cadence of 5–10 minutes between images, or 12 minutes in 2004. During these observing sequences approximately 330 CMEs were observed and cataloged. Of these CMEs, 85 have velocities less than 300 km s<sup>-1</sup>. Many of these slower CMEs have diffuse or irregular structures that make it difficult to sample the optical flow field average velocity. We finally chose 10 CMEs for further analysis because they had large, well-defined structures.

To remove as many instrumental effects as possible and to minimize intensity variations (i.e., varying background), we used excess mass images instead of the raw intensity files. Using Thomson scattering calculations, we converted the observed image intensities to masses (Vourlidis et al. 2000). Therefore, intensity changes in the images correspond to mass changes, which in turn correspond to mass movement. Thus the first assumption of the OFCE is satisfied. Working with mass images increases the contrast between the CME and the background and also increases the signal-to-noise ratio in the images. To isolate the mass in the CME, a pre-event image is subtracted from the sequence of images. We also removed stars and cosmic rays from the images because they cause large errors in the background velocity field.

Here we present the results from a CME that was first observed 2002 September 1 at 4 : 35 UT in the LASCO C2 field of view. The CME was observed in 20 images before its front left the field of view. Figure 3 shows three of the results of the optical flow estimation. The results from the full sequence of images can be viewed as an mpeg file in the online version of the paper. As described in the previous section, the images on the left are the mass images and the images on the right are the results of the optical flow estimation. The scale of the velocity magnitude is from 0 to 300 km s<sup>-1</sup>.

The dark feature in the mass image is the imprint of a pre-event streamer blown away during the event. Once again the arrows show a smooth front expanding radially outwards. The interior

of the CME reveals a complicated velocity field that is hard to interpret without detailed analysis of the particular event. We defer such analysis for a future paper. There is also a translational motion in the streamer next to the CME. Figure 4 is the HT plot for the height measurements in LASCO C2 recorded in the LASCO CME catalog. Again, to calculate the average optical flow velocity, we extracted the velocities along the PA and averaged the maximum velocity from the profile of each image. In Table 1, we compare the AOF velocity with the velocity from the HT plot. The AOF velocity is lower than the HT velocity by 12%. The observed CME is slower and has a higher image cadence than the simulated CME. Thus the optical flow estimation is expected to be more accurate for the observed CME.

The data for the other nine CMEs to which we applied our optical flow method are also listed in Table 1. Again, we compare the AOF velocity with the velocity from the LASCO CME catalog HT plot. The average optical flow velocity for all CMEs does not vary from the cataloged velocity by more the 30% for all the studied CMEs. The optical flow field for all CMEs can be viewed on our Web site.<sup>1</sup>

## 4. RESULTS FROM OPTICAL FLOW CME MAPS

Based on the optical flow movies, we chose to describe the observed motion of the CME as that of an expanding bubble along a radial trajectory defined by some position angle (PA). We define the translational motion as the bulk velocity, and the motion caused by the expansion of the bubble the expansion velocity. Based on this description, the velocity derived from a standard HT plot should be the sum of the bulk and expansion velocities. We have used the results of the optical flow method to separate and measure the bulk and expansion velocities of the CMEs in this paper.

### 4.1. Bulk Velocity

To derive the bulk velocity of a CME from the optical flow results, we first eliminate the background field by setting all velocities in the lower 90th percentile to zero. By setting the lower 90th percentile to zero, we are left with just the velocities in the CME. We visually verify that we have isolated the CME from

<sup>1</sup> <http://lasco-www.nrl.navy.mil/archive/colaninno>.




FIG. 2.—*Top*: Standard HT plot for the simulated CME front. *Middle*: HT of the center of motion generated from our analysis of the optical flow results. The intercept of the linear fit is different from the standard HT plot because the center of motion emergence later than the other features. *Bottom*: CME expansion relative to the bulk velocity in three directions.

the background and any other bright features in the image. Next we locate the center of motion of the CME with the following procedure:

Since we have the horizontal ( $x$ ) and vertical ( $y$ ) components of the velocity for every pixel, it is easy to select only those pixels

with a single component in the  $\pm x$  and  $\pm y$  directions. Then we calculate the average of the pixels' locations along each axis separately. The average  $x$  and  $y$  positions define the CME center of motion for the particular frame. In essence, this method is similar to the limb-fitting methods used for locating the center of the Sun in full disk images. The difference is that we use information from only two position angles.

We applied this analysis to the optical flow results for both the simulated and observed CMEs. In middle panels of Figures 2 and 4, we plot the height of the center of motion versus time for the simulated and the 2002 September 1 CMEs, respectively. Similar to the HT plot, we fit a line to the points in order to obtain the bulk velocity listed in Table 1. We also calculate a position angle (bulk PA) for the radial position of the height measurements. The bulk velocity for both the CMEs is less than that of the HT front velocity. From our definition of bulk velocity, this is what we expected. The bulk PA of the center of motion is also approximately the HT PA of the front. Our completely automated analysis finds the same PA of the CME as the user.

#### 4.2. Expansion Velocity

We can now extend the analysis of the optical flow results to calculate the expansion velocity. For each available optical flow image, we take the same set of pixels used for the derivation of the bulk velocity, but now we split them into four subgroups. We separate the vectors in the  $x$  direction into those left and right of the center of motion. Similarly, we separate the vectors in the  $y$  direction into top and bottom groups. We then take the average of the pixels' locations along each axis the four subsets of vectors and calculate the distance from the center of motion. We repeat this procedure for all available optical flow images. In the end, we obtain four plots (left, right, top, and bottom vectors) of the expansion of the front relative to the center of motion versus time (see bottom panels of Figures 2 and 4). By fitting a line to these data, we finally obtain the expansion velocities in each of the four directions. Note that these plots are not standard HT plots because the distance plotted is not the height above the limb of the Sun.

In Table 1, we list the expansion velocity of the CME front. The sum of expansion and bulk velocities is approximately the HT front velocity, as expected. For the simulated CME, the sum of the bulk and expansion velocity for the front is about  $313 \text{ km s}^{-1}$ . This velocity is only 4% larger than the HT front velocity. For the observed CMEs, the sum of bulk and expansions velocities are also all very close to the front velocity, ranging from 3% to 18%, with an average of 9%. We believe that these results show that we are accurately separating the bulk and expansion velocities of CMEs and are providing confidence in optical flow methods.

### 5. DISCUSSION AND CONCLUSIONS

We believe that these preliminary results show that our algorithm provides reliable measurements of the velocity fields of CMEs, albeit with some restrictions. First, the intensity of the various CME structures must remain relatively constant. Second, the CME must retain its shape as it expands. If the front, for example, breaks apart, the optical flow algorithm will fail in that frame and possibly in the subsequent ones. Third, the various CME features need to be above the background noise. Real CME data are not amenable to precise error analysis, but our experiments led us to the empirical conclusion that reliable velocity maps can be derived if the velocity vector does not exceed  $\sim 30$  pixels. In the case of LASCO C2, this restriction translates

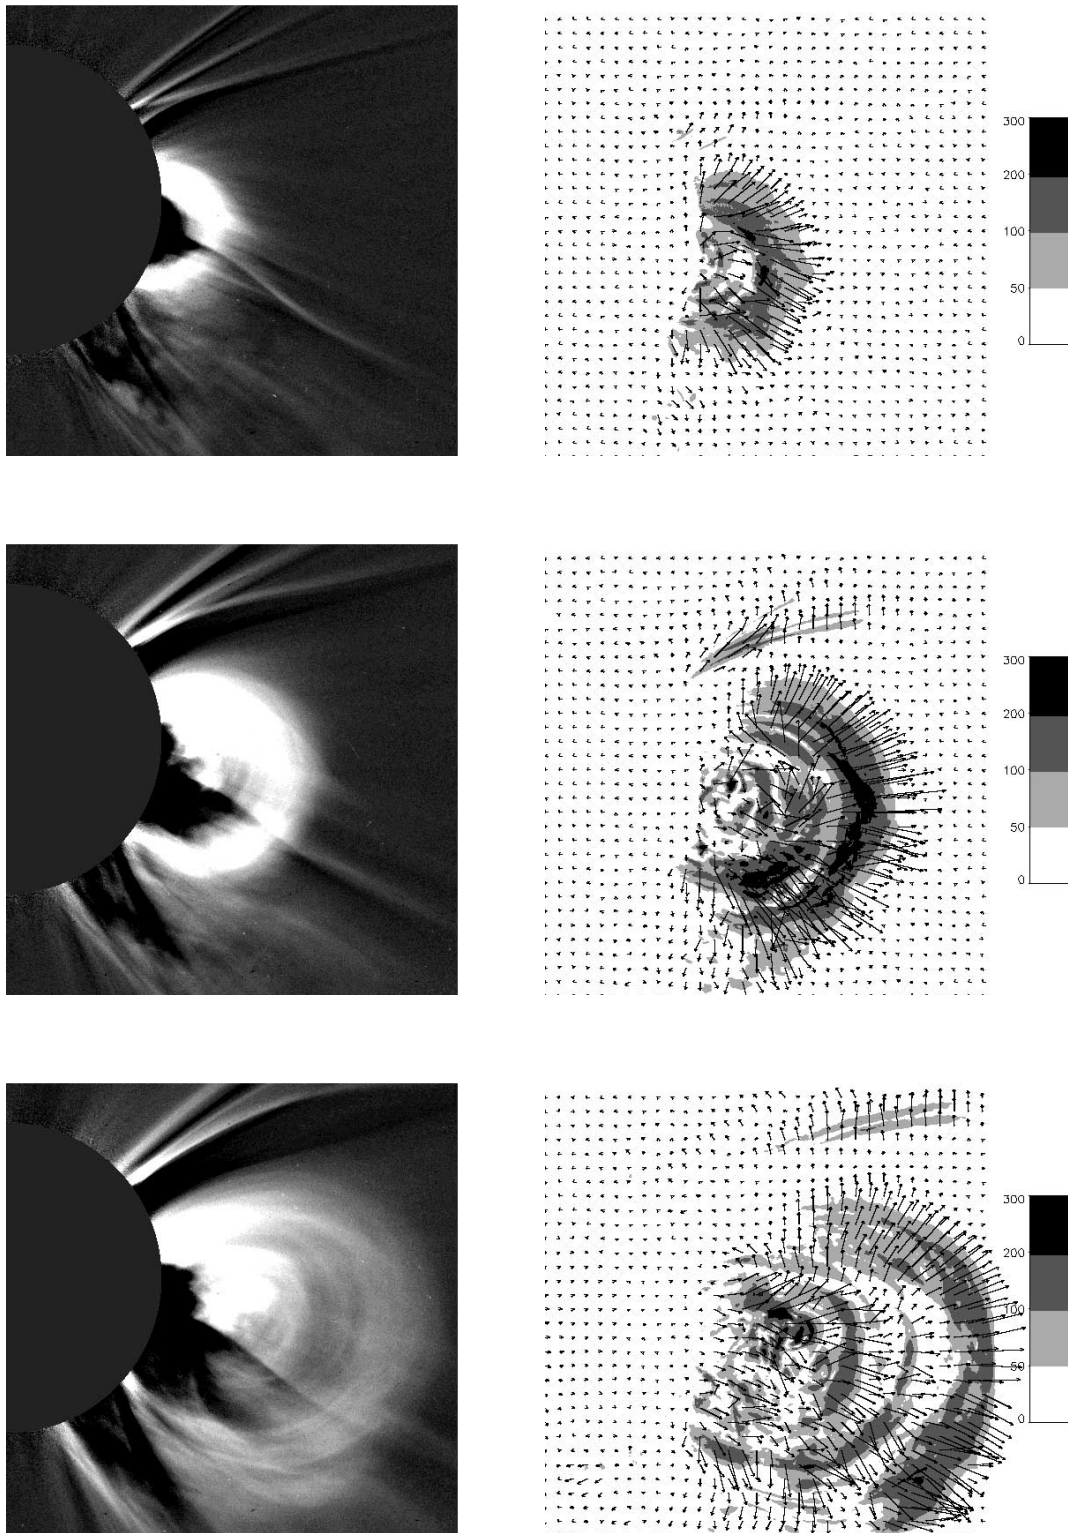


FIG. 3.—LASCO CME on 2002 September 1, showing the calculated mass images (*left*) and optical flow results (*right*). The shading represents the velocity magnitude. The arrows show the direction of the velocities sampled over an  $8 \times 8$  grid in the image. The scale bars are in units of kilometers per second. [*This figure is available as an mpeg file in the electronic edition of the Journal.*]

to a minimum cadence of about 15 minutes for a CME velocity of  $300 \text{ km s}^{-1}$ .

On the upside, the algorithm is very fast and can easily be applied to all available images of a CME. The resulting optical flow movies provide much more information than can be ob-

tained from a single velocity measurement. The velocity results can be better appreciated in the movies available in the online version of the paper.

One of the most obvious advantages of the optical flow movies is that the viewer can easily visualize the expansion of the CME.




FIG. 4.—*Top*: Standard HT plot for the LASCO CME front. *Middle*: HT of the center of motion generated from our analysis of the optical flow results. The intercept of the linear fit is different from the standard HT plot because the center of motion emergence is later than the other features. *Bottom*: CME expansion relative to the bulk velocity in the right (outward) direction.

The velocity vectors within the CME front (Fig. 3) show a radial orientation, pointing toward a common origin at the center of the CME. This behavior is especially obvious at the CME front, but not so much in other substructures. In other words, the superposing of the velocity vectors reveals immediately whether

a given CME is expanding systematically or not. We take these observations one step further in § 4 and separate the observed velocity of a CME into its two components: the bulk and expansion velocities. By subtracting the bulk velocity from the movie of the CME, we can also see the expansion-only motion of the CME. Such analysis is advantageous for studying the physics of CMEs because each component might be driven by distinct processes. For example, the bulk velocity of the CME might be influenced more by the local environment, such as the local solar wind speed, whether there is a preceding CME, the interaction with local streamers, and so forth. The expansion velocity, on the other hand, might be driven by the magnetic content of the ejected material. Some models predict that in the case of a flux rope CME, the expansion velocity is related to the reconnection rate at the current sheet behind the CME (Lin 2004). In those cases, we expect to measure different acceleration profiles for the two component velocities. We plan to investigate this point further in the future. Another very interesting result from the optical flow movies concerns structures away from the CME. An inspection of the movies reveals motions of streamers that are not obviously connected to the ejecta. The velocity vectors point away from the CME, and we see no plasma between the CME and the affected streamer that could act as the agent for the observed motion. The obvious interpretation is that the agent is none other than the CME-driven shock. Streamer deflection has been interpreted as an indirect proxy of shocks, and a recent shock analysis demonstrates a direct connection between the shock and the streamer deflection (Vourlidas et al. 2003). It is very encouraging that the optical flow maps can not only identify these distinct motions but also provide quantitative measurements of the wave propagation speeds within these structures.

We plan improvements for our optical flow algorithm. To reduce the effects of oversmoothing, we plan to change the term  $\alpha$  in equation (3) to a robust estimator. Then the smoothing function will be applied dynamically, depending on the intensity change within the neighborhood system. This improvement should lead to smooth results without reducing the estimated velocity. We will also try to extend the range of CME velocities that can be estimated by including a multiresolution scheme. We have not implemented a multiresolution scheme so far because we first wanted to obtain a good understanding of the operation of optical flow methods on CMEs. Multiresolution schemes involve warping the image by the velocity vectors at a lower image resolution. This scheme can very quickly propagate errors in the optical flow estimation and create erroneous results.

Another potential application of optical flow is the measurement of the solar wind speed. Previous measurements of coronal images have involved correlation methods either by sight (Sheeley et al. 1997) or automated (Tappin et al. 1999). The solar wind speed is measured by identifying the speeds of the so-called blobs often seen in LASCO C2 images. This type of measurement requires considerable image processing and significant personnel resources. It has never been applied to study the variations of speeds at different position angles (and structures) in the images. Therefore, it is not clear what the precise relationship between the blobs and the actual solar wind is. Optical flow does not rely on shapes to estimate velocity. Any changes in the intensity caused by the solar wind should be detected if the intensity change is above the image noise level. If the blobs are moving at velocities different from the solar wind, we will be able to detect them. If they are moving at the same speed as the solar wind, they will be indistinguishable from it. Such work will have important implications for heliospheric and space weather studies.

We thank A. F. Thernisien for providing the simulated CME images. *SOHO* is an international collaboration between NASA and ESA. LASCO was constructed by a consortium of institutions: the Naval Research Laboratory (Washington, DC), the Max-Planck-Institut für Sonnensystemforschung (Katlenburg-Lindau,

Germany), the Laboratoire d'Astronomie Spatiale (Marseille, France), and the University of Birmingham (Birmingham, UK). The LASCO CME Catalog is generated and maintained at the CDAW Data Center by NASA and The Catholic University of America in cooperation with the Naval Research Laboratory.

## REFERENCES

- Gissot, S. F., et al. 2003, in *Solar Variability as an Input to the Earth's Environment*, ed. A. Wilson (ESA SP-535; Noordwijk: ESA), 853  
Heitz, F. Pérez, P., & Bouthemy, P. 1994, *CVIU*, 59, 125  
Horn, B., 1986, *Robot Vision* (Cambridge: MIT)  
Horn, B., & Schunck, B. 1981, *Artificial Intelligence*, 17, 185  
Lin, J. 2004, *Sol. Phys.*, 219, 169  
Low, B. C., & Hundhausen, A. J. 1987, *J. Geophys. Res.*, 92, 2221  
Mémin, E., & Pérez, P. 1998, *IEEE Trans. Im. Proc.*, 7(5), 703  
November, L. J., & Simon, G. W. 1988, *ApJ*, 333, 427  
Robbrecht, E., & Berghmans, D. 2004, *A&A*, 425, 1097  
Sheeley, N. R., et al. 1997, *ApJ*, 484, 472  
Tappin, S. J., Simnett, G. M., & Lyons, M. A. 1999, *A&A*, 350, 302  
Thernisien, A. F., Howard, R., & Vourlidas, A. 2006, *ApJ*, in press  
Vourlidas, A., et al. 2000, *ApJ*, 534, 456  
———. 2003, *ApJ*, 598, 1392