

## Derivation of time-dependent two-dimensional velocity field maps for plasma turbulence studies

T. Munsat

*Center for Integrated Plasma Studies, Department of Physics, University of Colorado, Boulder, Colorado 80309*

S. J. Zweben

*Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543*

(Received 18 July 2006; accepted 21 August 2006; published online 2 October 2006)

Time-resolved two-dimensional (2D) velocity maps have been derived for fluctuation measurements in the edge plasma of the National Spherical Torus Experiment (NSTX). The maps have been derived from time sequences of 2D images recorded with the gas puff imaging diagnostic. A hybrid technique combining optical flow and local pattern matching has been implemented to overcome the individual limitations of each when used with data of limited temporal and/or spatial resolution. Local flow velocities of up to  $\sim 8$  km/s and average poloidal flow velocities of up to  $\sim 5$  km/s are found. Results are compared to previous velocity extraction techniques and NSTX results. © 2006 American Institute of Physics. [DOI: [10.1063/1.2356851](https://doi.org/10.1063/1.2356851)]

### I. INTRODUCTION

It is widely believed that small-scale turbulence is a primary cause of anomalous thermal and particle transport in toroidal magnetically confined plasmas. While turbulence measurements of tokamak plasmas, especially in the edge region, have been performed for decades, they have been primarily limited to single-point or one-dimensional (1D) measurements.

It is also known that, in fluid dynamics, the two-dimensional (2D) visualization of turbulent structures has greatly enhanced the level of understanding of the underlying physical mechanisms. In plasma physics, the past decade has seen a tremendous advance in computer simulations, which produce elegant and detailed 2D and three-dimensional (3D) results. Unfortunately, the experimental evidence to corroborate many of the most interesting theoretical results is lacking, due most often to the limitations of the diagnostic techniques.

To address this need, a number of diagnostics have been recently developed for the visualization of turbulent structures, including beam emission spectroscopy (BES),<sup>1,2</sup> electron cyclotron emission imaging (ECEI),<sup>3,4</sup> microwave imaging reflectometry (MIR),<sup>5,6</sup> and gas puff imaging (GPI).<sup>7,8</sup> ECEI makes use of focusing optics and frequency selection based on the gradient in the strength of the toroidal magnetic field to form a 2D image (in the poloidal plane) of electron temperature. MIR makes use of similar microwave frequencies and collection optics to form an image of the plasma cutoff layer, enabling localized and precise measurements of density fluctuations. Both BES and GPI are based on the detection of emission from neutral particles near the plasma edge. BES has higher time resolution ( $>1$  MHz, compared to 250 kHz for GPI with the present camera), while GPI has higher spatial resolution ( $\sim 4$  mm, compared to  $\sim 10$  mm for BES) for similar spatial coverage ( $\sim 10 \times 10$  cm<sup>2</sup> for BES,  $\sim 20 \times 20$  cm<sup>2</sup> for GPI).

The determination of the poloidal flow pattern in the plasma edge can contribute significantly to the understanding of turbulence and its relationship to edge transport. In particular, the derivation of time-resolved velocity maps can establish a phase relationship between the fluctuating density and velocity, which can then be used to derive an estimate for intermittent plasma flux from structures near the separatrix. Significantly, it has recently been observed via probe measurements at the plasma edge that an appreciable fraction of the observed radial transport in tokamaks can be attributed to intermittent phenomena, rather than purely diffusive transport.<sup>9,10</sup>

Recently there have been efforts by several groups to derive velocity fields from spatially/temporally resolved density fluctuation data. Of particular relevance are the time-resolved velocity maps produced by the DIII-D group using BES data<sup>11,12</sup> and the spatially resolved velocity maps derived by the C-Mod group using a GPI system similar to the National Spherical Torus Experiment (NSTX) system.<sup>13</sup>

The primary challenge of the BES measurements has been the limited spatial resolution ( $5 \times 6$  pixels), though the data are highly time resolved (1 MHz sampling rate). The BES group has employed two methods for deriving velocity fields, the first of which employed the “orthogonal dynamic programming” method, wherein the data are alternately broken down into vertical and horizontal strips then are iteratively deformed using velocity estimates of increasing resolution until the flow field reflects the evolution of the images to the required level of detail.

A complementary study on the BES data employed the “time delay estimation” (TDE) method, in which short subsections of the wave forms from each channel were correlated with similar wave forms from other channels, shifted by a variable time delay. The delay is adjusted until the maximum cross correlation is achieved, producing a velocity point from the spatial separation of the channels and the time

delay of maximum correlation. This method is particularly well suited to the BES data set, given its high time resolution and limited spatial resolution, which enabled reasonably time-resolved velocity fields even with 64 points per subsection in the correlated wave forms.

The hybrid technique presented here is somewhat similar in nature to this technique, except that the GPI data is subdivided into spatial rather than temporal subsets, and correlation was done between two single time frames rather than extended wave forms. Indeed, the comparison between these two approaches illustrates the general approach of cross-correlating subsets of the data, taking advantage of subsampling over whichever dimension (temporal, spatial) is recorded with the highest resolution compared to the size of the structures of interest.

The C-Mod study used the method of time delay cross correlation (similar to the TDE approach described above) applied to GPI data. While this method has the advantage that it is relatively direct, this technique has the disadvantage that the derived flow fields are not time resolved, with only a single field mapping produced per sequence of images. The results of this study have shown clear flow layers within different regions of the plasma edge in both *L*-mode and *H*-mode discharges. In *L* mode, the observed flow pattern was predominantly poloidal with velocities of up to 1 km/s inside the separatrix, sharply changing to radial flows of similar magnitude outside the separatrix. Similar flow patterns were observed in edge localized mode (ELM)-free *H* modes, average velocity magnitudes decreased by approximately the factor of two.

Algorithms developed for robot vision and weather tracking have been able to take advantage of extremely high spatial and/or temporal resolution, which allows for relatively accurate tracking with either optical flow or pattern-matching techniques. That said, the development of robust algorithms for producing highly detailed velocity maps is an area of active research. The application to GPI is limited by the time resolution of the instrument. Namely, even with the use of a state of the art camera, the motion of the features from frame-to-frame is significant, and most established algorithms fail to accurately track the plasma motion.

In this article we present a velocity extraction technique which overcomes many of the limitations present in common approaches (optical flow and pattern matching) when applied individually to data of this resolution. Specifically, optical flow techniques (when used independently) tend to underestimate plasma flow, for reasons described below. Pattern-matching techniques, on the other hand, tend to produce spurious results, primarily due to the arbitrariness of the initial conditions fed into the algorithms. The hybrid combination of the two techniques, in which the output from the optical flow is used as an input to a pattern matching algorithm, enables both of these limitations to be largely overcome.

In Sec. II, we briefly describe the gas puff imaging experimental technique. In Secs. III and IV, we present the optical flow and pattern-matching components of the velocity extraction algorithms. In Sec. V, the analysis of test data

sets is presented, followed by the analysis of several NSTX data sets in Sec. VI. A summary/discussion is included in Sec. VII.

## II. GAS PUFF IMAGING TECHNIQUE

The GPI technique is a fast, high resolution measurement of neutral gas line emission from the edge region of hot magnetically confined plasmas. The visible line emission from a gas cloud puffed in at the edge of the torus (for example, the He I line at 587.6 nm or  $D_\alpha$  line at 656.2 nm) is imaged with a fast framing camera and linear detector arrays. These measurements are performed with pixel resolution of  $\sim 4$  mm, spatial resolution of the emission cloud of 1–2 cm (due to averaging along sightlines), and time resolution of 4  $\mu$ s, shorter than the autocorrelation time of the turbulent structures. A detailed description of the technique can be found in Refs. 7 and 8, and only a brief overview is provided here.

The intensity of the line emission (under the collisional radiative approximation) depends on the local electron density  $n_e$  and temperature  $T_e$  as

$$S = n_0 f(n_e, T_e) A, \quad (1)$$

where  $n_0$  is the local neutral density,  $A$  is the radiative decay rate for the observed line, and  $f(n_e, T_e)$  is a function providing the density ratio between neutrals in the upper and lower transition states. The dependence on  $n_e$  and  $T_e$ , through the function  $f$ , modulates the line emission according to the fluctuations in both plasma parameters. Calculations using the DEGAS 2 neutral transport code, with NSTX geometry and typical edge parameters as inputs, have shown this dependence to be approximately  $S \propto n_e^{0.5} T_e^{0.7}$  near the center of the emission cloud, which corresponds to the region where  $5 \text{ eV} \leq T_e \leq 50 \text{ eV}$ .

For these experiments the images were recorded with a Princeton Scientific Instruments PSI-5 camera, coupled to the vessel via a coherent fiber-optic bundle inserted into a reentrant vacuum port. The PSI-5 camera uses a custom charge coupled device (CCD) chip capable of storing 300 image frames of  $64 \times 64$  pixels each at a rate of 250 000 frames/s. The CCD chip has a light-sensitive pixel size of 100  $\mu$ m, net quantum efficiency of over 15%, readout noise of  $< 32$  electrons/pixel, and saturation level of  $> 24$  000 electrons/pixel. An image intensifier was used at the front end of the camera, and the images were read out using a low-noise 14 bit analog-to-digital converter (ADC). The response of the image intensifier is not linear, though that does not affect the analysis presented here.

The GPI diagnostic layout in NSTX is illustrated in Fig. 1(a). A 30 cm long gas manifold is mounted on the outer wall behind the rf limiter shadow and oriented to produce a cloud of gas extending along the local poloidal direction at the plasma edge. Note that the magnetic field in NSTX and other spherical tori is inclined by  $\sim 35^\circ$ – $45^\circ$  with respect to the toroidal direction at the outer midplane. The gas puff which leaves this manifold is viewed from along the edge magnetic field lines through a reentrant port about 70 cm

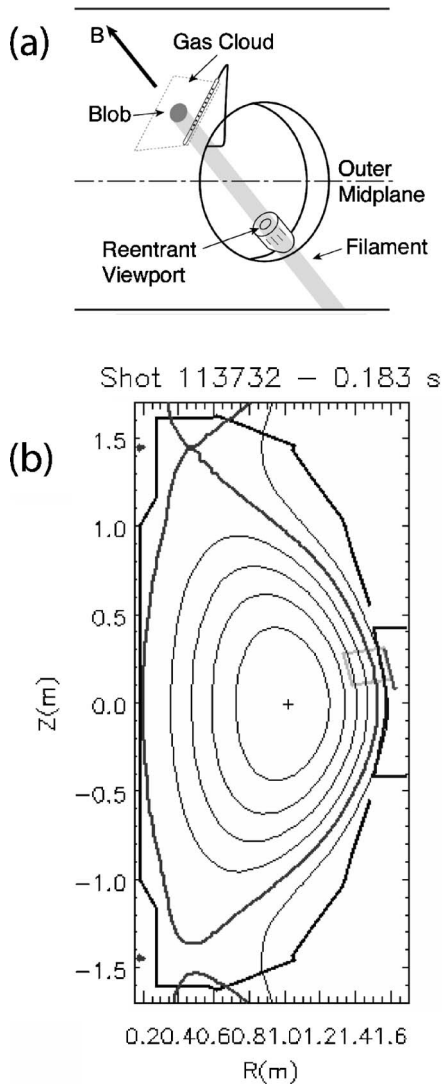


FIG. 1. Schematic layout of the GPI instrument geometry. The neutral gas cloud is puffed from a manifold oriented perpendicular to the local magnetic field. The gas cloud is viewed along a sight line parallel to the local magnetic field at the cloud, enabling an approximately 2D view of the line emission. The GPI view in the context of the NSTX poloidal cross section is indicated in (b) as the gray rectangle just above the outer midplane.

away. The gas puff is thus used to localize the light emission in order to obtain a cross section of the turbulence in the radial versus poloidal plane.

The camera image is centered near the magnetic separatrix (the boundary between open and closed field lines) about  $17^\circ$  above the outer midplane, as shown in Fig. 1(b). The camera has a  $\sim 23 \times 23$  cm<sup>2</sup>, (radial  $\times$  poloidal) viewing area perpendicular to the magnetic field at the gas cloud.

It should be noted that a fundamental constraint on the use of the GPI diagnostic is that imposed by the line emission profile of the neutral gas. Namely, the brightness of any given emission line has a limited radial extent in regions of steep temperature and density gradients, necessarily resulting in a degradation of the signal in some regions of the image. While different gases can be chosen to slightly alter the temperature dependence of the brightness profile, it is in all cases very difficult to use the GPI technique very far from the separatrix in NSTX.

### III. OPTICAL FLOW TECHNIQUE

#### A. Basic approach

The first technique employed for this study is an implementation of the “optical flow” approach, a method of deriving velocity fields from sequences of still images which has been used for a number of years for the analysis of fluid dynamical motion.<sup>14–19</sup> A detailed description of the optical flow technique can be found in Ref. 18, and a simple overview is provided here.

Optical flow is a technique for measuring motion within image sequences by determining the velocity field which maps one image to the next. The fundamental hypothesis (or constraint) is that the intensity structures of local time-varying regions within the image are approximately constant over several image frames. Namely, for small values of  $\Delta t$ ,

$$I(x, y, t) = I(x + v_x \Delta t, y + v_y \Delta t, t + \Delta t), \quad (2)$$

where  $I$  is the image brightness and  $\mathbf{v} \equiv (v_x, v_y)$  is the local 2D velocity vector. A detailed, time-dependent velocity map can be constructed through the derivation of local translation vector fields for each section of each image frame and knowledge of the time delay between each frame. The formulation used in this study is adapted from that presented in Ref. 20.

Equation (2) can be differentiated to produce

$$\frac{\partial I}{\partial t} + v_x \frac{\partial I}{\partial x} + v_y \frac{\partial I}{\partial y} = \frac{\partial I}{\partial t} + \mathbf{v} \cdot \nabla I = 0, \quad (3)$$

which is simply the continuity equation for image brightness,

$$\frac{\partial I}{\partial t} + \nabla \cdot (\mathbf{v}I) = 0, \quad (4)$$

with the assumption of 2D divergence free flow ( $\nabla \cdot \mathbf{v} = 0$ ). Thus the optical flow constraint is equivalent to the assumption of “constant brightness.”

In order to derive a velocity field which satisfies the optical flow condition, it is useful to first define the field components in terms of a family of basis functions, the coefficients of which can then be used to recast Eq. (3) into a matrix equation which can be solved by inversion.

Let  $\phi(x, y)$  be a family of 2D velocity basis functions so that

$$v_x = \sum_{k=0}^{K-1} v_{x,k} \phi_k \quad \text{and} \quad v_y = \sum_{k=0}^{K-1} v_{y,k} \phi_k, \quad (5)$$

where  $K$  is the number of basis functions required to describe  $v_x$  and  $v_y$ .

Equation (3) can now be written as

$$\frac{\partial I}{\partial t} + \sum_{k=0}^{K-1} v_{x,k} \phi_k \frac{\partial I}{\partial x} + \sum_{k=0}^{K-1} v_{y,k} \phi_k \frac{\partial I}{\partial y} = 0, \quad (6)$$

which must be solved for each location in the image. In the case of a digital image, this corresponds to solving Eq. (6) for each image pixel.

Using the following definitions:

$$\begin{aligned} X &= [X_0, X_1, X_2, X_3, X_4, X_5, \dots] \\ &\equiv [v_{x,0}, v_{y,0}, v_{x,1}, v_{y,1}, v_{x,2}, v_{y,2}, \dots], \end{aligned} \quad (7)$$

$$A_{p,0} \equiv \phi_0 \left[ \frac{\partial I}{\partial x} \right]_p, \quad A_{p,1} \equiv \phi_0 \left[ \frac{\partial I}{\partial y} \right]_p,$$

$$A_{p,2} \equiv \phi_1 \left[ \frac{\partial I}{\partial x} \right]_p, \quad A_{p,3} \equiv \phi_1 \left[ \frac{\partial I}{\partial y} \right]_p, \dots, \quad (8)$$

$$B_p \equiv - \left[ \frac{\partial I}{\partial t} \right]_p, \quad (9)$$

where  $p$  represents the index of a particular pixel, Eq. (6) can be rewritten as a matrix equation to represent the full ensemble of pixels,

$$AX = B. \quad (10)$$

Here the size of matrix  $A$  is  $M \times 2K$ , the size of  $X$  is  $2K \times 1$ , and the size of  $B$  is  $M \times 1$ , where  $M$  is the number of pixels in the image (4096 for the images studied here) and  $K$  is the number of terms in the 2D velocity field expansion (also 4096 in this study). Each row of the matrices  $A$  and  $B$  represents a single pixel in the image.

Equation (10) is solved for  $X$ , which is then used to reconstruct  $v_x(x, y, t)$  and  $v_y(x, y, t)$  through Eq. (5) and the definition of  $X$  in Eq. (7).

## B. Implementation details

First, each frame in the raw-data sequence is filtered with a 2D median filter (using a three-point width) to remove noise spikes from the image frame. In some cases the image frames are additionally filtered with a 2D low-pass filter (with wave number limit of 10/frame), which eliminates pixel noise while retaining the primary intensity patterns. This step can be quite important in practice, as the velocity routines are quite sensitive to spurious spikes in the data.

As mentioned above, the technique relies on the constraint or assumption of constant brightness, and thus the local brightness must be normalized to the average temperature-dependent emission profile of the gas cloud. This is done by simply dividing the brightness map of each frame by the brightness map of the 300-frame time average.

Because matrix  $A$  in Eq. (10) is not square, some additional steps are required to solve for the array  $X$ , the coefficients of the velocity field components. Following Ref. 20, a second matrix  $\hat{A}$  can be created from the smoothed image set, which can then be transposed and applied to Eq. (10) to form

$$\hat{A}^T A X = \hat{A}^T B. \quad (11)$$

The square matrix ( $\hat{A}^T A$ ) can then be inverted to solve for  $X$ ,

$$X = (\hat{A}^T A)^{-1} \hat{A}^T B. \quad (12)$$

The family of basis functions for the velocity maps was chosen to reflect the typical shape of the structures observed in the GPI data. Wavelet basis functions were chosen (rather than, for instance, a Fourier basis) to most completely describe the localized GPI structures using the lowest order components. Additionally, the 2D wavelet basis was to be composed of separable 1D functions. Namely,

$$\phi_k(x, y) = \phi_i(x) \phi_j(y), \quad (13)$$

where  $\phi_i(x)$  and  $\phi_j(y)$  are the 64-member families of 1D basis functions,  $\phi_k(x, y)$  is the 4096-member family of 2D basis functions, and  $k=64i+j$ .

With these restrictions in mind, the discrete second-order Coiflet wavelet basis, shown in Fig. 2(a), was chosen. On a 64-pixel grid in each dimension, the 1D basis set consists of 64 discrete wavelet functions, derived from the shape in Fig. 2(a) through spatial scaling, spatial translation, and magnitude scaling. The 2D basis set is then formed from the 1D basis set using Eq. (13). Two example 2D wavelets are shown in Figs. 2(b) and 2(c).

Although the optical flow routine derives a “dense” flow field which is resolved to the image resolution (i.e.,  $64 \times 64$  velocity vectors for each frame), it is generally easier to visualize the flow field with somewhat fewer vectors shown. All results in this study are thus subsampled into a  $16 \times 16$  vector field for each frame.

Because the optical flow algorithm necessarily derives velocity vectors only perpendicular to the intensity gradients (and because the algorithms should derive smooth velocity distributions over real, noisy data), it is important to choose a suitable smoothing technique and apply it at an appropriate level. Namely, as in any analysis of experimental data, one should smooth the results such that they wash out spurious noise while retaining the observation of the physical phenomena. For this study, we have preconditioned the data first with a three-pixel median filter and then with a mild low-pass filter, as described above.

Once the velocity map has been derived, it is further smoothed both spatially and temporally, by convolution with a Gaussian kernel. The Gaussian is defined by a separate full width at half maximum (FWHM) in space (pixels) and time (frames). Typical results presented here have been smoothed spatially with a Gaussian of FWHM=5 pixels (out of 64), and temporally with a Gaussian of FWHM=3 frames. In tests of artificial data sets and real NSTX data sets with easily identifiable motion, these settings have provided substantial correction of local aperturing (described below) while maintaining the spatial and temporal details of the motion. The velocity results for any particular data set, however, are significantly dependent on the smoothing characteristics, leading to the eventual conclusion that the optical flow technique when used on its own is unsatisfactory.

## C. Limitations of optical flow technique

The primary limitation of the optical flow technique is that the algorithm can only detect motion along an intensity gradient. This is commonly referred to as the “aperture problem,” presents a critical limitation on the determination of plasma velocity from sequences of intensity images, and must be taken into account when interpreting the velocity results derived from diagnostics of this type. The aperture problem as applied to the GPI diagnostic will be explored further in the following sections, but can be essentially understood with a simple example. Consider a rotating barber pole containing helical red and white stripes. While the actual physical motion of the stripes is horizontal (across the






FIG. 2. (a) 1D 64-point Coiflet wavelet basis function, from which the 4096 2D basis functions were derived. [(b) and (c)] Example 2D wavelets  $\phi_{4,4}$  and  $\phi_{0,27}$ .

field of view), a sequence of images that did not contain the edges of the structure would identify the velocity direction as diagonal, perpendicular to the direction of the stripes. (This particular example contains a further subtlety, in that the horizontal boundaries of the visual image lead to a false cor-

rection to the perceived direction of the motion, and we intuitively interpret the stripes to be moving vertically.)

The aperture problem has historically been confronted using a variety of techniques, a common method being to simply smooth the velocity field over an appropriate region. The smoothing method has been used in this study, and the choice of smoothing parameters is discussed in Sec. V.

With proper smoothing of the derived velocity fields, this technique can excel at the identification and tracking of “blobs” or identifiable moving structures. On the other hand, the technique can conclude nothing about plasma velocity perpendicular to the intensity gradient, for example, laminar poloidal flow in a quiescent plasma with static radial density and temperature gradients.

Secondly, an important limitation on this technique is that the data must be sufficiently time resolved in order for the structures to be successfully tracked. In short, the algorithm can fail if the translation from frame to frame is so great as to cause the routines to “lose track” of the motion. Namely, for large frame-to-frame translations, the algorithm can interpret the structure as having simply disappeared from the first frame and reappeared in the second, rather than to have traveled from one location to the other.

To test this effect as it applies to the GPI data sets, a 1D version of the optical flow algorithm was used to measure the velocity of a translating Gaussian “pulse.” The array of test cases was set up with a range of values for both Gaussian FWHM and imposed velocity. Intuitively, one might expect that the derived velocity would no longer reflect the actual imposed velocity when the frame-to-frame translation exceeds the scale size of the structure shape.

The results of this test are plotted in Fig. 3, which shows the ratio of  $v_{\text{derived}}/v_{\text{imposed}}$  as a function of imposed velocity (measured in pixels/frame) and Gaussian FWHM (measured in pixels). Ideally (for perfect tracking) this ratio is 1.0, and in practice the optical flow algorithm reaches a ratio exceeding 0.9 (i.e., relatively accurate reconstruction) for frame-to-frame translations of  $\leq 40\%$  of the FWHM of the structure size. It is also notable that this ratio does not reach values substantially greater than 1.0. Namely, when inaccuracies do occur, the algorithm consistently underestimates the actual velocity of the structures. This effect is compounded by the aperture effect (which is absent in 1D tests such as those presented in Fig. 3), which tends to underestimate the velocity further by deriving only the velocity component perpendicular to the gradient of the image intensity. The combined effect is demonstrated in Sec. V. The slightly oscillatory behavior seen in Fig. 3 is due to a pixel-aliasing effect in the 1D test.

## IV. PATTERN-MATCHING TECHNIQUE

### A. Basic approach

Based on the limitations of the optical flow technique, a second approach was adopted, in which subsections of each image are correlated with subsections of previous and subsequent images to form a map of local translation or velocity. This approach (pattern matching) is subject to its own limitations and indeed is found to be complementary to the op-




FIG. 3. Ratio of ( $v_{\text{derived}}/v_{\text{imposed}}$ ) for 1D test case of optical flow algorithm. Results are plotted against imposed velocity ( $v_{\text{imposed}}$ ), measured in pixels/frame. Contours represent various structure widths ( $w_{\text{FWHM}}$ ), measured in pixels. Typical NSTX “blobs” have characteristic width  $\geq$  ten pixels and identifiable velocities of up to ten pixels per frame (9 km/s for data recorded at 4  $\mu$ s per frame), shaded in the figure.

tical flow approach, with each overcoming a deficiency of the other. Namely, the pattern-matching technique provides a more accurate final result than optical flow but is dependent on a reasonable “initial guess.” When used in sequence, wherein the optical flow output is used as a starting point for the pattern-matching algorithm, the most accurate velocity reconstructions are achieved. The pattern-matching formulation used in this study are adapted from Ref. 21, though with several modifications as described below.

In this technique, shown schematically in Fig. 4, each image frame is divided into subsections, or “tiles,” and a single velocity vector is derived for each tile within each frame and each time point within the sequence. Each image [Fig. 4(b)] is broken into tiles (16 in this example). Figure 4(a) shows an overlay of two contour plots, representing two sequential image frames (first gray, second black). Each tile within the image is compared to a corresponding tile in the subsequent time frame, as shown in Fig. 4(c). An offset ( $\Delta x, \Delta y$ ) is then found which maximizes the similarity in the two corresponding tiles [Fig. 4(d)]. Each frame from the GPI camera is made up of a  $64 \times 64$  grid of pixels, which is typically divided up into a  $16 \times 16$  grid of tiles,  $4 \times 4$  pixels each. The local velocity vector is simply the  $x$  and  $y$  offset of the matching tile exhibiting the highest similarity, divided by the  $\Delta t$  between time frames.

Additionally, the algorithm incorporates a weighting scheme which weakly rewards the “smoothness” of the velocity fields. In practice, this preserves the velocity compo-




FIG. 4. Schematic representation of the pattern-matching technique. Each image (b) is broken into tiles (16 in this example). Diagram (a) shows an overlay of two contour plots, representing two sequential image frames (first gray, second black). Each tile is compared to a corresponding tile in the subsequent time frame (c). An offset ( $\Delta x, \Delta y$ ) is found which maximizes the similarity in the two corresponding tiles (d).

nents derived from strongly correlated regions, while preventing weakly correlated regions from producing spurious or arbitrary field vectors that obscure the strongly correlated fields.

In practice, the maximization of similarity between image subsections is determined by simply minimizing the absolute difference between the intensity, calculated pixel by pixel and integrated over each subsection, as follows.

Consider a series of image frames, defined by the local intensity  $I(x, y, t)$ . Each image is broken up into a grid of subsections (tiles), the resolution of which defines the final resolution of the velocity field. For each tile and each time point, a displacement vector is found which minimizes the function

$$F_D = \int_{\text{tile}} |I(x, y, t) - I(x + \Delta x, y + \Delta y, t + \Delta t)| dx dy, \quad (14)$$

where  $\Delta t$  is the time delay between each image frame,  $\Delta x$  and  $\Delta y$  are the horizontal and vertical displacements, respectively, and the integral is taken over the tile area (or summed over all pixels within the tile area). The local velocity field is then given by

$$\mathbf{v}(x, y, t) = v_x \hat{x} + v_y \hat{y} = \frac{\Delta x(x, y, t)}{\Delta t} \hat{x} + \frac{\Delta y(x, y, t)}{\Delta t} \hat{y}. \quad (15)$$

As with optical flow, this technique relies on the assumption or constraint of constant brightness, namely, that the brightness of a given structure within the image remains con-

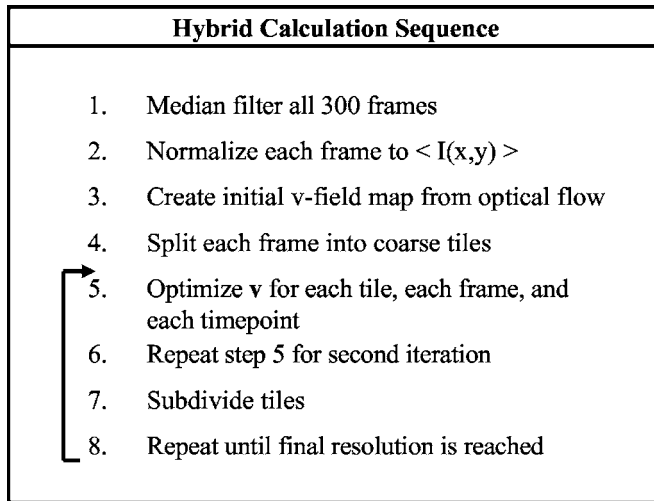


FIG. 5. Flow chart of the hybrid velocity field derivation technique used in this study.

stant during short displacements. In other words, flow can only be tracked if some brightness structure translates from one part of the image to another, rather than simply appearing or disappearing.

A smoothness constraint on the velocity maps is imposed by minimizing the function

$$c_s = \sqrt{\left(\frac{dv_x}{dx}\right)^2 + \left(\frac{dv_x}{dy}\right)^2 + \left(\frac{dv_y}{dx}\right)^2 + \left(\frac{dv_y}{dy}\right)^2}. \quad (16)$$

A “soft penalty” imposition of this constraint is used to allow for some natural degree of abrupt variation and is controlled through the use of a free parameter  $\tau_s$  in the Leclerc penalty function<sup>16</sup> so that the smoothness component of the final minimized quantity for the full image frame is given by

$$F_s = \sum_{\text{tiles}} [1 - \exp(-\tau_s c_s^2)], \quad (17)$$

where the summation is taken over the full grid of velocity vectors (one vector per tile) for each time point.

Finally, the relative weighting of the correlation function and the smoothness constraint is controlled by a free parameter  $\lambda$ , so that the final minimized function for each field is given by

$$F = \left(\sum_{\text{tiles}} F_D\right) + \lambda F_s. \quad (18)$$

Because of the generalized implementation of the weighted minimization scheme, it is relatively straightforward to introduce functions other than  $F_D$  and  $F_s$  used here, and in fact several functions were considered and rejected for various practical reasons.

## B. Implementation details

A schematic flow chart is shown in Fig. 5, outlining the most important details of the pattern-matching implementation.

First, each frame in the sequence is filtered for spikes in the data and normalized to the average emission profile of the gas cloud, just as was done for the optical flow technique.

Each frame is then split into subsections (tiles) and the optimization process for determining the v-field maps begins. For each tile in each frame at each time point, the local v vector is optimized by successively revising approximations to the magnitude and direction to finer and finer precision, each time calculating the error function in Eq. (18) until a local minimum is reached. This process necessarily begins with some initial guess, which can be randomly imposed or derived by some other method. As described below, the use of the optical flow results as an initial condition for this process produces the best results.

A local v-field vector is determined one at a time for each tile within each frame and then each frame within the sequence. Because of the finite weighting of the smoothing constraint [Eq. (17)], each tile depends on immediate neighboring tiles. Successive iterations of optimization process for the full data set will therefore provide further revisions to the final v-field values. In practice, it was found that the revisions resulting from a third or higher optimization cycle were insignificant, so the process was limited to two iterations.

It was also found that the initial conditions for the optimization process affected the final derived field, since the optimization depends on finding local minima to an error function [Eq. (18)]. Indeed, it was this arbitrary limitation which motivated the use of the optical flow output (which tends to be approximately correct at large scales) as a starting point for the pattern-matching algorithm. With this hybrid use of the two algorithms, the convergence of the final velocity maps has been very robust.

Once the v-field map for the entire data set has been optimized for a given spatial resolution of tiles, each tile is subdivided by a factor of two in each direction, and the process is repeated. The primary objective of this approach to finer and finer resolution is to save computation time, since lower resolution maps can often illuminate the global velocity behavior in a relatively short amount of time, providing the finer maps with a starting point which is closer to the final solution. For example, it was found that if the process was initiated with a grid of  $32 \times 32$  tiles, it took three iterations to achieve the level of convergence equivalent to that of only two iterations with the successively increasing resolution method.

A final detail is that because of the finite pixel size of the real images, only integer-pixel displacements are possible results of Eq. (14), which effectively imposes a precision limit to the derived velocity vectors. It is of course possible for the movement of structures to be accurately characterized to resolutions lower than a single pixel, which was achieved in practice by supersampling each image by a factor of 4 in each direction (using simple bilinear interpolation) prior to calculating Eq. (14) in the optimization process.

## C. Limitations of pattern-matching technique

The primary caveats of this technique are also exhibited by all techniques which derive flow from image sequences without additional constraints. As described in Sec. III C, these include the aperture effect and the limitations imposed by the finite time resolution of the data set.






FIG. 6. Intensity map from artificial test data set, typical of NSTX GPI images.

Significantly, however, the pattern-matching technique has an advantage over the optical flow technique in that the local velocity vectors are not constrained to lie along the local intensity gradient. Also, the pattern-matching technique is not subject to the limitations shown in Fig. 3, in which the maximum derivable velocity was limited to frame-to-frame translations of  $\leq 40\%$  of the structure size. Both of these differences act to alleviate the systematic underestimation of velocities exhibited by optical flow and represent a significant improvement over that technique used alone.

## V. ANALYSIS OF TEST DATA SETS

As a rudimentary test of the velocity routines, an artificial data set was constructed to simulate the response of the camera to typical NSTX plasma fluctuations. This data set was then prescribed with various known velocity fields and the resulting derived velocity fields compared to the known input.

The data set itself, shown in Fig. 6, is made up of a spectrum of spatial modes given by

$$I = \sum_{p=1}^M \sum_{q=1}^M \delta_{pq} \cos(p\kappa_x x) \cos(q\kappa_y y + \phi_{pq}), \quad (19)$$

with relative weights defined by

$$\delta_{pq}^2 \propto \exp\{-[(p\kappa_x - k_{x0})/\Delta k_x]^2 - [(q\kappa_y - k_{y0})/\Delta k_y]^2\}. \quad (20)$$

The spectrum is made up of discrete wave numbers  $p\kappa_x$  and  $q\kappa_y$  ( $\kappa_x$  and  $\kappa_y$  are constants),  $\phi_{pq}$  are random phases for each spectral component, and  $\Delta k_{x,y} = \kappa_{x,y} M/2$ . This is simply the summation of a 2D Gaussian spectrum of modes centered at  $(k_{x0}, k_{y0})$  with spectral widths  $\Delta k_x$  and  $\Delta k_y$ . For the test data set,  $k_{x0} = k_{y0} = 0.16 \text{ pixel}^{-1}$ ,  $\Delta k_x = \Delta k_y = 0.18 \text{ pixel}^{-1}$  and the spectral resolution was defined by  $M = 32$ . This spectrum

was chosen to represent the typical spatial characteristics of the NSTX GPI images.

To illustrate the performance of the hybrid technique and to demonstrate the improvement gain over the optical flow method, the test data set shown in Fig. 6 was translated horizontally, vertically, and diagonally over a range of velocities. Figure 7(a)–7(c) shows the derived velocity field maps corresponding to a translation of  $(v_x, v_y) = (-4, -4)$  pixels/frame for the optical flow method, the optical flow method smoothed with a Gaussian kernel of three pixels (of 64) FWHM, and the hybrid technique (unsmoothed), respectively. Ill-defined gradients at the very edge of the frame lead to poorly derived velocity vectors there, so Fig. 7 does not include these vectors. In the NSTX data sets, the emission intensity at the very edge of the frame is typically so low as to eliminate these velocity vectors in any case, so this does not impose an additional limitation on the derivation.

While the optical flow technique does reproduce the coarsest characteristics of the flow behavior, it is clearly inaccurate at smaller scales, and in particular suffers from spurious components in the flow field. This is primarily a result of the optical flow limitation that even in the ideal case, only the component of the velocity along the intensity gradient can be recovered. Further, for areas of the image with relatively weak gradients, the method fails altogether. While smoothing does help to “wash out” the spurious velocity components (and helps to interpolate over areas with insufficient gradients), it of course also acts to wash out real details in the flow. The smoothing parameter of three pixels FWHM [used in Fig. 7(b)] was chosen to strike a balance between these competing factors but is in general arbitrary. For a set of five frames [including that shown in Fig. 7(b)] within a sequence representing a constant velocity field of  $(v_x, v_y) = (-4.0, -4.0)$  pixels/frame, the mean values of the derived velocity components are  $(\langle v_x \rangle, \langle v_y \rangle) = (-0.53, -0.62)$  pixels/frame. The standard deviations of the derived  $v_x$  and  $v_y$  values over the spatial ensemble for this set are  $(\sigma_{v_x}, \sigma_{v_y}) = (0.41, 0.38)$  pixels/frame. Clearly the optical flow method is largely failing for this case, with the most significant failure in the underestimation of the velocity.

The hybrid technique [Fig. 7(c)] recovers the flow very uniformly and also very accurately. Significantly, the recovery does not depend on specific regions of stronger or weaker gradients, does not produce significant spurious velocity components, and accurately reproduces the overall magnitude of the flow. For a set of five similar velocity fields including that shown in Fig. 7(c), the mean derived velocity values are  $(\langle v_x \rangle, \langle v_y \rangle) = (-4.01, -3.99)$  pixels/frame, [to be compared to  $(v_x, v_y) = (-4.0, -4.0)$  pixels/frame]. The standard deviations for this set are  $(\sigma_{v_x}, \sigma_{v_y}) = (0.47, 0.47)$  pixels/frame.

Similar comparisons were performed over a range of imposed velocity values, including horizontal, vertical, and diagonal translations. The diagonal translation was tested to check whether the optical flow technique would perform better for motion perpendicular to the typical elongation of the structures within the image, which in this case would corre-



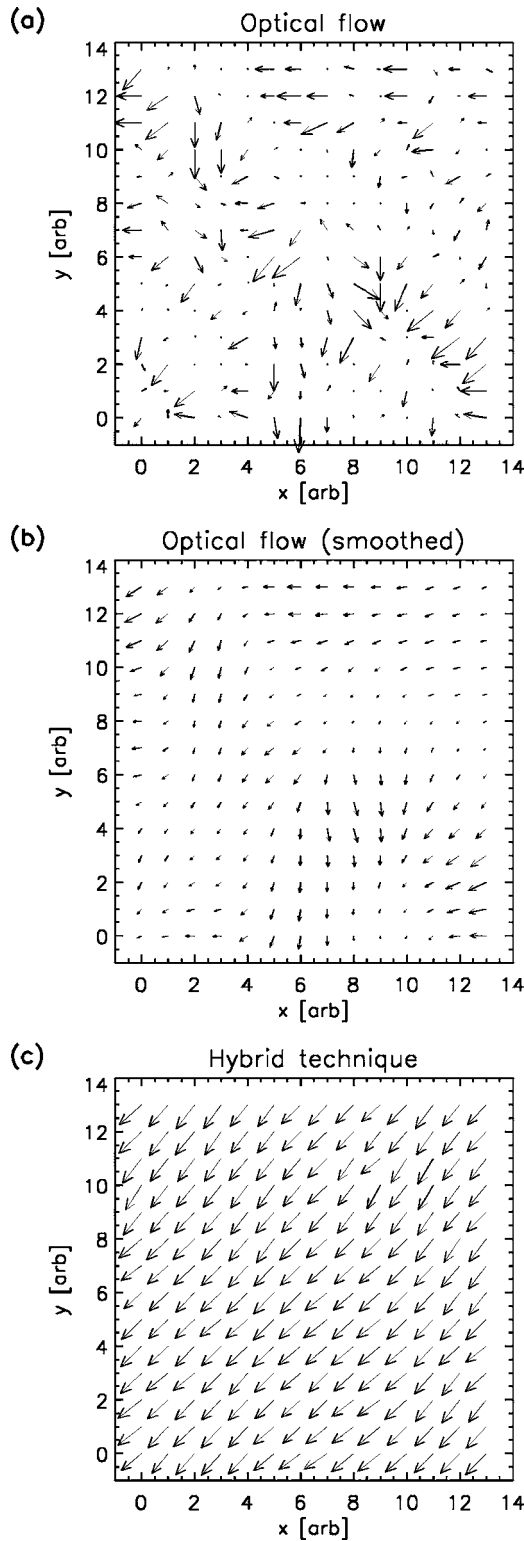


FIG. 7. Derived velocity vector maps from the diagonal translation of data set shown in Fig. 6, using the optical flow algorithm (a), the optical flow algorithm smoothed with three pixel FWHM Gaussian kernel on a  $64 \times 64$  pixel grid (b), and the hybrid optical flow/pattern-matching algorithm (c).

spond to motion in the negative  $x$  and  $y$  directions (see Fig. 6). The results of this series of tests is shown in Fig. 8 for velocity values up to  $|v|=9$  pixels/frame.

In Fig. 8(a), the horizontally translated data set, the derived mean values of  $\langle v_x \rangle$  and  $\langle v_y \rangle$  are plotted versus the

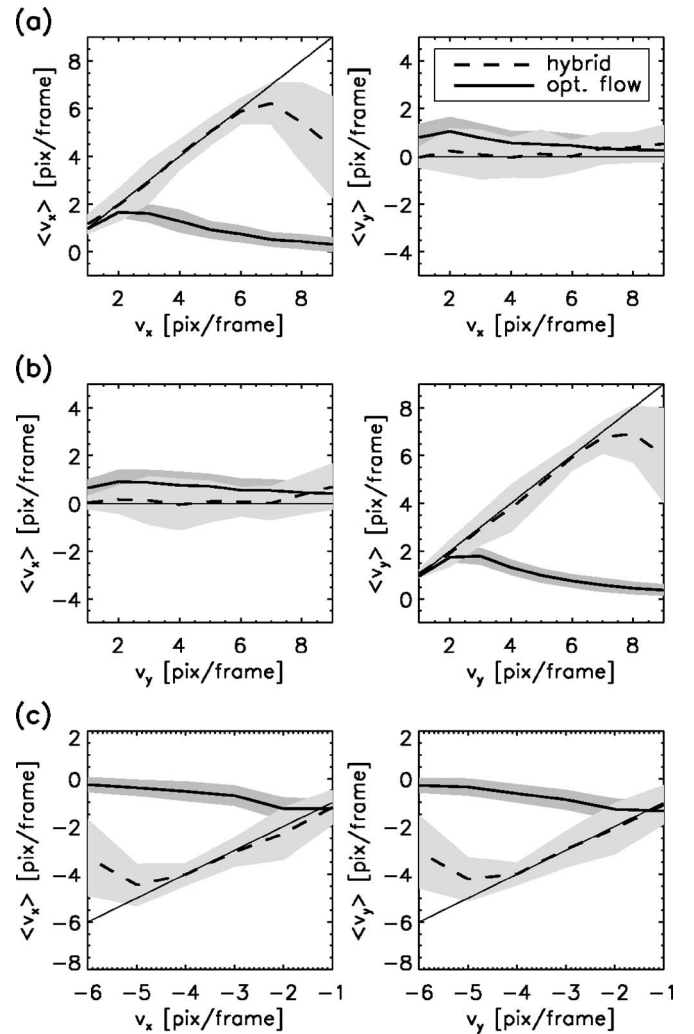


FIG. 8. Derived velocity values from horizontally (a), vertically (b), and diagonally (c) translated test cases, presented both for the smoothed optical flow technique (solid line) and the hybrid technique (dotted line).  $\langle v_x \rangle$  and  $\langle v_y \rangle$  are the spatial averages of the derived  $v_x$  and  $v_y$  values, and the shaded regions represent the standard deviation in the derived  $v_x$  and  $v_y$  values. The thin solid line represents perfect velocity recovery. The hybrid technique is very accurate up to translations of  $\sim 7$  pixels/frame (6.4 km/s), while the optical flow technique degrades at much lower velocities.

imposed or “true” values of  $v_x$  ranging from 1–9 pixels/frame ( $v_y=0$  for this set). The solid line represents the smoothed optical flow data [similar to Fig. 7(b)], and the dotted line represents the unsmoothed hybrid-technique data [similar to Fig. 7(c)]. In each plot, the shaded region represents  $\pm\sigma_{v_x}$  or  $\pm\sigma_{v_y}$  about the mean value, and the thin solid line represents perfect velocity recovery. Figure 8(b) represents the vertically translated data set, with  $\langle v_x \rangle$  and  $\langle v_y \rangle$  plotted versus the imposed values of  $v_y$  ( $v_x=0$  for this set).

In Fig. 8(c), the diagonally translated data set,  $\langle v_x \rangle$  and  $\langle v_y \rangle$  are plotted versus the imposed values of  $v_x$  and  $v_y$ , respectively, up to  $-6$  pixels/frame (here  $v_x=v_y$ ).

The hybrid technique is extremely accurate up to translations of  $\sim 7$  pixels/frame (corresponding to  $\sim 6.4$  km/s for NSTX data) and degrades for higher velocities. The optical flow technique degrades at much lower imposed velocities and shows significant noise even at the lowest velocities, for reasons described in Sec. III C.

It is notable that the 2D optical flow algorithm performs significantly worse than the 1D version (i.e., Fig. 3), due primarily to the smoothing made necessary by the pronounced aperture problem (which is not present in the 1D implementation). In particular, the optical flow test cases result in a consistent underestimation of the velocity magnitude. On the other hand, even when significantly underestimating the magnitude of the velocity, a smoothed optical flow result tends to get the overall direction of the flow correct. Indeed, this is precisely why the optical flow is useful as a starting point for the minimization process used in the pattern-matching technique and why the hybrid approach leads to more accurate solutions than the pattern-matching technique used alone with random initial conditions.

We have also applied the technique to test data sets subject to sheared flow, which may be centrally important to flow studies on NSTX. For this series of tests, we used a simple 2D sinusoidal intensity pattern  $I(x,y) = I_0 \sin(k_x x) \sin(k_y y)$  with  $k_x = k_y = 4\pi/D$ , where  $D$  is the image width or height (i.e., two maxima and two minima in each direction). This pattern was subjected to a sheared velocity  $v_y(x)$  composed of a uniform background velocity added to a “sawtooth” pattern of varying magnitude. For example, curves representing uniform background velocities of 0.0, 3.65, and 7.3 km/s added to a sawtooth pattern with magnitude of 3 km/s are shown as the gray curves in Fig. 9(b). The uniform velocity component was varied from 0–7.3 km/s, similar to the tests shown in Fig. 8. The maximum velocity within the sawtooth pattern, which had a spatial period of 32 pixels, was varied from 0 to 3 km/s, corresponding to a maximum velocity shear of  $\nabla v_y = 1.4 \times 10^5 \text{ s}^{-1}$ . The range over which to test the shear response was chosen based on the shear estimates from Ref. 22 which reached a maximum of  $\nabla v_{\text{pol}} \sim 1.0 \times 10^5 \text{ s}^{-1}$ .

A summary of the results is shown in Fig. 9. A representative velocity field derived from imposed  $\langle v_y \rangle = 0.0 \text{ km/s}$  and  $\nabla v_y = 1.4 \times 10^5 \text{ s}^{-1}$  is shown in Fig. 9(a). The spatial axes in this and following figures are scaled to duplicate the field of view in NSTX. The gray curves in Fig. 9(b) represent the imposed velocity, averaged over the same four-pixel tiles as used in the velocity derivation routine. As described above, three representative  $\langle v_y \rangle$  values of 0.0, 3.65, and 7.3 km/s are shown, all with shear of  $\nabla v_y = 1.4 \times 10^5 \text{ s}^{-1}$ . [The lower curve corresponds to the set shown in Fig. 9(a).] The black squares in Fig. 9(b) represent the derived  $v_y$  value, averaged over the vertical column of vectors and over the six frames of the test sequence. Error bars represent the standard deviation within each averaged ensemble.

The velocity shear does not appear to explicitly degrade the performance of the hybrid routine, and for local  $|v_y| \leq 6\text{--}7 \text{ km/s}$  the routine is able to successfully recover the imposed velocity. Similar to the test cases without shear (Fig. 8), the performance does degrade for local  $v_y$  values over  $\sim 6\text{--}7 \text{ km/s}$ , with a resulting underestimation of local velocity and increased scatter in the vector fields. Within the parameters of these tests, chosen to encompass the range of velocity shear exhibited in NSTX data, it appears that velocity shear does not impose any additional limitations or problems with the velocity recovery techniques.

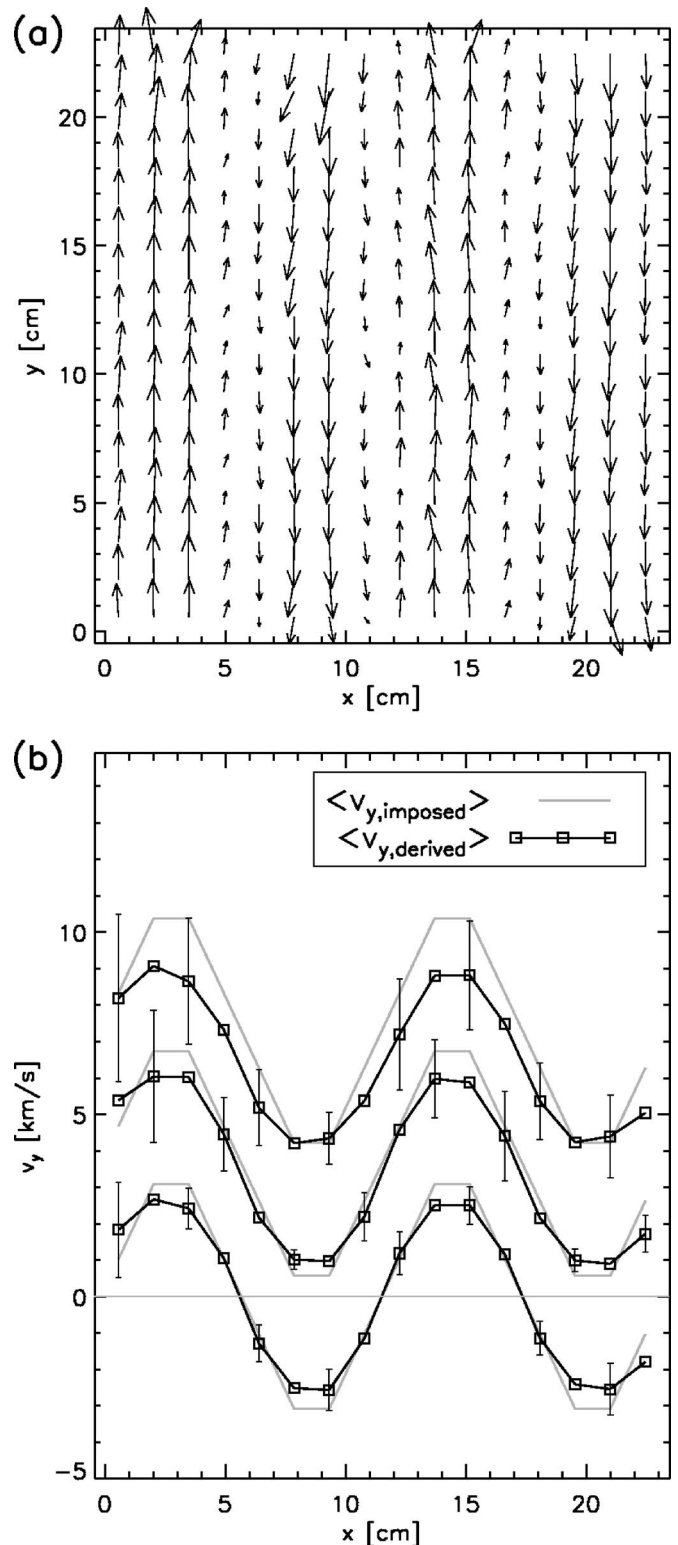


FIG. 9. Analysis of test data sets with sheared flow. Representative derived velocity field with  $\langle v_y \rangle = 0.0 \text{ km/s}$  and  $\nabla v_y = 1.4 \times 10^5 \text{ s}^{-1}$  is shown in (a). Imposed and derived velocity profiles for  $\langle v_y \rangle = 0.0, 3.7,$  and  $7.3 \text{ km/s}$  and  $\nabla v_y = 1.4 \times 10^5 \text{ s}^{-1}$  are shown in (b). The velocity shear does not explicitly degrade the velocity recovery, though the results do degrade beyond local  $v_y$  values of  $\sim 7 \text{ km/s}$ , consistent with Fig. 8.

## VI. MEASUREMENTS OF NSTX DATA

As an example of the technique applied to NSTX data, Fig. 10 shows the results of NSTX shot 113487, in which a

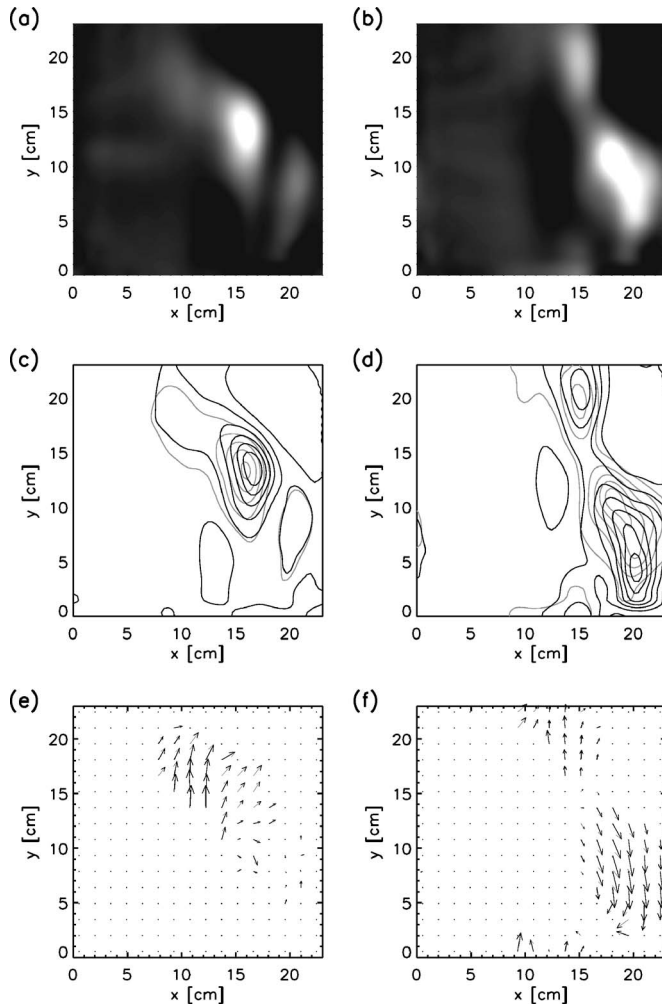


FIG. 10. Derived velocity fields for NSTX shot 113487 at  $t=300.196$  ms [(a), (c), and (e)] and  $t=300.216$  ms [(b), (d), and (f)]. Intensity maps are shown in (a) and (b). Contour plots for two successive time points (separated by  $4 \mu\text{s}$ ) overlaid in (c) and (d). Derived velocity maps shown in (e) and (f). The velocity of the brightest plasma blob visible in (a) is  $1.8$  km/s.  $20 \mu\text{s}$  later (b), this blob has split into two blobs: one moving upward at  $2.5$  km/s and one moving downward at  $5.2$  km/s.

blob of plasma is created, ejected from the plasma edge, and then bifurcates into two blobs which travel vertically away from each other. The exposure set for this shot was recorded at  $250\,000$  frames/s or  $4 \mu\text{s}$  per frame. Figures 10(a) and 10(b) show the intensity maps for two frames separated by  $20 \mu\text{s}$  ( $t=300.196$  ms and  $300.216$  ms, respectively). In these and subsequent figures, the negative (downward) direction is the ion diamagnetic drift velocity direction. Figure 10(c) is a contour map of the intensity from  $t=300.196$  ms (gray), overlaid with the subsequent time frame  $t=300.200$  ms (black) showing the relative frame-to-frame motion of the plasma blob. Figure 10(d) corresponds to time points  $t=300.216$  ms (gray) and  $t=300.220$  ms (black). Figure 10(e) shows the derived velocity fields for  $t=300.196$  ms, showing the quantifiable upward/outward motion of the first blob at  $(\langle v_x \rangle, \langle v_y \rangle) = (1.2, 1.3)$  km/s. Figure 10(f) shows the vector field for  $t=300.216$  ms, in which the blob has bifurcated into two oppositely moving parts, one moving upward at  $(\langle v_x \rangle, \langle v_y \rangle) = (0.2, 2.5)$  km/s and one moving downward at  $(\langle v_x \rangle, \langle v_y \rangle) = (0.3, -5.2)$  km/s. In Figs.

10(e) and 10(f) a brightness threshold is set, below which the velocity vectors are not shown (i.e., vectors are displayed only where there are identifiable structures in the intensity).

## VII. COMPARISON TO PREVIOUS NSTX ANALYSIS

It is instructive to perform a direct comparison of these results to previous velocity analysis of NSTX GPI data, in particular, the study presented in Ref. 22. In that analysis, the poloidal velocity at each radius was derived by comparing the intensity wave form (intensity versus  $t$ ) of the central pixel of each column within the image to the wave forms of the other pixels within that column. A time lag was imposed to maximize the cross correlation between each pair of wave forms, and a linear fit was made to the ensemble of spatial separations ( $\Delta y$ ) and time lags ( $\Delta t$ ) for the column of pixels.

In many cases, especially for  $H$ -mode shots, a particular column of data is dominated by the single brightest structure, and the results of the lagged-correlation analysis tend to represent the velocity of that structure. Indeed, that can often represent the quantity which is most physically interesting: the motion of the brightest structures.

In cases where a particular column of data contains more than one bright structure, and especially in cases where there are many structures, the lagged-correlation analysis tends to either track the brightest structure within the group, or else does not find a good linear fit to the  $\Delta y$  vs  $\Delta t$  for the ensemble of wave form comparisons.

Thus when comparing the “average” poloidal velocity from the lagged-correlation technique to the hybrid technique, one expects to find a good match if (a) a single bright structure is dominant and (b) the average of the hybrid-technique output is limited to include only those vectors corresponding to brightness and cross correlation above some threshold; in other words, if we limit the comparison to velocity vectors corresponding to “good tracking” (highly correlated) and bright structures. On the other hand, in cases where there are many structures in the column of data (which may be moving at different speeds and/or directions), one expects a poorer match between the average poloidal velocity from the two techniques, and furthermore the comparison can depend on which vectors one chooses to include in the average (i.e., which vectors represent the same structures as those represented by the lagged-correlation technique).

This is illustrated in Fig. 11, which shows a comparison of the average poloidal velocity derived by the two techniques.

Figure 11(a) is the brightness data from shot 113733, column 39, plotted against  $y$  and  $t$ . This data set contains one dominant structure (at  $\sim 0.75$  ms between  $5$ – $15$  cm) and several subdominant structures. The single dominant structure dominates the lagged-correlation analysis, leading to an “average velocity” measurement for this column of  $4$  km/s. Again, while this may not be a true average of the structure velocity within these data, it may indeed represent the quantity of most physical interest (the motion of the brightest structure).

Figure 11(b) is the brightness data from column 26 (also






FIG. 11. Comparison of lagged-correlation analysis to hybrid technique for NSTX shot 113733. (a) and (b) show the image intensity vs  $y$  and  $t$  for columns 39 and 26, respectively. (c) shows the derived poloidal velocity profile (solid lines) for two different thresholding values, compared to the lagged-correlation results (squares).

plotted against  $y$  and  $t$ ), which contains dozens of small moving structures, all at comparable brightness. Structures in this column are moving in both the positive and negative directions. The lagged-correlation analysis for this column did not produce a good linear fit to the  $\Delta y$  vs  $\Delta t$  values for the lagged correlations, and so no average velocity was found for this column.

Figure 11(c) is a plot of the average poloidal velocity profile derived from the lagged-correlation technique (squares) plotted versus  $x$ , the relative horizontal position in the NSTX field of view. Here the error bars represent the error from the linear fit to the data within each column. The black curve is the average velocity from the hybrid technique, including only points above a high threshold for normalized brightness and correlation. As discussed above, the two techniques agree well around column 39.

The gray curve in Fig. 11(c) is the average velocity from the hybrid technique, including velocity vectors above a significantly lower threshold for brightness and correlation. Notably, the absolute value of the gray curve is lower than the black (highly thresholded) curve at both column 39 and column 26 and indeed is close to zero at column 26, representing the fact that column 26 contains many structures which are moving in both the positive and negative directions [as seen in Fig. 11(b)].

Thus we can see that the degree to which the two techniques give the same average velocity results depends on the degree to which the “thresholding” of the vectors from the hybrid technique matches the dominance of the individual structures tracked by the lagged-correlation technique.

To further illustrate this point, the poloidal velocity profiles derived with the two techniques are compared in Fig. 12, for representative  $L$ -mode and  $H$ -mode shots 113744 and 113748. The  $L$ -mode shot [Fig. 12(a)] contains a large number of structures which do not move together, and thus the average poloidal velocity from the hybrid technique does not precisely match the derived velocity from the lagged-correlation technique, which primarily represents the motion of the brightest structure. The  $H$ -mode shot [Fig. 12(a)], on the other hand, is dominated by a small number of bright structures, which are similarly tracked by both techniques, producing a nearly identical derived velocity profile.

## VIII. DISCUSSION AND FUTURE WORK

We have implemented a novel technique for extracting time-resolved, 2D velocity field maps for use in plasma turbulence studies, in particular, for analyzing data from the gas puff imaging diagnostic instrument on NSTX. We have overcome certain limitations inherent to other velocity extraction algorithms, and in fact we have found that the most effective algorithm is a staged combination of optical flow and pattern-matching techniques.

We find that the primary limitations of the optical flow algorithm, when used alone, are an unsatisfactory upper limit on the derivable velocity and a pronounced aperture effect, wherein the algorithm can only extract the velocity component along the intensity gradient. It should be emphasized that the velocity limit is formally a limit on the recoverable frame-to-frame translation and thus could improve linearly with an increase in the framing rate of the camera.

While the tiled pattern-matching technique is not susceptible to either of these limitations, the accuracy of the results are dependent on the initial conditions fed to the minimization error function. Namely, if *no* velocity information is known *a priori* and a purely random field is used as input to

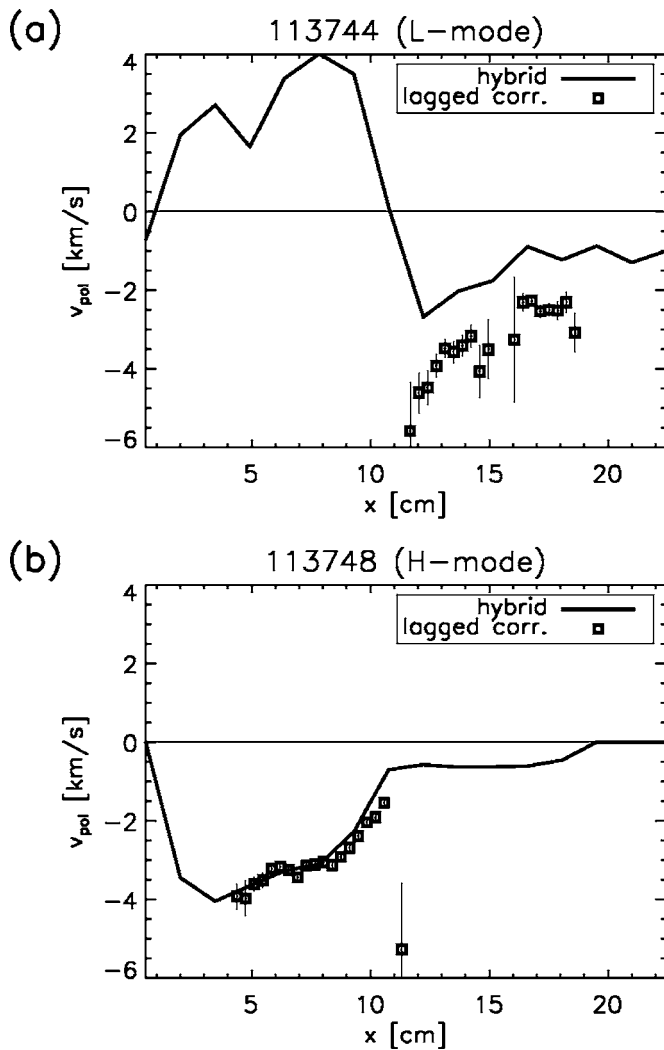


FIG. 12. Comparison of poloidal velocity profiles derived with lagged-correlation analysis and hybrid techniques, for representative *L*-mode and *H*-mode shots. A close match is found for the *H*-mode shot, which is dominated by a small number of bright structures. For the *L*-mode shot, which contains a larger number of structures which do not move together, the average velocity from the hybrid technique does not represent the poloidal velocity from the lagged-correlation technique.

the pattern-matching technique, it is susceptible to spurious solutions corresponding to local minima in the error function.

Each of these limitations can be largely overcome if the optical flow output is used as an input to the pattern-matching algorithm. Such a hybrid technique can successfully recover velocities up to  $\sim 6$  km/s for test cases duplicating NSTX data.

We also find that when the results of the hybrid technique are compared to lagged-correlation analysis of wave forms from pairs of individual pixels, an approach which has been used in a number of previous studies, the degree of

similarity depends on the characteristics of the structures within the data set. Namely, the degree to which the two techniques produce the same results reflects the degree to which the “thresholding” of the hybrid data included in the poloidal average matches the dominance of the brightest structures measured by the lagged-correlation analysis. This typically occurs for *H*-mode shots, which are dominated by a few bright structures.

The fundamental advantage of the hybrid technique presented here is that it produces a “dense” velocity field, resolved both in space and time. These data sets can be used in future studies to investigate flow details such as divergence and vorticity and to investigate higher statistical moments of the flow. Additionally, characteristics of the radial flow and its relationship to intermittent radial transport will be investigated in future articles.

## ACKNOWLEDGMENTS

The authors would like to thank D. D’Ippolito, C. Holland, R. Maqueda, G. R. McKee, J. Myra, T. Stoltzfus-Dueck, D. Stotler, and the NSTX team for many helpful discussions. This work was supported by DOE contract no. DE-AC02-76CHO3073.

- <sup>1</sup>S. Paul and R. Fonck, *Rev. Sci. Instrum.* **61**, 3496 (1990).
- <sup>2</sup>G. McKee, R. Ashley, R. Durst, R. Fonck, M. Jakubowski, K. Tritz, K. Burrell, C. Greenfield, and J. Robinson, *Rev. Sci. Instrum.* **70**, 913 (1999).
- <sup>3</sup>G. Cima, B. Deng, C. W. Domier, W. R. Geck, R. P. Hsia, C. Liang, F. Jiang, N. C. Luhmann Jr., and D. Brower, *Fusion Eng. Des.* **34**, 515 (1997).
- <sup>4</sup>B. H. Deng, C. W. Domier, N. C. Luhmann, Jr., D. L. Brower, A. J. H. Donné, T. Oyevara, and M. J. van de Pol, *Phys. Plasmas* **8**, 2163 (2001).
- <sup>5</sup>E. Mazzucato, *Nucl. Fusion* **41**, 203 (2001).
- <sup>6</sup>T. Munsat, E. Mazzucato, H. Park, B. H. Deng, C. W. Domier, N. C. Luhmann, Jr., A. J. H. Donné, and M. J. van de Pol, *Rev. Sci. Instrum.* **74**, 1426 (2003).
- <sup>7</sup>R. J. Maqueda *et al.*, *Rev. Sci. Instrum.* **74**, 2020 (2003).
- <sup>8</sup>S. Zweben *et al.*, *Nucl. Fusion* **44**, 134 (2004).
- <sup>9</sup>J. Boedo *et al.*, *Phys. Plasmas* **8**, 4826 (2001).
- <sup>10</sup>G. Y. Antar, S. I. Krasheninnikov, P. Devynck, R. P. Doerner, E. M. Hollmann, J. A. Boedo, S. C. Luckhardt, and R. W. Conn, *Phys. Rev. Lett.* **87**, 065001 (2001).
- <sup>11</sup>G. McKee, R. Fonck, D. Gupta, D. Schlossberg, M. Shafer, C. Holland, and G. Tynan, *Rev. Sci. Instrum.* **75**, 3490 (2004).
- <sup>12</sup>C. Holland, G. Tynan, G. McKee, and R. Fonck, *Rev. Sci. Instrum.* **75**, 4278 (2004).
- <sup>13</sup>J. Terry, S. Zweben, O. Grulke, M. Greenwald, and B. LaBombard, *J. Nucl. Mater.* **337**, 322 (2005).
- <sup>14</sup>B. Horn and B. Schunck, *Artif. Intell.* **17**, 185 (1981).
- <sup>15</sup>B. Horn and B. Schunck, *Artif. Intell.* **59**, 81 (1993).
- <sup>16</sup>T. Corpetti, E. Memin, and P. Perez, *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 365 (2002).
- <sup>17</sup>R. Larsen, K. Conradsen, and B. Ersbøll, *IEEE Trans. Geosci. Remote Sens.* **36**, 256 (1998).
- <sup>18</sup>S. Beauchemin and J. Barron, *ACM Comput. Surv.* **27**, 433 (1995).
- <sup>19</sup>D. Hitt, M. Lowe, and R. Newcomer, *Phys. Fluids* **7**, 6 (1995).
- <sup>20</sup>S. Srinivasan and R. Chellappa, *J. Opt. Soc. Am. A* **16**, 493 (1999).
- <sup>21</sup>R. Wildes, M. Amabile, A. Lanzillotto, and T. Leu, *Comput. Vis. Image Underst.* **80**, 246 (2000).
- <sup>22</sup>S. Zweben *et al.*, *Phys. Plasmas* **13**, 056114 (2006).