

Computation of finite-time Lyapunov exponents from time-resolved particle image velocimetry data

Samuel G. Raben · Shane D. Ross ·
Pavlos P. Vlachos

Received: 24 April 2013 / Revised: 19 November 2013 / Accepted: 22 November 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract This work presents two new methods for computing finite-time Lyapunov exponents (FTLEs) from noisy spatiotemporally resolved experimentally measured image data of the type used for particle image velocimetry (PIV) or particle tracking velocimetry (PTV). These new approaches are based on the simple insight that the particle images recorded during PIV experiments represent Lagrangian flow tracers whose trajectories lend themselves to the direct computation of flow maps, and related quantities such as flow map gradients and FTLEs. We show that using this idea we can improve the reliability and accuracy of FTLE calculation through the use of either direct pathline flow map (PFM) calculation, where individual particle pathlines over a fixed period of time are used to determine the flow map, or particle tracking flow map compilation (FMC), where instantaneous tracking results are used to estimate small snapshots of the flow map which are then compiled to describe the complete flow map. Comparisons of the traditional velocity field integration (VFI) method for computing FTLE fields with these new methods show that FMC produces significantly more accurate estimates of the FTLE field for both synthetic data and experimental data especially in cases where the particle number density is low. This is because the VFI estimates particle motion

while PTV directly measures particle motion and therefore generates a more accurate flow map. Overall, our results suggest that VFI is not always a reliable approach when applied to noisy experimental PIV data. For cases where particle loss between frames is minimal, the PFM can also produce better results, but the final field is susceptible to error due to the unstructured nature of the raw flow maps. When comparing the ability to match the true separatrix of a flow, FMC is shown to be a far superior method. The separatrix from FMC has an 80 % overlap with the true solution as compared to approximately 25 % for the PFM and only 1 % for the VFI method. FMC shows a significant advantage when the particle seeding is low, which is particularly relevant for applications to environmental or biological flows where adding seed particles is not always practical, and investigation of Lagrangian flow structures must rely on naturally occurring flow tracers.

1 Introduction

Finite-time Lyapunov exponents (FTLEs) can provide information on the mixing and transport mechanisms in a flow (Brunton and Rowley 2010; Shadden 2011), and they are applicable in both turbulent and laminar flows (Haller 2001). FTLEs are a measure of the exponential rate of divergence or convergence of Lagrangian particle trajectories over a finite time. The calculation of FTLEs has been used previously in both experimentally (Shadden et al. 2007; Peng and Dabiri 2009) and computationally generated (Haller 2001; Wilson et al. 2009) flow fields.

The ridges or high-magnitude locations in FTLE fields are referred to here as Lagrangian coherent structures (LCSs). Ridges can be defined precisely by appealing to differential geometric quantities (Shadden et al. 2005).

S. G. Raben · P. P. Vlachos (✉)
Department of Mechanical Engineering,
Virginia Tech, Blacksburg, VA, USA
e-mail: pvlachos@vt.edu

S. G. Raben
e-mail: sraben@vt.edu

S. D. Ross
Department of Engineering Science and Mechanics,
Virginia Tech, Blacksburg, VA, USA
e-mail: sdross@vt.edu

According to Haller (2011), LCS is defined more restrictively as hyperbolic material surfaces with extreme finite-time normal repulsion or attraction. This eliminates spurious LCS such as those due to shear or stretching. Thus, it is more correct to consider ridges of FTLE as candidate LCS, and while they provide insight into the Lagrangian skeleton of the flow, further criteria must be satisfied to classify a candidate LCS as a hyperbolic LCS. Nevertheless, FTLE ridge features can be used to determine the underlying transport structure in complex flow fields (Holmes et al. 1996; Lekien and Ross 2010; Senatore and Ross 2011), revealing mixing barriers that inhibit transport, or when there is a high density of rapidly moving LCS, regions of increased mixing (Shadden et al. 2005; Tallapragada and Ross 2008; Peng and Dabiri 2009).

While the calculation of LCS from FTLE fields has been used to better understand fluid dynamics behavior both numerically (Haller and Yuan 2000; Lekien and Ross 2010; Tallapragada and Ross 2013) and experimentally (Shadden et al. 2006; Shinneeb et al. 2006; Mathur et al. 2007; Charonko et al. 2013), their application has been limited. Part of the limitation is due to the high computational cost in calculating the FTLE fields. Currently, in order to calculate these fields, artificial flow tracers are numerically advected in time to determine the flow map for a flow field (Haller 2002). Determining this flow map using conventional methods is computationally costly (Brunton and Rowley 2010) and can yield incorrect values near boundaries (Ruiz et al. 2010). Additionally, methods have been developed for decreasing the computation cost of determining and tracking the motion of LCSs for a given flow field (Lipinski and Mohseni 2010).

The current work aims to efficiently and accurately calculate flow maps and the associated FTLE fields from experimentally measured particle image data of the type used for particle image velocimetry (PIV) or particle tracking velocimetry (PTV). Both are techniques for non-invasively measuring a fluid flow experimentally where flow tracers are added to a flow to probe the fluid's motion (Adrian 1991, 2005; Raffel et al. 1998). When properly matched with the experimental flow conditions, these particles are assumed to act as perfect flow tracers (Raffel et al. 1998), with zero response delay to the flow fluctuations, i.e., similar to the numerical particles that are artificially added to the flow during the numerical calculation of the flow maps.

Previous work has shown that the computational cost can be decreased when investigating successive FTLE fields from a single data set by reducing the number of redundant particle integrations (Brunton and Rowley 2010) by numerically integrating particle trajectories over small time intervals to create short snapshots of the flow map. By using compositions of these flow map snapshots, FTLE

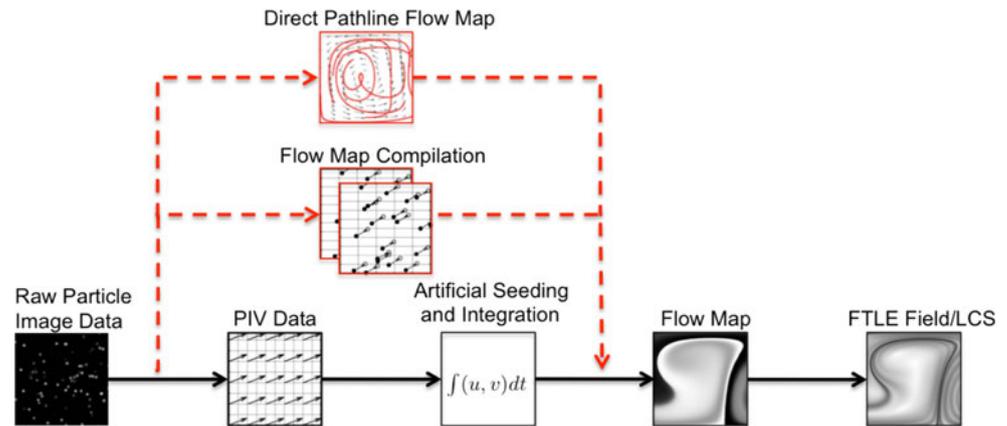
fields over different initial and final integration times can be computed quickly and efficiently (Brunton and Rowley 2010). This procedure allows for different time intervals to be linked together providing a flow map at varying time intervals with reduced computational cost and has been shown to work with experimental data (Green et al. 2011).

While this approach reduces the computational cost, it is still expensive and moreover does not take advantage of the flow map information contained in experimental flow fields derived from particle images, i.e., the Lagrangian motion of the particle flow tracers. When PIV or PTV images are acquired, they inherently contain information about the flow map over a short-time snapshot, i.e., the discrete particle motion from one frame to the next. This simple insight suggests that if properly extracted and compiled, these measured particle trajectories can be used to more accurately and directly measure the flow map, and therefore the FTLE/LCS field, from experimental PIV data.

This work presents a new method for FTLE calculation by simply measuring the flow map directly from the flow tracer particles present in the field. Compared with the conventional velocity field integration (VFI) method, this new method offers an alternative that is better suited for use with particle image data because it exploits information already inherent in the recorded images to deliver increased accuracy and robustness. A previous study has looked at using particle tracking information to compute FTLEs, but this method still required a numerical integration step (Voth et al. 2002). While this study claimed poor performance of the tracking-based method, improvements in the tracking procedure as well as the overall methodology will increase the performance of these methods. By measuring the motion of each of these particles from one time instant to the next, through the use of PTV, a direct measure of the flow map can be produced. This new procedure, illustrated in Fig. 1, eliminates the need for costly numerical integration. In addition, using PTV also has the additional benefit of better resolving near-wall flows (Kahler et al. 2012), which can increase the accuracy of the FTLE/LCS near boundaries.

Two different forms of particle tracking can be utilized in order to determine the required flow maps for a given field. The first form is to use direct pathline flow map calculation (hereinafter referred to as PFM), where the pathline for individual particles located in the images is constructed over a series of consecutive frames throughout the entire time of interest. The second method is to use the particle tracking between two adjacent frames and then combine these successive flow maps via interpolation. A compilation of this fashion was first proposed for numerical data by Brunton and Rowley (2010). The advantage of PFM is that the particle pathlines are direct measures of the flow map as they provide the exact path of a given particle

Fig. 1 Flow chart showing the common procedure for computation of the FTLEs from experimental data and our novel procedure denoted with the red dashed line



for a specific location. The disadvantage to this method is that experimentally it can be challenging to follow a particle over an extended period of time. Particles can be “lost” due to improper matching in the tracking algorithm, or in the case of planar PIV/PTV, they can be lost due to out-of-plane motion. Using tracking information only between two frames has the advantage over PFM that it can provide a higher density of successfully tracked particles. When limited to only two frames, the tracking algorithm will not be as heavily influenced by losses and when properly compiled should provide a very good estimate of the true flow map.

The following provides a detailed analysis of using both Lagrangian particle tracking and instantaneous particle tracking for flow map determination of FTLE fields from experimental data. First, a synthetic data set was created to test the methods as well as perform error analysis. This was followed by a demonstration of the methods on an experimental data set of a vortex ring.

2 Flow map determination

This section will describe the two newly developed methods for flow map estimation which are based on direct particle tracking of the flow tracer images, as well as the classical velocity field integration approach, that is based on numerical integration of trajectories of artificial fluid elements based on cross-correlation-based PIV velocity measurements and is used here as a benchmark for comparison.

2.1 Flow map computation using particle tracking methods

This work uses a multicomponent particle tracking algorithm that has been developed for single and multiphase flows (Cardwell et al. 2010). The process works by

calculating unique particle identifiers such as particle size, shape, and maximum intensity and uses this information to parametrically track the particles. By changing the relative weighting of these factors, along with the interparticle distance of the possible matched particles, the algorithm can be tuned to work in highly turbulent flows with particles that may not follow the flow. Even nonflow tracers can lead to informative FTLE analysis (Tallapragada and Ross 2008; Peng and Dabiri 2009). When the particles follow the flow path, the algorithm can also use information produced from PIV to more efficiently predict particle motion and pair the particles. Using this method has shown great improvements in both the probability to match particles and in the accuracy of those matches (Cardwell et al. 2010). However, there are several other efficient PTV methods that can also be used to produce accurate PTV velocity field estimations (Ohmi and Li 2000; Mikheev and Zubtsov 2008; Ohmi and Panday 2009).

2.1.1 Flow map compilation from instantaneous tracking

Because two-frame tracking results are direct measures of incremental flow maps over the interframe time, i.e., the particle motion between frames, it is possible to achieve more accurate measurements of the total flow map by combining these results together. This method was first proposed for numerical results by Brunton and Rowley (2010). They proposed that the complete flow map could be calculated from a compilation of small time step flow maps as in,

$$\Phi_{t_0}^{t_0+T} = \mathcal{I} \Phi_{t_0+(k-1)\Delta t}^{t_0+T} \cdots \mathcal{I} \Phi_{t_0+\Delta t}^{t_0+2\Delta t} \Phi_{t_0}^{t_0+\Delta t} \quad (1)$$

where $\Phi_{t_0}^{t_0+T}$ is the flow map calculated from time t_0 to time $t_0 + T$, etc., Δt is the time step between successive frames, and $T = k\Delta t$. Because these flow maps are computed on a discrete grid, interpolation, \mathcal{I} , is needed to compile the flow maps. Evaluating small time step flow maps consecutively allows for the total flow map to be computed. As

the tracking results are obtained on an unstructured grid, they must be interpolated onto the points coincident with the compiled flow map at a given time step. To this end, a thin-plate smoothing spline, which provides a high-accuracy interpolation while reducing the effects of noise in the tracking results, was used (Karri et al. 2009). Particle tracking results are inherently noisier than their PIV counterparts, and while no smoothing was applied to the tracking field directly, the smoothing spline helped to suppress any noise in the field. For a given location in the flow map, 50 closest tracks centered about the given location were used in the interpolation. This number was chosen in an effort to balance the competing effects of having sufficient information to produce a proper interpolation and not over-smoothing the data due to the inclusion of distant points in the interpolation. While the results presented here use a rectilinear flow map, this procedure is still valid if the initial flow map grid was on an unstructured mesh (Lekien and Ross 2010). A flow chart illustrating the traditional procedure along with the newly proposed method for calculating FTLE fields is shown in Fig. 1.

2.1.2 Flow map computation from particle pathline calculation

Particle pathlines were constructed from the instantaneous PTV results by linking particle locations from the final position in time step t to initial particle positions in time $t + \Delta t$. Only complete paths over the entire time of interest, excluding any broken or short-time particle paths, were used for computation of the FTLE field. Particle pairing can fail, thus terminating a trajectory, due to a number of factors including, but not limited to, out-of-plane motion and incorrect pairing with the previous step of the tracking procedure. This process produces a flow map that is regularly sampled in time, but irregularly sampled in space. In order to perform the FTLE computation, the flow maps were first interpolated onto a rectilinear grid. Again, a thin-plate spline using the 50 closest particles was used to perform the interpolation. After interpolation, the FTLE was computed using the method described in the “FTLE computation” section below.

2.2 Flow map computation using a velocity field integration method

To perform the VFI, the software package “Newman” was used (Du Toit 2010). Newman determines the flow map by numerically integrating a grid of massless tracer particles evenly spaced through the field. In cases where particles

may leave the velocity field domain, their trajectory is terminated. For the synthetic case, the flow field is such that flow tracers never leave the domain, and for the experimental data, the grid of massless particles was inset into the velocity field to help minimize this problem. To perform this integration, a Runge–Kutta 4–5 with an adaptive time step algorithm with tunable absolute and relative initial tolerances was used. For the synthetic data results, the absolute and relative tolerances were set to 10^{-12} and 10^{-14} , respectively, with an initial time step set to 10^{-7} . For the experimental vortex ring data, the absolute and relative tolerances were set to 10^{-10} and 10^{-9} , respectively, with an initial time step of 10^{-3} . After this numerical integration, the FTLE is computed using the method described below.

3 FTLE computation

Computation of FTLEs begins first by computing the right Cauchy–Green deformation tensor, C_{jk} , which is given by,

$$C = (\nabla \Phi_{t_0}^{t_0+T})^* \cdot \nabla \Phi_{t_0}^{t_0+T} \quad (2)$$

where $*$ denotes transpose and $\Phi_{t_0}^{t_0+T}$ is the flow map generated by either particle integration or particle tracking. Next, the largest eigenvalue, λ_{\max} , from the deformation tensor is identified and used to yield the forward FTLE field σ at each point in the flow field at time t_0 as in

$$\sigma_{t_0}^{t_0+T} = \frac{1}{|T|} \ln \left(\sqrt{\lambda_{\max}(C)} \right) \quad (3)$$

4 Synthetic data

In order to test this method on a field with a known solution, the double-gyre flow field was selected (Solomon and Gollub 1988a, b). This flow field is a standard test case in the FTLE/LCS literature (Shadden et al. 2005; Brunton and Rowley 2010; Lekien and Ross 2010; Tallapragada and Ross 2013) as well as in the fluids community as it is a close approximation to the flow field found in 2D Rayleigh–Benard convection (Solomon and Gollub 1988a, b). This field can be most easily thought of as a pair of counterrotating vortices confined on all sides as shown in Fig. 2. The time dependence of this flow is controlled by the asymmetric expansion and contraction of the vortices inside the domain. The flow field is described by the stream function:

$$\Psi(t) = A \sin(\pi f(t)) \sin(\pi y) \quad (4)$$

where

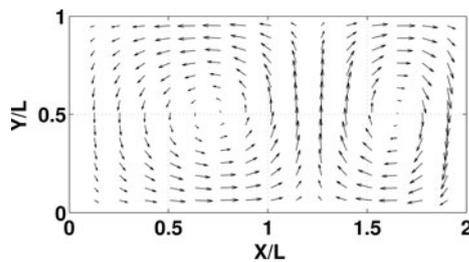


Fig. 2 Vector field snapshot of the double-gyre flow field

$$f(t) = a(t)x^2 + b(t)x \quad (5)$$

$$a(t) = \varepsilon \sin(\omega t) \quad (6)$$

$$b(t) = 1 - 2\varepsilon \sin(\omega t) \quad (7)$$

This formulation is valid over the domain of x from 0 to 2 and y from 0 to 1. The velocity field is given by

$$u = -\pi A \sin(\pi f) \cos(\pi y) \quad (8)$$

$$v = \pi A \cos(\pi f) \sin(\pi y) \frac{df}{dx} \quad (9)$$

where A is the scaling parameter for the magnitude of the velocity and ε determines how far the vortices grow and shrink. The original formulation of this comes from Solomon and Gollub (1988a, b), while this specific derivation was performed by Shadden et al. (2005). It is important to note that this solution does not satisfy the Navier–Stokes equations but rather is an approximation of a time-varying flow field with fixed boundaries. If $\varepsilon = 0$, then the flow becomes time independent. When $\varepsilon \neq 0$, then ε provides a measure of how much the separation line between the gyres moves in the x -direction over a period of $2\pi/\omega$. For the work herein, ε was set to 0.25, ω was set to $\pi/5$, and A was set to 0.1. For results pertaining to these data, the spatial domain is normalized using the characteristic length, L , of the domain $x, y = 1$.

4.1 Artificial image generation

To generate images simulating an experimental environment, the benchmark velocities must be integrated at finite locations so that these locations can be used to simulate particles as would typically be seen in PIV/PTV experiments. To perform this integration, 30,000 random tracers were seeded throughout the domain and their trajectories were integrated using MATLAB's ode45 solver, which is a Runge–Kutta-based numerical differential equation solver that is fourth-order accurate with a fifth-order check. The equations were integrated using a time step that would allow for 500 evenly spaced realizations over one complete period of the flow, $2\pi/\omega$. Because the particle locations were randomly distributed in space, the gradient calculations used radial basis functions that allow high-accuracy

derivative estimation on unevenly sampled and unstructured domains (Karri et al. 2009).

From this larger set of particles, randomly sampled subsets were extracted to generate images with specified seeding densities. Each seeding density was simulated ten times (with a different subset of particles) in an effort to produce statistically significant results. The seeding densities chosen for the simulations were 500, 1,000, 2,500, 5,000, 7,500, 10,000, and 20,000 particles per image. The image size was 1,024 by 512 pixels, which produced seeding concentrations of 0.001–0.038 particles per pixel for the 500 and 20,000 particle cases, respectively. For comparison, in a typical PIV experiment there are approximately ten particles in a 32 by 32 pixel window corresponding to a seeding density of 0.01 particles per pixel (Raffel et al. 1998; Adrian and Westerweel 2011).

The image generation was performed using the Gaussian approximation for light scatter from a particle illuminated from a thin sheet of laser light (Brady et al. 2009). The laser light distribution was also assumed to be Gaussian with the highest intensities at the center of the light sheet. As in an actual experiment, the particles were assumed to be nonuniform in size, with diameters assumed to continuously vary from 2 to 15 pixels.

4.2 PIV and PTV procedure

The PIV processing was performed using an in-house-developed code¹ employing the robust phase correlation technique (Eckstein and Vlachos 2009a, b) in conjunction with the multigrid discrete window offset method (Scarano and Riethmuller 1999) and image deformation (Scarano 2002). The initial window size was determined using the $1/4$ rule, which suggests that the window size should be four times larger than the maximum particle displacement within. The second effective window resolution was 16 by 16 pixels with a 4 by 4 grid resolution, creating 75 % overlap between interrogation windows and producing a constant total of 8,001 vectors for all of the different particle seeding cases. Two iterations were performed at each window size for a total of four passes. To improve the window deformation procedure, outliers in the intermediate results were removed using the universal outlier detection (UOD) (Westerweel and Scarano 2005) and then smoothing was performed. For the final pass only, the UOD was applied.

In order to increase the accuracy and robustness of the PTV results, PIV information was used to predict the most likely location of particles in future frames. A weighting of 75 % was given to the PIV estimation while 25 % was based on the previous tracking information as this code can

¹ PRANA software is freely distributed as open source <http://sourceforge.net/projects/qi-tools/>.

use previous tracks to better estimate a particle's future position. A two-pass median validation was also used to remove erroneous tracks (Duncan et al. 2010). After these tracks were removed, the new estimate on the next particle position was determined by a weighted average of its surrounding neighbors. Once these new estimates were established, the tracking algorithm was rerun to find more accurate matches. Lagrangian particle tracks were acquired by linking the instantaneous particles tracked in successive frames using the termination locations from the first frame as the starting location of particles in the second frame.

5 Experimental data

In order to test this method on experimental flow fields, data from a laminar vortex ring in a semi-infinite domain were chosen (Stewart 2012). These data provided a good test bed for the present work as previous studies have predicted the geometry of the FTLE/LCS from vortex rings (Shadden et al. 2006; Olcay et al. 2010). The experimental setup consisted of a piston-cylinder arrangement with the fluid being ejected into a semi-infinite domain. The stroke-to-diameter ratio (L/D) for the piston-cylinder was 1.2, which produced a circulation-based Reynolds number of approximately 1,800.

To measure the fluid motion, neutrally buoyant hollow glass spheres with average diameters of 85 microns were added to the fluid. These particles were illuminated by a Nd:YAG dual head laser (New Wave Pegasus) firing at a constant rate of 500 Hz. The images were captured using an IDT XS-3 CMOS camera with a resolution of 1,280 by 1,024 pixels and a magnification of 63 microns per pixel, providing a resolution of 2.9 diameters in the vertical direction and 6.1 diameter along the length. The spot diameter of the particles in the images was between 2 and 4 pixels. For more information on the experimental setup, see the work of Stewart (2012).

The PTV data were computed using the multiparametric particle tracking method (Cardwell et al. 2010). Particles were identified using a dynamic thresholding method with a lower pixel intensity limit of 50 (these were 8-bit images with a maximum intensity value of 255). After identification, the particles were sized using a least-squares Gaussian method (Brady et al. 2009) with the diameter assigned at 4 standard deviations. Hybrid PIV-PTV tracking was used to increase the efficiency of the PTV algorithm. The multiparametric tracking algorithm used a combination of particle size, intensity, and estimated position to determine the optimal particle match. Position estimation was performed using both previous track information and PIV results. The PIV was processed using the same in-house code utilizing robust phase correlation (RPC) (Eckstein and Vlachos

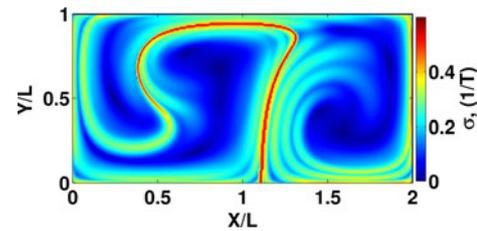


Fig. 3 Resulting FTLE field from using the benchmark flow field. This field will be used as the basis of comparison for the synthetic data analysis

2009a, b) as the synthetic images (Eckstein and Vlachos 2009a, b). Due to the high shear rate produced by the vortex, window deformation was also performed (Scarano 2002). An initial PIV window size of 32 by 32 pixels was selected. This window size was held constant for three iterations to allow the window deformation to converge (Scarano 2002). A final window size of 16 by 16 pixels was used and again allowed to converge over three iterations, resulting in a total of six passes. Because the deformation method can be very sensitive to outliers in the intermediate results, outliers were identified using the UOD method and smoothed, while only outlier detection was applied to the final output field (Westerweel and Scarano 2005). Once the data were fully processed, the PIV results were filtered using proper orthogonal decomposition (POD) (Sirovich 1987) preserving 95 % of the fluctuating energy. This step helped to reduce any noise that may have been present in the data.

6 Results and discussion

6.1 Synthetic data

Figure 3 shows the benchmark FTLE field computed at $t_0 = 0$ over one full period, $T = 2\pi/\omega$, of the double gyre. We consider this FTLE field the benchmark against which our methods should be compared. The dominant feature of this field is the ridge located in the center extending to the upper left. This ridge is the primary separatrix and will serve as a criterion to determine the effectiveness of the different methods to accurately resolve the FTLE field.

For brevity, only the 500, 5,000, and 20,000 particle cases will be shown in the following example figures, although all cases are included in the analysis. Figure 4 provides a qualitative comparison between three FTLE fields calculated using different methods from a set of 500 particles. Figure 4a was generated using VFI, Fig. 4b using PFM, and Fig. 4c using FMC. For the VFI method, massless tracer particles were placed at every pixel, which produced a grid eight times finer than the PIV field. For the

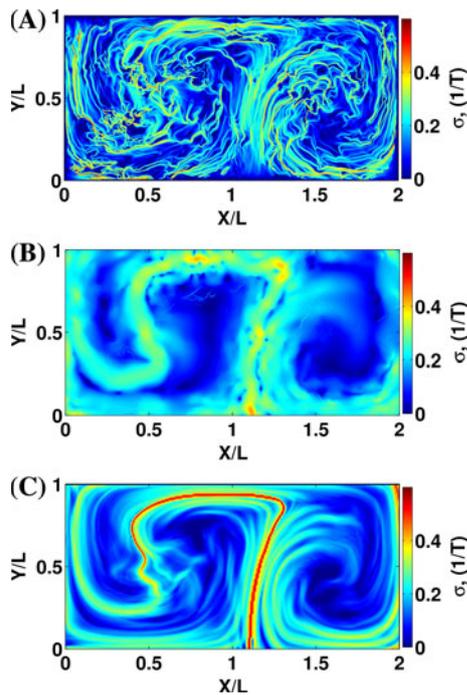


Fig. 4 FTLE results for the 500 particle case where **a**, **b**, and **c** represent velocity field integration of the PIV data, direct pathline flow map calculation, and the particle tracking flow map compilation method, respectively

FMC method, the massless particles were placed at every fourth pixel, which when using the 50 closest tracks produced an average window size of 95, 29, and 19 pixels for the 500, 5,000, and 20,000 particle cases, respectively. These correspond to the average window sizes relative to the characteristic length, L , of 19 %, 5.7 %, and 3.7 %, for the double-gyre flow field. For this seeding density, the PIV results contain a large amount of noise due to the small number of particles in each interrogation region, and as a result, the VFI also contains a large amount of error. While increasing the window size would improve the PIV results, the reduction in resolution would still yield compromised results (Olcay et al. 2010). Finite-time Lyapunov exponent fields can handle high error in the velocity field over a limited duration (Haller 2002), but long-duration errors can strongly distort the field (Olcay et al. 2010; BozorgMagham et al. 2013). For the PFM, there are a small number of pathlines due to the low seeding density. As a result, the FTLE field appears to be under-resolved, leading to a diffused and difficult-to-discern ridge. Unlike the VFI at this seeding density, the PFM resolves the separatrix that is present in the benchmark solution, although the ridge appears to have a reduced peak value and is broader than in the benchmark. For the FMC, even with this low seeding density, the method returns a field that is sufficiently similar to that of the solution. Some discrepancies can be seen which again are most likely due to the limited number

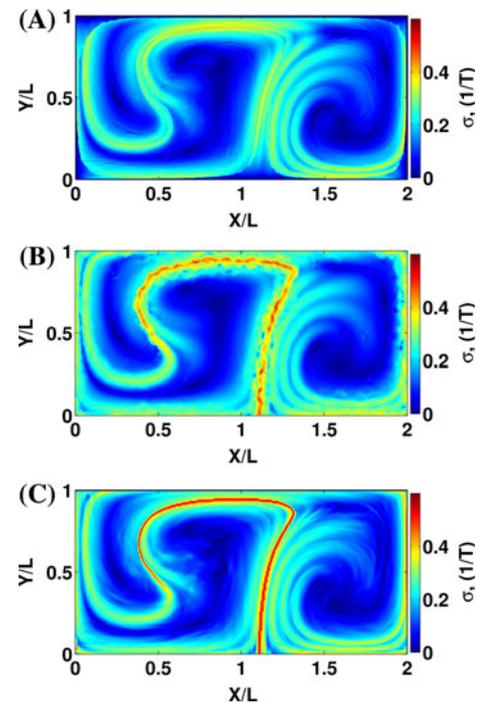


Fig. 5 FTLE results for the 5,000 particle case where **a**, **b**, and **c** represent velocity field integration of the PIV data, direct pathline flow map calculation, and the particle tracking flow map compilation method, respectively

of tracks used during the flow map interpolation. This field is able to capture not only the main separatrix but also some of the smaller ridges seen in the FTLE field.

Figure 5 shows results from the case with a seeding density of 5,000 particles where Fig. 5a, b, c corresponds to the VFI, PFM, and FMC methods, respectively. With an increased number of particles, the VFI does an improved job at resolving the field. The method, however, does struggle to capture the attachment of the separatrix to the lower wall. This inability stems from the limitation of resolving the velocity near the boundaries of the domain and the poor performance of the VFI method at these near-wall locations. While increasing the domain of the double gyre would separate these important trajectories for the near-wall region, it would not change the fact that PIV coupled with VFI will struggle to resolve near-wall behavior. Because the measured velocity from PIV is, in a sense, an average of all the particle velocities in the interrogation region, it is difficult to properly measure the velocity near a wall or near a location with zero velocity (Kahler et al. 2012). For the PFM, the field also shows improved estimation of the benchmark solution. Again, error in the flow map due to the discrete locations of the tracks appears as spurious ridges in the FTLE field. For the FMC method, the error in the tail of the separatrix appears to be gone, and the noise that was seen in the previous case appears to be diminished.

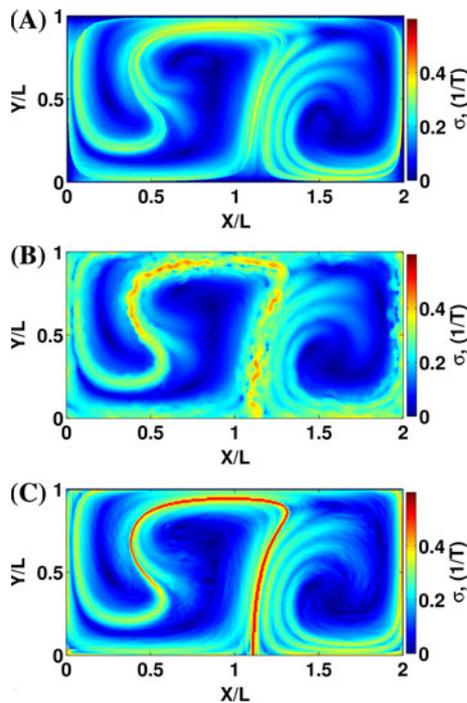


Fig. 6 FTLE results for the 20,000 particle case where **a**, **b**, and **c** represent velocity field integration of the PIV data, direct pathline flow map calculation, and the particle tracking flow map compilation method, respectively

Finally, Fig. 6 shows the results for the seeding density case of 20,000 particles where again Fig. 6a, b, c corresponds to the VFI, PFM, and FMC methods. This case is noteworthy as it is most similar to the seeding density that would be found in an ideally seeded PIV/PTV fluid experiment. Again, for the VFI it can be seen that the secondary structures appear to be well resolved (the multiple fingerlike structures in the right side of the image) while the main separatrix is still not fully captured. This shows that even with highly seeded data, VFI results cannot adequately capture the near-wall features due to the limitations of PIV measurements in this region. While decreasing the PIV window size would help improve the PIV resolution near the wall, PIV will inherently still struggle to resolve the near-wall velocities (Kahler et al. 2012). Assuming that adequate seeding density is achieved, particle tracking does not have this problem and performs better at resolving these velocities (Kahler et al. 2012). For consistency between cases, the window size was held constant, but it can be assumed that as the resolution is increased, the VFI field should approach the benchmark field. For the PFM case, again the separatrix appears to be captured, but the noise is still significantly affecting the resolution of the FTLE features. However, the FMC result shows good agreement with the benchmark solution. The separatrix appears to be well resolved with little noise

present in the field. However, it is worth noting that in Figs. 5c and 6c thin hairlike structures, which are spurious random ridges in the FTLE field, appear, which are not present in the benchmark solution.

For a quantitative analysis of these results, the normalized average error and normalized root-mean-squared (RMS) error were calculated using the following formulas:

$$\text{Normalized average error} = \frac{1}{N} \sum_{i=1}^N \left| \frac{\sigma(i)_{\text{Measured}} - \sigma(i)_{\text{Solution}}}{\sigma(i)_{\text{Solution}}} \right| \quad (10)$$

$$\text{Normalized RMS error} = \sqrt{\frac{1}{N} \sum_{i=1}^N \left(\frac{\sigma(i)_{\text{Measured}} - \sigma(i)_{\text{Solution}}}{\sigma(i)_{\text{Solution}}} \right)^2} \quad (11)$$

where σ is calculated from t_0 to $t_0 + T$ for all of the cases using the solution σ as a benchmark for comparison and the index i labels the pixel locations for which the FTLE is calculated. The normalized average error, Eq. (10), weights each point equally, while the normalized RMS error, Eq. (11), places increased emphasis on points with larger deviations from the solution. Both of these equations are important as the average error Eq. (10) provides an error metric for the entire field while the RMS error Eq. (11) will place more emphasis on the high values, or the location of what is most likely the separatrix. These values were computed for each of the ten cases at each seeding density. Figure 7 shows the average of all of these cases for each seeding density with the uncertainty bars showing the one standard deviation value computed from the ten cases. For the average error, Fig. 7a, it can be seen that the error decreases as the seeding density is increased. The VFI consistently contains higher error than the other methods reaching a minimum value of 0.136 for the 20,000 particle case. Direct particle pathline flow map calculation reaches a minimum average error of 0.0733 at the 10,000 particle case. There is an increase in the error for the PFM method in the 20,000 particle case to a value of 0.1085. This increase is due to the fact that as the seeding density increases, and the interparticle distance decreases, there is an increased likelihood that particle mismatch could occur which would introduce error into the flow map and therefore the FTLE field. For the FMC method, the error continues to decrease as the seeding density is increased, reaching a minimum at the 20,000 particle case of 0.0579, which is the lowest value compared to any of the other cases. The variance between the cases at a given seeding density also diminishes as the seeding density increases. For the 5,000 particle case and above, the variance is less than 2.5 % of the error on average.

For the normalized RMS error, as shown in Fig. 7b, the error in the FTLE field for the different cases again

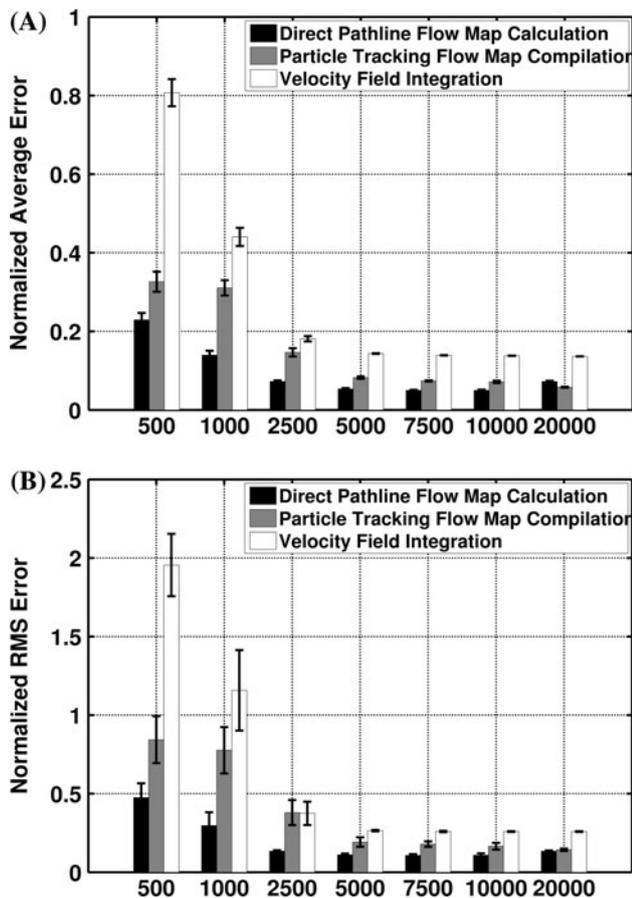


Fig. 7 Normalized average (a) and root-mean-squared (RMS) error for all three of the methods tested

decreases and levels off as the particle number reaches the 5,000 particle case. Again for the highest seeding density, the FMC method has the smallest RMS error of 0.1432. The VFI consistently has the highest RMS error with very large values at the lower seeding densities. The PFM method does well throughout the seeding densities until the 20,000 particle case where the error again rises as seen in the normalized average error. It can be understood from these results that the FMC proves, under these imaging conditions, to be a superior method when compared to VFI and comparable to PFM.

The FTLE difference calculations shown above may lead one to infer that the PFM method is almost always delivering the best results, however, which would be misleading since Figs. 4, 5, and 6 clearly show that it fails to properly resolve the FTLE structure and is only producing comparable results in a spatially averaged sense, where one-point spatial FTLE is weighted equally with the next. Therefore, in addition to investigating the total FTLE field errors, we also investigated the ability of the given methods to properly determine the separatrix seen in the benchmark solution. To this end, the amount of overlap between the

benchmark separatrix and the calculation method was computed. This analysis is similar to that completed by Olcay et al. (2010) with their analysis of the sensitivity of LCS identification on the flow field resolution. Figure 8 shows the benchmark separatrix (solid gray line) along with the separatrix from the three methods (colored lines) for the seeding density of 500 particles. For this study, the separatrix was defined as the region above a threshold value of 0.40, or about 75 % of the maximum value, which contained all of the primary ridge information. As these fields should produce the same results, a constant threshold was used. Here, the FMC method clearly shows the best agreement with the benchmark solution. While the PFM produces FTLE values close to that of the benchmark and thus performs well in a statistical sense when we calculate overall errors, here it can be seen that the separatrix is not well resolved and little overlap occurs.

Figure 9 shows the overlap for the seeding density of 5,000 particles. Again, because of the poor near-wall interpolation and velocity resolution, the separatrix is not clearly defined for the VFI. It has been shown that PTV methods have better near-wall performance (Kahler et al. 2012), which contributes to these methods performing

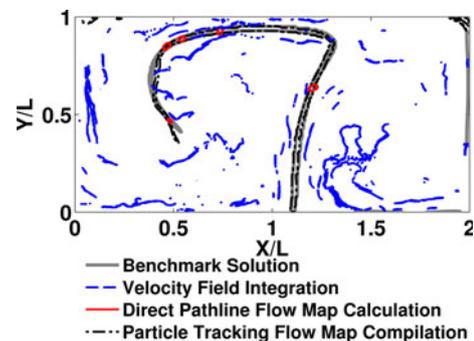


Fig. 8 Separatrix analysis for the case of 500 particles. The *solid gray line* represents the separatrix from the benchmark solution

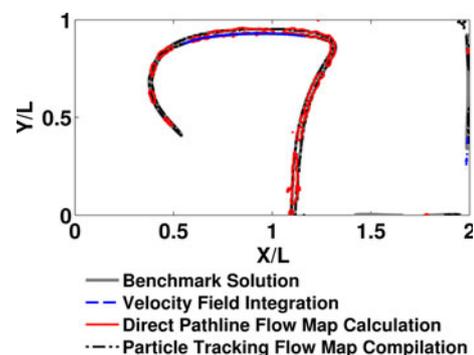


Fig. 9 Separatrix analysis for the case of 5,000 particles. The *solid gray line* represents the separatrix from the benchmark solution

better than the VFI method. For the PFM method, the separatrix appears noisy and not well resolved. The FMC method still shows strong overlap with the benchmark solution throughout the length of the separatrix. For the final 20,000 particle case, Fig. 10, the results have a similar appearance. The increased noise in the PFM can be seen in the noisy separatrix values while the FMC maintains a good correlation with the benchmark.

Figure 11 shows the overlap percentages for all of the cases and methods. Overlap was calculated by determining the number of FTLE pixel values above the given threshold that occupied the same position as those in the benchmark solution. It can be seen from this figure that the FMC consistently does a better job at resolving the separatrix than the other methods. The VFI is not able to capture the separatrix due to the sharp near-wall flow gradient and near-wall integration. The PFM performs better than VFI method but underperforms when compared to the FMC

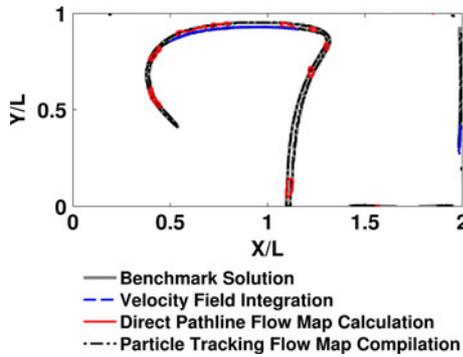


Fig. 10 Separatrix analysis for the case of 20,000 particles. The solid line represents the separatrix from the benchmark solution

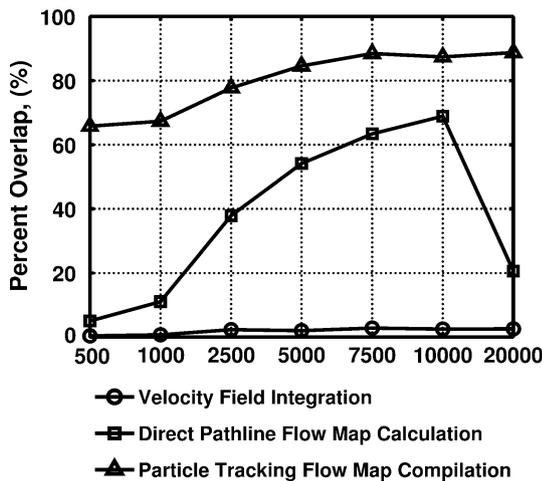


Fig. 11 Percentage overlap between the measure and true separatrix for the different methods

method. Also, the increased error that is present in the PFM when increasing the seeding density from 10,000 to 20,000 leads to decreased overlap for the highest seeding density case.

Another method for measuring the performance of these methods is to use the Hausdorff distance (Dubuisson and Jain 1994). Hausdorff distances have been used in the imaging community for pattern recognition (Zhao et al. 2005). The Hausdorff distance computes the minimum distance between all of the points in two different sets and then returns the largest of these minimum distances as given by

$$d_H(X, Y) = \max \left\{ \sup_{x \in X} \inf_{y \in Y} d(x, y), \sup_{y \in Y} \inf_{x \in X} d(x, y) \right\} \quad (12)$$

where sup is the supremum and inf is the infimum. This metric helps to describe quality of a match between the two sets; the closer to zero, the better. For the purpose of this study, each set will be composed of the largest continuous section of the FTLE above the threshold used in the previous overlap analysis and will be compared to the benchmark solution location above the same threshold. The use of only the largest continuous section is appropriate since, without a priori knowledge, this would be chosen as the dominant feature of the field. The results are normalized using the characteristic length, L . The result of this comparison is shown in Fig. 12. It can be seen from this figure that the FMC method significantly outperforms the other two methods. While qualitatively it could be seen that the FMC was producing a strong match with the solution, the Hausdorff distance allows for a more quantitative comparison. Again, it can be seen from this figure that the PFM improves as the seeding density is increased but then digresses as the seeding density is increased to 20,000. The

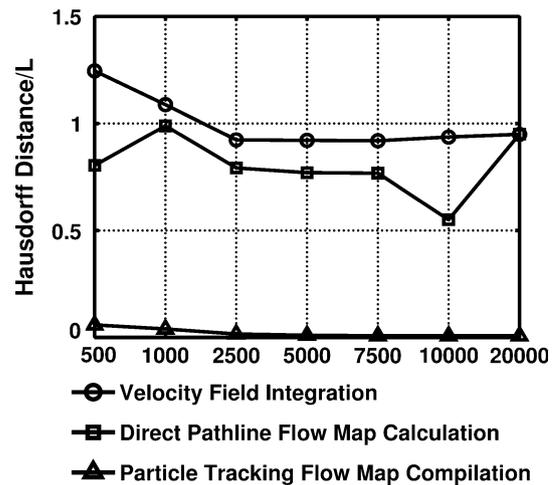


Fig. 12 Hausdorff distance for all three methods compared with the solution

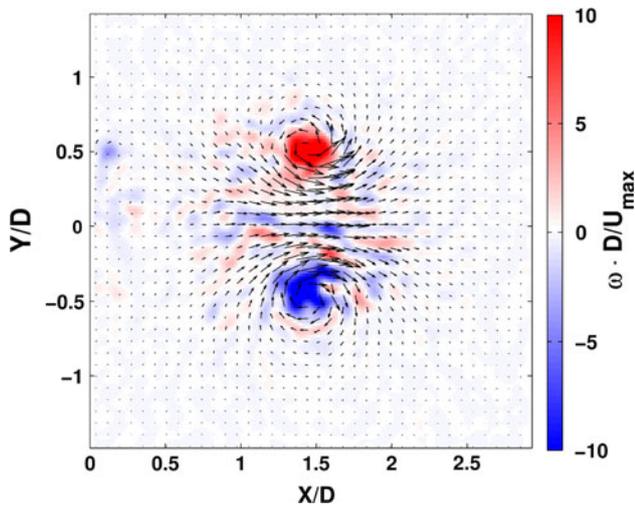


Fig. 13 Velocity field with the corresponding vorticity field magnitude calculated from PIV data for the experimental laminar vortex ring

FMC shows continuous improvement as the seeding density increases, reaching its smallest value of 0.0082 at a seeding density of 20,000.

6.2 Experimental data

Figure 13 shows the velocity field along with the nondimensional vorticity for the laminar vortex ring experimental case. As typically seen with a vortex ring, the field has maximum vorticity in the two counterrotating cores of the vortex ring cross section. The seeding density was measured from the particles used during the matching of the particle tracking; it was 0.0172 particles per pixel, which is similar to the 0.0191 particles per pixels seen in the 10,000 particle case for the synthetic data. Some noise in the velocity field can be seen in the wavy pattern present in the vorticity magnitude inside the vortex ring. Noise in the velocity field can become amplified when derivatives are calculated (Etebari and Vlachos 2005); thus, the vorticity field may show noise that is not immediately apparent in the velocity field.

Calculating the forward-time FTLE fields for all three methods yields the results shown in Fig. 14. These results show the FTLE field calculated using 75 frames. The VFI method produces a strong separatrix around the vortex core along with smaller noisier features in the far field. The shape of the separatrix is to be expected as it represents the boundary between fluid parcels that entrained into the vortex and those that are simply swept around the vortex. The PFM method, Fig. 14b, fails to capture the separatrix around the vortex ring. In experimental data, particles can be lost during the particle tracking process for a number of

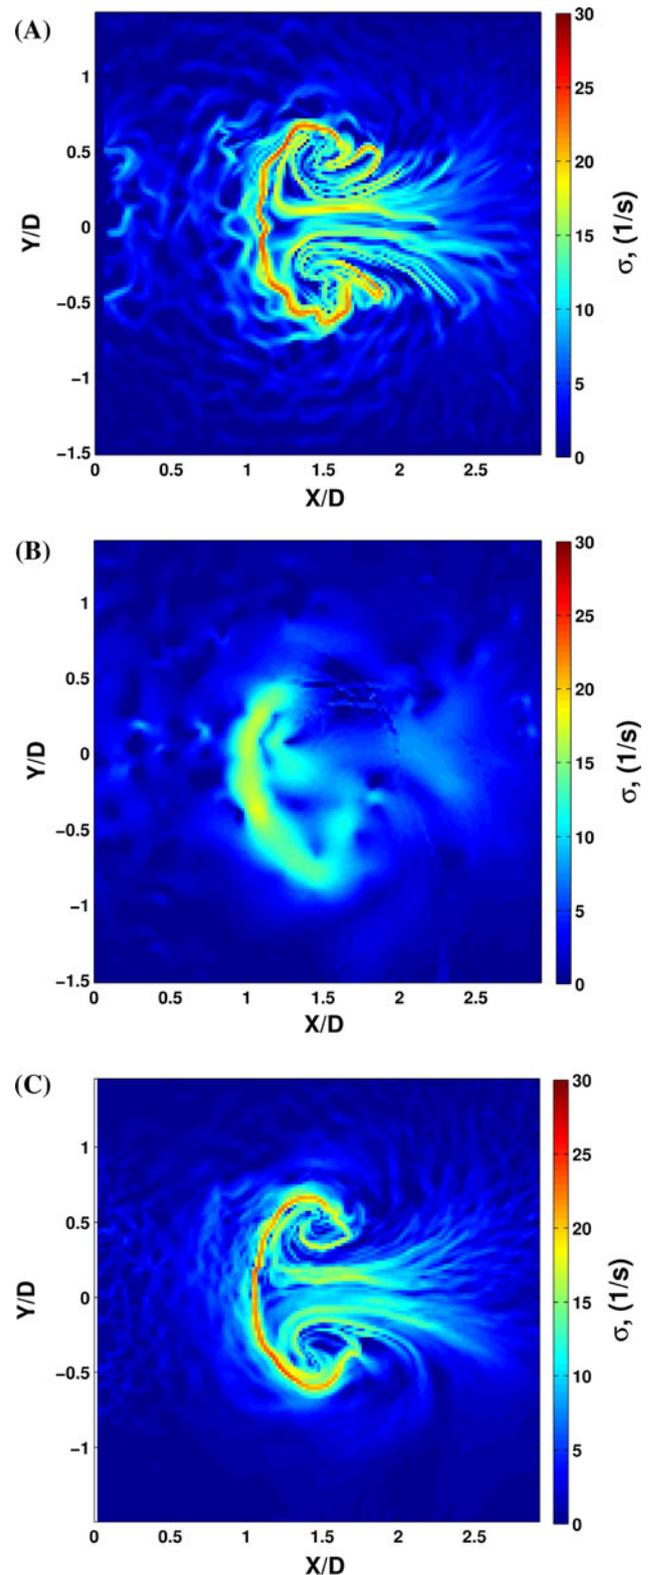


Fig. 14 Results for the experimental data from a laminar vortex ring. **a**, **b**, and **c** show the FTLE results for the velocity field integration of the PIV data, direct pathline flow map calculation, and the particle tracking flow map compilation method, respectively

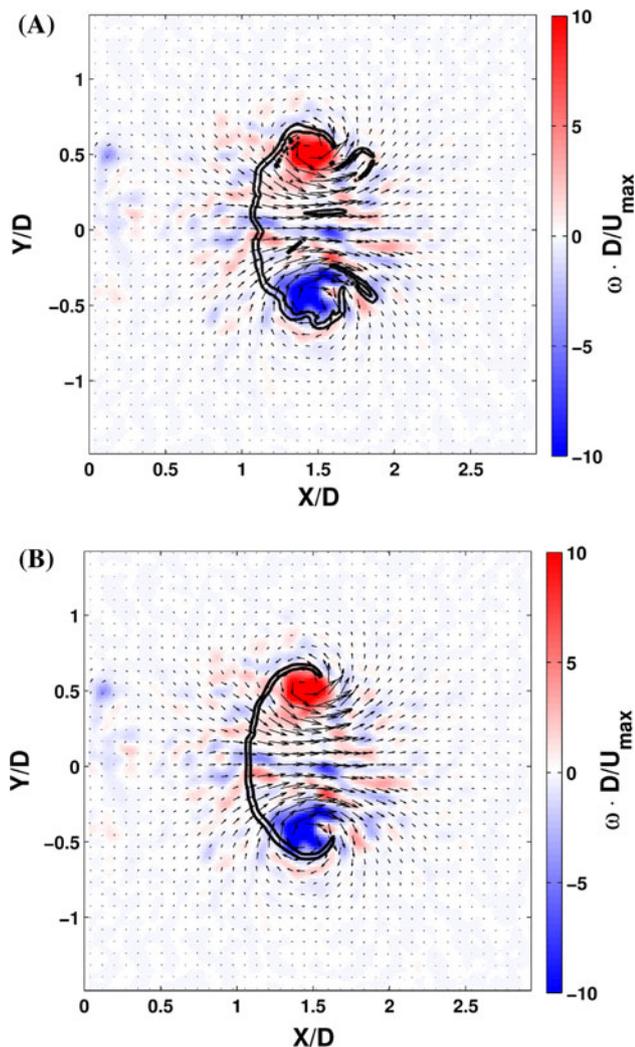


Fig. 15 Separatrix overlaid with the nondimensional vorticity field for the velocity field integration (a) and the particle tracking flow map compilation (b)

reasons, with the most notable being the out-of-plane motion. This loss of particles reduces the number of complete particle pathlines that can be used to calculate the FTLE and therefore produces a field with poor resolution of the separatrix. Figure 14c shows the results from the FMC method. This field also shows a strong separatrix, similar to the VFI method, but with a smoother shape and less noise in the far field. By comparison, this method produces an FTLE field around the vortex ring that most strongly resembles the fields that are seen in the literature (Shadden et al. 2006; Olcay et al. 2010).

To better illustrate the relationship between the separatrix and the vortex ring vorticity field, Fig. 15 shows the vorticity field overlaid with the separatrix from the VFI and FMC methods. Again, it can be seen that the FMC method appears less susceptible to noise in the velocity field as it

produces a smoother separatrix. While the FMC method has less “structure” in front (right side) of the vortex, this is attributed to noise and it is not actually present in the flow field.

7 Conclusions

This work presents two new methods for computing flow maps, and flow map derived quantities such as FTLE fields, based on PFM and FMC, and compares them against the traditional VFI of velocity field data. The focus is on application to noisy fluid measurements using PIV- or PTV-generated data where we take advantage of the fact that the flow map information is inherently contained in the Lagrangian motion of the particle flow tracers. Using this fundamental principle, we show that the FTLE fields can be calculated with increased accuracy directly from the measured particle trajectories.

We show that FMC produces the most accurate estimates of the FTLE field for both synthetic data and experimental data. For cases where particle loss between frames is minimal, the PFM method can produce reasonable results but can also be significantly affected by the unstructured nature of the data on which the flow maps are based upon. Prior to this study, it could have been assumed that a method capitalizing on direct measurement of the particle pathlines present in the field would produce the highest accuracy results. However, this work clearly shows that this is not the case as the FMC method outperforms the PFM method at low and high seeding densities. When comparing the ability to resolve the true separatrix of a flow field, the FMC was shown to be far superior with an average overlap percentage of 80 % as compared to ~ 25 and ~ 1 % for the PFM and VFI, respectively.

The FMC method shows advantage over the other methods, which becomes significantly apparent when the particle seeding is low. This can be particularly important for applications to environmental or biological flows where adding seed particles is not practical and investigation of Lagrangian flow structures must rely on naturally occurring flow tracers to produce the flow map. These results also demonstrate that when analyzing experimental data, it is advantageous to use the FMC for the computation of FTLE/LCS over VFI. The method not only produces strong measures of the FTLE field but also outperforms VFI of data with near-wall flow fields. FMC is also more computationally cost effective as each snapshot of the flow map only needs to be interpolated once whereas VFI can require multiple iterations. The FMC method also requires far fewer flow tracers (16 times less for the synthetic data) and thus fewer computations than the VFI procedure.

In future studies, an optimization of the interpolation radius should be undertaken. Using information such as seeding density and the spatial scales of interest in the flow, the interpolation radius used for the FMC method might be better optimized to further improve the results. Additionally, optimization of the massless flow tracers could also be undertaken. A procedure has been developed for numerical data (Lekien and Ross 2010), but it requires multiple computations of the flow map with refined grids, which can be computationally very costly. Future studies may be able to take advantage of information in the velocity snapshots such as the variation of strain rate and rate of rotation to optimize particle placement prior to the flow map computation. A procedure such as this should improve accuracy without being too computationally expensive.

Future work should also investigate these methodologies on three-dimensional data sets. While not shown here, the methods described are directly applicable to three-dimensional data. This approach should increase the efficiency of the tracking-based methods as extension to 3D will increase the number of possible matches in the tracking procedure while numerically integrating velocity fields in 3D will become significantly more computationally expensive. Finally, it should be noted that the majority of previous works have used FTLE/LCS primarily in a qualitative fashion. Improving the accuracy and robustness of these measurements would enable a more quantitative use and objective comparisons of future investigations.

Acknowledgments The authors would like to thank Roderick La Foy, Amir Bozorg Magham, Shibabrat Naik, and John Charonko for their helpful conversations and for commenting on the manuscript as well as Kelley Stewart and Cassie Niebel for providing the experimental data. SDR gratefully acknowledges partial support by the National Science Foundation under Grant No. CMMI-1150456. PPV gratefully acknowledges the partial support of NIH Grant 1R21HL106276-01A1.

References

- Adrian RJ (1991) Particle-imaging techniques for experimental fluid-mechanics. *Annu Rev Fluid Mech* 23(1):261–304
- Adrian RJ (2005) Twenty years of particle image velocimetry. *Exp Fluids* 39:159–169
- Adrian R, Westerweel J (2011) Particle image velocimetry. Cambridge University Press, Cambridge
- BozorgMagham AE, Ross SD et al (2013) Real-time prediction of atmospheric Lagrangian coherent structures based on forecast data: an application and error analysis. *Phys D* 258(1):47–60
- Brady MR, Raben SG et al (2009) Methods for digital particle image sizing (DPIS): comparisons and improvements. *Flow Meas Instrum* 20(6):207–219
- Brunton SL, Rowley CW (2010) Fast computation of finite-time Lyapunov exponent fields for unsteady flows. *Chaos* 20:1–12
- Cardwell ND, Vlachos PP et al (2010) A multi-parametric particle pairing algorithm for particle tracking in single and multiphase flows. *Meas Sci Technol* 22:105406
- Charonko J, Kumar R et al (2013) Vortices formed on the mitral valve tips aid normal left ventricular filling. *Ann Biomed Eng* 41(5):1049–1061
- Du Toit PC (2010) Transport and separatrices in time dependent flows. Ph.D., California Institute of Technology
- Dubuisson MP, Jain AK (1994) A modified Hausdorff distance for object matching. *Pattern Recognit* 1:566–568
- Duncan J, Dabiri D et al (2010) Universal outlier detection for particle image velocimetry (PIV) and particle tracking velocimetry (PTV) data. *Meas Sci Technol* 21(5):057002
- Eckstein A, Vlachos PP (2009a) Assessment of advanced windowing techniques for digital particle image velocimetry (DPIV). *Meas Sci Technol* 20(7):075402
- Eckstein A, Vlachos PP (2009b) Digital particle image velocimetry (DPIV) robust phase correlation. *Meas Sci Technol* 20(5):055401
- Etebari A, Vlachos PP (2005) Improvements on the accuracy of derivative estimation from DPIV velocity measurements. *Exp Fluids* 39(6):1040–1050
- Green MA, Rowley CW et al (2011) The unsteady three-dimensional wake produced by a trapezoidal pitching panel. *J Fluid Mech* 685:117–145
- Haller G (2001) Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Phys D* 149:248–277
- Haller G (2002) Lagrangian coherent structures from approximate velocity data. *Phys Fluids* 14(6):1851–1861
- Haller G (2011) A variational theory of hyperbolic Lagrangian coherent structures. *Phys D* 240(7):574–598
- Haller G, Yuan G (2000) Lagrangian coherent structures and mixing in two-dimensional turbulence. *Phys D* 147:352–370
- Holmes P, Lumley JL et al (1996) Turbulence, coherent structures, dynamical systems and symmetry. Cambridge University Press, Cambridge
- Kahler CJ, Scharnowski S et al (2012) On the uncertainty of digital PIV and PTV near walls. *Exp Fluids* 52(6):1641–1656
- Karri S, Charonko J et al (2009) Robust wall gradient estimation using radial basis functions and proper orthogonal decomposition (POD) for particle image velocimetry (PIV) measured fields. *Meas Sci Technol* 20(4):045401
- Lekien F, Ross SD (2010) The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds. *Chaos* 20(1):017505
- Lipinski D, Mohseni K (2010) A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures. *Chaos* 20(1):017504
- Mathur M, Haller G et al (2007) Uncovering the Lagrangian skeleton of turbulence. *Phys Rev Lett* 98:1–4
- Mikheev AV, Zubtsov VM (2008) Enhanced particle-tracking velocimetry (EPTV) with a combined two-component pair-matching algorithm. *Meas Sci Technol* 19(8):085401
- Ohmi K, Li H-Y (2000) Particle-tracking velocimetry with new algorithms. *Meas Sci Technol* 11:603–616
- Ohmi K, Panday SP (2009) Particle tracking velocimetry using the genetic algorithm. *J Vis* 12(3):217–232
- Olcay AB, Pottebaum TS et al (2010) Sensitivity of Lagrangian coherent structure identification to flow field resolution and random errors. *Chaos* 20(1):017506
- Peng J, Dabiri JO (2009) Transport of inertial particles by Lagrangian coherent structures: application to predator-prey interaction in jellyfish feeding. *J Fluid Mech* 623:75–84
- Raffel M, Willert CE et al (1998) Particle image velocimetry. Springer, Berlin
- Ruiz T, Boree J et al (2010) Finite time Lagrangian analysis of an unsteady separation induced by a near wall wake. *Phys Fluids* 22(7):075193
- Scarano F (2002) Iterative image deformation methods in PIV. *Meas Sci Technol* 13:R1–R19

- Scarano F, Riethmuller ML (1999) Iterative multigrid approach in PIV image processing with discrete window offset. *Exp Fluids* 26(6):513–523
- Senatore C, Ross SD (2011) Detection and characterization of transport barriers in complex flows via ridge extraction of the finite time Lyapunov exponent field. *Int J Numer Methods Eng* 86:116301174
- Shadden SC (2011) Lagrangian coherent structures. In: Grigoriev R (ed) *Transport and mixing in laminar flows: from microfluidics to oceanic currents*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim
- Shadden SC, Lekien F et al (2005) Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Phys D* 212:271–304
- Shadden SC, Dabiri JO et al (2006) Lagrangian analysis of fluid transport in empirical vortex ring flows. *Phys Fluids* 18:1–11
- Shadden SC, Katija K et al (2007) Transport and stirring induced by vortex formation. *J Fluid Mech* 593:315–331
- Shinneeb A, Balachandar R et al (2006). Quantitative investigation of coherent structures in a free jet using PIV and POD. ASME Joint US—European Fluids Engineering Summer Meeting, Miami, FL, pp 1–8
- Sirovich L (1987) Turbulence and the dynamics of coherent structures I–III. *Q Appl Math* 45(567–571):573–590
- Solomon TH, Gollub JP (1988a) Chaotic particle-transport in time-dependent Rayleigh-Benard convection. *Phys Rev A* 38(12):6280–6286
- Solomon TH, Gollub JP (1988b) Passive transport in steady Rayleigh-Benard convection. *Phys Fluids* 31(6):1372–1379
- Stewart K, Niebel C et al (2012) The decay of confined vortex rings. *Exp Fluids* 53(1):163–171
- Tallapragada P, Ross SD (2008) Particle segregation by Stokes number for small neutrally buoyant spheres in a fluid. *Phys Rev E* 78:1–9
- Tallapragada P, Ross SD (2013) A set oriented definition of finite-time Lyapunov exponents and coherent sets. *Commun Nonlinear Sci Numer Simul* 18(5):1106–1126
- Voth GA, Haller G et al (2002) Experimental measurements of stretching fields in fluid mixing. *Phys Rev Lett* 88(25)
- Westerweel J, Scarano F (2005) Universal outlier detection for PIV data. *Exp Fluids* 39(6):1096–1100
- Wilson MM, Peng J et al (2009) Lagrangian coherent structures in low Reynolds number swimming. *J Phys Condens Matter* 21:204105
- Zhao C, Shi W et al (2005) A new Hausdorff distance for image matching. *Pattern Recognit Lett* 26(5):581–586