

The Signal in the Noise: Understanding and Mitigating Decorrelation in Particle Image Velocimetry

Matthew N. Giarra

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University in
partial fulfillment of the requirements for the degree of

Doctor of Philosophy
In
Mechanical Engineering

Pavlos P. Vlachos, Co-Chair

John J. Socha, Co-Chair

Rafael V. Davalos

Sunghwan Jung

Mark A. Stremler

January 10, 2017

Blacksburg, VA, USA

Keywords: Particle image velocimetry, cross correlation, image processing, insects

The Signal in the Noise: Understanding and Mitigating Decorrelation in Particle Image Velocimetry

Matthew N. Giarra

Abstract

Particle image velocimetry (PIV) has become one of the most important tools for experimentally investigating the physics of fluid flows. In PIV, image-processing algorithms estimate flow velocity by measuring the displacements of flow-tracer particles suspended in a fluid. The fundamental operation in PIV is the cross correlation (CC), which measures the displacement between two similar patterns. These measurements can fail under circumstances that arise due to the nature of the underlying flow field (e.g., vortices and boundary layers, where particle patterns not only translate but also rotate, stretch, and shear) or of the images (e.g., X-ray images, with comparatively low signal to noise ratios). Despite these shortcomings, fairly little attention has been paid to fundamentally improving measurements at the level of the CC. The objective of this dissertation is to demonstrate specific modifications to the correlation kernel of PIV that increase its accuracy and in certain cases extend its utility to classes of flows and image types that were previously unresolvable. First, we present a new PIV correlation algorithm called the Fourier-Mellin correlation (FMC) that reduces velocity errors by an order of magnitude in rotating flows (chapter 1). Second, we develop a model of PIV cross correlations that explains the fundamental sources of several major drivers of error in these measurements. We show how the shapes of the tracer particles and the distributions of their individual displacements affect the correlation signal to noise ratio (SNR), whose effects have previously been described only heuristically. We use this insight to create an algorithm that automatically creates a Fourier-based weighting filter, and demonstrate that our algorithm reduces bias and RMS errors in multiple types of PIV experiments (chapter 2). Finally, we apply principles from our insights to measure blood flows in the hearts of grasshoppers using X-ray PIV, and discovered flow kinematics that were unexpected according to the current prevailing understanding of the heart as a peristaltic pump that produces directional flows. Our results suggest that flow production in insect hearts may be more complex than once thought (chapter 3).

The Signal in the Noise: Understanding and Mitigating Decorrelation in Particle Image Velocimetry

Matthew N. Giarra

General Audience Abstract

Particle image velocimetry (PIV) is a tool for measuring the motion of fluid flows. In PIV, reflective particles are suspended in a flowing fluid, and cameras record their motion. Computer algorithms measure the motion of the particles in those images to estimate the velocity of the fluid. This dissertation is about the theory, algorithms, and experiments of particle image velocimetry. We explain from a theoretical standpoint the reasons that PIV can fail to provide reliable measurements for several types of flows that are commonly encountered in the research of fluid physics and engineering, such as swirling vortices or eddies, jets, turbulence, and microscopic flows. We apply this understanding to create new algorithms that improve PIV measurements in these kinds of challenging scenarios. Lastly, we use PIV and high-speed X-ray imaging to measure flow patterns within the tubular hearts of living grasshoppers. From these experiments, we discovered flow behaviors that were strikingly different from what we expected according to the current prevailing presumption that the insect heart is a peristaltic pump. If the heart is proven to function other than by peristalsis, then this could imply that a previously overlooked flow mechanism could in fact be among the most prevalent among animals.

Acknowledgements

I owe this dissertation and my degree to the family, friends, neighbors, and colleagues whose support has kept me afloat throughout the years, and especially during graduate school. My co-advisors, Pavlos Vlachos and Jake Socha, provided me with invaluable guidance (and funding) throughout my tenure, and allowed me near unlimited freedom to choose my own research path. Jake infected me with his curiosity, and taught me the importance of story-telling in science. Pavlos' discerning eye helped to train my own healthy sense of skepticism, and reinforced the important lesson that even more important than questioning results that look wrong is to question results that look right. I was lucky to have the members of my committee, Rafael Davalos, Sunghwan "Sunny" Jung, and Mark Stremler, to provide fresh perspectives and insightful feedback about my work. John Charonko's patient willingness to teach, and his genuine interest in my work, gave me the confidence to develop my ideas as a new graduate student. Don Mullins introduced me to insects from the perspective of an entomologist, and kept me supplied with grasshoppers and cockroaches at Virginia Tech. Jon Harrison enabled my experiments at Argonne National Laboratory (ANL) by mailing me boxes full of grasshoppers from Arizona State University. Kamel Fezzaa welcomed me into his laboratory at ANL and was instrumental to performing X-ray experiments using the Advanced Photon Source. Steven "Doctor" Day introduced me to the world of research when he invited me (an undergraduate student) to lend a hand with an experiment in his fluids lab. That experiment grew into my master's thesis. After serving as my advisor for that degree, Dr. Day was the first to encourage me to consider pursuing a Ph.D., and since then has spent years encouraging me to finish one. My friends helped me to stay positive during graduate school. Artur Wolek was not only a wonderful room mate in Blacksburg, but also a paragon of hard work and focus, whom I am lucky to have as a role model. Jim Cezo is one of the most creative, patient, and unpretentious guys I know, and has been a permanent source of inspiration and levity ever since we were room mates as college freshmen. Ben Jackson and I started out as office mates and became fast friends over our shared values. He taught me the courage to have ambition, but also to maintain perspective when a plan doesn't come together right away. And, of course, my parents and family, whose love, support, and encouragement have been the fundamental constants of my life.

Contents

Contents	v
Attributions.....	vii
Chapter 1: Measurement of fluid rotation, dilation, and displacement in particle image velocimetry using a Fourier–Mellin cross-correlation.....	1
Abstract.....	1
Introduction	1
Background.....	3
Fourier–Mellin Transformation	4
FMC Algorithm	6
Performance Assessment.....	9
Monte Carlo error analysis	9
Synthetic vortex ring.....	11
Experimental vortex ring.....	13
Results	15
Monte Carlo error analysis	15
Synthetic Vortex Ring.....	18
Experimental vortex ring.....	22
Conclusions	23
References.....	24
Chapter 2: Adaptive Spectral Filtering of Particle Image Velocimetry Cross Correlations.....	27
Abstract.....	27
Introduction	27
Background.....	28
Introducing the Spectral Domain	30
Modeling the correlation in the absence of noise	32
The Spectral Correlation	34
The Adaptive Phase Correlation (APC) Algorithm	43
Performance Assessment.....	47
Error Analysis: Poiseuille flow (synthetic images).....	47
Experimental Images (PIV Challenge 2014 Case A)	54
Conclusions	56

References.....	57
Chapter 3: High-Speed Flow Visualization Reveals Rapid Flow Reversals in Insect Hearts ..	60
Abstract.....	60
Introduction	60
Results and Discussion.....	62
The flow of hemolymph in the hearts of cockroaches and grasshoppers reversed directions multiple times per second.....	62
The flows we measured in the hearts of grasshoppers and cockroaches appear qualitatively to exhibit some features of non-peristaltic pumping	69
The species we studied were more than an order of magnitude larger than other animals whose cardiac flows have been observed directly.	70
The species of insects we studied were not closely related to one another.....	70
Materials and Methods.....	71
Animal Care.....	71
Fluorescent Dye injection	71
Phase contrast synchrotron X-ray Imaging (PCXI)	72
Image Pre-Processing.....	74
PIV Processing.....	75
Spatial Ensemble Cross Correlation	76
Acknowledgement.....	78
References.....	78
Chapter 4: Conclusions and Significance.....	81
Conclusions and Significance of Chapter 1: Measurement of fluid rotation, dilation, and displacement in particle image velocimetry using a Fourier–Mellin cross-correlation	81
Conclusions and Significance of Chapter 2: Adaptive Spectral Filtering of Particle Image Velocimetry Cross Correlations	82
Conclusions and Significance of Chapter 3: High–Speed Flow Visualization Reveals Rapid Flow Reversals in Insect Hearts	85
References.....	87

Attributions

The three main chapters of this dissertation consist of manuscripts that have either been published or are in preparation to be published in archival journals. The co-authors of these manuscripts are acknowledged here.

Pavlos P. Vlachos (Professor, Department of Mechanical Engineering, Purdue University)

Dr. Vlachos served as the co-chair of my committee. He contributed in the development of concepts, design of studies, interpretation of results, and writing of the work presented in each chapter. He is a co-author of each chapter in this dissertation.

Jake Socha (Associate Professor, Department of Biomedical Engineering and Mechanics, Virginia Tech)

Dr. Socha served as the co-chair of my committee. He contributed to the development of concepts, design and execution of studies, interpretation of results, and writing of the work presented in chapter 3. He is a co-author of that chapter.

John J. Charonko (Research & Development Scientist, Los Alamos National Laboratory)

Dr. Charonko contributed to the development of concepts, design of studies, interpretation of results, and writing of the work presented in chapter 1. He is a co-author of that chapter.

Chapter 1: Measurement of fluid rotation, dilation, and displacement in particle image velocimetry using a Fourier–Mellin cross-correlation

Abstract

Traditional particle image velocimetry (PIV) uses discrete Cartesian cross correlations (CCs) to estimate the displacements of groups of tracer particles within small subregions of sequentially captured images. However, these cross-correlations fail in regions with large velocity gradients or high rates of rotation. In this paper, we propose a new PIV correlation method based on the Fourier–Mellin transformation (FMT) that enables direct measurement of the rotation and dilation of particle image patterns. In previously unresolvable regions of large rotation, our algorithm significantly improves the velocity estimates compared to traditional correlations by aligning the rotated and stretched particle patterns prior to performing Cartesian correlations to estimate their displacements. Our algorithm, which we term Fourier–Mellin Correlation (FMC), reliably measures particle pattern displacement between pairs of interrogation regions with up to ± 180 degrees of angular misalignment, compared to 6–8 degrees for traditional correlations, and dilation/compression factors of 0.5–2.0, compared to 0.9–1.1 for a single iteration of traditional correlations.

We apply our FMC algorithm to synthetic computer-generated PIV images with known velocity and vorticity fields, and to an experimentally measured flow field. Our results show that combining FMC with discrete window offset (DWO) or iterative image deformation (IID) algorithms decreases the mean and variance of displacement and vorticity errors compared to traditional correlations, and that FMC accelerates the convergence of IID.

Introduction

Traditional particle image velocimetry (PIV) can fail in the presence of strong velocity spatial gradients because of the inability of Cartesian cross correlations (CCs) to capture the shearing, stretching, and rotation of particle image patterns. PIV estimates the fluid velocity from the linear translation of groups of flow tracer particles between consecutive frames. However, particle patterns not only translate, but they can also rotate, stretch (scale), and shear between frames. Because traditional PIV does not resolve these additional deformations, they act as sources of

error for the velocity estimation by decreasing the signal-to-noise ratio (SNR) of the cross correlation [1-3].

The fundamental reason that velocity gradients degrade PIV velocity estimates is that cross correlations measure rigid-body shifts, and there is no rigid-body shift that can align rotated, dilated, sheared, or otherwise deformed image pairs. Previous researchers have developed a variety of techniques to reduce the adverse effects of velocity gradients on PIV measurements [1, 3, 4]. Several effective methods operate by iteratively deforming the original pair of photographs based on the previous iteration's velocity field estimate, which reduces the apparent misalignment of the particle patterns and thereby increases the SNR of subsequent correlations. These iterative image deformation (IID) algorithms improve velocity estimates in the presence of gradients up to 1.1 pixels/pixel [5]. However, because the measured velocity fields precondition the image deformations, their accuracy affects the accuracy of the entire algorithm. As a result, iterative image deformation algorithms that use Cartesian correlations remain subject to the same limitations, and sidestep rather than address the fundamental limitations of cross correlation-based PIV methods.

Other groups have proposed to overcome the inability of traditional correlations to deal with rotation by modifying the correlation process itself. Ruan *et al.* [6] increased the accuracy of cross correlations in rotating flows by resampling the raw PIV interrogation regions (IR) onto polar (ρ, θ) coordinates prior to performing cross correlations. By this image transformation, rotation in the original (Cartesian) image appears as linear translation along the angular θ axis in the transformed (polar) image. Then, the rotation angle relating the Cartesian image pair can be easily measured by cross-correlating their corresponding polar transforms. Finally, the Cartesian images are brought into alignment by rotating them through the measured angle, and the translation relating the rotated image pair is measured by cross-correlating them. However, if the original images not only rotate but also translate and/or dilate, or if they do not rotate about their geometric centroids, then their polar transforms are no longer related purely by translations. In this case, translation or dilation in the original images corrupts cross correlations between polar images similarly to how rotation and dilation corrupt cross correlations between Cartesian images, and the method fails. Although polar correlations do improve velocity measurements for purely rotating flows, they merely substitute one problematic class of particle pattern motions for another, and are therefore not globally applicable to flows that both rotate and translate.

In this chapter, we propose a new PIV algorithm that overcomes these limitations by replacing the underlying Cartesian CC with a different correlation algorithm that precisely measures and then corrects for both the rotation and the isotropic dilation between individual pairs of IRs, and does so even in the presence of large translations. Our method drastically improves the principal measurement of particle motion, at the fundamental level of the correlation, and thereby mitigates the heavy reliance of PIV on post-processing of error-ridden velocity fields in strongly rotating flows. Because our algorithm handles combinations of translation, rotation, and dilation, it can be applied globally to flows that exhibit any or all of these deformations. Furthermore, our algorithm's intrinsic measurements of particle pattern rotation and dilation provide direct estimates of fluid vorticity and stretching within individual PIV IRs. More broadly, our work overcomes the traditional assumption in PIV that particle motion is purely translational, and is a step toward directly measuring the full deformation tensor of fluid mechanics using PIV.

Our algorithm is based on the Fourier–Mellin Invariant (FMI) image transformation, which is widely used in computer vision and pattern recognition to align images that differ in scale and orientation angle [7, 8], but has been overlooked by the PIV community. We therefore call our algorithm the Fourier–Mellin Correlation, or FMC. The following sections describe the theory and use of FMC to improve velocity estimates and to directly measure rotation and scaling in individual PIV subregions. Then, we provide quantitative error analysis of FMC compared to traditional methods using computer-generated synthetic PIV images with known ground-truth solutions. Finally, we show how FMC compares to traditional methods for experimental data using PIV photographs of a vortex ring in water. Our results will show that FMC can tolerate arbitrary angular misalignment (± 180 degrees) between IR pairs compared to ± 6 – 8 degrees for traditional correlations, and dilation/compression factors of 0.5–2.0 compared to 0.9–1.1, for single iterations of Cartesian correlations. Furthermore, substituting FMC into iterative image deformation algorithms reduces the random error of velocity estimates by an order of magnitude compared to traditional CCs, and converges in fewer than half the number of iterations for the tested flows.

Background

Traditional PIV uses cross-correlations over a Cartesian grid to estimate the motions of groups of tracer particles as rigid translations. In contrast, our FMC algorithm begins by estimating the rotation and isotropic scaling of a particle pattern using the Fourier–Mellin (FM)

image transformation, which decouples rotation from isotropic scaling and is invariant to translation [7]. The FM transformation of an image amounts to the resampling of the magnitude of its Fourier transform (which is invariant to translation) from a grid that is evenly spaced in Cartesian coordinates onto a rectilinear grid that is evenly spaced in log-polar $(\log(\rho), \theta)$ coordinates. The two perpendicular axes of the resampled grid describe the angular coordinate and the logarithm of the radial coordinate in the original data, so that rotating and isotropically scaling an image in the Cartesian spatial domain causes perpendicular translations in the FM domain. Therefore, the 2-D cross-correlation between two images in the FM domain provides a direct estimate of the rotation and the isotropic scaling that relate them in the Cartesian spatial domain.

Fourier–Mellin Transformation

The Fourier Shift Theorem (Equation (1.1)) states that shifting a 2-D signal has no effect on the magnitude (and only affects the phase) of its Fourier transform (FT) [9]:

$$\mathcal{F}\{f(x-s_x, y-s_y)\}(u, v) = e^{-i2\pi(s_x u + s_y v)} F(u, v) \quad (1.1)$$

In Equation (1.1), $F(u, v) = \mathcal{F}[f(x, y)](u, v)$ is the Fourier transform of $f(x, y)$; s_x and s_y are the shifts along the perpendicular spatial axes x and y (e.g., in pixels), and u and v are the wave numbers in the directions of x and y (e.g., in pixels⁻¹). Meanwhile, the Affine Theorem states that rotating, scaling, and shearing an image also transforms its FT magnitude: the FT magnitude is rotated by the same angle, scaled by the reciprocal of the image scaling factor, and sheared in the opposite direction along perpendicular axes [9]. Combining the Affine and Shift theorems leads to the important conclusion that the rotation, scaling, and shearing that relate a pair of images can be decoupled from their corresponding translations through the magnitudes of their Fourier transforms.

The log-polar transform of the FT magnitude, known as the *Fourier–Mellin Transform* (FMT) of the original image [7], further decouples rotation from isotropic scaling by transforming orthogonal (u, v) Cartesian coordinates into orthogonal $(\log(\rho), \theta)$ log-polar coordinates [10]. If

$F'(\log(\rho),\theta)$ is the log-polar representation of the Cartesian FT magnitude $|F(u,v)|$, then rotating and scaling $|F|$ by α and m produces perpendicular shifts in F' by Equation (1.2):

$$F'(\log(m\rho),\theta+\alpha) = F'(\log(m)+\log(\rho),\theta+\alpha) \quad (1.2)$$

Because $|F|$ is invariant under translation by the Shift theorem, and because the shifts in F' are perpendicular on a rectilinear grid, cross-correlating the FMTs of two images provides a direct, translation-invariant estimate of the rotation angle and scaling factor that relate them. The row and column shifts of this correlation peak, $(\Delta r, \Delta c)$, indicate α and m according to Equation (1.3).

$$\Delta r = -\frac{\alpha N_r}{2\pi\left(1 - \frac{1}{N_r}\right)} \quad (1.3)$$

$$\Delta c = -N_c \frac{\log(m)(N_c - 1)}{\log\left(\frac{R}{r_o}\right)}$$

Here, N_r and N_c are the number of rows and columns in the resampled log-polar data, and r_o and R are the minimum and maximum radii (in pixels) of the log-polar resampling. The cross correlation between the Fourier–Mellin Transforms of images is the foundation of our new PIV algorithm.

An interesting and important consequence of the translation invariance of the FT magnitude is that Equation (1.3) holds when the particle pattern rotates and scales (dilates) about a point that is not coincident with the geometric centroid of the IR, and even when the center of rotation/scaling is located entirely outside of the IR. This is because the transformation that describes "off-center" rotation and scaling (that is, rotating and scaling about some point other than the image origin) is equivalent to a rotation and scaling about the origin followed by a rigid-body translation.

To demonstrate this equivalence, consider the fact that rotating a two-dimensional image through an angle α and scaling it by m about a point shifted from the origin by (x_c, y_c) carries an original coordinate vector \mathbf{x} to a transformed coordinate vector \mathbf{x}' according to

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} m \cos \alpha & -\sin \alpha \\ \sin \alpha & m \cos \alpha \end{bmatrix} \begin{bmatrix} x - x_c \\ y - y_c \end{bmatrix} \quad (1.4)$$

Since x_c , y_c , m , and α are all constants, Equation (1.4) can be written as

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} m \cos \alpha & -\sin \alpha \\ \sin \alpha & m \cos \alpha \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} \quad (1.5)$$

or compactly as $\mathbf{x}' = \mathbf{T}\mathbf{x} + \mathbf{t}$, where \mathbf{T} is the transformation matrix expressing combined rotation and dilation about the origin, and $\mathbf{t} = [t_x, t_y]^T$ is an equivalent translation vector whose components are given by Equation (1.6):

$$\begin{bmatrix} t_x \\ t_y \end{bmatrix} = \begin{bmatrix} -x_c m \cos \alpha + y_c \sin \alpha \\ -x_c \sin \alpha - m y_c \cos \alpha \end{bmatrix} \quad (1.6)$$

In the next section, we will show that transformations in the form of Equation (1.5) are precisely congruent with the measurement framework of our algorithm.

FMC Algorithm

The primary innovation of our new FMC-based PIV algorithm is to replace the underlying cross correlation between IRs with the cross correlation of their FMTs. The FMC algorithm (Figure 1.1) estimates the rotation, scaling, and translation that relate pairs of PIV IRs in three main steps. First, we estimate the rotation angle and scaling factor that relate two IRs by calculating the phase-correlation between their FMTs. Because the FMTs themselves are merely real-valued arrays, we effectively treat the pair of FMTs as a new pair of interrogation regions and measure their translation offset in the traditional way of conjugate-multiplying their

complex Fourier transforms. This means that we take the Fourier Transform of the (log-polar) Fourier Transform magnitude.

The sub-pixel location of the phase correlation peak is estimated in the traditional way of fitting a Gaussian function to the points surrounding it [11], and the rotation angle and scaling factor corresponding to this location are calculated from Equation (1.3). We empirically found good performance using FMTs generated from 128×128 -pixel FT magnitudes, which were themselves generated from IRs apodized by Gaussian filters and then symmetrically zero-padded to dimensions of 128×128 pixels prior to calculating their FTs [12]. For the log-polar resampling, we used cubic interpolation with angular and radial resampling rates of 256 and 64 points, respectively, with a radial resampling domain between 2 and 63 pixels measured from the zero-frequency pixel in the FT magnitude.

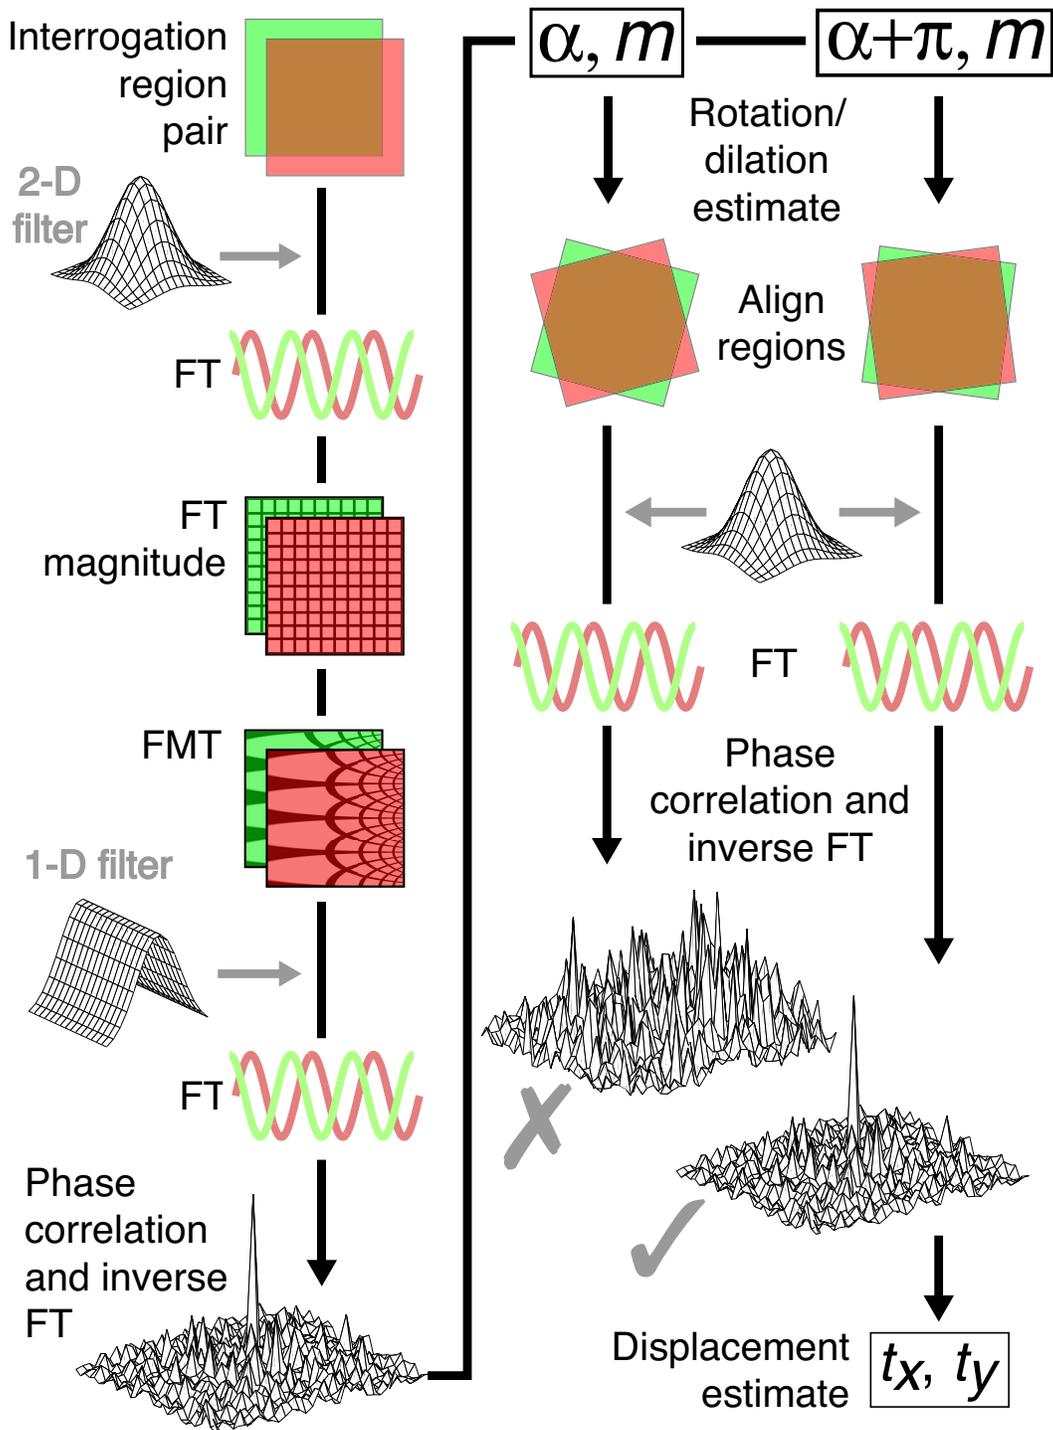


Figure 1.1: Illustration of our Fourier–Mellin Correlation (FMC) algorithm as applied to a single pair of interrogation regions (IRs). FT is Fourier transform, and FMT is Fourier–Mellin transform. α and m are the rotation angle and isotropic scaling (dilation) factor relating a pair of IRs, and t_x and t_y are the horizontal and vertical components of the particle pattern displacement.

Next, we rotate and scale each original IR to bring the pair into alignment up to a rigid-body translation. This "pre-alignment" of the IRs can be performed in the forward or backward first-order differencing sense by rotating and scaling only the first or second IR, or in the second order central-difference sense by rotating each IR by $\pm\alpha/2$ and scaling by $m^{\pm 1/2}$. We measure this remaining translation by calculating the phase correlation between the two aligned (rotated and scaled) IRs in the spatial domain [13].

Finally, the symmetry of the FFT magnitude causes an ambiguity of $\pm\pi$ in the rotation angle measured from the pair of FMTs. We resolve this ambiguity by repeating the translation measurement as described above using a rotation estimate of $\alpha + \pi$, and choosing the best estimate as the angle that results in the larger maximum value of the two spatial correlation planes. This chosen angle represents half the vorticity in units of frames⁻¹, and the corresponding displacement represents the FMC best-estimate of the velocity (in pixels per frame) of the group of particles in the IR.

By measuring first the rotation and scaling and then the rigid translation (displacement) relating IRs, our algorithm adheres to the framework of Equation (1.5). In the relevant example of a vortex center lying outside an IR, the particle pattern appears (in the IR reference frame) to rotate about a point that is not coincident with the IR's origin. This is precisely the situation described by Equation (1.4), and consequently by Equation (1.5).

The FMC algorithm is readily integrated into iterative predictor-corrector processing schemes. For this work, we implemented central-difference discrete window offset (DWO) [14-16] and iterative image deformation (IID) [5] without modification.

The computational cost of FMC is greater than for standard Fourier Transform-based correlation algorithms due to its use of additional forward and inverse FTs (eight forward FTs and three inverse for FMC compared to two forward and one inverse for traditional methods) and resampling operations (six for FMC and none for traditional correlations). In practice, this resulted in approximately a 3.5-fold increase in computation time for FMC compared to RPC using codes written in Matlab.

Performance Assessment

Monte Carlo error analysis

Monte Carlo error analysis was used to quantify the accuracy of FMC's estimates of the rotation, scaling (dilation), and displacements relating pairs of synthetic PIV IRs (64×64 and 128×28 pixels). We compared FMC to Cartesian correlations using robust phase correlation

(RPC), which is an analytically-filtered phase-only correlation algorithm that has been shown to reduce bias and random errors of PIV measurements compared to standard cross correlations [17, 18]. For FMC, each Monte Carlo simulation held one of three transformation parameters constant (zero rotation, unity scaling, or zero translation) and simultaneously varied the other two. The values of the varying parameters were drawn from pseudorandom uniform distributions with scaling $0.5 \leq m \leq 2.0$ (i.e. $2^{\pm 1}$) and rotation angle $-180 \leq \alpha \leq 180$ degrees. Displacements varied by $-8 \leq t_x, t_y \leq 8$ pixels for 64×64 -pixel regions, and $-16 \leq t_x, t_y \leq 16$ pixels for 128×128 -pixel regions (i.e., $\pm 1/4$ the effective diameter [12] of each region size after Gaussian apodization). Particle image concentration was randomly varied between 0.010 and 0.020 particles per pixel (10-20 particles per 32×32 -pixel region), and one million pairs of IRs were used for each simulation.

For RPC, the same region sizes, particle image concentrations, and numbers of region pairs were used. The displacement was always held fixed at zero pixels so that any loss of correlation was due purely to rotation and/or dilation and not due to in-plane displacement, while the range of rotation angles was reduced to ± 12 degrees and scaling factors to $1.25^{\pm 1}$.

Gaussian-shaped particle images (average diameter $d_p = \sqrt{8}$ pixels) were rendered based on the theory of Olsen and Adrian [19], and the specific details of our implementation can be found in Brady *et al.* [20]. Uncorrelated Gaussian white noise (mean value of 5% of the image white value) was added to each image to simulate camera noise.

We used our FMC algorithm to measure α , m , t_x , and t_y relating the particle patterns in each pair of IRs, and used RPC to measure t_x and t_y only. For both algorithms, a Gaussian apodization function (“window”) was applied to each IR to reduce aliasing in its Fourier Transform. The effective diameter [12] of the Gaussian window was half the IR size (a “50%” window, e.g., effective diameter 32×32 pixels for 64×64 -pixel IRs). For the FMC algorithm, each FMI-transformed image was apodized by a 50% Gaussian window in the $\log(\rho)$ direction, and was not apodized in the periodic θ direction.

We computed the average error magnitude of both algorithms’ estimates of translation, and FMC’s rotation and scaling estimates, over binned ranges of each varied transformation parameter. Fifty bins were used across the range of each parameter, so that 2,500 IR pairs contributed to the average error calculated for each bin.

Synthetic vortex ring

We generated and analyzed 1,000 pairs of synthetic full field particle images (1024×1024 pixels) of two counter-rotating vortex cores with Gaussian vorticity distributions. The ground truth velocity field $\mathbf{u}(x, y)$ was calculated as the superposition of two vortex cores separated by 200 pixels, each with velocity and vorticity distributions given by Equation (1.7) (adopted from Davenport *et al.* [21]).

$$\begin{aligned}
 u_{\theta}(r) &= \pm u_{\theta_0} \left(1 + \frac{1}{2\beta} \right) \left(\frac{r_c}{r} \right) (1 - e^{-\beta r^2 / r_c^2}), (r > 0) \\
 u_r &= u_z = 0 \\
 \omega_z(r) &= \pm \frac{2\beta u_{\theta_0}}{r_c} \left(1 + \frac{1}{2\beta} \right) e^{-\frac{\beta r^2}{r_c^2}}, (r \geq 0)
 \end{aligned} \tag{1.7}$$

Here, u_{θ} and u_r are the tangential and radial velocities (pixels/frame) in the reference frame of each core in polar coordinates, and ω_z is the out-of-plane component of the vorticity (in frames⁻¹); r is the radial distance from the center of each core (pixels); $u_{\theta_0} = 25$ pixels/frame is the maximum tangential velocity due to each core; $r_c = 100$ is the core radius (the radial distance from $r = 0$ to the location where $u_{\theta} = u_{\theta_0}$); and $\beta \approx 1.25643$ is a constant that satisfies $u_{\theta}(r = r_c) = u_{\theta_0}$. The large maximum displacement was chosen so that the rotation of particle patterns near the vortex cores was sufficient to clearly demonstrate the difference in performance between Cartesian correlations and FMC in rotating flows. Furthermore, FMC is subject to loss of correlation due to in-plane particle pattern displacement [22] in the same way that Cartesian correlations are, so large but measurable displacements should not preferentially bias our results in favor of either algorithm.

Figure 1.2 shows the ground truth velocity and vorticity fields and the radial velocity and vorticity profiles described by Equation (1.7). We define the “near-core” region as the area of the flow field where $r < r_c$ of each core, and compared the different PIV algorithms based on the absolute errors of their measurements in this region compared to the ground truth Eulerian solution.

Initial particle positions were generated randomly every ten images and then advected by forward integrating the velocity field $\mathbf{u}(x,y)$ with ode45 in Matlab. Particle images were rendered with the same method as for Monte Carlo simulations, with 5% uncorrelated Gaussian noise added to each image. We specified a particle image density of 0.025 particles per pixel (or about six particles per 16×16 -pixel area), which, based on the work of Keane and Adrian [22], should correspond to about a 95% probability of detecting valid peaks in the spatial correlation plane (i.e. the translation estimate) at the smallest interrogation regions we used. We also tested our algorithm using images with up to 0.100 particles per pixel (or 25.6 particles per 16×16 -pixel area) to elucidate any unexpected behavior at higher seeding densities.

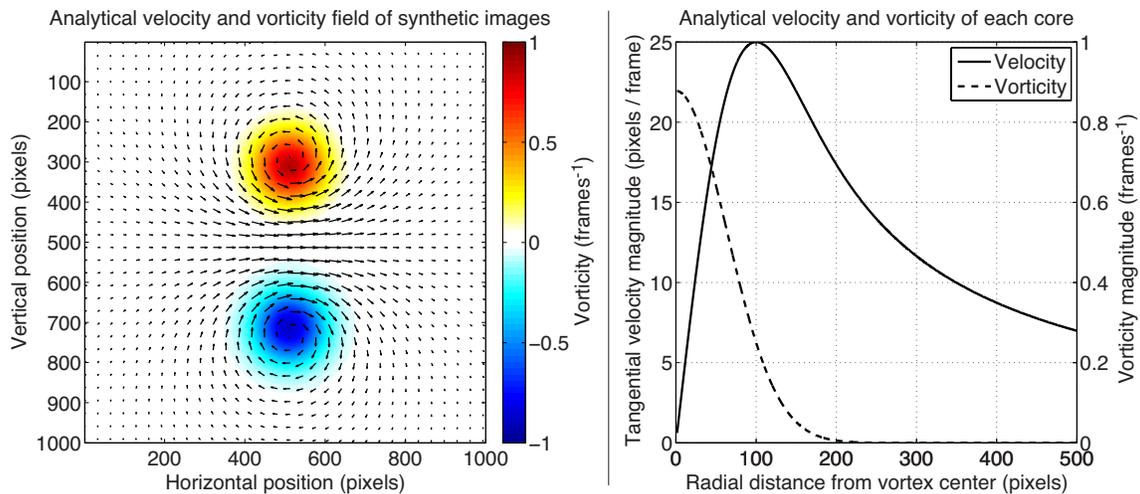


Figure 1.2: Left: ground truth velocity and vorticity fields used to generate PIV images for comparing correlation algorithms. Right: radial velocity and vorticity profiles of each vortex core for synthetic images.

We used FMC and RPC to calculate the velocity and vorticity fields of these images. Both algorithms used three-iteration discrete window offset (DWO) [14] as well as iterative image deformation (IID) [5]. The spacings between IRs at each grid refinement iteration were 32×32 , 16×16 , and 8×8 pixels, and the dimensions of the corresponding IRs were 128×128 , 64×64 , and 32×32 pixels. IRs were apodized with Gaussian windows in the same way as for the Monte Carlo simulations [12]. Statistical outlier rejection and replacement ("universal outlier detection", or UOD [23]) was performed after each iteration, and the central-difference DWO window shift was performed based on this "validated" vector field.

For IID, the UOD-validated velocity field was also smoothed prior to deforming the particle images. We used a Gaussian smoothing kernel with a standard deviation equal to half the IR size, encoded in a 7×7 grid-point stencil (i.e., approximate area of two IRs). Both FMC and

RPC were conservatively allowed ten image deformation iterations. The grid spacing and IR dimensions were refined only between the first, second, and third iterations, after which seven more iterations were performed at the final resolution. We defined the processing to have converged to its lowest achievable error when the change in average near-core displacement error between iterations (the “error residual”) fell below 0.01 pixels.

Vorticity was estimated from spatial gradients of the UOD-validated displacement fields (second order central difference) prior to any smoothing. For FMC with DWO, the "direct estimate" of vorticity was calculated as twice the measured angle of rotation. However, FMC's direct vorticity estimate is lost with IID because the rotation of particle patterns in deformed images no longer represents the vorticity in the original flow field.

Experimental vortex ring

We performed a vortex ring PIV experiment to assess the performance of our FMC algorithm with experimental images. A laminar vortex ring ($Re \approx 2300$) was generated in water using a piston-cylinder arrangement to drive flow through an orifice plate (Figure 1.3). The diameter of the cylindrical piston was 138 mm, and the diameter of the orifice was 25.4 mm. The piston was driven hydraulically by pumping water into the space behind it with a gear pump. The volumetric flow rate and total volumetric flow of water through the orifice was controlled nominally by varying the speed and duration of the driving pump. Flow rates were verified by measuring the displacement of the front edge of the piston in images recorded by a second high-speed camera (IDT XS-5). The water was seeded with 10 μm hollow glass spheres (Potters industries) and 27 μm red fluorescent polystyrene spheres (Duke Scientific Corporation catalog number 36-5B) as flow tracer particles, and the flow was illuminated using a thin light sheet produced by a dual-head Nd:YLF 527 nm pulsed laser (New Wave Pegasus). Both heads of the laser were pulsed simultaneously at 1000 Hz in synchrony with the exposure of a single high-speed camera, whose sensor was oriented approximately parallel to the light sheet (Photron APX-RS, 1024×1024 pixels). A Nikon 85 mm $f/1.4$ camera lens resulted in a field of view of approximately 75×75 mm.

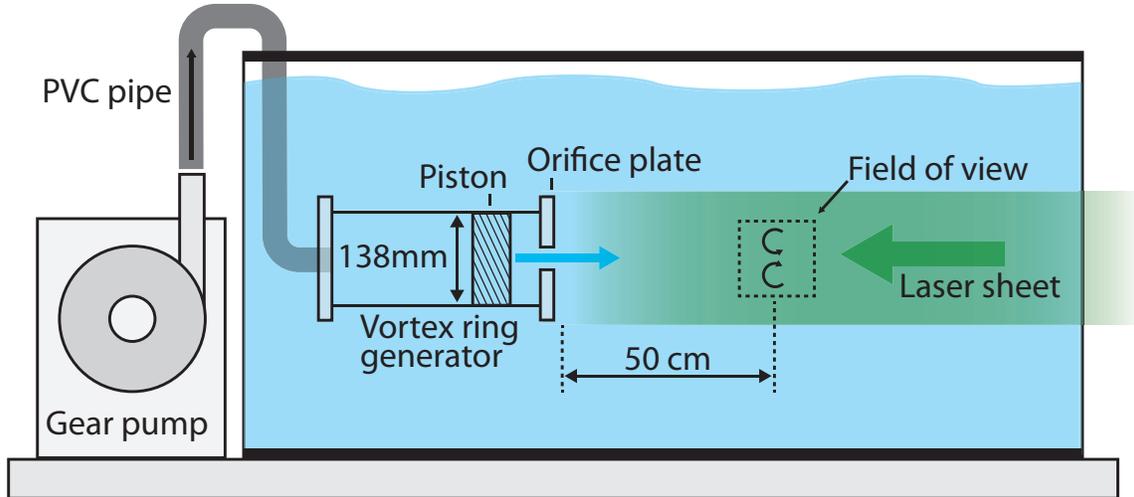


Figure 1.3: Illustration of the experimental apparatus used to capture PIV images of a vortex ring. The camera axis is perpendicular to the page, and the camera itself is not shown. The diameter of the vortex generator orifice was 25.4 mm.

We measured velocity fields in these images using single iterations of FMC and RPC in order to demonstrate the difference in performance between the methods independently of iterative image deformation methods. For both methods, the interrogation region dimensions were 128×128 pixels, and were apodized by a Gaussian window whose effective dimensions were 64×64 pixels. The spacing between grid points was 16 pixels in both the row and column directions.

Because the images were recorded at a high frame rate (1000 Hz) and were not double-pulsed, we were able to systematically increase the apparent strength of the vortex by increasing the inter-frame separation time between pairs of processed images (Δt). To compare the performances of FMC and RPC under increasing vortex strength, we used both methods to process image pairs separated by Δt of 1 to 12 ms. The corresponding maximum particle displacements near the vortex cores were about 2 – 24 pixels. As an additional point of comparison, we calculated the λ_2 vortex identification criterion [24] from spatial gradients (second order central difference) of the UOD-validated velocity fields measured by each correlation method.

Results

Monte Carlo error analysis

RPC displacement estimates

The RPC algorithm recovered particle pattern displacement to within one pixel only for rotation angles below about 6 – 8 degrees and scaling (dilation) factors between about 0.85 and 1.15 (i.e. $1.15^{\pm 1}$) for 64×64 -pixel regions. For 128×128 -pixel regions, this range was reduced to ± 4 –6 degrees of rotation and dilation factors of $1.1^{\pm 1}$. Not surprisingly, the error of the displacement estimates decreased drastically for rotation angles near zero and near-unity scaling factors: for rotation angles of about 1–2 degrees, the average displacement errors were below 0.1 pixels. Figure 1.4 shows the average RPC displacement errors for 64×64 -pixel particle image pairs that rotated and dilated, with displacement held constantly at zero (results for 128×128 -pixel regions are not shown). Zero displacement represents the “best case” scenario for the RPC algorithm because there is no loss of correlation due to in-plane displacement [22].

These results suggest that the FMC algorithm should measure displacement accurately when its rotation estimate is accurate to within 1–2 degrees and its scaling estimate is accurate to within a few percent of the true value, since it uses these estimates to bring the pair of IRs into alignment before estimating displacement with the RPC algorithm.

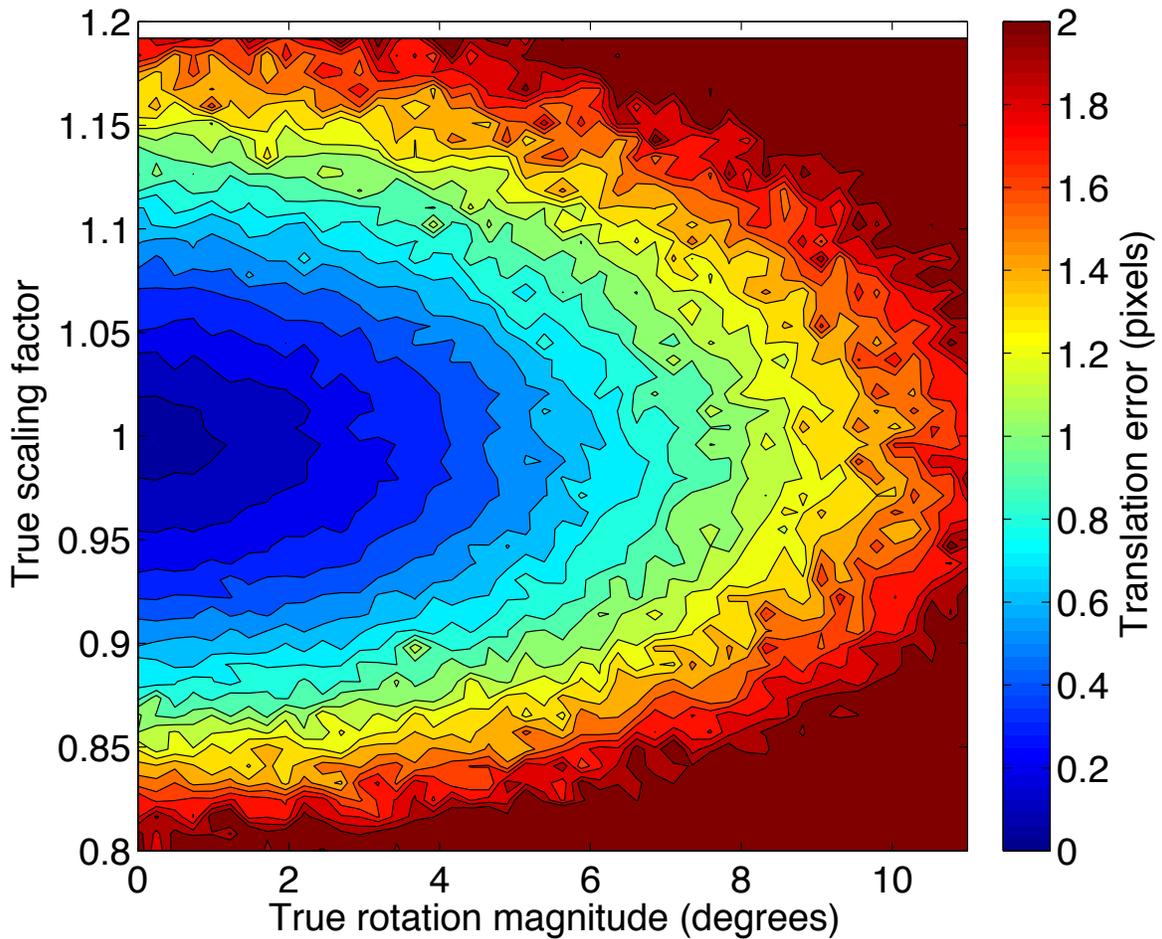


Figure 1.4: Monte-Carlo calculated average errors of single-iteration RPC translation estimates in the presence of rotation and scaling (and zero translation) for one million pairs of 64×64 -pixel interrogation regions.

FMC estimates of displacement, rotation angle, and dilation factor

Figure 1.5 shows the rotation, scaling (dilation), and translation errors for 64×64 -pixel regions; errors for the 128×128 -pixel regions were always lower than these, and are not illustrated. For each plot in this figure, parameters not shown on either axis were held fixed at zero (for rotation and translation) or unity (for dilation). For true scaling factors of less than about 1.5, the average errors of FMC displacement estimates were below 0.1 pixels across the full range of rotation angles (± 180 degrees) and displacement magnitudes (0 to $1/4$ the effective diameter of the Gaussian-apodized region) for both region sizes. This accuracy was nearly independent of the true rotation angle. For larger scaling factors of 1.75–2.0, the displacement error increased to almost one pixel for the 64×64 regions, but remained below 0.06 pixels in all cases for the 128×128 regions (not shown). The higher error with increasing scaling factor is

due to particles exiting the frame (similar to the in-plane loss of correlation due to translation, denoted by Keane and Adrian as F_i [22]), which is mitigated in the larger regions by the greater number of particle images they contain (denoted as N_i [22]). In 64×64 -pixel regions, FMC measured rotation angles to within 0.2 degrees of the true value on average for scaling factors less than 1.5 and across all tested angles and displacements. At the highest scaling factors, the maximum average rotation error was 5 degrees. Meanwhile, the rotation error was always less than 0.1 degrees on average for 128×128 -pixel regions for the full range of tested parameters. The average error of the FMC dilation estimate in 64×64 -pixel regions was below 0.5% (i.e., an absolute error of 5×10^{-3} in the dilation factor) for true dilation factors less than 1.5, and 7% (7×10^{-2}) at the larger dilation factors. For 128×128 regions, the largest average error was $0.6\% \times 10^{-3}$ across the full range of tested parameters.

FMC Errors, 64 x 64 pixel interrogation regions (1 million region pairs each)

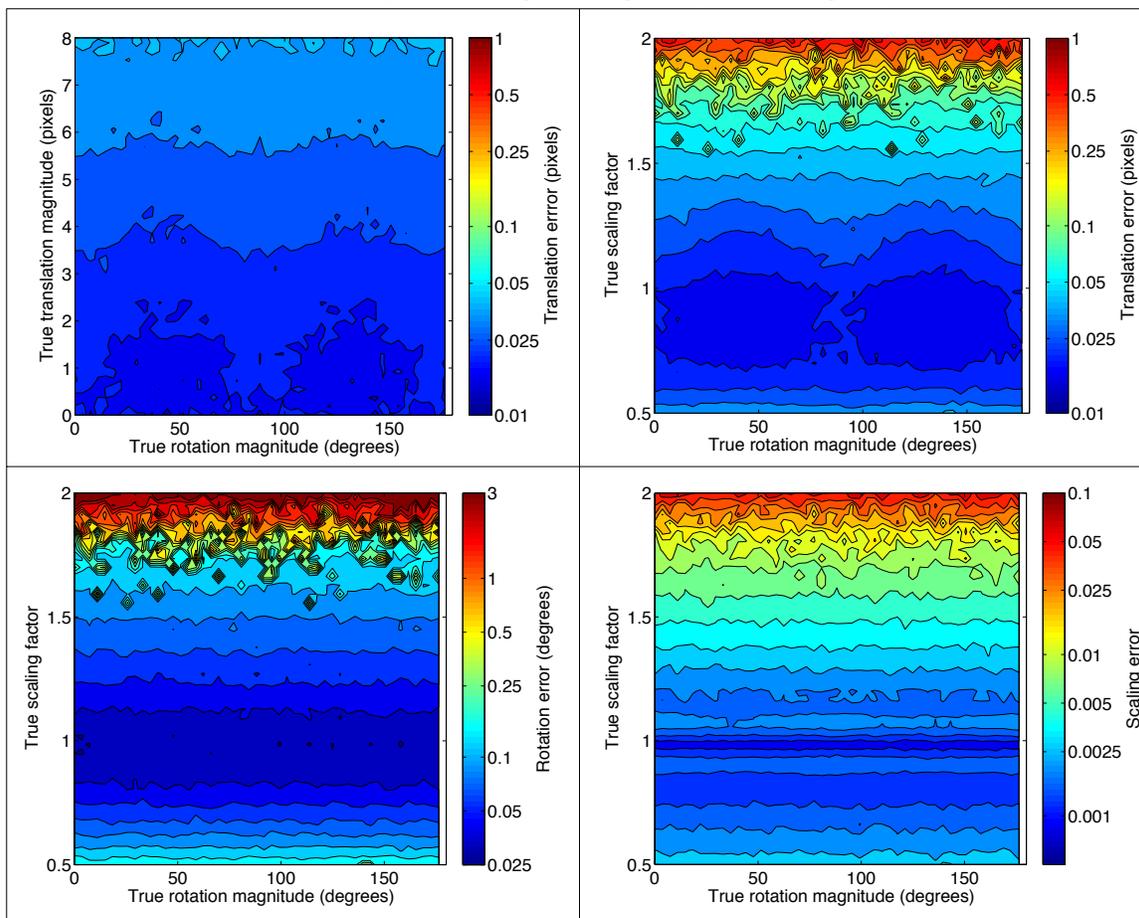


Figure 1.5: Average errors of FMC estimates of translation (top row), rotation (bottom left), and scaling (bottom right) for 64×64 -pixel synthetic interrogation regions. Parameters not shown on either axis were held fixed at zero (rotation and translation) or unity (dilation).

Synthetic Vortex Ring

We compared velocity and vorticity calculated with RPC and FMC near the cores of the synthetic vortex images. Figure 1.6 (left) shows a qualitative comparison between the FMC and RPC-measured displacement fields near one core for a single identical pair of synthetic images using discrete window offset (DWO). Here, the top row of vector plots show the initial measured field, before UOD or DWO was performed. The bottom row shows the fields measured in the third DWO iteration. FMC appears qualitatively to capture the core's displacement field on the first iteration without any outlier replacement (Figure 1.6A), while the first iteration of the RPC-calculated field contains clearly spurious vectors (Figure 1.6B). Subsequent UOD and DWO iterations only partially mitigated the bad RPC vectors (Figure 1.6D). In contrast, FMC captures the core velocity field throughout all three DWO iterations (Figure 1.6C). Moreover, when FMC vectors in this flow did fail, they tended not to fail in groups, so spurious vectors could be reliably identified and replaced using UOD before applying the discrete window offset.

The right side of Figure 1.6 shows the average errors of the displacement estimates calculated using RPC and FMC with DWO and IID compared to the ground-truth Eulerian displacement field near the vortex cores (the region where $r < r_c$ in Equation (1.7)). These results show that the mean and standard deviation of displacement errors were always lower for FMC with DWO than for RPC with DWO, and that the standard deviation of FMC with DWO was always lower than for RPC with IID, even after the latter method converged.

Combining FMC with IID reduced the standard deviation of the converged near-core displacement error by an order of magnitude compared to RPC with IID (0.06 pixels compared to 0.60 pixels for RPC), reduced the average error magnitude by 15% (0.11 pixels compared to 0.13 for RPC), and accelerated convergence from eight iterations (RPC) to three (FMC). Figure 1.6 illustrates the average and standard deviation of errors for both methods, including their converged values, and Figure 1.7 and Figure 1.8 show the convergence behavior of each method. For both methods, the average and standard deviation of displacement error unsurprisingly decreased slightly at higher particle seeding densities of 0.100 particles per pixel (0.10 ± 0.06 pixels for FMC).

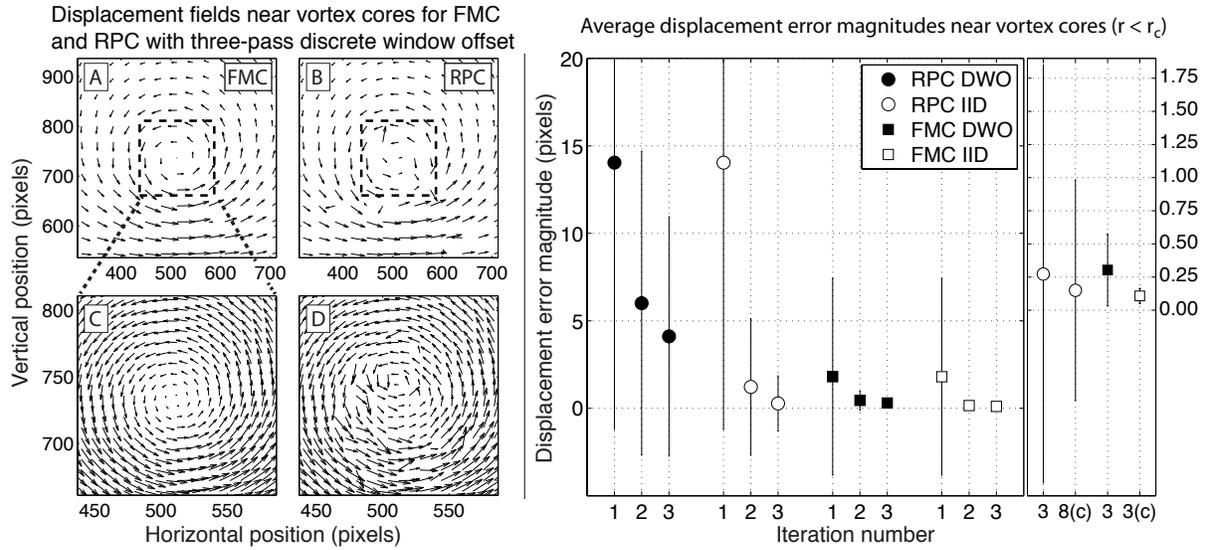


Figure 1.6: Left: displacement vectors near the top core of a pair of synthetic PIV images of counter-rotating vortices. Top and bottom rows show un-validated (no UOD) results from first and third DWO iterations for FMC (A and C) and RPC (B and D). Right: average displacement errors near vortex cores for 1000 pairs of synthetic images. Rightmost axes show detail of third iterations and IID converged results (indicated by "(c)" and the number of the iteration by which convergence occurred). Whiskers represent one standard deviation.

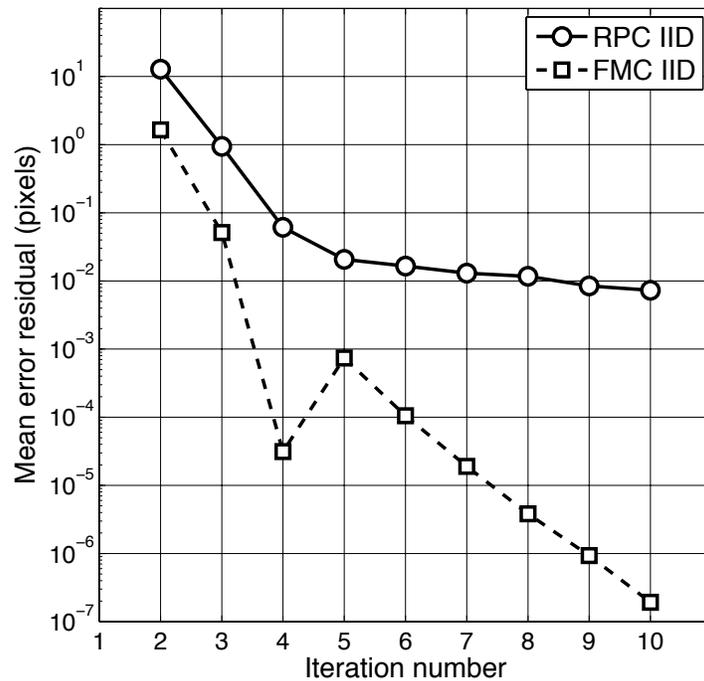


Figure 1.7: Inter-pass residuals (differences) of average displacement errors near vortex cores ($r < r_c$) for FMC and RPC with IID. FMC and RPC residual fall below the convergence criterion of 0.01 pixels after three and eight iterations, respectively.

The FMC measurement converges by its third iteration (Figure 1.7), with nearly 100% of the errors falling below 0.3 pixels (Figure 1.8). In comparison, the RPC IID error is 0.21 ± 1.3 pixels by its fourth iteration, and decreases to 0.13 ± 0.65 pixels at convergence after eight iterations. The relatively large error of the FMC displacement estimate in this flow compared to those from our Monte Carlo simulations is probably due to two primary factors. First, the vortex of equation (1.7) exhibits shearing, while the Monte Carlo simulations did not, and shearing degrades the FMC algorithm's estimates of rotation and translation. Secondly, our "ground truth" solution for the synthetic vortex was the Eulerian velocity field of a curvilinear flow, while correlation-based algorithms like FMC and RPC fundamentally estimate the Lagrangian displacements of particle patterns (our Monte Carlo ground-truth deformations were Lagrangian). When curvilinear motion is significant on the time-scale of the PIV image separation, the difference between the Eulerian and Lagrangian velocities is not negligible, causing a systematic error in the PIV measurement [25]. This is a fundamental limitation of two-frame PIV algorithms, and our FMC algorithm is therefore subject to it. Others have proposed multi-frame methods to overcome this limitation and increase the accuracy of Eulerian velocity estimates in curvilinear flows [4, 25], which we did not implement. However, our algorithm uniquely succeeds in overcoming the recognized difficulty of calculating valid vectors in regions of both large rotation and translation [25] using only single pairs of images.

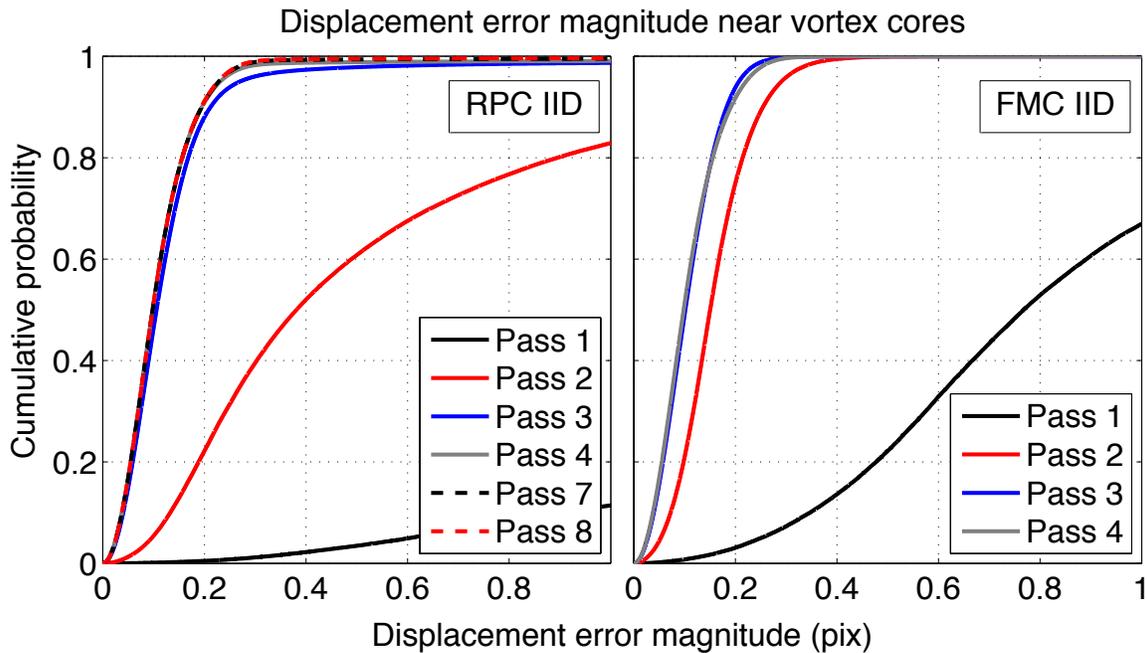


Figure 1.8: Cumulative probabilities of displacement errors near the cores of 1000 synthetic vortex images for RPC (left) and FMC (right), both with iterative image deformation.

Figure 1.9 compares the errors of each method's vorticity estimates. Here, gradients of velocity were calculated from the validated (post-UOD) velocity fields, but the FMC direct estimate is reported prior to UOD. The mean and variance of the FMC direct estimate's error was lower at each iteration than the RPC DWO gradient-based estimate, but generally higher than the other gradient-based estimates. The lowest vorticity errors were again due to FMC combined with IID. Taking spatial gradients of the converged FMC IID displacement field resulted in vorticity errors (mean \pm one standard deviation) of $6.0 \times 10^{-3} \pm 4.5 \times 10^{-3}$ frames⁻¹, compared to $7.1 \times 10^{-3} \pm 25.9 \times 10^{-3}$ frames⁻¹ for RPC with converged IID and $2.4 \times 10^{-2} \pm 2.4 \times 10^{-2}$ frames⁻¹ for the third iteration of FMC DWO.

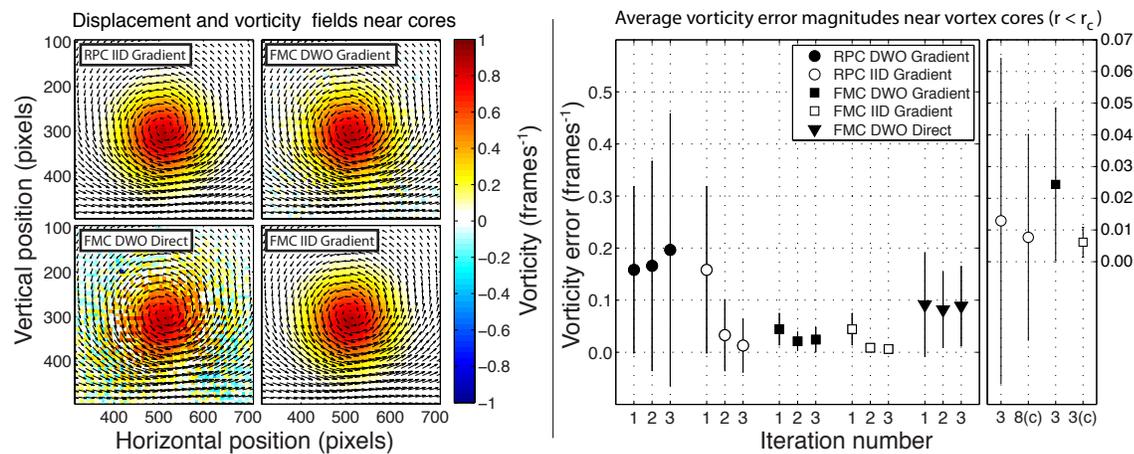


Figure 1.9: Left: vorticity fields calculated with gradient-based and FMC-direct methods, with displacement fields overlaid. DWO results shown after three iterations, and IID shown after convergence. RPC with DWO not shown. Right: average vorticity errors near vortex cores for 1000 pairs of synthetic images. Rightmost axes show detail of third iterations and IID converged results (indicated by "(c)" and the number of the iteration by which convergence occurred). Whiskers represent one standard deviation.

We speculate that all the gradient-based vorticity estimates (except for RPC DWO) outperformed the FMC direct measurement in part because the direct method equates vorticity with rotation, so the contribution of shearing to vorticity is lost. This is a fundamental shortcoming of our method. The increased performance of FMC when combined with IID suggests that perhaps the best use of the FMC algorithm is to complement existing iterative methods by reducing the occurrence of spurious vectors and increasing the fidelity of measured velocity fields for the accurate gradient-based calculation of vorticity or other parameters of interest in regions of rotating flow.

Experimental vortex ring

The same qualitative improvements of FMC over RPC were evident for experimental images as they were for the synthetic vortex tests. Figure 1.10 shows displacement fields near the lower vortex core of our experimental images for Δt of 4, 7, and 10 ms. For clarity, vectors are plotted with statistical outliers removed but not replaced, while the λ_2 criterion was calculated following outlier replacement. Even at relatively small Δt of 4 ms, where the highest near-core particle displacements were only about 8 pixels, the RPC-measured fields contained obviously erroneous vectors that escaped statistical outlier detection but nonetheless failed to capture the vortical motion in the core. The corresponding λ_2 distributions were nonphysical, and this failure was only exacerbated at larger values of Δt . This indicates that the vortex detection algorithm failed using RPC-measured vectors. Meanwhile, FMC captured the vortical motion in the vortex core and maintained physically reasonable distributions of λ_2 for Δt up to 10 ms. The FMC displacement measurement began to fail in this flow for Δt greater than about 10 ms, as evidenced by the greater number of statistical outliers removed in the outer regions of the core where the displacement is relatively large (> 20 pixels per frame) compared to cases with smaller Δt . This experiment sufficiently stressed RPC to the point of obvious failure, while FMC was able to measure displacements and locate vortex cores without the assistance of iterative deformation methods. This illustrates the robustness of the FMC correlation algorithm compared to traditional correlations in the presence of significant rotation.

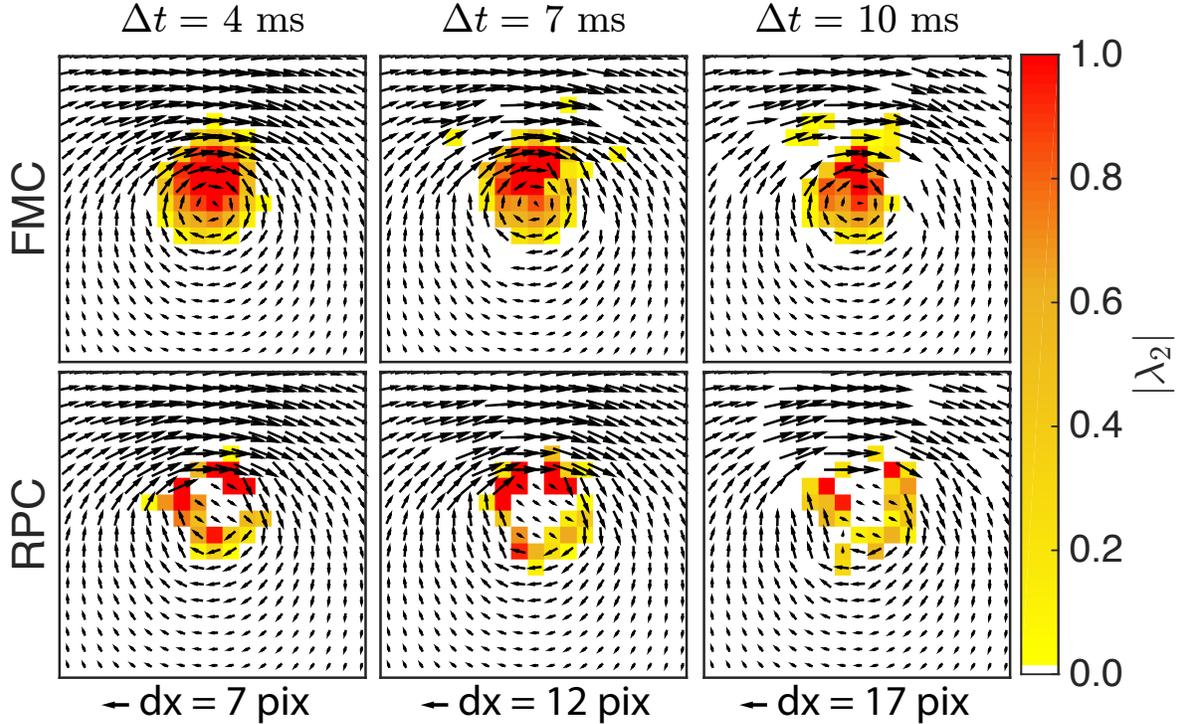


Figure 1.10: Displacement fields and magnitudes of the λ_2 vortex identification criterion [24] near one core of an experimental vortex ring using single iterations of the FMC and RPC algorithms. The height and width of each plot is 300 pixels (about 22 mm). Vector lengths are scaled by each value of Δt for consistent visualization, and $|\lambda_2|$ is normalized by its maximum value in each plot. For clarity, vectors identified as statistical outliers are not shown, but were replaced using UOD [23] prior to calculating λ_2 .

Conclusions

Traditional PIV cross correlations are fundamentally limited by their inability to correctly measure particle pattern motion between IRs that are misaligned simultaneously by translation, rotation, and dilation. Our FMC algorithm addresses the root cause of this limitation and overcomes it by replacing the usual correlation kernel with one that intrinsically measures rotation and dilation of the pattern, even in the presence of large displacements. In this way, we significantly increase the accuracy of PIV measurements of velocity and vorticity in strongly rotating or dilating flows.

Specifically, FMC extends the range of measurable rotational misalignments of particle patterns by over an order of magnitude to ± 180 degrees compared to less than ± 10 degrees for traditional methods, and dilation factors from $1.15^{\pm 1}$ to $2.0^{\pm 1}$. Replacing RPC with FMC in iterative image deformation algorithms accelerates their convergence (3 compared to 8

iterations) and significantly reduces the errors in the displacement measurement (FMC 0.11 ± 0.06 pixels, RPC 0.13 ± 0.65 pixels) and the vorticity calculation (FMC $6.0 \times 10^{-3} \pm 4.5 \times 10^{-3}$ frames⁻¹, RPC $7.1 \times 10^{-3} \pm 25.9 \times 10^{-3}$ frames⁻¹).

Additionally, FMC's direct estimate of rotation may be useful when paired with trajectory-based PIV methods [4, 25], which typically use multiple pairs of IRs to iteratively build estimates of particle pattern rotation. More broadly, FMC overcomes a primary assumption of PIV that the motions of tracer particles in fluid flows are purely translational, and instead permits the class of transformations that include rotation and isotropic dilation. Measuring these higher-order deformations is a step towards directly measuring the full deformation tensor of fluid mechanics within individual PIV interrogation regions.

References

1. Huang, H.T., H.E. Fiedler, and J.J. Wang, *Limitation and improvement of PIV*. Experiments in Fluids, 1993. **15**(3): p. 168-174.
2. Huang, H., D. Dabiri, and M. Gharib, *On errors of digital particle image velocimetry*. Measurement Science & Technology, 1997. **8**(12): p. 1427-1440.
3. Scarano, F. and M.L. Riethmuller, *Advances in iterative multigrid PIV image processing*. Experiments in Fluids, 2000. **29**: p. S51-S60.
4. Lynch, K. and F. Scarano, *A high-order time-accurate interrogation method for time-resolved PIV*. Measurement Science and Technology, 2013. **24**(3): p. 035305.
5. Scarano, F., *Iterative image deformation methods in PIV*. Measurement Science & Technology, 2002. **13**(1): p. R1-R19.
6. Ruan, X., Song, X., and Yamamoto, F., *Direct measurement of the vorticity field in digital particle images*. Experiments in Fluids, 2001. **30**: p. 696-704.
7. Chen, Q.S., M. Defrise, and F. Deconinck, *Symmetric phase-only matched filtering of Fourier-Mellin transform for image registration and recognition*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1994. **16**(12): p. 1156-1168.
8. Reddy, B.S. and B.N. Chatterji, *An FFT-based technique for translation, rotation, and scale-invariant image registration*. Trans. Img. Proc., 1996. **5**(8): p. 1266-1271.
9. Bracewell, R.N., *Fourier analysis and imaging*. 1 ed. 2003: Kluwer Academic/Plenum Publishers.

10. Wolberg, G. and S. Zokai, *Robust image registration using log-polar transform*. in *Proceedings of the International Conference on Image Processing, 2000*. 2000.
11. Nobach, H. and M. Honkanen, *Two-dimensional Gaussian regression for sub-pixel displacement estimation in particle image velocimetry or particle position estimation in particle tracking velocimetry*. *Experiments in Fluids*, 2005. **38**(4): p. 511-515.
12. Eckstein, A. and P.P. Vlachos, *Assessment of advanced windowing techniques for digital particle image velocimetry (DPIV)*. *Measurement Science & Technology*, 2009. **20**(7).
13. Wernet, M.P., *Symmetric phase only filtering: a new paradigm for DPIV data processing*. *Measurement Science & Technology*, 2005. **16**(3): p. 601-618.
14. Scarano, F. and M.L. Riethmuller, *Iterative multigrid approach in PIV image processing with discrete window offset*. *Experiments in Fluids*, 1999. **26**(6): p. 513-523.
15. Wereley, S.T. and C.D. Meinhart, *Second-order accurate particle image velocimetry*. *Experiments in Fluids*, 2001. **31**(3): p. 258-268.
16. Westerweel, J., D. Dabiri, and M. Gharib, *The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings*. *Experiments in Fluids*, 1997. **23**(1): p. 20-28.
17. Eckstein, A. and P.P. Vlachos, *Digital particle image velocimetry (DPIV) robust phase correlation*. *Measurement Science & Technology*, 2009. **20**(5).
18. Eckstein, A.C., J. Charonko, and P. Vlachos, *Phase correlation processing for DPIV measurements*. *Experiments in Fluids*, 2008. **45**(3): p. 485-500.
19. Olsen, M.G. and R.J. Adrian, *Out-of-focus effects on particle image visibility and correlation in microscopic particle image velocimetry*. *Experiments in Fluids*, 2000. **29**(1): p. S166-S174.
20. Brady, M.R., S.G. Raben, and P.P. Vlachos, *Methods for digital particle image sizing (DPIS): comparisons and improvements*. *Flow Measurement and Instrumentation*, 2009. **20**(6): p. 207-219.
21. Devenport, W.J., et al., *The structure and development of a wing-tip vortex*. *Journal of Fluid Mechanics*, 1996. **312**: p. 67-106.
22. Keane, R.D. and R.J. Adrian, *Theory of cross-correlation analysis of piv images*. *Applied Scientific Research*, 1992. **49**(3): p. 191-215.
23. Westerweel, J. and F. Scarano, *Universal outlier detection for PIV data*. *Experiments in Fluids*, 2005. **39**(6): p. 1096-1100.
24. Jeong, J. and F. Hussain, *On the identification of a vortex*. *Journal of Fluid Mechanics*, 1995. **285**: p. 69-94.

25. Hain, R. and C.J. Kähler, *Fundamentals of multiframe particle image velocimetry (PIV)*. Experiments in Fluids, 2007. **42**(4): p. 575-587.

Chapter 2: Adaptive Spectral Filtering of Particle Image Velocimetry Cross Correlations

Abstract

Particle image velocimetry (PIV) cross correlations (CCs) are usually thought of in the spatial domain, where the location of a correlation's maximum value represents the most probable average displacement of a group of flow tracer particles. In this paper, we instead consider the correlation in the Fourier (spectral) domain, and argue that doing so reveals characteristics and behaviors about PIV CCs that are obscured in the spatial domain. We show how the shapes of the tracer particles and the distributions of their individual displacements cause the signal-to-noise ratio of the spectral correlation to vary across spectral wave numbers, and thereby explain the previously heuristic observation that the effects of "noise" become more severe at higher wave numbers. We use this insight to create a spectral weighting filter that attenuates noisy wave numbers, and present an algorithm to automatically determine the parameters of this filter based on quantities measured from the spectral correlation planes themselves. We call our algorithm the "adaptive phase correlation" or APC, and compare its performance to the standard cross correlation (SCC) and a statically filtered phase correlation. We demonstrate that our APC algorithm significantly reduces bias and RMS errors of displacement estimates from synthetic PIV images whose particles are subject to significant random displacements (e.g., turbulence or Brownian motion), and also accelerates the convergence of ensemble correlation measurements for steady or quasi-steady flows. Finally, we analyze a set of real PIV images that required significant "pre-processing" to yield physically realistic results using prior correlation methods, and show that our algorithm provided comparable results automatically, without any such ad-hoc, subjective image manipulation.

Introduction

The mental model of a "correlation plane" interspersed with "peaks" of different "heights" has become part and parcel of PIV analysis, and lends itself to various "signal to noise ratio" (SNR) metrics related to the method's performance. For instance, the uncertainty of an individual PIV displacement measurement can be estimated by the ratio of heights of the largest and second-largest correlation peaks (the "peak height ratio" or PPR). The paradigm of estimating PIV displacements from the spatial correlation plane is pervasive for good reason. It

is fairly robust to noise and other sources of decorrelating error; it provides for estimation of statistical metrics such as the probability density of displacements (e.g., velocity gradients, Brownian motion); and moreover its topology is arguably intuitive.

However, like in all realms of signal processing, in PIV analysis the spatial domain is only part of the picture, and can be complemented by analysis in the Fourier (spectral, frequency) domain. With a few exceptions, cross correlations are calculated in the spectral domain for computational efficiency, but are rarely visualized there. In this paper, we will show that visualizing the correlation in the spectral domain reveals structure that is obscured in the spatial domain, and argue that understanding this structure provides important insight into the processes that drive error in PIV displacement estimates. To preface our argument, Figure 2.1 shows a representative PIV cross correlation in the spatial domain (the usual depiction), along with the same correlation visualized in the spectral domain.

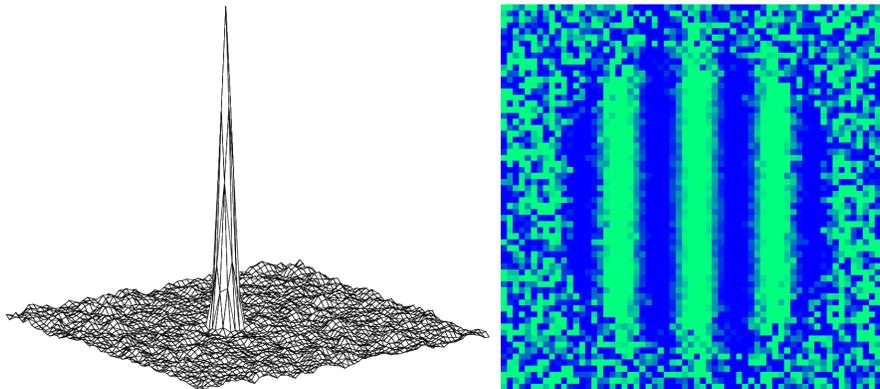


Figure 2.1: Representative cross correlation (CC; average of ten) visualized in the usual spatial domain (left), and the corresponding real part of the phase of the CC in the Fourier (spectral, frequency) domain. The axes of the spatial correlation correspond to horizontal and vertical spatial shifts, while those of the spectral correlation represent horizontal and vertical wave numbers (or spatial frequencies). The prominent “banding” structure in the spectral domain represents a cosine wave whose orientation and wavelength indicate the direction and magnitude of the pattern displacement (in this case, a horizontal displacement of five pixels over a 64x64 pixel domain).

Background

The success or failure of any cross correlation-based PIV measurement depends primarily on the correct identification of the correlation peak. Under ideal circumstances, the displacement peak is prominent within its domain and can be located reliably and precisely. In practice, PIV measurements are subject to myriad sources of "noise" or "decorrelating effects" that can interfere with detecting the displacement peak or obscure its location, which ultimately

reduce the accuracy or precision of the measurement. Some examples of common decorrelating effects include thermal image noise, velocity gradients within the correlated interrogation regions [1], and the disappearance of particles between snapshots due to their in-plane and out-of-plane motion [2]. Figure 2.2 illustrates the effects of an in-plane velocity gradient on the spatial correlation.

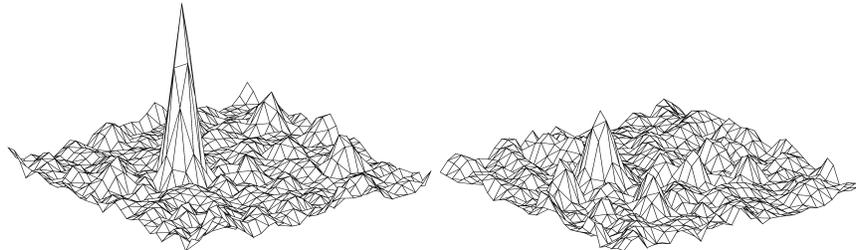


Figure 2.2: Representative examples of cross correlation planes for a perfectly correlated pair of PIV interrogation regions (left) and for a pair of IRs subject to in-plane gradients of displacement due to shearing. In this example, the presence of displacement gradients has diminished the height of the primary correlation peak.

In discussions about the accuracy and uncertainty of PIV measurements vis-à-vis these processes, it is common to think of a "signal to noise ratio" (or SNR) of the correlation plane (or volume), where the "signal" refers to its features that contribute to correctly locating the displacement peak, and the "noise" is the manifestation of everything that obscures it.

While the concept of a correlation SNR is straightforward, quantifying it is somewhat subjective, and several different metrics have been proposed for PIV [2-5]. Regardless of the specific metric used, the SNR is typically thought of as having a single value, where the strength of the "signal" is usually related to the magnitude of the displacement peak relative to some other quantity like the correlation RMS or the magnitude of the second-largest peak. When the SNR is sufficiently small (i.e., when the correlation is dominated by "noise"), often times the displacement estimate can yield a spurious vector that is incongruent with the velocity field under investigation. When this happens, the measurement can be said to have failed. Keane and Adrian [2] proposed that a PIV measurement was susceptible to failure when the ratio of its largest to second-largest correlation peak (the "primary peak ratio," or PPR) fell below a threshold value, and demonstrated that the probability of exceeding this threshold PPR was related to the number of individual particles common to both interrogation regions in the correlated pair. In addition to predicting object failure, PPR-based SNR metrics have been shown to reliably estimate the uncertainty of PIV-measured pattern displacements under a variety of conditions ([3, 4, 6-8]).

Under certain circumstances, the correlation SNR may be improved using clever interrogation schemes. For example, Keane & Adrian [2] identified that a common source of depressed SNR is the motion of particles into and out of the interrogation region between images, as well as the presence of in-plane gradients of displacement within the regions. These effects manifest as bias errors and are typically mitigated by implementing one of several iterative predictor-corrector algorithms. Image deformation iteratively warps the photographs according to the measured displacement field, which ideally reduces the appearance of spatially correlated velocity gradients in the deformed images [9, 10] and thereby increases the correlation SNR of subsequent correlations. Discrete and continuous window offset (DWO, CWO) reduce the loss of correlation due to particles leaving the interrogation regions between snapshots by shifting the locations of interrogation spots based on previously calculated vectors [11, 12].

Other situations are not helped by deforming the images or their sampling grids. For example, image thermal noise and stochastic particle motion within the interrogation regions (e.g., Brownian motion, turbulence) are detrimental to the correlation SNR, but cannot be mitigated by any spatially adaptive procedures. When these effects are significant, and if the flow under consideration is steady or quasi-steady, then the correlation SNR can be increased by averaging multiple correlation planes calculated at the same spatial location from independent pairs of images (the “ensemble correlation” of Meinhart *et al.* [13]).

These techniques have all been shown to increase the accuracy of PIV measurements, and indeed have become instrumental to the entire gamut of PIV applications. However, they side step rather than address the fundamental problem of low SNR to begin with. By treating the SNR of the cross correlation as a property whose value depends uniquely on its input data, their principles of operation are limited to modifying or augmenting the photographs under consideration. Other work (by many of the same researchers) has complemented these techniques by recognizing that the cross correlation itself is a mutable entity on which operations can be performed to increase its SNR and ultimately reduce errors in the pattern displacement estimate [5, 14, 15].

Introducing the Spectral Domain

With a few exceptions, most computer programs calculate the cross correlation in the Fourier (spectral) domain because doing so is computationally more efficient but mathematically identical to directly convolving the arrays that represent the images. But the PIV and broader

image processing communities have realized that performing cross correlations in the spectral rather than spatial domain has advantages that extend beyond speed [15-18]. In particular, the Fourier domain reveals structure in the cross correlation that is obscured in the spatial domain, and this structure can be leveraged to both understand and mitigate sources of error in the displacement estimate.

Although the effects of noise appear to contaminate the correlation randomly across the entire spatial domain, they are often localized to certain ranges of frequencies or spectral wave numbers in the Fourier domain [5, 15, 18]. This suggests that describing the correlation SNR with a single value is an over simplification, and in fact the correlation SNR varies across spectral wave numbers. In fact, suppressing noise-dominated wave numbers of the cross correlation in the spectral domain (the “spectral correlation”) has been shown to increase the accuracy of the peak-detection-based displacement estimate in the spatial domain for both PIV and general image registration [5, 14, 15, 17, 19].

This is usually accomplished by multiplying the spectral correlation by a weighting matrix that attenuates the noise-dominated wave numbers before returning to the spatial domain via the inverse Fourier transform (FT) and locating the displacement peak. Preferentially weighting certain wave numbers in the Fourier domain is known as spectral filtering, and the array containing the weights is known as the spectral filter.

Spectral filters are often designed based on assumptions about the spectral characteristics of the input data [5, 15]. Eckstein *et al.* were the first to introduce a filter based on the specific characteristics of PIV image formation [15]. They assumed that the SNR of the spectral correlation should follow the power spectrum of the input images – namely, a Gaussian distribution, greatest at the lowest wavenumber and decaying toward higher wavenumbers. They subsequently designed a spectral filter to attenuate the contribution of the noise-dominated wave numbers to the PIV displacement estimate. Their “robust phase correlation” (or RPC) filter reduced bias and random errors in PIV displacement estimates compared to the standard unfiltered cross correlation (SCC).

Eckstein’s model and filter inherently assumed that the SNR of the PIV cross correlation was governed purely by the input images themselves, but neglected any effects due to the relationships *between* the correlated pairs of images (i.e., the particle pattern displacements). This assumption seems to be incongruent with the multitude of work demonstrating the many and varied effects of the pattern displacements on the cross correlation in the spatial domain [1, 2, 20]. It is reasonable to presume that the SNR of the spectral correlation should depend not

only on the contents of the correlated interrogation regions themselves, but also on the particle displacements relating them. But because the particle displacements are not known *a priori* (indeed, this is what PIV is designed to measure), it is not obvious how this information can be incorporated into the design of spectral correlation filters.

The goal of this paper is to unify these two perspectives by explaining why and how the SNR of PIV cross correlations depends on the characteristics of not only the images but also on the displacements relating them. First, we will derive a model of the cross correlation that predicts the variation of correlation SNR across spectral wave numbers based on both the number and shapes of the correlated tracer particles and on the displacements of the patterns. We will illustrate how these effects manifest in both spectral and spatial domains and ultimately affect the peak-location-based displacement estimate. Next, we use the insight from our model to introduce a spectral filter that extends that of Eckstein to account for the effects of pattern displacements, and show that such a filter can increase the prominence of the displacement peak in the spatial correlation when displacement gradients and image noise are significant. We present an algorithm to calculate this filter using measured properties of the cross correlation plane, which relaxes assumptions about the sizes of the tracer particles, accounts for the effects of particle displacements, and adapts to spatially or temporally varying correlation SNR throughout the flow field. Our algorithm overcomes the seemingly circular logic of requiring *a priori* knowledge of the pattern displacements to design filters by measuring the SNR directly from the magnitude of the cross correlation. We call our algorithm the Adaptive Phase Correlation (APC), and demonstrate that it reduces bias and random errors in PIV displacement estimates compared to the un-filtered approach and compared to the RPC filter.

Modeling the correlation in the absence of noise

We model each PIV interrogation region $I(\mathbf{x})$ as a collection of N Dirac-delta impulse functions located at image coordinates \mathbf{x}_n , and each subject to an identical point-spread function $\tau(\mathbf{x})$:

$$I(\mathbf{x}) = \tau(\mathbf{x}) * \sum_n^N \delta(\mathbf{x} - \mathbf{x}_n) \quad (2.1)$$

If all of the N particles are assumed to remain within the field of view during recording frames, then the cross correlation between two PIV interrogation regions is given by

$$R(\mathbf{x}) = \tau(\mathbf{x}) * \tau(\mathbf{x}) * \sum_n^N \sum_m^N \delta(\mathbf{x} - \mathbf{d}_{n,m}) \quad (2.2)$$

where $\mathbf{d}_{n,m} = \mathbf{x}_m - \mathbf{x}_n$ is the displacement relating the positions of particle n in the first frame with particle m in the second frame. If $R(\mathbf{x})$ is separated into the contribution from corresponding particles (i.e., the displacements between images of the same particle recorded in the two consecutive photographs) and non-corresponding particles (all of the other combinations), then (2.2) becomes

$$R(\mathbf{x}) = R_D(\mathbf{x}) + R_C(\mathbf{x}) = a(\mathbf{x}) * \left(\sum_n^N \delta(\mathbf{x} - \mathbf{d}_n) + \sum_n^N \sum_m^N (1 - \delta_{m,n}) \delta(\mathbf{x} - \mathbf{d}_{n,m}) \right) \quad (2.3)$$

where $\delta_{n,m}$ is the Kronecker delta function, and $a(\mathbf{x}) = (\tau * \tau)(\mathbf{x})$. The “corresponding correlation” $R_D(\mathbf{x})$ expresses the displacement that PIV algorithms are concerned with measuring, and is all that contributes to the primary “displacement peak” in the familiar cross correlation plane. Meanwhile, $\tilde{R}_C(\mathbf{k})$ represents the “random” correlations between all of the other permutations of pairs of particles, as well as the sources of random noise, and does not contribute to the primary correlation peak. In equation (2.3), the sum $\sum_n^N \delta(\mathbf{x} - \mathbf{d}_n)$ is identically the histogram of displacements between the corresponding particles. Therefore, the displacement correlation $R(\mathbf{x})$ can be written as

$$R_D(\mathbf{x}) = N \left[a(\mathbf{x}) * p_{\bar{\mathbf{d}}}(\mathbf{x}) * \delta(\mathbf{x} - \bar{\mathbf{d}}) \right] \quad (2.4)$$

where $\bar{\mathbf{d}}$ is the average displacement of the particles, and $p_{\bar{\mathbf{d}}}(\mathbf{x})$ is the probability density function of the displacements of the particles about their mean displacement $\bar{\mathbf{d}}$. Equation (2.4) represents the familiar and ideal “correlation plane” of PIV. The characteristic “peak” of the

correlation plane is located at $\mathbf{x} = \bar{\mathbf{d}}$, and the shape of the peak is governed by the particle shape (expressed in $a(\mathbf{x})$) and the distribution of displacements of the corresponding particles $p_{\bar{\mathbf{d}}}(\mathbf{x})$. The convolution by $p_{\bar{\mathbf{d}}}(\mathbf{x})$ describes the well-known effect of the primary correlation peak “spreading out” and decreasing in magnitude due to the presence of velocity gradients within the field of view [1]. Absent any gradients, and with ideal point-particles, $a(\mathbf{x}) = \delta(\mathbf{x})$ and $p_{\bar{\mathbf{d}}}(\mathbf{x}) = \delta(\mathbf{x})$, and $R_D(\mathbf{x})$ thus reduces to the Dirac- δ function located at $\mathbf{x} = \bar{\mathbf{d}}$.

The Spectral Correlation

Considering $R_D(\mathbf{x})$ in the Fourier domain, however, will elucidate important characteristics about the correlation that are not obvious in the spatial domain. The Fourier transform of $R_D(\mathbf{x})$ is given by

$$\tilde{R}_D(\mathbf{k}) = NA(\mathbf{k})P_{\bar{\mathbf{d}}}(\mathbf{k})e^{-j\omega_{\bar{\mathbf{d}}}\mathbf{k}} \quad (2.5)$$

In equation (2.5) (which we refer to as the “spectral correlation”), the displacement correlation has been transformed into a complex sinusoid in the Fourier domain. This sinusoid has a single frequency $\omega_{\bar{\mathbf{d}}} = 2\pi\bar{\mathbf{d}}$, and it is modulated (multiplied) by the FT of the particle autocorrelation $A(\mathbf{k}) = \mathcal{F}\{a(\mathbf{x})\}(\mathbf{k})$ and by the FT of the PDF of the particles’ displacements, $P_{\bar{\mathbf{d}}}(\mathbf{k}) = \mathcal{F}\{p_{\bar{\mathbf{d}}}(\mathbf{x})\}(\mathbf{k})$.

The consequences of this modulation become apparent if we follow Eckstein in defining a signal-to-noise ratio of the spectral correlation that varies proportionally to the magnitude of $\tilde{R}_D(\mathbf{k})$:

$$\text{SNR}(\mathbf{k}) = \frac{|\tilde{R}_D(\mathbf{k})|}{\beta} \quad (2.6)$$

where β is a constant that characterizes the noise as approximately uniform across the Fourier spectrum, as he justified for the FT of digital images [15]. Substituting (2.5) into (2.6),

$$\text{SNR}(\mathbf{k}) = \frac{NA(\mathbf{k})P_{\mathbf{d}}(\mathbf{k})}{\beta} \quad (2.7)$$

The specific shape of the SNR modulation function, or “envelope” $E(\mathbf{k}) = NA(\mathbf{k})P_{\mathbf{d}}(\mathbf{k})$, is elucidated by calculating the Fourier transforms of the particle autocorrelation and the displacement distribution. If the particles are Gaussian-shaped, then the particle auto correlation is given by

$$A(\mathbf{k}) = \frac{\pi d_{\tau}^2 I_o^2}{8} e^{-\pi^2 \mathbf{k}^2 d_{\tau}^2 / 4} \quad (2.8)$$

Here, I_o is the maximum brightness of the representative particle and $d_{\tau} \equiv 4\sigma_{\tau}$ is the particle diameter, with σ_{τ} the standard deviation of its Gaussian intensity profile. Similarly, if the particles' displacements follow a Gaussian normal distribution, then

$$P_{\mathbf{d}}(\mathbf{k}) = e^{-2\pi^2 \mathbf{k}^2 \sigma_{\mathbf{d}}^2} \quad (2.9)$$

where $\sigma_{\mathbf{d}}$ is the standard deviation of the displacement distribution. Alternatively, if the particle displacements are uniformly distributed across a range $\Delta \mathbf{d} = \mathbf{d}_{\max} - \mathbf{d}_{\min}$ (e.g., shearing), then

$$P_{\mathbf{d}}(\mathbf{k}) = \text{sinc}(k\Delta \mathbf{d}) \quad (2.10)$$

where

$$\text{sinc}(x) = \begin{cases} 1 & \text{for } x = 0 \\ \frac{\sin(x)}{x} & \text{otherwise} \end{cases} \quad (2.11)$$

Equations (2.7)–(2.10) lead to an important conclusion about the spectral cross correlation. Namely, the SNR of the cross correlation varies across spectral wave numbers purely as a consequence of the particles' shapes, and of the fact that their individual displacements are not

all equal but rather are drawn from some distribution. Furthermore, when the particles are Gaussian shaped, the SNR decays toward higher wave numbers *even when the particle patterns are perfectly correlated*. This effect is the fundamental reason that the effects of noise are more prominent at higher wave numbers than lower wave numbers in PIV cross correlations.

Figure 2.3 illustrates a 1-D example of the modulation of the spectral correlation predicted by equation (2.5) for the commonly encountered situation of a constant displacement gradient in one direction (shearing) with Gaussian-shaped particles, with a mean displacement of 5 pixels. In this example, the ratio of the displacement distribution width to the particle diameter, $\Delta d / d_\tau$, was equal to 2.0. For particles with $d_\tau = 3.2$ pixels and a domain 64 pixels wide, this corresponds to a displacement gradient of $\Delta d / L = 0.1$ pixels per pixel.

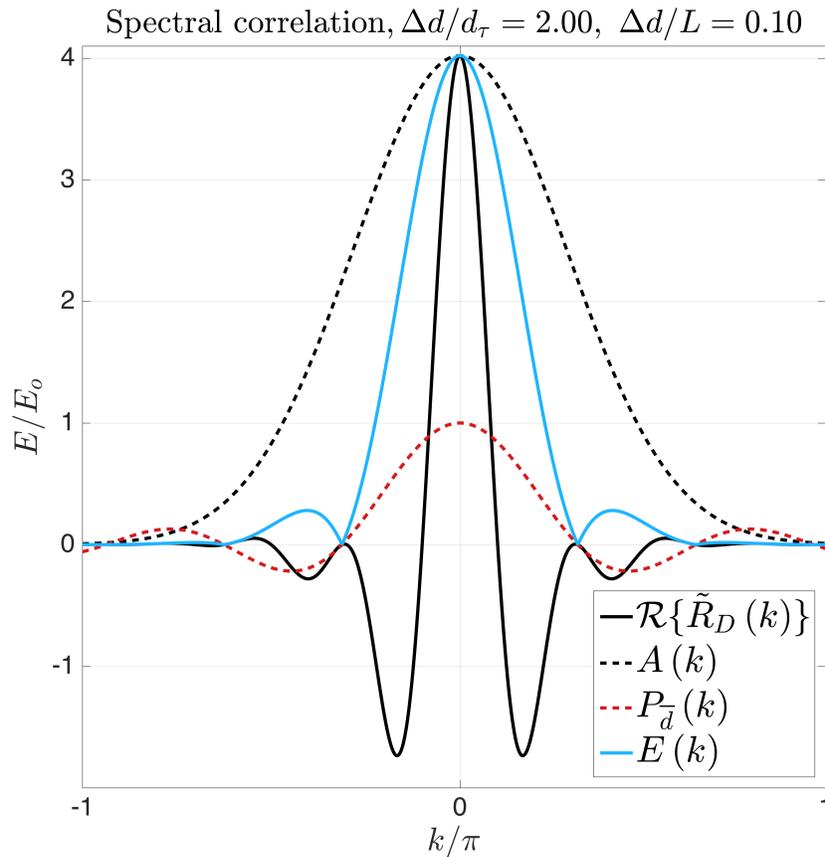


Figure 2.3: Illustration of the spectral SNR predicted by Equation (2.7) for a uniform distribution of displacements. The displacement PDF of a constant displacement gradient is a square pulse, and the FT of a square pulse is the sinc function whose frequency increases with the width of the distribution.

Here, the complex displacement correlation $\tilde{R}_D(\mathbf{k})$ (whose real part is shown) exhibits the characteristic oscillatory behavior described by the complex exponential of equation (2.5). The amplitude of that cosine wave is governed by that of $E(\mathbf{k})$, which itself is the product of $A(\mathbf{k})$ and $P_{\bar{d}}(\mathbf{k})$, and its frequency is $\omega_{\bar{d}}$ represents the average displacement of the particles. The quantity E_o is a constant scalar value that represents the contribution to the magnitude of the complex cross correlation of the convolution of a single pair of ideal unity-amplitude Dirac- δ point particles, and E can be interpreted as the contribution due to a pair of particles with the same maximum intensity but finite diameter d_{τ} . Given that ideal unity amplitude point particles produce $|A(\mathbf{k})|=1$ for all \mathbf{k} (across the entire spectrum), the fact that the ratio $E/E_o > 1$ for $A(\mathbf{k})$ at some low wave numbers and $E/E_o < 1$ at high wave numbers reveals that the finite size of the particles *amplifies* the SNR at low wave numbers and *suppresses* it at high wave numbers compared to the correlation of ideal point particles. The fact that the SNR is amplified by large particles compared to smaller ones provides an additional perspective about why PIV images with larger particles are more robust to the effects of image noise than images with smaller particles are.

In the absence of any noise, the decay of the spectral correlation toward higher wave numbers would not itself be problematic: simply dividing $\tilde{R}_D(\mathbf{k})$ by its own magnitude would normalize its amplitude to become unity everywhere, and its single frequency $\omega_{\bar{d}}$ could be easily recovered in the usual way of taking the inverse FT. This is the rationale behind the “phase correlation,” which is widely used in image processing [18] and has been reported for use in PIV [5, 15], where the phase of the cross correlation is given by

$$\phi(\mathbf{k}) = \frac{\tilde{R}(\mathbf{k})}{|\tilde{R}(\mathbf{k})|} \quad (2.12)$$

In practice, however, noise is *not* negligible, but rather contaminates the correlation approximately uniformly across all wave numbers, with an RMS described by the constant β in equation (2.6). In this situation, the decay of $\tilde{R}_D(\mathbf{k})$ implies that the correlation can become

dominated by noise in regions of the spectrum where $\tilde{R}_D(\mathbf{k})$ becomes small compared to β . In our analysis, we refer to the “noise dominated” regions of $\tilde{R}(\mathbf{k})$ as those in which $\omega_{\bar{a}}$ cannot be reliably estimated from $\tilde{R}(\mathbf{k})$, and therefore do not contribute to the correct estimate of $\omega_{\bar{a}}$. Instead, these noise-dominated regions reduce the portion of the spectral correlation that can be used to estimate $\omega_{\bar{a}}$. This effect can be shown conspicuously by visualizing the phase of the spectral correlation, $\phi(\mathbf{k})$ (as in Figure 2.1), wherein the single-frequency waveform of equation (2.5) is prominently distinguishable across signal-dominated (high SNR) wave numbers, but is obscured by noise in the noise-dominated regions. In the spatial domain, these effects manifest as “noise” in the form of additional “peaks,” as well as a reduction in magnitude of the displacement peak. Together, these effects increase the uncertainty of the measurement [3, 4] or can cause it to fail completely. Figure 2.4 illustrates how increasing displacement distributions (in-plane shearing in this case) affect the correlation plane in both the spatial and spectral domains.

The salient feature of Figure 2.4 is that larger distributions of displacement decrease the height of the displacement peak in the spatial correlation, and concomitantly reduce the extent of the signal-dominated region of the spectral correlation plane. The prominent banding structure in the spectral phase correlation again represents the waveform predicted by the complex exponential in equation (2.5), whose spatial frequency $\omega_{\bar{a}}$ indicates the average displacement of the particle pattern. As the width of the displacement distribution increases, the correlation SNR (i.e., the modulation of the spectral correlation) decays more rapidly with spectral wave number. This drives the onset of noise-dominated wave numbers to occur earlier in the spectrum (toward lower wave numbers), which results in the appearance of the noise encroaching on the phase correlation plane “from the outside inward.”

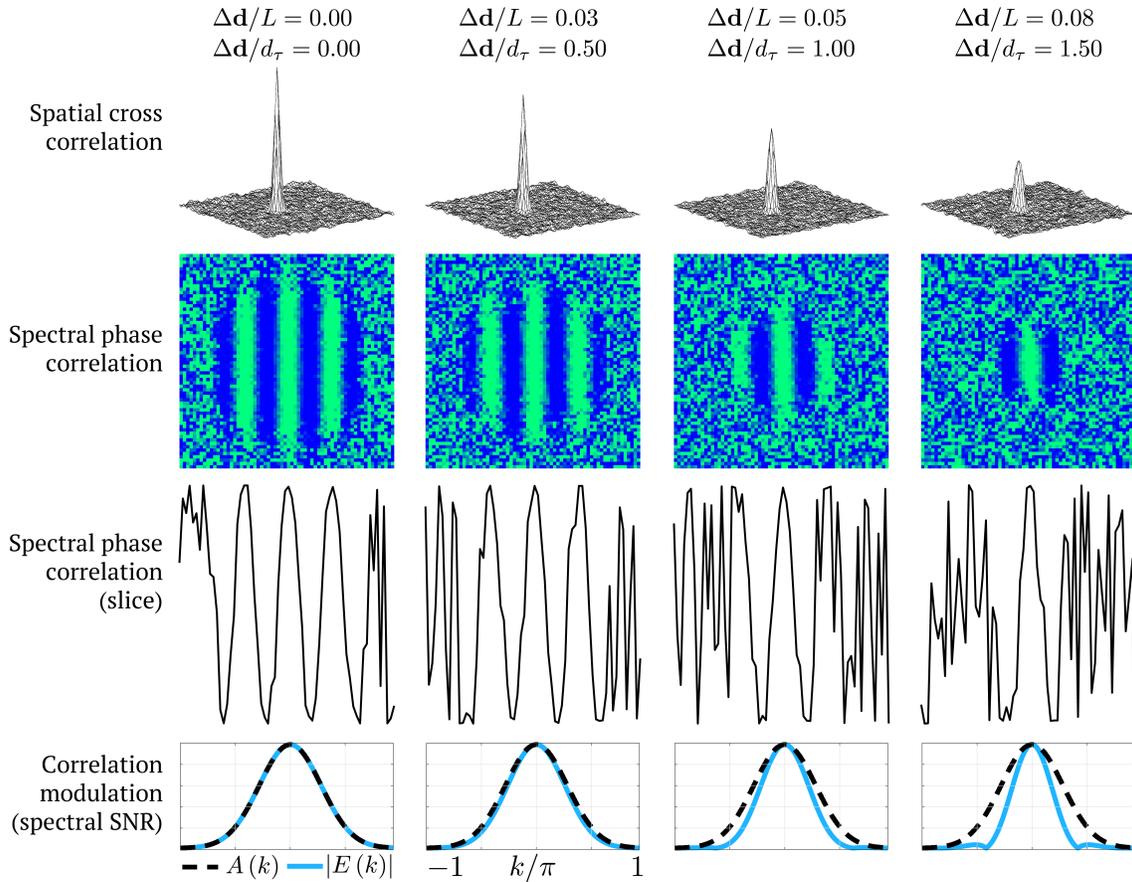


Figure 2.4: This figure depicts the effects of displacement gradients due to shearing on the cross correlation of PIV interrogation regions in the spatial domain (top row) and the spectral domain (second row from top). Third row: one-dimensional slices extracted from the centers of the spectral correlation planes in the second row. Bottom row: the theoretically derived modulation function on the spectral correlation (i.e., the spectral SNR) (blue) compared to the modulation due solely to the Gaussian shapes of the particle images (dashed black line) for each level of shearing (shown as 1-D slices from the 2-D correlation plane). The side-lobes visible in rightmost plot are due to the fact that the modulating function is proportional to the Fourier transform (FT) of the probability distribution of displacements. In shearing, this distribution is uniform, and its FT is the sinc function.

It is interesting to consider this behavior in the context of a long-standing design guideline for PIV analysis. The relationship between the particle size and the distribution of displacements (the “displacement distribution ratio”) has been shown to be a driver of PIV measurement error. Quoting Westerweel [1],

“...local gradients can be ignored provided that the variation of the displacement within the interrogation domain does not exceed the (mean) particle-image diameter.”

Revisiting this statement in the context of the spectral SNR provides an additional perspective about why it is correct. Equation (2.7) and the plots of spectral SNR in Figure 2.4 show that the

shape of the SNR is governed entirely by the particle shape in the absence of displacement gradients (i.e., when $p_d(\mathbf{x}) = \delta(\mathbf{x} - \bar{\mathbf{d}})$ and $P_d(\mathbf{k}) = 1$). For displacement distributions that are smaller than the particle diameter (e.g., $\Delta\mathbf{d}/d_\tau = 0.50$ in Figure 2.4), $E(\mathbf{k})$ deviates only slightly from $A(\mathbf{k})$; in this case, the displacement distribution barely affects the spectral SNR. However, when the breadth of the displacement distribution approaches the particle diameter (near $\Delta\mathbf{d}/d_\tau = 1.00$), $E(\mathbf{k})$ begins to deviate noticeably from $A(\mathbf{k})$. As $\Delta\mathbf{d}/d_\tau$ grows, $E(\mathbf{k})$ and $\text{SNR}(\mathbf{k})$ become increasingly dominated by $P_d(\mathbf{k})$, which consequently reduces the proportion of the correlation that is useful to the displacement estimate.

We can quantify this effect by defining an SNR ratio to describe the degree to which displacement gradients reduce the available signal compared to the situation of perfectly correlated pairs. We define the SNR ratio Φ_{SNR} as the ratio of the areas (or volumes in 2D) beneath $E(\mathbf{k})$ and $A(\mathbf{k})$, or

$$\Phi_{\text{SNR}} = \frac{\int_{\mathbf{k}} E(\mathbf{k}) d\mathbf{k}}{\int_{\mathbf{k}} A(\mathbf{k}) d\mathbf{k}} \quad (2.13)$$

The SNR ratio for uniform and Gaussian normal displacement distributions is illustrated in Figure 2.5 for a range of displacement distribution ratios.

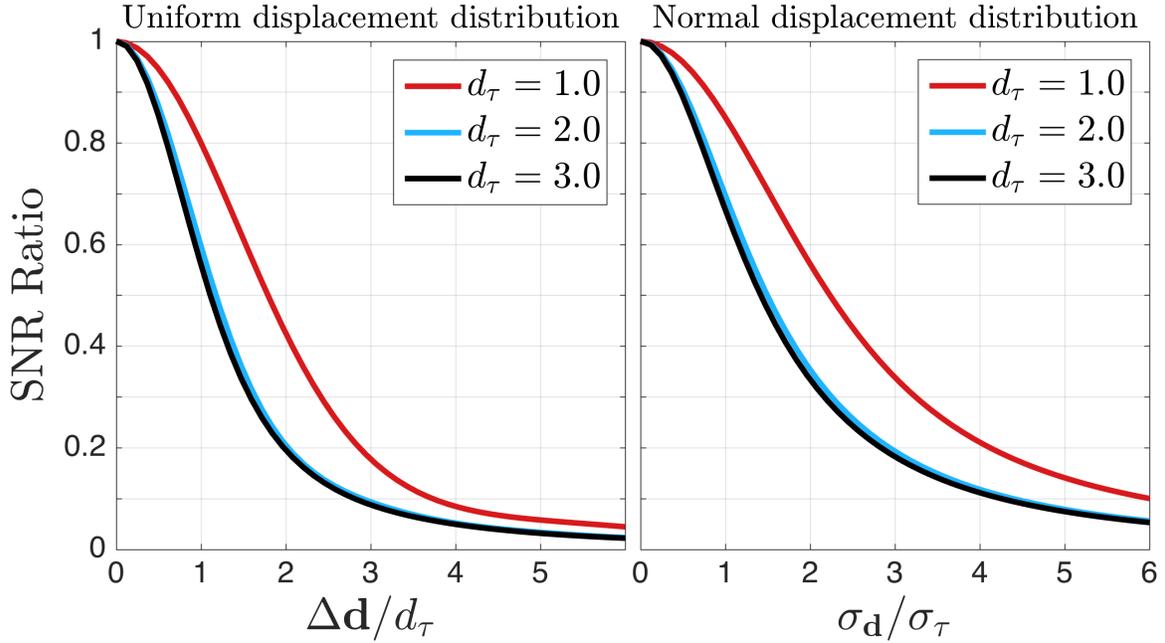


Figure 2.5: SNR ratio for uniform (left) and Gaussian normal displacement distributions. For both distributions, the SNR drops precipitously as the displacement distribution ratio increases. For the normal distribution, the horizontal axis could equivalently be written in terms of a displacement distribution “diameter ratio” d_d / d_τ , with $d_d = 4\sigma_d$.

From Equations (2.7)-(2.10), $\text{SNR}(\mathbf{k})$ cannot be written purely in terms of the ratio $\Delta\mathbf{d} / d_\tau$; the actual particle diameter matters. However, Φ rapidly converges to a single curve around $d_\tau \approx 3$ pixels, which is a typical particle size in PIV experiments.

From Figure 2.5, the second column of Figure 2.4 ($\Delta\mathbf{d} / d_\tau = 0.50$) corresponds to an SNR ratio of about $\Phi \approx 0.9$. In other words, for a displacement distribution ratio of 0.5, about 90% of the spectral SNR is available for estimating $\bar{\mathbf{d}}$ compared to the perfectly correlated situation ($\Delta\mathbf{d} / d_\tau = 0$). Increasing $\Delta\mathbf{d} / d_\tau$ leads to a rapid decrease in SNR ratio: the cases of $\Delta\mathbf{d} / d_\tau = 1.00$ and $\Delta\mathbf{d} / d_\tau = 1.50$ correspond to $\Phi \approx 0.6$ and $\Phi \approx 0.3$. Eckstein *et al.* demonstrated that these deleterious effects of the spectral SNR decay could be mitigated, and the errors of PIV measurements reduced, by suppressing the contribution of noise-dominated wave numbers to the displacement estimate. They achieved this by designing a filter in the spectral domain that preferentially weighted wave numbers according to their SNR as predicted by equation (2.6), subject to the assumption that the shape of the modulation of $\tilde{R}_D(\mathbf{k})$ was

due entirely (and was exactly equal) to that of the particle autocorrelation $A(\mathbf{k})$. These assumptions resulted in a filter that faithfully predicts the SNR envelope when the pair of patterns is perfectly correlated, in which case $P_d(\mathbf{k})=1$, and $E(\mathbf{k})=NA(\mathbf{k})$. But this situation represents only the best-case scenario for PIV cross correlations. As shown in the previous examples, even fairly narrow distributions of displacements can cause $E(\mathbf{k})$ to deviate significantly from the prediction of $A(\mathbf{k})$ alone. It is reasonable to suspect that an appropriate spectral filter, accounting not only for the particle shape but also the displacement distribution, would help to mitigate the accelerated decay in spectral SNR due to the presence of non-uniform displacements, and subsequently yield favorable performance compared to the unfiltered correlation (SCC) or filtering based only the particle shape (RPC).

Along these lines, Figure 2.6 shows an example of correlation planes derived a single pair of synthetic PIV IRs, subject to a Gaussian normal PDF of displacements, each filtered differently in the spectral domain. The leftmost correlation is unfiltered, and the middle is filtered using the particle shape only (the RPC filter). Meanwhile, the filter applied to the rightmost correlation was created using both the particle shape and the ground-truth displacement distribution that perfectly matched the simulation. Even though the displacement distribution is generally unknown *a priori*, Figure 2.6 is meant to illustrate anecdotally the effect that such a filter would have on the correlation plane when displacement gradients are not negligible.

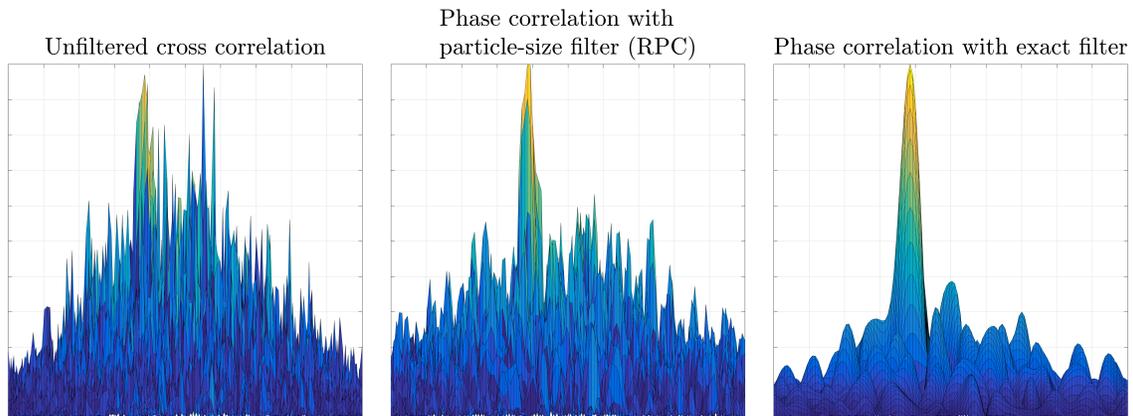


Figure 2.6: The spatial cross correlation of a pair of PIV interrogation regions with significant random motion, shown with no spectral filtering (left), and with spectral filter considering only the particle shape (middle) and both the particle shape and the displacement distribution (right).

Designing a spectral correlation filter that reliably estimates the true SNR requires estimates of both $P_a(\mathbf{k})$ and $A(\mathbf{k})$. While the particle shape can be reasonably estimated *a priori* (which was Eckstein's approach), the same is generally not true for the displacement distribution (rather, this is often the information that the experiment was performed to elucidate). In other words, properly filtering the correlation plane requires an estimate of the displacement distribution of the particles within the IR, and yet proper filtering is often required in order to reliably estimate the displacements in the first place.

In the following section, we will demonstrate how to overcome this apparent Catch-22 by introducing an algorithm to estimate the displacement distributions $P_a(\mathbf{k})$ and the particle autocorrelation shape $A(\mathbf{k})$ directly from PIV interrogation regions in the spectral domain. We subsequently use these estimates to create a spectral filter that accounts for both $A(\mathbf{k})$ and $P_a(\mathbf{k})$, and apply this filter to the spectral correlation prior to returning to the spatial domain and locating the correlation peak. We will show that this filter helps to suppress the appearance of spurious peaks and noise in the spatial correlation plane, and reduces errors in the final peak-detection-based displacement estimate. We call our algorithm the Adaptive Phase Correlation, or APC.

The Adaptive Phase Correlation (APC) Algorithm

In this section, we introduce an algorithm to design a spectral filter that extends Eckstein's RPC filter by accounting not only for the particle shape but also the distribution of displacements within the IRs, and does so without making assumptions about the particle size as Eckstein's algorithm does. Our APC algorithm is depicted graphically in Figure 2.7, and described subsequently.

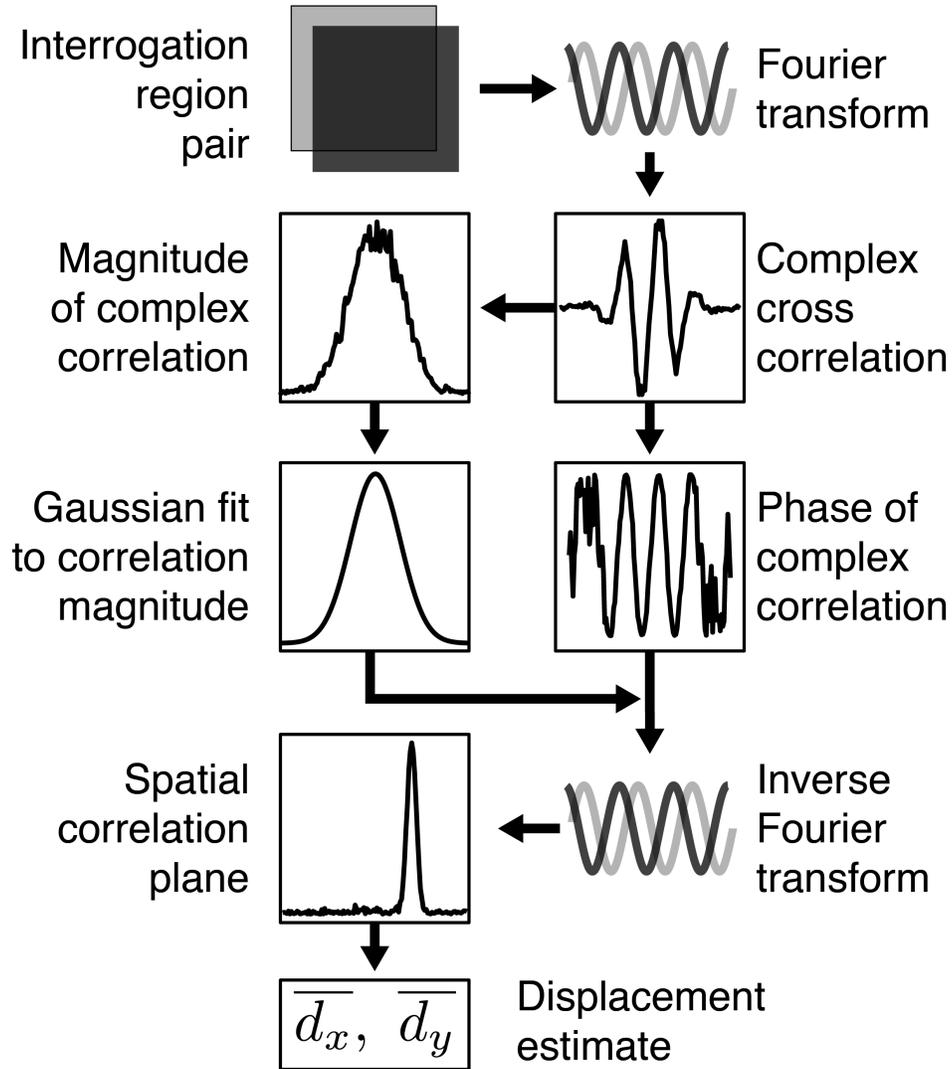


Figure 2.7: Our APC algorithm. $\overline{d_x}$ and $\overline{d_y}$ are the algorithm's estimates of the horizontal and vertical components of the average displacement of the particle image pattern.

From equations (2.5)–(2.7), it is evident that the SNR of the complex correlation is given by the magnitude of $\tilde{R}_D(\mathbf{k})$. When the interrogation regions under consideration are mean-subtracted prior to correlating, the correlation of the non-corresponding particles $\tilde{R}_C(\mathbf{k})$ manifests approximately as zero-mean, uniformly distributed “noise” in $\tilde{R}(\mathbf{k})$, in which case $\tilde{R}(\mathbf{k}) \approx \tilde{R}_D(\mathbf{k})$. Consequently, the magnitude of $\tilde{R}_D(\mathbf{k})$, and thus the spectral SNR, can be

estimated as the magnitude of $\tilde{R}(\mathbf{k})$. Given that $\tilde{R}(\mathbf{k})$ is complex valued, its magnitude can be calculated in the usual way of taking the quadratic sum of its real and imaginary parts:

$$\text{SNR}(\mathbf{k}) \approx |\tilde{R}(\mathbf{k})| = \sqrt{(\mathcal{R}\{\tilde{R}(\mathbf{k})\})^2 + (\mathcal{I}\{\tilde{R}(\mathbf{k})\})^2} \quad (2.14)$$

Equation (2.14) indicates that the shape of the correlation SNR is embedded in the correlation itself. All that is required to estimate the true SNR, then, is to calculate the magnitude of the complex correlation.

As shown in Figure 2.7, our algorithm first splits the spectral cross correlation into its real magnitude and complex phase (which has unity magnitude). Normally, the traditional phase-only correlation or the RPC algorithm would do away with the magnitude of the correlation. In contrast, we use this magnitude to estimate the underlying SNR. We do this by assuming that the true SNR is approximately Gaussian shaped, and subsequently fit a 2D elliptical Gaussian function to the 2D correlation magnitude in the least squares sense. Next, we use this best fit Gaussian function as our correlation spectral filter by multiplying it with the phase of the complex correlation (this is similar to Eckstein's method, with our best-fit Gaussian filter taking the place of his analytically calculated filter). Finally, we calculate the inverse FT of the filtered phase correlation, which yields the familiar spatial correlation plane. The mean displacement of the particle pattern is estimated in the usual way of locating the peak of this spatial correlation.

The key innovation in our algorithm is recognizing that the correlation magnitude contains information about the spectral SNR, and then using it to construct a spectral filter. Although we assumed in this work that the SNR was Gaussian-shaped, this need not be the case: if some other particle shapes or distributions can be assumed a priori (e.g., for experiments where the typical assumptions of PIV images do not hold), then the appropriate functions can be fit to $|\tilde{R}(\mathbf{k})|$ and the spectral filter calculated accordingly.

Figure 2.8 demonstrates the response of the spectral correlation and our APC filter to PIV image pairs subject to a mean displacement of five pixels with three different levels of normally distributed random motion. In these cases, the standard deviations of the random displacements in the horizontal direction were twice those in the vertical direction. These non-isotropic random motions are reflected by the asymmetric (elliptical) appearance of the correlation magnitude and their best-fit Gaussian functions. In other words, the correlation SNR

"shrinks" along each axis separately in proportion to the widths of the distributions for each component of displacement.

Because $E(\mathbf{k})$ depends in part on $P_{\tilde{\mathbf{q}}}(\mathbf{k})$, which is itself a statistical quantity, the estimation of the true SNR from R naturally becomes more accurate as the number of corresponding particle pairs contributing to the correlation increases. We therefore expect our algorithm to complement the well-established ensemble correlation of Meinhart *et al.* [13]. In the next section, we support this claim by demonstrating quantitatively that the performance of our algorithm exceeds both the standard cross correlation (SCC) and the RPC correlation in estimating flow velocities in synthetic computer generated PIV images subject to significant random displacements and high levels of image noise.

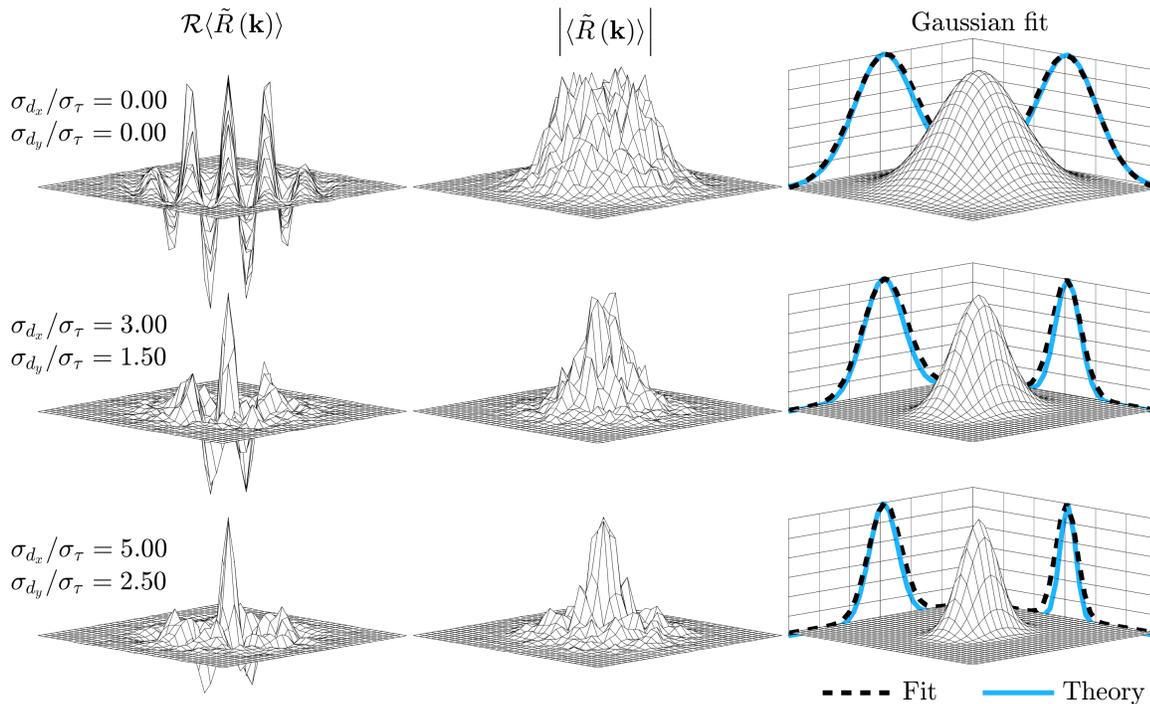


Figure 2.8: Response of spectral planes and APC filter to increasing degrees of normally distributed random particle displacements (32x32 pixel regions; average of 20). Left and middle columns are the real part and magnitudes of the spectral correlation; right column is the Gaussian fit to the spectral correlation magnitude; this serves as the spectral filter in our algorithm. The projections in the third column compare the measured fit with theory along the zero-frequency coordinates. In these simulations, diffusions were non-isotropic: the standard deviation of velocity in the horizontal direction was twice that in the vertical direction.

Performance Assessment

Error Analysis: Poiseuille flow (synthetic images)

To estimate the accuracy of our APC algorithm, we generated synthetic images of PIV particles in Poiseuille flow, with an additional fluctuating displacement applied to each component of the displacement field. For each pair of images, particles were randomly seeded throughout the domain with a particle image concentration of 0.02 particles per pixel in the first snapshot, and then displaced in the second snapshot according to equation (2.15):

$$\begin{aligned}u &= u_o \left(1 - \frac{y^2}{h^2} \right) + u' \\v &= v' \\w &= w'\end{aligned}\tag{2.15}$$

Here, u and v are the horizontal and vertical in-plane displacements, and w is the out-of-plane displacement; u_o is the maximum inter-frame displacement, which we chose to be 10 pixels per frame; h is the channel height, and y the vertical coordinate (perpendicular to the flow direction). The fluctuating displacement components u' , v' , and w' were each drawn from a Gaussian normal distribution of equal variance σ_d^2 . We considered three different levels of isotropic random displacements, with $\sigma_d / \sigma_\tau = 2.0$, 4.0 , and 6.0 ; these cases correspond to theoretical SNR ratios (from equation (2.13)) of about 0.5, 0.1, and 0.05.

Synthetic image generation

For each of these three cases, we generated one thousand pairs of images whose dimensions were 2048 x 2048 pixels. Particle images were rendered according to the theory of Olsen & Adrian [21], and specific details of the implementation are described by Brady et al. [22]. In our simulations, the diameter of each particle was drawn from a Gaussian normal distribution with a mean of 3.0 pixels and a standard deviation of 1.0 pixel, where the diameter $d_\tau = 4\sigma_\tau$ (σ_τ refers to the standard deviation of the Gaussian particle shape in this analysis). The brightness of each particle was proportional to the intensity of a simulated light sheet whose intensity was uniform in the horizontal and vertical directions and followed a Gaussian profile in the “out-of-plane” direction, with a standard deviation equal to the length of about 16 pixels.

Once rendered, each noise-free image was contaminated with uncorrelated Gaussian additive noise whose standard deviation was 15% of the maximum brightness of any particle in the noise-free image. Finally, the dynamic range of each image was truncated from 64 bits to 8 bits such that the maximum intensity of any particle was 95% of the saturation brightness.

PIV Processing

The purpose of our error analysis was to quantify the effect on PIV displacement estimates of our filtering algorithm alone. Therefore, we applied no iterative or multi-pass schemes for our comparisons with other methods. We compared the performance of our APC algorithm to the standard cross correlation (SCC) and Eckstein's robust phase correlation (RPC), each implemented in the ensemble sense. The processed images, and all parameters besides the correlation algorithms themselves, were identical among the three methods. The dimensions of the interrogation regions were 128×128 pixels. A Gaussian apodization window was applied to each region to reduce aliasing in the spectral domain, resulting in an effective interrogation resolution of 64×64 pixels [23]. The grid spacing was 64×64 pixels, and each interrogation region was centered about its grid point (specifically, the grid point was located at $x = L/2$, $x \in [1, L]$ within the square IR of side length L).

For SCC and RPC, the ensemble correlation was carried out in the normal way of adding together the spatial correlation planes [13]. For APC, complex spectral correlation planes were averaged, and the APC filter was calculated anew after each additional pair of IRs was added to the ensemble. Displacements were estimated each time the ensemble correlation was incremented, and sub-pixel precision was achieved by fitting a Gaussian function (three point fit [24]) to the values surrounding the largest peak in the spatial correlation. In order to compare the performance of the correlation algorithms themselves, no outlier detection or replacement was performed, and the displacements fields were not smoothed.

Quantification of error

We compared the performances of SCC, RPC, and APC on the basis of the mean and standard deviation of displacement error magnitude across all interrogation regions for each ensemble increment. We quantified error as the quadratic sum of the differences between the PIV-measured displacements in the x and y directions and the ground-truth solution of equation (2.15) evaluated at the locations of the PIV grid points. In addition to comparing error of the

different algorithms at each ensemble iteration, we also considered the rates of convergence of their estimated displacement fields. The metric used to quantify convergence was the average magnitude of the change of each component of displacement between ensemble increments. The convergence criterion was that both components changed by less than 0.05 pixels when averaged over 25 ensemble pairs; we chose 0.05 pixels because that value is around the achievable lower bound of the RMS error of PIV measurements [25-27].

Results: Synthetic images, Poiseuille flow

Figure 2.9 compares the errors of estimated displacements produced by the SCC, RPC, and APC algorithms for our simulations. The behaviors of these algorithms' errors exhibit four salient features. First, for all algorithms, greater levels of random displacements increase the error of displacement estimates (this is not surprising). Second, in each case, APC ultimately reaches the lowest displacement error, followed by RPC and then SCC. Third, the difference in performance between the three algorithms grows as the level of random displacement increases. Finally, APC produces estimates of lower error using fewer image pairs than the other methods.

Consider the results for the three different cases shown in Figure 2.9. At the lowest level of random displacements we considered ($\sigma_d / \sigma_r = 2.0$; $\sigma_d = 1.5$ pixels per frame; thin solid lines), the reduction in error of APC compared to RPC is on the order of 0.01 - 0.02 pixels early in the ensemble, and nearly diminishes as the ensemble size increases. Both of these algorithms outperform SCC by about 0.05 pixels of error for this case. These modest differences in performance suggest two conclusions about this case.

First, the clear distinction in performance between SCC (the unfiltered correlation) and the two filtered methods (RPC and APC) implies that attenuating certain wave spectral wave numbers does indeed affect the spatial cross correlation in a way that reduces error in the peak-detection-based displacement estimates. This is consistent with the results reported by Eckstein *et al.* when they first introduced the RPC filter [15]. Second, the closeness in performance between APC and RPC suggests that the theoretically derived RPC filter, which is based only on the size and shape of the tracer particles, is sufficiently close to the actual measured spectral SNR (the APC filter) that the two filters ultimately have similar effects on the spatial cross correlation for this level of random displacements. Taken together, these two conclusions imply a third: in these images, the particle shape is the primary driver of the spectral SNR.

In contrast to this situation, there is a clear distinction between the algorithms' performances for the other two cases, whose random displacement levels were twice and three times greater than that of the first. For $\sigma_d / \sigma_\tau = 4.0$ ($\sigma_d = 3.0$ pixels per frame; dashed lines in Figure 2.9), the average errors of RPC and SCC were at best about 0.1 and 0.25 pixels greater than for APC. At $\sigma_d / \sigma_\tau = 6.0$ ($\sigma_d = 4.5$ pixels per frame), the smallest differences increased to 0.5 pixels and 0.9 pixels respectively.

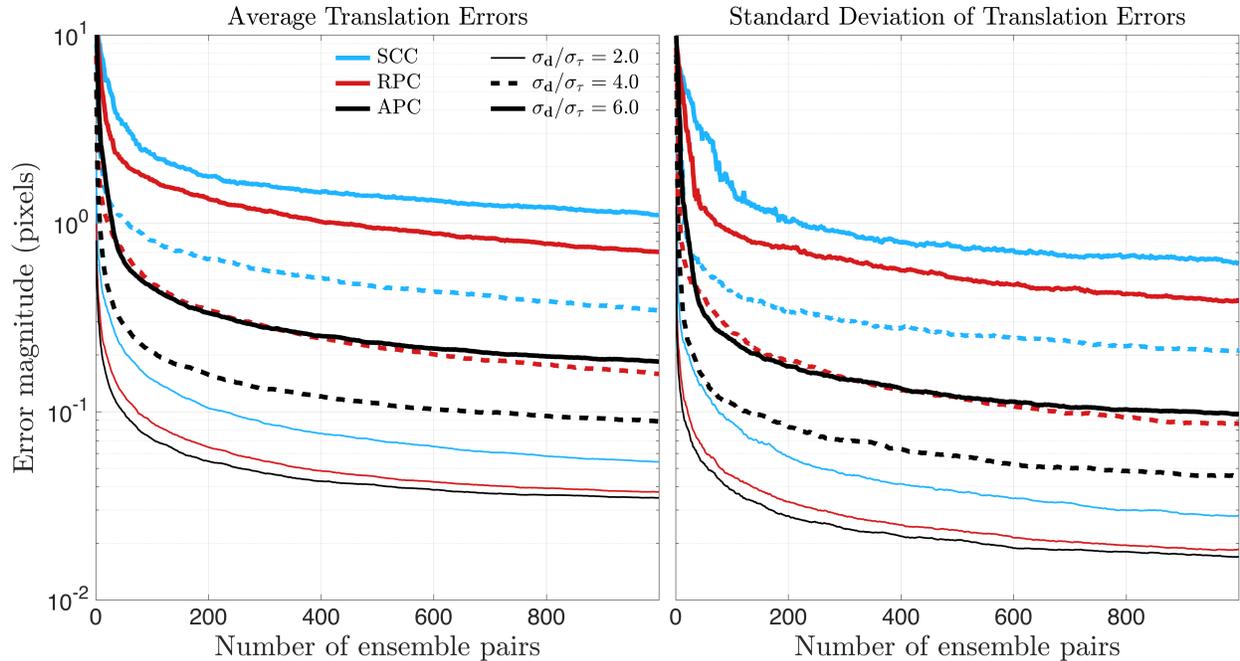


Figure 2.9: Displacement errors for synthetic images of Poiseuille flow with random displacements.

In addition to estimating the average displacements of particle patterns with lower error than SCC and RPC, our APC algorithm does so using fewer ensemble image pairs. As shown in Figure 2.10, at $\sigma_d / \sigma_\tau = 6.0$ (the highest level of random displacements tested; $\sigma_d = 4.5$ pixels per frame), the ensemble-correlation displacement estimates of SCC, RPC, and APC converge (according to our definition of convergence) after 599, 332, and 63 pairs of images, respectively. In these cases, APC accelerates the convergence of displacement estimates by factors of about 5 compared to RPC and 10 compared to SCC. The rates of convergence of the algorithms at all three levels of random displacement tested are summarized in Table 2.1.

Convergence of displacement estimates
 Poiseuille flow synthetic images, $\sigma_d/\sigma_\tau = 6.0$
 Convergence: $\mathbf{R} < 0.05$

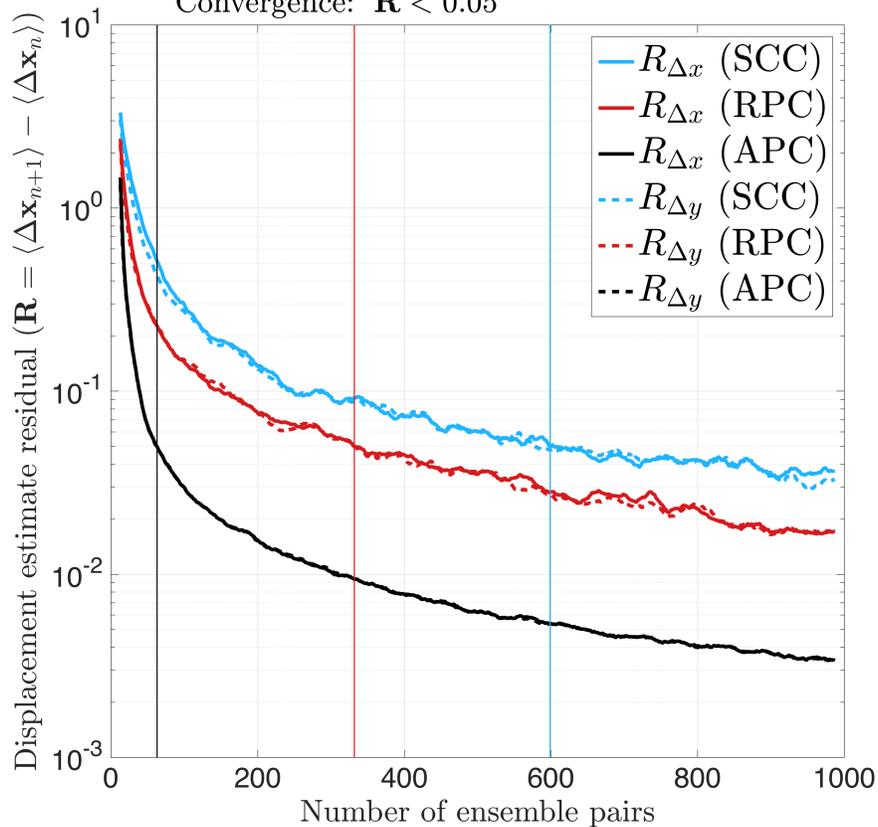


Figure 2.10: Change in displacement estimates between ensemble increments (i.e., residuals) throughout the ensemble correlation for SCC, RPC, and our APC correlation methods for synthetic PIV images of Poiseuille flow with the highest level of random displacements tested ($\sigma_d/\sigma_\tau = 6.0$). Vertical lines indicate convergence of each method according to our definition; the colors of the convergence lines correspond to those of the error estimate residuals.

		At convergence of method...		
Error of method...	σ_d/σ_τ	2.0		
		<i>APC</i>	<i>RPC</i>	<i>SCC</i>
	<i>APC</i>	0.18 ± 0.10	0.17 ± 0.09	0.13 ± 0.07
	<i>RPC</i>	0.23 ± 0.12	0.22 ± 0.11	0.16 ± 0.09
	<i>SCC</i>	0.40 ± 0.23	0.39 ± 0.23	0.30 ± 0.18
Error of method...	σ_d/σ_τ	4.0		
		<i>APC</i>	<i>RPC</i>	<i>SCC</i>
	<i>APC</i>	0.38 ± 0.20	0.24 ± 0.12	0.17 ± 0.09
	<i>RPC</i>	0.87 ± 0.48	0.57 ± 0.33	0.37 ± 0.19
	<i>SCC</i>	1.21 ± 0.67	0.92 ± 0.51	0.68 ± 0.35
Error of method...	σ_d/σ_τ	6.0		
		<i>APC</i>	<i>RPC</i>	<i>SCC</i>
	<i>APC</i>	0.56 ± 0.30	0.27 ± 0.14	0.22 ± 0.11
	<i>RPC</i>	1.94 ± 1.08	1.12 ± 0.61	0.88 ± 0.47
	<i>SCC</i>	3.01 ± 2.79	1.55 ± 0.85	1.33 ± 0.72

Table 2.1: Errors of different correlation methods at convergence of each method. Units are in pixels.

Figure 2.11 shows vector fields and displacement profiles calculated using the different correlation methods as they each converge, as well as those for an arbitrarily-selected point comparatively early in the ensemble (ten image pairs), for the highest level of random displacements tested ($\sigma_d/\sigma_\tau = 5.0$; $\sigma_d = 4.5$ pixels per frame). In this figure, the mean displacement profiles were calculated as the average of all the measured horizontal displacements along the horizontal direction (the "flow" direction), and the width of each shaded areas is two standard deviations about this mean. The vector plots were sub-sampled for ease of visualization. For both plots, the raw measured displacements are shown, without any outlier rejection or replacement. Dashed lines show the ground-truth mean displacement profile imposed to create the synthetic images, given by equation (2.15).

The differences between the measured displacement fields are striking. Within a few (ten) ensemble pairs, our APC method appears to approximate the true displacement field with reasonable fidelity. Meanwhile, the measurements of SCC and RPC both vary widely from the true solution. As each method converges, its measured displacement profile more closely

approaches the true solution; however, none appears to match the APC-estimated displacement profile at any level of convergence. These results demonstrate that our APC algorithm provides significant quantitative and qualitative improvements in measured displacement fields in situations where particle pattern displacements deviate significantly from uniformity on the scale of individual interrogation regions.

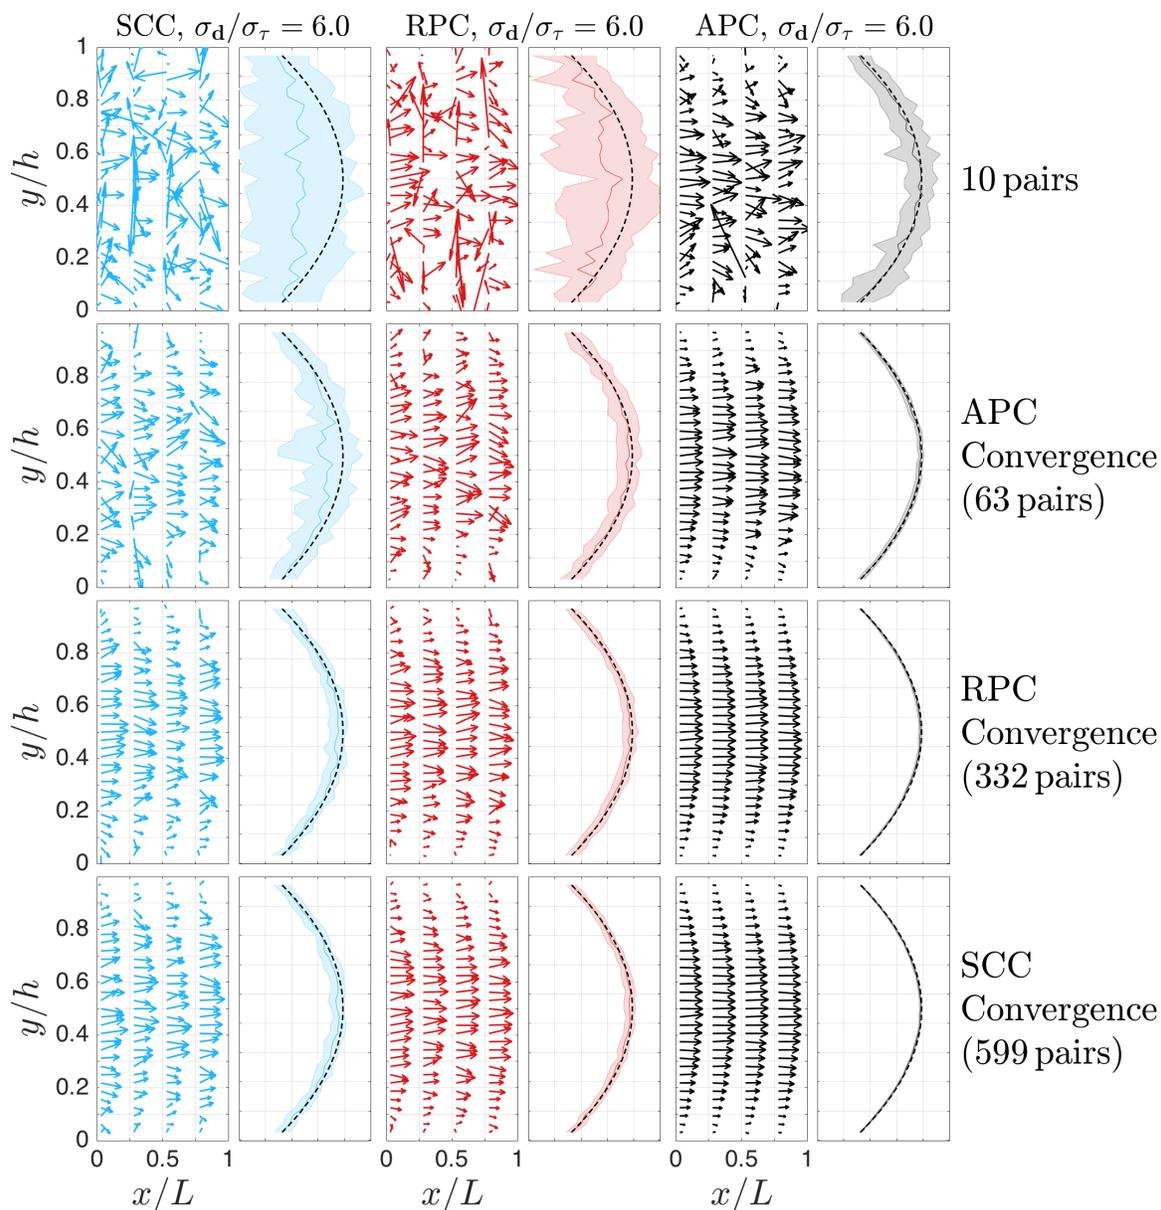


Figure 2.11: Representative velocity vectors and velocity profiles for synthetic Poiseuille flow with significant random displacements using the standard cross correlation (SCC; left), Eckstein's RPC filter (middle), and our APC algorithm (right).

Experimental Images (PIV Challenge 2014 Case A)

Procedure

We tested our algorithm on experimental PIV images from the 2014 PIV challenge (case A). The dataset contained 600 pairs of double-pulsed images of flow through a micro-fluidic nozzle. Details of the experiment are found in Kähler *et al.* [28]. This flow field exhibited strong velocity gradients (it contained a jet), a large dynamic range of velocities, and image artifacts where intensities "leaked" between frames in each pair. The "intensity leakage" artifact tended to strongly bias PIV results toward zero-displacement measurements because it caused a strong auto-correlation peak in the cross correlation plane. Most groups in the PIV challenge performed image processing prior to their PIV analysis to mitigate these effects. These "pre-processing" algorithms were largely ad-hoc and relied on the judgments and experiences of the researchers, and as such, they varied widely between groups.

For our analysis, we compared our APC algorithm with SCC and RPC correlations using these same images without any image pre-processing (the "raw" images), as well as on images that were pre-processed using the algorithm described in Kähler *et al.* [28] by the group from Los Alamos National Laboratory. For both cases, we used an ensemble correlation of all 600 pairs, and also a shorter ensemble correlation of 10 pairs. We used an iterative algorithm of five PIV passes, with image deformation [9] performed between passes. Before deforming each image, outlying vectors were identified and replaced using universal outlier detection [29] and fields were smoothed using a Gaussian kernel of 7×7 grid points and standard deviation of one grid point. The sizes of the interrogation regions were 128×64 pixels (width x height) on the first pass, and 32×32 pixels on the final pass.

Results: PIV Challenge 2014, case A

Figure 2.12 shows the velocity fields produced by the final iterations of SCC, RPC, and APC ensemble correlation algorithms for both the raw (unprocessed) and pre-processed images. It is immediately apparent that performing image pre-processing significantly improved the results for all algorithms, to the extent that neglecting to do so yielded obviously unphysical results over much of the flow field for SCC and RPC correlations. The final pass of SCC yielded near-zero displacement measurements in large portions of the nozzle's throat and in the jet, and RPC did so over nearly the entire flow field. Our APC algorithm appeared significantly more robust against the features of the raw images that caused the other algorithms to fail.

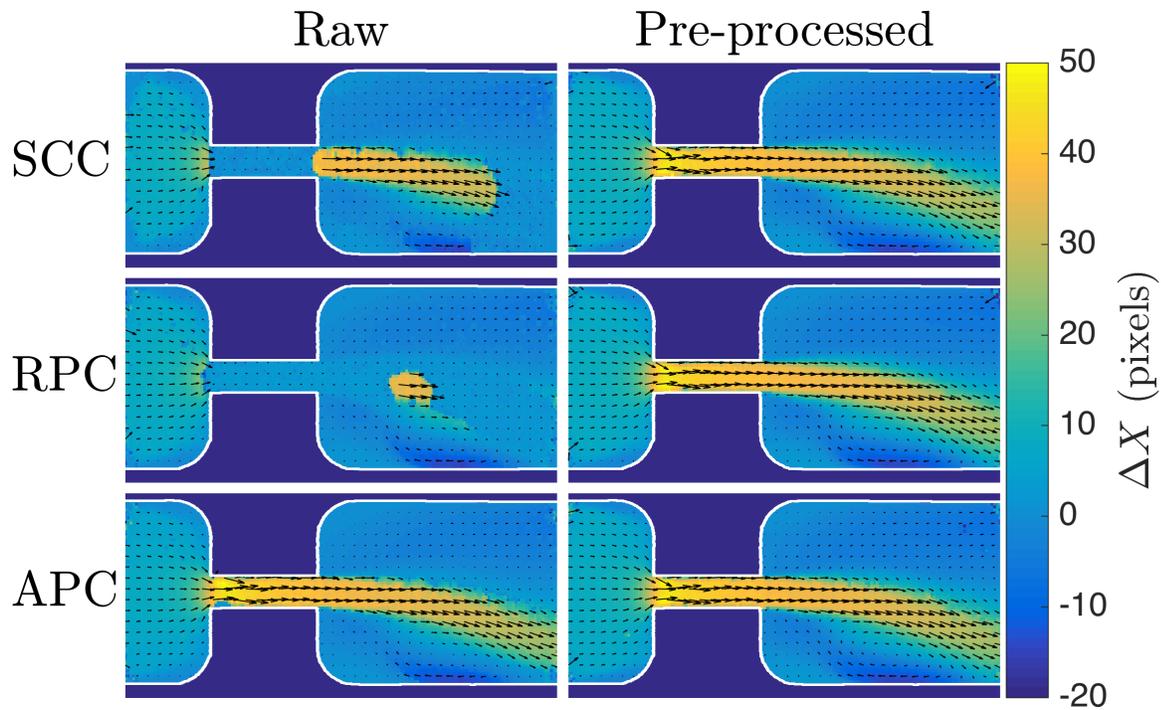


Figure 2.12: Displacement fields for the final iteration of ensemble PIV processing of real PIV images of flow through a microchannel (PIV Challenge 2014, case A [28]). Vectors are shown without any outlier replacement or smoothing, but both of these operations were performed between PIV iterations prior to performing image deformation.

To provide insight into the reasons underlying these differences, Figure 2.13 shows vector fields of all three algorithms for the first processing iteration, prior to any vector validation or smoothing. In iterative PIV methods like image deformation, the first iteration must resolve the displacement field with sufficient accuracy in order for the subsequent iterations to succeed. Figure 2.13 shows a clear distinction between the results of the first iteration of APC versus the other correlations. Both SCC and RPC exhibited large contiguous regions where the estimated displacements were nearly zero. In contrast, in most of the flow, APC captured the underlying field with sufficient fidelity that the iterative image deformation algorithm could proceed successfully after outlier detection and replacement. The reason for the strong bias toward zero of SCC and RPC, but not APC, is illustrated by the representative spatial cross correlation planes shown in the rightmost column of Figure 2.13, which were taken from a grid point near the shear layer of the jet. Both the SCC and RPC planes contained two prominent peaks: one corresponding to a physically reasonable pattern displacement, located away from the center of the plane, and a second, close to the center, representing a near-zero displacement. This "zero peak" was present in most of the correlation planes we inspected. The displacement peaks appears broadened, which is consistent with the large displacement gradients in the jet's shear

layer. The near-zero vectors in the SCC and RPC fields correspond to planes where the height of the zero-peak exceeded that of the displacement peak. This occurred most prominently in the nozzle throat and near the shear layer of the jet, where displacement gradients were significant. However, the APC plane showed only a single, wider peak, representing a physically reasonable pattern displacement.

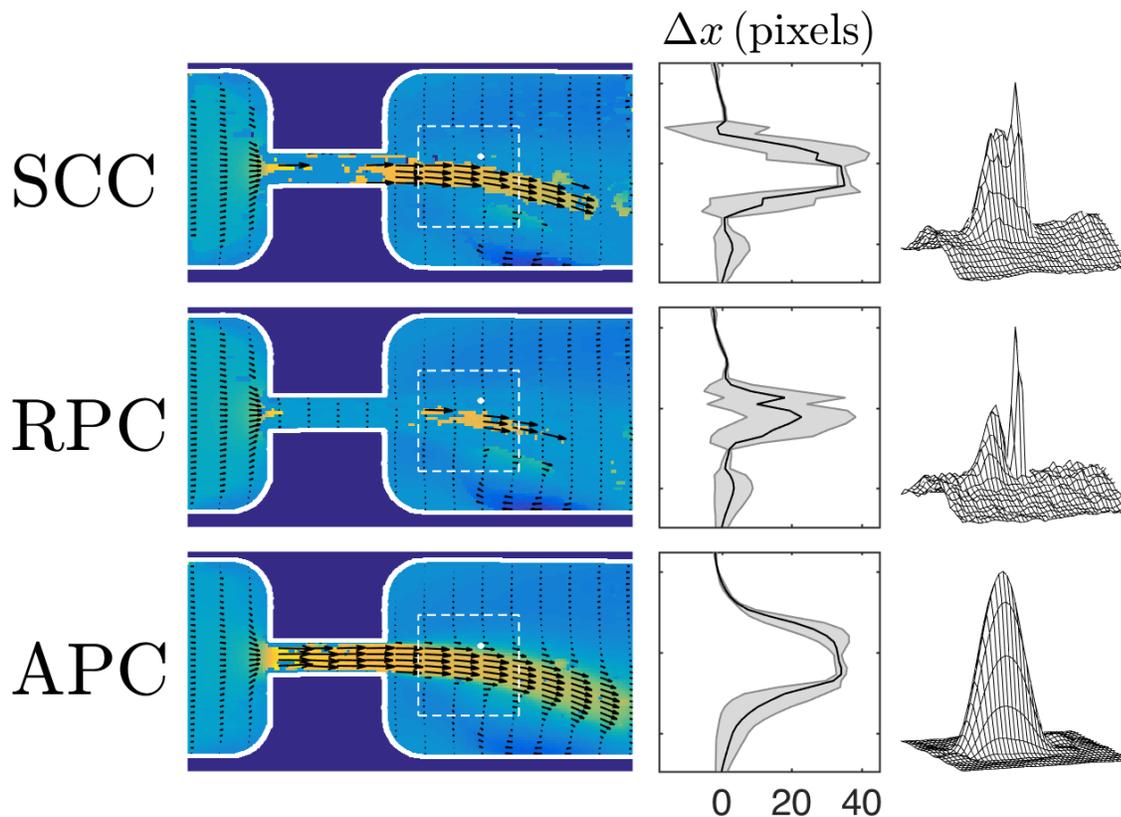


Figure 2.13: Left: Single-pass ensemble displacement fields for raw images (no image pre-processing) shown without any outlier replacement or smoothing (PIV Challenge 2014, case A). Middle: Profiles of horizontal displacement, averaged in the horizontal direction, for vectors within the dashed square region. Right: Ensemble cross correlation planes in the spatial domain, extracted at the location of the white dot within the vector fields. Color scale is the same as for Figure 2.12.

Conclusions

The main insight of this paper is to explain theoretically why the noise characteristics of PIV cross correlations behave as they do, and into why previously proposed filters work [5, 15]. Our results show that accounting for these characteristics and addressing them appropriately (via spectral filtering) can significantly increase the accuracy and reliability of PIV algorithms, and that failing to do so can result in failed measurements or decreased measurement accuracy.

Our algorithm provides a means for creating a spectral filter that responds to the characteristics of the specific images under consideration, rather than picking one heuristically (e.g., as in Wernet [5]), and without making assumptions about the sizes of the particles in the images (e.g., as in Eckstein *et al.* [15]). Moreover, our filter responds to local variations of image and flow characteristics (like spatially-varying displacement gradients), rather than supposing that a single filter can faithfully represent the SNR characteristics over the entire imaged flow field.

Our APC algorithm is not without limitations. As described by the theory we presented, the spectral SNR of a PIV cross correlation is governed in part by the probability distribution of the particles' displacements. The number of particle images present in a given interrogation region is often insufficient for the distribution of their displacements to faithfully represent that of their parent population. Therefore, our algorithm's ability to estimate the correlation SNR (i.e. the filter) improves with the number of particles under consideration, which can be artificially increased by considering an ensemble of multiple pairs of images. Our APC algorithm is therefore well suited for situations where the ensemble correlation is appropriate (steady or quasi-steady flows). Future developments may overcome this limitation.

References

1. Westerweel, J., *On velocity gradients in PIV interrogation*. Experiments in Fluids, 2007. **44**(5): p. 831-842.
2. Keane, R.D. and R.J. Adrian, *Theory of cross-correlation analysis of piv images*. Applied Scientific Research, 1992. **49**(3): p. 191-215.
3. Xue, Z., J.J. Charonko, and P.P. Vlachos, *Particle image velocimetry correlation signal-to-noise ratio metrics and measurement uncertainty quantification*. Measurement Science and Technology, 2014. **25**(11): p. 115301.
4. Charonko, J.J. and P.P. Vlachos, *Estimation of uncertainty bounds for individual particle image velocimetry measurements from cross-correlation peak ratio*. Measurement Science and Technology, 2013. **24**(6): p. 065301.
5. Wernet, M.P., *Symmetric phase only filtering: a new paradigm for DPIV data processing*. Measurement Science & Technology, 2005. **16**(3): p. 601-618.
6. Andrea, S., et al., *Collaborative framework for PIV uncertainty quantification: comparative assessment of methods*. Measurement Science and Technology, 2015. **26**(7): p. 074004.

7. Timmins, B.H., et al., *A method for automatic estimation of instantaneous local uncertainty in particle image velocimetry measurements*. Experiments in Fluids, 2012. **53**(4): p. 1133-1147.
8. Boomsma, A., et al., *A comparative experimental evaluation of uncertainty estimation methods for two-component PIV*. Measurement Science and Technology, 2016. **27**(9): p. 094006.
9. Scarano, F., *Iterative image deformation methods in PIV*. Measurement Science & Technology, 2002. **13**(1): p. R1-R19.
10. Scarano, F. and M.L. Riethmuller, *Advances in iterative multigrid PIV image processing*. Experiments in Fluids, 2000. **29**: p. S51-S60.
11. Gui, L. and S.T. Wereley, *A correlation-based continuous window-shift technique to reduce the peak-locking effect in digital PIV image evaluation*. Experiments in Fluids, 2002. **32**(4): p. 506-517.
12. Westerweel, J., D. Dabiri, and M. Gharib, *The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings*. Experiments in Fluids, 1997. **23**(1): p. 20-28.
13. Meinhart, C.D., S.T. Wereley, and J.G. Santiago, *A piv algorithm for estimating time-averaged velocity fields*. Journal of Fluids Engineering, 2000. **122**(2): p. 285-289.
14. Eckstein, A.C., J. Charonko, and P. Vlachos, *Phase correlation processing for DPIV measurements*. Experiments in Fluids, 2008. **45**(3): p. 485-500.
15. Eckstein, A. and P.P. Vlachos, *Digital particle image velocimetry (DPIV) robust phase correlation*. Measurement Science & Technology, 2009. **20**(5).
16. Anuta, P.E., *Spatial registration of multispectral and multitemporal digital imagery using fast fourier transform techniques*. IEEE Transactions on Geoscience Electronics, 1970. **8**(4): p. 353-368.
17. Kim, S.P. and W.Y. Su. *Subpixel accuracy image registration by spectrum cancellation*. in *1993 IEEE International Conference on Acoustics, Speech, and Signal Processing*. 1993.
18. Stone, H.S., et al., *A fast direct Fourier-based algorithm for subpixel registration of images*. IEEE Transactions on Geoscience and Remote Sensing, 2001. **39**(10): p. 2235-2243.
19. Jun, B., H. , et al., *Nanoparticle flow velocimetry with image phase correlation for confocal laser scanning microscopy*. Measurement Science and Technology, 2016. **27**(10): p. 104003.
20. Theunissen, R., *Theoretical analysis of direct and phase-filtered cross-correlation response to a sinusoidal displacement for PIV image processing*. Measurement Science and Technology, 2012. **23**(6): p. 065302.

21. Olsen, M.G. and R.J. Adrian, *Out-of-focus effects on particle image visibility and correlation in microscopic particle image velocimetry*. Experiments in Fluids, 2000. **29**(1): p. S166-S174.
22. Brady, M.R., S.G. Raben, and P.P. Vlachos, *Methods for digital particle image sizing (DPIS): comparisons and improvements*. Flow Measurement and Instrumentation, 2009. **20**(6): p. 207-219.
23. Eckstein, A. and P.P. Vlachos, *Assessment of advanced windowing techniques for digital particle image velocimetry (DPIV)*. Measurement Science & Technology, 2009. **20**(7).
24. Westerweel, J., *Fundamentals of digital particle image velocimetry*. Measurement Science and Technology, 1997. **8**(12): p. 1379.
25. Huang, H., D. Dabiri, and M. Gharib, *On errors of digital particle image velocimetry*. Measurement Science & Technology, 1997. **8**(12): p. 1427-1440.
26. Prasad, A.K., et al., *Effect of resolution on the speed and accuracy of particle image velocimetry interrogation*. Experiments in Fluids, 1992. **13**(2): p. 105-116.
27. Westerweel, J., *Theoretical analysis of the measurement precision in particle image velocimetry*. Experiments in Fluids, 2000. **29**(1): p. S003-S012.
28. Kähler, C.J., et al., *Main results of the 4th International PIV Challenge*. Experiments in Fluids, 2016. **57**(6): p. 97.
29. Westerweel, J. and F. Scarano, *Universal outlier detection for PIV data*. Experiments in Fluids, 2005. **39**(6): p. 1096-1100.

Chapter 3: High-Speed Flow Visualization Reveals Rapid Flow Reversals in Insect Hearts

Abstract

Insects are by far the most numerous and successful animals on planet earth. Most insects have hearts, but the mechanism by which these tube-shaped organs pump blood is unknown. Several recent studies [1-3] have shown that other tubular hearts, once presumed to act as peristaltic pumps, might actually produce flow by other means. These developments, together with at least one report of unexplained blood flow patterns in insects [4], motivate the pursuit of a deeper understanding of how the insect heart works. To this end, we measured blood flows within the hearts of living grasshoppers (*Schistocerca americana*) using synchrotron X-ray particle image velocimetry (PIV), and cockroaches (*Periplaneta americana*) by injecting fluorescent flow-tracer dye. We discovered previously unknown flow kinematics in both species – namely, that their heart flows rapidly reverse direction multiple times per second, on the time scale of individual heart beats. Our results were unexpected according to the current prevailing picture of the heart as a producer of directional flows, and may support the hypothesis that these insects' hearts produce flow by some means other than peristalsis.

Introduction

Insects have hearts that are different in structure and function than vertebrate hearts, but serve the analogous function of moving hemolymph (insect blood) within the body. Because insects deliver oxygen directly to their tissues using a network of air-filled tubes (tracheae) rather than indirectly through the blood, their circulatory systems are structurally relatively simple compared to those of vertebrates. In contrast to the closed-loop circulatory systems of vertebrates, which consist of a central pump (the heart) that drives blood through an interconnected network of veins and arteries, insects possess an "open" circulatory system made up of numerous pumps and vessels that are generally not connected to one another by vascular tissue. The largest circulatory organ is the dorsal vessel, which is a long, slender tube aligned with the anteroposterior axis ("head to tail"), situated along the dorsal surface, and extending nearly the length of the animal from the posterior end of the abdomen to the head [5]. The posterior portion of the dorsal vessel is muscularized, and is referred to as the "heart." The heart terminates near the intersection of the abdomen and the thorax, and the vessel continues

through the thorax and into the head as an un-muscularized tube called the aorta [6]. Radially-oriented groups of muscles called alary muscles attach the outer diameter of the heart to the surrounding tissues, but their specific role (if any) in flow production is unclear. Hemolymph flows into the heart through valves called ostia, which are positioned laterally in pairs along its length, and can flow in or out of the distal ends of the vessel [7].

Many insects across multiple phylogenetic orders have dorsal vessels [5, 7], which suggests that the organ is basal in origin. Despite a relatively mature body of literature describing its morphology [6-9], very little is known about the mechanism by which it produces flow. The heart produces wave-like contractions that travel along its length by shortening muscles that are wound helically around its diameter [6]. These contractions undoubtedly play a role in flow production, and much of the literature describes them as "peristaltic waves" [5, 7, 10-14], but the hypothesis that the heart is a peristaltic pump has never been tested. Only recently (in mosquitoes) was the hemolymph conclusively observed to flow in the same direction that the contractile waves traveled [6]. In some species, the contractions reverse direction periodically, switching from anterograde (toward the head) to retrograde (toward the posterior), with periods of unidirectional flow of about 5-15 seconds in mosquitoes [6], 15-30 seconds in other flies [9, 15], 1-3 minutes in moths [14], and 1-6 minutes in butterflies [16].

The majority of previous work has reported on the heart's superficial behavior such as the frequency and direction of contractions [9, 14, 16, 17], but few researchers have directly observed the flow of hemolymph. Glenn *et al.* used fluorescent flow tracer particles to observe the flow of hemolymph through the hearts of mosquitoes (*Anopheles gambiae*), and discovered that those animals' hearts produced net directional flow that reversed direction about five times per minute [6]. Lee *et al.* [4] used phase contrast synchrotron X-ray imaging (PCXI) to record the movements of individual flow tracers (air-filled bubbles) into and through the hearts of grasshoppers (*Schistocerca americana*). Although the small number of tracers limited their analysis measuring a few (< 10) localized flow patterns, they reported that the flow kinematics appeared to be more complicated than would be expected according to the presumption of peristaltic pumping – namely, that the tracers appeared to meander and oscillate back and forth rather than traveling unidirectionally through the lumen, which they would be expected to do according to the definition of peristaltic flow [18]. Those unexpected observations naturally motivate the question of whether the insect heart might produce flow by some other mechanism besides peristalsis.

The objective of this present work is to elucidate details about the flow of hemolymph within the insect heart as a first step to understanding how it produces flow. To this end, we extended the method of Lee *et al.* [4] and used phase contrast synchrotron X-ray imaging (PCXI) to record high-speed video of hemolymph flow in the hearts of grasshoppers (*S. americana*) via the movements of hollow glass tracer particles. We used particle image velocimetry (PIV) to quantitatively measure spatially and temporally resolved velocities of the flow tracers through the heart, as well as a custom PIV algorithm to estimate the net flow velocity at each instance in time. Additionally, we used light microscopy to image fluorescent flow tracer dye injected into the hearts of cockroaches (*Periplaneta americana*), and estimated the motion of the blood flow by tracking the motion the dye throughout the field of view.

Results and Discussion

The flow of hemolymph in the hearts of cockroaches and grasshoppers reversed directions multiple times per second.

Our measurements of the flows in the hearts of cockroaches and grasshoppers revealed behaviors that were unexpected according to the current understanding of the insect heart. In fluorescent dye injection trials in cockroaches, dye was drawn into the heart through a single ostium (valve). For 5–20 seconds following injection, the motion of the bolus of dye was clearly discernable against the dark image background. We observed that the flow rapidly and repeatedly reversed direction multiple times per second (Figure 3.1). We quantified this behavior by defining the flow reversal period as the time elapsed during a single cycle of the flow oscillation. To estimate the duration of flow periods, we measured the image brightness (intensity) within sub-regions of the images that encompassed portions of the heart located near the injection point, and interpreted changes in intensity through time to indicate the flow of dye into and out of the measurement regions. These data showed that local maxima of intensity in regions of interest located anterior to the injection site corresponded to local minima in posterior regions, and vice-versa (Figure 3.2); this pattern was consistent with our visual observation that the bolus of dye was drawn repeatedly back and forth between the pair of measurement regions. We calculated flow reversal periods from these data by measuring the time elapsed between intensity maxima in the anterior region of interest. In total, we measured 153 flow reversal periods over 58 seconds in 3 animals, and calculated the mean period of flow reversal in these trials to be 0.35 ± 0.22 seconds (95% C.I.).

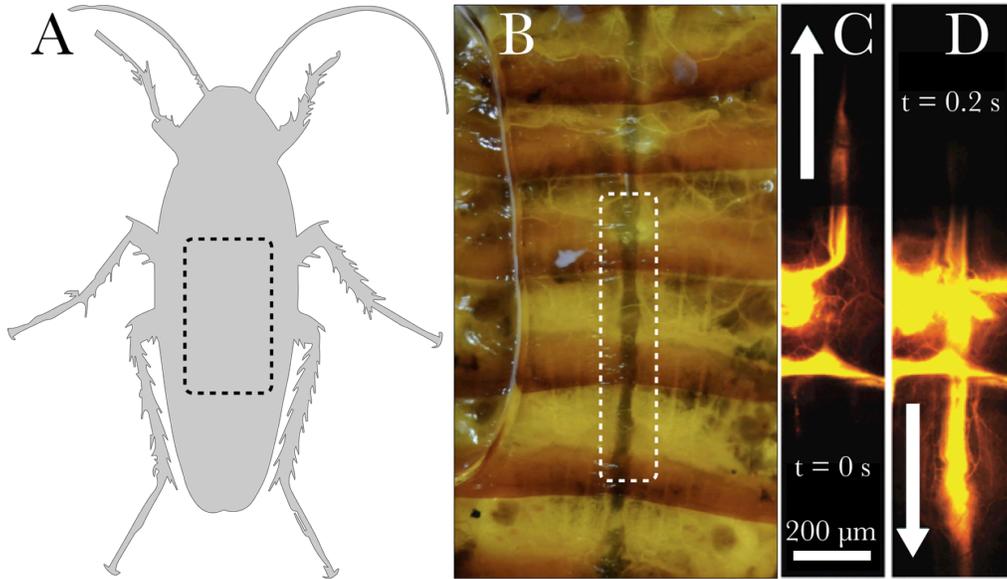


Figure 3.1: Fluorescent dye movement in the heart of a cockroach (*Periplaneta americana*). A: dashed rectangle indicates the approximate location of the field of view shown in B. B: Dorsal view of the abdomen; anterior direction is upward. The thin, dark, vertical stripe in the center of the image outlines the heart. Dashed rectangle indicates the fields of view in C and D. C: dye is drawn into the heart from a needle inserted to the left of the vessel. The dye was advected toward the head immediately after it entered the heart. D: the same bolus of dye is advected toward the posterior 0.2 seconds following the photograph in C, leaving behind it a region of low intensity.

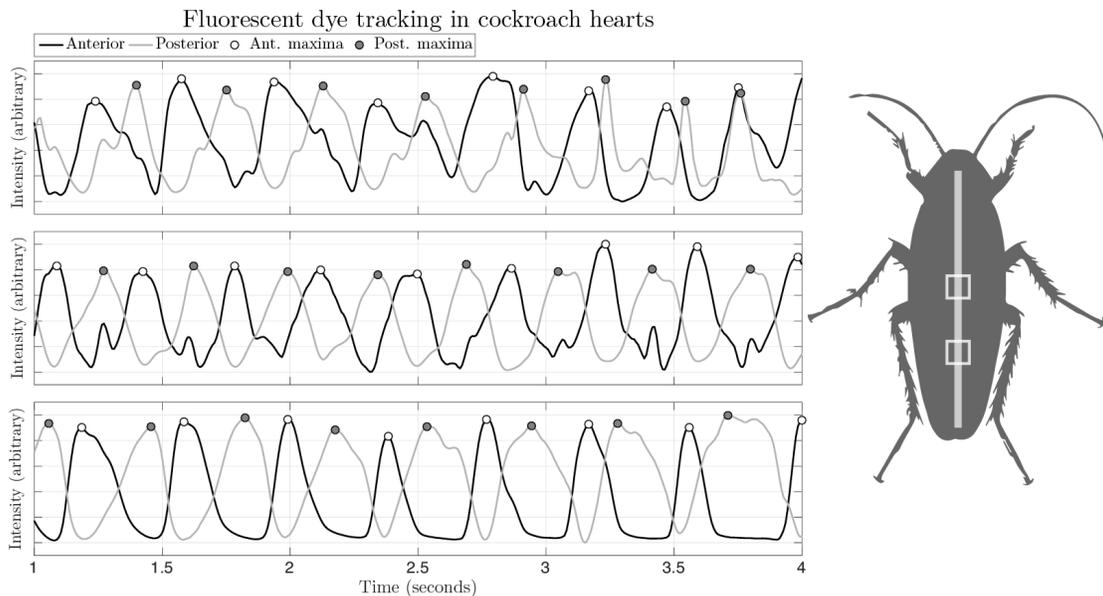


Figure 3.2: Average intensities of subregions in images of fluorescent dye injected into the hearts of cockroaches (*P. americana*). Right: approximate locations of the anterior and posterior regions within which intensity was measured. We interpreted the time between intensity peaks in the anterior region as the period of flow reversal.

In X-ray images of grasshoppers, the spatially-resolved velocity vector fields we measured using PIV appeared to exhibit regions of roughly parabolic-like flow oriented along the longitudinal axis of the animal. These regions contrasted against the remainder of the field where the calculated velocities were close to zero and exhibited no obviously discernable flow patterns. Parabolic-like velocity fields such as these are characteristic of flows through tubes and other channels, where the flow velocity diminishes to zero at the walls of the tube and is greatest somewhere in between. (The location of maximum velocity of fully-developed flow lies exactly at the radial center of perfectly straight tubes, but can deviate from the center in curved or tortuous tubes). We therefore interpreted the regions of parabolic-like flow to lie within the lumen of the heart, and will subsequently refer to the flow in that region as the heart flow. Based on this flow pattern, the heart exhibited a diameter of $\sim 600 \mu\text{m}$, measured dorsoventrally. Figure 3.3 illustrates a time series of spatially-resolved heart flow fields measured in one animal, as well as spatially averaged profiles of axial velocity over a slightly longer duration. The primary salient feature of these data was that the direction of flow appeared to alternate from anterograde to retrograde and back multiple times per second. These rapid reversals were prominent in every grasshopper we tested, and were consistent with our qualitative visual observations of the flow throughout the duration of each X-ray video we recorded as well as our dye injection measurements in cockroaches.

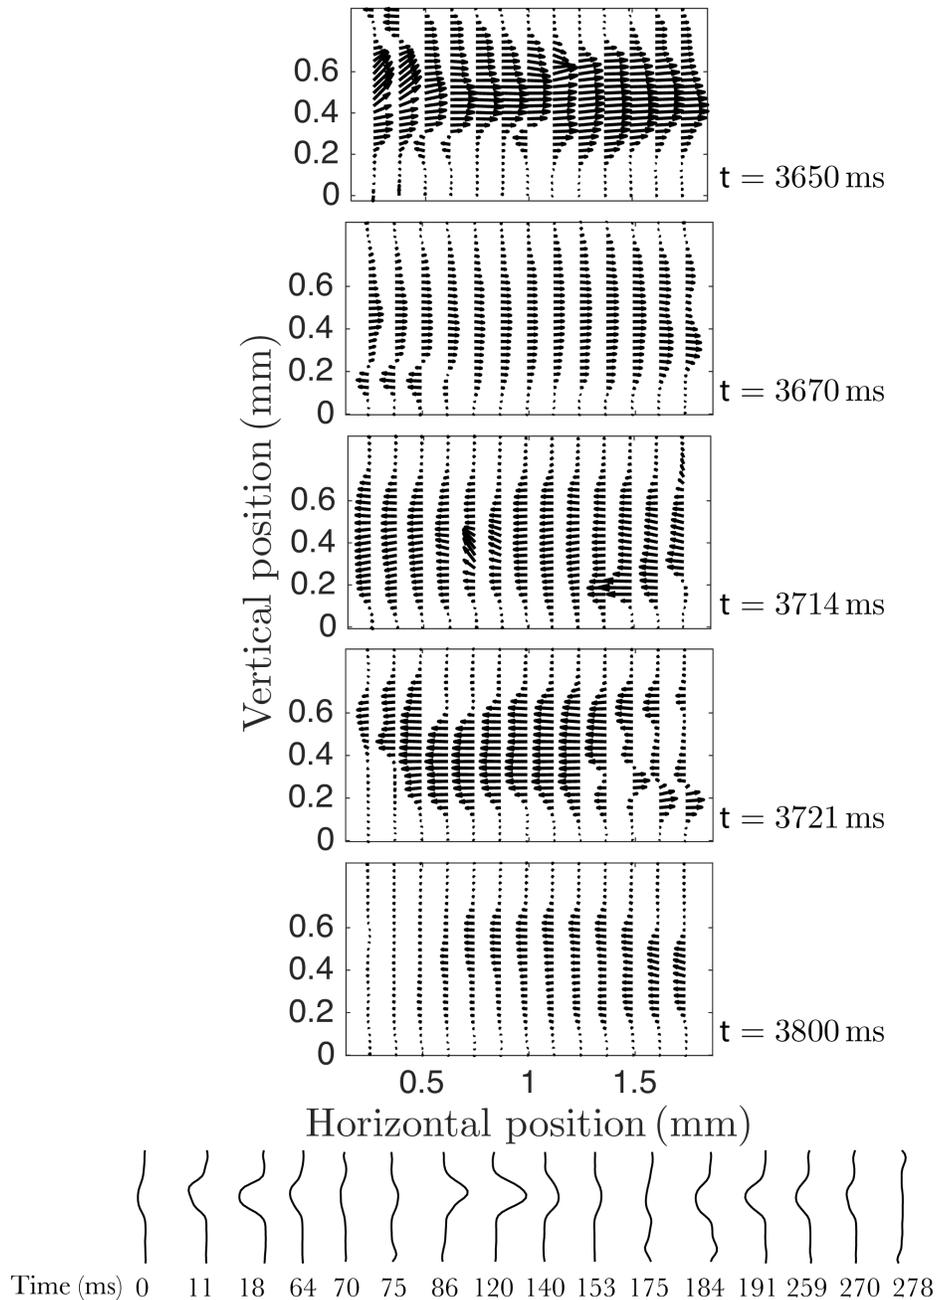


Figure 3.3: X-ray PIV-measured velocity fields and profiles in a grasshopper (Animal 2 in Figure 3.5). Anterograde flow is to the left. The vector field sequence (top five rows) shows a single flow reversal that occurs over 150 ms, with times corresponding to those in Figure 3.5. Vector fields are cropped in the vertical direction to highlight flow in/near the heart. Bottom row: profiles of axial component of flow velocity in the same animal, averaged in the axial direction; time relative to first profile shown, which corresponds to 3530 ms in Figure 3.5. The vertical direction of these profiles extends the entire field of view (about 2 mm). The prominent, rapid flow reversals in the vector fields and velocity profiles were congruent with our visual observations from X-ray videos, as well as dye injections in cockroaches.

We quantified the directions and magnitudes of the net flows in grasshoppers using estimates of the spatially-averaged velocities within the heart, which were measured using our custom spatial-ensemble PIV algorithm. As illustrated in Figure 3.5, these measurements indicate the presence of rapidly oscillating flow that was consistent with our observations from spatially-resolved PIV velocity measurements as well as from dye injection measurements in cockroaches. From these data, we located the times at which the flow velocity crossed the zero axis with increasing slope, and quantified the duration of reversal periods as the amount of time elapsed between these events. From these data, we calculated an average flow reversal period of 0.35 ± 0.35 seconds (95% C.I.) in four animals over 34 seconds; this result agreed well those from our dye injection measurements (Figure 3.4 and Table 3.1). Our observations of fast sub-second, oscillatory flow reversals in grasshoppers and cockroaches contrast starkly with reports that other species' hearts produce unidirectional flows [6]. Moreover, the time scale of the flow reversals we observed is one or more orders of magnitude shorter than those described in other species, whose directions reverse on the time scale of seconds to minutes [5-7, 9, 11, 14-17]. Our results support the preliminary work of Lee & Socha [4], who manually tracked flow tracer particles in X-ray images of grasshopper hearts and reported that the flow sometimes appeared to oscillate back forth, without producing any clearly discernable net flow.

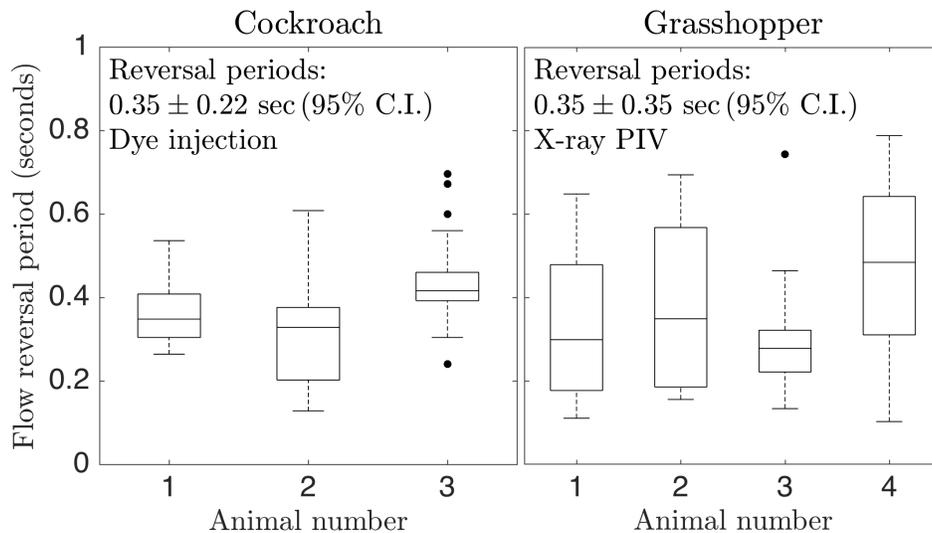


Figure 3.4: Times between flow reversals in the hearts of cockroaches and grasshoppers. Boxes indicate the 25th through 75th percentile of data; horizontal lines show median values; whiskers indicate a 99.5% confidence interval; dots show data points outside this confidence interval.

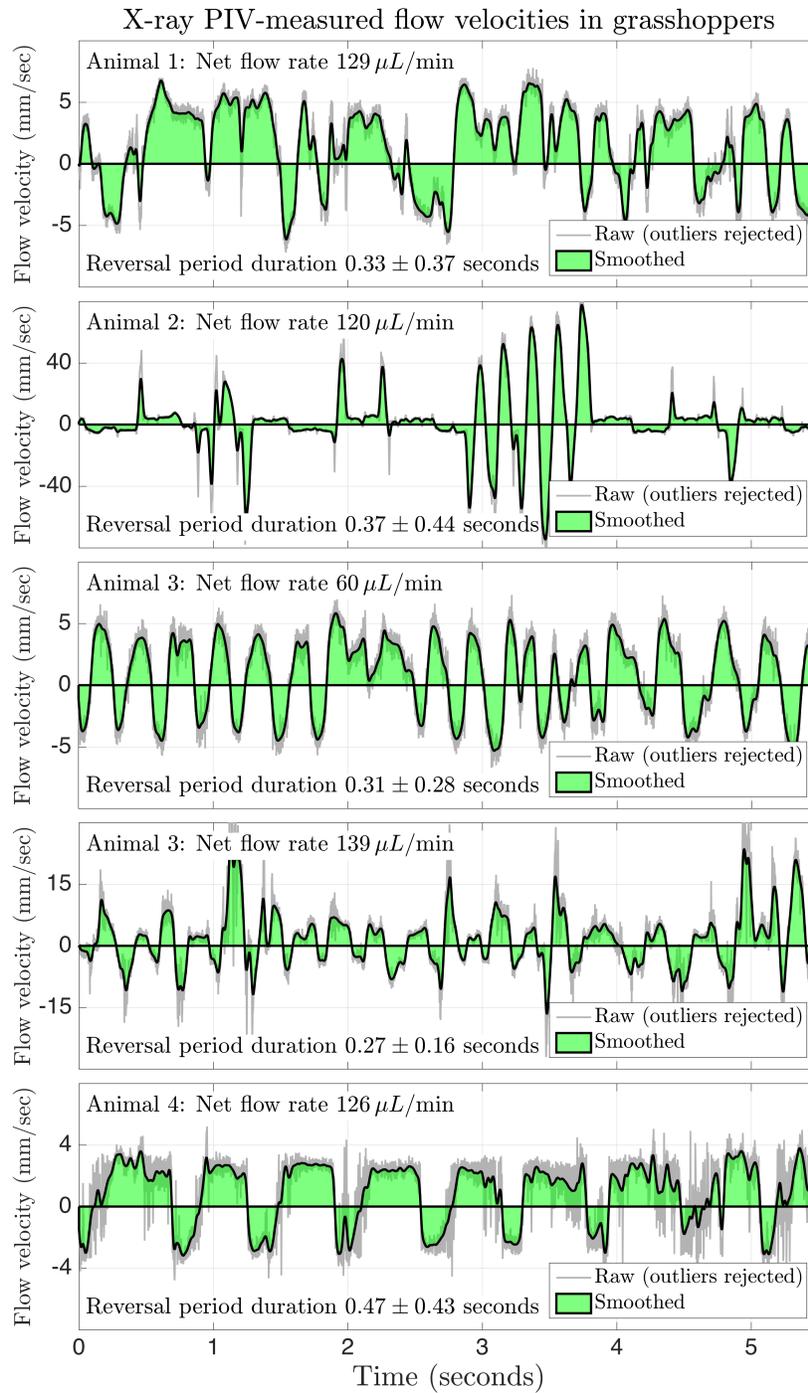


Figure 3.5: Spatially averaged flow velocities in the hearts of grasshoppers (*Schistocerca americana*). We used our spatial ensemble PIV correlation method to measure time resolved, spatially averaged flow velocities from synchrotron phase-contrast X-ray videos (1000 frames per second) of hollow glass tracer particles in four animals' hearts. The flow prominently switched direction multiple times per second; this is consistent with spatially resolved PIV measurements in these animals and dye injection in cockroaches. Green regions show the areas that were integrated to calculate time averaged net volumetric flow rates. Note the different scales on the vertical axes.

Species	No. animals	Duration (sec)	N_R	\bar{T} (sec)	σ_T (sec)
<i>P. americana</i>	3	58	153	0.35	0.11
<i>S. americana</i>	4	34	86	0.35	0.18

Table 3.1: Summary of flow reversal periods measured in cockroaches and grasshoppers. N_R number of reversals; \bar{T} average flow reversal period; σ_T standard deviation of flow reversal period.

The net flow rate in grasshopper hearts was 60-140 μL per minute (about 25%–75% their total blood volume) toward the head

We estimated time-averaged volumetric flow rates in the grasshopper heart by numerically integrating the PIV-measured velocities through time. The diameter of the heart's lumen was estimated from spatially-resolved PIV velocity fields to be 600 microns. By these calculations, the net flow rates through these animals' hearts were between 60-140 microliters per minute, and in all cases the direction of the net flow was toward the head. To compare these flow rates to the total volume of hemolymph in the body, we used volume measurements from Lee [19], who calculated that the mass-specific volume of hemolymph in the adult desert locust *S. gregaria* was about $144 \mu\text{L} / \text{g}$. Assuming similar values in *S. americana*, and given the masses of our animals were measured between 1.3–1.5 g, we estimate the volume of hemolymph in the individuals we tested was 190-220 μL . Using the upper and lower bounds of our estimated pumping capacities and hemolymph volumes, we estimate the time required for the heart to pump the body's volume of blood to be about 1-4 minutes. As a point of comparison, mice pump their body's volume of blood in about seven seconds, and humans and elephants do so in about 1 and 2.5 minutes [20]. Vogel points out that smaller animals (at least mammals) generally circulate their blood more rapidly (in body volumes per unit time) than larger animals do [20]. However, the flows we measured in the grasshopper heart appear to violate that trend. This might be explained by differences in their metabolisms, or by the different functions of their blood: vertebrate blood is primarily important for transporting oxygen, whereas insect blood functions only marginally in gas exchange [5, 7]. Instead, a primary role of hemolymph is the transport of nutrients and hormones throughout the body [5, 7]. The oscillatory flow in the heart seems to suggest that it is not particularly effective at producing directional flow, but an alternative explanation could be that its main function is related to mixing these substances throughout the blood rather than transporting them to different parts of the body.

The flows we measured in the hearts of grasshoppers and cockroaches appear qualitatively to exhibit some features of non-peristaltic pumping

Tubular pumps are widespread in nature, and many of them create flow via peristalsis [1, 5, 18, 21, 22]. The insect heart, with its tubular structure and pulsatile, wave-like contractions, has long been presumed to be among these [5, 7]. Alternative mechanisms for generating flow in valveless tubes have been described for years [18, 23], but only recently have they garnered much attention in the biological community. Most notably, Forouhar [1] showed that the embryonic zebrafish heart is probably not a peristaltic pump, but rather an impedance pump (also referred to as a dynamic suction pump [1, 18, 22], or a Liebau pump after the researcher who first described the mechanism [18, 22-24]). The evidence they presented to support this claim was that they observed relationships between the dynamics of the myocardium and of the blood flow that were unexpected according to the technical definition of peristaltic pumping, but were not inconsistent with previous analytical, numerical, and experimental models of impedance pumping. Namely, according to the technical definition, peristaltic flow is continuous (not oscillatory) and unidirectional; the maximum velocity of the flow does not exceed the velocity of the traveling waves; and the flow velocity is proportional to the frequency of the peristaltic contractions [18, 22]. In contrast, Forouhar observed that flows in the embryonic zebrafish heart were oscillatory; the maximum velocity of the fluid exceeded that of the traveling waves; and the relationship between the frequency of contractions and the flow velocity was nonlinear. Alternatively, Waldrop & Miller [22] proposed that the technical definition of peristalsis is overly restrictive, and used numerical simulations to show that such flow characteristics could indeed arise in peristaltic pumping under an expanded definition of the mechanism. Their results motivate the question of whether some animals' hearts are not peristaltic pumps, or if they are indeed peristaltic pumps but fall outside a definition of peristalsis that is unnecessarily narrow.

Our measurements of cardiac flows in these insects are by themselves insufficient to address the question of whether or not their hearts are peristaltic pumps. To do so would require simultaneously quantifying the kinematics of the myocardium, which we have not done. However, several characteristics of the flows we measured appear to bear some qualitatively similarity to existing data on the kinematics of flow in impedance pumps. Bringley *et al.* [25] derived analytically the volumetric flow rate through a valveless, tubular impedance pump, and compared these predictions with experimental data they collected using a bench-top physical model. Their theory and experiments agreed that the direction of the volume flux oscillated

through time on the time scale of the imposed pulsations, but under certain conditions resulted in a nonzero net flow rate. Similarly, Jung *et al.* [26] numerically simulated the flow through a valveless impedance pump, and they too reported volume fluxes that oscillated on the time scale of the imposed pulsations, but were nonzero on average. Moreover, they showed that the direction of the time-averaged net flux could be positive, negative, or nearly zero depending on the forcing frequency; this fact is notable because the direction of net flow in some insects' hearts is known to reverse direction. The apparent similarity between our data and the flows described by these other researchers warrants further investigation, and is especially motivating in light of the recent developments that others have contributed to our understanding of the mechanisms of flow production in tubular hearts.

The species we studied were more than an order of magnitude larger than other animals whose cardiac flows have been observed directly.

Glenn *et al.* [6] used fluorescent flow tracer particles to observe the flow of hemolymph through the hearts of mosquitoes (*Anopheles gambiae*). They found that the mosquito heart produced net directional flow, but did not report any behaviors similar to the rapid flow reversals that we observed in grasshoppers and cockroaches. One plausible explanation for this discrepancy is that the patterns of contraction of the heart muscles may differ between the species. Alternatively, it is conceivable that such reversals are present in mosquitoes but were not observable using standard image magnification and video frame rates. Size may play a role as well. Adult mosquitoes are ~5 mm in length, with heart diameters of about 30 μm [6]. In comparison, adult grasshoppers and cockroaches are several centimeters in length with heart diameters of ~500 μm . The order-of-magnitude difference in size could lead to qualitatively different pumping characteristics. Grasshoppers in particular could be well-suited for studying the differences in flow behavior across body size because they do not undergo metamorphosis as they age: nymphs are about 4-5 mm in length but are anatomically similar to adults. An interesting future study could repeat the measurements we have performed using grasshopper nymphs of different instars.

The species of insects we studied were not closely related to one another

It is worth pointing out the evolutionary relationship between the two species of insects that we studied. According to a recent, comprehensive phylogenetic analysis of insects, the lineages of *Orthoptera* and *Blattodea* diverged about 250 million years ago [27], indicating that

the species used in this study are not closely related. This suggests that the use of rapid flow reversals in the heart must either have evolved independently within the different orders, or are basal with respect to their two lineages. Given the virtual ubiquity of hearts among insects, and the appreciable phylogenetic distance between the animals we used, it seems plausible that that the flows we observed could also be present in other species across multiple orders of insects. If future work confirms this hypothesis, and if the heart is proven to function other than by peristalsis, then this could imply that a previously overlooked flow mechanism could in fact be among the most prevalent among animals.

Materials and Methods

Animal Care

Grasshoppers (*S. americana*) were obtained from a colony maintained by the Arizona State University (ASU) School of Life Sciences (Phoenix, AZ). The animals were mailed in a double-boxed container from ASU to Argonne National Laboratory via overnight shipping (USDA permit number P526P-13-03508). Immediately upon receipt, the animals were transferred to a wire mesh cage in the laboratory (ambient temperature 25°C). They were allowed to feed *ad libitum* on lettuce, and fresh lettuce was provided daily. Additional heat was provided using an incandescent light bulb directed inward toward the cage. In accordance with USDA regulations, the animals were sacrificed once experiments had concluded.

Cockroaches were housed indoors in a ventilated plastic terrarium with a moistened substrate of potting soil, with cardboard structures for shade. The animals fed *ad libitum* on dry dog food (Dog Chow Complete, Nestlé Purina PetCare, St. Louis, MO). Water was provided using wetted balls of cotton.

Fluorescent Dye injection

The American cockroach *P. americana* was amenable for visualization of hemolymph flow via the injection of fluorescent dye because its heart is visible optically through the semi-transparent cuticle that covers the dorsal surface of its abdomen. Animals (N = 3) were fixed dorsal-side up to glass microscope slides by restraining their legs using mounting putty (Scotch Brand, 3M Company, Maplewood, Minnesota), and their wings were rotated to the side and pinned to reveal the dorsal surface of the abdomen. We used a 30-gage hypodermic needle (nominally 311 μm outer diameter, 159 μm inner diameter) to pierce the dorsal cuticle between the 3rd and 4th tergites, and hand-actuated a 1 mL syringe to inject fluorescent rhodamine dye

(Kingscote Chemicals, Miamisburg, OH) into the abdomen. During injections, the tip of the needle was placed to the left of the heart, adjacent to an ostium (anatomical valve through which hemolymph enters the heart) that was conspicuously visible under a stereomicroscope (Zeiss Stemi 2000-C). This placement allowed dye to flow out of the tip of the needle and immediately into the lumen of the heart. The dorsal surface was illuminated using a green laser (Hercules 532 nm, 375 mW, Laserglow Technologies, Toronto, Canada), and viewed through the stereomicroscope equipped with a long-pass optical filter (550 nm cutoff wavelength, ThorLabs, Newton, NJ) to attenuate the incident green laser light but not the dye's fluorescent red wavelengths. During each injection, a high-speed color camera (TS3-100, Fastec Imaging, San Diego, CA) recorded images at 125 frames per second. The field of view was 20 mm x 16 mm on a 1280 x 1024 pixel sensor (image magnification 15.6 μm per pixel).

To quantify the cardiac flow reversals, we measured the average intensities of the resulting images within two interrogation regions at anterior and posterior locations approximately 0.5 mm from the injection point. The raw intensity data were smoothed using a five-point moving average filter prior to detecting peaks using an automated peak-detection algorithm programmed in MATLAB (The MathWorks, Inc., Natick, MA). We calculated the expected flow reversal period as the 95% confidence interval about the mean of all reversal periods in all animals.

Phase contrast synchrotron X-ray Imaging (PCXI)

We used phase-contrast synchrotron X-ray imaging (PCXI) at beamline 32-ID-C of the Advanced Photon Source at Argonne National Laboratory (Lemont, IL) to visualize the flow of hemolymph within the hearts of four living grasshoppers. Hollow glass spheres of 8-12 μm outer diameter (product number 10089, TSI, Inc., Shoreview, MN) were used as flow tracers. Particles were suspended in water and injected into the dorsal side of the abdomen using a 1 mL syringe equipped with a 30-gage hypodermic needle. The needle was inserted between the 4th and 5th tergites, 2-3 mm left of the heart, and about 100 μL of the particle suspension was injected into each animal. Imaging was performed within 1-2 minutes following the injection of particles, limited by the required safety search of the beamline.

Each grasshopper was restrained and aligned with the X-ray beam by gripping the wings with small rubber-padded clamps (Figure 3.6). To keep the abdomen from moving out of the field of view, we restrained it by either inserting it loosely into a Kapton tube, or gently fixing it to a thin platform using adhesive-backed Kapton film. So that they did not enter the camera's field

of view, the hind legs were pivoted above the abdomen and held in this position using mounting putty (Scotch Brand, 3M Company, Maplewood, Minnesota). To roughly mimic the animals' natural thermal environment of 25-35°C [28], we irradiated them with a ceramic heater, and measured their body temperatures using a thermocouple as well as a non-contact infrared temperature sensor. The range of measured body temperatures was 30-35°C.

We used an X-ray beam energy of 25 keV based on earlier work by Socha *et al.* [29]. After passing through the specimen, the beam intersected a scintillating crystal, which emitted a visible light image. The scintillator image was viewed via a mirror using a high-speed camera (Photron SA-Z) equipped with a 10x-power microscope objective lens. The resulting field of view was 1.98 mm x 1.98 mm over 1024 × 1024 pixels (1.94 μm / pixel magnification), and the precise image magnification (in microns per pixel) was calculated by measuring the distance in pixels between markings on the face of an X-ray imaging calibration target (400 line-per-inch mesh). Images were captured at a frame rate of 1000 frames per second with a shutter speed of 1/1000 second. Specimens were mounted on motorized translation stages that permitted control of the camera's field of view within the animal. The field of view was adjusted for each specimen so that particles in the heart were clearly visible and could be seen conspicuously to flow. About 5.5 seconds of high-speed images were recorded for each trial.

The brightness of the X-ray images we recorded varied across the field of view and in time because the intensity of the X-ray beam was spatially and temporally non-uniform. We captured this effect by recording images with no specimen in the field of view (“white field images”), which were subsequently used to reduce the non-uniformity of the background in our experimental trials (described in the following section). A new white field image was recorded moments before imaging each new grasshopper specimen.

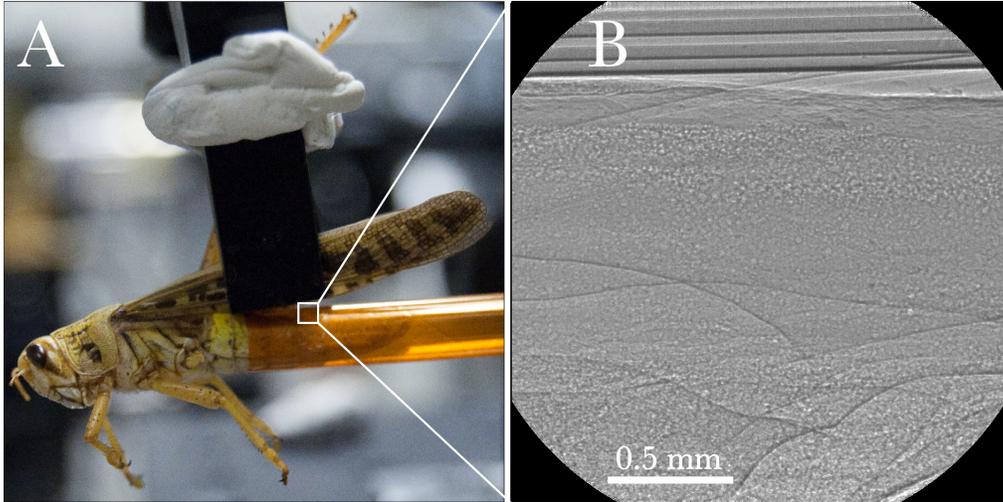


Figure 3.6: Preparation of a grasshopper (*Schistocerca americana*) for X-ray imaging. A: clamps hold the animal by its wings, and the abdomen is inserted into a Kapton tube to limit its movement. White box shows the approximate location of the camera field of view shown in B. B: X-ray image of a grasshopper after injection with hollow glass tracer particles, viewed through a 10x objective lens. Head is to the left; field of view is approximately 2mm x 2mm.

Our experimental methods were inherently invasive in that we physically restrained the animals and injected them with foreign flow tracers by piercing their exoskeletons with hypodermic needles. Moreover, in X-ray trials, the animals were subject to ionizing radiation within the imaged region for periods up to several minutes. Socha *et al.* demonstrated that another species of grasshopper, *Schistocerca gregaria*, could tolerate synchrotron X-ray irradiation of its head and abdomen for multiple hours without observable effects on gross behavior or respiratory characteristics [29]. Similarly, the animals used in our X-ray experiments did not exhibit any obvious qualitative changes in behavior over the duration of imaging. In both types of experiments, cockroaches and grasshoppers exhibited lively behavior after imaging, and continued to live normally in captivity until being sacrificed.

Image Pre-Processing

Compared to more typical PIV images, our X-ray images exhibited bright backgrounds of varying intensity, large flow tracer particles, and low image contrast. Moreover, in some trials, the animals' bodies moved within the field of view during recording, resulting from normal abdominal pumping behaviors. This motion would be expected to bias the PIV measurements and result in flow velocities measured in a frame of reference fixed to the camera rather than to the animal's body. Because these characteristics complicated the standard application of PIV

algorithms to the raw images, several image pre-processing steps were prerequisite to performing PIV measurements.

First, to mitigate the effects of the varying image background intensity due to non-uniformity across the X-ray beam, we divided each pixel in each image by the white field. Second, in trials where we judged body motion to be significant, we performed image registration (stabilization) prior to PIV measurements to reduce the apparent motion of the abdomen within the field of view and to shift the measurement frame of reference from that of the camera to that of the body. We registered images by tracking the motions of salient textures or features visible on the body using cross correlations, and then shifting each video frame so that the location of the tracked feature in the shifted image was coincident with its location in the first frame of the video. The displacements of tracked features were measured with sub-pixel precision by locating the maximum value of the Gaussian function that best fit the points surrounding the maximum value in the 2-D cross correlation plane. We identified spurious measurements by applying Chauvenet's criterion to the change of the tracked features' positions between consecutive frames, and replaced them using linear interpolation. After outlier replacement, the time series of the tracked features' positions were smoothed by convolution with a Gaussian kernel. Images were registered by shifting each frame (using cubic interpolation) according to the smoothed estimate of the tracked features' positions.

Finally, to reduce the contribution to the correlation of anatomical features that moved slowly compared to the flow tracer particles, we subtracted from each image the average intensity at each pixel calculated over a short time series spanning the nine preceding and subsequent images in the time series ("moving mean subtraction").

PIV Processing

Particle image velocimetry (PIV) is a technique used to experimentally measure velocity fields in fluid flows [30]. In traditional PIV, the fluid is seeded with small tracer particles (1 – 100 μm) that are either highly reflective or fluorescent. The region of interest in the flow is illuminated with a laser formed into a thin sheet, and one or more digital high-speed cameras image the illuminated area. The resulting images are grouped into subsequent pairs, and cross-correlations are used to estimate the displacements of groups of tracer particles between photographs. Modern PIV algorithms can reliably estimate particle displacements with sub-pixel accuracy [31]. Achieving this accuracy typically requires (among other considerations) images that exhibit high contrast between the particles and the background; a dark, uniform background

with a relatively low level of intensity noise; small particle images (2-3 pixels in diameter); and an appropriate level of particle seeding (about 10 particles per 32×32 -pixel interrogation region, or IR). In contrast to these characteristics, the X-ray photographs obtained during our experiments contained low-contrast images of relatively few large particles against a background of bright and varying intensity. Even after the image pre-processing steps described in the previous section, these images yielded PIV cross correlations of low signal-to-noise ratios (SNR), which in turn reduced the reliability with which traditional PIV algorithms could measure realistic displacements of the particle patterns. These reasons for effects were described in Chapter 2 of this dissertation, and as such, we mitigated them using insight from the theory and Fourier filtering algorithm that were developed in that chapter. Specifically, we used our adaptive phase correlation (APC) algorithm to measure the SNR distributions in cross correlations of X-ray PIV interrogation regions. However, our APC algorithm sometimes failed to measure reasonable SNR distributions (and therefore Fourier filters) for certain individual cross correlations. To address this, we chose not to use Fourier filters that were calculated independently for each IR, but instead calculated an average APC-estimated filter size for each X-ray trial, and enforced that filter for all interrogation regions. The effective particle diameters of these average filters ranged from 12-15 pixels; this nomenclature is explained in Chapter 2 and by Eckstein *et al.* [32].

We measured spatially and temporally-resolved flow fields from X-ray images in grasshoppers following the image pre-processing procedure described in the previous section. The dimensions of the interrogation regions were 256×128 pixels (width-by-height), and the grid spacing of was 32×16 pixels. Because cross correlations were performed in the spectral (Fourier) domain, the pixel intensities in each interrogation region were multiplied by a Gaussian-shaped apodization function to reduce aliasing of their Fourier transforms. The effective resolution [33] of the apodized regions was 192×96 pixels. Spurious vectors were removed and replaced using universal outlier detection, and the resulting vector fields were smoothed by convolving them with a unity-gain Gaussian-shaped kernel whose extent was 7×7 grid points and whose standard deviation was one grid point.

Spatial Ensemble Cross Correlation

In addition to measuring spatially-resolved velocity vector fields, another parameter of interest in our experiments was the velocity of the net flow (i.e., the spatially-averaged velocity) at each instance in time. One way to calculate this would be to simply average each component

of the spatially-resolved PIV-measured vector field for each snapshot. This would be similar in principle to calculating a temporally-averaged velocity vector at each spatial location in a quasi-steady flow field by averaging the vectors from each interrogation spot over multiple snapshots. In such situations, however, other researchers have shown that the accuracy of PIV velocity estimates can be improved by instead averaging the cross correlation planes themselves rather than averaging the vectors estimated from individual correlation planes, and subsequently estimating a single time-averaged vector from each time-averaged correlation plane [34]. This approach of adding together multiple correlation planes through time to estimate a single time-averaged vector is known as the "ensemble correlation" or "sum of correlations," which improves estimates of spatially resolved, time-averaged flow fields at the expense of temporal resolution.

We adopted the principle of the ensemble correlation to calculate a single spatially-averaged flow velocity vector at each instance in time. In this approach, we averaged correlation planes spatially (across multiple interrogation spots from each single pair of images) rather than temporally (at each single interrogation spot across multiple pairs of images). Our "spatial-ensemble correlation" (or SEC) algorithm sacrifices spatial resolution and returns a single spatially averaged correlation plane for each pair of images. The best estimate of the spatially-averaged displacements of the tracer particles between two consecutive frames is then calculated from the position of the maximum value of this spatially-averaged correlation plane.

We used our SEC algorithm to measure spatially-averaged, time-resolved flow velocities from X-ray videos of four grasshoppers. The interrogation region size, image apodization filter, and Fourier filters were identical to those used in spatially-resolved PIV processing, and the grid spacing of interrogation regions was 32 x 32 pixels. A hand-drawn polygonal-shaped "grid mask" was used to restrict the interrogation spots to a region of the images that we judged to lie within the heart, and a different mask was used for each trial.

We estimated the velocity of the flow along the longitudinal axis of the dorsal vessel by manually measuring the vessel's orientation within the X-ray images, and projecting the measured PIV velocities onto its axis. Outlying measurements were rejected from each time series by applying Chauvenet's criterion to the inter-frame change in axial flow velocity, and were replaced using linear interpolation; this decision was based on our interpretation that sudden, drastic changes in flow velocity corresponded to spurious PIV measurements. After outlier replacement, the time series was smoothed by convolution with a unity-gain Gaussian-shaped kernel with a standard deviation of 10 ms and a kernel length of 250 ms.

Acknowledgement

This research used resources of the Advanced Photon Source, a U.S. Department of Energy (DOE) Office of Science User Facility operated for the DOE Office of Science by Argonne National Laboratory under Contract No. DE-AC02-06CH11357.

References

1. Forouhar, A.S., et al., *The embryonic vertebrate heart tube is a dynamic suction pump*. Science, 2006. **312**(5774): p. 751-753.
2. Butcher, J.T., et al., *Transitions in early embryonic atrioventricular valvular function correspond with changes in cushion biomechanics that are predictable by tissue composition*. Circ Res, 2007. **100**(10): p. 1503-11.
3. Hu, N. and E.B. Clark, *Hemodynamics of the stage 12 to stage 29 chick embryo*. Circ Res, 1989. **65**(6): p. 1665-70.
4. Lee, W.K. and J.J. Socha, *Direct visualization of hemolymph flow in the heart of a grasshopper*. Integrative and Comparative Biology, 2009. **9**(1).
5. Chapman, R.F., *The Insects - Structure and Function*. 4 ed. 1998, Cambridge: Cambridge University Press.
6. Glenn, J.D., J.G. King, and J.F. Hillyer, *Structural mechanics of the mosquito heart and its function in bidirectional hemolymph transport*. Journal of Experimental Biology, 2010. **213**(4): p. 541-550.
7. Jones, J.C., *The circulatory system of insects*. 1977, Springfield, Ill.: Thomas.
8. Pass, G., *Accessory pulsatile organs: evolutionary innovations in insects*. Annu Rev Entomol, 2000. **45**: p. 495-518.
9. Wasserthal, L.T., *Drosophila flies combine periodic heartbeat reversal with a circulation in the anterior body mediated by a newly discovered anterior pair of ostial valves and 'venous' channels*. Journal of Experimental Biology, 2007. **210**(21): p. 3707-3719.
10. Slama, K. and J. Lukas, *Myogenic nature of insect heartbeat and intestinal peristalsis, revealed by neuromuscular paralysis caused by the sting of a braconid wasp*. Journal of Insect Physiology, 2011. **57**(2): p. 251-259.
11. Slama, K., *Extracardiac versus cardiac haemocoelic pulsations in pupae of the mealworm (Tenebrio molitor L.)*. Journal of Insect Physiology, 2000. **46**(6): p. 977-992.

12. Andereck, J.W., J.G. King, and J.F. Hillyer, *Contraction of the Ventral Abdomen Potentiates Extracardiac Retrograde Hemolymph Propulsion in the Mosquito Hemocoel*. Plos One, 2010. **5**(9).
13. Gerould, J.H., *Orders of insects with heart-beat reversal*. The Biological Bulletin, 1933. **64**(3): p. 424-431.
14. Wasserthal, L.T., *Oscillating haemolymph circulation and discontinuous tracheal ventilation in the giant silk moth *Attacus atlas* L.* Journal of Comparative Physiology B: Biochemical, Systemic, and Environmental Physiology, 1981. **145**(1): p. 1-15.
15. Thon, B., *Influence of the cardiac phase on the latency of a motor response to a visual stimulus in the blowfly*. Journal of Insect Physiology, 1982. **28**(5): p. 411-416.
16. Wasserthal, L.T., *Oscillating haemolymph 'circulation' in the butterfly *Papilio machaon* L. revealed by contact thermography and photocell measurements*. Journal of comparative physiology, 1980. **139**(2): p. 145-163.
17. Wasserthal, L.T., *Heartbeat reversal and its coordination with accessory pulsatile organs and abdominal movements in lepidoptera*. Experientia, 1976. **32**(5): p. 577-579.
18. Männer, J., A. Wessel, and T.M. Yelbuz, *How does the tubular embryonic heart work? Looking for the physical mechanism generating unidirectional blood flow in the valveless embryonic heart tube*. Dev Dyn, 2010. **239**(4): p. 1035-46.
19. Lee, R.M., *The variation of blood volume with age in the desert locust (*Schistocerca gregaria* Forsk.)*. Journal of Insect Physiology, 1961. **6**(1): p. 36-51.
20. Vogel, S. and R.A. Calvert, *Vital circuits: on pumps, pipes, and the workings of circulatory systems*. 1993: Oxford University Press on Demand.
21. Vogel, S., *Living in a physical world X. Pumping fluids through conduits*. J Biosci, 2007. **32**(2): p. 207-22.
22. Waldrop, L. and L. Miller, *Large-amplitude, short-wave peristalsis and its implications for transport*. Biomechanics and Modeling in Mechanobiology, 2016. **15**(3): p. 629-642.
23. Liebau, G., *Über ein ventilloses Pumpprinzip*. Naturwissenschaften, 1954. **41**(14): p. 327-327.
24. Auerbach, D., W. Moehring, and M. Moser, *An analytic approach to the liebau problem of valveless pumping*. Cardiovascular Engineering: An International Journal, 2004. **4**(2): p. 201-207.
25. Bringley, T.T., et al. *An experimental investigation and a simple model of a valveless pump*. 2008.

26. Jung, E. and C.S. Peskin, *Two-Dimensional Simulations of Valveless Pumping Using the Immersed Boundary Method*. SIAM Journal on Scientific Computing, 2001. **23**(1): p. 19-45.
27. Misof, B., et al., *Phylogenomics resolves the timing and pattern of insect evolution*. Science, 2014. **346**(6210): p. 763-767.
28. Stower, W.J. and J.F. Griffiths, *The body temperature of the desert locust (schistocerca gregaria)*. Entomologia Experimentalis et Applicata, 1966. **9**(2): p. 127-178.
29. Socha, J.J., et al., *Real-time phase-contrast x-ray imaging: a new technique for the study of animal form and function*. BMC Biology, 2007. **5**(1): p. 6.
30. Westerweel, J., *Fundamentals of digital particle image velocimetry*. Measurement Science and Technology, 1997. **8**(12): p. 1379.
31. Westerweel, J., *Theoretical analysis of the measurement precision in particle image velocimetry*. Experiments in Fluids, 2000. **29**(1): p. S003-S012.
32. Eckstein, A. and P.P. Vlachos, *Digital particle image velocimetry (DPIV) robust phase correlation*. Measurement Science & Technology, 2009. **20**(5).
33. Eckstein, A. and P.P. Vlachos, *Assessment of advanced windowing techniques for digital particle image velocimetry (DPIV)*. Measurement Science & Technology, 2009. **20**(7).
34. Meinhart, C.D., S.T. Wereley, and J.G. Santiago, *A piv algorithm for estimating time-averaged velocity fields*. Journal of Fluids Engineering, 2000. **122**(2): p. 285-289.

Chapter 4: Conclusions and Significance

Conclusions and Significance of Chapter 1: Measurement of fluid rotation, dilation, and displacement in particle image velocimetry using a Fourier–Mellin cross-correlation

The focus of this dissertation has been to improve or carry out measurements of fluid flow velocities. The algorithms we developed, however, are agnostic to the physical system under investigation, and are instead concerned more generally with the motions of patterns. They should therefore be applicable to any situation where the assumptions under which they were developed remain valid. One area of research to which we anticipate our FMC algorithm can immediately add value is the experimental study of solid mechanics. In addition to fluid mechanics, the technique of measuring the displacements of patterns using cross correlations is also widely used to experimentally investigate the mechanics of solid materials. When applied to solids, the technique is generally referred to as "digital image correlation" [1] or DIC, wherein the measurement of interest is the spatial distribution of strain. While specific details of DIC measurements differ from those of PIV (e.g., the correlating pattern can be created by spattering paint onto the surface of a solid material), many of their processing algorithms are identical. An important similarity between the two methods is that the materials often exhibit deformations that are poorly approximated by pure translation. A common example of this in DIC is the stretching of materials under large deformations (e.g., polymers and biological materials [2]). In these situations, the cross correlation-based measurement can fail for the same reasons as those we discussed for PIV measurements of rotating or dilating fluid flows. It seems reasonable, then, to anticipate that our FMC algorithm, which directly measures isotropic dilation in addition to rotation, could help to mitigate the decorrelation of DIC images due to stretching. Potential avenues of inquiry could include examining the degree to which the FMC algorithm can tolerate non-isotropic stretching (which is common in uniaxial strain tests, for example), as well as quantifying a relationship between characteristics of the correlating pattern in DIC (often created using paint) and the performance of the algorithm.

In addition to improving the performance of experimental measurements, the broader implication of this work was that a fundamental assumption of cross-correlation-based image velocimetry -- namely, that the deformation tensor is approximated purely by translation -- can be relaxed by modifying the cross correlation. Our specific algorithm extended the class of

directly measurable deformation tensors from rigid body translations to angle preserving "similarity transformations." By doing so, we increased the number of degrees of freedom permitted to the interrogation region from two to four (one additional for isotropic scaling and one for rotation). This development naturally motivates the investigation of other (possibly related) modifications to the CC that would further generalize the measurable class of deformations to include "affine" transformations, which preserve parallel relationships, and assume only that the fluid deformation tensor is locally linear. The assumption of linearity is fundamental to typical formulations in continuum mechanics, and would provide a more direct congruency between experiments and theory than current measurement methods provide. Other researchers have proposed methods to measure piecewise affine transformations relating pairs of photographs [3, 4], as well as modifications to the cross correlation to directly measure shearing [5] or non-isotropic scaling [6]. Anecdotally, we have had some limited success in preliminary tests of using a nonlinear least squares approach to solve for the full affine transformation relating pairs of interrogation regions by initializing the solution using our FMC estimates of rotation, scaling, and translation; we found that gradient descent methods were prone to becoming stuck in local minima because PIV images are generally sparse, although this could be less of an issue for DIC images. These lines of inquiry all warrant further investigation.

Conclusions and Significance of Chapter 2: Adaptive Spectral Filtering of Particle Image Velocimetry Cross Correlations

We anticipate that the work presented in this chapter represents a starting point for multiple practical applications and theoretical explanations in image-based flow velocimetry. Perhaps the most straightforward additional application of our theory would be to measure the distribution of particle displacements in PIV images directly in the Fourier domain. The plausibility of this idea is evident from equation (2.5): the FT of the particles' displacement distribution is equal to the magnitude of the cross correlation divided by the magnitude of the particle autocorrelation, and the autocorrelation is simply the cross correlation of each interrogation region with itself. Measuring the displacement distribution in the spectral domain would overcome an important limitation of doing so in the spatial domain. Current algorithms estimate the displacement distribution by measuring characteristics (e.g. the width) of the largest peak in the spatial CC [7-9]. These algorithms depend on the successful identification of the displacement peak, which shrinks and becomes increasingly obscured by noise as the

displacement distribution grows (this effect was discussed in chapter 2). If the displacement peak shrinks to the degree that it becomes indistinguishable from the noise, then these methods fail. Conversely, estimating the displacement distribution from the spectral correlation would require no peak identification because the FT of the distribution function modulates the entire domain. In cases where the type of function describing the displacement distribution can be reasonably assumed a priori (e.g., a Gaussian distribution due to Brownian motion), then all that is required to estimate the distribution would be to fit that function to the quotient of the CC and the autocorrelation magnitudes. Several interesting comparisons between spatial and spectral estimation of displacement distributions could include the measurable range of distributions, and the effects of particle size, seeding density, and the number of ensemble pairs on the accuracy of each method.

We anticipate that our APC algorithm may be extended to take into account not only the magnitude but also the phase of the spectral cross correlation. As we showed in Figure 2.4, regions of high spectral SNR in the phase correlation appear "smooth" while regions of low SNR appear increasingly "noisy." Other researchers [10] have quantified this effect by introducing the notion of "phase quality," which is a measure of the local standard deviation of the magnitude of the gradient of the phase angle $\phi(k, m)$, where

$$\phi(k, m) = \tan \left(\frac{\mathcal{R}(\tilde{R})}{\mathcal{I}(\tilde{R})} \right) \quad (4.1)$$

In the absence of any noise, equation (4.1) represents a plane wrapped to the range $\pm\pi$. Since the gradient of a plane is constant, the phase quality (in an unusual use of the word) is minimized for the noise-free case, and grows to a maximum of 4π as noise increases. In this way, the phase quality can be thought of as representing the reciprocal of a measure of the SNR of the phase correlation. It therefore seems reasonable that this metric might be used to complement the magnitude-based SNR estimate that we introduced.

We envision at least two approaches for complementing our APC method with the phase quality metric. The first is to use the phase quality to derive some weighting function for the least squares fit of the Gaussian (or other appropriate function) to the correlation magnitude. In contrast, our current method uses an unweighted least squares fit. This can be problematic because the method of least squares is sensitive to outliers, and yet we expect significant levels

of noise in regions of CC where the spectral SNR is small. Weighting the least squares fit by a metric derived from the phase quality would reduce the contribution of noise-dominated regions of the CC to the estimate of spectral SNR from the correlation magnitude and could conceivably improve the reliability of the filter estimation.

Another approach could be to estimate the spectral SNR or weight the spectral phase correlation directly from the phase quality. One potential advantage of this idea over our current magnitude-based method is that it would directly quantify the correlation SNR, whereas our method estimates the shape of the SNR but not its actual value. This subtle difference could prove to be important if the algorithm is applied to other types of flow velocimetry images (e.g., Schlieren) where the assumptions of our SNR model do not hold. For instance, our APC method inherently assumes that noise is uniformly distributed across the Fourier spectrum of the CC, which may not be the case with other types of images. A simple method for estimating SNR directly from the phase quality might be to suppress regions of the phase correlation in which the value of the phase quality meets some threshold criterion.

Our model of the cross correlation process in the Fourier domain helped to explain the fundamental reasons that velocity gradients reduce the SNR of PIV measurements. In a similar way, we speculate that future developments along these lines might quantitatively explain the reasons behind other PIV rules of thumb that have previously been described only heuristically. For example, others have experimentally found that the RMS error of PIV measurements is minimized for particle diameters of about 2.8 pixels [11], but we have found no explanation in the literature for why this is true. It is interesting to consider this result in the context of our explanation of the effect of the particle size on the SNR of the spectral correlation: as the particle size grows, the SNR increases at low wave numbers but decreases at high wave numbers. It seems plausible that the "amount of signal" (according to a metric yet to be defined) could be maximized for some particle size.

Another guideline in PIV that has been determined heuristically but not explained theoretically relates to the important experimental parameter of particle seeding density. Keane & Adrian determined experimentally that displacement peaks in the spatial CC tended to be "detectable" (peak height ratios greater than 1.2) in the absence of velocity gradients when the number of corresponding particles between IRs was greater than about 10 per 32 x 32 pixel region [12]. However, our model also predicts that particle size and the correlation of non-corresponding particles (which manifest as "noise" in the correlation) should play a role as well.

Revisiting this problem from the perspective of the spectral correlation formation process could prove to be insightful.

Conclusions and Significance of Chapter 3: High-Speed Flow Visualization Reveals Rapid Flow Reversals in Insect Hearts

Our measurements represent the first high-speed recordings of the flows in the dorsal vessel of insects and represent a major step toward understanding how these species produce heart flow. Many biological tubular hearts have been asserted to pump either by peristalsis or impedance pumping. Oftentimes peristaltic pumping is asserted as the mechanism as the default position, but largely without evidence; this is true of the insect heart (e.g., mosquitoes [13] and other flies [14], butterflies [15], moths [16], and beetle larvae [17]). Recently, researchers have argued that certain animals' hearts violate the definition of peristaltic pumping and appear instead to exhibit characteristics of other types of tubular pumps; the most notable example of this is the embryonic zebrafish heart [18]. However, the rejection of the peristalsis hypothesis is sometimes based on a technical definition of the mechanism that may be inappropriate for many biological pumps (an idea first espoused by Waldrop and Miller [19]). The tenets of this definition include that peristaltic pumps produce continuous flow whose maximum velocity is equal to that of the traveling contractile waves, and that the flow rate is proportional to the frequency of the contractions [19, 20]. This definition of peristalsis arose from analytical models that assumed that the amplitude of the contractile waves was small, or that the wavelengths were long relative to the tube diameter. However, some biological pumps exhibit large amplitude, short wavelength contractions. This discrepancy motivates the question of whether some animals' hearts are not peristaltic pumps, or if they are indeed peristaltic pumps but fall outside a definition of peristalsis that is unnecessarily narrow.

Waldrop and Miller [19] sought to address this question by performing numerical simulations of peristaltic tube pumping with large amplitude, short wavelength traveling waves, which mimic the apparent kinematics of heart tissue in some tunicates (sea squirts). They found that the resulting flow kinematics contrast with some of the tenets of the technical definition of peristalsis: namely, their simulations produced oscillatory flow, whose maximum velocities exceeded that of the compression waves, and the relationship between flow rate and the frequency of contractions was nonlinear. This result is significant because it suggests that flow kinematics that have been cited as evidence against the hypothesis of peristalsis in other

animals' hearts could actually be due to peristaltic pumping, and may warrant expanding and refining its definition.

It is tempting to speculate about the relationship between such results and our observations in insect hearts. As with the simulations of Waldrop *et al.* [19] and the measurements of Forouhar [18] in the embryonic zebrafish heart, our data suggest that the flows in insect hearts may violate at least one aspect of the technical definition of peristalsis—namely, they appear to produce oscillatory, not continuous, flow. More information is needed to determine whether other aspects of the current definition are violated as well. To this end, additional experiments should be performed to measure the kinematics of the heart wall as it beats. Quantifying the pattern of contractions and the speed of contractile waves in the dorsal vessel, along with simultaneous measurements of the hemolymph flow velocity, would answer the question of whether they are equivalent, as the technical definition of peristalsis would demand [20]. Similarly, manipulating and measuring the frequency of heartbeats during flow measurement experiments would reveal whether the relationship between them is linear (which, under the current definition of peristalsis, it should be). Numerical experiments could be performed to replicate the pattern of contractions and explore relationships between the relevant parameters (e.g., beat frequency, wave amplitude, wave speed, fluid properties, structural properties, and flow velocity). Relating such simulations to the physical system would require knowledge of the viscosity of the hemolymph (which we have begun to measure), and possibly the structural properties of the dorsal vessel. Furthermore, we have anecdotally observed that contractile waves sometimes appear to originate in multiple locations along the dorsal vessel, and do not always propagate unidirectionally along the vessel. If true, then the pattern of heart contraction would seem to be inconsistent with both peristalsis (even under an expanded definition proposed by Waldrop *et al.* [19]) and impedance pumping. This would suggest that the heart might produce flow by some other mechanism that has not previously been described.

In addition to quantifying additional metrics related to heart flows in insects, future experiments should include additional species. The two species in this study are distantly related, with lineages that diverged about 250 million years ago [21], which suggests that the similar behaviors we observed between them result either from characteristics that are basal to both or represent convergent evolution. Observing the presence or absence of similar characteristics across a broader phylogenetic range of insects would help to determine which of these alternatives is true, and would permit inferences that apply across multiple orders of insects, and perhaps the class as a whole. The most exciting speculation, in my opinion, is that

the recent enthusiasm for biological tubular pumps leads to a new theory of flow production that is neither peristalsis nor impedance pumping. If such a mechanism is found to be common among insects, then this could imply that a previously overlooked type of pump could in fact be among the most prevalent on Earth. It is my hope that this work contributes to a deeper understanding of the insect heart, and, more generally, about how organisms pump.

References

1. Michael A., M.A., J.-J. Orteu, and H.W. Schreier, *Digital Image Correlation (DIC), in Image Correlation for Shape, Motion and Deformation Measurements: Basic Concepts, Theory and Applications*. 2009, Springer US: Boston, MA. p. 1-37.
2. Pan, B., et al., *Two-dimensional digital image correlation for in-plane displacement and strain measurement: a review*. Measurement science and technology, 2009. **20**(6): p. 062001.
3. Seetharaman, G., G. Gasperas, and K. Palaniappan. *A piecewise affine model for image registration in nonrigid motion analysis*. in *Proceedings 2000 International Conference on Image Processing (Cat. No.00CH37101)*. 2000.
4. Zitová, B. and J. Flusser, *Image registration methods: a survey*. Image and Vision Computing, 2003. **21**(11): p. 977-1000.
5. Gonzalez, R. *Registration of sheared images using phase correlation*. in *2013 28th International Conference on Image and Vision Computing New Zealand (IVCNZ 2013)*. 2013.
6. Gonzalez, R. *Fourier based registration of differentially scaled images*. in *2013 IEEE International Conference on Image Processing*. 2013.
7. Breuer, K.S., *Microscale diagnostic techniques*. 2005: Springer.
8. Hohreiter, V., et al., *Cross-correlation analysis for temperature measurement*. Measurement Science and Technology, 2002. **13**(7): p. 1072.
9. Olsen, M.G. and R.J. Adrian, *Out-of-focus effects on particle image visibility and correlation in microscopic particle image velocimetry*. Experiments in Fluids, 2000. **29**(1): p. S166-S174.
10. Ghiglia, D.C. and M.D. Pritt, *Two-dimensional phase unwrapping: theory, algorithms, and software*. Vol. 4. 1998: Wiley New York.
11. Adrian, R.J. and J. Westerweel, *Particle image velocimetry*. 2011: Cambridge University Press.
12. Keane, R.D. and R.J. Adrian, *Theory of cross-correlation analysis of piv images*. Applied Scientific Research, 1992. **49**(3): p. 191-215.

13. Glenn, J.D., J.G. King, and J.F. Hillyer, *Structural mechanics of the mosquito heart and its function in bidirectional hemolymph transport*. Journal of Experimental Biology, 2010. **213**(4): p. 541-550.
14. Wasserthal, L.T., *Drosophila flies combine periodic heartbeat reversal with a circulation in the anterior body mediated by a newly discovered anterior pair of ostial valves and 'venous' channels*. Journal of Experimental Biology, 2007. **210**(21): p. 3707-3719.
15. Wasserthal, L.T., *Oscillating haemolymph 'circulation' in the butterfly Papilio machaon L. revealed by contact thermography and photocell measurements*. Journal of comparative physiology, 1980. **139**(2): p. 145-163.
16. Wasserthal, L.T., *Oscillating haemolymph circulation and discontinuous tracheal ventilation in the giant silk moth Attacus atlas L.* Journal of Comparative Physiology B: Biochemical, Systemic, and Environmental Physiology, 1981. **145**(1): p. 1-15.
17. Slama, K., *Extracardiac versus cardiac haemocoelic pulsations in pupae of the mealworm (Tenebrio molitor L.)*. Journal of Insect Physiology, 2000. **46**(6): p. 977-992.
18. Forouhar, A.S., et al., *The embryonic vertebrate heart tube is a dynamic suction pump*. Science, 2006. **312**(5774): p. 751-753.
19. Waldrop, L. and L. Miller, *Large-amplitude, short-wave peristalsis and its implications for transport*. Biomechanics and Modeling in Mechanobiology, 2016. **15**(3): p. 629-642.
20. Männer, J., A. Wessel, and T.M. Yelbuz, *How does the tubular embryonic heart work? Looking for the physical mechanism generating unidirectional blood flow in the valveless embryonic heart tube*. Dev Dyn, 2010. **239**(4): p. 1035-46.
21. Misof, B., et al., *Phylogenomics resolves the timing and pattern of insect evolution*. Science, 2014. **346**(6210): p. 763-767.