DEVELOPMENT OF ROBUST CORRELATION ALGORITHMS FOR IMAGE VELOCIMETRY USING ADVANCED FILTERING

by Adric Eckstein

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

**Master of Science
In
Mechanical Engineering**

**Dr. P.P. Vlachos (Chair)
Dr. M.R. Paul
Dr. A. Duggleby**

**December 7, 2007
Blacksburg, Virginia**

Keywords: Digital Particle Image Velocimetry, DPIV, image processing, time delay estimation, phase correlation, generalized cross correlation, digital filtering, nonlinear least squares regression, digital signal decomposition

DEVELOPMENT OF ROBUST CORRELATION ALGORITHMS FOR IMAGE VELOCIMETRY USING ADVANCED FILTERING

by Adric Eckstein

(ABSTRACT)

Digital Particle Image Velocimetry (DPIV) is a planar measurement technique to measure the velocity within a fluid by correlating the motion of flow tracers over a sequence of images recorded with a camera-laser system. Sophisticated digital processing algorithms are required to provide a high enough accuracy for quantitative DPIV results. This study explores the potential of a variety of cross-correlation filters to improve the accuracy and robustness of the DPIV estimation. These techniques incorporate the use of the Phase Transform (PHAT) Generalized Cross Correlation (GCC) filter applied to the image cross-correlation. The use of spatial windowing is subsequently examined and shown to be ideally suited for the use of phase correlation estimators, due to their invariance to the loss of correlation effects.

The Robust Phase Correlation (RPC) estimator is introduced, with the coupled use of the phase correlation and spatial windowing. The RPC estimator additionally incorporates the use of a spectral filter designed from an analytical decomposition of the DPIV Signal-to-Noise Ratio (SNR). This estimator is validated in a variety of artificial image simulations, the JPIV standard image project, and experimental images, which indicate reductions in error on the order of 50% when correlating low SNR images. Two variations of the RPC estimator are also introduced, the Gaussian Transformed Phase Correlation (GTPC): designed to optimize the subpixel interpolation, and the Spectral Phase Correlation (SPC): estimates the image shift directly from the phase content of the correlation. While these estimators are designed for DPIV, the methodology described here provides a universal framework for digital signal correlation analysis, which could be extended to a variety of other systems.

TABLE OF FIGURES

TABLE OF TABLES

NOMENCLATURE

ACRONYMS

DPIV	Digital Particle Image Velocimetry
DWO	Discrete Window Offset
FFT	Fast Fourier Transform
GCC	Generalized Cross Correlation
GTPC	Gaussian Transformed Phase Correlation
$GTPC^C$	Bias Corrected Gaussian Transformed Phase Correlation
IFFT	Inverse Fast Fourier Transform
ML	Maximum Likelihood
MRI	Magnetic Resonance Imaging
PHAT	Phase Transform
PCM	Phase Correlation Matrix
PID	Particle Image Distortion
PIP	Particle Image Pattern
PIV	Particle Image Velocimetry
RMS	Root Mean Square Error
RPC	Robust Phase Correlation
SCC	Standard Fourier Based Cross Correlation
SCOT	Smoothed Coherence Transform
SNR	Signal-to-Noise Ratio
SPC	Spectral Phase Correlation
SPE	Single Pixel Evaluation
SVD	Singular Value Decomposition
TRDPIV	Time Resolved DPIV
VVD	Valid Vector Detection
μPIV	Micro PIV

VARIABLES

B	number of bits in the memory storage of a number
$C(x,y)$	window normalization function
$C_{ij}(k)$	cross-spectrum of input signals i and j
d_c	correlation peak diameter
d_e	effective particle image diameter
d_{diff}	diffraction limited diameter
d_x, d_y	PIP image displacement in the specified orthogonal direction
d_p	physical particle diameter
$E(k)$	spectral energy function
E_{tot}	L2 norm of velocity error vector

e_ψ	error in the correlation value
e_x	error in the position value
F_i	factor describing in plane loss of particle image pairs
F_o	factor describing out-of-plane loss of particle image pairs
$f\#$	f-number
$G(k)$	Gaussian spectral filter
H	conjugate transpose
$H(k)$	prewhitening Spectral Filter
$I(k)$	Fourier transformed discrete image intensity distribution
$I(x)$	discrete image intensity distribution
I_0	maximum particle-image intensity
i	complex modulus
$J_1(x)$	Bessel function of the 1st kind
k_x, k_y	orthogonal wavenumbers
L_x, L_y	orthogonal window sizes
M	Magnification
M_r	median absolute deviation of the residuals
N_i	particle source density of image window
N_o	number of operations in FFT command
n	decomposition of noise component
P, Q	left and right singular matrices
p, q	left and right singular vectors
R	range of the quantized input levels
r_i	residual
s	decomposition of signal component
t_x, t_y	group delay in the specified orthogonal direction
x, y	orthogonal spatial directions
u, v	velocity estimate
U, V	actual velocity
$W(x)$	window mask
$W^0(x)$	zeropadded window mask
$W(k)$	arbitrary spectral filter
$W^T(k)$	weighting function with thresholding
$\alpha(k)$	factor representing the aliasing contributions in the SNR
β	factor representing the random noise contributions in the SNR

$\partial(x-x_0)$ delta function about x_0

ϕ — phase angle of complex number

$\gamma_{12}(k)$ — magnitude squared coherence

η — weighting function threshold

λ — wavelength of light

σ_n — thermal noise within image

ψ — correlation value

* — complex conjugate

$\otimes$ — convolution

Time-Resolved Digital Particle Image Velocimetry (TRDPIV) is a non-invasive flow visualization tool designed to resolve the spatio-temporal dynamics associated with complex flows. This technique is capable of high-frequency planar flow measurements through the use of a high-speed camera that captures the time evolution of the displacement of micron-size flow tracer particles illuminated by a high-frequency laser. In order to quantify these measurements, sophisticated processing algorithms correlate the locations of particle clusters within the image to estimate the velocity [1]. Minimizing the systematic errors associated with the image correlation algorithm is essential to resolve the dynamics of these flow fields accurately.

Over the past 20 years, many efforts have been directed towards the refinement of the basic digital algorithms developed by Willert and Gharib [2]. Among the most substantial contributions are the use of discrete window offsets [3], which can be further expanded upon through the use of iterative image window deformations [4]. Additionally, techniques such as the ensemble correlation [5] and the single pixel correlation [6] can offer greater accuracy when evaluating certain steady or periodic flows. With such a vast amount of research dedicated to the development of these algorithms, many are of the opinion that DPIV image evaluation has reached its full maturity [7]. However, the foundation of these digital evaluation techniques still rely upon a convolution of the image pairs evaluated with Fourier-based cross-correlation techniques[1]. The errors inherent within the Fourier based cross-correlation impose a fundamental limit on the accuracy of these techniques.

This research offers a modern approach to the DPIV processing methodology by exploring the potential of filtered phase correlation techniques, directed toward reducing this fundamental error limitation in DPIV image evaluation. In addition, digital image masks and filters are designed and optimized to increase the correlation strength in DPIV image shift estimation. This study demonstrates that the proper implementation of these filters with the phase correlation estimators significantly extends the dynamic range of measurements by providing a stronger correlation in a variety of degraded image conditions.

The striking improvements provided by these robust estimators offer the potential to extend the capabilities of DPIV as a flow visualization tool. High-speed DPIV systems can be substantially improved through the larger dynamic range offered by these algorithms. In addition, these algorithms are capable of substantially reducing peak-locking effects [8], a discretization error caused by the subpixel diameter of the flow tracers. The robustness of these estimators are also capable of reducing the effects of background illumination, large particle-image diameter, and white noise, which currently limit micro-PIV (µPIV) to time-averaged measurements.

Although these robust correlation estimators were designed for DPIV velocity estimation, these ideas and designs can also be extended to a variety of other systems such as topographical and map analysis, fingerprint identification, radar time-delay estimation, and MRI imaging. These systems all rely upon resolving signals from a finite duration output, often in the presence of extensive noise. The methodology in this research provides a framework for increasing the robustness and resolution of such systems.

---

[1] Although the Fourier-based cross-correlation is far more common, many still employ a direct correlation via the formal definition of the convolution, which is far more computationally expensive. This method is addressed further in Chapter 4.

Digital Particle Image Velocimetry (DPIV) has matured into a well characterized and versatile measurement technique over the past 20 years. However, there still exist some fundamental limitations to this technique. As we extend DPIV toward more complex flow systems, the systematic errors associated with the digital processing algorithms become more dominant and ultimately limit the measurement capabilities.

DIGITAL IMAGE EVALUATION TECHNIQUES

Early PIV systems used photographic recordings because of the high spatial resolution (which currently still provides the highest resolution). However, the low frame rate of photographic recordings limited higher speed flows to multiple pulsed PIV systems [8], which rely upon the autocorrelation of single images. Digital Particle Image Velocimetry (DPIV) showed that lower resolution high-speed cameras were still capable of a high accuracy velocity estimates, using digital processing algorithms [2]. These basic algorithms used a cross-correlation between a windowed image pair, from which a Gaussian curve was fit to the maximum correlation peak to obtain subpixel resolution on the image displacement estimate.

THE FOURIER-BASED CROSS-CORRELATION

The use of Fast Fourier Transforms allows for calculating the cross-correlation quickly due to the orthogonal nature of the frequencies. However, the estimation of spectral content is limited by pixel discretization and the finite signal length. As a result, the Fourier cross-correlation contains errors associated with the estimation of spectral content as well as aliased signals due to the periodic boundary condition. Alternatively, the cross-correlation can be directly evaluated within the spatial domain, but is computationally expensive, and suffers from loss of correlation effects [9].

SUBPIXEL ESTIMATORS

The use of curve fitting on the correlation peak to obtain subpixel resolution was the critical step that allowed for a digital implementation of PIV. The PIV measurement accuracy changed from the 1/2 pixel discretization error up to $1/100^{th}$ of a pixel. Generally, the subpixel fit is taken over the primary correlation value and the neighboring pixel on either side. The most common three-point estimators are the centroid fit, parabolic fit and Gaussian fit [8]. The Gaussian fit is generally accepted as the standard estimator due to the near Gaussian shape of the particle images (since the correlation of 2 Gaussian signals is also a Gaussian).

MEASUREMENT ACCURACY AND VALID VECTOR DETECTION PROBABILITY

The measurement accuracy of the DPIV processing algorithm can be defined by three metrics: the valid vector detection (VVD) probability, biases, and RMS errors. The most crucial factor in designing a digital algorithm is the VVD probability [7], which is a direct reflection of the signal strength. The VVD probability should be on the order of 90-100% to provide an accurate measurement. Invalid measurements, subsequently known as erroneous vectors, occur when the correlation is dominated by noise, aliasing, background correlation, or autocorrelations. The erroneous vectors are removed through the process of validation, for which there is also extensive method development [10]. However, validation does not provide additional

measurements and cannot be used to recover large sections where the PIV correlation has failed. Therefore it is critical to maintain a high VVD probability, which requires a very carefully designed and executed experiment and processing schemes that are robust against varying image quality.

The VVD probability is most often characterized using the effective number of correlated image pairs [11]. This metric is directly related to the image density, window size, image displacement, and out-of-plane motion. Naturally, other factors can compromise the VVD probability such as background signals and image noise. Provided the measurement is valid, the accuracy is then described by bias and RMS errors. The measurement RMS is a reflection of the noise floor, which is clearly related to the image noise and signal density. Several sources of bias error can exist within the measurement.

The loss of correlation bias is a result of the decreased energy density at higher image shifts. As a result, the correlation peak is biased to lower image displacements [8]. Peak-locking is a discretization error in which measurements are biased toward integer pixel values [8]. This effect is more dominant for smaller particle-image diameters and increased image noise. Fluid acceleration can introduce a bias due to the assumption of constant velocity between image pairs. Similarly, shear and rotation within the velocity field can lead to large biases due to the assumption of constant displacement over the window domain. The iterative correlation methods described above have shown substantial capabilities to minimize these errors.

While erroneous vectors and RMS are relatively easy to detect through a visual inspection of the vector fields, these bias errors can be more subtle. Therefore, it is important to quantify these errors and to remove them when possible.

ITERATIVE CORRELATION METHODS

One solution to increase measurement accuracy is through the use of iterative correlation methods. Taking advantage of the low subpixel displacement error [8], the initial correlation can be used as an estimate of the displacement. This estimate is used as a discrete window offset (DWO), from which the second correlation is ideally less than one pixel displacement [3]. Several extensions have been made to this technique through the use of iterative image deformations [4]. The simplest extension is the continuous window offset, from which the windows must be interpolated onto another grid on a second pass. Further improvements to this technique utilize higher order estimates from the initial velocity estimate to deform the window into a nonrectangular domain, first developed as the Particle Image Distortion (PID) technique [12]. Although these methods have been shown to have a substantially better estimate over the standard DWO, they are substantially more computationally expensive and are heavily dependent upon the interpolation scheme used to return to a rectilinear grid.

AVERAGING TECHNIQUES

Another common method to increase the measurement accuracy is through the use of averaging. This has the obvious detriment of a loss of either spatial or temporal resolution. However, this process is often necessary in order to obtain a high enough VVD probability. One method of temporal averaging is to use an average of the vector values obtained over several realizations. While this method is often attributed as the lowest accuracy averaging technique, it is often used for its computational simplicity. Alternatively, the ensemble averaged correlation is able to resolved the correlation peak in the presence of far more dominant noise sources, provided enough averaging is used [5]. A simple extension to this technique is the Single Pixel

Ensemble (SPE) evaluation [6]. The same correlation strength is shown to be present in a single pixel averaged over multiple realizations as a larger domain taken over one realization. However, these techniques rely upon steady or periodic flows, and all temporal information is lost. The Hart correlation [13], is a method of spatial averaging, whereby the correlation is averaged over one pixel offsets in order to obtain a stronger correlation strength. Although there is a very small loss of spatial resolution with this technique, it is far more computationally expensive, requiring 9 times the processing time. In addition, the Hart correlation is very robust against white noise and random errors, but offers no benefit in the presence of background or autocorrelation signals.

FILTERING TECHNIQUES

Several techniques have attempted to provide a stronger correlation through the use of filters. A normalized cross-correlation was introduced in order to account for the average signal-to-noise ratio within an image pair [14]. This technique showed a reduction in RMS and bias errors, but the performance decreases substantially for noisy images since the normalization could only be computed using the average image signal. The generalized cross-correlation (GCC) was introduced as a dynamic filtering technique which greatly increased the signal to noise ratio and was qualitatively demonstrated to be more robust against background noise [15].

EXPERIMENTAL DESIGN CONSIDERATIONS

With a carefully designed experiment, the standard digital evaluation algorithms are quite sufficient to provide accurate measurements. However, more complex flow systems do not always allow one to design the experiment around the digital evaluation techniques. In order to design a more robust estimator, it is important to characterize the development of digital processing errors as the experimental design approaches the measurement limitations.

FLOW TRACERS

When selecting tracer particles for a DPIV experiment, identifying the proper particle size is essential. From a physical perspective, this means ensuring that the particles accurately follow the fluid motion. Specifically, the particles should match the density of the fluid to minimize the buoyancy forces. Also, the particles should be small enough to minimize the time scale of the particle response to flow fluctuations, characterized by a Stokes number approximating zero.

From an imaging perspective, the particle-image diameter will be determined by the magnification of the optical system [16]. The magnification is generally defined for the experimental setup in order to resolve the desired details of the flowfield. In order to maximize the correlation between images, the windowed image should contain a sufficient number of well illuminated particles [11]. Scattering theory dictates that larger particles are able to scatter more of the incident light, increasing their illumination. Additionally, the particle-image diameter influences the correlation by modifying the spectral content of the displacement signal [17].

As we will subsequently demonstrate, the increase in particle-image diameter results in a shift in the spectral content of the images, which in turn compromises higher frequency components. When the diameter becomes too small, the pixel discretization of the particle-image can no longer preserve the true particle shape, and much of the spectral content of the particles is again lost. As a result, there is an optimal particle-image diameter range imposed by the correlation technique, usually falling between 2 and 3 pixels [8].

Taking these factors into account, the particle size must be selected to balance the physical requirements and the imaging effects. However, the physical considerations almost always dominate and the loss of measurement accuracy is taken as a necessary effect.

HIGH-SPEED TIME RESOLVED DPIV (TRDPIV)

High-speed Time Resolved DPIV (TRDPIV) images are often degraded due to experimental limitations. One of the most dominant limitations is the selection of tracer particles. Submicron particles are generally required in order to resolve the high levels of turbulence in the flow. The size requirement is far more dominant for gas flows, due to the substantial difference in densities, increasing the buoyancy force. The imaging of these tracers is therefore given by a subpixel image diameter, which leads to substantial increases in the RMS error [8].

Another consequence of the subpixel particle-image diameter is the peak-locking effect [8]. Peak-locking is a discretization error which is most dominant for subpixel diameters due to an undersampling of the particle-image. As a result, the displacement estimate is shown to be biased toward integer pixel values. Additive noise also serves to amplify this effect. Although there have been several correction techniques proposed to minimize this error [18-20], the variety of sources that effect peak locking limit the robustness of these methods. The root of the problem stems from modeling the image recording process, which is further complicated due to the variety of imaging devices.

High-speed TRDPIV also requires very high repetition rates on both the camera and laser systems, typically limited to the kHz rate (ms image or pulse separation). The use of dual-head lasers and frame-straddled cameras allows one to capture two image frames with μs image separation. For supersonic flows, the limit on the imaging rate leads to large displacement between image pairs. In addition, because of large time separation, the fluid acceleration, curvature, shear, and out-of-plane motion are all amplified. The combination of these factors can lead to substantial errors in the velocity estimation, far greater than the generally quoted 0.1 pixel accuracy [8], and even degrade the correlation strength to the point where valid measurements are no longer possible.

MICRO-PIV

Micro-DPIV (μPIV), an extension of traditional PIV through the use of a microscope to achieve higher spatial resolution [21], is subject to many of these challenges. For high magnification, several limitations are introduced which compromise the image quality and therefore the flow measurement. Achieving high spatial resolution requires that tracer particles are more densely seeded, which is not always possible. In addition, the measurement volume of PIV is smaller than the minimum thickness of a laser plane. As a result, μPIV experiments must illuminate the entire volume and rely upon the focal depth of the camera to provide images within the desired measurement plane of the illuminated in-focus particles. The illuminated out-of-focus particles contribute to a background glow which greatly increases the image noise [22]. The added noise from the volume illumination effect often limits the density of tracer particles within the flow. Conversely, the measurement accuracy of the digital processing algorithms drops substantially as the number of correlated particles within an image decreases [11].

Due to the optical resolution, the micron sized tracer particles create large particle-image diameter images. The measurement accuracy of standard DPIV interrogation algorithms is optimal around 2.5 pixels, and increases considerably for larger diameters [8]. In contrast, if

smaller physical sizes are selected for tracer particles to meet the optimal particle-image diameter, the position of the particle motion is no longer reliable. For sub-micron particles, Brownian motion can become a dominant factor in the particle [23]. This Brownian motion effect is more dominant for small pulse separations between images, limiting the temporal resolution and the dynamic range of velocity measurements.

The images are further compromised by optical aberrations and decreased numerical apertures of the high resolution optics. These combined factors lead to considerable losses in the correlation strength between image pairs. Resolving these measurements accurately requires the use of an ensemble correlation [5]. With the ensemble correlation, the system is capable of only time-averaged measurements. A time-resolved μDPIV methodology has not yet been demonstrated due to the optical and experimental limitations described above.

Artificial image simulations, or Monte Carlo simulations, are used as a tool to evaluate the systematic error generated in DPIV as a result of the digital processing algorithms. In this method, images are generated with a known velocity distribution, particle-image, time separation, background noise, etc. from which the accuracy of the DPIV velocity estimation can be evaluated [24]. Ultimately, the performance of a digital processing algorithm must be evaluated with actual DPIV image data. However, this method provides a useful tool to describe the way in which these systematic errors occur and to quantify them.

The most critical aspect of artificial image simulations is the generation of the particle-image. To a good approximation, the Fraunhofer diffraction pattern, which governs the particle intensity distribution, can be approximated by a Gaussian function, with a width of d_e:

$$I = I_0 \exp\left(\frac{-8x^2}{d_e^2}\right) \qquad (1)$$

where the particle image diameter corresponds to 4 standard deviations [8]. Because of its simplicity, this model is used to describe the particle-image intensity distribution. The digital particle-image signal is obtained by integrating the intensity distribution over the pixel area. However, this integral is dependent upon the manufacturing of the camera, where the fill factor and quantum efficiency will drastically affect the particle-image [8]. For the simplicity of these simulations, it is sufficient to discretize the Gaussian distribution of particle-image intensity, sampled at each pixel value.

In addition to the shape of the particle-images, the maximum intensity value, I_0, must be modeled as a function of position along the optical axis. This energy distribution is often modeled as either a Gaussian or Tophat (uniform intensity distribution) function, where most laser energy distributions fall roughly between the two [25]. For this analysis, the laser energy distribution is assumed to be Gaussian, and the source density is to include particles within 95% of the cumulative energy.

In addition, these simulations used the maximum exposure discretized into an 8 bit/pixel image. Thermal noise was modeled as additive Gaussian noise at a given standard deviation of noise. Unless otherwise specified, the images were displaced uniformly, and there was no out-of-plane motion. These assumptions allow us to quantify the dependence of correlation estimators upon specific image conditions, but do not represent the actual performance in most experimental conditions. In order to more fully evaluate these techniques under a more rigorous simulation, the Japanese Standard Image Project [26] and PIV Challenge image simulations [27, 28] are subsequently examined.

The Fourier based cross-correlation is the standard method used to evaluate the particle pattern translation in DPIV images, approximating the direct spatial cross-correlation between a windowed image pair. The phase correlation is another technique which uses a phase-transform filter of the generalized cross-correlation [29]. This technique filters the cross-spectrum with the inverse of its magnitude, leaving only the phase correlation matrix of the cross-spectrum.

THE GENERALIZED CROSS CORRELATION

The Generalized Cross-Correlation (GCC) has successfully demonstrated its potential for robust time delay estimation in SONAR applications using finite length, noisy signals [30]. More recently, the GCC was introduced as a possible technique for processing low signal-to-noise ratio DPIV images [15].

The GCC is designed to provide a smoothed cross-correlation through the use of dynamic filter operations constructed from the autospectra of the input signals. These filters are necessary in order to reduce the errors in the estimation of the cross-correlation for finite length signals. The implementation of these filtering functions for the cross-correlation is presented in Figure 1. Because of the linearity of the cross-correlation operation, the filter functions can be described as either prefiltering operations or smoothing operations on the cross-correlation. Subsequently, the smoothed cross-correlation implementation will be used to describe the GCC operations due to its simplicity.
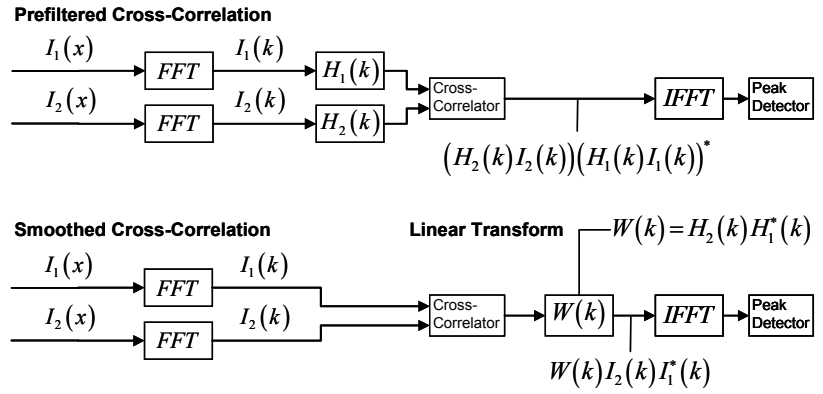
Figure 1 GCC Filter Implementations. H is a prewhitening filter applied before the correlation. W is a smoothing filter applied after the cross correlation

Previous research has resulted in the development of a range of functions offering various advantages and capabilities [29]. The GCC filters are functions of the power spectrums of the input signals, making them dynamic filters. In order to define these filters, we first define the power spectrums:

$$C_{ij} = \left| I_i(k) I_j^*(k) \right| \tag{2}$$

where C_{11} and C_{22} are the auto-power spectrums and C_{12} is the magnitude of the cross-power spectrum. One should note that these power spectrums are a reflection of the energy content within the two signals, whereas the phase content of their Fourier transforms carries the displacement information that we are attempting to estimate. In addition, these power spectrums can only be estimated as described in Equation 2. The actual power spectrums of these signals can only be defined with an infinite time series, and some assumptions may be required of the

input signals in order to implement these filters, namely that Equation 2 is an accurate representation of the infinite time-series signal.

Several of the more notable filters are described by the following Equations:

$$W_{Roth} = \frac{1}{C_{11}(k)} \tag{3}$$

$$W_{SCOT} = \frac{1}{\sqrt{C_{11}(k)C_{22}(k)}} \tag{4}$$

$$W_{PHAT} = \frac{1}{C_{12}(k)} \tag{5}$$

$$W_{ML} = \frac{1}{C_{12}(k)}\left(\frac{\gamma_{12}(k)}{1-\gamma_{12}(k)}\right) \tag{6}$$

$$\gamma_{12}(k) = \frac{C_{12}^2}{C_{11}C_{22}} = \frac{W_{PHAT}^2}{W_{SCOT}^2} \tag{7}$$

where W is the filter applied to the smoothed cross-correlation in Figure 1. All of these filters attempt to weight the cross-spectra inversely with some form of the input energy. This prevents high energy background signals from dominating the cross-spectrum.

The Roth filter in Equation 3 is the most basic of the GCC filters which assumes that the power spectrum of the first window can be used to represent the energy content of both input signals. The SCOT filter (Smoothed Coherence Transform), given in Equation 4, provides a better estimation of the energy content in the cross-spectrum by taking into account the energy content of both input signals. However, this window requires estimating two autospectrums, whereas the Roth filter required only one autospectrum.

The PHAT filter (Phase Transform), has the effect of removing all energy content from the cross-spectrum. As a result, the conversion of the smoothed cross-correlation back to the spatial domain produces a very sharp correlation peak, similar to a delta function. Using the estimation of the power spectrum from Equation 2, and for a linear cross-correlator, the PHAT filter is computationally equivalent to the SCOT filter, while only requiring one calculation of the cross-power spectrum. Although the PHAT filter produces a very sharp correlation peak, indicative of a strong SNR, it may be very sensitive to additive noise [29].

The ML filter (Maximum Likelihood) filter attempts to improve the PHAT filter by weighting frequencies according to their magnitude-squared coherence, given in Equation 7. This effectively weights the phase content of the cross-spectrum with the inverse of its standard deviation. This filter is shown to be Cramer-Rao bound on error [31].

However, all of these filters rely on an accurate estimation of the power spectra of the input signals, which is often not possible. In order to acquire a more accurate representation of the power spectra, one would require multiple realizations of the input signals, through both windowing and averaging. However, this would lose the dynamic nature of the filtering operation, reducing the GCC to an average filter technique. An alternative technique is to model the input signals, while retaining the dynamic nature of the GCC filters. This will be the subject of the Robust Phase Correlation (RPC) presented in Chapter 5.

Using the definition of the Power Spectra from Equation 2, the SCOT filter, the PHAT filter, and the ML filter are all equivalent. The equivalence of the SCOT and PHAT filters can be easily derived through:

$$W_{SCOT} = \frac{1}{\sqrt{I_1(k)I_1^*(k)I_2(k)I_2^*(k)}}$$

$$= \frac{1}{\sqrt{\left(I_1^*(k)I_2(k)\right)\left(I_1(k)I_2^*(k)\right)}}$$

$$= \frac{1}{\sqrt{\left|I_1^*(k)I_2(k)\right|^2}}$$

$$= \frac{1}{\left|I_1^*(k)I_2(k)\right|} = W_{PHAT}$$

(8)

and due to their equivalence, Equation 7 is unity. The maximum likelihood estimator is therefore the PHAT filter multiplied by a constant, which goes to infinity. However, because the coherence is constant with respect to frequency, the multiplication of a constant will not effect the cross-spectrum peak detection, and can therefore be removed. The three filters given in Equations 4, 5, and 6 are then equivalent, where each filters the cross-spectrum between the signals by removing the magnitude content, leaving only a matrix of phase differences between the signals. This will therefore be referred to as the PHAT filtering operation, which produces the Phase Correlation Matrix (PCM).

There are several unique properties of the phase correlation matrix. Primarily, because the energy in each frequency is unity, the phase correlation is more robust against varying energy content within the signals. Additionally, there is a close relationship between this phase correlation and the Fourier shift theorem, given as:

$$I(x) \otimes I(x+d_x) = \partial(x-d_x) = \exp\left(ik\left(\frac{2\pi x_0}{L_x}\right)\right)$$

(9)

where d_0 is a linear shift of the input signal, I, and L_x is the length of the Fourier transform. This is the ideal correlation from DPIV measurements, where the second image window is a linear shift of the first image window. The correlation is a delta function centered around x_0, and its Fourier transform is a uniform magnitude function with linear phase shifts [32]. The slope of the linear phase shift is known as the group delay, which is directly proportional to the image displacement. This property can be exploited to directly estimate the displacement in the Fourier domain, discussed further in Chapter 6.

The PHAT filter transforms the cross-spectrum into an approximation of Equation 7, giving the phase correlation a very sharp correlation peak. The phase correlation is often very sensitive to additive noise [30]. This is because certain frequencies with very low SNR are given the same weighting as high signal content windows. The aim of the different GCC filters is to minimize the effect of these frequencies. However, as stated earlier, this requires a more precise

estimation of the power spectra or some a priori knowledge of the signals that comprise the inputs.

Despite its sensitivity to noise, the phase transform is useful to remove the dependence of the cross-spectrum upon the energy content of the windows. Using the PHAT filter will provide a smoothed cross-correlation which is far more robust to the input signals and easier to manipulate with further filters designed to attenuate the negative effects of additive noise.

Although the cross-correlation is relatively easy to implement in the Fourier domain, pixel discretization and finite window lengths introduce errors within the estimation of the spectral content, inhibiting the performance of Fourier based cross-correlation. Additionally, the periodic boundary condition leads to aliasing effects, which severely compromise measurements at higher image shifts.

ALIASING AND SPECTRAL LEAKAGE

Aliasing is a result of the periodicity of the discrete Fourier transform for finite signals, which assumes the signal to be periodic over its length, *L*. The periodic replications, outside the *-L/2* to *L/2* domain, can overlap with the valid regions of the signal to create an aliased signal in the correlation. These aliased signals can be seen as the correlation of particles exiting image window 1 with particles entering image window 2.

The second concern with the Fourier based cross-correlation is spectral leakage. When wavenumbers higher than the sampling rate are present within the image signal, these values are aliased onto lower spatial modes. One should remember that spectral leakage aliasing is different from the wraparound aliasing that occurs as a result of the periodicity of the discrete Fourier transform. The wraparound errors are a sort of spatial domain aliasing, where as spectral leakage is aliasing within the spectral domain [1]. This spectral leakage contributes added noise to the estimation of the Fourier transform of the image window.

One cause of spectral leakage is an oversampling of the image signal. For instance, when the particle-image diameter becomes sufficiently small, the higher frequency energy is aliased. A second source of spectral leakage is a discontinuity at the window boundary. The endpoints of the signal can cause a discontinuous jump for the periodic signal. The discontinuity is described by high frequency modes, which again leads to spectral leakage. The spectral leakage of a discontinuous jump is known as the Gibbs phenomenon, which can be observed as a ringing effect near the discontinuity [33]. The additive noise from these discontinuities can severely compromise the estimation of the cross-correlation, especially for particles on the edge of the boundary.

ZEROPADDING

Zeropadding is a windowing technique used to remove the wraparound aliasing caused by the periodic boundary condition. Zeropadding buffers the rectangular $L_x x L_y$ window with zeros such that the zeropadded window is $2L_x x 2L_y$. By zeropadding the rectangular window, the periodic boundary condition is pushed outward so that the overlapping regions of the circular correlation are removed.

A second effect of the zeropadded window is the discontinuity at the image boundaries. The discontinuity is far more severe when a nonzero background signal is present. One method to minimize the discontinuity caused by a nonzero background signal is through background subtraction and byte scaling [34]. In this method, the image is scaled such that the discontinuity is minimal. However, this method is not able to compensate for particles located at the edge of the window domain. Therefore, while the wraparound aliasing can be removed, the spectral leakage can become a substantial detriment to the cross-correlation.

While computationally efficient, the use of Fourier transforms for the cross-correlation estimation is a detour to the formal definition of the direct cross-correlation [1]. Alternatively, with the use of zeropadding, the Fourier based cross-correlation is computationally equivalent to the direct correlation[2], since the periodic boundary effects have been removed.

THE EXTENDED WINDOW DOMAIN MASK

Ideally, the window should be selected so that the wraparound aliasing is removed, while at the same time providing a smooth transition at the image boundaries, so that the spectral leakage is minimized. In order to retain the benefits of zeropadding without the degradation of the spectral content, a new windowing technique is now introduced. This window is based on the zeropadded window, but with a smooth transition to attenuate the spectral leakage. Instead of buffering the window with zeros, zeropadding can alternatively be thought of as an extended window domain with periodic image replications, of size $2L_x x2L_y$, which is then masked with a tophat profile. The spectral leakage can then be attenuated by using a smooth tapering window as an alternative to the tophat window.

While a variety of tapering functions can be implemented in this method, we examine the use of the Blackman apodization function, a common window with very good spectral characteristics [35]. Apodization functions, which go to zero at the boundaries, are required to remove the discontinuity at the image boundary. The extended Blackman window is presented in Figure 2. This window can be thought of as a tradeoff between wraparound aliasing and spectral leakage. The extended Blackman mask attenuates wraparound aliasing while preserving the spectral content of the image. However, because this mask does not eliminate all of the periodic signal, wraparound aliasing can be a significant source of error for high image shifts where the overlap is more dominant.

THE LOSS OF CORRELATION ERROR AND WINDOW NORMALIZATION

The loss of correlation at increased image shifts, known as the out-of-pattern effect [8], results in a decreased signal to noise ratio. As a result the maximum correlation at a particular image shift, *(x,y)*, is limited by:

$$C(x, y) = \frac{(L_x - x) \cdot (L_y - y)}{L_x \cdot L_y} \tag{10}$$

where *C(x,y)* is normalized to the maximum correlation at zero image displacement [34]. The periodic boundary condition of the Fourier based cross-correlation causes the overlapped regions to correlate, creating aliased signals. Although the strength of the true correlation is theoretically bounded by Equation 10, the maximum value of the Fourier based cross-correlation is uniform over all displacements due to the aliased signals. As a result of the loss of correlation, a correction to the correlation peak is often implemented to reduce bias errors based on Equation 10 [8, 34].

---

[2] While theoretically equivalent to the direct correlation, the zeropadded Fourier based cross-correlation will include the round-off errors associated with the discrete Fourier transforms. However, these round-off errors are substantially lower than the discretization error of a typical digital image.
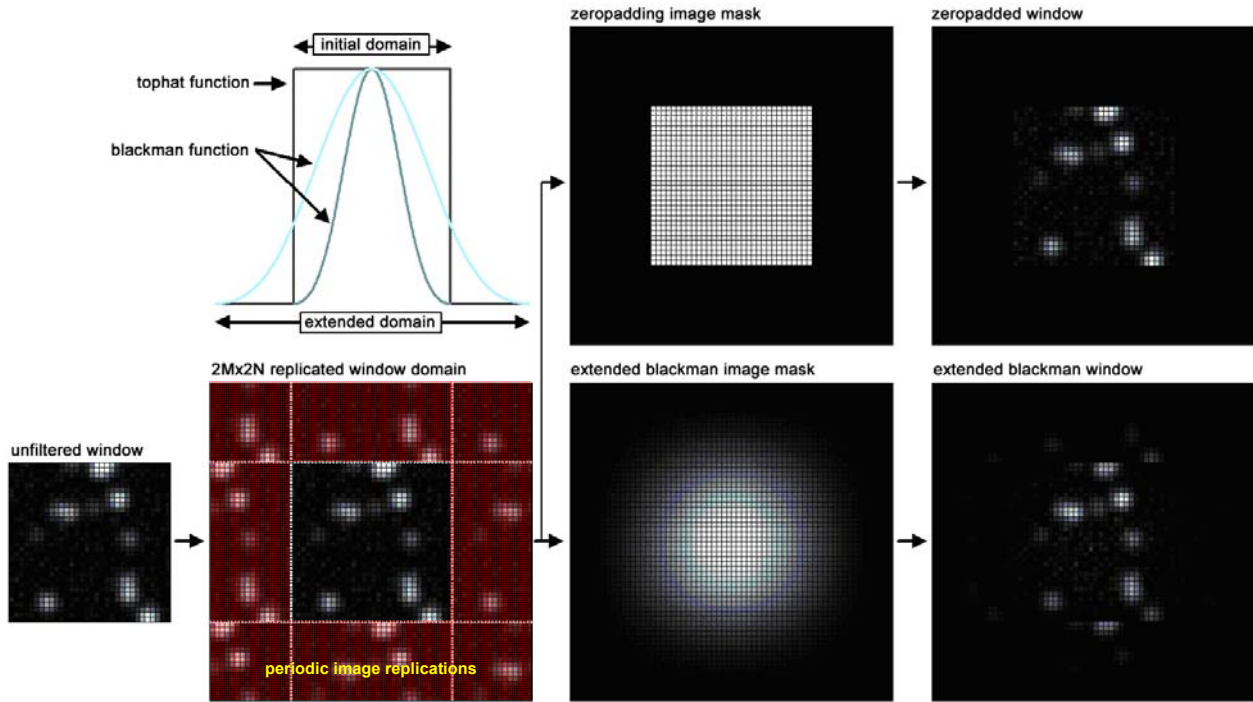
Figure 2   Representation of various spatial filtering techniques.

RECTANGULAR AND ZEROPADDED WINDOWS

When image masking is used, such as zeropadding, the energy content of the periodic signal is modified. As a result of zeropadding, the periodic contributions to the correlation are removed, from which the correlation peak is no longer biased. However, the correlation plane now resembles the triangular shape given by Equation 10. The displacement peak, while no longer affected by the aliased contributions, may become secondary to lower displacement peaks where the maximum correlation strength is greater. One method is to use background subtraction and byte scaling to minimize the triangular shape of the correlation plane. However, this method does not select the maximum correlation peak, but is rather biased to selecting lower displacement peaks. In order to identify the correct peak, it is often necessary to normalize the correlation, before the peak detection. One downside of this process is that larger image shifts then contain substantially lower signal strength, but a uniform distribution of random noise. This random noise is then amplified for the low SNR regimes, often dominating over the true correlation peak.

Of the two methods for measuring windowed displacements, both have their downsides. Without normalizing, the estimator is less robust to image conditions and will provide erroneous measurements when the particle-image signal is not ideal, even with background subtraction and byte scaling [34]. Alternatively, with normalization, the correlation is far more robust against background noise and is able to pick up larger displacements with low correlation strength, but can be highly compromised for increased levels of white noise due to the non-uniform SNR.

TAPERED AND NON-UNIFORM WINDOWS

The normalization, whether applied to the peak detection or subpixel estimation, is defined uniquely for the rectangular and zeropadded windows by Equation 10. When a tapering

window is used to window the image, the energy content of the periodic signal is distorted, so that Equation 10 no longer holds. Without a proper normalization, tapered windows are often subject to high bias errors for large displacements. However, Equation 10 is a specific set of the more general Equation for determining the maximum correlation strength. The maximum correlation for a given window, *W(x,y)*, can be defined by:

$$C(x,y) = W^0(x,y) \otimes W^0(x,y) \tag{11}$$

which is the autocorrelation of the zeropadded window, $W^0(x,y)$.

ARTIFICIAL IMAGE SIMULATIONS: THE WINDOWED PHASE CORRELATION

Ideally, we would like to keep the advantages of tapering or zeropadding, without the negative effects caused by the normalization requirement. As explained in Chapter 3, the phase correlation is useful in the presence of background signals and is relatively invariant to the energy content of the input signals. As such, the phase correlation appears ideal for removing these issues that arise in windowing.

Figure 3 presents the valid vector detection probability for the SCC and PHAT processors without the use of spatial filtering. Normally distributed (white) noise was introduced to the images, described as the ratio of noise standard deviation to the image dynamic range. Previous studies [15] qualitatively demonstrated that the PHAT is capable of significant attenuation of background noise and increased signal-to-noise ratio, but no detailed error analysis was presented. One would expect that these results would lead to increased valid vector detection. However, the analysis presented in this paper uncovers that the standard PHAT implementation, despite its promising outlook, actually has a reduced probability of valid vector detection. However, these are artificial images generated with minimal background noise, which was the key factor which the PHAT was able to minimize.

In order to further understand the poor performance of the PHAT, we must examine the details of the cross-correlation. The Fourier based cross-correlation contains both amplitude and phase information. The phase information contains the shift of spectral content between the images, whereas the amplitude information conveys the amount of energy within a particular frequency. The SCC is more biased toward the lower frequency content of the correlation, where the energy is dominant. The PHAT, however, contains only the phase information of the correlation. As a result, the PHAT correlation produces a very sharp pronounced peak [15], in contrast to the SCC which contains a broader correlation peak, surrounded by low frequency noise.
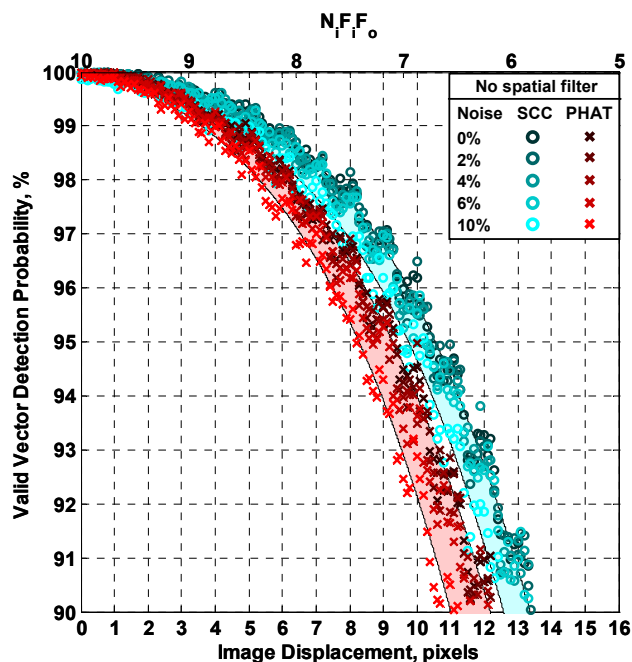
Figure 3 Valid Vector Detection Probability for the SCC and GCC processors presented without spatial filtering.

However, valid vector detection depends on the ability to resolve the displacement peak in the presence of false correlation peaks caused by uncorrelated noise and aliasing. While the PHAT is able to remove uniform or low spatial-frequency noise caused by light flare or uneven illumination, aliased signals and uncorrelated white noise often disrupt the phase content of the correlation, leading to erroneous displacement vectors. The SCC shows improved vector detection over the unfiltered PHAT due to its bias toward lower frequencies within the correlation.

The zeropadded correlations are presented in Figure 4. The valid vector detection probability is drastically different from the case without spatial filtering, presented in Figure 3. Substantial improvements are noted for the PHAT, extending the 90% valid vector detection probability from about 11.5 pixels (SCC – unfiltered) to 14 pixels. The SCC, however, shows a substantial reduction in vector detection probability. In addition, the SCC detection probability is much more sensitive to white noise within the image.

Spatial filtering has often met with mixed results when used for the SCC, for the reasons discussed earlier. Specifically, the SCC must be normalized to account for the loss of correlation, given by an explicit function applied to each window. However, DPIV contains sparse amounts of data, making the energy content non-uniform. As a result, the normalization function can only be described by the average loss of signal [14], making the SCC much more sensitive to white noise and the energy distribution within a window.

Since the PHAT contains phase-only data, each window is normalized with respect to its specific energy content. As a result, there is no average normalization required for the PHAT. The PHAT is therefore capable of dynamic signal filtering, making the processor

Figure 4   Valid Vector Detection Probability for the SCC and PHAT processors using the zeropadded filter.
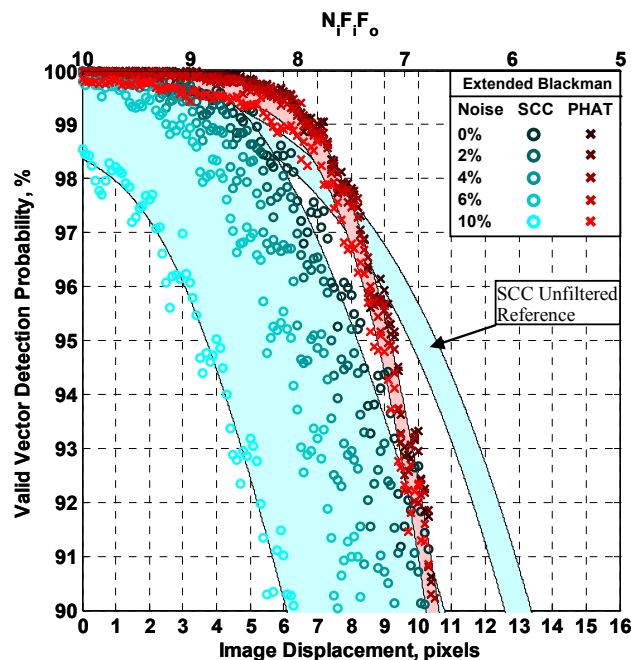
Figure 5   Valid Vector Detection Probability for the SCC and PHAT processors using the extended Blackman filter.

16

superior to cross-correlation normalization [14]. The limitations that prevent the use of spatial filtering for the SCC, through averaged normalization functions, are no longer present within the PHAT, increasing its versatility.

This behavior can also be observed in Figure 6, which shows the correlation plane for a given windowed image pair. Figure 6a and Figure 6b are the correlation planes for the unfiltered image pair. The PHAT shows a very strong correlation peak, where as the SCC shows a high level of background noise. In Figure 6c shows the zeropadded SCC window, which removes the aliased signals, but one can easily see the non-uniform correlation strength for varying image shifts, which would require normalization. Figure 6d gives the zeropadded PHAT window, which is immune to the background noise as well as the non-uniform correlation plane, removing the need to normalize. This illustrates the robustness of the phase correlation method to both image conditions and windowing techniques.

Since the PHAT has demonstrated the ability to better incorporate the use of spatial filtering, the extended Blackman filter is also examined. Figure 5 presents the valid vector detection probability of the Blackman spatial filter. Both the SCC and the PHAT show improved vector detection over the 0-8 pixel range and a quick decrease in vector detection probability after 8 pixels. However, the SCC is still subject to the same errors as with zeropadding, where vector detection is very sensitive to noise. Again, the PHAT is not subject to these errors, resulting in a dramatically improved performance, extending the valid vector detection probability of 99% from about 5 pixels (SCC – unfiltered) to about 7 pixels, and appears largely invariant to the white noise that compromises the filtered SCC. This immunity to noise is a crucial factor in providing a more robust correlation processor.

This behavior can be explained by examining the benefits of the extended apodization filters. The discontinuity within the zeropadded filter often contaminates the spectral content of the original window, outweighing the benefits of removing the aliasing. The extended Blackman filter provides a much smoother transition that is able to preserve spectral content while attenuating aliased signals, without excessive filtering of the initial signal domain. However, because of the design of the window, the aliasing is reduced, but not removed. As a result, two regimes are apparent for the Blackman filter. For image shifts below a quarter of the window size, aliasing is not a significant factor, and so preserving the frequency content appears more important. In contrast, image shifts above a quarter of the window size are subject to increased aliasing. As a result the benefits of the Blackman filter no longer dominate and the zeropadded filter provides better vector detection.

Figure 6   Correlation plane shown without normalization for: SCC unfiltered (a), PHAT unfiltered (b), SCC zeropadded (c), PHAT zeropadded (d).

In addition, one should remember that these images were generated to represent a well executed experiment. As such, there are no background signals or poor illumination which are common to DPIV experiments. For these real conditions, we expect the PHAT to perform even better than the SCC due to its immunity to background noise [15]. Subsequent experimentation confirms this behavior, presented in the experimental validation of Chapter 5.

The PHAT filter, as discussed in Chapter 3, is used to dynamically filter the cross-spectra of the two image windows into the phase correlation. However, it was noted that the phase correlation may be sensitive to additive noise in the input signals, due to the amplification of low SNR frequencies. More robust estimators, such as the ML filtered GCC, attempt to modify the phase correlation, applying a weighting to minimize low SNR frequencies. One of the primary problems with the ML filter is the requirement of an accurate estimation of the input power spectra, which is difficult to achieve for the finite length inputs.

A second method to attenuate these low SNR frequencies is by modeling the power spectra of the input signals. In order to define these modified filters, the image signals are divided into its signal and noise components:

$$I(x) = s(x) + n(x) \tag{12}$$

where s is the desired particle-image signal and n represents the sum of the noise signals. The desired smoothing filter to be applied to the cross-spectra is then given by:

$$W_{RPC} = \frac{1}{C_{12}(k)} \frac{|s_1(k)| \, |s_2(k)|}{|n_1(k)| \, |n_2(k)|} = \frac{1}{C_{12}(k)} \left(\frac{|s(k)|}{|n(k)|} \right)^2 \tag{13}$$

which is the product of the phase transform filter with the square of the SNR, assuming that the power spectrums of the two image windows are identical. Equation 13 is now introduced as the Robust Phase Correlation (RPC), where the PHAT filter is used to transform the cross-spectrum into the phase correlation, and then an energy filter, defined by an analytical model of the SNR, is used to mask the phase correlation matrix [36]. The combined effect of these filtering operations replaces the energy content of the cross-spectra with the theoretical energy distribution, making the correlation very robust to a variety of image conditions. The RPC estimator in Equation 13 subsequently follows the same path as the SCC estimator, using the IFFT followed by the subpixel estimator.

THE DPIV SIGNAL-TO-NOISE RATIO

The determination of an analytical Signal-to-Noise Ratio (SNR) is essential to defining the filter in Equation 13 for the RPC correlation. The SNR can be analytically defined by examining the spectral energy contributions from the particles and random noise signals. The spectral energy filter is given by the ratio of signal energy to that of the noise spectrum and aliased signals:

$$W(k) = \frac{E_{PIP}}{E_{aliased} + E_{quantized} + E_{FFT} + E_{noise}} = \left(\frac{|s(k)|}{|n(k)|} \right)^2 \tag{14}$$

where E_{PIP} and $E_{aliased}$ are functions of wavenumber. $E_{quantized}$, E_{FFT}, and E_{noise} signify uncorrelated random errors, and are constants. While other sources may exist for certain experimental conditions, these can also be appropriately filtered out by including additional terms in Equation 14.

THE PARTICLE IMAGE PATTERN (PIP)

In order to model the power spectrum of the Particle Image Pattern (PIP), we consider the image acquisition of the particle-images by the camera. The true intensity from a point source is given by the Fraunhofer diffraction pattern (in the absence of lens aberrations):

$$I(x) = I_0 \left(\frac{2J_1(x)}{x} \right)^2 \tag{15}$$

which is known as the Airy distribution, given by a central disk surrounded by rings [37]. The first minima of this function corresponds to the diffraction-limited diameter [38]:

$$d_{diff} = 2.44 f_\# (M+1) \lambda \tag{16}$$

giving the minimum resolvable size of an image for a given magnification, M, f-number, $f\#$, and wavelength, λ. When imaging a spherical particle, the diameter can be approximated using the effective particle diameter [39]:

$$d_e = \sqrt{\left(M d_p \right)^2 + d_{diff}^2} \tag{17}$$

where d_p gives the physical size of the particle. To a good approximation, the Airy distribution can be approximated by a Gaussian function, with a width of d_p:

$$I(x) = I_0 \exp\left(\frac{-8x^2}{d_e^2} \right) \tag{18}$$

where the particle image diameter corresponds to 4 standard deviations. Because of its simplicity, this model is used to describe the particle-image intensity distribution. Figure 7 shows a comparison of the actual intensity distribution of Equation 15, with the simplified Gaussian function in Equation 18.

The digital particle-image signal is obtained by integrating the intensity distribution over the pixel area. However, this integral is dependent upon the manufacturing of the camera, where the fill factor and quantum efficiency will drastically effect the particle-image. For the
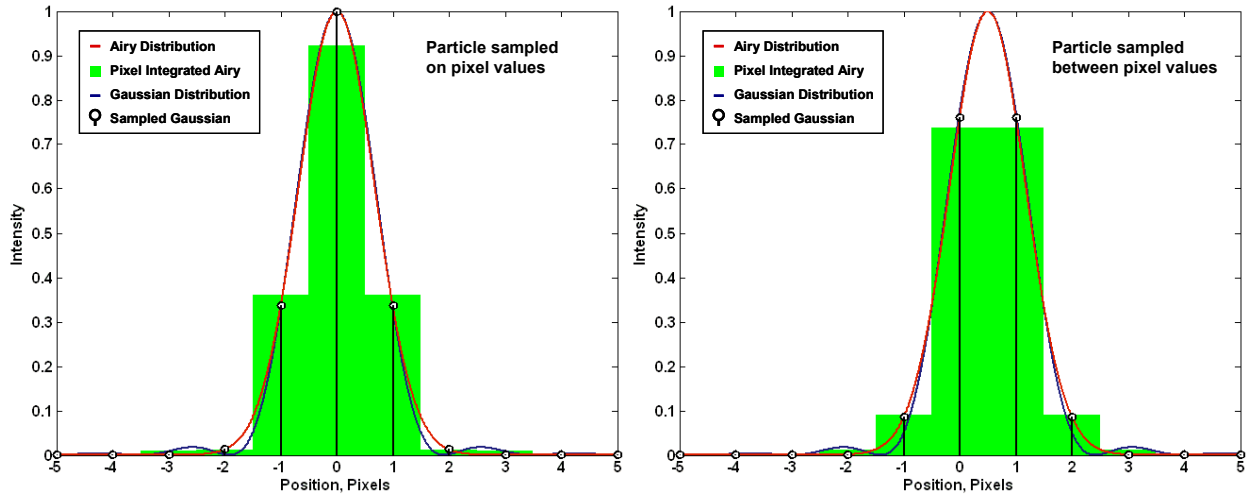
Figure 7   Particle-intensity image distributions with a particle-image diameter of 3 pixels.

simplicity of this model, it is sufficient to discretize the Gaussian distribution of particle-image intensity, sampled at each pixel value.

The energy spectrum of a particle-image is given by the autocorrelation of its intensity distribution, according to the Weiner-Khintchine theorem [33]. This can be analytically calculated using the Fourier transform of the Gaussian signal in Equation 18:

$$E_{particle} = \frac{\pi^2 I_0^2 d_e^4}{64} \exp\left(-\frac{d_e^2 k^2}{16}\right)$$

(19)

where k is the normalized wavenumber between $-\pi$ and π.

The PIP power spectrum will be dependent upon the actual distribution of particles within a window. As such, it cannot be generically defined without knowledge of the particle locations, which is exactly what the DPIV processor is trying to estimate. In order to simplify the estimation of the power spectrum, we use the model of the particle-image power spectrum defined in Equation 19:

$$E_{PIP} = \frac{N_i \pi^2 I_0^2 d_e^4}{64} \exp\left(-\frac{d_e^2 k^2}{16}\right)$$

(20)

where N_i is the source density, signifying the average number of particles within the image window. Equation 20 makes the assumption that the PIP is a linear superposition of the individual particle-images, which is not necessarily true. The wavefunctions of the light scattered from each particle may interfere causing interference fringes [37]. However, for most planar DPIV experiments, the particle density is sufficiently low that particle interference rarely occurs. Therefore, it is sufficient to assume that Equation 20 is an accurate approximation of the PIP. In the case of µPIV, the particle interference is more dominant and may need to be accounted for in order to optimize the SNR given in Equation 14.

PARTICLE-IMAGE ALIASING

Aliasing occurs in the analog to digital conversion process when the signal spectral content exceeds the Nyquist sampling frequency. For DPIV images, the spectral content is then limited by the resolution of the digital imaging system. Using the PIP power spectrum in Equation 20, the aliased spectral energy of the PIP will be given by:

$$E_{aliased} = \frac{\pi^2 I_0^2 d_e^4}{64} \exp\left(-\frac{d_e^2 (k \pm 2\pi)^2}{16}\right)$$

(21)

where the higher frequencies are folded about the Nyquist sampling frequency. Because the spectral content of a Gaussian particle is not bandlimited, some of the particle spectral content is always aliased onto the frequency range of the digital signal. However, for large particle-images, the majority of the spectral content is below the Nyquist sampling frequency, and the small amount of aliased energy is below the noise floor of the system. For smaller particle-images, the aliased energy becomes more dominant and will affect the measurement accuracy of the system.

It is important to note that there is a substantial assumption in Equation 21 regarding the recording process. Namely, since we have assumed the particle-image as a discretely sampled Gaussian, we have disregarded the effects of quantization efficiency, fill factors, and Poisson

statistics for the photons incident upon the imaging device. However, Equation 21 is used for simplicity to provide a reduced order model, which is subsequently shown to be dominated by random noise for most cases, making this approximation reasonable.

RANDOM NOISE SIGNALS

The image quantization error and FFT round-off error are caused by the discretization of the image into a finite number of intensity levels. The image quantization error is a nonlinear process which rounds or truncates the signal to the nearest quantized level. The spectral distribution of energy between quantized levels is related to the image acquisition process and will be controlled by fill factors, sensitivity, and coupling between neighboring pixels [40]. However, it is often adequate to model the uncertainty by a uniform distribution between quantized levels, for which the spectral energy per pixel is given by:

$$E_{quantization} = \frac{R^2 q^2}{12} = \frac{R^2 2^{-2B}}{12}$$

(22)

which is independent of frequency, for a given quantized range, R, and memory storage bits, B.

DFT round-off errors are closely related to the image quantization error. The uncertainty for each operation is given by Equation 22. The DFT error is then the compounded quantization errors generated by each operation. The number of operations for a DFT will depend upon the specific FFT algorithm implemented. For a radix-2 FFT algorithm, the errors are given by:

$$E_{FFT} = (N_o - 1)\frac{R^2 2^{-2B_M}}{12}$$

(23)

where N_o denotes the number of FFT operations and B_M is the number of bits in the mantissa of the storage format [33]. When images are processed, the discrete intensity levels are usually converted to a double precision memory format. Although the quantization errors are compounded by a series of operations, the double precision quantization error is much less than the image quantization error. Assuming a 32x32 pixel 8-bit image window converted to a double precision matrix with 52 bits stored in the mantissa, the FFT errors are 8 orders of magnitude smaller than the image quantization error. For this reason, the FFT errors can generally be neglected.

The third type of random error considered in this analysis is thermal noise, or dark current noise [8]. Thermal noise is caused by the random motion of electrons within the energy wells of the imaging device, which is indistinguishable from electrons triggered by the photoelectric effect. Thermal noise is highly dependent upon the temperature of the recording device, leading to the use of intercooled camera technology [8]. However, this noise is typically the most dominant source of random errors within DPIV images, typically on the order of 5% maximum exposure. The thermal noise energy spectrum will be simply represented as:

$$E_{noise} = \sigma_n^2$$

(24)

which will vary typically from 0-10% normally distributed "white" noise, at a given standard deviation, σ_n.

THE DPIV SIGNAL-TO-NOISE RATIO

Using these energy spectra defined in Equations 20, 21, 22, 23, and 24, the spectral energy filter in Equation 14 can be rewritten as:

$$W(k) = \frac{\exp\left(-\dfrac{d_e^2 k^2}{16}\right)}{\alpha(k) + \beta} \tag{25}$$

$$\alpha(k) = \exp\left(-\frac{d_e^2 (k \pm 2\pi)^2}{16}\right) \tag{26}$$

$$\beta = \frac{16 N_1 N_2 \left(R^2 2^{-2B} + \sigma_n^2\right)}{3\pi^2 N_i I_0^2 d_e^2} \tag{27}$$

where $\alpha(k)$ is a parameter describing the aliased spectral signal and β is a constant parameter describing the random errors. Although β is dependent upon numerous image properties which may not be clearly defined for a given image set, this constant can be easily tuned to a specific set of DPIV images to determine the optimal value. Furthermore, for $\beta \gg \alpha(k)$, the random errors dominate above any aliased signals, and the SNR reduces simply to a Gaussian function:

$$W(k) = \exp\left(-\frac{d_e^2 k^2}{16}\right) \tag{28}$$

where β is removed since it is a constant and will not effect the peak estimation.

Figure 8 illustrates several properties of the spectral energy filter from Equation 25. Figure 8b gives the 1D spectral energy filter as a function of diameter for a zero noise image window of 64x64 pixels with 10 Gaussian particles at the maximum 8 bit intensity. This instance provides the minimal noise and consequently shows the largest contributions that aliasing can have to the signal-to-noise ratio. The respective spectral energy filters for only random errors and only aliasing are shown in Figure 8a and Figure 8c. For diameters greater than 2.8 pixels, aliasing can be neglected and the spectral energy filter is dominated by random errors. In contrast, diameters below 1.9 pixels see only contributions from the aliased particle-image. The transition point between aliased errors and random errors is dependent upon the relation of α_k to β. For non-ideal images, σ_n will increase β, pushing the transition to smaller diameters. For images with diameters above this transition point, the weighting function reduces to Equation 28 which is only a function of the particle-image diameter, which greatly reduces the complexity of the spectral energy filter.
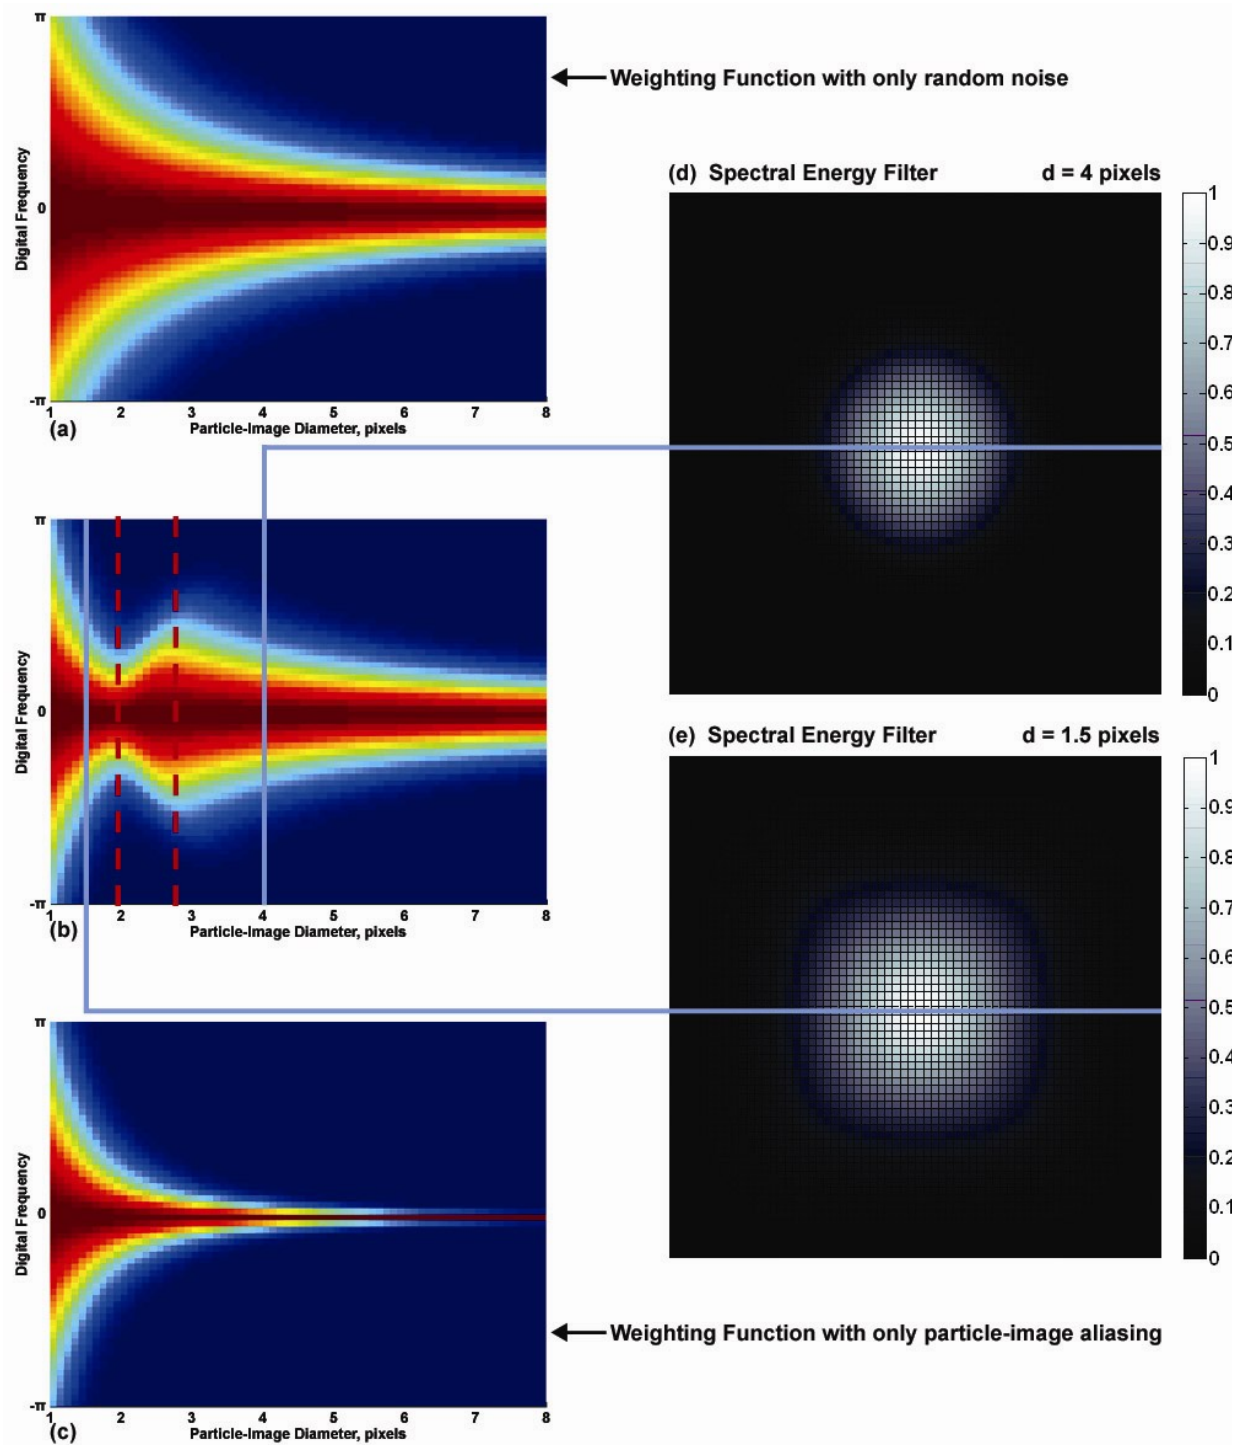
Figure 8   Spectral energy filter with aliasing and random errors.

The 2D spectral energy filters for 4 pixel and 1.5 pixel diameters are shown in Figure 8d and Figure 8e, respectively. Because the 4 pixel spectral weighting is above the transition, aliasing has no effect. The spectral energy filter is then given by Equation 28, which is a radially symmetric Fourier transformed Gaussian distribution The 1.5 pixel diameter spectrum is reflective of an aliasing dominated weighting. Because the DFT is taken over a rectangular domain, the aliasing is more dominant along the orthogonal window axes. As a result, the spectral energy filter is no longer radially symmetric. The measurement accuracy will be most directly affected by this asymmetry since standard 3-point estimators assume a Gaussian shaped correlation peak.

ARTIFICIAL IMAGE SIMULATIONS: PARTICLE-IMAGE DIAMETER PERFORMANCE

Since the spectral energy filter is primarily a function of diameter, it is important to understand the relation between measurement accuracy and the particle-image diameter. Figure 9 presents the RMS uncertainty for image pairs at a 0.3 pixel displacement, corresponding to the linear error regime with respect to displacement. The images contained no white noise. The unfiltered SCC provides a standard estimate of the measurement accuracy presented in other studies [8, 25]. The zeropadded image mask shows a substantial increase in RMS errors, especially for large particle image diameters. Although the zeropadded window has been shown to remove the biasing effects of aliasing, the RMS errors are substantially increased. The extended Blackman mask shows reductions in RMS consistent with previous studies due to the preservation of the spectral energy and reduction in window aliasing. However, image masking of the SCC has shown to severely compromise the valid vector detection probability [41].

The RMS errors of the RPC correlation are substantially reduced for all image masking methods at higher particle image diameters. Phase correlation has also been demonstrated to incorporate image masking without the spectral contamination present in the SCC [41]. This advantage is seen through the vast differences in the zeropadded and extended Blackman windows for the two correlation techniques.

At large particle-image diameters, the RMS reductions of the RPC correlation match the errors found for the SPC correlation [42]. For small diameters, the RPC error is reduced due to the incorporation of particle-image aliasing into the spectral energy filter. However, the RMS is still higher than the SCC for diameters less than 1.75 pixels. This increased rate of RMS for smaller
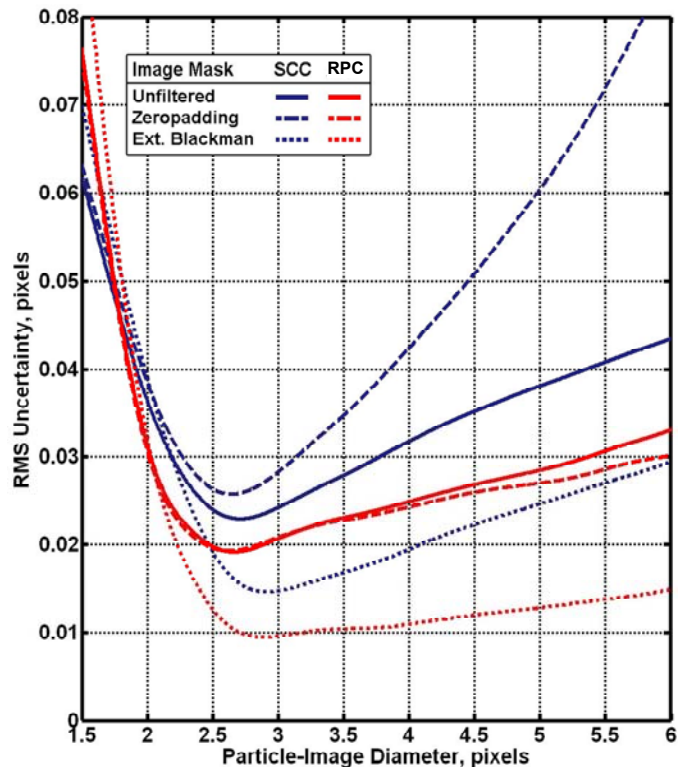
Figure 9   RMS uncertainty presented for the SCC and RPC correlation evaluated at a 0.3 pixel image displacement.

diameters is most likely caused by the loss of radial symmetry in the spectral energy filter. As a result, the subpixel resolution of the correlation peak cannot be as easily determined with Gaussian estimators.

The optimal particle image diameter for both the SCC and RPC correlation falls between 2.5 and 3 pixels. An examination of the spectral energy filter from Figure 8 reveals the cause of this minimum error. The optimal diameter corresponds to the local maximum distribution of spectral energy. For smaller diameters, the PIP is degraded by particle aliasing. For larger diameters, the spectral energy is contained within a narrow band of frequencies, where the higher spectral phase shifts are uncorrelated and contribute to the random errors. In addition, the predicted transition will move to smaller diameters for higher levels of white noise, shifting the position of the optimal particle size.

ARTIFICIAL IMAGE SIMULATIONS: PARTICLE-IMAGE DISPLACEMENT PERFORMANCE

The valid vector detection probability represents the most important aspect of DPIV data processing [7]. Furthermore, increasing the valid vector detection of the current correlation methods is essential in order to measure the temporal dynamics of microfluidic devices. Defining a valid vector probability requires first defining a validation method to identify erroneous vectors. For this analysis, erroneous vectors were defined as median absolute deviations greater than 0.5 pixels. Because these artificial images contain no noise, shear, or out-of-plane motion, the RMS errors are small enough to contain valid measurements below the 0.5 pixel threshold. The deviations were selected about the median displacement vector since it is more robust than the mean vector to outliers. In addition, large biases do not contribute to erroneous vectors as they would by selecting the true displacement vector.

Figure 10 presents the valid vector detection probability for the SCC and RPC correlation

Figure 10 Valid vector detection probability presented for the SCC and RPC correlation for several particle-image diameters. Images contain 8 particles / 32x32 pixel window with no noise. The RPC correlation was filtered with an extended Blackman image mask.

over a range of particle-image diameters at a density of 8 particles / 32x32 pixel window. The RPC correlation was filtered with the extended Blackman window image mask. The corresponding mask filtered SCC is not presented due to the substantially reduced valid vector detection [41]. Below about 0.5 pixels, the probability is 100% since the images contained uniform displacements with no image noise. For higher displacements, the out-of-pattern effect causes a loss of correlation which compromises the vector detection [8]. For all of the particle-image diameters, the RPC correlation shows a dramatic improvement, reducing the bad vectors by 50% on average. Examining the 6.4 pixel diameter case, the valid vector detection sees a drastic decrease after 1 pixel image displacements. In contrast, the RPC correlation provides a smoother transition, extending the 99% valid vector detection probability from 1 pixel to 4 pixel image displacement. For the 8 pixel image displacement, the error reductions for the RPC become less pronounced due to the properties of the extended Blackman window. Past a quarter of the window displacement, the extended Blackman window is no longer effective, and zeropadding must be employed [41].

Figure 11 presents the valid vector detection probability for the same images at a density of 4 pixels / 32x32 pixel window. Lower image densities at high particle-image diameters are more consistent with a typical micro-DPIV experiment. Because the seeding density is reduced by a factor of two, the average erroneous vectors increase by approximately a factor of two. In addition, the correlation strength is greatly reduced, causing the image aliasing and random errors to become more dominant. The large decrease in probability at 1 pixel image displacement for the SCC is far more pronounced at a lower image density. Again, the RPC correlation provides substantially improved vector detection, reducing the erroneous vectors by approximately 60% at a 2 pixel diameter. Even for the 8 pixel image displacement where the extended Blackman image mask losses its effectiveness, the RPC correlation still provides notable increases in the detection probability.
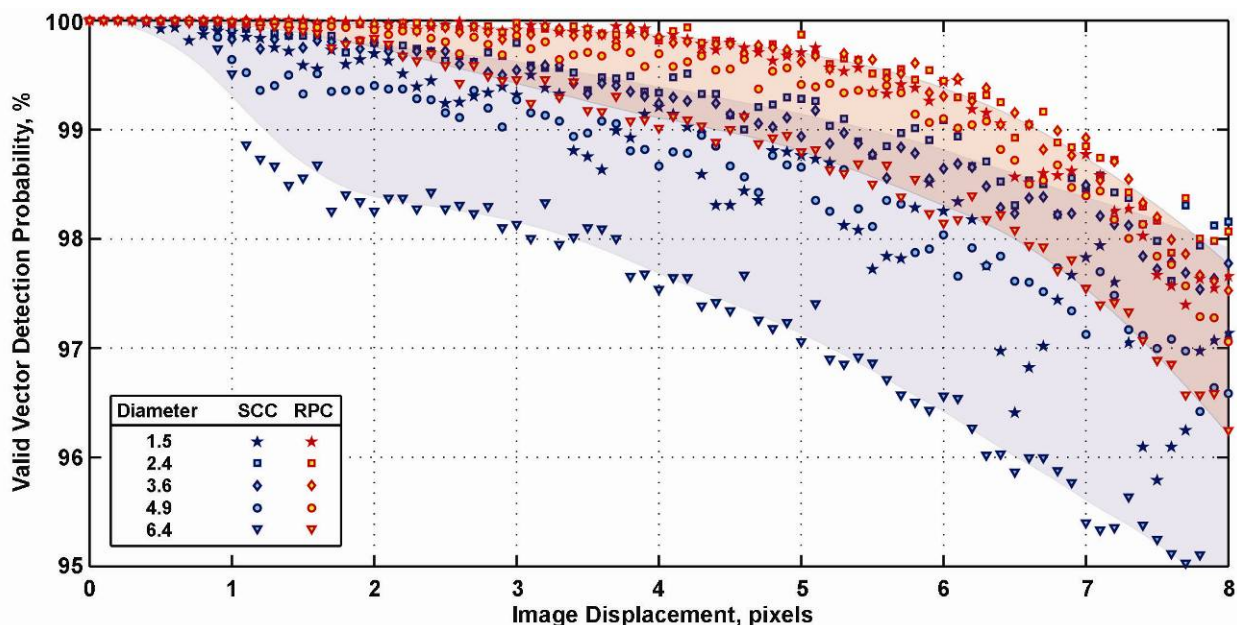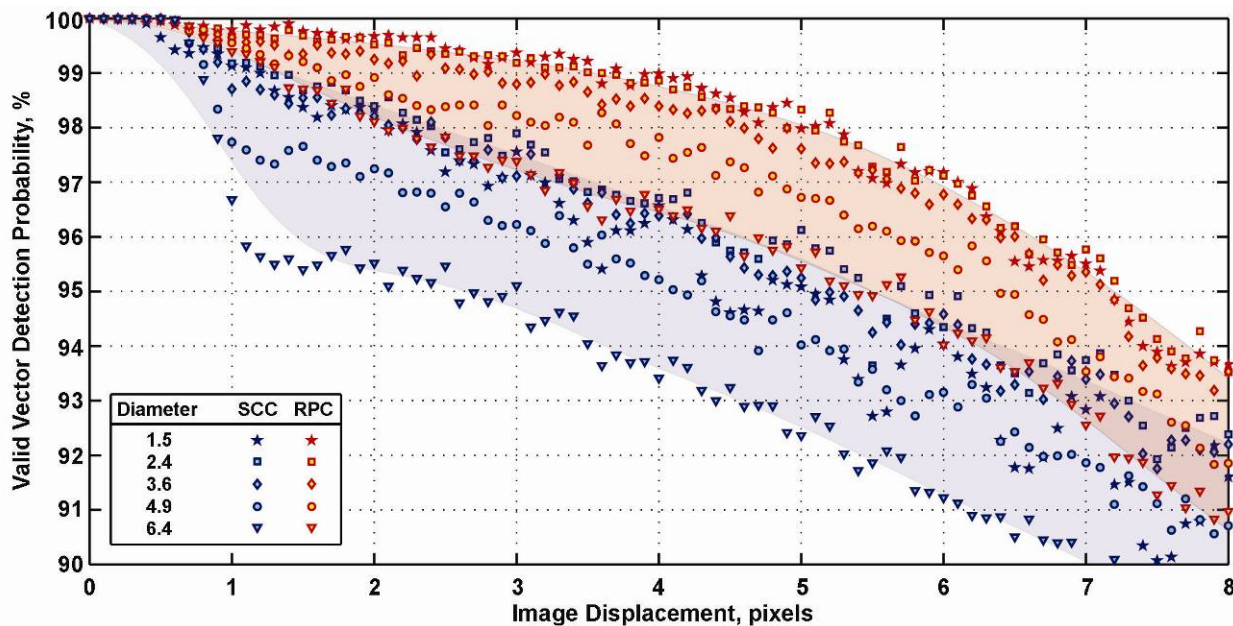

Figure 11 Valid vector detection probability presented for the SCC and RPC correlation for several particle-image diameters. Images contain 4 particles / 32x32 pixel window with no noise. The RPC correlation was filtered with an extended Blackman image mask.

Standard images were developed as a tool to be used in the evaluation of DPIV image processing techniques [26]. Unlike previously noted image conditions, these images examine the performance of an image processor for a variety of factors including shear, rotation, out-of-plane motion, and polydispersed particle-image sizes. However, these images do not include any thermal or background noise, which considerably affect the measurement accuracy of the image processing system.

The SCC and RPC are now applied to these standard images to evaluate their performance. The extended Blackman image mask was utilized with the RPC correlation. Images were correlated with a 64x64 pixel window followed by a 32x32 pixel window after a second-order discrete window offset. The spatial resolution was set to 8 pixels, corresponding to the resolution of the supplied actual displacement field. The large displacement standard images were correlated using an initial discrete image shift given by the average displacement vector over the entire image. Although this displacement could be captured by using a sufficiently large window, the initial discrete window shift allowed for many more independent statistical samples for an error analysis.

Table 1 presents a summary of the standard flow error analysis using the SCC and RPC processors. The total error is defined from:

$$E_{tot} = \sum_{ij} \sqrt{\left(u_{ij} - U_{ij}\right)^2 + \left(v_{ij} - V_{ij}\right)^2} \tag{29}$$

where U and V are the actual displacement field. The total error is calculated without validation. The results indicate strong improvements for many of the more difficult DPIV image conditions. Notably, the large displacement field, the large particle-image diameter field, and the densely seeded field show drastic improvements in the measurement accuracy. The small displacement field is the only field for which the RPC correlation shows increased error, a 2.6% increase in median total error (while also a 2.6% decrease in mean total error). In contrast, the relative error for the large displacement field was reduced by 61.1% (mean) and 35.8% (median).

Figure 12 presents the error fields of the standard flow #2 for the SCC and RPC processors. The SCC histogram shows two distinct regions: a primary error peak and a uniform distribution of erroneous vectors. The RPC histogram shows a greatly reduced region of erroneous vectors. This reduction in erroneous vectors is the primary source of error reduction for the large displacement standard flow in Table 1. In addition to the removal of erroneous vectors, the median total error is reduced by 35.8%. This indicates both a better VVD probability and a higher accuracy. The dominant total error within the SCC contours appears to be centralized around regions of high shear. The highest reductions in error for the

Table 1   Summary of Standard Flow Error Analysis.

Std Flow (shear)	Mean E_{total} (%)			Median E_{total} (%)		
	SCC	RPC	Reduction	SCC	RPC	Reduction
1 reference	5.6	4.0	28.6	4.9	3.5	28.4
2 large displacement	22.1	8.6	61.1	7.3	4.7	35.8
3 small displacement	10.6	10.3	2.6	9.8	10.1	-2.6
4 dense seeding	10.7	6.4	40.6	7.0	5.6	20.7
5 sparse seeding	7.6	6.3	16.8	6.4	5.4	15.4
6 constant diameter	6.0	5.1	15.9	5.5	4.5	17.4
7 large diameter	13.0	9.0	31.0	9.4	8.0	15.3
8 out-of-plane motion	6.6	5.5	15.5	6.0	4.9	18.2

(a) SCC Processor

(b) RPC Correlation (with extended Blackman image mask)
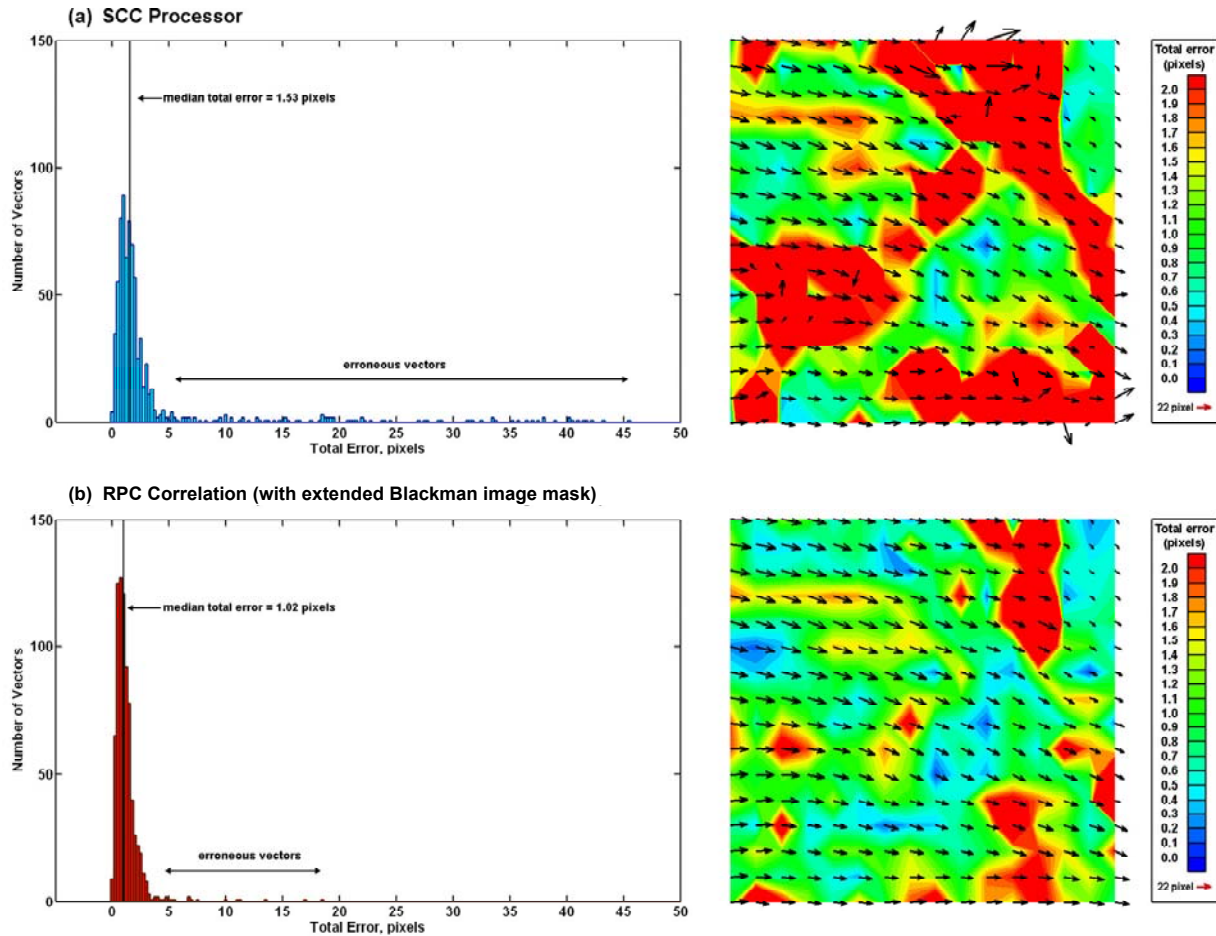
Figure 12 Histograms of total error presented for the 3 image pairs of standard flow #2. The error field contours are presented for the first image pair of this standard flow.

RPC correlation occur within these regions. These results are consistent with the notion that the phase correlation peak is more immune to the effects of shear than the SCC processor, due to the modified shape of the correlation peak [15].

EXPERIMENTAL VALIDATION

Although artificial simulations are useful in identifying and quantifying various types of systematic errors, these methods must ultimately be tested in actual experimental data to ensure the proper behavior of the estimation. However, there is obviously a limit to the amount of quantitative assessment that can be performed upon actual experimental data since the systematic error cannot be easily separated from the experimental errors.

STENTED VESSEL IN PULSATILE FLOW

The high valid vector detection of the RPC correlation provides a tool for examining the dynamics of more complex flowfields. For this study, DPIV images were generated from a 4mm compliant stented vessel subject to pulsatile flow, simulating arterial flow [43]. The stented vessel is shown in Figure 13a. Time-Resolved Digital Particle Image Velocimetry (TRDPIV) is required to accurately resolve the unsteady and anisotropic nature of the flow. TRDPIV relies upon the accuracy of the cross-correlation to provide a flowfield from a single pair of DPIV

images. Accordingly, the highest measurement accuracy and valid vector detection are required for such a system.

TRDPIV measurements were taken at a Reynolds number of 300 with a Womersley number of 2.73, corresponding to strong unsteady effects. Single-exposure double-pulsed image pairs were recorded with a resolution of 1.85 μm/pixel and a pulse separation of 40 μs taken at 500 Hz. The corresponding peak velocity is approximately 12 pixels.

Several challenges were introduced through the compliable walls and stent bars which inhibited the acquisition of well illuminated images. Primarily, the walls of the stent obstruct the camera's field-of-view and the laser's illumination of the flowfield. These obstructions are illustrated in Figure 13b and Figure 13c. These factors greatly affect the intensity of the particles as they pass through the field of view. In addition, the diffuse transmission through the vessel walls caused substantial optical distortion.

The images were processed using the SCC and RPC with a spatial resolution of 29.6 μm. Images were correlated using a 128x64 pixel window followed by a 64x16 pixel window after a second-order discrete window offset. Figure 14 presents the flowfield from a single frame pair during the pulsatile flow. Contours give the axial velocity at full spatial resolution, while vectors are shown for full resolution in the radial direction and half-resolution along the axial direction. Measurements are presented without validation in either the first or second window passes.

Measurements using the SCC processor reveal large areas of erroneous vectors. Specifically, the near wall regions subject to shadows and obstructions show an absence of correlated information. In addition, numerous groups of erroneous vectors are seen throughout the field, as one would expect for a low density shear flow with large particle-image diameters in the presence of background noise. The RPC correlation of the same image pair shows drastic improvements over the SCC processor. The RPC correlation was filtered with extended Blackman image mask and a particle-image diameter of 7 pixels was used to construct the spectral energy filter. The poorly illuminated regions show a substantial increase in the valid vectors through the much more uniform axial velocity. In addition, many of the stray patches of erroneous vectors are reduced or removed.

Measurements of vorticity and shear require the use of velocity derivatives. As a result, the measurement errors for these flow variables will increase dramatically. Only the superior

(a) 4mm stented vessel

(b) model of PIV setup

(c) DPIV image with shadows and obstructions

vessel wall
camera obstructions
laser plane obstructions

1 mm
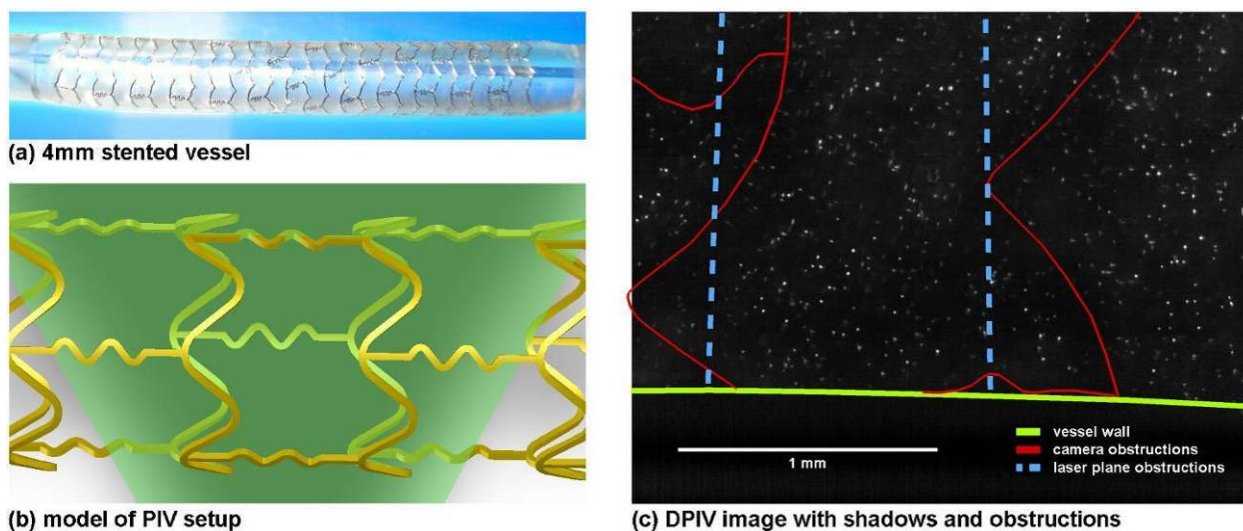
Figure 13 Stented vessel TRDPIV experimental design and image acquisition.

(a) **SCC Processor**

(b) **RPC Correlation (with extended Blackman image mask)**
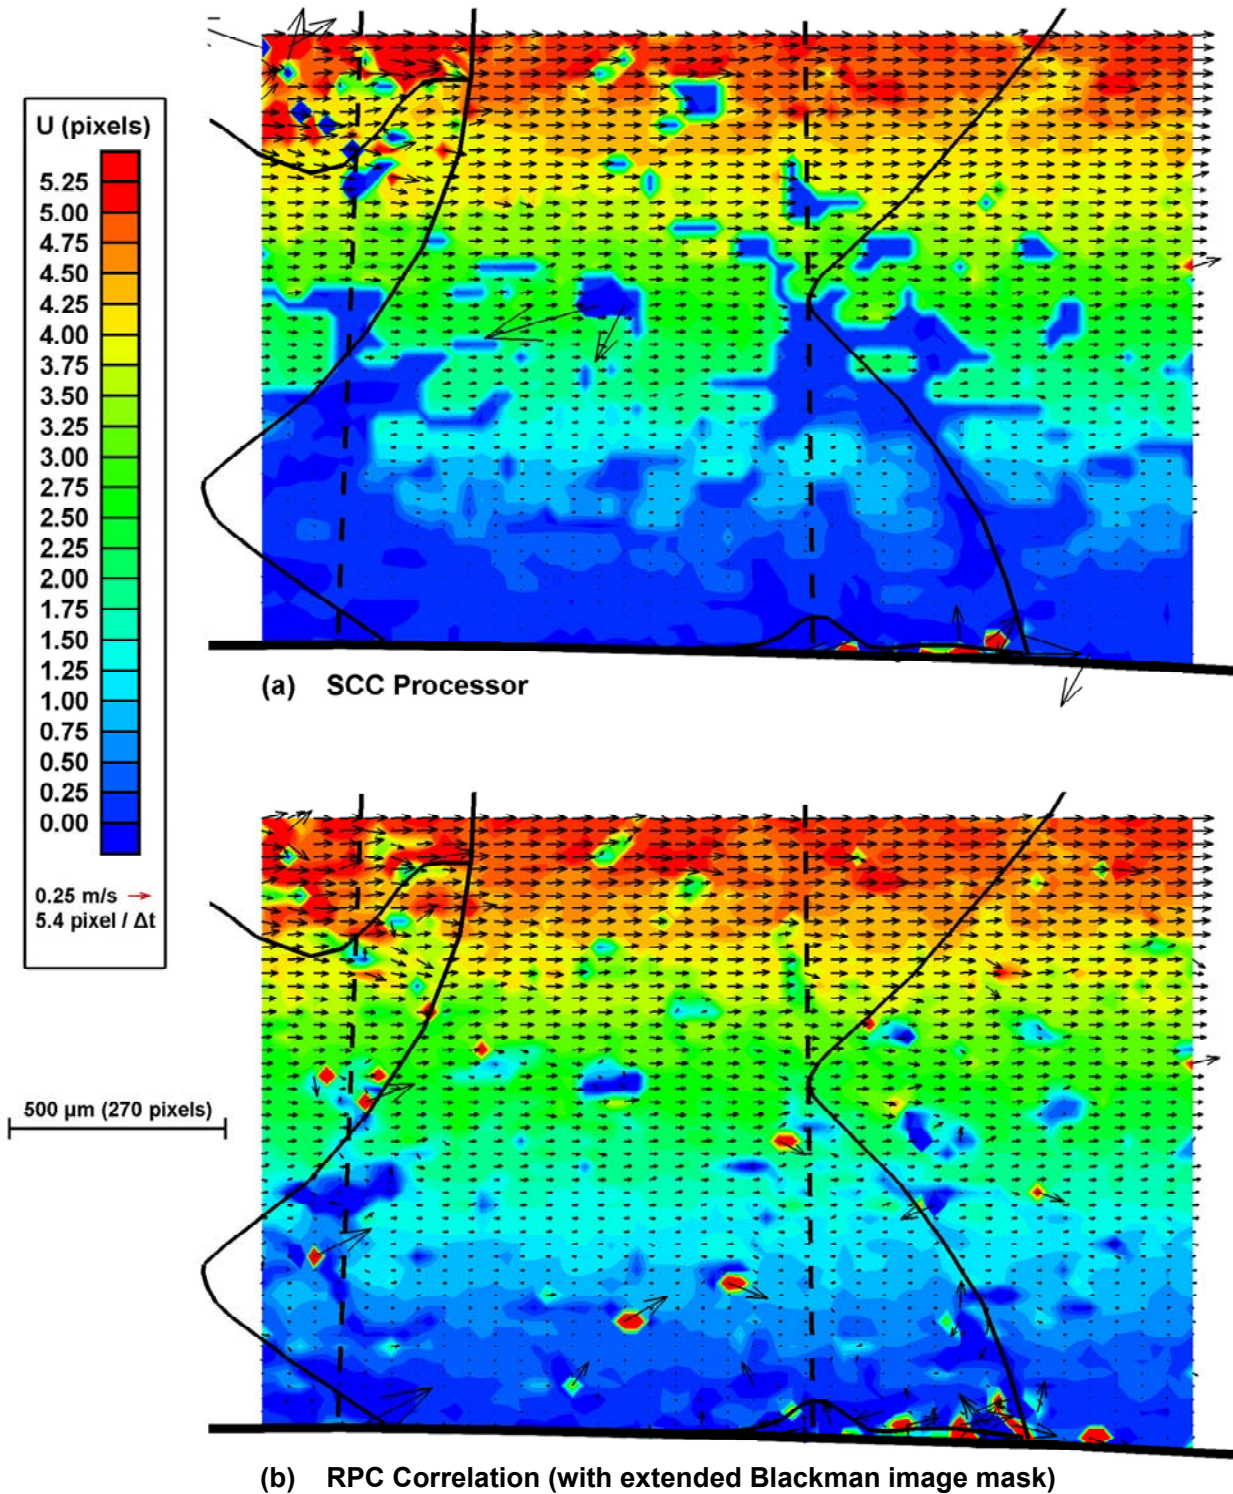
Figure 14 Single image pair flowfield generated with a two-pass discrete window offset correlation: 128x64 and 64x16 pixels. Flowfields are presented without validation.

flowfield given by the RPC correlation would be capable of reliable measurements of these variables. The wall shear stress, calculated from the velocity shear, is essential to the control over such an arterial flow. Accurate measurements of wall shear stress require high spatial resolution near the surface. In addition to the decrease in erroneous vectors, the RPC correlation provides better near wall measurements, especially in areas with low amounts of signal.

TURBULENT RIB-ROUGHENED CHANNEL FLOW

Another major challenge to DPIV processing algorithms is the measurement of micron size flow tracers in a macro-scale flow experiment. As a result of the small particle-image diameter, peak locking effects can become a very dominant source of error within the estimation. In order to examine the peak locking behavior of the RPC estimator, we examine the turbulent flow through a rib-roughened channel [44]. This channel is a 5x scaled model of a cooling channel for a turbine blade, with a diameter of 3.81 cm.

TRDPIV measurements were performed with a pulse separation of 450 μs and a magnification of 84.5 μm/pix. For a Reynolds number of 2500, the bulk velocity is 1.1 m/s, corresponding to a bulk displacement of 5.9 pixels between frames. Seed particles were generated using 1 μm sebacate droplets, making the particles-images dominated by the diffraction-limited image and the pixel resolution.

Figure 15 presents a mean image sequence obtained from this experiment. The scattered light off the surface of the ribs causes substantial background illumination, making the displacement hard to resolve in these regions. The rest of the image contains a relatively high seeding density, allowing for a large number of particle-image pairs to contribute to the correlation. As such, the VVD probability is high enough that little validation is necessary. The dominant source of error will then be the peak locking effects, and bias errors cause by the high shear.

Figure 16 presents the velocity magnitude contours from a sample image pair for both the SCC processor and the RPC processor. These flowfields were processed with a 64x32 pixel

Figure 15 Mean of a 4 image TRDPIV sequence of a rib-roughened channel (diameter = 3.81 cm). 1 μm sebacate seeding particles were imaged at a magnification of 84.5 μm / pixel.

window, followed by a 32x16 pixel window after a discrete window offset. The SCC processor was evaluated with rectangular windows, where as the RPC was processed with the extended Blackman image mask due to the benefits observed in Chapter 4. From these flowfields, it is apparent that the flow is highly turbulent. Large separation regions are observed behind each rib. Additionally, the deviation of the bulk flow indicates strong turbulent interaction caused by the ribs. Separated flow regions can be observed over the sharp corners after each rib.

The RPC valid vector detection is shown to be about 50% greater for regions of high shear and background noise, however, the overall VVD probability is very high for both estimations. The high reflection regions off of the surfaces results in a loss of measurement for a small region just above each rib. For the RPC, this region is slightly smaller that the SCC, by about 6 pixels, but is still present due to the very high glare, which saturates the image.

Histograms of these flowfields are also presented in Figure 16. From the SCC histogram, there is clearly a large amount of peak locking within the flowfield, reflecting a high amount of bias error within the velocity estimation. The RPC histogram shows a dramatic reduction in the peak locking, only slightly apparent for some of the higher velocities. The reduction of the peak locking in the histogram is indicative of either a reduction in peak locking biases, or a substantial increase in the RMS errors. However, had the RMS errors increased to level required to smear out the peak locking effects, this would be easily visible in the velocity contours, which is not observed. Instead, the separation regions and the bulk flow appear to have equally high resolutions and accuracy, indicating a substantial reduction in peak locking biases for the RPC.
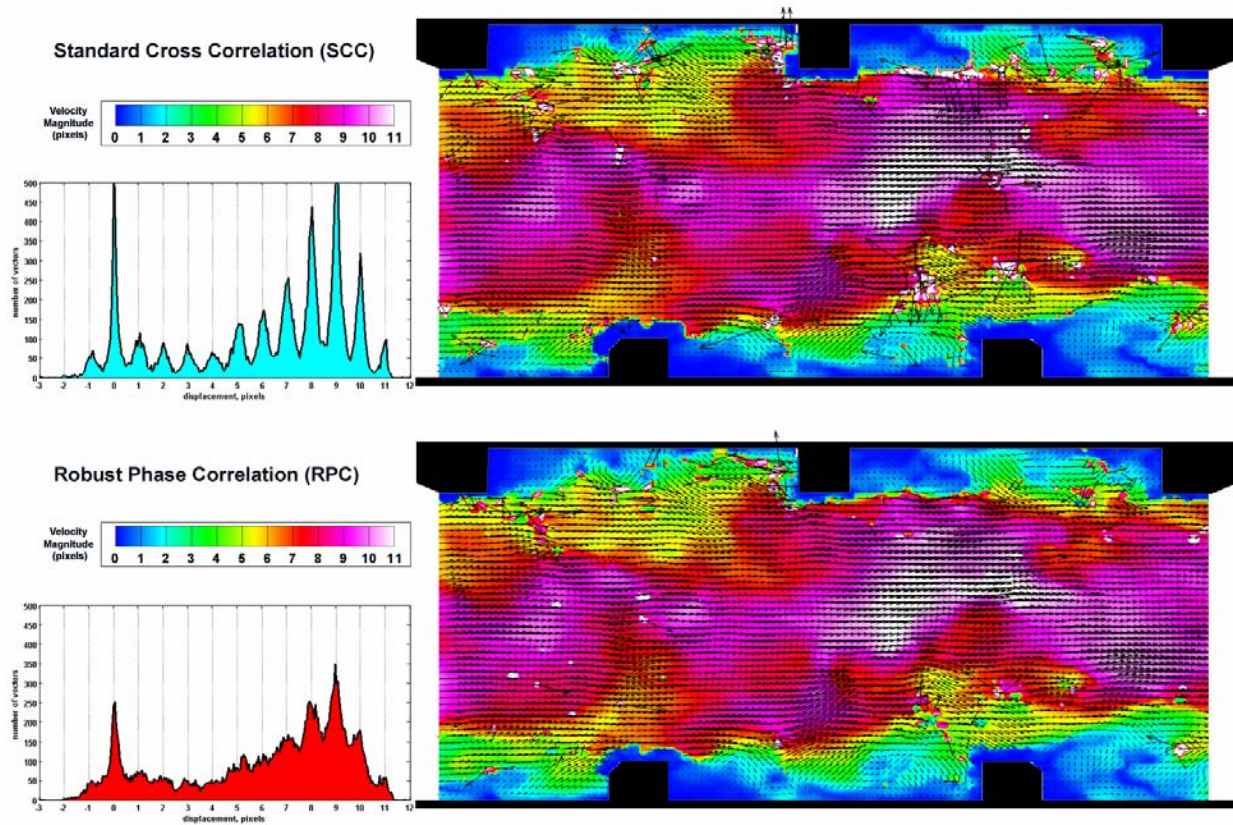
Figure 16 Comparison of the SCC and RPC processors for a rib-roughened channel at a Reynolds number of 2500, giving a bulk velocity of 1.1 m/s (5.9 pixels / Δt).

In the RPC estimator, the phase correlation was optimized such that it provided the strongest SNR. This has the advantage of providing a very high VVD probability. However, we are concerned not only with VVD, but also high measurement accuracy. Therefore, we now examine a variation of the RPC estimator, where the phase correlation is designed to minimize the error associated with the subpixel estimation.

THE GAUSSIAN SUBPIXEL ESTIMATOR

After the identification of the correlation peak, a Gaussian curve is fit to the 3 central pixels to identify the displacement with subpixel accuracy [2]. Although the PHAT is capable of superior vector detection, much of the subpixel resolution is lost due to the nature of the PHAT correlation peak. Since the cross-correlation contains phase-only data with unity amplitude, the spatial correlation peak is no longer Gaussian, but more closely resembles a delta function. As a result, the neighboring pixels to the central correlation peak are often buried in the noise floor. Since the subpixel resolution depends heavily on the shape of the correlation peak and the signal to noise ratio of the neighboring pixels, subpixel resolution is often lost.

Figure 17 presents histograms of vectors calculated from a uniform distribution of image displacements, using only valid vectors. The SCC histogram demonstrates the effect of peak locking, where measurements are biased toward integer pixel locations. Because peak-locking can be very dependent on image conditions, this often represents a substantial source of error
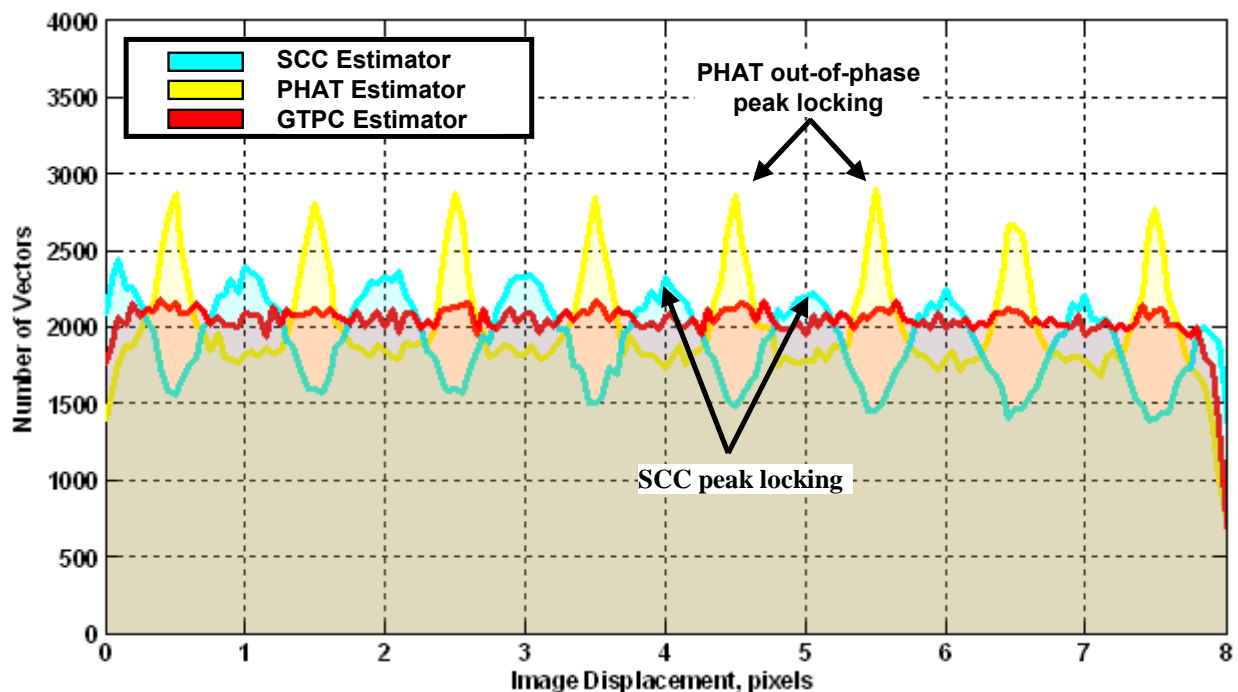
Figure 17 Histograms of a uniform distribution of image displacements between 0 and 8 pixels evaluated for the SCC, the PHAT, and GTPC. Images contained 10% white noise and were correlated using the Blackman spatial filter for the PHAT and GTPC estimators.

Gaussian Transformed PHAT

Figure 18   GTPC estimator diagram.

when processing DPIV images [18, 45].

The PHAT histogram indicates a more pronounced peak-locking error, where instead of being biased toward integer pixel values, measurements are biased away from integer pixel locations. This behavior can be explained by examining the delta function. For a delta function centered at an integer pixel, the neighboring pixels contain only information from the noise floor. The Gaussian fit is very sensitive to these pixels, increasing the RMS uncertainty. For delta functions centered between pixel locations, the correlation peak will fall on one of the two pixels, and the other will have nearly the same correlation value. Since the signal to noise ratio for this neighboring point is high, the Gaussian fit is better able to resolve displacements between pixels.

An alternative approach, the Gaussian Transformed Phase Correlation (GTPC), is now introduced to recover the subpixel resolution of the PHAT, using a Gaussian transform filter in the spectral domain, shown in Figure 18. This spectral filter is designed to transform the filtered cross-correlation, approximated by a delta function, into a Gaussian function. This reduces the measurement error associated with the 3-point Gaussian estimator. As a result, the peak locking effect is substantially reduced, as shown in Figure 7 by the Gaussian transformed GCC.

Since the correlation matrix approximates a delta function, transformations take the form of simple convolutions. By the definition, convolving any function, $f(x)$, with the delta function, $\delta(x-x_0)$, gives:

$$f(x) \otimes \delta(x-x_0) = \int_0^\tau f(\tau)\delta(x-x_0-\tau)d\tau = f(x-x_0) \tag{30}$$

where the original function has been shifted to the location of the delta function. The spectral filter is then only a function of the desired input to the subpixel estimator.

The 3-point Gaussian estimator fits a Gaussian curve to the three primary pixels in each orthogonal direction. The desired input function should then also be a Gaussian, with a width selected to minimize the estimator error. The error associated with the displacement estimation, e_x, can be approximated by a first order error propagation:

$$e_x \frac{\partial \Psi}{\partial x} = e_\Psi \tag{31}$$

for a given correlation uncertainty, e_Ψ. By Equation 31, e_x is minimized for large gradients in the correlation. The width of the Gaussian filter is then optimized when the maximum of the correlation peak occurs at the neighboring pixels.

Figure 19 shows this representation of the optimal Gaussian correlation peak, $G(x_1,x_2)$, where the inflection point occurs at a one pixel offset. The spectral filter, $\hat{G}(k_1,k_2)$, is defined as the inverse Fourier transform of $G(k_1,k_2)$. Using a discrete Fourier transform, the spectral filter is given by:

$$G\left(k_x, k_y\right) = \exp\left(-\frac{1}{2}\left(\frac{\pi k_x}{L_x}\right)^2\right)\exp\left(-\frac{1}{2}\left(\frac{\pi k_y}{L_y}\right)^2\right) \tag{32}$$

where k_i is the wavenumber in the i^{th} direction. Since the spectral filter is only dependent upon the subpixel estimator, Equation 32 is given explicitly and the computational requirements are minimal.

ARTIFICIAL IMAGE SIMULATIONS: PARTICLE-IMAGE DISPLACEMENT PERFORMANCE

The histogram of the GTPC is also presented in Figure 17. The peak-locking appears to be entirely removed, providing a uniform distribution of vectors. In order to more closely examine the differences in measurement error between the PHAT processor and the GTPC processor, bias and RMS errors are presented in Figure 20. Significant non-integer peak locking is observed for the PHAT, indicated by increased RMS and bias errors. Furthermore, the bias errors are clearly varying with white noise levels.

When the Gaussian transform filter is applied to the PHAT, a remarkable reduction in measurement uncertainty is observed. The GTPC does not display any peak-locking within the RMS errors and indicates very minimal effects within the bias errors. Additionally, there is a dramatic reduction in average RMS errors over the entire displacement range. Error reductions can be attributed to increased signal-to-noise ratio for neighboring pixels and the peak's relation to the Gaussian 3 point estimator, allowing for a more accurate detection of the subpixel displacement. Clearly, the implementation of a Gaussian filter is essential when correlating DPIV images with the PHAT processor.

Figure 21 gives a comparison of RMS uncertainty for the SCC and GTPC for several window and noise conditions. Similarly to the valid vector detection performance, the use of spatial windowing is essential for the GTPC in order to reduce the RMS. Most importantly, there is a substantial reduction in RMS errors for both correlation techniques when using the Blackman filter, demonstrating the advantage of extended window tapering.

As indicated before, the GTPC exhibits minimal peak-locking effects, whereas the SCC becomes severely compromised by these effects at higher noise levels. Furthermore, since discrete window offsets [1] are used to increase the accuracy, measurement uncertainty is critically determined in the 0-1 pixel displacement regime. Over this range, the GTPC measurement uncertainty is on the same order as the SCC, but without the substantial errors due to peak-locking effects.
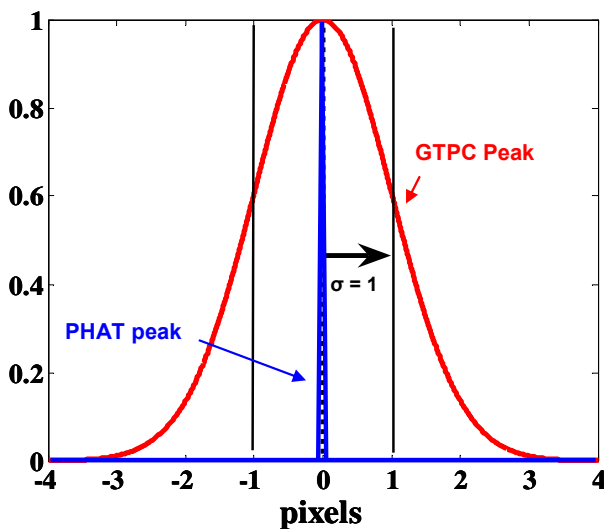
Figure 19 One dimensional correlation peak to optimize the 3 point Gaussian estimator.
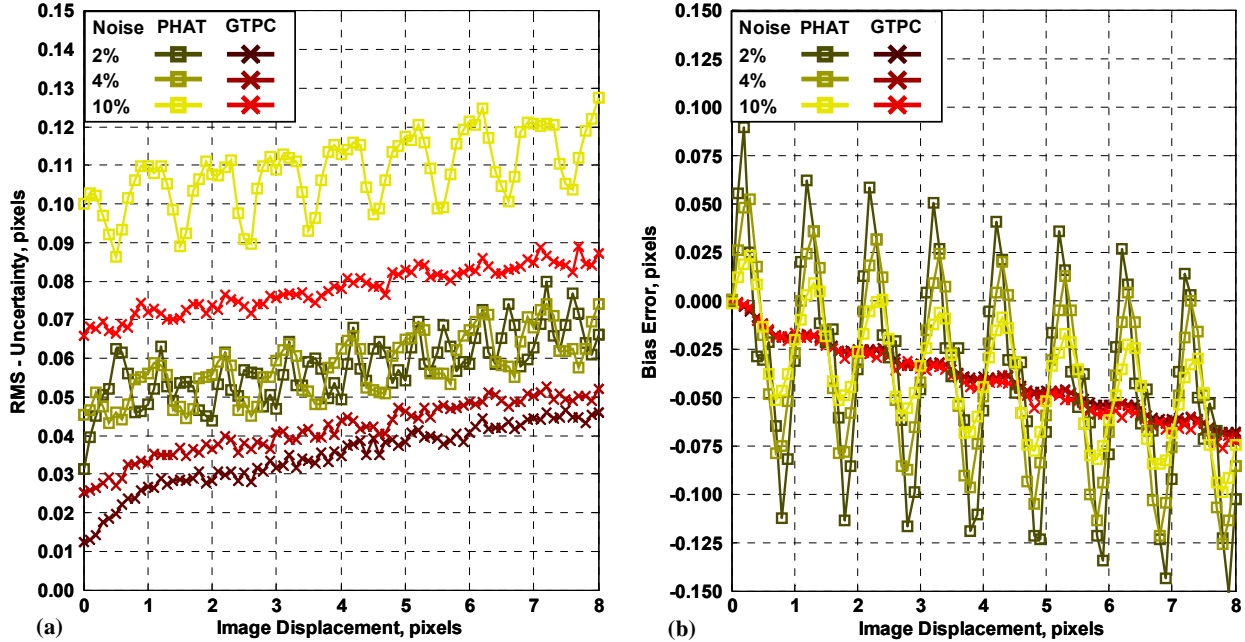
Figure 20 Bias and RMS errors presented at several image noise conditions. Results are presented for the PHAT and GTPC using the Blackman spatial filter. Errors were calculated from valid detected vectors only.

The corresponding bias errors are presented in Figure 22, for the same image conditions. The bias errors presented for the SCC without spatial filtering are typical of standard bias error results, composed of peak locking effects and a negative bias associated with the loss of correlation [25]. The GTPC shows the same loss of correlation errors, but no apparent peak-locking effects. As a result the bias errors of the GTPC processor are invariant to noise.

The bias errors presented for the Blackman window show substantially different results. The SCC shows a reduced bias associated with the loss of correlation due to the normalization of spatial windowing introduced in Equation 10. However, because the Blackman window contains some aliased signals, the required normalization does not represent the actual loss of correlation.

The GTPC does not require a normalization and as a result this processor is unable to account for the loss of correlation. Bias correction [8] can be incorporated by defining the loss of correlation for the Blackman window, and applying the loss of correlation mask to the subpixel processor. As a result, the bias corrected generalized cross correlation, $GTPC^C$, has superior performance in both the reduction in magnitude of the bias errors, and significantly diminished peak-locking effects. The difference in error between the GTPC and $GTPC^C$ can therefore be attributed to the loss of correlation at increased image displacements.

Zeropadding is able to better attenuate the bias errors due to the loss of correlation [34], than tapered spatial windows due to the removal of aliased signals. However, the zeropadded window exhibits far greater RMS uncertainty. As a result the $GTPC^C$ processor using the Blackman window appears to be the most robust against noise, while providing increased measurement accuracy.

37

Figure 21 RMS uncertainty plotted against particle displacement at varying levels of white noise. Correlations shown for no spatial filter (a) and the Blackman spatial filter (b). The errors were calculated using only valid vector measurements previously presented.

LIMITATIONS OF THE GTPC ESTIMATOR

The GTPC represents a special case of the RPC correlation. In order to increase the measurement accuracy of the PHAT, a Gaussian transformed filter was designed to optimize the 3-point Gaussian estimator [41]. Incidentally, the Gaussian transform filter corresponds to the spectral weighting function of the RPC estimator at a correlation peak diameter of 4 pixels (standard deviation of 1 pixel). The shape of the correlation peak is directly related to the shape of the particle signal. The spectral energy filter defines the diameter of the Gaussian correlation peak:

$$d_c = \sqrt{2}d_e \tag{33}$$

in the absence of particle aliasing, which is a commonly derived result [1]. The optimal particle-image diameter for the Gaussian subpixel estimator is then 2.83 pixels. Again, this is related to the optimal particle-image diameter to minimize the measurement error. In the spectral domain, the optimal particle-image diameter is defined by the maximum spread of spectral energy without aliasing. For smaller diameters, the higher spectral energy is aliased onto the lower frequencies. At larger diameters, the energy content is contained within a narrow band of the digital spectrum. In the spatial domain, smaller diameters result in a correlation peak that is too narrow to be resolved due to pixel discretization. Larger diameters result in a broad correlation peak where the neighboring pixels in the Gaussian subpixel estimator are dominated by noise. Therefore the GTPC is optimized to the subpixel estimator, where as the RPC is optimized to the image SNR.

38
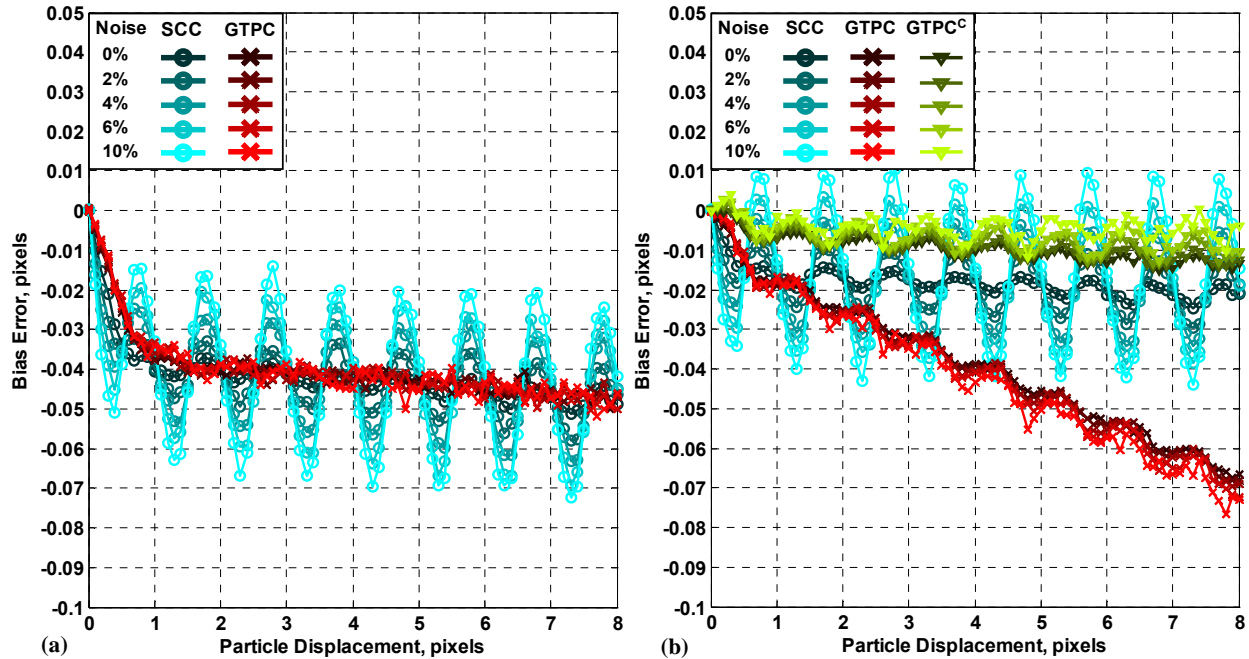
Figure 22 Bias errors plotted against particle displacement at varying levels of white noise. Correlations shown for no spatial filter (a) and the Blackman spatial filter (b). The errors were calculated using only valid vector measurements previously presented.

The most obvious limitation of the GTPC is its valid vector detection. Since this technique is designed to optimize the subpixel estimation, it may no longer provide the strongest SNR (namely, if the particle-image diameter deviates substantially from the 2.83 particle-image diameter listed above. However, provided the correlation signal is strong, this estimator minimized the error in the subpixel estimator. Therefore, its primary benefits would be for clean data sets, ensemble averaged measurements, or iterative correlation offsets.

The traditional DPIV estimators use the Fourier cross-correlation only as an estimation of the direct evaluation of the cross-correlation [1]. However, the use of Fourier transforms allows for another technique to estimate the displacement. This method is through a direct estimation of the group delay of the phase correlation matrix. While more complex than the spatial evaluation, this method offers several potential benefits over traditional DPIV estimators.

THE PHASE CORRELATION METHOD

The Phase Correlation Method is commonly used to measure translational displacements of MRI images [32]. Since MRI images are taken directly within the Fourier domain, image processing using the Phase Correlation Method was natural. The Spectral Phase Correlation (SPC) is now introduced as an extension of the Phase Correlation for DPIV displacement estimation. For DPIV images recorded in the spatial domain, the phase content must be extracted from the Fourier cross-correlation. Figure 23 shows a flow diagram for the SPC. There is a clear relation between the Generalized Cross-Correlation (GCC) processor [30] and the SPC, as they both operate on the phase-filtered cross-correlation. However, the GCC processor converts back to the spatial domain to estimate the image shift, while the SPC directly estimates the image shift in the Fourier domain using the phase data.

The SPC is derived from the Fourier shift theorem, which states that the displacement in the spatial domain will correspond to linear phase changes (constant group velocity) in the Fourier domain [46]. For a DPIV image pair, I_1 and I_2, the Fourier shift theorem is given by:

$$I_2(k_x,k_y) = I_1(k_x,k_y) \cdot \exp(-j(k_x t_x + k_x t_y)) \tag{34}$$

where t_x and t_y represent the two orthogonal group velocities between the images. In order to determine the displacements directly within the Fourier domain, t_x and t_y must be identified. This function can be estimated by using the phase correlation matrix:

$$\frac{I_2(k_x,k_y)I_1^*(k_x,k_y)}{\left|I_2(k_x,k_y)I_1^*(k_x,k_y)\right|} = \exp(-j(k_x t_x + k_y t_y)) \tag{35}$$

where the term on the left hand side is the symmetric phase-only filter of the cross-correlation, hereafter referred to as the phase-correlation matrix. The phase-correlation matrix contains the phase shift between the Fourier transforms of the image window pair at each frequency, (k_x,k_y). The phase-correlation matrix must be solved to identify t_x and t_y, which can be translated into the spatial image shift by:


Cross-Correlation ⟶ Phase Transform ⟶ Dominant Singular Value ⟶ Phase Angle Matrix ⟶ Unwrapped Phase Plane

Figure 23 Flow diagram of SPC processor.

$$d_i = \frac{t_i L_i}{2\pi} \tag{36}$$

where L_i is the length of the matrix in the i^{th} direction [32]. By identifying the group velocities, and utilizing Equation 36, solving for the displacement does not require an inverse Fourier transforms, thus simplifying computations.

DECOMPOSITION OF THE IMAGE SIGNALS WITH SINGULAR VALUE DECOMPOSITION

The phase correlation matrix will contain the displacement information as well as information from image noise, pixel discretization, and FFT based spectral contamination. In order to separate linear phase changes from spectral contamination, Singular Value Decomposition (SVD) is used to identify the dominant mode within the phase correlation matrix [47]. SVD is a method of Eigenvalue decomposition, which decomposes a matrix among its most energetic modes given by:

$$\Omega = P\Sigma Q^H \tag{37}$$

where Σ is a diagonal matrix corresponding to the nonnegative singular values of the phase correlation matrix, Ω represents the phase-correlation matrix and H denotes the conjugate transpose [48]. P and Q are the unitary matrices containing the corresponding singular vectors p_i and q_j, respectively. Equation 37 is evaluated with the most dominant singular value, which recovers a filtered form of the phase-correlation matrix. The phase angles of the complex phase-correlation matrix are then computed and unwrapped [49] into a phase plane, shown in Figure 23.

Because of the decomposition, the phase plane is separable into the dominant right and left singular vectors of the unitary matrices. The group velocities are found using linear least squares regression on each of the dominant singular vectors, and the corresponding image shifts are determined by Equation 36.

ADVANCED REGRESSION TECHNIQUES

One of the fundamental assumptions when using linear least squares regression is that the errors are homoscedastic, meaning that each sample within the vector exhibits the same finite variance. When this assumption is invalid, a significant source of error is introduced to the least squares estimation. However, standard least squares approaches remain dominant due to the complexity and varying success of nonlinear regression techniques for heteroscedastic signals [50]. Stone demonstrated that the higher frequencies of the phase-correlation matrix show a larger variance since they are more sensitive to errors, especially in the presence of white noise [32]. He suggested masking the phase correlation matrix to limit the regression to frequencies below a certain threshold, η. Resulting analysis demonstrated very high accuracy in displacement estimation for MRI images.

An alternative approach to improve the linear phase estimation is now introduced using two nonlinear regression techniques. These methods are based on the use of a weighted least squares regression which minimizes the error:

$$E = \sum_i w_i (y_i - \hat{y}_i) \tag{38}$$

where w_i is a weighting function to account for the heteroscedasticity of the signal. A common method is to apply a weighting given by the inverse of the variance [50]. Through this, the standard linear least squares regression becomes a subset of the weighted least squares regression for which the weighting function is unity.

Defining the weighting function usually requires *a priori* knowledge of the signal variance, which is not always available or easily determined. One method of overcoming this requirement is to use a robust least squares algorithm, designed to remove data outliers. This method is similar to the weighted least squares regression, except the weighting function can be determined directly from the data, without any prerequisite knowledge of the signal variance. Several models have been suggested to describe w_i as a function of the residuals [51]. For this analysis, the robust algorithm uses a bisquare weighting function:

$$W_i = \begin{cases} \left[1 - \left(\dfrac{r_i}{KM_r} \right)^2 \right]^2 & r_i \leq K \\ 0 & r_i > K \end{cases} \tag{39}$$

where K is an empirically determined tuning constant, equal to 4.685 for the bisquare distribution [51]. M_r provides normalization, defined to be the median absolute deviation of the residuals. The weighting function essentially gives less weight to points which exhibit high residuals. The weighted least squares fit is then performed, resulting in a new regression curve and the associated residuals. Then the weighting is iteratively recomputed until the solution converges. Because this is an iterative technique, the computation time is significantly higher than the standard linear regression. However, the robust method provides a more accurate estimation than standard least squares regression by removing data outliers.

Our subsequent analysis showed that the robust least squares technique is still subject to significant errors from spectral contamination. In order to overcome this source of error we developed a weighting function based upon the signal-to-noise ratio. In this technique, a weighting function is defined from the signal-to-noise ratio of the spectral content. The spectral content of the noise signal should be uniform with respect to frequency, since the phase-correlation contains primarily uncorrelated noise. In contrast, the displacement signal of the PIP will be given by the spectral content of the tracer particles. For a given particle-image diameter size, and using the assumption that particle-images can be approximated by a Gaussian intensity distribution[3], we can estimate the one-dimensional spectral content of DPIV images (and in turn the weighting function):

$$W_k(k) = \left| FFT \left\{ \exp\left(\frac{8(x-x_0)^2}{d_e^2} \right) \right\} \right| \approx \exp\left(\frac{-\pi^2 d^2 k^2}{8L_x^2} \right) \tag{40}$$

---

[3] The Gaussian approximation of the particle-image diameter is addressed further in the development of the RPC estimator in Chapter 5

where d_e is the particle diameter and x_o defines the particle location. Only the one-dimensional spectrum is examined since the phase-correlation matrix is separable into singular vectors, and the regression is performed in each orthogonal direction. Since the amplitude of the spectral content is defined by the Gaussian particle shape and the phase information determines the particle location, W is roughly independent of the location of x_o, provided the particle is not located near the edge of a window. This method then defines the weighting function using the *a priori* knowledge of the particle shape and is defined only by the particle diameter. Figure 24 shows the weighting function for several particle diameters.

Examining the weighting function in Equation 40 as the inverse of the variance, the signal variance is noted to increase exponentially with frequency, indicative of the heteroscedasticity of the signal. Although the weighting function can provide some correction for the loss of signal, frequencies carrying sufficiently high phase variance often appear as outliers in the data for finite length discrete signals. Because of the regression's sensitivity to outliers, these points can still compromise the estimation, even if they are assigned a low weighting.

Figure 24 Regression weighting function for several particle diameters. The threshold is also shown which defines a frequency cutoff for the weighting.

Examining the bisquare weighting given in Equation 39, points considered to be outliers are given zero weight to ensure that they do not contribute to the regression. Similarly, since the variance increases exponentially, a threshold is defined, η, above which the phase content is discarded. The weighted regression then removes the outlier tendencies of these high variance frequencies.

Defining the threshold requires balancing the contributions of each point to the regression. The DC phase shift is identically zero, meaning that the regression must intersect this point. Contributions to the regression can then be estimated by defining an error as the uncertainty in the slope given by the ratio of phase standard deviation, $\Delta\varphi$, to the frequency, Δk:

$$e_{GD}\left(k\right)=\frac{\Delta\phi}{\Delta k}=\frac{\sigma_k}{k}=\frac{1}{\sqrt{W_k}\,k}=\frac{\exp\left(\dfrac{\pi^2 d^2 k^2}{16 L_x^2}\right)}{k} \tag{41}$$

where the weighting function is used to define the variance. This error essentially balances the advantages of increased regression points to the variance of additional points. Since the error increases monotonically with frequency, the minimum of e_{GD} defines the threshold to be used in the weighting function. Finding the minimum of Equation 39 defines the threshold, η, of the weighting function:

$$\eta = \frac{\sqrt{8}L_x}{\pi d_e} \tag{42}$$

where the corresponding weighting value, $W_k(\eta)$, is $1/e$, below which the weighting is defined as zero, also depicted in Figure 24. Since higher frequency information is discarded, the SVD need only be performed on frequencies below the threshold, greatly decreasing computational expenses. Rewriting Equation 40 using the threshold, the weighting function becomes:

$$W^T(k) = \begin{cases} \exp\left(\dfrac{-\pi^2 d_e^2 k^2}{8L_x^2}\right) & k \leq \dfrac{\sqrt{8}L_x}{\pi d_e} \\ \\ 0 & k > \dfrac{\sqrt{8}L_x}{\pi d_e} \end{cases} \tag{43}$$

where $W^T(k)$ denotes the threshold filtered weighting function. The threshold filtered weighting functions are also shown in Figure 24.

Finally, one should remember that the weighting function given by Equation 43 was constructed using several assumptions. Namely, the spectral content of the noise is dominated by uncorrelated phase information and the spectral content of the signal can be described by the spectral content of a single particle. Signal aliasing, polydispersed particle distributions, edge effects, and spatial windows will also influence the heteroscedasticity of the spectral content. However, including these effects significantly increases the complexity of the weighting function, making them beyond the scope of this study. A more in depth analysis could extend this weighting function to account for more complex systems and additional image factors.

ARTIFICIAL IMAGE SIMULATIONS: PARTICLE-IMAGE DIAMETER PERFORMANCE

Figure 25a shows the total error for the SPC using three regression models: standard least squares, weighted least squares, and a robust algorithm using bisquare weights, for an image density of 10 particle pairs / 32x32 pixel window. The pixel displacement was taken to be 0.3 pixels, corresponding to the linear regime where error is proportional to image displacement [25]. The standard least squares method performs poorly when compared with other regression techniques because it assumes homoscedasticity. The robust algorithm shows improved measurement accuracy over the standard least squares technique, but still demonstrates a similar rate of increase with respect to diameter. This is also the result of the heteroscedasticity of the data discussed earlier. By examining the threshold shown in Figure 24, one observes that over half of the frequency content is assumed to be contaminated by noise at image diameters greater than about 3 pixels.

Figure 25 RMS errors against varying particle diameter. Performance of TPC for 3 regression processors is presented for no spatial filtering (a). A comparison of standard and phase correlations is presented for several spatial filters (b). Errors are averaged over a 0-1 displacement; 10 particles / 32x32 window.

The robust algorithm uses the residuals found from a standard least squares regression to find an initial weighting function. If the signal is significantly contaminated, the standard regression is unable to lock onto the true regression curve, and a reasonable weighting function cannot be constructed. The robust algorithm is comparable to the weighted regression, until the spectral contamination becomes too significant, and departs from the weighted curve at about 2.5 pixels. Alternatively, since the threshold described in Equation 42 identifies the contaminated frequencies, the weighted least squares algorithm is clearly superior and subsequent results will be presented using this algorithm.

The RMS errors of the SPC processor are presented in Figure 25b as a function of varying particle-image diameter. Three spatial windows were considered for this technique: the zeropadded window [34], the Blackman window [41], and the standard unfiltered window. The traditionally examined case of the SCC without spatial filtering matches well with previously reported results [8, 25]. The zeropadded window is able to remove aliased signals, reducing bias errors, but has been shown to compromise the spectral content which leads to an increase in the more dominant random errors [52]. As a result, the zeropadded correlation does not perform as well, especially as the diameter is increased. The Blackman window uses the same extended image domain as the zeropadded window, but instead of masking the domain with a top-hat profile, the overlapped image region is tapered with the Blackman window [41]. As a result, the Blackman window attenuates the aliased signals, but avoids compromising the spectral content. The benefits of this windowing technique can be observed in the decreased measurement uncertainty of both the SCC and SPC processors.

The SPC processor shows a drastically different behavior with respect to particle-image diameter. The minimum occurs at roughly the same location between 2.5 and 3 pixels. However, for diameters greater than 3 pixels, the SPC shows a very slow rate of error increase. The use of the spatial windows with the SPC shows a similar trend, but with reduced uncertainty, as expected from the benefits of spatial filtering for phase correlations [41]. Because of this decrease, the SPC is shown to be more flexible with respect to higher particle diameters. To the best of the author's knowledge, no other correlation technique has been able to demonstrate reductions of this order with respect to higher particle-image diameters.

In contrast, for diameters below 2.5 pixels, the SPC shows an increased amount of error. This increased error demonstrates a limitation of the weighting function described in Equation 43. For smaller particle diameters, the discrete illumination of the digital image sensor is no longer able to preserve the Gaussian intensity profile. As a result, the signal to noise ratio can no longer be estimated using this Gaussian energy density function.

ARTIFICIAL IMAGE SIMULATIONS: PARTICLE-IMAGE DISPLACEMENT PERFORMANCE

In order to further examine the performance of this processor, the measurement uncertainty was calculated with respect to image displacement. Figure 26 presents the RMS uncertainty for the SCC and SPC, at a particle-image diameter of 5 pixels. Percentages of noise are expressed as a ratio of the standard deviation of normally distributed (white) noise to the dynamic range of the digital image. The noise for each pixel was uncorrelated with its neighbor, spatially and temporally. Phase-correlation has already shown the ability to remove correlated noise such as background illumination, non-uniform illumination, and signal flare [15]. However, white noise, usually introduced by thermal effects, is much more difficult to remove from correlations because the noise is entirely uncorrelated.


Figure 26 RMS Uncertainty evaluated for the SCC (a) and SPC (b) for the Blackman window filtered and unfiltered correlations at varying levels of white noise. Images density of 10 particles / 32x32 pixel window.

As expected from Figure 24, the RMS uncertainty is substantially increased in comparison to 3 pixel particle-image diameter images [8]. Specifically, the RMS errors are noted to increase linearly below about 0.5 pixel displacements, and then taper off as the displacement increases [25]. Again, the Blackman window shows a definitive reduction in error for all cases examined, clearly demonstrating superior performance for both the SCC and SPC processors.

Further examination of the SCC processor indicates a more pronounced peak-locking effect [25], especially for increased levels of image noise. In contrast, the SPC processor has no apparent peak-locking effects. The SPC also shows a substantially decreased slope below 0.5 pixels, which is responsible for the error reductions given in Figure 25 that were evaluated at an image displacement of 0.3 pixels. However, where the SCC error tapers off quickly for displacements above 1 pixel, the SPC error continues to slowly increase.

The corresponding bias errors for these image conditions are presented in Figure 27. Previous studies indicate that the bias error is dominated by peak-locking and loss of correlation effects [25]. For the unfiltered case, the SCC shows the common bias error trend resulting from these two dominant effects. The SPC bias errors follow the same trend, but do not include any of the peak-locking errors. This indicates that the bias errors are dominated by only loss of correlation effects, and as a result the SPC is more robust against image noise. Additionally, the errors within the linear regime at low displacements have a substantially lower slope. At 0.5 pixels, the SPC reduces the bias by 50% when compared with the SCC. As the displacement increases, the bias errors approach the same limit.

The implementation of the Blackman window results in a substantially different effect. Applying spatial windows to the SCC processor requires the use of normalization functions, which result in a clear reduction of the bias errors [8], but can often compromise vector detection
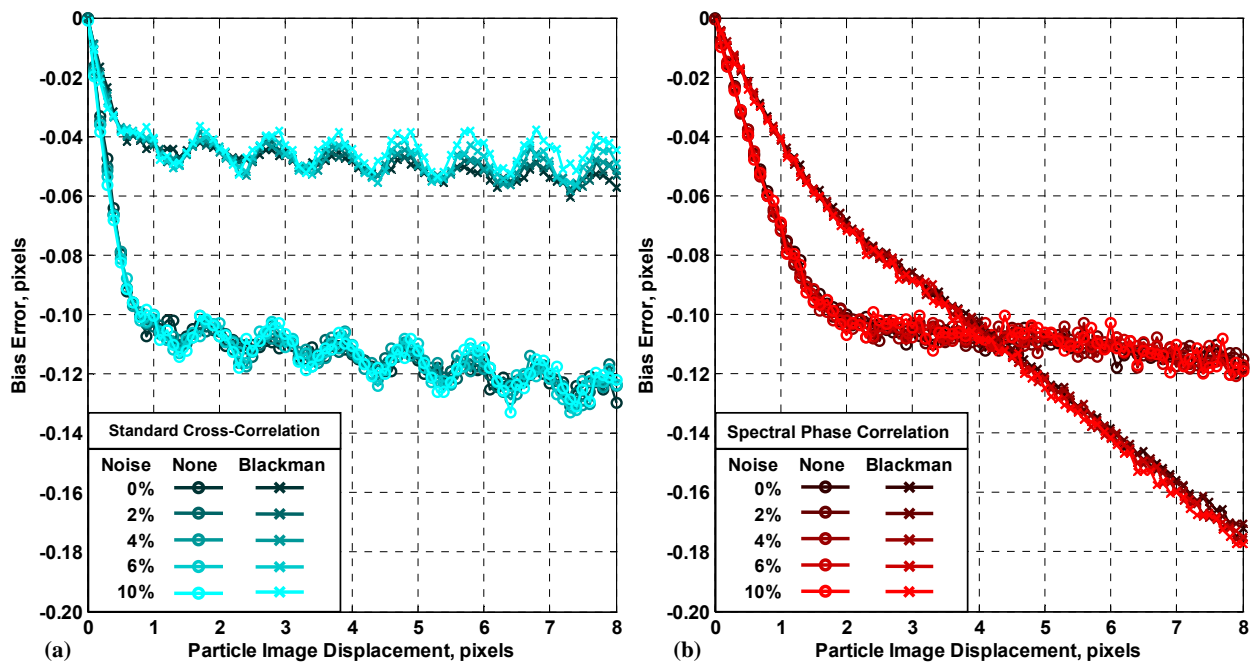

Figure 27 Bias errors evaluated for the SCC (a) and SPC (b) for the Blackman window filtered and unfiltered correlations at varying levels of white noise. Images density of 10 particles / 32x32 pixel window.

[41]. These normalization functions are used to account for the loss of correlation which biases displacement estimation. The spatially filtered GCC was able to incorporate these normalization functions into subpixel detection algorithms [41]. However, since the SPC computes the displacement directly from phase information these correction functions cannot be used, leading to the increased errors shown in Figure 27b.

The use of the SPC reveals a different approach for examining the causes of measurement error. The displacement is estimated from the linearity of the phase with respect to digital frequency. Figure 28 shows a typical phase distribution for an image pair displacement of 3.6 pixels at 4% white noise for 3 and 5 pixel particle-image diameters, calculated without spatial filtering. The mean phase curve was calculated from the ensemble average over multiple realizations of the SPC dominant singular vectors. The decrease in signal to noise ratio at higher frequencies results from the Gaussian distribution of the particle-image, which is more apparent for larger diameters.

The true phase curve corresponds to a constant group velocity calculated from Equation 36. The mean phase curves depart from the true phase curve roughly at the theoretically determined threshold from Equation 42. After this departure, the particle signal is no longer resolvable within the noise, and the resulting phase angles become completely uncorrelated, leading to zero group velocity. This departure from the ideal linear phase curve contributes to the SPC measurement error, creating a negative bias. The use of threshold filtering was shown in Figure 25 to drastically improve the SPC performance, which can now be attributed to the removal of these severely detrimental uncorrelated frequencies.

The use of discrete window offsets is able to substantially reduce measurement uncertainty since errors are smaller in the lower displacement range. Figure 29 shows the RMS and bias errors over the 0-1 pixel displacement range for the Blackman window, with image conditions identical to Figure 26. The linear increase in RMS uncertainties is significantly reduced for the SPC at low image noise levels. For high image noise, the SPC is on the same order of error as the SCC, but indicates a far more uniform distribution, due to the SPC's invariance to peak-locking effects. The invariance to peak-locking is also apparent in the bias errors. Specifically, the SCC shows larger bias, and is more affected by image noise. Because the SPC does not have any of these peak-locking effects, the bias errors are insensitive to image noise, and are approximately linear with respect to displacement. These advantages make the SPC far more robust than the SCC processor for large particle-image diameters.
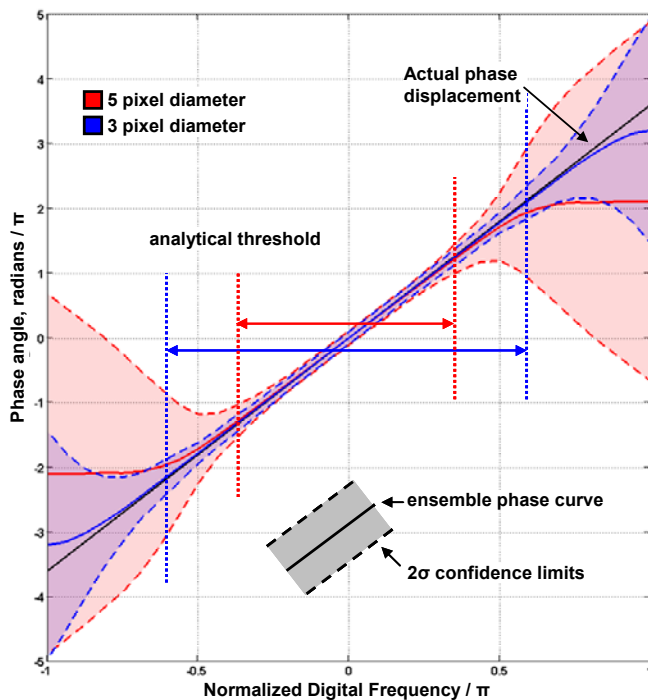
Figure 28 Mean phase curves and 2σ confidence limits shown for a 3 and 5 pixel particle-image diameter image pair with a displacement of 3.6 pixels. Theoretical thresholds are also plotted as well as the true phase curve.

Figure 29 RMS (a) and bias (b) errors within the 0-1 displacement range. Results are presented for the Blackman window filtered SCC and SPC processors. Images density of 10 particles / 32x32 pixel window.
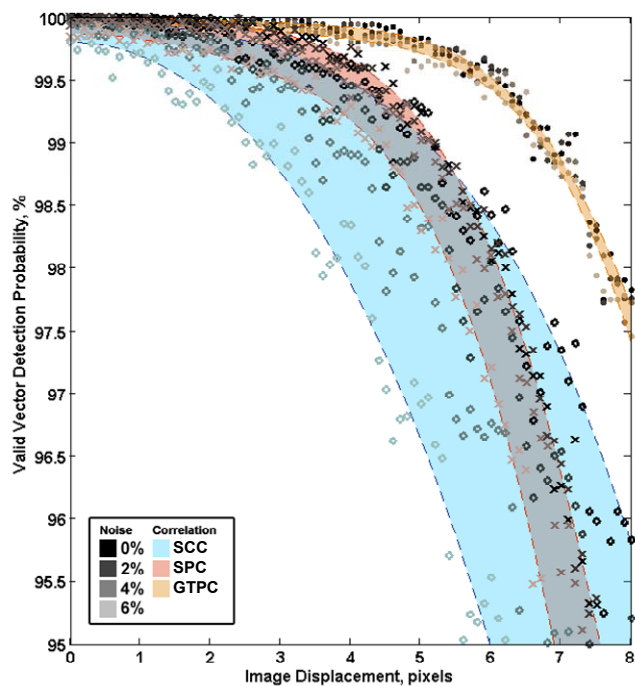
Figure 30 Valid vector detection probability for the SPC processor. Correlations were computed using the Blackman Window filter.

LIMITATIONS OF THE SPC ESTIMATOR

Although the SPC random errors can be reduced with the use of weighted least squares regression, the identification of the phase content within singular vectors is not always possible. The linear phase information must be unwrapped to account for the 2π periodicity of the phase angles for complex exponentials [49]. Unwrapping commands generally use a threshold to identify these discontinuities and add the corresponding 2π factor. When random errors dominate, differentiating between the discontinuities and noisy data is not always possible. If the unwrapping command is unsuccessful, the whole phase vector is compromised, and the measurement becomes invalid. This error is similar to the incorrect identification of the correlation peak for the SCC.

Figure 30 shows the valid vector detection for the SPC processor, where

erroneous vectors were defined to be absolute deviations greater than 0.5 pixels from the median displacement vector. Results are compared using the Blackman spatial filter. The use of phase correlation in the GCC showed a superior vector detection capability, due to its dynamic filtering of the image window pair [41]. The GCC and SPC analyze the displacement from the same phase correlation matrix, yet show dramatically different results. Because of the limitation in unwrapping the phase angle, the SPC valid vector detection is more sensitive to increased levels of white noise. Although these factors increase the erroneous vectors for the SPC, this error is still on the same order as the SCC processor. Since phase correlation is shown to be highly sensitive to the spectral content and aliased signals, the use of spatial filtering is essential to enhance the performance of the SPC.

With proper filtering, the SPC is capable of significantly higher measurement accuracy. The peak-locking effects can be entirely avoided through this spectral evaluation, as well as substantially reducing the RMS errors. However, because of the limitation in unwrapping the phase vectors, high VVD is limited to subpixel displacements. This makes the SPC an ideal tool for iterative correlation passes, where the image displacement should be in the subpixel regime, and high subpixel accuracy is desired.

This thesis has explored a variety of processing techniques directed toward an accurate estimation of Digital Particle Image Velocimetry (DPIV) measurements. Among these efforts, several novel techniques were introduced based upon the combination of windowing, filtering, and phase correlation estimators. While many of the methods presented here have been explored independently and for separate applications, this is the first time they have been demonstrated toward resolving DPIV measurements more accurately. Furthermore, while each of these filters can be applied separately to the estimation, this study demonstrated the advantages of their combined use in minimizing the systematic errors associated with this digital estimation.

A summary of the DPIV processing methodology is presented in Figure 31. The Standard Cross-Correlation (SCC) estimator gives the operational map of the traditional DPIV estimation.

The Robust Phase Correlation (RPC) estimator follows a similar path, but incorporates the use of dynamic filtering operations in the spectral domain of the cross-correlation. In this technique, the Fourier based cross-correlation is first transformed into the phase correlation via the phase transform (PHAT) generalized cross-correlation smoothing filter. This operation removes the energy content from the cross-spectrum, making the estimation immune to many of the input conditions. This immunity gives the phase correlation the capability to incorporate window masks without the requirement of normalization. Therefore, the RPC is able to provide an unbiased estimation, which is more robust against wraparound aliasing and spectral leakage.

Subsequently, the RPC phase correlation is masked with a spectral energy filter designed to attenuate noisy frequencies according to an analytical determination of the Signal-to-Noise


Figure 31 Summary of the various correlation techniques.

Ratio (SNR). The form of this analytic filter is based upon assumptions about the particle-image, namely that it may be approximated by a Gaussian function. In addition, the aliasing and noise signals are modeled in order to generate the spectral energy filter. This sequence of filtering operations for the RPC has the effect of removing the highly vulnerable energy content of the cross spectrum with an analytic energy distribution of the desired output signal. As a result, the RPC estimation was shown to increase the Valid Vector Detection (VVD) probability by an average of 50% in comparison to the SCC, taken from a series of uniform displacement artificial image simulations. The VVD probability is the most crucial metric in constructing a DPIV estimator, since it is a direct reflection of the correlation strength and will govern the dynamic range and capabilities of the system.

Next, the RPC estimator was compared against the SCC utilizing the PIV standard image project, a series of carefully constructed artificial simulations designed to comprehensively test the capabilities of a processor. For the most ideal image set, there was no notable increase in either the measurement accuracy or the VVD, as both processors were resolving the velocity field with the highest possible estimate. However, for the remaining 7 sets of more intense image conditions, the RPC outperformed the SCC in both measurement accuracy and VVD probability. In the most difficult set, with very large image displacements and shear, the RPC showed a reduction in total error of 79%. This was due primarily to a removal of the large quantity of erroneous vectors present in the SCC estimation. In addition, for the valid vectors, the RPC median error was reduced by 30% in comparison to the SCC processor.

While artificial image simulations are useful to quantitatively characterize the systematic errors of DPIV, these estimators must also hold up under actual experimental image conditions. Therefore, the RPC estimator was tested against the SCC in two experiments selected for their challenges to DPIV measurements. First, a stented vessel in pulsatile flow was considered due to its high magnification, low seeding, and background noise. Many of the obstacles of this experiment were similar to the challenges associated with micro-PIV. In addition, the stent walls created an obstruction to both the laser plane and the camera, which resulted in regions of the image where the particle-image signal was very weak. The SCC estimator showed a complete loss-of correlation in this region, resulting in an unusable measurement. Again, the RPC estimator was able to recover a substantial portion of these measurements, due primarily to its robustness against background noise and filters tuned to the image SNR. Additionally, the RPC estimator showed a higher near-wall resolution, which was made difficult both by the obstructions as well as distortion caused by the imperfect optical qualities of the vessel wall.

The second experiment considered in this study was the turbulent gas flow through a rib-roughened channel. This flow was highly turbulent with several shear layers, making high spatial resolution essential. As such, the seeding density was very high, which allowed for a high VVD probability, even for small windows. However, because of the high levels of turbulence and since the experiment was conducted in air, it was necessary to use very small flow tracers. As a result, the particle-image diameter was on the order of 2 pixels, which leads to domineering peak locking errors. Histograms of the velocity obtained with the SCC estimator confirmed the high levels of peak locking, despite a relatively clean velocity field. The RPC estimator was able to significantly attenuate these peak locking effects, since the poorly estimated energy spectra of the cross-correlation were replaced with the analytical model.

Two additional techniques were developed, which demonstrate substantial advantages for certain processing conditions. The first of these was the GCC estimator. This technique used the same dynamic filtering operation to transform the cross-spectrum into the phase correlation.

Next, an analytical filter was again masked to this phase correlation. However, the RPC designed the spectral energy filter to optimize the SNR within an image pair. The GCC estimator designs this filter to minimize the error associated with the subpixel estimation. As such, the GCC is more ideally suited to increase the measurement accuracy for well conditioned images, where the VVD is high. Interestingly, this Gaussian transform filter is identical to the RPC spectral energy filter evaluated at a particle-image diameter of 2.83 pixels. This relation is more apparent when considering the Gaussian subpixel estimator's relation to the particle-image.

With the use of the GCC estimator over the SCC, there were substantial reductions in the peak locking effect, similarly to the performance of the RPC estimator. In addition, the use of the extended Blackman image masks was shown to reduce the RMS error by about 50%, for both the SCC and GCC processors. However, this window mask was far more ideal for use with the GCC than the SCC due to the phase correlation's immunity to noise. In addition, it was demonstrated that the large bias errors commonly associated with masking could be removed with the proper normalization of the GCC correlation peak.

The final technique was developed from the phase correlation method, used to evaluate image shifts in MRI data. However, this was the first time it had been applied to the DPIV data, where the phase correlation matrix was approximated from the PHAT filtered cross-spectrum. In this method, the Spectral Phase Correlation (SPC), the displacement is estimated directly within the spectral domain using a least squares regression of the phase vectors. Singular Value Decomposition (SVD) is first used to decompose the signals in the phase correlation matrix, similar to the way the IFFT decomposes this matrix into several spatial correlation peaks. From this decomposition, the most dominant mode is selected, from which a weighted least squares regression is performed. The weighting for the regression was derived from the inverse of the phase variance at each frequency in the phase correlation.

The SPC was shown to have the most substantial benefits for large particle-image diameters. With the conjunction of this estimator with the extended Blackman image mask, the RMS errors were again substantially reduced in comparison to the SCC. More interestingly, the bias errors were shown to be virtually devoid of any peak locking errors. Since this estimator uses a regression over the entire frequency domain to generate its estimation, it is not as easily susceptible to peak locking errors as the standard 3-point subpixel estimators. However, this method was shown to have a decreased VVD probability for large image displacements. The erroneous vectors generated in this estimation were primarily the result of a failure in unwrapping the periodic phase vector, where the variance exceeded the unwrapping threshold. The best use of this estimator is then for subpixel displacements, where unwrapping the phase vector is unnecessary. As such, the SPC estimator is primarily considered as an iterative pass correlation technique, where image displacements are low and the correlation strength is high.

This study has achieved the overall goal of designing a robust correlation estimator capable of extending the dynamic range and capabilities of DPIV measurements. In addition, this estimator was shown to be ideally suitable for window masking, from which stemmed the majority of the error reductions. However, this study may prove more beneficial through its insight into the nature of the DPIV estimation. This analysis was able to reveal characteristics of the correlation through a spectral analysis, which in turn led to the development of better filtering techniques. While designed for DPIV in this study, the methodology described here provides a universal framework for digital signal correlation analysis, which could be extended to a variety of other systems such as topographical and map analysis, fingerprint identification, radar time-delay estimation, and MRI imaging.

REFERENCES

1. Westerweel J, *Fundamentals of digital particle image velocimetry.* Meas Sci Tech, 1997. **8**(12): p. 1379-1392.
2. Willert C and Gharib M, *Digital Particle Image Velocimetry.* Exp Fluids, 1991. **10**(4): p. 181-193.
3. Westerweel J, Dabiri D, and Gharib M, *The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings.* Exp Fluids, 1997. **23**(1): p. 20-28.
4. Scarano F, *Iterative image deformation methods in PIV.* Meas Sci Tech, 2002. **13**: p. R1-R19.
5. Meinhart CD, Wereley ST, and Santiago JG, *A PIV algorithm for estimating time-averaged velocity fields.* Journal of Fluids Engineering, 2000. **122**: p. 285-289.
6. Westerweel J, *Single-pixel resolution ensemble correlation for micro-PIV applications.* Exp Fluids, 2004. **37**: p. 375-384.
7. Adrian R, *Twenty years of particle image velocimetry.* Exp Fluids, 2005. **39**: p. 159-169.
8. Raffel M, Willert C, and Kompenhans J, *Particle Image Velocimetry: A Practicle Guide.* 3rd ed. 1998: Springer.
9. McKenna S and McGillis W, *Performance of digital particle image velocimetry processing techniques.* Exp Fluids, 2002. **32**: p. 106-115.
10. Westerweel J, *Universal outlier detection for PIV data.* Exp Fluids, 2005. **39**: p. 1096-1100.
11. Keane R and Adrian R, *Theory of cross-correlation analysis of PIV images.* App Sci Res, 1992. **49**: p. 191-215.
12. Huang, H.T., H.F. Fielder, and J.J. Wang, *Limitation and improvement of PIV, part II. Particle image distortion, a novel technique.* Exp Fluids, 1993. **15**: p. 263-273.
13. Hart D, *The Elimination of Correlation Errors in PIV processing.* 9th Int Symp App Laser Tech Fluid Mech, 1998.
14. Huang H, Dabiri D, and Gharib M, *On errors of digital particle image velocimetry.* Meas Sci Tech, 1997. **8**: p. 1427-1440.
15. Wernet M, *Symmetric phase only filtering: a new paradigm for DPIV data processing.* Meas Sci Tech, 2005. **16**(3): p. 601-618.
16. Adrian R and Yao C, *Pulsed laser technique application to liquid and gaseous flows and the scattering power of seed materials.* App Opt, 1985. **24**(1).
17. Adrian R, *Dynamic ranges of velocity and spatial resolution of particle image velocimetry.* Meas Sci Tech, 1997. **8**: p. 1393-1398.
18. Gui L and Wereley ST, *A Correlation-based continuous window-shift-technique to reduce the peak-locking effect in digital PIV image evaliation.* Exp Fluids, 2002. **32**: p. 506-517.
19. Nogueira J, Lecuona A, and Rodriguez PA, *Identification of a new source of peak locking, analysis and its removal in conventional and super-resolution PIV techniques.* Exp Fluids, 2001. **30**: p. 309-316.
20. Cholemari M, *Modeling and correction of peak-locking in digital PIV.* Exp Fluids, 2007. **42**: p. 913-922.
21. Meinhart CD, Wereley ST, and Santiago JG, *PIV measurements of a microchannel flow.* Exp Fluids, 1999. **27**: p. 414-419.
22. Meinhart C, Wereley ST, and Gray M, *Volume illumination for two-dimensional particle image velocimetry.* Meas Sci Tech, 2000. **11**(6): p. 809-814.
23. Olsen M and Adrian R, *Brownian motion and correlation in particle image velocimetry.* Opt Laser Tech, 2000. **32**(7-8): p. 621-627.
24. Charonko J, *Personal communication, artificial image simulations.* Virginia Polytechnic Institute and State University, 2007.
25. Westerweel J, *Theoretical analysis of the measurement precision in particle image velocimetry.* Exp Fluids, 2000. **29**: p. S3-S12.
26. Okamoto K, Nishio S, Saga T, and Kobayashi T, *Standard images for particle-image velocimetry.* Meas Sci Tech, 2000. **11**: p. 685-691.
27. Stanislas M, Okamoto K, and Kahler C, *Main results of the first international PIV challenge.* Meas Sci Tech, 2003. **14**: p. R63-R89.
28. Stanislas M, Okamoto K, Kahler C, and Westerweel J, *Main results of the second international PIV challenge.* Exp Fluids, 2005. **39**: p. 170-191.

29.   Nikias C and Petropoulou A, *Higher order spectral analysis: A nonlinear signal processing framework.* 1993: PTR Prentice Hall.

30.   Knapp C and Carter GC, *The Generalized Correlation Method for Estimation of Time Delay.* IEEE Trans Acous Speech Sig Proc, 1976. **24**(4): p. 320-327.

31.   Carter GC, *Coherence and Time Delay Estimation.* Proc. IEEE, 1987. **75**: p. 236-255.

32.   Stone H, Orchard MT, Chang EC, and Martucci SA, *A fast direct fourier-based algorithm for subpixel registration of images.* IEEE Trans Geo Remote Sens, 2001. **39**(10): p. 2235-2243.

33.   Proakis JG and Manolakis DG, eds. *Digital Signal Processing.* 4th ed. 2007, Prentice Hall.

34.   Fore L and Tung A, *Nonlinear Temporal Filtering of Time-Resolved Digital Particle Image Velocimetry Data.* Exp Fluids, 2005. **39**: p. 22-31.

35.   Nuttall AH, *Some windows with very good sidelobe behavior.* IEEE Trans Acous Speech Sig Proc, 1981. **ASSP-29**: p. 84-91.

36.   Eckstein, A. and P. Vlachos, *A robust phase correlation DPIV processing algorithm for time resolved measurements*, in *PIV symposium.* 2007: Rome, Italy.

37.   Born, M. and E. Wolf, *Principles of Optics.* 5th ed. 1975, New York: Pergamon Press.

38.   Meinhart CD and Wereley ST, *The theory of diffraction-limited resolution in microparticle image velocimetry.* Meas Sci Tech, 2003. **14**: p. 1047-1053.

39.   Olsen M and Adrian R, *Out-of-focus effects on particle image visibility and correlation in microscopic particle image velocimetry.* Exp Fluids, 2000. **Suppl.**: p. S166-S174.

40.   Desaubry C and Gervias P, *Characterization of a standard CCD video camera devoted to the design of an efficient particle image velocimetry instrument.* Flow Meas. Inst., 2000. **11**: p. 133-141.

41.   Eckstein A, Charonko J, and Vlachos P, *Phase Correlation Processing for DPIV measurements: Spatial Domain Analysis*, in *Proc ASME Fluids Conf.* 2007: San Diego, Ca.

42.   Eckstein A, Charonko J, and Vlachos P, *Phase Correlation Processing for DPIV measurements: Spectral Domain Analysis*, in *Proc ASME Fluids Conf.* 2007: San Diego, Ca.

43.   Satyaprakash K, *Personal Communication, DPIV experiment of pulsatile flow in stented vessel.* Virginia Polytechnic Institute and State University, 2007.

44.   Cardwell N, *Personal Communication, TRDPIV of turbulent rib-roughened channel flow.* Virginia Polytechnic Institute and State University 2007.

45.   Gui L, Longo J, and Stern F, *Biases of PIV measurement of turbulent flow and the masked correlation-based interrogation algorithm.* Exp Fluids, 2001. **30**: p. 27-35.

46.   Foroosh H, Zerubia J, and Berthod M, *Extension of phase correlation to subpixel registration.* IEEE Trans Image Proc, 2002. **11**(3): p. 188-200.

47.   Hoge W, *A subspace Identification to the phase correlation method.* IEEE Trans Med Img, 2003. **22**: p. 277-280.

48.   Loan C, *Generalizing the Singular Value Decomposition.* J Num Anal, 1976. **13**(1): p. 76-83.

49.   Hoge W and Westin C, *Identification of translational displacements between N-dimensional data sets using the high-order SVD and phase correlation.* IEEE Trans Image Proc, 2005. **14**(7): p. 884-889.

50.   Beal S and Sheiner L, *Heteroscedatic Nonlinear-Regression.* Technometrics, 1988. **30**(3): p. 327-338.

51.   Fox J, *Appendix*, in *Applied Regression.* 2002.

52.   Gui L, Merzkirch W, and Fei R, *A digital mask technique for reducing the bias error of the correlation-based PIV interrogation algorithm.* Exp Fluids, 2000. **29**: p. 30-35.