

Particle image velocimetry

Ajay K. Prasad

Department of Mechanical Engineering, University of Delaware, Newark, DE 19716-3140, USA

Particle image velocimetry (PIV) is the newest entrant to the field of fluid flow measurement and provides instantaneous velocity fields over global domains. As the name suggests, PIV records the position over time of small tracer particles introduced into the flow to extract the local fluid velocity. Thus, PIV represents a quantitative extension of the qualitative flow-visualization techniques that have been practised for several decades. The basic requirements for a PIV system are an optically transparent test-section, an illuminating light source (laser), a recording medium (film, CCD, or holographic plate), and a computer for image processing. This review article addresses the basics of the PIV technique such as PIV algorithms, optical considerations, tracer particles, illuminating lasers, recording hardware, errors in PIV measurements, and PIV vector processing.

SOME of the earliest quantitative velocity measurements in fluid flows were obtained using Pitot-static tubes. The subsequent introduction of hot-wire anemometers in the 1920s was a significant advance, especially in terms of probe miniaturization, frequency response, and the ability to measure multiple velocity components. However, both these techniques require the insertion of a physical probe which can intrude on the flow itself. The invention of the laser in the 1960s led to the development of the laser-Doppler anemometer which uses a laser probe to enable non-intrusive velocity measurements. Despite the rapid strides in the design of such systems, and the great sophistication of the associated electronics, one cannot escape the fact that all these techniques are at best *point-wise*, i.e. the velocity information is obtained only at the point occupied by the probe. While these techniques continue to retain an important position in an experimentalist's arsenal, the ability to make *global* velocity measurements has elevated particle image velocimetry (PIV) to a special status in fluid mechanics.

PIV can provide *instantaneous* velocity measurements over global (2D or 3D) domains with high accuracy. In this sense, PIV is a logical extension – from qualitative to quantitative – of the classical flow-visualization techniques that have been practised for decades. According to the classification proposed by Hinsch¹, a measurement system can be labelled as (k, l, m) , where $k = 1, 2, 3$ indicates the number of velocity components measured,

$l = 0, 1, 2, 3$ indicates the number of spatial dimensions of the measurement domain, and $m = 0, 1$ indicates instantaneous or continuous time recording, respectively. Accordingly, even the best point-wise techniques can only attain a $(3, 0, 1)$ status. In contrast, the simplest form of PIV provides two-dimensional velocity data on a planar domain at discrete time instants, i.e. $(2, 2, 0)$. The majority of PIV systems in use today belong to this category. At the other end of the spectrum, cinematic holographic PIV would belong to the $(3, 3, 1)$ category, although the author is not aware of any groups that have successfully developed such a system. Even holographic PIV systems conforming to $(3, 3, 0)$ have been pursued by a mere handful of groups^{2,3} owing to their high cost and complex implementation. In contrast, stereoscopic PIV systems $(3, 2, 1)$ which provide 3D velocity data on planar domains are becoming increasingly popular⁴⁻⁶.

PIV requires four basic components (Figure 1): (1) An optically transparent test-section containing the flow seeded with tracer particles; (2) A light source (laser) to illuminate the region of interest (plane or volume); (3) Recording hardware consisting of either a CCD camera, or film, or holographic plates; (4) A computer with suitable software to process the recorded images and extract the velocity information from the tracer particle positions.

Commercial vendors such as TSI, Dantec, and others, nowadays sell entire PIV systems comprising items (2)–(4), however, they can be expensive. It is not difficult to build such a system in-house, and this paper will attempt to provide some directions towards this objective. Obviously, a number of options are available to the experimentalist to create a new PIV system: For example, based on simplicity, cost, and the nature of the experiment, one might prefer a conventional single camera PIV, or stereo-

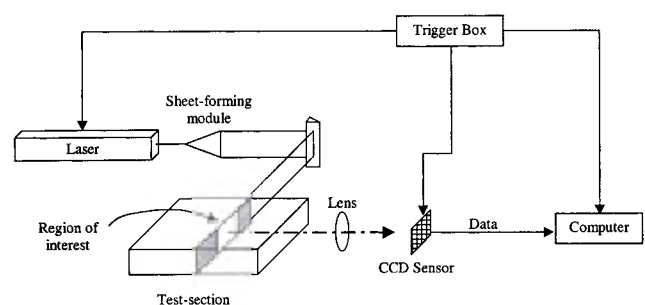


Figure 1. Basic requirements for a PIV system.

scopic PIV, or even holographic PIV. Similarly, one must decide between a continuous wave laser and a pulsed laser. Further, the recording hardware could be film or CCD, although CCD offers a number of advantages and is the preferred choice today. Finally, there is a choice of computer platforms, and algorithms.

A review of PIV literature reveals a great variety of PIV algorithms that have been developed over the years. The scope of this article is too limited to delve into greater details. The primary focus of this article will be conventional, single-camera correlation-based PIV. However, as an introduction, it is useful to first examine an earlier, simpler algorithm based on particle tracking.

Particle tracking velocimetry

Particle tracking velocimetry (PTV) is a direct descendent of flow visualization using tracer particles in fluid flows. If the particles are illuminated by two successive bursts of light, each particle produces two images on the same piece of film. Subsequently, the distance between the images can be measured to approximately determine the local Eulerian velocity of the fluid. Early measurements of particle displacements were made by hand using blow-ups of photographs of particle-laden flow fields. As technology advanced, these photographs could be electronically digitized and the determination of particle displacements was automated. Current practitioners prefer to directly obtain digital measurements using CCD cameras, thus bypassing the photographic film and the darkroom.

The PTV algorithm has to be sophisticated enough to discern a genuine particle image from noise, and to correctly match the twin images belonging to the same particle. The typical process consists of thresholding the raw image field to eliminate noise, and computing centroids of bright pixels. Difficulties arise due to overlapping or fragmented particle images, or if the second image of an image-pair is lost when the particle exits the light sheet between exposures due to out-of-plane motion. Typically, PTV performs better when particle concentration is somewhat low, i.e. when the inter-particle distance is large compared to the displacement of the same particle between exposures. Under this condition, the probability that two neighbouring images belong to the same particle is higher, and matching pairs is simplified. The implication is that the data density is somewhat lower. Secondly, velocity vectors are located randomly in the region of interest, as it is impossible to control particle positions in the flow field. Consequently, it is necessary to resort to interpolation to obtain vectors on a uniform grid in order to calculate vorticity and other such information.

Correlation-based PIV

In contrast to particle-tracking algorithms, correlation-based PIV does not require the matching of individual

images belonging to a pair. Instead of determining the displacement of individual particles, correlation-based PIV determines the *average* motion of *small groups* of particles contained within small regions known as *interrogation spots*. Essentially, the overall frame is divided into interrogation spots, and the correlation function is computed sequentially over all spots providing one displacement vector per spot (Figure 2). Typically, interrogation spots are square-shaped and therefore, the velocity map obtained from PIV presents vectors arranged on a uniform grid. More importantly, the process of averaging over multiple particle pairs within an interrogation spot makes the technique remarkably noise-tolerant and robust in comparison to PTV. The determination of the average particle displacement is accomplished by computing the spatial auto-correlation, or preferably the spatial cross-correlation of the particle images.

Auto-correlation is performed when images from both laser pulses are recorded on the same sensor, i.e. the sensor shutter stays open during the time that both laser bursts occur. Such a recording is called single-frame/double-pulse⁷. The auto-correlation function $R(\mathbf{S})$, of the intensity pattern, $I(\mathbf{X})$, of the interrogation spot is

$$R(\mathbf{S}) = \int_{\text{spot}} I(\mathbf{X})I(\mathbf{X} + \mathbf{S}) d\mathbf{X}. \quad (1)$$

A direct computation of the auto-correlation function by evaluating eq. (1) is prohibitively expensive. Instead, the auto-correlation function is computed via a two-dimensional Fast Fourier Transform (FFT) of the digitized intensity pattern.

$$\begin{aligned} R(\mathbf{S}) &= I(\mathbf{X}) \star I(\mathbf{X}) \\ &= I(\mathbf{X}) * I(-\mathbf{X}), \end{aligned}$$

where \star represents the correlation operation, and $*$ the convolution operation. Using the convolution theorem:




Figure 2. Correlation-based PIV.

$$\begin{aligned} F\{R(S)\} &= F\{I(X)\} \cdot F^*\{I(X)\} \\ &= |F\{I(X)\}|^2, \end{aligned}$$

where \mathcal{F} denotes the Fourier transform. Therefore,

$$R(S) = F^{-1}\{|F\{I(X)\}|^2\}. \quad (2)$$

Figure 3 *a* shows a computer-generated 64×64 interrogation spot containing about 15 particle image pairs. Images are Gaussian ($\sigma = 2$ pixels) and the chosen displacement is 8 pixels in both the x and y directions. An arrow indicates the displacement vector. It is apparent that some particles overlap, and yet others are truncated at the edge of the interrogation region. The spatial auto-correlation of the particle field displayed in Figure 3 *b* reveals that the tallest peak is located at the origin, and corresponds to the self-correlation peak. In addition, there exist two smaller signal peaks, S^+ and S^- ; the location of the signal peak with respect to the self-correlation peak provides the x and y components of displacement. The signal peak can be located with sub-pixel accuracy ($\approx 1/10$ pixel) by fitting a parabola or a Gaussian curve to the discrete data.

From Figure 3 *b* it is apparent that the maximum displacement that can be measured by the auto-correlation technique is half of the interrogation spot dimension. In reality, the signal-to-noise ratio (SNR) begins to degrade for particle displacements exceeding about one-fourth of the interrogation spot size due to a phenomenon called in-plane loss-of-pairs: At larger particle separations, the probability that both images of a particle pair reside within a given interrogation spot begins to decrease, leading to a drop in SNR⁸. A suitable time separation (Δt) between exposures for a given experiment can be approximated by knowing (roughly) the expected flow velocity in the object plane (U), the recording magnification (M), and the interrogation spot size in the image plane (d_I). A general guideline can be proposed for the auto-correlation technique:

$$\Delta t \leq 0.25 \frac{d_I}{MU}. \quad (3)$$

As seen in Figure 3 *a* the auto-correlation function is rotationally symmetric. The unfortunate consequence is that, a signal located at (S_x, S_y) in the auto-correlation plane is replicated by an identical signal located at $(-S_x, -S_y)$ leading to *directional ambiguity* of the velocity vector: Because both signal peaks (S^+ and S^- in Figure 3 *b*) are of identical height, one cannot select between them. Consequently, while the *magnitude* of the particle displacement is known (S^+ and S^- are equidistant from the origin), an incorrect peak choice implies that the *direction* of the displacement would be opposite to the true direction. Basically, such a double-exposed PIV frame

does not contain explicit information indicating whether the particle images were moving from left to right or vice versa.

This feature can be tolerated if the general direction of the flow were known *a priori*, for example, in turbulent channel flow where reversed flow is non-existent; in this case, it is possible to select S^+ and discard S^- (or vice versa) for every interrogation spot. But even in the




Figure 3. Spatial auto-correlation of I (note the rotational symmetry of R). *a*, Particle image field; *b*, Auto-correlation field.

absence of reversed flow, the auto-correlation technique cannot resolve particle displacements smaller than the particle image diameter, because the self-correlation peak can encroach upon, and even overwhelm the signal peak.

A variety of schemes exist to resolve such directional ambiguity, the simplest being to translate the flow field (or the camera⁶) during recording such that the translation speed exceeds the speed of any reversed flow. In this way, the entire flow is forced to become uni-directional; the translation speed is then subtracted as the final step to resolve regions of reversed flow. Additional schemes include placing an oscillating mirror between the test-section and the camera^{9,10}. The mirror rotates during the exposure, thereby causing the test-section to 'fly' at a speed exceeding any back-flow. Other schemes exploit the birefringent property of certain materials like calcite¹¹. For this technique, the illuminating laser pulses must have opposite polarizations; a calcite plate placed in front of the camera directs the light scattered by the tracer particles (which must retain their opposite polarizations) along different paths which provides the required shift. However, such schemes are cumbersome at best.

The trend now is to eschew single-frame/double-pulse recording in favour of double-frame/single pulse recording⁷ according to this sequence: (1) The shutter on the sensor opens to admit the first pulse, (2) the image is stored in a buffer, (3) the shutter opens again to admit the second pulse, and (4) both frames are transferred to memory. Obviously, the laser, camera and computer must be triggered in the correct sequence. In this case, each frame contains images from either the first pulse or the second pulse but not both. Because the identity of the frame containing the first pulse is known, the process of cross-correlating the first frame with the second yields a unique signal peak without directional ambiguity. Typical sensors used for PIV perform step (2) very rapidly, such that the pulse separation may be shortened to allow the measurement of even high-speed aerodynamic flows.

The cross-correlation function $C(\mathbf{S})$ of the intensity patterns $I_1(\mathbf{X})$ and $I_2(\mathbf{X})$ of interrogation spots 1 and 2 is

$$C(\mathbf{S}) = \int_{\text{spot}} I_1(\mathbf{X}) I_2(\mathbf{X} + \mathbf{S}) d\mathbf{X}. \quad (4)$$

Using the FFT, it is shown that

$$C(\mathbf{S}) = F^{-1}\{F\{I_1(\mathbf{X})\} \cdot F^*\{I_2(\mathbf{X})\}\}.$$

Figure 4 *a* and *b* shows the same particle image field as Figure 3 *a*, but now the particles' first and second images are separated into different frames according to the double-frame/single-pulse recording. The arrow is retained to indicate the displacement vector. The spatial cross-correlation of the two-particle field is shown in Figure 4 *c*. It is obvious that only one signal peak is produced; the self-correlation peak has disappeared, and so has the duplicate signal peak. This implies that directional

ambiguity is eliminated, and furthermore, even zero displacements can be recorded. A third benefit of cross-correlation is an improvement in the SNR compared to auto-correlation. For these reasons, cross-correlation is the preferred choice. A thorough treatment of the theoretical aspects of correlation-based PIV may be found in ref. 12.

The desire to resolve the *Kolmogorov* scale in turbulent flows implies that the interrogation spot should be as small as possible. At the same time, the desire to capture the *integral* length scales in a turbulent flow implies that the overall view be as large as possible. Obviously the two requirements are in conflict with each other and a suitable compromise may become necessary in some situations.

For example, consider a PIV image acquired using a commonly available $1\text{ k} \times 1\text{ k}$ pixel CCD sensor. A typical interrogation spot would correspond to a 32×32 pixel square region. The lower bound on interrogation spot size is determined by the requirement that it should contain at least 5 to 10 particle images with an ideal particle image diameter of about 2 pixels. As a result, the range of spatial scales that can be typically resolved is about 30. It should be mentioned that $2\text{ k} \times 2\text{ k}$ sensors are now becoming affordable and algorithms can be fine-tuned to allow 16×16 pixel interrogation spots, which push the range of scales to over 100. Ref. 13 provides additional details regarding the dynamic ranges of velocity and spatial scales in PIV.

Optical considerations

Particle image diameter

With reference to Figure 5 some useful relations may be written for the object distance d_o and the image distance d_i .

$$d_o = (1 + M^{-1})f,$$

$$d_i = (1 + M)f,$$

$$M = d_i / d_o,$$

where M is the image magnification and f is the focal length of the lens.

In order to determine the particle image diameter at the image plane, one must consider both geometric and diffraction effects. Geometric considerations indicate that the image diameter should be Md_p where d_p is the particle diameter in the object plane. However, for typical recording situations, diffraction is important. Due to diffraction, a point source in the object field at O' will spread out into an Airy function on the image plane centred at O'' . The diffraction limited spot-size, d_s , of a particle is equal to the diameter of the Airy disk (Figure 5) and is governed by the diameter of the lens aperture D , the focal length of the lens f , and the wavelength λ of the illuminating laser:




Figure 4. Spatial cross-correlation between I_1 and I_2 . *a*, Particle image field 1; *b*, Particle image field 2; *c*, Cross-correlation field.

$$d_s = 2.44(1 + M) \frac{f}{D} \lambda.$$

The ratio f/D is the f -number of the lens. When the two effects are combined, the effective particle image diameter d_e is given by

$$d_e = (M^2 d_p^2 + d_s^2)^{1/2}.$$

In situations where $Md_p \gg d_s$, $d_e \approx Md_p$; conversely, if $Md_p \ll d_s$, $d_e \approx d_s$. For example, using $M = 0.2$, $\lambda = 532$ nm for frequency-doubled Nd : YAG laser light, and $f/8$, one obtains $d_s = 12.5$ μm . For $d_p = 8$ μm , we get $Md_p = 1.6$ μm , and $d_e = 12.6$ μm . Thus, for these values, diffraction dominates.

Depth-of-field

A second important relation involves the depth-of-field δz of the recording optics. The depth-of-field is the thickness of the region over which particles will be in acceptable focus (see Figure 5).

$$\delta z = 4(1 + M^{-1})^2 (f/D)^2 \lambda.$$

For the values used above, one obtains, $\delta z = 4.9$ mm. It is seen that the depth-of-field increases rapidly with $1/M$ and the f -number. Ideally, δz should not be smaller than the thickness of the light sheet in order to avoid imaging out-of-focus particles.

Creation of a light sheet

The simplest method to create a light sheet is to use a cylindrical lens in combination with a spherical lens as shown in Figure 6. The cylindrical lens causes the laser beam to expand in one direction only, i.e. it 'fans' the beam out. The spherical lens causes the expanding beam to focus along the perpendicular direction, at a distance of one focal length downstream, to its beam waist. In Figure 6 (side view) it is seen that the light sheet continues to expand, albeit at a reduced rate, downstream of the spherical lens; it is possible with a judicious choice of focal lengths to obtain a collimated sheet, i.e. one that continues with a constant beam height H . This prevents a dilution of light intensity with distance. The intensity of light illuminating the particles depends on the laser power and inversely on H and the light sheet thickness Δz_0 . Because laser beams have a Gaussian intensity profile, the Gaussian variation is retained in both the in-plane and the thickness directions.

For example, if one wished to measure a 10 cm \times 10 cm region in a flow, one could employ a – 20 mm focal length cylindrical lens with a 1000 mm focal length spherical lens. The light sheet thickness for this arrangement is

about 1 mm. Ref. 14 provides a detailed explanation of the various aspects of PIV recording.

Particles for PIV

Tracer particles for PIV must satisfy two requirements: (1) They should be able to follow the flow streamlines without excessive slip, and (2) they should be efficient scatterers of the illuminating laser light. While the first requirement is fairly obvious, the second requirement significantly impacts the illuminating lasers and recording hardware. For example, if a given particle scatters weakly, then one would have to employ more powerful lasers or more sensitive cameras, both of which can drive up costs, as well as the associated safety issues. Although the search for ideal particles may seem somewhat trivial, it can potentially provide immense benefits.

A simple way to evaluate the first requirement is to determine the settling velocity of the particle under gravity. Assuming that the process is governed by Stokes drag, the settling velocity u_∞ is given by:

$$u_\infty = \frac{gd_p^2(\rho_p - \rho_f)}{18\mu},$$

where d_p and ρ_p are the particle diameter and density respectively, and μ and ρ_f are the fluid viscosity and density respectively. Particles are suitable as long as u_∞ is negligible compared to the actual flow velocities. When working with low-speed liquid flows, it is beneficial to use particle densities close to liquid density; for example, polystyrene and other plastic particles which have densities within a few per cent of water are good candidates for water flows. Even here, one must ensure that d_p is not too large. Unfortunately, most materials with specific gravity ≈ 1 also have a refractive index similar to that of water, therefore these materials scatter weakly. Some vendors supply silver-coated particles which can be expensive. A low-cost and highly effective alternative is to use hollow glass spheres. These are readily available in sizes ranging from a few μm to hundreds of μm from suppliers who sell them in bulk as fillers for lightweight composite parts. The




Figure 5. Diffraction effects during particle image recording.

glass/air interface *inside* the glass bubble gives rise to strong scattering due to the large refractive index mismatch.

For gaseous flows, typically oil droplets are used (silicone oil or organic oils such as olive oil). Owing to the very large density difference between oil and the gaseous medium, it becomes essential to use very small droplets, typically $< 1 \mu\text{m}$, to minimize settling velocities. (Even then, it is difficult to prevent centrifuging of particles out of vortex cores.) The usual method of producing such fine oil droplets is with a Laskin nozzle (Figure 7), in which air is bubbled under oil inside a bottle. As the bubbles rise through the surface they burst and produce fine droplets. The smallest bubbles may then be selected by forcing the droplets to navigate the small annular gap around a concentric baffle plate; larger droplets are unable to survive the turn due to inertia and are removed by impaction.

The ultimate goal in particle selection is to obtain the most efficient scattering and maximize the light intensity on the sensor. When operating in the diffraction limit, $Md_p \ll d_s$; consequently, one can increase the particle size to increase the scattered light while maintaining the same particle image diameter on the sensor. The implication is that the image *intensity* will increase as the particle diameter increases as long as one operates within the diffraction limit. Under this condition, it is advantageous to use the largest particle which can follow the flow without excessive slip. The reader is referred to ref. 15 for additional details of PIV tracer particles.

Illuminating laser

PIV can be accomplished using continuous wave (CW) lasers or, more optimally, pulsed lasers. Typical CW lasers used for PIV are argon-ion lasers producing in the range of a few watts; for some less demanding applications, even low-powered helium-neon lasers might suffice. Typical pulsed lasers are frequency-doubled Nd:YAG (neodymium:yttrium aluminum garnet) lasers producing 0.1 to 0.3 joules/pulse, at a repetition rate of tens of Hz.




Figure 6. Schematic of a simple sheet-forming module.

Pulses can be created from a CW laser using a chopper, or by rapidly sweeping the entire laser beam through the test-section by reflecting it off a rotating polygonal mirror with all but two facets covered (Figure 8). In the case of the chopped beam, a light sheet must still be created as described earlier using a sheet-forming module. In the case of the swept-beam, a sheet-forming module is unnecessary because the sweeping action of the beam creates the 'sheet'; although the light sheet is not an instantaneous entity, it constitutes an adequate substitute if the beam is swept rapidly when compared to the unsteadiness of the flow.

The advantage of pulsed lasers is the short duration of the laser pulse, typically a few to several nanoseconds. As a consequence, a particle travelling at even very high speeds is essentially 'frozen' during the exposure with minimal blurring (for example, a particle travelling at 100 m/s will move only 0.06 μm during a 6 nanosecond pulse). Pulsed lasers operate by discharging energy stored in capacitor banks at discrete time intervals to the flash-lamp followed by the laser pulse. Pulsed lasers are ideally suited for PIV because they store and deliver all of the laser power at exactly the desired instant. Pulsed lasers for PIV operate at a constant repetition rate of a few tens of Hz. Therefore, in order to obtain arbitrary Δt 's, one requires two identical lasers firing in tandem. As shown in Figure 9, the pulse train from the second laser can be suitably staggered in time with respect to the first laser to produce arbitrarily small Δt 's. A delay box producing the required trigger signals must be incorporated to ensure that the lasers, camera, computer, and other hardware are synchronized. With suitable hardware, one can obtain cinematic PIV data with a frame-rate equal to the repetition rate of the lasers. Of course, when working with two lasers, extreme care must be exercised to ensure that the two laser sheets exactly overlap inside the test-section. A lack of such overlap can lead to poor correlation and unsatisfactory results.




Figure 7. Schematic of Laskin nozzle to generate micron-sized oil droplets.

In contrast, CW lasers deliver energy on a continuous basis and 'pulsing' is obtained by either chopping or sweeping the beam. In the case of a chopped laser beam, the light energy incident on a particle increases with the duration of the pulse so created. However, if the pulse duration is too long the particle will produce streaks rather than crisp circular images. (To a small extent, streaky images can be tolerated in correlation-based PIV.) Unfortunately, the brightness of the particle image on the recording medium reaches a saturation level if the particle travels more than one diameter during the exposure owing to the fact that the scattered energy is being smeared over a greater area on the sensor. For these reasons, it is not beneficial to increase the pulse duration beyond some limit determined by the flow speed. A similar restriction applies with swept-beam systems, though the exact mechanism differs slightly. In the latter case, the intensity is higher because the beam is not 'fanned out', however, the duration that the beam illuminates a given particle may be reduced because the beam must be swept rapidly enough to maintain a quasi-instantaneous recording.

In general, the amount of CW laser energy incident upon a particle decreases as the flow velocity increases, which places an upper bound on the flow velocity. In contrast, in the case of pulsed lasers, the quality of particle images is immune to the flow velocity. While pulsed lasers are optimally suited for PIV, many fluid mechanics laboratories may already possess argon-ion lasers which could be implemented in a low cost illumination system for low-speed flows.

A chopped CW laser was used for PIV measurements in non-penetrative turbulent convection¹⁶. A description of

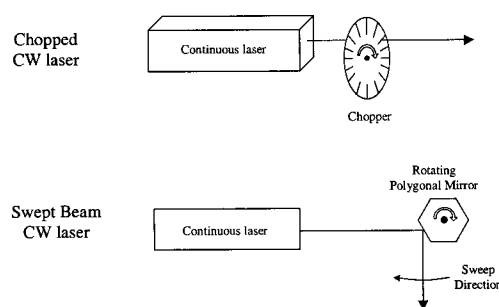


Figure 8. Obtaining a pulsed laser sheet from continuous wave lasers.




Figure 9. Staggered pulse trains from twin pulsed lasers for PIV.

the swept-beam technique can be found in Rockwell *et al.*¹⁷; the technique was applied to a stalled delta wing subjected to small amplitude pitching oscillations¹⁸ with water as the working fluid.

Recording and processing hardware

PIV images may be recorded on a film or on CCD (charge-coupled device) sensors, whereas holographic PIV requires the use of holographic plates. The popular trend nowadays is to use CCD cameras for PIV recordings. However, photographic film may still be a viable choice when very high resolutions are required. For example, the commonly available 8-bit $1\text{ k} \times 1\text{ k}$ CCD sensors can store only 1 Mbyte per frame. In contrast, high resolution films such as Kodak Technical Pan (300 line pairs per mm) can potentially store about 80 Mbytes on a 35 mm frame; large format film (100 mm \times 120 mm) can potentially store over one Gbyte. It must be noted however, that the high density of film storage may not be fully exploitable during the process of electronic retrieval, because typical scanners have a resolution of about 10 μm . (Higher scanning resolutions may not yield corresponding benefits in the final results especially if $d_e/d_{\text{pix}} \geq 2$.) Consequently, 35 mm film will yield about 8 Mbytes of usable data.

While the film outscores CCD sensors in terms of resolution, it loses out severely in terms of convenience. The turnaround time involved in recording, developing the film and finally digitizing it can take a long-time. Such delays are most inconvenient when setting up a PIV experiment, when several trips to the darkroom may be needed to optimize experimental conditions. In contrast, CCD sensors allow the same set-up process to be accomplished in a few seconds. Digital PIV results can be viewed in almost real-time with high-speed on-board computational hardware dedicated towards computing correlations.

A second advantage with CCD sensors is that they are far more light-sensitive than the film. Superior sensitivity allows one to perform the same experiment with a far less powerful laser, producing cost benefits and a safer working environment. CCDs also provide a linear response to light intensity in contrast to the film which responds logarithmically. Owing to these features, our experience suggests that CCD sensors allow interrogation with smaller spots (in terms of pixels) than the photographic film. Consequently, the resolution advantage of the film is somewhat blunted, except in the case of large format film.

CCD sensors have been developed specifically for PIV (such as the $1\text{ k} \times 1\text{ k}$ pixel Kodak ES 1.0). The dual channel version of this camera allows double-frame/single-pulse recording which enables cross-correlation between frames. For certain applications, Δt may be very small (microseconds). It is not possible to read out 1 Mbyte of data to the computer memory in such a short time, therefore, the first frame is stored with extreme rapidity into a buffer pixel

located immediately adjacent to every sensor pixel. The second PIV frame is subsequently recorded, and the two frames are read out sequentially to the computer memory in time for the next double-frame capture. The cycle is completed in time T which is the period of laser pulsing.

The most popular choice for PIV platforms is the PC. Several vendors provide frame-grabbers to interface the CCD camera to the computer. The processing of PIV frames can also be performed efficiently on the PC, especially using dedicated array processing hardware. The basic requirement is that the FFT operation be performed extremely rapidly. Similarly, the hardware must also support high-speed graphics to enable the visualization of the particle field, and the plotting of vector fields in almost real time.

Errors in PIV

PIV measurements contain errors arising from several sources: (1) *Random* error due to noise in the recorded images; (2) *Bias* error arising from the process of computing the signal peak location to sub-pixel accuracy; (3) *Gradient* error resulting from rotation and deformation of the flow within an interrogation spot leading to loss of correlation; (4) *Tracking* error resulting from the inability of a particle to follow the flow without slip; (5) *Acceleration* error caused by approximating the local Eulerian velocity from the Lagrangian motion of tracer particles.

Certain errors can be minimized by careful selection of experimental conditions (for example, tracking error). However, other sources are inherent to the nature of the correlation in PIV and cannot be eliminated. For example, even if the recorded images are free from noise, the location of the correlation peak can be influenced by random correlations between particle images not belonging to the same pair (such random correlations are visible as small peaks in Figures 3 *c* and 4 *b*). In addition, bias errors result from a phenomenon called pixel-locking in which the signal peak location is biased towards the nearest pixel while using a curve-fit or centroiding schemes to locate the discretized signal with sub-pixel accuracy. Similarly, gradient errors will occur in turbulent flow. Finally, acceleration error cannot be eliminated because of the very principle of PIV which uses the Lagrangian motion of particles to approximate the instantaneous Eulerian flow velocity.

Previous work has shown that random errors in PIV usually scale with the particle image diameter⁷.

$$\sigma_{\text{random}} = c d_e,$$

where d_e is the effective particle image diameter, and c is a constant whose value is between 0.05 and 0.10 depending upon experimental conditions^{19,20}.

Bias errors arise during the process of calculating the particle displacement to sub-pixel accuracy. Essentially, the correlation field is available on a discretized grid (typically 32×32). The location (S_x, S_y) of the maximum

value in the correlation field will correspond to the particle displacement, but obviously, such a displacement will be an integer pixel value. In order to reduce measurement error, one attempts to locate the peak with sub-pixel accuracy using either a curve-fit method or a centroiding scheme. Unfortunately, any such scheme causes both random errors and bias errors. Our experience shows that while centroiding schemes perform poorly on both counts in comparison to a parabolic or Gaussian curve-fitting, the latter can also cause significant bias errors.

The term bias error was coined to describe the phenomenon of pixel-locking, i.e. during the process of determining the displacement to sub-pixel accuracy, the resulting value is always biased towards the nearest integer-valued pixel¹⁹. Bias error is zero if the particle is displaced exactly n pixels or $n + 0.5$ pixels, where n is an integer. For displacements $n < \Delta X < n + 0.5$ pixels, the measured displacement is biased towards n , and for $n + 0.5 < \Delta X < n + 1$ pixels, the measured displacement is biased towards $n + 1$. The situation is depicted in Figure 10 where a linear velocity profile (Couette flow) actually gets measured as a 'staircase' profile. The true and measured values coincide at integer and at half-pixel values.

It is of course possible to predict the amount of bias error using modelling or measurement, implying that it should be possible to develop a look-up table to retrieve the correct displacement from the measured displacement just by subtracting the appropriate value of the bias error. Unfortunately, any measured result will also include random, gradient, and acceleration errors in addition to bias error. By their very nature, these other errors are not deterministic. Bias error may therefore only be corrected if it greatly dominates over the other errors. Bias errors, if present, become very apparent when histograms of displacements are plotted. Pixel-locking is seen when data-points gather in clusters near integer pixel values. Bias errors become large when the particle image diameter approaches *small* values (≈ 1 pixel). On the other hand,

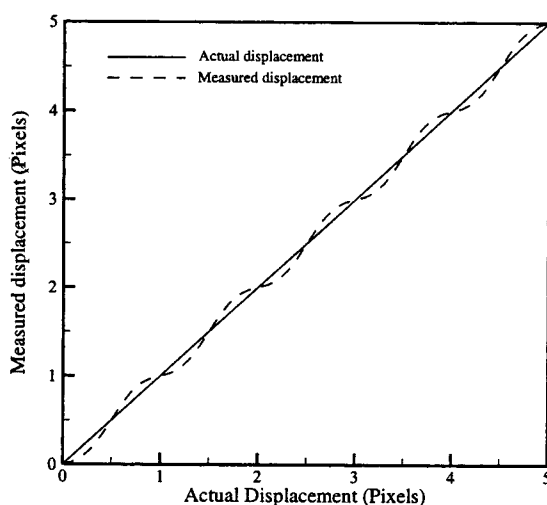


Figure 10. Bias error due to pixel-locking.

random errors increase when d_e increases. A suitable compromise, as verified by several practitioners, is to keep d_e at about 2 pixels.

In typical flows, particle streamlines are not straight, but curved. Consequently, an error is incurred when the Eulerian velocity is approximated using Lagrangian particle displacements as depicted in Figure 11. It is easy to see that if the pulse separation Δt is larger, then so is ΔX , and greater is the deviation of the measured velocity from the true velocity. This suggests that one must reduce Δt . On the other hand, if one were to reduce Δt too much, the measurement of ΔX becomes difficult; at very small values of Δt , it may become impossible to distinguish ΔX from random error. Obviously, simple error considerations suggest that the fractional random error can be reduced only by increasing ΔX , and correspondingly Δt . So what is the value of Δt at which the overall error is minimized? (Note that the guideline suggested by eq. (3) may be removed by cross-correlating between *shifted* interrogation spots^{20,21}, which implies that some other mechanism must decide Δt .) The optimal pulse separation was derived²⁰ as:

$$\Delta t_{\text{opt}} = \sqrt{\frac{2\sigma_{\text{random}}}{Ma}} = \sqrt{\frac{2cd_e}{Ma}},$$

where a is the local acceleration of the particle. Of course, a is unknown until the flow is measured, therefore, a two-step measurement may be required, wherein the first measurement is conducted with a guessed Δt in order to determine a , followed by a second measurement with an optimized Δt . It should be mentioned that not all regions of a flow field possess the same value of a , hence, additional optimization may be necessary.

Vector processing

As mentioned earlier, PIV processing typically provides vectors on a uniform grid. A $1\text{ k} \times 1\text{ k}$ pixel frame interrogated using 32×32 pixel spots will yield a 32×32 field of *independent* vectors. The term independent refers to the fact that the data contributing to a particular vector is not shared with any other vector. Very often however, interrogation is performed with an 'overlap' between adjacent spots. The commonly accepted value of overlap is 50%. Accordingly, the centre of each interrogation spot is displaced $d_1/2$ with respect to its neighbours, in both x




Figure 11. Acceleration error: approximation of the local Eulerian velocity using Lagrangian particle displacement.

and y directions, where d_I is the size of an interrogation spot. As a result, the total number of vectors increases by a factor of 4. It should be noted that overlapped vectors are no longer independent, but share 50% of the contributing particle pairs with the neighbouring vectors. Some early workers used an overlap of 75% giving an increase in vector count by a factor of 16, however, it is now accepted that this is excessive.

Even when experimental conditions are ideal, a PIV vector map will contain 'bad' vectors. The so-called bad vectors (also referred to as false or spurious vectors) are readily identifiable when the vector field is replotted after subtracting the mean; bad vectors have magnitudes and/or directions which are substantially different from their neighbours. Bad vectors result from interrogation spots in which the SNR is less than unity, i.e. a noise peak is higher than the signal peak. Typically, less than 2% of vectors will be bad. The usual causes are the lack of particle pairs in the interrogation spot due to inadequate seeding density, or excessive out-of-plane motion such that the particle exits the light sheet between laser pulses. Occasionally a piece of debris, a large bubble or flare from physical boundaries or objects protruding into the measured region will overwhelm the interrogation spot leading to a bad vector.

Obviously, PIV data must be validated (i.e. spurious vectors must be filtered out) before computing vorticity and other derived quantities. The process consists of identifying the bad vector and replacing it by interpolation from the neighbouring vectors. It is not difficult to automate the search. The human eye detects bad vectors by comparing a vector with its immediate neighbours. Similarly, most automated detection schemes implement a procedure whereby the mean and standard deviation are calculated for a 3×3 box surrounding a given vector, and checking if the given vector lies within some neighbourhood of the mean; the size of the neighbourhood is proportional to the local standard deviation. Most often, it is preferable to use the median rather than the mean, because the mean can be thrown off by the presence of a large bad vector in the 3×3 box. Some schemes implement multiple passes with the criterion for an acceptable vector being tightened with each pass. Ref. 22 discusses such data validation techniques in detail.

The PIV algorithms described so far are rather basic. It is not difficult to increase the level of sophistication to improve resolution (i.e. increasing the total number of vectors from a frame), and to also reduce the number of bad vectors. For the latter, some image pre-processing can be applied (for example, thresholding). Improving resolution from 32×32 to say 16×16 pixel interrogation spots requires some extra effort. As the spot size decreases by a factor of 2, the number of contributing pairs drops by a

factor of 4 leading to decreased SNR. The usual procedure is to use results from an earlier pass at coarse resolution to fine-tune the peak search routine during the second pass at higher resolution. In the limit, it may even be possible to match *each* particle image with its pair, leading ironically to a convergence between correlation-based PIV and particle tracking. Such a limit is called super-resolution²³.

Summary

A review of particle image velocimetry has been presented. The review examines the basic components of PIV such as PIV algorithms, optical considerations, tracer particles, illuminating lasers, recording hardware, typical errors associated with PIV and PIV vector processing.

1. Hinsch, K. D., *Meas. Sci. Technol.*, 1995, **6**, 742–753.
2. Barnhart, D. H., Adrian, R. J. and Papen, G. C., *Appl. Opt.*, 1994, **33**, 7159–7170.
3. Meng, H. and Hussain, F., *Phys. Fluids*, 1995, **7**, 9–11.
4. Prasad, A. K., *Exp. Fluids*, 2000 (to appear).
5. Prasad, A. K. and Jensen, K., *Appl. Opt.*, 1995, **34**, 7092–7099.
6. Prasad, A. K. and Adrian, R. J., *Exp. Fluids*, 1993, **15**, 49–60.
7. Adrian, R. J., *Annu. Rev. Fluid Mech.*, 1991, **23**, 261–304.
8. Keane, R. D. and Adrian, R. J., *Meas. Sci. Technol.*, 1990, **1**, 1202–1215.
9. Hopkins, L. M., Kelly, J. T., Wexler, A. S. and Prasad, A. K., *Exp. Fluids*, 2000, **29**, 91–95.
10. Landreth, C. C., Adrian, R. J. and Yao, C. S., *Exp. Fluids*, 1988, **6**, 119–128.
11. Landreth, C. C. and Adrian, R. J., *Appl. Opt.*, 1988, **27**, 4216–4220.
12. Westerweel, J., *Meas. Sci. Technol.*, 1997, **8**, 1379–1392.
13. Adrian, R. J., *Meas. Sci. Technol.*, 1997, **8**, 1393–1398.
14. Stanislas, M. and Monnier, J. C., *Meas. Sci. Technol.*, 1997, **8**, 1417–1426.
15. Melling, A., *Meas. Sci. Technol.*, 1997, **8**, 1406–1416.
16. Prasad, A. K. and Gonuguntla, P. V., *Phys. Fluids*, 1996, **8**, 2460–2470.
17. Rockwell, D., Magness, C., Towfighi, J., Akin, O. and Corcoran, T., *Exp. Fluids*, 1993, **14**, 181–192.
18. Cipolla, K. M. and Rockwell, D., *AIAA J.*, 1995, **33**, 1256–1262.
19. Prasad, A. K., Adrian, R. J., Landreth, C. C. and Offutt, P. W., *Exp. Fluids*, 1992, **13**, 105–116.
20. Boillot, A. and Prasad, A. K., *Exp. Fluids*, 1996, **21**, 87–93.
21. Westerweel, J., Dabiri, D. and Gharib, M., *Exp. Fluids*, 1997, **23**, 20–28.
22. Nogueira, J., Lecuona, A. and Rodriguez, P. A., *Meas. Sci. Tech.*, 1997, **8**, 1493–1501.
23. Keane, R. D., Adrian, R. J. and Zhang, Y., *Meas. Sci. Technol.*, 1995, **6**, 754–768.

ACKNOWLEDGEMENTS. I thank Prof. Vijay Arakeri at the Indian Institute of Science, at whose suggestion I started to write this article during my sabbatical in Bangalore.

Received 7 April 2000; accepted 6 May 2000