Aquatic Ecology (2006) 40:463–479
DOI 10.1007/s10452-005-6567-z

Basics and principles of particle image velocimetry (PIV) for mapping biogenic and biologically relevant flows

Eize J. Stamhuis
Department of Marine Biology, University of Groningen, P.O. Box 31, 9750 AA Haren, The Netherlands; (e-mail: e.j.stamhuis@biol.rug.nl; phone: +31-50-3632078)

Received 9 September 2004; accepted in revised form 26 April 2005

Key words: Benthic boundary layer, BIOFLOW, Biogenic flows, Flow analysis, PIV

Abstract

Particle image velocimetry (PIV) has proven to be a very useful technique in mapping animal-generated flows or flow patterns relevant to biota. Here, theoretical background is provided and experimental details of 2-dimensional digital PIV are explained for mapping flow produced by or relevant to aquatic biota. The main principles are clarified in sections on flow types, seeding, illumination, imaging, repetitive correlation analysis, post-processing and result interpretation, with reference to experimental situations. Examples from the benthic environment, namely, on filter feeding in barnacles and in bivalves, illustrate what the experiments comprise and what the results look like. Finally, alternative particle imaging flow analysis techniques are discussed briefly in the context of mapping biogenic and biologically relevant flows.

Introduction

Particle image velocimetry (PIV) has been developed from the early 1980's onwards to map complete flow fields instantaneously (Adrian 1991). The novelty and great advantage of PIV are that it delivers high resolution flow velocity vector information of a whole plane in the flow in one time. All other methods known until the 1980's delivered point values, or low resolution vector diagrams of whole planes after repeating the experiment several times. Consequently, PIV has given new impulse to fluid dynamics research, in particular at unsteady flows that were difficult to map instantly beforehand (Raffel et al. 1998).

PIV has been developed in the fields of fundamental and applied fluid dynamics research and mechanical engineering. It has therefore mainly been used to study experimentally induced flows, such as water flow in confined areas or around streamlined objects, air flow around wing profiles and plane models, etc. (Adrian 1991; Stanislas et al. 2000). Since 1993, PIV has been applied to biogenic or biologically important flows, using systems that were either developed for this purpose (e.g. Stamhuis and Videler 1995; Müller et al. 1997) or adapted or tuned to the particular circumstances (e.g. Drucker and Lauder 1999; Zirbel et al. 2000). Due to its non-intrusive character, PIV can be characterised as an ideal tool for studying animal-generated flows, restricting disturbances and thereby non-natural behaviour to a minimum.

For the benthic zone, information on methods for flow studies, on flow facilities and on results are being exchanged within the European BIOFLOW network. This paper contributes to this exchange by explaining the basics and principles of PIV, being one of the more recently developed flow mapping methods that can well be applied to study flows in the benthic zone. It will focus on 2-dimensional (planar) PIV and is geared to scien-

tists that have no real experience with PIV and do not have a degree in physics, fluid dynamics or mathematics. For detailed background information on PIV, the following references may prove valuable: *Particle Image Velocimetry, a practical guide* (Raffel et al. 1998), *Particle Image Velocimetry* (Hinsch 1993), and the *Special Issue on PIV* of Measurement Science and Technology, Volume 8 Nr 12 (Jones 1997), in particular the first paper in this volume, on the *Fundamentals of Digital PIV* (Westerweel 1997). *Digital PIV, theory and application* (Westerweel 1993) is recommended for the theoretical backgrounds of Digital PIV (DPIV). Details on DPIV applied to animal generated flows and on the specific problems that one can run into when working with animals in flow studies, can be found in Stamhuis and Videler (1995), and Stamhuis et al. (2002). The nature of flow phenomena around biota in general and the basic relationships between organisms and their fluid environment are very well explained by Steven Vogel in his book *Life in Moving Fluids, the physical biology of flow* (Vogel 1994) which is recommended for basic background reading.

The following subjects will be covered successively: (1) 2D-PIV in a nut-shell, describing what 2D-PIV basically comes down to; (2) The basics and principles of 2D-PIV, with explanatory sections on every step when applying 2D-PIV; (3) Biological examples, shortly covering two examples from our studies on biogenic flows, (4) Other PIV-techniques, with an overview and discussion of recent developments such as Micro-PIV, Stereo-PIV, Holographic PIV, Super Resolution PIV and 3D Particle Tracking Velocimetry (3D-PTV); and (5) Conclusion. The treatment of sections 1, 2 and 3 is based mainly on of the author's own experience with reference to the literature and section 4 is a brief overview of the relevant literature.

PIV in a nut-shell

2D-PIV can be defined as 'Mapping of average displacements of groups of tracer particles over a short time interval in a fluid flow by correlating sub-images of two successive images of one illuminated plane of that flow'. In essence, 2D-PIV requires a flow seeded with particles, bright illumination of one plane of that flow and imaging of the particles in the plane. Successive or double-exposed images are analysed for displacements of groups of particles to derive vector flow fields. The flow fields can be post-processed to derive additional parameters describing e.g. the local character of the flow. A generalised PIV set-up is shown in Figure 1.

To find the displacements of groups of particles over a short time interval, two successively recorded images are compared. A small sub-area of the first image, usually called interrogation areas (IA), is compared with a sub-area at the same location in the second image using a cross-correlation technique. This results in the most probable displacement vector for that particular particle pattern (Figure 2). The process is repeated for all interrogation areas of the image pair resulting in a complete vector diagram of the flow studied. (Adrian 1991; Willert and Gharib 1991; Stamhuis and Videler 1995).

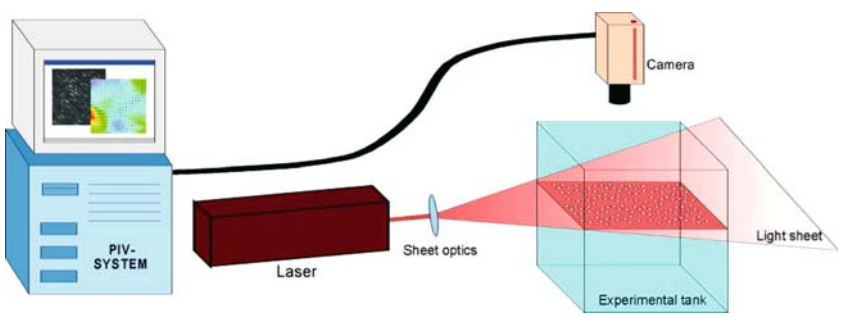

Figure 1. Diagram of a generalised PIV setup showing all major components: the experimental tank with the particle seeded fluid, a laser producing a light sheet that illuminates only one plane in the fluid, a camera imaging the particles in the sheet in the area of interest, and a PIV system that analyses the images to derive a vector representation of flow field.

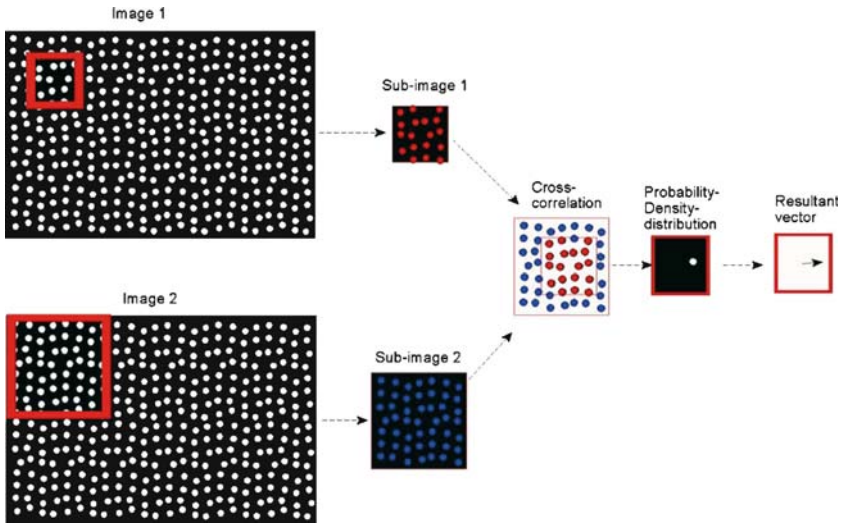

Figure 2. Diagram of the steps in PIV analysis of successively recorded particle patterns in a flow: two sub-images from the same location of two frames are compared in a cross-correlation procedure, resulting in a 2D probability density distribution which shows a peak at the most probable displacement. The peak is located with high precision and a velocity vector, representing the average displacement of the particles in sub-image 1 compared to sub-image 2 is calculated.

Basics and principles of PIV

Types of flow

Particle image velocimetry can be applied to virtually any kind of flow, as long as the fluid is transparent to enable imaging of suspended particles. The fluid may be moving, as in a flow tank or wind tunnel, or standing still, as in e.g. an aquarium in which an organism itself is moving. The fluid may be liquid or gaseous. Liquids such as water, but also mineral oil or glycerine, are easily seeded with neutrally buoyant particles and sinking of particles may only become relevant at very low flow speeds. Liquids allow a broad Reynolds' range (Stamhuis et al. 2002), ranging from 10^{-2} for small structures in still water, or even lower in oil or glycerine, to 10^{7} for intermediate speed flow tanks (Vogel and LaBarbera 1978; Vogel 1994). Gaseous media such as air are of course necessary when the flow around life organisms normally living in air is to be studied, as in e.g. flying birds (Spedding et al. 2003).

Seeding

To visualise flow for PIV purposes, the fluid has to be seeded with particles. To follow the flow accurately, these particles need to be neutrally buoyant and small with respect to the flow phenomena studied, (Raffel et al. 1998). For liquids, a whole series of polymer powders or particulate half-products for plastics with a range of densities are available, and can be matched to the fluid density. In water, e.g. polystyrene (unexpanded), polythene, nylon, PVC or pliolite particles work well and can be purchased from chemical companies (e.g. BASF, Dow, Dupont, Goodyear) in a range of diameters. The typical particle size for fluid flow applications ranges from 5 to 200 μm, depending on flow speed and camera magnification. Also hollow glass balloons in micrometer sizes (e.g. Nortek A.S, Potters Industries), sometimes coated with a reflective layer, can be used in aqueous flows.

Seeding is more difficult in low speed gas flows because particles tend to sink due to the relatively high density difference. The seeding particles thus need to be very light and small, e.g. very small polystyrene (expanded) beads (diameter ca. 10–5 μm). Alternatively, Helium filled soap bubbles of a few mm in diameter have been used to map flow around flying birds (e.g. Spedding 1986). In high velocity air flow, e.g. in wind tunnels, vegetable oil or water droplets produced with an aerosol generator in diameters between 1 and 10 μm are commonly used (Raffel et al. 1998; Spedding et al. 2003). For biological applications, the particles should not be toxic, or clog up the respiratory

structures of the organisms involved (see also Stamhuis et al. 2002). In all cases the particles should preferably be highly reflective, to yield good particle images on the PIV recordings.

Illumination

In PIV flow studies, illumination is usually provided by a laser light sheet. A narrow sheet of light is used to define the plane of investigation. It prevents blurry imaged particles just above or below the focussed plane, which would pollute the PIV images and thereby adversely affect the analysis result. A laser is used because it can deliver a bright sheet with almost constant thickness without aberration or diffusion, due to the coherent and monochromatic character of the emitted light. The slightly diverging light beam produced by a laser is usually transformed into a sheet by converging it with a weak positive lense and subsequently make the beam fan out in one plane to a sheet by an additional cylindrical lense (Figure 3). This results in a sheet with a slightly converging thickness, enabling the experimenter to select a certain sheet thickness tuned to the experimental conditions. Sometimes additional optics are added to further condition the light sheet to the application (Raffel et al. 1998). As a consequence of the sheet illumination, the main flow direction and the light sheet have to be aligned. This prevents the particles that are illuminated and pictured in the first image from disappearing from the illuminated plane and being replaced by other, differently arranged particles in the second image.

The types of lasers used for PIV illumination can be divided in two categories: continuously emitting lasers, usually called CW (= continuous weight) lasers, and pulsating lasers. The CW lasers, e.g. Helium–Neon and Argon lasers, produce relatively low power light (Pmax = 0.01 – 50 W) of good beam quality and are easy to set up and maintain. The light sheet is continuously visible to the experimenter and to the experimental animals. Pulsed lasers, e.g. Nd:YAG lasers, however, can produce much higher illumination levels per pulse with very short intervals between two pulses. They are, therefore, highly recommended in high speed applications where short illumination times and a high pulse frequency are mandatory. Pulsed lasers such as Nd:YAG lasers are usually more expensive and more difficult to set up due to the added timing and synchronisation equipment compared with CW lasers.

Imaging

The illuminated plane in the particle seeded flow is imaged by a camera with its optical axis perpendicular to the plane, in order to record the whole plane in focus. High quality lenses are generally necessary because the light levels are usually relatively low and large apertures have to be used. In normal 2D-PIV, a black-and-white (B/W) camera suffices.

Until about 10 years ago, photographic cameras were used to record the PIV images, so that evaluation and analysis could only be done after the film was developed. The analysis included scanning of all interrogation areas of the (double-exposed) image with a separate light source and an electronic camera, taking a considerable amount of time. Today, electronic cameras, with a CCD or CMOS pick-up device, are the main imaging media. They provide immediate feedback on image quality and

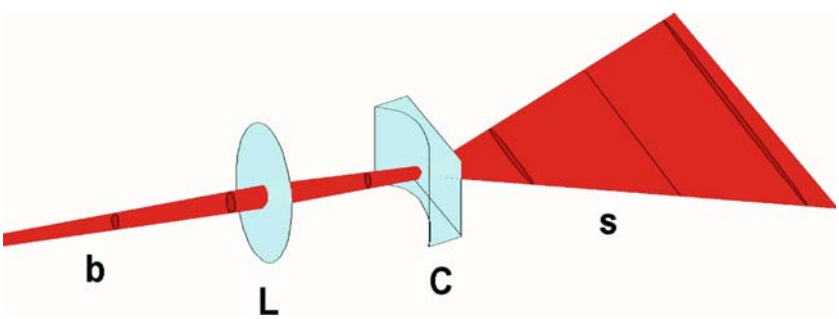

Figure 3. Basic optical arrangement to produce a light sheet from a laser beam: The mildly divergent laser beam coming from the left (b) is collimated by a weak positive lense (L) and subsequently fanned out in one plane only by a cylindrical lense (C). This results in a sheet (s) with a slightly converging thickness.

Table 1. Recommended camera system depending on the speed and the repeatability of the event studied, and the illumination intensity level.

Event speed	Illumination int.	Repeatability	Camera type
Fast	High	Good	Double-exposure
		Poor	HS-video
Slow	Low	Good	Progr. scan video (normal/MS)
		Poor	Progr. scan video (normal/MS)

enable instant analysis (Willert and Gharib 1991). Electronic cameras are available in a whole range of resolutions, sensitivities, shutter speeds, frame rates and image transfer protocols (analog or digital). Experimental characteristics such as event speed, repeatability and illumination intensity need to be evaluated for choosing a camera for a specific application, often in combination with the illumination system. Table 1 gives a summary of the most occurring experimental situations and the recommendable camera system. The choices are in general: progressive scan normal video rate camera, medium speed progressive scan camera, high speed video system, or double exposure camera.

A progressive scan normal video rate camera provides a stream of video images at a rate of 25 – 30 fps. The images are usually recorded in (dedicated) memory on a host computer which also controls the camera. Each pair of images out of this stream can be used for flow analysis. Due to the relatively long inter-frame period (0.03 – 0.04 s), this camera can only be used for low speed PIV applications. A medium speed progressive scan camera can record a stream of video images at a rate up to (ca.) 100 fps. Because such a camera is often an enhanced model of the normal progressive scan camera, recording and control are usually performed in the same way. Frame intervals of 0.01 s are possible, enabling higher event speeds than the normal camera. A high speed video system can usually record up to 1000 fps at full sensor resolution and much higher rates at reduced image sizes. This camera is suitable for fast PIV events, especially those with poor repeatability, which is often the case for animals. Any image pair of a series can be used for flow analysis, enabling mapping of unique flow velocity distributions in both space (2D) and time. High speed video in combination with a PIV analysis system is also known as 'time-resolved PIV'.

A double-exposure camera is capable of recording two frames with a very short inter-frame period. After recording, both frames are transferred to the host system and the camera is prepared for the next image pair. Inter frame times of below 500 ns are possible, but inter pair times can be quite substantial, in the order of 0.1 – 0.5 s (!), depending on the image transfer system. Consequently, double-exposure cameras are well suited for high speed repeatable flow events, or when only snap-shots of a flow are desired. Due to the architecture of the camera shuttering system, a double-exposure camera has to be used in combination with a pulsed laser.

Most of the camera types mentioned are available with sensor resolutions of 512×512 pixels or $1K \times 1K$, or higher. An image with a high resolution usually results in vector flow field with a high resolution or high accuracy. However, alternative analysis schemes, such as discrete window offsetting (Westerweel et al. 1997) or recursive local correlation (Hart 1999), can greatly enhance the results from a lower resolution recording system.

A low-cost alternative for a double-exposure camera + pulsed laser, combining the advantages of double-exposure, continuous video stream and a CW laser, may be to use an external fast shutter. The shutter should allow illumination of every first image just before the imaging device (e.g. CCD) is blocked and the second image immediately after the imaging device is opened again after frame transfer. Such a system needs external triggering tuned to and synchronised with the normal or medium speed progressive scan camera, making it somewhat more complicated. The shuttering may result in very low image illumination levels and therefore a powerful CW laser is essential.

PIV-analysis

The principle of PIV-analysis of a pair of images recorded successively from one illuminated plane of a particle seeded flow is depicted in Figure 2. The steps in the analysis process are: (1) selection of interrogation areas or sub-images for analysis, (2) correlation analysis, (3) finding the displacement peak, (4) calculating the velocity

vector, and (5) repeat for next set of sub-images (back to 1). These steps will be explained in detail below.

Selection of sub-images for analysis

To map local flow velocities, sub-images (interrogation areas) are selected containing a relatively small number of particles. In principle, an area containing only one single particle would be ideal, but it is hard to find that single same particle back in the second image, especially in a densely seeded flow as for PIV-recordings. The probability for a good analysis result is highest when the interrogation area contains about 8 – 10 particle images (Keane and Adrian 1992; Hinsch 1993). Such an area is selected from frame 1 (IA-1), and compared with an area at exact the same location in frame 2 (IA-2). Because the particles have moved in the inter-frame period and some of them may have left area 1, area 2 may be selected larger, e.g. twice the size, depending on the mathematical procedure used (see later).

Correlation analysis

To get the most probable displacement of the particle pattern in the interrogation area, a mathematical correlation procedure is applied. One can imagine this procedure as 'moving area 1 over area 2 until the best matching is found'. The expression 'best matching' is used because in practice there is never a 100% matching due to particles that have left or entered the imaged area in the second image compared with the first.

Mathematically, there are two methods used for image correlation analysis: Fourier transformations and convolution filtering. The results of both methods are comparable, but their applicability and input formats vary, causing a somewhat less accurate result from the Fourier path. Details on the mathematics of both methods can be found elsewhere (e.g. Raffel et al. 1998). Here only the essence of each method is covered briefly. Convolution filtering comes closest to the description of 'moving the images over one another for the best match'. A new correlation image of the same size as the interrogation area is built up by calculating a new value for each pixel from the sum of the products of all pixel values from interrogation area 1 and interrogation area 2 for that particular location (Figure 4). After normalisation, the new image is a visible representation of the 2D probability-density function of the level of matching between the 2 sub-images (Figure 5a).

The procedure for Fourier analysis of the image pairs is in short: each sub-image is transformed from the real to the complex domain using fast Fourier transforms (FFT). In the complex domain a conjugate multiplication between the transform results from both images takes place, and the product is transformed back to the real domain using an inverse FFT. This yields after normalisation about the same image representation of the 2D probability–density function of the level of matching between the 2 sub-images (Figure 5b).

The two methods differ in that in general the 'Fourier-pathway' is faster, especially when the FFT routines are optimized for this kind of analysis, compared with the 'convolution-pathway'. Fourier analysis needs interrogation area formats of a power of 2, e.g. 16×16, 32×32 and 64×64 pixels and both input images must have the same format when applying FFT's. While using convolution filtering, any interrogation area format is possible, enabling fine tuning of interrogation window size to e.g. number of particle images or average particle displacement. Also, the second sub-image can be chosen larger (see (Figures 2, 4 and 5a), to make sure that as many particles in the interrogation area (or IA-1) as possible are also found back in the sub-image (or IA-2), resulting in a more reliable probability–density distribution. This prevents the implementation of 'repair' routines such as window-offsetting in which a first FFT-cycle is used to get a first estimate of the particle displacement and a second cycle at the estimated displaced position is used to get an accuracy comparable to convolution filtering.

Finding the displacement peak

The high intensity peak in the PIV-analysis result image has to be located to know the particle displacement. Because the peak is usually larger than one pixel, all pixel information of the peak can be used to calculate the peak position, and obtain sub-pixel accuracy. Methods used so far are centroiding, centre of gravity B/W, centre of gravity weighed to grey value, Gaussian curve fitting and parabolic curve fitting. Centroiding, in which all the pixels of the peak outline are used to calculate it's centre, works well for larger peaks but is not

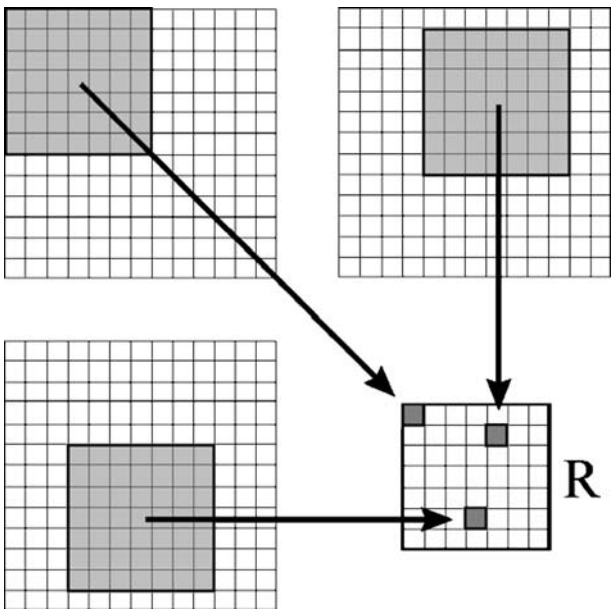

Figure 4. In convolution filtering, the resultant image (bottom right) has the same number of pixels as the interrogation area from the first image. Each new pixel is calculated from the sum of the products of all individual overlapping pixels from the interrogation area from image 1 (Area 1, light grey) and the sub-image from image 2 (Area 2, white). The sum of products of the top-left position results in a new value for pixel 1 in the result image (dark grey pixel). E.g. in formula for the first result pixel that is: $R(1,1) = A1(1,1) \times mA2(1,1) + A1(2,1) \times A2(2,1) + A(3,1) \times B(3,1) + \cdot + A(7,7) \times B(7,7)$. The new (virtual) image is normalised to the same grey value range as the original images and can then be displayed and/or used for further processing.

very accurate for small ones, and the same is true for calculation of the centre of gravity B/W (= after thresholding). Calculation of the centre of gravity weighed to grey value, in which lighter pixels contribute more than darker ones, gives good and accurate results for average size peaks (3 pixels diameter) and even better for larger peaks. Gaussian and parabolic curve fitting work well and accurate on average size peaks (3 pixels) but are less reliable on smaller or larger peaks. Depending on the peak size and the method used, accuracies of less than one tenth of a pixel are possible. With displacements of 3 pixels or more, a confidence level of much higher than 95% can thus be achieved.

Calculating velocity vectors

The displacement information in pixels is translated to real world units (metres or millimetres) after calibration. Subsequently, velocity vectors are calculated by dividing the calibrated displacements by the inter-frame or inter-pulse time. The resulting vectors can each be represented by $\boldsymbol{u}_i$ and $\boldsymbol{v}_i$, being the components in X- and Y-direction of the ith vector, respectively, or by a magnitude, and an angle with the horizontal with respect to a defined origin.

Repeat for next set of sub-images

The process of selection of an interrogation area, cross-correlation, peak-finding and calculation of the velocity vector is repeated over the whole image, resulting in a complete vector representation of the flow field (Figure 6). To obtain more velocity information out of the images, adjacent interrogation areas are often overlapping. Any overlap is in principle possible, but one should keep in mind that a high degree of overlap also means that the same information is sampled multiple times, resulting in over-sampling. An overlap of 50% is recommendable because of a good balance between information density and over-sampling (Hinsch 1993).

Post-processing

Post-processing of velocity vector arrays resulting from PIV analysis comprises two main components:

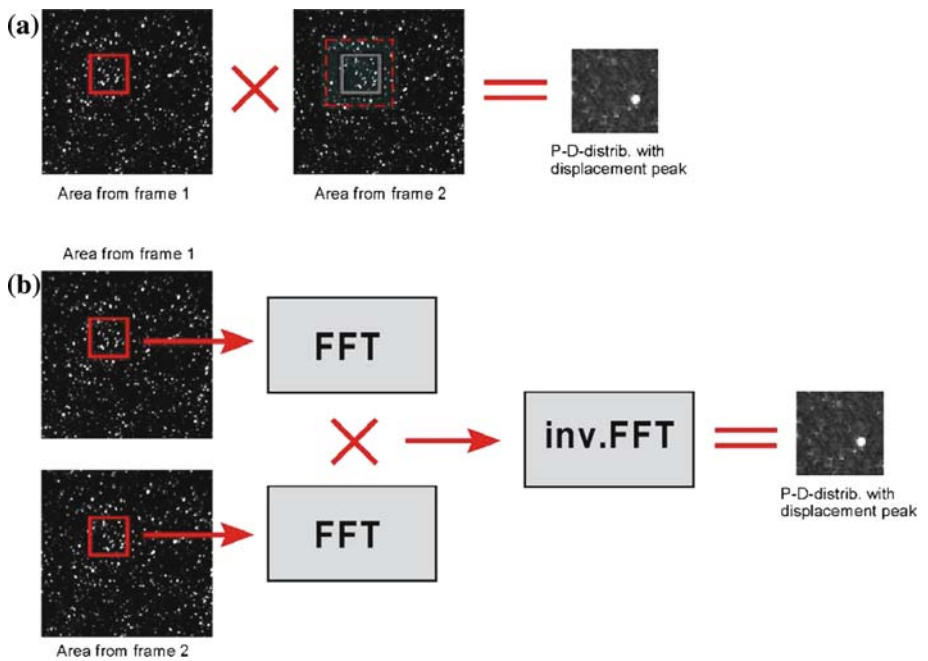

Figure 5 (a) Cross-correlation by convolution filtering in practice, being the filter product of 2 sub- images from 2 successively recorded images displayed as mathematical product, resulting in a probability density distribution with the displacement peak. (b) Cross-correlation by Fourier analysis in practice: the sub-images from the 2 successively recorded images are transformed to the complex domain using FFT's, the products are subjected to a complex conjugate multiplication of which the product is translated back to the real domain with an inverse FFT, eventually resulting in a probability density distribution with the displacement peak.

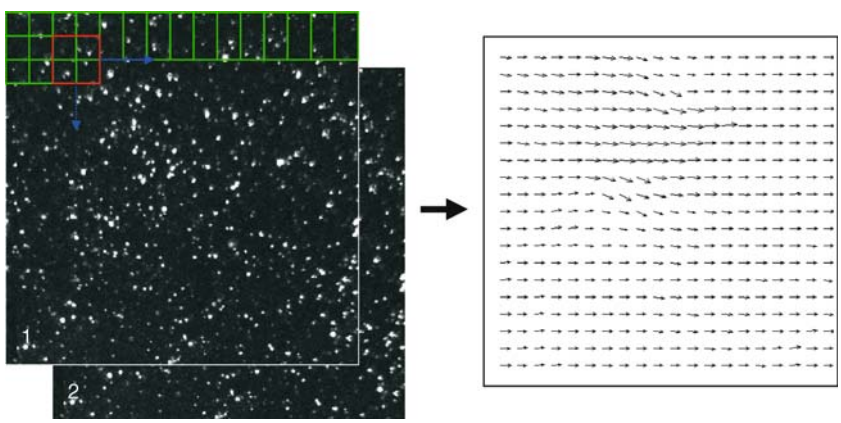

Figure 6. A set of two successive images is analysed by repetitive sampling and subsequent PIV analysis, often with overlap between neighbouring interrogation areas (in this example the overlap is 50% of the interrogation area size in X and Y direction) to get an accurate result without over- or under-sampling.

(1) data-validation and replacement of erroneous vectors, and (2) analysis and visualisation of velocity gradients. They are both essential to get reliable and well-understandable results that help to quantify, comprehend and interpret the flow phenomena studied, the ultimate goal of a PIV flow study.

Data validation and error correction

The resultant vector set of automated PIV analysis often includes a certain number of incorrect vectors that are usually obvious in the vector diagram. This type of erroneous vectors are usually caused by imperfections in the input images, caused

mainly by local variations in seeding density, local over-illumination due to an object or wall in the light sheet, strong out-of-plane-flow, local low illumination close to the image borders, or crippled interrogation areas next to the image border. These problems cause lack of correlation in the normal way, and a background (lower) peak is then recognized as if it were the displacement peak, resulting in an erroneous vector. Incorrect vectors usually stand out clearly with respect to surrounding vectors, due to their different orientation and often different – read 'much higher' – magnitude (Figure 7a). This enables relatively easy recognition to the observer as well as to any automated validation routine. Automated validation routines usually compare every vector with the surrounding one's in e.g. a 3×3 matrix, and use statistical parameters (e.g. the standard deviation or a histogram distribution function) or fluid dynamical parameters (e.g. vorticity or divergence) to discriminate between correct and erroneous vectors. The details of some of these validation procedures are discussed elsewhere (e.g. Raffel et al. 1998). Sometimes, especially close to the image border or when erroneous vectors show up in clusters, manual removal is necessary because automated validation does not work due to lack of reliable neighbours. This has to be done with great care, of course, and often needs an experienced evaluator.

Once the incorrect vectors are recognized and removed, empty places should preferably be filled by new vectors for continuous calculation of local derivatives and gradient parameters (see later). The new vector should represent the local flow velocity as close as possible. The most common way to do this is by 2D interpolation (Figure 7b). Although simple linear or polynomial interpolation yields reasonable results, the best and also principally the most reliable results have been obtained with 2D cubic natural spline interpolation (Spedding and Rignot 1993; Stamhuis and Videler 1995). Cubic natural spline fits leave the original neighbouring data unaffected and find the missing points according to 'the most fluent line or plane', i.e. with minimisation of the local curvature (second derivative). In non-compressible or subsonic compressible fluids, gradients are hardly ever sudden, which illustrates why spline fits give the best results. But even in shock waves in transonic flow, spline interpolation may perform very well.

Analysis and visualisation of velocity gradients

The vector flow field resulting from the PIV analysis and validation + interpolation procedure gives an enormous amount of information on local velocity and direction of the flow. Sometimes it is difficult to see the actual flow phenomena in the vector diagrams. Rotationality, for example, may be masked by the overall flow speed, and so may be other gradients. A first tool to reveal hidden flow structures is to subtract the average or background flow from all vectors. This is often done when studying the flow phenomena caused by a still or moving object, e.g. a wing section, in a flow tank or wind tunnel (Figure 7c vs. b). Additionally, the velocity magnitude (absolute or relative) shown as a colour code superimposed on the vector diagram is also quite revealing (Figure 7d).

To visualise other gradients in the flow, a number of parameters can be derived, of which vorticity, rate of shear, rate of strain, divergence and a vortex locator are most common and most clarifying. These parameters are calculated from the local velocity derivatives (see Table 2) (Vollmers et al. 1983; Lighthill 1986; Stamhuis and Videler 1995; Vollmers 2001). Two of them are illustrated in Figure 7e and f, showing the same jet pulse with accompanying circular flow areas (without background flow) as in Figure 7b, but now with visualisation of the vorticity (Figure 7e) and the vortex cores (Figure 7f) through colour coding. The rate of shear often helps to assess the velocity profiles and pressure/force distributions close to objects or walls. The rate of strain indicates acceleration or decelerations of the flow. High strain rates in a 2D analysis are often indicative for a significant 3D flow component, which can be checked with the divergence. Because they are computed from the velocity derivatives, these parameters are insensitive to masking effects of a high average flow speed or background flow. This makes them very useful in revealing the real but sometimes visually hidden flow morphology, e.g. as in turbulence in high speed flows or impinging jets.

The distribution of vorticity and other derived parameters give insight in the relative importance and the distributions of velocity gradients in the flow field. That is above all necessary to assess and understand the morphology of the flow phenomena studied. Moreover, velocity vectors and derived parameters can be used for further

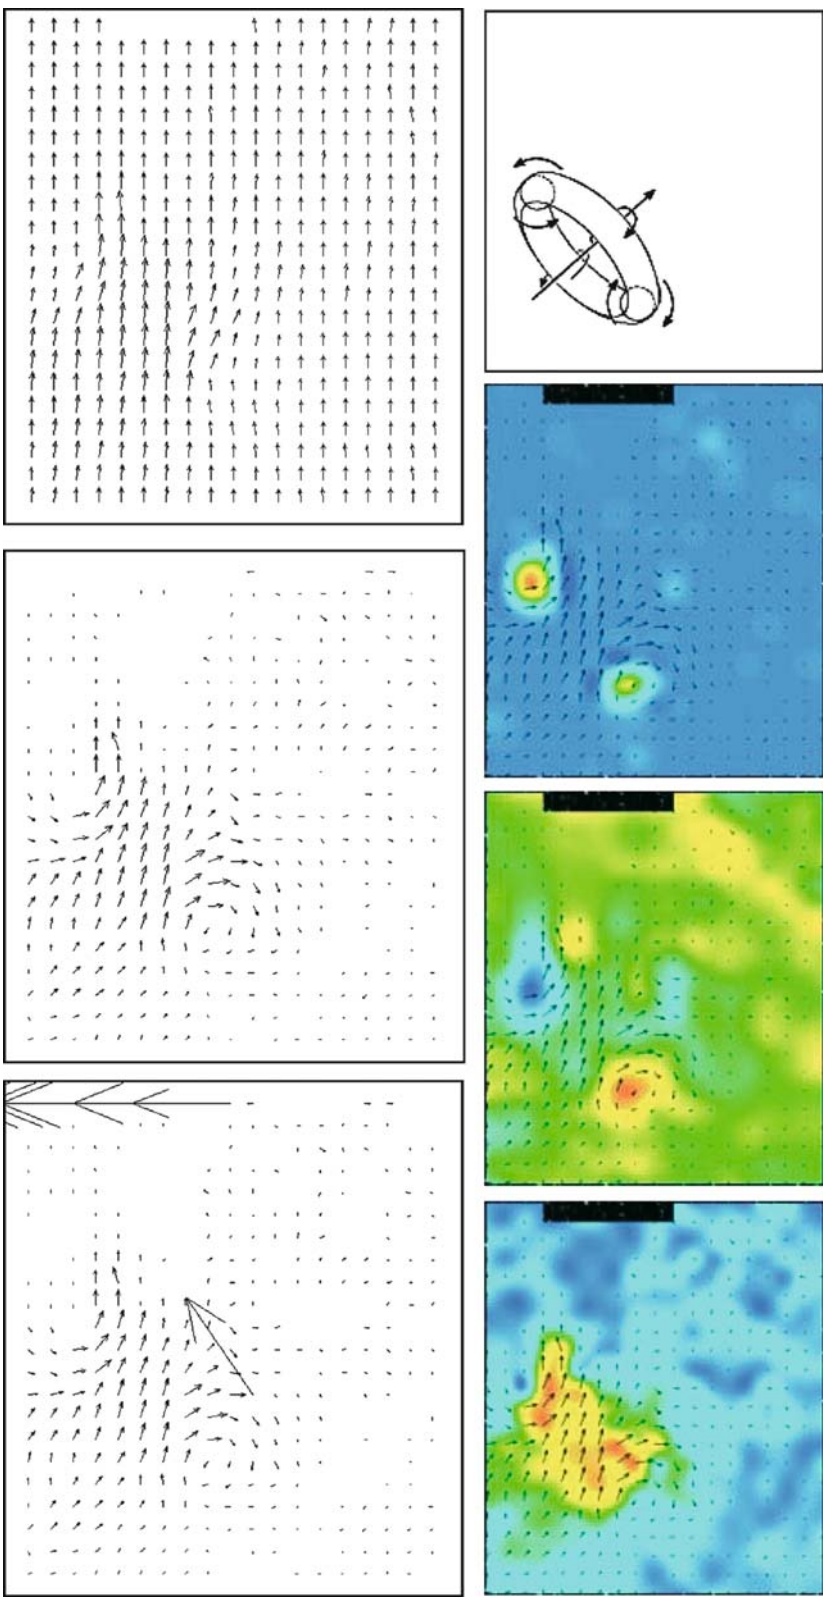

Figure 7. Post-processing PIV-results from a pulse-jet. *top row*: Vector diagram before (left, a) and after (middle, b) validation and interpolation, and when masked by a background flow (right, c). *lower row*: Colour coded vector diagrams showing the velocity magnitude (left, d), the vorticity (middle, e) and the vortex locator (right, f), and a 3D interpretation of the results from this analysis (far-right, g).

Table 2. Gradient parameters that can be derived from the flow velocity information resulting from 2D-PIV, and aid in quantification and interpretation of the flow phenomena studied.

Parameter	Stands for:	Equation
Vorticity	Rotationality of the fluid	$\omega = \partial v/\partial x - \partial u/\partial y$
Shear rate	Strength of velocity gradient perpendicular to the local velocity ('sliding of adjacent water layers')	$h = \partial v/\partial x + \partial u/\partial y$
Strain rate	Strength of velocity gradient in the direction of the local velocity ('acceleration within one water layer')	$e = \partial u/\partial x - \partial v/\partial y$
Divergence	Strength of out-of-plane flow rate (water leaving or entering the illuminated plane)	$\Theta = \partial u/\partial x - \partial v/\partial y$
Vortex locator	Finds centre of vortex (e.g. through discriminant for complex eigenvalues)	

processing or comparison, e.g. between successive analyses, between comparable setups, or to validate theoretical predictions.

Interpretation

To fully understand PIV results, proper interpretation is essential. A first challenge in interpreting 2D PIV results (and of other types of PIV) is that we are basically looking at one 2D plane or one cross-section of a 3D flow phenomenon. On top of that, the flow pattern may be changing in time, which adds an extra dimension to our problem. For example, the mass of water travelling from the upper-left corner towards lower-right with its neighbouring elevated vorticity areas as displayed in Figure 7a–f is in fact a cross-section of a jet pulse with accompanying vortex ring (Figure 7g). The exact shape and morphology of the jet + ring combination can only be found when using a real 3D visualisation technique or when repeating the experiment and making a series of parallel and perpendicular cross sections. That may not always be possible, especially when working with biota.

For a good interpretation a good background knowledge of the fluid dynamics involved, and experience are important. The first of these two can be found in fluid dynamics text books and in the literature. Experience can only be built up by directly working with the technique. When studying 3D free flow structures, flows close to objects or boundaries, sudden gradients, etc., one should be aware of the three-dimensionality of the flow and know the limitations of 2D representations.

Examples from Biology

The application of 2D-DPIV to animal generated flows is illustrated with two examples, one on acorn barnacle filter feeding and one on filter feeding in four different bivalve species. All these animals are benthic and were chosen because of the emphasis of the BIOFLOW network on benthic flow phenomena. These experiments serve as examples here, and the results will be discussed only briefly because that is not in the scope of this paper.

The normal habitat of benthic animals is often dominated by the external flow regime (waves and tides), but these animals may also generate their own flow phenomena. Here we concentrate on the flow patterns generated by the animals themselves, with some detail on the PIV setup, recording and analysis (for details see Stamhuis et al. 2002; Stamhuis and Vos in prep.). All PIV analyses were performed using SWIFT 4.0 (Dutch Vision Systems), a commercially available package derived from the system originally developed in our laboratory, using convolution filtering for the cross-correlation analysis. Illumination was provided by a red CW Krypton laser at a wavelength of 647 nm (Pmax = 1 W), delivered as a sheet at the recordable area through an armed glass fibre and a light sheet probe. The probe allows adjustment of the sheet thickness in the range 0.1–5 mm.

Acorn barnacle filter feeding

Acorn barnacles from the North Sea area are relatively small crustaceans living in protective calcareous tests, that colonize hard substrates such as rocks, dikes, breakwaters and bivalve shells. They have relatively long leg-like structures called cirri that can be extended and swept through the water to filter it for food, such as detritus particles and small planktonic organisms. Due to the small size of the cirri (lengths around 6 mm, diameter < 0.2 mm) and their low beat frequencies (around 1 beat/s), they are confronted

with viscous as well as inertial flow aspects. The barnacles used in the experiments (*Balanus crenatus*) are only actively beating at very low ambient flow velocities, and were studied while feeding in still water.

To map the flow around and caused by the beating cirri, single animals that had settled on mussel shells were mounted in a small aquarium with sea water, free from any wall or the surface within 5 cm. To reduce reflections and diffusion by the white calcareous tests, the animals were covered with a small black cloth cloak, leaving only the test orifice open for protrusion of the cirri. The flow was recorded in a series of horizontal and vertical sheets to get a good 3D impression of the flow morphology. The water was seeded with 5 μm hollow glass spheres (Nortek A.S.) which were neutrally buoyant in sea water. The particle seeded flow was illuminated with a light sheet of 0.4 mm thickness produced by the krypton laser with sheet probe. Images were recorded using an Adimec MX12 camera equipped with a Nikon macro-objective at 25 frames/s and 512×512 pixels resolution. The images were digitized and stored uncompressed in real time. PIV analysis and evaluation was performed after the recording was finished. Figure 8 shows a diagrammatic representation of the setup.

The results from kinematic and PIV analyses showed that the cirri accelerate fast during the first part of the power stroke and that water moves through the cirri at that time at a cirral Reynolds number of about 3 (based on cirrus diameter). During the 2nd half of the power stroke, the water catches up with the cirri, the Reynolds number drops to 0.3, and the cirral filter starts to function as a paddle. During the recovery stroke, the cirri are folded back towards the test opening and rubbed along the mouthparts to collect and ingest retained particles, and then the cirri flick back to the initial starting position without much interference with the water. Figure 9a shows the flow pattern of an active barnacle in side view halfway the power stroke (to the left in the figure). The volume of water accelerated and filtered during the previous beat cycle is still visible as a big yellow–red blob in the left part of the image. The results indicate that the barnacle actively filters the water for food particles during the first half of the power stroke. Because the cirral function changes from a filter to a paddle approximately halfway the power stroke, the amount of water that has just been filtered is swept away from the barnacle, probably to make sure that it does not filter the same water again. This seems to guarantee a good and continuous food supply to the animal.

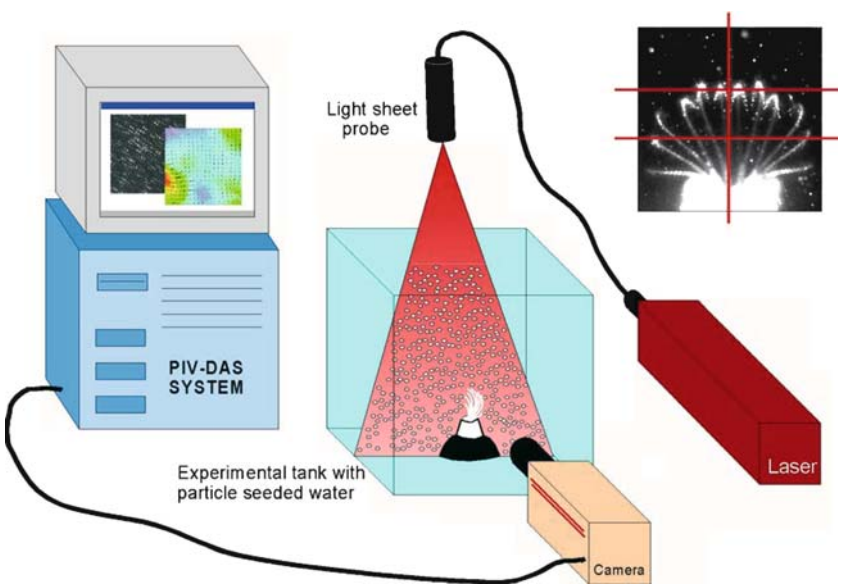

Figure 8. PIV setup for mapping the flow patterns produced by the beating 6 mm long cirri of the acorn barnacle *Balanus crenatus*. The picture in the upper-right corner shows the horizontal and vertical planes that have been illuminated and analysed.

Bivalve filter feeding strategies

On the coastal intertidal flats of the Dutch Wadden Sea, several species of bivalves can be found buried in the mud. Some species, e.g. the cockle (*Cerastoderma edule*) and the soft-shelled clam (*Mya arenaria*) are found close to the gullies with relatively fast tidal currents, other species, such as the blue mussel (*Mytilus edulis*) and the Baltic tellin (*Macoma balthica*) are found higher up the flats in areas with slower currents and higher detritus sedimentation rates. All four species, like all bivalves, use their internal gills to filter the water for food particles, letting the water in and out through inhalant and exhalant siphons. From these four species, the Baltic tellin is regarded as a facultative suspension/deposit feeder, able to inhale particles suspended in the overlaying water as well as deposit from the sea floor, with its long and flexible siphons (Kamermans 1994; De Goey and Luttikhuizen 1998) although it recently has been considered to be an exclusive deposit feeder with consequent low pumping rates (Riisgaard and Kamermans 2001).

The flow velocity and flow morphology just outside of the inhalant and exhalant siphons were studied with PIV in a setup similar to Figure 8. Animals of comparable size (1.5–2.0 cm) were placed upright on a rubber ring on the bottom of the aquarium and illuminated with a vertical laser sheet of 0.5 mm thickness. The water was seeded with 25 μm Pliolyte (BASF) particles and images were recorded and stored digitally in real time at 25 fps with a 512×512 resolution using an Adimec MX12 camera with a Nikon macro-lense. Suspension feeders sometimes reduce their pumping rate in response to high particle loading. All animals were offered the same conditions to make the results of this experiment comparable with one another. The pumping rates may, however, in general be lower than in natural situations. The seeding particles did not seem to hamper the experimental animals except for occasional backwash of the filter (Stamhuis et al. 2002).

The results for the comparable size specimens of all four species are summarized in Table 3. The Baltic tellin showed very low feeding current velocities and pumping rates. Both the cockle and the clam displayed diffuse inflow patterns and jet-like outflow patterns and processed relatively large volumes of water. Their filtration efficiency appeared to be so high that the middle part of the outflow jet showed up as a dark band, because most particles had been filtered out. The outflow patterns were therefore partially reconstructed

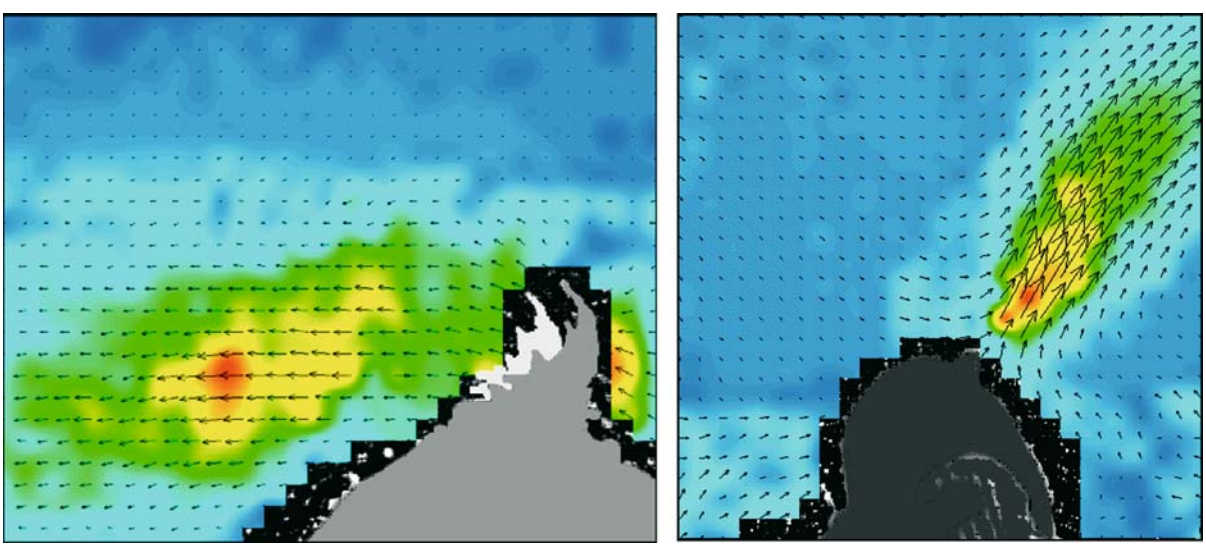

Figure 9. Samples of the PIV-results from the biological examples. Left panel (a): The vector flow field with colour coded velocity magnitude from a lateral view (with a vertical light sheet) of the acorn barnacle *Balanus crenatus*, showing the moving 6 mm long cirri, the water accelerated by the cirri, and the water volume on the left that was accelerated and shed during the previous beat cycle; imaged area = 24×30 mm; cirri beat towards left. Right panel (b): The diffuse outflow pattern of the exhalant siphon of an actively filter feeding mussel *Mytilus edulis*; imaged area = ca. 30×0 mm. colour coding (a + b): red = rel. high velocity, blue = rel. low velocity.

using spline interpolation (for more details see Stamhuis et al. 2002). The mussel was intermediate in both pumping rate and outflow pattern (Figure 9b). Thus, the two species living close to the gullies had a diffuse inflow pattern, a jet-like outflow and high filtered volumes. The two species living higher up the flats had lower flow rates. Cockle and clam apparently optimized for high filtering rates to compensate for low particle food loading of the dilute water of the gullies, whereas the mussel lives and feeds higher up the flats with a higher detrital loading, and the tellin 'grazes' from an almost slurry-like food source very high up the flats and therefore does not need high pumping rates. This may support the conclusion that the Baltic tellin is mainly a deposit feeder. In conclusion, filtering rate seems inversely related to particle concentration in the water in their natural habitat.

Alternative PIV techniques

In a number of cases, e.g. with animal induced flow phenomena that cannot be reproduced, or with a highly variable 3D flow component, other flow analysis techniques may be necessary. Stereo-PIV, Holographic PIV (H-PIV) or modern 3D-PTV variants are potential candidates when it comes to mapping 3D flow phenomena. Plain PTV may be valuable when very low seeding densities are required, and Micro-PIV may assist in resolving flow phenomena at the very small scale.

Stereo-PIV is often also called 3C-PIV, the 3C standing for 'three components' or 'three velocity components'. This emphasizes the main differences between 2D and Stereo PIV: Stereo-PIV enables one to resolve the third velocity component in the illuminated plane. In Stereo-PIV a second camera is added and both cameras and their imaging optics are arranged in such a way that the whole area of interest within the illuminated plane is projected sharply onto the imagers. Additionally, the laser sheet is made somewhat thicker, to prevent the particles that travel in the out-of-plane 3rd dimension from leaving the illuminated area before the next picture is recorded. Finally, after calibration to correct for image distortion, two pairs of images recorded in the stereo setup yield the distribution of flow in X, Y and Z direction ($\boldsymbol{u}$, $\boldsymbol{v}$ and $\boldsymbol{w}$, respectively) in one 2D plane. The big step forward in comparison with 2D-PIV is thus resolution of the out of plane flow. This information is, however, available for one plane at a time only.

In Holographic PIV, a certain volume in the 3 dimensional space is illuminated as uniformly as possible and imaged together with a reference beam at two successive moments onto a hologram (Hinsch 1993). This hologram holds in principle all information of particle displacements in all three dimensions. The way to retrieve this information from the hologram is to read virtual planes from the hologram with a CCD camera, and perform regular 2D PIV on the resulting image. By combining parallel and perpendicular planes afterwards, the whole volume can be mapped (Hermann et al. 2000). The advantage and great power of HPIV is that all 3D fluid movements are recorded instantaneously with a high resolution and are available for later evaluation. A disadvantage, at least from the experimenter's point of view, is that setting up an HPIV experiment and performing HPIV analyses needs a lot of experience as well as expensive equipment such as special optics and usually a powerful laser. High illumination levels or high energy laser pulses sometimes invoke avoidance reactions or may even be harmful to experimental animals. Up to now only a few laboratories have specialised in HPIV, mainly for fundamental fluid dynamics research. Recently, the 3D flow field of a swimming copepod has been mapped using HPIV (Malkiel et al. 2003).

Table 3. Summary of the results from a comparative PIV study at 4 species of bivalve to quantify their filtration activity with respect to their normal habitat on the Dutch coastal flats.

Species	Flow speed (mm/s)	Flow rate(ml/h)	Pattern	Habitat
Balthic tellin (*M.balthica*)	2.2	6.2	diffuse	flats
Cockle (*C.edule*)	11	76	jet	gullies
Clam (*M. arenaria*)	13	330	jet	gullies
Mussel (*M.edulis*)	ca. 5	ca. 50	intermediate	intermediate

Three-dimensional particle tracking velocimetry (3D-PTV) is a technique that has already been used for quite some time and has also been applied to complex biological flows (e.g. Spedding 1986). In early 3D-PTV studies, the experimenter had to manually trace particles in e.g. 2 stereo images from the same scene, allowing only low seeding densities and involving labourious analysis. The resulting 3D vector diagrams showed a resolution high enough to get a good impression of the flow phenomena and local flow velocities, but too low to derive additional descriptive and analytical parameters. In the recently developed 3DPTV techniques, often more than 2 views from the same scene are taken, the seeding density of the fluid has been increased to get higher spatial resolutions and the analysis has been automated. Three views from the same scene yield 3 sets of stereo images instead of 1 set, and, when placing the cameras in a fixed frame of reference, triple triangulation can be performed to locate each particle with a high accuracy . Newly developed image analysis algorithms, which are mainly based on statistical procedures and forecasting (Cowen and Monismith 1997; Van der Plas and Bastiaans 1998), may manage to track 1500 particles or more in one volume. The resulting vector diagrams show a reasonable to good spatial resolution, high enough for derivation of e.g. vorticity. 3D-PTV may see a revival with these new recording and analysis methods, and the technique seems to be very suitable for biogenic flows, providing that the recording system has the proper speed and resolution.

When the seeding density is not high enough for automated (3D-) PTV or PIV analysis due to limitations on the animals side, one can revert to plain PTV as it has already been used for decades (e.g. Strathman 1973; Joergensen 1975; Sponaugle 1991; Larsen and Riisgård 2002). Here, the particles are usually located visually and manually, and the result is a flow field with irregular spaced vectors at a relatively low density. The velocity vectors may be very informative, but the low vector density makes it difficult to derive gradient parameters (shear, vorticity) without a more complex mathematical procedure. The accuracy depends highly on the analysis procedure used, which may range from 'drawing on a foil on the monitor screen' to 'centroiding or peak curve fitting at particle images in HR digital pictures'.

At the very small scale (ca. 1 mm), the limited focal depth of the camera- microscope-combination may be used to define the plane of interest in PTV or PIV studies (e.g. Klank et al. 2001). Such a system is usually classified as a Micro-PIV system. However, a very thin lightsheet (e.g. 0.1 mm thickness) may still help to prevent blurred particle images just above or below the plane of interest (e.g. Van Duren et al. 1998), that increase the noise and thereby lower the accuracy. Micro-PIV (or PTV) may undergo some improvements in optical as well as analytical aspects which increase the accuracy and make this technique more applicable for biological micro-applications.

One of the latest developments in pursuit of higher resolutions and higher accuracy in PIV analyses is called Super Resolution PIV (SR-PIV). In SR-PIV, two routes have been followed to achieve higher resolution: a PIV route and a PTV route. The PIV route can be summarized as first performing normal PIV analysis with a normal interrogation window size, and subsequently taking smaller and smaller segments of the interrogation window using the displacement information from former steps, to increase accuracy and resolution, almost down to the single particle image size (Hart 1999). The other route combines PTV algorithms, sometimes similar to those used in the latest 3D-PTV methods, with classic PIV. A PIV step provides a first estimation of the flow velocity, followed by PTV which provides information on gradients within the interrogation area using the velocity information from the PIV step to improve particle pairing probabilities (Keane et al. 1995; Cowen and Monismith 1997; Bastiaans et al. 2002). Both routes yield comparable results, the vector density may increase more than 10-fold, but the PIV route with recursive correlation seems capable of even more spectacular increases (Hart 1999).

Conclusion

Two-dimensional digital particle image velocimetry appears to be a very useful and versatile technique that can be applied to biogenic or biologically relevant flows relatively easily. Even irreproducible and very fast flow events can be tackled, when using the right illumination and recording techniques. Application depends on tolerance to the relatively high level of (laser)

illumination, and to the presence of densely seeded particles in the flow, that may hamper the organism or prevent it from behaving naturally. PIV flow analysis can deliver 2D vector flow fields in an instant, and a number of derivable parameters assist in understanding and describing the flow phenomena studied in a qualitative as well as a quantitative way. Interpretation of the 3D flow patterns from the 2D PIV cross sections is, however, not trivial and has to be done with care. When doing so, 2D DPIV can greatly assist in understanding how an organism interacts with a fluid and what it actually does with that fluid, or e.g. how it manages to stand up to external currents.

Acknowledgements

Jan Drent assisted in recording of the bivalve flow fields, Michiel Vos recorded the barnacle flow fields, Sandra Nauwelaerts recorded the vortex ring flow fields (cooperation with the University of Antwerp), Luca van Duren and Per Jonson are acknowledged for the invitation to the BIO-FLOW-2002 symposium. The comments of Dr H.U. Riisgaard and 2 anonymous referees enhanced the paper greatly.

References

Adrian R.J. 1991. Particle imaging techniques for experimental fluid mechanics. Ann. Rev. Fluid Mech. 23: 261–304.

Anderson J.D. 2000. Introduction to Flight (4th int. ed.). McGraw-Hill, Boston.

Bastiaans R.J.M., der Plas G.A.J. and Kieft R.N. 2002. The performance of a new PTV algorithm in super-resolution PIV. Exp. In Fluids 32: 346–356.

Cowen E.A. and Monismith S.G. 1997. A hybrid digital particle tracking velocimetry technique. Exp. In Fluids 22: 199–211.

Drucker E.G. and Lauder G.V. 1999. Locomotor forces on a swimming fish: three-dimensional wake dynamics quantified using digital particle image velocimetry. J. Exp. Biol. 202: 2393–2412.

Goey P. de and Luttikhuizen P. 1998. Deep-burying reduces growth in intertidal bivalves: field and mesocosm experiments with *Macoma balthica*. J. Exp. Mar. Biol. Ecol. 228: 327–337.

Hart D.P. 1999. Super-resolution PIV by recursive local-correlation. J. Visualisation 10: 1–10.

Hermann S., Hinrichs H., Hinsch K. and Surmann C. 2000. Coherence concepts in holographic particle image velocimetry. Exp. In Fluids 29(Suppl.): S108–S116.

Hinsch K.D. 1993. Particle image velocimetry. In: Sirohi R.S. (ed.), Speckle Metrology. Marcel Dekker Inc, New York, pp. 235–342.

Joergensen C.B. 1975. On gill function in the mussel *Mytilus edulis* L. Ophelia 13: 187–232.

Jones J.D.C. (ed.), 1997. Particle image velocimetry. Meas. Sci. Technol. 8(12): 1379–1583.

Kamermans P. 1994. Similarity in food source and timing of feeding in deposit- and suspension-feeding bivalves. Mar. Ecol. Prog. Ser. 104: 63–75.

Keane R.D. and Adrian R.J. 1992. Theory of cross-correlation analysis of PIV images. Appl. Sci. Res. 49: 191–215.

Keane R.D., Adrian R.J. and Zhang Y. 1995. Super-resolution particle image velocimetry. Meas. Sci. Technol. 6: 754–768.

Klank H.G., Goranovic G., Kutter J.P., Gjelstrup H., Michelsen J. and Westergaard C.H. 2001. Micro PIV measurements in micro cell sorters and mixing structures with three dimensional flow behavior. CD-ROM proceedings 4th Int. Symp. on PIV, Göttingen, Paper, Paper 1161, pp. 1–12

Larsen P.S. and Riisgård H.U. 2002. On ciliary sieving in pumping bryozoans. J. Sea Res. 48: 181–195.

Lighthill J. 1986. An Informal Introduction to Theoretical Fluid Mechanics. Oxford University Press, Oxford.

Malkiel E., Sheng J., Katz J. and Strickler J.R. 2003. The three-dimensional flow field generated by a feeding calanoid copepod measured using digital holography. J. Exp. Biol. 206: 3657–3666.

Müller U.K., den Heuvel B.L.E., Stamhuis E.J. and Videler J.J. 1997. Fish foot 23 prints: morphology and energetics of the wake behind a continuously swimming mullet (*Chelon labrosus*- Risso). J. Exp. Biol. 200: 2893–2906.

Raffel M., Willert C. and Kompenhans J. 1998. Particle Image Velocimetry, A Practical Guide. Springer-Verlag, Berlin.

Riisgaard H.U. and Kamermans P. 2001. Switching between deposit and suspension feeding in coastal zoobenthos. In: Reise K. (ed.), Ecological Comparisons of Sedimentary Shores. Ecological Studies Vol. 151. Springer Verlag, Berlin, pp. 73–101.

Spedding G.R. 1986. wake of a Jackdaw (*Corvus monedula*) in slow flight. J. Exp. Biol. 125: 287–307.

Spedding G.R. and Rignot E.J.M. 1993. Performance analysis and application of grid interpolation techniques for fluid flows. Exp. In Fluids 15: 417–430.

Spedding G.R., Hedenström A. and Rosén M. 2003. Quantitative studies of the wakes of freely flying birds in a low-turbulence wind tunnel. Exp. In Fluids (in press).

Sponaugle S. 1991. Flow patterns and velocities around a suspension-feeding gorgonian polyp: evidence from physical models. J. Exp. Mar. Biol. Ecol. 148: 135–145.

Stamhuis E.J. and Videler J.J. 1995. Quantitative flow analysis around aquatic animals using laser sheet particle image velocimetry. J. Exp. Biol. 198: 283–294.

Stamhuis E.J., Videler J.J., Van Duren L.A. and Müller U.K. 2002. Applying digital particle image velocimetry to animal-generated flows: Traps, hurdles and cures in mapping steady and unsteady flows in Re regimes between 10^{-2} and 10^{5}. Exp. In Fluids 33: 801–813.

Stanislas M., Kompenhans J. and Westerweel J. (eds) 2000. Particle Image Velocimetry, Progress Towards Industrial Application. Kluwer Academic Publishers, London.

Strathmann R. 1973. Function of lateral cilia in suspension feeding of lophophorates (Brachiopoda, Phoronida, Ectoprocta). Mar. Biol. 23: 129–136.

der Plas G.J. and Bastiaans R.J.M. 1998. Accuracy and resolution of a fast PTV- algorithm suitable for HiRes-PV. In: Grant I. and Carlomagno G.M. (eds), Proceedings of the 8th International Symposium on Flow Visualisation. Sorrento, Italy, pp. 1/1–1/12.

Van Duren L.A., Stamhuis E.J. and Videler J.J. 1998. Reading the copepods personal ads: increasing encounter probability with hydromechanical signals. Phil. Trans. R. Soc. Lond. B 353: 691–700.

Vogel S. and LaBarbera M. 1978. Simple flow tanks for research and teaching. Bioscience 28: 638–643.

Vogel S. 1994. Life in Moving Fluids, the Physical Biology of Flow (2nd ed.). Princeton Press, New Jersey.

Vollmers H., Kreplin H.-P. and Meier H.U. 1983. Separation and vortical-type flow around a prolate spheroid – evaluation of relevant parameters. NATO-AGARD Conference Proceedings No. 342, pp. 14/1–14/14.

Vollmers H. 2001. Detection of vortices and quantitative evaluation of their main parameters from experimental velocity data. Meas. Sci. Technol. 12: 1199–1207.

Westerweel J. 1993. Digital Particle Image Velocimetry – Theory and Application. Delft University Press, Delft.

Westerweel J. 1997. Fundamentals of digital PIV. Meas. Sci. Technol. 8(12): 1379–1392.

Westerweel J., Dabiri D. and Gharib M. 1997. The effect of a discrete window offset on the accuracy of cross-correlation analysis of PIV recordings. Exp. In Fluids 23: 20–28.

Willert C.E. and Gharib M. 1991. Digital particle image velocimetry. Exp. In Fluids 10: 181–193.

Zirbel M.J., Veron F. and Latz M.I. 2000. The reversible effect of flow on the morphology of *Ceratocorys horrida* (Peridiniales, Dinophyta). J. Phycol. 36: 46–58.