Measurement Science and Technology

PAPER

Volumetric particle image velocimetry with a single plenoptic camera

To cite this article: Timothy W Fahringer *et al* 2015 *Meas. Sci. Technol.* **26** 115201

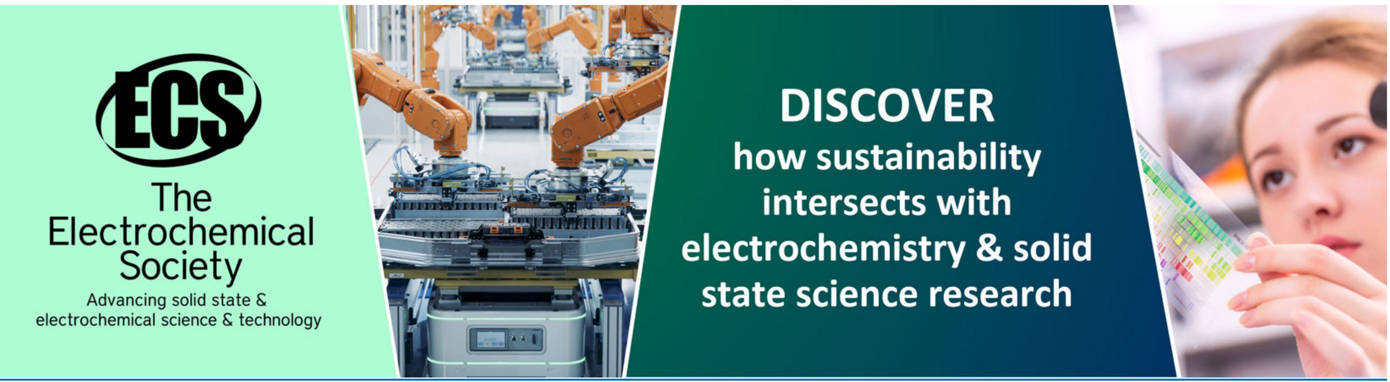

Meas. Sci. Technol. **26** (2015) 115201 (25pp) doi:10.1088/0957-0233/26/11/115201

Volumetric particle image velocimetry with a single plenoptic camera

Timothy W Fahringer[1], Kyle P Lynch[2] and Brian S Thurow[1]

[1] Auburn University, Auburn, AL 36849, USA
[2] TU Delft, Mekelweg 2, 2628 CD Delft, The Netherlands

E-mail: twf0001@auburn.edu

Received 13 April 2015, revised 25 June 2015
Accepted for publication 27 July 2015
Published 29 September 2015

Abstract

A novel three-dimensional (3D), three-component (3C) particle image velocimetry (PIV) technique based on volume illumination and light field imaging with a single plenoptic camera is described. A plenoptic camera uses a densely packed microlens array mounted near a high resolution image sensor to sample the spatial and angular distribution of light collected by the camera. The multiplicative algebraic reconstruction technique (MART) computed tomography algorithm is used to reconstruct a volumetric intensity field from individual snapshots and a cross-correlation algorithm is used to estimate the velocity field from a pair of reconstructed particle volumes. This work provides an introduction to the basic concepts of light field imaging with a plenoptic camera and describes the unique implementation of MART in the context of plenoptic image data for 3D/3C PIV measurements. Simulations of a plenoptic camera using geometric optics are used to generate synthetic plenoptic particle images, which are subsequently used to estimate the quality of particle volume reconstructions at various particle number densities. 3D reconstructions using this method produce reconstructed particles that are elongated by a factor of approximately 4 along the optical axis of the camera. A simulated 3D Gaussian vortex is used to test the capability of single camera plenoptic PIV to produce a 3D/3C vector field, where it was found that lateral displacements could be measured to approximately 0.2 voxel accuracy in the lateral direction and 1 voxel in the depth direction over a $300 \times 200 \times 200$ voxel volume. The feasibility of the technique is demonstrated experimentally using a home-built plenoptic camera based on a 16-megapixel interline CCD camera and a 289×193 array of microlenses and a pulsed Nd:YAG laser. 3D/3C measurements were performed in the wake of a low Reynolds number circular cylinder and compared with measurements made using a conventional 2D/2C PIV system. Overall, single camera plenoptic PIV is shown to be a viable 3D/3C velocimetry technique.

Keywords: 3D PIV, light field imaging, plenoptic camera, tomography, volume reconstruction

(Some figures may appear in colour only in the online journal)

1. Introduction

Particle image velocimetry (PIV) is a well-established measurement technique used extensively to resolve planar velocity fields in a variety of flow environments. Due to PIV being an image-based technique, the measurements obtained have been traditionally limited to two dimensions. Consequently, traditional PIV is not capable of capturing the full three-dimensional (3D), three-component (3C) velocity field instantaneously, which is important for quantifying the topology and extent of flow structures which pervade most turbulent flows. Moreover, turbulence is inherently 3D in nature, and a full description requires a measurement of the 3D velocity field and derivative quantities such as the stress tensor and vorticity vector.

1.1. Current 3D PIV techniques

These limitations have led to a number of efforts to develop 3D, 3C PIV-based measurement techniques. Advances such as stereoscopic PIV (Arroyo and Greated [1], Willert [2]) extend traditional PIV to allow 3C measurements within a 2D plane, and dual plane stereoscopic PIV (Kähler and Kompenhans [3]) applies this technique to two planes which allows the derivative quantities of each dimension and each component to be calculated. Since these techniques only acquire 3C data within a single plane or two planes, these techniques are more appropriately labeled 2D/3C. An extension of the aforementioned techniques is scanning PIV (Brücker [4]), where high-repetition-rate lasers and camera systems are used to illuminate and capture images at multiple planes throughout the measurement volume. The advantage of these systems are the intuitive setup and data processing steps; however, even with kHz-rate lasers the volume scanning time is often large compared to the characteristic timescales of the flow under consideration and prevents the technique from being applied to most practical flows. The use of MHz-rate laser systems (Lynch and Thurow [5], Thurow *et al* [6]) has the potential to improve scan rates; however, the complexity and expense of the laser and camera systems are currently too prohibitive for broad application.

Four other techniques that have recently received attention for their ability to perform 3D/3C measurements are defocusing PIV (Willert and Gharib [2], Pereira *et al* [7]), holographic PIV (Hinsch [8], Herrman and Hinsch [9]), tomographic PIV (Elsinga *et al* [10]), and synthetic aperture PIV (Belden *et al* [11]). Defocusing PIV is based on the use of a specialized aperture near the camera lens to encode the depth of a particle into a uniquely sized and shaped blur spot when the particle is located off of the focal plane. Computational algorithms use the knowledge of the aperture shape to determine the particle position and depth. More recently, this process has been adapted for a multi-camera experimental arrangement, where the image formed by each camera represents a different point on the overall system aperture. The strength of this technique is the relative simplicity of the equipment required, particularly in the single camera case; however, the particle density is limited to relatively low levels since the location of individual particles must be resolved. Also, in the case of the single camera system, the use of a coded aperture greatly reduces the amount of collected light. The combination of these factors typically restricts the application of the technique to water tunnels, where particle density can be precisely controlled and relatively large particles can be used.

Holographic PIV is based on the recording of the interference pattern, or hologram, generated by a reference light beam passing through a volume. The volumetric light intensity distribution is then reconstructed by illuminating the hologram with the same reference light beam or a synthetic reference beam. The resulting volume represents the light intensity field, which can then be evaluated to determine particle positions or perform cross-correlation. Current efforts focus on digital holographic PIV utilizing CCD sensors and digitized reconstruction algorithms, most notably a study on wall-bounded turbulence (Sheng *et al* [12]). Nevertheless, these techniques are limited to small measurement volumes, while maintaining a high optical complexity, thus precluding the wide spread adoption of the technique in the near future.

Tomographic PIV has seen rapid development and maturation, and is now offered as a commercially available system; for a comprehensive review of tomographic PIV see Scarano [13]. Briefly, in this technique, four or more high-resolution CCD cameras are used to image a particle field illuminated by a thick laser sheet. Tomography algorithms are used to reconstruct the 3D light intensity distribution discretized over voxels, after which cross-correlation algorithms are used to determine the particle displacement. This technique has been demonstrated in a variety of flows including turbulent boundary layers (Schroder *et al* [14]), cylinder wakes (Scarano *et al* [15]), and shock-wave/turbulent boundary layer interactions (Humble *et al* [16]). It has also been adapted to kHz rates using high-speed cameras for aeroacoustic studies (see Violato *et al* [17]). Tomo-PIV, however, has some rather significant restrictions that limit its use in many situations. These include the relatively thin (∼10 mm depth) volume over which a measurement can typically be made, errors in the volume reconstruction process due to the limited number of viewing angles (e.g. the generation of image artifacts known as ghost particles), the limited particle number density, the complexity of the experimental arrangement and the expense of the overall system. Nonetheless, tomo-PIVs success in obtaining 3D, 3C velocity measurements in a multitude of facilities is notable and has revitalized recent research in 3D flow diagnostics.

Synthetic aperture PIV (SAPIV) is another multi-camera 3D PIV technique, described by Belden *et al* [11]. This technique uses a large camera array (eight or more cameras) to capture multiple views of the measurement volume simultaneously. In contrast to tomo-PIV, the map-shift-average algorithm is used to construct synthetically refocused images from the individual views by projecting each view onto a common focal surface. In the resulting image, particles that lie on the focal surface are sharp and in-focus, whereas particles off of the surface are blurred. By thresholding the refocused images, the 3D intensity field is compiled and is used as the input to cross-correlation algorithms. The technique is limited by many of the same restrictions as tomo-PIV, and unfortunately uses an even greater number of cameras.

1.2. Light field imaging

The field of light field imaging has experienced significant growth over the last couple decades and has evolved into a rich and active area of research. In this section, we attempt to provide a basic overview of the history and fundamental concepts of light field imaging; however, the reader is encouraged to consult other sources, such as Adelson *et al* [18, 19], Levoy *et al* [20, 21], Ng *et al* [22], and Lumsdaine and Georgiev [23], for more detailed information.

Historically, the notion of a light field is over a century old, with its roots outlined in Lippman [24]. The modern definition of a light field comes from Adelson and Bergen [18], where space is described as being filled with a dense array of light rays of varying intensities. These light rays contain information about our world and can be described in a systematic manner using the plenoptic function. The plenoptic function refers to the parameterization of the light field, where each

Figure 1. Computationally refocused images generated from a single exposure, focused: (left) on an alarm clock that is in front of the nominal focal plane, (center) at the nominal focal plane, and (right) on a student behind the nominal focal plane.

light ray is represented by its 3D position in space (x, y, z) and its angle of propagation (θ, ϕ), thus forming a 5D function[3] representing all light rays traveling through space. Assuming the constant intensity, or more precisely irradiance, of a light ray along its path of propagation, the plenoptic function is typically reduced to a 4D function, denoted as $L(x, y, \theta, \phi)$. In this context, a conventional photograph or image can be thought of as a 2D projection of the 4D light field where the angular dimensions have been integrated out at the sensor plane.

Adelson and Wang [19] utilized this concept to estimate the depth and shape of objects by measuring the plenoptic function with a single camera, referred to as a plenoptic camera. The camera utilized a specialized optical design to encode both the spatial (x, y) and angular (θ, ϕ) components of the light field onto a 2D image sensor. In a conventional camera, a main lens collects light across a range of input angles bound by the size of the aperture and focuses the light directly onto the image sensor, which records the total light intensity at each pixel regardless of the angle of incidence. In contrast, in a plenoptic camera the main lens focuses the entire angular distribution of light onto an array of microlenses. Each microlens covers a small number of pixels on the image sensor and can be thought of as forming a macropixel. In this configuration, the microlenses capture the spatial information contained in the light field, while the pixels contained under the microlens record the angular distribution. This relationship will be described in greater detail in the following section. Adelson and Wang's version of the plenoptic camera utilized a 500×500 pixel CCD with a microlens array of 100×100 microlenses. This results in a camera with a spatial resolution of 100×100 pixels with an angular sampling of 5×5.

Capturing and altering the light field is not limited to using a plenoptic camera. Levoy [20, 21] describes several methods of obtaining the light field in order to computationally generate an image or rendering of an object. One method places the object of interest at the centre of a sphere, then, using a spherical gantry, thousands of images can be taken at different positions along the sphere's surface. The resulting collection of 2D images taken at discrete angles is a representation of the 4D light field. Another method is to mount multiple cameras, Levoy [21] used 128, in an array allowing an instantaneous light field to be acquired. These techniques utilize multiple 2D images to build the 4D light field. In this vein, we note that defocus PIV, tomo-PIV and SAPIV are implicitly measuring the light field, albeit with relatively low angular resolution. In contrast, the plenoptic camera directly captures the 4D light field on a single image sensor in a single snapshot, with a fairly dense angular sampling over a limited angular range.

As camera and microlens technology has improved, the interest in plenoptic cameras has grown. Of the more recent developments we particularly note the work of Ng [25, 26], who designed a hand-held plenoptic camera for digital photography. The camera consisted of a modified DSLR with a 16 megapixel image sensor and a microlens array of 296×296 microlenses. Ng's research focused on computationally rendering conventional 2D images from the light field data collected by the plenoptic camera in a single snapshot. They demonstrated the ability to computationally generate, after the fact, photographs with a different focal position or a shift in the perspective. Examples of refocused images acquired with our plenoptic camera (described later) are shown in figure 1. The three images represent the focus shifted toward the camera, stationary, and shifted away from the camera relative to the nominal focal plane. In figure 2, the perspective of the observer is shifted with one image showing a 'left' view and the other showing a 'right' view. These images serve to illustrate the unique information obtained by a plenoptic camera and how it can be used for computational imaging. Recently, commercial variants of plenoptic cameras have become available. For consumer photography Lytro (Founded by Ng) offers a point-and-shoot plenoptic camera with built in refocusing capabilities. In the field of machine vision Raytrix offers a 'plenoptic 2.0' camera that offers a similar ability to change the perspective of an image after the fact.

More recently, Levoy *et al* [27, 28] developed a light field microscope based on the plenoptic camera. The fundamental principle remains the same; however, their work focused on additional challenges associated with microscopic imaging. For one, wave optics and diffraction must be considered in a microscopic environment, whereas geometrical optics is sufficient for macroscopic imaging. In addition, a typical microscope objective functions differently from a normal camera lens, producing orthographic rather than perspective views. Next, most objects in microscope images are partially transparent, whereas the previous effort had focused on scenes with opaque objects.

The objective of the current work is to utilize the information obtained about the light field by a plenoptic camera to obtain 3D/3C PIV measurements. Section 2 describes the

[3] In a general sense, one can also include the wavelength, polarization and time dependency of light in space such that the full light field may be considered as an 8D function. This is known as the radiance function.

Figure 2. Computationally rendered image where the viewpoint of the observer has been changed to (left) the left side of the aperture and (right) the right side of the aperture.

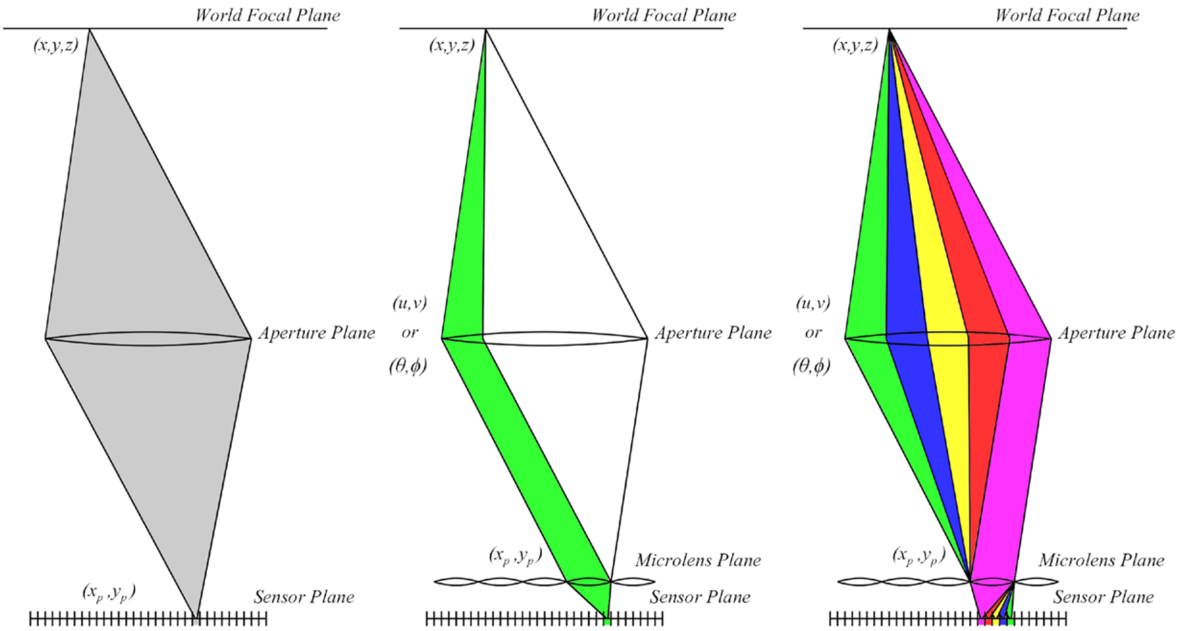

Figure 3. Illustration of the differences between a conventional camera (left) and a plenoptic camera (center, right) in how they sample the light field. The centre figure shows a single pixel's line of sight as captured by the plenoptic camera, collecting all light rays emanating from a point on the focal plane and propagating with angles contained in the green area. The left figure shows all of the pixels beneath a microlens' line of sight.

fundamentals of plenoptic imaging as well as the details of a prototype plenoptic camera built in our laboratory. Then, the process for building the light field from the data obtained by the camera, computational refocusing and the generation of perspective views are detailed in section 3. In order to facilitate the development of the reconstruction algorithms, a synthetic plenoptic image generation tool was developed. This tool allows the reconstructed volumes to be compared against a true solution, and is discussed in section 4. Section 5 discusses the coupling of the 4D light field to tomographic algorithms in order to reconstruct a volume of particles. In so doing, the unique relationship between the image and the volume, known as the weighting matrix, is defined. The weighing matrix is shown to be an evolutionary step forward from the computational rendering algorithms. In section 6, the reconstruction algorithms are tested with synthetic data (generated with the plenoptic simulation tool) in a number of cases. Finally, experimental results are presented in section 7.

2. The plenoptic camera

Figure 3 schematically illustrates the fundamental concept of a plenoptic camera by contrasting it with the familiar picture of a conventional camera. The main function of a plenoptic camera is to measure both the position and angle of light rays collected by an imaging lens. This is in contrast to a conventional camera which records only spatial information about incident light rays through the integration of the angular information at the sensor plane. In both cases, geometric optics can be used to map an arbitrary location (x, y, z) on the world focal plane to a corresponding location on the image plane (x_p, y_p). In a conventional camera (figure 3, left), the angular information is integrated, and therefore lost, as the image sensor records the amount of light striking that position, but not the angle. The separation between the imaging lens and sensor plane is typically chosen, via the thin lens equation, such that all rays converge to the same point leading to an in-focus image. If the sensor plane is not coincident with the image plane, the image will be out-of-focus with point sources on the world focal plane forming blur spots that are dependent on the size of the lens aperture and the position of the image sensor. This leads to a loss of spatial resolution in the image and the familiar concept of depth-of-field, where reducing the lens aperture leads to increased depth of field.

In a plenoptic camera, on the other hand, a microlens array is positioned at the image plane with the image sensor shifted back by one microlens focal length. The function of the microlens array is to direct light incident on the microlens at a particular angle onto one of the pixels located behind the microlens. This is depicted in figure 3, centre, which shows the point-of-view of a single pixel. Neighbouring pixels contained under the same microlens are exposed to light at different incident angles, as shown in figure 3, right where each colour represents a different subset of angles captured by each pixel. As such, each microlens in the array determines the position (x, y) of the light rays collected by the main lens and each pixel determines the angle (θ, ϕ) of the light rays striking that particular microlens. Alternatively, each microlens can be thought of as forming a micro-image of the main lens aperture. By considering the full array of micro-images formed under each microlens, the resulting 2D image recorded on the image sensor represents a multiplexed sampling of the 4D light field captured by the camera lens.

Figure 4. Prototype plenoptic camera.

Table 1. Prototype plenoptic camera parameters.

Parameter	Symbol	Value
Microlens pitch	p_l	0.125 mm
Microlens focal length	f_l	0.5 mm
Number of microlenses: *X*-direction	n_{l_x}	289
Number of microlenses: *Y*-direction	n_{l_y}	193
Pixel pitch	p_p	0.0074 mm
Number of pixels: *X*-direction	n_{p_x}	4904
Number of pixels: *Y*-direction	n_{p_y}	3280
Microlens array material		BK7/epoxy

2.1. Prototype camera

As part of the development of the plenoptic PIV technique, a prototype plenoptic camera has been constructed using an Imperx Bobcat ICL-B4820 camera as its base. This camera uses the Kodak KAI-16000 image sensor, which at the time of fabrication was the highest resolution commercially available interline CCD. The choice of an interline CCD is motivated by the need to perform a double exposure similar to traditional PIV cameras. Figure 4 shows a photo of the camera without a lens attached, and a US quarter to provide scale. The compact design of the camera is evident.

The microlens array was fabricated by Adaptive Optics Associates, a subsidiary of Northrup Grumman. Specifically, the microlenses are manufactured using a proprietary process, where an epoxy-filled mold is used to print the microlenses onto the glass surface. The primary challenge faced in constructing the prototype camera was fabricating a custom mounting device for the microlens array to position it accurately over the sensor. A custom mount was designed by Light Capture, Inc. and manufactured in-house. The mount consists of a series of positioning screws to adjust the height of the microlens array above the sensor and to adjust the orientation of the array with respect to the sensor. To align the camera, we follow a similar procedure to that outlined in Ng *et al* [22]. In this procedure, the main lens of the camera is removed and the microlens and image sensor are exposed to an approximately collimated beam of light (a point source at a distance). In this configuration, each microlens forms a small spot on the image sensor with a diameter determined by its distance from the image sensor. For proper alignment (image sensor at the focal plane of the microlenses), the microlens mount is adjusted until the spot size reaches a minimum value. To determine this, the image captured is displayed on a monitor while adjustments are made to the mount. This is accomplished in a few iterations. The alignment procedure is accurate, for this microlens array, within a range of 36 μm [22]. The full parameter list for the CCD and microlens array are shown in table 1.

To test the camera, a simple scene was set up using some objects set up at different depths. As with a conventional camera, the field of view was adjusted by changing the main lens attached to the camera body. This is seen later in this work where we focus nominally on a 1 : 1 magnification. The objects were set up with one in the foreground (alarm clock), some in focus (camera box, hat), and one in the background

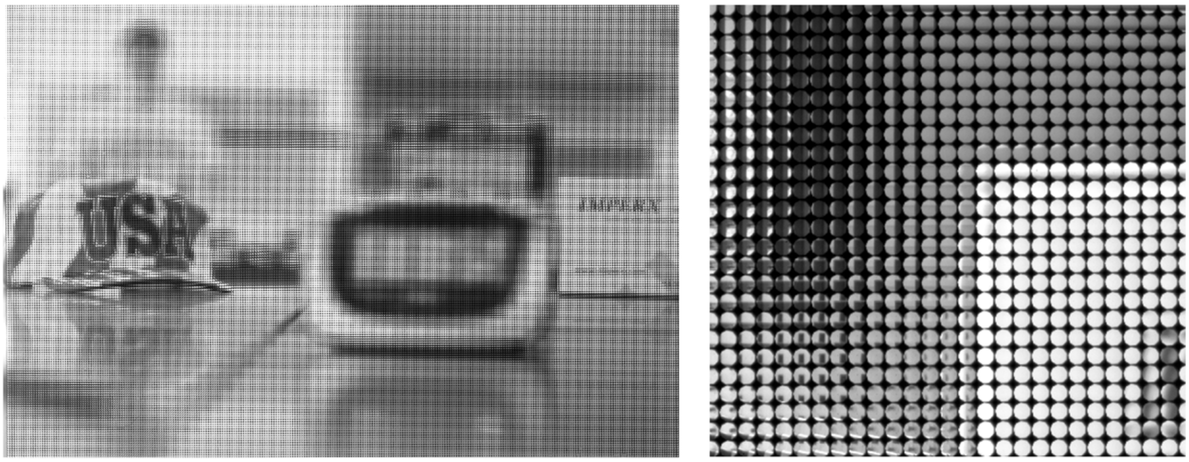

Figure 5. Raw image taken with prototype plenoptic camera (left) and a close up of the same raw image, showing individual microlenses (right).

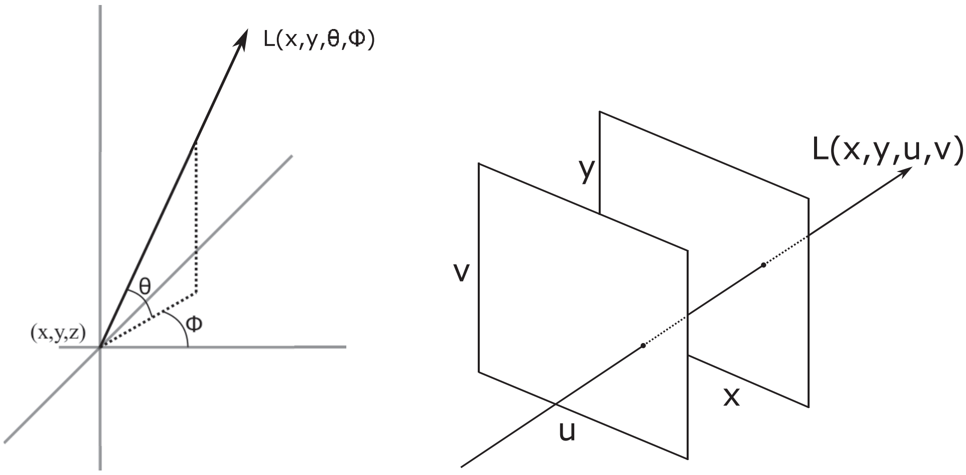

Figure 6. Two geometric representations of a light ray. The first parameterizes the light ray by its position and angle of propagation (left) and the second parameterizes the same light ray by a pair of points on two planes (right). Adapted from Levoy [21].

(graduate student). A raw image of this scene is shown in figure 5, left. From the raw image it is hard to discern the differences from a conventional photograph as each microlens image is vanishingly small when viewed in such a small format. A close look at the top left corner of the Imperx box, as shown in figure 5, right, shows the individual microlens images that comprise the larger image.

3. Light field rendering

3.1. Two-plane parameterization

The preceding discussion parameterizes a light ray by its position on the world focal plane and angle of propagation. An alternative, and often more convenient, way to parameterize the light field is known as the two-plane parameterization. The discussion herein is derived from Levoy [20]. Figure 6, left describes a light ray by its position (x, y, z) and its angle of propagation (θ, ϕ). Figure 6, right shows a light ray that is defined by pairs of points, (x, y) and (u, v), located on two planes separated by a known distance. These two descriptions of the light ray are equivalent, since they can be derived from each other using simple trigonometric relations.

The plenoptic camera lends itself to this type of parameterization due to it inherently having two primary planes that light rays intersect: the microlens plane and the aperture plane, separated by a fixed distance, s_i. As discussed previously, the microlenses are responsible for discretizing the spatial location of all the incoming light rays. The second plane, the aperture, represents the angular information, where each microlens is effectively forming an image of the aperture on the image sensor. Therefore, each pixel of the image sensor is associated with a discretized point on the microlens plane (x, y coordinate) as well as a point on the aperture plane (the u, v coordinate) separated by the image distance of the main lens.

The two-plane parameterization offers a more straightforward and convenient representation of the light field as the upper and lower bounds of the aperture plane are fixed and constant for every microlens. This is in contrast to the angular parameterization, where the range of sampled angles varies with each microlens.

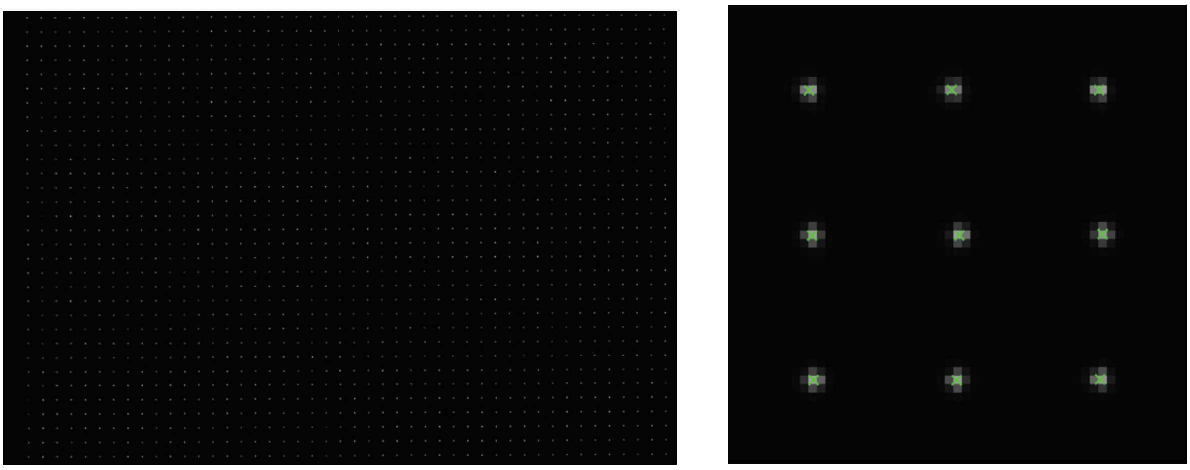

Figure 7. Subset of an experimental calibration image (left) and corresponding centroid fit (right).

3.2. Building the light field

Using the above-mentioned two-plane parameterization, the recorded light field can be fully described through determination of the (x, y, u, v) position of each pixel. For experimentally obtained images, the exact locations of the microlenses relative to the image sensor are not known. As such, a calibration procedure was developed to determine the positions of the microlenses and the pixels beneath them. This procedure begins by taking a calibration image. This image is obtained by minimizing the aperture of the camera (i.e. increasing the f-stop to its maximum value) and imaging a uniformly illuminated white surface, such as a piece of paper, while keeping the focal position of the camera constant. The last statement is very important as the positions of the microlens images on the CCD shift depending on the main lens configuration. A sample calibration image is shown in figure 7, left. The white dots are the centres of the reduced aperture image formed by each microlens. In terms of the two-plane parameterization these dots represent the centre of the aperture (x, y, u_0, v_0). Since the aperture is not closed to a perfect point and the centre of a microlens may not fall directly on a single pixel, the exact location of each microlens is calculated to sub-pixel accuracy using a simple centroid fit. An example of the centroid fit is shown in figure 7, right, where the centre of each group of pixels is shown as a green 'x'.

The calibration procedure uses *a priori* knowledge of the microlens array, specifically we assume that they are arranged in a rectilinear fashion with a pitch of 125 microns. According to the manufacturer specifications the pitch of the microlens is subject to a ±3% non-cumulative error. From the calibration procedure, the (x, y) values for each pixel can be assigned using the value of the microlens in front of them. To determine the angular components for each pixel the distance between the microlens array and main imaging lens, s_{i}, must be determined. To do this an image of a ruler is taken at the nominal focal plane, allowing for the calculation of the nominal magnification. From this, and the definition of magnification, the distance s_{i} can be determined. The u and v values are locations on the main lens that correspond to the pixels themselves. Their calculation takes the difference between their position and the centre of the microlens (x, y) and converts it to a distance away from the centre of the aperture using similar triangles. This expression is given for the u component by

$$u_i = (x - x_i)\frac{p_{\mathrm{p}} s_{\mathrm{i}}}{f_{\mathrm{l}}} \tag{1}$$

where the subscript i represents the current pixel. A similar expression using y values is used for the v component.

Once the centres of the microlenses are known, the (x, y, u, v) values for each pixel can be determined. The x and y values for each pixel are the microlens centre (in mm) determined earlier in the calibration process. The u and v values are the position on the main lens aperture that each individual pixel is imaging. To calculate these positions the distance from the centre of the microlens (centre of the aperture) is measured to the centre of each pixel. Then, using similar triangles, the measured distance is converted to a distance away from the centre of the aperture in millimetres. This fully parameterizes each pixel recorded in the raw image.

3.3. Computational refocusing

A simple introduction into manipulating a light field is to resample the light field at a new focal plane. This process, termed computational refocusing, has been adapted from the work of Ng [25] and relies on the two-plane parameterization of the light field. One consequence of the refocusing algorithm was the insight needed to construct a physically accurate weighting function for tomography, to be discussed in a later section.

Conceptually, the rendering of a traditional 2D image from the 4D light field is achieved by selecting a subset of rays from the complete 4D light field and integrating out the two angular dimensions for a pre-determined focal plane. Using the two plane approach a simple interpolation scheme can be applied to re-sample the light field inside the camera at a virtual image sensor location creating a refocused image. An illustration of the geometry used in the refocusing process is shown in figure 8.

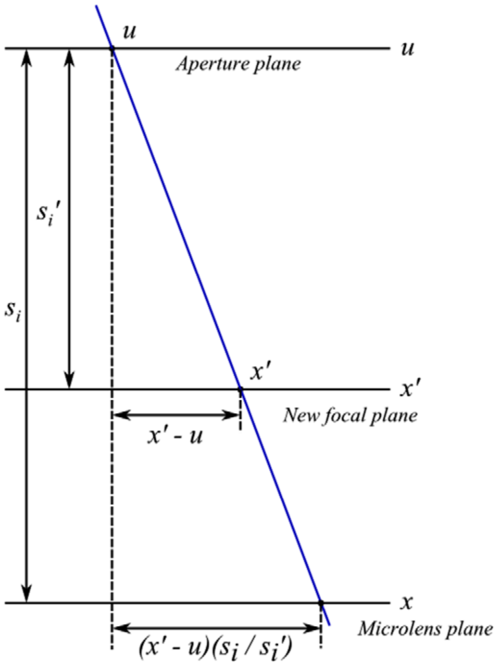

Figure 8. Illustration of interpolation for refocusing using the two-plane parameterization. Adapted from Ng [25].

To generate a refocused image, the light field is resampled at a virtual image sensor x' located at a distance s_i' from the aperture plane. The virtual light field L' can be written in terms of the original light field L through a linear projection operator, as shown graphically in figure 8, where the desired virtual light field being resampled at (x', u) is projected onto the original sensor yielding the point (x, u) in the recorded light field. Mathematically, the location of this projection from x' onto x for a single u value, denoted x_find is given by

$$x_\mathrm{find} = u\left(1 - \frac{1}{\alpha}\right) + \frac{x'}{\alpha} \tag{2}$$

where $\alpha = s_\mathrm{i}'/s_\mathrm{i}$. Substituting x_find into the plenoptic function results in an equation for the light field located at a virtual image sensor (x', y') expressed in terms of the original light field, and is given by

$$L'(x', y', u, v) = L\left(u\left(1 - \frac{1}{\alpha}\right) + \frac{x'}{\alpha}, v\left(1 - \frac{1}{\alpha}\right) + \frac{y'}{\alpha}, u, v\right). \tag{3}$$

To generate a refocused image at the synthetic image sensor plane, the angular information contained in the light field is integrated such that the final value for each microlens is the sum of all its angles. This is expressed in equation form by

$$I(x', y') = \int\int L\left(u\left(1 - \frac{1}{\alpha}\right) + \frac{x'}{\alpha}, v\left(1 - \frac{1}{\alpha}\right) + \frac{y'}{\alpha}, u, v\right)\mathrm{d}u\mathrm{d}v. \tag{4}$$

Due to both the non-uniformity of the u, v sampling caused by each microlens being displaced differently from the optical axis, and that the projection x_find may not necessarily coincide with a single microlens, a 4D interpolation scheme is required to determine the contribution of each pixel. An example of the refocusing algorithm applied to actual image data was shown in figure 1.

3.4. Perspective shift

Another benefit of capturing the entire light field is the ability to change the perspective of the scene, or in other words to change the angle at which the scene is presented. These images are generated by only considering a single angle (i.e. aperture position) in the light field. Similar to the refocused image, a single value is used to represent a microlens; however, instead of summing the angular information into a single value, a specific angle (u, v) is chosen and that value is used. As the u, v plane corresponds to the aperture plane, we can generate perspectives where the viewer is located at different points across the aperture. Some sample images of this effect are shown in figure 2.

4. Synthetic image generation

While the previous discussion about manipulating the light field is useful for understanding its unique capabilities, it does not directly apply to 3D fluid velocimetry measurements. To develop this technique synthetic data is needed to test the overall accuracy of the particle reconstruction algorithm. Specifically, synthetic data allow the reconstructed volumes to be compared against a known solution, whereas experimental data do not allow for such a comparison. To do this a plenoptic camera simulator has been developed and is detailed herein.

4.1. Overview

At the core of the simulation is the use of linear (Gaussian) optics to geometrically trace the path of light through space and the various optical elements that comprise the plenoptic camera. The application of Gaussian optics used in this work is similar to ray tracing from computer graphics. Briefly, ray tracing is a rendering technique in which a large number of light rays from a scene are used to form an image at arbitrary locations or viewpoints. Rays of light are initialized at the light source by specifying an initial position and direction. Ray transfer matrices are used to simulate optical elements and the propagation of light through free space [29]. The intersection that each ray makes with a sensor plane or designated viewpoint defines the generated image. An extension of Gaussian optics, known as affine optics (Georgeiv and Intwala [30]), was developed to apply linear optics to light field imaging. The synthetic plenoptic image tool discussed herein uses ray transfer matrices, with the affine optics extension, to simulate 3D particle fields imaged by a plenoptic camera.

As mentioned previously, the optical configuration for a plenoptic camera differs from a conventional camera with the addition of a microlens array. In order to construct the simulator, the following variables and relationships are defined in figure 9. Due to the nature of the ray transfer matrices all parameters are measured relative to the optical axis in both the x- and y-directions. The origin of the z-axis is defined at

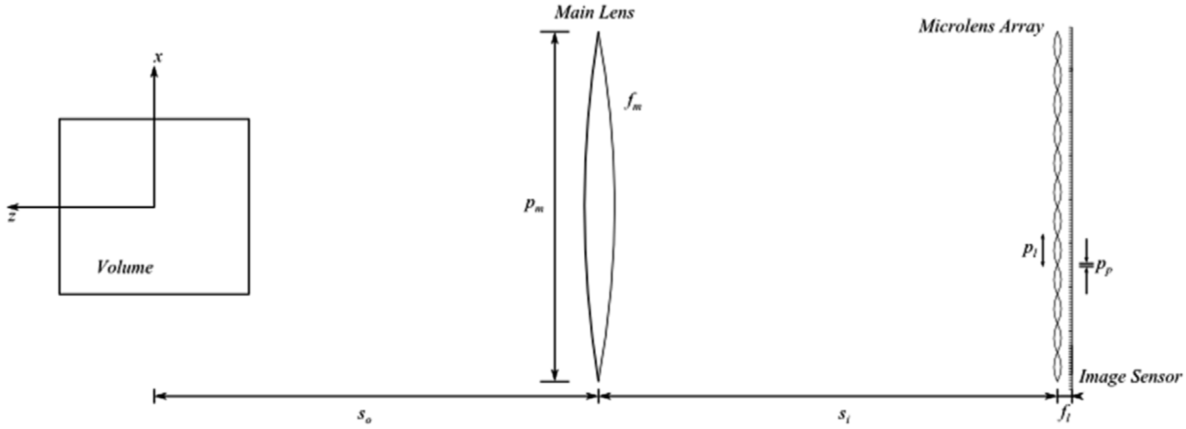

Figure 9. Optical configuration of the plenoptic camera.

the nominal focal plane of the camera with positive z pointing away from the camera.

Particle positions are defined by their position relative to the centre of a volume positioned at the nominal focal plane of the main lens, where the main lens is modeled as a thin lens with focal length, f_m, and an aperture with diameter, p_m. Similarly, the microlenses are defined by their focal length, f_l, and pitch, p_l. The physical image sensor is defined by a pixel pitch, p_p, which denotes the size of a pixel. The distances separating the elements are the object distance, s_o, which separates the focal plane of the camera and the main lens and the image distance, s_i, which separates the main lens and the microlens array. The image and object distances are related by the thin lens equation, as shown in equation (5), which makes the assumption that the thickness of the lens is negligible relative to the length of the optical system itself.

$$\frac{1}{s_i} + \frac{1}{s_o} = \frac{1}{f_m}. \tag{5}$$

We note that modern camera lenses, which typically contain multiple lens elements, can be approximated by a thin lens where s_i and s_o are measured relative to the principal planes of the lens. While not considered here, the present framework also allows for more detailed modeling of these additional lens elements. s_i and s_o are related to the magnification of the imaging system through equation (6).

$$M = -\frac{h_i}{h_o} = -\frac{s_i}{s_o}. \tag{6}$$

In combination with equation (5), this equation allows for the calculation of s_i and s_o knowing only the magnification, which can be obtained by imaging a ruler, and the focal length of the main lens.

The optical parameters are now divided into two categories: input and fixed parameters. The input variables can change with each experiment and include the main lens focal length, aperture diameter and magnification. The object and image distances are also variable; however, as shown previously, they are dependent on the main lens focal length and magnification. The second class of parameters are set through hardware design and cannot be modified once the camera has been assembled. These include microlens pitch, microlens focal length, pixel pitch and the number of pixels. These parameters can be modified in the simulator to accommodate testing and camera design, but are not varied in this work.

Table 2. Variable parameters for plenoptic camera simulation.

Parameter	Symbol	Value
Main lens focal length	f_m	50 mm
Main lens F-number	$(f/\#)_m$	2
Magnification	M	−1
Number of microlenses: *Y*-direction	n_{l_y}	193
Object distance	s_o	100 mm
Image distance	s_i	100 mm

One consequence of the microlens parameters being fixed is a forced condition known as f-number matching. This condition, recognized by Ng *et al* [25], states that the image-side f-number of the main lens must be equal to or greater than the f-number of the microlenses. This condition prevents any overlap between adjacent microlens images, which would otherwise cause ambiguity in the light field parameterization. The equation for calculating the image-side f-number, as described by Ng [22], is shown in equation (7), where f is the focal length, and $f/\#$ is the f-number, which is defined as the focal length divided by the size of the aperture.

$$(f/\#)_m = (f/\#)_l/(1 - M). \tag{7}$$

In this work, we simulate a nominal 1 : 1 imaging magnification such that $h_i = h_o$ and $M = -1$. In the future, the parameterization of the plenoptic camera as a function of magnification needs to be considered; however, in order to keep the number of variables used in this work manageable only a single magnification is used. The fixed parameters used in the present simulation are shown in table 1 and are used throughout this work unless otherwise noted. The input or variable parameters used throughout this work are shown in table 2.

In this regard, it is worth commenting that the degree of parallax observed in the perspective views is limited by the size of the lens aperture used to form the image and the object's

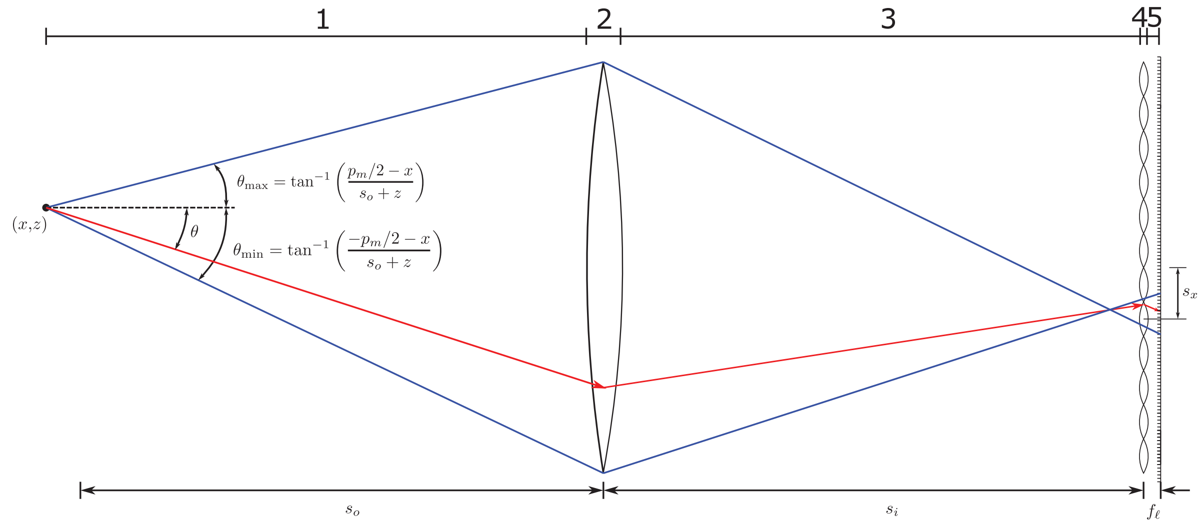

$$\underset{1}{\begin{bmatrix} 1 & 0 & s_o+z & 0 \\ 0 & 1 & 0 & s_o+z \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}} \quad \underset{2}{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ -1/f_m & 0 & 1 & 0 \\ 0 & -1/f_m & 0 & 1 \end{bmatrix}} \quad \underset{3}{\begin{bmatrix} 1 & 0 & s_i & 0 \\ 0 & 1 & 0 & s_i \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}} \quad \underset{4}{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ -1/f_\ell & 0 & 1 & 0 \\ 0 & -1/f_\ell & 0 & 1 \end{bmatrix} + \begin{pmatrix} 0 \\ 0 \\ s_x/f_\ell \\ s_y/f_\ell \end{pmatrix}} \quad \underset{5}{\begin{bmatrix} 1 & 0 & f_\ell & 0 \\ 0 & 1 & 0 & f_\ell \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}$$

Figure 10. Schematic of the ray-tracing process for a plenoptic camera.

location relative to the main lens. Ultimately, the aperture size is limited by the requirement that the f-number of the microlenses must be matched to the image-side f-number of the main imaging lens. In the work described herein, we focus on 1 : 1 imaging with $f/4$ microlenses. Under these conditions, the f-stop of the imaging lens is set to $f/2$ with a nominal working distance equal to $2f$ (i.e. 1 : 1 magnification is achieved at a working distance equal to twice the focal length of the imaging lens).

The process of the ray-tracing simulation is shown schematically in figure 10. For each synthetically generated particle, represented as a point source located at (x, y, z), a large number of rays (typically >10 000) are used to simulate the light emanating from that point. Each light ray is given an initial position, determined from the particle's location, as well as an initial angle. The angle is generated as a random number between $\theta_{\min}$ and $\theta_{\max}$, which are determined based on the distance to the lens and the aperture size. In figure 10 the maximum angles are shown as the outermost blue rays, and the expressions for the maximum and minimum angles are given. From this initial state the ray is propagated to the main lens using the first ray-transfer matrix, labeled as 1. From there the use of a lens ray-trace matrix, number 2, is used to model the main lens, then the light ray is propagated to the microlens array using matrix 3. Once at the microlens array, the individual microlens that the ray has struck is determined. From there the affine optics adaptation of the lens ray-trace matrix is used to model the microlens, as shown in matrix 4, which also includes a matrix addition term. Finally, the ray is propagated to the image sensor using the final matrix, 5. Once at the image sensor the pixel which the ray hits is determined and its value is increased.

It should be noted that the simulator takes into account diffraction effects by randomizing the spatial coordinate of each light ray at the microlens plane and sensor plane through a normally distributed random number generator, set in a manner that the standard deviation is equal to the diffraction-limited spot size. For both the microlens array as well as the main lens the diffraction-limited spot size is 5.2 μm. Analysis at the condition presented here indicates that diffraction does not result in a substantial change in the simulator results. This is due to the large f-number of the main lens and the microlenses, where the diffraction limited spot size is smaller than the characteristic spatial dimensions (microlens and pixel pitch) of the camera.

4.2. 1D simulations

A 1D simulator was constructed as a simple means to evaluate basic camera concepts without requiring a full image simulation, and is far easier to visualize. A detailed description of the simulator construction in Lynch [31], but the results are shown here to illustrate the ray tracing process. Figure 11 (top) shows a particle simulated at the focal point of the optical system. The red lines represent the ray propagation from the particles position through the entire optical system culminating at the image sensor. The blue line, shown behind the CCD, is the integrated signal resulting from the ray tracing procedure. In this case the rays converge onto a single microlens, then spread out onto the image sensor. In figure 11(a)–(d), the particle is moved in the volume illustrating the unique signal patterns formed by the plenoptic camera. In figure 11(a), all of the light rays converge in front of the microlens plane in a manner that is consistent with the image plane moving closer to the main lens as the object plane moves further away. After passing through this focal point, the rays spread out and intersect several microlenses. Depending on the incident angle, the microlenses

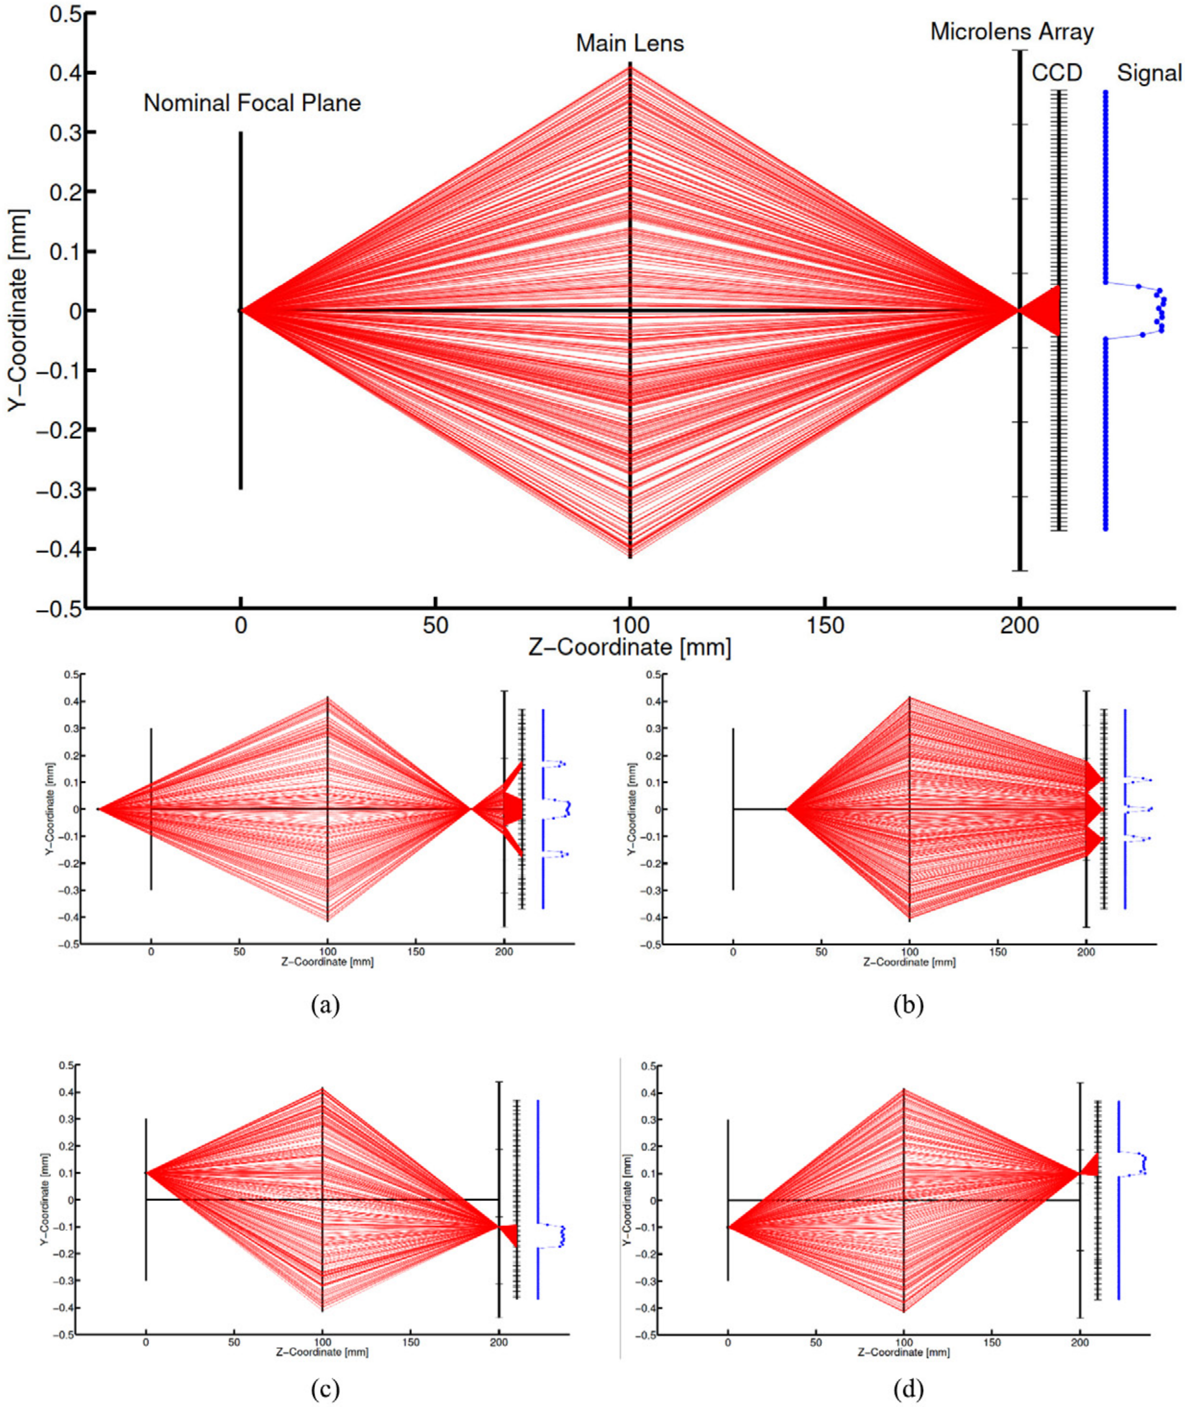

Figure 11. 1D simulations at different lateral positions. 1 out of every 100 rays shown. The integrated signal is shown in blue. (a) dz = +30 mm. (b) dz = −30 mm. (c) dy = +0.1 mm. (d) dy = −0.1 mm.

redirect the incident light to different pixels on the image sensor forming a unique image pattern corresponding to the particle positions. Conversely, in figure 11(b), the light rays are intersected by the microlens array prior to reaching their focal point forming a distinctly different image pattern. Figures 11(c) and (d) show the effect of shifting the particles position in the y-direction. The effect shown is that the signal is simply shifted. This fundamental relationship between position in the volume and the pattern formed on the image sensor will be described later in terms of a weighting matrix for tomographic reconstruction.

4.3. 2D simulations

A sample of the full 2D simulator is shown in figure 12, left. The image provided is a subset of a full image whose size is set in accordance with the KAI16000 image sensor to 4872 × 3248 pixels. This image was generated using a particle

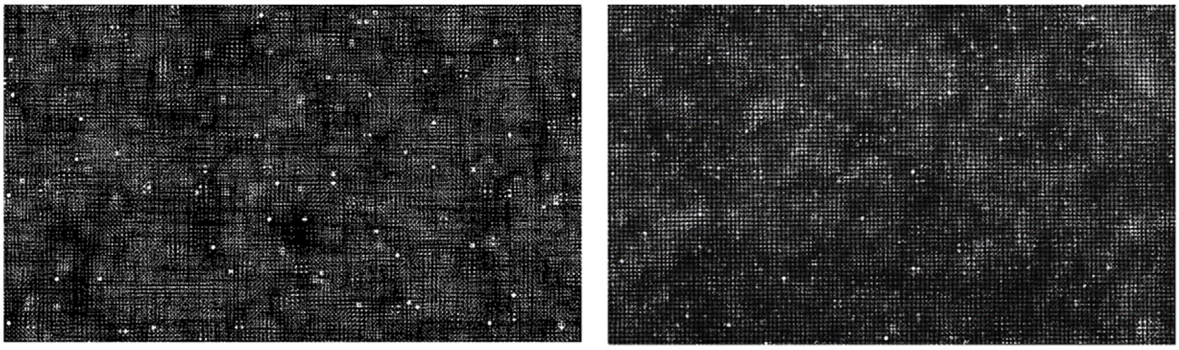

Figure 12. Example plenoptic image generated using the ray-tracing simulator (left) as well as an experimental image taken with the prototype plenoptic camera (right).

volume ranging from $z = -10$ mm to $+10$ mm and a particle density of 0.5 particles per microlens (ppm) or 0.0017 particles per pixel (ppp) resulting in a particle concentration of 2.32 part mm^{-3}. Upon a visual inspection of the image, the particles that lie near the focal plane produce nearly circular images that stand out from the rest of the field. The remaining particle images are distributed across multiple microlenses and are difficult to distinguish. As a comparison an experimental image taken with the prototype plenoptic camera is provided in figure 12, right.

5. Tomographic reconstruction

To reconstruct a volumetric intensity field useful for PIV, tomo-PIV principles are used with appropriate modifications. The working principle of tomo-PIV as detailed in Elsinga *et al* [10] is used to reconstruct a volume of particles based on a finite number of 2D projections of the volume. Plenoptic PIV is similar in that the recorded angular information is equivalent to a projection image acquired in a tomo-PIV experiment. A significant difference is that plenoptic PIV records a much higher density of projections on a single camera, albeit at the expense of spatial resolution.

5.1. Basic concept

For plenoptic PIV the reconstruction of particle fields is in general an ill-posed problem whose system of equations is underdetermined leading to ambiguity in the solution. A special class of reconstruction algorithms are better suited for these problems and are known as algebraic methods as described by Herman and Lent [32]. These methods rely on iteratively solving a system of linear equations which model the imaging system. As with conventional tomo-PIV the 3D volume to be reconstructed is discretized into cubic voxel (volume equivalent of a pixel) elements, with intensity $E(x, y, z)$. The size of the voxel was chosen to be similar to that of a microlens, since they nominally govern the spatial resolution of a plenoptic camera. The problem can be stated as the projection of the volumetric intensity distribution $E(x, y, z)$ onto a pixel located at (x_i, y_i) yields the known intensity of that pixel, $I(x_i, y_i)$. In equation form this is given by

$$\sum_{j \in N_i} w_{i,j} E(x_j, y_j, z_j) = I(x_i, y_i) \qquad (8)$$

where N_i represents the number of voxels in the line-of-sight of the ith pixel. The weighting function $w_{i,j}$ describes the relationship between the recorded image (ith pixel) and the 3D volume of interest (jth voxel), and is detailed in the next section. In order to solve this set of equations, iterative techniques have been developed that update the current solution for E based on the previous solution. For additive techniques such as the algebraic reconstruction technique (ART [32]) the update is based on the difference between the image intensity data and the projection of the volume such that when they are equal the update added to the solution is zero. For multiplicative techniques such as the multiplicative algebraic reconstruction technique (MART [32]) the update is based on the ratio of the image intensity data to the projection of the volume such that when they are equal the update multiplied to the solution is unity.

The algorithm used in this work is the standard MART algorithm, which was shown by Elsinga *et al* [11] to work well in multi-camera tomo-PIV. Starting from an initial guess of the volume $E(x_j, y_j, z_j)^0 = 1$ MART is updated via the following expression

$$E(x_j, y_j, z_j)^{k+1} = E(x_j, y_j, z_j)^k \left(\frac{I(x_i, y_i)}{\sum_{j \in N_j} w_{i,j} E(x_j, y_j, z_j)^k} \right)^{\mu w_{i,j}} \qquad (9)$$

where k is the number of iterations and μ is a relaxation parameter which must be less than or equal to one. The exponent restricts updates to parts of the volume affected by the ith pixel by raising the argument to 0, therefore multiplying the current voxel by 1, if the voxel is not affected by the ith pixel.

5.2. Calculation of the weighting function

In order to use tomographic reconstruction, a weighting function describing the unique relationship between the plenoptic camera and the volume must be determined. In techniques such as tomo-PIV, the weighting function is based on a straight line projection of a pixel through the volume. The weighting coefficients are calculated as the overlapping volume between

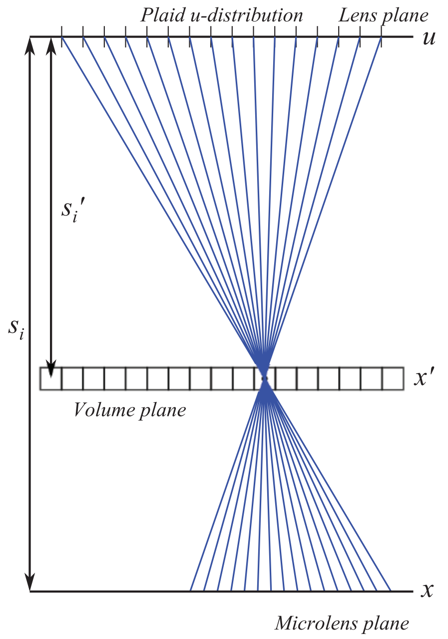

Figure 13. Demonstration of two-plane projection of x' and u in two dimensions.

the pixel's line-of-sight and the voxel's elements normalized by the volume of a voxel. This weighting function works well when the entire volume is in focus, such that the line-of-sight of the pixel is a decent approximation for the formation of the image. Due to the unique point spread function of the plenoptic camera as well as the fact that the volume is ideally out-of-focus (at least in a conventional sense), this method of calculating the weights is not applicable. With this in mind, a new method for determining the weighting function was developed by considering the unique nature of the plenoptic camera. This approach, inspired by our experience with the computational refocusing algorithm [33], is based on interpolating a distribution of light rays passing through a virtual point within image space.

The method begins by defining the discretized volume in object space that we wish to reconstruct. In this work, we assume a conventional Cartesian grid with uniform spacing between all the volume elements. The coordinates in object space are then transformed into image space (i.e. inside the camera), where the light field was measured. Each voxel element (x_o, y_o, z_o), with the subscript o referring to a location in object space, uses the following transformation for z:

$$s_o' = s_i + z_o \tag{10a}$$

$$s_i' = f_m \cdot s_o' / (s_o' - f_m) \tag{10b}$$

$$\alpha = s_i'/s_i \tag{10c}$$

where s_o' is the distance from the main lens to the voxel and s_i' is its image space counterpart, calculated using the thin lens equation. The term α, the ratio of the voxel's image space location to the nominal image distance, is used instead of the actual location in image space. For x and y the following transformations are used: $M' = -s_i'/s_o'$, $x' = x_o \cdot M'$, and $y' = y_o \cdot M'$, where M' is the magnification of the voxel. The result is a voxel in image space whose position is given by (x', y', α).

Once the discretized volume has been transformed into image space, each slice of the volume in the depth direction can be treated like a focal plane for refocusing, except instead of considering the distribution of rays as converging toward a voxel, they are considered as emanating away from the voxel towards the plenoptic camera. Figure 13 shows a distribution of rays passing through a particular voxel, denoted by x', that are defined by first specifying their position on the (u, v) plane. The number of rays intersecting the (u, v) plane and projected through the volume is chosen such that the resulting spacing at the x-plane is less than one microlens dimension. For x' planes that are relatively distant from the x-plane, this results in an oversampling of the (u, v) plane relative to the nominal angular sampling rate (i.e. the number of pixels under each microlens) of the camera. Additional oversampling on the (u, v) plane results in additional computational expense, but does not contribute additional information to the calculation of the weighing matrix.

In contrast to refocusing, where we are interested in interpolating the light ray's intensity through a particular voxel, we utilize the interpolation coefficients themselves as a measure of the weighting between the voxel and the image pixels. For a single light ray passing through a voxel (x_2, y_2, u_2, v_2), where the subscript 2 refers to the point of interpolation in interpolation space, there are sixteen coefficients that are used to interpolate the irradiance of the light ray from the measured light field. These coefficients can be more easily visualized by considering the interpolation process as a series of two-dimensional interpolations, one for each plane. First, we consider the intersection of the light ray with the (x, y) plane to determine the distribution of the light ray on the nearest four microlenses. This is represented schematically in figure 14(a), where the green 'x' is the point where the projection strikes the microlens plane, the blue dots represent the centre of each microlens, and the shaded area enclosed by the dotted lines is the interpolation domain. In this representation, each ray is implicitly assumed to have a finite width equal to the size of one microlens, which is consistent with the physical function of the microlenses within the camera. The surrounding microlens positions are determined, in microlens coordinates, by using the floor and ceiling operators, where the subscript 0 is associated with the floor operator and the subscript 1 with the ceiling operator. This allows the relative position of the light ray to the neighbouring microlens centres to be easily calculated, and it has the benefit of auto-normalizing the coefficient since the separation is equal to one (i.e. $\mathrm{ceil}(x_2) - \mathrm{floor}(x_2) = 1$). Once the interpolation coefficients for the four microlenses have been calculated the u, v interpolation can take place. Figure 14(b) shows the discretization of the aperture plane as viewed from the pixel behind microlens (x_1, y_0). The green 'x' refers to where the projection strikes the aperture plane, in this case one of the designated plaid (u, v) values. The red dots represent the centres of each (u, v) location on the aperture. As with the (x, y) interpolation (u_2, v_2) is expressed in terms of pseudo-pixel coordinates using the floor/ceiling operators.

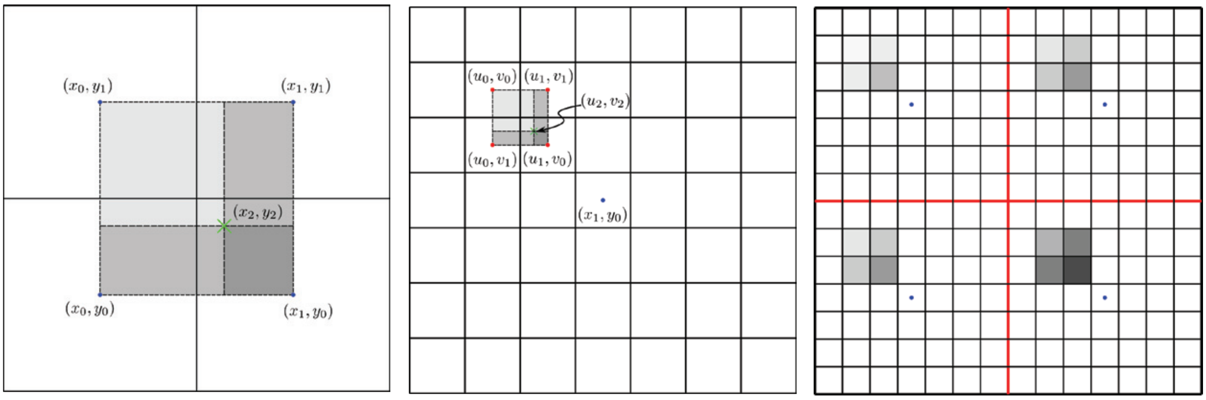

Figure 14. Determination of the weighting function coefficients via linear interpolation for the x, y plane (left), the u, v plane (center). The final result of the interpolation procedure is shown on the right.

Once the sixteen locations for which we need to calculate a coefficient have been found, the value of the coefficient must be determined. To do this we employ a simple linear interpolation scheme in which the coefficient is the combined value of the (x, y) and (u, v) interpolation steps. The distance from the $(0, 0)$ point in both interpolation schemes is all that is needed to calculate the coefficient. The relative distances, t, are given by

$$t_x = x_2 - x_0 \quad t_y = y_2 - y_0 \quad t_u = u_2 - u_0 \quad t_v = v_2 - v_0 \tag{11}$$

Using these and simple geometry the sixteen coefficients can be calculated. The interpolation coefficients, N_{xyuv}, have subscripts that represent their location relative to the voxel to be interpolated. For example, N_{0000} is the coefficient for point (x_0, y_0, u_0, v_0). The coefficients are calculated by using the normalized distances and are shown to be

$$\begin{aligned} N_{0000} &= (1 - t_x)(1 - t_y)(1 - t_u)(1 - t_v) \\ N_{0001} &= (1 - t_x)(1 - t_y)(1 - t_u)(t_v) \\ &\vdots \\ N_{1111} &= (t_x)(t_y)(t_u)(t_v). \end{aligned} \tag{12}$$

The result of this procedure can be seen in figure 14(c), where the red border represents the four microlenses shown in figure 14(a) with the (u, v) distribution behind it. The sixteen interpolation coefficients are shown as the shaded squares with intensity depending on their weight (white = 0, black = 1). In other words, figure 14(c) shows the relative distribution of intensity on the image sensor that results from a single light ray.

In the formulation of the weighting function no consideration was made for the aperture edges that are contained in the overall microlens image. Since the edge of the aperture image will fall partially on a pixel, its weight toward the reconstruction is diminished. To account for this a sequence of images are taken of a white background with the aperture open such that the intensity should be constant for all pixels under a microlens. These images are averaged together and normalized such that if the pixel falls completely inside the microlens image it yields a one, thus not affecting the weight, to zero if the pixel falls completely outside the microlens image. This is shown schematically in figure 15, where the green 'x' represents the centre of the microlens, and the green circle is the outer edge of the microlens image. Once the corrective image has been normalized it is multiplied by the weights to correct for the boundaries.

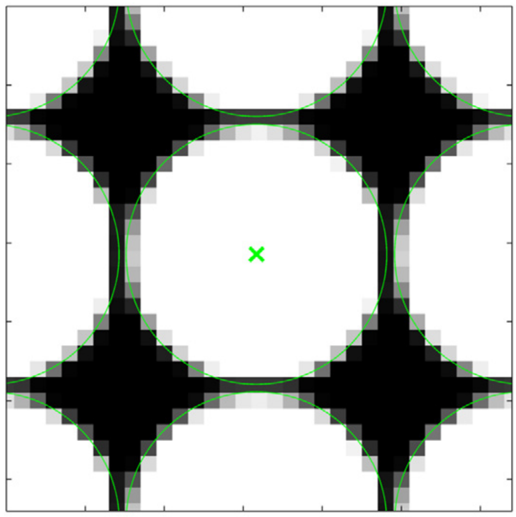

Figure 15. Schematic of the white image weighting correction.

The final step necessary for the calculation of the weighting function is to normalize the weights for each voxel by the sum of the weights for that voxel. This is done so that the intensity contained in a voxel is conserved. In equation form the normalization process is given by

$$\overline{w}_{i,j} = \frac{w_{i,j}}{\sum_i w_{i,j}}. \tag{13}$$

This forces the condition $\sum_i \overline{w}_{i,j} = 1$, such that all the light emanating from the voxel j must strike the image sensor.

To validate the weighting function, a comparison is drawn from that of a particle simulation. The particle simulator, which treats a particle as a point source of rays, simulated 400 particles distributed uniformly within the boundaries of a single voxel. This was determined to be the best comparison since the weighting coefficient should be representative of the entire voxel not just the centre. The simulation of 400 particles within a voxel does produce an accurate weighting function; however, it has considerable computational costs that prevent it from being used for the computation of the entire weighting matrix. As an illustration, consider that each voxel is simulated

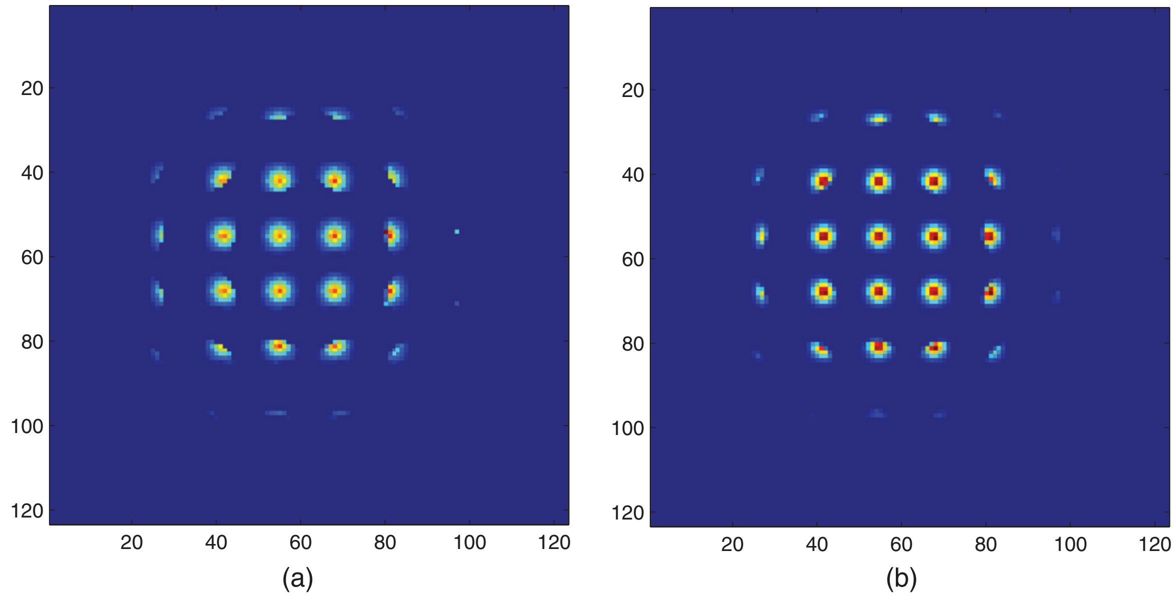

Figure 16. Weighting function comparison to particle simulation. (a) Weighting matrix. (b) Particle simulation.

in this way, using 10 000 rays for each particle such that the distribution at the image sensor is accurate and continuous. In contrast, the interpolation process uses 244 u and v values to represent the same data. Figure 16 shows both the weighting coefficients of the affected pixels as well as the particle simulation described previously. It can be seen that the weights are in fact representative of the particle simulation and are taken to be accurate, with the exception of some minor discrepancies. Specifically, the simulation produces higher weights in the centre of each microlens when compared to the weighting matrix. Differences can be seen along the boundaries as well with the weighting matrix showing contributions from pixels which are not illuminated by the simulation. This is attributed to the interpolation process blurring the ray across extra pixels.

5.3. Implementation and computational considerations

The implementation of the above mentioned MART algorithm was not as simple as the single equation seems. First, the weighting matrix is calculated on a per voxel basis, which is reversed from conventional tomography. This necessitates a pre-calculation of the summation term in the denominator of the MART equation, essentially adding an additional iteration of computational time. In conventional tomography the weighting matrix is stored on a per pixel basis making a summation over all voxels affected by a pixel straightforward. Compounding this, is the size of the weighting matrix. For a weighting matrix of size $300 \times 200 \times 200$ voxels the weighting matrix is 350 GB, storing only non-zero values. This makes storing the weighting matrix in memory impractical, therefore the data is stored on a hard disc in slices (1 slice per z location), in this case 200 slices, allowing for smaller chunks to be read into memory. The algorithm was implemented in $C++$, and uses binary files for faster processing. Using a 12 core workstation (Intel Xeon E5-2697 v2 at 2.7 GHz) with 64 GB of RAM and a RAID 0 array with 2 1 TB solid state disks, the weighting matrix takes approximately 1 h to complete and the MART algorithm takes approximately 30 min per image (running 50 images simultaneously takes approximately 1 day).

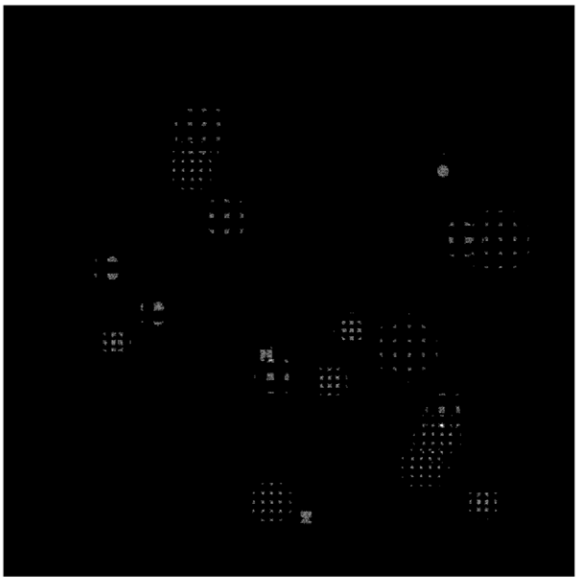

Figure 17. Synthetic raw image.

5.4. Sample reconstruction results

As a qualitative illustration of the capability of the reconstruction algorithm to reconstruct particles, a small group of particles were simulated using the aforementioned plenoptic simulator. For this exercise a smaller version of the prototype camera was used to cut down on computational time. Specifically, the synthetic camera has an image sensor of 850×850 pixels behind a 50×50 microlens array, and all other parameters were kept constant. Twenty particles were randomly generated inside a $5 \times 5 \times 5$ mm volume. The raw image is shown in figure 17.

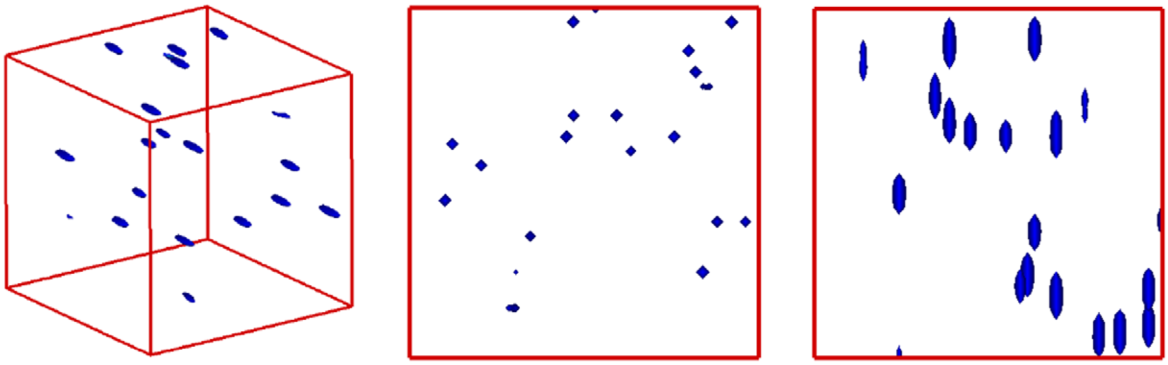

(a) (b) (c)

Figure 18. Tomographic reconstruction of a synthetic particle field. (a) Isometric view. (b) Front. (c) Top.

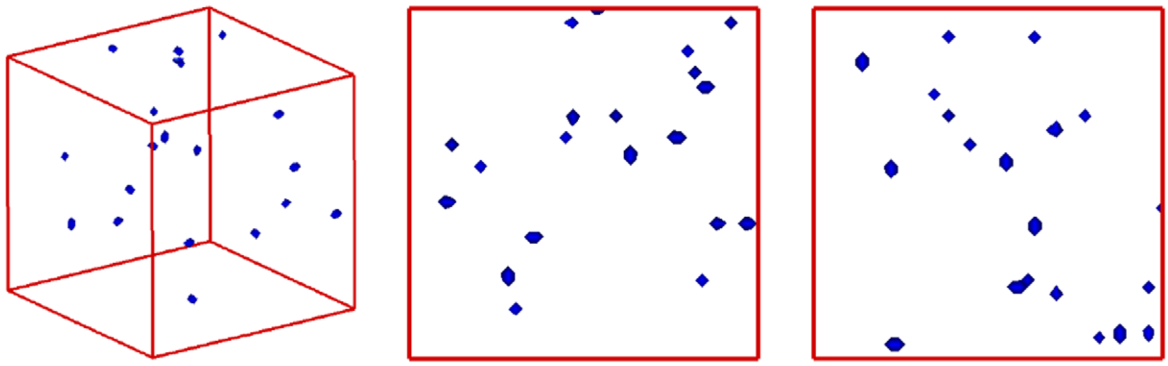

(a) (b) (c)

Figure 19. Actual particle positions. (a) Isometric view. (b) Front. (c) Top.

As a means of comparison a volume using the actual particle positions was generated using a $3 \times 3 \times 3$ voxel Gaussian blob fit to the particle positions. The final reconstruction of the particles is shown in figure 18 and the true particle positions are shown in figure 19. Figure 18(b) shows a front view of the reconstructed volume. When compared to the actual particle positions (figure 19(b)) the reconstructed particles are shown to match the actual particles in both size and location. Alternatively, when the reconstructed particles are compared to the actual particles in depth (figures 18(c) and 19(c)) they are shown to match locations, but the reconstructed particles are elongated in depth. This can be attributed to the limited range of angles that a plenoptic camera measures. Fortunately, the intensity in depth is not a constant. Figure 20 shows a single reconstructed particle iso-surface as well as a slice through the centre of the particle on the *YZ*-plane. The particle has a 'hot' centre with decreasing intensity at the front and back, as shown in figure 20(b). This allows for the resolution of the location of the centre of the particle in depth, where a constant intensity would create a large ambiguity. The lateral spatial resolution of this particle's reconstruction is limited to a single voxel. For other particles this may be four voxels or larger depending on their location spatially as well as in depth. In particular, the reconstruction of a particle far away from the focal plane is more elongated in depth and blurred spatially.

6. Experiments with synthetic images

In order to test the accuracy of the algorithm detailed above, we consider several cases starting with the best case scenario: a single particle. This test gives the upper limit of accuracy with the current reconstruction method and is useful in defining the accuracy as a function of depth. An extension to this test is the multiple particle test, where 500 particles are simulated inside a volume. Using the same metric as the single particle tests, the accuracy is determined in a non-ideal scenario (a random distribution in the presence of other particles). The final group of tests are full simulations, where the accuracy measured is in terms of the velocity, not particle position. These tests include a uniform flow field as well as a Gaussian ring vortex.

6.1. Single particle reconstructions

Using the synthetic image generation technique mentioned previously 40 particles are simulated (generating 40 different images) 1 mm (8 voxels) apart from each other in depth along the optical axis of the camera. The volume for each reconstruction was kept constant, such that the weighting matrix was the same for each reconstruction. The volume of size $6.125 \times 6.125 \times 50.125$ mm was discretized into a grid of $50 \times 50 \times 402$ voxels, creating cubic voxel elements with sides of length 0.125 mm. For the reconstruction, a relaxation

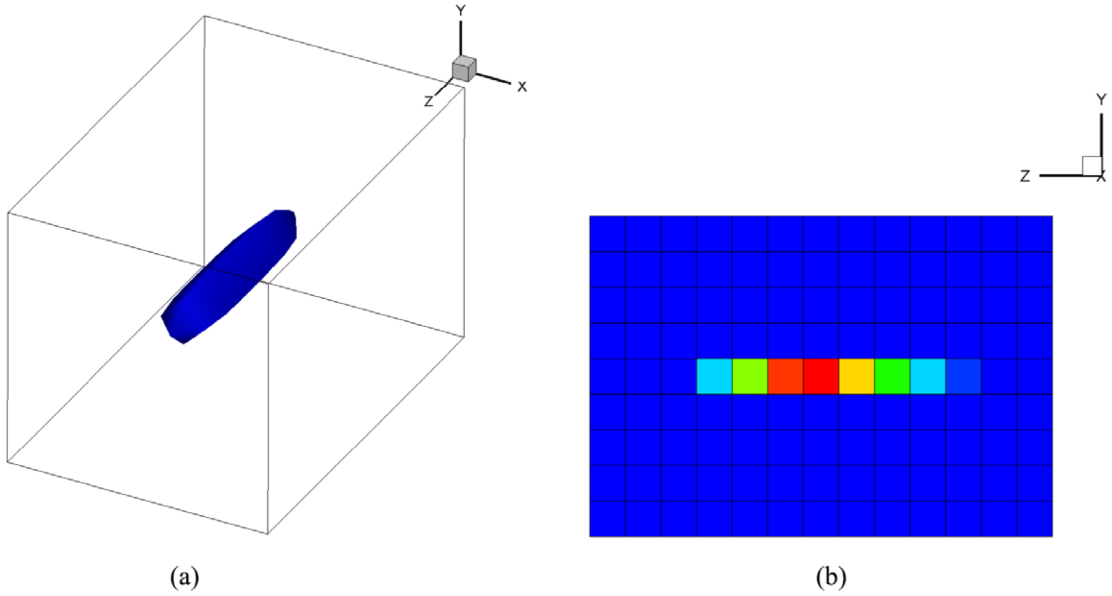

Figure 20. View of a single particle reconstruction. (a) 3D isometric view of a single particle. (b) Slice through center of particle.

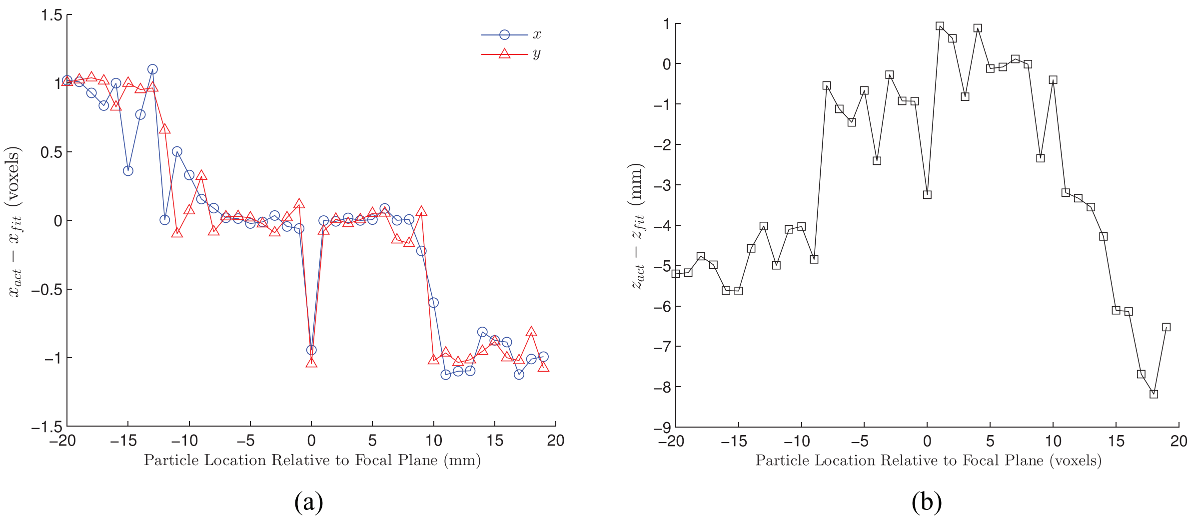

Figure 21. Error in reconstruction accuracy via the Gaussian fit of 40 particles spaced 1 mm apart along an optical axis. (a) *X* and *Y* (lateral). (b) *Z* (depth).

parameter of 0.5 was used and the MART algorithm was run for 5 iterations. Since the particle locations are known, the error in the reconstructed particles can be calculated. To precisely determine the particle location with sub-voxel accuracy, a 3D Gaussian function was fit to the reconstructed intensity data and the peak location was taken to be the location of the reconstructed particle. The results are shown in figure 21 with the absolute error (in voxels) on the *y*-axis and the relative position of the particle to the focal plane of the camera (100 mm away from the lens plane) on the *x*-axis. The results shown use a nominal magnification of −1, it is noted that the results will vary for other magnifications, although those are not considered in this work.

Figure 21(a) reveals the lateral accuracy of the algorithm as a function of depth for this optical configuration. In this case the particle position was perfectly aligned with a voxel, representing the best-case scenario. For the region near the focal plane [−10, 10], the error is minimal and nearly zero, with a notable exception being at the focal plane. This is due to ambiguity in a 1 mm region around the focal plane caused by the nominal depth of field of our camera. More specifically, in this region light emanating from a particle strikes a single microlens, whereas in other locations the light is spread across multiple microlenses. Thus, the algorithm does not have the information to 'interpolate' between microlenses. The MART algorithm spreads the intensity throughout this region, often leaving two peaks: one before and one after the focal plane. This results in the 1 voxel error shown. Further away from the focal plane the algorithm is shown to be less accurate, although the absolute error is only 1 voxel. There is some noticeable peak locking occurring causing the solution to be forced into a single voxel. The depth accuracy is shown in figure 21(b) as a function of depth. In the region near the focal plane [−10, 10] the error in depth was shown on average

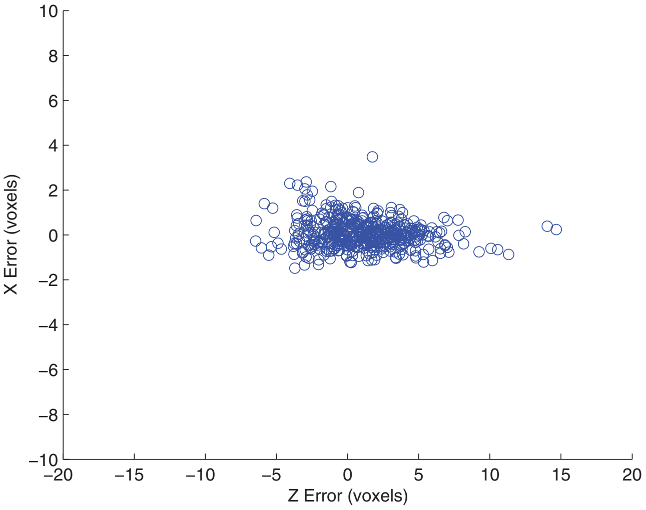

Figure 22. *X* and *Z* absolute errors in a reconstruction of 500 simulated particles.

to be 1 voxel, with a standard deviation of 1.5 voxels. Outside of this region the average error is five voxels. It is noted that the depth accuracy is worse than the spatial accuracy, as is to be expected.

An extension to the single particle test is to calculate the reconstruction error of multiple particles simultaneously. For this test a volume of size $30 \times 20 \times 20$ mm discretized into $300 \times 200 \times 200$ voxels was used. Inside the volume 500 particles were randomly positioned and an image was generated. This is still a relatively small particle density, although the purpose of this test is to obtain the accuracy of individual particles in the presence of additional particles. To determine the error in the reconstruction a sub-volume around the area of a known particle location was extracted (sub-volume was of size $6 \times 6 \times 30$ voxels), and fit with a Gaussian blob yielding the peak location, resulting in the absolute reconstruction error of the particles. A plot of the absolute X error versus the absolute Z error is shown in figure 22. The absolute error in X has a mean of 0.0658 voxels and a standard deviation of 0.7990 voxels. The absolute error in Z has a mean of 1.0392 voxels and a standard deviation of 2.9782 voxels. This is consistent with the single particle data in the range of depths used.

6.2. Full simulation reconstruction quality

For a dense simulation of particles calculating the particle position error for each individual particle is compromised due to the presence of other particles affecting the Gaussian fit. Therefore, to determine the accuracy of the reconstruction process, a statistical measure, known as the reconstruction quality factor, is used. This work utilizes the zero-mean reconstruction quality factor Q^* defined in La Foy and Vlachos [34], where the term zero-mean specifies that the volumes have a mean of zero, which is done by subtracting the mean from the original volume. They demonstrated that as the particle density increased the zero-mean reconstruction quality factor became a more accurate measure than the standard reconstruction quality factor defined in Elsinga *et al* [10]. The zero-mean quality factor is defined as:

$$Q^* = \frac{\sum \tilde{E}(x, y, z) \cdot \tilde{E}_0(x, y, z)}{\sqrt{\sum \tilde{E}(x, y, z)^2 \cdot \sum \tilde{E}_0(x, y, z)^2}} \tag{14}$$

where $\tilde{E}(x, y, z)$ and $\tilde{E}_0(x, y, z)$ are the zero-mean reconstructed intensity field and the zero mean exact intensity field, respectively. The exact intensity volume was created using a Gaussian blob consistent with the shape of an average single particle reconstruction.

The first study conducted was on the convergence of the MART algorithm using the plenoptic weighting function based on the residual after each iteration. The normalized residual, as shown in figure 23(a), is shown to decrease as the iterations increase and converge onto a solution. It is noted that the convergence of the residual does not guarantee that the algorithm converged onto the correct solution. In order to test the convergence onto the correct solution the reconstruction quality factor is used and is plotted against iterations, as shown in figure 23(b). The results of this study illustrate the notion that the algorithm converges towards the correct solution since the quality of the solution increases as the number of iterations increase. The solution is known to diverge with a high relaxation parameter, therefore a suitable ($\mu \sim 0.5$) relaxation parameter is used.

The next study conducted was on the effect of particle density on the reconstruction accuracy measured by the reconstruction quality factor. Conventionally, particle density is defined as the number of particles per pixel, although for a plenoptic camera the spatial resolution is governed by the microlens array, therefore the results are presented as number of particles per microlens (ppm). For completeness the results are also presented as particles per pixel (ppp). The focus of this section is to show the change in reconstruction quality as the particle density increases, therefore the results are normalized by the single particle reconstruction quality, $Q_0^* = 0.38$. The results are shown in figure 23(c), with the reconstruction quality on the y-axis and the particle density on the x-axis. The results show that there is little variance with respect to particle density until a ppm value of 1 or greater. As the particle density approaches 3 ppm the image becomes uniformly white and the algorithm has trouble producing accurate results.

6.3. Cross-correlation algorithm

This paper has, until now, discussed a method for obtaining particle fields from an image. However, the purpose of this technique is to obtain the velocity of the fluid being measured. To do this a method to extract the displacement, and therefore velocity, from the reconstructed particle fields is needed. Consequently, a method for obtaining the displacement was developed using a cross-correlation-based technique, whose implementation is based on Adrian and Westerweel [35] and Scarano and Riethmuller [36]. Briefly, each reconstructed volume pair is divided into several interrogation volumes defined by a size in the number of voxels and a percentage overlap. For each interrogation volume pair, a fast Fourier transform (FFT)-based cross-correlation is computed and the location of the maximum correlation peak is estimated by a

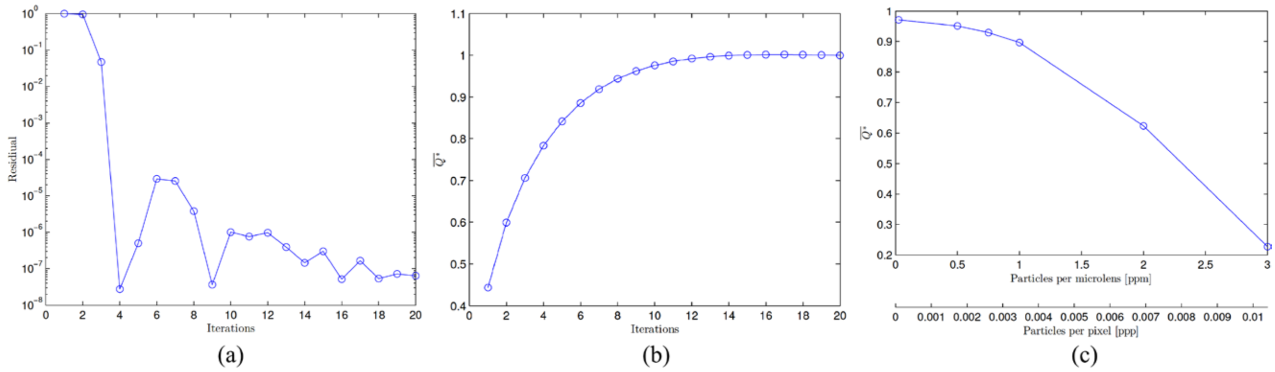

Figure 23. Reconstruction quality metrics, (a) normalized residual, (b) trend in quality factor as a function of iterations. Values are normalized by the result at iteration 20 ($Q_0^* = 0.36$), (c) quality factor as a function of particle density. Values are normalized by the quality of a single particle ($Q_0^* = 0.38$).

Gaussian peak fit to sub-pixel accuracy. From this location, as well as the time between exposures, the velocity can be calculated.

A more advanced version of this basic concept uses a multi-pass, multi-grid window deformation technique known as VODIM (Scarano and Poelma [15]). Each iteration begins by defining the interrogation volumes for cross-correlation, based on the sizes and overlap for that iteration. This allows for grid refinement in the later iterations. Next, the FFT-based cross-correlation is performed and the displacement for each interrogation volume is calculated. The displacements are then validated using a median test with the displacement data in a $5 \times 5 \times 5$ neighbourhood. If the displacement exceeds a pre-determined threshold (usually 2 [35]), the displacement is replaced by either a secondary peak or an interpolated value of the valid neighboring displacements. For subsequent iterations, the new interrogation volumes are displaced/deformed based on the displacements in the previous iteration. The deformation is calculated using a cardinal interpolation function on a $7 \times 7 \times 7$ stencil [37]. The final velocity is calculated as the location of the correlation peak plus the predicted displacement divided by the time between exposures.

6.4. Simulated Gaussian ring vortex

The final test for the accuracy of the reconstruction of a 3D particle field is to test the effect of the errors on the final velocity data. In order to test the velocity error a synthetic displacement is applied to a randomized particle field (0.5 ppm) and two synthetic images are acquired. The displacement field used is a Gaussian ring vortex. The vortex core's centre-line is aligned with the y-axis and is located at the centre of the volume forming a ring with an 8 mm diameter. The tangential displacement (in voxel coordinates) is given by

$$V_\theta = -\frac{\Gamma}{2\pi r}\left(1 - \mathrm{e}^{-\frac{r^2}{c_\theta}}\right) \tag{15}$$

where r is the distance from the particles location to the centre of the vortex. The other parameters Γ and c_θ are defined as

$$\Gamma = \frac{2\pi r_c^2}{v_{r_c}} \qquad c_\theta = \frac{r_c^2}{\gamma} \tag{16}$$

where r_c, the radius of the vortex, is 40 voxels, the tangential displacement at r_c, v_{r_c}, is 8 voxels, and γ is a constant equal to 1.256 431.

To compare the reconstructed velocity field, a synthetic particle field was generated using the actual positions of the particles with a $3 \times 3 \times 3$ Gaussian blob representing each particle. The synthetic volumes were then run through the same cross-correlation algorithms, thus providing an accurate baseline to test the reconstruction algorithm. Both volume pairs were run through the cross-correlation algorithm with final window sizes of $16 \times 16 \times 16$ voxels with 75% overlap. The results of the reconstructed velocity field are presented in figure 24 and the actual field in figure 25. Figure 24(a) shows a cross-section of the vortex in the XY plane at $Z = 100$ voxels (centre of the volume) with velocity vectors and contours of the z component of vorticity. Figure 24(b) shows a cross-section in the YZ plane at $X = 150$ voxels with a contour showing the X component of vorticity. When compared to the exact solution (figure 25) the solution matches well, but has some spurious vectors in the vortex ring (figure 24(a)), and some issues in capturing the motion in depth (figure 24(b)). The errors in depth are attributed to both the elongation of the particles, causing issues with the cross-correlation algorithm, as well as the increased variance in the reconstruction accuracy. The overall RMS error of the velocity field is 1.02 voxels, with each component having an RMS error of 0.165, 0.23 and 0.9833 for the u, v and w components, respectively.

7. Experimental assessment

To complement the synthetic image results, experiments were conducted with the plenoptic camera in Auburn University's water tunnel described in [38]. The test section is made out of acrylic with dimensions of $0.61 \times 0.61 \times 2.44$ m. The flow under consideration was that of the wake behind a cylinder at $\mathrm{Re_D} = 185$, with a cylinder of diameter 4.7625 mm and a freestream velocity of 0.039 m s^{-1}. For PIV measurements the flow was seeded with silver-coated hollow glass spheres with a mean diameter of 10 μm. Illumination was provided by a New Wave solo III Nd-Yag laser system with a maximum output energy of 50 mJ per pulse at 532 nm and a pulse duration of 3–5 ns.

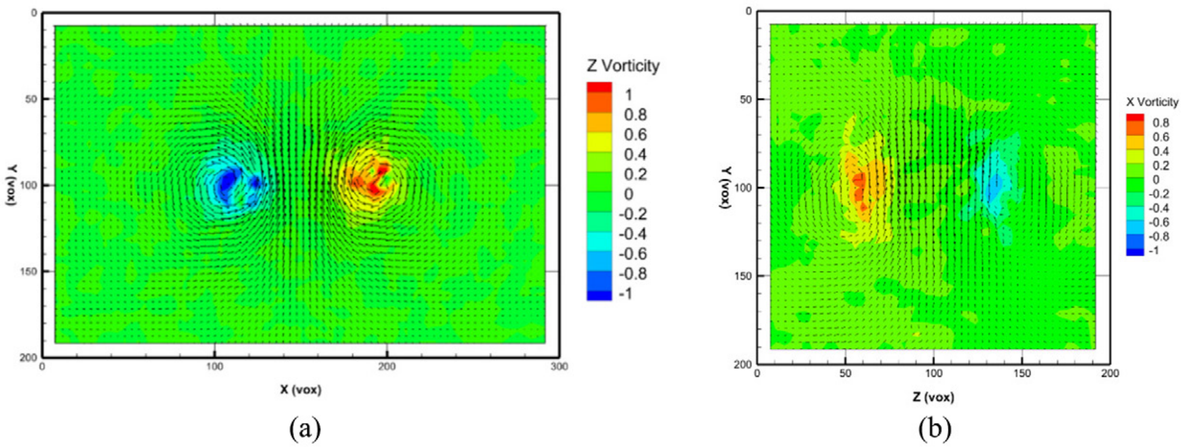

Figure 24. Cross-sections of reconstructed Gaussian ring vortex. (a) Cross section of velocity in *XY* plane at $Z = 100$ voxels. Contours show the vorticity in the z-direction. (b) Cross section of velocity in *YZ* plane at $X = 150$ voxels. Contours show the vorticity in the x-direction.

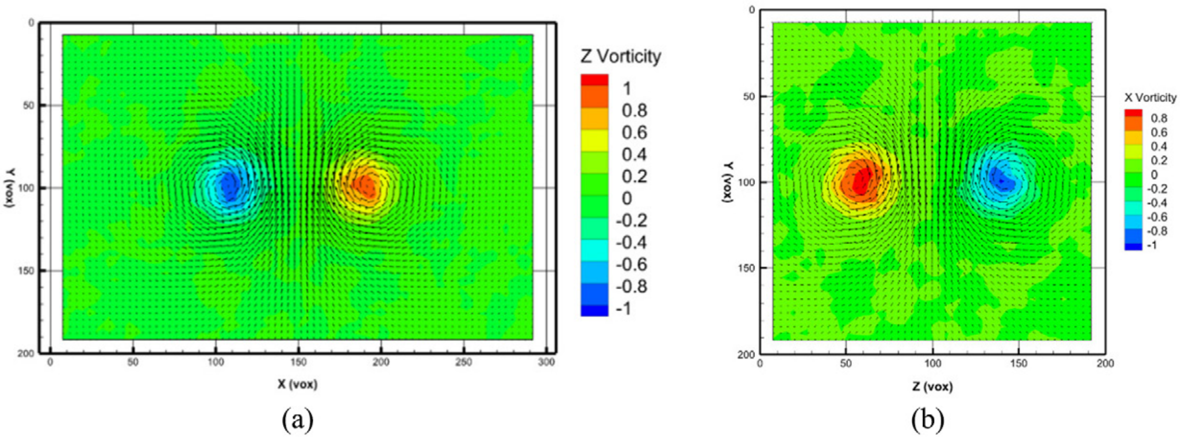

Figure 25. Cross-sections of simulated Gaussian ring vortex. (a) Cross section of velocity in *XY* plane at $Z = 100$ voxels. Contours show the vorticity in the z-direction. (b) Cross section of velocity in *YZ* plane at $X = 150$ voxels. Contours show the vorticity in the x-direction.

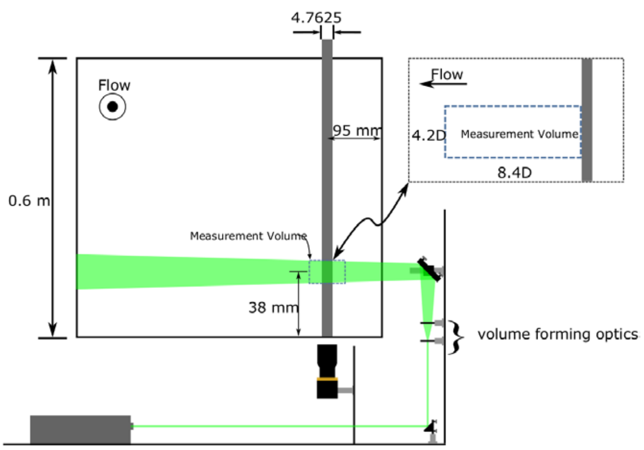

Figure 26. Experimental arrangement: camera mounted beneath the tunnel such that vortex shedding occurs in the *XY* plane with laser volume illumination entering from the side of the wind tunnel.

The experimental arrangement, as shown in figure 26, allows the vortex shedding in the wake of the cylinder to occur in the *XY* plane yielding a best case scenario for the plenoptic system, since the motion in the third dimension (Z) should be minimal. The volume of interest spanned $40 \times 26.7 \times 20$ mm and was discretized into $300 \times 200 \times 150$ voxels. A time separation of 27 msec was used to generate an 8-voxel displacement based on the freestream velocity. Using a final cross-correlation window size of $32 \times 32 \times 32$ ($4.3 \times 4.3 \times 4.3$ mm) with a vector spacing of $16 \times 16 \times 16$ created $34 \times 22 \times 15$ vectors. 350 vector fields were acquired and processed for statistical measures. It is noted that to achieve the desired field of view for the plenoptic system, 1 : 1 magnification using a 60 mm macro lens, the centre of the measurement volume was located just 38 mm above the bottom of the tunnel floor such that the wall boundary layer influenced the measurements. Longer working distances while maintaining 1 : 1 imaging can be achieved using longer focal length $f/2$ lenses, but this was not pursued in the current work.

For comparison, planar PIV measurements were taken in a similar configuration to that shown in figure 26. In this case, an Imperx Bobcat B4020 11 Mp (4032×2688 pixels) CCD camera was used to image the wake of the cylinder. Utilizing a 60 mm macro lens at a magnification of 0.37, the field of view (FOV) of the 2D measurement system was 99 × 66 mm. The laser-pulse time separation was 5 msec creating an 8-pixel displacement in the free stream and a total of 500 image-pairs were recorded. The velocity field was processed using a 2D version of the cross-correlation process described earlier with a final window size of 32 × 32 (0.8 × 0.8 mm) with a 16 pixel separation resulting in 218 × 165 vectors.

An example of an instantaneous 3D/3C velocity measurement for plenoptic PIV is shown in figure 27 (top and centre). The top and middle figure show the streamwise (x) and transverse (y) components of velocity along 5 streamwise-transverse (x–y) planes extracted from the full volume of data. The z-axis is stretched to facilitate viewing of each of the 5 planes. The cylinder is located along the z-axis at $Y/D = 0$ and the x-axis is aligned with the flow direction. The increased velocity of the flow around the top and bottom of the cylinder along with the reduced velocity in the wake is clearly present in all the slices. In addition, vortex shedding in the form of a Karman vortex street is also clearly evident. This is more easily visualized in figure 27, bottom, which shows 3D iso-surfaces of constant vorticity with red showing positive vorticity and blue showing negative vorticity. One observation from the 3D measurement is that although the shedding extends across the entire depth of the volume, it is not uniform in the spanwise direction (i.e. phase of shedding varies in the spanwise direction). Rather, the vortices display a slight tilt relative to the cylinder. This observation is consistent with observations made in Williamson [39], where they noted that the transition from a 2D to a 3D unsteady wake occurred around $\mathrm{Re_D} = 150$. In addition, the influence of the boundary layer on the shedding could be a source of three-dimensionality.

For comparison to the planar PIV data, 2D slices were extracted from the volumetric data along the same 2D plane as the 2D PIV measurements and are displayed in figure 28. The top row of figure 28 is the ensemble averaged 2D planar PIV measurements trimmed down to the field of view of the plenoptic cameras. The bottom row shows the same measurement region, but is extracted from the 3D plenoptic data. The first two columns show the average of 500 and 350 vector fields, respectively, with the third column showing an instantaneous velocity field (the same as in figure 27 for the plenoptic data). Qualitatively, the average velocity field is in very close agreement with the planar data in structure. For the u component, the acceleration of the flow around the cylinder as well as the wake deficit are clearly visible. In the v component an asymmetry of the velocity in the wake is shown for both cases, an effect which might be due to the proximity of the cylinder to the tunnel wall. The Karman vortex street is clearly visible in the instantaneous data with similar characteristics as the planar data, clearly showing the ability of plenoptic PIV to produce comparable measurements to traditional 2D PIV.

Quantitatively there are some discrepancies between the 2D and 3D techniques. For the u component of velocity the

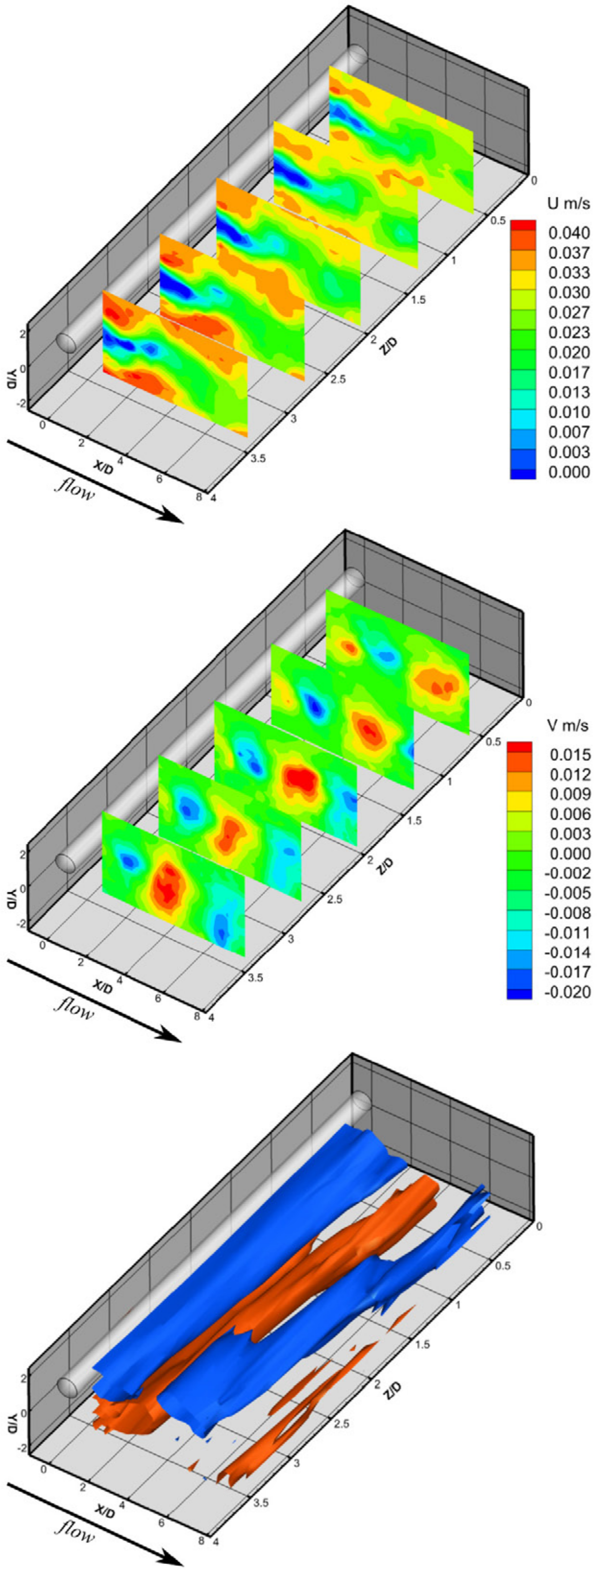

Figure 27. (top) 3D instantaneous u-velocity distribution showing five u velocity contours. (centre) 3D instantaneous u-velocity distribution showing five v velocity contours. (bottom) 3D instantaneous z-vorticity distribution blue = $-0.17\ s^{-1}$, red = $0.17\ s^{-1}$. Note: The z, or depth, direction is stretched to make the figure easily readable.

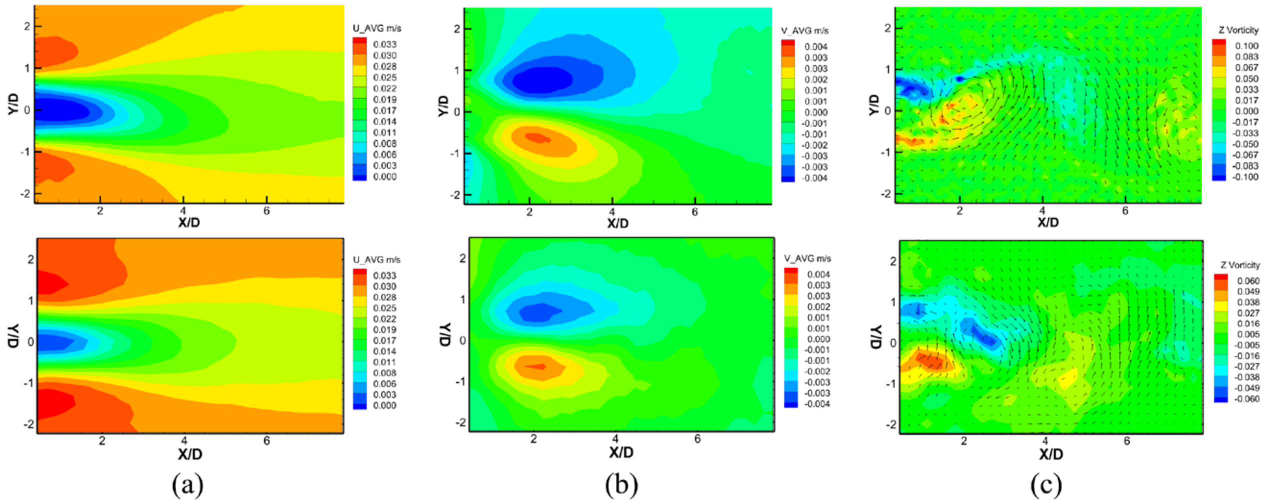

Figure 28. The top row contains 2D planar PIV data, the bottom row 2D slices of 3D data, taken at centre plane, showing the average velocity field as well as instantaneous velocity vectors with z-vorticity contour. (a) $\bar{u}$ contour. (b) $\bar{v}$ contour. (c) ω_z contour, vectors $(u - \bar{u}, v)$.

planar data shows the characteristic velocity deficit associated with the wake region extending further than in the plenoptic data. In the planar data, the location where the velocity has recovered to half of the freestream (0.015 m s^{-1}) is located at $X/D = 3.75$, whereas it is located at $X/D = 2.9$ in the plenoptic data. In the acceleration region around the cylinder, the plenoptic measurements have a peak velocity of 0.033 m s^{-1}, whereas the planar datas peak velocity in the same region is 0.029 m s^{-1}. The region of positive v component in the wake of the cylinder extends from $X/D = 1$ to $X/D = 4.5$ for both cases, with the 2D case being noticeably smoother. For the region of negative v velocity, located above the centre-line (positive Y/D), the magnitudes are larger (0.0047 m s^{-1} compared to 0.0034 m s^{-1}) than those below the centre-line for the 2D case. This is also the case in the plenoptic data (0.0037 m s^{-1} compared to 0.0034 m s^{-1}), although the magnitudes are different. Both data sets show the line of symmetry in the negative y-direction, and the region of negative v velocity being larger than the positive, although the planar data shows a larger region. Measuring from half the peak height before and after the peak the region of negative v velocity extends from $X/D = 1$ to $X/D = 5$, whereas the plenoptic data show the same region extending from $X/D = 1$ to $X/D = 4.2$. In the instantaneous data, the overall structures look similar, but the 2D data shows finer detail of the vortical structures.

Further insight can be gained by extracting a single line of data (at $X/D = 2$) from both the planar and plenoptic PIV data. Statistical data is extracted in the form of the average velocity and the fluctuating component of the velocity field and the results are shown in figure 29. The top left-hand figure shows the average streamwise component of velocity. Here the plenoptic PIV data match the main trends of the planar data but with a higher velocity throughout the profile. In the acceleration region around the top and bottom of the cylinder the difference is approximately 0.0025 m s^{-1} corresponding to a 0.5 voxel average displacement error. In the wake, the error is approximately 0.005 m s^{-1} or a 1 voxel displacement error. Some sources of this error can be attributed to misalignment between the two independent measurement systems, variations in the day-to-day operation of the tunnel, as well as a spatial averaging of the velocity field due to the lower spatial resolution of the plenoptic system. In the plenoptic data, each measurement corresponds to an $4.3 \times 4.3 \times 4.3$ mm^3 volume, whereas for the planar PIV data, it is $0.8 \times 0.8 \times 1.0$ mm^3. A more likely source of the error, however, is attributed to image distortions caused by the use of a real lens, which can cause pincushion and barrel-type distortions in conventional imaging systems, as well as optical imperfections in the acrylic water tunnel test section walls. A robust 3D calibration procedure that accounts for these effects within the image processing framework of a plenoptic camera is currently under development. For the streamwise velocity fluctuations (top right) the plenoptic data displays the same trends as the planar PIV data, but underestimates the magnitude of the fluctuations. This underestimation of the velocity fluctuations is also consistent with a spatial averaging effect due to the larger physical size of the correlation windows used to process the plenoptic PIV data.

Similar observations are made in the evaluating the transverse component of velocity, although the data appear to provide a better overall match. In the average profile (centre left) the results show very close agreement, with slight errors, approximately 8.6×10^{-4} m s^{-1} corresponding to a 0.17 voxel displacement, at $Y/D > 1$. The transverse fluctuating quantities (centre right) also show close agreement with errors of approximately 1×10^{-3} m s^{-1} (approximately 0.2 voxels) starting at $Y/D = 1$.

Although there are no comparison data, the out-of-plane component of velocity measured by the plenoptic camera is shown in the bottom row of figure 29. In the average velocity field (bottom left) the flow field is shown to be moving towards the wall in the wake region and away from the wall in the freestream. The velocity fluctuations in the z-direction are fairly uniform throughout the volume with a magnitude on

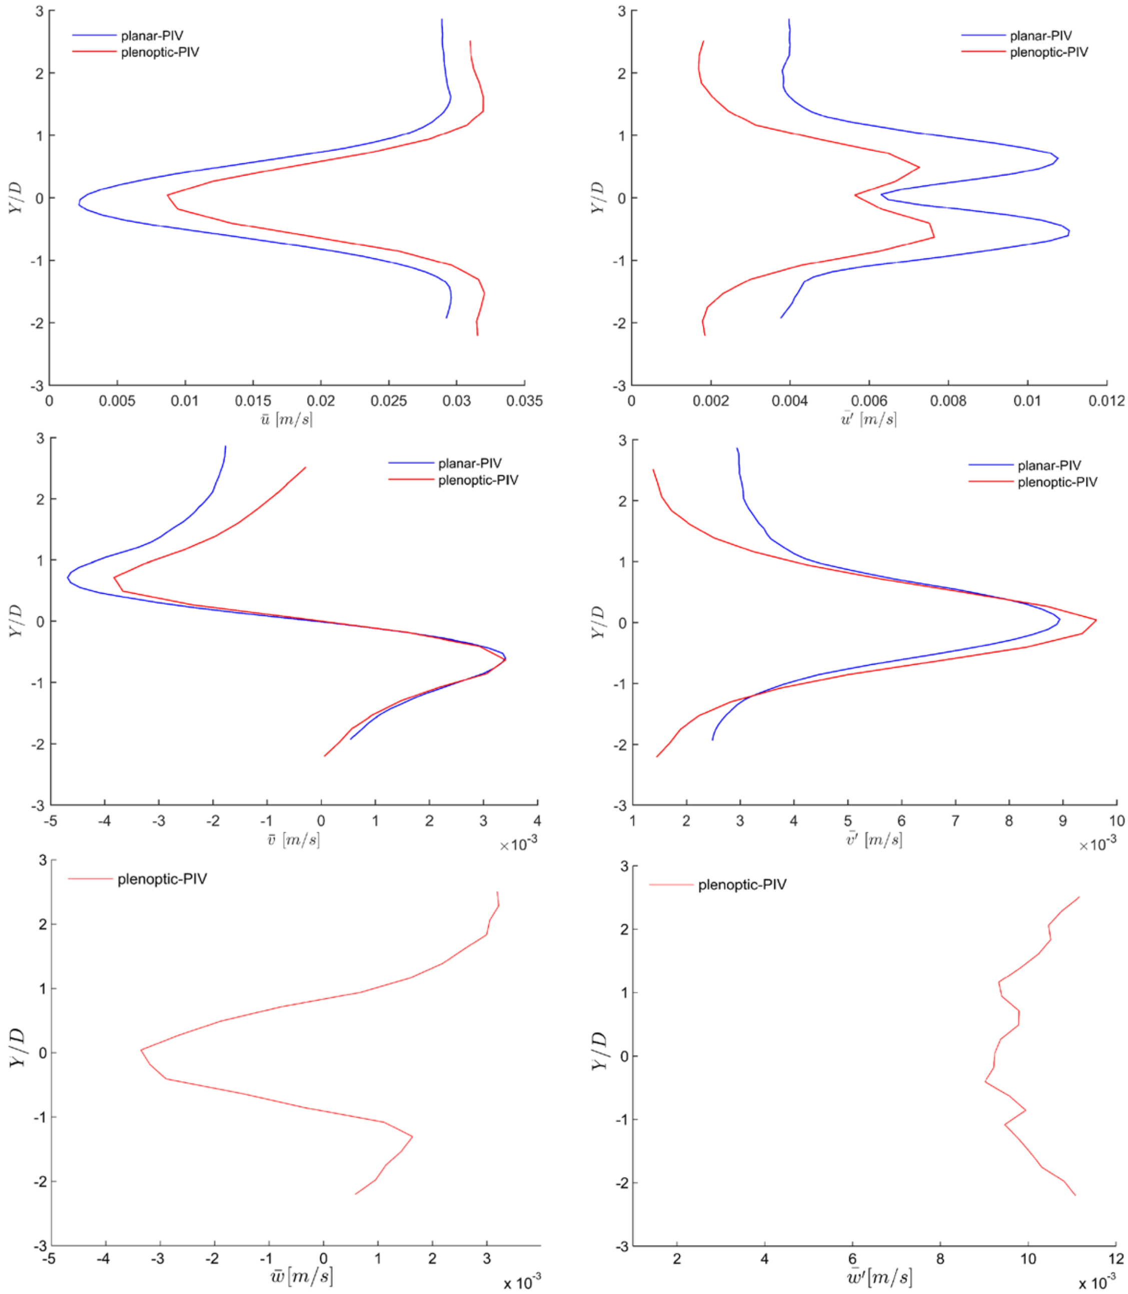

Figure 29. Line of data extracted at $X/D = 2$ for both planar and plenoptic PIV. The left column is the average velocity field, and the right column contains the average velocity fluctuations. The top row is the u component, the centre row is the v component of velocity, and the bottom row is the w component of velocity.

the order of 0.01 m s^{-1}, which corresponds to approximately 2 voxels of displacement.

8. Conclusion

A new 3D/3C PIV measurement technique based on light field imaging with a plenoptic camera (i.e. plenoptic PIV) has been presented. Plenoptic cameras are most well known for their ability to computationally refocus or change the perspective of an image after it has been acquired. These capabilities are enabled by the unique manner in which incident light rays are densely sampled by the camera as both the position and angle of the light rays are resolved through the placement of a microlens array near a conventional 2D image sensor. In the present work, the MART algorithm was adapted for use with plenoptic image data in order to reconstruct a volume of particles. The 3D/3C velocity field is then estimated from a pair of volume reconstructions using cross-correlation. The validity of this approach was shown using both synthetic data and experimental data acquired with a home-built plenoptic camera.

Overall, the main strength of plenoptic PIV is that it is a single camera technique that operates in a very similar fashion (e.g. same light source, timing etc) as conventional 2D/2C PIV while enabling 3D/3C velocity measurements. In fact, the home-built plenoptic camera used in this work was a modified version of an interline CCD camera commonly supplied with commercial PIV systems with the result that the frame rate, minimum inter-pulse timing and triggering of the camera are unaffected by the modification. Similar modifications can be made to other cameras such that higher resolutions or framing rates can be expected in the future. In fact, we have already built several higher resolution (by a factor of 3) plenoptic cameras that are based on a 29 megapixel image sensor and hexagonally arranged microlens array. The single camera nature of the technique offers a significant advantage over multi-camera methods in that the experimental arrangement is simplified and potentially more cost effective. Perhaps, and more importantly, the technique requires less optical access allowing for applications in facilities where optical access is insufficient for the application of other methods.

The two main trade-offs associated with plenoptic cameras are that (1) spatial resolution is sacrificed for angular resolution and (2) the resolution in the depth direction is worse than the two lateral directions. For the prototype plenoptic camera design considered here, the spatial resolution of the light field measurement is characterized by the microlens pitch, while the angular resolution is determined by the number of pixels located behind each microlens. In this work, an array of 290×194 microlenses was coupled with a 16-megapixel image sensor yielding an overall light field resolution of $290 \times 194 \times 16 \times 16$. Under nominal 1 : 1 magnification imaging, a volume size of $300 \times 200 \times 200$ voxels ($37.5 \times 25.0 \times 25.0$ mm) was reconstructed from synthetic data. The synthetic data showed the potential of MART to resolve particle locations to better than 1 voxel (0.125 mm) in the two lateral directions and 3 voxels (0.375 voxels) in the depth direction. Particle displacements calculated using a $16 \times 16 \times 16$ voxel cross-correlation window were accurate to within 0.2 and 1.0 voxels, respectively. The experimental data obtained in the wake of a cylinder clearly demonstrated the ability of the technique to be used for 3D flow visualization, although further work is needed to better quantify the experimental uncertainty. The 3D features of these flows were clearly identified and comparisons with 2D PIV data suggest that the synthetic image data offer a reasonable, although idealized, approximation of the measurement uncertainty.

In the near future, several steps can be taken to further develop and improve plenoptic PIV, such as designing higher resolution plenoptic cameras. While higher resolution cameras will improve overall resolution, the disparity between lateral and depth resolution is likely to remain as it is primarily limited by the size of the aperture of the main imaging lens, which forms the baseline for the 3D estimation and should be as large as possible. In this respect, the development of a two camera plenoptic PIV system has significant potential to improve depth resolution albeit at the expense of added experimental complexity. Beyond improvements in hardware, the continued development of volume reconstruction algorithms and the associated calibration procedure are likely to improve the overall accuracy of the technique. In particular, the development of a robust 3D calibration procedure that more accurately models the mapping of object space to image space and accounts for the possible misalignment of the microlens array position relative to the image sensor as well as the development of more computationally efficient reconstruction algorithms for the processing of large data sets are needed.

Looking more broadly, plenoptic imaging has the potential to be adapted for 3D variations of a wide variety of optical flow diagnostics including non-particle-based measurements such as laser-induced fluorescence and background-oriented Schlieren imaging. In general, most multi-camera techniques represent a sparse angular sampling of the light field such that specific assumptions (e.g. sparse particle field) are needed in order to produce a 3D reconstruction of the volume of interest. Due to these assumptions, these methods cannot easily be extended to provide equivalent 3D information using other types of measurement. Plenoptic cameras, on the other hand, provide a dense sampling of the angular space that allows for the development and implementation of novel algorithms that are not subject to the same constraints and can thus be more easily adapted to provide 3D information for other types of measurements. Computational refocusing is a clear example of one such algorithm that might be exploited for such a purpose such that plenoptic cameras are likely to find a multitude of applications in the future.

Acknowledgments

This work has been supported through funding provided by the Air Force Office of Scientific Research specifically grant FA9550-100100576 (program manager: Dr D Smith). The authors would like to gratefully acknowledge M Levoy from Stanford University for permission to use a template for manufacturing our microlens array, and for a variety of helpful discussions. Additionally, the authors thank S Reeves from Auburn University for continued discussions that have led us to consider the direct tomographic approach.

References

[1] Arroyo M and Greated C 1991 Stereoscopic particle image velocimetry *Meas. Sci. Technol.* **2** 1181–86
[2] Willert C and Gharib M 1992 Three-dimensional particle imaging with a single camera *Exp. Fluids* **12** 353–8
[3] Kahler C J and Kompenhans J 2000 Fundamentals of multiple plane stereo particle image velocimetry *Exp. Fluids* **29** S070–77
[4] Brucker C 1995 Digital-Particle-Image-Velocimetry (DPIV) in a scanning light sheet 3-D starting flow around a short cylinder *Exp. Fluids* **19** 255–63
[5] Lynch K P and Thurow B S 2012 3-D flow visualization of axisymmetric jets at Reynolds number 6,700 and 10,200 *J. Vis.* **15** 309–19
[6] Thurow B, Lynch K, Williams S and Melnick M 2010 3D flow imaging using a MHz rate puls burst Laser system *15th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics (Lisbon, Portugal)*
[7] Pereira F, Gharib M, Dabiri D and Modarress D 2000 Defocusing digital particle image velocimetry: a 3-component 3-dimensional DPIV measurement technique. Application to bubbly flows *Exp. Fluids* **29** S078–84

[8] Hinsch K D 2002 Holographic particle image velocimetry *Meas. Sci. Technol.* **13** R61–72
[9] Herrmann S F and Hinsch K D 2004 Light-in-flight holographic particle image velocimetry for wind-tunnel applications *Meas. Sci. Technol.* **15** 1–9
[10] Elsinga G E, Scarano F, Wieneke B, Oudheusden B W and Van Oudheusden B W 2006 Tomographic particle image velocimetry *Exp. Fluids* **41** 933–47
[11] Belden J, Truscott T T, Axiak M C and Techet A H 2010 Three-dimensional synthetic aperture particle image velocimetry *Meas. Sci. Technol.* **21** 125403
[12] Sheng J, Malkiel E and Katz J 2006 Digital holographic microscope for measuring three-dimensional particle distributions and motions *Appl. Opt.* **45** 3893–901
[13] Scarano F 2013 Tomographic PIV: principles and practice *Meas. Sci. Technol.* **24** 012001
[14] Schröder A, Geisler R, Elsinga G E, Scarano F and Dierksheide U 2008 Investigation of a turbulent spot and a tripped turbulent boundary layer flow using time-resolved tomographic PIV *Exp. Fluids* **44** 305–16
[15] Scarano F and Poelma C 2009 Three-dimensional vorticity patterns of cylinder wakes *Exp. Fluids* **47** 69–83
[16] Humble R A, Elsinga G E, Scarano F and van Oudheusden B W 2009 Three-dimensional instantaneous structure of a shock wave/turbulent boundary layer interaction *J. Fluid Mech.* **622** 33
[17] Violato D and Scarano F 2011 Three-dimensional evolution of flow structures in transitional circular and chevron jets *Phys. Fluids* **23** 124104
[18] Adelson E H and Bergen J 1991 The Plenoptic Function and the Elements of Early Vision *Computational Models of Visual Processing* (Cambridge, MA: MIT Press) pp 3–20 (preprint 10.1.1.2.9848)
[19] Adelson E H and Wang J Y A 1992 Single lens stereo with a plenoptic camera *IEEE Trans. Pattern Anal. Mach. Intell.* **14** 99–106
[20] Levoy M and Hanrahan P 1996 Light Field Rendering *Proc. 23rd Annual Conf. on Computer Graphics and Interactive Techniques—SIGGRAPH* pp 31–42
[21] Levoy M 2006 Light Fields and Computational Imaging *Computer* **39** 46–55
[22] Ng R, Levoy M, Duval G, Horowitz M and Hanrahan P 2005 Light field photography with a hand-held plenoptic camera *Informational* 1–11
[23] Lumsdaine A and Georgiev T 2009 The focused pleoptic camera *IEEE Int. Conf. on Computational Photography*
[24] Lippman G 1908 Epreuves revesibles donnant la Sensation du relief *J. Phys.* **7** 821–5
[25] Ng R 2006 Digital light field photography *PhD Thesis* Stanford University, CA
[26] Ng R 2005 Fourier slice photography *ACM Trans. Graph.* **24** 735
[27] Levoy M, Ng R, Adams A, Footer M and Horowitz M 2006 Light field microscopy *ACM Trans. Graph.* **25** 924
[28] Levoy M, Zhang Z and McDowall I 2009 Recording and controlling the 4D light field in a microscope using microlens arrays *J. Microsc.* **235** 144–62
[29] Gerrard A and Burch J 1975 *Introduction to Matrix Methods in Optics* (New York: Wiley)
[30] Georgiev T and Intwala C 2008 Light field camera design for integral view photography *Technical Report* unpublished
[31] Lynch K 2011 Development of a 3D fluid velocimetry technique based on light field imaging *PhD Thesis* Auburn University
[32] Herman G T and Lent A 1976 Iterative reconstruction algorithms *Comput. Biol. Med.* **6** 273–94
[33] Fahringer T W and Thurow B 2012 Tomographic reconstruction of a 3D flow field using a plenoptic camera *42nd AIAA Fluid Dynamics Conf.*
[34] La Foy R and Vlachos P 2013 *10th Int. Symp. on Particle Image Velocimetry (Delft, The Netherlands, 1–3 July 2013)*
[35] Adrian R J and Westerweel J 2011 *Particle Image Velocimetry* (Cambridge: Cambridge University Press)
[36] Scarano F and Riethmuller M L 1999 Iterative multigrid approach in PIV image processing with discrete window offset *Exp. Fluids* **26** 513–23
[37] Scarano F 2001 Iterative image deformation methods in PIV *Meas. Sci. Technol.* **13** R1–19
[38] Rifki R, Khan M, Anwar A and Zafar B 2005 Effect of aspect ratio on flow field of surface-mounted obstacles *Applied Aerodynamics Conf. (Toronto, Ontario: AIAA)*
[39] Williamson C H K 1996 Vortex dynamics in the cylinder wake *Annu. Rev. Fluid Mech.* **28** 477–539