Meas. Sci. Technol. **21** (2010) 125403 (21pp) doi:10.1088/0957-0233/21/12/125403

Three-dimensional synthetic aperture particle image velocimetry

Jesse Belden[1], Tadd T Truscott[2], Michael C Axiak[1] and Alexandra H Techet[1]

[1] Department of Mechanical Engineering, MIT, Cambridge, MA, USA
[2] Department of Mechanical Engineering, BYU, Provo, UT, USA

E-mail: jbelden@mit.edu

Received 9 June 2010, in final form 5 September 2010
Published 15 November 2010
Online at stacks.iop.org/MST/21/125403

Abstract
We present a new method for resolving three-dimensional (3D) fluid velocity fields using a technique called synthetic aperture particle image velocimetry (SAPIV). By fusing methods from the imaging community pertaining to *light field imaging* with concepts that drive experimental fluid mechanics, SAPIV overcomes many of the inherent challenges of 3D particle image velocimetry (3D PIV). This method offers the ability to digitally refocus a 3D flow field at arbitrary focal planes throughout a volume. The viewable out-of-plane dimension (Z) can be on the same order as the viewable in-plane dimensions (X–Y), and these dimensions can be scaled from tens to hundreds of millimeters. Furthermore, the digital refocusing provides the ability to 'see-through' partial occlusions, enabling measurements in densely seeded volumes. The advantages are achieved using a camera array (typically at least five cameras) to image the seeded fluid volume. The theoretical limits on refocused plane spacing and viewable depth are derived and explored as a function of camera optics and spacing of the array. A geometric optics model and simulated PIV images are used to investigate system performance for various camera layouts, measurement volume sizes and seeding density; performance is quantified by the ability to reconstruct the 3D intensity field, and resolve 3D vector fields in densely seeded simulated flows. SAPIV shows the ability to reconstruct fields with high seeding density and large volume size. Finally, results from an experimental implementation of SAPIV using a low cost eight-camera array to study a vortex ring in a $65 \times 40 \times 32$ mm^3 volume are presented. The 3D PIV results are compared with 2D PIV data to demonstrate the capability of the 3D SAPIV technique.

Keywords: synthetic aperture refocusing, 3D particle image velocimetry, light field imaging, volume reconstruction, camera array, geometric optics, vortex ring

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Efforts for resolving three-dimensional velocity fields are justified by the need to experimentally resolve flows that are highly three-dimensional and to validate numerical simulations of complex flows. The ability to spatio-temporally resolve flow features from small to large scales in arbitrarily large volumes is the goal of any 3D PIV system. Of course, there have been many roadblocks to achieving all of these goals with a single system, and compromises must be made. Two-dimensional particle image velocimetry (2D PIV) is the most pervasive method for resolving velocity fields, thus it is not surprising that recent efforts to resolve 3D flow fields have extended many of the fundamentals of 2D PIV to the third dimension.

Several methods exist for resolving 3D particle fields, or any 3D scenes for that matter, but the methods of data acquisition seem to fall into three broad categories: multiple-viewpoints, holography and internal optics alteration. The technique described herein falls into the multiple-viewpoint category and makes use of an algorithm known as *synthetic aperture refocusing* to examine the imaged volume and is thus referred to as synthetic aperture PIV (SAPIV). Herein, we focus on the application of the principles of synthetic aperture imaging to develop a measurement system for resolving three-dimensional fluid velocity fields. The system performance is evaluated theoretically and numerically. The practical utility of SAPIV is demonstrated through an experimental study of a canonical vortex ring.

The evolution of 2D PIV is described by Adrian [1] and is not reviewed here. One of the earliest, but still frequently utilized, methods for 3D PIV is two camera stereoscopic PIV, which is primarily used to resolve the third component of velocity within a thin light sheet [2]. Maas *et al* [3, 4] and Malik *et al* [5] use a 3D particle tracking velocimetry (PTV) method which resolves the location of individual particles imaged by two, three or four cameras in a stereoscopic configuration. They report measurements in a large volume (e.g. 200 $\times$ 160 $\times$ 50 mm^3), but with very low seeding density ($\approx$1000 particles). Through precise calibration and knowledge of the imaging geometry, the particle field can be reconstructed. More recently, improvements to PTV methods are presented by Willneff and Gruen [6]. In general, low seeding density is a typical limitation of PTV, yielding low spatial resolution in the vector fields.

Another technique which makes use of multiple viewpoints is defocusing digital particle image velocimetry (DDPIV) [7–9]. In theory, DDPIV capitalizes on the defocus blur of particles by placing an aperture with a defined pattern (usually pinholes arranged as an equilateral triangle) before the lens, which is a form of coded aperture imaging [10]. The spread between three points generated by imaging a single particle corresponds to the distance from the camera along the *Z*-dimension. In practice, the spread between particles is achieved using three off-axis pinhole cameras which causes a single point in space to appear at separate locations relative to the sensor of each camera. As described in [11, 12], the images from all three camera sensors are superimposed onto a common coordinate system, an algorithm searches for patterns which form an equilateral triangle, and based on size and location of the triangle the 3D spatial coordinates of the point can be resolved. A main limitation of this technique appears to be seeding density, because the equilateral triangles formed by individual particles must be resolved to reconstruct the particle field. Pereira and Gharib [8] have reported simulations with seeding density of 0.038 particles per pixel (ppp) in a volume size of 100 $\times$ 100 $\times$ 100 mm^3, and experiments with seeding density of 0.034 ppp in a volume size of 150 $\times$ 150 $\times$ 150 mm^3. The technique has also been efficiently implemented with a single camera using an aperture with color-coded pinholes, to measure velocity fields in a buoyancy driven flow in a 3.35 $\times$ 2.5 $\times$ 1.5 mm^3 volume with seeding density $\approx$0.001 ppp [13].

Tomographic-PIV also uses multiple viewpoints (usually three–six cameras) to obtain 3D velocity fields [14, 15]. Optical tomography reconstructs a 3D intensity field from the images on a finite number of 2D sensors (cameras); the intensity fields are then subjected to 3D PIV cross-correlation analysis. The seeding density for tomographic-PIV seems to be the largest attainable of the existing techniques. Simulations by Elsinga *et al* [14] show volumetric reconstruction with a seeding density of 0.05 ppp, and recent tomographic-PIV experiments typically have seeding density in the range of 0.02–0.08 ppp [15, 16]. The viewable depth of volumes in tomographic-PIV is typically three to five times smaller than the in-plane dimensions [14, 15, 17]. Elsinga *et al* [14, 15] thoroughly characterize the performance of tomographic-PIV, and we have adopted many of the same metrics in evaluating the synthetic aperture PIV method presented herein since both methods are based on 3D cross-correlation of reconstructed intensity fields.

Holographic PIV (HPIV) is a technique in which the three-dimensional location of particles in a volume is deduced from the interference pattern of the light waves emanating from particles and the coherent reference wave that is incident upon the field [18]. The nature of the interference pattern is used to back out information about the phase of light diffracted from objects in the volume, which is related to the distance of the objects from the sensor (i.e. depth in the volume) [19]. Holographic PIV makes use of this principle to image particle-laden volumes of fluids, and extract information about location of particles in the volume. Meng *et al* [20] provide an extensive review of film and digital Holographic PIV techniques. In holography, the size of the observable volume is ultimately limited by the size and spatial resolution of the recording device. Zhang *et al* [21] have reported very high resolution measurements of turbulent flow in a square duct using film-based HPIV, where particles were seeded to a reported density of 1–8 particles mm^{-3} in a volume measuring 46.6 $\times$ 46.6 $\times$ 42.25 mm^3. Although film has much better resolution and is larger than digital recording sensors, Meng *et al* [20] extensively cite the difficulties of film-based holographic PIV, which have likely prevented the method from being widely utilized. In contrast, digital holographic PIV is more readily usable, but is often limited to small volumes and low seeding density [20]. A digital hybrid HPIV method has been implemented by Meng *et al* [20] which allows for measurement in volumes with larger depth, but the sizes of the in-plane dimensions are limited by the physical size of the digital sensor, and seeding density remains low. Sheng *et al* [22] have presented recent results of measurements in a turbulent boundary layer with increased seeding density (0.014 ppp) in a volume measuring 1.5 $\times$ 2.5 $\times$ 1.5 mm^3.

The synthetic aperture PIV technique is implemented using an array of synchronized CCD cameras distributed such that the fields of view overlap. Images are recombined in software using a refocusing algorithm, commonly applied in synthetic aperture applications [23, 24]. The result is sharply focused particles in the plane of interest (high intensity), whereas particles out-of-plane appear blurred (low intensity). Due to the multiple camera viewpoints and the effective reduction of signal strength of out-of-plane particles in image recombination, particles that would otherwise be occluded can in fact be seen. The 3D intensity field of particle-laden flows can be reconstructed by refocusing throughout the entire volume and thresholding out particles with lower intensities. Typical 3D PIV techniques can then be applied to the intensity fields to extract velocity data. This technique enables larger volumes to be resolved with greater seeding density, yielding higher spatial resolution than prior 3D PIV methods. Additionally, the algorithms are simple and robust and build on established image processing techniques. Results of simulated particle fields show the ability to reconstruct 3D volumes with seeding densities of 0.17 ppp (6.68 particles mm^{-3}) when the ratio of X–Y to Z dimension is 5:1 (50 × 50 × 10 mm^3 volume), and 0.05 ppp (1.08 particles mm^{-3}) when the ratio of X–Y to Z dimension is 4:3 (40 × 40 × 30 mm^3 volume). A vortex ring flow field is imposed on each of these simulated volumes, and 3D PIV analysis yields highly resolved vector fields.

Results are presented from an experimental implementation of SAPIV using a custom-built camera array to study a vortex ring in a 65 × 40 × 32 mm^3 volume. Design considerations for experimental 3D SAPIV implementation are discussed throughout the paper. The experimental data presented are benchmarked with 2D PIV and demonstrate the ability of SAPIV to resolve 3D flow fields, providing a useful and flexible tool for making 3D PIV measurements.

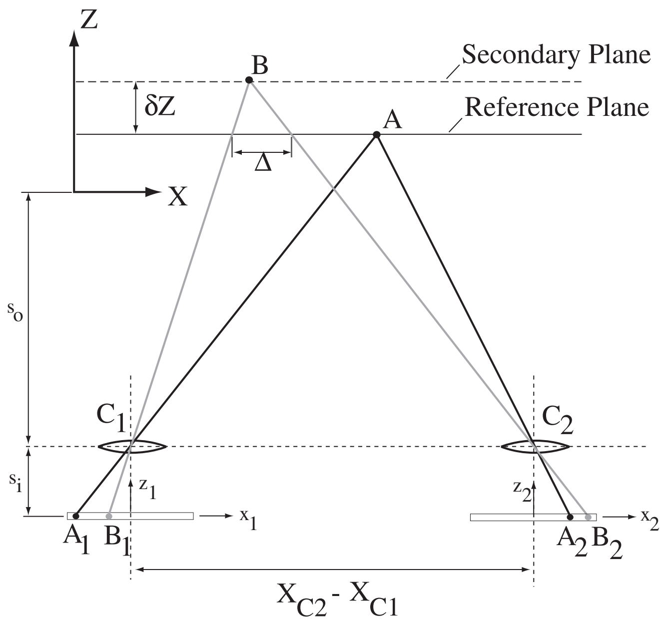

Figure 1. Schematic showing the optical arrangement and the concept of parallax. X–Z is the global coordinate system and x–z are local image coordinates.

2. Synthetic aperture methodology

2.1. Light field imaging

Synthetic aperture PIV is based on the concept of *light field imaging*, which involves sampling a large number of light rays from a scene to allow for scene reparameterization (Isaksen *et al* [23]). In practice, one method used by researchers in the imaging community for sampling a large number of rays is to use a camera array [23–26]. The novelty of the approach presented herein is the application of the reparameterization methods to 3D PIV, and the development of the technique into a measurement system, including the generation of algorithms to reconstruct 3D particle intensity fields from the refocused images. The technique is broken down into sequential components in the list below.

(i) **Image acquisition.** Image capture is performed using an array of cameras typically arranged in a multi-baseline stereo configuration, which view the scene from different viewpoints. The cameras can be placed at arbitrary locations and angles as long as the desired refocused planes are in the field of view (FOV) of each camera. The depth of field of each camera is large enough such that the entire volume of interest is in focus. The multiple viewpoints array captures many more light rays than can be seen with one camera (i.e. *light field imaging* [23]).

(ii) **Synthetic aperture refocusing.** The light fields are reparameterized using synthetic aperture refocusing [23, 24, 26].

(iii) **3D intensity field reconstruction.** 3D intensity fields generated by a fluid seeded with flow tracers are extracted from images refocused using the synthetic aperture refocusing method.

(iv) **3D cross-correlation.** Typical 3D intensity field cross-correlation methods similar to those used in tomographic-PIV [14, 15] operate on the reconstructed 3D intensity fields to extract velocity fields.

2.2. Synthetic aperture refocusing

In synthetic aperture techniques, images captured by an array of cameras each with large depths of field are post-processed to generate one image with a narrow depth of field on a specific focal plane [23, 25]. Through synthetic aperture refocusing, further post-processing allows the location of the focal plane to be arbitrarily placed within the imaged volume [23, 24, 26]. This technique provides many desirable implications for 3D PIV; namely, a particle-laden volume can be captured at one instant in time and synthetic aperture refocusing allows for the reconstruction of particles at known depths throughout the volume post-capture.

In general, the post-processing for synthetic aperture refocusing involves projecting all images onto a focal surface (planar or otherwise) in the scene on which the geometry is known, averaging the projected images to generate one image, and repeating for an arbitrary number of focal planes [23]. The working principle of the synthetic aperture technique is demonstrated with a simplified example. Consider the case where two cameras at different X locations view the same portion of the reference plane in figure 1. If the images from each camera are mapped to the reference plane, and then all images (now in reference plane coordinates) are averaged, the point A will be in sharp focus in the averaged image. However, points not on this plane (i.e. at different depths) will appear out of focus because of the parallax between cameras. In the average of the reference plane aligned images, the images

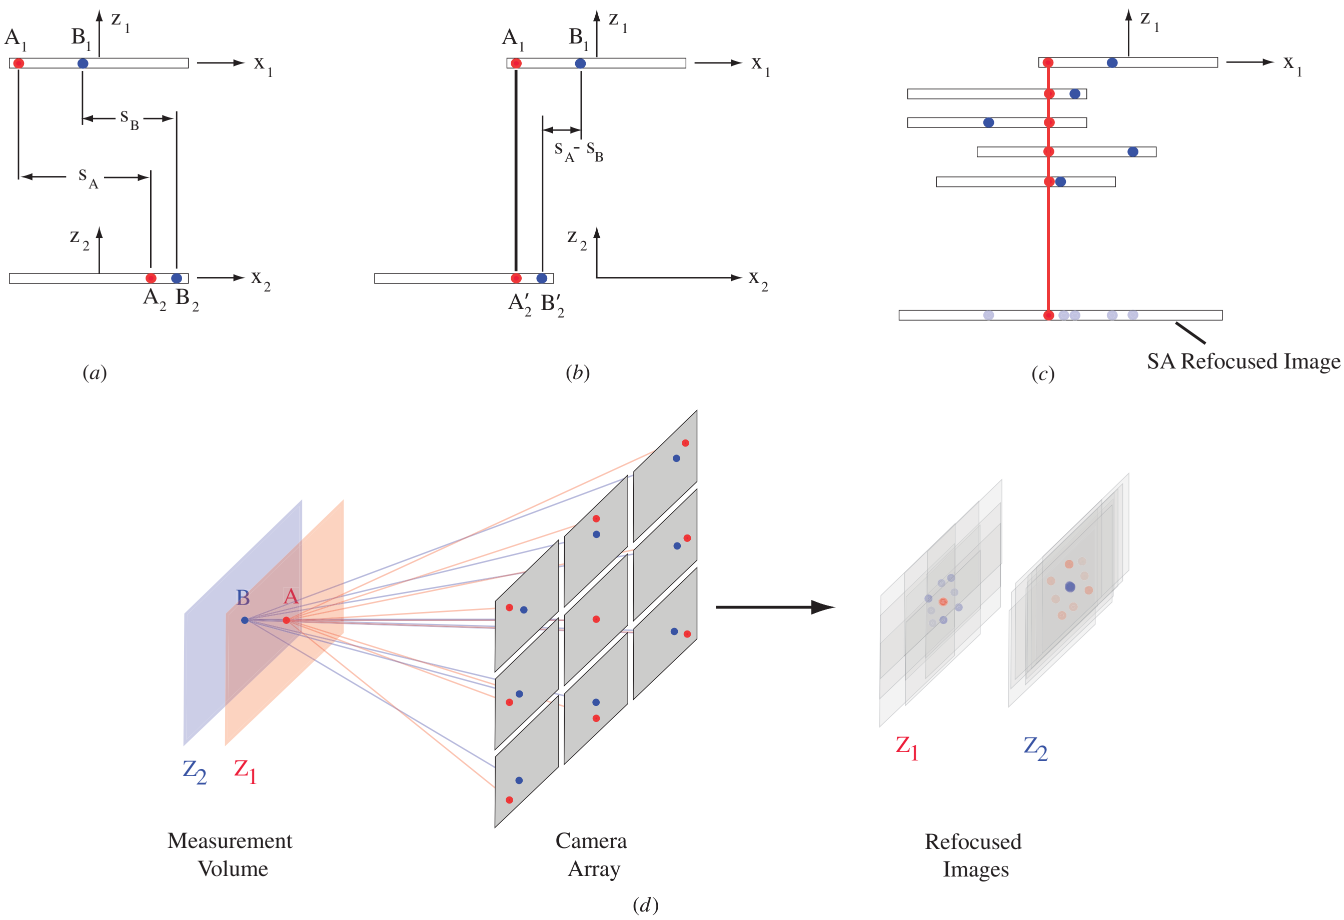

Figure 2. Schematic demonstrating the synthetic aperture refocusing method. In (*a*) the *z* axes of the image planes are aligned and in (*b*) image plane 2 has been shifted to align the images of point A. In (*c*) multiple images are shifted to align on the plane containing point A, and are averaged to generate a refocused image. The schematic in (*d*) shows a 3D depiction of image capture by a nine-camera array, and subsequent refocusing on two planes.

of point B from the two cameras will appear at the points where the gray rays in figure 1 intersect the reference plane, which are separated by Δ. By adding more cameras, mapping images to the reference plane and averaging, the signal of point A will grow increasingly larger than the 'noise' of the off-plane points. Note that this 'noise' does not refer to actual image noise, but the signal from particles not on the plane of interest. The concept is shown schematically in 3D in figure 2(*d*), which shows image capture by a nine-camera array, and subsequent refocusing on two planes. By positioning the cameras on a sufficiently large baseline (larger separation between camera centers of projection (COP)), some of the cameras can see particles which are occluded in other images [23–26]. Therefore, the partially occluded particles retain a high signal in the refocused image.

2.3. Three-dimensional volume reconstruction

The aim of the synthetic aperture PIV technique is to reconstruct 3D particle intensity fields which are suitable for cross-correlation-based 3D PIV processing. The starting point for volume reconstruction is the implementation of the synthetic aperture algorithm to generate refocused images on planes throughout the volume. Thereafter, the actual particle field must be extracted from the refocused images and organized into a volume with quantifiable locations.

To implement synthetic aperture refocusing, relationships between the image coordinates and focal planes in world coordinates must be established (i.e. a set of mapping functions). In the simulations presented herein, cameras are represented with a pinhole model, and the mapping functions can be generated by an algorithm presented in [24], which we refer to as the *map-shift-average algorithm*. This algorithm is suitable for pinhole cameras with no image distortion and no changes in optical media in the line of sight (e.g. air–glass–water transition). In practice, mapping functions that can account for distortion and changes in optical media can be generated by more sophisticated calibration techniques, as is done for the experiment discussed in section 3. For example, as employed by [16], a planar calibration grid can be placed at several planes throughout the volume to generate mapping functions, and error in these functions is reduced using a volume self-calibration technique [27]. The map-shift-average algorithm represents one of the most straightforward refocusing algorithms that can be implemented in SAPIV; it is a good illustration of how the method works and is implemented in the simulations presented in section 2.4.

2.3.1. Map-shift-average algorithm. The first step in the map-shift-average algorithm is to align all images on a reference plane [24]. The reference plane is an actual plane in the view of the cameras which, in practice, is defined using a planar calibration grid. Images are aligned on the reference plane by applying a *homography* which, as described in [28], is a central projection mapping between two planes given by

$$\begin{pmatrix} bx' \\ by' \\ b \end{pmatrix}_i = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix}_i \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}_i , \tag{1}$$

where b is a constant, $h_{33} = 1$, and the vectors $\mathbf{x}'$ and $\mathbf{x}$ are the homogeneous coordinates of points in the reference plane and the image plane, respectively. In this way, the raw image from the ith camera, I_i, is mapped to a new image in reference plane coordinates, I_{RP_i}. Hartley and Zisserman [28] describe robust algorithms for estimating the homography based on the images of several points on the reference plane.

Light field imaging with camera arrays offers the ability to refocus on arbitrarily shaped focal surfaces [23, 24, 26]. For the simulations herein, we restrict the situation to refocusing on fronto-parallel planes where the raw images are captured with a camera array where all camera COP lie on the same plane. For this restricted case, the mapping that must be applied to refocus on each fronto-parallel plane is simply a shift of the reference-plane aligned images by an amount proportional to the relative camera locations [26]. The coordinates of image points on the new focal plane are given by

$$\begin{pmatrix} x'' \\ y'' \\ 1 \end{pmatrix}_i = \begin{bmatrix} 1 & 0 & \mu_k \Delta X_{C_i} \\ 0 & 1 & \mu_k \Delta Y_{C_i} \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix}_i , \tag{2}$$

where μ_k is a constant that determines the location of the kth focal plane of the refocused image, and ΔX_{C_i} and ΔY_{C_i} are the relative camera locations of the ith camera. This portion of the algorithm transforms the ith reference plane aligned image, I_{RP_i} to an image on the kth focal plane, $I_{FP_{ki}}$. The relative camera locations can be found using a simple calibration method presented in Vaish *et al* [26] which operates on the images from the camera array of points on planes at several depths. No *a priori* knowledge of the camera locations or optical configurations are required to implement this algorithm.

In the final step in the refocusing algorithm, the mapped and shifted images (now all aligned on the kth focal plane) are averaged. The resultant image is referred to as the refocused image, and is given by

$$I_{SA_k} = \frac{1}{N} \sum_{i=1}^{N} I_{FP_{ki}}, \tag{3}$$

where images are combined from N cameras [24]. Particles that lie on the plane of the refocused image appear sharp and have high intensity, while particles not on this plane are blurred and have lower intensity. The only difference between the methodology described here and the practical realization comes in the mapping functions used to align all images on the focal planes. Rather than using the linear homography and homology, a higher order mapping will likely be required to deal with distortions and changes in optical media [16, 27].

2.3.2. Focal plane spacing and viewable depth. Theoretical analysis of the focal plane spacing and viewable depth is presented using the two-camera model shown in figures 1 and 2. Here, we define the focal plane spacing, δZ, as the distance between refocused planes throughout the volume, where Z is the distance from the front of the experimental volume. The viewable depth is the total depth dimension (Z_o) of the reconstructed volume. The simple two-camera model is sufficiently general to establish the theoretical limits of the system, and the results can be applied to design and predict the performance of various arrangements. A full-scale simulated model is implemented in section 2.4 to examine the effects of parameters not considered in this theoretical treatment, such as particle seeding density, number of cameras and mapping error.

In the model of the imaging system shown in figure 1, the image sensors are parallel to each other and the COP lie on the same plane. The geometric imaging optics are described by four parameters: the focal length (f), distance from the lens to the front of the imaged volume (s_o), distance from lens to image plane (s_i) and magnification ($M(Z) = -s_i/(s_o + Z)$). Also, the focal length is related to s_o and s_i through the imaging condition: $1/f = 1/s_i + 1/s_o$.

To examine the effect of camera layout on focal plane spacing, we start by considering the relationship between points in physical space and the images of these points. The coordinates of a general point, A (see figure 1), projected onto the imaging plane of one of the cameras is given by

$$x_A = \left[\frac{-s_i}{s_o + Z_A}\right](X_A - X_C) - d_C, \tag{4}$$

where X_C is the X coordinate of the COP in global coordinates, and d_C is the displacement of the image sensor from the center of projection (zero as shown). We define the value of 'initial' magnification as $M(Z = 0)$. The dimension of each pixel projected into physical space is given by

$$\delta X = \frac{-p}{M(Z)}, \tag{5}$$

where p is the pixel pitch (width of a pixel).

As described earlier, the first step in the reconstruction is to align the images on a reference plane. In the model in figure 1, the reference plane is chosen (arbitrarily) to be the plane in physical space on which point A lies. In figure 2, the image sensors are shown with the centers aligned along the z-axes of the local image coordinates. To align the images on the reference plane would require a shift of the image sensor (or equivalently, the digital image) of camera 2 by an amount equal to $s_A = x_{A_2} - x_{A_1}$ (the disparity of the images of point A) in the negative x-direction of the camera 2 coordinate system. Since point B does not lie on the reference plane, $s_B = x_{B_2} - x_{B_1}$ does not equal s_A, and thus points at different depths are disambiguated.

The focal plane spacing, δZ, is dictated by the image plane shift required to refocus on another plane. To move from the reference plane to the secondary plane shown in the model, the image plane must be shifted by an amount equal to $s_A - s_B$, which is given by

$$s_A - s_B = \Delta X_C \left[\frac{-s_i}{s_o + Z_B} - \frac{-s_i}{s_o + Z_A}\right], \tag{6}$$

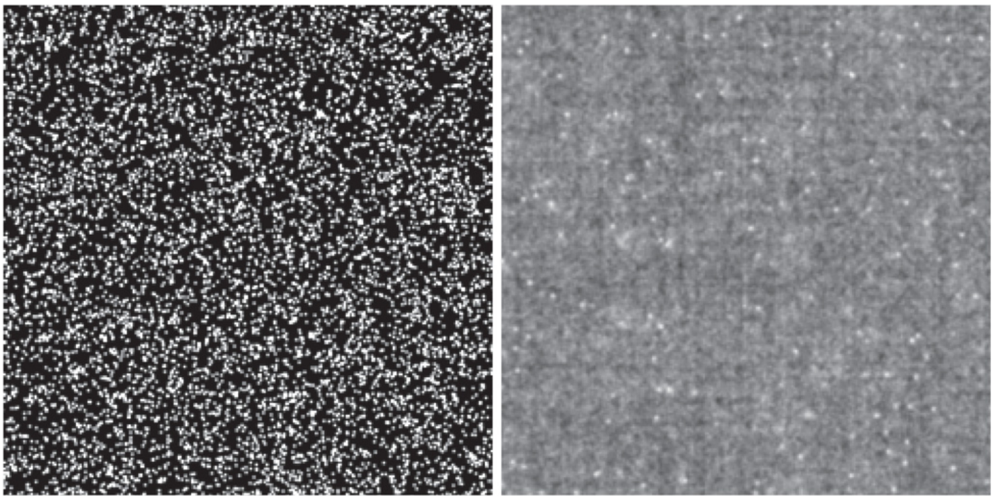

(*a*) (*b*)

Figure 3. Zoomed views (250 × 250 pixels) of (*a*) a simulated image from the central camera of the array and (*b*) a refocused image using all of the simulated images of the array (25 images).

where $\Delta X_C = X_{C2} - X_{C1}$ is the separation between the camera COPs. The minimum amount that the image sensor can be shifted by is the width of one pixel (assuming no spatial interpolation); therefore, $s_A - s_B \geqslant p$. Letting $Z_A = Z$ and $Z_B = Z + \delta Z$ and imposing the minimum shift requirement yields

$$\Delta X_C \left[\frac{-s_i}{(s_o + Z) + \delta Z} - M(Z) \right] = p. \tag{7}$$

Solving equation (7) for δZ gives

$$\delta Z = -s_i \left[\frac{\Delta X_C}{p + \Delta X_C \cdot M(Z)} - \frac{1}{M(Z)} \right]. \tag{8}$$

Dividing the top and bottom of the first term in the bracket by s_o yields

$$\delta Z = -s_i \left[\frac{\left(\frac{\Delta X_C}{s_o}\right)}{\frac{p}{s_o} + \left(\frac{\Delta X_C}{s_o}\right) \cdot M(Z)} - \frac{1}{M(Z)} \right]. \tag{9}$$

Equation (9) contains the convenient geometric parameter, $\frac{\Delta X_C}{s_o}$, which is the ratio of the camera spacing to the distance from the front of the imaged volume. For convenience, we will let $D \equiv \frac{\Delta X_C}{s_o}$. This parameter characterizes the baseline of the system. Equation (9) can be further simplified by applying equation (5) at Z to replace p and rearranging,

$$\frac{\delta Z}{\delta X} = \frac{\frac{1}{D} + \frac{Z}{Ds_o}}{1 - \frac{\delta X}{Ds_o}}. \tag{10}$$

For typical PIV applications, it is reasonable to assume $\delta X \ll Ds_o$; applying this approximation yields

$$\frac{\delta Z}{\delta X} \cong \frac{1}{D} + \frac{Z}{Ds_o}. \tag{11}$$

Therefore, the ratio $\delta Z/\delta X$ is a linear function of Z with intercept defined by the camera baseline and slope defined by the camera baseline and the imaging system optics. The depth of field in the reconstructed volume is directly related to $\delta Z/\delta X$. When $\delta Z/\delta X$ is small, the camera lines of sight are at large angles to one another, and the physical depth over which particles are in focus is small (small depth of field). Conversely, larger $\delta Z/\delta X$ leads to a larger depth of field, which is manifested as reconstructed particles which are elongated in Z.

In theory, the overall viewable range in X, Y and Z is limited only by the field of view (FOV) and depth of field (DOF) of a single camera of the array. In reality, images from the outer cameras of the array must be shifted with respect to the central camera image to refocus at different depths. The outer edges of the refocused images have a lower signal-to-noise ratio than regions where all images contribute to the average. This effective limitation on the FOV can be characterized by the number of image shifts required to refocus through an entire volume, which is given by

$$N = \frac{Z_o}{\delta Z}, \tag{12}$$

where Z_o is the depth of the volume. Since images are shifted in integer pixel increments (assuming no interpolation), it is possible to calculate the region of each refocused image to which all images will contribute. As will be seen in section 2.4, this effect does degrade the performance near the outer edges of the larger reconstructed volumes. This effect could be mitigated by excluding the portions of the shifted images which do not contribute from the average, but this adjustment to the refocusing algorithm has not been included here. Nonetheless, the implications of the technique are that the observable range of the system is highly scalable, with the ability to trade-off depth of field for viewable range much as one would trade-off X–Y resolution for FOV with a single camera.

2.3.3. 3D particle intensity field extraction. Once the refocused images have been generated, the images must be processed to extract the actual particles from the blurred, lower intensity background. Figure 3(*a*) shows a zoomed view (250 × 250 pixels) of a simulated image from the central camera of the array and figure 3(*b*) shows a refocused image

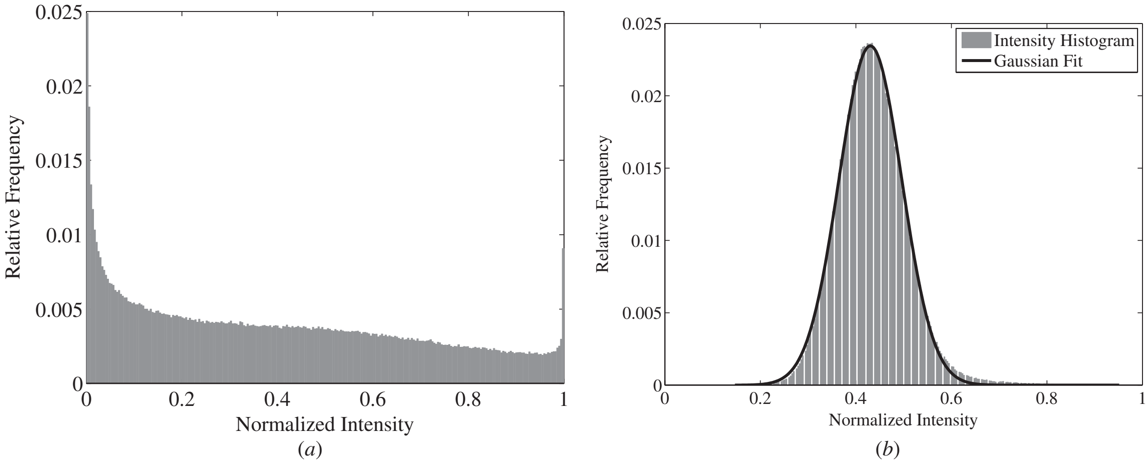

Figure 4. Intensity histograms for a single simulated image aligned on a given focal plane (*a*) and for the refocused image (from figure 3(*b*)) on that focal plane (*b*).

from the 50 × 50 × 10 mm^3 volume simulation which will be described in detail in section 2.4. The refocused image has a higher background 'noise' level than the raw image or a typical 2D PIV image due to the averaging of multiple images; however, the 'noise' is probabilistic. If we consider the intensity fields of the N images aligned on a given focal plane to be independent and identically distributed random variables, $I_1, I_2, ..., I_N$, with means μ and standard deviations σ, then the central limit theorem states that the distribution of the average of the random variables will be Gaussian with mean μ and standard deviation of $\sigma/\sqrt{N}$. Therefore, the intensity distribution of a refocused image (which is an average of the focal plane aligned images) can be modeled as Gaussian. Figure 4(*a*) shows the intensity histogram for a reference plane-aligned image from one camera of the array and the histogram for the refocused image in figure 3(*b*) is shown in figure 4(*b*). Clearly, the distribution of intensity for the single camera image is not Gaussian, but the shape of the distribution of the refocused image follows a Gaussian distribution quite well, as indicated by the model fit. As more particles are added, the mean of the individual images becomes larger, and thus the mean of the refocused image becomes larger. Actual particles appear with high intensity values and are thus outliers with respect to the distribution of the refocused image. Intensity thresholding can be applied to retain actual particles and eliminate background 'noise' from the images. It was found that a threshold value around three standard deviations above the mean intensity of each refocused image yielded acceptable reconstruction. Figure 5 shows the refocused image from figure 3(*b*) now thresholded to reveal the true particles. By refocusing throughout the volume and thresholding the refocused images, the three-dimensional intensity field is reconstructed. Although the thresholding method may require more optimization, it appears that by detecting the outliers in the refocused images the true particle field can be reconstructed, attesting to the simplicity of the SAPIV technique.

Figure 5. Zoomed view (250 × 250 pixels) of the thresholded refocused image from figure 3(*b*) revealing the particles and removing the background noise.

2.4. Simulated camera array

A 5 × 5 camera array model is simulated to investigate the system performance as a function of particle seeding density, size of measurement volume and error in the mapping function. The effect of array layout and camera number on reconstruction performance is also investigated by changing the spacing between cameras and removing certain cameras from the array, respectively. Cameras are arranged with COPs on the same plane and equal spacing along X and Y between all camera COPs (unless otherwise noted). In order to overlap the fields of view, the cameras are angled such that the ray passing through the center of the measurement volume intersects the center of the image sensor. For this arrangement, the map-shift-average algorithm applies. The perspective due to the

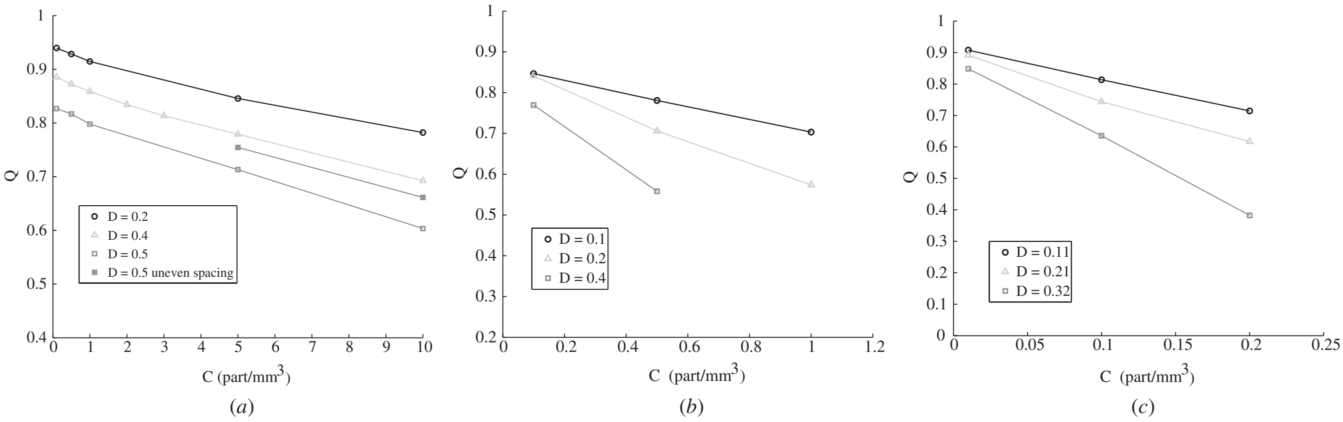

Figure 6. Effect of particle seeding density (C, particles mm^{-3}) on reconstruction quality, Q, for various camera baselines in a 50 × 50 × 10 mm^3 volume (*a*), a 50 × 50 × 50 mm^3 volume (*b*) and in a 100 × 100 × 100 mm^3 volume (*c*).

angling of the cameras is compensated for when the reference plane homographies are applied to the images.

The ability of the system to resolve the particle field in four different measurement volume sizes is examined. The volume sizes ($X \times Y \times Z$) are 50 × 50 × 10 mm^3, 40 × 40 × 30 mm^3, 50 × 50 × 50 mm^3 and 100 × 100 × 100 mm^3. For each volume, the system performance is evaluated for several different camera baselines, D, and for each camera baseline the particle seeding density is varied. For all simulations presented herein, the cameras are modeled with 85 mm focal length lenses and the imaging sensors are 1000 pixels × 1000 pixels with a pixel pitch of 10 μm. For the 100 × 100 × 100 mm^3 volume, the initial magnification is set to $M(Z = 0) = -0.1$ and for the other three volumes the magnification is $M(Z = 0) = -0.2$.

Reference plane homographies are calculated for each camera from calibration images of known points on several Z planes with the central camera of the array as the reference camera. The camera positions relative to the central camera are established using the calibration method described in [26]. The calibration images are used to establish the shift required to refocus at each depth in order to define the conversion between voxels and physical units in the Z dimension. Herein, we define a voxel as having the dimensions of a pixel in X and Y and dimension equal to the focal plane spacing in Z. Because integer pixel shifts are used in the map-shift-average algorithm, a given calibration depth may not correspond exactly to any of the refocused images. Therefore, the actual voxel to Z calibration is approximated by fitting a Gaussian curve to the summed intensity from refocused images surrounding each calibration plane and finding the voxel corresponding to the peak of the fit.

Particles are randomly seeded within the volume and imaged using the camera array. Once the image plane coordinates of a point are known, a realistic model of the intensity distribution must be applied. A standard method of simulated image generation is to apply a Gaussian intensity profile as described in [2]. The distribution is applied to each camera image for each pixel location, which forms an image similar to the one presented in figure 3(*a*).

After simulated images have been formed by each camera, the map-shift-average algorithm, followed by intensity thresholding and three-dimensional field reconstruction, is carried out for each numerical experiment. In order to quantify how well the intensity field is reconstructed, the same measure of reconstruction quality, Q, used in Elsinga *et al* [14, 15] is applied here:

$$Q = \frac{\sum_{X,Y,Z} E_r(X,Y,Z) \cdot E_s(X,Y,Z)}{\sqrt{\sum_{X,Y,Z} E_r^2(X,Y,Z) \cdot \sum_{X,Y,Z} E_s^2(X,Y,Z)}}, \qquad (13)$$

where E_r is the reconstructed intensity field and E_s is a synthesized intensity volume based on the known particle locations.

Baseline spacing affects the Z dimension of the voxels such that they represent larger physical sizes than the $X-Y$ dimensions; therefore, the intensity distribution in the synthesized field is scaled in Z in voxel space such that in physical space the intensity distribution is spherically symmetric. This ensures that a perfectly reconstructed particle would yield a Q value of 1 when compared to the synthesized field.

The value of Q is calculated for each numerical experiment conducted, and we use the same requirement as in [14] of $Q \geqslant 0.75$ for the reconstruction to be considered adequate. In all cases other than the 40 × 40 × 30 mm^3 measurement volume, the outer 50 edge pixels were cropped prior to calculating Q because of the effective loss in FOV as described earlier (in the 40 × 40 × 30 mm^3, the images do not fill the entire image sensor and thus the outer pixels of the reconstructed images contain no particles). Figures 6(*a*)–(*c*) present the reconstruction quality as a function of particle seeding density for various camera baselines in each volume. The number of seeded particles, maximum particle seeding density (C) and number of particles per pixel (ppp) in each case are summarized in table 1. To find the maximum seeding density, simple linear interpolation of the data was used to find the seeding density corresponding to $Q = 0.75$.

In the case of the 50 × 50 × 10 mm^3 measurement volume, the reconstruction quality falls off with increasing seeding density, and also decreases with increasing camera baselines.

Table 1. Summary of resolution and particle seeding density for simulated measurement volumes.

Measurement volume (mm^3)	D	δX (mm)	$\delta Z/\delta X$	C (particles mm^{-3})	particles/ pixel (ppp)	No of particles
$50 \times 50 \times 10$	0.2	0.05	5.00	>10	>0.25	>250 000
	0.4	0.05	2.50	6.68	0.17	166 960
	0.5	0.05	2.00	3.26	0.08	81 537
	0.5 (uneven)	0.05	2.00	5.24	0.13	131 040
$40 \times 40 \times 30$	0.4 (uneven)	0.05	2.51	1.08	0.05	51 988
$50 \times 50 \times 50$	0.1	0.05	10.24	0.70	0.09	87 371
	0.2	0.05	5.13	0.37	0.05	46 229
	0.4	0.05	2.56	0.14	0.02	17 159
$100 \times 100 \times 100$	0.11	0.1	9.40	0.16	0.16	164 240
	0.21	0.1	4.70	0.10	0.10	95 028
	0.32	0.1	3.13	0.05	0.05	51 659

For larger seeding density, the reduction in reconstruction quality is expected; the occurrence of off-plane particle overlap is increased, and the overall signal-to-noise ratio decreases in the refocused images. The reason for reduced reconstruction quality with increasing baseline is less obvious. Investigation of the data reveals that the reason for the degradation in reconstruction quality with increasing baseline is due to the more extreme warping of the particle images imposed by the homography which maps images to the reference plane. The particles become elongated in the mapped image which raises the background noise in the refocused images. This may be mitigated by placing the outer cameras of the array a normalized distance $D/2$ from the inner cameras (which determine the focal plane spacing of the system and would still be placed a normalized distance D from the central camera) but requires interpolation when shifting images from the outer cameras of the array. This has been implemented in the case labeled $D = 0.5$ uneven spacing, and indeed the reconstruction quality is improved for the same seeding density. For this configuration, the achievable seeding density is $C = 5.24$ particle mm^{-3} and the resultant particles per pixel is 0.13.

Simulations in the $50 \times 50 \times 50$ mm^3 volume show a similar trend as for the $50 \times 50 \times 10$ mm^3 volume, with Q decreasing with increasing C more rapidly for a larger camera baseline. However, the achievable seeding density is lower than for the volume with smaller depth (e.g. $C = 0.37$ particle mm^{-3} for $D = 0.2$). The reasons for the lower seeding density and the lower actual number of particles that can be seeded as compared to the $50 \times 50 \times 10$ mm^3 volume are fourfold. First, the larger depth of field of each camera requires a larger f-number which increases the particle image diameter; this results in a larger mean intensity in the refocused image. Second, the depth of the volume creates more likelihood of particle overlap on each simulated image, which can, in a sense, decrease the dynamic range of the system by increasing the likelihood of saturated pixels. Third, the larger depth creates a higher likelihood of overlapping of many different out-of-focus particles in the refocused images. These false particles may be retained in the thresholding if enough images overlap. Finally, the limitation on the field of view imposed by the image shift contributes to the loss in reconstruction quality because some images contribute zero values to the average toward the outer regions of the refocused images; thus, true particles have a lower intensity value. This can be mitigated by averaging only the portions of the image which are known to contribute to the refocused image, but that technique has not been implemented here. Tuning the camera array and reconstruction algorithms to enable more seeding in the very large volumes is the subject of ongoing work. By decreasing the dimensions of the volume somewhat, more particles can be seeded, even in volumes where the Z dimension approaches that of the X–Y dimensions, and all are relatively large. Simulations in the $40 \times 40 \times 30$ mm^3 volume are carried out at only one array configuration—$D = 0.4$ uneven spacing—and it was found that a seeding density of $C = 1.08$ particle mm^{-3} corresponding to 0.05 particles per pixel could be achieved. As shown section 2.5, this seeding density allows 3D PIV measurements to be made in this volume with reasonable spatial resolution.

Finally, the camera magnifications are reduced to accommodate the large $100 \times 100 \times 100$ mm^3 measurement volume. The trend for reconstruction quality as a function of seeding density and baseline is similar to that observed for the other volumes studied. The total number of particles that can be seeded is, however, larger than for the $50 \times 50 \times 50$ mm^3 volume for comparable camera baselines. Thus, trading off X–Y resolution allows for more particles to be seeded even with increasing depth dimension. Overall, these results indicate that the synthetic aperture PIV technique is capable of imaging extremely densely seeded volumes where the depth dimension is somewhat reduced, and still quite densely seeded volumes when the Z dimension approaches that of the X–Y dimensions.

The effect of the camera number on reconstruction quality is investigated in the $50 \times 50 \times 10$ mm^3 volume by using only some of the cameras in the array. Figure 7(*a*) shows Q as a function of camera number with seeding densities of 2, 3 and 5 particle mm^{-3}. Clearly, the reconstruction quality reaches a point of diminishing returns as more cameras are added, and the most efficient number of cameras seems to be in the range of 10–15. Finally, to determine an upper limit to the allowable error in calibration, simulations in the $50 \times 50 \times 10$ mm^3 volume with 13 cameras and 3 particle mm^{-3} are carried out imposing error in the reference plane homography. This results in the misalignment of mapped images on each focal plane. The total misalignment error is

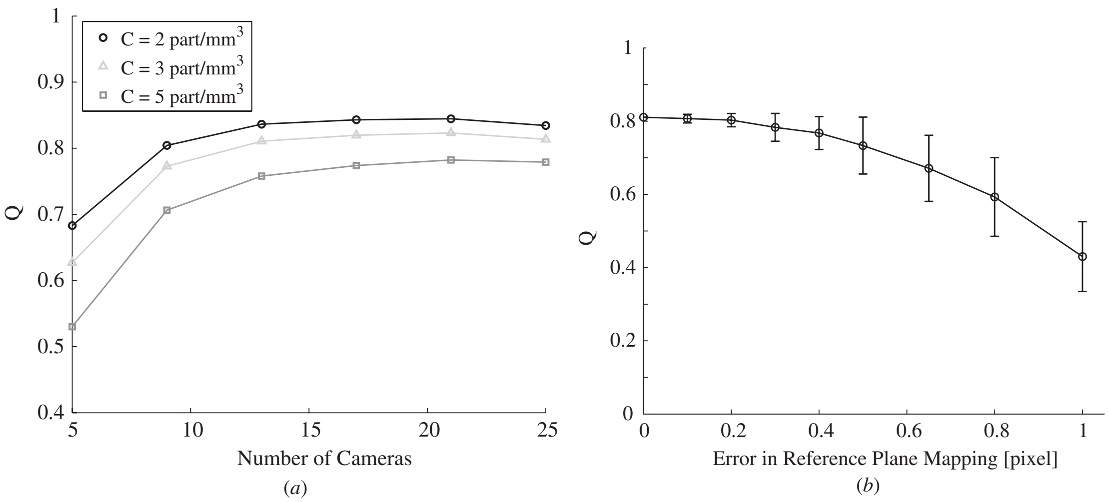

Figure 7. Reconstruction quality, Q, as a function of camera number for various particle seeding densities (*a*) and error in reference plane mapping (*b*) in the $50 \times 50 \times 10$ mm^3 volume

defined as $\epsilon_{RP} = \sqrt{\epsilon_x^2 + \epsilon_y^2}$ and is the same for all cameras, but ϵ_x and ϵ_y are randomized for each camera so as not to introduce any bias. For each value of ϵ_{RP}, ten simulations are carried out such that different random values of ϵ_x and ϵ_y are applied to each mapping for each simulation; the mean value of Q is then found for each ϵ_{RP}. Figure 7(*b*) shows the mean value of Q as a function of ϵ_{RP}; the error bars represent three standard deviations from the mean. The value of ϵ_{RP} corresponding to a mean value of $Q = 0.75$ was found to be 0.45 pixels. Volume self-calibration techniques such as that described in [27] are capable of reducing mapping errors to less than 0.1 pixels; therefore, we expect the synthetic aperture PIV technique to be robust to the error levels that will exist in mapping functions applied in actual experiments.

2.5. Synthetic 3D flow fields

Two synthetic flow fields are simulated to assess the ability of the synthetic aperture PIV method in reconstructing 3D intensity fields that are suitable for cross-correlation-based 3D PIV analysis. In each simulation, the fluid motion is prescribed by the same equation for a vortex ring as used in Elsinga *et al* [14, 15], where the velocity magnitude is given by

$$|V| = \frac{8KR}{l} e^{-\frac{R}{l}}, \tag{14}$$

where R is the radial distance of a point from the toroidal axis of the vortex ring and l is a length scale determining the size of the vortex ring (chosen to be 4 mm). The constant, K, is a conversion factor from voxel to physical units and is required since, in the present case, the synthetic particles are seeded in physical space and imaged by the model camera array. The toroidal axis forms a circle of diameter 20 mm. The first synthetic experiment is carried out in a $50 \times 50 \times 10$ mm^3 volume where the central axis of the toroid is parallel to the Z-axis. The volume is seeded with 125 000 particles ($C = 5$ particle mm^{-3}), and images are simulated in a 21-camera array with spacing $D = 0.5$ (uneven spacing, see table 1) and magnification $M(Z = 0) = -0.2$. This results in a reconstruction quality of $Q = 0.76$. The maximum displacement in this flow field is 0.3 mm which corresponds to 6 voxels in X and Y and 2.9 voxels in Z.

Since the reconstructed volumes are intensity fields, a cross-correlation-based PIV calculation is suitable for calculating vector fields. In this study, we have adapted an open-source 2D PIV code, matPIV [29], for 3D functionality. A multipass algorithm with a final interrogation volume containing $32 \times 32 \times 16$ voxels and 75% overlap generates 327 448 vectors ($122 \times 122 \times 22$ vectors). The Z dimension of the interrogation volumes in voxel units is half that of the X–Y dimension because the focal plane spacing is twice the pixel size for this camera configuration. Each interrogation volume contains approximately 20 particles, based on the gross particle seeding density.

Figure 8(*a*) shows the resultant velocity field with every tenth vector plotted in X and Y and every vector plotted in Z. The Z voxel and velocity values are multiplied by $\frac{\delta Z}{\delta X}$ to create the correct aspect ratio for plotting purposes. The fields are plotted in voxel units; if converted to physical units, the data set would not be cubical. Two slices are shown with normalized velocity magnitude contours revealing the vortex ring structure and symmetry. The maximum velocity magnitude in the exact known field is used to normalize the velocity magnitude in the processed field. Figure 8(*b*) shows the vector field and a vorticity iso-surface (0.15 voxels/voxel) with every sixth vector plotted in X and Y and every vector plotted in Z.

To quantitatively evaluate the performance, both the reconstructed 3D intensity fields and the synthesized 3D intensity fields are processed using the 3D adaptation of matPIV, and each is compared to the exact velocity field. The error is defined as the difference between the processed and exact field at every vector location. By comparing the PIV results for both fields, error due to the PIV algorithm itself can be identified. Both the synthesized and reconstructed volumes are processed using exactly the same window sizes, PIV code and filtering routines. We will refer to the vector fields resulting from PIV processing of the reconstructed 3D intensity fields and the synthesized 3D intensity fields as

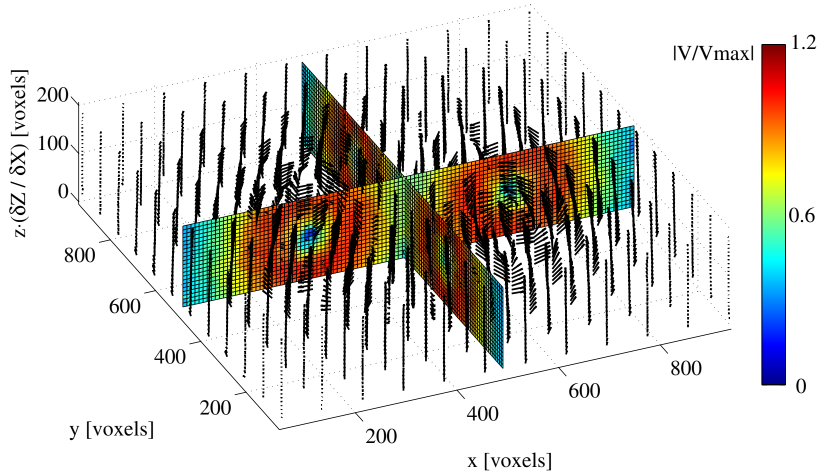

(*a*)

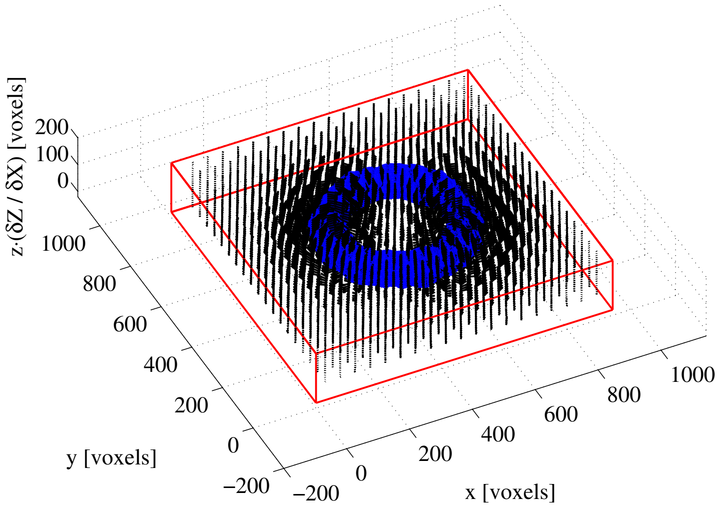

(*b*)

Figure 8. Three-dimensional vector field resulting from PIV processing of the reconstructed intensity volumes using SAPIV in the 50 × 50 × 10 mm^3 volume. (*a*) Two cuts with normalized velocity magnitude contours; (*b*) the vector field and a vorticity iso-surface (0.15 voxels/voxel).

Table 2. Summary of the PIV error in three simulated flow fields. The fourth column is the 90% precision interval for error between the reconstructed and exact field, and the fifth column is for the error between the synthesized and the exact field.

Measurement volume (mm^3)	Component	$\|V_{max}\|$ (voxels)	90% precision interval (voxels)–recon.	90% precision interval (voxels)–syn.
50 × 50 × 10	u	6	±0.37	±0.33
	v	6	±0.38	±0.34
	w	2.9	±0.20	±0.14
50 × 50 × 10	u	18	±1.69	±1.62
	v	18	±1.72	±1.62
	w	8.7	±0.90	±0.80
40 × 40 × 30	u	7.3	±0.54	±0.48
	v	7.3	±0.43	±0.36
	w	2.9	±0.19	±0.14

the reconstructed vector field, and synthesized vector field, respectively.

Figure 9 shows scatter plots of the error in each vector component for the reconstructed vector field ((*a*) and (*b*)) and the synthesized vector field ((*c*) and (*d*)). The error is further characterized by calculating the 90% precision interval for the error in each vector component. This is summarized in table 2 along with the maximum magnitude of velocity in the exact field. The fourth column is the 90% precision interval for error between the reconstructed and exact field,

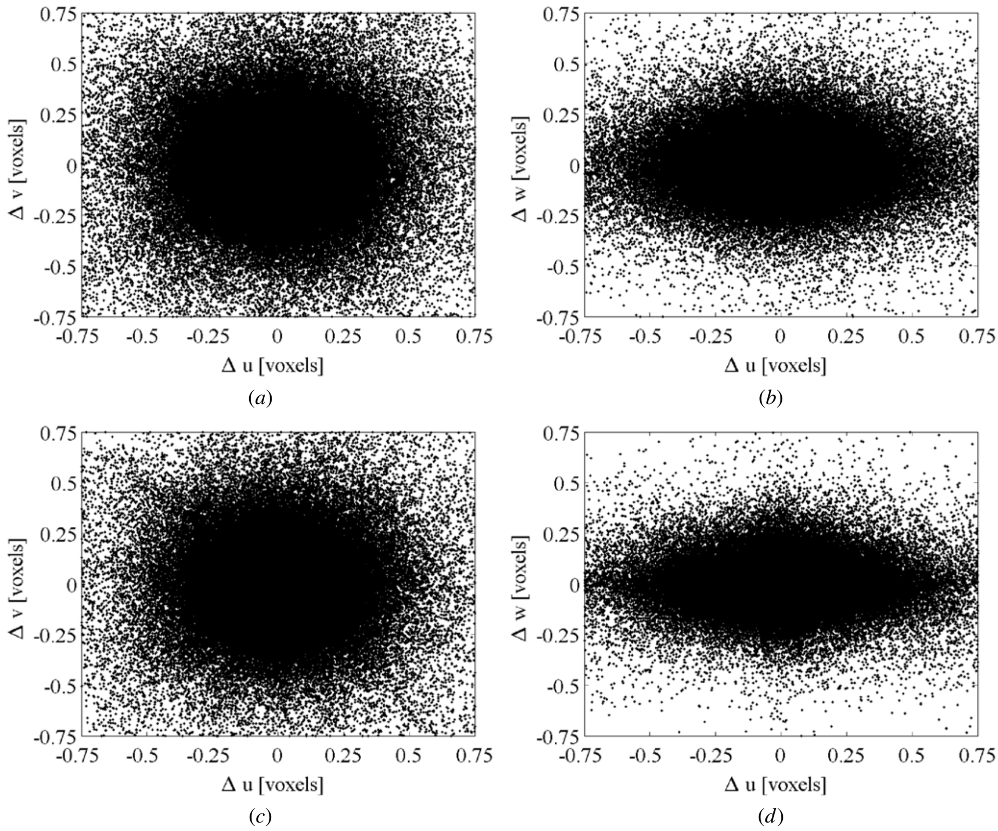

Figure 9. Scatter plots of the error in the reconstructed vector field ((*a*) u, v components and (*b*) u, w components), and error in the synthesized vector field ((*c*) u, v components and (*d*) u, w components). All plots are for the vortex ring in the 50 × 50 × 10 mm^3 volume.

and the fifth column is the 90% precision interval for the error between the synthesized and the exact field. The precision intervals seem rather large, but comparison between the error in the reconstructed and synthesized vector fields shows that the difference between the precision interval values are small compared to the actual error magnitude. This indicates that most of the error is due to the PIV algorithm itself, and a much smaller percentage is due to the actual intensity field reconstruction. This is supported by the error scatter plots in figure 9, which show only slightly more spread in the error for the reconstructed velocity fields. To examine this further, the numerical experiment was repeated with the velocity magnitude increased by three fold everywhere in the flow field. As shown in table 2, the precision interval magnitude increases for both the reconstructed and synthesized vector fields, but the difference between the precision intervals essentially remains the same. This further points to the PIV algorithm as the largest source of error in the vector field. This is not surprising since the 3D PIV algorithm is not very sophisticated, and we would expect a reduction in error magnitude with a more advanced 3D PIV algorithm.

In the second simulated flow, a vortex ring of the same size is oriented with the central axis of the toroid parallel to the X-axis, such that the ring spans deeper into the flow in the Z dimension. The 5 × 5 model camera array is used with spacing $D = 0.4$ (uneven spacing, see table 1) and the magnification is set to $M(Z = 0) = -0.2$. The maximum displacement in the flow field is 0.37 mm which corresponds to 7.3 voxels in X and Y and 2.9 voxels in Z. A particle seeding density of $C = 1$ particle mm^{-3} in a 40 × 40 × 30 mm^3 volume results in a distribution of 48 000 particles (resulting in $Q = 0.76$). The lower seeding density requires larger interrogation volume sizes in order to contain an appropriate number of particles; therefore, a final interrogation volume containing 60 × 60 × 24 voxels is used with each containing 27 particles based on the seeding density. Using 50% overlap in the multipass 3D PIV calculation yields 18 432 vectors (32 × 32 × 18), which includes the imaged area that contained no seeding particles (with a magnification of −0.2, images of the seeded volume did not span the entire imaging sensor). Figure 10 shows the vector field results, with figure 10(*a*) revealing velocity magnitude on a slice in an X–Y plane with every second vector plotted in X and Y and every vector plotted in Z. Figure 10(*b*) shows the vector field with every third vector plotted in X and Y and every vector plotted in Z and 0.15 voxels/voxel vorticity iso-surface.

Figure 11 shows the error scatter plots resulting from the difference in velocity components between the reconstructed and exact vector fields ((*a*) and (*b*)) and the synthesized and exact vector fields ((*c*) and (*d*)). Table 2 summarizes the 90% precision interval for this experiment as well. The data again indicate that the largest percentage of the error is due to the PIV algorithm. The reason for the difference in errors in the X and Y dimensions for this case is likely due to the orientation of the vortex ring, which induces larger spatial gradients in the X–Z plane than in the Y–Z plane.

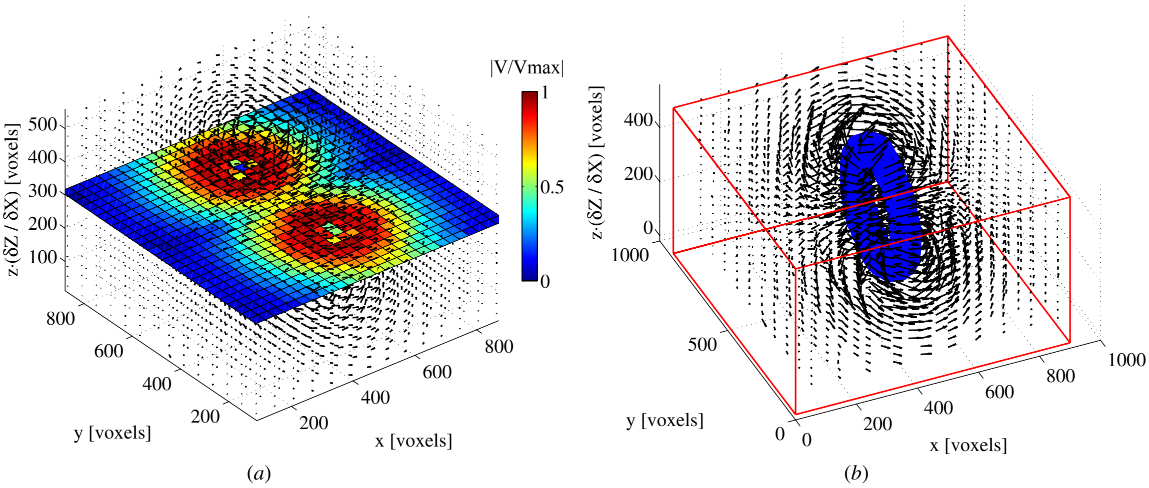

Figure 10. Three-dimensional vector field resulting from PIV processing of the reconstructed intensity volumes using SAPIV in the $40 \times 40 \times 30$ mm^3 volume. (*a*) Slice in an X–Y plane with normalized velocity magnitude contours; (*b*) the vector field and a vorticity iso-surface (0.15 voxels/voxel).

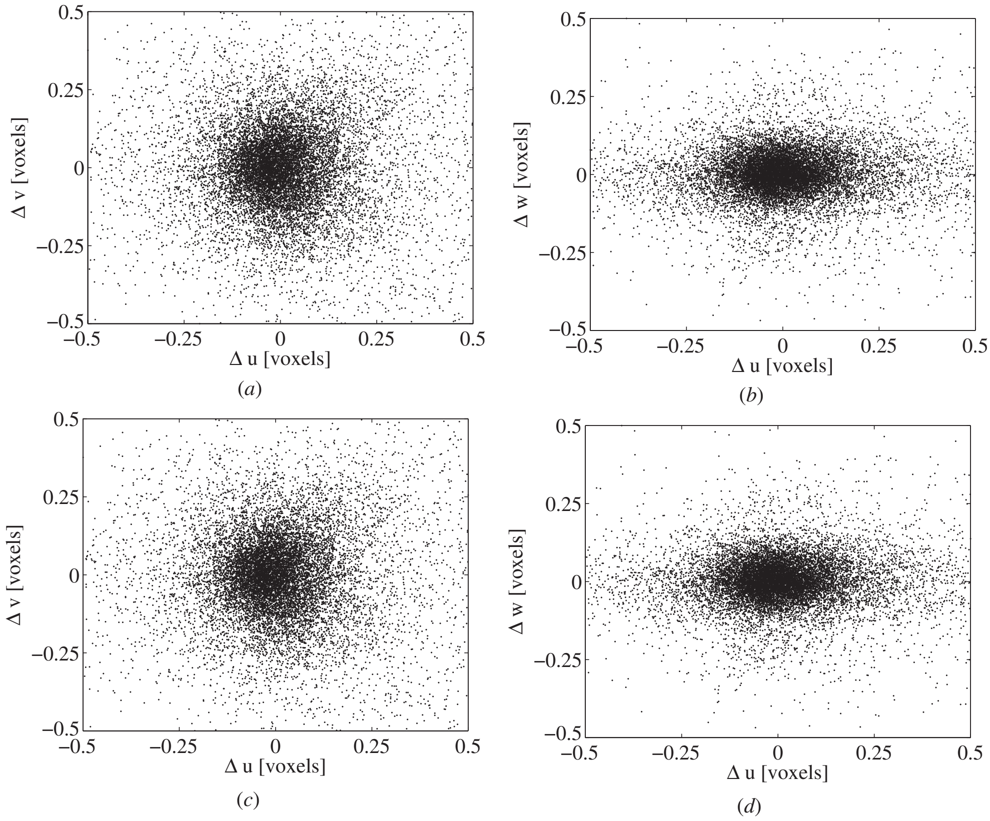

Figure 11. Scatter plots of the error in the reconstructed vector field ((*a*) u, v components and (*b*) u, w components), and error in the synthesized vector field ((*c*) u, v components and (*d*) u, w components). All plots are for the vortex ring in the $40 \times 40 \times 30$ mm^3 volume.

3. SAPIV experimental implementation

3.1. Experimental apparatus

To illustrate the capabilities of the SAPIV technique in practice, a canonical 3D flow field is captured experimentally using an array of eight cameras. Instantaneous 3D SAPIV and classic 2D PIV velocity data of a piston-generated vortex ring are acquired for comparison.

The experiment is conducted in a glass tank ($504 \times 254 \times 280$ mm^3) seeded with neutrally buoyant particles (50 μm diamter) yielding an estimated seeding density of $C = 0.23$ particles mm^{-3}. A cylindrical piston-driven vortex ring generator (40 mm inner diameter) is mounted in the center of the tank. The outlet orifice diameter of the cylinder is 30 mm and the piston stroke length and velocity are 10 mm and 15 mm s^{-1} respectively. Based on the total stroke length

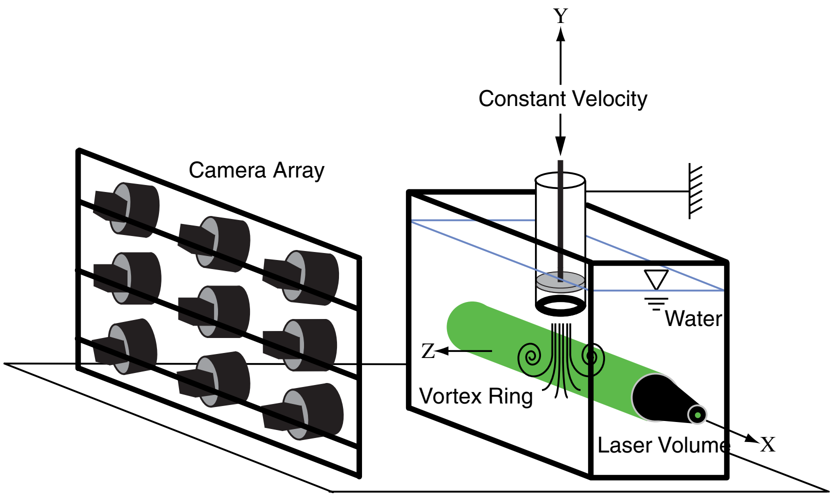

Figure 12. Experimental setup with camera array imaging a vortex ring illuminated with a laser volume.

Figure 13. Photograph of the camera array. Nine cameras are mounted with 50 mm lenses on an extruded aluminum frame. Only eight cameras are used in the study.

(as opposed to integral of piston velocity over time), the ratio of stroke length to outlet orifice diameter is $L/D_o = 0.33$. The camera array images from the side of the tank (figure 12) with all cameras mounted on a common frame. Initially, the array included nine cameras; however, the camera in the upper left corner of figure 12 had outdated firmware that did not support the appropriate triggering and therefore is not used in this study.

The cameras are Point Grey Research, Inc., Flea2 cameras (FL2-08S2M/C), arranged in a rectangular array mounted on 80/20® aluminum rails (figure 13), with ≈150 mm spacing between cameras. Cameras are angled in order to overlap the fields of view. The array is placed approximately 760 mm from the center of the water tank. As the number of cameras in the array increases, the hardware cost associated with the array will increase. Thus, it is important to consider the minimum number of cameras necessary for successful volume reconstruction. Simulations showed that the reconstruction quality reaches a point of diminishing returns as more cameras are added, and that the most efficient number of cameras seems to be in the range of 10–15, as shown in figure 7(*a*).

All eight cameras in the array capture 1024 × 768 pixels, 8 bit, monochromatic images at 10 frames s^{-1}. Each camera is equipped with a Nikon Nikkor 50 mm lens. At the given frame rate, each camera requires 8 megabytes per second (MBps) of bandwidth, totaling 64 MBps to capture the entire event. The cameras use the IEEE 1394b serial bus interface standard which is rated to support 400 MBps for each bus, which provides ample bandwidth for the eight cameras. The cameras are connected to a single computer which records the data onto three hard drives. A custom C++ program using PointGrey's SDK libraries was developed in-house to interface with the cameras. The data are stored on one of three hard drives in AVI files with a raw encoding to maintain information integrity (i.e. no image compression). After completion of the capture sequence, the software reorganizes the data to place all of the necessary files into a single location.

While the cameras record at 10 frames per second (fps), classic PIV frame straddling timing is used to obtain appropriate image pair time spacing. A Litron 532 nm, 180 mJ/pulse, dual cavity Nd:YAG laser is used to illuminate the particle volume. The laser pulse duration is 4 ns, and inter-frame straddling time is 8 ms. A 10× beam expander creates a cylindrical laser volume with a 40 mm diameter, and a 10 mm fanning optic spreads the laser beam into a 1 mm thick laser sheet for the 2D PIV images. A Berkeley Nucleonics Corporation timing box is used to synchronize the cameras and laser. While the focus of this effort is to resolve 3D velocity fields using SAPIV, 2D PIV images of the vortex ring are used for benchmarking and comparison.

3.2. Particle volume reconstruction

To achieve proper focus in synthetic aperture images, accurate mapping between image and world coordinates is required. The mapping is found by imaging a precision machined calibration plate traversed through the target volume with Z location increments of 2 mm. Since the SAPIV technique involves reprojecting the images onto several planes

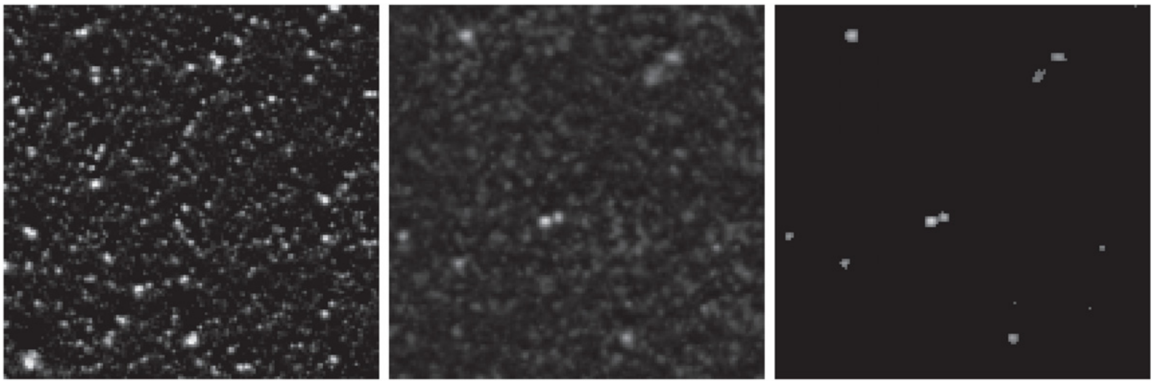

Figure 14. Zoomed views (140 × 140 pixels) of a preprocessed 3D PIV image from single camera of the array (*a*), a refocused image at $Z = 5$ mm (*b*) and thresholded image at the same depth (*c*).

throughout the volume, a suitable coordinate system must be established on the calibration plate. Here, we use the average calibration in pixels mm^{-1} from the center camera image of the plate at the Z location farthest from the cameras to convert the reference geometry of the calibration plate from mm to pixels. Second-order polynomial fits are used to map image coordinates to reference coordinates on each Z calibration plane, and linear interpolation is used to find the polynomial fits on Z planes between each calibration plane. This approach follows that of prior tomographic-PIV studies where polynomial fits are used to deal with the distortion introduced when imaging through an air–glass–water transition [16, 27]. As described in [27], even with higher order fits, the errors in mapping functions can be in excess of 0.5 pixels, but volume self-calibration can be used to reduce the errors to 0.1 pixels or less. As discussed in section 2.4, the error in the mapping functions should be less than 0.45 pixels for adequate reconstruction in SAPIV. Volume self-calibration is not implemented in the present experiment, yet reconstruction still yields a volume which is suitable for 3D PIV. We expect that implementing volume self-calibration will improve greatly the particle yield rate in the reconstruction.

The synthetic aperture images are formed with focal plane spacing of 0.2 mm in Z. Theoretically, the number of focal planes that can be generated within an illuminated volume is infinite, but the information on each plane will not necessarily be unique, as the sharpness of a refocused object is determined by the degree to which mapped images of the object overlap. For example, it is possible to generate focal planes very close to each other by shifting all images by a very small amount. However, if the shift size is much smaller than the refocused object, the two refocused images will be essentially indistinguishable, in which case the information is redundant. Therefore, focal plane spacing (which determines voxel size) should be made large enough so as not to retain redundant information. The depth over which objects are in focus can be controlled by changing the camera baseline, D; for the present experiment, $D = 0.2$. Ultimately, smaller focal plane spacing should yield better resolution in the Z dimension of the reconstructed fields, and thus the vector fields. The influence of focal plane spacing and camera baseline on the accuracy and resolution of 3D PIV vector fields is the subject of ongoing work.

Other volumetric PIV studies have discussed the need for image preprocessing to deal with non-uniformities in laser profiles and pulse intensities, as well as to remove background noise [27, 30]. Prior to refocusing, images in this study are subjected to the following preprocessing steps:

(i) subtract sliding minimum (window size = 10 pixels);
(ii) convolve with 3 × 3 Gaussian kernel;
(iii) equalize histograms to histogram of image with highest contrast;
(iv) increase contrast by trimming the bottom and top 0.1% of intensity values;
(v) subtract sliding minimum (window size = 10 pixels).

After preprocessing, the images are mapped to each plane throughout the volume, averaged to form the synthetic refocused image, and thresholded to retain particles to generate an intensity volume for each time instant. Because the mapping functions are not simple linear homographies, interpolation is required to re-project the images; here, a bilinear interpolation is used. Figure 14 shows a preprocessed image from the array, a synthetic aperture refocused image at one depth ($Z = 5$ mm) and a thresholded image at the same depth. It is estimated from the preprocessed images that the seeding density is 0.026 partices/pixel (ppp) ($C = 0.23$ particles mm^{-3}), and the actual number of particles in the reconstructed volumes is 9863 particles for the first time instant and 9890 particles for the second time instant ($C = 0.13$ particles mm^{-3}). The reason the yield rate is low in the reconstructed volumes is due to errors in the calibration which have not been reduced through the use of volume self-calibration. Images of particles that span two pixels or less in the the original camera images tend to not be properly reconstructed in 3D space because the calibration error can be on the same order as the particle size.

Figure 15(*a*) shows the intensity histogram and Gaussian fit for a refocused image from the eight-camera array experiment. The mean intensity value is very low, but actual particles in the refocused image have high intensity;

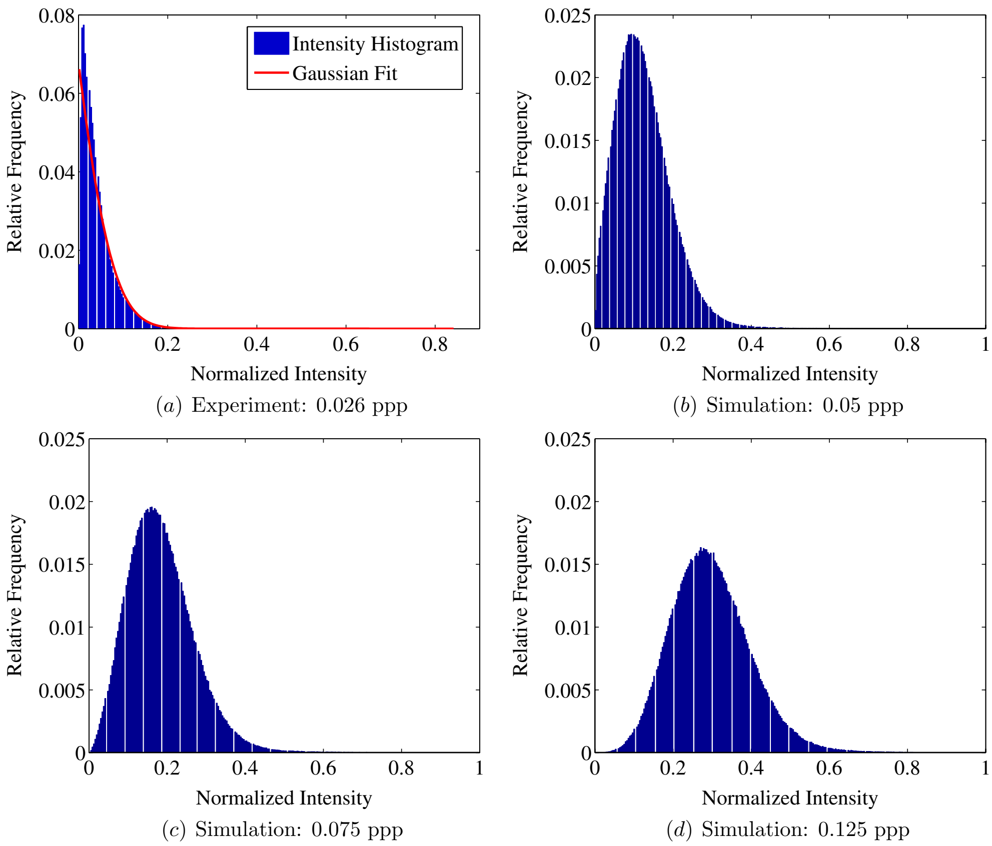

Figure 15. Intensity histograms and Gaussian fit for a refocused image from from the 3D SAPIV experiment with eight cameras (*a*), and intensity histograms for refocused images from simulations in 50 × 50 × 10 mm^3 volume with nine cameras and seeding density of 0.05 ppp (*b*), 0.075 ppp (*c*) and 0.125 ppp (*d*).

therefore, the signal-to-noise ratio in the refocused image is very large and the in-focus particles can readily be determined (figure 14(*c*)). Figure 15(*b*)–(*d*) show the histograms from simulations in a 50 × 50 × 10 mm^3 volume imaged with nine cameras and seeding density of 0.05 ppp (*b*), 0.075 ppp (*c*) and 0.125 ppp (*d*). With increasing seeding density, the mean and standard deviation of the intensity of the refocused images increases, which reduces the signal-to-noise ratio of actual particles. Despite the reduced signal-to-noise ratio, the simulations produced adequate reconstruction of the intensity volume for the 0.075 ppp case (reconstruction quality, $Q = 0.77$) and also provided reasonable reconstruction for the 0.125 ppp case ($Q = 0.71$). The signal-to-noise ratio of actual particles for each of these simulated cases is much lower than for the SAPIV experiment, indicating that seeding density can be greatly increased in future experimental studies.

Once reconstructed, the intensity volumes are ready for cross-correlation-based 3D PIV analysis; the adapted version of matPIV is again employed. A multi-pass algorithm with one pass at an initial interrogation volume size of 128 × 128 × 64 voxels and two passes at an final interrogation volume size of 64 × 64 × 32 voxels and 50% overlap generates a 23 × 31 × 11 vector field. Each 64 × 64 × 32 voxels interrogation volume contains approximately 15 particles. The resultant vector field resolution is 2.1 mm in *X* and *Y* and 3.2 mm in *Z*. Post-processing consists of a filter based on the signal-to-noise ratio of the cross-correlation peak, a global filter which removes vectors five standard deviations above the mean of all vectors, and a local filter which removes vectors which deviate by more than three standard deviations from the median of a 3 × 3 × 3 vector window. The filtered field is interpolated using linear interpolation and smoothed with a 3 × 3 × 3 Gaussian filter. At this point some mention should be made of the overall processing time. The time required to reconstruct the two volumes used to generate the 3D vector field is 18% of the time required for the 3D PIV processing of the fields. Therefore, the limiting time factor in processing is the 3D PIV analysis, which demonstrates the relative efficiency of the synthetic aperture refocusing technique.

Processing of the 2D PIV data is also performed with matPIV; a multi-pass algorithm with one pass using initial interrogation window size of 64 × 64 pixels and two passes with a final window size of 32 × 32 pixels and 50% overlap is used to create a vector field with vector spacing of 1.03 mm. Post-processing for the 2D vector fields also consists of the signal-to-noise, global and local filters, as well as linear interpolation and smoothing with a 3 × 3 Gaussian kernel.

3.3. Experimental results

Experimental results for the instantaneous 3D velocity data of the vortex ring are shown in figure 16(*a*); the resultant 3D vector field and an iso-vorticity contour (magnitude 9 s^{-1}) are plotted at one time instant. For ease of comparison between

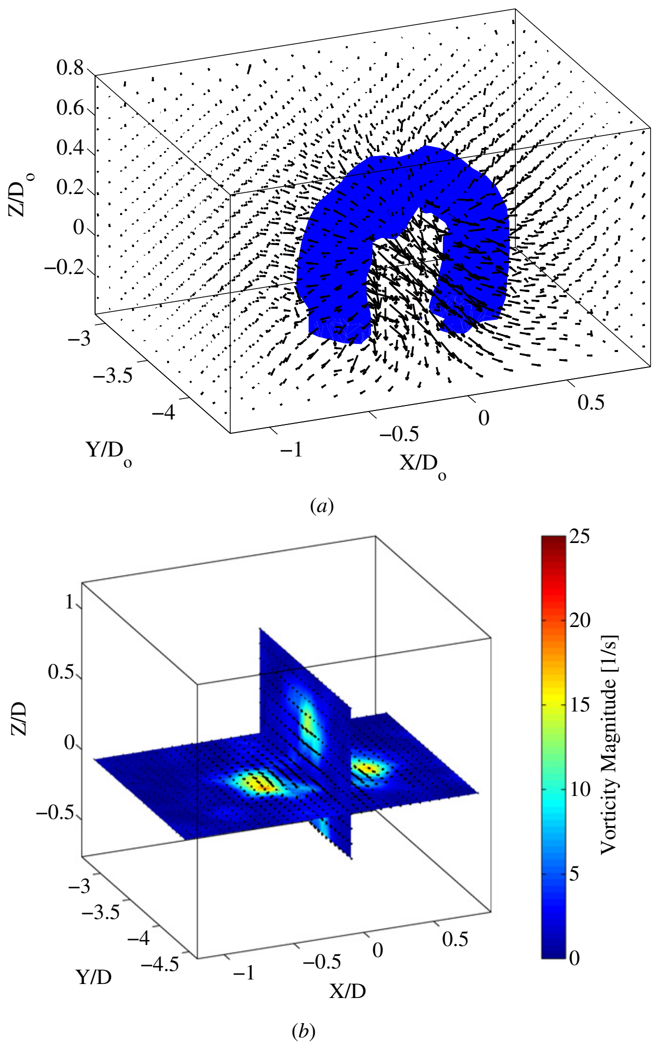

Figure 16. (*a*) Experimental SAPIV velocity vector field for the vortex ring with an iso-vorticity contour of magnitude 9 s^{-1} plotted in blue. (*b*) *X*–*Y* and *Y*–*Z* cross-sectional cuts of vorticity through the vortex ring center, with superimposed velocity vectors.

the 3D SAPIV and 2D PIV data, we normalize all lengths by the orifice diameter ($D_o = 30$ mm), as is done in Weigand and Gharib [31]. The origin of the *X*–*Y*–*Z* global coordinate system is placed approximately at the center of the outlet orifice of the vortex generator, with *Y* positive up and *Z* decreasing in the direction away from the cameras. In figure 16(*a*), the ring has propagated to a distance $Y/D_o \approx 3.72$ below the orifice outlet. Every vector in *Z* is plotted and every second vector is plotted in *X* and *Y* directions. The measured SAPIV volume is only limited by the volume which is illuminated by the laser. The iso-vorticity contour shows an incomplete ring, due to the fact that the ring is not centered in the laser volume and part of the ring is outside of the illuminated region. Cross sectional slices of vorticity, with planes at $Z/D_o = 0.003$ and $X/D_o = 0$, are plotted in figure 16(*b*). As expected, isolated regions of vorticity are located where the ring passes through the planar cut in the 3D volume. The vorticity magnitude on the *X*–*Z* slice is slightly lower than for the *X*–*Y*, which is likely due to the lower resolution of the vector field in the *Z*-dimension. By locating the peaks in the normal vorticity component in each core cross-section on the plane passing through $Z/D_o = 0.003$, the normalized diameter of the vortex is found to be $D/D_o = 0.77$.

Data from the 2D PIV experiments are used to compare and benchmark the 3D PIV results. Figure 17(*a*) shows the velocity and vorticity on an *X*–*Y* plane of the 3D SAPIV data at $Z/D_o = 0.11$, and figure 17(*b*) shows velocity and vorticity resulting from a 2D PIV experiment, where the laser sheet is located at approximately $Z/D_o = 0.15$. The slice of the 3D data at $Z/D_o = 0.11$ is the location in the vector field nearest to the 2D laser sheet. Although the 2D and 3D experiments could not be performed simultaneously, the experimental vortex ring generator is designed to be repeatable, and the vortex ring is at approximately the same *Y*-location for the 2D and 3D vector fields shown in figure 17. Qualitatively, the topology of the vortex ring compares well between the slice from the 3D and 2D data. The distance between the positive and negative cores of the vortex ring is $D/D_o = 0.82$ for the 2D data, and $D/D_o = 0.77$ for the 3D plane at $Z/D_o = 0.11$.

For quantitative comparison, the velocity, vorticity and circulation are examined and compared for the 3D and 2D data. In figure 18(*a*), u and v velocity profiles from the 2D vector field are plotted against X/D_o at $Y/D_o = -3.75$. The profiles show good qualitative agreement with those presented in [31]. In addition to the 2D profiles, two profiles from the 3D data are plotted. The u and v velocity profiles are plotted against X/D_o at $Y/D_o = -3.72$ and $Z/D_o = 0.11$ (*X*–*Y* plane), and the v and w velocity profiles are plotted against Z/D_o at $Y/D_o = -3.72$ and $X/D_o = 0$ (*Y*–*Z* plane). The *Y*–*Z* velocity profile contains no data for $Z/D_o < -0.32$, due to the fact that the ring is not entirely in the illuminated volume and thus not resolved beyond $Z/D_o < -0.32$. Profiles from the 3D data show good quantitative agreement with the 2D profiles. The maximum negative v velocity from the 3D *X*–*Y* and *Y*–*Z* profiles are 6% and 8% below the maximum negative v velocity for the 2D profile, respectively. The profiles from the 3D data also capture the sign reversal in v velocity as the profile moves through the core, and the spatial decay in velocity moving away from the core.

Along each of the same profiles, the component of vorticity normal to the plane containing the profile is plotted in figure 18(*b*). The trends in all profiles agree very well, but the vorticity calculated from the 3D data underestimates the peak magnitude as compared with the vorticity from the 2D data. This is likely due to the lower spatial resolution in the vector fields (2D data were processed with smaller interrogation windows), which results in a spatial averaging of velocity gradients during the cross-correlation calculation. Also, the smoothing implemented in the post-processing of the vector field is more likely to remove gradients in the 3D data, because the Gaussian kernel is three dimensional, and thus the smoothing is based on more neighboring vectors.

Finally, to serve as another quantitative measure for benchmarking the 3D SAPIV system the circulation, Γ, is calculated on a variety of planes. The circulation is calculated by taking the line integral of velocity around a rectangular contour on a particular plane. Here, we calculate the circulation in the 2D data, as well as on several planes for

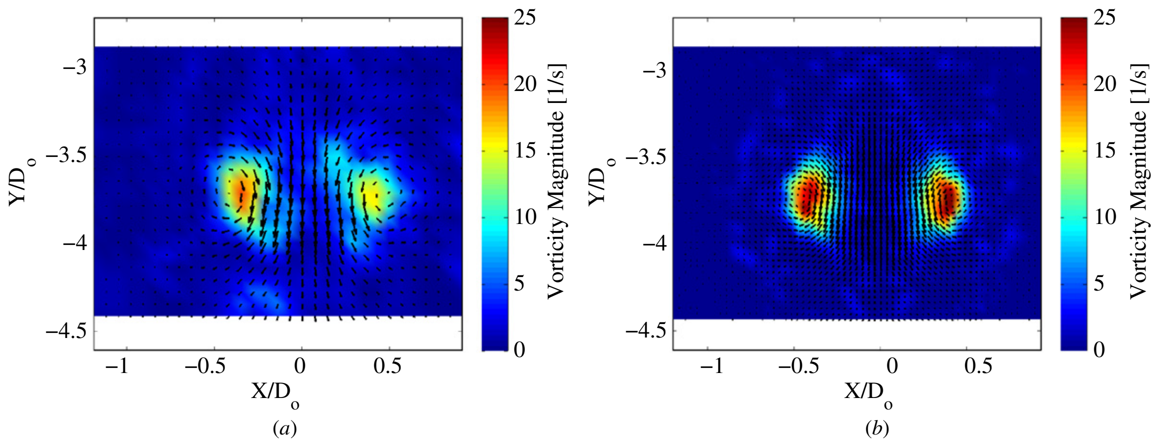

Figure 17. Vorticity contours and vector fields from (*a*) 3D PIV cut and (*b*) 2D PIV instantaneous slice through the vortex ring.

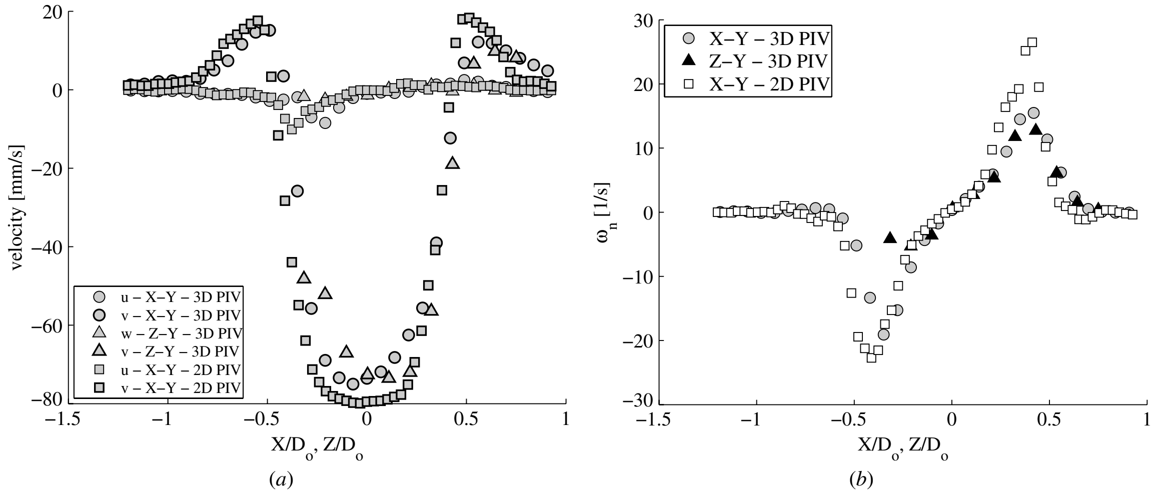

Figure 18. (*a*) Velocity profiles (u, v and w) along axes passing through the center of the vortex ring for 2D and 3D cuts. (*b*) Vorticity profiles along same planes used in (*a*) for 2D and 3D cuts normalized by the diameter of the vortex generator orifice (D_o).

the 3D data, using rectangular contours of increasing size to encompass up to one half of any 2D slice (i.e. encompassing one vortex core). Each contour is centered around the location of maximum surface-normal vorticity on the plane under consideration. The inset of figure 19 shows the location and angle of each plane on which circulation is computed. For the 3D data, nine planes are chosen with angles varying between 0° and 180°, as well as one additional plane located at $Z/D_o = 0.11$, on which the circulation is calculated for both cuts through the vortex ring core. Figure 19 shows the circulation plotted against the area enclosed by the integration contour for each plane.

For a symmetric vortex ring, the circulation on half of any one plane containing the axis passing through the ring center should be constant. From figure 19, it can be seen that the magnitude of circulation remains relatively constant regardless of plane angle, for a given integration contour size. The maximum difference in peak circulation is 2.69 cm^2 s^{-1} (13.5% of maximum) for the angled planes. On the plane passing through $Z/D_o = 0.11$, as well as for the 2D data (laser plane at $Z/D_o = 0.15$), the peak circulation is reached at a smaller integration contour size than for the planes passing

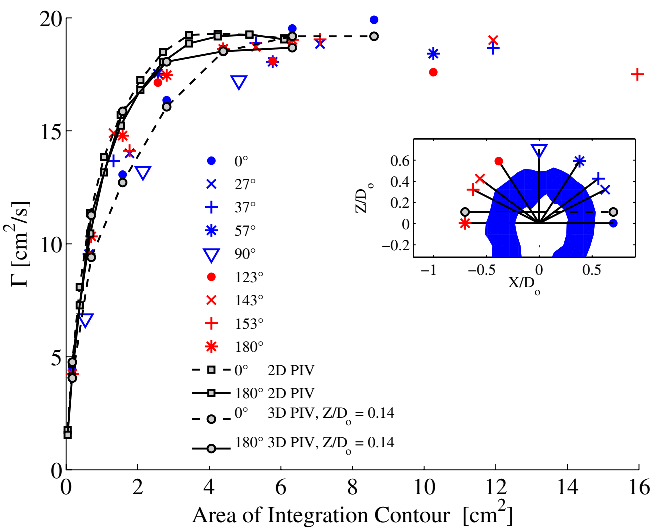

Figure 19. Distribution of circulation as a function of area enclosed by integration contour on planar slices for several angles of the vortex ring. Planes of interest and their positions are shown in the inset and labeled in the legend.

Table 3. Comparison of efficiency of some 3D PIV techniques.

Technique	Volume size (mm^3)	No of cameras	No of particles	Particles/ pixel	Particles/ total pixels #
SAPIV (simulation)	50 × 50 × 10	13	125 000	0.125	0.01
SAPIV (experiment)	65 × 40 × 32	8	9 860	0.015	0.002
Tomo-PIV [14]	35 × 35 × 7	4	24 500	0.05	0.013
Tomo-PIV [15]	37 × 36 × 8	4	114 500	0.08	0.02
Tomo-PIV [16]	100 × 100 × 20	4	100 000	0.024	0.006
HPIV [32]	10 × 10 × 10	1	2 000	0.0015	0.0015
HPIV [22]	1.5 × 2.5 × 1.5	1	56 250	0.014	0.014
DDPIV [8]	100 × 100 × 100	3	40 000	0.038	0.013
DDPIV [33]	150 × 150 × 150	3	35 000	0.034	0.011
PTV [3]	200 × 160 × 50	3	3 000	0.0038	0.0013

through the central ring axis, because the bisected cross-sections are offset from the center of the ring. Nonetheless, the maximum circulation magnitude on the offset planes for both the 2D and 3D data is within 7% of the maximum found on the planes which pass approximately through the central axis of the ring.

The quantitative agreement between the 3D SAPIV data and the 2D PIV data confirms the viability of the SAPIV technique for making accurate measurements in 3D volumes. Although the simulations show the ability to reconstruct very densely seeded fields, the seeding density was kept rather low (0.026 ppp in the raw images) in this experiment to ensure proper reconstruction. However, it is expected that increased seeding density can be achieved in practice.

4. Conclusion

Synthetic aperture PIV offers a novel and exciting method for imaging complex 3D flow-fields with high seeding densities and partial occlusions. SAPIV draws on the concept of light field imaging, which involves sampling and encoding many rays from a 3D scene, and is practically implemented with an array of cameras. Recombining the camera array images using synthetic aperture refocusing provides many desirable capabilities for 3D fluid flow measurement; namely, the ability to digitally refocus on isolated planes post-capture, to effectively 'see-through' partial occlusions by exploiting the multiple camera viewpoints and to capture volumetric information at a single time instant. We expect the capabilities of the synthetic aperture system to be flexible enough to measure in other flow environments, such as multi-phase and bubbly flows or flows with partial obstructions.

Simulations showed that a single array arrangement allowed for measurement within volumes with depth ranging from 10 mm to 50 mm. Altering the optics on the cameras enables further scalability of the measurement range, as was shown in the simulation of the 100 × 100 × 100 mm^3 volume. In this manner, the behavior of the camera array is similar to the behavior of a single-lens camera: we have control over the viewable depth for a given magnification and can change the FOV by changing the magnification. Two simulated flow fields demonstrated the performance of the technique in resolving vector fields with high resolution and in a relatively large volume. The focal plane spacing of the system in the Z dimension, which is related to the depth of field, was theoretically derived for the simple model of two coplanar image sensors. The observed focal plane spacing in the simulations agreed extremely well with that predicted by the theory, despite the fact that the camera image sensors in the simulated model were not coplanar (the cameras were angled). This shows that the concise theory derived is an accurate and useful tool for predicting the depth of field of the refocused images as a function of camera baseline and optics.

The results of the 3D PIV experiment indicate that SAPIV is a viable technique for efficiently and accurately resolving a range of 3D flow fields. In practice, the hardware implementation successfully captured an instantaneous 3D velocity vector field for a vortex ring, with only eight cameras, in a volume with an aspect ratio (Z:X–Y) that is comparable to some of the largest found in the literature. 3D SAPIV results compared well with the 2D PIV experimental data for a similar vortex ring. The signal-to-noise ratio of actual particles for each of the simulated cases was much lower than for the SAPIV experiment, indicating that seeding density can be greatly increased in future experimental studies, which will allow for increased vector resolution.

The simulations indicate that synthetic aperture PIV can resolve particle fields with seeding density and volume sizes competitive with the state of the art; however, we have made no mention of the efficiency with which the system performs. We have reported the number of particles per pixel based on the number of pixels in one camera, perhaps a measure of efficiency is to divide the particle number by the *total* number of pixels (e.g. 13 cameras × 1000 pixels × 1000 pixels). Table 3 compares this measure of efficiency for several volumetric studies, including the $D = 0.4$ configuration for the 50 × 50 × 10 mm^3 simulation using 13 cameras and the 8 camera experiment from this study, as well some representative results from tomographic-PIV, holographic-PIV, defocusing DPIV and PTV. The comparison is not intended to give a concrete answer as to the most efficient method, rather to test whether the efficiency of the synthetic aperture PIV technique lies in the realm of other techniques, as it does.

Regardless of the notion of efficiency, performance is the bottom line for a 3D PIV system, if the cost is acceptable. Thus, a practical consideration that naturally arises is how to deal with such a high number of cameras needed for synthetic aperture PIV. Certainly, with a large number of cameras the cost of a system becomes a concern. We have developed an eight camera array which performs double-pulsed frame straddling 3D PIV at 10 frames per second with a total system cost (excluding laser) of less than \$15 000, in present-day dollars. An array of high speed cameras for fully time-resolved experiments will be more costly, but SAPIV offers the ability to trade-off individual camera sensor size for a technique that synthesizes smaller sensors, potentially making a high-speed system affordable for 3D PIV. Another form of cost is the required computation time; all data processing was performed on a Macintosh Power Mac G5 computer with a 3 GHz Quad-Core Intel Xeon Processor and 2 GB of RAM. Reconstruction and 3D PIV analysis was implemented in Matlab, and the codes are not optimized at this point. However, the computation time to implement the map-shift-average algorithm, refocus and threshold the images, and assemble them into the reconstructed volume for two timesteps in the simulated $40 \times 40 \times 30$ mm^3 volume required 15% of the time taken to compute the vector fields with three passes and 50% overlap (67 minutes to reconstruct the two timesteps, and 446 minutes to perform the PIV processing). For the SAPIV experiment, the time required to reconstruct the two volumes used to generate the 3D vector field (62 minutes) was 18% of the time required for the 3D PIV processing of the fields (414 min). This attests to the relative simplicity of the refocusing algorithm. Therefore, the actual 3D PIV calculations will dominate the computation time for synthetic aperture PIV.

In order to fully realize the capabilities of SAPIV, further work is required to address several practical issues and challenges, and is ongoing. For example, volume self-calibration can greatly improve the image reconstruction quality and allow for increased seeding densities. The ability of the SAPIV technique to reconstruct the intensity fields without the use of volume self-calibration in this study underscores the capability of the method. In addition, increasing the camera baseline spacing is expected to increase Z resolution. By increasing the baseline, the depth of field can be reduced allowing for more distinction between particles in the Z direction, which we expect will yield higher resolution in Z. In practice, this requires further investigation to determine acceptable minimum and maximum baseline spacing limits, as well as the maximum flow volume size that can be resolved for a given baseline spacing.

Ultimately, synthetic aperture PIV (SAPIV) provides a new and novel method for 3D quantitative flow velocimetry, which offers the ability to reconstruct very dense flow fields in relatively large volumes for a wide range of applications.

Acknowledgments

The authors would like to thank Dr Ramesh Raskar for introducing us to the subject and for many fruitful discussions and Roderick R LaFoy for his help with initiating the project. We gratefully acknowledge funding provided by the Office of Naval Research (ONR), Contract N00014-09-1-1167 managed by Steven J Russell, and ONR In-house Laboratory Independent Research funds, as well as the support of the Naval Undersea Warfare Center, Newport, RI.

References

[1] Adrian R J 2005 Twenty years of particle image velocimetry *Exp. Fluids* **39** 159–69

[2] Raffel M, Willert C E and Kompenhans J 1998 *Particle Image Velocimetry: A Practical Guide* (Berlin: Springer)

[3] Maas H-G, Gruen A and Papantoniou D 1993 Particle tracking velocimetry in three-dimensional flows *Exp. Fluids* **15** 133–46

[4] Maas H-G and Gruen A 1995 Digital photogrammetric techniques for high-resolution three-dimensional flow velocity measurements *Opt. Eng.* **34** 1970–6

[5] Malik N A, Dracos Th and Papantoniou D A 1993 Particle tracking velocimetry in three-dimensional flows *Exp. Fluids* **15** 1432–114

[6] Willneff J and Gruen A 2002 A new spatio-temporal matching algorithm for 3D-particle tracking velocimetry *9th Int. Symp. on Transport Phenomena and Dynamics of Rotating Machinery (Honolulu, HI, February 2002)*

[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] Pereira F and Gharib M 2002 Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows *Meas. Sci. Technol.* **13** 683–94

[9] Yoon S Y and Kim K C 2006 3d particle position and 3d velocity field measurement in a microvolume via the defocusing concept *Meas. Sci. Technol.* **17** 2897–905

[10] Fenimore E E and Cannon T M 1978 Coded aperture imaging with uniformly redundant arrays *Appl. Opt.* **17** 337–47

[11] Kajitani L and Dabiri D 2005 A full three-dimensional characterization of defocusing digital particle image velocimetry *Meas. Sci. Technol.* **16** 790–804

[12] Grothe R L and Dabiri D 2008 An improved three-dimensional characterization of defocusing digital particle image velocimetry (DDPIV) based on a new imaging volume definition *Meas. Sci. Technol.* **19** 065402

[13] Tien W-H, Kartes P, Yamasaki T and Dabiri D 2008 A color-coded backlighted defocusing digital particle image velocimetry system *Exp. Fluids* **44** 1015–26

[14] Elsinga G, Scarano F, Wieneke B and van Oudheusden B 2006 Tomographic particle image velocimetry *Exp. Fluids* **41** 933–47

[15] Elsinga G, Wieneke B, Scarano F and Schröder A 2008 Tomographic 3D-PIV and applications *Particle Image Velocimetry (Topics in Applied Physics* vol 112*)* (Berlin: Springer) pp 103–25

[16] Scarano F and Poelma C 2009 Three-dimensional vorticity patterns of cylinder wakes *Exp. Fluids* **47** 69–83

[17] Micahealis D, Wieneke B, Tiedemann C and Henke R 2008 Visualization of small three dimensional turbulent structures using thin volume tomographic PIV *13th Int. Symp. on Flow Visualization (July)*

[18] Hinsch K D 2002 Holographic particle image velocimetry *Meas. Sci. Technol.* **13** R61–72

[19] Goodman J W 2005 *Introduction to Fourier Optics* 3rd edn (Greenwood Village, CO: Roberts & Company)

[20] Meng H, Pan G, Ye P and Woodward S H 2004 Holographic particle image velocimetry: from film to digital recording *Meas. Sci. Technol.* **15** 673–85

[21] Zhang J, Tao B and Katz J 1997 Turbulent flow measurement in a square duct with hybrid holographic PIV *Exp. Fluids* **23** 373–81

[22] Sheng J, Malkiel E and Katz J 2008 Using digital holographic microscopy for simultaneous measurements of 3d near wall velocity and wall shear stress in a turbulent boundary layer *Exp. Fluids* **45** 1023–35
[23] Isaksen A, McMillan L and Gortler S J 2000 Dynamically reparameterized light field *SIGGRAPH '00: Proc. 27th Ann. Conf. on Computer Graphics and Interactive Techniques* (New York: ACM Press/Addison-Wesley) pp 297–306
[24] Vaish V, Garg G, Talvala E, Antunez E, Wilburn B, Horowitz M and Levoy M 2005 Synthetic aperture focusing using a shear-warp factorization of the viewing transform *Proc. 2005 IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR05')—Workshops (June)* vol 3 (Los Alamitos, CA: IEEE Computer Society Press) p 129
[25] Wilburn B, Joshi N, Vaish V, Talvala E-V, Antunez E, Barth A, Adams A, Horowitz M and Levoy M 2005 High performance imaging using large camera arrays *ACM Trans. Graph.* **24** 765–76
[26] Vaish V, Wilburn B, Joshi N and Levoy M 2004 Using plane + parallax for calibrating dense camera arrays *Proc. 2004 IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR04') (June–July 2004)* vol 1 (Los Alamitos, CA: IEEE Computer Society Press) pp 2–9
[27] Wieneke B 2008 Volume self-calibration for 3D particle image velocimetry *Exp. Fluids* **45** 549–56
[28] Hartley R and Zisserman A 2000 *Multiple View Geometry in Computer Vision* (Cambridge, UK: Cambridge University Press)
[29] Sveen J K 2004 An introduction to MatPIV v.1.6.1 *Eprint no.* 2, Department of Mathematics, University of Oslo, http://www.math.uio.no/~jks/matpiv
[30] Elsinga G, Wieneke B, Scarano F and van Oudheusden B 2005 Assessment of tomo-piv for three-dimensional flows *6th Int. Symp. on Particle Image Velocimetry*
[31] Weigand A and Gharib M 1997 On the evolution of laminar vortex rings *Exp. Fluids* **22** 447–57
[32] Salazar J P L C, De Jong J, Cao L, Woodward S H, Meng H and Collins L R 2008 Experimental and numerical investigation of inertial particle clustering in isotropic turbulence *J. Fluid Mech.* **600** 245–56
[33] Pereira F and Gharib M 2004 A method for three-dimensional particle sizing in two-phase flows *Meas. Sci. Technol.* **15** 2029–38