



AFRL-OSR-VA-TR-2013-0619

**DEVELOPMENT OF A COMPACT AND EASY-TO-USE 3-D
CAMERA FOR HIGH-SPEED TURBULENT FLOW FIELDS**

BRIAN THUROW

AUBURN UNIVERSITY

**12/05/2013
Final Report**

DISTRIBUTION A: Distribution approved for public release.

**AIR FORCE RESEARCH LABORATORY
AF OFFICE OF SCIENTIFIC RESEARCH (AFOSR)/RSA
ARLINGTON, VIRGINIA 22203
AIR FORCE MATERIEL COMMAND**

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

Development of a Compact and Easy-to-Use 3-D Camera for High Speed Turbulent Flow Fields

Final Performance Report for AFOSR Award No. FA9550-10-1-0476

Prepared by: Dr. Brian S. Thurow

Abstract

A methodology, termed plenoptic particle image velocimetry (PIV), was developed and demonstrated for the measurement of three-dimensional, three-component velocity fields in high-speed turbulent flow fields. The concept is based on light-field imaging of particles where a plenoptic camera is used to measure both the position and angle of light rays scattered by particles contained in the flow field. Tomographic algorithms are applied to plenoptic camera data to determine the 3D position of particles contained in the flow field and cross-correlation algorithms are used to determine the displacement of particles in two successive images. A prototype camera was designed, built and used to demonstrate the viability of the technique in practice. A synthetic image generation tool was developed to simulate the camera and investigate the technique's accuracy. Particle positions are found to be accurate within 0.1 mm in the lateral directions and within 0.3 mm in the depth direction for a volume with dimensions 30 mm x 20 mm x 20 mm. Demonstration experiments included 3D velocity measurements of a turbulent boundary layer formed on a wind tunnel wall and a heated, supersonic axisymmetric jet. Overall, plenoptic PIV is shown to be a viable 3D flow measurement technique with its main strengths being its compact form factor, simple experimental arrangement, limited need for optical access and overall ease of use.

Table of Contents

Abstract.....	1
I. Introduction.....	4
A. Current 3D PIV techniques.....	4
B. Light Field Imaging.....	6
II. The Plenoptic Camera.....	9
A. Prototype Camera.....	10
III. Light Field Rendering.....	13
A. Two-plane parameterization.....	13
B. Building the Light Field.....	13
C. Computational Refocusing.....	15
D. Perspective Shift.....	16
IV. Synthetic Image Generation.....	17
A. Overview.....	17
B. 1D Simulations.....	20
C. 2D Simulations.....	21
V. Tomographic Reconstruction.....	22
A. Basic Concept.....	22
B. Calculation of Weighting Function.....	23
C. Implementation & Computational Considerations.....	29
D. Sample Reconstruction Results.....	30
VI. Experiments with Synthetic Images.....	32
A. Single Particle Reconstructions.....	32
B. Cross-correlation Algorithm.....	34
C. Uniform Flow Field.....	35
D. Oseen Vortex.....	37
VII. Experimental Assessment.....	39
A. Turbulent Boundary Layer on a Wind Tunnel Wall.....	39

B.	Heated, Supersonic Jet	40
VIII.	Frequency-Domain Deconvolution-Based Volume Reconstruction	44
A.	Overview	44
B.	Imaging Model	44
C.	Frequency-Domain Deconvolution	45
D.	Deconvolution Applied to Simulated Plenoptic Data	47
E.	Artifacts and Improvements	53
1.	Banding in the object image response	53
2.	PSF Asymmetry	55
F.	3-D Fourier Slice Refocusing	58
G.	2.4 Iterative Algorithm	62
1.	Iterative Gradient Descent Applied to Plenoptic Imaging	62
H.	Fourier Projection-Slice Theorem	63
1.	2-D/3-D of the Fourier Projection-Slice Theorem	63
I.	Future Work	66
IX.	Conclusions & Future Work	67
X.	Publications	69
A.	Theses	69
B.	Journal Publications	69
C.	Conference Papers & Other Publications	69
XI.	References	70

I. Introduction

Particle image Velocimetry (PIV) is a well-established measurement technique used extensively to measure the velocity field in a variety of flow environments. As PIV is an image based technique, the measurements traditionally have been limited to two components (2C) of velocity measured across a two-dimensional (2D) plane. Ideally, however, researchers desire three dimensional (3D), three component (3C) velocity field instantaneously, which is important for quantifying the topology and extent of coherent flow structures that 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.

A. Current 3D PIV techniques

The desire for 3D/3C velocity measurements has led to a number of efforts being made over the years. Advances such as stereoscopic PIV [1] extend traditional PIV to allow 3C measurements within a 2D plane, and dual plane stereoscopic PIV [2] 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, the out-of-plane spatial resolution is much lower than the in-plane resolution. For this reason, these techniques are not considered truly three dimensional. An additional extension of the aforementioned techniques is scanning PIV [3], where high-repetition-rate laser 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 [4, 5] has potential to improve scan rates; however, the complexity and expense of the laser and camera systems are currently too prohibitive for broad application.

Four techniques that have recently received attention for their ability to conduct 3D, 3C measurements are defocusing PIV [6], holographic PIV [7-10], tomographic PIV [11, 12], and synthetic aperture PIV [13]. The former is based on the use of specialized apertures near the camera lens or multiple cameras which eliminates the ambiguity in particle depth that occurs when a particle is not located within the focal plane. Computational algorithms use the knowledge of the aperture shape or camera positions to determine the particle position and depth. The strength of this technique is the relative simplicity of the equipment required and ease of analysis; however, the particle density is significantly limited since the location of individual particles must be resolved. Also, in the case of a single camera system, the use of an 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 volume 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. Traditional holographic PIV setups utilize specialized holographic films enabling the hologram to be densely sampled and a large number of vectors to be generated. At the same time, the use of films is a disadvantage due to the time-consuming reconstruction and wet processing steps required. Incidentally, progress has been made recently in digital holographic PIV using CCD sensors and digitized reconstruction algorithms, most notably a study on wall-bounded turbulence [Sheng_2006]. 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 [12]. 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 volume, 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 [14] cylinder wakes [15], and shock-wave/turbulent boundary layer interactions [16]. It has also been adapted to kHz rates using high-speed cameras for aeroacoustic studies (see Violato et al. [17, 18]. 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), limited particle number density, complexity of the experimental arrangement, and the expense of the overall system. Nonetheless, tomo-PIV's success in obtaining 3D, 3C velocity measurements in a multitude of facilities has revitalized recent research in 3D flow diagnostics.

Synthetic aperture PIV (SAPIV) is another multi-camera 3D PIV techniques described by Belden et al. [13]. The 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 used as the input to cross-correlation algorithms. The technique is limited by many of the same restrictions as tomographic PIV and uses an even greater number of cameras.

B. Light Field Imaging


Light field imaging encompasses methods for measuring the intensity and direction of light propagating through space, which within the past decade has gained attention due to advances in camera technology and computing resources. 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 is given; however, the reader is encouraged to consult other sources, additional information is available in the works of such as Adelson et al. [19, 20], Levoy et al. [21-23], Ng et al. [24], and Georgiev [25-27] for more detailed information.

Historically, the notion of a light field is over a century old as outlined in Lippmann [28]. The modern definition of a light field comes from Adelson and Bergen [19] where the intensity and direction of light rays is parameterized by the *plenoptic function*. Each light ray is represented by its 3D position in space (x, y, z) and its angle of propagation (θ, ϕ) , thus forming a 5D function¹ representing all light rays traveling through space. Assuming 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 $L_F(x, y, \theta, \phi)$. In this context, a conventional photograph represents a 2D projection of the 4D light field.

Adelson and Wang [20] 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 sensor. In contrast to a conventional camera that collects light across a range of input angles and focuses all angles to individual spatial locations on the sensor, the plenoptic camera focuses all angles onto an array of pinholes or 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 x 500 pixel CCD with a microlens array of 100 x 100 microlenses. This results in a camera with a spatial resolution of 100 x 100 pixels and an angular sampling of 5 x 5.

¹ 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.

As camera and microlens technology has improved the interest in plenoptic cameras has grown. Ng et al. [24] presented 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 also offers a variant of plenoptic camera technology that offers a similar ability to change the focus or perspective of an image after the fact .



a) Near-field

b) In-focus

c) Far-field

Figure 1: Computationally refocused images generated from a single exposure, focused: a) on a figure that is in front of the nominal focal plane, b) at the nominal focal plane, and c) on figure behind the nominal focal plane.



a) Left side of aperture



b) Right side of aperture

Figure 2: Computationally rendered image where the viewpoint of the observer has been changed.

Capturing and altering the light field is not limited to using a plenoptic camera. Levoy [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 center 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 [29] used 48, 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.

More recently, Levoy et al. [23] 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. They consider the effect of wave optics and diffraction in a microscopic environment whereas geometrical optics is sufficient for macroscopic imaging. In addition, a typical microscope objective functions differently than a conventional 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 this work is to utilize the information obtained by a plenoptic camera about the light field to measure 3D particle fields, which have direct applicability to 3D PTV and PIV. This paper first describes the design and construction of a prototype plenoptic camera and the process for building the light field from the data obtained by the camera. The geometry describing the 4D light field is introduced, and its relationship to tomographic algorithms is explored. In particular, the structure of the tomographic weighting matrix for plenoptic cameras is detailed. The 4D light field is then coupled to tomographic algorithms 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 weighting matrix is shown to be an evolutionary step forward from the computational rendering algorithms. The algorithms are tested synthetically to determine their performance in reconstruction as well as their accuracy in obtaining a velocity field. Finally, experimental results are presented as a preliminary gauge of the performance of plenoptic PIV in a practical setting.

II. The Plenoptic Camera

Figure 3 schematically illustrates the fundamental concept of a plenoptic camera by contrasting it with a conventional camera. The 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 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) in object space to a corresponding location on the image plane (x_p, y_p) . In a conventional camera (Fig 3a), there is no mechanism for discriminating the angle of the light incident on a pixel; therefore, all angles are integrated for each spatial position and the angular information is lost. 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, a microlens array is positioned at the image plane with the image sensor shifted back by one the 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 3b, which shows the point-of-view of a single pixel. Neighboring pixels contained under the same microlens are exposed to light at different incident angles as shown in Figure 3c where each color 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 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 captured by the image sensor represents a multiplexed sampling of the 4D light field captured by the camera lens.




Figure 3: Illustration of the differences between a conventional camera and a plenoptic camera in how they sample the light field.

A. Prototype Camera

A prototype plenoptic camera has been constructed using an Imperx Bobcat ICL-B4820 camera equipped with a Kodak KAI-16000 interline CCD image sensor. The choice of an interline CCD is motivated by the need to perform a double exposure similar to traditional PIV cameras. Figure 4a shows a photo of the camera without a lens attached, and a U.S. quarter to provide scale. Immediately the compact design of the camera is evident, compared to the complex arrangements required for tomographic PIV or synthetic aperture PIV.

The microlens array was manufactured by Adaptive Optics Associates and mounted to the camera with a custom designed provided by Light Capture, Inc. The mount was manufactured in-house and allows for precise positioning of the microlens array at a distance of the microlens focal length (500 microns) from the CCD sensor. It 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. An exploded view is shown in Figure 4b. To align the camera, we follow a similar procedure to that outlined in Ng et al. [24] 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

by the camera is displayed on a monitor while adjustments are made to the mount. This is accomplished in a few iterations. The accuracy of the alignment procedure is estimated to be ~ 30 microns. The full parameter list for the CCD and microlens array are shown in Table 1.





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

As an initial test of the camera, a simple scene was setup using Lego Star Wars mini-figures. Three objects were setup, one in the foreground (storm trooper), one in focus (Darth Vader), and one in the background (C-3PO) all spaced approximately 25 mm apart in depth. A raw

image of this scene is shown in Figure 5a. From the raw image it is hard to discern the differences from a conventional photograph as the signal contained under each microlens is naturally integrated by the eye to form an apparent image with a narrow depth-of-field. A close look at the region inside the green square, shown in Figure 5b, shows the individual microlens images from which the angular information can be extracted.



a) Raw image taken with plenoptic camera




b) Close up of raw image, showing individual microlenses

Figure 5: Raw image taken with prototype plenoptic camera.

III. Light Field Rendering

A. Two-plane parameterization

The preceding discussion parameterizes a light ray by its position on the world focal plane and angle of propagation. An alternative is the two-plane parameterization developed by Levoy [21]. Figure 6a describes a light ray by its position (x, y, z) and its angle of propagation (θ, ϕ) . Figure 6b defines a light ray by specifying its point of intersection on two planes separated by a known distance. These points of intersection are labeled (x, y) and (u, v) and are equivalent to the other representations of the light field.



a) Light ray parameterized by its position and angle of propagation. Adapted from Levoy [21] b) Light ray parameterized by a pair of points on two planes. Adapted from Levoy [21].

Figure 6: Different representations of a light ray and the final representation used for the plenoptic camera

The plenoptic camera lends itself to the two-plane parameterization due to it inherently having two physical planes involved in the light-capture process: the microlens plane and the aperture plane, separated by a fixed distance, s_i . As discussed previously, the microlenses discretize the spatial location of all incoming light rays, while the aperture fixes the domain over which the angular information is captured. 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.

B. Building the Light Field

The recorded light field can be fully described through determination of the (x, y, u, v) position of each sensor pixel. For experimentally obtained images, the exact positions of the microlenses relative to the image sensor are not known. A calibration procedure was developed to determine the positions of the microlenses, and the pixels beneath them. A calibration image of a diffuse illuminated surface is acquired with the aperture of the camera minimized (i.e.

increasing the f-stop to its maximum value). Note that the calibration is defined for a specific value of the main lens focus; variations in the main lens focus lead to a shift of the position of the microlens images on the sensor. A sample calibration image is shown in Figure 7a. The white dots are the centers of the reduced aperture image formed by each microlens. In terms of the two-plane parameterization these dots represent the center of the aperture (x, y, u_0, v_0) . The location of each microlens is roughly estimated using a 2-D peak finder and refined to sub-pixel accuracy by calculating the 2-D intensity centroid. This is indicated in Figure 7b where the centroid is shown as a green "x".




Figure 7: Example calibration (synthetic) image and corresponding centroid fit.

The calibration procedure then 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 provided specifications, the pitch of the microlens is subject to a $\pm 3\%$ non-cumulative error. Based on this assumption, the (x, y) values for each pixel are assigned as corresponding to the position of the microlens in front of them. To determine the angular coordinates (u, v) for each pixel, the distance between the microlens array and main imaging lens, s_i , must be determined. This is determined by taking an image of a scale at the nominal focal plane, allowing for the calculation of the nominal magnification of the system. From the definition of magnification and the focal length of the main lens, the distance s_i can be determined. The magnification is used to convert the pixel coordinates of the microlens images to physical coordinates on the main lens aperture plane. In pixel coordinates, the u and v values are taken as the distance between the pixel and the center of the microlens. This expression is given for the u component by

$$u_i = (x - x_i) \frac{p_p s_i}{f_l} \quad (1)$$

where the subscript i represents the current pixel. A similar expression using y values is used for the v component. In this fashion, each pixel on the camera is assigned a unique (x, y, u, v) coordinate and the intensity of the light recorded at the pixel value corresponds to the irradiance of the light ray defined by the (x, y, u, v) coordinate.

C. Computational Refocusing

An introduction into light field manipulation is rendering a focused image of the light field at a shifted focal plane from the focal plane in which the image was originally acquired. This process, termed computational refocusing, has been adapted from the work of Ng [Ng] and relies on the two-plane parameterization of the light field.

Rendering a 2D image from the 4D light field involves selecting a subset of rays from the complete 4D light field and integrating the angular dimensions for a pre-determined focal plane. Using the two plane approach an interpolation can be applied to re-sample the light field inside the camera at a virtual image sensor. 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 [Ng]

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'_F can be written in terms of the original light field L_F through a linear projection operator, 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 original light field. Mathematically, the location of this projection from x' onto x for a single u value, denoted x_{find} is given by

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

Where $\alpha = s'_i/s_i$. 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

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

To generate a refocused image at the synthetic image sensor, 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

$$I(x', y') = \iint L_F \left(u \left(1 - \frac{1}{\alpha} \right) + \frac{x'}{\alpha}, v \left(1 - \frac{1}{\alpha} \right) + \frac{y'}{\alpha}, u, v \right) du dv \quad (4)$$

In equation 4, the light field is queried at fractional positions. Therefore, a 4D interpolation scheme (in this work, linear interpolation) 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.

D. Perspective Shift

Another benefit of capturing the entire light field is the ability to render scenes from a different angle than the optical axis of the system. 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.

IV. Synthetic Image Generation

Particle imaging is a unique imaging scenario that to date has not been considered by researchers in plenoptic photography. To develop this technique synthetic data is needed to test the overall accuracy of the algorithm. To do this a plenoptic camera simulator has been developed and is detailed herein.

A. Overview

The generation of synthetic plenoptic camera images from given particle positions in object space is nontrivial due to the complex optical configuration of the camera; the mapping function from object space is not only nonlinear, but exhibits step-function behavior as rays impinge on different microlenses. Therefore, a ray-tracing approach is adopted, where linear (Gaussian) optics is used to geometrically trace the path of light through space and the various optical elements that comprise the plenoptic camera. 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 [30]. 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 [Georgiev,Intwala]), was developed to extend the concepts to light field imaging.

The simulator is configured via the variables and relationships as defined in Figure 9. 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 the nominal focal plane of the camera with positive z pointing away from the camera.




Figure 9: Optical configuration of the plenoptic camera.

Particle positions are defined by their position relative to the center 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, shown in equation 3, which makes the assumption that the thickness of the lens is negligible relative to the length of the optical system itself.

$$\frac{1}{s_i} + \frac{1}{s_o} = \frac{1}{f_m} \quad (5)$$

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 principle 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.

$$M = -\frac{h_i}{h_o} = -\frac{s_i}{s_o} \quad (6)$$

In combination with eq. 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 focal length of the main lens.

The optical parameters are now divided into two categories: input and fixed parameters. The former are adapted for specific experimental conditions: main lens focal length, aperture diameter, and magnification. The object and image distances are a result of the the main lens focal length and magnification. The fixed parameters are set through hardware configuration and cannot be modified after assembly: the microlens pitch, microlens focal length, pixel pitch, and the number of pixels.

A constraint on the input parameters is a condition of f-number matching between the main lens and the microlenses. It was recognized by Ng et al. [Ng] 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 image-side f-number, as described by Smith [smith_2007], is shown in equation 5 where, f is the focal length, and d is the diameter of the lens aperture.

$$(f/\#)_m = (f/\#)_l / (1 - M) \quad (7)$$

In this work, we simulate a nominal 1:1 imaging magnification such that $h_i = h_o$ and $M = -1$. The parameters used in the present simulation are shown in Table 1.

Note that the achievable parallax shift is limited by the size of the lens aperture used to form the image and the distance of the object from the main lens. The aperture size is limited by f-number matching condition; f/4 microlenses are used in this work, which allow for a maximum main lens aperture of f/2 at 1:1 imaging conditions.

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

The process of the ray-tracing simulation is shown schematically in Figure 9. For each synthetically generated particle, 10,000 rays are used to simulate the light emanating from that point. The initial angle of each light ray is randomized between θ_{min} and θ_{max} , which are determined based on the distance to the lens and the aperture size. In Figure 9 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 using 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. More details about the simulator can be found in Lynch [31].




Figure 9: Schematic of ray-tracing process for a plenoptic camera.

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. 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.

B. 1D Simulations

Figure 10 shows the simulation of two particles, one displaced by 1 mm behind the nominal focal plane (Fig. 10a) and one displaced 1 mm in front of the nominal focal plane (Fig. 10b). In Figure 10a, all of the light rays converge in front of the microlens plane in a manner which 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 microlens redirect the incident light to different pixels on the image sensor forming a unique image pattern corresponding to the particle positions. Conversely, in Figure 10b, the light rays are intersected by the microlens array prior to reaching their focal point forming a distinctly different image pattern.





Figure 10: 1D simulations at various particle depths. 1 out of every 100 rays shown. Integrated signal shown in blue.

C. 2D Simulations

The image provided is a subset of a full image, whose size is set in accordance with the KAI-16000 image sensor to 4872 x 3248 pixels. This image was generated using a particle volume ranging from $z = -10$ mm to $+10$ mm and a particle density of 0.001 particles per pixel (ppp). For the plenoptic camera the calculation of ppp is simply the number of particles divided the number of pixels on the raw image sensor (4872 x 3248). Upon visual inspection of the image, 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 visually.



V. Tomographic Reconstruction

Tomographic PIV is based on the reconstruction of a volume of particles based on a limited number of 2-D projections of the volume captured by independent cameras. Plenoptic PIV behaves in a similar fashion; a volume of particles is reconstructed by a dense collection of views given by the multiple samples of the angular distribution of light at a position. First, an overview of the MART procedure is given, followed by details necessary for its application to plenoptic data sets.

A. Basic Concept

Tomographic reconstruction of particle fields in fluid measurement is unique in the field of tomography due to the high-frequency content of the reconstructed signal and the limited number of views available. An overview of tomographic algorithms for particle imaging is given by Elsinga [11]. In particular, algebraic reconstruction methods (see Herman and Lent [32]) are well-suited for reconstruction given the above constraints. These methods form a linear system of equations relating the intensity distribution within the reconstructed volume to the projections formed on the camera(s). Thus an inverse problem is formed, where the projections are known, however the reconstructed volume is unknown. In general, the linear system is highly underdetermined, and therefore requires iterative, approximate solution methods.

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

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

where N_i represents the number of voxels in the line-of-sight of the i th pixel. The weighting function $w_{i,j}$ describes the relationship between the recorded image (i th pixel) and the 3D volume of interest (j th 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) 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 to work very well in 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

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

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 i th pixel by raising the argument to 0, therefore multiplying the current voxel by 1, if the voxel is not affected by the i th pixel.

B. Calculation of Weighting Function

In tomo-PIV, the weighting function is calculated based on a projection of a pixel through the volume. The weighting coefficients are calculated by the integral of the pixel line-of-sight and a voxel element (typically via a intersection of a cylinder and sphere; for a review of methods see Scarano [12]). This definition requires that the line-of-sight diameter of the pixel is constant along the volume depth, i.e., the entire volume must lie within the depth of focus of the optical system. An alternative view is that the optical point spread function (PSF) for a tomographic-PIV reconstruction must not exceed the particle image diameter over the illuminated depth.

The plenoptic camera has a more complex PSF due to the microlens arrangement; therefore, this method of calculating the weights is not applicable. A new method for determining the weighting function is proposed to take into account the plenoptic PSF. The approach 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 volume elements. The coordinates in object space are then transformed into image space 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 :

$$s'_o = s_o + z_o \quad (10)$$

$$s'_i = f_m * s'_o / (s'_o - f_m) \quad (11)$$

$$\alpha = s'_i / s_i \quad (12)$$

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:

$$M' = -s'_i/s'_o \quad (13)$$

$$x' = x_o * M' \quad (14)$$

$$y' = y_o * M' \quad (15)$$

where M' is the magnification of the voxel. The result is a voxel in image space whose position is given by (x', y', α) .

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. As shown in Figure 12, a distribution of rays passing through a particular voxel, denoted by x' , 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.




Figure 12: Demonstration of two-plane projection of x' and u in two dimensions.

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 13a where the green "x" is the point where the projection strikes the microlens plane, the blue dots represent the center 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. Expressing (x_2, y_2) in terms of microlens coordinates yields the following relations for the surrounding microlens positions.

$$x_0 = \text{floor}(x_2) \quad x_1 = \text{ceil}(x_2) \quad (16)$$

$$y_0 = \text{floor}(y_2) \quad y_1 = \text{ceil}(y_2)$$

This allows the relative position of the light ray to the neighboring microlens centers to be easily calculated, and it has the benefit of auto-normalizing the coefficient since the separation is equal to one (i.e. $\text{ceil}(x_2) - \text{floor}(x_2) = 1$). Once the interpolation coefficient for the four microlenses have been calculated the u, v interpolation can take place. Figure 13b shows the discretization of the aperture plane as viewed from the pixel behind each 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 centers 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. The surrounding pixel values are given by

$$u_0 = \text{floor}(u_2) \quad u_1 = \text{ceil}(u_2) \quad (17)$$

$$v_0 = \text{floor}(v_2) \quad v_1 = \text{ceil}(v_2)$$




Figure 13: Determination of the weighting function coefficients via linear interpolation.

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 a 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

$$t_x = x_2 - x_0 \quad t_y = y_2 - y_0 \quad t_u = u_2 - u_0 \quad t_v = v_2 - v_0 \quad (18)$$

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

$$N_{0000} = (1 - t_x)(1 - t_y)(1 - t_u)(1 - t_v)$$

$$N_{0001} = (1 - t_x)(1 - t_y)(1 - t_u)(t_v) \quad (19)$$

$\vdots$

$$N_{1111} = (t_x)(t_y)(t_u)(t_v)$$

The result of this procedure can be seen in Figure 14, where the red boarder represents the four microlenses shown in Figure 13 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).




Figure 13: Illustration of the sixteen interpolation coefficients found using the weighting function. The red lines represent the edges of microlenses with the blue dots representing microlens centers.

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 center 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 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 so that the intensity contained in a voxel is conserved. In equation form the normalization process is given by

$$\bar{w}_{i,j} = \frac{w_{i,j}}{\sum_i w_{i,j}} \quad (20)$$

This forces the condition $\sum_i \bar{w}_{i,j} = 1$, 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 center. 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 computation of the entire weighting matrix. As an illustration, consider that each voxel is simulated 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 near the boundaries.




Figure 16: Weighting function comparison to particle simulation.

C. Implementation & Computational Considerations

The implementation of the aforementioned 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 weighing matrix of size 300 x 200 x 200 voxels and using a $f/2$ aperture (a larger aperture increases the u,v sampling) 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 (algorithm is multi-threaded), with a RAID 0 array with 3 128 GB SSDs, the weighting matrix takes 4 hours to complete and the MART algorithm takes 4 hours for the first image and 1 hour per additional image.

D. Sample Reconstruction Results

As an 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 16.







Figure 17: Synthetic raw image

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 particles position. The final reconstruction of the particles is shown in Figure 18 and the true particle positions are shown in Figure 19. Figure 18b shows a front view of the reconstructed volume. When compared to the actual particle positions (Fig 19b) 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 (Figs 18c & 19c) they are shown to match locations, but the reconstructed particles are elongated in depth. This can be attributed to the limited range of angles collected by the plenoptic camera. Fortunately, the intensity in depth is not constant. Figure 20 shows a single reconstructed particle iso-surface as well as a slice through the center of the particle on the YZ plane. The particle has a “hot” center with decreasing intensity at the front and back as shown in Figure 20b. This allows for resolution of the location of the center of the particle in depth, where a constant intensity would create a large ambiguity. The lateral spatial resolution of this particles 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 laterally.



VI. 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 provides a nominal measure of accuracy, as well as defines 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 an Oseen vortex.

A. 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 x 6.125 x 50.125 mm was discretized into a grid of 50 x 50 x 402 voxels, creating cubic voxel elements with sides of length 0.125 mm. For the reconstruction, a relaxation 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, however those are not considered in this work.




Figure 21: Error in reconstruction accuracy via Gaussian fit of 40 particles spaced 1 mm apart

along optical axis

Figure 21a 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 accuracy 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, however 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 21b as a function of depth. In the region near the focal plane $[-10\ 10]$ the error in depth was shown on average 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 in 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, however 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 v. 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.




Figure 22: Scatter plot of X and Z absolute error in reconstruction of 500 simulated particles

B. Cross-correlation Algorithm

This paper has, until this point, discussed a method for obtaining particle fields from an image, however the ultimate 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 a pair of reconstructed particle fields is needed. The method employed here was developed in house and is a cross-correlation based technique, whose implementation is based on Adrian and Westerweel [33] and Scarano and Riethmuller [34]. Briefly, each reconstructed volume pair is divided into several interrogation volumes defined by a size in number of voxels and a percent 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 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 VOLDIM [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 $5 \times 5 \times 5$ neighborhood. If the displacement exceeds a pre-determined threshold (usually 2 [33]), 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 function interpolation function on a $7 \times 7 \times 7$ stencil. The final velocity is calculated as the location of the correlation peak plus the predicted displacement

divided by the time between exposures. In this work four passes are used starting with a window size of $32 \times 32 \times 32$ and ending with a $16 \times 16 \times 16$ window. All passes use a 50% overlap as well as a 2 pass median test validation.

C. Uniform Flow Field

This test involves randomly generating 10 velocity fields with a uniform displacement. The purpose of such a test, is to determine if there is any systematic error in the measurement associated with position. Five of the fields were displaced in the x-direction, and 5 in the z-direction. For each case, the volume contained 25,000 randomly positioned particles corresponding to 0.0023 particles per pixel (according to the definition provided earlier). The number of particles was chosen to correspond to ~ 10 particles per interrogation volume for a volume divided into $16 \times 16 \times 16$ voxel windows. The particles were displaced 1.1682 mm or 11.682 voxels in the x and z directions respectively.

The results of the test cases are presented in Figures 23-26. The data is presented by averaging all 10 test cases into a single result. Furthermore, the results are presented as 2D images where the third dimension is averaged into the result. The first series of plots (Fig 23) display the mean of the absolute error in displacement as a function of lateral spatial location (the XY plane). In this case, the slices of the data along the z-dimension were averaged into the result. The mean of the absolute error shows the systematic error present in the system, and is plotted for each component of velocity in Figure 23. In Figure 23a, the average error in the calculated displacement in the direction of motion (for five cases) is generally less than 0.1 voxels indicating that the systematic error is quite small. The results are even better in the y-direction (Fig 23b, perpendicular to the direction of motion). Figure 23c shows the error in the z-component to be at most 1 voxel, which is similar to the single particle results. The standard deviation, or uncertainty, of the measurement is shown in Figure 24 using the same XY representation. In the lateral directions the uncertainty is generally less than 0.2 voxels indicating a small uncertainty. The z-component shows an uncertainty of approximately 0.8 voxels. Generally speaking the uncertainty should decrease from the single particle case in proportion to $1/\sqrt{N}$, where N is the number of particles averaged. Since each interrogation region consists of roughly 10 particles the uncertainty should decrease from 3 voxels (single particle) to 0.9 voxels. Figures 25 and 26 show the mean and standard deviation as a function of depth and are presented in the YZ plane. The results for the lateral displacements remain unchanged however; the results for the z-component (Figs 25c & 26c) indicate that most of the error occurs in the extremes of the volume in depth. Most of the inner regions of the volume have less than 0.5 voxels mean error and a standard deviation of 0.6 or less. This was to be expected since the reconstruction error increases more dramatically at these outside positions.




Figure 23: Mean of 10 velocity fields shown as a function of lateral spatial location. The three plots shown are for the three components of velocity.




Figure 24: Standard deviation of 10 velocity fields shown as a function of lateral spatial location. The three plots shown are for the three components of velocity.





Figure 25: Mean of 10 velocity fields shown as a function of depth. The three plots shown are for the three components of velocity.



a) X-displacement

b) Y-displacement

c) Z-displacement

Figure 26: Standard deviation of 10 velocity fields shown as a function of depth. The three plots shown are for the three components of velocity.

D. Oseen Vortex

The final synthetic test is of a more complicated flow field: an Oseen vortex. An Oseen vortex was chosen because of its 2D nature, therefore allowing for isolation of one of the components of velocity. The Oseen vortex is described by a maximum tangential velocity, $V_{\theta, \max}$ and a core radius, r_c . The following equation defines the Oseen vortex in cylindrical coordinates

$$V_{\theta}(r) = V_{\theta, \max} \left(1 + \frac{0.5}{\alpha} \right) \frac{r_c}{r} \left[1 - \exp \left(-\alpha \frac{r^2}{r_c^2} \right) \right] \quad (21)$$

where r is the radius to the current location, and a scale factor $\alpha = 1.25643$. For this test a maximum tangential velocity of 1 mm/s or a ~ 10 voxel displacement was used with a core radius of 10 mm. Two flow fields were simulated, one with the vortex rotating about the z-axis (optical axis) highlighting the lateral spatial accuracy of the camera as well as rotating about the x-axis such that the vortex is rotating in depth. For these tests, only one velocity field was obtained for each case and the number of particles and other parameters remain unchanged from the uniform displacement tests. The results are presented in a similar fashion to the uniform displacement case, where the mean of the velocity component is presented. In this case, the mean is taken from a single velocity field, and averaged in the axial direction of the Oseen vortex. In order to test the results an ideal case was generated by creating two volumes from the exact particle solutions, assuming a Gaussian shape for the particles, and running the volumes through the cross-correlation algorithm. This provides a benchmark of the highest possible accuracy for the reconstruction. The results of the z-axis rotation case are shown in Figures 27 & 28 and the results of the x-axis rotation are shown in Figures 29 & 30. The first figure (a) is the reconstructed velocity field, (b) is the ideal simulation, and (c) is the absolute error between the two. Since in an Oseen vortex only two of the components are displaced (axial velocity is zero) only two are shown. It is shown that the velocity field derived from the reconstructions is a very good match for the vortex rotated around the z-axis. The absolute error in both the x and y displacement is shown to be about 1 voxel or less. The w-component is captured; however it is less representative of the true flow-field. This stems from the greater inaccuracy in depth, as detailed previously.




Figure 27: Oseen vortex $\bar{u}$ velocity contours rotating about z-axis




Figure 28: Oseen vortex $\bar{v}$ velocity contours rotating about z-axis




Figure 29: Oseen vortex $\bar{w}$ velocity contours rotating about x-axis




Figure 30: Oseen vortex $\bar{w}$ velocity contours rotating about x-axis

VII. Experimental Assessment

To complement the synthetic image results, physical experiments were conducted with the plenoptic camera in two different facilities. These experiments consisted of 1) wall bounded flow of a turbulent boundary layer and 2) jet flow from a heated, supersonic axisymmetric nozzle. These experiments are simple proof-of-concept demonstrations of plenoptic PIV in different settings.

In addition, experiments were also conducted in a water tunnel along with traditional PIV for the purposes of validation. At the time of this report, the data was still being analyzed and not available.

A. Turbulent Boundary Layer on a Wind Tunnel Wall

The first example is 3D measurement of an incompressible boundary layer with adverse pressure gradient. The boundary layer was formed on the test section wall of an open loop wind tunnel with free stream velocity of ~ 15 m/s. The plenoptic camera was positioned to image the boundary layer through a window looking up in the direction of shear (depth direction of the camera). The flow was seeded through an upstream slit with alumina particles and illuminated with a dual pulse laser outputting 50 mJ/pulse and formed into a 50 mm thick sheet. The Reynolds number based on momentum thickness, Re_θ , was 7,239 and the adverse pressure gradient ($\beta=10.1$) was imposed using a Stratford Ramp mounted on the opposite wall. Figure 31 shows the 3D velocity field determined using the plenoptic camera with dimensions highlighted in the figure. The streamwise velocity is indicated by the color of the vector; the shear of the boundary layer is quite clear and the observed boundary layer thickness qualitatively agrees with that measured using traditional 2D PIV.

Additional details and examples of 3D velocity fields obtained in this flow can be found in Melnick et al. [35]




Figure 31: Preliminary experimental data showing 3D velocity field of a turbulent boundary measured using prototype plenoptic camera. Camera was oriented to look vertically up through the boundary layer illustrating ability to resolve shear along the optical axis.

B. Heated, Supersonic Jet

The second example is taken from experiments recently conducted at the National Center for Physical Acoustics (NCPA) located at the University of Mississippi. These experiments were conducted as a proof-of-concept of the technique's viability for performing 3D velocity measurements in high Reynolds number, supersonic jet. The facility consists of a heated ($T = 1005$ K), 50.8 mm diameter, Mach 1.74 supersonic jet exhausting into an anechoic chamber. The jet nozzle is constructed from conic shaped converging and diverging sections that result in the production of shock-expansion cells even when the jet is operated at nominally ideally expanded conditions. The jet was seeded with submicron alumina particles injected through ports contained in the stagnation chamber. A volume of approximately 61 mm (streamwise) x 91 mm x 100 mm was illuminated using a pulsed Nd:YAG laser with pulse energy of approximately 200 mJ/pulse. Figure 32 shows a preliminary result obtained from a single day of experiments. The color indicates the streamwise component of velocity with the y-axis (streamwise direction) stretched to show cross-sections of the jet at different downstream locations. The cross-sections spans approximately from $x/D = 1.5$ to 2.5. The ability of the camera to resolve the rough circular shape of the jet and the relatively thin shear layer is apparent although further work is needed to validate the small scale features observed around the jet periphery. LES performed in the same flow also indicates that variations in streamwise velocity within the jet core are to be expected in this flow.




Figure 32 Sample 3D velocity field obtained in a 2" diameter supersonic jet seeded with alumina particles. Y-axis is stretch by factor of 5 to illustrate different cross-sections of jet flow. Color corresponds to streamwise (y) component of velocity.

To further verify the feasibility of this approach for this flow field, 24 image pairs were processed and used to determine mean and rms quantities. Figure 33 shows the average velocity field calculated from this limited set of data. A key feature of the average velocity field is the presence of a low velocity region in the center of the jet, a feature which is expected to occur in conic, supersonic nozzles such as the one used in this study. The general shape and magnitude of the velocity is also consistent with expectations. Figure 34 shows the RMS fluctuations of the velocity field. The jet core and ambient both show low levels of fluctuations whereas the shear layer shows the highest levels of fluctuations.




Figure 33 – Average velocity field calculated using 24 image pairs. 5 cross-sections are shown with the lower velocity jet core apparent near the center of the jet. Velocity scale ranges from 200 (blue) to 800 (red) m/s.




Figure 34 – RMS velocity field calculated from 24 image pairs. 5 cross-sections are shown with the fluctuations in the shear layer being apparent.

VIII. Frequency-Domain Deconvolution-Based Volume Reconstruction

While the implementation of the MART algorithm discussed previously is quite promising, its main disadvantage is that it is computationally expensive. As such, a parallel line of work was initiated to develop computationally more efficient reconstruction algorithms. As will be discussed, our approach is centered on the development of FFT based deconvolution algorithms, which are known to be significantly more computationally efficient for problems in similar fields, such as deconvolution microscopy. This work is currently being pursued through a collaboration with Dr. Stan Reeves, a faculty member in the Department of Electrical Engineering at Auburn University. A Ph.D. student, Paul Anglin, is expected to defend his dissertation in early 2014 based on the work initiated as part of this grant. The purpose of this chapter is to summarize the approach adopted, to identify some of the anticipated challenges and to report on the progress made to date.

A. Overview

Deconvolution is a frequency-domain inversion technique based on the Fourier transform property relating convolution to point-by-point multiplication in the frequency domain. The focal stack generated by digitally refocusing the acquired data can be modeled as a linear process whereby the system point spread function (PSF) is convolved with the imaged volume. It follows that the imaged volume can then be estimated by point-by-point division of its spectrum by the spectrum of the PSF. This is beneficial as calculation of the frequency spectrum of a signal can be done efficiently via the FFT. Where volume reconstruction may have taken hours using tomographic methods, solutions utilizing deconvolution can be obtained in minutes or even seconds. To truly understand the impact that such a drastic reduction in processing time can have, one must consider that PIV relies on not just a single reconstructed volume. To fully describe the flow field, the volume must be imaged and subsequently reconstructed multiple times to allow particles to be tracked as they move through the medium. Reducing the time to perform a single reconstruction to minutes significantly decreases the time required to produce usable results and image the flow field using PIV.

B. Imaging Model

In order to apply deconvolution to light-field imaging and PIV, a generalized model for the imaging volume must be established. The imaging operation can be modeled as the convolution of the object with an appropriate point spread function (PSF) that describes the blur as the focal plane moves away from the object and any blur due to optical aberrations. Defining the object in 3-space as $f(x, y, z)$, the PSF as $h(x, y, z)$, and additive noise as $\eta(x, y, z)$, the acquired image $g(x, y, z)$ is given by,

$$g(x, y, z) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x', y') h(x - x', y - y', z - z') dx' dy' dz' + \eta(x, y, z)$$

$$g(x, y, z) = f(x, y, z) * h(x, y, z) + \eta(x, y, z) \quad (22)$$

where the second step is achieved by denoting the convolution operator as $*$.

The PSF is simply the system impulse response. The PSF that describes the typical imaging system will resemble a double cone where the apex of each describes the point where the system is optimally focused. An example PSF is given in Figure 35, which shows a particle in object space (left figure) and two orthogonal slices of the PSF (center and right figures). From Figure 35, it can be seen that as the system is focused in front of or beyond a point, the resulting image will become increasingly blurred. Further, when multiple points are present in the imaged volume the blur resulting from out-of-plane points will contribute to the in-plane energy as well. This can reduce the resolution of the system, and in extreme cases, completely obscure points.




Figure 35: Impulse (left) followed by orthogonal cross sections in the (x, y) plane (center) and the (y, z) plane (right) of an example 3-D impulse response.

The 3-D PSF completely characterizes the system [36], and knowledge of the PSF hints at the ability to deblur the image with proper processing. In general, the PSF can be estimated analytically, observed experimentally, or created from some combination of the two. However, processing techniques that rely on the PSF are typically sensitive to errors in the PSF, or mismatch between the estimated PSF and the actual PSF. This mismatch can occur for several reasons such as noise and/or imaging artifacts due to optical aberrations, but a significant source of PSF mismatch is the result of the system not being shift invariant as is the case with plenoptic imaging.

C. Frequency-Domain Deconvolution

The deconvolution algorithm presented here is a frequency-domain processing technique that relies on the well-known convolution theorem of the Fourier transform. This states that

convolution in the spatial domain is equivalent to point-by-point multiplication in the frequency domain, or

$$f(x,y,z)*h(x,y,z) \xrightarrow{\mathfrak{F}} F(\omega_x, \omega_y, \omega_z)H(\omega_x, \omega_y, \omega_z) \quad (23)$$

Applying this to the previous result gives

$$\begin{aligned} g(x,y,z) &= f(x,y,z)*h(x,y,z) + \eta(x,y,z) \\ f(x,y,z)*h(x,y,z) + \eta(x,y,z) &= \mathfrak{F}^{-1} \left[F(\omega_x, \omega_y, \omega_z)H(\omega_x, \omega_y, \omega_z) + N(\omega_x, \omega_y, \omega_z) \right] \end{aligned} \quad (24)$$

In the absence of noise, the imaged space can be estimated as

$$\hat{F}(\omega_x, \omega_y, \omega_z) = \frac{G(\omega_x, \omega_y, \omega_z)}{H(\omega_x, \omega_y, \omega_z)} \quad (25)$$

Unfortunately, the system is rarely, if ever, noise free. The general case is then,

$$\hat{F}(\omega_x, \omega_y, \omega_z) = F(\omega_x, \omega_y, \omega_z) + \frac{N(\omega_x, \omega_y, \omega_z)}{H(\omega_x, \omega_y, \omega_z)} \quad (26)$$

This relationship offers some insight into both the advantages and disadvantages of inverse filtering. The primary advantage is speed of computation. Direct convolution for an N-length signal requires N² operations while convolution via FFT requires O(N log N). However, perfect knowledge of the transfer function no longer guarantees a perfect reconstruction of the original volume once noise is introduced. Furthermore, a transfer function containing small values will result in an amplification of the noise term in the estimate of the imaged volume. As these values approach zero, the noise term will dominate the response. This case is common because the system typically represents an infinite bandwidth scene with band-limited data. The consequence is a transfer function with zero crossings.

The impacts of this noise amplification and PSF mismatch can be mitigated to some degree with proper filtering. One approach is to minimize the expected value (E) of the squared error

$$e^2 = E \left\{ (f - \hat{f})^2 \right\}. \quad (27)$$

The well-known result of this minimization is the Wiener filter given by

$$\hat{F}(\omega_x, \omega_y, \omega_z) = \left[\frac{H^*(\omega_x, \omega_y, \omega_z)}{|H^*(\omega_x, \omega_y, \omega_z)|^2 + \frac{S_\eta(\omega_x, \omega_y, \omega_z)}{S_f(\omega_x, \omega_y, \omega_z)}} \right] G(\omega_x, \omega_y, \omega_z) \quad (28)$$

where S_η and S_f are the power spectrum of the noise and the original image volume respectively. When the power spectra are not known, which is often the case, the ratio of the power spectra are replaced by a regularization parameter K . Substituting this into the above equation gives

$$\hat{F}(\omega_x, \omega_y, \omega_z) = \left[\frac{H^*(\omega_x, \omega_y, \omega_z)}{|H^*(\omega_x, \omega_y, \omega_z)|^2 + K} \right] G(\omega_x, \omega_y, \omega_z) \quad (29)$$

For the noise free case with perfect PSF agreement, the regularization factor K can be set to zero, which reduces Equation 29 back to Equation 25. These relationships can be applied directly to plenoptic data.

D. Deconvolution Applied to Simulated Plenoptic Data

In order to determine the efficacy of deconvolution applied to plenoptic data, simulations begin with 2-D cases. This simplification is only necessary to simplify visualization of the results, and the results are directly applicable to the 3-D case. The system was simulated for a plenoptic camera with the following parameters:

Lenslet Focal Length	f_l	0.500mm
Lenslet Pitch	p_l	0.125mm
Pixel Pitch	p_p	0.0074mm
Number of Pixels	n_p	1424
Sensor Size		10.5mm
Number of Lenslets	n_l	89
Main Lens Focal Length	f_m	50mm

Table 3: 2-D Simulation Parameters

Selection of a PSF is critical to the performance of deconvolution algorithms. There are many parameters that can affect the impulse response of an optical system such as optical aberrations. Furthermore, a plenoptic camera has unique characteristics that result in a system that is spatially variant. Characteristics such as microlens to sensor misalignment and quantization errors resulting from particles overlapping multiple microlenses cause the impulse response to differ throughout the imaging volume. With this in mind, the PSF chosen for these

simulations is generated by placing a point on the optical focal plane along the optical axis. The point is an ideal impulse in that all the simulated rays originate from an infinitesimally small point. This ensures that the point does not overlap multiple lenslets. Furthermore, simulations are performed with the microlens and sensor perfectly aligned.




Figure 36: Simulated 2-D PSF

The focal stack for the test PSF is shown in Figure 36 and demonstrates an important aspect of plenoptic imaging. The PSF appears more pixelated than one might expect. This is not a simulation artifact but is the result of the reduced resolution inherent in the system. Again, the spatial resolution of the image is set by the size of each microlens and each ‘virtual’ pixel is one lenslet pitch, p_l , in width.

With a PSF chosen, the simulation proceeds with a test case where three points are placed within the simulated imaging volume at (3mm, 103mm), (0mm, 100mm), and (–3mm, 88mm), where the second point is again located on the optical focal plane along the optical axis. The second point is included in order to provide a reference, as it should perfectly match the test PSF. As with the PSF, each point is simulated using 5000 rays originating from an infinitesimally small point. Figure 37 shows the resulting focal stack for this case where the figures are normalized to the peak value to aid in clarity.




Figure 37: Focal stack for three simulated points

There are two aspects of Figure 1.3 that are immediately apparent and ultimately related. Specifically, the three points have significantly different intensities and varying sizes. This is noteworthy as each point is simulated using the same number of rays, and therefore has the same energy, as well as being simulated as an infinitely small point. Ideally, each point would have the same or similar intensity and size at its optimal focal plane. However, the plenoptic camera is not able to focus equally well at all possible focal planes as the resolution at synthesized focal planes is limited by the bandwidth of the camera. For the purposes of this research, this is a contributor to the spatially variant nature of the PSF. The result is a variation in the PSF throughout the imaging volume.

Applying deconvolution per Equation 29 with no regularization ($K = 0$) gives the results shown in Figure 38. This appears to have little in common with the original image shown in Figure 37, but these results are to be expected. Closer inspection shows that the point at the origin is present in the deconvolved image. Again, the central point was chosen as it perfectly matches the test PSF.




Figure 38: Deconvolved image with $K=0$

Applying regularization to the deconvolution algorithm produces more desirable results. Figure 39 shows the results when a $K = 0.00015$ regularization is applied. This is a small amount of regularization, but the impact is significant. The three points are clearly present in the reconstruction. However, the central point dominates the response, and the two other points appear to have a much lower intensity. Further increasing the regularization to $K = 0.01$ yields the response shown in Figure 40. This is an increase of two orders of magnitude in the regularization, and again the impact is significant. While each of the points becomes more clearly identifiable, the deconvolution has become less effective at reducing the blur.




Figure 39: Deconvolved image with $K=0.00015$




Figure 40: Deconvolved image with $K=0.01$

Noting that the response in the deconvolved image is dominated by the central point, an additional case is created in order to provide more insight into the results. For this case, three points are again simulated, but the central point is now moved to (2.6mm,106mm). The resulting focal stack is shown in Figure 41. Initially, this focal stack appears quite similar to that in Figure 37 above. However, upon closer inspection, two important characteristics of the system can be observed.




Figure 41: Case 2 focal stack for three simulated points.

First, it is clear that the plenoptic camera loses resolution as the points move away from the optical focal plane. Again, this is seen as a loss in peak intensity and a blurring of the point, and is due to the limited angular resolution of the imaging system. Some loss in resolution will occur as points move away from the optical axis, but this is a result of points being imaged by multiple micro lenses, and is much less significant when compared to displacements away from the focal plane. Next, a “banding” phenomena is clearly present for each point. This artifact resulting from computational refocusing has not been addressed in the current literature and will be investigated further in subsequent research.

Deconvolving the image using the test PSF described above with a regularization parameter $K = 0.00015$ yields the image shown in Figure 42. The resulting image is a significant improvement over the blurred focal stack shown in Figure 41, but artifacts do remain. These artifacts are primarily the result of mismatch between the PSF and the simulated response of each object. However, one artifact of note is the vertical line near the middle of the deconvolved image. This is also an issue of PSF mismatch, but this can be mitigated to some extent with proper signal processing.




Figure 42: Deconvolved image with $K=0.00015$

E. Artifacts and Improvements

Significant artifacts remain in the volume reconstruction, and an efficient method of reducing these is desirable. Two significant artifacts have been identified that do not appear to have been previously published. These are banding in the object response and a vertical line at the origin of the deconvolved image, both of which degrades the reconstruction.

Furthermore, while deconvolution significantly reduces the computational burden when reconstructing a volume from plenoptic data compared to tomographic methods, a significant bottleneck remains in generating the object focal stack and the PSF and the efficiency of the algorithm can be improved by reducing these operations as well. Once again, Fourier processing, and specifically Fourier-based refocusing, offers a means of further reducing the computational requirements. For the purposes of this work, this approach can be leveraged into a frequency-domain iterative algorithm intended to further reduce remaining artifacts.

1. Banding in the object image response

Figure 41 shows the focal stack for three ideal points spread throughout the imaged volume. This image shows an artifact inherent to the imaging system that has not been previously identified. That is, each object exhibits banding in its response. To more clearly illustrate the artifact, an additional case is shown in Figure 42 where a single point is located at (0mm, 118mm). It should be noted that the banding is more prominent in Figure 42 only due to scaling of the image where, again, the peak intensity is normalized to 1.




Figure 43: Example focal stack illustrating banding in a refocused image

The source of this artifact stems from the tradeoff inherent in plenoptic imaging. A plenoptic camera sacrifices spatial resolution in order to obtain angular information, and one consequence is that the system no longer records each ray location in space with infinite resolution. Its position must therefore be quantized to within one microlens. When an image is refocused, each ray collected by a microlens is assumed to originate from the center of the microlens. The result is a region where no data is present in the refocused plane. Similarly, a finite number of angular samples are collected, which is determined by the number of pixels underneath each lenslet. While each pixel captures a range of angles, the reconstruction assumes a discrete angular. Again, this discretization results in regions where no data is present in the refocused plane.

Further complicating the issue is microlens/sensor misalignment and the case where the microlens and pixel pitch are not integer multiples of one another. This is shown graphically in Figure 43. To accommodate these conditions, pixels around the perimeter of each microlens are removed from the reconstruction in order to compensate for overlapping information from adjacent microlenses. This effectively removes u plane samples and results in larger regions with no data.




Figure 44: Simplified diagram showing pixels overlapping the boundary between two microlenses

This artifact is significant due to its asymmetry and the fact that it appears in all points not falling on the optical focal plane and optical axis. This artifact alone confirms that the response is spatially variant, and multiple PSFs would be required to perfectly describe the system.

2. PSF Asymmetry

The mismatch between the selected PSF and an object focal stack leads to another interesting artifact. To demonstrate, the focal stack used to generate Figure 43 above is deconvolved using the same PSF shown earlier. The result is shown in Figure 45, and the artifact of interest is the vertical line at the origin. With the exception of a single point at the origin, this vertical line appears in all the reconstructions, regardless of the point distribution. The source of this artifact is the asymmetry of the object focal stack. Specifically, the test PSF is centered at, and symmetrical about, the origin. However, because the object is not centered in the focal stack, the response is asymmetric due to the presence of more focal planes on one side of the object compared to the other. Figure 46 shows a cross section of both the PSF and the object focal stack that clearly shows the mismatch in question.




Figure 45: Deconvolved image showing a single point offset from the origin




Figure 46: Cross section of PSF and object focal stack

The relationship between the object focal stack asymmetry and the vertical line is not immediately obvious. However, considering the process in the forward direction can be useful in clarifying the issue. The deconvolved image can be thought of as a series of impulses, where

convolution between the impulse and the PSF explains some portion of the object focal stack. The zero lag case, that is, the case where the PSF is centered over the object focal stack is the only case where every point in the PSF has a corresponding point in the object focal stack. A shift in any direction results in points in the PSF with no corresponding point in the object focal stack. The result is an impulse at the origin that attempts to explain the error resulting from the asymmetry of the object focal stack.

A more precise explanation, although less intuitive, can be found through inspection of the spectral content of the component signals. Referring to Figure 47, an oscillation or ringing can be noted along the horizontal axis. This is the result of taking the Fourier transform of a spatially limited signal. Specifically, the PSF and object focal stack do not decay to identically zero values before reaching the edge of the focal stack. This discontinuity results in a sinc-like response in the frequency domain. However, because the PSF and the object focal stack are truncated differently, the resulting spectra differ.



Figure 47: Fourier transform of PSF (left) and the object focal stack (right)

Comparing the Fourier transform of the deconvolved image $\hat{F}(k_x, k_z)$ shown in Figure 48 illustrates the differences between the spectra. Of note is the significant ringing along the horizontal axis. Because the PSF and the object focal stack do not perfectly match, the ringing is amplified in the result of the deconvolution. Recalling that the spectrum of a single point (or impulse) is unity, where the phase of the response determines the location of the points in space, it is clear that the resulting deconvolved spectrum will yield artifacts.




Figure 48: Fourier transform of deconvolved image with $K=0.00015$

F. 3-D Fourier Slice Refocusing

Fourier slice refocusing, as presented by Ng [37], provides a method of refocusing an image to a specified distance based on the Fourier slice theorem. The Fourier-slice theorem, or projection-slice theorem, was first developed by Ron Bracewell [38] while studying radio astronomy but has been applied to many fields with medical imaging being a significant benefactor [39, 40]. The advantage of this approach over integral-based refocusing techniques is seen by comparing the number of operations required to generate a refocused image. Integral-based refocusing requires $O(n^4)$ operations for each focal plane, where n samples are assumed in each dimension. Compare this to Fourier-slice refocusing which, requires only $O(n^4 \log n)$ for the initial 4-D FFT, $O(n^2)$ for the slicing operation, and $O(n \log n)$ for the inverse 2-D FFT. As the number of samples increases, this offers a significant computational savings. Furthermore, when cast in the context of deconvolution, multiple focal planes must be generated. These can each be calculated from the 4-D FFT of the plenoptic data, which further improves the efficiency.

Existing derivations for Fourier slice refocusing are predicated on generating a single 2-D image. However, in the context of deconvolution, a complete 3-D focal stack is required; more specifically, the FFT of the 3-D focal stack is desired. Obtaining this without requiring an interim conversion to the spatial domain further improves the efficiency of this technique. To achieve

this, Fourier slice refocusing can be used to generate an array of 2-D FFTs corresponding to each focal plane in the final focal stack. A final 1-D FFT is performed across the third dimension of the array, which gives the 3-D FFT of the focal stack.

Establishing the framework for Fourier slice refocusing begins with the light field at the lenslet plane given by $L(x, y, u, v)$. For an arbitrary point in 3-space, the light field is given by

$$L_\alpha(x', y', z') = L\left(u\left(1 - \frac{1}{\alpha}\right) + \frac{x'}{\alpha}, v\left(1 - \frac{1}{\alpha}\right) + \frac{y'}{\alpha}, u, v\right). \quad (30)$$

Defining the shearing function as B ,

$$B_\alpha = \begin{pmatrix} \frac{1}{\alpha} & 0 & 1 - \frac{1}{\alpha} & 0 \\ 0 & \frac{1}{\alpha} & 0 & 1 - \frac{1}{\alpha} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (31)$$

and arranging the coordinates (x', y', u, v) into the column vector a , the function can be rewritten as,

$$L_\alpha(a) = L(Ba) \quad (32)$$

The intensity at any given point on the $z' = z \propto$ plane is obtained by integrating the light-field with respect to u and v as

$$E_\alpha(x', y') = \frac{1}{\alpha^2 F^2} \iint L\left(u\left(1 - \frac{1}{\alpha}\right) + \frac{x'}{\alpha}, v\left(1 - \frac{1}{\alpha}\right) + \frac{y'}{\alpha}, u, v\right) \quad (33)$$

A traditional 2D image is obtained by evaluating the integral for a fixed z' , and by extension, the volume is obtained by evaluating the integral at all points within the volume. This is a computationally intensive process. However, the projection-slice theorem offers a more efficient means of obtaining the 3D focal stack directly through Fourier processing.

It can be shown that for an arbitrary N -D space and M -D projection (where $M < N$), the Fourier projection-slice theorem is

$$F^M \circ I_M^N \circ B = S_M^N \circ \frac{B^{-T}}{|B^{-T}|} \circ F^N \quad (34)$$

where F^M represents the M-D Fourier transform, I_M^N represents the projection from M to N, S_M^N represents the slicing operation whereby the last N-M dimensions of a function are set to 0, and B is the shearing operator.

When generating a 2D image ($M = 2$) from 4D plenoptic data ($N = 4$), Equation 34 becomes

$$F^2 \circ I_2^4 \circ B = S_2^4 \circ \frac{B^{-T}}{|B^{-T}|} \circ F^4 \quad (35)$$

Using the right-hand side of Equation 35, the 2D FFT corresponding to a particular focal plane can be selected from the 4D FFT of the plenoptic data. Each 2D FFT is then arranged into a 3D array, and a final 1D FFT along the third dimension of this dataset results in the 3D FFT of the corresponding focal stack. While taking the 3D IFFT of this data set would yield the focal stack, for the purposes of deconvolution, it is more desirable to operate entirely in the frequency domain.

A test case is again simulated in order to compare integral-based refocusing and FFT-based refocusing. However, FFT-based refocusing has proven to be more sensitive to edge discontinuities as well as limited angular samples. To limit the impact of this, the simulated microlens array is increased to 150 microlenses, and subsequently, the number of sensor pixels increases as well.

Lenslet Focal Length	f_l	0.500mm
Lenslet Pitch	p_l	0.125mm
Pixel Pitch	p_p	0.0074mm
Number of Pixels	n_p	2400
Sensor Size		10.5mm
Number of Lenslets	n_l	150
Main Lens Focal Length	f_m	50mm

Table 4: 2-D Simulation Parameters

A comparison of the focal stack resulting from integral-based refocusing and FFT-based refocusing is provided in Figure 49, where only the central 89 lenslet pixels in the vertical direction are considered for clarity. Clearly, FFT-based refocusing compares favorably and has the added advantage of a reduced computational burden. For the purposes of this work, an FFT-based solution also reduces the initial computational overhead associated with an FFT-based iterative algorithm. Additionally, each focal stack exhibits a spatial distortion where the

PSFs appear to ‘lean.’ This is believed to be related to the geometric distortion that occurs when the imaged object space is compressed into the image space residing between the main lens and the lenslet array.





Figure 49: Comparison of integral-based and FFT-based refocusing

G. 2.4 Iterative Algorithm

Establishing a method of generating the spectral content of an object focal stack and system PSF directly leads to the development of an efficient FFT-based iterative algorithm. The algorithm proposed here is based on a gradient descent method, which attempts to minimize the squared error of the estimate at each iteration. To further utilize the research presented here, deconvolution can be used to rapidly generate the initial estimate in order to speed convergence.

1. Iterative Gradient Descent Applied to Plenoptic Imaging

The imaging process and subsequent refocusing can be described as a linear process. Let g be the plenoptic data acquired by the imaging system sensor, f be the 3D object, n be additive noise, and D be the matrix that describes the mapping from object space to the data. The system is then modeled as,

$$g = Df + n. \quad (36)$$

The squared error is

$$\phi(f) = \|g - Df\|^2 \quad (37)$$

Taking the first derivative and setting to zero gives the desired result of minimizing the squared error.

$$\begin{aligned} \nabla \phi(f) &= D^T g - D^T Df \\ 0 &= D^T g - D^T Df \\ \hat{f} &= (D^T D)^{-1} D^T g \end{aligned} \quad (38)$$

Equation 38 gives an estimate of the 3D space based on the plenoptic data acquired by the system. The gradient descent algorithm can then be established by moving in the opposite direction of the gradient of the error term and using this to update the previous estimate as

$$\begin{aligned} f_{k+1} &= f_k + \beta_k D^T (g - Df_k) \\ f_{k+1} &= f_k + \beta_k (D^T g - D^T Df_k) \end{aligned} \quad (39)$$

Considering these relationships in more detail gives a more intuitive interpretation. The $D^T g$ term generates a focal stack from the plenoptic data, $D^T D$ is a blurring operation which uses the system PSF to blur the current estimate of the object space, and β_k is a scalar used to adjust

the step size at each iteration. By repeating this process, the algorithm can iteratively reduce the error between the focal stack and the object space estimate. Furthermore, this can be done efficiently because the entire operation can be carried out in the frequency domain using point-by-point multiplications rather than convolutions necessary in the spatial domain. Unfortunately, limitations remain.

First, the quality of the estimate is limited by the accuracy of the PSF. Because the PSF is used to blur each estimate in the error calculation, PSF mismatch will lead to errors in the estimate. Plenoptic imaging systems are spatially variant, and therefore a single PSF cannot completely characterize the system. This indicates that the use of a single PSF will result in some level of error that cannot be removed through application of the iterative algorithm, or multiple PSFs must be implemented in order to obtain a better estimate.

Second, termination criteria must be established. The algorithm in Equation 39 has no inherent means of termination due to the significant number of unknowns. Termination can be initiated by several means including user intervention, limiting the number of iterations, establishing a threshold for the error criterion, etc. Further, due to PSF mismatch and potential edge artifacts, the algorithm can begin to amplify artifacts in the estimate if allowed to progress without termination criteria or without modifications to the algorithm.

H. Fourier Projection-Slice Theorem

The previous sections have proposed a method of reconstructing the imaged volume through the use of deconvolution. However, another frequency-domain reconstruction technique exists that may provide yet another efficient means of estimating the space. Plenoptic imaging can be thought of as taking 2-D projections from a 3-D object over a range of angles. This transformation is an extension of the 2-D Radon transform to 3-D space, i.e., the 3-D Radon transform. It is proposed that the 3-D Radon transform also has an inverse as does the 2-D transform through the use of the projection-slice theorem.

While the projection-slice theorem has been extended to higher dimensions, plenoptic imaging poses a unique situation in that the imaging sensor is on a fixed plane, and only a limited number of angular samples can be collected. The typical geometry is one where the sensor is rotated around the object to be imaged, allowing each projection to be taken normal to the imaging surface over a full range of angles.

1. 2-D/3-D of the Fourier Projection-Slice Theorem

The projection-slice theorem is the basis for Fourier-based tomographic reconstruction. Traditionally, this is used to reconstruct a 2-D image $f(x, y)$ from its 1-D projections $g(l, \theta)$ [41] and is formally given by:

$$G(\rho, \theta) = F(\rho \cos \theta, \rho \sin \theta) \quad (40)$$

This relationship states that the Fourier transform of the 1-D projection $g(l, \theta)$ is equivalent to a slice from the 2-D Fourier transform of the object $f(x, y)$. The projection-slice theorem has been extended to higher dimensions [42-44], but typically the geometry is such that the line integrals necessary to calculate the projection are taken normal to the surface. However, the sensor in a plenoptic camera is fixed, and the image of the 3-D object is projected onto this sensor. That is, the line integrals used to calculate the projection would intersect the sensor plane at some arbitrary angles θ and ϕ . The projection-slice theorem must be established for this new geometry.

Defining an arbitrary object in a 3-dimensional space by $f(x, y, z)$, the projection $g(x', y', z')$ resulting from rays parallel to $\hat{v}$ onto a 2-dimensional plane at $z = 0$ is given by:

$$g(x', y', z') = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y, z) \delta(\bar{r} - \bar{r}_0 - t\bar{v}) dx dy dz \quad (41)$$

Here, $\bar{r} = \bar{r}_0 - t\bar{v}$ is the equation of a vector in three-space as shown in Figure 50. The vector $\bar{v}$ is a unit vector in the direction of $\bar{a}$, and it can be shown that a number t exists such that $t\bar{v} = \bar{a}$.




Figure 50: Geometry for the equation of a line in three-space

The unit vector $\bar{v}$ can be written in polar coordinates as

$$\hat{\mathbf{v}} = \begin{pmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{pmatrix} \quad (42)$$

and the equation of a line parallel to the unit vector in three-space is given by

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} + t \begin{pmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{pmatrix}$$

$$t = \frac{z}{\cos \theta}$$

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x' \\ y' \end{pmatrix} + \begin{pmatrix} z \tan \theta \cos \phi \\ z \tan \theta \sin \phi \end{pmatrix} \quad (43)$$

The geometry of the system is shown in Figure 51.




Figure 51: Geometry of the projections through $f(x,y,z)$ on to the sensor plane

Taking the 2-D Fourier transform of $g(x', y', \vec{v})$.

$$\begin{aligned}
G(k_x, k_y, \hat{v}) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x', y', \hat{v}) e^{-j2\pi(k_x x' + k_y y')} dx' dy' \\
G(k_x, k_y, \hat{v}) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y, z) \delta(r - r_0 - t\hat{v}) e^{-j2\pi(k_x x' + k_y y')} dx dy dz dx' dy' \\
G(k_x, k_y, \hat{v}) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y, z) \delta(x - x' - z \tan(\theta) \cos(\phi), y - y' - z \tan(\theta) \sin(\phi)) \\
&\quad e^{-j2\pi(k_x x' + k_y y')} dx dy dz dx' dy' \\
G(k_x, k_y, \hat{v}) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y, z) e^{-j2\pi(k_x(x - z \tan(\theta) \cos(\phi)) + k_y(y - z \tan(\theta) \sin(\phi)))} dx dy dz
\end{aligned} \tag{44}$$

where the last line follows from the sifting property of the Dirac delta function. Regrouping terms yields

$$G(k_x, k_y, \hat{v}) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y, z) e^{-j2\pi(k_x x + k_y y - z \tan(\theta)(k_x \cos(\phi) + k_y \sin(\phi)))} dx dy dz \tag{45}$$

This is the Fourier projection-slice theorem relating the 2-D projections from an object to its 3-D Fourier transform for the geometry shown in Figure 51. This relationship allows the imaged volume to be reconstructed directly from the plenoptic data without the need for deconvolution. Furthermore, because this relationship is based on the frequency-domain representations of the signals, efficient FFT processing can be used.

I. Future Work

The preceding discussion of FFT based deconvolution algorithms sets the stage for continuing and promising work in the area of 3D image reconstruction with a plenoptic camera. In the context of plenoptic PIV, these algorithms are expected to provide an order of magnitude improvement in computational efficiency. It is clear, however, that significant work is still necessary to implement these algorithms in a practical and efficient manner. The most notable challenges are related to the spatial invariance of the plenoptic camera's PSF, which can cause reconstruction artifacts. Solutions to these challenges have been proposed and work is currently ongoing.

IX. Conclusions & Future Work

A methodology, termed plenoptic particle image velocimetry (PIV), was developed and demonstrated for the measurement of three-dimensional, three-component velocity fields in high-speed turbulent flow fields. The concept is based on light-field imaging of particles where a plenoptic camera is used to measure both the position and angle of light rays scattered by particles contained in the flow field. Tomographic algorithms are applied to plenoptic camera data to determine the 3D position of particles contained in the flow field and cross-correlation algorithms are used to determine the displacement of particles in two successive images. A prototype camera was designed, built and used to demonstrate the viability of the technique in practice. A synthetic image generation tool was developed to simulate the camera and investigate the technique's accuracy. Particle positions are found to be accurate within 0.1 mm in the lateral directions and within 0.3 mm in the depth direction for a volume with dimensions 30 mm x 20 mm x 20 mm. Demonstration experiments included 3D velocity measurements of a turbulent boundary layer formed on a wind tunnel wall and a heated, supersonic axisymmetric jet.

Plenoptic PIV's chief advantages are:

- Single, Compact Camera – Plenoptic PIV only requires a single camera with a small form factor that is virtually identical to that of conventional PIV systems. As such, plenoptic PIV can be used in facilities where optical access and nearby floor space is limited (e.g. gas turbine flows, combustors, pressurized flow facilities, etc.). This is in contrast to tomo-PIV, which generally requires 4 or more cameras spread over a fairly large baseline that prevents its application in space and optical access restricted environments.
- Easy to set-up and calibrate – Calibration of a plenoptic camera is a fairly simple procedure that involves imaging a white background with the main lens aperture set to a minimum size and imaging a ruler to determine nominal magnification. Multi camera systems, on the other hand, require a rigorous and complex calibration procedure in order to align and register the images produced by different cameras viewing the scene from different positions. Experience has shown that plenoptic PIV can be set-up to acquire data in a short period of time.
- Economical – The cost of a plenoptic PIV system is expected to be only slightly increased over that of a conventional PIV system, of which hundreds of systems have already been installed around the world.

The main shortcoming of the present work are two-fold:

- Limited spatial resolution – The ability to acquire 3D data about the flow comes with the tradeoff of reduced spatial resolution. This is directly tied to the concept of the

plenoptic camera where a microlens array is used to sample the angular content of the light-field and record it on the pixels located behind the microlens array.

- Computational Expense – The current implementation of the MART algorithm for volumetric reconstruction of particle fields is computationally intense, requiring on the order of several hours per reconstructed image.

Taken together, the future of plenoptic PIV appears to be very bright. The ability to acquire 3D velocity data from a compact and relatively inexpensive camera with minimal optical access cannot be matched by any other existing diagnostics. As such, plenoptic PIV is expected to find utility in a wide variety of applications. In addition, the main shortcomings of plenoptic PIV are only expected to improve with time. In particular, the resolution of image sensors continues to grow and is not currently constrained by technical capability, but rather, by consumer demand. As plenoptic cameras grow in popularity (e.g. Lytro), higher resolution sensors are expected to become available. In fact, the PI is currently constructing a new plenoptic camera (via a DURIP grant jointly funded with Florida State University) based on a 29 megapixel image sensor. Computationally, the seeds for the development of fast, FFT based deconvolution algorithms was outlined here. Combined with advances in GPU programming and clusters, reconstruction algorithms are expected to be reduced to the order of minutes or less per image such that large sets of data can be collected and analyzed in a short period of time.

Lastly, plenoptic imaging also holds tremendous potential for development of 3D variants of other flow diagnostics, such as 3D background oriented schlieren, 3D laser induced fluorescence and 3D photogrammetry. This ability is due to the fact that plenoptic camera provide a dense sampling of the angular space of a light-field in contrast to multi-camera techniques, which generally are characterized as a sparse sampling of the angular space. For example, tomo-PIV is capable of measurements of particle fields as particles are sparsely located in space; however, the multi-camera tomo-PIV system is not capable of generating ‘photographs’ with new focal planes or perspectives.

Overall, the future work of the PI is expected to focus on:

- Continued refinement of plenoptic PIV for improved resolution and accuracy.
- Development of high repetition rate plenoptic PIV systems for time-resolved 3D flow measurements.
- Application of plenoptic PIV to various flow fields.
- Extension of plenoptic imaging to 3D versions of other flow diagnostics.

X. Publications

The following publications resulted from the present work and provide more details about the work performed as part of this grant. They are all publically available and can be furnished upon request.

A. Theses

- Fahringer, T., “Volumetric particle image velocimetry with a plenoptic camera”, M.S. Thesis, under preparation for completion in 2013.
- Lynch, K., “Development of a 3-D Fluid Velocimetry Technique based on Light Field Imaging”, M.S. Thesis, Auburn University, 2011.

B. Journal Publications

- Fahringer, T., Lynch, K. and Thurow, B. “Plenoptic Particle Image Velocimetry,” submitted to *Measurement Science and Technology*, November 2013

C. Conference Papers & Other Publications

- Thurow, B. and Fahringer, T., “Recent development of volumetric PIV with a plenoptic camera,” Proceedings of the 10th International Symposium on Particle Image Velocimetry, Delft, The Netherlands, July 1-3, 2013.
- Fahringer, T. and Thurow, B., “The effect of grid resolution on the accuracy of tomographic reconstruction using a plenoptic camera”, AIAA Paper 2013-0039, 51st AIAA Aerospace Sciences Meeting, Grapevine, TX, January 2013
- Melnick, M.B., Thurow, B., Fahringer, T. and Brock, B., “Experimental investigation of three-dimensional structures in an adverse pressure gradient turbulent boundary layer”, AIAA Paper 2012-2850, 42nd AIAA Fluid Dynamics Conference, New Orleans, LA, June 2012
- Fahringer, T. and Thurow, B., “Tomographic reconstruction of a 3-D flow field using a plenoptic camera,” AIAA Paper 2012-2826, 42nd AIAA Fluid Dynamics Conference, New Orleans, LA, June 2012
- Lynch, K., Fahringer, T., and Thurow, B., “Three-dimensional particle image velocimetry using a plenoptic camera,” AIAA Paper 2012-1056, 50th AIAA Aerospace Sciences Meetings, Nashville, TN, January 2012.
- Lynch, K. and Thurow, B., “Preliminary Development of a 3-D, 3-C PIV Technique using Light Field Imaging,” AIAA Paper 2011-3729, 41st AIAA Fluid Dynamics Conference, Honolulu, HI, June 2011.

XI. References

1. Arroyo, M.P. and C.A. Greated, *Stereoscopic Particle Image Velocimetry*. Measurement Science & Technology, 1991. **2**(12): p. 1181-1186.
2. Kahler, C.J. and J. Kompenhans, *Fundamentals of multiple plane stereo particle image velocimetry*. Experiments in Fluids, 2000. **29**(7): p. S70-S77.
3. Brucker, C., D. Hess, and J. Kitzhofer, *Single-view volumetric PIV via high-resolution scanning, isotropic voxel restructuring and 3D least-squares matching (3D-LSM)*. Measurement Science & Technology, 2013. **24**(2): p. 024001.
4. Lynch, K. and B.S. Thurow, *Three-Dimensional Flow Visualization Using a Pulse Burst Laser System*. 2009: American Institute of Aeronautics and Astronautics, 1801 Alexander Bell Dr., Suite 500 Reston VA 20191-4344 USA.
5. Thurow, B., N.B. Jiang, and W. Lempert, *Review of ultra-high repetition rate laser diagnostics for fluid dynamic measurements*. Measurement Science & Technology, 2013. **24**(1): p. 012002.
6. Pereira, F., et al., *Defocusing digital particle image velocimetry: a 3-component 3-dimensional DPIV measurement technique. Application to bubbly flows*. Experiments in Fluids, 2000. **29**(1): p. S78-S84.
7. Hinsch, K.D., *Holographic particle image velocimetry*. Measurement Science & Technology, 2002. **13**(7): p. R61-R72.
8. Herrmann, S.F. and K.D. Hinsch, *Light-in-flight holographic particle image velocimetry for wind-tunnel applications*. Measurement Science & Technology, 2004. **15**(4): p. 613-621.
9. Katz, J. and J. Sheng, *Applications of Holography in Fluid Mechanics and Particle Dynamics*. Annual Review of Fluid Mechanics, 2010. **42**(1): p. 531-555.
10. Trolinger, J.D., M. Rottenkolber, and F. Elandaloussi, *Development and application of holographic particle image velocimetry techniques for microgravity applications*. Measurement Science & Technology, 1997. **8**(12): p. 1573-1583.
11. Elsinga, G.E., et al., *Tomographic particle image velocimetry*. Experiments in Fluids, 2006. **41**(6): p. 933-947.
12. Scarano, F., *Tomographic PIV: principles and practice*. Measurement Science & Technology, 2013. **24**(1): p. 012001.
13. Belden, J., et al., *Three-dimensional synthetic aperture particle image velocimetry*. Measurement Science & Technology, 2010. **21**(12).
14. Schroder, A., et al., *Investigation of a turbulent spot and a tripped turbulent boundary layer flow using time-resolved tomographic PIV*. Experiments in Fluids, 2008. **44**(2): p. 305-316.
15. Scarano, F. and C. Poelma, *Three-dimensional vorticity patterns of cylinder wakes*. Experiments in Fluids, 2009. **47**(1): p. 69-83.

16. Humble, R.A., et al., *Three-dimensional instantaneous structure of a shock wave/turbulent boundary layer interaction*. Journal of Fluid Mechanics, 2009. **622**: p. 33-62.
17. Violato, D. and F. Scarano, *Three-dimensional vortex analysis and aeroacoustic source characterization of jet core breakdown*. Physics of Fluids, 2013. **25**(1): p. 015112.
18. Violato, D. and F. Scarano, *Three-dimensional evolution of flow structures in transitional circular and chevron jets*. Physics of Fluids, 2011. **23**(12).
19. Adelson, E.H. and J.R. Bergen, *The plenoptic function and the elements of early vision*. Computational models of visual processing, 1991. **91**(1): p. 3-20.
20. Adelson, E.H. and J.Y.A. Wang, *Single Lens Stereo with a Plenoptic Camera*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1992. **14**(2): p. 99-106.
21. Levoy, M., *Light fields and computational imaging*. Computer, 2006. **39**(8): p. 46-+.
22. Levoy, M. *The digital Michelangelo project*. in *3-D Digital Imaging and Modeling, 1999. Proceedings. Second International Conference on*. 1999.
23. Levoy, M., et al., *Light field microscopy*. ACM Transactions on Graphics, 2006. **25**(3): p. 924-934.
24. Ng, R., et al., *Light Field Photography with a Hand-Held Plenoptic Camera*, 2005.
25. Georgiev, T. *New results on the Plenoptic 2.0 camera*. in *Signals, Systems and Computers, 2009 Conference Record of the Forty-Third Asilomar Conference on*. 2009.
26. Georgiev, T.I., C., *Light Field Camera Design for Integral View Photography*, A.S.T. Report, Editor 2006.
27. Lumsdaine, A. and T. Georgiev. *The focused plenoptic camera*. in *Computational Photography (ICCP), 2009 IEEE International Conference on*. 2009.
28. Lippmann, G., *Epreuves reversibles donnant la sensation du relief*. J. Phys. Theor. Appl., 1908. **7**(1): p. 821-825.
29. Wilburn, B., et al., *High performance imaging using large camera arrays*, in *ACM SIGGRAPH 2005 Papers* 2005, ACM: Los Angeles, California. p. 765-776.
30. Gerrard, A. and J.M. Burch, *Introduction to matrix methods in optics*. 2012: Courier Dover Publications.
31. Lynch, K., *Development of a 3-D Fluid Velocimetry Technique based on Light Field*

Imaging, in *Aerospace Engineering* 2011, Auburn University.

32. Herman, G.T. and A. Lent, *Iterative reconstruction algorithms*. Comput Biol Med, 1976. **6**(4): p. 273-94.
33. Adrian, R.J.a.W., Jerry, *Particle Image Velocimetry*. 2011, New York, New York: Cambridge University Press.
34. Scarano, F. and M.L. Riethmuller, *Advances in iterative multigrid PIV image processing*. Experiments in Fluids, 2000. **29**(7): p. S51-S60.
35. Melnick, M., et al., *Experimental Investigation of Three-Dimensional Structures in an Adverse Pressure Gradient Turbulent Boundary Layer*, in *42nd AIAA Fluid Dynamics Conference* 2012: New Orleans, LA.
36. Williams, C.S. and O.A. Becklund, *Introduction to the optical transfer function*. 1989: Wiley New York etc.
37. Ng, R. *Fourier slice photography*. in *ACM Transactions on Graphics (TOG)*. 2005. ACM.

38. Bracewell, R.N., *Strip integration in radio astronomy*. Australian Journal of Physics, 1956. **9**(2): p. 198-217.
39. Macovski, A., *Medical imaging systems*. Vol. 20. 1983: Prentice-Hall Englewood Cliffs, NJ.
40. Nishimura, D.G., *Principles of magnetic resonance imaging*. 1996: Stanford University.
41. Deans, S.R., *The Radon transform and some of its applications*. 2007: DoverPublications.com.
42. Averbuch, A. and Y. Shkolnisky, *3D Fourier based discrete Radon transform*. Applied and Computational Harmonic Analysis, 2003. **15**(1): p. 33-69.
43. Chen, G.-H., S. Leng, and C.A. Mistretta, *A novel extension of the parallel-beam projection-slice theorem to divergent fan-beam and cone-beam projections*. Medical physics, 2005. **32**: p. 654.
44. Liang, Z.-P. and D.C. Munson Jr, *Partial Radon transforms*. Image Processing, IEEE Transactions on, 1997. **6**(10): p. 1467-1469.