


**Assessment of radial image distortion and
spherical aberration on three-dimensional
synthetic aperture particle image velocimetry **ARCHIVES**
measurements**

by

Daniel Mark Kubaczyk



S.B., Mechanical Engineering, Massachusetts Institute of Technology
(2011)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2013

© Massachusetts Institute of Technology 2013. All rights reserved.

Author
.....
Department of Mechanical Engineering
January 18, 2013

Certified by
.....
Douglas P. Hart
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by
.....
David E. Hardt
Ralph E. and Eloise F. Cross Professor of Mechanical Engineering
Chairman, Department Committee on Graduate Students

**Assessment of radial image distortion and spherical
aberration on three-dimensional synthetic aperture particle
image velocimetry measurements**

by

Daniel Mark Kubaczyk

Submitted to the Department of Mechanical Engineering
on January 18, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

This thesis presents a numerical study of the effects of radial image distortion and spherical aberration on reconstruction quality of synthetic aperture particle image velocimetry (SAPIV) measurements. A simulated SAPIV system is used to image a synthetic particle volume. An idealized pinhole camera model is used for image formation with distortion and spherical aberration being added with a polynomial model and a Fourier waveform model, respectively. Images from a simulated 5 x 5 camera array are taken, distorted and/or aberrated, realigned and averaged to form synthetic aperture images at a set of depths within the particle volume. These images are thresholded to recover three-dimensional particle locations and a reconstructed three-dimensional intensity field is formed. This reconstructed field is then evaluated according to intensity data and a signal-to-noise ratio (SNR). Results show that even small amounts of image distortion and spherical aberration can lead to degradation of SNR and information loss.

Thesis Supervisor: Douglas P. Hart
Title: Professor of Mechanical Engineering

Acknowledgments

- Professor Doug Hart for providing the inspiration, encouragement and tutelage for this thesis.
- Tom Milnes for technical guidance and for having the best fashion sense in the Department.
- Jesse Belden for technical support and for building the tool that made this work possible.
- Eric Shields for guidance on implementing spherical aberration in MATLAB, saving several weeks of my life.
- Professor Alex Techet, Barry Scharfman and Leah Mendelson for technical guidance, support and for showing me just how difficult (and fun!) experimental work can be.
- The Bubs and Bubettes for preserving sanity and destroying it.
- Rodzice za wszystko, nawet jeśli nie zrozumiecie nic z tego.
- Kris and Mark for being the giants you are and for always making me feel that I'm in the middle.
- Sophie for severely diminishing my productivity and sneakily eating my food.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

Contents	7
List of Figures	9
List of Tables	15
1 Introduction	17
1.1 Particle image velocimetry methods	19
1.1.1 Tomographic PIV	20
1.1.2 Holographic PIV	21
1.1.3 Defocusing digital PIV	22
1.1.4 Particle tracking velocimetry	23
1.1.5 Synthetic aperture particle image velocimetry	24
1.2 Accuracy of PIV systems	27
2 Synthetic aperture fundamentals	31
2.1 History	31
2.2 Theory	32
2.2.1 Image acquisition	35
2.2.2 Image refocusing	35
2.2.3 Three-dimensional reconstruction	37
2.3 Implementation in a PIV system	38
2.3.1 Equipment	38
2.3.2 Calibration	40

2.3.3	Image acquisition, refocusing and thresholding	40
3	Implementation of distortion and spherical aberration in a SAPIV system	49
3.1	Radial image distortion	50
3.2	Spherical aberration	52
3.3	Simulated SAPIV system	55
4	Results	61
4.1	Correlation	61
4.2	Rank correlation	66
4.3	Histogram analysis	67
5	Conclusion	77
	Bibliography	79
A	Distortion parameter calculation	83
B	Point spread function derivation	87

List of Figures

- 1-1 A flow visualization—similar to those of Prandtl’s—of separation past a wing using tracer particles in water. 18
- 1-2 A common two-dimensional PIV experimental setup with a single camera imaging particle displacement with two images at different times (t & t'). Note how laser sheet illumination determines the plane of particles which are visible in the image. 19
- 1-3 A two-dimensional schematic of tomographic reconstruction. Given enough views of an object space an accurate reconstruction can be determined from the intensity summation along each line on a particular image sensor. 21
- 1-4 A simple schematic of a holographic PIV system. Spatially coherent light generates predictable diffraction patterns (hologram) that are modified to reconstruct objects at different depths. 22
- 1-5 Comparison of a typical optical system (a), and a defocusing optical system (b) used to determine the depth of objects along the optical axis. Notice that light rays from B focus at C but generate a defocused blur at the image plane. The size of this blur and how it varies with placement of the reference plane can allow for determination of depth. 23
- 1-6 A schematic of an experimental PTV system. Cameras are mounted in different positions and orientation to obtain a stereoscopic effect. . 24

1-7	A schematic of an experimental setup of a synthetic aperture PIV system. Cameras are typically confined to one side of volume of interest such that their perspectives overlap enough to compare when refocusing.	25
1-8	A schematic of how synthetic aperture refocusing works. Since both objects in the volume of interest are visible to at least four out of the five cameras, both objects can be clearly seen after refocusing. Objects at different depths fade out as they are located well outside the narrow depth of field.	26
1-9	A simple example of how a synthetic aperture PIV system can image objects that are partially occluded. The unobstructed object (left) is difficult to make out when occluded (center) with a conventional image but becomes much clearer after refocusing (right).	27
2-1	A unique perspective of a simple light field. The viewpoint, (X_i, Y_i, Z_i) , and the UV -plane define a bundle of light rays which are sampled to generate an image. If an entire light field is known, then any image from any perspective can be generated.	33
2-2	The effect of parallax when using multiple viewpoints (eyes or cameras). Nearby objects require the viewpoints to converge more than objects located further away (i.e. $\beta > \alpha$).	34
2-3	A schematic showing how image views can be shifted using a homography, H . The object of interest, x_π , is viewable by both camera views, C and C' .	36
2-4	An image of a nine-camera array along with the supporting 80/20 [®] structure and camera mounts. The orientation of the cameras is such that their respective fields of view intersect at the volume of interest.	39
2-5	Two images of planar calibration grids used in SAPIV experiments. Calibration grids are typically made up of a consistent geometrical pattern, whether it be an alternating two-dimensional checkerboard (a) or an evenly-spaced grid of circles (b).	41

2-6	A sampling of four images from a camera array at a specific point in time. The experiment sought to better understand turbulent flow sheet breakup for water in air.	42
2-7	An example of a raw SAPIV image from a vortex ring quantification experiment. The section of particles being illuminated by the laser is clearly visible.	42
2-8	A refocused view of the multiphase flow experiment in figure 2-6. The ligament in the lower-left quadrant of the image is the only identifiable object in focus at this plane.	43
2-9	A refocused view of the PIV experiment from figure 2-7. An underlying low-intensity blur makes up most of the image but many particles of high intensity and high contrast are visible.	44
2-10	A normalized intensity histogram of an individual refocused image. Due to the large number of out-of-focus particles, the distribution of intensity in any particular refocused image converges towards a normal distribution at the low-intensity range of the distribution.	46
2-11	A thresholded refocused PIV image. This image represents a best estimate as to which particles are located where at this particular focal plane.	47
3-1	Radial image distortion in synthetic images of a rectangular particle volume with a square front face. An undistorted view (<i>a</i>) produces an image which accurately depicts a symmetric square. A view distorted with 25% barrel distortion (<i>b</i>) shifts pixel intensity values radially inward relative to their undistorted pixel locations. A view distorted with 25% pincushion distortion (<i>c</i>) shifts pixel intensity values radially outward relative to their undistorted pixel locations.	52
3-2	Lens with spherical aberration focusing incident, non-paraxial rays (red) along a range of points. The imperfect focusing of rays causes an image point to spread and reduce its peak intensity.	54

3-3	An unaberrated, diffraction-limited PSF (<i>a</i>) with higher concentrated intensity pattern and an aberrated PSF (<i>b</i>) ($RMS = 0.15$) showing intensity spread away from the central peak to the second annular lobe. Assumes a point source of light with unit intensity.	56
3-4	A non-aberrated image (<i>a</i>) of a simulated calibration target and an aberrated image (<i>b</i>) of a simulated calibration target. Note the resolution degradation and loss of high spatial frequency information in the aberrated image.	57
3-5	Synthetic images taken of 50mm x 50mm x 10mm particle volume from four corner cameras of the square 5 x 5 array. Image from top-left camera (<i>a</i>); image from top right-camera (<i>b</i>); image from bottom-left camera (<i>c</i>); image from bottom-right camera (<i>d</i>).	59
3-6	Refocused synthetic image of a particular depth within the imaged volume. The low-intensity background blur is due to unfocused particles at various other depths in the volume.	60
3-7	Thresholded version of image 3-6 retaining only pixels of relatively high intensity. Finding the proper threshold is critical for maximizing inclusion of desired, at-depth particles and minimizing inclusion of false particles.	60
4-1	Reconstruction quality as provided by correlation as a function of seeding density and distortion D . It shows that significant reduction in reconstruction quality occurs with more than a few tenths of a percent of radial image distortion.	63
4-2	Reconstruction quality as provided by correlation as a function of seeding density and spherical aberration SA . It shows that a much smaller reduction in reconstruction quality occurs with the inclusion of spherical aberration effects than with distortion effects.	64

4-3	Reconstruction quality as provided by correlation as a function of seeding density, distortion D and spherical aberration SA . It shows that the effects of distortion and spherical aberration are mostly independent.	65
4-4	Reconstruction quality as provided by rank correlation as a function of seeding density and distortion D . Results show that distortion has only a small effect at 2.0% and a negligible effect for distortions below that. The range of correlation values is greater than 0.97, which is much higher than the values given for standard correlation.	68
4-5	Reconstruction quality as provided by rank correlation as a function of seeding density and spherical aberration SA . Results show that spherical aberration has no discernible effect.	69
4-6	Mean of the normal fit of the intensity distribution as a function of seeding density and distortion D . No significant effect is evident; this agrees with the theoretical prediction that the mean intensity of distorted images should not vary.	70
4-7	Mean of the normal fit of the intensity distribution as a function of seeding density and spherical aberration SA . No significant effect is evident; this agrees with the theoretical prediction that the mean intensity of aberrated images should not vary.	71
4-8	Standard deviation of the normal fit of intensity data as a function of distortion D and particle density. The standard deviation decreases for distortion percentages 0.5% and above and also increases with higher particle density as more particles create much more non-uniform background noise.	72
4-9	Standard deviation of the normal fit of intensity data as a function of spherical aberration SA and particle density. The standard deviation decreases for SA values 0.02 and above and also increases with higher particle density as more particles create much more non-uniform background noise.	73

4-10 Normalized signal-to-noise ($nSNR$) ratio as a function of distortion D and particle density. All SNR values are normalized with respect to the zero-distortion ($D = 0\%$) case. 74

4-11 Normalized signal-to-noise ($nSNR$) ratio as a function of spherical aberration SA and particle density. All SNR values are normalized with respect to the zero-aberration ($SA = 0$) case. 75

A-1 A schematic of how pixel intensity value shift according to the amount of distortion applied to an image. Positive distortion represents a radially-inward shift while negative distortion represents a radially-outward shift. 84

B-1 The front half of the optical system collimates the incoming wavefronts scattered off the object into plane waves (ideally). This is where the wavefront takes on the aberrated form seen in equation 3.3 since it can deviate from a strict plane wave. 89

List of Tables

- 3.1 Pixel shift for various distortion values at various normalized image radii for a 1,000 x 1,000 pixel image. The normalized radius r_n is defined as the image radius divided by the maximum image radius for a rectangular image. 53
- 3.2 A representation of how the intensity distribution from the diffraction-limited, aberration-free PSF from figure 3-3a would be measured on an image sensor with 10 μm pixel pitch. This assumes unit input intensity; the sum of all values equals 1. 56
- 3.3 Comparison of PSFs from each of the four *RMS* values used to generate PSF with different amounts of spherical aberration. As the amount of the spherical aberration increases, intensity moves away from the center pixel, where all intensity should be localized, to the first and second rings. 57

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

The study of fluid flow visualization was born in 1904 when Ludwig Prandtl, a German fluid mechanician, inserted a suspension of mica particles into the water flowing through his experimental fluid tunnel. An example of flow visualization can be seen in figure 1-1. Although only able to analyze qualitative changes in flow properties, Prandtl founded the modern method of flow visualization known as particle image velocimetry (PIV). [1]

Particle image velocimetry is an optical imaging method of fluid flow visualization and is used to characterize and quantify relevant flow properties such as velocity, vorticity and turbulence. To visualize a fluid continuum, discrete tracer particles are seeded within the flow. These particles are small enough such that their characteristic inertial force is relatively small compared to the dynamic forces of the fluid flow. As a result, these fluid particles follow the flow field as they move through the fluid volume of interest. High-intensity illumination is often applied to the fluid volume to enhance contrast between tracer particles and the fluid. Images are typically taken with digital cameras that utilize high-quality, photographic lenses and light-sensitive image sensors to record intensity information about the scene. Because the intensity information stored in these images is recorded digitally, images can be analyzed computationally to determine fluid flow properties. An example experimental setup is shown in figure 1-2. Typically, cross-correlation methods are used to compare images of the fluid volume in two time steps separated by a small fraction of a second. The




Figure 1-1: A flow visualization—similar to those of Prandtl’s—of separation past a wing using tracer particles in water [1].

peak fit of the cross-correlation between two images corresponds to the most likely displacement in a set of tracer particles. Given the difference in time from when the first and second images were taken as well as the most likely distance traveled by the particle, the fluid flow velocity can be determined for either two or three spatial dimensions. PIV is advantageous to other methods such as hot wire anemometry [2] or laser doppler velocimetry [3]. Unlike hot wire anemometry, PIV is almost completely non-obtrusive to the fluid flow. Measurements are taken visually on a charge-coupled device (CCD) or complementary metal-oxide semiconductor (CMOS) sensor from tracer-particle-scattered electromagnetic radiation. Neither the particle nor the intense illumination affect the fluid flow in any significant way. Unlike laser doppler velocimetry, PIV has the potential to measure an entire velocity field at points in two or three spatial dimensions with two or three velocity vector components at each point. This has the potential to characterize all of the available kinetic information about the flow field (three velocity vector components for each point in three spatial dimensions). This information can allow for accurate particle trajectory estimation as well as calculation of kinematic flow field properties like vorticity and circulation, which cannot be measured directly.

The following subsections will provide brief descriptions of a range of PIV methods




Figure 1-2: A common two-dimensional PIV experimental setup with a single camera imaging particle displacements with two images at different times (t & t'). Note how laser sheet illumination determines the plane of particles which are visible in the image [1].

to help understand the trade-offs and limitations of PIV systems.

1.1 Particle image velocimetry methods

The system shown in figure 1-2 represents the setup for a two-dimensional PIV system. These systems, while easy to implement and simple to operate, can typically only image a particular planar surface within the volume. This precludes capturing a three-dimensional velocity field without using an array of cameras focused on different depths within the flow. Also, given that a camera's depth of field has a discrete thickness, it can be problematic to determine the exact depth at which particles are located. More complex systems require use of multiple connected cameras but can effectively image the entirety of complex flow.

PIV is an experimental technique used to characterize flows that are too complex to model analytically. As a result, the ability to measure flows in multiple dimensions with multiple components is critical. The systems described in the following

subsections describe useful methods for carrying out three-dimensional PIV analyses. They vary mostly in the method of visual information capture and subsequent reconstruction. The algorithms used to determine velocity fields all employ some form of cross-correlation and require image pre- and post-processing techniques to ensure images are normalized for exposure intensity and contrast. Understanding and measuring complex three-dimensional flows is relevant in resolving complex dynamical fluid behavior in applications ranging from fuel combustion [4] to the mechanics of bat flight [5].

1.1.1 Tomographic PIV

Tomographic PIV is a three-dimensional PIV method that utilizes multiple views of a volume of interest to reconstruct that three-dimensional volume using optical tomography [6]. Tomography is an imaging technique that can visualize the cross-section of an object by transmitting radiation through the object along a line at varying angles. This is the same method that is used in computed tomography (CT) scans. The intensity of the radiation is attenuated differently based on what objects lie in the path of the radiation. By summing each of these intensity attenuation records at every angle, a two-dimensional, cross-sectional view can be reconstructed, as seen in figure 1-3.

Similarly, if radiation is transmitted through a volume and captured along a plane, such as a CCD sensor, then by imaging multiple views simultaneously, a three-dimensional volume can be reconstructed. The constituent pieces of this reconstructed volume are known as “voxels” and can be thought of as pixels with discrete thickness. Once volumes from two different time instants are reconstructed, they can be analyzed using cross-correlation techniques of what are known as interrogation windows, or small sub-volumes within the overall volume, that are small enough to be amenable to the assumptions required for determination of particle displacement.




Figure 1-3: A two-dimensional schematic of tomographic reconstruction. Given enough views of an object space an accurate reconstruction can be determined from the intensity summation along each line on a particular image sensor [6].

1.1.2 Holographic PIV

Holographic PIV [7, 8] is an imaging technique that can reconstruct a volume that is generated from the interference of coherent light that has been scattered by non-uniform objects within a volume of interest as seen in figure 1-4. The interference pattern contains information that can reliably reconstruct the original volume by illuminating the interference pattern (known a hologram) with a coherent light source that is similar to the one used to generate the interference pattern. Typically, interference patterns for arbitrary objects can be rather complex and reconstructing the intensity field from a hologram can prove problematic. However, given the consistent geometry and size of tracer particles used to seed a flow, a fairly deterministic interference pattern emerges which is straightforward to reconstruct into the original intensity field.

Holographic techniques typically utilize just a single camera but can still reconstruct volumes in three spatial dimensions. Given that the size of an interference




Figure 1-4: A simple schematic of a holographic PIV system. Spatially coherent light generates predictable diffraction patterns (hologram) that are modified to reconstruct objects at different depths [7].

pattern from an object varies with distance from the camera, depth information can be encoded in just a single two-dimensional image. However, the volume must be relatively sparse in order for reconstruction to be successful. A dense volume would preclude reconstruction of occluded particles in the volume. Holographic PIV systems also require a complex and delicate optical set-up and are sensitive to environmental factors. [9]

1.1.3 Defocusing digital PIV

Defocusing digital PIV [9] is an imaging technique that uses an asymmetric non-standard aperture, known as a mask, to obtain multiple images of an object. If the object is in focus then it forms a clear image at the focal plane. If the object is out of focus, then the light rays emanating from any point on the object will contact the image sensor at different points as seen in figure 1-5. The distance between these out-of-focus points is determined by the distance away from the plane that is in

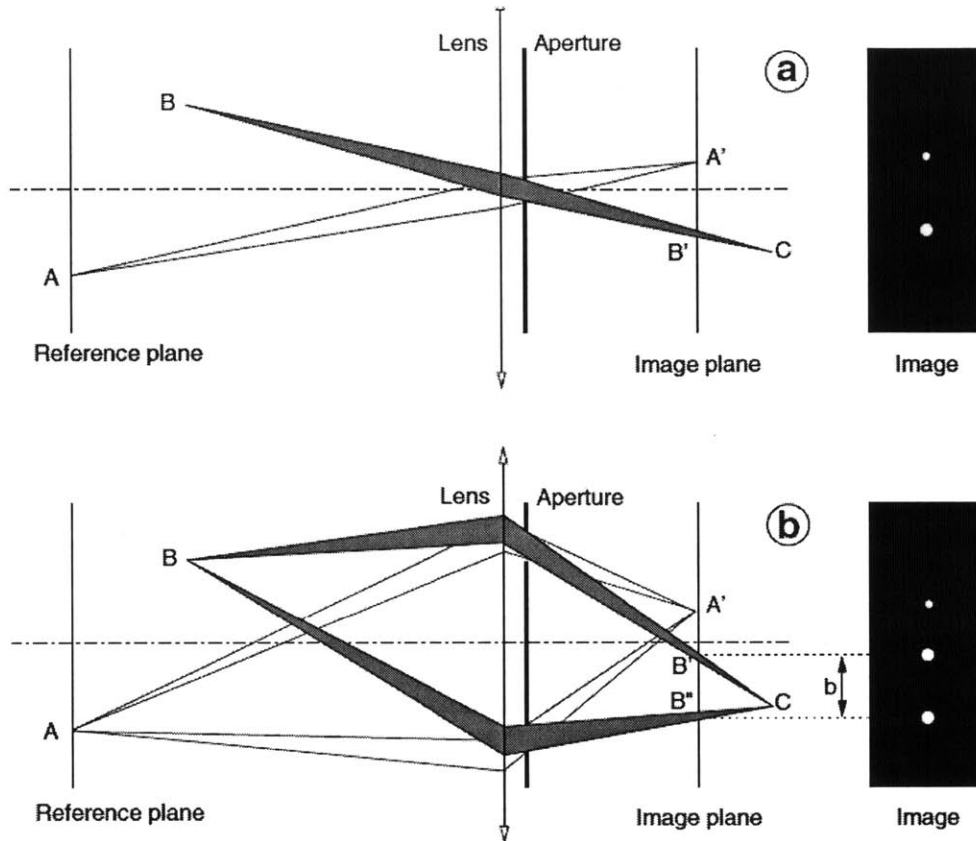


Figure 1-5: Comparison of a typical optical system (a), and a defocusing optical system (b) used to determine the depth of objects along the optical axis. Notice that light rays from B focus at C but generate a defocused blur at the image plane. The size of this blur and how it varies with placement of the reference plane can allow for determination of depth [9].

focus. Thus an object that appears to be more out of focus is further away from the focal plane. Using geometric optics, the depth of that object can be determined by measuring how blurred it is over a range of focal planes.

1.1.4 Particle tracking velocimetry

Particle tracking velocimetry [10] is an imaging technique that utilizes a stereoscopic recording of tracer particles to record flow paths as seen in figure 1-6. Unlike PIV methods, PTV determines the actual trajectory that tracer particles follow. This requires reliable identification of individual particles to ensure accurate trajectory recording is maintained. This method is limited to relatively low particle densities




Figure 1-6: A schematic of an experimental PTV system. Cameras are mounted in different positions and orientation to obtain a stereoscopic effect [10].

since having too many particles can make it prohibitively difficult to track particle between recorded images. Particle trajectories are interpolated from recorded data via post-processing techniques.

Particle tracking velocimetry can be implemented effectively as long as particle seeding density is relatively low. This prevents characterization of flows where high resolution is needed [9].

1.1.5 Synthetic aperture particle image velocimetry

Synthetic aperture particle image velocimetry (SAPIV) is an imaging technique that utilizes an array of cameras to simultaneously record multiple images of a volume of interest as seen in figure 1-7. SAPIV is a practical and effective measurement tool for visualizing self-occluding, densely-seeded flows [11]. An array of cameras oriented at a volume of interest captures images from different perspectives allowing for calculation of depth from parallax and accurate whole volume three-dimensional reconstruction. Captured images are then digitally shifted using pre-defined calibration data that solves for the relative position and orientation of all cameras in the array. Images can be shifted such that they appear to have been taken from the same position and orientation.

Two significant advantages are apparent: first, particles occluded to some cameras




Figure 1-7: A schematic of an experimental setup of a synthetic aperture PIV system. Cameras are typically confined to one side of volume of interest such that their perspectives overlap enough to compare when refocusing [11].

and imaged by others will likely be recorded in the reconstructed image as seen in figure 1-8 and second, use of many cameras in many different positions creates a large, effective (synthetic) aperture which drastically reduces depth of field and allows for accurate imaging of thin slices through the volume and recording of particle depth information. This large synthetic aperture avoids the depth of field issues that arise in two-dimensional PIV. A real world synthetic aperture image can be seen in figure 1-9. It shows how objects that are partially occluded can be made out clearly after computational refocusing.

SAPIV performs comparably to other three-dimensional PIV systems. An analysis by Belden [11] showed that SAPIV performed similarly to tomographic PIV, [6] holographic PIV, [7] defocusing digital PIV [9] and particle tracking velocimetry (PTV) [10] with respect to number of particles imaged normalized to the sum of all pixels in the imaging system. While SAPIV typically requires 8 - 13 cameras for optimal reconstruction and alternative techniques generally require 5 or fewer, the ability to resolve highly-seeded, occluded flows is unique to SAPIV [11]. SAPIV has the capability for effective functionality in confined spaces. Skupsch and Brücker [12] showed SAPIV can be implemented using a single camera and lens array. Such com-




Figure 1-8: A schematic of how synthetic aperture refocusing works. Since both objects in the volume of interest are visible to four out of the five cameras, both objects can be clearly seen after refocusing. Objects at different depths fade out as they are located well outside the narrow depth of field.




Figure 1-9: A simple example of how a synthetic aperture PIV system can image objects that are partially occluded. The unobstructed object (left) is difficult to make out when occluded (center) with a conventional image but becomes much clearer after refocusing (right) [11].

pact setups can be advantageous, especially in applications where perspectives of the scene are confined to small viewing windows such as portholes for nuclear reactors.

SAPIV is of particular interest for the purposes of this report. Proceeding chapters will focus further analysis and discussion on this method alone.

1.2 Accuracy of PIV systems

Though three-dimensional PIV systems have the capability and flexibility to image a wide range of flows they still must use real-world optics for imaging and must therefore be calibrated in spite of the characteristic aberrations and distortions of multi-element glass lenses. These aberrations arise not only from manufacturing flaws but also from the imperfect focusing performance of spherical lenses. Most high-quality photographic lenses incorporate many individual glass elements that seek to minimize many different forms of aberration. Ultimately, photographic lenses are very well designed and manufactured but cannot achieve diffraction-limited imaging performance. Furthermore, many PIV experimental setups have a non-air optical medium between

the cameras and the object being imaged. These can include the glass wall of a water tank and even the water itself. These additional optical “elements” distort light rays through refraction and render the linear pinhole model inadequate for high-accuracy mapping. This is especially true for off-axis light rays in wide-angle lenses [13].

Three-dimensional PIV systems typically require sub-pixel mapping accuracy and in some cases, such as tomographic PIV, less than a tenth of a pixel [6]. Imperfect three-to-two dimensional mapping can occur as a result of resolution-degrading aberrations (i.e. spherical aberration, coma) as well as image distortions. Use of the linear pinhole model is often used to determine a camera’s defining geometric parameters even though it fails to model aberrations and distortions. Many techniques exist to correct for image distortion [14, 15, 16] in a single camera setup but imaging with multiple cameras requires full volume calibration to avoid excessive mapping error [17]. Wieneke [17] showed that full volume self-calibration can improve accuracy for tomographic PIV systems as well as stereoscopic PIV and PTV systems. However, volume self-calibration does not work nearly as well in densely-seeded flows and the improved measurement accuracy seen in these systems underscores the need to improve planar calibration techniques used for SAPIV [11].

Planar techniques typically involve recording many images of a planar calibration grid—a checkerboard pattern is commonly used—at various configurations with a volume. These configurations attempt to “fill” the volume with points that are directly calibrated to the image planes of cameras but ultimately fall short of achieving total volume calibration. Interpolation is required to map all points within the volume and this interpolation is often inaccurate in light of the complex interaction of multi-camera image distortion.

Mapping errors significantly degrade reconstruction quality and a major source of these errors is optical aberrations. The primary set of aberrations for optical systems is the Seidel aberrations: spherical aberration, coma, astigmatism, field curvature, chromatic aberration and image distortion. All but image distortion are resolution-degrading aberrations (distortion is a mapping-degrading aberration); we have chosen to examine spherical aberration here due to its prevalence and ease of implementation

in a simulated SAPIV system. Image distortion is also included in the analysis as its effect on mapping accuracy has historically proven to be significant. The current work seeks to quantify the effects of radial image distortion and spherical aberration in terms of intensity histogram data and a signal-to-noise ratio.

This thesis focuses on an analysis of the effects of optical aberrations on the reconstruction capability of SAPIV. Chapter 2 covers fundamentals of synthetic aperture methods including a brief history and overview of its uses and underlying theory. Chapter 3 focuses on the implementation of a simulated SAPIV system for investigation of spherical aberration and radial distortion. The theory of optical aberrations and wavefront imaging will be summarized. Chapter 4 presents the simulation results and provides analysis in light of commonly used reconstruction metrics like correlation, rank correlation as well as histogram analysis. Finally, Chapter 5 will provide discussion and conclusions on the results presented as well as direction for future work.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Synthetic aperture fundamentals

This chapter introduces the synthetic aperture method as a form of *light-field imaging* and the manner in which it was developed and is currently used. This information will provide context for the analysis and results described in chapters 3 and 4.

2.1 History

The idea of a synthetic aperture originated at Mullard Radio Astronomy Observatory at Cambridge University in the 1950s [18]. Martin Ryle, the co-winner of the 1974 Nobel Prize in Physics, developed a method to combine radio signals from different points of observation in what was called “aperture synthesis” [19]. The large wavelength range of radio frequencies makes it difficult to image faint, astronomical objects with very low angular breadth. To give a sense of scale, a human with 20/20 vision can resolve an object with angular resolution of roughly one arc-minute. The moon is roughly 30 arc-minutes when viewed from Earth. For objects being imaged with radio wavelengths (on the order of tens of centimeters), an angular breadth of 1 arc-minute would require a telescope with an 800 meter diameter. The cost of such an instrument would be prohibitive for most research institutions; therefore, a clever way to synthesize signals from many small telescopes was developed by Ryle and his colleagues.

Radio waves emanating from distant objects reach Earth as coherent waveforms

and the interaction of those waveforms with the atmosphere produces some scattering which gives a non-ideal interference signature [20]. Both amplitude and phase information on the incoming wavefront is required; therefore it is imperative that the distance between telescopes is known and calibrated. Furthermore, Ryle developed a way to use the rotation of the Earth to generate additional perspectives; calibration now incorporated the geometry of the telescopes as well as the rotation and geometry of the Earth. With the use of multiple perspectives from each telescope in the system, signal noise from the atmosphere could be reduced in the form of a more concentrated point-spread function. In addition to reducing noise, each perspective could add a finite amount of signal intensity which strengthens the incoming signal. The computational power to carry out advanced Fourier waveform transform methods did not arise until the late 1960s and as a result, telescope systems with effective apertures of 5 kilometers were developed [20].

2.2 Theory

The theory underlying the development of synthetic aperture imaging using visible light for human and machine vision applications originated from the work of Edward H. Adelson [21]. Adelson developed the *plenoptic function* which parameterizes all the variables in a *light field*. A light field is a description of a spatial domain that contains visible electromagnetic radiation traveling in varying directions at various points in time. It describes underlying visual properties of a three-dimensional scene like motion, color, orientation and brightness. The plenoptic function,

$$P = P(X, Y, Z, U, V, \lambda, t) \tag{2.1}$$

describes a field of light rays and can generate unique images at each defined view point in three-dimensional space (X, Y, Z) . Each light ray can be fully defined with three parameters for space (X, Y, Z) , two parameters for orientation (U, V) , one parameter for wavelength λ and one parameter for time t . An example derivation of




Figure 2-1: A unique perspective of a simple light field. The viewpoint, (X_i, Y_i, Z_i) , and the UV -plane define a bundle of light rays which are sampled to generate an image. If an entire light field is known, then any image from any perspective can be generated. Conversely, many images taken inside a light field could theoretically be used to reconstruct the entire field.

a unique perspective on a simple light field is shown below in figure 2-1. In most cases, the full set of parameters of a light field are not known and must be determined through rigorous reconstruction. The methods used in this thesis carry out this process of light field reconstruction.

Use of the plenoptic function [21] spawned a great deal of research into capturing light fields as well as rendering from three-dimensional animations. Adelson and Wang developed the first known *light field camera* by using a thin multi-lens array positioned in front of a camera's image sensor [22]. This allowed for analysis of multiple visual perspectives from within the housing of a single camera body. Typical cameras with a circularly-symmetric aperture behave like a single human eye: able to focus on objects over a great range of distances but not easily able to perceive depth. Adelson and Wang's camera behaved like a two-dimensional array of low-resolution eyes, able to perceive depth through the parallax effect as in figure 2-2. Humans can perceive depth because the angle of convergence of human eyes determines the distance from a viewable object. Objects located an effectively infinite distance away—stars, for example—orient a pair of eyes in a parallel configuration such that the angle of convergence is zero. Others have also developed image-based rendering systems for digital animation of visually-rich, three-dimensional environments [23, 24].

The application of aperture synthesis to plenoptic imaging for applications us-

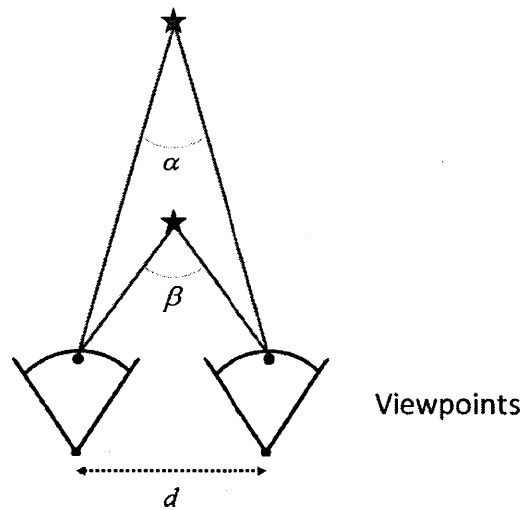


Figure 2-2: The effect of parallax when using multiple viewpoints (eyes or cameras). Nearby objects require the viewpoints to converge more than objects located further away (i.e. $\beta > \alpha$). If the distance d between the center of the two viewpoints is known, then distance from the viewpoints to the objects can be determined. This effect allows cameras with multiple perspectives, or systems of multiple cameras, to image objects and reconstruct their geometry and location in three-dimensional space.

ing visible light was first explored by Marc Levoy and has been further developed by numerous colleagues including Aaron Isaksen and Ramesh Raskar. The driving principle behind the synthetic aperture method is the same as that which drove the development of aperture synthesis for radio telescopes: the extension of an imaging aperture to access more information while avoiding prohibitively large imaging instruments and their associated costs. Therefore, an array of traditional cameras with overlapping fields of view can be used to image a particular object or volume of interest to drastically increase the visual information provided relative to a single typical camera. However, use of many cameras requires a method of linking the cameras such that their relative positions and orientations are known as well as the computational software to make this link and modify images to achieve a desired reconstruction effect [25, 26, 27, 28].

The process of the application of a synthetic aperture can be broken up into three components:

- Image acquisition

- Image refocusing
- Three-dimensional reconstruction

2.2.1 Image acquisition

Image acquisition, as seen in figure 1-8 can be carried out using common, off-the-shelf imaging equipment including consumer-level photographic cameras and optics. It is critical that the images taken of the object of interest have a wide enough depth-of-field such that the entire object is in focus. Reconstruction results will be suboptimal if parts of the objects are not in focus in any of the raw images. High-resolution images will enhance the information throughput in the refocusing and reconstruction phase. Typically, if the object being imaged is time-varying, these images are taken at the same instant in time using electronic triggering equipment. However, a single moving camera imaging the same static object at multiple points in time can achieve the same effect. Maintaining consistency is critical for comparing images of the same object in the refocusing phase. Before images can be recombined, the camera array must be calibrated. Calibration can take many forms including imaging an object of known geometry [25, 27, 28], calibrating from the object itself [17] or controlling the position and orientation of the cameras using electronic sensors and/or mechanical stops such that calibration need not require image information.

2.2.2 Image refocusing

Image refocusing, as seen in figure 1-8, takes the raw images from the acquisition phase and along with an accurate calibration, reorients and averages the images together along a particular focal plane [25, 27, 28]. The first step, reorientation, requires an image transformation known as a *homography* [29]. Figure 2-3 shows schematically how an image reorientation is carried out. A homography is defined by a square matrix, h_{ij} , which maps image coordinates from one view to another




Figure 2-3: A schematic showing how image views can be shifted using a homography, H . The object of interest, x_π , is viewable by both camera views, C and C' [29].

$$\begin{bmatrix} bx' \\ by' \\ b \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (2.2)$$

where x and y represent pixel coordinates in the raw images and x' and y' represent the shifted pixel coordinates. A unique homography is required for each unique view that is shifted to the *reference view*. The reference view is the perspective of the final refocused image and is the view to which all raw images are shifted to. The scalar b represents a scaling factor that allows for magnification of the shifted view relative to that of the original view. The values in the homography matrix h_{ij} are obtained through calibration. The reference view is the single view to which all of the raw images are shifted to generate a single refocused view. The reference view can have arbitrary orientation such that it could be at an arbitrary angle, translation or magnification to any of the views of the cameras [30]. This refocused view has the benefit of the large effective aperture of the multi-camera system: it can visualize objects occluded to some cameras and it can focus on a very thin plane within the volume with significantly reduced out-of-focus noise as seen in figure 1-9.

Once united in the reference view, the images can be combined to yield the refocused image [25, 27, 28]. The simplest method for image combination is additive averaging [31]

$$I_{R_k} = \frac{1}{N} \sum_{i=1}^N I_{S_{ki}} \quad (2.3)$$

where I_{R_k} is the refocused image combined from all i shifted images $I_{S_{ki}}$ at a particular focal plane k . The focal plane that is imaged is specified beforehand and based on the mapping derived from the calibration, the relevant pixels are picked from the raw image and shifted to the perspective of the reference view [25, 27, 28]. Objects which fall along the focal plane should generate a high intensity value as that object should show up in all or most of the shifted images. After averaging, that object should retain all or most of its original intensity. Items which are not located at the focal plane will form an out-of-focus blur but because the orientation of each view is different, the blur between images should be mostly uncorrelated and will thus generate a relatively low intensity value relative to the in-focus, on-plane objects. Other methods for combining images exist; a commendable alternative is multiplicative averaging [31].

$$I_{R_k} = \left(\prod_{i=1}^N I_{S_{ki}} \right)^n, \quad 0 < n < 1 \quad (2.4)$$

where multiplicative averaging retains an additional degree-of-freedom in the exponent, n . The value set for the exponent will vary between systems depending the quality of object-to-scene mapping, the visual layout of the object of interest, the number of cameras in the system and other factors. In theory, the exponent should take on the value of $\frac{1}{N}$ but systems with imperfect calibration will require a larger exponent to retain information. Multiplicative refocusing yields a higher signal-to-noise ratio than additive refocusing provided that the system is well-calibrated. [31]

2.2.3 Three-dimensional reconstruction

Refocused images can be filtered using a variety of techniques to eliminate out-of-focus noise while avoiding elimination of in-focus objects. Once filtering of refocused

images is complete, the set of refocused images can be brought together to form a reconstructed volume where each refocused image represents a “slice” through the volume at a discrete plane. This set of discrete images can be combined into volume through three-dimensional smoothing and post-processing to form continuous three-dimensional solids.

2.3 Implementation in a PIV system

2.3.1 Equipment

An experimental synthetic aperture PIV (SAPIV) system generally requires the following equipment:

- Camera array
- Trigger and wiring
- High-intensity illumination
- Computer with image processor

A camera array, seen in figure 2-4, can contain any number of cameras. Arrays for synthetic aperture methods have had up to 100 cameras [32]. Cameras are typically selected to have focal lengths in the 50-110 mm range which provides an effective combination of zoom and field of view. Shorter focal length lenses also tend to induce more distortions and aberrations. Cameras are mounted on adjustable mounts such that all relevant spatial degrees-of-freedom can be controllably modified including pitch, yaw and horizontal/vertical translation. Camera mounts must be able to be locked securely to eliminate any camera movement during calibration and experimentation. Any movement, even a minimal change in orientation of a single camera or a slight nudge of the camera array structure, can completely ruin a data set and require the calibration process to be redone. Electronic triggering is required to sync all cameras together such that they take images at the same point in time. Cameras are



Figure 2-4: An image of a nine-camera array along with the supporting 80/20[®] structure and camera mounts. The orientation of the cameras is such that their respective fields of view intersect at the volume of interest.

typically wired into a computer and a software program enables in-sync camera actions. However, focusing as well as lens zoom and orientation adjustment is typically done by hand.

High-intensity illumination is typically synced with the image exposure of cameras in the array. Illumination is typically supplied by a laser or light-bank pulse. The duration of the pulse is much shorter than the duration of the image exposure; this allows for a much smaller effective exposure time. Laser illumination is used more frequently for PIV applications [11] where resolution of fine tracer particles is critical. Experiments using light-banks of LEDs tend to have larger volumes of interest and are more interested in bubbly flows and bubble size characterization [33].

A computer is necessary to facilitate high-speed image capture and processing. Cameras used for PIV tend to be small to maximize mounting flexibility and thus outsource processing duties to the mainframe processor from which the experiment is being run. This avoids the hassle of having n distinct processors and memory cards in each camera and then having to download all those images onto a single machine.

Use of a single computer also allows for real-time visualization of the camera feeds which can be useful for diagnosing problems and debugging.

2.3.2 Calibration

Multi-camera calibration is necessary to identify the relative location and orientation of all cameras within an array which allows object-to-image correspondences to be derived between cameras. This process is carried out with an object of strictly-defined geometry in the location at which the volume of interest is to be imaged. This means that if a volume of interest is inside of a water-filled glass tank then the calibration object needs to be put in that tank at the location which is to be imaged by all of the cameras. If it is a planar calibration object, as seen in figure 2-5, then multiple images from each camera must be taken. The planar grid must be swept through the volume at varying depths and orientations to “fill” the volume with points of known correspondence. The grid spacing is well-defined and used to determine dimensions of objects throughout the volume. This is especially useful for identifying bubble sizes in multi-phase flows or in measuring displacements between images.

Auto-calibration helps speed up the calibration process significantly. The grid in figure 2-5a is used for auto-calibration whereas the grid in figure 2-5b requires manual selection of calibration points in each calibration image. Manual point correspondence is time consuming as every image from every camera (sometimes over 100 images in total) in the calibration process must have an operator select 4-6 points (typically corners or the center of circles) to map points between images. Auto-calibration requires no manual selection and allows for quick evaluation and optimization of the calibration process.

2.3.3 Image acquisition, refocusing and thresholding

Images are acquired of a volume of interest that is supplied with pulses of high-intensity illumination. Most of the content in these images overlap but some does not; typically, cameras are oriented such that their field of view is large enough




Figure 2-5: Two images of planar calibration grids used in SAPIV experiments. Calibration grids are typically made up of a consistent geometrical pattern, whether it be an alternating two-dimensional checkerboard (*a*) or an evenly-spaced grid of circles (*b*). Because the entirety of the grid is visible in (*a*), the pattern can be uniform as any point on the grid can be considered a known origin. The grid in (*b*) has a square and a triangle which can both act as origins.

to image beyond the volume of interest. Example images from a multiphase flow experiment are shown in figure 2-6 and an image from a SAPIV experiment with laser-illuminated tracer particles are shown in figure 2-7.

Generating a set of refocused images allows for “stepping through” the volume of interest by viewing each image in an ordered sequence. Theoretically, only objects located at a specific focal plane will come into focus. Figure 2-8 shows a water ligament that is in focus separate from the rest of the fluid volume. The in-focus object is identified by its higher contrast with the surrounding image. Objects not in focus are noticeably more blurred. Figure 2-9 shows the volume seen in figure 2-7 but refocused at a particular focal plane. As a result, there are far fewer discrete objects visible since the effective depth-of-field is much smaller in refocused images.

For PIV experiments, refocusing image post-processing is typically necessary. The goal of post-processing is to remove false objects known as “ghost particles” to leave only the objects in the image which were actually located at the particular focal plane’s depth. To filter out the background out-of-focus blur, Gaussian thresholding is used to separate noise from the signal. In this case, “noise” is really just out-of-




Figure 2-6: A sampling of four images from a camera array at a specific point in time. The experiment sought to better understand turbulent flow sheet breakup for water in air. Common features such as water droplets and water ligaments are seen in some or all of the images.



Figure 2-7: An example of a raw SAPIV image from a vortex ring quantification experiment. The section of particles being illuminated by the laser is clearly visible.




Figure 2-8: A refocused view of the multiphase flow experiment in figure 2-6. The ligament in the lower-left quadrant of the image is the only identifiable object in focus at this plane. Its high contrast makes it easy to spot against the low-intensity blur that surrounds it.



Figure 2-9: A refocused view of the PIV experiment from figure 2-7. An underlying low-intensity blur makes up most of the image but many particles of high intensity and high contrast are visible.

focus particles and not random noise. However, this blur degrades refocused image quality just the same. Volumes with higher density of tracer particles will have greater noise due to the increased ratio of out-of-focus to in-focus particles. Thresholding is carried out by fitting a Gaussian normal distribution to the normalized intensity distribution as seen in figure 2-10. The fit generates two parameters: a mean and a standard deviation. The intensity cutoff for what are considered “signal” particles is three standard deviations—a heuristic derived from optimization of reconstruction of refocused images in previous experiments [11].

Applying this thresholding to the refocused PIV image in figure 2-9 yields the thresholded image in figure 2-11. The inherent trade-off in thresholding is to try to reduce noise while avoiding the loss of information of objects in the volume. If the thresholding limit is set too high, then real particles which may have been occluded could be lost. However, if the limit is set too low then false particles will be included in the reconstruction. There is no particular limit which ideally separates false and real particles, both will be present at any reasonable thresholding limit set. The critical objective is to maximize the ratio of signal to noise in the volume reconstruction.

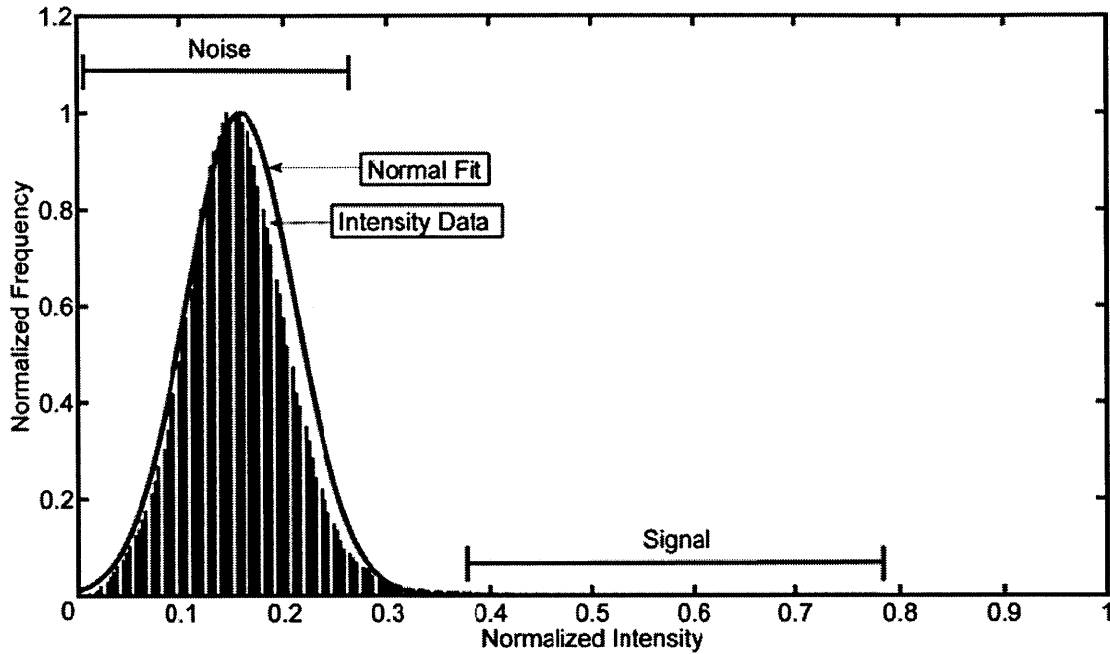


Figure 2-10: A normalized intensity histogram of an individual refocused image. Due to the large number of out-of-focus particles, the distribution of intensity in any particular refocused image converges towards a normal distribution at the low-intensity range of the distribution. These intensity buckets are considered to be “noise”. Particles with high intensity located sufficiently far enough away from this distribution are considered to be statistically significant. These particles make up the “signal” portion of the distribution. A typical thresholding limit is three standard deviations; in this case, that limit corresponds to a normalized intensity of 0.32. Any pixels having a normalized intensity greater than this value are retained in the image, pixels below this value are considered noise and set to zero.



Figure 2-11: A thresholded refocused PIV image. This image represents a best estimate as to which particles are located where at this particular focal plane. A set of these images constitutes a reconstructed volume and can be compared to the images taken of the actual volume—given the known geometry of the volume of interest—to assess reconstruction quality.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Implementation of distortion and spherical aberration in a SAPIV system

This chapter describes models that approximate radial distortion and spherical aberration effects in real optical systems. These models are then implemented in a simulated SAPIV system in order to determine their effect on reconstruction quality.

SAPIV systems require accurate, sub-pixel calibration to reconstruct 3D volumes. Calibration is often based on an idealized, pinhole model that fails to incorporate the error-inducing effects of optical aberrations. In order to quantify the effects of these aberrations in a SAPIV system, we select radial distortion and spherical aberration and try to replicate their detrimental effects via mathematical models. Radial distortion is chosen because it is typically the most common type of distortion in photographic optics [34] and spherical aberration is chosen because it is among the first order of optical aberrations (known as *Seidel aberrations*) and is typically more significant than others [35]. Its symmetrical form also allows for intuitive understanding and relatively simple implementation.

They are modeled in a simulated SAPIV system to measure how each affects the accuracy of volume reconstruction. A simulated system is chosen because it allows for precise application of different amounts of radial distortion and spherical aberra-

tion and it can directly compare the reconstructed volume with the true synthetic volume. The MATLAB computing environment is used because of its image processing capability as well as ease of use in programming. All synthetic images are generated with a pinhole camera model with distortion and/or spherical aberration then applied according to models given in sections 3.1 and 3.2, respectively. These distorted/aberrated images are used for calibration and reconstruction of the imaged volume. Reconstruction accuracy is then evaluated in light of the amount of distortion/aberration applied.

Knowledge about how distortion and spherical aberration affect the imaging capability of a single camera have been well-cataloged in the literature, but how they affect an entire camera *system* is a problem which has yet to be explored. Given that each camera in a typical system has a unique perspective on the volume being imaged, how distortions and aberrations propagate to refocused images and the final reconstructed volume is a complex and non-trivial problem.

3.1 Radial image distortion

Radial distortion is applied to sets of synthetically-generated images. This includes images generated of synthetic calibration targets during calibration which are used to determine the three-to-two dimensional mapping function as well as synthetic images generated of the particle volume which are ultimately used for reconstruction. The distortion should have two degrading effects; first, the calibration will be inaccurate because pixel locations from the calibration targets will not be consistent among the images taken which will give erroneous object-to-image mappings. Second, the images taken of the particle volume will not line up as well as they would without image distortion. In both cases, this mode of imperfect alignment will result because images are distorted in each respective image plane, but the orientation of each image taken of the volume is different. Therefore, the image distortion will affect each image differently and hence will make it seem as if each image taken is of a slightly different particle volume.

Image distortion is a mapping aberration and therefore only degrades the location at which light rays focus without degrading the focus itself. Essentially, this means the distortion does not induce blur but rather warping. No image resolution is lost due to distortion but rather the relative shapes of objects are warped as a result of the process of mapping a spherical wavefront to a planar image sensor [36]. This leads to variable magnification that is the result of imperfect optics, distortion from the fluid medium and any non-optical glass surface between the volume of interest and the camera (such as a tank wall). Specifically, radial image distortion is mostly a product of radially-symmetric optical elements. Images can be distorted in modes other than radial, including tangential distortion, but radial distortion typically dominates and has incited a flurry of research for its correction [14, 15, 16]. A polynomial model for radial distortion correction [14] is given as

$$r_u = r_d(1 + K_1 r_d^2 + K_2 r_d^4) \quad (3.1)$$

where r_u denotes the undistorted image point and r_d denotes the distorted image point and K_1 and K_2 are the second and fourth-order radial distortion coefficients. For this investigation, we drop the fourth-order term as it is typically only required for wide-angle lenses [37].

Radial distortion can be positive (barrel distortion) or negative (pincushion distortion) as seen in figure 3-1. A second-order radial distortion model can also be parameterized by denoting the distortion of the furthest radial points in percentage terms relative to their undistorted location. This has the added benefit of non-dimensionalizing the amount of distortion within an image. Typical “portrait” style lenses, which are often used for PIV applications and have focal lengths between 50-110 mm, will have distortion on the order of 0 - 1% [34].

All radially-distorted images, including calibration images, used in these simulations are distorted according to a certain percent value within a reasonable range for lenses that are typically used for fluid imaging applications. Distortion of individual images was carried out over a range of 0% - 2% percent distortion (where percent




Figure 3-1: Radial image distortion in synthetic images of a rectangular particle volume with a square front face. An undistorted view (*a*) produces an image which accurately depicts a symmetric square. A view distorted with 25% barrel distortion (*b*) shifts pixel intensity values radially inward relative to their undistorted pixel locations. A view distorted with 25% pincushion distortion (*c*) shifts pixel intensity values radially outward relative to their undistorted pixel locations. It should be noted that 25% distortion is exceptionally large and used here simply for demonstration. Distortions in lenses used for PIV will be significantly smaller; distortions of this scale are typically seen only in very short focal length “fisheye” lenses.

distortion corresponds to the movement of pixel values located at the image extrema (corners for rectangular images)). For example, at 2 percent positive image distortion the corner pixel of an image with 1,000 x 1,000 pixels (located at images coordinates $u = 500$, $v = 500$) has a pixel radius ($r^2 = u^2 + v^2$) of 707 pixels, so 2% distortion will correspond to an inward shift of that pixel’s intensity value of about 14 pixels. All other pixels are shifted according to equation 3.1 with K_1 determined from the given percent distortion at the maximum image radius. A sampling of pixel shifts in images at various image radii with various amounts of distortion is shown in table 3.1. A derivation of the second-order radial distortion parameter (K_1) in terms of percent distortion is shown in Appendix A.

3.2 Spherical aberration

Spherical aberration is applied to a set of synthetically-generated images. This includes images generated of synthetic calibration targets during calibration which are

Percent Distortion	Pixel Shift		
	$r_n = 0.1$	$r_n = 0.5$	$r_n = 1.0$
0%	0	0	0
0.1%	0.01	0.18	0.71
0.5%	0.03	0.88	3.54
1.0%	0.07	1.77	7.07
2.0%	0.14	3.54	14.14

Table 3.1: Pixel shift for various distortion values at various normalized image radii for a 1,000 x 1,000 pixel image. The normalized radius r_n is defined as the image radius divided by the maximum image radius for a rectangular image. Even at 0.1% distortion, a distortion amount not visually recognizable, pixel shift is large enough to induce pixel errors that are significant for accurate multi-camera calibration. Distortions on the order of 1-2% can generate very significant shifts towards the periphery of images that can significantly degrade mapping accuracy.

used to determine the object-to-image mapping function as well as synthetic images generated of the particle volume which are ultimately used for reconstruction.

Spherical aberration is known as the first mode of the Seidel aberrations which are the first set of optical distortions not accounted for in paraxial Gaussian optics. Gaussian optics assumes that all light rays are paraxial and are therefore refracted at small angles. This allows for the small-angle approximation ($\sin \theta \approx \theta$) to be made in paraxial ray-tracing analyses. While this linearization greatly simplifies calculation of image formation it fails to adequately model Snell's Law, which governs the refraction of light and is given below

$$n_1 \sin \theta_1 = n_2 \sin \theta_2 \quad (3.2)$$

Under the paraxial approximation, the refraction of light at an air-lens interface on a spherical lens is equivalent to the refraction on an elliptical lens. This is significant because an elliptical lens theoretically focuses light to a singular point—free of spherical aberration. A spherical lens also focuses light to a singular point, but only in the paraxial region. Non-paraxially, light rays diverge as shown in figure 3-2. Lens elements in photographic optics tend to be made with spherical surfaces because they




Figure 3-2: Lens with spherical aberration focusing incident, non-paraxial rays (red) along a range of points. The imperfect focusing of rays causes an image point to spread and reduce its peak intensity. Note that more light rays contact the image plane (black) close to the ideal image point than further away.

are significantly cheaper to manufacture; a spherical surface is easy to grind with a tool that has a constant curvature whereas an elliptical surface requires much greater effort.

Spherical aberration arises from the non-singular focus of off-axis light rays leading to imperfect image formation of points. Spherical aberration induces a blur that is constant over the image field and the magnitude of this blur scales with the fourth power of the pupil size [36]. This means that doubling the diameter of the pupil will increase the amount of aberration in the entire image by a factor of 16. The precise formulation of spherical aberration as a distinct Zernike mode of aberration [38] is given as

$$Z = (RMS)\sqrt{5}(6\rho^4 - 6\rho^2 + 1) \quad (3.3)$$

where ρ is the normalized pupil coordinate and RMS denotes the amount of spherical aberration given in waves and is the root-mean-square deviation of the wavefront from an unaberrated plane wave. This higher-order dependence highlights the strong dependence of spherical aberration as well as other higher-order aberrations on pupil

size. Typical values for RMS vary depending on the optical system but values usually range from 0.01 to 0.10 [35, 39, 40, 41]. This aberration from a plane wavefront gives a point-spread function (PSF) that is wider than one obtained from a diffraction-limited system (where aberrations are negligible) as seen in figure 3-3 and quantified in table 3.3. The point-spread function is a two-dimensional impulse response where the input is a point source of light. In this case, since spherical aberration is uniform over the image field, we can find the PSF for our given aperture and wavefront and convolve it with the unaberrated image produced by the simulation tool, including calibration images, as seen in figure 3-4, and images of the particle volume. For a system with a finite circular aperture, the PSF in image coordinates (u, v) is, [42]

$$PSF(u, v) = \iint_A circ(\rho(x, y)) e^{i2\pi Z(\rho(x, y))} e^{-i2\pi(ux+vy)} dx dy \quad (3.4)$$

where the Fourier transform represents a wavefront transformation with a focusing lens and,

$$circ(\rho(x, y)) = \begin{cases} 1 & \text{if } \rho \leq 1, \\ 0 & \text{if } \rho > 1. \end{cases} \quad (3.5)$$

the "circ" function is commonly used to describe circular apertures in pupil spatial filtering analyses. More information on this is given in Appendix B.

3.3 Simulated SAPIV system

Radial distortion and spherical aberration are applied to synthetic images within a simulated SAPIV system. This system, based on the work of Belden et al. [11], creates a 50 mm x 50 mm x 10 mm volume and populates it with a randomly-generated field of static particles at a rate specified by a particle seeding density parameter (*particles/mm³*). The MATLAB computing environment, along with the image processing toolbox, is used for this exercise. Particles have a diameter of 20 μm . The camera array is a 5 x 5 camera grid evenly spaced in the X and Y dimensions




Figure 3-3: An unaberrated, diffraction-limited PSF (a) with higher concentrated intensity pattern and an aberrated PSF (b) ($RMS = 0.15$) showing intensity spread away from the central peak to the second annular lobe. Assumes a point source of light with unit intensity. There is noticeable spread in the aberration-free PSF; this is a result of diffraction from a finite circular aperture (here, $f/2.8$) and is known as a diffraction-limited PSF.

0.002	0.001	0.003	0.000	0.002
0.001	0.004	0.012	0.005	0.001
0.003	0.012	0.908	0.014	0.003
0.000	0.005	0.014	0.005	0.001
0.002	0.001	0.003	0.000	0.002

Table 3.2: A representation of how the intensity distribution from the diffraction-limited, aberration-free PSF from figure 3-3a would be measured on an image sensor with $10 \mu\text{m}$ pixel pitch. Each table entry above represents a single pixel on the square synthetic sensor. This assumes unit input intensity; the sum of all values equals 1. As indicated by the center entry, 90.8% of all intensity would be measured by the central pixel. Non-symmetry

<i>RMS</i>	Relative Intensity		
	Center Pixel	First Ring	Second Ring
0	0.908	0.070	0.022
0.02	0.903	0.075	0.022
0.05	0.877	0.098	0.026
0.10	0.786	0.175	0.039

Table 3.3: Comparison of PSFs from each of the four *RMS* values used to generate PSF with different amounts of spherical aberration. As the amount of the spherical aberration increases, intensity moves away from the center pixel, where all intensity should be localized, to the first and second rings. The first ring is the eight pixels that surround the center pixel in table 3.2; the second ring is the remaining 16 pixels.




Figure 3-4: A non-aberrated image (*a*) of a simulated calibration target and an aberrated image (*b*) of a simulated calibration target. Note the resolution degradation and loss of high spatial frequency information in the aberrated image. The blurring of pixel intensities is uniform across the image; no part of the image of the calibration target is more blurred than any other.

with a camera spacing dimension of 15.3 cm. All camera centers of projection are coplanar. Cameras are modeled as having square 1,000 x 1,000 pixel sensors with an 85mm focal length and magnification of -0.2. Pixels are 10 μm x 10 μm . Simulated images are project from the synthetic object to synthetic images using a linear pinhole camera model. Four images of the same volume from the corner cameras of the 5 x 5 array can be seen in figure 3-5. Prior to particle seeding, calibration is carried out using pre-defined simulated calibration targets, as seen in figure 3-4 that define the relative position and orientation of each camera in the array as well as how each camera's pixels map to coordinates in the object space. Since depths in the Z dimension are defined discretely at reference planes, particle locations are found by fitting the intensity values of calibration points to a normal distribution. The peak of the distribution then corresponds to the voxel location of the calibration point. Once a calibration is successful in returning the 3D locations of the calibration targets, it can be used to reconstruct synthetic particle fields.

These particle fields are generated, as seen in figure 3-5, and are realigned with the reference camera (central camera) with homographies obtained from the calibration process. After averaging, a synthetic image is formed at a particular focal plane corresponding to a given depth, as seen in figure 3-6. Focal planes are spaced in 0.2 mm increments such that 201 refocused images span the 10 mm depth of the volume. To remove background noise (particles not located at the particular focal plane of interest) thresholding is required in order to generate an accurate reconstruction. Due to the relatively large number of particles, the background noise (when plotted according to pixel intensity) takes on the shape of a normal distribution with a long tail corresponding to in-focus points. The typical threshold cutoff used is three standard deviations above the normally-fitted mean for acceptable reconstruction using a histogram just like that in figure 2-10 [11]. A thresholded image is shown in figure 3-7; these synthetic images makeup the discrete slices that form the reconstructed volume.




Figure 3-5: Synthetic images taken of 50mm x 50mm x 10mm particle volume from four corner cameras of the square 5 x 5 array. Image from top-left camera (*a*); image from top right-camera (*b*); image from bottom-left camera (*c*); image from bottom-right camera (*d*).




Figure 3-6: Refocused synthetic image of a particular depth within the imaged volume. The low-intensity background blur is due to unfocused particles at various other depths in the volume. In-focus particles can still be seen standing out from the background blur.




Figure 3-7: Thresholded version of image 3-6 retaining only pixels of relatively high intensity. Finding the proper threshold is critical for maximizing inclusion of desired, at-depth particles and minimizing inclusion of false particles.

Chapter 4

Results

The results of the simulations described in section 3.3 are derived from the comparison of the “true” synthetically-generated particle volume with the reconstructed volume which was built from the calibrated mapping of two-dimensional synthetic images. This analysis will utilize correlation and rank correlation in sections 4.1 and 4.2, respectively. Further analysis in section 4.3 also describes the performance of the SAPIV system in terms of a signal-to-noise ratio determined from the intensity histogram of refocused images that make up the reconstructed volume. These results attempt to characterize how and to what extent image distortion and spherical aberration affect a SAPIV system.

4.1 Correlation

Correlation is a measure of statistical dependence between any two sets of data. The statistical significance of any correlation between two variables of interest is determined by the size of the samples and the degree to which these samples agree. This can be applied to the “true” and reconstructed particle volume to determine how much they agree. Given that both volumes contain enormous data sets—1,000,000 pixels per refocused image and 201 refocused images gives 201,000,000 data points—the significance of the correlation will be based on how well they agree.

A normalized three-dimensional correlation metric used in PIV experiments [6] is

given as

$$Q = \frac{\sum_{X,Y,Z} V_R(X, Y, Z) \cdot V_S(X, Y, Z)}{\sqrt{\sum_{X,Y,Z} V_R^2(X, Y, Z) \cdot \sum_{X,Y,Z} V_S^2(X, Y, Z)}} \quad (4.1)$$

where Q returns a value between -1 and 1 where -1 represents a perfectly negative correlation, 0 represents no correlation and 1 represents a perfectly positive correlation. V_R is the reconstructed volume being evaluated and V_S is the randomly-generated synthetic particle volume with known particle locations. A value of $Q = 0.75$ or above indicates an adequate reconstruction according to Elsinga et al. [6]. This cutoff value was shown to correspond to a mapping error of 0.45 pixels in a SAPIV system [11].

Results for reconstruction quality as given by the normalized correlation value are presented in figures 4-1, 4-2 and 4-3 as a function of distortion, spherical aberration and both, respectively.

Correlation presents an attractive and convenient way of evaluating reconstruction quality but certain drawbacks exist which preclude it from providing a definitive quantification of the damaging effects of distortion and spherical aberration. Correlation only has informational value when making relative comparisons between data sets, as is done in figures 4-1, 4-2 and 4-3. However, this relative comparison only has value if all of these data sets are being compared to a single ground truth data set. For example, if you want to compare two reconstructions of a particle volume (A & B) with the ground truth C then a comparison would be made between the correlation of (A & C) with that of (B & C). However, if A and B are reconstructions of C and D , respectively, then the comparison of the correlation of (A & C) and (B & D) conflates the change in reconstruction parameters with the change in ground truth parameters.

Furthermore, reconstructed volumes with more particles have more data to sample from. A high-correlation value from data sets with more particles carries more weight than a high-correlation value from a data set with very low data. The amount of data in the figures below varies over an order of magnitude (0.5 to 4 *particles/mm*³) and thus comparing data in the same trace for a given amount of distortion makes




Figure 4-1: Reconstruction quality as provided by correlation as a function of seeding density and distortion D . It shows that significant reduction in reconstruction quality occurs with more than a few tenths of a percent of radial image distortion. At 1% and beyond, none of the simulated cases recorded a reconstruction quality above the $Q = 0.75$ adequacy criterion. Such a degradation is to be expected when dealing with three-to-two dimensional mapping systems that rely heavily on accurate ray-tracing to generate quantitatively accurate three-dimensional reconstruction. Some dependence between particle density and reconstruction quality is evident. Even at zero distortion, a perfect reconstruction is not observed. This is due to the imperfect calibration procedure, finite spacing of discrete reference planes, the probabilistic nature of occlusions and the heuristic of thresholding. Higher reconstruction quality can be obtained with additional computational cost. Voxels have an absolute size of $1/8000 \text{ mm}^3$.

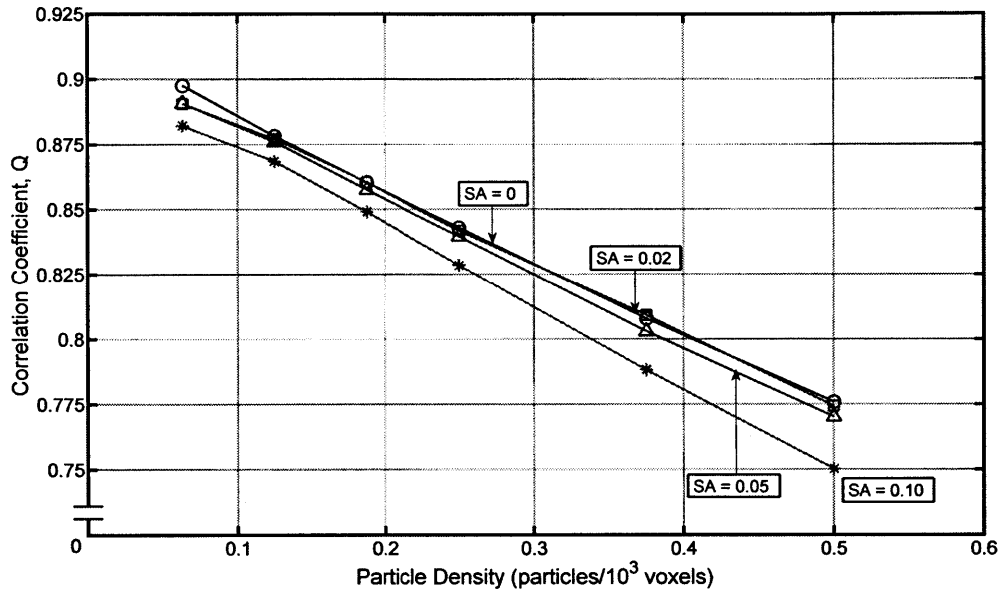


Figure 4-2: Reconstruction quality as provided by correlation as a function of seeding density and spherical aberration SA . It shows that a much smaller reduction in reconstruction quality occurs with the inclusion of spherical aberration effects than with distortion effects. Even at the highest value of spherical aberration, corresponding typically to a simple, uncorrected one or two-lens system, degradation in image quality only resulted in a reduction in Q of no more than 0.04 at the highest particle density. Only in the worst case scenario ($SA = 0.10$) with the highest particle seeding density did the quality come close to the adequacy criterion. Spherical aberration induces no mapping errors, but the spreading of light intensity from a given point over a larger area reduces the intensity of the true point and increases intensity around that point. The results shown here indicate that spreading is not too significant relative to the pixel size of the synthetic sensor.

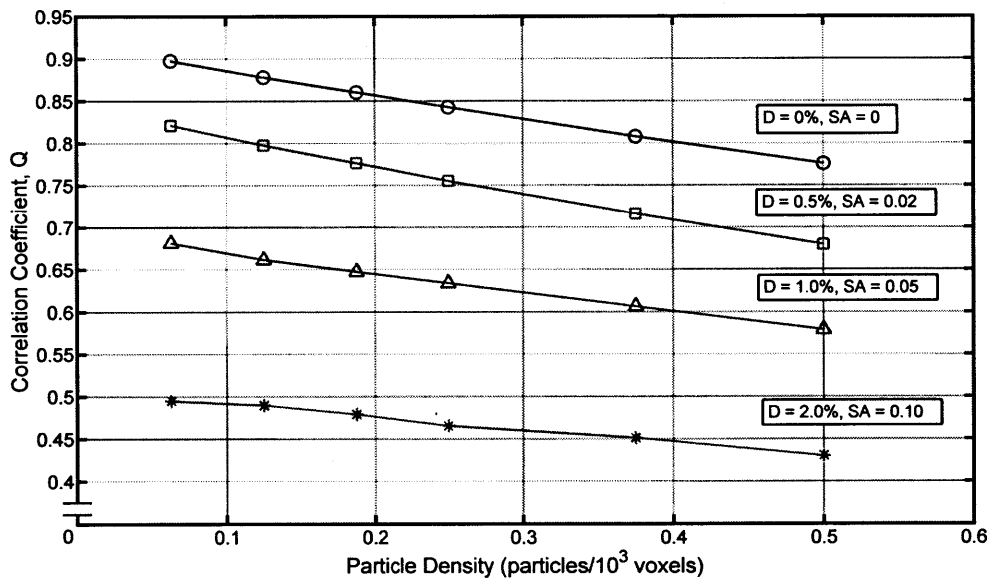


Figure 4-3: Reconstruction quality as provided by correlation as a function of seeding density, distortion D and spherical aberration SA . It shows that the effects of distortion and spherical aberration are mostly independent. Figure 4-3 mirrors figure 4-1 except for a small reduction in Q for all traces very similar to that seen in the difference between the control ($SA = 0.00$) and spherically aberrated simulation sets in figure 4-2. In light of the fact that distortion and spherical aberration are theoretically independent modes of higher-order image aberrations, the results seen here are not surprising.

little sense. For example, imagine an image used in PIV cross-correlation contained just one particle and was correlated with another image at a different time point and yielding a correlation value of $Q = 1$. Now imagine an image with ten distinct particles also yielding a correlation value of $Q = 1$; the statistical significance of the image (data set) with 10 points in perfect correlation is much greater than the image with just one distinct point. The same analysis can be applied for particle volumes of varying particle densities.

Particle size also affects the correlation value of a particular reconstruction. Two volumes filled with small and large particles cross-correlated with themselves at a later point in time will yield different correlations. Small particles will tend to overlap less in a second image than large particles will, thus giving a lower correlation value. Even though this analysis keeps particle size constant, these factors show that a correlation value being used as a metric for reconstruction quality has significant shortcomings in practice. Its use in this analysis may be plausible but in a less contrived, more experimental setting, these same observations may not be as evident.

These issues can be allayed by using rank correlation which is the topic of section 4.2.

4.2 Rank correlation

Standard correlation is used to determine the quality of a reconstruction but the ultimate objective is to accurately determine particle location and velocity. In that sense, achieving the highest correlation value does not correspond to reconstructing the most accurate velocity field.

Rank correlation is an alternative method of comparing data sets where the relative magnitude of data points is still important but only in an ordinal sense. Rank correlation has been shown to successfully indicate the statistical significance of PIV images [43]. Rank correlation takes all the elements of a data set and assigns them a number from 1 to N where N represents the number of elements within a data set. In this case, each element is a voxel in a three-dimensional volume with some intensity

value. If each voxel in the volume $V(x, y, z, I)$ is ranked according to its intensity value I in relation to all voxels within this volume then this voxel will be assigned an ordinal rank value J . This voxel now contains the rank value instead of the intensity value as does the rest of the volume $G(x, y, z, J)$. The rank number is assigned by order of increasing magnitude; thus when correlation is carried out it is not the voxel intensity values that are correlated but rather the rank of those intensities. Voxels with the same intensity value receive an average of the rank they would have received had their values been different. This gives us a rank correlation metric

$$Q' = \frac{\sum_{X,Y,Z} G_R(X, Y, Z) \cdot G_S(X, Y, Z)}{\sqrt{\sum_{X,Y,Z} G_R^2(X, Y, Z) \cdot \sum_{X,Y,Z} G_S^2(X, Y, Z)}} \quad (4.2)$$

where G_R is the ranked reconstructed volume being evaluated and G_S is the ranked randomly-generated synthetic particle volume with known particle locations.

Figures 4-4 and 4-5 show that rank correlation drastically increases the ability to match elements between the synthetic and reconstructed volumes.

While rank correlation provides an effective way of correlating PIV images and volumes, it doesn't help us better understand how image distortion and spherical aberration affect reconstruction quality of three-dimensional volumes. What is needed is a method of evaluation which is relevant to the method in which reconstruction is carried out. Since the process of averaging shifted images generates distinct, high-contrast particles (signal) and out-of-focus, low-contrast particles (noise), it would be sensible to evaluate reconstruction according to a signal-to-noise metric. Evaluation of such a metric is carried out in section 4.3

4.3 Histogram analysis

The quality of volume reconstruction is evaluated according the intensity histogram data of the reconstructed volume. As seen in figure 2-10, an intensity histogram plots the frequency of discrete ranges of voxel intensity that appear within a volume.

The normal fit of the intensity distribution yields two fitted parameters - a mean

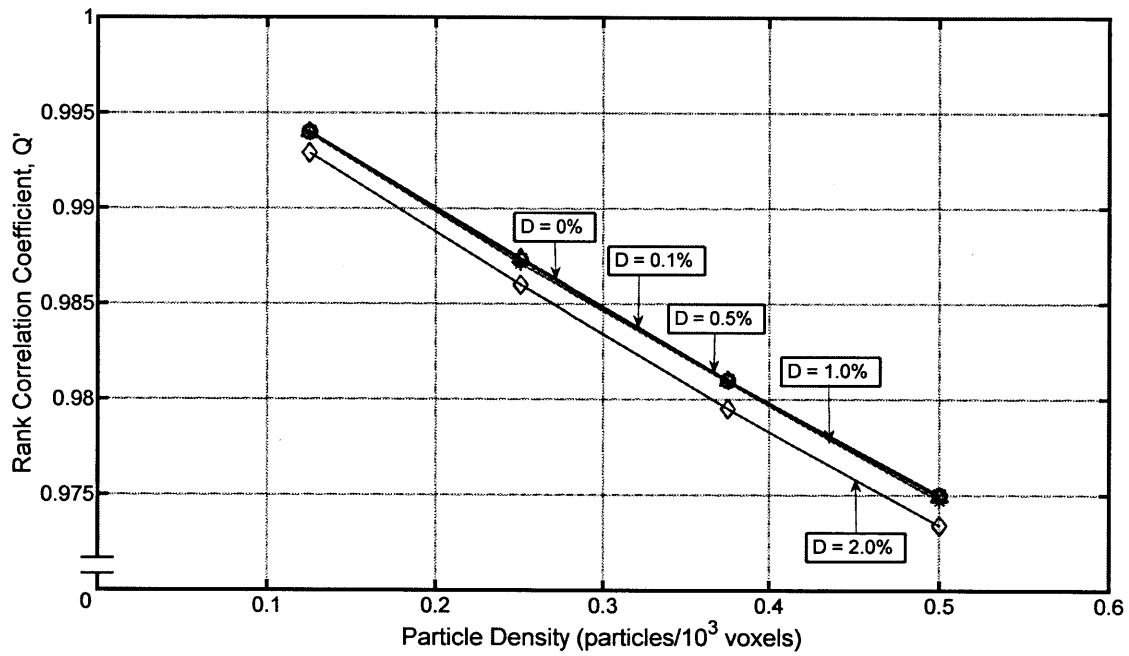


Figure 4-4: Reconstruction quality as provided by rank correlation as a function of seeding density and distortion D . Results show that distortion has only a small effect at 2.0% and a negligible effect for distortions below that. The range of correlation values is greater than 0.97, which is much higher than the values given for standard correlation.




Figure 4-5: Reconstruction quality as provided by rank correlation as a function of seeding density and spherical aberration SA . Results show that spherical aberration has no discernible effect. The range of correlation values is greater than 0.97, which is much higher than the values given for standard correlation. The lack of correlation degradation can be attributed to the nature of spherical aberration. A uniform “spreading” of intensity would affect a standard correlation as voxel intensities are changing, but since this spreading is uniform, ranked values will not see significant deviation.

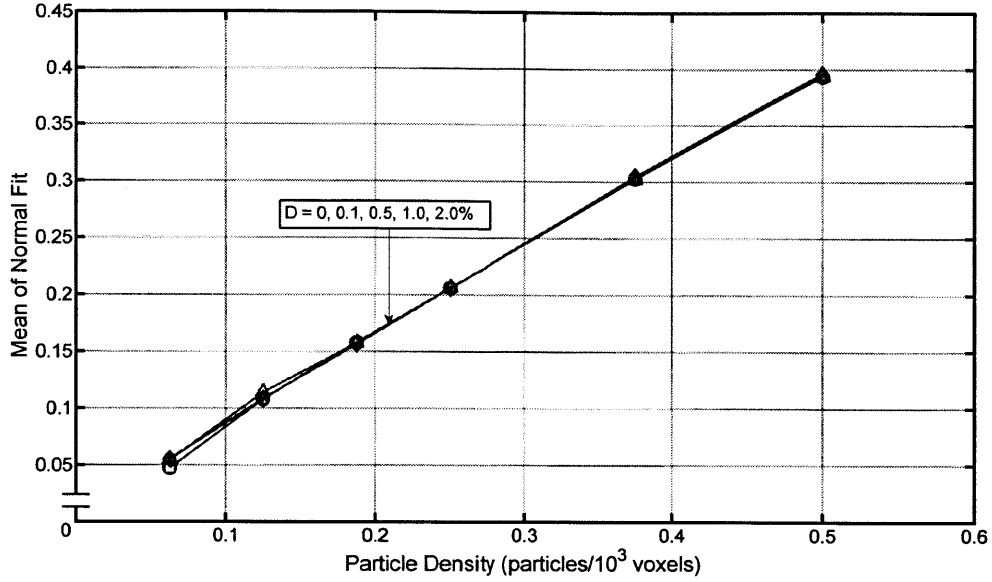


Figure 4-6: Mean of the normal fit of the intensity distribution as a function of seeding density and distortion D . No significant effect is evident; this agrees with the theoretical prediction that the mean intensity of distorted images should not vary. Distortion does not add or remove intensity from averaged, shifted images, it simply displaces intensity values.

and standard deviation. The mean intensity does not vary with different amounts of distortion and spherical aberration, as seen in figures 4-6 and 4-7 but the standard deviation does, as seen in figures 4-8 and 4-9. Given the normal fit parameters, the informational content of the reconstructed volumes can be compared quantitatively. The sets of refocused images that make up each reconstruction of the imaged volume are evaluated according to their signal-to-noise ratio,

$$SNR(V_R) = \frac{1}{N} \sum_j^N \frac{\max(I_j)}{\mu(I_j) + \sigma(I_j)} \quad (4.3)$$

Where V_R is the reconstructed volume with j reconstructed focal planes, I_j is a reconstructed image, $\mu(I_j)$ is the mean of the normal fit of the intensity histogram and $\sigma(I_j)$ is the standard deviation of the normal fit.

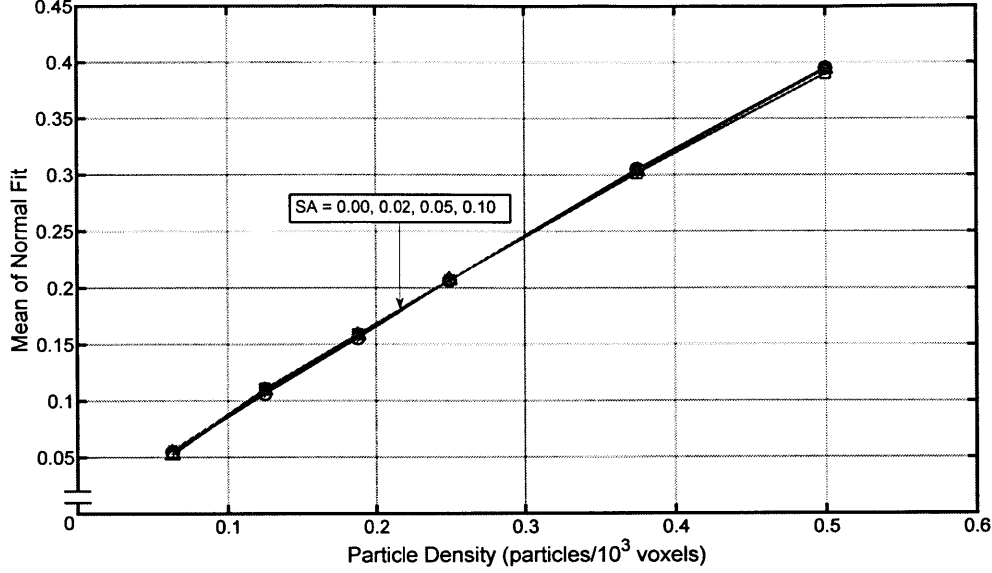


Figure 4-7: Mean of the normal fit of the intensity distribution as a function of seeding density and spherical aberration SA . No significant effect is evident; this agrees with the theoretical prediction that the mean intensity of aberrated images should not vary. Spherical aberration does not add or remove intensity from images, it simply displaces intensity values.

The inclusion of the standard deviation of the normal fit adjusts for the spread of the distribution. A distribution with more spread will mark a greater intensity range as noise. The calculation of SNR from each reconstructed image within a volume is meant to dampen the statistical variation in the measurement of maximum intensity. Figures 4-10 and 4-11 show the normalized SNR as a function of distortion/spherical aberration, respectively, and particle density. Normalized SNR ($nSNR$) is given as

$$nSNR(V_R(D, \gamma), V_R(D = 0, \gamma)) = \frac{SNR(V_R(D, \gamma))}{SNR(V_R(D = 0, \gamma))} \quad (4.4)$$

where D is the amount of distortion induced in the images used to reconstruct the volume and γ is the particle density.

This metric provides an effective comparison between volumes with different levels of distortion and particle density. Normalization weeds out the effects of particle density so that comparison of the signal-to-noise ratio can be made for different levels of distortion and particle density.

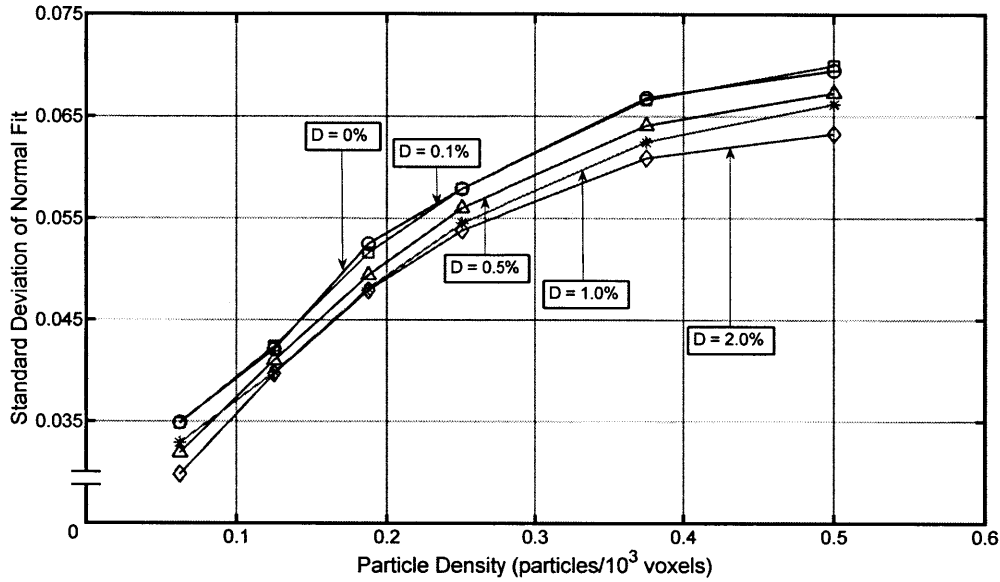


Figure 4-8: Standard deviation of the normal fit of intensity data as a function of distortion D and particle density. The standard deviation decreases for distortion percentages 0.5% and above and also increases with higher particle density as more particles create much more non-uniform background noise. At low particle densities, there is less out-of-focus particle overlap, so the distribution of noise intensities is tighter. The tightening of the normal distribution as distortion increases corresponds to the greater spreading of intensity in reconstructed images due to misaligned images. Instead of a distinct, high-intensity voxel surrounded by very low or zero-intensity voxels, distortion reduces resolution of reconstructed particles which leads to many more non-zero, low-intensity voxels. In the limiting case where distortion is very large, then the intensity of all voxels will converge to a single value - the mean intensity - and all information will be lost to background noise.




Figure 4-9: Standard deviation of the normal fit of intensity data as a function of spherical aberration SA and particle density. The standard deviation decreases for SA values 0.02 and above and also increases with higher particle density as more particles create much more non-uniform background noise. As more SA is input to the system, the standard deviation decreases. This tightening of the normal distribution corresponds to the greater spreading of intensity in reconstructed images due to the improper focusing from spherical aberration. Instead of a distinct, high-intensity voxel surrounded by very low or zero-intensity voxels, spherical aberration reduces resolution of reconstructed particles which leads to many more low-intensity voxels and thus lower variance and standard deviation.




Figure 4-10: Normalized signal-to-noise ($nSNR$) ratio as a function of distortion D and particle density. All SNR values are normalized with respect to the zero-distortion ($D = 0\%$) case. Normalization removes the strong dependence of SNR on particle density and allows for relative comparison. $nSNR$ decreases slightly for higher values of distortion. However, a clear trend between $nSNR$ and particle density is not evident thus showing that normalization is at least somewhat effective in helping to make useful comparisons. $nSNR$ degradation is observed even for distortion levels as low as 0.1%.




Figure 4-11: Normalized signal-to-noise ($nSNR$) ratio as a function of spherical aberration SA and particle density. All SNR values are normalized with respect to the zero-aberration ($SA = 0$) case. Normalization removes the strong dependence of SNR on particle density and allows for relative comparison. $nSNR$ decreases for higher values of distortion. $nSNR$ degradation is observed even for distortion levels as low as 0.02.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Conclusion

This thesis has shown the detrimental effects that image distortion and spherical aberration can have on an uncompensated optical system used for synthetic aperture applications. Chapter 1 introduced particle image velocimetry as a useful and practical measurement tool that relies on precise optical equipment for proper function. Chapter 2 explained the fundamentals of the synthetic aperture method as applied to making PIV measurement and the sub-pixel accuracy which is required from the optical imaging elements in the system. Chapter 3 explained two dominant modes of imaging aberration: image distortion and spherical aberration as well as mathematical models that explain their deleterious effects. These models were implemented in a simulated SAPIV environment in order to observe and measure the extent of their influence on three-dimensional volume reconstruction. Chapter 4 described three sets of metrics for the evaluation of quality of reconstructed three-dimensional PIV volumes: correlation and rank correlation of the reconstructed and “true” volumes as well as histogram analysis that culminated in the use of a normalized signal-to-noise ratio to weed out effects of particle density. Both distortion and spherical aberration were shown to have significant effects on correlation values as well as the standard deviation of the histogram normal fit and the normalized SNR. Only image distortion showed an effect with the rank correlation metric.

Means to correct for radial as well as tangential and decentering distortion have existed long enough to be incorporated into most 3D imaging systems, but only for

individual cameras. Nevertheless, even a relatively small amount of uncompensated radial distortion can have significant consequences for accurate mapping of an object in space. Spherical aberration is just one of several Seidel aberrations and here is shown to have a non-negligible effect on object reconstruction.

Future work will need to model additional, more complex forms of aberrations and distortions over a wider range of optical setups—including wide-angle lenses where distortions are much more pronounced and more difficult to model—and imaging conditions. Also, inclusion of typical PIV experimental apparatuses such as glass tank walls as well as fluid mediums with non-unity indices of refraction will illuminate different modes of image distortion. Further analyses that include all Seidel aberrations as well as higher-order aberrations to fully characterize the departure of real optical systems from ideal imaging conditions are the next logical step. Experimental work that validates these analyses will be necessary to motivate future work on the subject of imaging aberrations in three-dimensional systems. Furthermore, physical testing of lenses will be required to enhance the precision of future synthetic aperture system simulations.

Focus will need to be put on image distortion, since it has a much greater capability for thwarting high-accuracy calibrations. Instead of correcting out distortion within a single camera, generating a system-wide, effective point-spread function for aberrated optical systems could lead to more effective calibration and understanding of these multi-camera systems. Furthermore, analysis of distortions and aberrations in other three-dimensional PIV systems could help shed light on the relative robustness of these methods to imaging imperfections.

Bibliography

- [1] M. Raffel, C.E. Willert, and J. Kompenhans. *Particle image velocimetry: a practical guide*. Springer Verlag, 1998.
- [2] G. Comte-Bellot. Hot-wire anemometry. *Annual Review of Fluid Mechanics*, 8(1):209–231, 1976.
- [3] JB Abbiss, TW Chubb, and ER Pike. Laser doppler anemometry. *Optics & Laser Technology*, 6(6):249–261, 1974.
- [4] C. Willert, C. Hassa, G. Stockhausen, M. Jarius, M. Voges, and J. Klinner. Combined piv and dgv applied to a pressurized gas turbine combustion facility. *Measurement Science and Technology*, 17(7):1670, 2006.
- [5] X. Tian, J. Iriarte-Diaz, K. Middleton, R. Galvao, E. Israeli, A. Roemer, A. Sullivan, A. Song, S. Swartz, and K. Breuer. Direct measurements of the kinematics and dynamics of bat flight. *Bioinspiration & Biomimetics*, 1(4):S10, 2006.
- [6] GE Elsinga, F. Scarano, B. Wieneke, and BW Van Oudheusden. Tomographic particle image velocimetry. *Experiments in Fluids*, 41(6):933–947, 2006.
- [7] J. Sheng, E. Malkiel, and J. Katz. Digital holographic microscope for measuring three-dimensional particle distributions and motions. *Applied optics*, 45(16):3893–3901, 2006.
- [8] J. Sheng, E. Malkiel, and J. Katz. Using digital holographic microscopy for simultaneous measurements of 3d near wall velocity and wall shear stress in a turbulent boundary layer. *Experiments in fluids*, 45(6):1023–1035, 2008.
- [9] F. Pereira and M. Gharib. Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows. *Measurement Science and Technology*, 13(5):683, 2002.
- [10] HG Maas, A. Gruen, and D. Papantoniou. Particle tracking velocimetry in three-dimensional flows. *Experiments in Fluids*, 15(2):133–146, 1993.
- [11] J. Belden, T.T. Truscott, M.C. Axiak, and A.H. Techet. Three-dimensional synthetic aperture particle image velocimetry. *Measurement Science and Technology*, 21:125403, 2010.

- [12] C. Skupsch and C. Brücker. Single camera, multiple-plane piv by synthetic refocusing. 2012.
- [13] D.B. Gennery. Generalized camera calibration including fish-eye lenses. *International Journal of Computer Vision*, 68(3):239–266, 2006.
- [14] G.P. Stein. Lens distortion calibration using point correspondences. In *Computer Vision and Pattern Recognition, 1997. Proceedings., 1997 IEEE Computer Society Conference on*, pages 602–608. IEEE, 1997.
- [15] R. Hartley and S.B. Kang. Parameter-free radial distortion correction with center of distortion estimation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(8):1309–1321, 2007.
- [16] R. Cucchiara, C. Grana, A. Prati, and R. Vezzani. A hough transform-based method for radial lens distortion correction. In *Image Analysis and Processing, 2003. Proceedings. 12th International Conference on*, pages 182–187. IEEE, 2003.
- [17] B. Wieneke. Volume self-calibration for 3d particle image velocimetry. *Experiments in fluids*, 45(4):549–556, 2008.
- [18] PAG Scheuer. The development of aperture synthesis at cambridge. *The Early Years of Radio Astronomy-Reflections Fifty Years after Jansky's Discovery*, 1:249, 1984.
- [19] M. Ryle. The new cambridge radio telescope. *Nature*, 194:517–518, 1962.
- [20] R. Levanda and A. Leshem. Synthetic aperture radio telescopes. *Signal Processing Magazine, IEEE*, 27(1):14–29, 2010.
- [21] E.H. Adelson and J.R. Bergen. The plenoptic function and the elements of early vision. *Computational models of visual processing*, 1:3–20, 1991.
- [22] E.H. Adelson and J.Y.A. Wang. Single lens stereo with a plenoptic camera. *IEEE transactions on pattern analysis and machine intelligence*, 14(2):99–106, 1992.
- [23] L. McMillan and G. Bishop. Plenoptic modeling: An image-based rendering system. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 39–46. ACM, 1995.
- [24] M. Levoy and P. Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42. ACM, 1996.
- [25] Z. Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334, 2000.

- [26] A. Isaksen, L. McMillan, and S.J. Gortler. Dynamically reparameterized light fields. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 297–306. ACM Press/Addison-Wesley Publishing Co., 2000.
- [27] V. Vaish, B. Wilburn, N. Joshi, and M. Levoy. Using plane+ parallax for calibrating dense camera arrays. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 1, pages I–2. IEEE, 2004.
- [28] V. Vaish, G. Garg, E.V. Talvala, E. Antunez, B. Wilburn, M. Horowitz, and M. Levoy. Synthetic aperture focusing using a shear-warp factorization of the viewing transform. In *Computer Vision and Pattern Recognition-Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on*, pages 129–129. IEEE, 2005.
- [29] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*, volume 2. Cambridge Univ Press, 2000.
- [30] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan, and J. Tumblin. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Transactions on Graphics*, 26(3):69, 2007.
- [31] Jesse Belden. *Synthetic aperture imaging for three dimensional resolution of fluid flows*. PhD in mechanical engineering, Massachusetts Institute of Technology, Department of Mechanical Engineering, 2011. <http://hdl.handle.net/1721.1/67577>.
- [32] B. Wilburn, N. Joshi, V. Vaish, E.V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy. High performance imaging using large camera arrays. *ACM Transactions on Graphics*, 24(3):765–776, 2005.
- [33] J. Belden, S. Ravela, T.T. Truscott, and A.H. Techet. Three-dimensional synthetic aperture imaging and resolution of multi-phase flows. ASME, 2011.
- [34] Thom Hogan. Lens reviews, 2009.
- [35] T.A. Brunner. Impact of lens aberrations on optical lithography. *IBM Journal of Research and Development*, 41(1.2):57–67, 1997.
- [36] W.T. Welford. *Aberrations of optical systems*. Taylor & Francis, 1986.
- [37] G. Vass and T. Perlaki. Applying and removing lens distortion in post production. In *Proceedings of the 2nd Hungarian Conference on Computer Graphics and Geometry*, pages 9–16. Citeseer, 2003.
- [38] R.J. Noll. Zernike polynomials and atmospheric turbulence. *J. Opt. Soc. Am*, 66(3):207–211, 1976.

- [39] R. Kingslake and R.B. Johnson. Lens design fundamentals, (spie press book). 2010.
- [40] C.H. Lin, H. Chang, and H.Y. Lin. Application of pupil spherical aberration in tuning zoom lens design. *Journal of Optics*, 14(6):065401, 2012.
- [41] N.R. Farrar, A. Smith, D. Busath, and D. Taitano. In-situ measurement of lens aberrations. In *Proc. SPIE*, volume 4000, pages 18–29, 2000.
- [42] J.W. Goodman and S.C. Gustafson. Introduction to fourier optics. *Optical Engineering*, 35(5):1513–1513, 1996.
- [43] D.P. Hart. High-speed piv analysis using compressed image correlation. *Transactions-American Society of Mechanical Engineers Journal of Fluids Engineering*, 120:463–470, 1998.

Appendix A

Distortion parameter calculation

Typically, calculation of radial image distortion seeks to determine distortion coefficients in order to find *undistorted* image points using equation 3.1:

$$r_u = r_d(1 + K_1r_d^2 + K_2r_d^4) \quad (\text{A.1})$$

However, for the purposes of this analysis the converse problem is trying to be solved: given a prescribed amount of distortion, to what image coordinates will the distorted pixels shift? This can be visualized in figure A-1. Using the same equation form as in equation A.1 we can find r_d from r_u given K_1

$$r_d = r_u(1 + K_1r_u^2 + K_2r_u^4) \quad (\text{A.2})$$

If we neglect the fourth-order term we are left with

$$r_d = r_u(1 + K_1r_u^2) \quad (\text{A.3})$$

Now r_d can be solved for all pixels within a two-dimensional image if we have the image data from the undistorted image (r_u) and the distortion parameter. However, it is more convenient and intuitive to work with percent distortion for second-order image distortion. Percent distortion is a dimensionless parameter that does not vary with sensor size. A photographic lens induces a certain amount of radial distortion




Figure A-1: A schematic of how pixel intensity value shift according to the amount of distortion applied to an image. Positive distortion represents a radially-inward shift while negative distortion represents a radially-outward shift.

and using different camera sensors will change K_1 , but not percent distortion. We can define percent distortion as a parameter relating the pixel in the undistorted image at the maximum radius with its distorted counterpart

$$r_d = r_{u,max} \left(1 - \frac{p}{100}\right) \quad (\text{A.4})$$

where $r_{u,max}$ is the maximum image radius (located at the corner pixel of a rectangular image) and p is the percent distortion parameter. Positive values of p correspond to barrel radial distortion and negative values to pincushion radial distortion. The schematic in figure A-1 shows this. For example, a 1,000 x 1,000 pixel image has a corner pixel located at coordinates ($u = 500, v = 500$) with an image radius of 707. The number of pixels that the value in this pixels displaces defines the percent distortion. Substituting equation A.4 into equation A.3 gives

$$r_{u,max} \left(1 - \frac{p}{100}\right) = r_{u,max} (1 + K_1 r_{u,max}^2) \quad (\text{A.5})$$

$$1 - \frac{p}{100} = 1 + K_1 r_{u,max}^2 \quad (\text{A.6})$$

$$\frac{-p}{100} = K_1 r_{u,max}^2 \quad (\text{A.7})$$

$$K_1 = -\frac{p}{100r_{u,max}^2} \quad (\text{A.8})$$

Now that K_1 can be specified using given input information, distorted pixel coordinates can be determined.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

Point spread function derivation

The point spread function (PSF) given in equation 3.4 is derived using diffraction theory. Diffraction can be defined in two regimes: near-field Fresnel diffraction and far-field Fraunhofer diffraction. The Fresnel number

$$F = \frac{D^2}{z\lambda} \quad (\text{B.1})$$

determines the regime: $F > 1$ corresponds to Fresnel diffraction and $F \ll 1$ corresponds to Fraunhofer diffraction. Fraunhofer diffraction is a subset of Fresnel diffraction where the diffraction pattern, which is the intensity pattern produced when light is scattered by an object, is located a relatively large distance away from the object/aperture. Fraunhofer diffraction is determined from the electric field located a great distance z from the aperture

$$E(x', y') = \frac{e^{ikz}}{i\lambda z} e^{ik\frac{x'^2+y'^2}{2z}} \iint_A t(x, y) E(x, y) e^{\frac{-ik}{z}(x'x+y'y)} dx dy \quad (\text{B.2})$$

where $E(x', y')$ is the diffracted electric field, k is the wavenumber, λ is the wavelength, $t(x, y)$ is the aperture transmission function, x and y are the dimensions of the transmission plane and x' and y' are the spatial dimensions along which the diffraction pattern forms. However, the field that is observed on any image plane is an intensity field and intensity is determined by the amplitude of this function, so

equation B.2 can be simplified to

$$E(x', y') = \iint_A t(x, y)E(x, y)e^{-i(k_x x + k_y y)} dx dy \quad (\text{B.3})$$

where k_x and k_y are the spatial frequencies in the x and y direction. This gives us a spatial Fourier transform of the transmitted electric field

$$E(x', y') = F\{t(x, y)E(x, y)\} \quad (\text{B.4})$$

This shows that Fraunhofer diffraction produces a spatial Fourier transform of an incoming wavefront over a relatively long distance of empty space, but that distance is on the order of tens of meters for typical aperture sizes used in cameras with visible light. This would be highly impractical were it not for the properties of glass lenses. Lenses can serve to bring the properties of far-field Fraunhofer diffraction into the near-field. This is because when an object is placed one focal length in front of a lens an image of its spatial Fourier transform is produced at the image plane on the other side of the lens. This allows the same result that is given in equation B.4 to be obtained over a substantially reduced distance and also allows the use of Fourier optics for the analytical calculation of the PSF and the characterization of optical aberrations.

Equation B.4 can be applied for a typical optical camera system that focuses, collimates and filters light onto an image plane as seen in figure B-1. A typical optical system will have a circular aperture that defines the transmission function

$$t(x, y) = \text{circ}(\rho(x, y)) \quad (\text{B.5})$$

and the field being imaged onto the image sensor is not an electric field $E(x, y)$ but rather an intensity field $I(x, y)$ since the incoming wavefront are spatially incoherent. However, the relationship between the electric and intensity field is easily defined where the intensity field is the amplitude of the electric field and does not change the

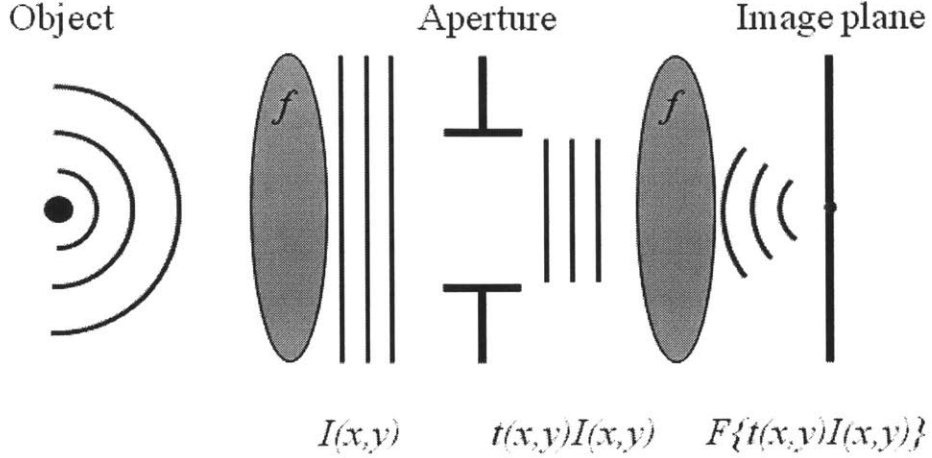


Figure B-1: The front half of the optical system collimates the incoming wavefronts scattered off the object into plane waves (ideally). This is where the wavefront takes on the aberrated form seen in equation 3.3 since it can deviate from a strict plane wave. The aperture filters out parts of the wavefront and in doing so induces diffraction effects that will materialize in a diffraction-limited PSF. Finally, the wavefront is focused down to image point by a lens which is analytically represented by a Fourier transform.

analytical determination of the point spread function

$$I(x, y) = E(x, y) \cdot E'(x, y) = e^{i2\pi Z(\rho(x,y))} \quad (\text{B.6})$$

where $E'(x, y)$ is the complex conjugate of the electric field, $E(x, y)$. Now for an system with spherical aberration, equations B.5, 3.3 and B.6 are input into equation B.3 to yield

$$PSF(u, v) = \iint_A circ(\rho(x, y)) e^{i2\pi Z(\rho(x,y))} e^{-i2\pi(ux+vy)} dx dy \quad (\text{B.7})$$

where $e^{i2\pi Z(\rho(x,y))}$ describes a wavefront of light in the aperture region of the optical system. However, if the system is unaberrated then Z approaches zero and the equation reduces to

$$PSF(u, v) = \iint_A circ(\rho(x, y)) e^{-i2\pi(ux+vy)} dx dy \quad (\text{B.8})$$

which is simply the Fourier transform of the transmission function and is equivalent to equation B.4. Equation B.7 can represent any number of aberrations with any arbitrary set of relative magnitudes; we use it here simply to describe spherical aberration. Any empirically-observed set of aberrations can be broken down into Zernike components [38] and described analytically. As a result, a PSF can be generated analytically as long as the aberrations are measured and grouped.