SINGLE CAMERA, THREE-DIMENSIONAL PARTICLE TRACKING VELOCIMETRY

by

Kevin Howard Peterson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in The University of Michigan
2012

Doctoral Committee:

Professor Volker Sick, Chair
Professor Steven L. Ceccio
Professor Theodore B. Norris
Research Scientist David L. Reuss
Michael C. Drake, General Motors Company

For Carole Jean

ACKNOWLEDGMENTS

I would like to thank Volker Sick for guiding my research over the last several years. His advice has been invaluable, and working with him has truly been a pleasure. Despite having many jobs in the past, I have not found anyone for whom I would rather work.

I appreciate my committee members giving feedback and sharing ideas regarding my project. I am very grateful that these experts in their fields would take the time to help me along my educational path.

Discussions with Dr. Mark Supanich provided the early inspiration for this project, and for that I am thankful.

Much of this project was performed in close collaboration with Dr. Boris Regaard from the Fraunhofer Center for Laser Technology, and I would like to thank him for his contributions to this work, and also for his friendship. The advice and assistance from Ben Lewis and Dr. Stefan Heinemann, also from the Franuhofer CLT, are much appreciated.

I am grateful to Todd Fansler, Mike Drake, Jerry Silvas, Scott Parrish, Cherian Idicheria, and Bill Tisler from General Motors for all the knowledge they have shared. In the summers I spent working at the General Motors Technical Center, I not only vastly increased my knowledge of optical diagnostics, but also saw the excitement that could come from working in a cutting-edge laboratory. I look forward to joining these gentlemen on a full-time basis.

I would like to thank my friends in the Autolab for all the hours we have spent together, especially Ali Alharbi, Brian Peterson, Mike Cundy, Mike Mosburger, Preeti Abraham, Louise Lu, Omar Almagri, and Philipp Schiffmann, with whom I have both suffered and celebrated. Working with you all has been wonderful. I also thank Philipp Sick and Kelsey Mork for their help in performing the simple flow experiments.

I am indebted to my friends in China, especially Chen Hao, Zhang Ming, and Zeng Wei, for their wonderful hospitality during my time studying there. I would have been lost without them. Professor Xu Min was instrumental in making my visit possible, and I am grateful for that.

My family and friends have supported me through all the ups and downs of my research, and for that I am extremely appreciative. Whether through phone-calls, walks in the woods, dinners, or a night on the town, I appreciate all their moral support, without which I would not have completed this program. Although I did not accept my grandmother's offer to do some of my homework for me, I appreciate the offer nonetheless.

I would especially like to thank my fiancée Jane for staying by me through it all. Through late nights at the lab, summers apart, and the stress and strain of homework, papers, and research, she has been there. In fact, she has treated me so well that I am going to marry her.

This work could not have happened without financial support from several sources. My early graduate studies were funded through a graduate student instructor position in the Department of Mechanical Engineering at the University of Michigan. I have received support from General Motors in the form of four student internships at the General Motors Technical Center. The University of Michigan Alternative Energy Technology for Transportation Program provided funding for the project developed in this thesis. I am extremely grateful that all of these sources invested in me, and am

honored to have been given the opportunities I have been given, which would not have been possible without this funding.

I could not have done this work without many generous gifts of time, knowledge, and financial support, for which I am extremely grateful. I feel humbled to have received the chance to perform this work, and truly appreciate the friendships I have made along the way.

TABLE OF CONTENTS

LIST OF FIGURES

ABSTRACT

The efficiency, power output, and pollutant formation of an internal combustion engine are strongly influenced by the air flow inside the engine's cylinder. Single-camera, three-dimensional particle tracking velocimetry (SC3D-PTV) has been developed to better study and understand 3D engine air flows.

SC3D-PTV is performed using a compact, sealed optical element in conjunction with a commercially available image sensor. SC3D-PTV requires only one point of optical access for imaging and one point of optical access for illumination, making this technique ideal for performing 3D flow measurements in an engine with minimal optical access.

SC3D-PTV allows the simultaneous capture of the 3D position of numerous tracer particles suspended in a non-stationary fluid, in one exposure, using a single camera. The 3D velocity profile of the particle field is then obtained by computing the displacement of each particle between consecutive exposures.

A novel PTV algorithm relying on the similarity of the particle images corresponding to a single, physical particle produces 3-component, volumetric velocity fields without the reconstruction of an instantaneous 3D particle field.

This thesis presents the development of both the physical optical element and the analysis algorithm. Proof-of-principle results from a simple flow are given, along with experimental results from an optically-accessible combustion engine.

CHAPTER 1

MOTIVATION

The aim of this thesis is to develop single-camera, three-dimensional particle-tracking velocimetry (SC3D-PTV), a method for measuring 3D air flow inside an optically-accessible internal combustion engine with many advantages over existing techniques. However, the motivation for studying 3D flow inside an engine is much more general, namely improving energy use within the transportation sector by aiding in the development of more efficient internal combustion engines.

Each year, tens of millions of automobiles are sold worldwide, the vast majority of which are powered by internal combustion engines utilizing hydrocarbon-based fuels. The chemical energy contained within the fuel is converted to kinetic energy of the vehicle by means of the engine, but the process of energy conversion is not 100% efficient, thus some fraction of the energy available in the fuel is simply wasted. A small improvement in the energy conversion efficiency of an engine leads to large energy savings for the transportation sector as a whole, since that small improvement is applied to tens of millions of vehicles.

Increased knowledge of air flow can aid in the development of many types of advanced engines, and increase the understanding of many aspects of the engine cycle. One example is the spray-guided, spark-ignition, direction-injection (SG-SIDI) engine,

which offers greater fuel economy than the port fuel-injection or homogeneous-charge, direct-injection engines commonly found in present-day vehicles, but suffers from occasional poor burns (Drake and Haworth, 2007). These poor burns increase the output of unburned hydrocarbons (a regulated emission) and prevent the SG-SIDI engine from realizing its full potential in terms of fuel economy. This project aims to increase the understanding of in-cylinder air flow, which plays a major role in many of the physical and chemical processes that occur during an engine cycle. To motivate this thesis, the following sections describe the means for increasing fuel efficiency with an SG-SIDI engine, the importance of air flow during a SG-SIDI engine cycle, and the shortcomings of current measurement techniques.

1.1 SG-SIDI Engines

The primary method by which SG-SIDI engines increase fuel efficiency is by using a stratified charge approach when the engine is operating at partial load (Zhao et al., 1997). Since gasoline can only burn when mixed with air in the proper proportions, homogeneous charge gasoline engines must fill the entire combustion chamber with air and fuel in the proper proportions to obtain an ignitable mixture. Homogeneous charge engines control the engine load by limiting the amount of air entering the engine using a throttle plate. By restricting the air entering the engine, less gasoline is required to attain an ignitable mixture and engine load is reduced. However the engine must perform work to pull air past the throttle plate, and this work (referred to as "pumping losses") is subtracted from the total useful work output of the engine, decreasing efficiency. Pumping losses can be reduced by employing a stratified charge approach.

During stratified operation, an ignitable cloud of air and fuel is formed near the spark plug. The mixture grows progressively leaner away from the spark plug, as seen in Figure 1.

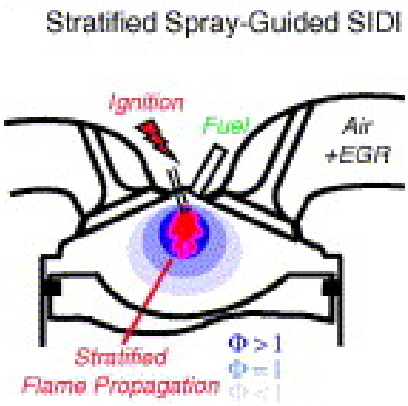

Figure 1: Stratified charge approach in a SG-SIDI engine. The air/fuel mixture is rich near the spark plug, and grows progressively leaner away from the plug (image taken from Drake and Haworth, 2007).

Although globally the mixture has too much air to burn, locally the fuel cloud near the spark plug is flammable. The engine load is controlled by injecting more or less fuel to form this cloud, rather than by limiting the amount of air entering the engine. The throttle plate can be removed, reducing pumping losses and improving efficiency.

When operating in stratified mode, engine performance is the result of many complex and inter-dependent processes, such as air flow, fuel injection, heat transfer, and combustion, which affect mixture formation, ignition, and flame propagation. Control of these processes for every engine cycle is necessary for reliable and optimum combustion in stratified SG-SIDI engines. The following paragraphs describe the role of flow field effects during a SG-SIDI cycle.

1.2 Flow Field Effects During a SG-SIDI Engine Cycle

During the intake stroke, the piston descends and flow around the intake valve takes the form of a ring-like jet which collides with the cylinder wall or piston top and sets up large scale flow structures. There is also recirculation behind the valve (Arcoumanis et al., 1987), as shown in Figure 2.

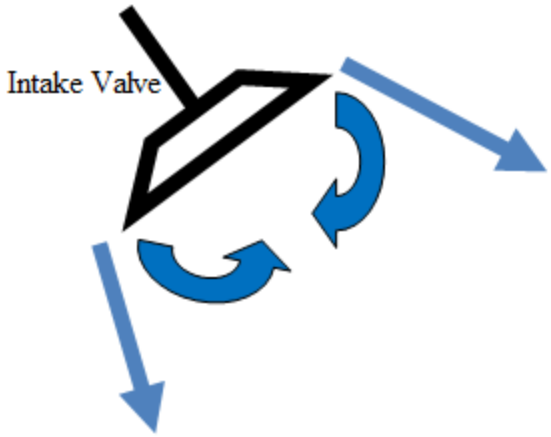

Figure 2: Schematic representation of flow past the intake valve and recirculation.

While the valve is open, the flow is largely dependent on the valve and port geometry (Jeon et al., 1998). After the intake valve closes and the piston begins to ascend, large flow structures are altered by the cylinder geometry or broken into smaller-scale turbulent eddies. Although under certain conditions the flow can approach solid body rotation, generally the in-cylinder flow during the intake stroke is 3D and highly turbulent.

As the piston nears the top of its travel, fuel is injected directly into the combustion chamber. The fuel injection drastically alters the flow field since the velocity

and the density of the fuel droplets are much higher than the in-cylinder air (2 and 3 orders of magnitude, respectively). It is primarily the interaction of the spray with the in-cylinder flow and internal surfaces of the combustion chamber that determines if an ignitable air/fuel mixture is created in the vicinity of the spark plug.

The rate of air/fuel mixing has been shown to increase with the bulk flow (Probst and Ghandi, 2003), and the orientation and geometry of the spark plug change the nature of the spray/plug interactions and can significantly affect the flow field (Peterson, 2009). The close injector/plug spacing used in spray-guided mixture formation strategies leads to high fuel spray velocities at the plug, liquid fuel impinging on the plug, and large gradients in velocity and fuel concentration, all of which can interfere with proper combustion (Fansler and Drake, 2008). Flow effects can also preferentially distribute fuel (Eaton and Fessler, 1994). For example, fuel droplets tend to be expelled from the center of a vortex, causing a locally lean region. Conversely, converging flows tend to accumulate fuel. The fuel distribution at the spark gap in a stratified engine can vary significantly prior to ignition, and variations in equivalence ratio at the spark plug have also been linked to variations in burn rate, affecting overall engine operation (Fansler et al., 1995, Cho et al., 1998).

Near the time of the fuel injection, the spark is triggered. Rather than an instantaneous point process, the high energy spark system typically employed in a SG-SIDI engine creates a spark that can last for several milliseconds. The air flow can cause the spark to be stretched several millimeters out of the spark gap, as in Figure 3.

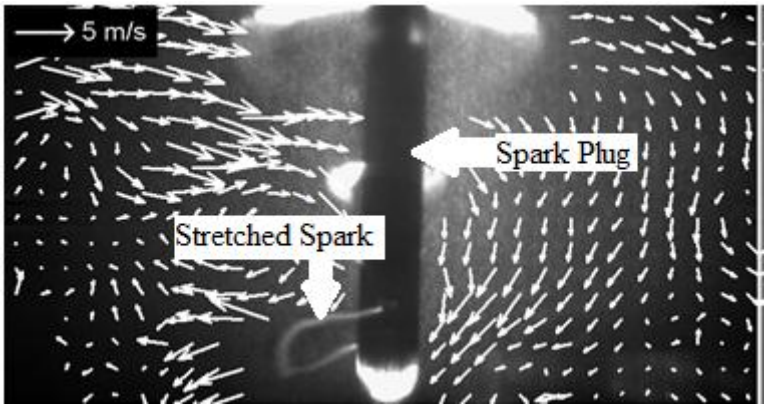

Figure 3: This image of the experimentally measured, in-cylinder flow near the spark plug shows that flow field effects can stretch the spark several millimeters out of the spark gap (Peterson, 2009).

Although some amount of spark stretch is favorable since it increases the surface area available for heat transfer to the air/fuel mixture, excessive spark stretch can cause spark blow-out and multiple spark re-strikes, which quickly drain the available spark energy and reduce spark duration (Fansler et al., 2006). The spark ignites the air/fuel cloud forming a flame kernel.

The in-cylinder flow has long been recognized as a critical factor in controlling combustion in internal combustion engines (Heywood, 1986). The presence of some flow at the spark gap is desirable since it convects the early flame kernel away from the spark plug electrodes, reducing flame kernel heat loss to the electrodes and promoting combustion (Pischinger and Heywood, 1990). Mie-scattering experiments by Fansler et al. (2006) have shown that in a SG-SIDI engine, downward flow near the spark gap is useful for convecting the developing flame kernel into the piston bowl toward the majority of the fuel.

After ignition, the flame kernel grows into a turbulent flame sheet which propagates throughout the combustion chamber. The speed with which this flame travels

is governed primarily by the mixture composition and the local 3D flow features, and the power produced during an engine cycle in general increases with the burn rate of the fuel. Some amount of turbulence increases flame speed and promotes high power output, because vortices and flow-induced flame stretch can increase the surface area of the flame, thereby increasing the burning rate. However, if the increase in surface area becomes excessive the heat loss associated with this greater surface area will cause local extinction (Mueller et al., 1998). These vortices evolve on a millisecond time scale, and span a range of sizes from less than a millimeter up to tens of millimeters (Peterson, 2009). This turbulence must be controlled so that the burn time for each cycle is approximately equivalent (Heywood, 1986). Large variations in burn time are manifested as discrepancies in power output for different cylinders and result in a rough-running engine and poor drivability.

To summarize, 3D flow field effects play a large role throughout a SG-SIDI engine cycle. To gain a greater understanding of the physical and chemical processes underlying a cycle, knowledge of the turbulent, 3D flow field is necessary.

1.3 Necessity of a New Flow Measurement Technique

Although 3D flow measurements have been performed in optical engines previously (e.g. Brücker, 1997, Calendini et al., 2000, Konrath et al., 2002), these were either done at low spatial or temporal resolution, or were done in an area of the engine cylinder with good optical access, such as through a quartz cylinder. However, many important parts of the in-cylinder air flow happen in regions with poor optical access, such as near the engine head or near the chamber wall. To capture all pertinent fluid

characteristics in a SG-SIDI engine, 3D flow measurements with millisecond temporal resolution and millimeter spatial resolution that can be performed with limited optical access are needed.

The development of a 3D flow measurement technique suited to optical engine diagnostics is also motivated in part by recent work. Extensive measurements of two-dimensional flow fields and fuel distributions in a SG-SIDI engine were performed (Peterson, 2010), and although this study drew a variety of conclusions about poor burn behavior, the 3D nature of in-cylinder flow was confirmed and it was suggested that ideal engine measurements would capture information in all three spatial dimensions instead of only two.

The SC3D-PTV method developed in this thesis can perform high-speed, high-resolution 3D flow measurements with very limited optical access over a measurement volume on the order of tens of cubic millimeters, and lead to a greater understanding of in-cylinder flow. Although the entire combustion chamber is not measured, the measurement volume captured by SC3D-PTV is still sufficient to study local flow field effects during important engine processes, such as ignition and early flame growth.

The necessity of a new flow measurement technique for optical engine flows can be demonstrated by examining the unique experimental challenges presented by optical engines, as well as the shortcomings of existing flow measurement techniques. The next section provides an overview of optical engine experiments, as well as the applicability of existing flow measurement techniques inside optical engines.

CHAPTER 2

IMAGE-BASED FLOW MEASUREMENTS IN OPTICAL ENGINES

Image-based flow measurements are attractive because they are minimally intrusive, and can capture velocity fields in multiple dimensions with high spatial and temporal resolution. However, several difficulties are associated with performing image-based measurements inside an engine.

2.1 Optically-accessible Engines

To perform image-based experiments inside an engine, the engine must be modified to be optically accessible. Maximum optical access is desirable to allow a wide range of experiments to be performed, but much of the engine is made of metal to provide the structural integrity necessary to withstand the forces generated by the engine. Endoscopes can be used to view the combustion chamber with minimal structural alteration, but offer limited light collection. Greater optical access is granted by using transparent materials such as quartz to form a portion of the combustion chamber. The trade-off between structural integrity and optical access must be balanced to provide an accessible experimental environment without risking catastrophic failure of the engine. Also, the locations of optical access cannot be arbitrarily chosen since the engine design must accommodate some functional hardware, such as valves and a spark plug. An

example of an optically accessible engine, the single-cylinder, fourth generation, optical SG-SIDI engine used at General Motors Research and Development in Warren, Michigan is shown below in Figure 4. This engine provides optical access through two quartz windows (one window is shown) in the engine head and a window in the piston bowl.

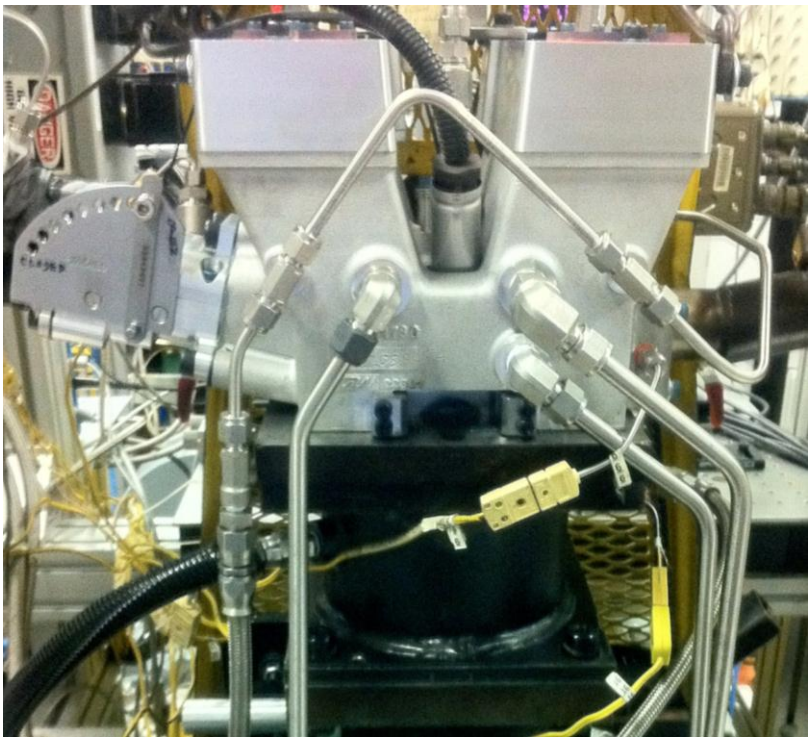

Figure 4: The single-cylinder, fourth generation, optical SG-SIDI engine in use at GM Research and Development in Warren, Michigan. Optical access is provided through two quartz windows (only one is shown) located diametrically about the engine cylinder and a piston window (shown in Figure 5).

To study the ignition processes that happen when the piston is near the top dead center position, the piston used in this engine has been modified with a transparent bowl and a side cut-out to permit viewing over the full extent of piston travel, as shown in Figure 5.

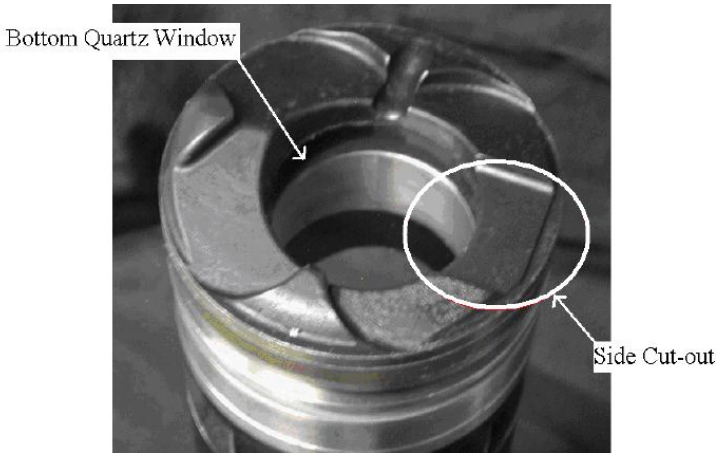

Figure 5: A piston with a quartz window in the piston crown and a side cut-out is used to provide optical access in the SG-SIDI optical engine at GM R&D.

In addition to optical access limitations, image-based optical engine measurements are prone to low signal-to-noise ratios. The windows quickly become dirty as the engine runs, decreasing image quality. If laser illumination is used, the laser light can be scattered multiple times from the internal surfaces of the combustion chamber, causing stray light to be imaged. Figure 6 shows a raw in-cylinder image both before and after image processing to demonstrate the degree to which background noise can affect image quality, and the processing which can be used to decrease the effects of background light. A sliding background subtraction scheme was applied to the image to increase the contrast between the background noise and the particles.

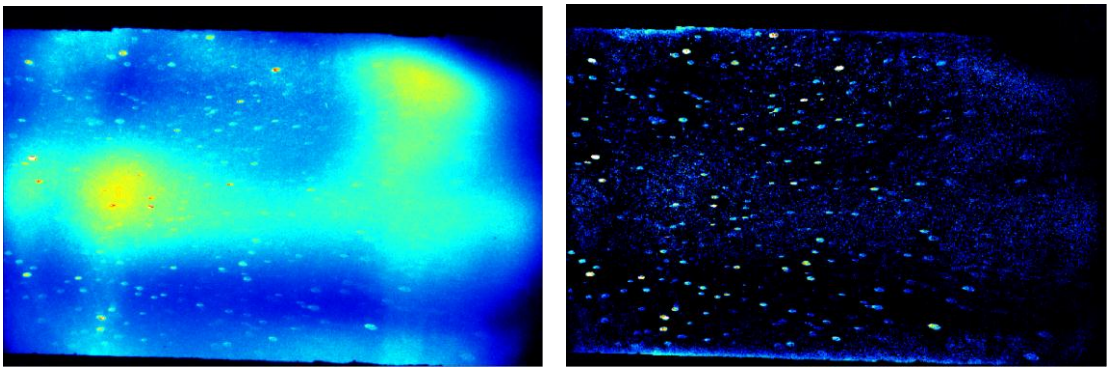

Figure 6: Raw images of the in-cylinder particle field show a large amount of background noise due to the laser light being scattered from the internal surfaces of the combustion chamber. After performing a sliding background subtraction routine, the contrast between the particle images and the background is much higher.

Despite these limitations, a number of image-based flow measurement techniques can be performed in optical engines.

2.2 Review of Image-Based Flow Measurements

In the 1930's, image-based flow measurements were applied to an engine with a glass cylinder. The intake air was mixed with chopped goose feathers and high speed images were collected. Individual feathers were tracked through successive images by their unique shape, and dividing the measured feather displacement by the known temporal image spacing gave the velocity of individual feathers, and an indication of the flow field in which the feathers were suspended (Lee, 1939, cited by Lumley, 2001). Although in-cylinder flow diagnostics have since progressed considerably, the basic principle of measuring a flow by tracking objects suspended in the flow through successive images remains the same.

Modern image-based flow measurements typically use laser light to provide focused, high-intensity illumination of tracer particles suspended in the flow. Small

particles are used since they can follow the flow more faithfully than something as large and heavy as a goose feather. For example, the particles used in the experiments described later in this document are silicon oil droplets with a nominal diameter of 1 micron. A camera images the illuminated particle field, and successive images are analyzed to determine the motion of the particles, and hence the flow field.

Image-based flow measurements can be roughly divided into two forms according to the method used to evaluate particle images: particle tracking velocimetry (PTV) and particle image velocimetry (PIV). During PTV analysis, individual particles are tracked between successive images based on spatial location or other identifying characteristics (e.g. shape, intensity, etc.), similar to the goose feather example. PIV differs from PTV since groups of particles are tracked, rather than individual particles, by sub-dividing the images into interrogation windows and performing a cross-correlation calculation on successive images. The correlation peak indicates the most likely displacement of the group of particles within that interrogation window. By tracking groups of particles, the signal-to-noise ratio and the accuracy of the measurement are increased, but the Lagrangian information available in PTV is lost.

The following paragraphs describe several types of image-based flow measurements that can be applied to optical engines.

2.2.1 Two-Component, Two-Dimensional PIV

Most of the techniques described here are based on two-component, two-dimensional (2C2D) PIV. The phrase “two-component” indicates that a two-component projection of the three-component velocity vector is measured, and the term “two-

dimensional" is used since these two-component velocities are measured within a plane. 2C2D PIV is performed by directing a laser sheet into a flow seeded with tracer particles, as seen in Figure 7. 2D projections of the particle field are captured by a camera, and successive images are cross-correlated to find the velocity field.

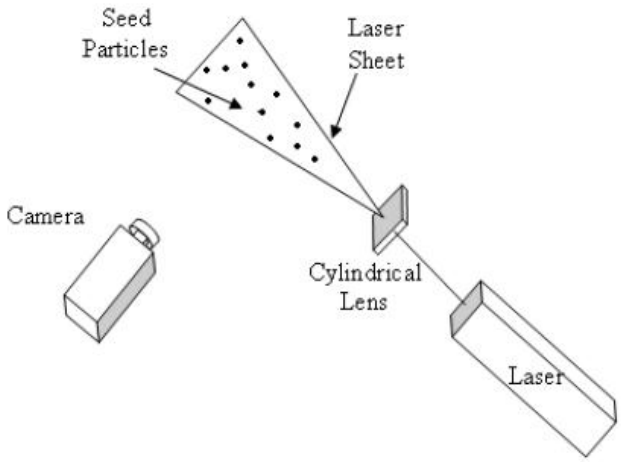

Figure 7: Schematic representation of two-component, planar PIV. The flow is seeded with tracer particles, and then illuminated with a laser sheet. A camera images the particles and velocity fields are computed by analyzing successive images (image adapted from LaVision Flowmaster Manual).

No information about the distance between the particle and the camera is retained, so only the 2C projections of the velocity components in the image plane can be determined. An example of a 2C2D PIV velocity field is shown in Figure 8.

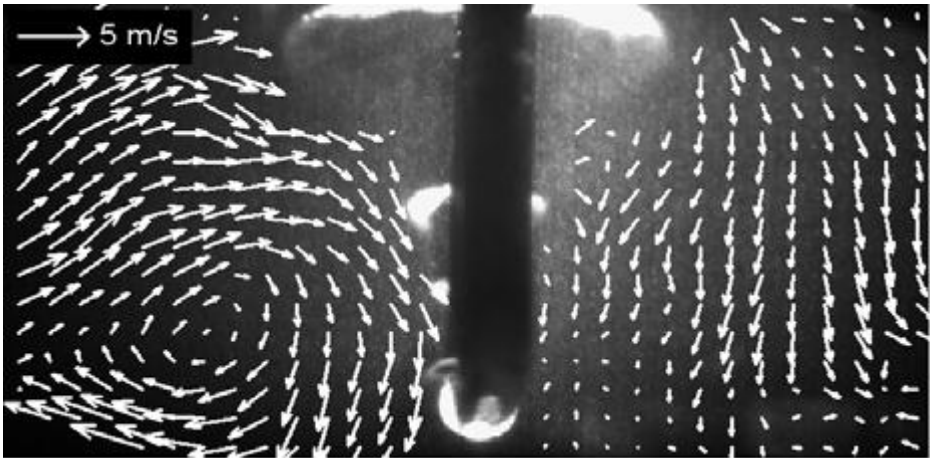

Figure 8: A velocity field showing the result of a two-component, planar PIV measurement inside an optical engine (image from Peterson, 2009).

The primary advantages of 2C2D PIV are its relative simplicity, its high spatial and temporal resolution, the minimal amount of optical access required, and the reasonable amount of computation time needed to compute vector fields. In flows without a large out-of-plane velocity component, 2C2D PIV can provide most relevant flow information, including gradient quantities within the plane. However for a highly 3D flow, such as is found in a SG-SIDI engine, the flow cannot be fully characterized without 3C3D velocity measurements.

2.2.2 Stereoscopic PIV

Stereoscopic PIV (Arroyo and Greated, 1991) is a direct extension of 2C2D PIV. Stereoscopic PIV is performed using two cameras viewing the same plane of interest, as in Figure 9.

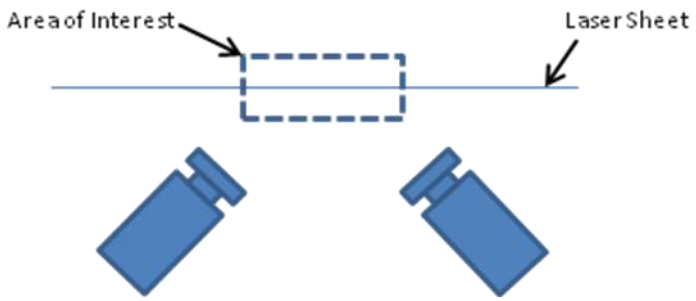

Figure 9: 3C2D stereoscopic PIV is performed using two cameras, both viewing the same area. Each camera produces a unique two-component, planar PIV measurement.

Each camera performs a unique 2C2D PIV measurement of the projections of the velocity parallel to the image sensor, as seen in Figure 10.

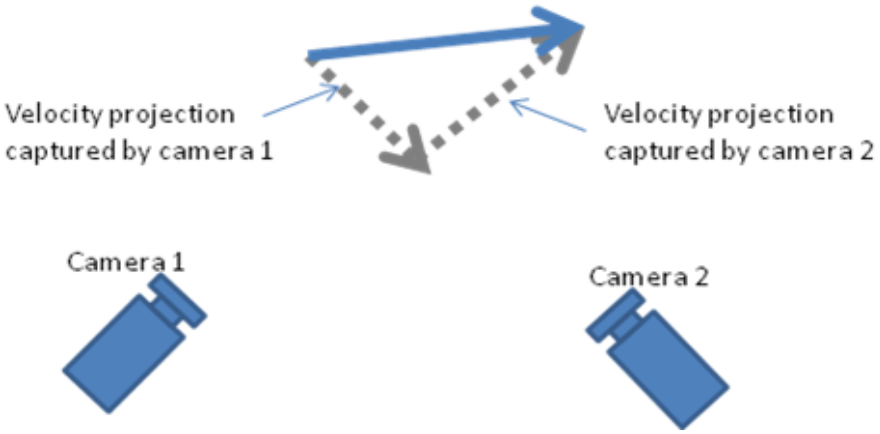

Figure 10: Each camera captures the projection of the velocity parallel to the camera's image plane.

Since the two image sensors are not parallel, the 3C velocity vector can be reconstructed by spatially matching the velocity projections from each camera. The accuracy of 3C2D stereoscopic PIV is greatest when the two cameras measure orthogonal projections of the flow velocity. This requires significantly more optical access to the measurement volume than 2C2D PIV, which uses only a single camera.

When performing 3C2D stereoscopic PIV, the two cameras typically keep the object plane in focus by satisfying the Scheimpflug condition, as seen in Figure 11, which requires that the image plane, lens plane, and object plane intersect at a point.

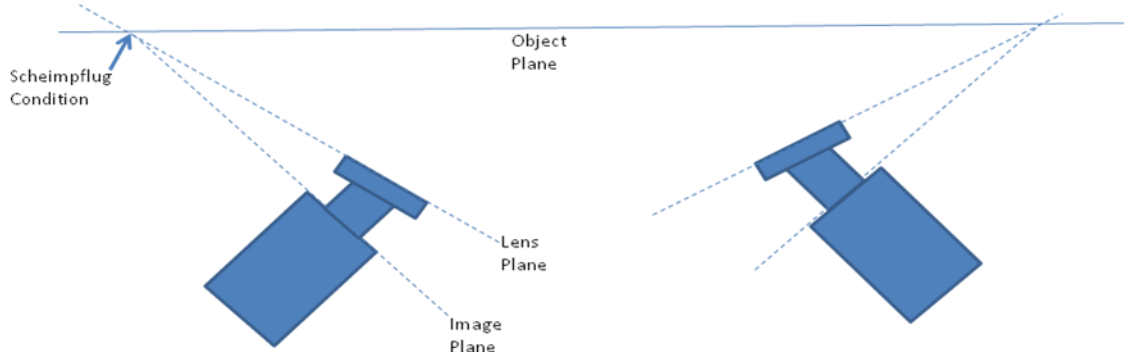

Figure 11: The cameras keep the object plane in focus by satisfying the Scheimpflug condition.

Perspective distortion results when the image and object planes are not parallel. When both cameras are on the same side of the laser sheet, this requires that a mapping function be applied to remove the distortion, as seen in Figure 12, before spatially matching the measurements from the two cameras.

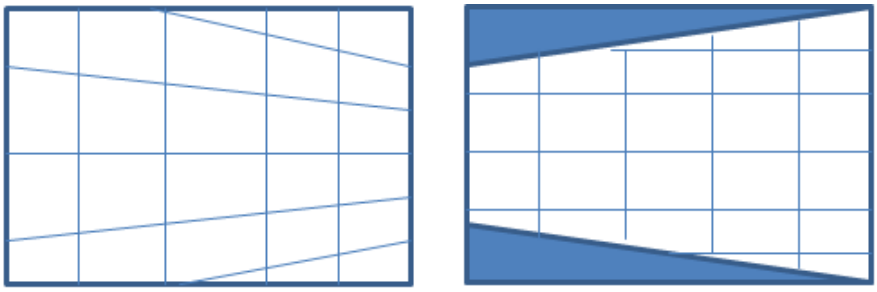

Figure 12: The distorted result of imaging a rectangular grid when the object plane and image plane are not parallel (left). For the experimental set-up shown in Figure 11, this distortion is different for each camera. A mapping function must be used to remove the distortion (right) before the contributions from each camera can be spatially matched.

The primary advantage of 3C2D stereoscopic PIV is that it measures all the components of the flow within the measurement plane. Also, the stereoscopic PIV and

2C2D PIV algorithms are similar, so data processing times are of the same order. Spatial and temporal resolutions between the two techniques are also comparable.

However, since 3C2D stereoscopic PIV measures only a single plane, this technique fails to fully quantify all fluid gradient terms (such as strain rate and vorticity) in the direction normal to the measurement plane. These gradients have been shown to affect flame growth (Mueller et al. 1998, Eichenberger and Roberts, 1999), and thus are desirable for understanding engine behavior.

2.2.3 Off-Axis Holographic PIV

Conceptually, off-axis holographic PIV is much different from single-camera or stereoscopic PIV. The recording medium used for holographic PIV does not capture a projection of the particle field, but rather the interference pattern created between a reference beam of light and the light reflected by tracer particles in the flow under study, as in Figure 13. The same beam can be used as both the reference and the source for the light scattered by the particles.

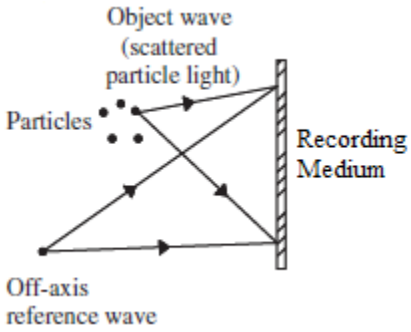

Figure 13: When performing off-axis holographic PIV, a recording medium stores the interference pattern created between a reference beam and light scattered from the tracer particles. (image adapted from Hinsch, 2002)

Using knowledge of the reference beam characteristics, the captured interference pattern is evaluated to determine the 3D particle positions. Consecutive sets of 3D particle position fields are cross-correlated, and the 3C3D velocity field is determined.

Holographic PIV can be performed using film or a digital image sensor. Although film provides significantly higher spatial resolution, manipulating the film is very time and labor intensive, and there are many opportunities to introduce error (Hinsch, 2002). Digital holographic PIV is less labor intensive, but the poor resolution of digital image sensors compared to film significantly limits spatial resolution of the reconstructed particle field (Meng et al., 2004).

Holographic PIV is superior to single-camera PIV and stereoscopic PIV since holography provides 3C3D velocity measurements, allowing all pertinent flow quantities to be resolved. Also, this technique can be performed using only a single camera. However, the reference beam must have a line-of-sight through the measurement volume, which constrains the environments that can be studied. Optical access requirements and poor spatial resolution (for digital holography), limit the applicability of this technique to engine flows.

2.2.4 Defocusing PIV

The experimental set-up for defocusing PIV has only two primary differences to that of 2C2D PIV. First, a volume is illuminated with laser light, rather than just a plane. Second, a custom aperture with off-center pinholes is used in conjunction with the image forming optics. These pinholes cause multiple images to be created for sources off the

reference plane, as seen in Figure 14, and the spacing of the multiple images can be used to determine source depth.

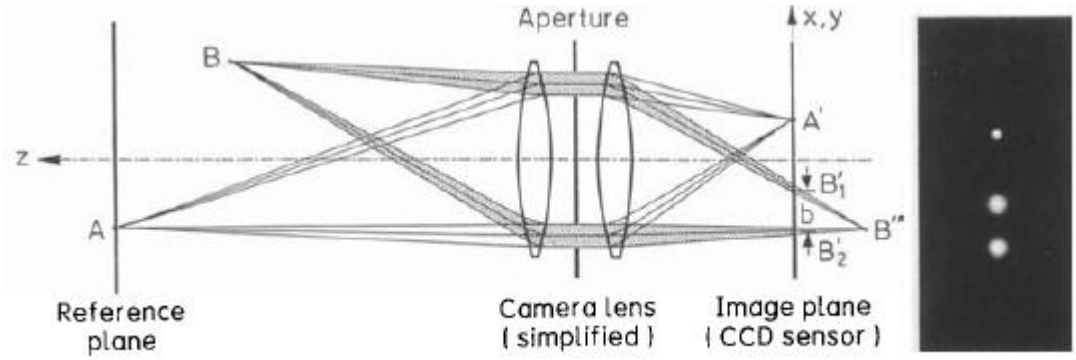

Figure 14: This simplified drawing shows how adding an aperture to an imaging system allows retrieval of depth information. Particle A is located at the reference plane and is imaged to a single point. Particle B is off the reference plane and is imaged to two points due to the aperture. The distance between the dual images of particle B can be used to determine the position of particle B along the z-axis. (image from Willert and Gharib, 1992)

When the aperture has only two pinholes, there is ambiguity between sources behind or in front of the reference plane. By introducing a third non-collinear pinhole, three images in the shape of a triangle are formed for sources off the reference plane, and the orientation of the triangle is inverted as the source crosses the reference plane. Image trios from multiple particles are imaged onto a single sensor and matched based on similarities in shape and intensity, as well as spatial location (Pereira and Gharib, 2002).

A sample image from the work of Willert and Gharib (1992) is shown below in Figure 15. Evaluating the images gives a 3D map of particle positions, and consecutive 3D particle maps are cross-correlated to give 3C3D velocity measurements.

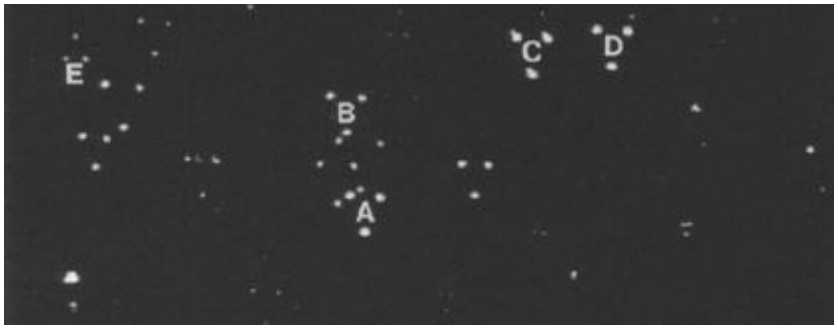

Figure 15: A sample image from defocusing PIV measurements. The characteristic triangular image pattern from sources off the reference plane is shown for labeled trios A-E. (image taken from Willert and Gharib, 1992)

As with holographic PIV, defocusing PIV also gives 3C3D velocity measurements, allowing full characterization of the flow. Image analysis is fully automated and this technique can be performed with the same amount of optical access as single-camera PIV.

However, defocusing PIV has two disadvantages of primary concern: low light collection efficiency and poor spatial resolution due to low seed density. The light collection efficiency is low since light enters the camera through only three pinholes. This requires high laser intensity to gather images of sufficient brightness. Since all particle images are captured on the same sensor, this can lead to many overlapping particles and ambiguities about which trio a given particle image belongs to, decreasing the accuracy of the measurement. Although the problem can be avoided by using a low seed density, reducing seed levels decreases the spatial resolution of the results.

2.2.5 Tomographic PIV

Tomographic PIV is performed by seeding a flow with tracer particles and illuminating a measurement volume with laser light, similar to defocusing PIV.

However, rather than a single camera, tomographic PIV employs multiple cameras to capture 2D projections of the scattered light from several different viewing angles. The original particle field is then iteratively determined, using any of a variety of reconstruction algorithms (e.g. MLOS-SART, MART) (Atkinson and Soria, 2009, Elsinga et al., 2006). Successive reconstructed particle fields are then cross-correlated to give 3C3D velocity measurements, similar to defocusing PIV or holographic PIV, but with vastly superior spatial resolution. Figure 16 below shows the key steps in performing tomographic PIV.

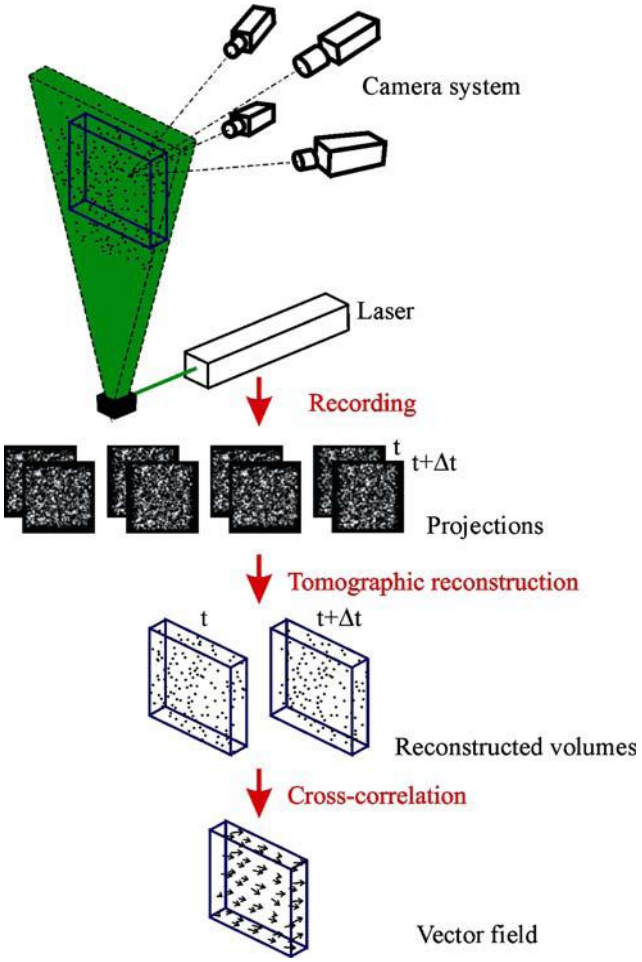

Figure 16: Schematic representation of tomographic PIV. Multiple simultaneous 2D projections of an illuminated particle field are captured, and then a reconstruction algorithm is applied. Successive reconstructed particle fields are cross-correlated to give 3C3D velocity measurements. (image from Elsinga et al., 2006)

The primary disadvantage of tomographic PIV is the large amount of optical access required to accommodate multiple (usually four) cameras and laser illumination, making this technique ill-suited to optical engine measurements.

Another disadvantage is processing times that can be orders of magnitude higher than for planar PIV. This is especially problematic given the large data sets (~10,000 velocity fields at a given condition) that are commonly used in studying poor burn behavior in SG-SIDI engines.

2.2.6 Stereoscopic Micro-PIV

Stereoscopic micro-PIV (μPIV) differs from the other measurement techniques described in this chapter since it was designed to study micro-scale flows. Micro-scale flows and flows inside optical engines both have limited optical access, and thus stereoscopic μPIV provides a conceptual basis for SC3D-PTV technique developed in this thesis. Some aspects of stereoscopic μPIV are similar to SC3D-PTV, although both the hardware and software of SC3D-PTV have distinct features not found in stereoscopic μPIV. The remainder of this section gives a brief introduction to the hardware and software used for stereoscopic μPIV. Details of the hardware and software used for SC3D-PTV are given in subsequent chapters.

Stereoscopic μPIV is based on the μPIV technique proposed by Santiago et al. (1998) for studying transport inside of micro-fluidic devices. Using coaxial illumination, flow data can be collected with limited optical access, and the data can be analyzed using cross-correlation-based PIV algorithms, resulting in 2D2C vector fields.

The capabilities of μPIV were expanded with the introduction of stereoscopic μPIV (Bown et al., 2005, Lindken et al., 2006), which uses two off-center optical paths through a single lens to create two imaging sub-systems with different viewing angles, but without perspective distortion, as seen in Figure 17. Typically, each viewing angle uses a separate camera, although a counter-example is found in Hagsater et al., 2008. Since the two viewing angles are not parallel, 3D particle positions can be calculated.

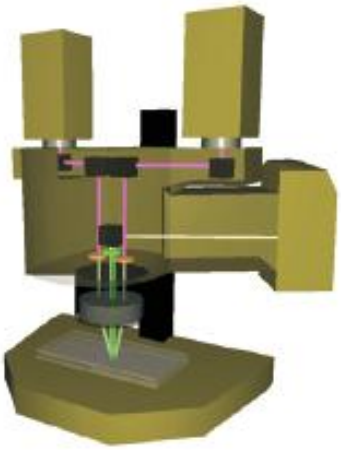

Figure 17: Two optical paths are collected through a single lens to create two views of the measurement volume without perspective distortion. (image from Bown et al., 2005)

For stereoscopic µPIV, four images are processed to obtain a single vector field. Two simultaneously acquired images constitute an image pair, and a total of two image pairs are captured successively. Early papers describing stereoscopic-µPIV (Bown et al, 2005, Lindken et al, 2006) use stereoscopic-PIV algorithms based on cross-correlation (Arroyo and Greated, 1991) to obtain 3C2D vector fields from the four images. Later papers (e.g. Bown et al., 2006, Lindken et al., 2008, Bao et al., 2011) use algorithms based on the tracking of individual particles to obtain 3C3D results from the same four-image data sets.

The main advantages of stereoscopic µPIV are its ability to collect 3C3D information with minimal optical access and relatively simple image analysis. The primary disadvantages of stereoscopic µPIV are the additional noise introduced by coaxial imaging (Olsen and Adrian, 2000) which often necessitates fluorescent tracer particles (allowing the back-reflections to be filtered), and the amount of space necessary for the experimental set-up.

2.3 Summary

To summarize the advantages and disadvantages of the techniques described above, Table 1 rates each technique on a several pertinent aspects on a qualitative and relative basis.

	Velocity Components Measured	Planar or Volumetric Measurement	Spatial Resolution	Processing Time	Optical Access Requirements
Planar PIV	2	Planar	Good	Good	Small
Stereoscopic PIV	3	Planar	Good	Good	Medium
Digital Holographic PIV	3	Volumetric	Poor	Decent	Small
Tomographic PIV	3	Volumetric	Good	Poor	Large
Stereoscopic μPIV	2 or 3	Volumetric	Good	Good	Small

Table 1: Comparison of image-based flow measurement techniques

The table shows that, except for stereoscopic μPIV, each technique has serious shortcomings that limit applicability to optical engine measurements. Although stereoscopic μPIV has many desirable features, the experimental arrangements found in the literature do not lend themselves to engine measurements. This motivates the need to develop a new three-component, volumetric flow measurement with good spatial resolution, reasonable processing time, and requires only a small amount of optical access. SC3D-PTV has these features.

CHAPTER 3

CONCEPTUAL DEVELOPMENT OF SC3D-PTV

Except for stereoscopic μPIV, each of the image-based flow measurements described in the previous chapter has inherent shortcomings with regard to optical engine diagnostics. Since stereoscopic μPIV was designed to measure environments significantly different from optical engines, SC3D-PTV employs concepts similar to those underlying stereoscopic μPIV, but uses both new hardware and new software specifically designed for optical engine measurements.

After a brief introduction to the experimental set-up of SC3D-PTV, this chapter examines the individual components comprising the imaging system and details how they are combined into a system that can measure the 3D motion of a particle. This chapter focuses on the conceptual basis of measuring particle position and motion using SC3D-PTV; further details of the hardware and software as used in practice are given in subsequent chapters.

3.1 Brief Overview of SC3D-PTV

Like stereoscopic μPIV, SC3D-PTV employs a large main lens and two smaller, off-center lenses to create two optical paths with different viewing angles, but without perspective distortion. An additional optic with an internal aperture can be used to crop

the images and prevent crosstalk between the two images. During SC3D-PTV data collection, two simultaneously acquired images are focused to a single image sensor using a mirror setup. This mirror arrangement rotates the images by 90° to use the sensor space most efficiently. Unlike stereoscopic μPIV, which uses coaxial illumination, SC3D-PTV uses volume illumination propagating in the object plane of the imaging system. A schematic representation of the volume illumination and imaging optics is seen in Figure 18.

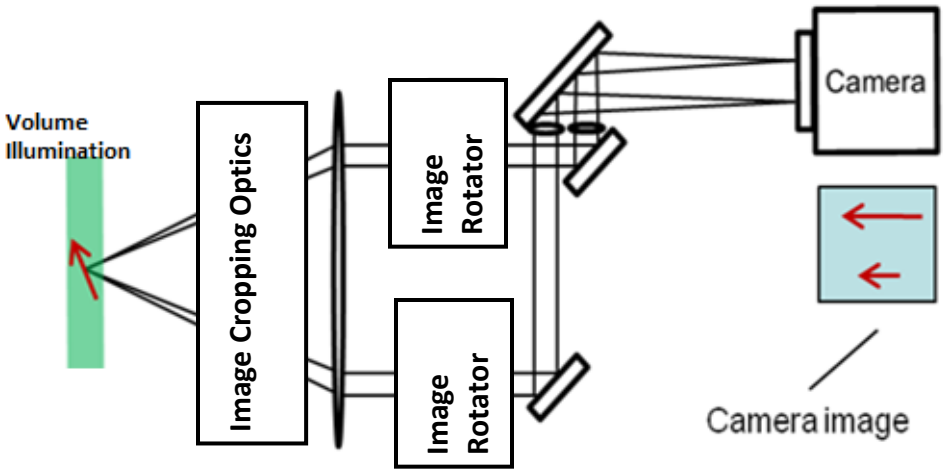

Figure 18: Two off-center apertures through a single, main lens create two views of the illuminated measurement volume. An optional image cropping optic prevents cross-talk between the two imaging sub-systems. A set of mirrors rotates the images by 90° to use the space of a single image sensor most efficiently.

Geometric optics describes the behavior of the lenses and mirrors that comprise the imaging system, so a brief explanation of basic geometric optics is given below, followed by a description of the geometric optics of the SC3D-PTV optical system.

3.2 Basic Geometric Optics

Figure 19 shows the front focal point, front focal plane, and optical axis of an ideal lens. Rays from a source at the focal point are made parallel to the optical axis after passing through the lens. Rays from a source on the focal plane (but not at the focal point) are made parallel to each other, but not the optical axis, after passing through the lens.

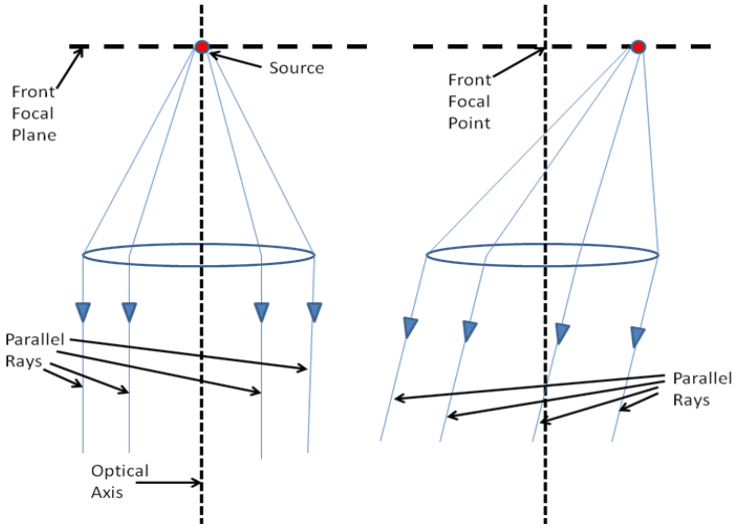

Figure 19: Rays from a source at the front focal point are made parallel to the optical axis after passing through the ideal lens (left). Rays from a source on the front focal plane but not at the focal point are made parallel to each other, but not parallel to the optical axis, after passing through the ideal lens (right).

Similarly, parallel rays entering the lens intersect at the rear focal plane. If the incoming rays are parallel to the optical axis, they intersect at the rear focal point, as seen in Figure 20.

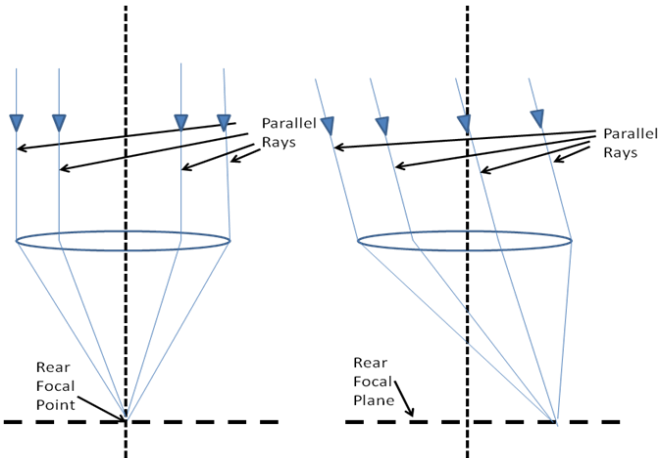

Figure 20: Rays parallel to the optical axis intersect at the rear focal point after passing through the lens. Parallel rays at an angle to the optical axis are focused at the rear focal plane, away from the rear focal point.

In Figure 21, an imaging system is created by combining two lenses and a beam blocker. Rays from a source at the front focal plane of 'Lens 1' are made parallel after passing through 'Lens 1.' These parallel rays are brought to a focus at the rear focal plane of 'Lens 2' after passing through 'Lens 2,' creating an image. This image can be recorded by placing an image sensor at the rear focal plane of 'Lens 2.'

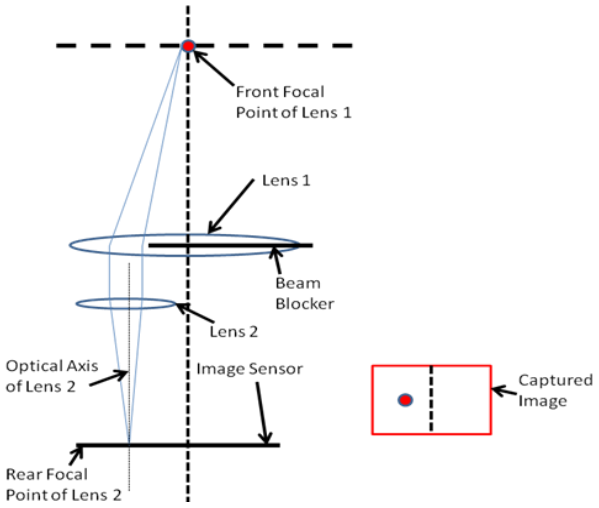

Figure 21: Rays from a source at the front focal point pass through 'Lens 1' and are made parallel, then pass through 'Lens 2' and come to a focus at the rear focal plane, producing an image of the source that is recorded by the image sensor. The beam blocker ensures that rays pass through both lenses before striking the image sensor.

Sources displaced from the front focal point of ‘Lens 1’, but still on the front focal plane of ‘Lens 1’, produce focused images at the rear focal plane of ‘Lens 2’. Sources displaced from the front focal plane produce unfocused images, as seen in Figure 22.

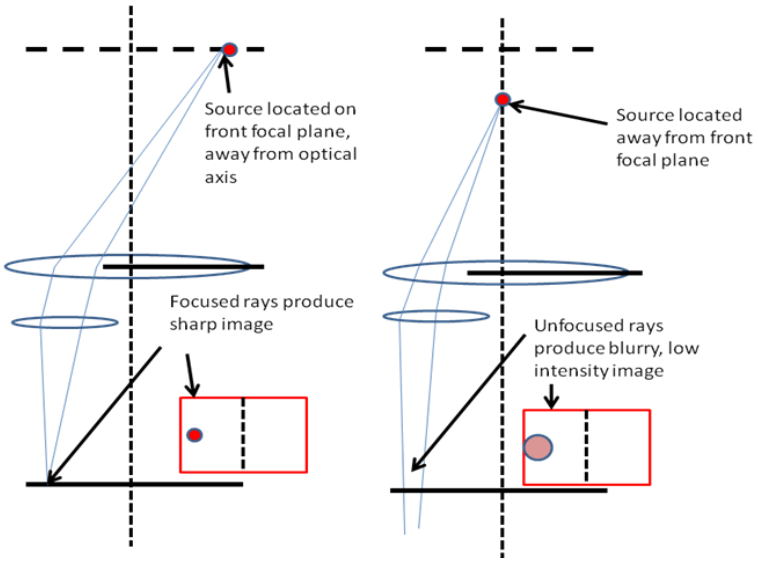

Figure 22: Sources on the front focal plane produce focused images, but sources off the front focal plane produce unfocused images.

As shown in Figure 23, a planar mirror added to the path of the rays after the second mirror changes the location at which the focused image is formed.

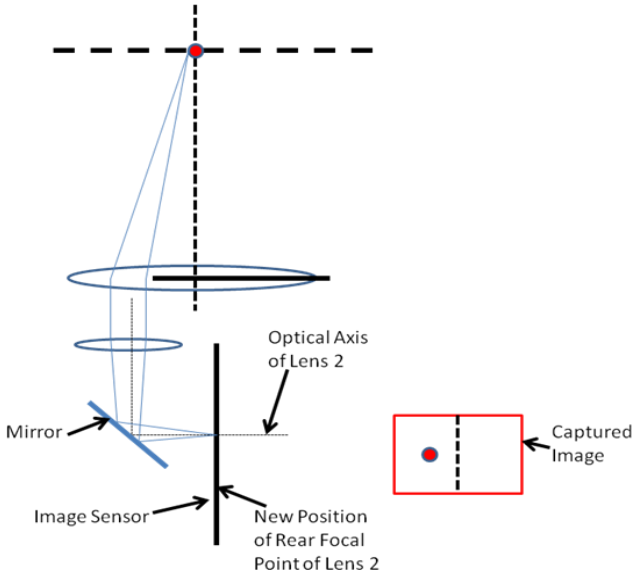

Figure 23: Adding a planar mirror after the second lens changes the position of the rear focal plane of the second lens. The focused image is now formed at the new position of the rear focal plant of the second lens.

The SC3D-PTV imaging system is made from a combination of lenses and planar mirrors in conjunction with a commercially available image sensor. The means by which SC3D-PTV can measure 3D position is shown in the next section.

3.3 Measuring 3D Position using SC3D-PTV

Imaging systems comprised of a single lens and image sensor are able to unambiguously determine two dimensions of a source's spatial location, but cannot resolve the distance between the source and the imaging system. As shown in Figure 24, for an imaging system employing two offset lenses and a beam blocker, an indication of source depth is given by the lateral displacement of the image for sources at different locations along the optical axis.

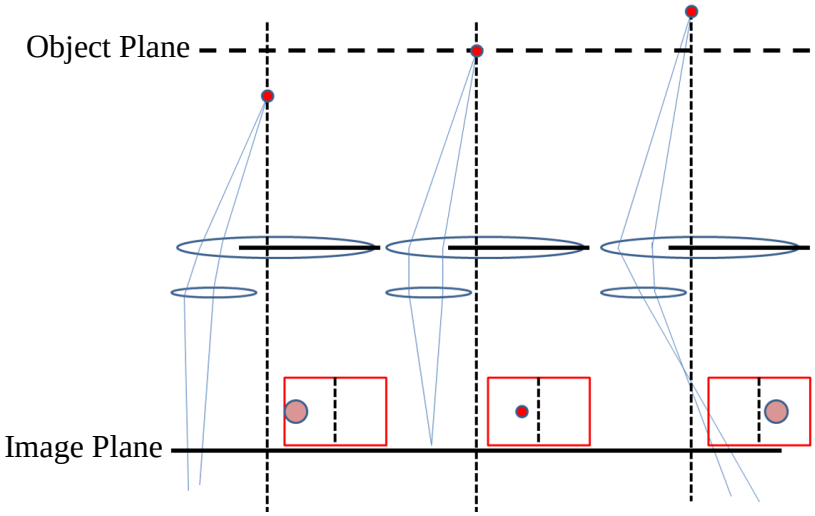

Figure 24: Sources at different points along the optical axis produce images at different lateral locations, indicating of the depth of the source.

But the images captured by this set-up do not uniquely indicate the 3D source position, as shown in Figure 25.

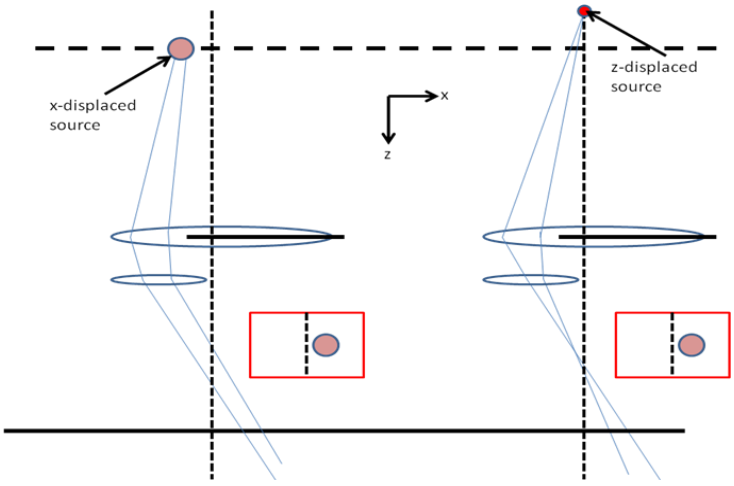

Figure 25: The 3D position of the source is not uniquely defined by the lateral position of the image. Both z-displacement and x-displacement of the source cause the image to be displaced laterally. The image from a small, intense, z-displaced source could be identical to that of a large, less intense x-displaced source.

The ambiguities present in Figure 25 can be resolved using the imaging system shown in Figure 26 (similar to Doh et al., 2007). Two imaging sub-systems are created using a main lens and two smaller lenses placed behind the main lens. A small central beam blocker ensures only light passing through the two smaller lenses reaches the sensor. Two images of a single source are recorded on the same image sensor using this configuration.

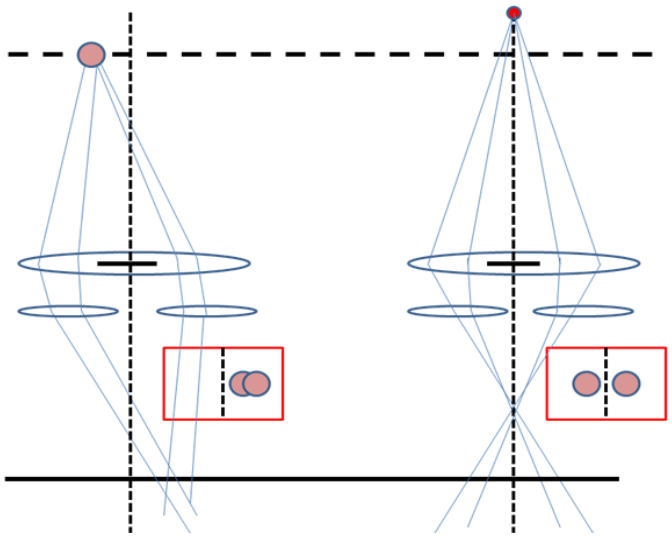

Figure 26: Using two small lenses behind the main lens creates two images of the source. The dual images provide sufficient information to resolve the ambiguities present in Figure 25 when imaging one source. The 3D spatial location of a single source can be unambiguously resolved using the set-up shown in this image, but the partially overlapping images in the system on the left can be problematic when measuring multiple sources.

The partially overlapping images seen in the imaging system on the left in Figure 26 are the result of cross-talk between the two imaging systems. This cross-talk is undesirable. Rays originating away from the focal point do not create distinct images in each half of the image sensor, rather they can create two images in one half of the sensor, or create images spanning both image halves. As seen below in Figure 27, although the

system on the left is recording two particles, and the system on the right is recording only one particle, the resulting images are identical due to cross-talk.

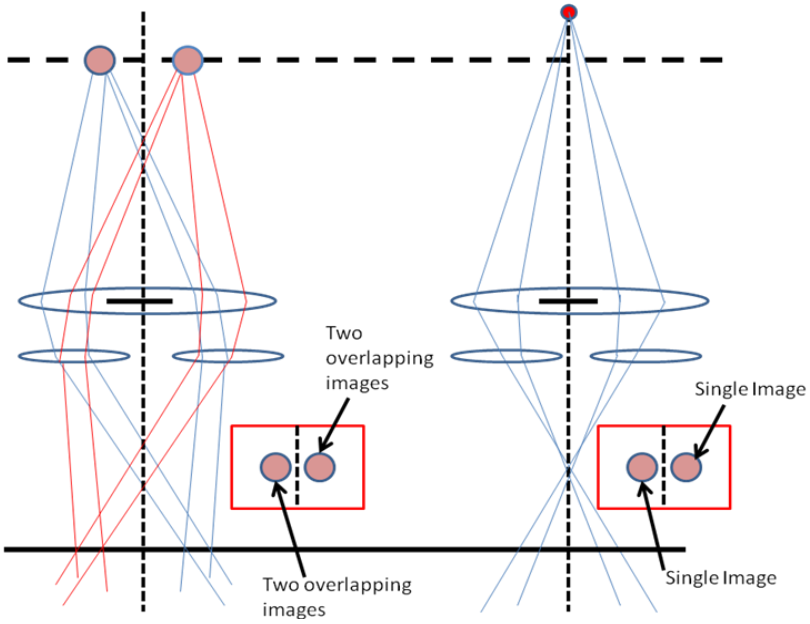

Figure 27: This imaging set-up can lead to ambiguities about the number of particles being imaged. The left and right systems record identical images, but the left system is recording two particles, and the right system is only recording one.

As seen in Figure 28, the rays originating within a small volume near the focal point do not cause cross-talk.

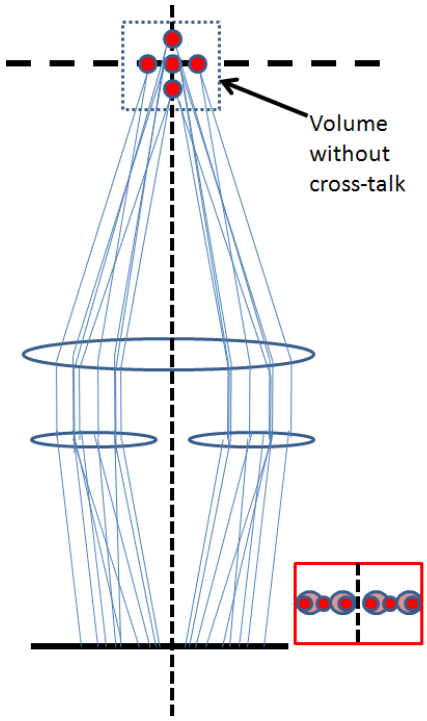

Figure 28: Rays originating from a small volume near the focal point do not cause cross-talk.

The problem of cross-talk can be solved by defining the small volume that does not cause cross-talk as the measurement volume, and employing an additional lens and beam blocker to block any rays originating from outside the measurement volume. A schematic of the additional optical components blocking rays from outside the measurement volume is seen in Figure 29. The measurement volume is centered on the objet plane of the first lens, and the beam blocker is located at the image plane of the first lens. The image formed by the first lens is at the front focal plane of the second lens.

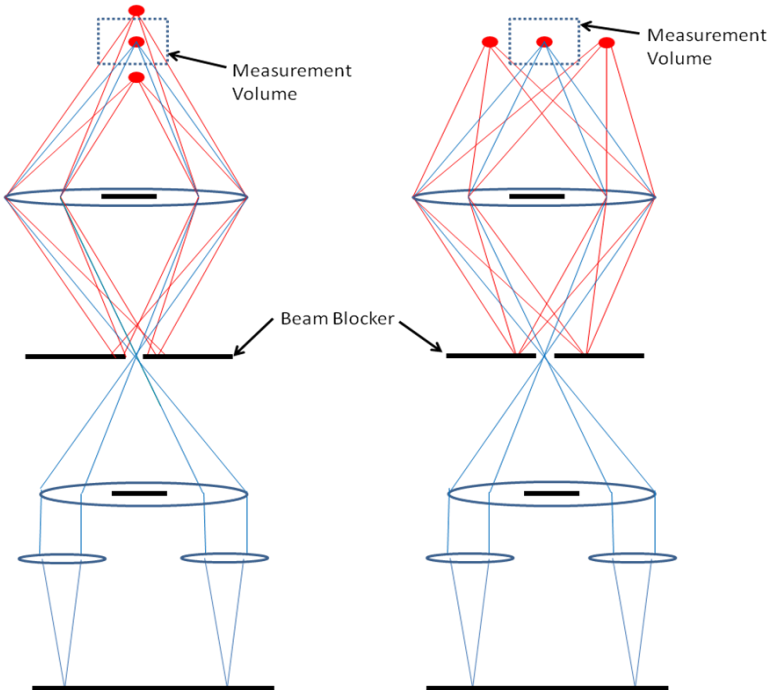

Figure 29: An additional lens and beam blocker block rays that originate outside the measurement volume, eliminating cross-talk.

Using set of planar mirrors (an "image rotator"), the image is rotated by 90° and repositioned, as seen in Figure 30.

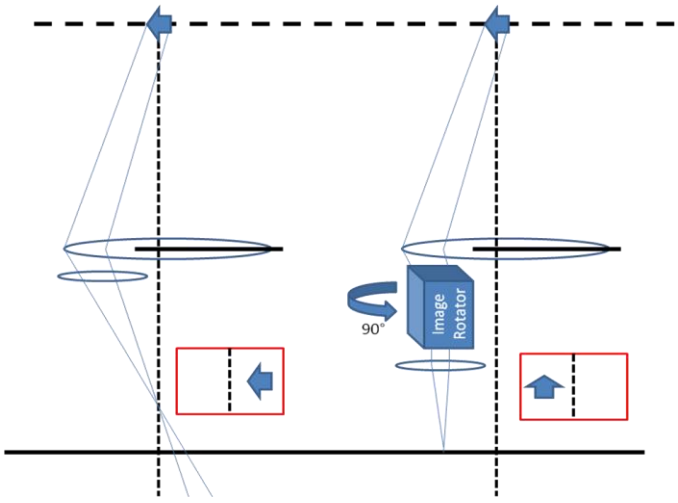

Figure 30: After adding a set of planar mirrors (an 'image rotator"), the images are rotated by 90° and repositioned so rays that enter the left (right) half of the large lens are imaged to the left (right) half of the image sensor.

Rather than the lateral image displacement being related to the source depth, the action of the image rotator causes the vertical displacement of the image to indicate the depth of the source, as in Figure 31.

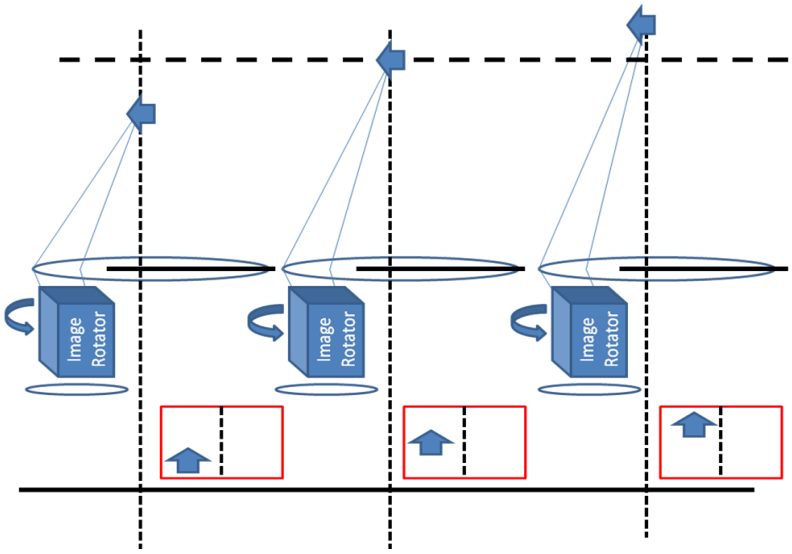

Figure 31: An indication of the depth of the particle is given by the particle image displacement, but the action of the image rotator causes this displacement to be in the vertical direction rather than the lateral direction, as in Figure 24.

Combining the image cropping optics, two smaller focusing lenses, and two sets of image rotators creates a system that unambiguously encodes the 3D position of a particle within the measurement volume into the relative position of two images of that particle, as seen schematically in Figure 32.

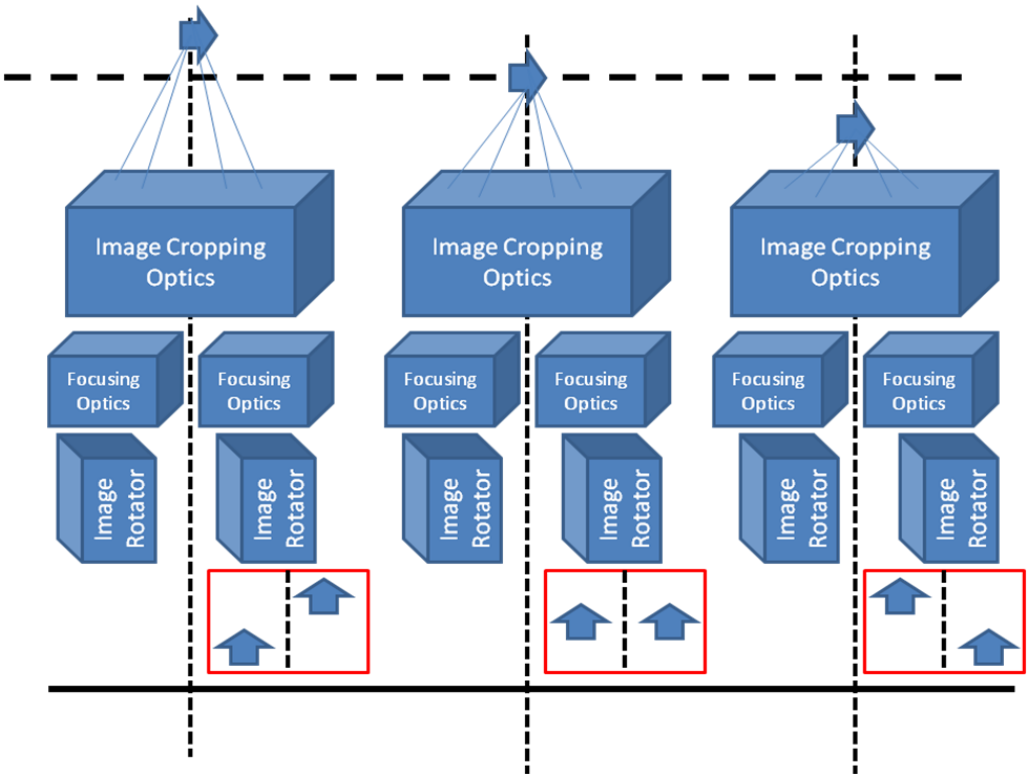

Figure 32: Combining image cropping optics, focusing optics, and two image rotators creates an imaging system that encodes the position of a source of rays into the relative position of two images of the source.

The next chapter describes the construction and characterizes the performance of a working prototype of the optic described in this chapter.

CHAPTER 4

SC3D-PTV OPTICAL ELEMENT

The previous chapter showed that using the SC3D-PTV imaging system, the 3D position of a particle within the measurement volume can be unambiguously measured. This chapter describes the steps taken to design and build a working prototype of the imaging optic used for SC3D-PTV. Although stereo-microscopes, which are similar in function, are commercially available (e.g. Carl Zeiss Optical, Inc., Leica Microsystems, GmbH), the commercial systems were not designed to project two simultaneous images onto a single sensor, nor were they designed for the specific parameters of engine experiments, so a custom SC3D-PTV prototype optic was designed and built.

The design and simulation of the prototype were performed by Dr. Boris Regaard from the Fraunhofer Center for Laser Technology, and all images in this chapter are courtesy of Dr. Regaard. The funding for the development of SC3D-PTV was provided through the Alternative Energy Technology for Transportation program (a joint program between the Fraunhofer Center for Laser Technology and the University of Michigan).

4.1 Optical Element Simulation

The first step toward bringing the SC3D-PTV concept to reality was to verify the design (without the additional image cropping optics) in simulation, as seen in Figure 33.

SC3D-PTV was simulated using Zemax, a commercially available, ray tracing optical system simulator.

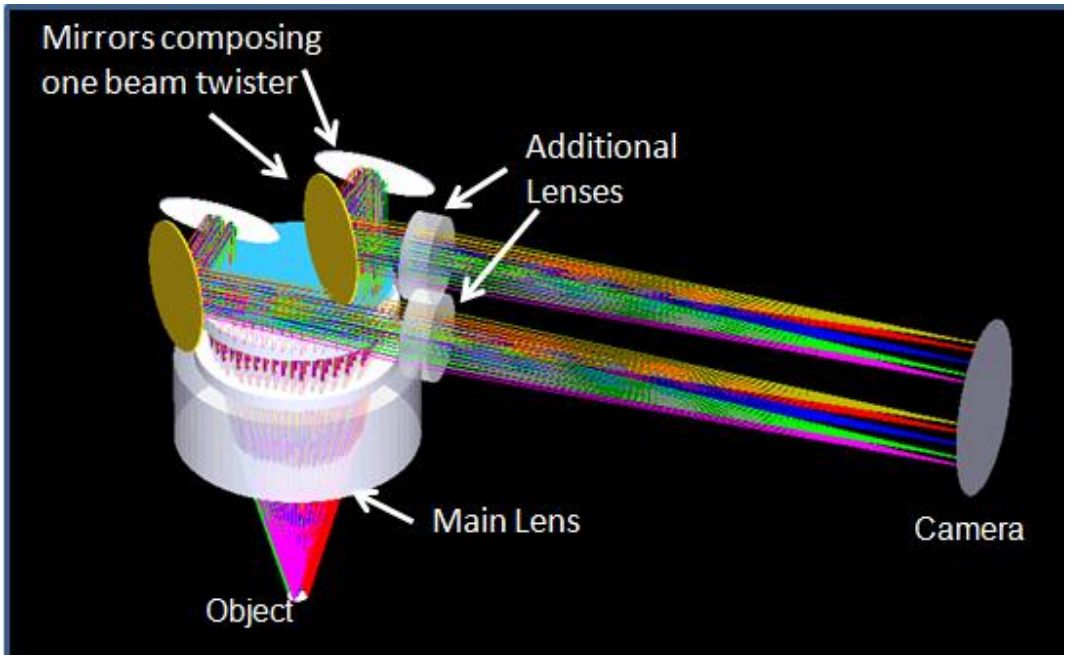

Figure 33: Simulation of SC3D-PTV using the commercial ray-tracing software Zemax.

The simulation does not account for all non-idealities (like lens aberrations or misalignments of the lenses and mirrors), but does indicate of the feasibility of the design and estimates the imaging properties. In anticipation of the experimental parameters of future engine tests, the SC3D-PTV system was designed to image a measurement volume approximately 6mm x 6mm x 4mm at a distance of a few centimeters from the main lens. The working distance was kept to a minimum to maximize resolution of the depth measurement, but the minimum working distance for engine experiments is bounded by the fact that the imaging system must be located outside the combustion chamber.

The 6mm x 6mm x 4mm volume was chosen since previous work demonstrated the importance of the two-component, planar flow field and planar fuel concentration in a 4mm x 4mm area downstream of the spark plug for poor burn behavior (Peterson, 2009).

Although this thesis does not include any misfire analysis, the natural extension of those 2D measurements is a series of 3D measurements over a volume centered on the same area.

As shown in Figure 34, the image of a grid within the measurement volume was simulated at different depths to demonstrate how object depth is related to the relative position of the two simultaneously acquired images. For a grid at the focal plane, the two images are at the same vertical position. As the depth of the grid within the measurement volume was changed, the images shifted relative to each other in the vertical direction, with the degree of the shift corresponding to the depth of the simulated grid. The focus of the grids was also affected.

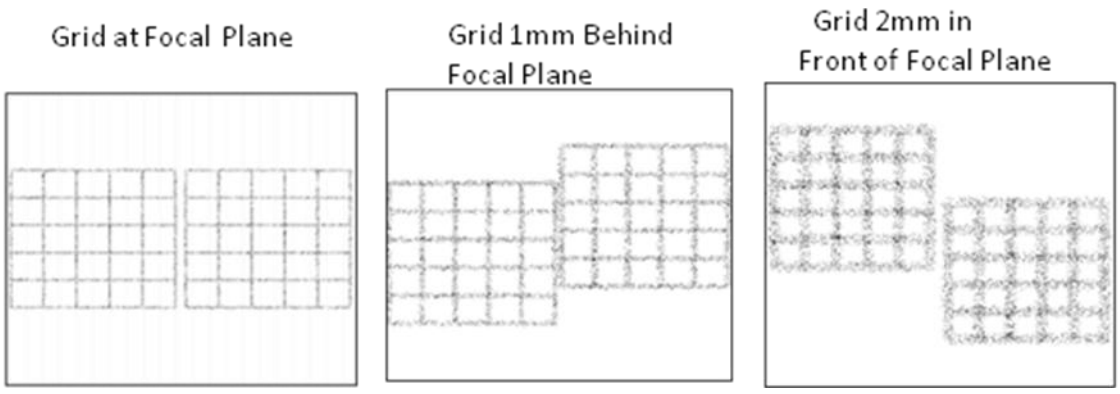

Figure 34: Simulated SC3D-PTV images of a grid show the expected displacement and defocusing for objects off the focal plane.

To determine spatial resolution, point sources were simulated at various locations in the measurement volume. These sources were displaced by small increments within the simulated measurement volume and their position was then calculated based on the resulting simulated images, as seen in Figure 35. To a first approximation, the spatial resolution of SC3D-PTV for these experimental parameters corresponds to a one pixel

shift of the image center. Perpendicular to the optical axis of the optic, the spatial resolution is 0.02 mm per pixel. Parallel to the optical axis (i.e. the out-of-plane component) a shift of one pixel corresponds to a displacement of 0.05 mm.

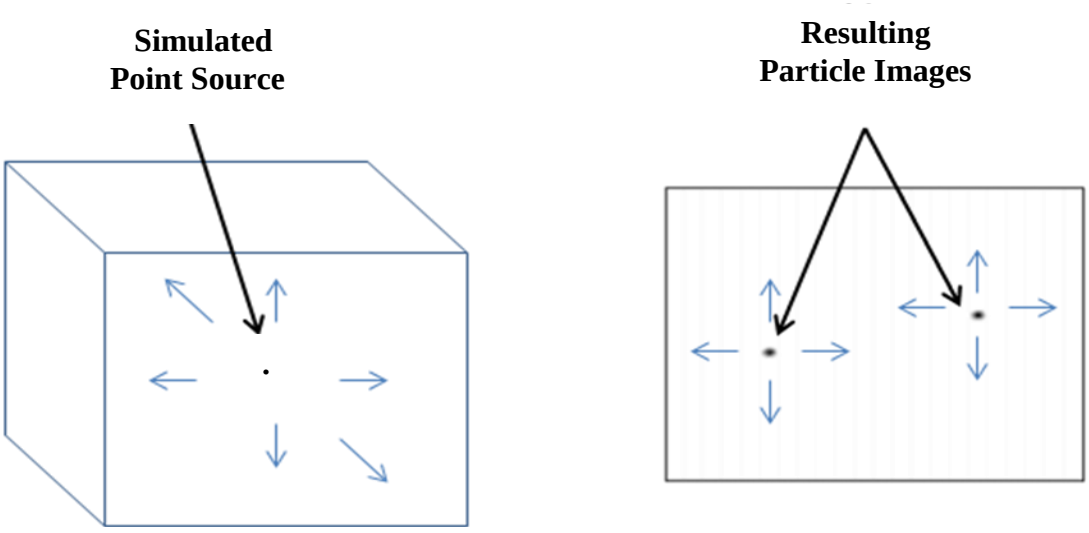

Figure 35: A simulated point source was translated within the measurement volume. The translation of the resulting images was used to determine the spatial resolution of SC3D-PTV for a set of parameters closely matching the experimental conditions in an engine.

In addition to the smallest measurable shift of a particle image, another pertinent measure of spatial resolution is the number of distinguishable line pairs per millimeter. This quantity affects the size of the particle images and ultimately the number of particles per unit image space that are resolvable. When particle images overlap, they are no longer distinguishable as single particle images, however, as detailed in Chapter 8, particle image density is also limited by factors related to the analysis algorithms.

As seen in Figure 36, the number of distinguishable line pairs per millimeter was simulated at different depths for both a point at the center of the measurement volume, and a point displaced laterally from the center by two millimeters. Resolution drops quickly away from the focal plane. Also, the focal plane is slightly curved, meaning that

the position of best-focus (maximum resolvable line pairs per millimeter) is not at the same depth position for the center point and the laterally displaced point.

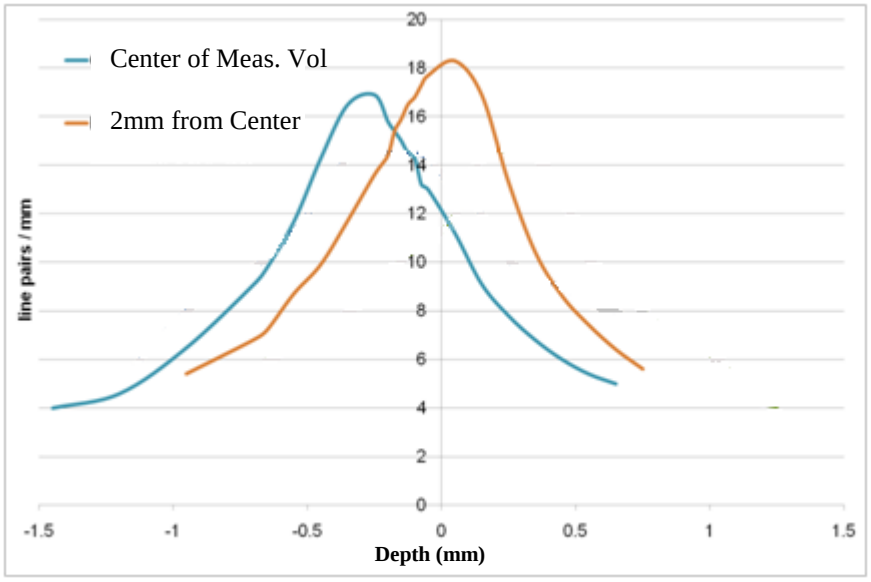

Figure 36: The number of line pairs per millimeter that can be resolved affects the number of particles that can be distinguished within a given area. Simulation data shows that the resolution of both the center and edge of the measurement volume peak at about 17 line pairs per millimeter, but decrease quickly away from the focal plane located at a depth of 0 mm.

4.2 Physical Prototype of Optical Element

After the design of the optic was verified in simulation, a physical prototype was built. The optic was produced in two pieces: a base unit and an additional one-to-one imaging optic that can be used to eliminate cross-talk in certain experimental configurations. For the base unit, the optics (off-the-shelf lenses and mirrors) are packaged in a sealed aluminum housing (approximately 6” x 6” x 8”, as seen in Figure 37) that mounts to a Phantom v7.1 (Vision Research) camera. The camera has an image sensor of 800x600 pixels with ~20 um pixel pitch. The SC3D-PTV optic has a focal length of 77 millimeters, and apertures of five millimeters each.

By mounting the optic directly to the camera and sealing the aligned optics, experimental complexity is reduced, the optics are protected from environmental contamination, and the amount of laboratory space required is minimized.

Figure 37: Base unit of the SC3D-PTV imaging system. After the performance of the optical element was simulated, a physical prototype was built. The optics are mounted inside a sealed housing that attaches directly to the camera.

Under certain experimental conditions, an additional image cropping optic is necessary to prevent cross-talk between the two imaging sub-systems. A schematic of the image cropping optic was shown previously in Figure 29. An example of an image with cross-talk is shown in Figure 38.

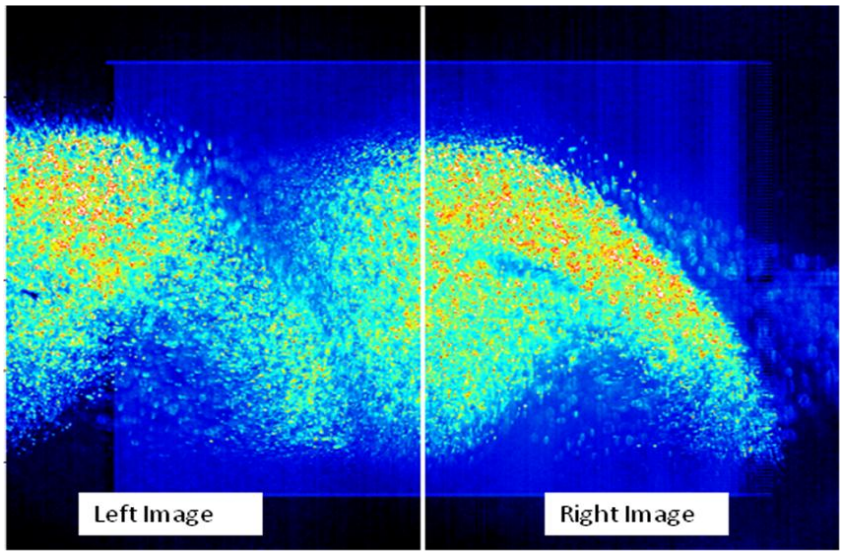

Figure 38: Example of cross-talk. The left half and right half of the image sensor should contain distinct views of the measurement volume, but this image shows the structure in the right half of the sensor extending into the left half.

If the size of the measurement volume can be limited by either the size of the seeded flow or the size of the laser illumination, the additional optic is not necessary. However, for engine experiments the measurement volume cannot be limited in this way, so the additional optic is needed.

The additional optic has a cylindrical aluminum housing approximately 12” long and 2” in diameter, as seen in Figure 39, and is threaded to fit directly onto the end of the base unit.

Figure 39: An additional one-to-one imaging optic can be threaded onto the base SC3D-PTV unit to eliminate cross-talk. The additional optic forms an image onto an aperture, and the aperture blocks all light that would cause cross-talk.

The additional one-to-one imaging optic creates an image on an aperture. The image formed on the aperture is then re-imaged by the base unit. This effectively blocks all light originating from sources outside the measurement volume, as described in Chapter 3. The aperture is shown in Figure 40.

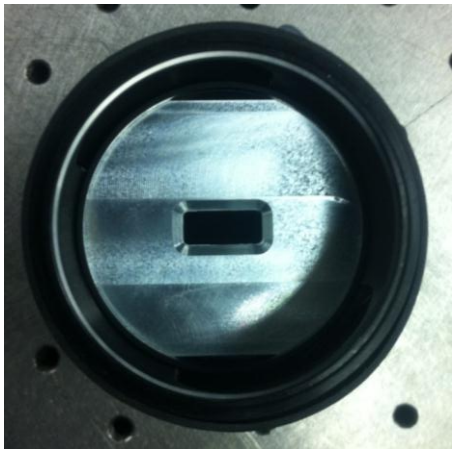

Figure 40: An aperture inside the additional one-to-one imaging optic blocks all light originating from sources outside the measurement volume, thereby eliminating cross-talk.

After building the SC3D-PTV optic, the performance of the optic was determined by imaging a USAF target at different depths. This test shows the resolution in line pairs/mm. Comparing the measurement and simulation, as in Figure 41, shows that the resolution for the prototype is below that predicted by the simulation, both at the center of the measurement volume and at an off-center point laterally displaced two millimeters from the center, likely because the simulation does not account for all real world non-idealities.

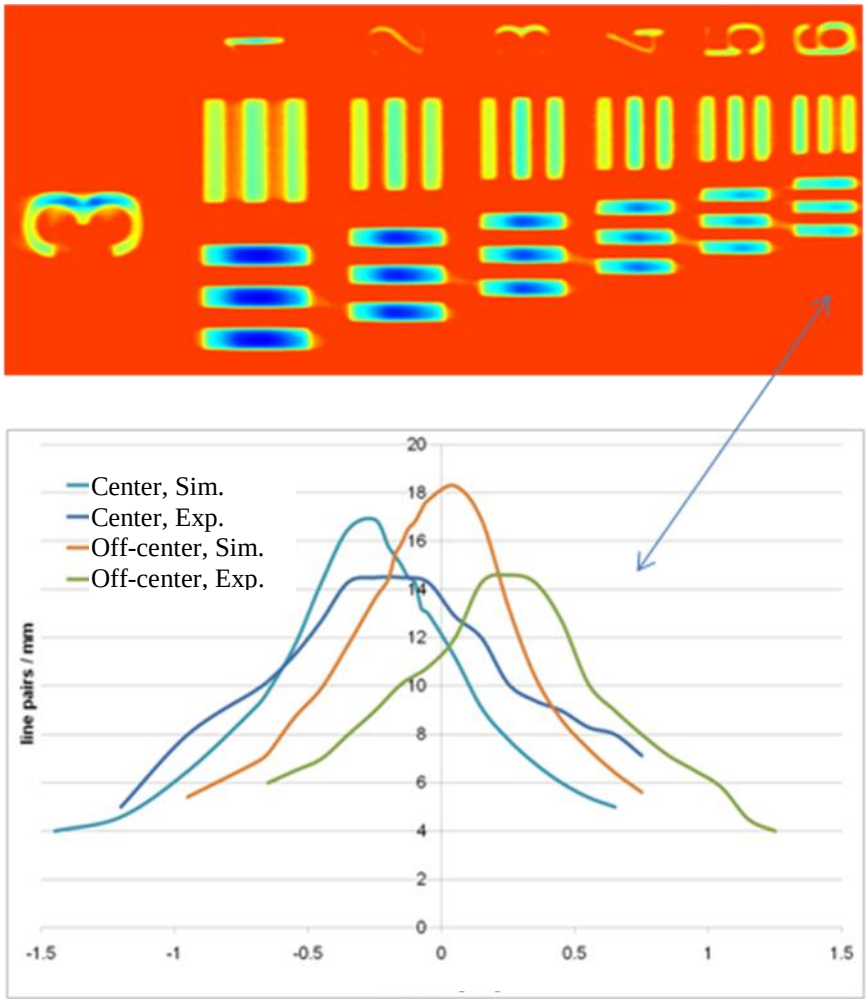

Figure 41: A USAF target (top) was imaged to experimentally determine the resolution of the SC3D-PTV optic in line pairs per millimeter, and compared to the values given by the simulation (bottom).

4.3 Summary

This chapter showed the process of verifying the concept of SC3D-PTV through simulation, and then using that simulation to construct a working prototype of the optic. The performance of the optic was also quantified. The next chapter describes the process for recovering calibrated, 3D motion of a single particle from multiple images of the particle.

CHAPTER 5

MEASURING THE DISPLACEMENT OF A SINGLE PARTICLE

Chapter 3 showed that using the images acquired with the SC3D-PTV imaging system, it is possible to unambiguously recover the 3D position of a particle within the measurement volume. Chapter 4 showed the SC3D-PTV concept being verified in simulation and the construction of the optic. This chapter describes in detail how the three spatial dimensions of particle position are encoded into the images acquired with the SC3D-PTV optic, and how multiple successive measurements of a single particle can be used to determine particle displacement.

5.1 Measuring 3D Particle Position

To motivate the geometric relations that permit the recovery of 3D particle position, Figure 42 and Figure 43 illustrate the relationship between particle translation perpendicular to the optical axis of the imaging system and particle image translation within the two simultaneously acquired images.

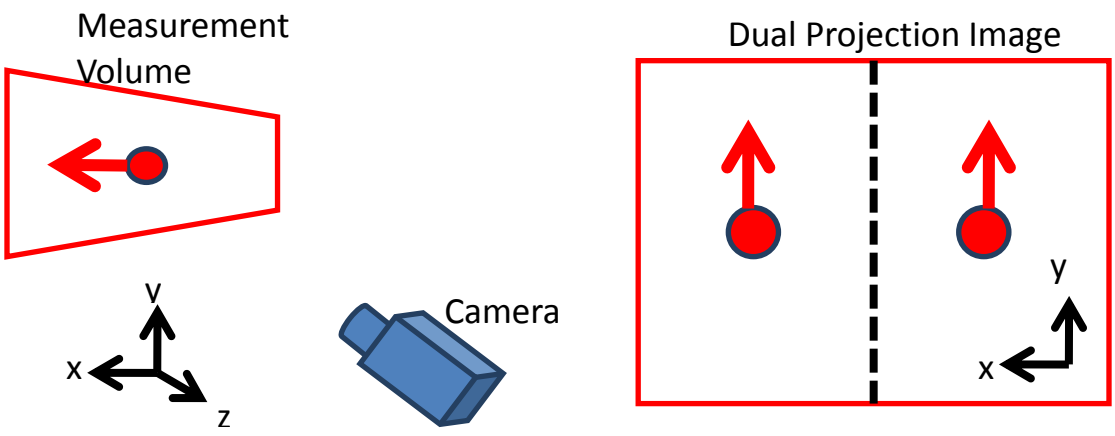

Figure 42: As the particle is displaced in the x-direction within the measurement volume, both particle images are displaced along the y-axis in the same direction and magnitude. X-displacement of the particle corresponds to y-displacement of the particle images due to the action of the beam twister.

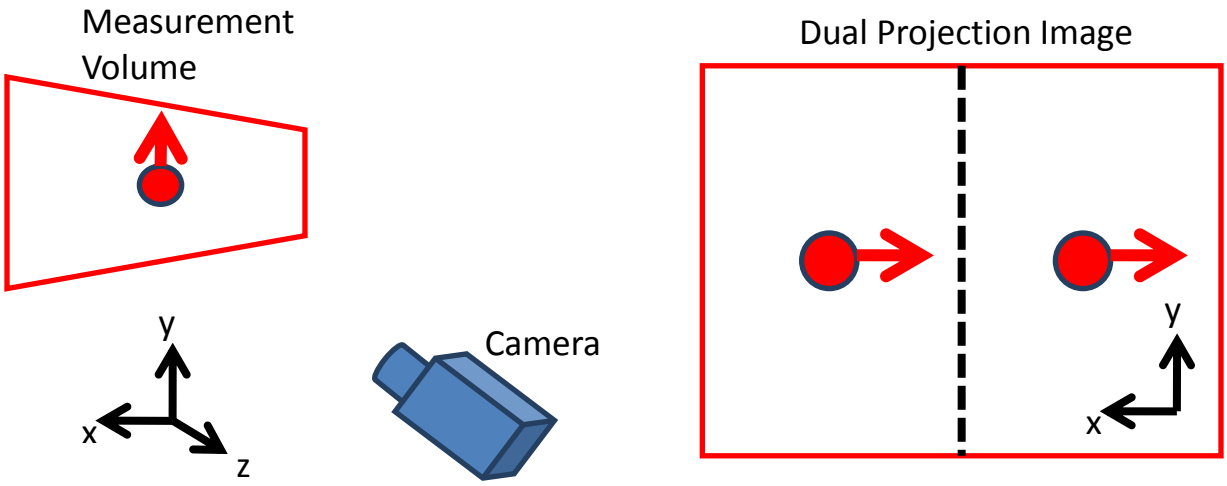

Figure 43: As the particle is displaced in the y-direction within the measurement volume, both particle images are displaced along the x-axis in the same direction and magnitude. Y-displacement of the particle corresponds to x-displacement of the particle images due to the action of the beam twister.

For any particle on the focal plane, the rays entering each imaging sub-system (as described in Chapter 3) will have the same angle with respect to the optical axis due to the collimating action of the lens. This results in the same y-position within the image for both images of the particle. However, rays from a particle off the focal plane are not

parallel after passing through the main lens, and the beam twisters cause this angular information to be encoded into the y-displacement of the images. This angular information can then be used to determine the z-position of the particle. Greater separation of the images in the y-direction indicates greater displacement of the particle from the focal plane, as shown in Figure 44.

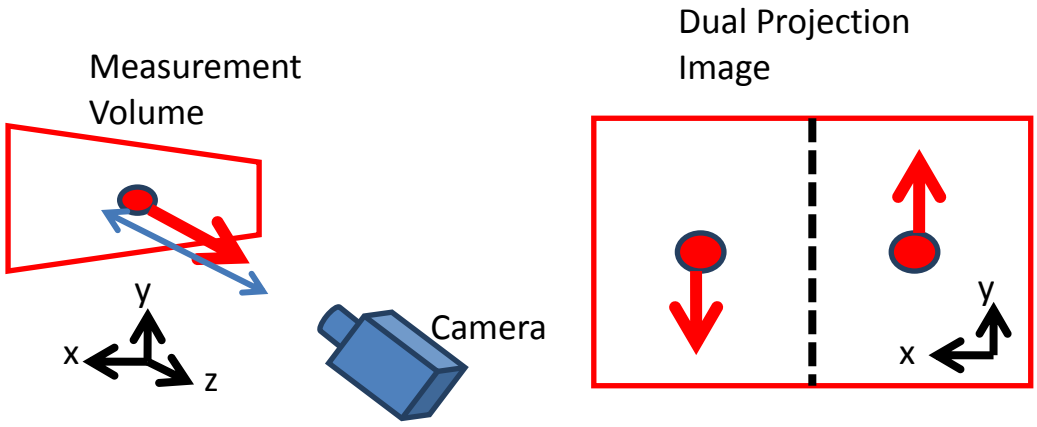

Figure 44: Particles at the focal plane produce two images at the same y-position. As the particle is displaced in the z-direction (toward or away from the camera) within the measurement volume, both particle images are displaced by the same magnitude, but the particle images are displaced in opposite directions along the y-axis. The direction of particle image displacement depends upon whether the particle is moving closer or further from the imaging optics.

Recovery of 3D spatial position is done through basic geometric considerations. Using the same coordinate system as Figure 42, the x-position of the particle is given as the average y-position of the images multiplied by a calibration constant. The y-position of the particle is equal to a calibration constant multiplied by the x-position of either of the images (since both particle images will be at the same x-position within their unique projection). The z-position is given by the difference in y-positions of the particle images multiplied by a calibration factor. The calibration factor for the z-position of the particle is not a constant, but rather a function that depends upon the z-displacement. The

calibration factor has z-dependence since the z-position is determined from the difference in angle between the rays entering each of the imaging sub-system, and the change in angle is not constant with the change in z-position. These relations are illustrated in Figure 45.

- $X_{particle} = C_1(Y_{image1}+Y_{image2})/2$
- $Y_{particle} = C_2(X_{image})$
- $Z_{particle} = C_3(Y_{image2}-Y_{image1})$

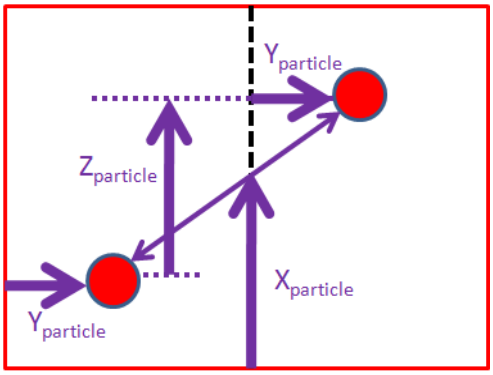

Figure 45: The 3D spatial location of the particle can be determined from geometric considerations using the equations shown above. $C_{1,2}$ are calibration constants which depend on the magnification of the images. C_3 is a function of the magnification, but is also a function of z-displacement.

The calibration constants are determined by imaging a calibration plate at a series of known depths within the measurement volume. The calibration plate was a thin aluminum plate drilled with 0.06mm holes on a grid of 0.3 x 0.3 mm and illuminated from the rear. An example calibration image from two different depths is shown in Figure 46.

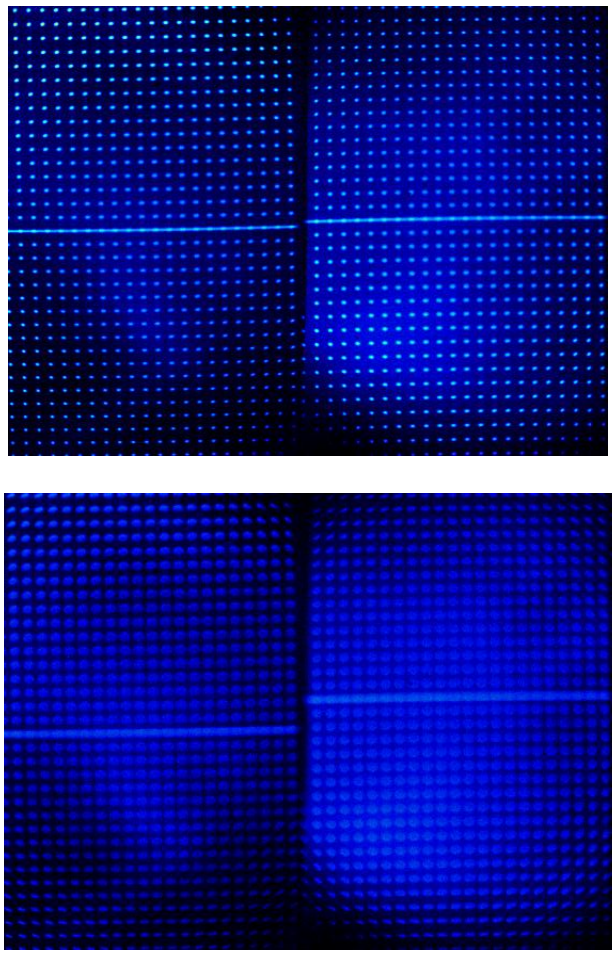

Figure 46: Calibration images taken at two different depths within the measurement volume: at the focal plane (top) and 1.5 mm behind the focal plane (bottom). Each calibration image is two different, simultaneously-acquired views (the left half of each image is one view, the right half is the other). As the calibration plate is translated away from the focal plane the images become defocused and shift vertically with respect to each other.

Using the formulas in Figure 45, together with the calibration data, the 3D spatial position of a particle can be calculated from the position of the two particle images. To emphasize that the two simultaneously acquired images are taken from separate viewing angles, although they are captured on a single image sensor, the remaining explanation of the algorithm uses the convention of showing the image captured from the left viewing

angle in a red border, and the image captured from the right viewing angle in a green border. Also, the images are rotated 90° to reflect the actual orientation of the measurement volum, as seen in Figure 47.

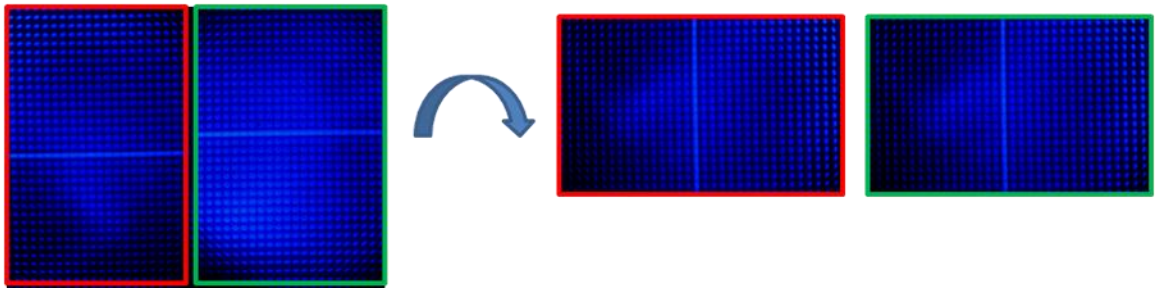

Figure 47: The convention of showing the image taken from the left viewing angle in a red border, and the image captured from the right viewing angle in a green border is used in the remaining discussion of the algorithms. Also, the images are rotated 90°.

5.2 Measuring 3-Component Particle Velocity

The 3D position of a particle can be determined from one set of simultaneously acquired images, but in order to determine the displacement of a particle a second set of simultaneously acquired images must be recorded at a later time. By comparing the relative and absolute positions of the four images of a single particle, the three-component magnitude and three-dimensional position of the velocity vector corresponding to that particle can be obtained. This process is shown in Figure 48 using four separate images of the particle, referred to as image L_0, R_0, L_1, and R_1. The letter (L or R) indicates whether the image was acquired from the left or right imaging sub-system, and the subscripted number (0 or 1) indicates whether the particle was acquired at time=t_0 or time=t_1. The relative and absolute positions of the particle images in L_0 and R_0 (L_1 and R_1) indicate the absolute 3D position of the particle at t_0 (t_1). The relative positions of the particle images between L_0 and L_1 (R_0 and R_1) indicate the particle's two

displacement components in the object plane for the left (right) imaging sub-system, respectively. Dividing the displacement by the temporal separation of the successively acquired images gives the particle velocity.

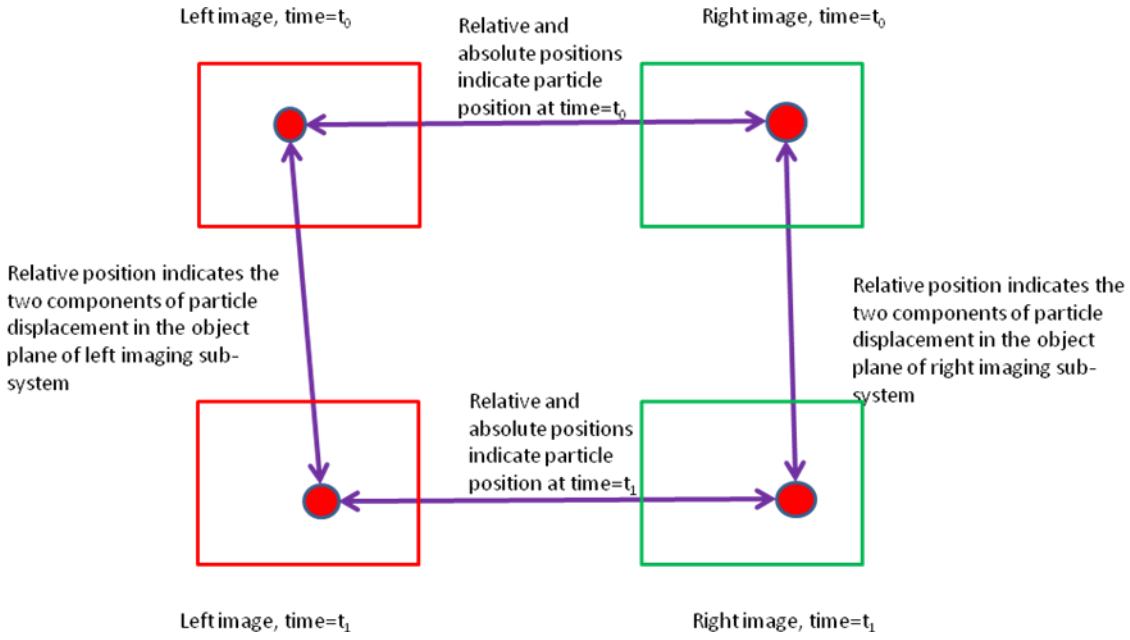

Figure 48: Four images of a single, physical particle are required to determine the velocity of that particle. The relative and absolute positions of the particle images in simultaneously acquired images indicate the particle position in 3D real space, and the relative position of particle images in successively acquired images indicates the two spatial displacement components in the object plane of that imaging sub-system.

The average of the two 3D particle positions is used to place the vector in 3D real space, and the two 2C velocity vectors determine the 3C velocity vector associated with that particle.

5.3 Summary

This chapter describes the process of using the relative and absolute positions of four particle images to determine a 3C3D velocity vector for that particle. When measuring the velocity of a single particle, this process is very straightforward. However,

when multiple particles are present within the measurement volume, the four particle images corresponding to a single physical particle must be correctly matched, or the resulting vector will be spurious. The process of determining the 3C3D velocity vectors of multiple particles within the measurement volume at the same time is the subject of the next chapter.

CHAPTER 6

MEASURING THE DISPLACEMENT OF MULTIPLE PARTICLES SIMULTANEOUSLY

As described in the previous chapter, the relative and absolute positions of four images of a particle provide enough information to determine the 3C3D velocity vector of a single particle. When measuring multiple particles at the same time, the particle images corresponding to a single, physical particle must be matched before displacements can be computed. This chapter describes existing algorithms for computing simultaneous and successive particle image matches, and details the algorithm used for analyzing SC3D-PTV images.

6.1 Particle Image Matching

When numerous particles are being measured at the same time, the particle images corresponding to a single physical particle must be matched between the simultaneous and successive images, as seen in Figure 49, prior to determining 3C3D velocity vectors.

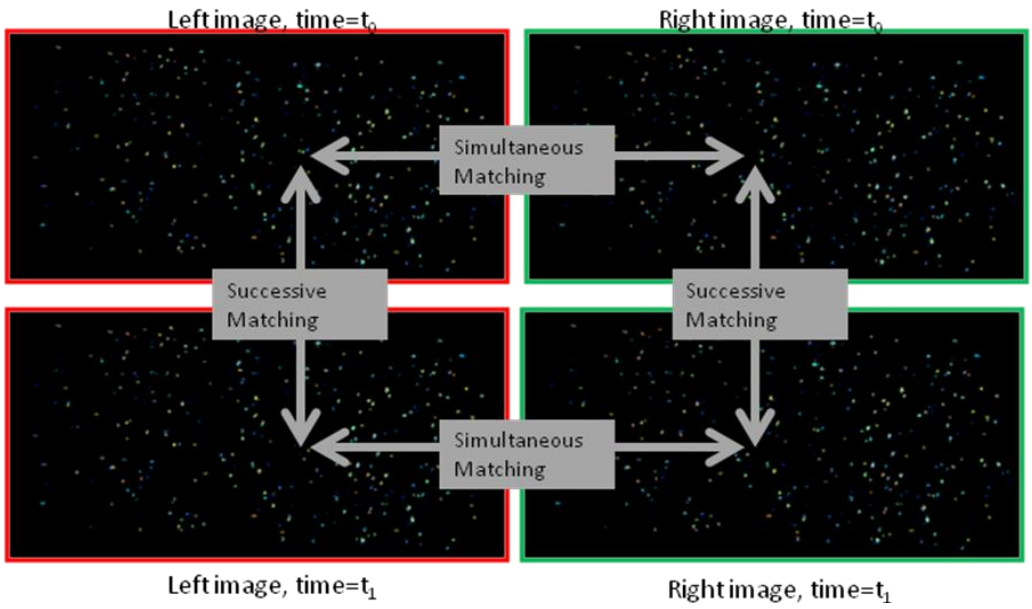

Figure 49: Particle images from simultaneous and successive images must be matched to a single, physical particle before 3C3D velocity vectors can be computed.

In the literature, many flow measurement techniques require some form of simultaneous and/or successive particle image matching. Simultaneous matching and successive matching are usually performed as two separate processes, where simultaneous matching is performed first to reconstruct the 3D particle field, and then successive reconstructions of the 3D particle field are analyzed to determine particle motion (e.g. Willert and Gharib, 1992; Maas et al., 1993; Elsinga et al., 2006). Some groups reverse this order, first determining two-dimensional particle paths, then matching the paths to determine the full 3D displacement of the particles (e.g. Herring et al., 1997; Nishimure et al., 2005).

6.1.1 Matching Simultaneously-Acquired Particle Images

In general, finding simultaneous particle image matches requires consideration of the particle image positions, and is based on the method of epi-polar lines (as described in Wieneke, 2008). Since the distance between the particle and the camera is not known, a single 2D image can only define a volume along the line-of sight of the camera in which the particle must be found. When using algebraic reconstruction (1976, Hermann and Lent, cited by 2006, Elsinga et al.), as in tomographic PIV, epi-polar lines are constructed for each simultaneously acquired image, from the position of the acquired image along the line-of-sight of each camera, and define a long, thin volume in which the particle must lie. Using multiple, simultaneously acquired images from different viewing angles (and multiple epi-polar lines) allows particle position to be determined through triangulation by finding the intersection of multiple epi-polar lines within the measurement volume, as seen in Figure 50 for a two-dimensional slice of the measurement volume.

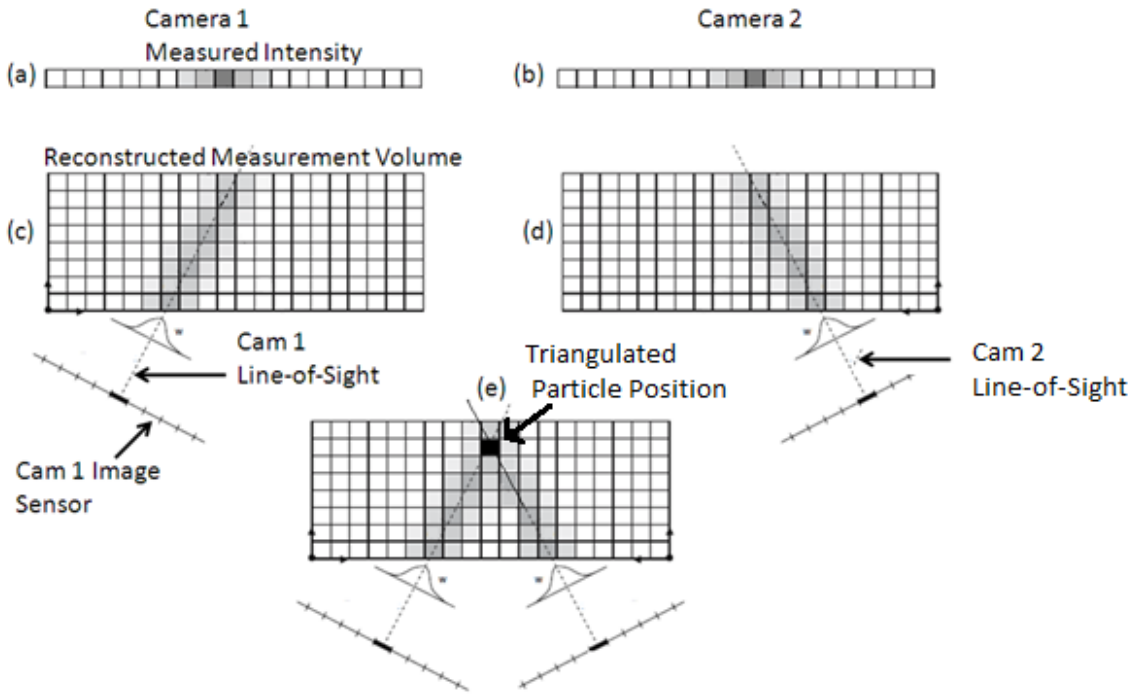

Figure 50: A 2D pictorial representation of the algebraic reconstruction algorithms used in tomographic PIV. (a) and (b) show one particle image within a single row from a camera image. A particle image captured by one camera indicates a region of the measurement volume in which the particle must reside, which extends along a line-of-sight from the camera, as shown in (c) and (d). Particle location is indicated by the intersection of volumes projected from different cameras, as indicated by the black grid point in (e). (images adapted from Elsinga et al., 2006)

However, rather than using epi-polar lines to reconstruct the measurement volume, similar considerations can be used to define a search area within the simultaneously-acquired images for possible matches, based on the position of a particle image within one of the cameras. The depth of the measurement volume, the relative angle between the cameras acquiring simultaneous images, and the position of a particle image in one of the cameras defines a subset of pixels in the other camera in which a simultaneous match must be found. A simplified illustration of this concept is shown in Figure 51.

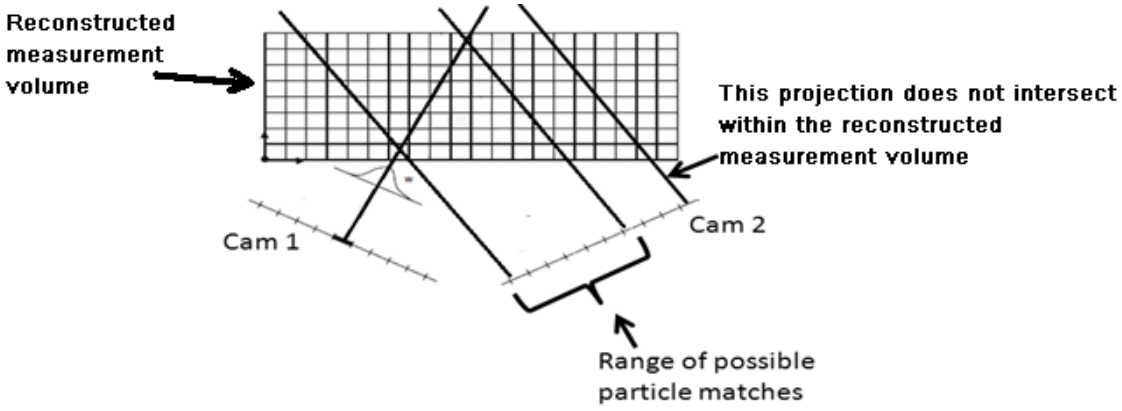

Figure 51: Since the measurement volume is of limited depth, for a given particle position in Camera 1, only a restricted number of the pixels in Camera 2 have the possibility of intersecting within the measurement volume. (image adapted from Elsinga et al., 2006)

Thus for every possible position of a particle image in Camera 1, the measurement volume dimensions and imaging parameters define a search area within Camera 2 where possible matches must reside (Wieneke, 2008).

Using epi-polar considerations to define a search area is easily implemented for the imaging arrangement used in SC3D-PTV. Since the two imaging sub-systems have co-planar optical axes, approximately horizontal slices of the measurement volume will be imaged onto the same row of pixels in both of the simultaneously acquired images. For clarity, this is shown in Figure 52 for two separate cameras, but also applies to the single camera set-up of SC3D-PTV. Under the constraint of co-planar optical axes, a particle in a given row in the image of Camera 1 must be found in approximately the same row of Camera 2.

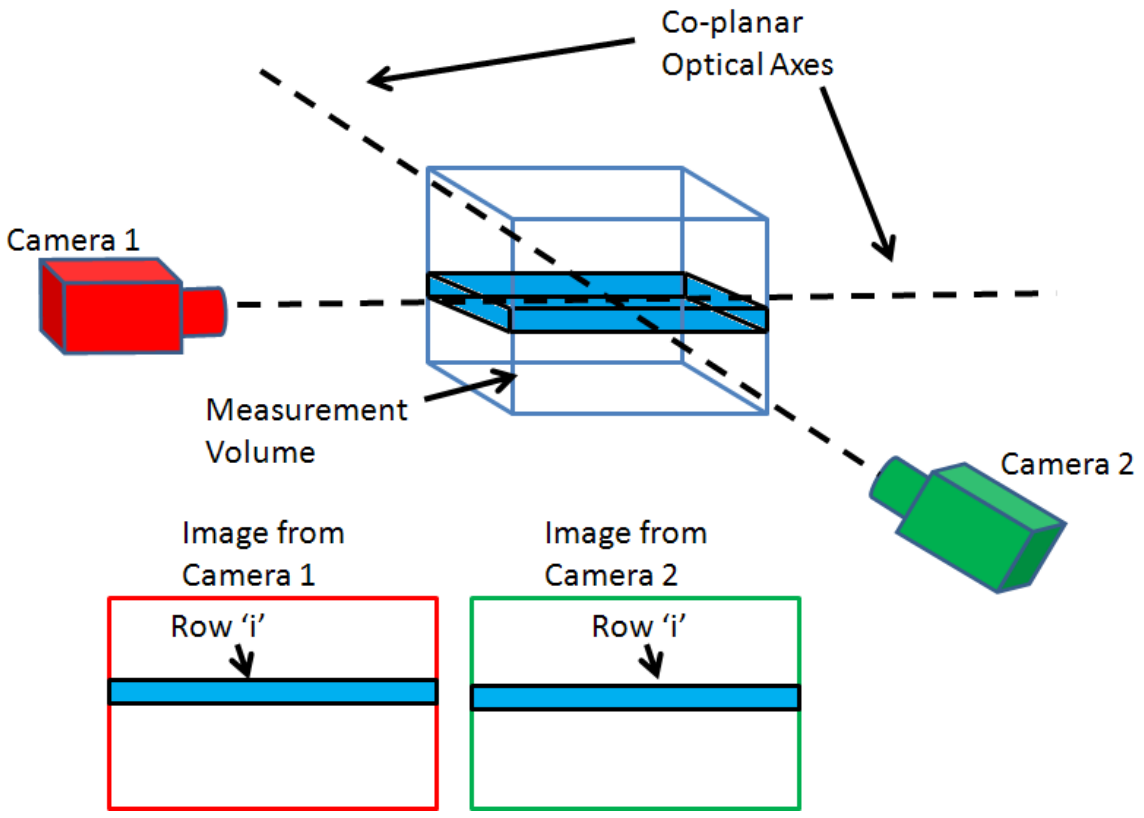

Figure 52: Two imaging systems with co-planar optical axes and identical magnification will image the same horizontal slice of the measurement volume onto the same row of pixels. SC3D-PTV uses a custom optic to create two imaging sub-systems with co-planar optical axes.

The number of columns comprising the search area for simultaneous matches can also be restricted using knowledge of the angle between the sub-imaging systems, since the measurement volume is of limited depth (as shown above in Figure 51). These two criteria mean that the search area for a simultaneous match in the SC3D-PTV data is a rectangle a few pixels in height and of limited width (where the width is a function of measurement volume depth and the relative angle between the imaging sub-systems), centered on the position of the original particle, as shown in Figure 53.

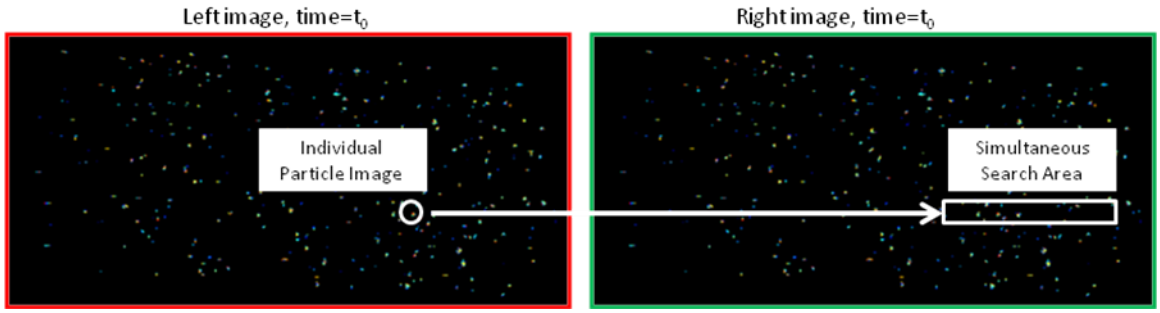

Figure 53: The position of a particle image defines a search area for the corresponding, simultaneously-acquired particle image based on the relative angle between the two viewing angles and the depth of the measurement volume.

When using two cameras, a method based on intersecting epi-polar lines can only tolerate low seed densities without ambiguity, so the accuracy of the simultaneous matching process is often increased by using more than two cameras (e.g. Maas et al., 1993, Elsinga et al., 2006).

Some techniques also use features of the particle image as a parameter in defining particle image matches. For example, in addition to epi-polar line considerations, Bao et al. (2011) require a similar degree of defocus of the particle images. Bown et al. (2006) supplement a search for intersecting epi-polar lines with a cross-correlation calculation to indicate the direction and magnitude of the relative displacement between likely simultaneous matches. In general, adding additional constraints to the method of intersecting epi-polar lines allows higher seed density to be tolerated without ambiguity.

The simultaneous particle image matching performed during SC3D-PTV analysis (detailed at the end of this chapter) is also based on the concept of search areas defined by epi-polar line constraints, with additional particle image similarity requirements. However, SC3D-PTV analysis performs both simultaneous and successive particle image matching at the same time.

6.1.2 Matching Successively-Acquired Particle Images

Like the search for simultaneous matches, the successive matching process also relies on a defined search area, but rather than being defined by epi-polar considerations, successive search areas are defined by the expected particle displacement. Generally, the time between successive image exposures is small, so the displacement of a particle is limited to a few (<10) pixels in image space. The search area can then be limited to a radius equal to the maximum expected pixel displacement, as seen in Figure 54 (although during analysis the search area is defined as a square for computational simplicity).

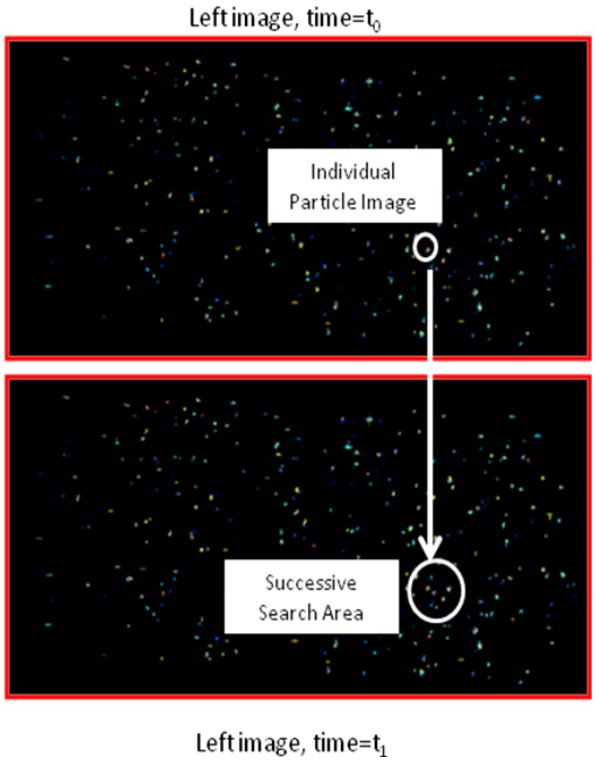

Figure 54: The search area for the corresponding, successively acquired particle image is defined by the position of the original particle image and the maximum expected particle displacement.

The simplest form of successive matching based on a search area is the nearest-neighbor algorithm. The nearest-neighbor algorithm assumes that particle displacement

is smaller than inter-particle spacing, thus the particle image in the second frame closest to the position of the particle image in the first frame is assumed to correspond to the same particle. This method is only accurate at low levels of seed density.

Higher seed densities are tolerated using the 3- or 4-frame method (e.g. Hassan and Canaan, 1991), which extrapolates all possible particle tracks based on two frames, then searches for a particle image within an area defined by the extrapolated tracks in the third or fourth frame. A smoothness or acceleration threshold on particle displacement is used to define the search area and eliminate unlikely particle paths.

A variety of other methods rely on the similar motion of groups of particles that are close in image space, such as the relaxation method (Ohmi and Li, 1999), the minimization of the velocity gradient (Ishikawa, 2000), or cross-correlation-based super-resolution PIV (Keane et al., 1995).

Tapia et al. (2006) extend the nearest-neighbors method by examining all neighbors within a defined radius, and then using the intensity and the shape individual particle images as additional criteria to aid in the successive matching process. Similarly, Mikheev and Zubtsov (2008) add size-similarity and intensity-similarity thresholds to a nearest-neighbors algorithm. These additional matching criteria allow higher seed densities to be tolerated.

The successive matching performed during SC3D-PTV analysis is based on the concept of a limited search area with additional particle image similarity constraints, but is performed at the same time as simultaneous particle image matching.

The next section describes SC3D-PTV analysis in detail.

6.2 SC3D-PTV Raw Image Analysis

Four raw images are processed to produce a single 3C3D vector field. Each particle image defines a set of search areas that are used to find possible simultaneous and successive particle image matches. The simultaneous and successive matching processes are then performed in tandem using constraints based on particle image similarity to identify valid matches.

6.2.1 Defining Particle Images

The raw images are initially scanned using a two-stage threshold to define particle images. Each particle image must contain at least one pixel above the upper threshold, and any 4-connected pixels that are above the lower threshold are also considered part of that particle image. An example particle image (magnified and shown in false color) on a black background is shown in Figure 55, and each colored square represents one pixel of the particle image. The typical size of a particle image is on the order of ones to low tens of pixels.

Figure 55: An example of a single particle image on a black background. The particle image intensity is shown in false color, and each colored square represents one pixel.

6.2.2 Defining Search Areas

Every particle image in each of the four raw images defines two search areas: a simultaneous search area and a successive search area. All possible simultaneous and successive matches are found within these search areas. For example, choosing a particle image from the L_0 image (the image acquired from the left viewing angle at time t_0) defines search areas in the R_0 and L_1 images. Each possible match found within the R_0 and L_1 search areas define another search area within the R_1 image. In general, every particle image in the L_0 image defines a set of search areas like those seen in Figure 56.

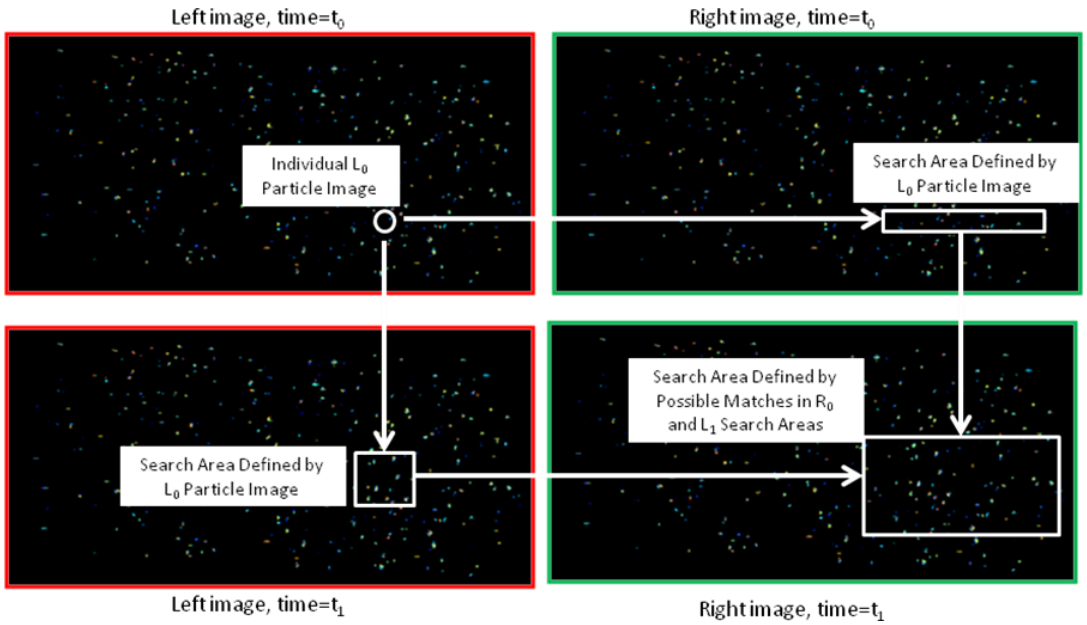

Figure 56: A particle image in the L_0 image defines search areas in the R_0 and L_1 images where possible matches can be found. Those possible matches then define a search area in the R_1 image.

6.2.3 Finding Complete Sets

To compute a vector, four particle images (one from each of the four raw images) must be determined to correspond to one physical particle. Using this constraint, all

possible simultaneous and successive matches (as defined by search area constraints) are searched to find 'complete sets.' A complete set is a group of two simultaneous and two successive matches that contains only four total particle images, as seen in Figure 57.

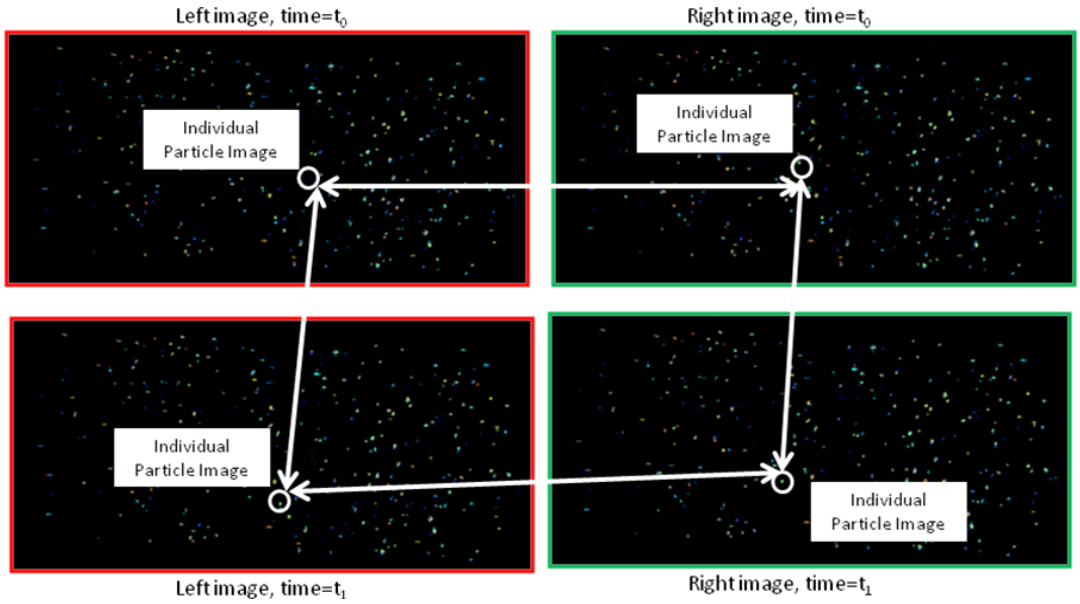

Figure 57: A "complete set" contains four particle images, one from each raw image, and contains two possibly-valid simultaneous matches and two possibly-valid successive matches.

It is likely that each particle image will be included in more than one complete set, as in Figure 58, but in the absence of overlapping particle images, each particle image can be included in only one valid complete set (i.e. a complete set that leads to a valid vector). Conversely, it is unlikely that all sets containing unique particles will be valid, due to overlapping particle images or edge effects.

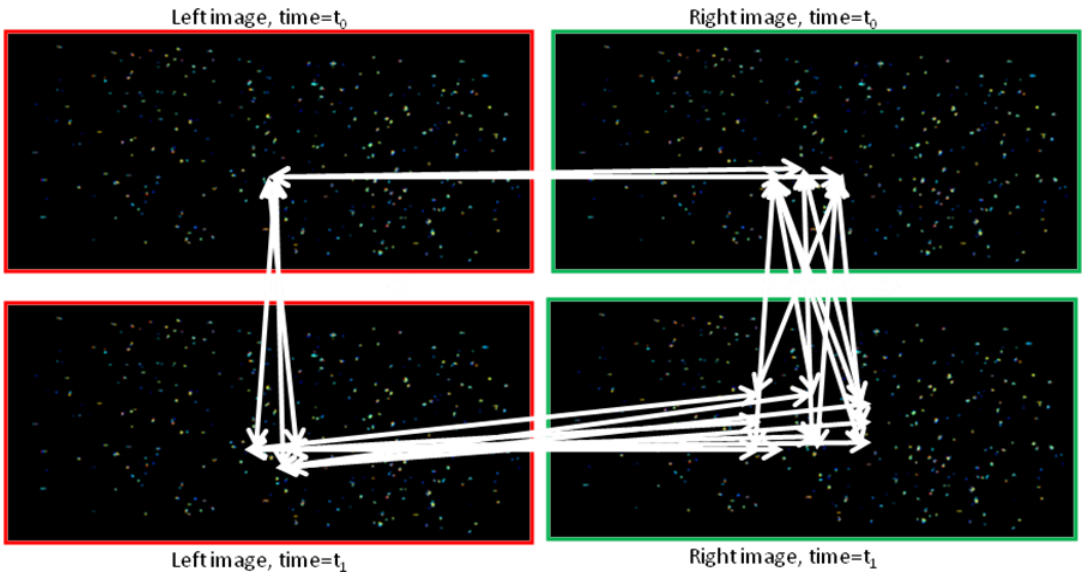

Figure 58: Each particle image is likely to be included in more than one complete set. This image shows a particle image in the L_0 raw image included in numerous complete sets.

After all sets have been found, valid and invalid sets are separated by quantifying the similarity of the particle images within the set using shape, intensity, and size information.

6.2.4 Particle Image Similarity

The similarity-based matching algorithms used in SC3D-PTV were developed using several ideal assumptions, including: perfectly spherical particles, a homogeneous light-sheet, no obfuscation by other particles, perfect alignment of optical elements, no camera noise, and no optical aberrations. Since the two imaging sub-systems share a focal plane without perspective distortion (2006, Lindken et al.) and the Mie scattering angle is approximately equal, under these assumptions the two simultaneous images of a particle will be nearly identical. Also, the particle displacement between successive image exposures is small compared to the size of the measurement volume, so the

imaging parameters will be similar for the first and second set of images, leading to similar particles images.

However, the ideal assumptions stated above do not hold rigorously, so the four particle images corresponding to a single particle will not be completely identical. To account for non-idealities, the SC3D-PTV matching algorithm does not require that the particle images be exactly identical to be considered a valid set, only that the images corresponding to one particle meet a lower threshold on similarity, and be more similar to each other than the images of other particles within the search area. Quantification of the similarity of the particle images is described in the next chapter.

After valid sets have been identified and invalid sets have been discarded, vector fields are computed using the geometric relations based on relative and absolute particle image positions given previously.

6.3 Summary

The analysis of SC3D-PTV data is based on ideas from the literature (i.e. search areas, feature matching), but also incorporates a process of finding complete sets using both simultaneous and successive search areas. Using the search area to find complete sets is equivalent to performing the simultaneous and successive matching processes at the same time. Of the possible sets, those with a high probability of being valid are found by comparing a range of particle image features for similarity. The development of the similarity-checking algorithm and the determination of appropriate threshold levels for particle image similarity are the subjects of the next chapter.

CHAPTER 7

ANALYSIS OF SPARSELY SEEDED IMAGES

As described in the previous chapter, although the relative positions of the particle images indicate possible matches, matching based on position alone can only tolerate low seed densities. Particle image similarity comparisons are used to increase the allowable particle image density during SC3D-PTV analysis.

Under ideal assumptions, particle images corresponding to a single, physical particle will be identical, but a range of real-world non-idealities contribute to their lack of similarity (such as lens aberrations, non-spherical particles, and camera noise). Since these non-idealities are difficult to quantify analytically, they are better quantified through experimentation. This chapter explains the procedure of determining appropriate thresholds for an algorithm that can separate valid sets from invalid sets based on the features of the particle images.

7.1 Defining Particle Image Similarity

To determine the appropriate thresholds, the features of a large set of likely-valid two-image matches were compared to the features of a large set of invalid matches. A series of sparsely-seeded data were acquired to provide complete four-image sets that could be unambiguously verified using position assumptions alone, without using particle

image similarity considerations. Because the images were sparsely seeded, most particle images were members of only a single set. All sets whose members were not part of any other set were considered valid regardless of the similarity of the particle images. The data were collected using an air jet seeded with silicon oil droplets (~1 μm diameter) over a measurement volume of 6mm x 6mm x 2mm. Further details of the experimental configuration can be found in the Chapter 9.

Although not all of the sets found using only position assumptions were valid (likely due to edge effects as particles leave or enter the measurement volume) the great majority (>95%) were found to be valid when the computed velocities were compared to the expected flow properties. An example of a sparsely seeded image set is shown in Figure 59. Only a single particle image resides within the search areas defined by the position of the original particle image.

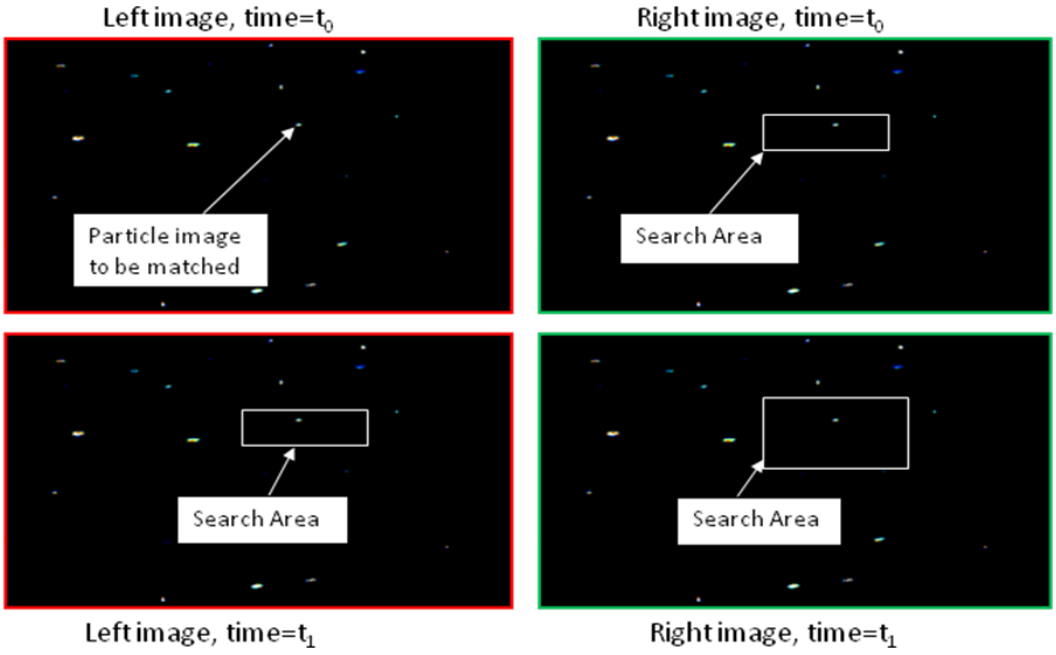

Figure 59: Four sparsely-seeded raw images. The position of a particle image in the top left raw image defines search areas in the other three raw images. For the analysis presented in this chapter, a complete set is considered valid if each search area contains only a single particle images, regardless of the similarity of the particle images.

A control group was formed by randomly matching particle images from non-consecutive raw images (so that the particle images were guaranteed to be mismatched).

This sparse set analysis assumes that the similarity of the particle images will not change significantly for the higher seed density cases that will ultimately be studied. However, at higher seed densities the similarity of the particle images is likely to change, at least minimally, because a higher seed density increases the background noise caused by out-of-focus particles, and decreases the particle image quality since a dense cloud of particles is located between the particle being imaged and the image sensor. Since no *a priori* information was available about which particle images truly correspond to a single particle, this sparse set analysis was used during algorithm development while noting that the similarity thresholds derived from it may change slightly when examining densely seeded cases.

For both the valid and invalid sparsely seeded data, the four particle images comprising each complete set were analyzed for similarity using seven criteria:

1) Total Intensity - the sum of all intensity values within the particle image
2) Maximum Intensity - the peak intensity value within the particle image
3) Total Size - the total number of pixels in the particle image
4) Y-Size - the height of the particle image, as measured in pixels
5) X-Size - the width of the particle image, as measured in pixels
6) Intensity Pattern - the maximum value of a normalized cross-correlation calculation between the zero-padded raw particle images
7) Shape - the maximum value of a normalized cross-correlation between the zero-padded, binarized particle images

During the sparse data analysis, the goal was to define appropriate thresholds for similarity such that a maximum number of valid sets, and a minimum number of invalid sets, would be retained.

7.2 Finding Particle Image Similarity Thresholds

Although the thresholds change slightly for different experimental arrangements and the simultaneous or successive comparisons, the general trends are quite similar. To illustrate typical results, the statistics of the simultaneous matches of one sparsely seeded data set are shown below (containing 8,000 valid matches, and 10,000 invalid matches). The results are presented in the form of cumulative distribution functions, where the value of the feature being analyzed is shown on the x-axis, and the percentage of particle image matches at, or below, a given feature level is shown y-axis.

Of the seven features examined, no single feature uniquely separates the valid matches from the invalid matches. For the difference in summed intensity, as seen below in Figure 60, at a threshold level that contains 90% of the valid matches, approximately 70% of the invalid matches are also included.

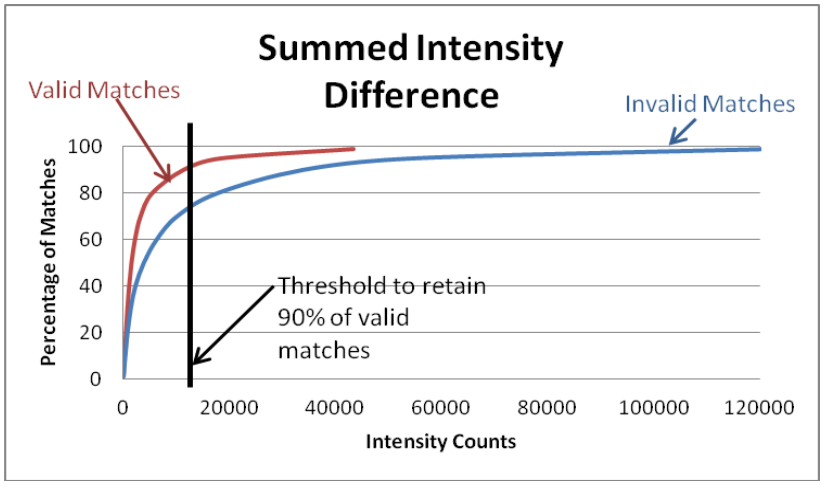

Figure 60: Comparing the summed intensity difference of valid and invalid matches, at a threshold level that keeps 90% of valid matches, 70% of invalid matches are also retained.

The difference in peak intensity between valid and invalid matches shows a similar result. As seen in Figure 61, at a threshold level that contains 90% of the valid matches, approximately 70% of the invalid matches are also included.

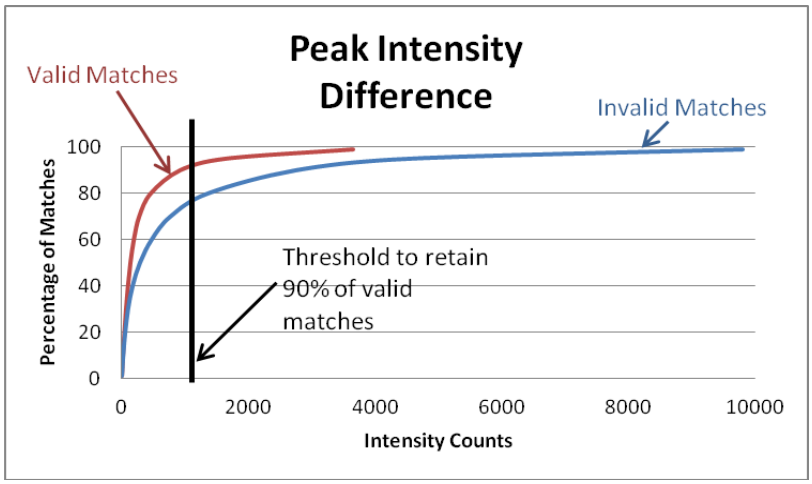

Figure 61: Comparing the difference in peak intensity of valid and invalid matches, at a threshold level that retains 90% of valid matches, 70% of invalid matches are also included.

Particle size better separates valid and invalid matches, but as seen in Figure 62, at a threshold level that contains 90% of the valid matches, greater than 50% of the invalid matches are still included.

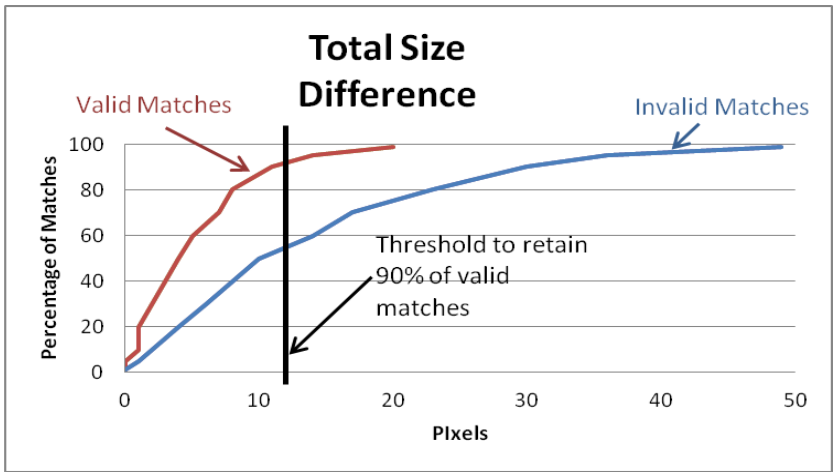

Figure 62: The difference in total particle image size, as measured in pixels, separates valid and invalid matches better than intensity comparisons. However, at a threshold level that includes 90% of valid matches, more than 50% of invalid matches are still kept.

Separating particle size into measures of both height and width reveals that particle width is a much better indicator of match validity than particle height. This is largely due to distortion caused by the optical elements that affects particle height much more than width. As seen in Figure 63, a sample simultaneously-acquired image pair demonstrates this distortion.

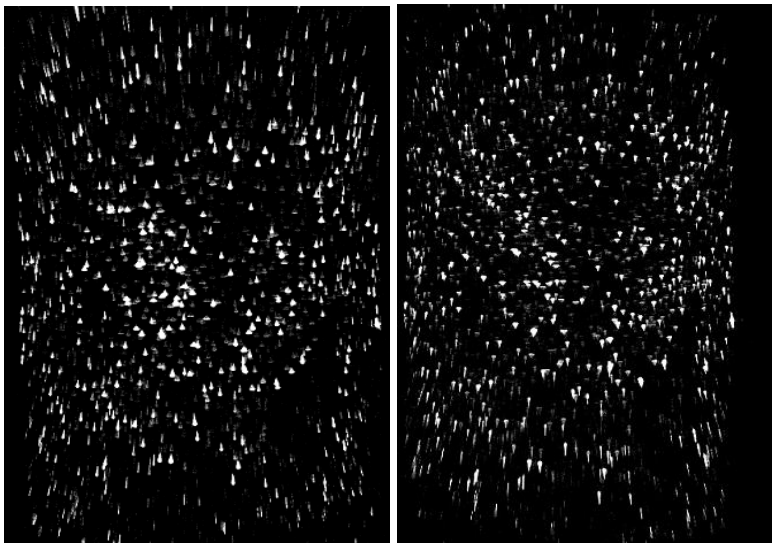

Figure 63: A simultaneously-acquired image pair showing distortion. Image distortion in the vertical direction is especially prevalent at the upper and lower edges of the images.

This difference in the degree of distortion can be seen in comparing the difference in height and width of the particle images. As seen in Figure 64, for the difference in particle height, a threshold level at which 90% of valid matches are retained also includes 60% of invalid matches.

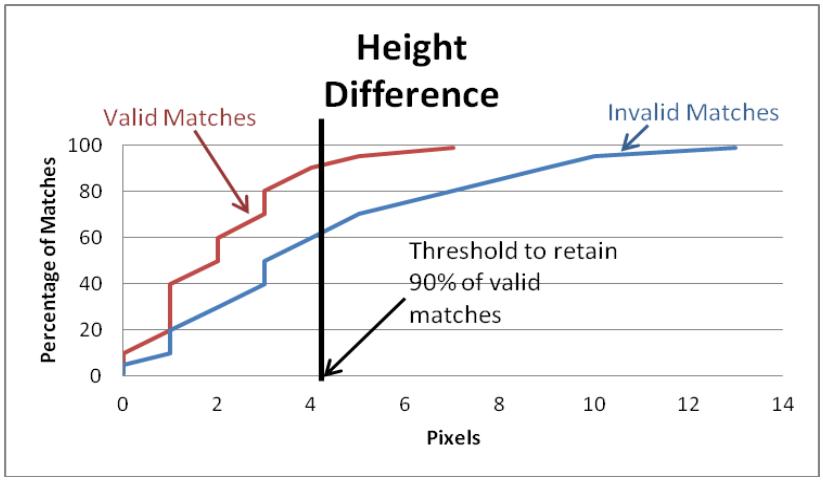

Figure 64: When examining the difference in particle height, at a threshold level that retains 90% of valid matches, 60% of invalid matches are also kept.

However, for the difference in particle width, a threshold level that retains 90% of valid particle matches retains only 30% of invalid matches.

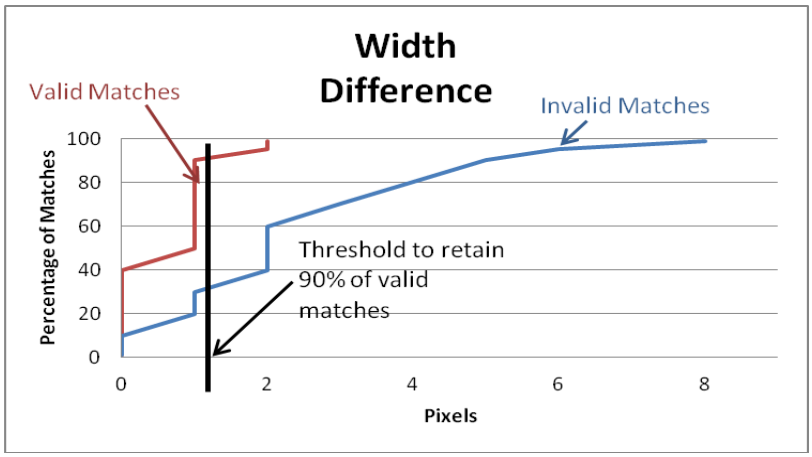

Figure 65: Difference in particle width is a better indicator of validity than particle height (Figure 64). When considering particle height, at a threshold level that includes 90% of valid matches, only 30% of invalid matches are retained.

The similarity in shape of the particle images, as quantified by performing a normalized cross-correlation calculation on a binarized version of the particle image, retains 40% of invalid matches at a threshold level that keeps 90% of valid matches, as seen in Figure 66.

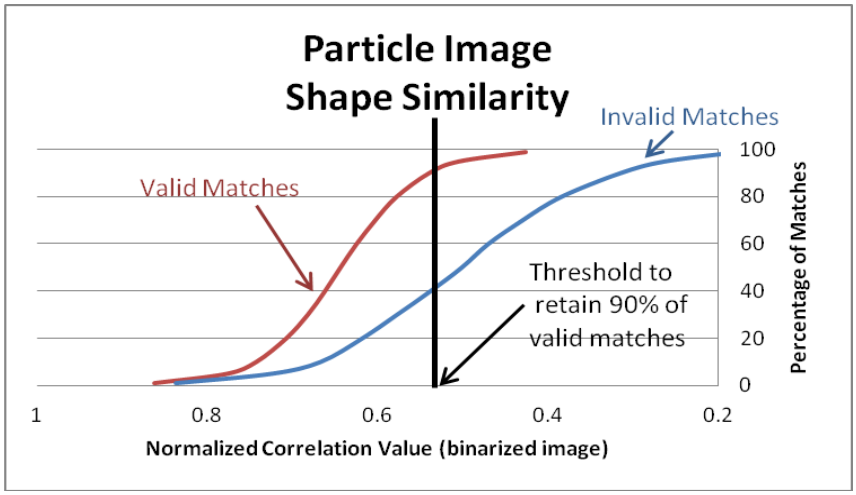

Figure 66: For particle image shape similarity, at a threshold level that retains 90% of valid matches, 40% of invalid matches are also retained.

The particle image intensity pattern similarity, as quantified by performing a normalized cross-correlation calculation on a the raw version of the particle image, retains approximately 50% of invalid matches at a threshold level that keeps 90% of valid matches, as seen in Figure 67.

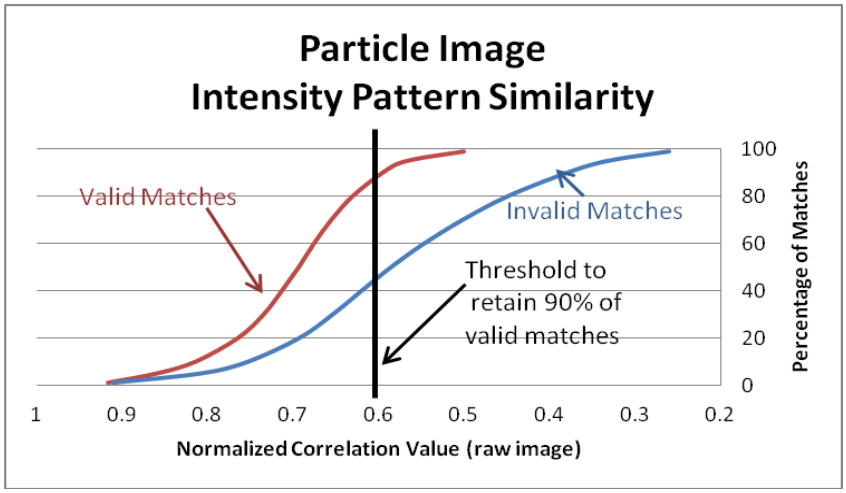

Figure 67: For particle image intensity pattern similarity, at a threshold level that includes 90% of valid matches, approximately 50% of invalid matches are also kept.

As seen in the preceding figures, by examining a single feature and retaining 90% of valid matches, at least 30% of invalid matches are also kept. To increase the number

of valid matches that are retained, compared to the number of invalid matches, multiple features must be examined simultaneously.

7.3 Separating Valid and Invalid Sets

For the data in Figure 68, the value at which 90% of valid vectors are retained is used as a threshold on validity for multiple features simultaneously. As more features are examined, the ratio of valid vectors retained to invalid vectors retained increases. When all seven features are examined, the ratio of valid matches to invalid matches (as measured on a percentage basis) is greater than 7:1.

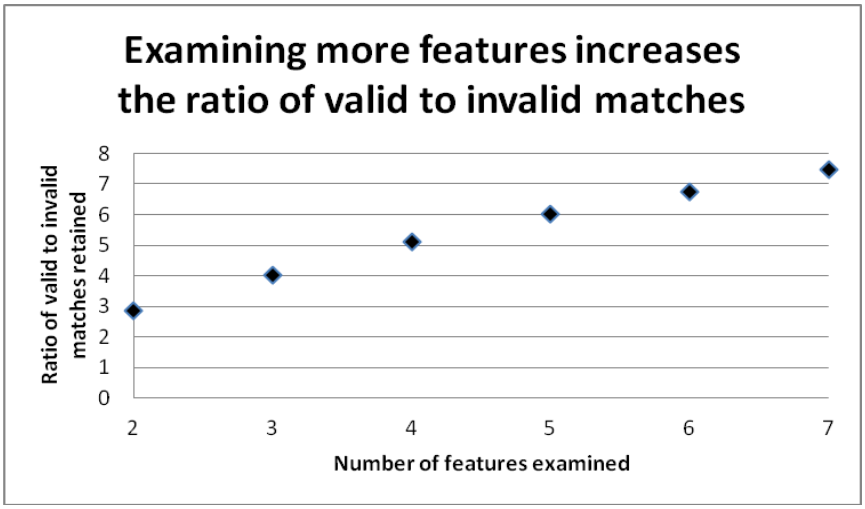

Figure 68: The ratio of valid to invalid matches increases (as calculated on a percentage basis) as more features are examined simultaneously.

Using the same data as Figure 68, Figure 69 shows that when all seven features are examined, almost 70% of the valid matches still remain, but less than 10% of the invalid matches are kept.

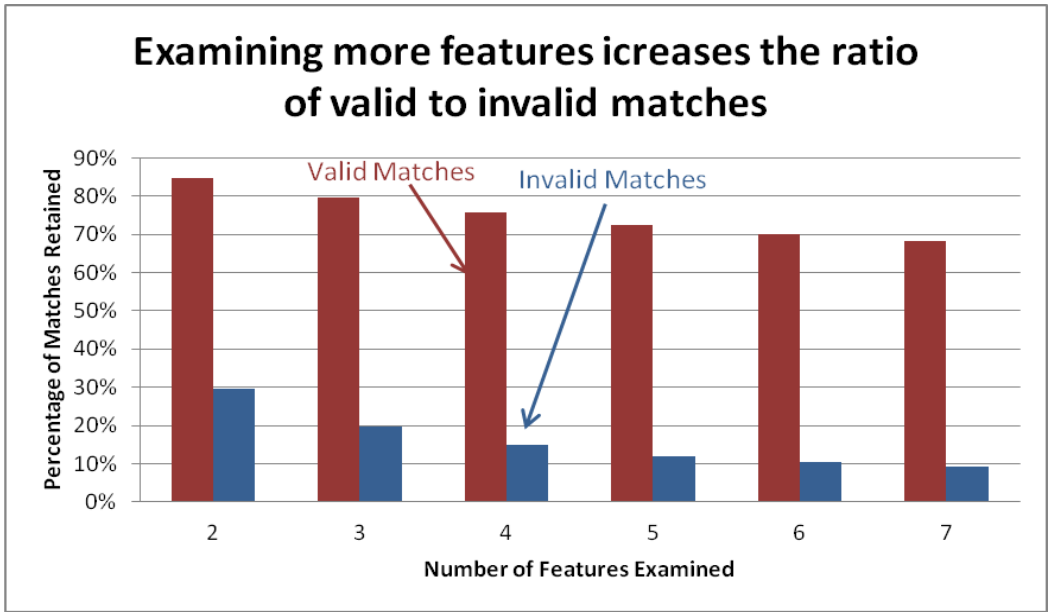

Figure 69: A threshold at which 90% of valid matches are retained is applied to multiple features simultaneously. As the number of examined features increases, the percentage of valid vectors that are retained decreases at a slower rate than the percentage of invalid matches.

The data presented above show the likelihood that an individual match (comprised of two particle images) will exceed a certain threshold, but each set is composed of four matches (two simultaneous and two successive). Thus the ratio of valid to invalid four-image sets retained using this comparison scheme is even greater than the ratio of valid to invalid two-image matches that are retained, because valid sets are composed of four valid matches (most of which have a high degree of similarity), but it is less likely that four highly similar invalid matches will be found in the same set. The likelihood of an invalid set meeting the criteria to be considered valid is explored more fully in Chapter 8.

Besides the particle image similarity considerations discussed above, one additional constraint must be met for a set to be considered valid:

1) The x-displacement of the particle image of both successive matches must be approximately equal.

This additional constraint is made because, due to the set-up of the imaging optics, the x-displacement of the particle is projected similarly onto both successively acquired images (as shown previously in Figure 52), and should therefore be equal in the images. The same assumption cannot be made about the y-displacement of the image.

Because seven features per match are examined, and there are four matches in a set, a total of 28 feature comparisons are made for each set. It is possible that a valid set may be above the threshold for most features, but not all, so rather than requiring that each valid set exceed the threshold for each feature, a tolerance is set on the number of features that must meet the 90% threshold, allowing a greater number of valid sets to be retained.

Figure 70 shows the results of combining the x-displacement similarity requirement with the particle image similarity comparisons, including a tolerance for particle image similarity. If all 28 features are required to be above the threshold, and the x-displacement must be similar, the ratio of valid to invalid sets is very high (35% of valid sets, but only 1% of invalid sets), but because only 35% of valid sets are retained the spatial resolution of the measurement suffers (since each valid set is equivalent to a measurement at one point in space). For comparison, if sets that have 26 features above the 90% threshold are retained, nearly 70% of the valid sets are kept, but only about 4% of the invalid sets. Although the percentage ratio of valid to invalid sets is worse, the

overall number of vectors is a factor of two larger, doubling the spatial resolution of the measurement while keeping a reasonable signal-to-noise ratio.

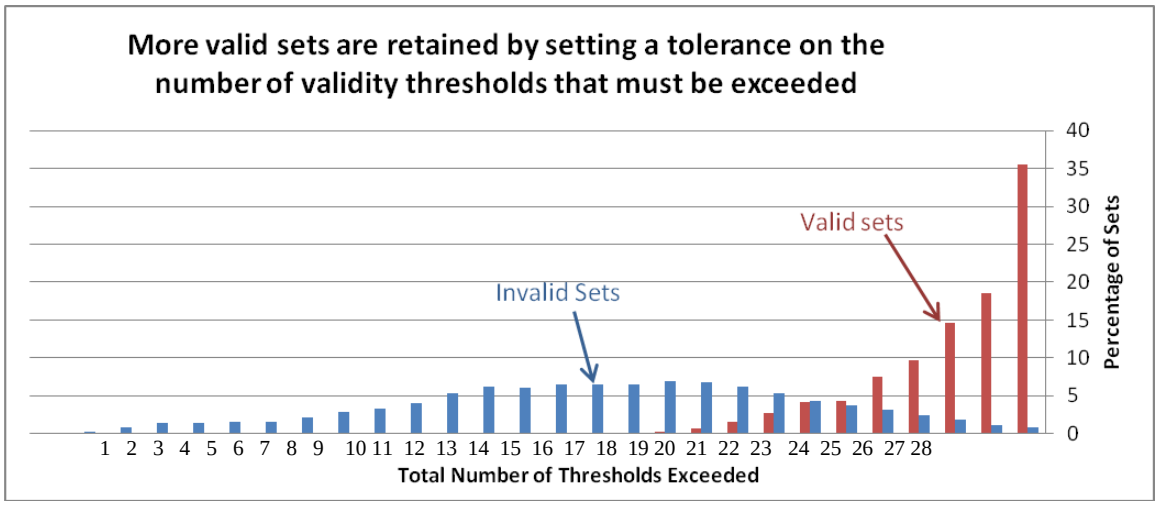

Figure 70: A histogram showing the percentage of valid and invalid sets that exceed a given number of 90% thresholds. By setting a tolerance on the number of features that must exceed the 90% threshold, significantly more valid sets can be retained, while keeping a reasonable ratio of valid to invalid sets.

7.4 Summary

The analysis of sparsely-seeded images presented in this chapter describes the SC3D-PTV algorithm in detail, and provides some statistical basis. This chapter is concisely summarized by Figure 70, which shows that grouping particle images into sets (using search areas) and applying similarity comparisons allows valid and invalid sets from sparsely-seeded images to be largely separated (but not entirely), at least on a percentage basis, while still retaining a large proportion of the valid sets. However, for the percentage-based values to have meaning, the absolute number of valid and invalid sets must be known. Also, because the data properties are expected to change somewhat for densely seeded images, and densely seeded images are of greater interest (since they

correspond to measurements with higher spatial resolution), the next chapter examines how the algorithms developed here behave when applied to densely seeded images and gives estimates of the numbers of valid and invalid sets for various experimental conditions.

CHAPTER 8

ANALYSIS OF DENSELY SEEDED IMAGES

The previous chapter described the algorithm underlying SC3D-PTV image analysis, and provided a statistical basis for that algorithm from an examination of sparsely-seeded data. This chapter analyzes densely-seeded data, in particular addressing the issues that arise as the particle density increases to the point that each particle image is a member of several possibly-valid four-image sets. As in the previous chapter, the analysis of a single data set is shown to illustrate typical results.

8.1 Advantages of a Set-Based Algorithm

Although the image similarity algorithms used during SC3D-PTV analysis could be applied to matching particle images between simultaneously-acquired and successively-acquired images separately, the algorithm returns more four-image sets of highly similar particle images by performing both matching processes in tandem. This is more apparent when analyzing densely seeded data rather than sparsely seeded data.

For comparison, the same data was analyzed using both a four-image set-based analysis (as described in Chapter 6), and a two-image match-based approach (where each particle is linked to its most similar simultaneous or successive match, as determined by the number of features above the 90% threshold, and the matches are then combined). As

seen in Figure 71, although the trend is the same for both the set-based approach and the match-based approach (the number of sets found increases with increasing particle image similarity), the set-based approach returns more complete sets with a high degree of similarity.

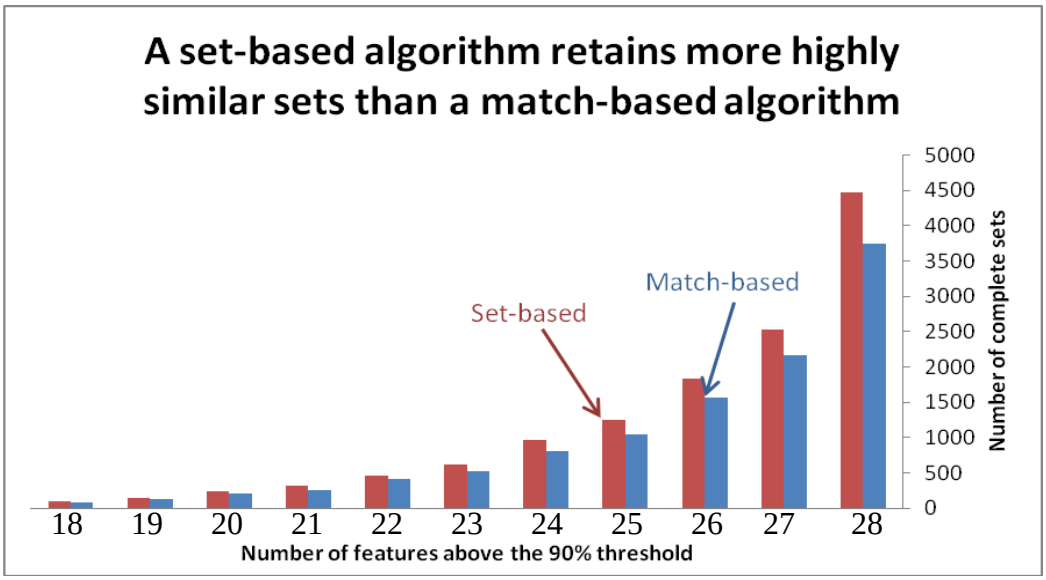

Figure 71: A series of densely-seeded images were analyzed using both a four-particle-image set-based approach and a two-particle-image match-based approach. Although both methods return an increasing number of vectors as more features are above the 90% threshold (using the same thresholds for both methods), the set-based approach returns a greater number of vectors.

If performing the simultaneous and successive matching separately on densely seeded images, a group of four matches may contain some, but not all of the same particle images, as seen in Figure 72. A valid vector cannot be computed from such a combination of matches, and instead the matches are discarded. Using a set-based analysis avoids this problem, and more highly similar sets are retained.

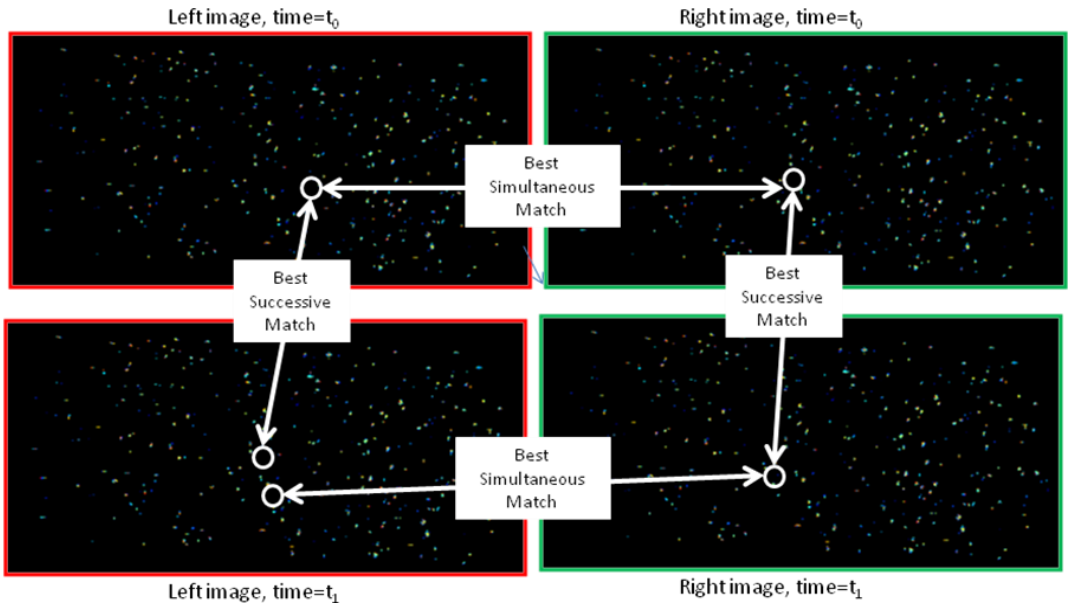

Figure 72: If the simultaneous matching and successive matching are performed separately, and four matches (two simultaneous and two successive) contain more than four particle images, a valid vector cannot be computed. By performing the SC3D-PTV analysis on sets consisting of four particle images, rather than matches consisting of two particle images, this problem is avoided.

For experimental data, it is desirable to maximize the number of computed vectors to maximize the spatial resolution of the results. As shown above, more vectors are returned using a set-based analysis than a match-based analysis, but another way to maximize the number of computed vectors is to image as many particles as possible. The maximum density of particles that can be tolerated by the SC3D-PTV algorithms is limited primarily by two factors. The first factor is the number of overlapping particle images, which complicates particle image matching and determination of particle image features. The second factor is that increasing particle density means an increasing number of invalid sets per particle.

8.2 Relative Number of Valid and Invalid Sets

8.2.1 Simulation Data

The previous chapter showed the likelihood that the set-based particle image similarity comparisons could separate valid and invalid sets on a percentage basis. However, to understand the significance of the percentage-based results, the relative numbers of valid and invalid sets must be known. There can be at most one valid set per particle image, but the number of invalid sets can vary significantly depending on experimental parameters, particularly seed density and search area size. To understand this behavior, the number of invalid sets per particle image was evaluated in simulation. The SC3D-PTV set-finding algorithms were applied to groups of 1000 images of randomly-placed synthetic particle images, and particle density, measurement volume depth, and search area size were varied. For experimental data, most particle images can be expected to be in at least one set (the valid set corresponding to that particle) but for a valid set the particle image placement is not random, and therefore will not be reflected in the simulation. Thus any sets that were found were assumed to represent invalid sets. These simulations do not account for all real world effects, like overlapping particle images and variations in laser profile, but do give estimates of the expected number of invalid sets for various experimental conditions.

During simulation, seed density was varied from approximately the level of a high seed density measurement technique, tomographic PIV (0.05 particles per pixel) down to the level of a sparsely-seeded PTV measurement (0.001 particles per pixel). Typical PTV measurements have a seed density of about 0.005 particles per pixel (Kinzel, 2010).

The simulation results are shown in Figure 73 and Figure 74 (Figure 74 is an enlarged view of the smallest seed densities found in Figure 73). The number of invalid sets per particle decreases with decreasing seed density, and at a seed density of 0.001 particles per pixel, the number of expected invalid sets per particle is much less than one.

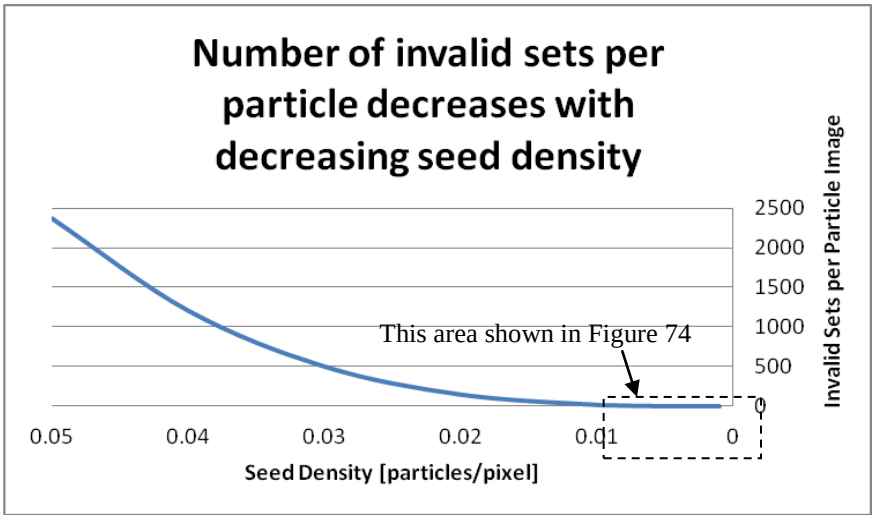

Figure 73: Monte Carlo simulations show that the number of random sets per particle decreases as a function of seed density, from more than 2000 random sets per particle at a seed density of 0.05 particles per pixel, to nearly zero at a seed density of 0.001.

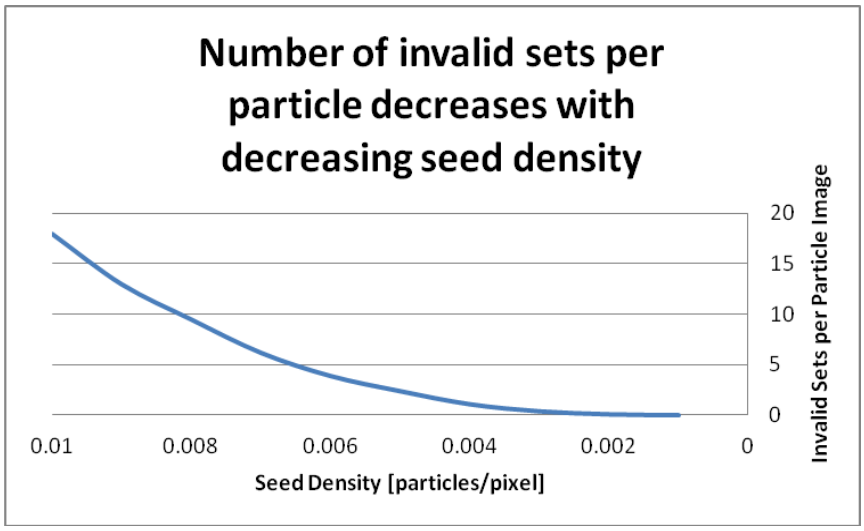

Figure 74: At the seed density level commonly found in PTV measurements (0.005 particles per pixel), the average number of random sets per particle is 2.4. At a seed density of 0.001, the average number of random sets per particle found by the simulation was only 0.01.

The number of invalid sets per particle image is also a function of the size of the simultaneous and successive search areas. The size of the simultaneous search area is a function of the depth of the measurement volume, and the size of the successive search area is a function of the maximum expected particle displacement (as measured in pixels). Figure 75 shows the simulated number of invalid sets per particle image as a function of measurement volume depth for a fixed successive search area size at a seed density of 0.005 particles per pixel.

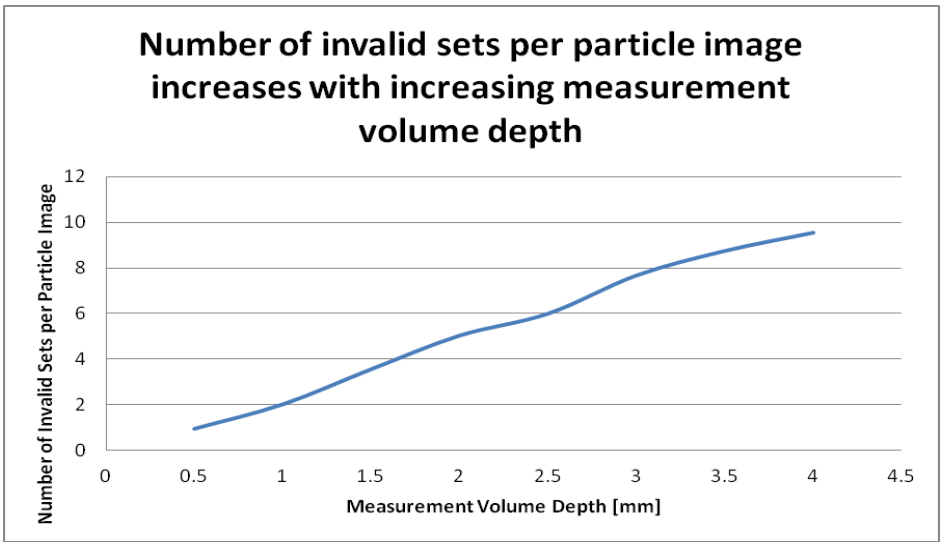

Figure 75: At a fixed seed density, the number of invalid sets per particle image increases with increasing measurement volume depth.

Figure 76 shows the expected number of sets per particle image as a function of maximum expected particle displacement at a fixed seed density of 0.005 particles per pixel for a measurement volume with a depth of one millimeter.

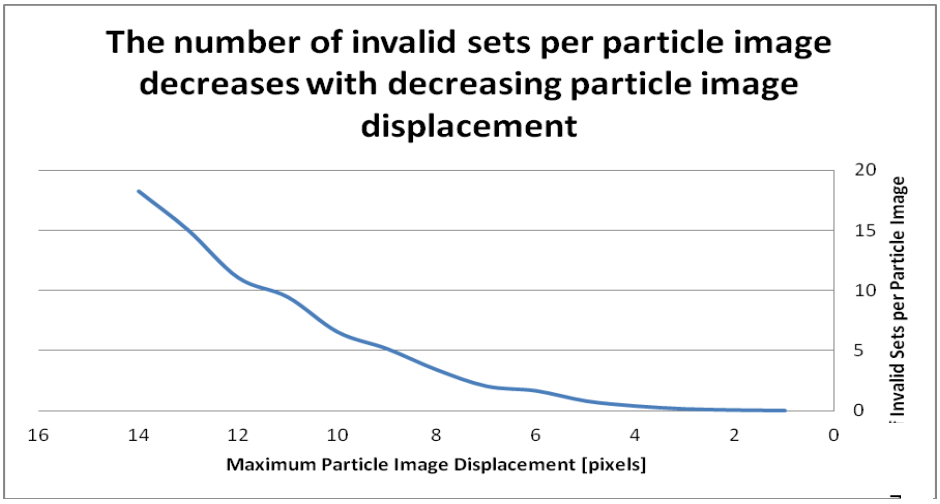

Figure 76: At a fixed seed density (0.005 particles per pixel shown above), the number of sets per particle image decreases with decreasing expected particle displacement. The measurement volume has a depth of one millimeter.

For simulated data, the approximate expected ratio of valid sets above the similarity threshold to invalid sets above the similarity threshold can be computed using the data presented at the end of the last chapter (at a given threshold value, 70% of valid sets are above the threshold, but only 4% of invalid sets) and the expected number of invalid sets per particle, as given above. For example, at a seed density of 0.005 particles per pixel (typical PTV density), a measurement volume depth of one millimeter, a maximum particle displacement of seven pixels, and assuming each particle image is a member of one valid set, the average particle will be a member of

$$1 \text{ valid set} + 2.4 \text{ invalid sets} = 3.4 \text{ sets} \quad (1)$$

There is a

$$70\% \times 1 \text{ valid set} = 70\% \quad (2)$$

chance that the valid set will exceed the validity thresholds, and a

$$4\% \times 2.4 \text{ invalid sets} = 9.6\% \quad (3)$$

chance that an invalid set will exceed the validity thresholds, meaning the valid set is more than seven times as likely to be chosen compared to the invalid sets.

8.2.2 Experimental Data

In addition to the simulations, the number of sets per particle was evaluated for 100 experimental images with approximately 250 particles spread over each 400x600 pixel image. The data is taken from the experiments described in Chapter 9. Although the global seed density was ~0.001 particles per pixel, the particles were distributed non-uniformly, as seen in Figure 77.

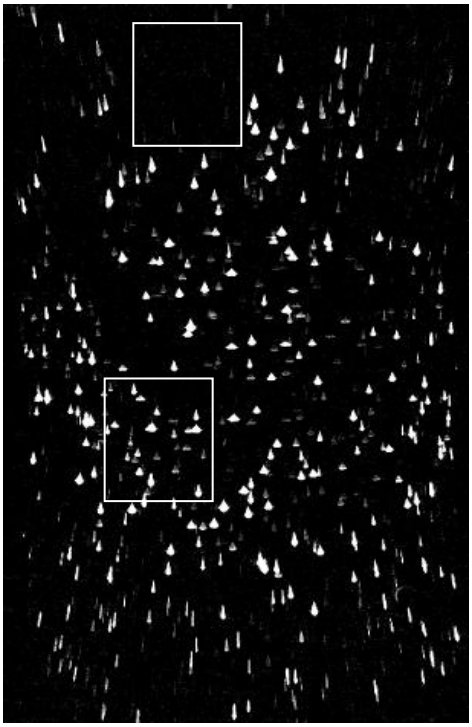

Figure 77: A raw experimental image shows that particles are distributed non-uniformly, resulting in high seed densities in some parts of the image, and low seed densities in other parts. The two rectangles within the figure indicate regions of vastly different seed density.

Figure 78 shows a larger view of the rectangles highlighted in Figure 77, and demonstrates the large degree to which seed density can vary within a single experimental image.

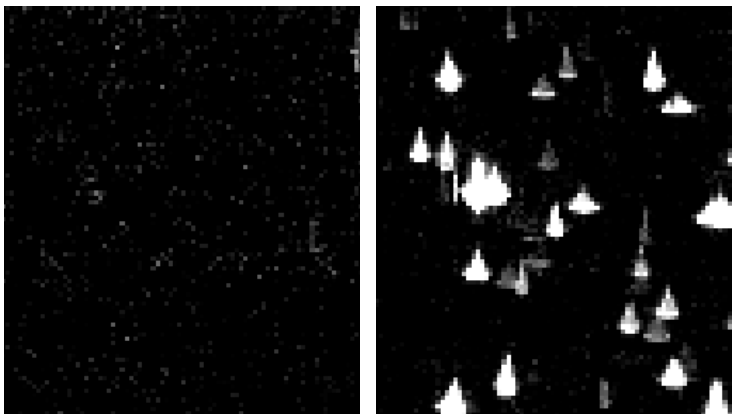

Figure 78: A higher magnification of the highlighted rectangles (~100x100 pixels) from Figure 77 shows that in the some areas the seed density approaches zero (left image, only background noise is visible), but in other areas the seed density is about 0.04 particles per pixel (right image).

When the number of sets per particle is displayed in a histogram, as in Figure 79 (based on 24,000 particle images), several features stand out. The first is that the peak of the histogram is at one set per particle image (presumably corresponding to the single correct set in most cases). This coincides with the fact that, as shown above in simulation, for a seed density of 0.001 particles per pixel, the number of random sets per particle is much less than one. But because the local seed density can vary significantly within the image, some areas have a much higher seed density than others. The particle images within those densely seeded areas will have a much higher number of sets per particle image, which is reflected in the long-tailed distribution of sets per particle seen in the histogram.

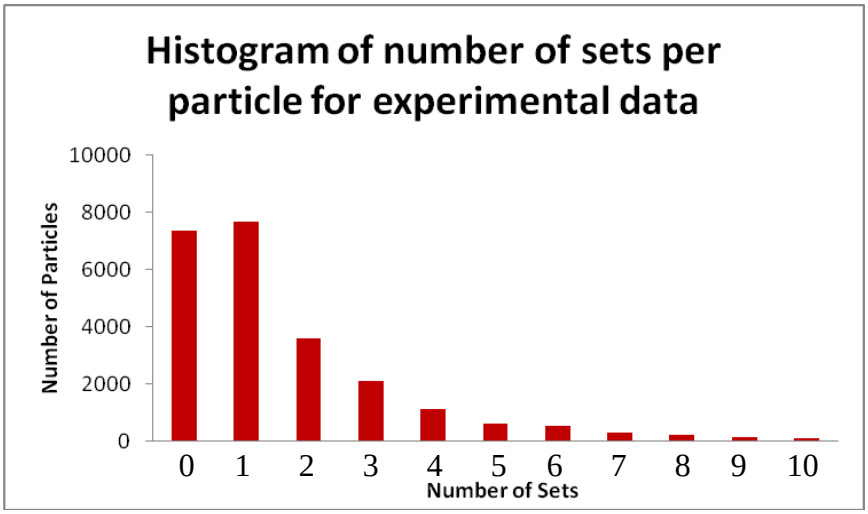

Figure 79: A histogram of the experimentally measured number of sets per particle image shows that the largest number of particle images are members of only one (presumably valid) set. A small number of particle images belong to many sets, and a large number of particle images are not part of any set.

An approximate measure of the experimental ratio of valid to invalid sets above the validity threshold can be obtained by using the percentage-based results from the previous chapter (at a given threshold value, 70% of valid sets are above the threshold, but only 4% of invalid sets) along with the data presented above in Figure 79. Of the 24,000 particle images analyzed in Figure 79, about 17,000 are in at least one set, but the total number of sets is approximately 40,000. Assuming each particle in at least one set is also member of one valid set, approximately 92% of the sets above the validity threshold will actually be valid, as shown below.

70% x 17,000 valid sets =11,900 valid sets above threshold (4)

4% x 23,000 invalid sets = 920 invalid sets above threshold (5)

11,900 valid sets + 920 invalid sets = 12,820 total sets (6)

11,900 /12,820 = 92% of sets above threshold are actually valid (7)

However this is an upper bound on the accuracy due to the large number of particle images that are not a part of a set. If it is assumed that 30% of the particles in at least one set are not a member of a valid set (because approximately 30% of the particle images analyzed in Figure 79 are not part of any set), the percentage of valid vectors above the validity threshold drops to 88%, as shown below.

$$70\% \times 11{,}900 \text{ valid sets} = 8{,}330 \text{ valid sets above threshold} \quad (8)$$

$$4\% \times 28{,}100 \text{ invalid sets} = 1{,}124 \text{ invalid sets above threshold} \quad (9)$$

$$8{,}330 \text{ valid sets} + 1{,}124 \text{ invalid sets} = 9{,}454 \text{ total sets} \quad (10)$$

$$8{,}330 / 9{,}454 = 88\% \text{ of sets above threshold are actually valid} \quad (11)$$

8.3 Particle Images without Sets

A large number of the particle images in the histogram in Figure 79 do not belong to any sets. There are several possible causes. One cause is the motion of the particles between successive exposures. It is possible for a particle to move into or out of the measurement volume between the first and second pair of simultaneously-acquired images, meaning that the images corresponding to that particle appear only in the first or second pair of images. This problem can be exacerbated if the laser sheet providing the illumination is not identical for the successive image pairs. Because the temporal separation of the two image pairs is small, the illumination for the first and second image typically comes from two different laser heads (reducing the necessary repetition rate for a single laser head). The two laser heads often have slightly different beam profiles, so the shape of the measurement volume changes slightly between the first and second exposure, and even for stationary particles, not all particles in the measurement volume at

the time of the first exposure will still be in the measurement volume at the time of the second exposure, as shown in Figure 80.

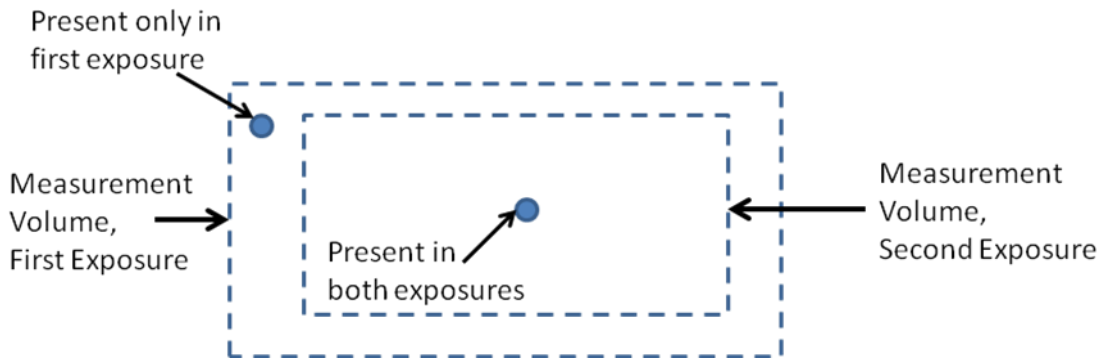

Figure 80: If the shape or size of the measurement volume changes between the first and second exposure (likely due to differences in laser beam profile), the same particles will not be captured in both exposures.

The difference in beam properties can also mean that the intensity with which a particle is illuminated does not stay constant for the first and second exposure. If the intensity changes, one particle image may be just above the intensity threshold necessary to be considered a valid particle image, while the other pair is just below the threshold.

Overlapping images from separate particles can also cause some particle images to not be placed in sets. The most likely cause is a shift in the apparent center of a particle image when it is partially overlapping another particle image. For example, Figure 81 shows two different views of two particle images. In one view the two particle images are clearly separated, meaning two distinct particle center positions will be calculated. In the other view, the two particle images overlap slightly, meaning that the two particle images will be interpreted as a single particle image, with only one center. The position of the center of these overlapped images is displaced several pixels from their true centers, meaning that it may no longer be within the search areas defined by the separated particle image centers.

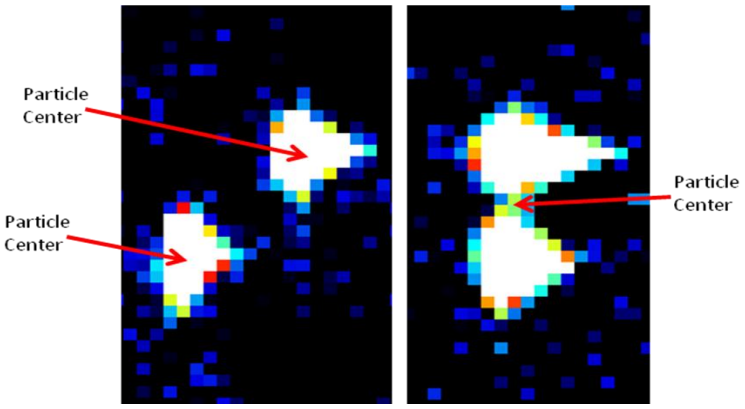

Figure 81: The left figure shows two particle images separated and the two particle centers. The right figure shows the same two particle images when viewed from a different angle, such that they are overlapping slightly, and causing them to be interpreted as a single particle image.

If images from the two separate particles are overlapping significantly, the single overlapped particle image will likely be placed in the two correct sets, but at most one of the correct sets will pass the validation procedure, since the SC3D-PTV algorithm does not allow a single particle image to be a member of more than one valid set. Because the size and shape of the overlapped particle images is now much different than for the separated particle images, possibly neither set will be considered valid.

The number of overlapping particle images grows linearly with the particle image size and the depth of the measurement volume, and grows with the square of the number of particle images within a single raw image (Stuer, 1999). Although some methods do exist for separating partially overlapping particle images (e.g. Maas et al., 1993), because SC3D-PTV analysis makes additional image similarity calculations not performed in the epi-polar based algorithms that typically employ image separating algorithms, they were not implemented for SC3D-PTV analysis. Rather, the number of overlapping particles is minimized by using moderate seed densities.

8.4 Summary

The sparse data analysis of Chapter 7 and the analysis of densely-seeded images presented in this chapter indicate that valid and invalid sets can be separated using the SC3D-PTV similarity-based algorithm. However, beyond a statistical basis for separating valid and invalid sets, it is important to show that the results given by SC3D-PTV accurately reflect the properties of a flow under study. To provide this verification, the following chapter applies the ideas of both the previous and current chapters to a simple flow that was studied using both stereoscopic PIV and SC3D-PTV. By comparing the results of the two techniques, greater confidence in the results of the SC3D-PTV analysis can be gained.

CHAPTER 9

SC3D-PTV VALIDATION: COMPARISON TO STEREOSCOPIC PIV

The previous two chapters showed how valid and invalid particle image sets could be separated using similarity comparisons, and estimated the success rate of the SC3D-PTV algorithm for identifying valid and invalid sets. This chapter aims to provide further validation of the algorithm using experimental data. Because no *a priori* information about which sets are truly valid is available for experimental flow studies, the algorithms cannot be validated in an absolute sense using experimental data.

Further validation of the SC3D-PTV technique can be gained from experimental data in two ways: either the results of SC3D-PTV can be compared to the properties of a known flow, or a flow can be studied with both SC3D-PTV and a known flow measurement technique, and the results compared. The latter approach, i.e. comparing SC3D-PTV results to the results of another technique when applied to the same flow, was chosen to further validate SC3D-PTV for two reasons. First, SC3D-PTV was designed around very specific parameters for working distance, measurement volume, and particle size. A known flow with the same parameters was not readily available. Second, because the SC3D-PTV optic functions as an image splitter with two separate viewing angles of the same measurement volume, the data can be analyzed with either the SC3D-PTV algorithm, or a stereoscopic PIV (SPIV) algorithm (a description of SPIV is given in

Chapter 2). This allows a single data set to be analyzed two different ways, and the results of the SC3D-PTV algorithm can be compared to the results of the SPIV algorithm. Because only a single raw data set is used, this removes the possibility of misalignments between two different sets of measurement apparatus and ensures that the flow is identical for both analyses.

9.1 Experimental Set-up

The flow studied in this chapter is the air exiting from an air hose (1/2" inner diameter) into ambient conditions. This experimental arrangement was chosen since the direction is nominally uniform and steady and the velocity profile is approximately parabolic (allowing both visual confirmation of results and temporal averaging). Also, many of the experimental parameters (e.g. measurement volume size, flow speed, particle size) could be matched to the parameters of engine measurements, while offering a much simpler experimental set-up.

The air was seeded with silicon oil droplets (nominal diameter: 1 micron) using a TSI 9306 atomizer. Pairs of temporally-separated images were collected at 1 kilohertz using a Phantom camera (Vision Research, v7.3) and the SC3D-PTV optic. Image pairs were separated by 20 microseconds using a frame-straddling approach (as is common for PIV experiments). Illumination was provided by a Quantronix Darwin Duo laser operating at 527 nanometers (frequency-doubled Nd:YLF) and gave 500 micro-joules of energy per pulse over a measurement volume of 6 millimeters in the horizontal and vertical direction, with a depth of 2 millimeters. The measurement volume was defined

by the thickness of the laser sheet and the size of the aperture inside the image cropping optic. (The SC3D-PTV optic and image cropping optic are described in Chapter 4.)

The experimental set-up can be seen in Figure 82.

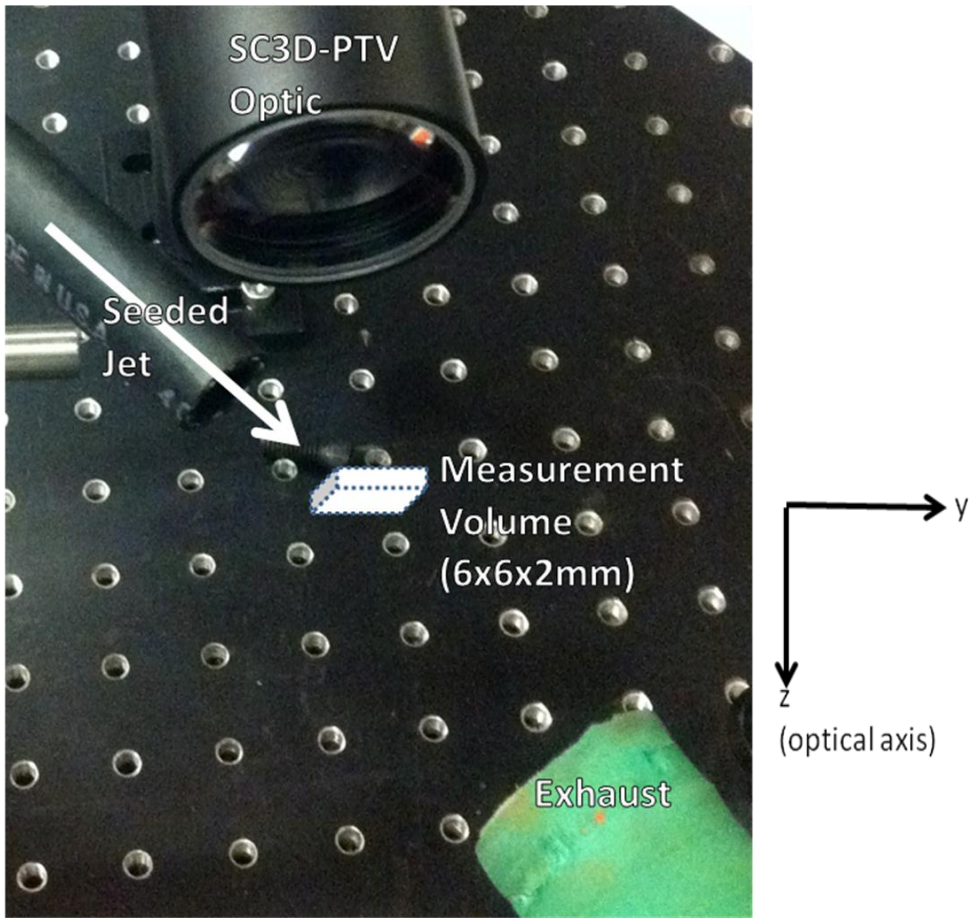

Figure 82: Experimental set-up. An approximately uniform, steady air jet seeded with silicon oil droplets was studied by illuminating the droplets with laser light and imaging the flow using the SC3D-PTV optic attached to a high-speed camera.

The measurement of the in-plane velocity component and out-of-plane velocity component is fundamentally different for both SC3D-PTV and SPIV. The resolutions of the two components are also different. To study both in-plane and out-of-plane motion, ten data sets of 100 image pairs each were collected for a flow with approximately equal

in-plane and out-of-plane velocity components. By altering the pressure supplied to both the seeder and the bypass air, the seed density of the flow was varied while maintaining approximately equal flow velocity. Although seed densities as high as 0.005 particles per pixel were studied, the maximum rate of returned vectors corresponded to a seed density of approximately 0.002 particles per pixel. Smaller particle images were also observed to correspond to higher levels of returned vectors, presumably because of a smaller amount of overlapping particle images.

9.2 Experimental Results

9.2.1 SC3D-PTV Data

Simultaneously-acquired raw images taken from the left and right view of the seeded flow for a densely seeded case are shown in Figure 83. Approximately 300 particle images are found within each raw image.

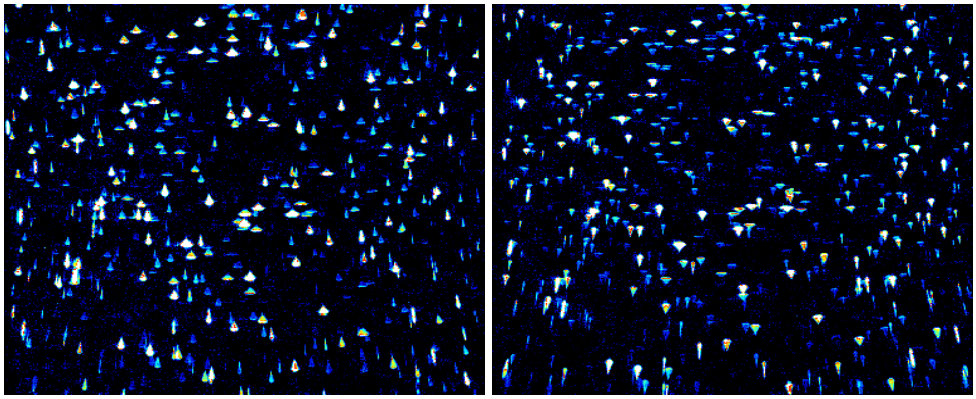

Figure 83: The simultaneously-acquired left and right views of the seeded flow.

Of the 300 particle images in the raw images above, approximately 100 were found to be part of a valid set. Each three-component, three-dimensional velocity fields

required approximately 20 seconds of processing time to compute from the raw images. Figure 84 shows this velocity field in the form of a cone plot. The position of the cone indicates the average position of a measured particle, the length of the cone indicates the velocity of that particle, and the direction of the cone indicates the three-component direction of the particle's displacement. A reference cone of known magnitude is also shown. Most of the particles represented have similar magnitude and direction, with about 5% deviating significantly from the mean.

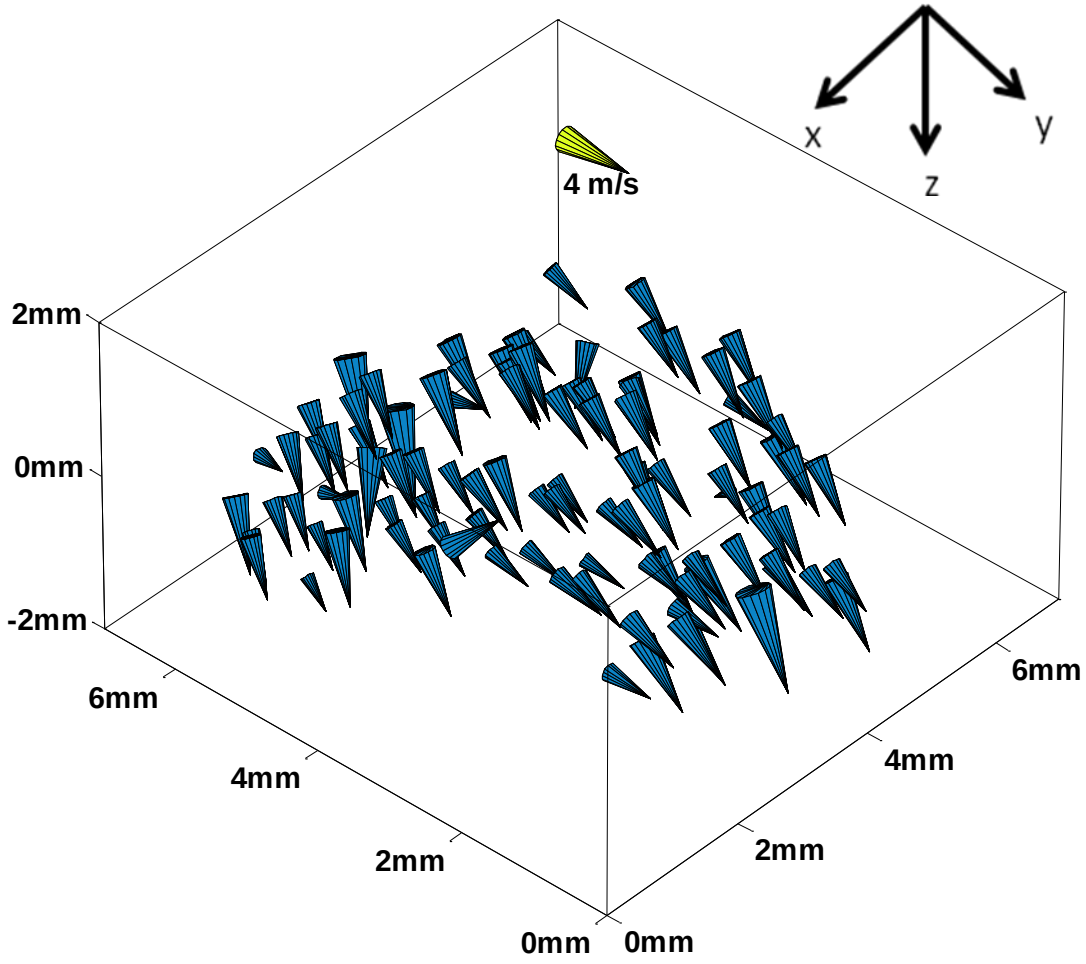

Figure 84: The velocity field resulting from SC3D-PTV analysis is shown with cones representing particle displacement. Cone location indicates particle position, cone height indicates displacement, and direction of the point of the cone indicates velocity direction.

Although the cone plot gives a sense of the flow field, visual interpretation is difficult for two reasons. First, the data is calculated in 3D space, but is projected onto a 2D plane, which results in some obscured data, depending on the projection angle. Second, the random positioning of the cones can make the characteristics of the overall flow hard to determine. The first difficulty can be alleviated somewhat by projecting the 3D information at several different angles. Figure 85 shows the same data as Figure 84, but projected along the x-axis. The combination of Figure 84 and Figure 85 gives a better sense of the data than either projection individually.

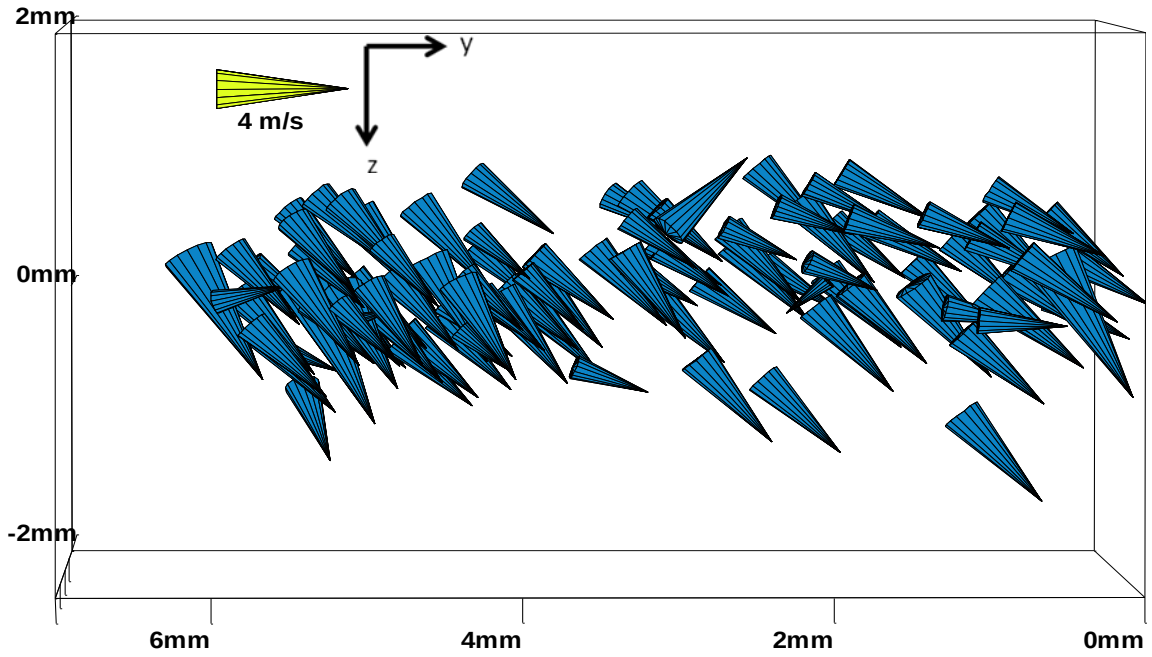

Figure 85: The same data shown in Figure 84, but projected along the x-axis.

The random positioning of the velocity measurements still hinders visual interpretation when multiple views are presented, and also limits other types of analysis of the results. Because the flow is nominally steady, it is desirable to average multiple measurements of the same flow to reduce the influence of occasional spurious vectors on

the results. Comparison to SPIV data is also difficult, since the SPIV vectors are computed on an evenly spaced grid in a single plane.

To facilitate both averaging of SC3D-PTV velocity fields and comparison of the SPIV and SC3D-PTV results, the randomly-located SC3D-PTV vectors were linearly interpolated onto a single-plane grid of the same spacing as the SPIV results (32x32 pixels). Interpolation has a smoothing effect on the data, since multiple raw vectors are averaged to produce each interpolated vector. For flows with characteristic lengths of the same order as the vector spacing, interpolation will cause information about these structures to be lost, but since the jet flow studied in this chapter is approximately uni-directional with a slightly parabolic velocity profile, the interpolation should only minimally affect the results. Figure 86 shows the data from Figure 84 after interpolation.

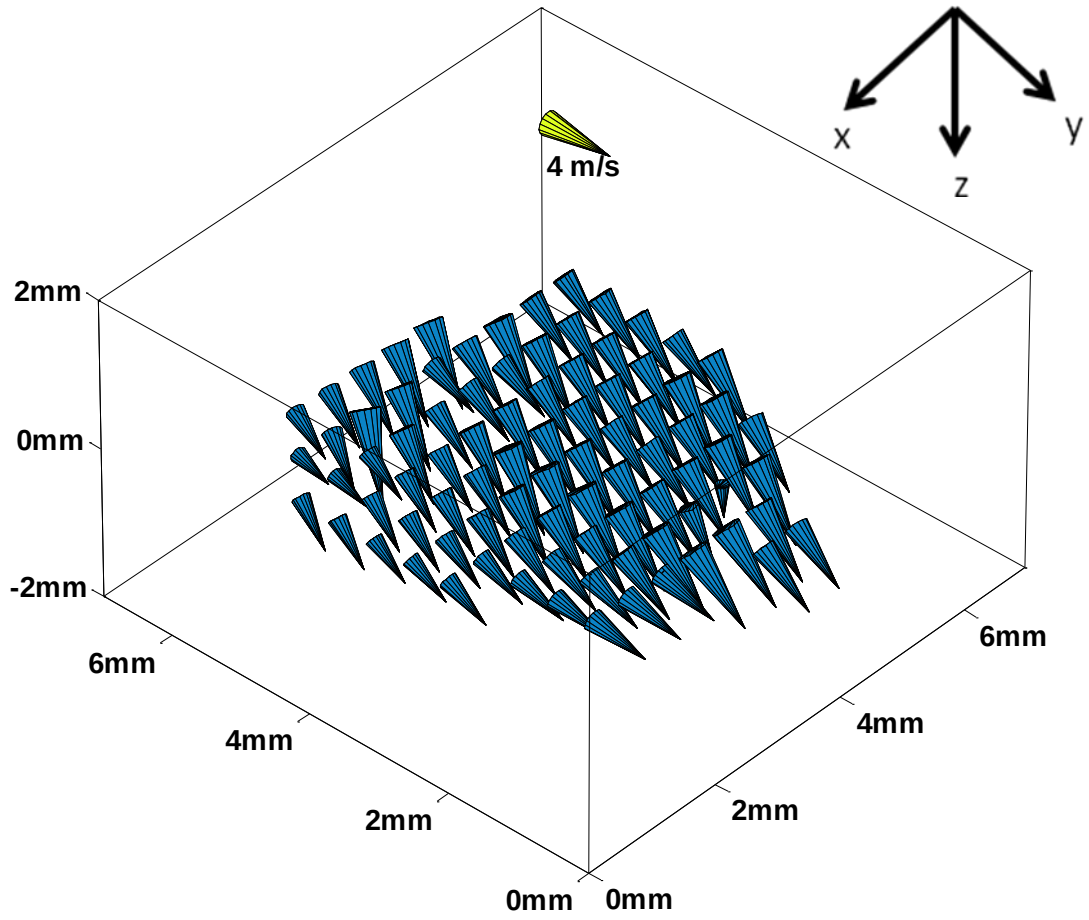

Figure 86: The flow shown in Figure 84 after interpolating onto an evenly spaced grid.

The interpolated flow shown in Figure 86 appears much more uniform than the flow in Figure 84, but a small number of outliers are still present. The impact of these random outliers was minimized by averaging multiple measurements of the same flow from different points in time. Figure 87 shows the averaged results of 100 SC3D-PTV measurements taken over a 100 millisecond time interval, and appears more coherent than the individual velocity field in Figure 86.

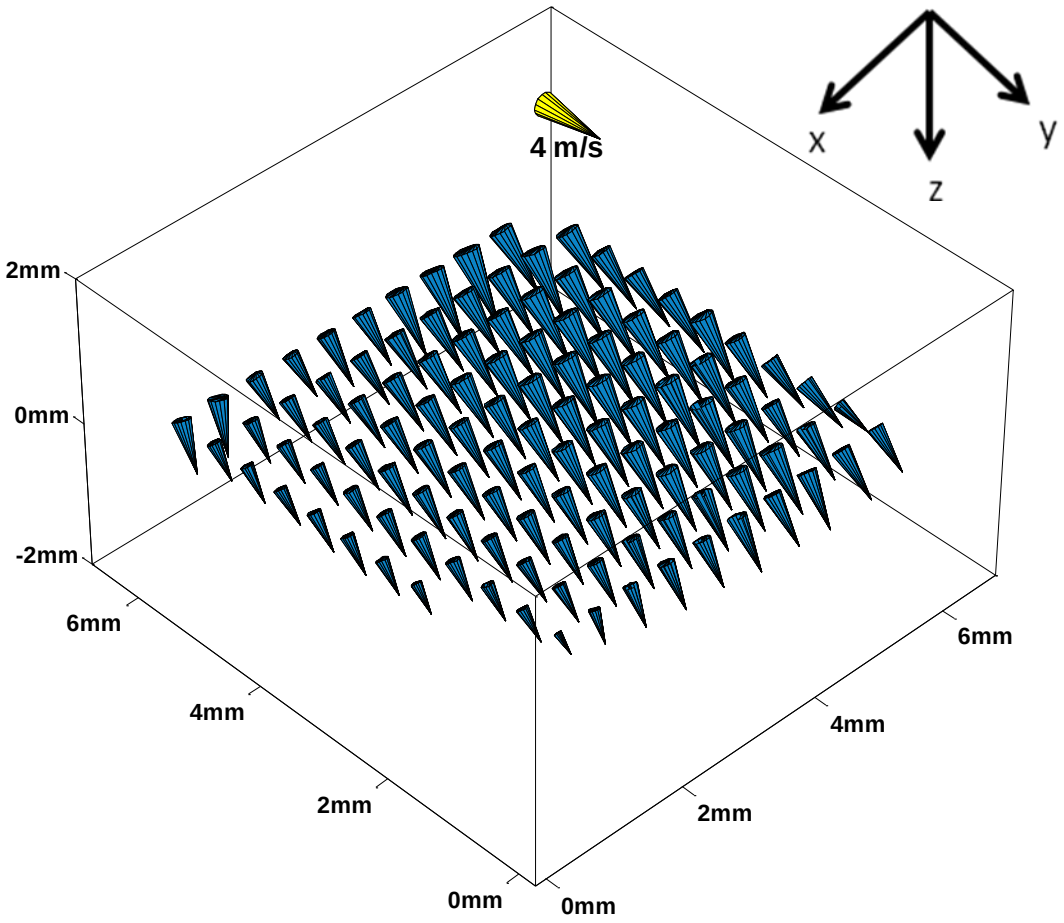

Figure 87: The average velocity field calculated from 100 velocity measurements of the same flow taken over an interval of 100 milliseconds shows a very coherent, uni-directional velocity field.

After interpolation is used to place the data on a single-plane grid of the same spacing as SPIV, and multiple velocity fields are averaged to minimize the impact of occasional spurious vectors, the SC3D-PTV data is in a form suitable for comparison to the results of a SPIV algorithm applied to the same set of data.

9.2.2 Stereoscopic PIV Data

Because the raw images collected using the SC3D-PTV optic can be analyzed using either the SPIV algorithm or the SC3D-PTV algorithm, it is possible to compare

the two algorithms using the same data set. The results will likely not be identical due to differences in the assumptions underlying the two algorithms, but the comparison can validate SC3D-PTV in a general sense and provide a bound on the uncertainty of SC3D-PTV measurements.

Unlike SC3D-PTV, which tracks individual particles in three dimensions, SPIV processing is performed by treating the two viewing angles as individual 2C2D PIV measurements, and three-component velocities are reconstructed from the two separate 2C2D flow fields that are computed using a cross-correlation-based PIV algorithm.

Cross-correlation can be interpreted as a statistical method of finding the most likely value of the average in-plane velocity for a small number of particles that are spatially clustered. A PIV cross-correlation calculation is computed by spatially matching successively-acquired sub-images (the sub-images for this analysis are 32x32 pixels) from the same viewing angle, shifting the sub-images with respect to each other, and multiplying the sub-images together pixel-by-pixel for each shift. The shift (both direction and magnitude) that yields the highest product is interpreted as the most likely displacement of the particles contained in that sub-image (a more extensive discussion of cross-correlation and PIV analysis can be found in Raffel et al., 1998). A 2C2D vector field is computed by performing cross-correlations on sub-images covering the entire image space.

After spatially matching the measurements taken from different viewing angles, there is sufficient information to reconstruct a single 3C2D velocity field using geometric considerations. All three components of velocity are computed, but they are only computed on a single 2D plane. No information about the depth of the particles is

obtained using SPIV analysis, unlike SC3D-PTV, which also provides the 3D position of each 3C vector.

The DaVis commercial flow analysis software package from LaVision GmbH was used to compute 2C2D PIV flow fields separately for each viewing angle, and the 3C2D flow fields were reconstructed using knowledge of the relative angle of the two viewing angles. Like the SC3D-PTV results, the SPIV analysis is the average of 100 measurements to minimize the influence of occasional spurious vectors. The SPIV algorithm was applied to the same data used to produce Figure 87 above, and the results are shown in Figure 88.

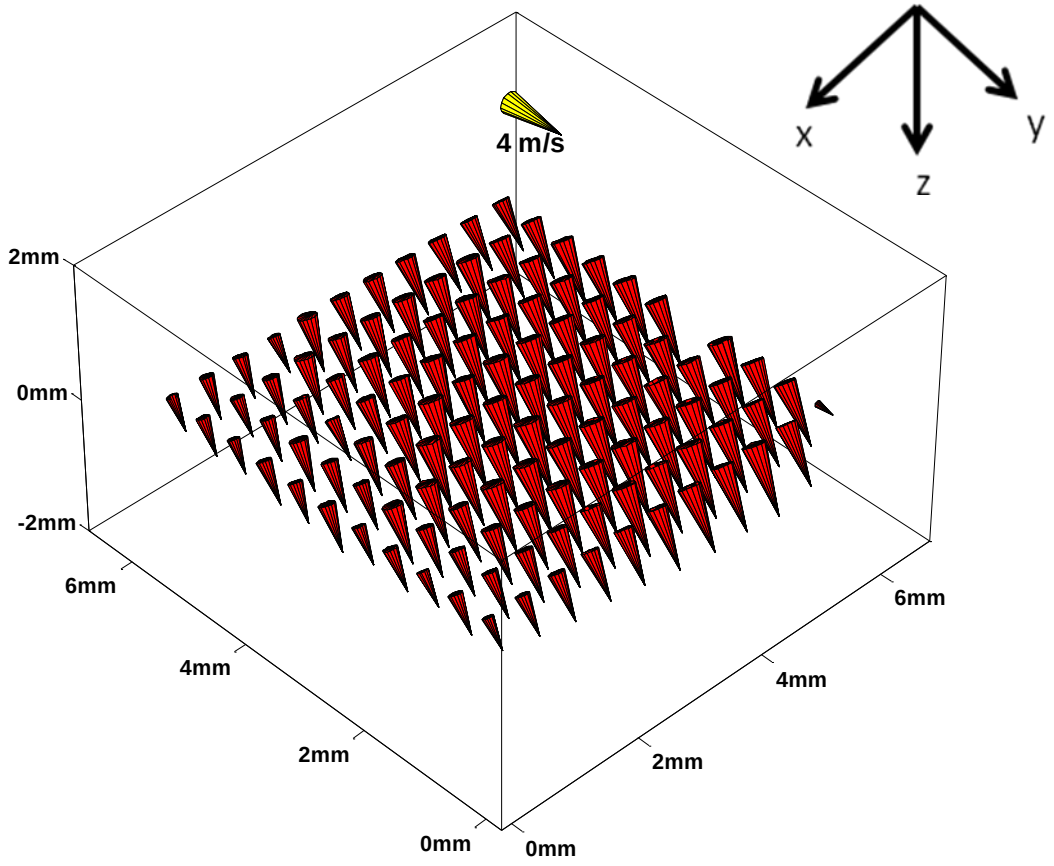

Figure 88: 100 SPIV velocity fields were computed using the same data as Figure 87, and the average velocity field was computed.

9.2.3 Comparison of SC3D-PTV Data and Stereoscopic PIV Data

Comparing Figure 87 (SC3D-PTV results) and Figure 88 (SPIV results) at a qualitative level, both velocity fields show very similar flows. The comparison can be made quantitative by looking at both the bulk flow and the velocity profile.

Bulk Flow Comparison

The bulk flow is the average displacement of the flow over the entire measurement volume. The average bulk flow values for the in-plane and out-of-plane components were calculated separately for each of 10 data sets (where each data set contains 100 velocity fields), and the results are shown in Figure 89 and Figure 90. Displacements were calculated in pixels to remove possible discrepancies in calibration. For both the in-plane and out-of-plane flows, the minimum percentage error is about 1%, and the maximum is about 10%. The average percentage difference for the 10 data sets is 4%. On an absolute basis, the average difference is between 0.1 and 0.2 pixels for both the in-plane and out-of-plane direction. In 9 of the 10 data sets, the SPIV bulk flow is higher for both the in-plane and out-of-plane flow, but the difference is not systematic on either an absolute basis or a percentage basis.

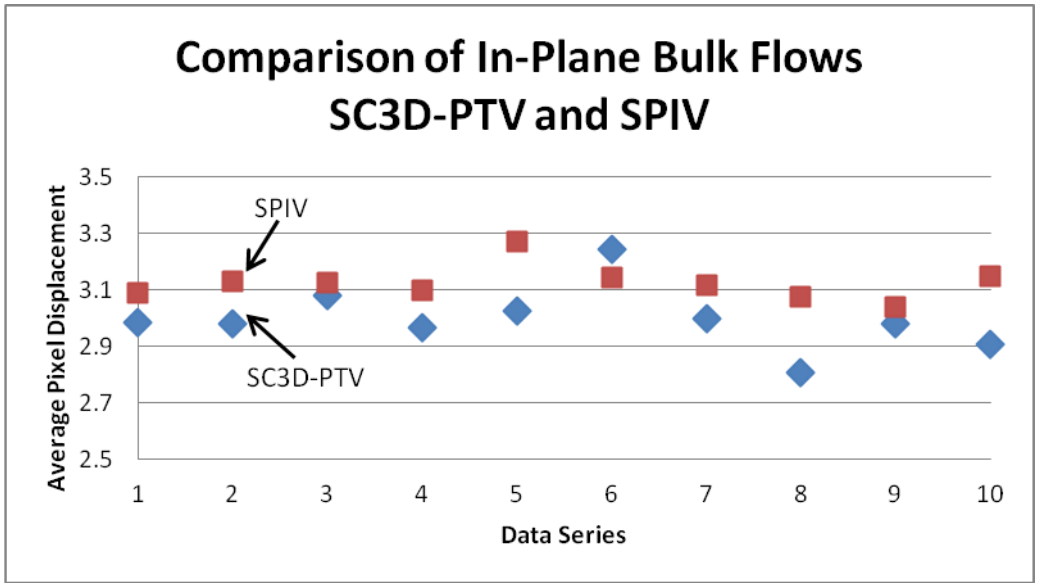

Figure 89: The average in-plane bulk flow was computed for 10 data sets using both the SC3D-PTV algorithm and a SPIV algorithm. The results agree within 4% on average.

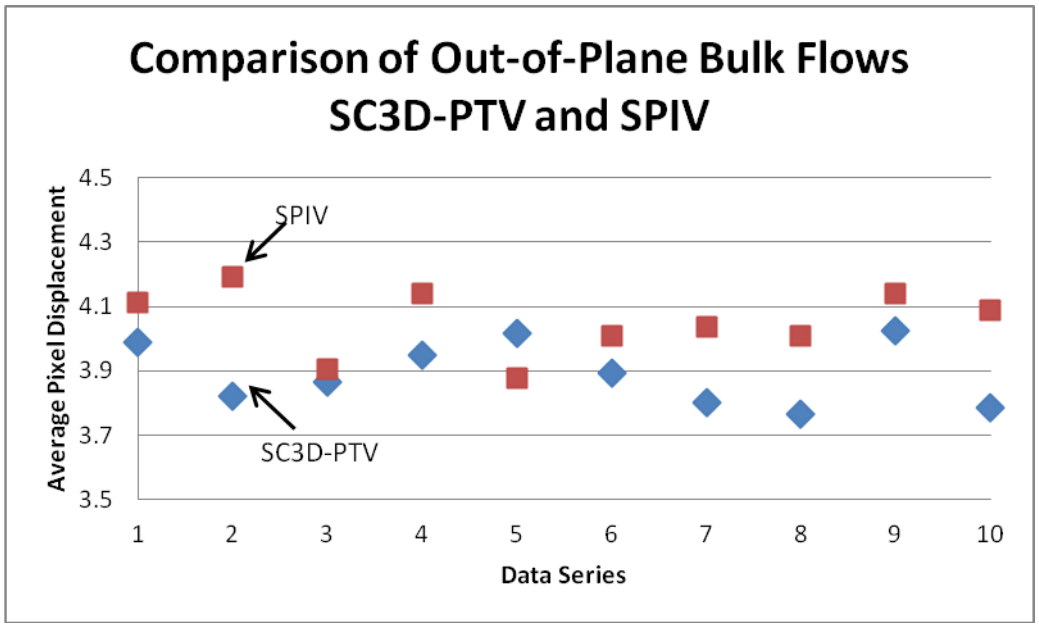

Figure 90: The average out-of-plane bulk flow was computed for 10 data sets using both the SC3D-PTV algorithm and a SPIV algorithm. On average, the results agree within 4%.

Velocity Profile Comparison

In addition to properly resolving the bulk flow, SC3D-PTV must be able to measure features within the flow. To make a quantitative comparison of flow features measured by SC3D-PTV and SPIV, the flow profile is computed for a single data set of 100 images, integrating along both the x- and y-axis. To serve as an upper bound on the uncertainty, the set with the largest difference in bulk flow is analyzed (data set #2 of the data presented in Figure 89 and Figure 90). An initial estimate of the profile similarity along the x-axis can be gained by examining the two flow fields when projected parallel to the x-axis. The flow field calculated using the SC3D-PTV algorithm is shown in Figure 91, and the flow field calculated using the SPIV algorithm is shown in Figure 92.

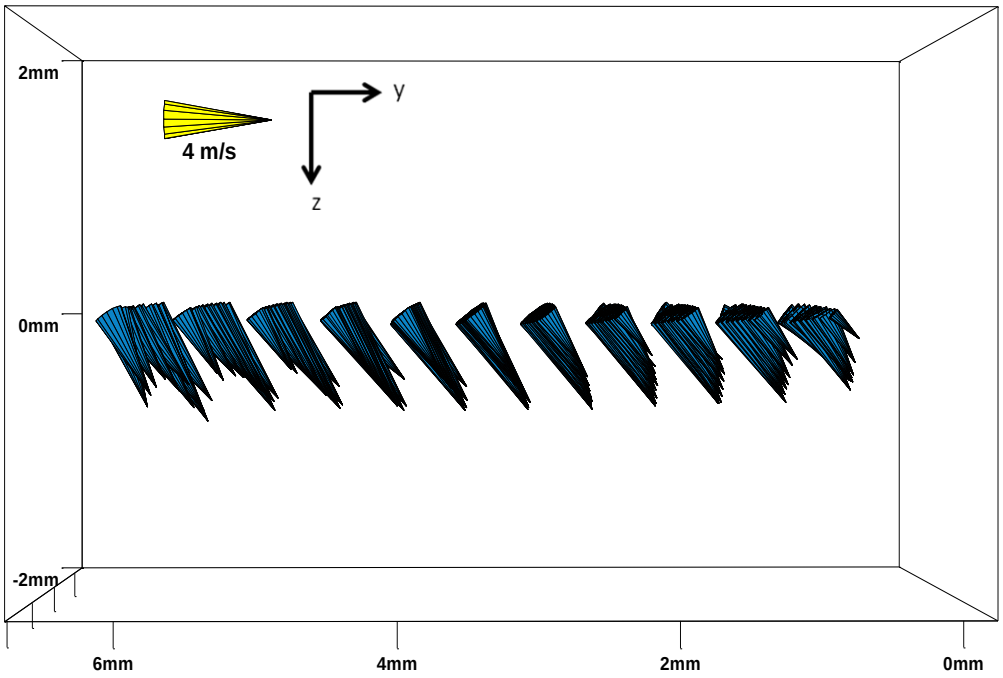

Figure 91: The average flow field, as computed using the SC3D-PTV algorithm, is projected along the x-axis.

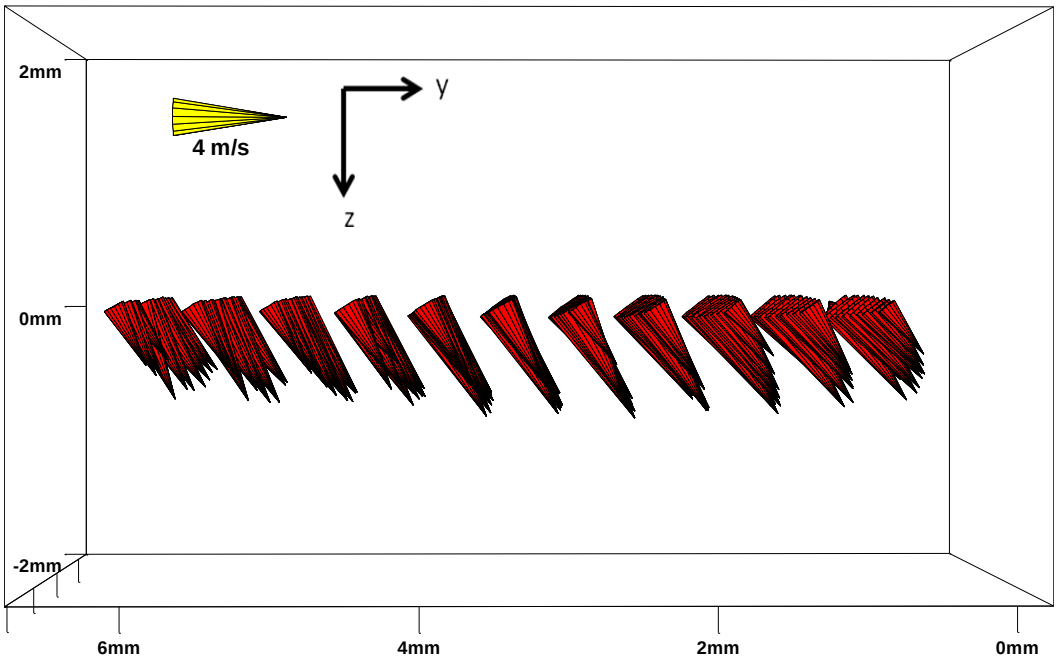

Figure 92: The average flow field, as computed using the SPIV algorithm, is projected along the x-axis.

The two profiles are qualitatively similar, and the quantitative profiles agree closely, as seen in Figure 93. The average velocity difference between the two results is 3% and the maximum difference at a single point is 9%.

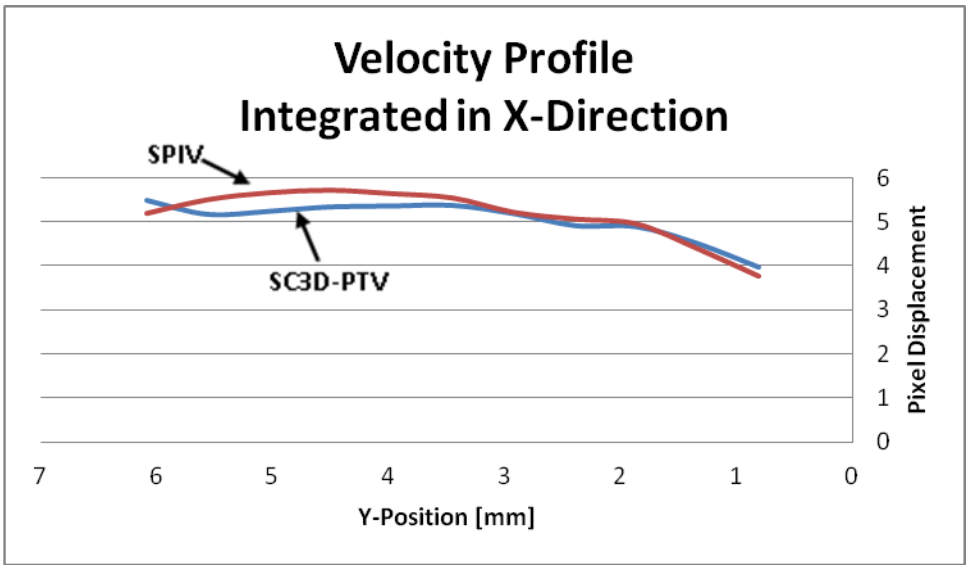

Figure 93: The velocity field is integrated along the x-axis for both the SC3D-PTV results and the SPIV results, and the profiles are compared. The general shapes of the profiles agree, the average velocity difference is only 3% and the maximum difference is less than 10%.

Figure 94 and Figure 95 show the y-projections of the flow fields computed using the SC3D-PTV and SPIV algorithms, respectively.

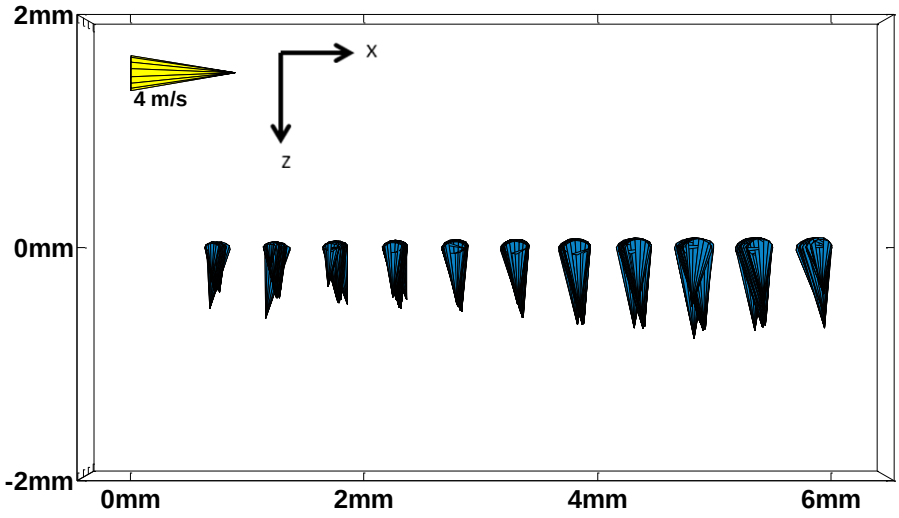

Figure 94: The average flow field, as computed using the SC3D-PTV algorithm, is projected along the y-axis.

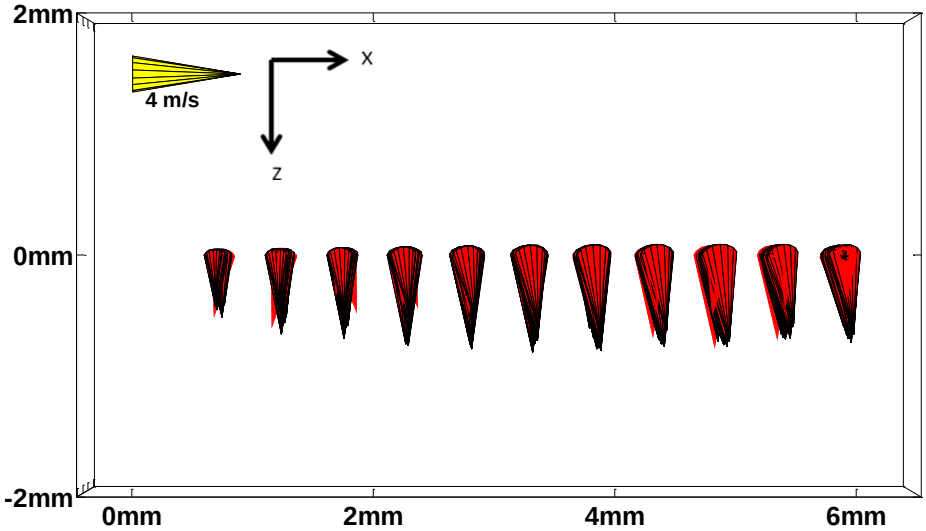

Figure 95: The average flow field, as computed using the SPIV algorithm, is projected along the y-axis.

As with the x-axis projections, the two profiles are qualitatively similar. The quantified results are slightly worse for the y-projections than the x-projections, as seen in Figure 93. For the velocity profiles integrated in the y-direction, the average difference is 7% and the maximum difference between the two results is 14%.

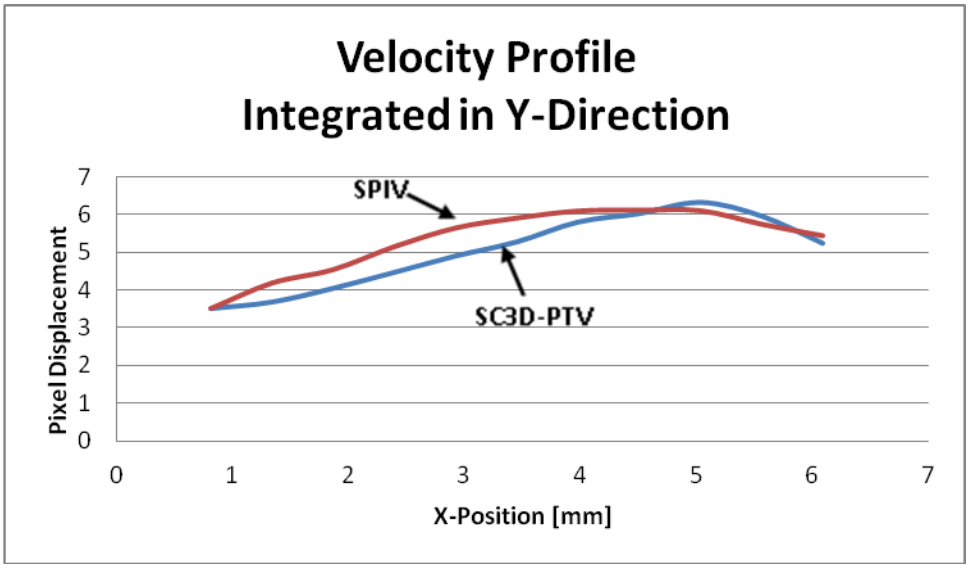

Figure 96: The velocity field is integrated along the y-axis for both the SC3D-PTV results and the SPIV results, and the profiles are compared. The general shapes of the profiles agree. The average difference is 7% and the maximum difference is 14%.

9.3 Summary

Analyzing the same data set with both the SC3D-PTV algorithm and a SPIV algorithm gives results that are qualitatively similar and agree quantitatively within a few percent, on average, for both bulk flow values and velocity profiles. The results are not identical, but some deviation is expected due to the difference in underlying assumptions behind the SC3D-PTV algorithm and the SPIV algorithm.

Furthermore, for the intended application of SC3D-PTV, i.e. engine measurements, an absolute error of a few percent is acceptable. Because of the cyclic

variability of engine flows, and the myriad unsteady boundary conditions that influence these flows (for example, intake and exhaust manifold pressure fluctuations, cylinder temperature variations) even the bulk flow values of tests repeated at nominally identical conditions can vary by several percent. Significant knowledge of the engine is gained by knowing the general properties of the flow, such as the direction of the flow and the characteristic length of the structures within the flow, even if the absolute values are miscalculated by a few percent. Toward this end, this chapter shows that SC3D-PTV fully meets these requirements and can be applied to engine flows.

The next chapter applies SC3D-PTV to the air flow inside a motored engine.

CHAPTER 10

IN-CYLINDER EXPERIMENTAL RESULTS

The previous chapter validated the SC3D-PTV technique developed in this thesis by comparing the results of a simple flow measured using both SC3D-PTV and stereoscopic PIV and showing that the velocity fields match within a reasonable uncertainty.

This chapter applies the SC3D-PTV technique to the air flow inside a motored, optically-accessible engine. The aim of this chapter is to present examples of the types of flows known to exist in engines, and show what insights are gained by having three-component, three-dimensional knowledge of the flow field, as is offered by SC3D-PTV, rather than two- or three-component planar information, like that provided by PIV or stereoscopic PIV.

These tests were performed at the General Motors Technical Center and the financial support from General Motors to perform these tests is gratefully acknowledged.

10.1 Experimental Set-up

All tests were performed in a single-cylinder General Motors fourth-generation SG SIDI optical engine. The measurement volume was located between the intake and exhaust valves, and near the edge of the piston bowl. The location of the measurement volume, fuel injector, spark plug, piston and head window, as well as the positioning of

the camera, can be seen in the schematic image in Figure 97. Optical access is through a quartz window (10 x 25 mm) in the engine head (for imaging) and a customized Bowditch piston with a bottom quartz window (for laser illumination). The piston also has a side cut-out in the piston bowl, to allow imaging when the piston is at top dead center.

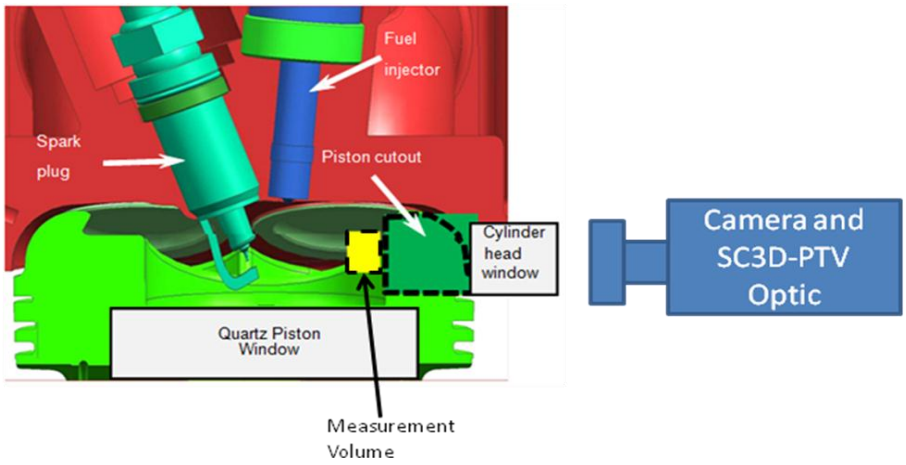

Figure 97: Schematic of In-Cylinder Imaging Experimental Set-Up. Laser light enters the cylinder through the quartz window that forms the bottom of the piston bowl and illuminates the measurement volume. The camera images through the quartz head window and piston cut-out (image courtesy GM R&D).

Figure 98 shows the SC3D-PTV optic in the same position depicted in the schematic image in Figure 97.

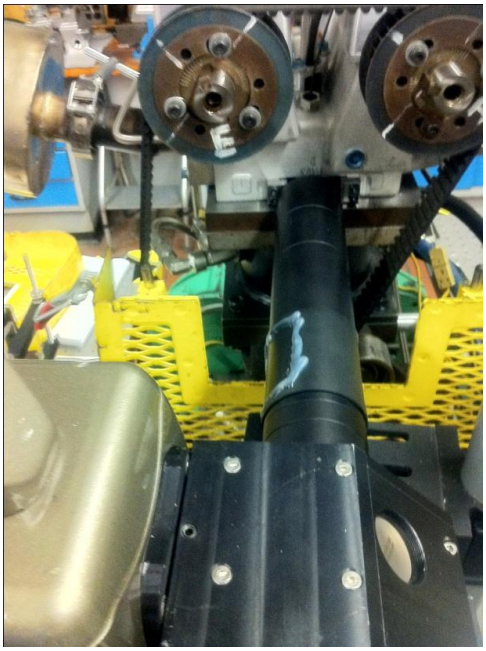

Figure 98: The SC3D-PTV optic is positioned to take data inside the engine.

The 0.5L displacement engine cylinder has a bore and stroke each measuring 86 millimeters, a compression ratio of 10.3:1, and a four-valve, pent-roof combustion chamber.

A bottom view of the engine head showing the valves, spark plug, and fuel injector is shown in Figure 99. The position of the measurement volume can also be seen. One intake valve is equipped with a variable throttle that controls the ratio of swirl and tumble flow, but otherwise the intake ports are symmetric.

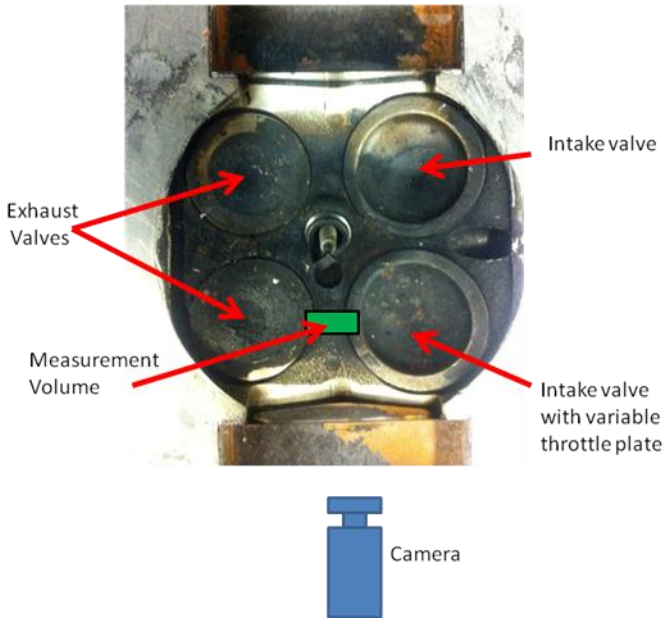

Figure 99: Bottom view of combustion chamber surface of engine head. The positions of the valves and measurement volume are shown. One intake valve has a variable throttle plate to control the degree of swirl flow. The laser sheet is aligned with the space between the intake and exhaust valve. The view of the camera can also be seen.

These experiments were performed at 600 RPM. Intake manifold pressure was 95 kPa, with oil and water at 50° C and the intake air at 25° C. The intake air flow was seeded with silicon oil droplets (nominal diameter: 1 micron) using an atomizer (TSI model 9306). The atomizer pressure (~1-5 psi) was adjusted to produce a seed density between 1 and 10 particles per cubic millimeter. The droplets were illuminated using a Quantronix Dual-Hawk II, frequency-doubled Nd:YAG laser operating at 532 nm. Each of the two laser heads operated at 1.8 kHz, producing approximately 10 Watts of average power each. Image pairs were captured at 1.8 kHz, which corresponds to one velocity measurement every other crank angle degree at 600 RPM. Motored conditions (no fuel injection, no spark, and no combustion) were studied during the compression stroke, from bottom dead center to top dead center. The intent was to examine the flows created by the valve and piston motion as a function of intake-valve-throttle position. Two sets of tests

were performed: one set with the variable throttle fully open, and one set with the throttle fully closed.

When the throttle is fully open, air enters the cylinder through both intake valves, and the flow structure is primarily tumble flow. When the throttle is fully closed, air enters through only one intake valve and creates primarily a swirl flow. Schematic representations of tumble and swirl flows are shown in Figure 100. Although the air flow at a given throttle position approaches the tumble or swirl on average, the individual cycles can still be quite three-dimensional and turbulent.

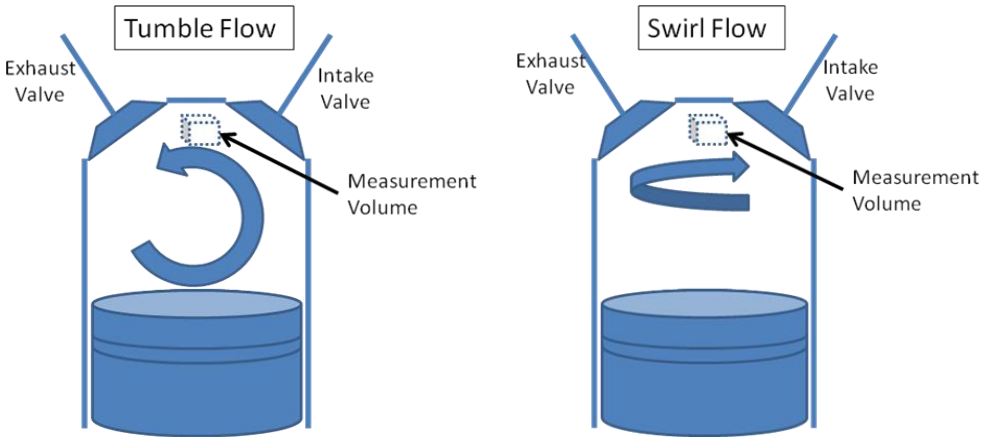

Figure 100: A variable throttle plate allows the dominant in-cylinder flow structure to be altered. When the throttle plate is open, a tumble flow (left image) develops. When the throttle plate is closed, a swirl flow (right image) forms.

10.2 Experimental Results

The air flow near the spark plug was studied previously in a very similar engine using two-component, planar PIV (Peterson, 2009), and those tests serve as a baseline to which the SC3D-PTV results can be compared. However, rather than comparing flow fields from the two different types of experiments quantitatively, the two-component, planar results are used a benchmark for the types of information that can learned from flow studies, and several examples are taken from the SC3D-PTV results to emphasize

the greater degree of insight provided by three-component, volumetric flow information compared to the two-component, planar results. All SC3D-PTV results are shown after linear interpolation onto a 0.5 x 0.5 mm grid.

10.2.1 Multi-Cycle Averages

Bulk Flows

As in Chapter 9, bulk flow values serve as an initial method of quantifying and comparing results. Sixty engine cycles of SC3D-PTV data were processed for both the swirl flow and tumble flow conditions, and the average bulk flow values for each condition were computed at every crank angle. The three components of the bulk flow are shown separately, to emphasize the information that is added by having all three components of the velocity, instead of only two. A schematic representation of the definitions of three flow components is shown in Figure 101.

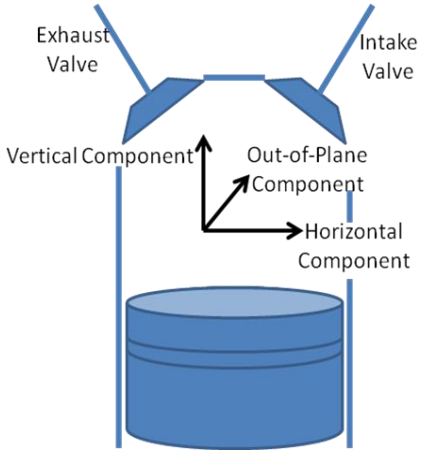

Figure 101: Schematic representation of the three flow components. Horizontal flow is parallel to the piston crown, toward or away from the intake side of the cylinder. Vertical flow is toward or away from the engine head, and the out-of-plane flow is toward or away from the center of the cylinder.

The horizontal component of the flow is expected to differ significantly for the swirl and tumble flows. For the tumble flow condition, the variable throttle is open and the flow near the measurement volume is directed toward the exhaust valve. The dominant tumble structure reinforces the horizontal component in this direction. For the swirl condition, the dominant structure is rotation about the axis of the cylinder. At the position of the measurement volume, the horizontal component is directed toward the intake side of the engine. However, at some crank angles, interactions between the dominant swirl flow and the contoured piston disrupt this flow.

Figure 102 shows that the average horizontal component of the tumble flow is oriented primarily in the direction of the exhaust valve throughout the entire compression stroke. The swirl flow has a minimal horizontal component, for much of the compression stroke, likely due to the interaction with the piston (further examples of this interaction are given in the next section). As the piston approaches top dead center, the horizontal component is directed toward the intake side of the cylinder.

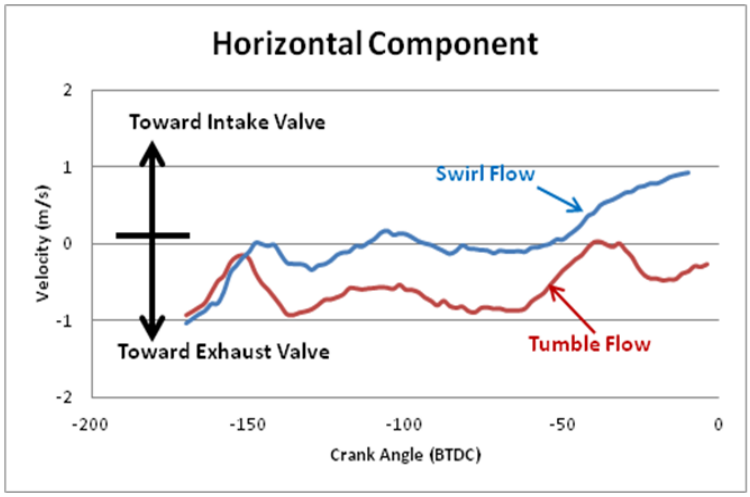

Figure 102: The average horizontal component of the tumble flow is toward the exhaust valve for most of the compression stroke. The horizontal component of the swirl flow is near zero, due to interactions with the piston, until close to top dead center, when the flow is directed toward the intake valve.

The swirl flow is expected to have a small vertical component of the flow, although some motion in the vertical direction is inevitable since the piston is moving vertically. Figure 103 shows that the average tumble flow has a larger vertical component directed toward the engine head for most of the compression stroke, until the piston nears top dead center. After 40° before top dead center both swirl and tumble flows have approximately equal vertical components directed toward the engine head.

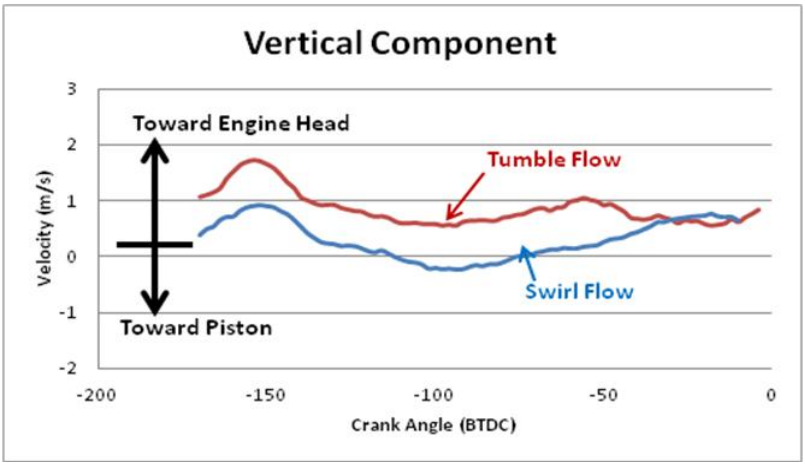

Figure 103: Throughout the compression stroke, the average tumble flow has a larger vertical flow component directed toward the engine head than the swirl flow, until 40° before top dead center, when the magnitudes become approximately equal.

When performing two-component PIV measurements, no additional bulk flow information beyond that presented in Figure 102 and Figure 103 is available. But because SC3D-PTV measures all three velocity components, the average out-of-plane bulk flows can also be compared. As seen in Figure 104, the difference is significant. The swirl flow has minimal average out-of-plane velocity, as expected for a flow rotating about the cylinder axis. However, during most of the compression stroke the average out-of-plane component of the tumble flow is the largest of the three components. The large out-of-plane component of the tumble flow toward the center of the cylinder is

likely due to precession of the primary tumble vortex as it is compressed. Depending on the relative positions of the measurement volume and center of rotation, the out-of-plane flow component can be significant.

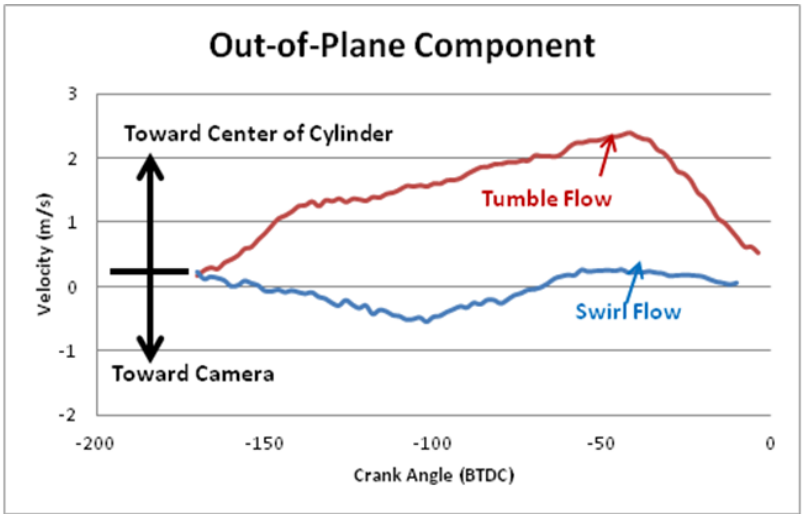

Figure 104: The average out-of-plane component of the swirl and tumble flows differ significantly. The average out-of-plane magnitude for the swirl flow remains near zero throughout the compression stroke, but the tumble flow, on average, maintains a large magnitude flow toward the center of the cylinder.

Averaged Flow Fields

In addition to the bulk flow values, multi-cycle-averaged flow fields from a single crank angle also add to the understanding of the in-cylinder flow field. Averages computed from the SC3D-PTV data show distinct differences for the swirl and tumble conditions. As the piston is near the middle of the compression stroke, the high-swirl case shows vorticity, but the tumble flow is largely uni-directional.

Figure 105 shows the direction of the swirl flow and the relative locations of the measurement volume and a projection from the piston crown. While the piston is near the middle of its travel, a vortex is visible in the averaged flow fields for the swirl

condition, likely due to interaction between the dominant flow structure and the contoured piston.

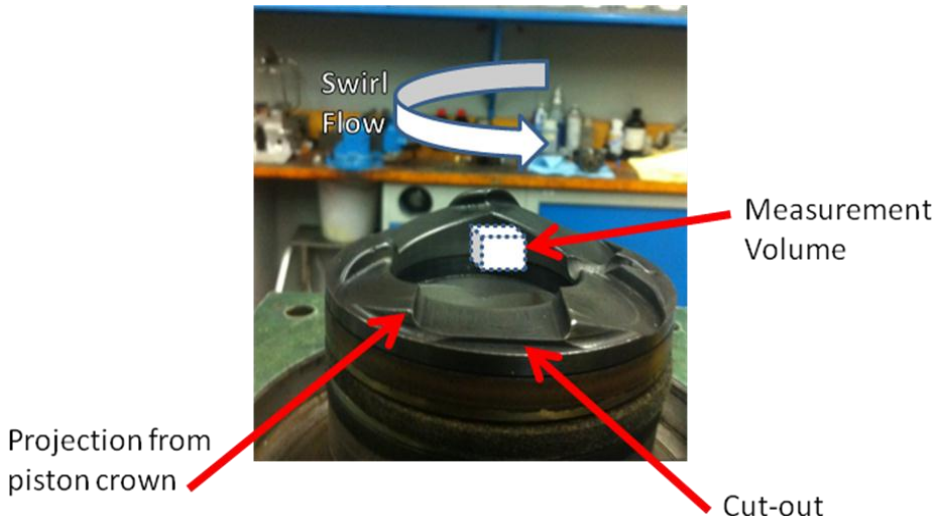

Figure 105: The swirl flow interacts with a projection from the piston crown in the vicinity of the measurement volume.

For the swirl case, horizontal translation of the vortex center with respect to the measurement volume is evident over the course of many crank angles. Figure 106 shows three flow fields spanning 18 crank angles (in 6 degree increments) from the high swirl case, and the approximate position of the vortex center. The flow in the out-of-plane direction is minimal.

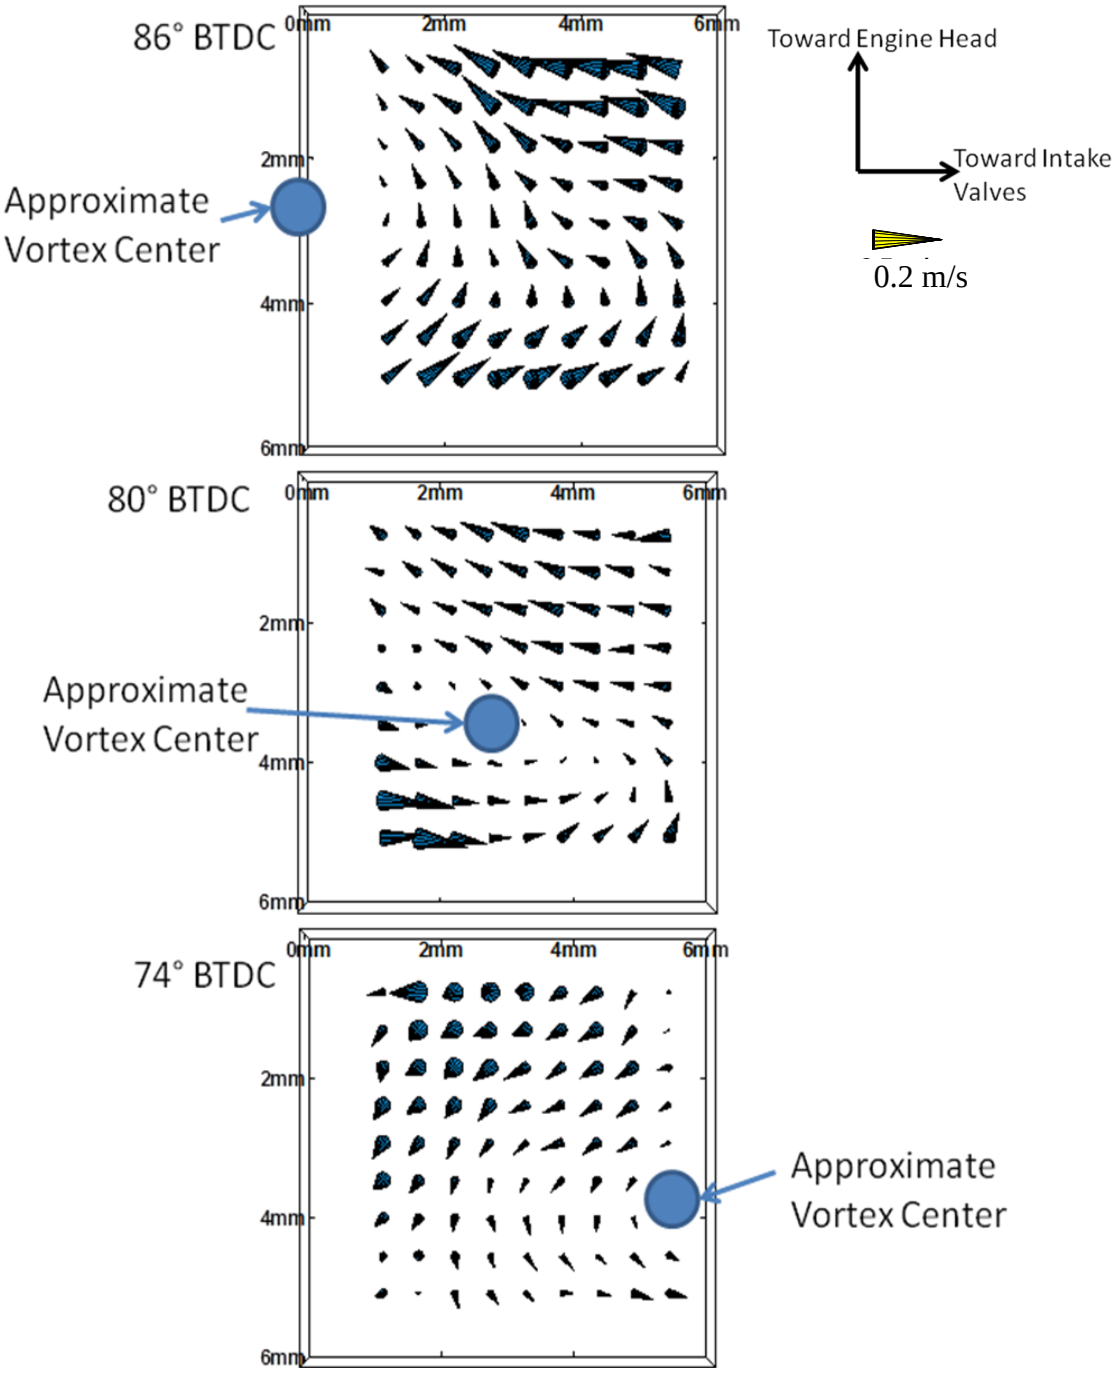

Figure 106: A vortex is seen to translate across the measurement volume over the course of numerous crank angles for the averaged swirl case.

Unlike the vortex seen in the swirl case, the tumble case shows a highly directed and uniform flow over the same crank angles, as seen in Figure 107. The velocity of the flow is also much higher. The main component of the flow is in the out-of-plane direction, toward the center of the cylinder, which would have been missed entirely if these measurements were performed with a 2D technique. The flow also has a horizontal component in the direction of the exhaust valves.

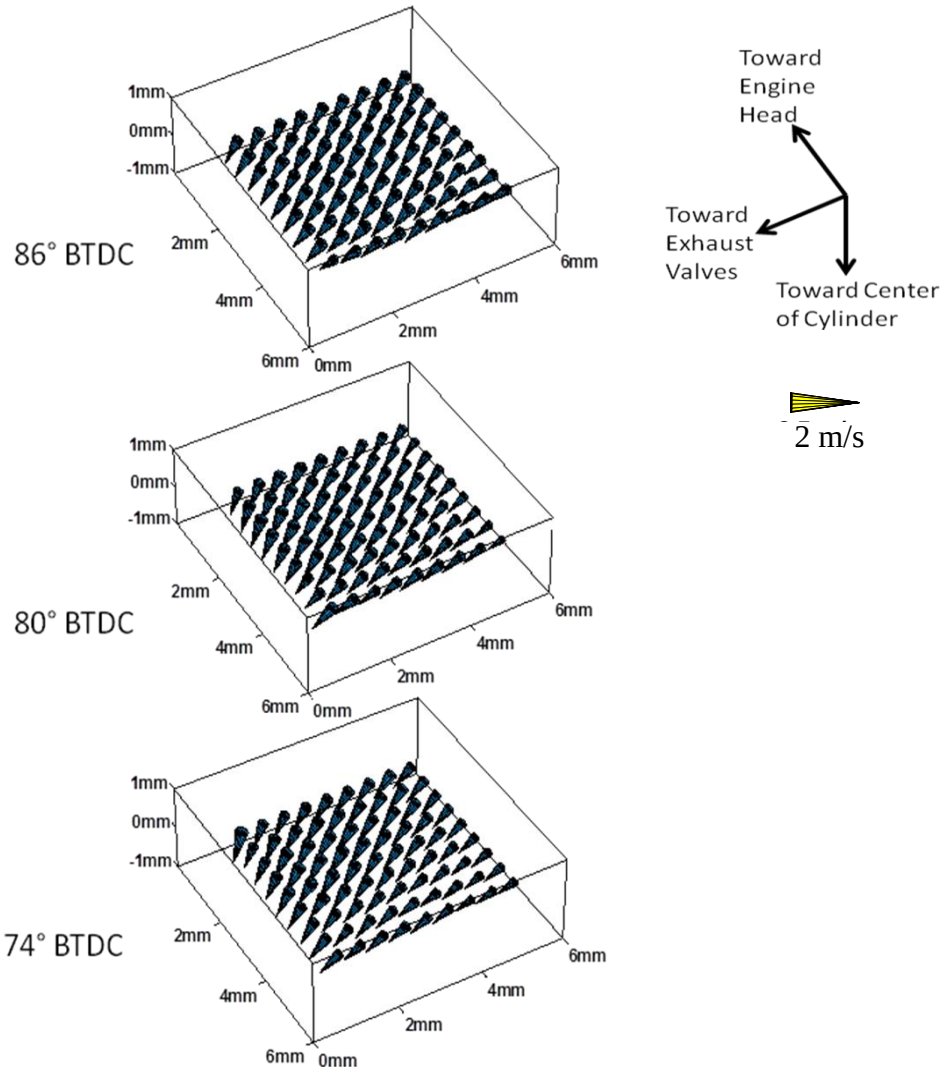

Figure 107: The average flow during the tumble case is highly uniform and directed toward the center of the cylinder during the same crank angles that the swirl flow shows vorticity.

10.2.2 Instantaneous flow fields

As shown above for multi-cycle averages, studying all three components of the flow, rather than only two, can elucidate flow features that are lost during two-component measurements. Capturing three velocity components is also useful when studying individual cycles.

Pseudo Point Source

A fluid point source is a mathematical idealization from which fluid flows outward in every direction. Point sources can be used to model certain phenomena, but are not found in reality. To demonstrate the shortcomings of two-component fluid measurements, Figure 108 shows the two-component projection of an instantaneous, three-component SC3D-PTV measurement. When only the 2D projection is viewed, the data appear to represent a pseudo-point source. The flow is outward from the indicated point in almost every direction, and the flow into the point is significantly less than the flow outward. Assuming the fluid is not locally compressing or expanding, such a flow is not physically possible.

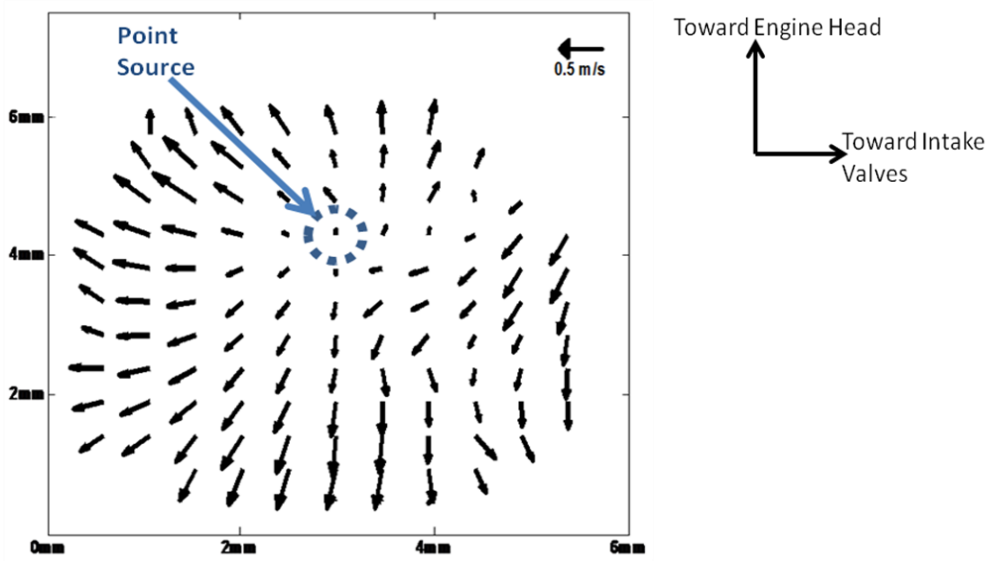

Figure 108: The 2D projection of instantaneous SC3D-PTV data appears to represent a point source, which is a physical impossibility.

Using a cone plot, viewed at the same angle, to visualize the third flow component reveals that some sections of the flow field have a large out-of-plane component, as seen in Figure 109.

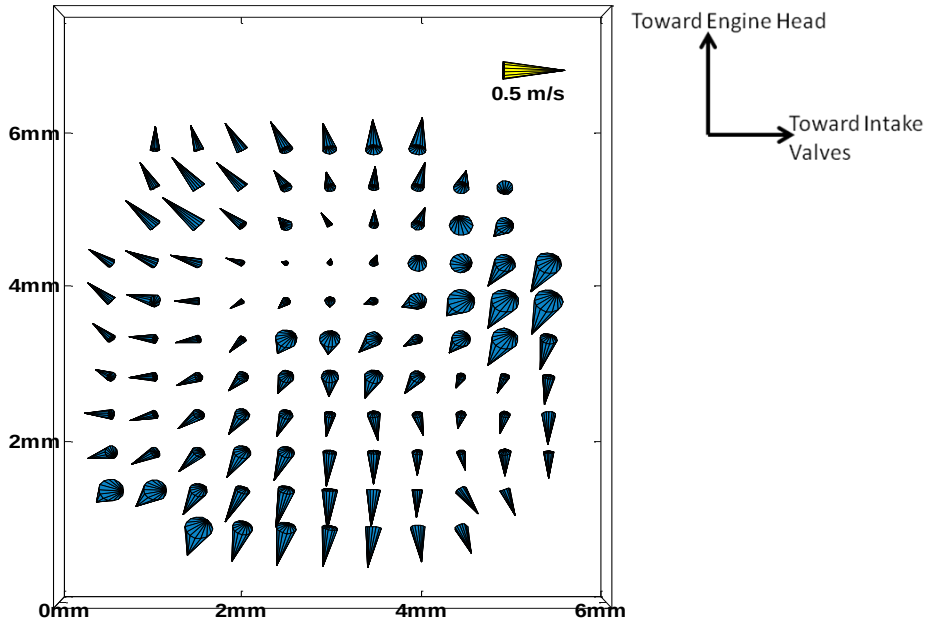

Figure 109: Using a cone plot to visualize the third component of the flow field shows that what appeared as a 2D point source in Figure 108 has large out-of-plane flows in some areas.

The three-dimensional nature of this flow field can be seen more clearly in an off-axis projection of the cone plot, as in Figure 110.

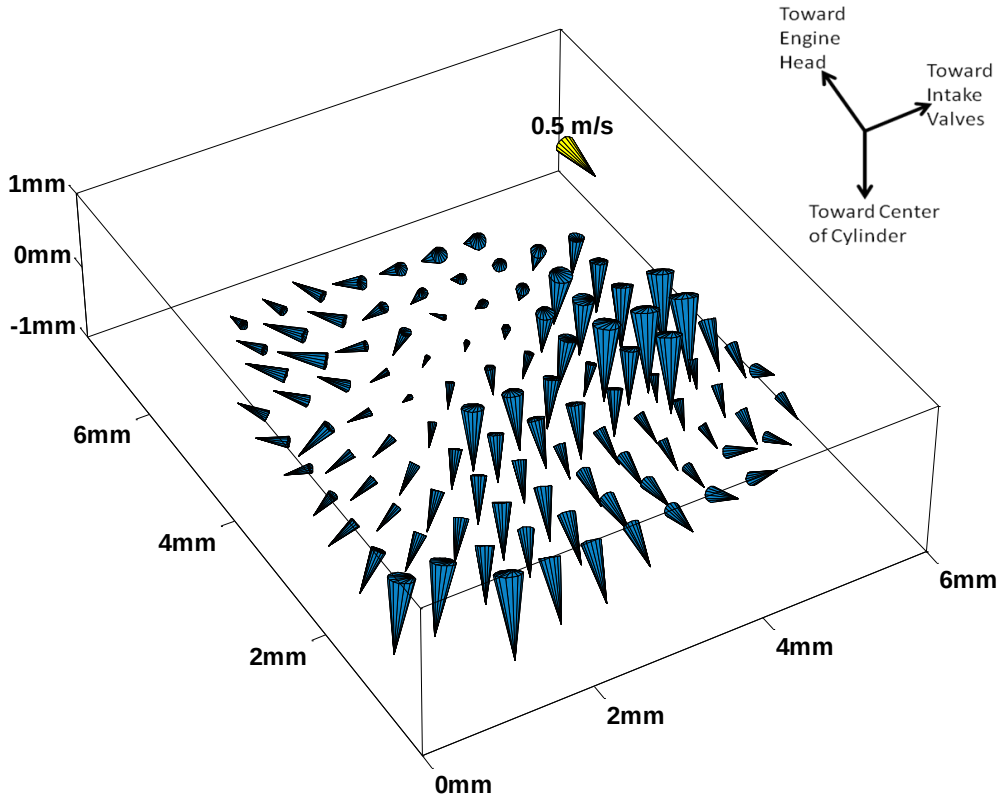

Figure 110: An off-axis projection of the same data shown in Figure 109 more clearly shows the instantaneous out-of-plane flow component.

Because SC3D-PTV not only measures all three components of the flow, but also measures them over a volume, multiple planes of the flow field can be examined. As seen in Figure 111, the data shown in Figure 110 is merely a planar slice of a larger, coherent flow field with an out-of-plane component and some vorticity. With the extra information added by having all three flow components over a volume, it is clear that what appeared to be a physically-impossible point source when viewed in two dimensions (Figure 108) is actually a perfectly reasonable flow field and does not violate any physical laws.

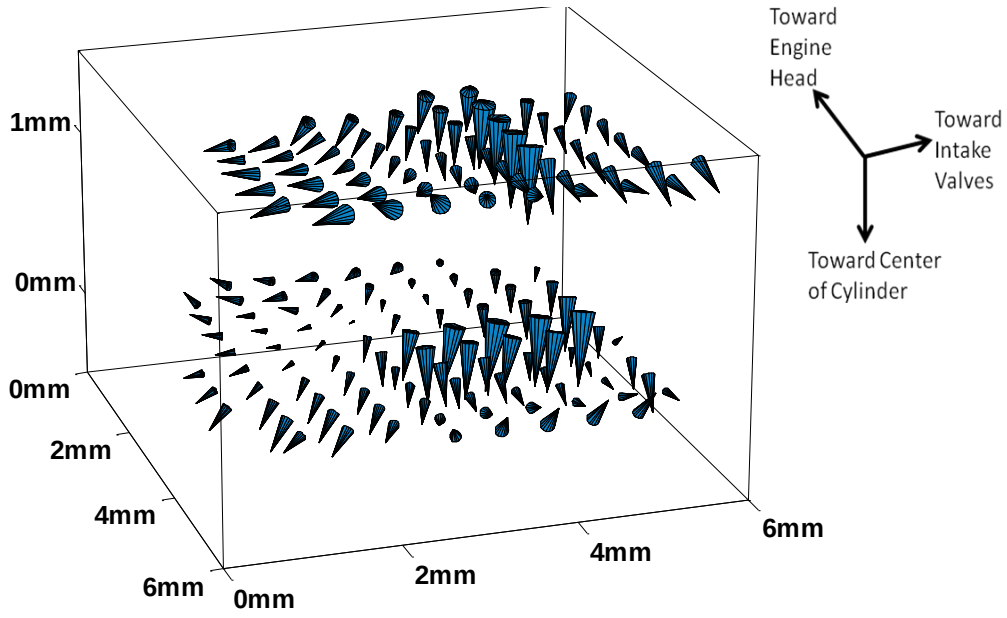

Figure 111: Viewing multiple planes simultaneously shows that what appeared as a physically-impossible point source in two dimensions is in fact a smooth and coherent flow when all three flow components are measured over a volume.

Three-Dimensional Vortex Structure

Quantifying vortices is an important part of fluid flow analysis since the size and velocity of vortex structures are often used to describe the statistics of turbulent flow fields. Figure 112 shows an instantaneous flow field captured using two-component PIV that contains a vortex with a diameter of about 5 millimeters.

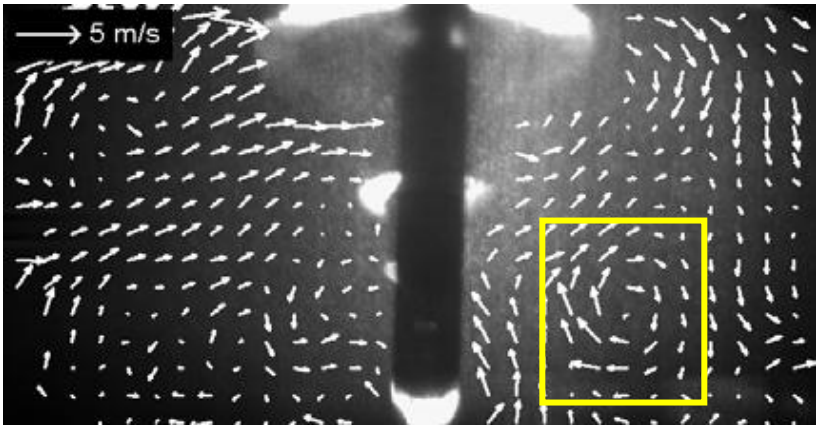

Figure 112: This instantaneous two-component PIV measurement shows a vortex with a diameter of approximately 5 millimeters. (image from Peterson, 2009)

Similar structures were seen in the SC3D-PTV data. Figure 113 shows the two component projection of a vortex captured inside the engine using SC3D-PTV. As in Figure 112, the diameter of the 2D projection of this vortex is about 5 millimeters, and the in-plane magnitudes are of the same order.

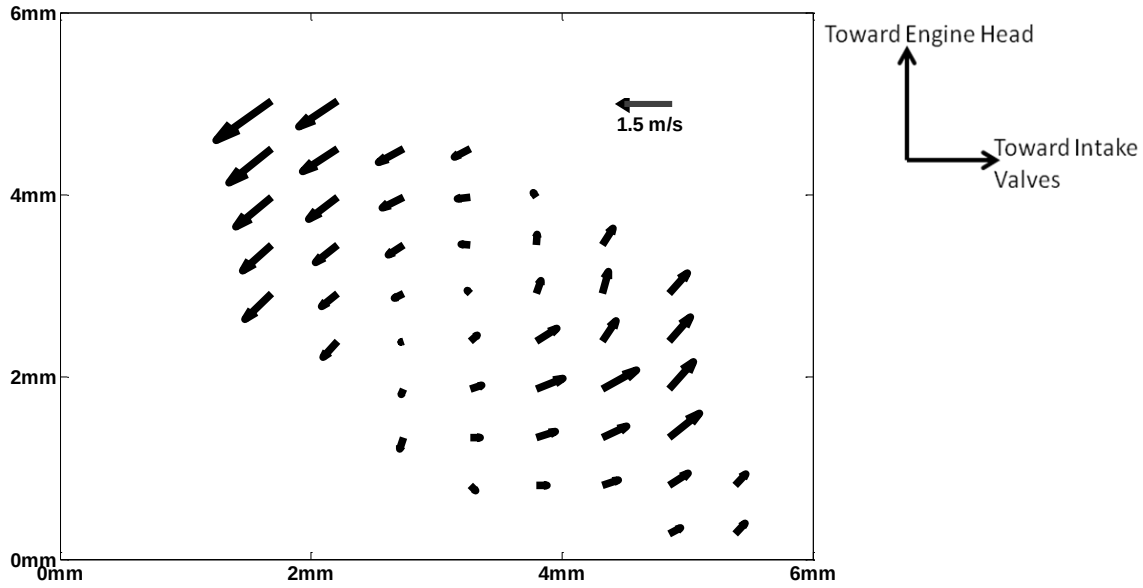

Figure 113: This 2D projection of instantaneous SC3D-PTV data shows a vortex similar to the one seen in Figure 112.

Figure 114 makes the third component of the flow visible using a cone plot, showing that some parts of the vortex have a positive out-of-plane component, and other parts have a negative out-of-plane component.

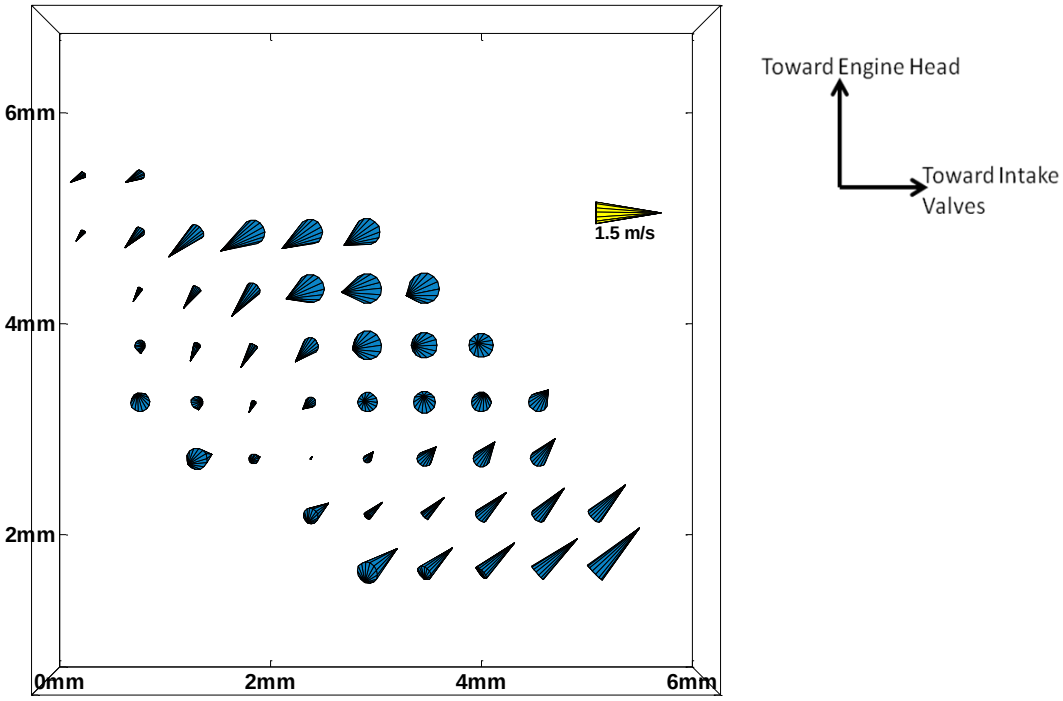

Figure 114: When all three flow components are visible, some parts of the instantaneous vortex are seen to have a positive out-of-plane component, while in other areas the out-of-plane component is negative.

The three-dimensionality of this vortex is made more apparent by viewing an off-axis projection of the cone plot, as in Figure 115. From this viewing angle, it is clear that the axis of rotation of the vortex is not parallel to the out-of-plane direction and the out-of-plane velocity is not negligible. When computed using only two components of the flow, the average magnitude of the vectors is about 1 meter per second, but when using all three velocity components, the actual average magnitude of the vortex vectors is calculated to be about 2.2 meters per second, more than double what was measured by the 2D projection.

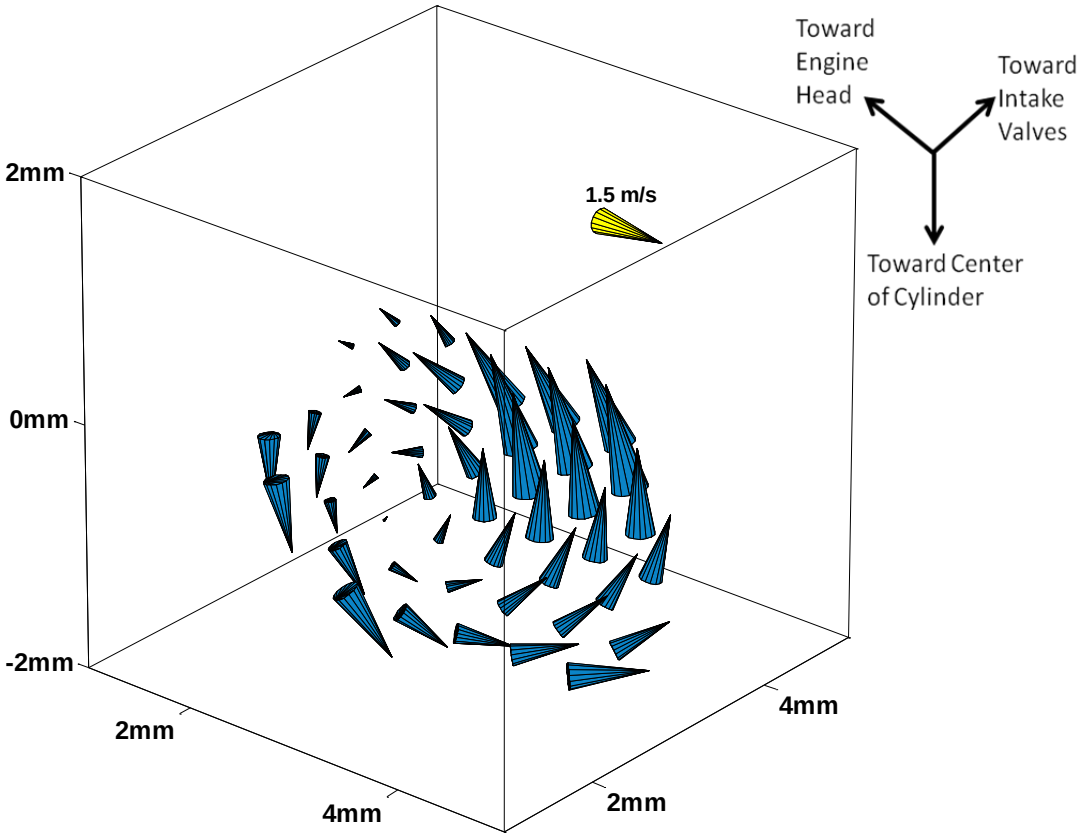

Figure 115: Viewing an off-axis projection of the data presented in Figure 114 shows that the axis of rotation of this vortex is not parallel to the out-of-plane flow component, and the velocity of the vortex is much higher than what was measured by the 2D projection of the same vortex.

With volumetric SC3D-PTV data, multiple planes within the measurement volume can be examined simultaneously. The measurement plane shown in Figure 116 is 2 millimeters above the plane shown in Figure 114 and Figure 115, in the out-of-plane direction. The vortex structure is still clearly visible, and the axis of rotation is also angled with respect to the out-of-plane direction.

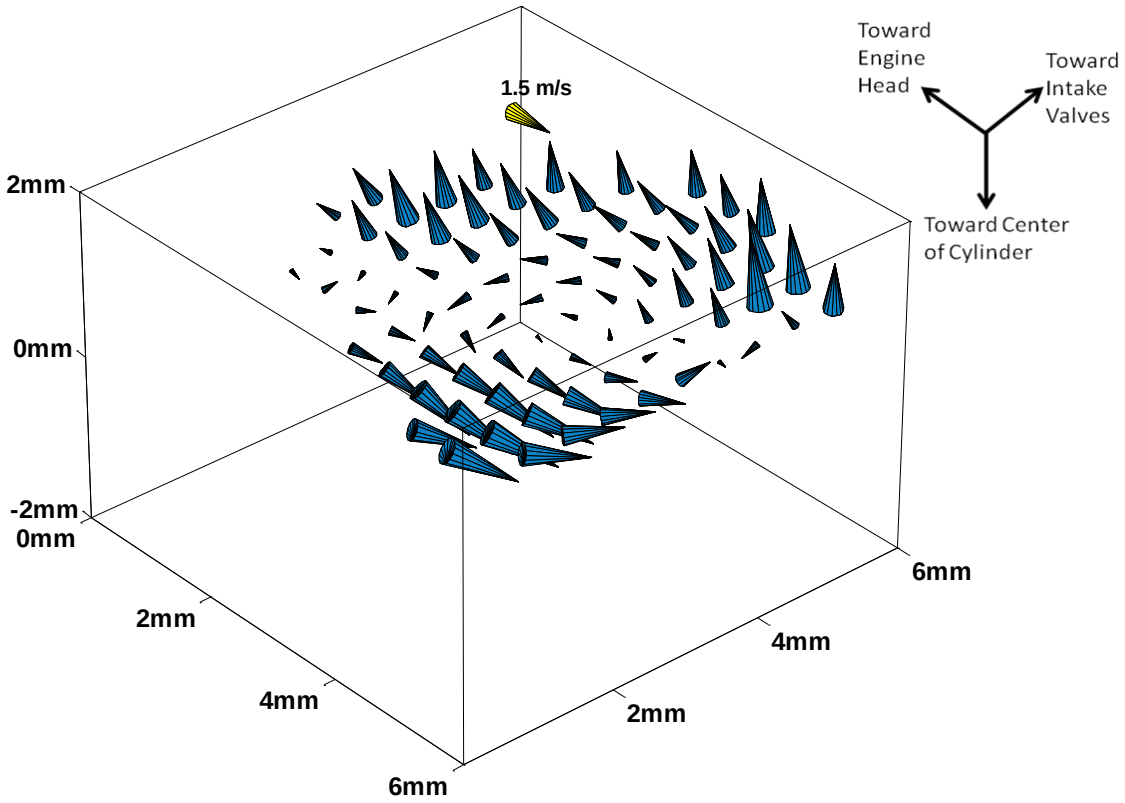

Figure 116: Because SC3D-PTV is a volumetric technique, multiple planes within the measurement volume can be viewed simultaneously, showing that the vortex structure seen in Figure 115 is also present at another plane 2 millimeters away.

When the information in Figure 115 and Figure 116 is shown on the same plot, as in Figure 117, it is clear that the vortices seen at the two different measurement planes are simply two different slices of a larger, coherent vortex structure extending over much of the measurement volume.

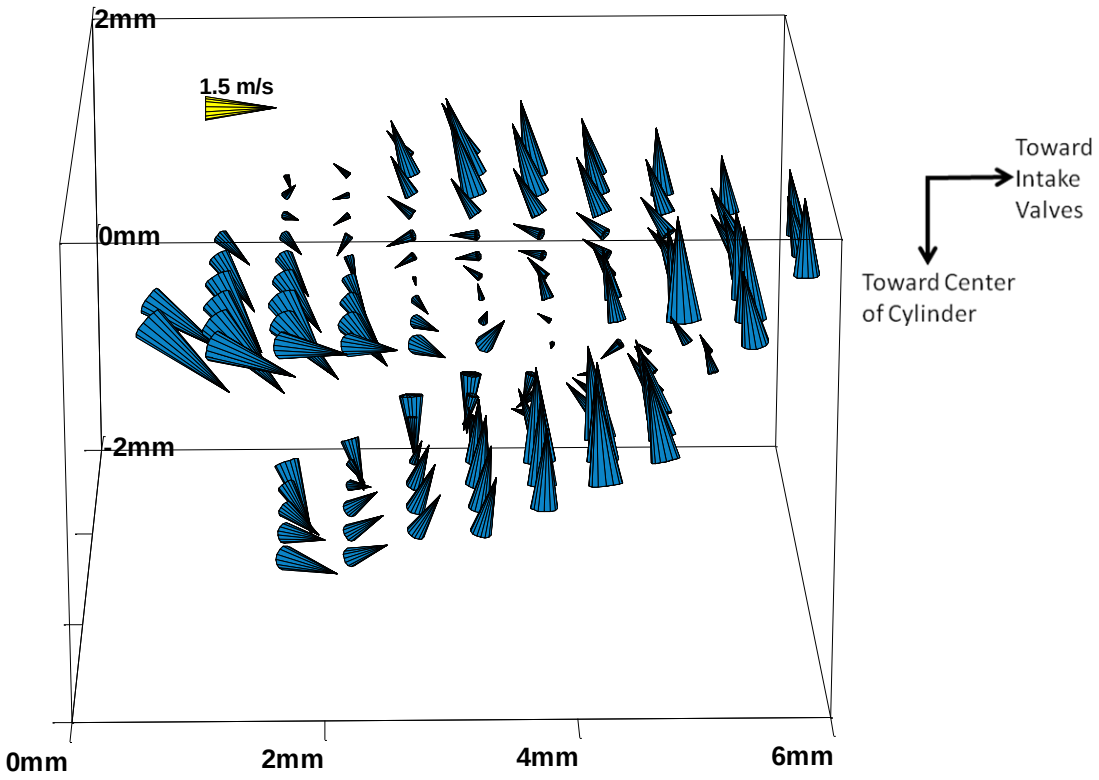

Figure 117: The vortices seen in Figure 115 and Figure 116 are seen to be slices of the same extended vortex structure when viewed simultaneously.

The volumetric, three-component results given by SC3D-PTV provide much more information than either two-component or three-component measurements within a single plane. The SC3D-PTV results show that both the size and the velocity of the vortex were underestimated by the two-component planar results shown above in Figure 113, and the two-component results were also unable to determine the axis of rotation of the vortex, which SC3D-PTV results are capable of showing.

10.3 Summary

The data presented in this chapter show the clear advantages of using SC3D-PTV to study engine flows when compared to a two-component, or even three-component, planar technique. The out-of-plane component of the bulk flow can be the largest component, and even in the average, smaller scale structures within the flow are possible, necessitating the measurement of all three flow components simultaneously for a full understanding to the in-cylinder flow field. Measuring all three velocity components is also important for studying instantaneous flow fields, since two-component techniques can give the impression of unrealistic results, or significantly misjudge the structures within a flow field. Capturing volumetric, rather than planar, data allows the three-dimensionality of common flow structures like vortices to be understood much more fully.

From the data presented in this chapter, it is clear that SC3D-PTV represents a significant step forward in the field of engine flow measurements.

CHAPTER 11

THESIS SUMMARY AND FUTURE WORK

11.1 Thesis Summary

This thesis presented the development of single-camera, three-dimensional particle-tracking velocimetry (SC3D-PTV), a method for measuring 3D air flow inside an optically-accessible internal combustion engine.

The development of this technique was motivated by a trio of factors. First, the widespread use of internal combustion engines consumes large amounts of energy and creates large amounts of pollution. Maximizing the efficiency and minimizing the environmental impact of engines is critically important. Second, in-cylinder air flow has a large impact on both engine efficiency and pollutant formation. Studying and understanding this air flow is necessary for improving engine operation. Third, because engine flows are three-dimensional, evolve quickly, and many areas of the combustion chamber have limited optical access, existing flow measurement techniques are ill-suited to the study of engine flows.

In response to these motivating factors, the SC3D-PTV method developed in this thesis can perform high-speed, high-resolution 3D flow measurements with very limited optical access, promote a greater understanding of in-cylinder flow, and lead to improved operation of internal combustion engines.

The optical element used for SC3D-PTV is similar to a stereo-microscope. A single large lens and two smaller lenses are used to create two parallel imaging systems within a single housing. An additional one-to-one imaging optic with an internal aperture can be used to avoid cross-talk. The two imaging sub-systems view the same measurement volume from different angles, but share a focal plane without perspective distortion. Objects within the measurement volume are imaged twice, once by each imaging sub-system, and under the proper assumptions the two images are identical (but in practice the images are highly similar without being completely identical). The relative positions of the two images indicate the 3D position of the object. 3D velocities are measured by imaging an object at successive points in time, and computing the change in 3D position for known time intervals.

Two images of a particle are necessary to determine the 3D position, and four images of a particle are needed to calculate the 3D velocity of a particle. When imaging multiple particles within the measurement volume, the four images corresponding to a single particle must be properly matched. The SC3D-PTV algorithm performs the matching of the simultaneously-acquired images and successively-acquired images in tandem, and as a result, returns more valid matches than an equivalent algorithm applied to simultaneous and successive matching separately.

The matching processes are performed by first using knowledge of the geometry of the imaging optics and the maximum flow velocity to define search areas within the raw images where matching particle images must reside. Using the search areas, the particle images are grouped into sets of four particles images each, where a set contains sufficient information to compute a single three-component velocity vector. The sets are

checked for validity by comparing the size, shape, and intensity of the particle images within the set.

The appropriate definitions for particle image similarity were determined by analyzing sparsely seeded data, and selecting valid particle image sets using relative position alone. These sets were compared to a control group of purposefully mismatched particle image sets. Thresholds on shape, size, and intensity for which the valid sets could be largely separated from the invalid sets were calculated.

The degree to which the algorithm separates valid and invalid particle image sets is largely a function of seed density and particle image size. A particle image can only be a member of one valid set, but as seed density increases, each particle image is a member of more and more invalid sets, increasing the likelihood that one of the invalid sets will contain highly similar particle images and exceed the validity threshold. As the size of the particle images increases, particle images overlap more frequently. For sufficient overlapping, the two particle images cannot be separated and both must be discarded, decreasing the number of vectors that can be computed.

Validation of the SC3D-PTV method was obtained by analyzing a single experimental data set from a simple flow with both the SC3D-PTV algorithm and a stereoscopic PIV algorithm. Stereoscopic PIV is a well-known and widely accepted flow measurement technique, and the quantitative results of the two analyses matched closely, despite fundamental differences in the assumptions behind the two algorithms.

After validation, the SC3D-PTV technique was applied to the air flow inside a motored engine. Because methods like PIV or stereoscopic PIV are commonly employed to study engines flows, the primary emphasis of the engine measurements was on the

advantages of the volumetric, three-component results given by SC3D-PTV compared to the planar, two- and three-component results given by PIV and stereoscopic PIV, respectively.

Engine flows are three-dimensional and turbulent. Volumetric, three-component measurements, like SC3D-PTV, are necessary to fully resolve these flows. When measuring only two flow components, very misleading results can be obtained. The dominant direction of the flow can be missed entirely, and the 2D results may even appear to represent physically-impossible flow fields. Measuring three velocity components within a plane lends additional insight, however the complex structures present within engine flows are not limited to a plane, but rather extend in all three dimensions. Thus, planar results still fail to capture the full nature of engine flows.

Because SC3D-PTV can fully quantify the three-dimensional flow structures found within engines, even with very limited optical access, the development of SC3D-PTV significantly advances the study of engine flows and offers the chance to gain insights that would be impossible with existing measurement techniques.

11.2 Future Work

Proof-of-principle has been solidly demonstrated for SC3D-PTV, but work remains with regard to increasing the flexibility of the technique, making the technique more user friendly, optimizing the results, and collecting useful data.

With the present hardware design, both the working distance and the size of the measurement volume are fixed. While fixed parameters were sufficient for initial testing, adding the ability to change these measurement parameters is necessary to apply SC3D-PTV to a wider variety of flows.

The hardware used for SC3D-PTV is already user friendly. The optical elements are contained within a single, sealed housing, and need only be aligned once. The housing is small, and thus requires minimal space in the laboratory. The pre-aligned system can be attached to a camera and used for data collection in a matter of minutes. Although greater functionality could be added, the current system has many positive features.

However, the SC3D-PTV analysis is currently performed with a command-line based code that requires knowledge of the code on the part of the user. Some aspects of the analysis, such as the determination of intensity and similarity thresholds, file handling, and some portions of the flow visualization are done manually. For the technique to be accepted widely, a graphical user interface and the automation of the majority of the analysis tasks would be necessary. This could be implemented either as a stand-alone program, or integrated with existing software.

Both the analysis code and the physical hardware are fully functional, but the possibility for optimization remains. Adding functionality to the code for separating overlapped particle images would increase the maximum possible spatial resolution of the technique. Additional vector validation could be added through local median filtering or checking for conservation of mass and momentum. Resolution could also be increased by minimizing the diffraction through the optical elements, thereby decreasing particle size and reducing the number of overlapped particle images.

The analysis code was written in the Matlab scripting language, and contains some redundancies that were included for error-checking while the code was written.

Moving the code to a compiled language and removing the redundancies would improve processing times.

Although improvements to the technique are possible, a measurement technique is not an end in itself. A measurement technique is only useful insofar as the results given by the technique are useful, therefore SC3D-PTV must be pushed beyond the proof-of-principle stage and applied rigorously to engine flows. Because engine flows have been studied extensively with two-component PIV, many aspects of the flow that are particularly relevant to engine operation have already been identified. The flow near the spark plug at the time of ignition can affect early flame growth, and the flows created by the interaction of the fuel spray and the spark plug vary significantly with spark plug orientation (Peterson, 2009). The flow downstream of the spark plug has been shown to be important in the study of misfires in SG-SIDI engines (Peterson, 2010). Boundary layer flow fields affect heat transfer rates in the cylinder (Alharbi and Sick, 2010). In each of these examples, measuring the 2D flow advanced the knowledge of engine operation, but in each of the cases, the flow was not strictly two-dimensional. Thus, potentially significant information was not collected simply due to the lack of an appropriate measurement technique. With the development of SC3D-PTV, the full 3D nature of all of these flow fields can be studied.

BIBLIOGRAPHY

Alharbi, A.Y., Sick, V., 2010, "Investigation of boundary layers in internal combustion Engines using a hybrid algorithm of high speed micro-PIV and PTV," Exp. In Fluids, **49**, 949-959

Arcoumanis, C., Whitelaw, J.H., 1987, "Fluid Mechanics of Internal Combustion Engines: A Review," Proc. I. Mech. Eng., **201**

Arroyo, M.P., Greated, C.A., 1991, "Stereoscopic particle image velocimetry," Meas. Sci. Tech., **2**, 1181-1186

Atkinson, C., Soria, J., 2009, "An efficienct simultaneous reconstruction technique for tomographic particle image velocimetry," Exp. In Fluids, **47**, 553-568

Bown, M.R., MacInnes, J.M., Allen, R. W. K., Zimmerman, W. B. J., 2006, "Three-dimensional, three-component velocity measurements using stereoscopic micro-PIV and PTV," Meas. Sci. Tech., **17**, 2175-2185

Brücker, C., 1997, "3D scanning PIV applied to an air flow in a motored engine using high-speed digital video," Meas. Sci. Tech., **8**, 1480-1492

Calendini, P.O., Duverger, T., Lecerf, A., Trinite, M., 2000 "In-Cylinder Velocity Measurements with Stereoscopic Particle Image Velocimetry in a SI Engine," SAE Paper 2000-01-1798

Cho, H.-S., Lee, J.-H., Lee, K-Y., 1998 'Measurement of HC Concentration Near Spark Plug and Its Effects on Combustion," SAE Paper 981431

Drake, M.C. and Haworth, D.C., 2007, "Advanced gasoline engine development using optical diagnostics and numerical modeling," *Proc. Combust. Inst.* **31** 99–124

Eaton, J.K., Fessler, J.R., 1994, "Preferential Concentration of Particles by Turbulence," Int. J. Multiphase Flow, **20**, 169-209

Eichenberger, D.A., Roberts, W.L., 1999, "Effect of Unsteady Stretch on Spark- Ignited Flame Kernel Survival," Combustion and Flame, **118**:469-478

Elsinga, G.E., Scarano, F., Wieneki, B., van Oudheusden, B.W., 2006, "Tomographic particle image velocimetry," Exp. in Fluids, **41**, 933-947

Fansler, T.D., French, D.T., Drake, M.C., 1995, "Fuel Distributions in a Firing Direct-Injection Spark-Ignition Engine using Laser-Induced Fluorescence Imaging," SAE Paper 950110

Fansler, T.D., Drake, M.C., Düwell, I., Zimmerman, F.P., 2006, "Fuel-spray and Spark-plug Interactions in a Spray-Guided Direct Injection Gasoline Engine," 7th Intl.Symp. on Internal Combustion Diagnostics, Baden-Baden.

Fansler, T.D., Drake, M.C., 2006, "'Designer Diagnostics' for developing direct-injection gasoline engines" J. Phys.: Conf. Ser. **45**, 1-17

Fansler, T.D., Drake, M.C., Böhm, B., 2008, "High-Speed Mie-Scattering Diagnostics for Spray-Guided Gasoline Engine Development," 8th International Symposium on Internal Combustion Diagnostics, June 10–11, 2008, Baden-Baden

Heywood, J.B., 1986, "Fluid Motion Within the Cylinder of Internal Combustion Engines-the 1986 Freeman Scholar Lecture," Journal of Fluids, Vol. 109

Hinsch, K.D., 2002, "Holographic particle image velocimetry," Meas. Sci. Tech., **13**, R61-R72

Jeon, C.-H., Chang, Y.-J., Cho, K.-B., Kang, K.-Y., 1998, "Effects of Intake Ports on In-cylinder Flow and Lean Combustion in a 4-Valve Engine," SAE Paper 981048

Konrath, R., Schröder, W., Limberg, W., 2002, "Holographic particle image velocimetry applied to the flow within the cylinder of a four-valve internal combustion engine," Exp. In Fluids, **33**, 781-793

Law, C.K., 2006, Combustion Physics, Cambridge University Press

Lee, D.W., 1939, "A Study of Air Flow in an Engine Cylinder," NACA Rep. No. 653

Lindken, R., Westerweel, J., Wieneke, B., 2006, "Stereoscopic micro particle image velocimetry," Exp. In Fluids, **41**, 161-171

Lumley, J.L., 2001, "Early Work on Fluid Mechanics in the IC Engine," Annu. Rev. Fluid Mech. **33**, 319-338

Maas, H.G., 1992, "Complexity analysis for the establishment of image correspondences of dense spatial target fields," Intern. Arch. Of Photogrammetry and Remote Sensing, **XXIX**, Part B5, 102-107

Maas, H.G., Gruen, A., Papantoniou, D., 1993, "Particle tracking velocimetry in three-dimensional flows. Part I: Photogrammetric determination of particle coordinates ," Exp. In Fluids, **15**, 133-146

Malik, N.A., Dracos, Th., Papantoniou, D.A., 1993, "Particle tracking velocimetry in three-dimensional flows. Part II: Particle tracking," Exp. In Fluids, **15**, 279-294

Meng, H., Pan, G., Pu, Y., Woodward, S.H., 2004, "Holographic particle image velocimetry: from film to digital recording," Meas. Sci. Tech., **15**, 673-685

Mueller, C.J., Driscoll, J.F., Reuss, D.L., Drake. M.C., Rosalik, M.E., 1998, "Vorticity Generation Attenuation as Vortices Convect Through a Premixed Flame," Comb. And Flame, **112**, 342-258

Pereira, F., Gharib, M., 2002, "Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows," Meas. Sci. Tech., **13**, 683-694

Peters, N., 2000, Turbulent Combustion, Cambridge University Press, Cambridge

Peterson, B., 2010, "High-Speed Fuel and Flow Imaging to Investigate Misfires in a Spray-Guided Direct-Injection Engine," Ph.D. Thesis, U. Michigan

Peterson, K. H., 2009, "The effect of spark plug orientation and geometry on flow development in close proximity to the spark plug in a spray-guided, stratified-charge spark-ignition direct injection engine," MS Thesis, U. Michigan

(Note that sections of text from Chapter 1 and Chapter 2 are taken directly from Peterson, 2009 with only minor modifications)

Pischinger, S., Heywood, J.B., 1990, "A Model for Flame Kernel Development in a Spark-Ignition Engine," Twenty-Third Symposium (International) on Combustion, 1033-1040

Probst, D.M., Ghandi, J.B., 2003, "An Experimental Study of Spray Mixing in a Direct Injection Engine," Int. J. Engine Res., **4**, No. 1

Willert, C.E., Gharib, M., 1992, "Three-dimensional particle imaging with a single camera," Exp. In Fluids, **12**, 353-358

Yu, C.H., Yoon, J.H., Kim, H.B., 2009, "Development and validation of stereoscopic micro-PTV using match probability," Journal of Mechanical and Science Technology, **23**, 845-855

Zhao, F.-Q., Lai, M.-C., Harrington, D., 1997, "A Review of Mixture Preparation and Combustion Control Strategies for Spark-Ignited Direct-Injection Gasoline Engines," SAE Paper 970627