FLOW IMAGING

NOEL T. CLEMENS

The University of Texas at Austin
Austin, TX

INTRODUCTION

Imaging has a long history in fluid mechanics and has proven critical to the investigation of nearly every type of flow of interest in science and engineering. A less than exhaustive list of flows where imaging has been successfully applied would include flows that are creeping, laminar, turbulent, reacting, high-temperature, cryogenic, rarefied, supersonic, and hypersonic. The wide range of applications for flow imaging is demonstrated by the recent development of techniques for imaging at micro- and macroscales. For example, (1) and (2) report imaging velocity fields in 100-μm channels, and (3) describes a schlieren technique for imaging density gradient fields around full-scale supersonic aircraft in flight for the study of sonic booms. Impressively, the range of flow length scales spanned by these techniques is more than six orders of magnitude.

Traditionally, flow imaging has been synonymous with "flow visualization," which usually connotes that only qualitative information is obtained. Examples of flow visualization techniques include the imaging of smoke that has been introduced into a wind tunnel or vegetable dye introduced into a water flow. Owing to the complex and often unpredictable nature of fluid flows, flow visualization remains one of the most important tools available in fluid mechanics research. Excellent compilations of flow visualization images captured in a number of different flows can be found in (4) and (5). Modern flow imaging, however, goes far beyond qualitative flow visualization. Advances in computer, laser, and digital camera technologies have enabled the development of imaging techniques for obtaining quantitative images of a large number of flow variables such as density, temperature, pressure, species concentration, and velocity. Image data of this type enable the computation of a number of quantities that are important in fluid mechanics research, including vorticity, strain rate, dissipation, and heat flux.

As an example of the power of flow imaging, consider Fig. 1, which shows a 3-D volume of the conserved scalar field in the far field of a turbulent water jet (6,7). The jet was seeded with a fluorescent dye, and the image volumes were captured by recording the fluorescence induced by a thin laser beam that was swept through the flow. The beam was swept rapidly in a raster fashion, and the fluorescent images were recorded by using a high-speed 2-D photodiode array. The resulting data volumes resolve the finest scales of mixing in three spatial dimensions and time, and when several such volumes are acquired sequentially, the data enable studying the temporal-evolution of the conserved scalar field. These data can yield details of the mixing process and even the complete 3-D, unsteady velocity vector field within the volume (7). This example shows that flow imaging is providing the type of multidimensional, multiparameter data that could be provided only by computational fluid dynamics not too long ago (8).

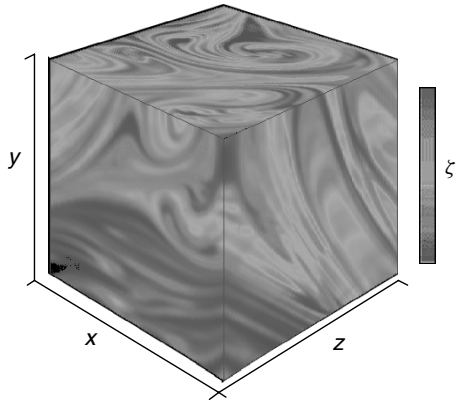

Figure 1. Three-dimensional rendering of the conserved scalar (ζ) field measured in a turbulent water jet using laser-induced fluorescence of a fluorescent dye seeded into the jet fluid. The cube is approximately 27 mm on each side, and the data resolve the finest scalar and vorticity scales in the flow. (Reprinted with permission from Quantitative Flow Visualization via Fully-Resolved Four-Dimensional Spatio-Temporal Imaging by W. J. A. Dahm and K. B. Southerland, in *Flow Visualization: Techniques and Examples*, A. J. Smits and T. T. Lim, eds., Imperial College Press, London, 2000.) See color insert.

Most imaging in fluid mechanics research involves planar imaging, where the flow properties are measured within a two-dimensional cross section of the flow. This is most often accomplished by illuminating the flow using a thin laser light sheet, as shown in Fig. 2, and then recording the scattered light using a digital camera. The laser light is scattered from either molecules or particles in the flow. The primary emphasis of this article will be on this type of planar laser imaging because it remains the cornerstone of quantitative imaging in fluid mechanics research. Furthermore, planar imaging is often a building block for more complex 3-D imaging techniques, such as that used to produce Fig. 1. Readers interested in

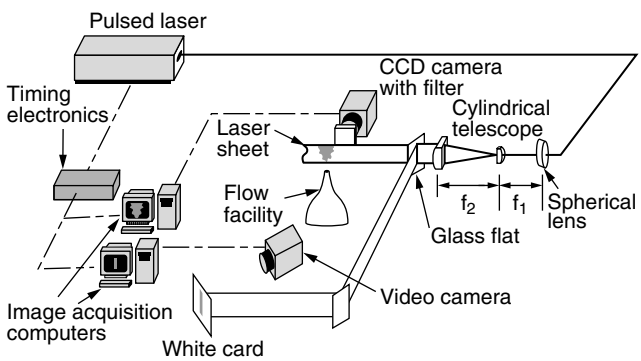

Figure 2. Schematic of a typical planar imaging experiment.

details of qualitative flow imaging techniques should note that several good references are available in the literature (5,9,10). Quantitative imaging is substantially more challenging than simple visualization because a greater degree of knowledge and effort are required before the researcher can ensure that the spatial distribution of the flow property of interest is faithfully represented in the image. The first part of this article will discuss some of the most important issues that need to be addressed in quantitative flow imaging. The article will end with a brief survey of primarily planar imaging techniques that have been developed. This survey will not be able to discuss all, or even most, of the techniques that have been developed, but hopefully readers will gain an appreciation for the wide range of techniques that can be applied to their flow problems.

BASIC PLANAR LASER IMAGING SYSTEMS

Lasers

Lasers are used almost universally in flow imaging, owing to their high brightness, coherence, excellent focusing properties, and the nearly monochromatic range of wavelengths at which they operate. Lasers can be either pulsed or continuous wave (CW); pulsed lasers are more commonly used because they provide high-energy pulses that are sufficiently short (e.g., 10 ns) to freeze the motion of nearly any flow. Most lasers used in flow imaging operate at visible or UV wavelengths (11). One of the main reasons for this is that until recently, there were few low-noise imaging arrays that operate outside of the UV-visible to near-IR wavelength range. Furthermore, some techniques, such as Rayleigh and spontaneous Raman scattering, increase in scattering efficiency as the frequency of the incident light increases, and therefore UV and visible lasers have a large advantage over IR sources. Furthermore, planar laser-induced fluorescence (PLIF) techniques typically involve the excitation of atomic/molecular electronic transitions, which occur primarily at UV and visible wavelengths for species of interest in fluid mechanics. The predominance of techniques in the visible/UV is by no means absolute, however, as recent advances in laser and camera technology have enabled the development of PLIF techniques that rely on the excitation of vibrational transitions at IR wavelengths (12).

The most widely used laser in flow imaging is the flashlamp-pumped neodymium: yttrium–aluminum garnet (Nd:YAG) laser, which emits in the infrared (1.06 μm), but whose output is usually frequency-doubled (532 nm), tripled (355 nm) or quadrupled (266 nm), using nonlinear crystals (13). Frequency-doubled Nd:YAG lasers are primarily used in particle image velocimetry (PIV), Rayleigh and Raman scattering, and for pumping tunable lasers. Nd:YAG lasers are essentially fixed frequency, but when injection seeded (a technique that is used primarily to obtain very narrow line width), they can be tuned across a narrow frequency range. This ability to tune is used extensively in a class of techniques called filtered Rayleigh scattering (described later). Flashlamp-pumped Nd:YAG lasers operate at repetition rates of a few tens of Hz and pulse energies of hundreds of millijoules at 532 nm. One drawback to flashlamp-pumped Nd:YAG lasers is that their repetition rates are typically much lower than the characteristic flow frequencies typical in most applications; the images are thus not temporally correlated and are effectively randomly sampled from the flow.

Excimer lasers provide high-energy pulses of UV light (e.g., hundreds of millijoules at hundreds of hertz) in a narrow range of frequencies that depend on the particular gas mixture that is used. The most commonly used wavelengths in flow imaging are 193 nm (ArF), 249 nm (KrF), 308 nm (XeCl), and 350 nm (XeF). Because Rayleigh and Raman scattering are more efficient at short wavelengths, excimers are particularly attractive for these techniques. Furthermore, versions are commercially available that have narrow line width and are tunable over a small range. These lasers can be used to excite the fluorescence from O_2 and NO (193 nm) and from OH (248 and 308 nm), without using a dye laser. Copper-vapor lasers are pulsed lasers that produce visible light simultaneously at two wavelengths (510 and 578 nm) and operate at high repetition rates (tens of kHz) but have relatively low pulse energies (a few mJ). Because of their high repetition rates, they have been used extensively for high-speed flow visualization (such as smoke scattering), but they are not as widely used as Nd:YAG lasers because of their relatively low pulse energies. Flashlamp-pumped dye lasers provide very high pulse energies (e.g., a few joules per pulse) but at repetition rates of just a few hertz. Because of their high pulse energies, they have been used primarily in imaging techniques where the signals are very weak, such as in spontaneous Raman or Rayleigh scattering imaging. For spectroscopic techniques, where it is necessary to tune the laser wavelength to coincide with an atomic/molecular absorption line, then laser-pumped dye lasers and more recently, optical parametric oscillators (OPO) are used. Both dye lasers and OPOs are typically pumped by Nd:YAG lasers, although dye lasers are also pumped by excimers.

The use of CW lasers is limited to low-speed flows (typically liquids) or to high-speed flows where only time-average measurements are desired. The reason is that they typically provide insufficient energy in times that are short enough to freeze the motion of most gas flows. For example, a 20-W CW laser provides only 0.02 mJ of energy in one microsecond, compared to a frequency-doubled Nd:YAG that can provide up to 1 J per pulse in 10 ns. The argon-ion laser is the most commonly used CW laser in flow imaging. The argon-ion laser has found a niche particularly for laser-induced fluorescence of dyes seeded into liquid flows.

Some techniques, such as cinematographic imaging, require high-repetition rate light sources such as copper-vapor or high-repetition rate diode-pumped Nd:YAG lasers. The latter achieve repetition rates up to hundreds of kHz by acousto-optic Q-switching of a continuously pumped Nd:YAG rod. The drawback of these high-repetition rate lasers is that they tend to have low energy per pulse (a few millijoules maximum), despite relatively high average power (e.g., 20–50 W). For

slower flows, electro-optically Q-switched diode-pumped Nd:YAG lasers can produce repetition rates of the order of a kilohertz and pulse energies of the order of tens of millijoules at 532 nm. Recently, a pulse-burst Nd:YAG laser has been developed that produces a train of up to 100 pulses at a rate as high as 1 MHz and individual pulse energies at 532 nm of about 25 mJ (14). In another technique, repeated Q-switching of a ruby laser (694 nm) was used to generate a train of 65 pulses at a rate of 500 kHz, where the energy for each of the 65 pulses was about 350 mJ (15). If this laser could operate continuously, its average power would be an impressive 175 kW. These laser systems are not currently available commercially, but they are particularly well suited for imaging very high-speed flows.

Optics

The focusing properties of laser beams are related to the mode structure of the beam, or specifically to the number of transverse electromagnetic modes (TEM) that characterize the energy flux field (16). A single-mode (TEM_{00}) laser beam has a Gaussian intensity distribution and is considered diffraction-limited. Note that in this article, the term "intensity" refers to the power density, or irradiance, of the laser beam (in units of W/m^2), whereas the term "fluence" refers to the energy density (in units of J/m^2). The focusing properties of diffraction-limited beams are described by Gaussian optics (16). Many laser beams, however, are not diffraction-limited because they contain many transverse modes. Multimode beams have higher divergence and poorer focusing characteristics than single-mode beams. The degree to which a beam is multimode is often specified by the M^2 value (pronounced "M-squared"), where the more multimode the beam, the higher the M^2 value, and where M^2 equals unity for a diffraction-limited beam. Many scientific lasers have M^2 values of 1 to 2, although many lasers, such as copper-vapor or high-power diode-pumped Nd:YAG lasers, can have M^2 values ranging from tens to hundreds. To see the effect of nonunity M^2, define the beam diameter d as twice the radius where the laser beam intensity drops to e^{-2} of the maximum. Assume that a laser beam whose initial diameter is d is focused by a spherical lens of focal length f. The resulting focal spot will have a diameter d_0 given by the relationship (17),

$$d_0 = \frac{4f\lambda M^2}{\pi d} \tag{1}$$

The focal spot diameter for a Gaussian (diffraction-limited) beam is $4f\lambda/(\pi d)$; thus Eq. (1) is the same as for a Gaussian beam, except that λ is replaced by λM^2. Equation (1) shows that the multimode focal spot diameter is M^2 times the diffraction-limited value for equal beam diameter at the focusing lens. Owing to this, a laser beam is often referred to as being "M^2 times diffraction-limited," meaning that it will have M^2 times the spot size. Equation (1) also shows that it is possible to get a smaller focal spot by using a shorter focal length lens or by increasing the initial beam diameter (by using a beam-expanding telescope).

When the beam diameter at the lens is the same for both diffraction-limited and multimode beams, then the far-field full-angle divergence, $\theta = d/f$, is the same for both beams. However, if the focal spot sizes (d_0) are made to be the same — because the multimode beam has a larger diameter at the lens — then the divergence will be M^2 times larger for the multimode beam. This is seen by considering the Rayleigh range, which is an important parameter in imaging because it is a measure of the distance across which the laser beam (or sheet) remains focused. The definition of the Rayleigh range x_R is the distance along the beam from the focus to the point where the beam diameter is $\sqrt{2}$ times the diameter at the focus. The relationship is

$$x_R = \frac{\pi d_0^2}{4\lambda M^2}, \tag{2}$$

which is the same as the Rayleigh range for a Gaussian beam, except that λ has been replaced by λM^2. Equation (2) shows that for equal spot size, as M^2 increases, the Rayleigh range decreases because of the greater divergence. It can be concluded from this that aberrated beams can be focused to as small a spot as a diffraction-limited beam (by expanding it before the focusing lens), but the focal spot cannot be maintained over as large a distance. Note that the M^2 value can usually be obtained from the laser manufacturer, but it can also be measured by passing the beam through a lens of known focal length and then measuring the beam diameter at several locations (17).

In planar imaging, the laser beam is formed into a thin sheet, which can be accomplished by several different techniques (11). One of the more common methods is shown in Fig. 2 where a spherical lens, which is typically plano-convex and has a focal length of 500 to 1000 mm, is used to focus the beam near the center of the field of view of the camera. Such long focal length lenses are used to increase the Rayleigh range, or the distance across which the beam remains focused. The larger Rayleigh range obtained from long focal length lenses does not come without a cost, however, because the longer focal length lenses also result in larger focal spots, or thicker sheets, in planar imaging.

Figure 2 also shows the use of a cylindrical telescope formed from a plano-convex lens of focal length f_1 and a larger plano-convex lens of focal length f_2. For high peak power laser beams, it may be best to use a negative (plano-concave) lens as the first lens to avoid a real focus and hence reduce the possibility of air breakdown. The reason for using two plano-convex lenses — where the convex sides are directed toward the collimated beam — is that this configuration minimizes the aberrations for a telescope formed from simple spherical lenses (18).

The cylindrical lenses expand the laser beam only in one direction, by a factor of f_2/f_1. Because the laser sheet height is determined by the height of the second cylindrical lens, producing large sheets (e.g., 100 mm) requires a large lens, which can be very expensive. Often, the second lens is omitted, and the sheet is allowed to diverge. The disadvantage is that the laser intensity varies in the propagative direction, which can make it harder to

correct the image of the scattered light for variations in intensity. Because a laser sheet is formed by expanding the beam in only one direction by using a cylindrical lens, the thickness of the sheet at the focus is approximately equal to the focal spot diameter given by Eq. (1). However, when the sheet thickness must be measured, this can be accomplished by using the scanning knife-edge technique. In this technique a knife-edge (e.g., a razor blade) is placed normal to the laser sheet and is translated across it so that the beam is progressively blocked by more of the knife-edge. The transmitted light is measured by a power meter as the knife-edge is translated. The derivative of the power versus distance curve is the mean sheet intensity profile. For example, if the laser sheet intensity profile is Gaussian, then the knife-edge intensity profile will be an error function.

Spreading the laser beam into a sheet results in a large reduction in the intensity (or fluence); thus, when the intensity must be maximized, such as in Raman scattering imaging, the laser sheet can be formed by using a multipass cell (19). In this case, the laser beam is reflected back and forth between two confocal cylindrical mirrors. The main problem in this technique is that the sheet intensity profile is very nonuniform, and the nonuniformity may be difficult to correct for on a single-shot basis. In this case, shot-to-shot fluctuations in the intensity distribution can be left as an artifact in the image. Another technique that can be used in low-velocity flows is the scanning method, where a CW laser beam is swept past the field of view by using a moving mirror (6). If time-resolved data are desired, then the sweep time must be short enough to freeze the motion of the flow. Because of this, the scanning technique is really useful only in liquid flows, which have relatively small characteristic flow timescales.

Cameras

The most commonly used cameras in quantitative imaging are based on charged-coupled device (CCD) arrays or image-intensified CCD arrays. Note that there are a few applications where film may be preferred to a digital camera, such as large field-of-view PIV (20) and high-framing-rate PIV (21,22). Nevertheless, CCD arrays have largely supplanted film and other detectors, including TV tubes, photodiode and charge-injection device (CID) arrays, owing to their low noise, excellent linearity, uniformity, and resistance to blooming. The operation of a CCD is based on the fundamental property that a photon incident on the CCD produces an electron–hole pair in a region of silicon that is biased to some potential. The electrons generated are called "photoelectrons," which migrate to the "potential well" of the CCD pixel where they are stored for later readout. Because the CCD stores charge, it is essentially a capacitor, whose charge is proportional to the number of incident photons. The quantum efficiency η is the ratio between the number of photoelectrons generated and the number of photons incident. Front-illuminated CCDs have quantum efficiencies of 10–50% at visible and near-IR wavelengths (peaking near 700 nm) but are virtually zero at UV and mid-IR wavelengths. Back-illuminated CCDs, although more expensive, provide quantum efficiencies up to 90% in the visible and can maintain good response (e.g., $\eta = 20\%$) well into the UV.

CCD arrays can be full frame, frame transfer, or interline transfer type (23). *Full frame* CCD arrays read out the charge by shifting it down through the entire array (like a "bucket brigade") into an output register where it is then read out serially. Because the array is used to shift the charge, the image will be blurred if the CCD is exposed during readout. Because readout can take several seconds, a mechanical shutter must be used. In contrast, *frame transfer* CCD arrays use a photosensitive array and an identical array that is masked off from any incident light. After an exposure, the charge of each pixel is shifted down through the array into the masked array, and the masked array is then read out in the same manner as a full frame CCD array. Frame transfer CCD arrays offer some level of electronic shuttering, but this is limited to a few milliseconds. The pixel area for both full frame and frame transfer CCD arrays is 100% photosensitive, thus the pixel width is the same as the pixel pitch (spacing). *Interline transfer* CCD arrays have nonphotosensitive storage registers located adjacent to the photosensors. This enables the rapid transfer of charge (in parallel) from the pixels into the storage registers. This makes it possible to rapidly shutter the array electronically, where exposure times of the order of microseconds or less are possible. The interline transfer arrays also enable "frame straddling," whereby two frames can be captured in rapid succession. For example, standard RS-170 format video cameras based on interline transfer arrays can acquire two video fields in less than 10 μs between frames (24). More expensive scientific grade interline transfer cameras report interframe times as short as 200 ns. Frame-straddling by video cameras is useful for double-pulse imaging in high-speed flows (25), whereas frame-straddling by higher resolution scientific/industrial cameras (e.g., Kodak ES1.0 and ES4.0) is now becoming the norm for PIV because it enables the use of cross-correlation processing algorithms. The main drawback of interline transfer imagers is that they tend to be noisier than either full frame or frame transfer imagers. The main reason for this is that the storage registers are located adjacent to the photosensitive sites; therefore the photosensitive area of the pixel is substantially smaller than the physical area of the pixel. The fraction of the pixel area that is photosensitive is called the "fill factor" and is typically 20–30% for an interline transfer CCD. As will be discussed later, the signal scales with the number of photons collected per pixel; thus low fill factors result in low signals. Some manufacturers mitigate this problem to some extent by using microlenses over each pixel to collect light across a larger area and can increase the fill factor to about 60%. If neither electronic shuttering nor frame straddling is required, then full frame or frame transfer imagers are desired to maximize the signal-to-noise ratio (SNR). Generally, the relatively long shutter times are not a problem when pulsed lasers are used because the laser pulse duration acts as the exposure time.

Intensified CCD cameras (ICCD) are used for low light-level imaging and for very short exposure times (e.g., as

low as a few nanoseconds). The most common type of image intensifier consists of a photocathode, a microchannel plate, a phosphor screen, and a mechanism to couple the screen to the CCD (26,27). Photons that are incident on the photocathode eject photoelectrons, which in turn are amplified in the microchannel plate. The amplified electrons contact the phosphor screen causing photon emission, and these photons are collected by the CCD. The phosphor screen is usually coupled to the CCD by a fiber optic bundle, although lens coupling is also used. Image intensifiers are shuttered by switching on and off, or "gating," the photocathode by a high-voltage pulse. The electron gain is a function of the voltage applied across the microchannel plate. Short duration gating is necessary to reject the background luminosity of very luminous flows, such as sooting flames or plasmas. Because the duration of the laser scattering signal is often of the order of several nanoseconds, short gates greatly reduce the background luminosity but do not affect the signal. One of the main drawbacks of intensifying CCD cameras is that the intensifiers tend to have both lower resolution and lower signal dynamic range than the bare CCD. The dynamic signal range is usually limited by saturation of the microchannel plate, particularly at high electron gain (26), rather than by saturation of the CCD itself. Furthermore, as will be shown later, it is unlikely that an ICCD camera will provide better SNR than a low-noise CCD camera under the constraint that a certain minimum SNR is required for an image to be useful for quantitative analysis. For these reasons, ICCD cameras are preferred to low-noise UV-sensitive CCD cameras only when fast gating is required, which is why they are primarily used for imaging high-temperature gases.

SIGNAL AND NOISE

One of the most critical issues in flow imaging is obtaining an adequate SNR. Imaging measurements that use laser light scattering are particularly susceptible to a low SNR because the laser beam must be spread out into a sheet; thus, signals are lower by hundreds to thousands of times, compared to a point measurement with the same laser energy. Figure 3 shows a generic camera system that views a region in the flow that is illuminated by a laser light sheet of height y_L and thickness Δz. Assume that the camera uses an array sensor and a lens of known focal length f and limiting aperture diameter D. Each pixel of the camera, of width δx and height δy, transforms to a region in the flow of dimensions, $\Delta x = \delta x/m$, $\Delta y = \delta y/m$, where $m = y_i/y_o$ is the magnification and y_i and y_o are as defined in Fig. 3. Each pixel also spatially integrates the signal in the z direction across a distance equal to the sheet thickness Δz. Note that usually in flow imaging, the image is inverted, and the magnification is typically less than unity, that is the object is minified. Now, assuming that a pulsed laser light sheet is used that has a local fluence F_L, then the number of photons collected by each pixel S_{pp} will be

$$S_{pp} = \frac{F_L}{h\nu}\Delta V \frac{d\sigma}{d\Omega} n \Delta\Omega \eta_t, \tag{3}$$

where h is Planck's constant, ν is the laser frequency, $\Delta V = \Delta x \Delta y \Delta z$ is the volume imaged by each pixel, $d\sigma/d\Omega$ is the differential scattering cross section, n is the number density of the scattering medium, $\Delta\Omega$ is the solid angle subtended by the lens, and η_t is the transmission efficiency of the collection optics (lens and spectral filters). For a CW laser, $F_L = I_L \Delta t$, where I_L is the laser intensity (power flux density) and Δt is the integration time. The solid angle, $\Delta\Omega = (\pi D^2/4)/z_o^2$ (where z_o is the distance from the object to the lens), is related to the magnification and f number ($f_\# = f/D$) of the lens by:

$$\Delta\Omega = \frac{\pi}{4}\frac{m^2}{(f_\#)^2(m+1)^2}. \tag{4}$$

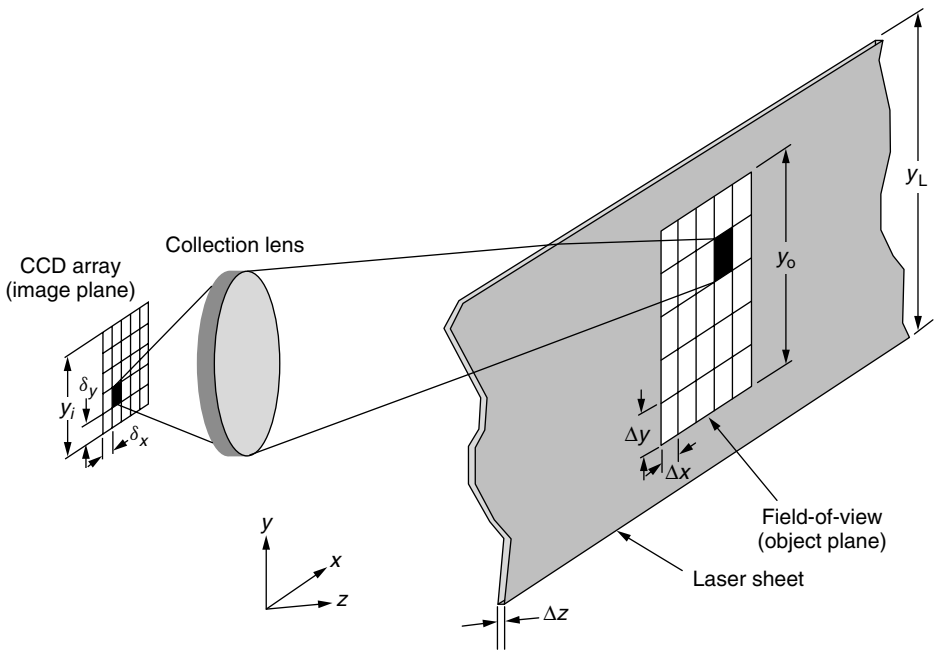

Figure 3. Planar laser imaging of a flow field using an array detector.

Assuming that the laser sheet is uniform (i.e., the fluence is constant), then the fluence can be approximated as $F_L = E_L/y_L$, where E_L is the laser energy. Now combining Eqs. (3) and (4), and substituting $\Delta x = \delta x/m$ and $\Delta y = \delta y/m$ gives

$$S_{pp} = \frac{E_L}{h\nu}\frac{\delta x \delta y}{y_L}\frac{d\sigma}{d\Omega}\left(\frac{\pi}{4}\frac{1}{(f\#)^2(m+1)^2}\right)n\eta_t. \quad (5)$$

Equation (5) shows that the photons collected per pixel actually increase as $m \to 0$, or as the camera is moved farther from the object plane. This may seem counterintuitive because the solid angle subtended by the lens progressively decreases. The reason for this is that Δx and Δy increase as the magnification decreases, which means that each pixel collects light from a larger region of the flow. This is correct as the problem has been posed, but is not realistic, because it assumes that the laser sheet has the same fluence, regardless of the field of view. However, in practice, as the camera is moved farther away, the laser sheet must be enlarged to accommodate the larger field of view. To see this effect, assume that the condition $y_L = y_o$ must be maintained as the magnification is changed; in this case $y_L = N_p \delta y/m$, where N_p is the number of pixels in one column of the array. Now, Eq. (5) reduces to

$$S_{pp} = \frac{E_L}{h\nu}\frac{y_i}{(N_p)^2}\frac{d\sigma}{d\Omega}\left[\frac{\pi}{4}\frac{m}{(f_\#)^2(m+1)^2}\right]n\eta_t. \quad (6)$$

This form of the equation is probably the most useful for seeing the effect of varying different parameters. For example, Eq. (6) shows that the signal depends only on the laser energy (actually, the term $E_L/h\nu$ represents the total number of incident photons) and is independent of Δz or on how tightly the sheet is focused. Although tighter focusing increases the fluence, this effect is counteracted by a decrease in the number of molecules that is available to scatter the light. In addition, as the number of pixels is increased (at fixed detector size y_i), the signal decreases because the pixels are smaller and thus collect light from a smaller area of the flow. This shows the importance of having large pixels (or small N_p at fixed y_i) to improve the SNR, albeit possibly at the expense of resolution. The trade-off between SNR and resolution is a fundamental one, whose manifestation in point measurements is the trade-off between SNR and bandwidth (or response time). Equation (6) also shows that $S_{pp} \sim m/(m+1)^2$, a dependence that is plotted in Fig. 4. Here, it is seen that the signal is maximized at a magnification of unity and that there is an abrupt decrease in signal as $m \to 0$. Equation (6) also shows that the signal is inversely proportional to $f_\#^2$, and thus it is essential in many imaging techniques to use lenses that have low f numbers. For several techniques such as PLIF, Rayleigh scattering, and Raman scattering in gas-phase flows it is difficult to obtain adequate SNRs using lenses whose f numbers are higher than $f/1.2$.

Equation (6) gives the number of photons incident on a pixel of a generic detector. The resulting *signal* then consists of the photoelectrons that are generated, whether by creating an electron–hole pair in a CCD or by ejecting an electron from a photocathode. The signal S_e (in units of electrons, designated as e^-) is given by

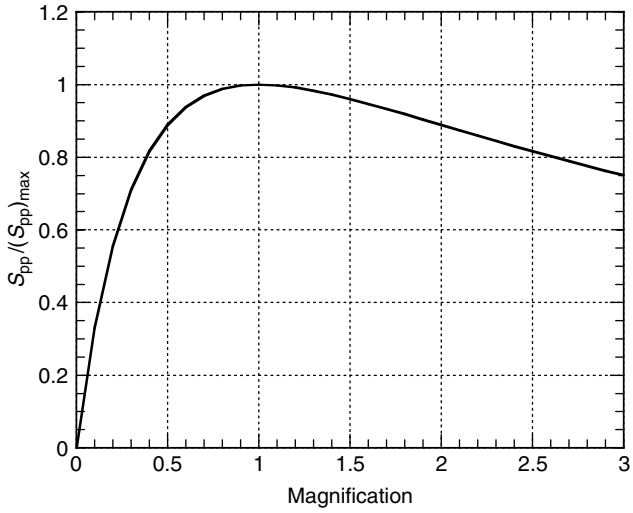

Figure 4. Relative variation of photons-per-pixel (S_{pp}) versus magnification for a typical planar imaging experiment.

$$S_e = \eta S_{pp} G, \quad (7)$$

where G is the overall electron gain from the photocathode to the CCD. For an unintensified CCD, $G = 1$.

The noise in the signal will have several sources, but the dominant sources in scientific grade CCD and ICCD cameras are shot noise and "read" noise. Shot noise results from statistical fluctuations in the number of photoelectrons generated at each pixel. The statistical fluctuations of photoelectrons and photons exhibit Poisson statistics, for which the variance is equal to the mean (28). Most of the shot noise arises from statistical fluctuations in the photoelectrons generated, although some noise is induced in the amplification process of image intensifiers. The shot noise (in units of e^-), which is the square root of the variance, is given by (29)

$$N_{shot} = G(\eta\kappa S_{pp})^{1/2}, \quad (8)$$

where κ is the noise factor. The noise factor quantifies the noise that is induced through the overall gain process between the photocathode and the array; for an ICCD, it is gain dependent and falls within the range of $1.5 < \kappa < 2.5$. In an unintensified CCD, $G = \kappa = 1$, and the shot noise is equal to $(\eta S_{pp})^{1/2}$, which is the square root of the number of photoelectrons collected per pixel during the integration period. One way of interpreting the shot noise in a detector array is to consider the case where the array is composed of identical pixels and is illuminated by a spatially uniform light source. If it is assumed that each pixel collects an average of 1000 photons during the integration time and if it is further assumed that $\eta = 0.1$, then, on average, each pixel will collect 100 photoelectrons. However, the actual number of photoelectrons collected will vary from pixel to pixel, and compiling a histogram of the pixel values will reveal that the variance of the distribution is equal to the mean number of photoelectrons collected per pixel.

The dominant noise source intrinsic to scientific grade CCD cameras is "read noise" (30). Read noise is incurred in the output registers in the process of converting the charge of each pixel into a voltage that can be read by an analog-to-digital converter. A pixel is read by transferring the charge of each pixel to a small capacitor, whose integrated charge is converted to a voltage by an on-chip amplifier. The dominant sources of read noise are dark-current shot noise, "reset noise," and output amplifier noise. Dark current is the current that is generated in the absence of incident light due to thermally induced charge carriers. Cooling a CCD greatly reduces the dark current. For example, an uncooled CCD might generate a dark current of 300 e^-/s at 20 °C, but only 1 e^-/s at −40 °C. Owing to the relatively short exposure and readout times that are typically used in flow imaging (of the order of 10 seconds or less), shot noise in dark current is not usually a large contributor to the noise in *cooled* CCD arrays. Reset noise is injected into the small capacitor by a switching transistor, whose job is to reset the capacitor to a reference voltage in preparation for reading the next pixel's charge. This switching transistor contaminates the capacitor charge with both "digital feedthrough" and thermal noise. Digital feedthrough noise is caused by capacitive coupling of the clock signals through the switching transistor. These noise sources can be greatly limited by slow (low-bandwidth) readout rates and correlated double sampling (30,31). Because means have been developed to reduce these noise sources, the intrinsic camera noise is typically limited by the on-chip output amplifier to a few electrons rms per pixel (typically 5–20 e^-).

When photoelectron shot noise is not the only noise source, then it is assumed that the noise sources are uncorrelated and therefore their variances add. In this case, the SNR is given by (29)

$$\mathrm{SNR} = \frac{\eta S_{\mathrm{pp}} G}{(\eta\kappa S_{\mathrm{pp}} G^2 + N_{\mathrm{cam}}^2)^{1/2}}, \tag{9}$$

where N_{cam} is the intrinsic background noise of the camera (in electrons rms) and includes contributions from amplifier noise, digital feedthrough noise, thermal noise, dark-current shot noise, and quantization noise from the analog-to-digital converter. There are several interesting implications of Eq. (9). The first is seen by considering the limit when the signal is dominated by shot noise, that is, when $\eta\kappa S_{\mathrm{pp}} G^2 \gg N_{\mathrm{cam}}^2$. This shot-noise-limited operation of the detection system occurs when either the read noise is small or when the signal is high. Equation (9) also shows that it is possible to obtain shot-noise-limited operation by increasing the gain until the first noise term dominates the other. This is the way an image intensifier works; it provides very high electron gain through the microchannel plate and thus causes the shot noise to overwhelm the intrinsic noise sources in the camera. It may seem odd that the goal is to increase the noise, but the signal is also increased as the gain increases, so the SNR either improves or remains constant. At low gain, the signal will be detector-noise-limited. As the gain is increased to arbitrarily high levels, the SNR continues to improve until it reaches the shot noise limit, beyond which the SNR is constant. This is seen in Eq. (9) by letting $G \to \infty$, in which case the SNR becomes independent of G. Because electron gains of 10^3 are possible by using single-plate microchannel intensifiers that are typical, it is possible to operate in the shot-noise-limited regime, even when the camera that stores the image has relatively high noise, such as a video format CCD camera.

The dynamic range of a CCD — defined as the ratio of the maximum to the minimum usable signals — is limited by the well depth, which is the total number of photoelectrons that can be stored in a CCD pixel, and the intrinsic noise of the camera. Specifically, the dynamic range D_{R} is given by (29)

$$D_{\mathrm{R}} = \frac{S_{\mathrm{e,sat}} - S_{\mathrm{dc}}}{N_{\mathrm{cam}}}, \tag{10}$$

where $S_{\mathrm{e,sat}}$ is the signal at saturation (full well) and S_{dc} is the integrated dark charge. For example, for a cooled slow-scan CCD array whose integration time is short (hence low S_{dc}) and has a well depth of 10^5 e^- and noise of 10 e^-, then $D_{\mathrm{R}} \approx 10^4$, which is much larger than can usually be obtained in single-shot planar imaging. The dynamic range of an ICCD can be much smaller than this because the electron gain from the photocathode to the CCD effectively reduces the well depth of the CCD (29). For example, if the overall electron gain is 10^2, then a CCD that has a well depth of 10^5 e^- will saturate when only 10^3 photoelectrons are generated at the photocathode. In addition, ICCD cameras may have an even lower dynamic range than that allowed by saturation of the CCD well because of saturation of the microchannel plate (26).

Figure 5 shows how the SNR varies as a function of the number of photons per pixel for cameras of high and low read noise, as might be found in video format and slow-scan CCD cameras, respectively. In this figure, it is assumed that $\eta = 0.7$ for the low-noise camera and $N_{\mathrm{cam}} = 10\ e^-$, and $\eta = 0.7$ and $N_{\mathrm{cam}} = 200\ e^-$ for the high-noise camera. Also shown is the case where the high-noise camera has been intensified. It is assumed that the intensified camera

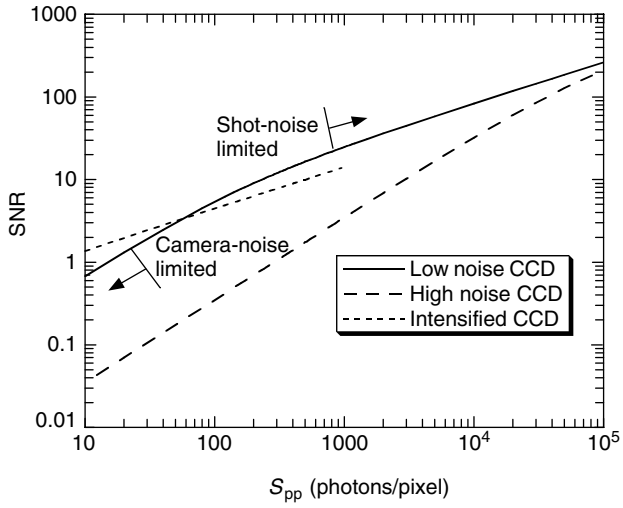

Figure 5. Variation of the SNR versus signal (S_{pp}) for three different camera systems.

has a lower quantum efficiency ($\eta = 0.2$) and $G = 500$. Dark charge has been neglected in all cases. The high-noise camera is camera-noise-limited for the entire range of S_{pp} (hence, the slope of unity on the log–log plot), whereas the low-noise camera is camera-noise-limited only for low S_{pp}. As expected, the SNR is substantially higher for the low-noise camera at all S_{pp}. At higher S_{pp}, the low-noise camera becomes shot-noise limited, as seen by the region where the slope is one-half on the log–log plot. By intensification, the high-noise camera reaches the shot-noise limit even at very low S_{pp}; thus results in a SNR that is even higher than that of a low-noise camera. However, for S_{pp} greater than about 60, the low-noise camera outperforms the intensified camera, owing to its higher quantum efficiency. Figure 5 also shows that at an overall electron gain of 500, if the well depth is 10^{-5} e^-, the intensified camera saturates the CCD when 1000 photons are incident per pixel.

One point to consider is that for flow imaging, it is usually not necessary or desired to intensify a slow-scan low-noise CCD camera, unless gating is required to reject a luminous background. The main reason is that if the signal is so low that read noise is a significant contributor to the total noise, then it is unlikely that single-shot images will be useful for quantitative purposes. For example, assume that a minimum SNR of 20 is desired for quantitative analysis and that the intensified slow-scan camera has $\kappa = \eta = 1$, is operated at high gain, and the CCD has 10 e^- rms of read noise. If 100 e^- are collected per pixel, then the high gain overwhelms the read noise, and the signal is shot-noise limited, that is, $\text{SNR} = (100)^{1/2} = 10$, which is well below our minimum value. Now, assuming that 500 e^- are collected, then the SNR based only on shot noise is $(500)^{1/2} = 22$. However, at these signal levels, the signal is nearly shot-noise-limited, even without the intensifier, because including the camera noise gives a SNR ≈ 20; thus there would be very little benefit in intensifying the CCD. The fact that the intensifier is likely to have a smaller dynamic signal range, worse resolution, lower quantum efficiency, and a larger noise factor than the CCD, makes intensification even less desirable.

It is also interesting to consider how the high-noise camera would perform with the signal of 500 e^-. In video format cameras, the read noise will be about 100–200 e^- rms. Using the lower value, the SNR for the video camera would be $500/100 = 5$. In this case, adding an image intensifier would be an advantage because high electron gain could be used to obtain shot-noise-limited operation, so that the $\text{SNR} = (500)^{1/2} = 22$ (assuming equal η with and without intensification).

IMAGE CORRECTIONS

Quantitative imaging always requires several correction steps so that the measured signal can be related to the flow property of interest and to ensure that the spatial structure of the object is faithfully represented by the image. First, consider corrections to the signal measured at each pixel of the array. Most planar imaging involves only relative measurements of signal intensity, from which absolute measurements can be obtained by calibrating a single point within the image. To obtain an image that represents quantitatively accurate relative intensity measurements requires making several corrections to the measured image. For example, let $S_e(x, y)$ represent the desired signal level at a given pixel or location on the array (x, y). By "desired" it is meant that $S_e(x, y)$ is proportional to the number of photons incident on that pixel originating from the scattering process of interest. The signal S_e can be related to the total signal (S_{tot}) recorded at that pixel by the imaging system through the relationship

$$S_{tot}(x, y, t_i, t_{ro}) = w(x, y)\left[L(x, y)S_e(x, y) + S_{back}(x, y, t_i)\right] + S_{dark}(x, y, t_{ro}), \tag{11}$$

where $L(x, y)$ is a function that is proportional to the laser sheet intensity (or fluence) distribution function, S_{back} is the signal resulting from unwanted background light, S_{dark} is the fixed pattern signal that occurs with no light incident on the detector, t_i is the exposure time, and t_{ro} is the array readout time (which includes the exposure time). The function $w(x, y)$ is the "white-field" response function, which accounts for variation in the signal across an image of a uniformly white object. It has been assumed that a pulsed laser is used as the light source, in which case the signal S_e is not a function of the exposure time. Furthermore, in general, all of the functions involved in the correction may vary from shot to shot. The desired scattering signal is obtained by solving for S_e in Eq. (11):

$$S_e(x, y) = \frac{S_{tot}(x, y, t_i) - [w(x, y)S_{back}(x, y, t_i) + S_{dark}(x, y, t_{ro})]}{w(x, y)L(x, y)}. \tag{12}$$

Equation (12) gives a means of obtaining the desired scattering signal image by arithmetic processing of the signal and correction images. $S_{dark}(x, y, t_{ro})$ is not noise because it is an offset that is nominally the same for each image that has the same exposure and readout time. The dark image is obtained by acquiring an image when the shutter is closed (or when the lens cap is on) and using the same integration and readout times as in the experiment. The background signal $S_{back}(x, y)$, is due to reflections of the laser from walls/windows, natural flow luminosity (as in combustion), fluorescence from windows or species not of interest, and external light sources. For nonluminous flows, a good approximation to the background can be obtained by acquiring an image when the laser beam is present but without the scattering medium (e.g., without the fluorescent species seeded into the flow). This is only an approximation of the actual background because the light itself that is scattered from particles/molecules in the flow can reflect from the walls and windows; therefore, an image obtained when the scattering medium omitted may not have the same background signal as during an actual experiment. There is usually no simple way around this problem, but fortunately, this effect is often negligible.

It is important to note that the background cannot be measured directly because it is the function wS_{back} that is actually measured when a background image is acquired. In fact, the background image is also affected by the dark signal; therefore, if the background image

is acquired by using the same exposure and readout times as the scattering signal image, then this yields the term $S_{\text{correction}} = (wS_{\text{back}} + S_{\text{dark}})$ in Eq. (12). In this case, the correction relationship is simply, $S_e = (S_{\text{tot}} - S_{\text{correction}})/(wL)$. Note also that to reduce the effect of noise on the correction procedure, the images $S_{\text{correction}}(x, y)$, $w(x, y)$, and $L(x, y)$, should be average images, unless the corrections are made on a single-shot basis.

If the flow is unsteady and luminous, the luminosity varies from shot to shot, and therefore, it is more difficult to correct for the background signal. In this case, it is useful to consider the signal-to-background ratio (SBR), S_e/S_{back}, which is sometimes confused with the SNR. Background luminosity is usually not random, and thus it is not noise (although it may appear so if one does not have an easy way to correct for it). One option for dealing with background luminosity is to reduce the luminosity incident on the array through gating, by using an intensified camera or by using spectral filters in front of the camera that pass the scattered light but reject the bulk of the luminosity. Another option is to use a second camera to capture an image of the flow luminosity a very short time before (or after) the laser fires. This assumes, of course, that the flow is essentially frozen for each camera image, which is unlikely to be the case for the millisecond shutter times used for full frame CCD cameras, but it is likely to be true when using microsecond gates and an intensified camera.

The laser sheet intensity distribution function, $L(x, y)$, is not easy to obtain, but it can be approximated in a few different ways. In general, the sheet intensity varies in both the x and y directions and from shot to shot. Figure 2 shows a technique, described in (32), for measuring $L(x, y)$ on a single-shot basis. For single-shot corrections, it is necessary to collimate the laser sheet, so that L is a function only of y. In this case, part of the laser sheet energy can be extracted, as done using the glass flat in Fig. 2, and directed onto a target. The glass flat reflects several percent of the laser light from each surface, depending on the angle of incidence (33). In Fig. 2 the target is a white card, although a cell containing fluorescent material could also be used (e.g., laser dye in water, or acetone vapor). The scattering (or fluorescence) from the target must obviously be linear in its response to the incident light intensity and must scatter the light uniformly. In Fig. 2, a video camera is used to image the laser sheet intensity profile. Rather than using a target, it is also possible to image the beam directly using a 2-D or linear array. The main drawback of this technique is the risk of damage to the array by the focused laser beam. The scattering image and the sheet profile can be registered by blocking the beam, before the optical flat, at two discrete vertical locations using two very thin wires. Both the scattering image and the profile image will include a shadow of the wires, which can be used to index the two images.

If the laser sheet is not collimated, but diverging, this makes it much more difficult to correct for the sheet on every shot. In this case, the laser energy and distribution must be sufficiently repeatable so that $L(x, y)$ can be obtained at a time different from that for the scattering image. The correction image is obtained by placing a uniform, linear scattering medium in the field of view. Sometimes, it is possible to use the Rayleigh scattering from the air itself, although it is more common to have to introduce a more efficient scattering medium, such as smoke or a fluorescent test cell. Care must be taken when using fluorescent materials, such as laser dyes or acetone vapor, because they will cause substantial absorption of the beam if the concentration is too high. Unless the absorption itself is corrected for, the sheet intensity distribution will be incorrect. Therefore, when using fluorescent media, it is best to use very low concentrations to keep the absorption to less than a few percent across the image. The low concentration may necessitate averaging the correction image over many shots to obtain sufficient SNR.

The white-field response function, $w(x, y)$, is obtained by imaging a uniformly white field, such as a uniformly illuminated white card. The signal of a white-field image will tend to decrease from the center of the image because the solid angle subtended by the lens is smaller for point sources located near the periphery of the field of view. The variation in intensity across an image formed by a circular aperture will theoretically follow the "cosine-to-the-fourth" law, or $I(\beta)/I(0) = \cos^4 \beta$, where β is the angle between the optical axis and a line connecting the center of the lens aperture and the given point on the object plane (18). The white-field response function will also enable correction for variable response of the pixels in the array. Note that the dark charge contribution to the signal must also be subtracted from the white-field image.

In some cases, it will be necessary to correct for geometric distortion. The distortion in an image is typically larger for points farther from the optical axis. For this reason, a square will be imaged as an object whose sides either bulge out (called barrel distortion) or in (called pincushion distortion). When using high quality photographic lenses, the maximum distortion is usually small (often less than a pixel). However, when it must be corrected for, this is usually accomplished by imaging a rectangular grid and then warping (or remapping) the image so that each point of the grid is consistent with its known geometry (34). The warping procedure involves finding a large number of "tie" points across the image such as the "points" where two gridlines cross and using these to solve for a set of polynomial coefficients required for the remapping. Pixels other than the tie points are remapped by interpolating among coefficients for the tie points.

IMAGING SYSTEM RESOLUTION

Even though the proper specification of the resolution of the imaging system is often critically important to a particular application, it is often neglected in flow imaging studies. For example, it is not unusual to find scalar imaging papers that quote the resolution in terms of the area that each pixel images in the flow. In many cases, however, this is not the factor that limits the resolution, particularly when using fast (low $f_\#$) optics. A somewhat better approach involves imaging a standard resolution target, such as the USAF or NBS targets (35), available

from major optics companies, which are composed of a periodic sequence of light and dark bars of varying spatial frequency. The user typically reports the resolution limit as the smallest set of bar patterns for which a contrast modulation can be distinguished. In some cases, this may give the user an idea of the limiting resolution of the imaging system, but this technique is subjective, can be misleading because of aliasing (discussed further below), and is inadequate as a measure of the limitations that finite resolution impose on the data.

The resolution is fundamentally related to the point-spread function (PSF), which is the intensity distribution at the image plane, $I_i(x, y)$, produced by imaging an infinitesimally small point source of light. The overall size of the PSF is referred to as the blur spot, whose diameter is denoted as d_{blur}. In the diffraction limit, the PSF will be the Airy function (33), which has a blur spot diameter that can be approximated as the Airy disk diameter, $(d_{\text{blur}})_{\text{dl}}$, given by the relationship (20)

$$(d_{\text{blur}})_{\text{dl}} = 2.44(m + 1)\lambda f_{\#}. \tag{13}$$

Most flow imaging experiments employ camera lenses designed for 35-mm film cameras. When used at high $f_{\#}$ and for magnifications that are not too far off design, these lenses give nearly diffraction-limited performance. The lenses have several lens elements that are necessary to correct for the many types of aberrations, including spherical, chromatic, coma, astigmatism, and distortion. However, chromatic aberrations are not usually a problem in flow imaging, because in most cases the scattered light is effectively monochromatic. In practice, such photographic lenses used at low $f_{\#}$ and off-design produce spot sizes that can be several times larger than the Airy disk. For example, Fig. 6 shows digitally sampled images of a point light source ($\lambda = 532$ nm) whose diameter is approximately two microns in the object plane, taken at unity magnification by a Nikon 105-mm Micro lens coupled to a Kodak ES1.0 1 k × 1 k CCD camera (9 μm × 9 μm pixels). For comparison, the length of the horizontal white bar below each image of Fig. 6 is equal to the diffraction-limited spot size computed from Eq. (13). Figure 6 shows that the spot size is approximately diffraction-limited at $f/22$ and $f/11$, but at $f/2.8$, the spot size is about 50 μm, which is substantially larger than the diffraction-limited value. The increase in the blur spot, relative to the diffraction limit, results from the greater aberrations of the lower $f_{\#}$.

The PSF directly affects the resolution because the image is the result of the convolution of the PSF with the irradiance distribution of the object. Therefore, the smallest objects that can be imaged are related to the size and shape of the PSF; worse resolution is associated with a broader PSF or larger blur spot. In addition to setting the limiting resolution, or the highest spatial-frequency structure that can be resolved, the imaging system also tends to blur increasingly smaller scale structures. Because of this, it is usually not sufficient to simply state the limiting resolution of the system. For example, it will be shown later that measurements of scalar gradients, such as derived from temperature or concentration fields, can exhibit substantial errors due to resolution limitations, even at frequencies substantially lower than the limiting resolution of the system.

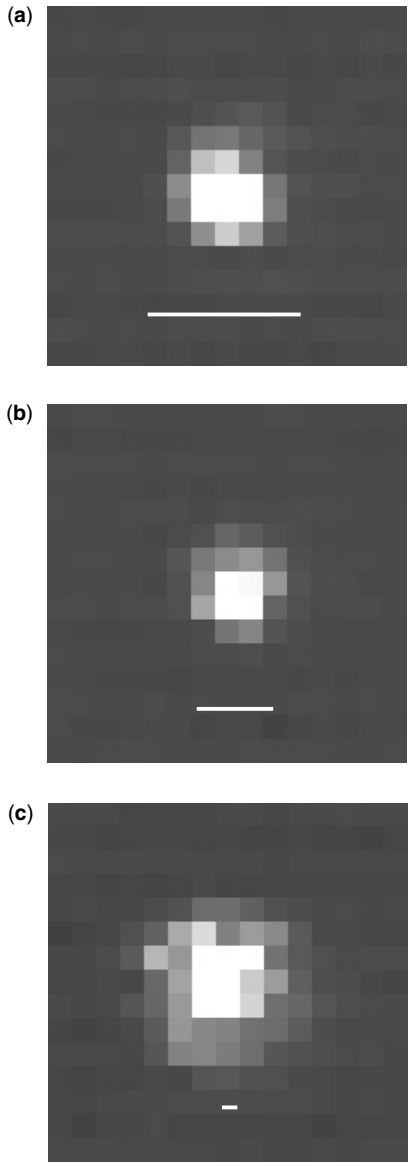

Figure 6. Digitally sampled point-spread functions acquired using a Kodak ES1.0 CCD camera (9 × 9 μm pixels) fitted with a Nikon 105-mm lens. The object imaged is a point source approximately 2 μm in diameter, and the magnification is unity. The three images are for three different aperture settings: (**a**) $f/22$, (**b**) $f/11$, and (**c**) $f/2.8$. The white line below each spot is the diameter of the diffraction-limited blur spot.

The blurring incurred by an imaging system that has finite resolution is essentially a result of the system's inability to transfer contrast variations in the object to the image. The accepted means of quantifying how accurately an imaging system transfers contrast is the optical transfer function (OTF) (18,35). The OTF, which is analogous to a linear filter in time-series analysis, describes the response of the imaging system to a sine wave contrast variation in the object plane. For example,

assume that the intensity distribution of the object is described by the equation

$$I_o(x) = b_0 + b_1 \cos(2\pi sx), \tag{14}$$

where I_o is the intensity of the object, b_0 and b_1 are constants, and s is the spatial frequency (typically in cycles/mm, or equivalently, line-pairs/mm). It can be shown that a linear system will image the object as a sine wave of the form (18)

$$I_i(x) = b_0 + c_1 \cos(2\pi sx - \phi), \tag{15}$$

where I_i is the intensity of the image, c_1 is a constant, and ϕ is a phase shift. Examples of functions are shown in Fig. 7, where the image exhibits both a reduction in the contrast (i.e., $c_1 < b_1$) and a phase shift, which corresponds to a shift in the location of the wave. Because the phase shift is associated with a shift in the position of the image, it is generally associated with geometric distortion. The OTF can be described mathematically by the relationship

$$\mathrm{OTF}(s) = \mathrm{MTF}(s)e^{i\mathrm{PTF}(s)}, \tag{16}$$

where MTF(s) is the modulation transfer function and PTF(s) is the phase transfer function. The MTF describes the contrast transfer characteristics of the imaging system, and the PTF describes the phase transfer characteristics. Equation (16) shows that the magnitude of the OTF is the MTF, that is, MTF(s) = |OTF|. The MTF is generally considered more important in describing the transfer characteristics of an imaging system because phase differences typically occur only at high spatial frequency where the MTF is very small (35).

The MTF is measured by imaging objects that have a sine wave irradiance variation of known spatial frequency. The maximum and minimum intensities are defined as $I_{\max}$, $I_{\min}$, respectively, and the contrast of the object is defined as $C_o = (I_{o_{\max}} - I_{o_{\min}})/(I_{o_{\max}} + I_{o_{\min}})$. The contrast of the image is defined similarly as $C_i = (I_{i_{\max}} - I_{i_{\min}})/(I_{i_{\max}} + I_{i_{\min}})$. The MTF is then defined as

$$\mathrm{MTF}(s) = \frac{C_i}{C_o} \tag{17}$$

For an imaging system that reproduces the contrast of an image perfectly, the MTF is equal to unity, but for all real imaging systems, MTF $\to 0$ as $s \to \infty$. For example, Fig. 8 shows the MTF of a diffraction-limited $f/8$ lens at a magnification of unity. The figure shows that the MTF immediately begins decreasing as spatial frequency increases, and implies that there are no nonzero frequencies that can be imaged without contrast distortion. This is different from ideal transfer functions in time-series analysis, which generally have a flat response over a wide range and then roll off only at high frequency. In imaging, it is virtually impossible to measure without some level of contrast distortion. The limiting resolution is often specified by a cutoff frequency s_{co}, where the MTF goes to zero. Note that all diffraction-limited MTFs have a universal shape and a cutoff frequency $(s_{co})_{dl}$ that is related to the numerical aperture (NA) on the image side of the lens and the wavelength of light (36). In the literature, it is common to see the cutoff frequency related to the lens $f_\#$ but assuming an infinite conjugate ratio (i.e., object at infinity). However, for noninfinite conjugate ratios and assuming that the image is formed in a medium whose index of refraction is unity, the cutoff frequency depends on the magnification per the relationship $(s_{co})_{dl} = [\lambda f_\#(m+1)]^{-1}$.

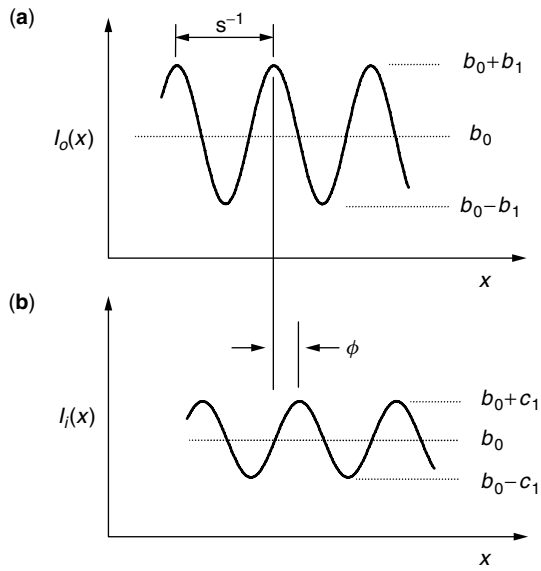

Figure 7. Effect of the imaging system on a sine wave object. (**a**) the irradiance distribution of the object; (**b**) the irradiance distribution of the image resulting from the convolution of the object sine wave with the LSF. The resulting image exhibits contrast reduction and a phase shift ϕ. (Adapted from W. J. Smith, *Modern Optical Engineering: The Design of Optical Systems*, 2e., McGraw-Hill, NY, 1990, with permission of The McGraw-Hill Companies.)

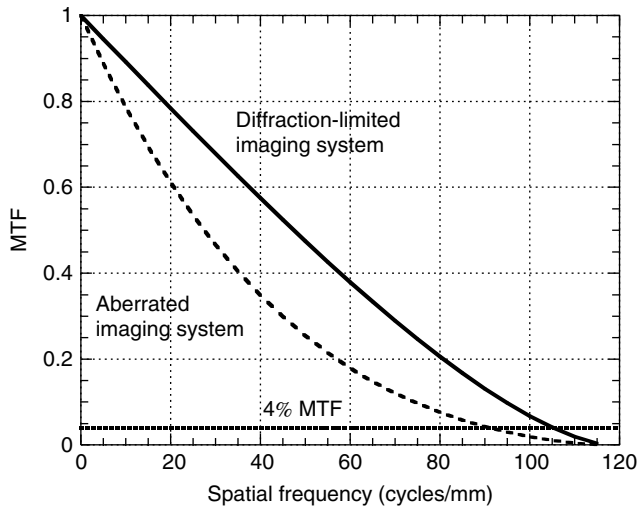

Figure 8. Diffraction-limited MTF for an $f/8$ lens operated at a magnification of unity. Also shown is a hypothetical MTF for an aberrated imaging system. The cutoff frequency for the diffraction-limited MTF is $(s_{co})_{dl} = 117$ cycles/mm.

The diffraction-limited MTF is given by (18)

$$\mathrm{MTF}(s) = \frac{2[\alpha(s) - \cos\alpha(s)\sin\alpha(s)]}{\pi} \tag{18}$$

where, $\alpha(s) = \cos^{-1}[s/(s_{\mathrm{co}})_{\mathrm{dl}}]$. The human eye can distinguish contrast differences of a few percent, and so the cutoff frequency, particularly for Gaussian MTFs, is sometimes specified as the frequency at which the MTF is 0.04, or 4% of the peak value. Figure 8 also shows a hypothetical MTF for an aberrated optical system. The aberrated system exhibits reduced contrast transferability across the entire frequency range and a lower cutoff frequency.

One of the main advantages of the concept of the MTF is that MTFs for different components of an optical system can be cascaded. In other words, the overall MTF is the product of the MTFs of each component. For example, the overall MTF for an intensified camera system is the product of the MTFs for the photocathode, microchannel plate, phosphor screen, optical fiber bundle, and CCD. Because virtually all MTFs exhibit rapid roll-off, the overall MTF is always worse than the worst MTF in the system.

It is enlightening to consider an example of how significantly the MTF can affect a certain type of measurement. Assume that it is desired to measure the irradiance gradient dI_{o}/dx, such as is necessary when computing diffusive fluxes. Consider an object that has a sine wave intensity distribution as given by Eq. (14). It can be shown that the image contrast is given by (18)

$$I_{\mathrm{i}}(x) = b_0 + b_1\mathrm{MTF}(s)\cos(2\pi sx - \phi) \tag{19}$$

The derivatives of both I_{o} and I_{i} are sine waves; for simplicity, consider only the maximum derivative, which occurs at $2\pi sx - \phi = \pi/2$. In this case, the relative error in the maximum gradient (derivative) is

$$\mathrm{Error} = \frac{1}{\dfrac{dI_{\mathrm{o}}}{dx}}\left(\frac{dI_{\mathrm{o}}}{dx} - \frac{dI_{\mathrm{i}}}{dx}\right) = 1 - \mathrm{MTF} \tag{20}$$

Equation (20) shows that the error in the gradient is very large (96%) at the 4% MTF point. If an error no larger than 10% is desired, then the MTF at the frequency of interest must be no less than 0.9. This can be a very stringent requirement for some imaging systems. For the diffraction-limited case shown in Fig. 8, the measurements would be limited to frequencies less than 10 cycles/mm or wavelengths greater than 100 μm. As exemplified in Fig. 8, the situation is typically much worse for an actual aberrated imaging system.

In practice, the MTF is a very difficult thing to measure directly because it is difficult to achieve a true sine wave contrast modulation in the object plane (35). It is relatively easy, however, to produce black-and-white bar patterns of varying frequency, which is why the MTF is often approximated by this method. The response of the system to a periodic black-and-white bar pattern is sometimes called the contrast transfer function (CTF) (also the square-wave transfer function). The CTF is relatively easy to measure, and several square-wave targets are available commercially. However, the CTF is not the same as the MTF, although they are related. Because the FT of a square wave is a sinc function, which exhibits a finite bandwidth of frequencies, the CTF is a reflection of the imaging system's ability to transfer contrast across a range of frequencies, rather than at just a single frequency as for the MTF. The CTF is related to the MTF by the relationship (28)

$$\mathrm{MTF}(s) = \frac{\pi}{4}\left\{\mathrm{CTF}(s) + \frac{\mathrm{CTF}(3s)}{3} - \frac{\mathrm{CTF}(5s)}{5} + \frac{\mathrm{CTF}(7s)}{7} - \frac{\mathrm{CTF}(11s)}{11} + \cdots\right\}. \tag{21}$$

The CTF generally has a shape similar to that of the MTF, but it will have higher values of the transfer function at a given spatial frequency; therefore, measuring the CTF tends to give the impression that the resolution is better than it actually is.

Despite the ease of measuring the CTF, it is not a recommended means of determining the resolution because it is not very accurate, particularly when using discrete sampling detectors, such as CCD arrays (35,37). An array detector can be thought of as a device that averages, owing to the finite size of the pixels (δx), and samples at a frequency that is the inverse of the pixel pitch (spacing) a. When the image, as projected onto the array detector, is sampled at too low a frequency, then aliasing can occur. Aliasing occurs when high-frequency components of the image are incorrectly sampled as lower frequency components and results in spurious contrast modulation in the sampled image. Aliasing can be avoided by ensuring that the image (before sampling) has no frequency content higher than the Nyquist frequency $s_{\mathrm{N}} = (2a)^{-1}$. When the spatial frequency content of the image is higher than the Nyquist frequency, then the resulting spurious frequency content can mislead the user into thinking that the resolution is higher than it actually is (38). In flow imaging, the input optics typically have a cutoff frequency that is higher than the Nyquist frequency of the array, and thus aliasing is often a potential problem. Furthermore, the broad range of frequencies in a square-wave target makes it very difficult to avoid aliasing effects. In fact, the avoidance of aliasing when measuring contrast transfer characteristics is imperative because the MTF of a discrete sampling detector is not even defined when aliasing is present (35,37). The reason is that for a device to have an MTF, it must be linear and isoplanatic. Isoplanatic means that the output image is insensitive to movement of the input image. Array detectors are typically sufficiently linear, but they are not necessarily isoplanatic. For example, consider the case where a white/black bar pattern is imaged at a magnification of unity and where the spacing of the bars is equal to the pixel pitch. In this case, the contrast modulation of the image will depend on whether the bars are "in-phase" (aligned with the pixels), or "out-of-phase" (straddling the pixels). Such nonisoplanatic behavior is mainly a problem at spatial frequencies near the Nyquist limit. For this reason, MTFs

for CCDs and other detectors can be considered "pseudo"-MTFs only, which have a limited range of applicability. For example, it has been shown that array detectors are approximately isoplanatic for frequencies lower than S_N (35).

From purely geometric considerations, the array MTF follows a sinc function,

$$\mathrm{MTF}(s) = \frac{\sin(\pi \delta x\, s)}{\pi \delta x s}, \tag{22}$$

which goes to zero at a frequency of $s = 1/\delta x$. In practice, the MTF will be smaller than given by Eq. (22), owing to the diffusion of photon-generated charge carriers, light scatter between detector elements, reflections between the array and the protective window, and nonideal charge-transfer efficiency. For video systems, the processing electronics and frame grabber will also reduce the quality of the MTF.

Several studies have shown that a useful means of inferring the MTF is by measuring the line-spread function (LSF). The LSF is the 1-D analog of the PSF because it is the intensity distribution at the image plane resulting from imaging an infinitesimally narrow slit at the object plane. The importance of the LSF is that its FT is the OTF (35). Furthermore, if the LSF is a symmetrical function, then the OTF is real, indicating that there is no phase distortion and the PTF is zero. If the intensity distribution of the PSF is given by $p(x, y)$, then the LSF irradiance distribution is

$$l(x) = \int_{-\infty}^{\infty} p(x, y)\, dy. \tag{23}$$

Consider the sampled PSF represented by the image of Fig. 6c. Because the LSF covers such a small range of pixels, it is not known how the actual LSF is affected by array sampling. For example, if the LSF contains spatial frequency content that is higher than the Nyquist frequency, then aliasing is present, and the sampled LSF may not reflect the true LSF.

There is, however, a superior technique for measuring the LSF that does not suffer from aliasing (39). In this technique, the object (whether sine wave or line source) is translated within the object plane (say in the x direction), and the output from a single pixel is monitored as a function of the x location. This technique is free from aliasing errors because the LSF is sampled at only a single point and the pitch of the measurement (i.e., the resolution) can be much finer than the pixel pitch. For example, it is not difficult to obtain 1-μm resolution on standard optical translation stages, which is substantially smaller than the pitch of most CCD arrays.

Because good sine wave and line sources may be difficult to generate in practice, a relatively easy technique is to measure the step response function (SRF), which is the intensity distribution at the image plane obtained by scanning a knife-edge across the object plane. In this case, the output of a single pixel is also measured as a function of the knife-edge position. The SRF irradiance distribution $k(x)$ is the convolution of a step function with the LSF. It necessarily follows that the derivative of $k(x)$ is the LSF

$$l(x) = \frac{dk(x)}{dx}. \tag{24}$$

Figure 9 shows example an setup for obtaining the LSF by scanning a knife-edge and monitoring the output from a single pixel. Figure 10 shows the SRF obtained using this same setup, for $m = 1$, $f/2.8$, where the knife-edge was translated in 2-μm increments. A single 9-μm pixel near the center of the field of view was monitored, and the resulting SRF was very well resolved. Figure 10 also shows an error function curve fit to $k(x)$, where the error function provides a reasonably good fit to the data. Also shown in Fig. 10 is the Gaussian LSF obtained by differentiating the error function curve fit. The LSF is seen to have a $1/e^2$ full width of about 40 μm, which corresponds to about 4.5 pixels. The point source images of Fig. 6c indicate a larger LSF, but the heavy quantization and the potential for aliasing makes this difficult to determine from these types of images. The MTF, which is the FT of the LSF (and

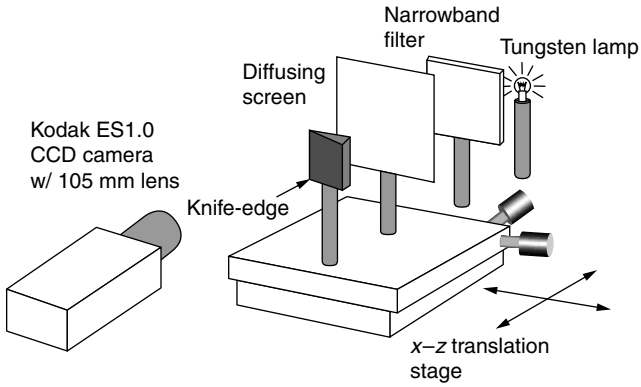

Figure 9. Schematic of the setup for measuring the step response function (SRF) for a single pixel of a CCD camera. The camera images the back-illuminated knife-edge, and the output of a single pixel is monitored as the knife-edge is translated across the field of view.

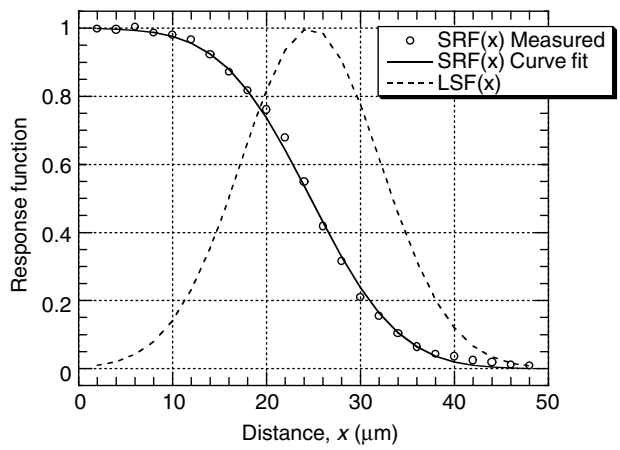

Figure 10. Measured SRF for an $f/2.8$ lens operated at unity magnification. The dashed line is the LSF computed from the derivative of the curve fit to the SRF.

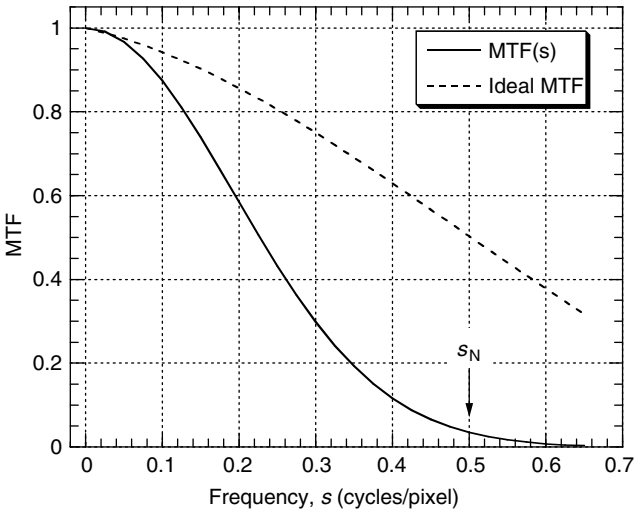

Figure 11. Comparison of MTFs for an $f/2.8$ lens and 9-μm pixel array operated at unity magnification. The Gaussian MTF was inferred from the measured LSF shown in Fig. 10, and the ideal MTF was computed assuming a diffraction-limited lens and a geometric sampling function for the CCD detector.

is also Gaussian) is shown in Fig. 11. From the figure, it is seen that the resolution of this system is really not very good because sine wave structures whose frequency is 0.2 cycles/pixel (or a wavelength of 5 pixels) will exhibit a 40% contrast reduction. The Nyquist frequency S_N is associated with an MTF of about 5% and emphasizes the danger of specifying the resolution in terms of the projection of a pixel into the field of view. An "ideal" MTF is also shown for comparison. The ideal MTF is the product of the MTFs for a diffraction-limited lens (at $f/2.8$, $\lambda = 532$ nm, and $m = 1$) and an ideal sampling detector whose pixel size is 9 μm [i.e., the product of Eqs. (18) and (22)]. The figure shows that the measured MTF is substantially worse than the ideal one, owing largely to aberrations in the lens.

Note that because taking a derivative is a noise-enhancing process, if the SRF cannot be fit to a relatively simple functional form, such as an error function, this makes the determination of the LSF much more difficult using this technique. In some cases, it may be worth the trouble of measuring the LSF directly by using a narrow slit rather than a knife-edge.

Paul (26) shows in a planar imaging experiment, that the MTF will be a function of the laser sheet thickness when the sheet thickness is greater than the depth of field of the imaging system. The *depth of field* δ_{df} is the distance that the object may be shifted in the direction of the lens and still maintain acceptable blur, whereas the *depth of focus* δ'_{df} is the distance that the detector can be shifted and maintain acceptable blur. Note that the two are related by the magnification, that is, $\delta'_{df} = m^2\delta_{df}$. If the laser sheet is larger than the depth of field, then the region across which the imaging system collects light will be a "bow-tie" shaped region, rather than the "box" region shown in Fig. 3. Therefore, near the tails of the laser sheet, the blur spot may be substantially larger than at best focus. The depth of field is related to the blur spot of the imaging system per the relationship (18)

$$\delta_{df} = \frac{f(m+1)d_{blur}}{m(D \pm d_{blur})}. \tag{25}$$

The ± sign in Eq. (25) indicates that the depth of field is smaller in the direction of the lens and larger away from it. The total depth of field δ_{tot} is the sum of the depths of field toward and away from the lens. When $d_{blur} \ll D$, which is so for most flow imaging cases, then the total depth of field simplifies to

$$\delta_{tot} \approx 2d_{blur}f_{\#}\frac{m+1}{m}. \tag{26}$$

For the diffraction-limited case, the blur spot size is given by Eq. (13). Equation 26 shows that the depth of field increases as blur spot size increases and decreases for increasing magnification. For example, the blur spot of Fig. 6c is about 50 μm, which at $f/2.8$ and $m = 1$ amounts to $\delta_{tot} = 560$ μm. This is somewhat larger than the typical laser sheet thicknesses that are used in planar imaging of scalars, and therefore, it is unlikely that additional blur at the edge of the sheet would be an issue.

In many cases, such as using faster optics, this effect will not be negligible. One way to account for the collection of light over the "bow-tie" shaped volume is given in (26), where the MTF was evaluated as the weighted sum of the MTFs of thin laminates (infinitesimally thin planes) parallel to the laser sheet but at different z locations. The weighting function used was the laser sheet energy distribution. This technique of weighting the MTFs by the energy distribution accounts for the fact that more energy will be collected from regions that have smaller blur spots. However, this technique requires either the assumption of ideal MTFs or detailed system MTF measurements at a number of z locations. Another approach is to measure the MTF by the knife-edge technique at best focus and at the edge of the laser sheet. To be conservative, the MTF at the edge of the sheet could be used as the primary measure of resolution, although a reasonable compromise might be to take the average of these two MTFs as the representative MTF of the entire system, including the laser sheet. It is also important to note that the MTF is measured at a given point in the image, but it may vary across the field of view. For this reason, it is also advisable to measure the MTF at the center and near the edges of the field of view.

RESOLUTION REQUIREMENTS IN FLUID FLOWS

One of the major difficulties in flow imaging is achieving adequate spatial and temporal resolution. This is particularly the case when flows are turbulent because the resolution requirements are typically very severe if it is desired to resolve the smallest scales at which fluctuations occur. Laminar flows, however, pose substantially less stringent requirements on resolution, compared to turbulent flows. The primary issue when considering the resolution requirements is the gradient of the flow property that is being measured because the gradient determines the amount of averaging that occurs across the

resolution volume. In many laminar shear flows including boundary layers, pipe flows, wakes, jets, and mixing layers, the maximum gradient is the same order of magnitude as the overall gradient. In other words, the maximum velocity and temperature gradients are approximately $(\partial U/\partial y)_{\max} \sim \Delta U/\delta$ and $(\partial T/\partial y)_{\max} \sim \Delta T/\delta$, where ΔU is the characteristic velocity difference, δ is the local width of the shear flow, and ΔT is the characteristic temperature difference across the flow. For example, in a boundary layer formed by the flow of air over a heated flat plate, the maximum velocity gradients in these flows scale as $(\partial U/\partial y)_{\max} \approx U_\infty/\delta \sim (U_\infty/x)Re_x^{1/2}$, where $Re_x = U_\infty x/\nu$, x is the downstream distance, and ν is the kinematic viscosity. Gradients in scalars, such as temperature or species concentration, will similarly scale with Reynolds number, but will also depend on the relative diffusivities for momentum and the scalar. For example, the maximum scalar gradient in the boundary layer will scale as $(\partial T/\partial y)_{\max} \approx [(T_\infty - T_w)/x](Re_x Pr)^{1/2}$, where T_∞ and T_w are the free-stream and wall temperatures, respectively, $Pr = \nu/\alpha$ is the Prandtl number, and α is the thermal diffusivity. The preceding relationships show that gradients become large at large Reynolds and Prandtl numbers (or Schmidt number, $Sc = \nu/\mathscr{D}$, where $\mathscr{D}$ is the mass diffusivity for mass transfer), which is the same as saying that shear flows become "thin" at high Re (and Pr).

Turbulent flows have substantially more severe resolution requirements than laminar flows, owing to the much larger gradients that occur at the smallest scales of turbulence. In turbulent flows, the spatial fluctuations in flow properties, such as velocity, temperature, or concentration range in scale from the largest physical dimension of the flow (e.g., the local width of the boundary layer or jet) to the scale at which diffusion acts to remove all gradients. The largest scales are often called the "outer scales," whereas the smallest scales are the "inner" or dissipation scales because these are the scales at which the energy of fluctuations, whether kinetic or scalar, is dissipated.

In classical turbulence theory, the kinetic energy dissipation scale is the Kolmogorov scale (40),

$$\eta \equiv \left(\frac{\nu^3}{\varepsilon}\right)^{1/4}, \tag{27}$$

where ε is the kinetic energy dissipation rate. Batchelor (41) argued that the smallest scale of scalar fluctuations λ_B, called the Batchelor scale, is related to the Kolmogorov scale and the ratio of the kinematic viscosity to the scalar diffusivity. For Sc(or Pr) $\gg 1$, he argued that $\lambda_B = \eta Sc^{-1/2}$. There is some disagreement in the literature about the scaling for fluids when $Sc \ll 1$ (42), but because most gases and liquids have Schmidt numbers of order unity or larger, this is of little practical concern. Generally, it is assumed that the $Sc^{-1/2}$ scaling applies at near unity Schmidt numbers, in which case $\lambda_B \approx \eta$. For liquids, it is typical that Pr, $Sc \gg 1$; thus the $Sc^{-1/2}$ scaling is appropriate, in which case $\lambda_B \ll \eta$.

Using scaling arguments, the Kolmogorov scale can also be related to outer scale variables through the relationship, $\eta \propto Re^{-3/4}$, where Re is the Reynolds number based on outer scale variables (such as ΔU, the maximum velocity difference, and δ, the local width of the shear flow). Buch and Dahm (43) make explicit use of such an outer scaling by defining the strain-limited scalar diffusion scale λ_D, as

$$\frac{\lambda_D}{\delta} = \Lambda Re_\delta^{-3/4} Sc^{-1/2} \tag{28}$$

where δ is the 5–95% velocity full width of the shear flow and $Re_\delta = \Delta U \delta/\nu$. Their planar imaging measurements of the finest mass diffusion scales in round turbulent jets suggest that $\Lambda \approx 11$. Similar measurements in planar jets suggest a value of $\Lambda \approx 14$ (32). The finest velocity gradient scale, analogous to the Kolmogorov scale, is the strain-limited vorticity scale, $\lambda_\nu = \lambda_D Sc^{1/2}$. The strain-limited diffusion scales can be related to the Kolmogorov scale by using measurements of the kinetic energy dissipation rate. For example, using the data for the decay of the kinetic energy dissipation rate for gas-phase round jets (44) and taking $\Lambda = 11$, it can be shown that $\lambda_D \approx 6\lambda_B$ and $\lambda_\nu \approx 6\eta$. If the mean kinetic energy dissipation scales are about 6η, then accurate measurements of the gradients will necessitate better resolution than this. This is consistent with thermal-wire measurements of temperature and velocity fluctuations, which suggest that a resolution of about 3η is sufficient for correct measurements of the smallest scale gradients (45–47). Therefore, it is recommended that the resolution of the imaging system be no worse than $\lambda_D/2$ and $\lambda_\nu/2$, if the smallest fluctuations in a turbulent flow are to be measured accurately. It cannot be emphasized enough that because of the nature of the MTF of the imaging system, it is too simplistic to speak of "resolving" or "not resolving" particular scales in the flow. Progressively finer scales will be increasingly affected by the imaging system, and any quantitative measurement of gradients must take this into account.

Another perspective on Eq. (28) is that it describes the *dynamic spatial range* that is required for measuring the full range of scales. Here, the *dynamic spatial range* (DSR) is defined as the ratio of the largest to the smallest spatial structures that can be measured. The *largest* spatial scale in turbulent flows is generally considered the local width of the shear flow (or in some enclosed flows, a characteristic dimension of the enclosing box). Therefore, for turbulent shear flows, δ/λ_D given by Eq. (28), is the range of scales of the flow. This also shows that the Reynolds (and Schmidt) number can be thought of as directly related to the DSR of the turbulent shear flow. Equation (28) also shows that the DSR for scalars is even larger for low scalar diffusivity (high Sc or Pr numbers). For example, fluorescein dye in water has a diffusivity of about 2000, and thus the finest mass diffusion scale is about 45 times smaller than the smallest vorticity scale (42). The other important point that Eq. (28) reveals is that the DSR is a strong function of the Reynolds number; thus, it is often not possible to resolve the full range of turbulent scales by using currently available camera systems. For example, assume that it is desired to obtain planar images of the jet fluid concentration in a turbulent round jet and to resolve the full range of scales 500 mm downstream of a 5-mm diameter nozzle. The jet velocity 5% full width grows at a rate of $\delta(x) = 0.44x$, where x is the distance downstream of the jet exit and the centerline

velocity decays as $U_c/U_0 = 6.2/(x/d_j)$, where U_c is the centerline velocity, U_0 is the jet exit velocity, and d_j is the jet exit diameter (48). In this case, the outer scale Reynolds number, $Re_\delta = U_c\delta/\nu = 1.9Re_d$, where Re_d is the source Reynolds number ($= U_0 d_j/\nu$). If we desire to study a jet where $Re_d = 20,000$, then the range of scales given by Eq. (28) is 150. If a smallest scale of $\lambda_D/2$ must be resolved, then our required DSR is $2\delta/\lambda_D = 300$. In planar imaging using the 1000×1000 pixel CCD camera whose measured MTF is shown in Fig. 11, it would not be possible to resolve the entire range of scales because substantial blurring occurs across 4–5 pixels.

Turbulent timescales are generally bounded by relatively low frequency outer scale motions and high-frequency inner scale motions (40). The largest scale motions are independent of viscosity and occur over a characteristic time that is of the order of $\tau_{os} \sim \delta/\Delta U$. This is also commonly referred to as the "large-eddy-turnover" time. The small-scale motions, however, occur over a substantially shorter timescale, which is of the order of $\tau_{is} \sim (\nu/\varepsilon)^{1/2}$ or $\tau_{is} \sim (Re_\delta)^{-1/2}\tau_{os}$, if based on outer scale variables. This latter relationship shows that, at high Reynolds numbers, the inner scale timescales can be orders of magnitude smaller than outer scale timescales.

The turbulent inner scale timescales may not be the shortest timescales that must be resolved, if the flow is convecting past the measurement volume. In this case, the shortest time may be the convective inner scale time $(\tau_{is})_{conv} = \lambda_\nu/U$, where U is the local velocity. For example, consider a mixing layer that forms between two parallel streams of air, where the streams have velocities of 100 m/s and 90 m/s. The range of turbulent spatial scales will depend only on the outer scale Reynolds number, which in turn depends only on the velocity difference of 10 m/s. The absolute velocities are irrelevant, except to the extent that they affect the local mixing layer thickness δ. If the imaging system is in the laboratory frame of reference, then the timescales will depend on both the velocity difference (which drives the turbulence) and the bulk convection of these spatial structures, which depends on the local velocity of the structures with respect to the imaging system. For the mixing layer conditions given before, if the mixing layer at the imaging location is 10 cm thick, then $\tau_{os} \approx 10$ ms, and $\tau_{is} \approx 40$ μs. However, if the small-scale structures convect by the measurement station at the mean velocity of $U_{conv} = (U_1 + U_2)/2 = 95$ m/s, then the timescale that needs to be resolved is $(\tau_{is})_{conv} = \lambda_\nu/U_{conv} = 3$ μs, which is considerably less than τ_{is}. It is clear that the smaller of the convective and turbulence timescales must be resolved.

FLOW IMAGING: SURVEY OF TECHNIQUES

The purpose of this section is to give the reader an idea of the wide range of flow imaging techniques that have been developed and applied in fluid mechanics research. Owing to space limitations, however, this survey must leave out many techniques that are certainly worthy of discussion. Hopefully, in most cases, a sufficient number of general references is provided for readers to learn about these omitted techniques on their own. Furthermore, excellent reviews of a number of qualitative and quantitative flow visualization techniques, including some that were omitted in this article, can be found in (5). The reader should keep in mind that the physical principles underlying each technique are usually not discussed in this article because they are covered in different sections of this encyclopedia and in the references cited in the bibliography.

This section is organized on the basis of the flow variable to be imaged, because in most cases the user starts with a need (say, for temperature imaging in an aqueous flow) and then must find the technique that best addresses that need. Because some techniques can be used to measure several flow variables, their use may be described under more than one category. Therefore, to avoid too much redundancy, a technique is described only the first time it is mentioned; thus, the uninitiated reader may need to read the article all the way through rather than skipping to later sections.

Density Gradients (Schlieren and Shadowgraph)

Two of the most widely used techniques for qualitative flow visualization, particularly in high-speed flows, are the schlieren and shadowgraph techniques. Although the main emphasis of this article is on quantitative planar imaging techniques, the shadowgraph and schlieren techniques will be briefly discussed because of their extensive use in gas dynamics. Furthermore, the mechanism of light ray deflection by index-of-refraction gradients, which is the basis for these techniques, is a potential source of error in quantitative laser imaging. In their most commonly used forms, the schlieren and shadowgraph techniques provide line-of-sight integrated information about gradients in the index-of-refraction field. Because the index of refraction is related to gas density, in fluid flows such as air whose composition is uniform, the schlieren technique is sensitive to variations in the first derivative of density, and the shadowgraph to the second derivative. Interferometry is a quantitative line-of-sight technique that enables imaging the density field, but it will not be discussed here because it is becoming increasingly supplanted by planar imaging techniques. Because these techniques are spatially integrated along the line of sight, they are limited in the quantitative information that can be inferred in complex, three-dimensional flows. Further details of these techniques can be found in several excellent references (5,9,10,49).

The physical basis for the shadowgraph and schlieren techniques is that spatial variations in the index of refraction of a transparent medium cause spatial variations in the phase of plane light waves (33). The index of refraction is defined as $n = c_0/c$, where c_0 is the speed of light in vacuum and c the speed of light in the medium. When traveling through a medium when $n > 1$, the phase of the transmitted wave undergoes a negative phase shift, owing to a lag in the oscillations of the induced dipoles within the medium. For this reason, an object that causes such a phase shift is termed a "phase object," and it can be contrasted with an "amplitude object," such as an opaque disk, which changes the amplitude of the light waves. Because the velocity of light is usually considered the

"phase velocity," which is the velocity of a point of constant phase on the wave, the phase shift can be interpreted as a change in the velocity of the transmitted wave.

Both the schlieren and shadowgraph techniques are analyzed by considering how a plane wave is affected by propagation through an index-of-refraction gradient. Consider the propagation of a plane wave in the z direction through a transparent medium that has a gradient of n in the y direction. It can be shown that the angular deflection θ_y in the y direction is given by (9)

$$\theta_y = \int_L \frac{1}{n}\frac{\partial n}{\partial y}\,dz, \tag{29}$$

where the integration is along the line of sight and over the path length L. Equation (29) shows that the angular deflection increases for increasing gradients and longer path lengths. The equation also shows that the light rays are bent in the direction of the gradient, that is, the rays are bent toward regions of higher index of refraction.

In gases, the index of refraction is related to the fluid density ρ by the Gladstone-Dale relationship (9),

$$n = 1 + K\rho, \tag{30}$$

where K is the Gladstone–Dale constant. For example, for 633-nm light and $T = 288$ K, $K = 2.26 \times 10^{-4}$, 1.57×10^{-4}, 1.96×10^{-4} m^3/kg, for air, argon and helium, respectively. In water, which is largely incompressible, the index of refraction varies primarily with temperature. For example, for 632.8-nm light, the index of refraction across the temperature range of 20–34 °C is given by (9)

$$n(T) = 1.332156 - 8.376 \times 10^{-5}(T - 20\,^\circ\mathrm{C}) - 2.644 \times 10^{-6}(T - 20\,^\circ\mathrm{C})^2 + 4.79 \times 10^{-8}(T - 20\,^\circ\mathrm{C})^3 \tag{31}$$

An example of a schlieren setup is shown in Fig. 12. For instantaneous imaging, the light source is usually a point source of short duration (typically a microsecond or lower); common sources are xenon flash lamps and lasers. In most cases, a flash lamp is preferred to a laser source because lamps are cheaper and the coherence and mode structure of most pulsed lasers causes a nonuniform image. In some cases, however, such as in plasmas where the background luminosity is very high, the high brightness of a laser is a necessity. In this case, the beam must be spatially filtered to improve its spatial uniformity (33). This is usually accomplished by tightly focusing the beam through a small pinhole by using a microscope objective lens. Note that it is typically very difficult to focus the laser beam onto such a small pinhole, and an integrated lens/pinhole mount that has three axes of translation is necessary. A further problem when using a pulsed laser is that it is difficult to keep from burning the pinhole material, owing to the very high peak intensity at the focus. This problem can be alleviated by substantially reducing the energy of the beam. Although such low laser energies will result in weaker signals, obtaining a sufficient signal is not usually a problem in the schlieren and shadowgraph techniques because the beam is usually directed into the camera (Fig. 12). When using a flash lamp that has an extended arc, the point source is approximated by imaging the arc onto a small aperture (e.g., submillimeter diameter) with a lens. The sensitivity of the schlieren system will be improved by a smaller point source, but the signals are reduced accordingly. For many flash lamps, the arc is small enough that it may not be necessary to use any spatial filter at all.

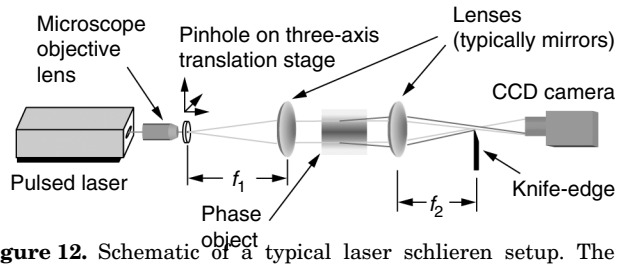

Figure 12. Schematic of a typical laser schlieren setup. The undeflected rays are shown in gray, whereas the deflected rays are shown in black. The knife-edge blocks rays deflected downward by negative gradients, which renders those gradients dark in the image. In contrast, the rays deflected upward by the positive gradients miss the knife-edge and are rendered light in the image.

The point light source is collimated by what is typically a large diameter spherical mirror, which is at least as large as the object that is being imaged. Lenses can also be used when smaller fields of view are desired, and this is the situation shown in Fig. 12. The mirror/lens is placed one focal length from the source, which collimates the beam. After the beam passes through the test section, it is then directed to a second mirror/lens (called the "schlieren head"), which refocuses the beam. In the conventional schlieren setup, a knife-edge (e.g., razor blade) is placed at the second focal spot, as shown in Fig. 12. The horizontal knife-edge shown produces an optical system that renders upward density gradients as light and downward gradients dark. This occurs because the rays are deflected up by the upward gradients and thus miss the knife-edge, whereas the knife-edge blocks the rays that are deflected down by the downward gradients. The analysis of the schlieren intensity variations is conceptually simpler for a line light source, which forms a line image at the focus. In this case, the relative intensity variations at the film plane are given by (9)

$$\frac{\Delta I}{I} = \frac{f_2}{a}\int_L \frac{1}{n}\frac{\partial n}{\partial y}\,dz, \tag{32}$$

where I is the intensity of the image when no gradient is present, $\Delta I = I^* - I$, I^* is the intensity when the gradient is present, f_2 is the focal length of the schlieren head, and a is the height of the focal image that is not blocked by the knife-edge. Equation (32) shows that a longer focal length schlieren head and decreasing height of the transmitted portion of the image at the focus increases sensitivity. Interestingly, increasing the distance between the phase object and the focusing lens/mirror does not affect the sensitivity.

The focus is found by traversing the knife-edge along the optical axis. When the knife-edge is upstream of the focus, the image reveals an inverted shadow of the knife-edge, whereas when the knife-edge is downstream of the focus, the image reveals an upright shadow. It is only at the focus that inserting the knife-edge into the focal spot results in a uniform reduction of intensity of the image and no shadow of the knife-edge. An example of a schlieren image of a supersonic helium jet issuing into room air is shown in Fig. 13 (from Ref. 50). In this figure the knife-edge was horizontal, and therefore the vertical n-gradients are visualized. Because helium has a very low index of refraction, the n gradients resulting from mixing are very distinct. Furthermore, even subtle features such as the Mach waves in the ambient air are visualized as the dark lines to the outside of the jet.

A useful extension of the schlieren technique is "color" schlieren, which uses a white light source combined with a transparency of varying color in place of the knife-edge (5,9,10). Because the eye is better able to distinguish colors than shades of gray, color schlieren is superior for visualizing the density gradients in flows. Although most color schlieren is used for flow visualization, it has also been used to obtain quantitative temperature data in flows by relating the color of the image to the angular deflection of the light rays (51). When used with an axisymmetric phase object, this technique enables the tomographic reconstruction of the three-dimensional temperature field (52).

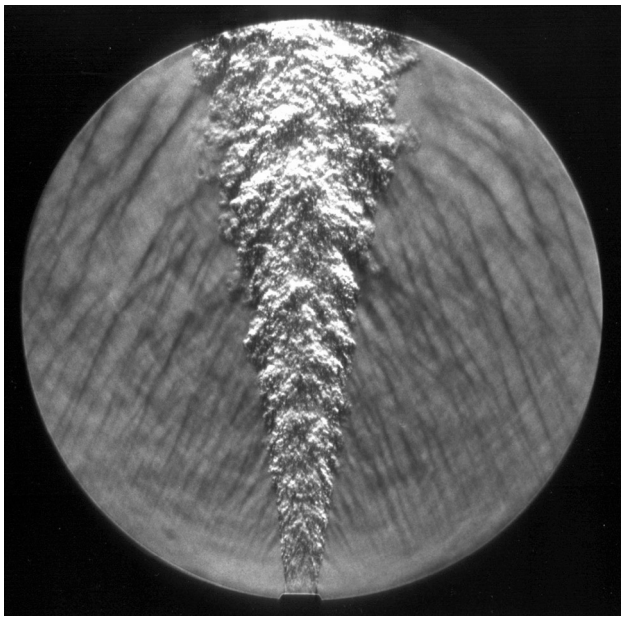

Figure 13. Sample schlieren image of a Mach 2 jet of helium exhausting into room air. The knife-edge is oriented horizontally, thus the vertical index of refraction gradients are visualized. The image reveals fine structures of the jet turbulence in addition to Mach waves that are generated by structures that travel at supersonic speeds with respect to the ambient. (Reprinted with permission from Mach Waves Radiating from a Supersonic Jet by N. T. Clemens and P. H. Paul, *Physics of Fluids A* **5**, S7, copyright 1993 The American Institute of Physics.)

An interesting way of looking at the function of the knife-edge is as a filter, which acts on the spatial frequencies in the phase object. This can be seen by considering that the second focus is called the "Fourier transform plane," because the intensity distribution at the focus is related to the spatial frequency content of the phase object (33,53). Higher spatial frequencies are associated with increasing radial distance from the center of the focal spot. The dc component is the neutral intensity background present when there is no phase object, and it can be filtered out if an opaque disk is used as the spatial filter. In this case, when the phase object has no high-frequency content, then the image will be uniformly dark. When higher spatial frequencies are present in the phase object, the disk will not block them, and they will be visualized as light regions in the image. It can be seen from this that the shape of the spatial filter can be tailored to visualize different frequencies in the phase object. This principle has been used to develop a system that directly measures the power spectrum of the line-of-sight integrated index of refraction fluctuations in turbulent flows (54,55).

Note that an alternative technique has been developed, named the "focusing schlieren" technique, which enables visualizing density gradients as in conventional schlieren, but the depth of field along which the signal is integrated can be just a few millimeters (56). The focusing schlieren technique can yield nearly planar images of the density gradient field at substantially lower cost than by planar laser imaging. In some cases, such as a large-scale wind tunnel where optical access is limited, it may be the only means of acquiring spatially resolved image data.

The shadowgraph effect can be understood from simple geometrical ray tracing, as shown in Fig. 14. Here a plane wave traverses a medium that has a nonuniform index-of-refraction gradient and is allowed to illuminate a screen. The rays traversing through the region that has no gradient are not deflected, whereas the rays traversing the region that has an the upward gradient are bent up. The resulting image on the screen consists of regions where the rays converge and diverge; these appear as regions of light and dark, respectively. It is this effect that gives the technique its name because gradients leave a shadow, or dark region, on the viewing screen.

It can be shown that the intensity variations on the screen follow the relationship (9)

$$\frac{\Delta I}{I} = L \int \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) (\ln n)\, dz. \tag{33}$$

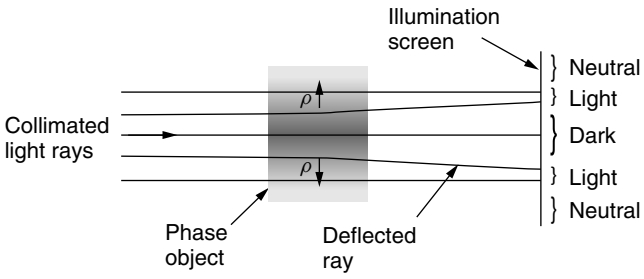

Figure 14. Illustration of the shadowgraph effect.

For gas flows, incorporating the Gladstone–Dale relationship into Eq. (33) shows that the shadowgraph technique is sensitive to the second derivative of the density along the line of sight of the light beam.

A shadowgraph system can be set up almost trivially by using an approximately collimated light source and a screen. For example, a shadowgraph system suitable for classroom demonstrations can be made by expanding the beam from a laser pointer using a short focal length lens and projecting the beam onto a wall a few meters away. This simple system will enable the visualization of the thermal plume rising from a candle flame. Despite the simplicity of this system, more sophisticated setups are typically desired. For example, the schlieren setup shown in Fig. 12 can be used for shadowgraph by simply removing the knife-edge. However, unlike schlieren, where the camera is focused on the phase object, the camera must be slightly defocused to produce sufficient divergence of the deflected rays on the image plane. This feature enables one to "focus out" the shadowgraph effect in a schlieren system. An obvious disadvantage to this technique is that any amplitude objects in the image (e.g., a bullet) will be slightly out of focus. The problem of slight defocus is generally tolerable, compared to the advantages of being able to alternate quickly between the schlieren and shadowgraph techniques.

Concentration/Density

Imaging the concentration of a particular type of fluid or chemical species is primarily of interest in studies of mixing and combustion. Concentration and density are related quantities in that they both quantify the amount of a substance per unit volume. Because most optical diagnostic techniques are sensitive to the number of scatterers per unit volume, rather than to the mass per unit volume, the concentration is the more fundamental quantity. Of course, density can be inferred from the concentration if the fluid composition is known. Concentration imaging is of interest in nonreacting mixing studies and in reacting flows for investigating the relationship between the chemical state of the fluid and the fluid mechanics.

Planar laser-induced fluorescence imaging is probably the most widely used technique for quantitative scalar imaging because it can be used in liquids and gases, it is species specific, and its high signals enable measuring even minor species in gas-phase flows (11,27,57). In laser-induced fluorescence (LIF), a laser is used to excite an atom or molecule from a lower energy state into a higher energy state by the absorption of a photon of light. The frequency of light required is related to the energy difference between the states through the relationship $E = h\nu$, where E is the energy per photon and ν is the frequency of light. The excited state is a state of nonequilibrium, and thus the atom/molecule will tend to return to equilibrium by transiting to a lower energy state. The return to the lower state can occur by several processes, including spontaneous emission of a photon of light (fluorescence); stimulated emission by the incident laser light; "quenching," that is, the transfer of energy to other atoms/molecules through molecular collisions; and by internal energy transfer, or the transfer of energy to other energy modes within the molecule. Because the probability of quenching depends on local thermodynamic conditions, the LIF signal is in general a function of several flow variables, including the concentrations of all species present, temperature, and pressure. Furthermore, the theoretical dependence of the LIF signal on the flow variables depends on the specific model of the energy-transfer physics. Because properly modeling the physics is an important part of quantifying PLIF measurements, PLIF can be a particularly challenging technique to use. The dependence of the signal on many variables presents both an opportunity and a disadvantage for making quantitative measurements. The opportunity is that PLIF can be used to measure a range of flow variables for a remarkable number of chemical species. However, it is generally very difficult to relate the LIF signal to a particular variable of interest (e.g., species concentration) because the signal depends on so many other flow variables, which may not be known. For example, in using PLIF for OH, which is commonly used in flames as an approximate marker of the reaction zone, the PLIF signal is a function of the OH mole fraction, the mole fractions of several other species, including N_2, O_2, H_2O, and CO_2; and the temperature. Because it is virtually impossible to measure all of these variables, the signal can be quantified only by assuming a certain level of knowledge about the thermochemical state of the flow (e.g., equilibrium chemistry).

Despite the caveat about the difficulties that can be encountered when using PLIF imaging, there are many cases where PLIF imaging is in fact relatively simple to implement. The first case is using PLIF in liquid flows. PLIF in liquids, particularly water, is achieved by seeding a fluorescent organic dye into the flow. Because many liquids are essentially incompressible and isothermal, the PLIF signal is usually a function only of the dye concentration and therefore is ideal for mixing studies (6,58,59). Fluorescent dyes absorb light across a very broad range of wavelengths, and thus they can be stimulated by using a number of different lasers. Some of the more popular dyes for aqueous flows include fluorescein, rhodamine B, and rhodamine 6G; all of their absorption bands overlap one or more emission lines of the argon-ion, copper-vapor, and doubled Nd:YAG lasers. Because of this and because liquid-phase PLIF tends to exhibit high signal levels (due to the high density of the fluid), excellent results can usually be achieved without highly specialized equipment.

It is important to note that some dyes, suffer from photobleaching effects at high laser intensity (or fluence), which can lead to significant errors in concentration measurements (60–63). Photobleaching is the reduction in the concentration of fluorescent molecules due to laser-induced photochemistry. Both fluorescein and rhodamine 110 are particularly problematic, and (60) even suggests abandoning the use of fluorescein in favor of rhodamine B.

Another important issue in using PLIF of organic dyes is that the high signals are often a result of the high absorption coefficient of the dye solution. In this

case, substantial laser beam attenuation is encountered when the optical path lengths are relatively large. Beam attenuation can be alleviated by reducing the dye concentration along the beam path or by reducing the optical path length; however, this is often not possible, owing to SNR considerations or other practical limitations. Alternatively, attenuation along the ray path can be corrected for by using the Beer-Lambert absorption law, provided that the entire path length of a given ray of the laser sheet is imaged (11,64).

PLIF is also relatively easy to implement in nonreacting gas-phase flows, where the flow can be seeded with a gas-phase tracer species. By far the most popular tracer to date is acetone, although biacetyl, NO, and I_2 have also been used to a more limited extent. Acetone (CH_3COCH_3) is an excellent tracer species in nonreacting flows because it is relatively nontoxic, fairly easy to seed into flows, generally provides good signals, and can be pumped at a range of UV wavelengths (65). A characteristic feature of polyatomic molecules, such as acetone, is that they have broad absorption bands. The absorption band of acetone ranges from about 225 to 320 nm, and thus it is readily pumped by quadrupled Nd:YAG (266 nm) and XeCl excimer (308 nm) lasers. Furthermore, although its fluorescence efficiency is not very high (about 0.1–0.2%), its high saturation intensity means that high laser energies can be used, which compensates for any limitation in fluorescence efficiency. A small sample of studies where acetone PLIF was used for concentration measurements includes jets (32,65), jets in crossflow (66), supersonic shear layers (67), and internal combustion engines (68).

Biacetyl ($CH_3(CO)_2CH_3$) is another low toxicity seed species that has been used in nonreacting flows and to a lesser extent in flames. Biacetyl vapor absorbs in the range 240–470 nm and exhibits blue fluorescence over the range 430–520 nm and green phosphorescence over the range 490–700 nm. The quantum yield, that is, the fraction of emitted photons to absorbed photons is 15% for phosphorescence, but is only 0.2% for fluorescence. For this reason, biacetyl phosphorescence has been used to produce very high SNR imaging (29). Several different lasers have been used for biacetyl pumping, including dye lasers, excimers, and frequency-tripled Nd:YAGs. One drawback to using biacetyl is that O_2 quenches the phosphorescence, which leads to a significantly lower SNR when biacetyl is seeded into air. Furthermore, the long lifetime of the phosphorescence (about 1 ms) can be severely limiting if high temporal resolution is required. Finally, biacetyl can be difficult to work with because it has a very strong odor (akin to butterscotch) that can rapidly permeate an entire building if not contained.

Other seed species that have been used for species mole fraction measurements, primarily in supersonic mixing flows, include I_2 and NO. Both species are difficult to work with because they are highly corrosive and toxic. One of the main advantages of diatomic molecules is that they tend to have many discrete absorption lines, in contrast to more complex polyatomic molecules, such as acetone, whose lines are very broad. Diatomic molecules give the user much greater ability to choose the temperature dependence of the LIF signal. This property has been used in several supersonic mixing studies where relatively temperature-insensitive transitions were used so that the resulting LIF signal was approximately proportional to the mole fraction of the fluorescent species (69–71).

A major issue in mixing studies is that unless the smallest scales of mixing are resolved, it is not possible to differentiate between fluid that is uniformly mixed at the molecular level or simply "stirred" (i.e., intertwined, but without interdiffusion). An interesting application of NO PLIF is in "cold chemistry" techniques, which can differentiate between mixed and stirred fluid on a scale smaller than can be resolved. These techniques use the fact that NO fluorescence is rapidly quenched by O_2 but is negligibly quenched by N_2. Cold chemistry has been used to obtain quantitative statistical mixing properties of high Reynolds number shear layers where the smallest mixing scales were not resolved (72,73). This technique has also been extended to enable direct imaging of the degree of mixing/stirring for each pixel by simultaneously imaging the fluorescence from a quenched (NO) and nonquenched (acetone) species (74).

PLIF has proven extremely useful in investigating mixing and supersonic flows by seeding a tracer, and it is also important for imaging naturally present species, such as occur in chemically reacting flows. Because PLIF is a highly sensitive technique, it enables the imaging of trace species, such as combustion intermediates. For example, PLIF has been used to image an astounding number of species in flames. A limited list of major and intermediate species that have been imaged in flames by PLIF include CH, OH, NO, NO_2, C_2, CN, NH, O_2, CO, C_2H_2, H_2CO, O, and H (11,57,75). The power of PLIF species imaging in combustion research is exemplified by Fig. 15, which shows a pair of simultaneously acquired images of CH and OH in a turbulent nonpremixed methane–oxygen jet flame (76). The CH was pumped at a wavelength of about 390 nm, and the fluorescence was collected across the range of 420–440 nm; the OH was pumped at about 281 nm, and the fluorescence was collected across the range of 306–320 nm. The laser excitation was achieved by using two Nd:YAG lasers, two dye lasers, frequency doubling, and mixing crystals; the images were captured

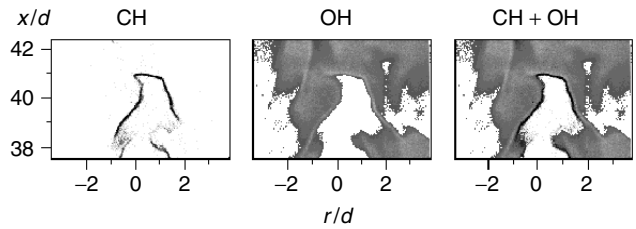

Figure 15. Sample of simultaneously acquired CH/OH PLIF images in a turbulent methane–oxygen jet flame. The CH field is shown at left, the OH field at center, and the superposition of the two at the right. The coordinates x and r refer to axial and radial distances, respectively, and d is the diameter of the jet nozzle. (Reprinted by permission of Elsevier Science from Reaction Zone Structure in Turbulent Nonpremixed Jet Flames — From CH-OH PLIF Images by J. M. Donbar, J. F. Driscoll and C. D. Carter, *Combustion and Flame*, **122**, 1–19, copyright 2000 Combustion Institute.) See color insert.

on two intensified CCD cameras. The CH field is shown at the left, the OH in the middle, and the two images are shown superimposed at the right.

Rayleigh scattering has also been used successfully to image the concentration field in a range of flows. Rayleigh scattering is defined as the elastic scattering from particles, including atoms and molecules, which are much smaller than the wavelength of the incident light (77). In molecular Rayleigh scattering, the differential Rayleigh scattering cross section at 90°, $(d\sigma_{\mathrm{Ray}}/d\Omega)$, is given by the relationship (78)

$$\frac{d\sigma_{\mathrm{Ray}}}{d\Omega} = \frac{4\pi^2(n-1)^2}{N_{\mathrm{d}}^2\lambda^4} \tag{34}$$

where N_{d} is the number density. Note that in a mixture of fluids, the Rayleigh scattering signal is proportional to the total cross section of the mixture, and thus it is not species specific (11), which greatly limits its utility for measuring concentration in reacting flows.

Equation (34) shows that the differential scattering cross section scales as λ^{-4}, which indicates a much greater scattering efficiency for short wavelengths of light. However, whether it is advantageous to work at UV rather than visible wavelengths is determined from an analysis of the entire electro-optical system. For example, is it better to measure using a frequency-quadrupled (266 nm) or a doubled (532 nm) Nd:YAG laser? To see this, it must first be considered that the signal recorded by a detector is directly proportional to the number of incident photons $(E_{\mathrm{L}}/h\nu)$, as shown in Eq. (5). Because $\nu = c/\lambda$, the number of photons per pixel for Rayleigh scattering scales as $S_{\mathrm{pp}} \propto (E_{\mathrm{L}}\lambda)\lambda^{-4} \propto E_{\mathrm{L}}\lambda^{-3}$; thus, the dependence of the signal on the wavelength is weaker on a per photon basis. Furthermore, the quantum efficiency of most detectors decreases in the UV, and there are few high-quality fast (low $f_{\#}$) photographic lenses that operate at UV wavelengths. For example, consider scattering measured by a Nd:YAG laser that produces 500 mJ at 532 nm and 125 mJ at 266 nm. In this case, the number of photons scattered will be only twice as large as that at 266 nm. After accounting for the likely reduced quantum efficiency and $f_{\#}$ of the collection optics, UV excitation may not be a means of improving the signal. UV excitation is more likely to be beneficial when using excimer lasers, which produce very high energies per pulse well into the UV. This example shows that it is necessary to account for all aspects of the electro-optical system, not just the scattering cross section, when determining the optimal wavelength to use.

One of the most common uses of Rayleigh scattering is in nonreacting mixing studies, where it is used as a passive marker of fluid concentration. For example, jet mixing can be studied by imaging the Rayleigh scattering when a jet fluid that has a high Rayleigh cross section issues into an ambient fluid that has a low cross section (32,43,79,80). In this case, the mole fraction of jet fluid can be related to the scattering signal through the relationship $\chi_{\mathrm{jet}} = [S_{\mathrm{e}} - (S_{\mathrm{e}})_{\infty}]/[(S_{\mathrm{e}})_0 - (S_{\mathrm{e}})_{\infty}]$, where $(S_{\mathrm{e}})_{0,\infty}$ are the signals obtained at the jet exit and ambient, respectively. An example of a Rayleigh scattering image is shown in

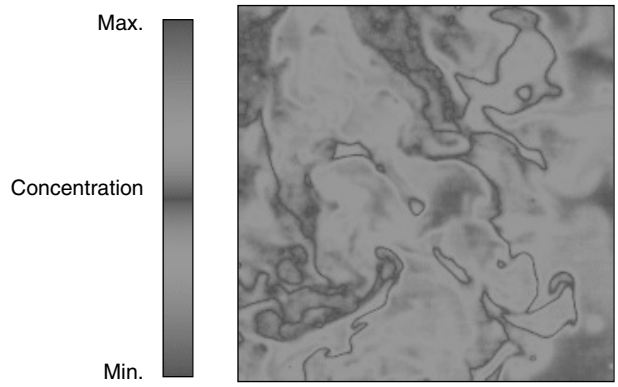

Figure 16. Example of a planar Rayleigh scattering image of a turbulent propane/acetone jet. The jet issued into a slow co-flow of filtered air, the field-of-view was 35 × 35 mm, and the local Reynolds number at the measuring station was 5600. The signal is proportional to the concentration of jet fluid. (Reprinted with permission from Planar Measurements of the Full Three-Dimensional Scalar Dissipation Rate in Gas-Phase Turbulent Flows by L. K. Su and N. T. Clemens, *Experiments in Fluids* **27**, 507–521, copyright 1999 Springer-Verlag.) See color insert.

Fig. 16 (from Ref. (32)), which was acquired in a planar turbulent jet of local Reynolds number 5600 at a distance of 100 slot widths downstream. The jet fluid was propane, which was seeded with about 5% acetone vapor, and the jet issued into a slow co-flow of air. The jet was illuminated by 240 mJ of 532-nm light produced by a Nd:YAG laser, and the images were captured by a slow-scan CCD camera that had a 58-mm focal length, $f/1.2$ lens and a laser line filter (50% maximum transmission) at a magnification of 0.28. In Fig. 16, the signal is proportional to the concentration of jet fluid, and the figure demonstrates that Rayleigh scattering can be used to obtain very high quality images of the jet concentration field. One of the main difficulties in such an experiment is the need to reduce all sources of unwanted elastic scattering, such as reflections from test section walls/windows and scattering from dust particles in the flow.

The background scattering from windows and walls is particularly problematic because it can easily overwhelm the weak Rayleigh signals. Although theoretically these background signals can be removed as part of a background correction, such as obtained by filling the test cell with helium (81), this can be done only if the shot-to-shot variation in the background is substantially less than the Rayleigh signals of interest. In many cases, however, such as in a relatively small test section, this is not the case, and background interference is unacceptably high. When the background due to laser reflections from walls/windows is high, increasing the laser energy does not improve the signal-to-background ratio because the signal and background increase proportionately. In this case, the only recourse is to increase the signal by using a higher cross section or number density or to lower the background by reducing reflections by painting opaque surfaces flat black and by using antireflection coatings on all windows.

In some cases, filtered Rayleigh scattering (FRS) can be used to reduce the background reflections greatly from walls and windows (5,82). In FRS, the Rayleigh scattering from a narrow bandwidth laser is imaged through a narrow line notch filter. For example, in one implementation used in high-speed flows, the Rayleigh scattering is induced by the narrow line light from an injection seeded frequency-doubled Nd:YAG laser, and the scattering is imaged through a molecular iodine absorption filter. If the laser beam and camera are oriented in the appropriate directions, the light scattered by the moving molecules will be Doppler-shifted, whereas the reflections from the stationary objects will not be shifted. Figure 17 illustrates the FRS technique. Because the scattering is imaged through an absorption filter, the unshifted light is absorbed by the filter, whereas the Doppler-shifted scattering is partially or completely transmitted. This same technique also forms the basis of a velocity diagnostic that will be discussed later.

Spontaneous Raman scattering has also been employed for quantitative concentration measurements in turbulent flows. It is particularly useful in combustion research because it is linear, species specific (unlike Rayleigh scattering), and enables measuring multiple species using a single laser wavelength (11). Spontaneous Raman scattering is caused by the interaction of the induced-dipole oscillations of a molecule with its rotational and vibrational motions. In other words, the incident laser beam of frequency ν_0 is shifted in frequency by the characteristic frequency of rotation/vibration. The frequency of Raman scattering is either shifted to lower frequencies (called Stokes-shifted) or to higher frequencies (called anti-Stokes-shifted). The photon that is Stokes-shifted has lower energy than the incident photon, and the energy difference is transferred to the energy of vibration/rotation of the molecule. Similarly, anti-Stokes-shifted photons have higher energy, and thus energy is given up by the molecule. In most flow imaging studies, vibrational Raman scattering is used because the lines for different species are fairly well separated. For example, for excitation at 532 nm, the Stokes-shifted vibrational Raman scattering from N_2, O_2, and H_2 occurs at wavelengths of 607, 580, and 683 nm, respectively. In contrast, owing to the smaller energies of rotation, the rotational Raman lines in a multispecies mixture tend to be grouped closely around the excitation frequency, thus making it very difficult to distinguish the scattering from a particular species.

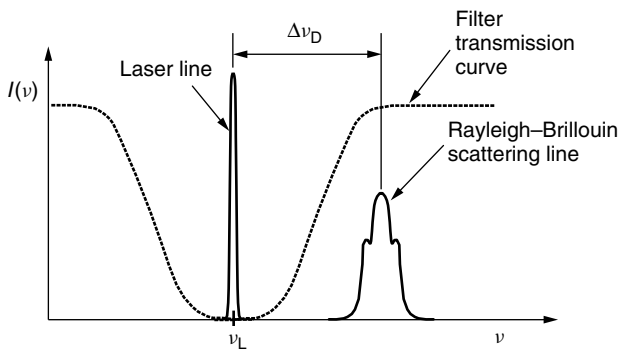

Figure 17. Illustration of the filtered Rayleigh scattering technique. The scattering from walls and windows has the same line shape and line center frequency as the laser itself. The scattering from the flow is shown as molecular (Rayleigh–Brillouin) scattering, which may be broader than the laser line, owing to thermal and acoustic motions of the molecules. If the scattering medium is particles rather than molecules, then the Rayleigh scattered light will have the same line shape as the laser. When the scattering is imaged through a notch filter (shown as the dotted line), then the Doppler-shifted light is partially or completely transmitted, whereas the scattering from stationary objects is not transmitted.

The main problem in spontaneous Raman scattering is that the signals tend to be very weak, owing to very small scattering cross sections. Typically, Raman scattering cross sections are two to three orders of magnitude smaller than Rayleigh cross sections (11). For example, for N_2 at STP, $(d\sigma/d\Omega)_{\mathrm{Ray}} = 7 \times 10^{-32}\mathrm{m}^2/\mathrm{sr}$, whereas the vibrational Raman cross section $(d\sigma/d\Omega)_{\mathrm{Ram}} = 4.6 \times 10^{-35}\mathrm{m}^2/\mathrm{sr}$, which is more than three orders of magnitude smaller than the Rayleigh cross section. The low signals that are inherent in Raman scattering make it applicable in only a few very specialized cases, such as when only major species are of interest and when very high laser energies can be generated. For example, methane concentration has been imaged in jets and flames; however, this required a high-energy flashlamp-pumped dye laser ($\lambda \approx 500$ nm, >1 J/pulse), combined with a multipass cell (19,83,84). The multipass cell resulted in an increase in laser fluence of about 30 times over that which could be achieved using only a cylindrical lens. A similar setup was used to image the Raman scattering from the C–H stretch vibrational mode in methane-air jet flames (85). Despite the use of very high laser energies and multipass cells, the relatively low SNRs reported in these studies demonstrate the great challenge in the application of Raman scattering imaging in flames.

Temperature/Pressure

Several techniques have been developed to image temperature in both liquid- and gas-phase flows. Most liquid-phase temperature imaging has been accomplished using either liquid crystals or PLIF of seeded organic dyes. For example, suspensions of small liquid crystal particles were used to image the temperature field in aqueous (86) and silicon oil flows (87,88). In these studies, the liquid crystal suspensions were illuminated by a white light sheet, and the reflected light was imaged by using a color CCD camera. The color of the crystals were then related to the local flow temperature using data from independent calibration experiments. The advantage of liquid crystals is that they can measure temperature differences as small as a fraction of a degree, but typically in a range of just a few degrees. Furthermore, they have a rather limited response time and spatial resolution that is not as good as can be achieved by planar laser imaging. PLIF thermometry offers an improvement in some of these areas, but the minimum resolvable temperature difference tends to be inferior. The simplest PLIF technique is single-line

thermometry, where a temperature-sensitive dye is uniformly seeded into the flow and the signals are related to temperature using data from a calibration experiment. For example, rhodamine B dye has relatively good temperature sensitivity because the LIF signal decreases 2–3% per °C. In (89), temperature fields were acquired by exciting rhodamine B by using a frequency-doubled Nd:YAG laser and imaging the fluorescence through a color filter. They report measurement uncertainty of about 1.7 °C.

A potential source of error in flows that have large index of refraction gradients, such as occur in variable temperature liquid- or gas-phase flows, is the variation in the intensity of the laser beam, owing to the shadowgraph effect. This can be a significant problem in liquid-phase flows where the temperature differences are of the order of several degrees or more or where fluids that have different indexes of refraction are mixed. In gas-phase flows, shadowgraph effects are less of a problem, but they may be significant when mixing gases such as propane and air that have very different indexes of refraction, and at high Reynolds numbers where gradients tend to be large. For example, careful viewing of the mixing of propane and air shown in Fig. 16 reveals subtle horizontal striations that are caused by shadowgraph effects. In principle, it is possible to correct for shadowgraph effects (64,89) — provided that the smallest gradients are resolved — by correcting the laser beam intensity along a ray path using Eq. (33). In the planar imaging of turbulent flow, however, it is not possible to correct for out-of-plane gradients, and thus the correction procedure will not be completely accurate. As an alternative to correcting for shadowgraph effects, two-line techniques have been developed where a mixture, composed of a temperature-sensitive dye and a temperature-insensitive dye, is seeded into the flow (63,90). If dyes are chosen that fluoresce at different wavelengths (when excited by the same wavelength of light), then the ratio of the two LIF signals is related to the temperature but is independent of the excitation intensity. In some cases, it is desired to remove shadowgraph effects, while maintaining density differences. In this case, it is possible to make a binary system of fluids, which have different densities but the same index of refraction (e.g., 91).

One of the simplest techniques for measuring temperature in gas-phase, constant-pressure flows is to measure density by schlieren deflectometry, interferometry, or Rayleigh scattering, from which the temperature can be inferred using an equation of state. For example, the rainbow schlieren (or deflectometry) technique discused previously (51,52) enables imaging the temperature field under certain conditions, such as in constant pressure, steady, two-dimensional, laminar flows. However, because this technique is spatially integrated, it has limited applicability to 3-D, unsteady flows, particularly where the composition and temperature (hence, index of refraction) vary in space and time. Unfiltered Rayleigh scattering techniques typically require a constant pressure flow that has a uniform index of refraction (hence, Rayleigh scattering cross section). In this case, variations in the Rayleigh scattering signal are due only to temperature variations. In general, however, mixing and reacting flows exhibit variations in fluid composition, which lead to variations in the index of refraction, even at constant temperature. It is for this reason that combustion researchers have used specialized fuel mixtures where the Rayleigh scattering cross section is approximately constant for all states of combustion, and thus the Rayleigh scattering signal is inversely proportional to temperature (92,93). The main drawback of this technique is that it assumes equal molecular diffusivities of heat and species, which is a rather dubious assumption in many cases.

FRS can be used for temperature imaging by relating changes in the scattered signal line shape to the temperature. In Rayleigh scattering from molecules, even if the incident light is essentially monochromatic, the scattered light will be spread over a range of frequencies due to thermal and acoustic motions, as illustrated in Fig. 17 (5,82). When the scattering combines thermal and acoustic broadening, it is sometimes called Rayleigh–Broullin scattering. The resulting scattered light line shape, which is sensitive to the temperature, pressure, and composition of the gas, can be used to measure those quantities. For example, if the Rayleigh scattering is imaged through a notch filter that has a known transmission curve and the theoretical Rayleigh–Brillouin line shape is known, then it is possible to infer the temperature field under certain conditions. Techniques using this procedure have enabled imaging the mean pressure and temperature fields in a Mach 2, free, air jet (94) and the instantaneous temperature field in premixed flames (95). In a related technique, Rayleigh–Brillouin scattering is imaged through a Fabry–Perot interferometer, which gives a more direct measure of the frequency and line shape of the scattered light (96). This technique has been used to measure temperature (and velocity) in high-speed, free, jet flows.

PLIF has also been extensively used for temperature imaging in gas-phase flows. The most commonly used technique is two-line PLIF of diatomic species (such as NO, I_2, and OH), where the ratio is formed from the fluorescence resulting from the excitation of two different transitions originating from different lower rotational levels (11,57). The advantage of the two-line technique is that the ratio of the signals is directly related to the rotational temperature but is independent of the local collisional environment, because the quenching affects the fluorescence from both lines similarly. The main difficulty in this technique is that if instantaneous temperature fields are desired, then two tunable laser sources and two camera systems are required. If only time-average measurements are required, then it is possible to use only a single laser/camera system. The two-line imaging technique has been used on a wide variety of flows for a range of fluorescent species. For example, measurements have been made in flows seeded by NO (97–99) and I_2 (70,100) and by naturally occurring species such as OH (101). An example of a mean temperature image of a Mach 3 turbulent bluff trailing-edge wake is shown in Fig. 18. This image was obtained by seeding 500 ppm of NO into the main air stream and then capturing the fluorescence that results from the excitation of two different absorption lines (99). The figure clearly reveals the structure of the wake flow field, including

the warm recirculation region behind the base, the cool expansion fans, and the jump in temperature across the recompression shocks.

Most PLIF thermometry has employed diatomic molecules, but the temperature dependence of acetone fluorescence has been used for single- and two-line temperature imaging in gas-phase flows (102,103). The advantage of using acetone is that it absorbs across a broad range of frequencies and thus tunable lasers are not required. In the single-line technique, which is applicable to flows that have a uniform acetone mole fraction and constant pressure, relative temperature measurements can be made up to temperatures of about 1000 K. For example, pumping by a KrF excimer laser at 248 nm can provide an estimated 1 K measurement uncertainty at 300 K. When the acetone mole fraction is not constant (such as in a mixing or reacting flow), a two-line technique can be used that is based on measuring the ratio of the LIF signals resulting from the excitation by two fixed-frequency lasers. For example, the ratio of PLIF images obtained from pumping by a XeCl excimer (308 nm) and quadrupled Nd:YAG (266 nm) can be used to achieve a factor of 5 variation in the signal ratio across the range 300–1000 K. Compared to the single-laser technique, the two-laser technique is considerably harder to implement (particularly if both images are acquired simultaneously), and it exhibits substantially lower temperature sensitivity.

All of the techniques that have been developed for measuring pressure do not measure it directly but instead infer its value from an equation of state combined with measurements of the fluid density and temperature. For this reason, pressure is very difficult to infer in low-speed flows because the pressure fluctuations result in only very small fluctuations in the density and temperature. PLIF has seen the most extensive use in pressure imaging, although one technique based on FRS (94) has been developed and was described earlier. For example, in (104), PLIF of seeded iodine and known absorption line shapes were used to infer first-order accurate pressure information for an underexpanded jet. This technique requires an isentropic flow assumption, which makes it inapplicable in many practical situations. In a related iodine PLIF technique, the pressure field was obtained by measuring its effect on the broadening of the absorption line shape (105). A limitation of this technique is that it is not very sensitive to pressure for moderate to high pressures (e.g., near 1 atm and above). In (106), NO PLIF was used to infer the 2-D pressure field in a high-enthalpy shock tunnel flow using the ratio of NO PLIF signals from a pressure-insensitive B-X transition and a pressure-sensitive A-X transition. A correction for the temperature measured in an earlier study then allowed them to infer the static pressure. This technique may be more practical than I_2 PLIF in some cases because NO occurs naturally in high-temperature air flows, but its disadvantages include the low fluorescent yield of the B-X transition and that accurate quenching rates are required. NO PLIF was also used to infer the pressure field of a bluff-body turbulent wake whose temperature field is shown in Fig. 18 (99). In this technique, trace levels of NO were seeded into a nitrogen-free stream. Because N_2 is very inefficient in quenching NO fluorescence, the LIF signal is directly proportional to the static pressure and to a nonlinear function of temperature. However, the temperature dependence can be corrected for if the temperature is measured independently, such as by the two-line method. The resulting mean pressure field obtained by this technique is shown in Fig. 19. This figure shows the low-pressure expansion fans originating from the lip of the splitter plate, the pressure increase across the recompression shock, and the nearly constant-pressure turbulent wake.

Velocity

The most widely applied velocity imaging technique in fluid mechanics is particle image velocimetry (PIV). PIV is a very robust and accurate technique, which in its

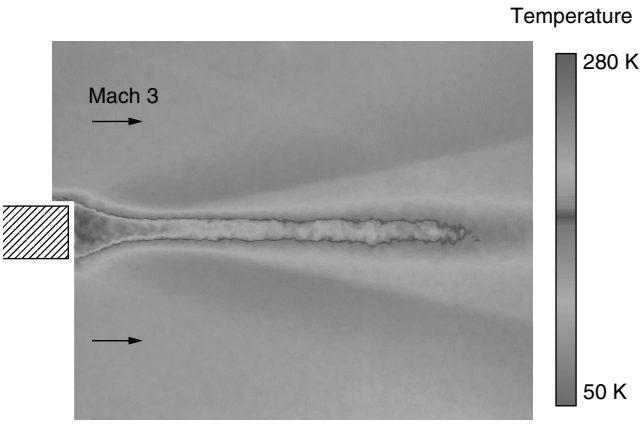

Figure 18. The mean temperature field of a supersonic bluff-body wake derived from two-line NO PLIF imaging. The field of view is 63 mm wide by 45 mm high. (Reprinted with permission from PLIF Imaging of Mean Temperature and Pressure in a Supersonic Bluff Wake by E. R. Lachney and N. T. Clemens, *Experiments in Fluids*, **24**, 354–363, copyright 1998 Springer-Verlag.) See color insert.

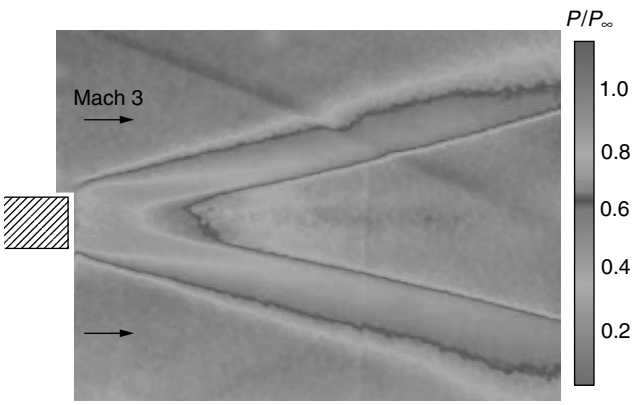

Figure 19. The mean pressure field of a supersonic bluff-body wake derived from NO PLIF imaging. The field of view is 63 mm wide by 45 mm high. (Reprinted with permission from PLIF Imaging of Mean Temperature and Pressure in a Supersonic Bluff Wake by E. R. Lachney and N. T. Clemens, *Experiments in Fluids*, **24**, 354–363, copyright 1998 Springer-Verlag.) See color insert.

most common implementation enables the imaging of two components of velocity in a cross-section of the flow. PIV measurements are now commonplace and they have been applied in a wide range of gas- and liquid-phase flows, including microfluidics, large-scale wind tunnels, flames, plasmas, and supersonic and hypersonic flows. Excellent introductions to PIV can be found in several references (5,20,23,107). At its simplest, PIV involves measuring the displacement of particles moving with the fluid for a known time. The presumption, of course, is that the particles, which are usually seeded, have sufficiently low inertia to track the changes in the motion of the flow (108). Even a cursory review of the literature shows that there are myriad variations of PIV, and therefore for brevity, only one of the most commonly used configurations will be discussed here. In a typical PIV experiment, two spatially coincident laser pulses are used where there is a known time between the pulses. The coincident beams are formed into thin sheets and passed through a flow seeded with particles. The lasers used are usually frequency-doubled Nd : YAG lasers, and the two pulses can originate from two separate lasers, from double-pulsing the Q-switch of a single laser, or from one of several dual-cavity lasers that were designed specifically for PIV applications. In two-component PIV, the scattering from the particles is imaged at 90° to the laser sheets using a high-resolution CCD camera, or less commonly today, a chemical film camera. The particle pairs can be imaged onto either a single frame (i.e., a double exposure) or onto separate frames. A major issue in PIV is that if it is not possible to tell which particle image of the pair came first, then there is an ambiguity in the direction of the velocity vector. This is one of the main advantages of the two-frame method because it does not suffer from directional ambiguity. Several CCD cameras on the market are ideal for two-frame PIV. They are based on interline transfer technology and can "frame-straddle," or allow the capture of two images a short time apart. An example of a two-frame particle field captured in a turbulent jet is shown in Fig. 20 (from Ref. 109). The camera used was a 1k × 1 k frame-straddling camera (Kodak ES1.0), the field of view was 33 × 33 mm, and the time between pulses was 8 μs.

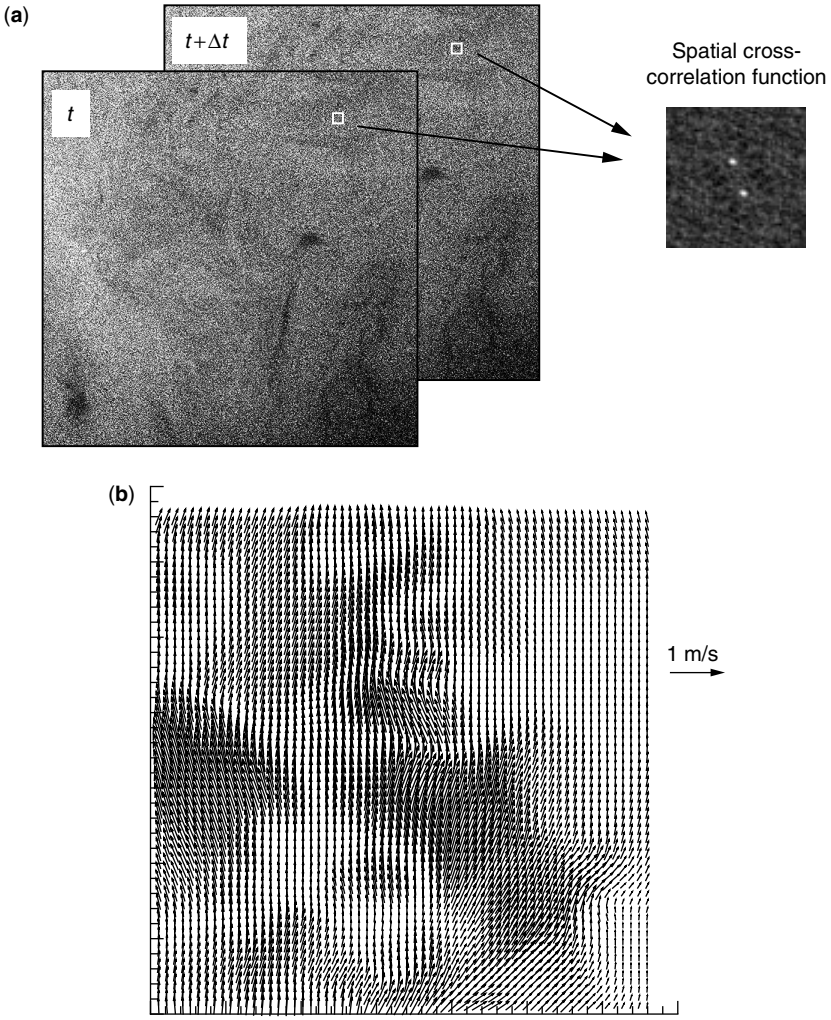

Figure 20. Sample PIV images. (**a**) Two-frame particle field images. The right image was captured 8 μs after the left image. (**b**) A two-component velocity vector field computed from a cross-correlation analysis of a two-frame particle image pair. (Reprinted with permission from Ref. 109.)

The particle displacements are obtained by dividing the image into smaller interrogation windows (usually ranging from 16×16 to 64×64 pixels), for which a single velocity vector is computed for each window. Examples of interrogation windows are shown as white boxes in Fig. 20a. The displacement is determined by computing the spatial cross-correlation function for the corresponding windows in each image of the pair, as shown in Fig. 20a. The mean displacement and direction of the velocity vector can then be determined from the location of the peaks in the cross-correlation function. This is then repeated for every interrogation window across the frame. A sample turbulent jet velocity field computed from this process is shown in Fig. 20b. For this vector field, the interrogation window size was 32×32 pixels, and the window was offset by 16 pixels at a time (50% overlap), which resulted in 62×62 vectors across the field.

Because the velocity is averaged across the interrogation window, PIV resolution and DSR are important issues. For example, typically cited values of the resolution are about 0.5 to 1 mm. Perhaps a bigger limitation though, is the DSR N_p/N_w, where N_p and N_w are the linear sizes of the array and the interrogation window, respectively. For example, a $1k \times 1k$ array that has a 32×32 window gives a DSR of only 32. If the minimum required resolution is 1 mm, then the maximum field of view that can be used is 32 mm. Limited DSR is one of the main reasons for using large format film and megapixel CCD arrays. Several algorithms have been developed that use advanced windowing techniques (110) or a combination of PIV and particle tracking (111–113) to improve both the resolution and DSR of the measurements substantially.

The PIV technique described can measure only two components of velocity; however, several techniques have been developed that enable measuring all three components. Probably the most widely used technique to date is stereoscopic PIV, which requires using two cameras, separated laterally, but share a common field of view (20,23). Particle displacement perpendicular to the laser sheet can be computed by using the particle images from the two cameras and simple geometric relationships. Although stereoscopic PIV is somewhat more difficult to implement than two-component PIV, much of the development burden can be avoided because entire systems are available from several different companies.

In another class of velocity imaging techniques, the scattered light signal is related to the Doppler shift imparted by the bulk motion of the flow. Both FRS and PLIF techniques have been applied that use this effect and may be preferable to PIV under some circumstances. For example, both FRS and PLIF velocimetry become easier to use in high-speed flows, owing to increasing Doppler shifts, whereas PIV becomes more difficult to use at high speeds because of problems in obtaining sufficient seeding density and ensuring small enough particles to track the fluid motion.

Rayleigh scattering velocimetry has seen substantial development in recent years, and different researchers have implemented closely related techniques, which go by the names of global Doppler velocimetry, filtered Rayleigh scattering, and planar Doppler velocimetry. Here, the less ambiguous term, planar Doppler velocimetry (PDV), will be used. A recent review of these techniques can be found in (114). All of these techniques operate on the basic principle that small changes in the frequency of the scattered light resulting from Doppler shifts can be inferred from the signal when the scattered light is imaged through a narrowband notch filter. Two Doppler shifts affect measurement. When molecules in the flow are illuminated by an incident laser beam, the radiation by the induced dipoles in the gas will be Doppler-shifted if there is a component of the bulk fluid velocity in the direction of the laser beam propagation. Similarly, the detector will perceive a further Doppler shift in the induced radiation if there is a component of the bulk fluid velocity in the direction of the detector. The result is that the perceived Doppler shift Δf_D measured by the detector is given by (82)

$$\Delta f_D = \frac{(\vec{s} - \vec{o}) \cdot \vec{V}}{\lambda}, \tag{35}$$

where $\vec{V}$ is the bulk fluid velocity, $\vec{o}$ is the unit vector in the laser propagation direction, and $\vec{s}$ is the vector originating from the probe volume and pointing toward the detector.

In PDV, the Rayleigh scattering is induced by a tunable narrow line width laser, and the flow is imaged through a notch filter. In the most common implementation, the laser source is an injection seeded, frequency-doubled Nd:YAG laser, which has a line width of about 50–100 MHz and can be tuned over a range of several GHz (114). The notch filter is usually an iodine vapor cell. In one technique, the laser is tuned so the non-Doppler-shifted light is centered on the edge of the absorption line, such as the right edge of the line shown in Fig. 17. Usually the scattering medium is an aerosol, such as a condensation fog, and thus the scattered line width is nearly the same as that of the laser. If the flow has a constant particle density, then the signal will increase as the Doppler shift increases. If the notch filter line shape is known, then the signal can be directly related to the velocity. In most cases, the density is not constant, and therefore a separate non-Doppler-shifted density measurement must be made. This can be accomplished by using another camera or a single-camera split-image configuration (114). Much of the recent work in this area has been in improving the accuracy of the technique and in extending it to enable measuring three components of velocity.

PLIF velocimetry is also a Doppler-shift-based technique, which is particularly applicable in high-speed reacting flows where seeding the flow with particles is not practical or where low gas densities preclude the use of Rayleigh scattering. In most PLIF velocimetry studies, the flow is seeded by a tracer, such as iodine or NO, although naturally occurring species, such as OH, have also been used successfully. PLIF velocimetry is accomplished by having the laser sheet propagate as much as possible in the direction of the bulk flow, which maximizes the Doppler shift seen by the absorbing molecules. The camera is usually oriented normally to the laser sheet, and the broadband fluorescence is collected (i.e., it is not spectrally resolved). Thus, unlike PDV, only the Doppler

shift induced by the flow/laser beam is relevant. The difficulty in PLIF is that in addition to velocity, the fluid composition, pressure, and temperature also affect the signal through number density, population, quenching, and line shape effects. Therefore, schemes have to be devised that enable isolating effects due to velocity alone. In an early PLIF velocimetry technique, a tunable CW narrow line laser (argon-ion) was used to scan in frequency across an absorption line of I_2 seeded in a high-speed flow, and several images were captured during the scan, which enabled reconstructing the line shape for each pixel (115). The measured Doppler-shifted line shapes were then compared to an unshifted line shape taken in a stationary reference cell. Although this technique worked well, it can provide only time-average measurements because it takes finite time to scan the laser. In another technique also employing I_2 PLIF, two discrete laser frequencies and four laser sheets were used to measure two components of mean velocity and pressure in an underexpanded jet (104).

In the techniques mentioned before, the laser line needs to be much narrower than the absorption line. It can be an advantage, however, when the laser line width is much larger than the absorption line width because it reduces the sensitivity of the signal to variations in the absorption line shape. For example, in (116), two counterpropagating laser sheets and two cameras were used to image one component of velocity in NO-seeded supersonic flows. The reason for using counterpropagating sheets is that the ratio of the LIF signals from the two sheets can be related to the velocity component but is independent of the local temperature and pressure. When the laser line is of the same order of magnitude as the absorption line, such two-sheet fixed-frequency techniques require modeling the overlap integral for the absorption and laser line shapes (117).

Future Developments

Although new quantitative imaging techniques will certainly continue to be developed, it is likely that the greatest effort in the future will be directed at simply improving existing techniques by making them easier and cheaper to implement and by improving the accuracy, precision, resolution, and framing rate. A good example of the improvement that can be achieved by better technology is to compare the quality of OH PLIF imaging from one of the first images captured by this technique in 1984 (118) to images that have been captured more recently (76). The difference in quality is dramatic, despite the use of the same technique in both cases.

A major trend that started in the past decade, but will no doubt continue, is the application of two or more "established" techniques to obtain simultaneous images of several flow parameters (81). Multiple-parameter techniques include the simultaneous acquisition of multiple flow variables, such as velocity and scalars. Multiple-parameter imaging also includes imaging the same flow variable with a short time delay between images, to obtain the rate of change of a property, and acquiring two images of the same flow variable where the laser sheets are placed a small distance apart to enable the computation of spatial gradients. Because multiple-parameter imaging usually involves established techniques, its implementation is usually limited by the availability of the required equipment and by optical access for all of the laser beams and cameras.

An obvious limitation of most of the techniques that have been discussed is that the framing rates are typically limited to a few hertz. This limitation is imposed by the laser and camera systems that are available now. Although there is no question that the laser power of high-repetition rate commercial lasers will continue to increase with time, limited laser power will remain an obstacle to kilohertz imaging for many of the techniques discussed in this article. For example, the Rayleigh scattering image of Fig. 16 required about 300 mJ of light from a frequency-doubled Nd:YAG operating at 10 Hz, which corresponds to 3 W of average power. If it was desired to acquire images that have the same SNR at 10 kHz, such as is likely to be required in even a moderate Reynolds number gas-phase flow, then this would require a doubled Nd:YAG laser whose average power is 3 kW. This amount of continuous average power might not be large compared to that required for metal cutting or ballistic missile defense, but it is an astounding amount of power by flow diagnostics standards, and handling such a laser would provide many practical problems for the user. High-framing rate imaging is also currently limited by camera technology; no camera is currently available that operates quasi-continuously at 10 kHz at 10–20 e^- rms noise per pixel as is necessary to obtain images of the quality of Fig. 16. The reason for this is that high framing rates require high readout bandwidths, which in turn lead to more noise. Thus to keep the noise low, either the framing rate or the number of pixels must be degraded. Despite this caveat, progress toward higher framing rate imaging for all of the techniques discussed here will continue as the necessary technologies improve.

Another major trend that will continue is the further development and refinement of three-dimensional techniques. The most commonly used three-dimensional techniques are classified as either tomography or reconstructions from stacks of planar images. Tomography is the reconstruction of a 3-D field of a fluid property from line-of-sight integrated data measured from several different directions through the flow. For example, both absorption (11) and interferometry (119,120) have been used, which enable reconstructing the 3-D concentration and index-of-refraction fields, respectively. A more popular technique is to reconstruct the 3-D field using a set of images that have been acquired by rapidly scanning a laser sheet through the flow and capturing several planar images during the sweep (5). This technique has been used effectively in many aqueous flows using PLIF excited by either continuous or pulsed lasers (6,59,63,121). However, because these techniques rely on sweeping a laser beam or sheet through the flow on a timescale that is shorter than the characteristic fluid timescales, such techniques are significantly more challenging in gas-phase flows. It is remarkable, however, that such experiments have been accomplished by sweeping a flashlamp-pumped dye laser sheet through the flow in only a few microseconds. In one case, the Rayleigh scattered light from a

freon-gas jet was imaged (122) and in another case the Mie scattering from a particle-laden supersonic mixing layer was imaged (123). Both studies used a high-speed framing camera that could acquire only a few frames (e.g., 10–20), and thus the resolution of the reconstructions was obviously quite limited.

The future of 3-D flow imaging is probably best exemplified by holographic PIV (HPIV), which provides accurate three-component velocity fields throughout a volume of fluid (124–127). HPIV is an intrinsically 3-D technique, which begins with recording a hologram of the 3-D double-exposure particle field onto high-resolution film. The image is then reconstructed, and the particle field is digitized by sequentially imaging planes of the reconstruction using a digital camera. HPIV enables the acquisition of an astounding amount of data, but because it is a challenging technique to implement and it requires using very high resolution large-format chemical film, the framing rates will remain low for at least the near future.

In conclusion, flow imaging is driving a revolution in fluid mechanics research that will continue well into the future. Continued advances in laser and digital camera technologies will make most of the imaging techniques described in this article possible one day at sufficient spatial resolution and framing rates to resolve virtually any flow spatial and temporal scale of interest. This is an exciting proposition as we enter a new century of experimental fluid dynamics research.

Acknowledgments
The author acknowledges the generous support of his research into flow imaging by the National Science Foundation, particularly under grants CTS-9319136 and CTS-9553124. In addition, the author thanks Michael Tsurikov and Yongxi Hou of UT-Austin for help in preparing this article.

ABBREVIATIONS AND ACRONYMS

2-D	two-dimensional
3-D	three-dimensional
CCD	charge-coupled device
CTF	contrast transfer function
CW	continuous wave
DSR	dynamic spatial range
FRS	filtered Rayleigh scattering
FT	Fourier transform
ICCD	intensified charge-coupled device
IR	infrared
LIF	laser-induced fluorescence
LSF	line spread function
MTF	modulation transfer function
NA	numerical aperture
Nd:YAG	neodymium: yttrium-aluminum garnet
OPO	optical parametric oscillator
OTF	optical transfer function
PDV	planar Doppler velocimetry
PIV	particle image velocimetry
PLIF	planar laser-induced fluorescence
PSF	point spread function
PTF	phase transfer function
SBR	signal to background ratio
SNR	signal to noise ratio
SRF	step response function
STP	standard temperature and pressure
TEM	transverse electromagnetic modes
TV	television
UV	ultraviolet

BIBLIOGRAPHY

1. P. H. Paul, M. G. Garguilo, and D. J. Rakestraw, *Anal. Chem.* **70**, 2459–2467 (1998).
2. J. G. Santiago et al., *Exp. Fluids* **25**, 316–319 (1998).
3. L. M. Weinstein, *High-Speed Research: 1995 Sonic Boom Workshop, Atmospheric Propagation and Acceptability Studies, NASA CP-3335, October, 1995.*
4. M. Van Dyke, *An Album of Fluid Motion*, The Parabolic Press, Stanford, 1982.
5. A. J. Smits and T. T. Lim, eds., *Flow Visualization: Techniques and Examples*, Imperial College Press, London, 2000.
6. W. J. A. Dahm and K. B. Southerland, in A. J. Smits and T. T. Lim, eds., *Flow Visualization: Techniques and Examples*, Imperial College Press, London, 2000, pp. 289–316.
7. L. K. Su and W. J. A. Dahm, *Phys. Fluids* **8**, 1,883–1,906 (1996).
8. M. Gharib, *J. Fluids Eng.* **118**, 233–242 (1996).
9. W. Merzkirch, *Flow Visualization*, 2nd ed., Academic Press, Orlando, 1987.
10. G. S. Settles, *AIAA J.* **24**, 1,313–1,323 (1986).
11. A. C. Eckbreth, *Laser Diagnostics for Combustion Temperature and Species*, Abacus Press, Cambridge, 1988.
12. B. J. Kirby and R. K. Hanson, *Appl. Phys. B* **69**, 505–507 (1999).
13. J. Hecht, *The Laser Guidebook*, 2nd ed., McGraw-Hill, NY, 1992.
14. P. Wu and R. B. Miles, *Opt. Lett.* **25**, 1,639–1,641 (2000).
15. J. M. Grace et al., *Proc. SPIE* **3642**, 133–141 (1999).
16. A. E. Siegman, *Lasers*, University Science Books, Mill Valley, CA, 1986.
17. M. W. Sasnett, in D. R. Hall and P. E. Jackson, ed., *The Physics and Technology of Laser Resonators*, Adam Hilger, Bristol, 1989, pp. 132–142.
18. W. J. Smith, *Modern Optical Engineering: The Design of Optical Systems*, 2nd ed., McGraw-Hill, NY, 1990.
19. M. B. Long, D. C. Fourguette, M. C. Escoda, and C. B. Layne, *Opt. Lett.* **8**, 244–246 (1983).
20. R. J. Adrian, *Ann. Rev. Fluid Mech.* **23**, 261–304 (1991).
21. A. Vogel and W. Lauterborn, *Opt. Lasers Eng.* **9**, 274–294 (1988).
22. J. C. Lin and D. Rockwell, *Exp. Fluids* **17**, 110–118 (1994).
23. M. Raffel, C. E. Willert, and J. Kompenhans, *Particle Image Velocimetry: A Practical Guide*, Springer, Berlin, 1998.
24. B. Lecordier et al., *Exp. Fluids* **17**, 205–208 (1994).
25. N. T. Clemens, S. P. Petullo, and D. S. Dolling, *AIAA J.* **34**, 2,062–2,070 (1996).
26. P. H. Paul, *AIAA Paper 91–2315*, June, 1991.
27. J. M. Seitzman and R. K. Hanson, in A. Taylor, ed., *Instrumentation for Flows with Combustion*, Academic Press, London, 1993.
28. RCA Staff, *Electro-Optics Handbook*, RCA, Lancaster, PA, 1974.
29. P. H. Paul, I. van Cruyningen, R. K. Hanson, and G. Kychakoff, *Exp. Fluids* **9**, 241–251 (1990).

30. J. R. Janesick et al., *Opt. Eng.* **26**, 692–714 (1987).
31. I. S. McLean, *Electronic Imaging in Astronomy: Detectors and Instrumentation*, John Wiley & Sons, NY, 1997.
32. L. K. Su and N. T. Clemens, *Exp. Fluids* **27**, 507–521 (1999).
33. E. Hecht, *Optics*, 3rd ed., Addison-Wesley, Reading, MA, 1998.
34. W. K. Pratt, *Digital Image Processing*, 2nd ed., Wiley, NY, 1991.
35. T. L. Williams, *The Optical Transfer Function of Imaging Systems*, Institute of Physics Publishing, Bristol, 1999.
36. C. S. Williams and O. A. Becklund, *Introduction to the Optical Transfer Function*, John Wiley & Sons, NY, 1989.
37. W. Wittenstein, J. C. Fontanella, A. R. Newbery, and J. Baars, *Optica Acta* **29**, 41–50 (1982).
38. R. N. Bracewell, *The Fourier Transform and its Applications*, 2nd ed., McGraw-Hill, NY, 1986.
39. F. Chazallet and J. Glasser, *SPIE Proc.* **549**, 131–144 (1985).
40. H. Tennekes and J. L. Lumley, *A First Course in Turbulence*, MIT Press, Cambridge, 1972.
41. G. K. Batchelor, *J. Fluid Mech.* **5**, 113–133 (1959).
42. K. A. Buch Jr. and W. J. A. Dahm, *J. Fluid Mech.* **317**, 21–71 (1996).
43. K. A. Buch Jr. and W. J. A. Dahm, *J. Fluid Mech.* **364**, 1–29 (1998).
44. C. A. Friehe, C. W. van Atta, and C. H. Gibson, in *AGARD Turbulent Shear Flows CP-93*, North Atlantic Treaty Organization, Paris, 1971, pp. 18-1–18-7.
45. J. C. Wyngaard, *J. Phys. E: J. Sci. Instru.* **1**, 1,105–1,108 (1968).
46. J. C. Wyngaard, *J. Phys. E: J. Sci. Instru.* **2**, 983–987 (1969).
47. R. A. Antonia and J. Mi, *J. Fluid Mech.* **250**, 531–551 (1993).
48. C. J. Chen and W. Rodi, *Vertical Turbulent Buoyant Jets: A Review of Experimental Data*, Pergamon, Oxford, 1980.
49. R. J. Goldstein and T. H. Kuen, in R. J. Goldstein, ed., *Fluid Mechanics Measurements*, Taylor and Francis, London, 1996.
50. N. T. Clemens and P. H. Paul, *Phys. Fluids A* **5**, S7 (1993).
51. P. S. Greenberg, R. B. Klimek, and D. R. Buchele, *Appl. Opt.* **34**, 3,810–3,822 (1995).
52. A. K. Agrawal, N. K. Butuk, S. R. Gollahalli, and D. Griffin, *Appl. Opt.* **37**, 479–485 (1998).
53. J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed., McGraw-Hill, Boston, 1996.
54. G. F. Albrecht, H. F. Robey, and T. R. Moore, *Appl. Phys. Lett.* **57**, 864–866 (1990).
55. D. Papamoschou and H. F. Robey, *Exp. Fluids* **17**, 10–15 (1994).
56. L. M. Weinstein, *AIAA J.* **31**, 1,250–1,255 (1993).
57. R. K. Hanson, J. M. Seitzman, and P. H. Paul, *Appl. Phys. B* **50**, 441–454 (1990).
58. M. M. Koochesfahani and P. E. Dimotakis, *J. Fluid Mech.* **170**, 83–112 (1986).
59. R. R. Prasad and K. R. Sreenivasan, *J. Fluid Mech.* **216**, 1–34 (1990).
60. C. Arcoumanis, J. J. McGuirk, and J. M. L. M. Palma, *Exp. Fluids* **10**, 177–180 (1990).
61. J. R. Saylor, *Exp. Fluids* **18**, 445–447 (1995).
62. P. S. Karasso and M. G. Mungal, *Exp. Fluids* **23**, 382–387 (1997).
63. J. Sakakibara and R. J. Adrian, *Exp. Fluids* **26**, 7–15 (1999).
64. J. D. Nash, G. H. Jirka, and D. Chen, *Exp. Fluids* **19**, 297–304 (1995).
65. A. Lozano, B. Yip, and R. K. Hanson, *Exp. Fluids* **13**, 369–376 (1992).
66. S. H. Smith and M. G. Mungal, *J. Fluid Mech.* **357**, 83–122 (1998).
67. D. Papamoschou and A. Bunyajitradulya, *Phys. Fluids* **9**, 756–765 (1997).
68. D. Wolff, H. Schluter, and V. Beushausen, *Berichte der Bunsen-Gesellschaft fur physikalisc* **97**, 1,738–1,741 (1993).
69. R. J. Hartfield Jr., J. D. Abbitt III, and J. C. McDaniel, *Opt. Lett.* **14**, 850–852 (1989).
70. J. M. Donohue and J. C. McDaniel Jr., *AIAA J.* **34**, 455–462 (1996).
71. N. T. Clemens and M. G. Mungal, *J. Fluid Mech.* **284**, 171–216 (1995).
72. N. T. Clemens and P. H. Paul, *Phys. Fluids* **7**, 1,071–1,081 (1995).
73. T. C. Island, W. D. Urban, and M. G. Mungal, *Phys. Fluids* **10**, 1,008–1,021 (1998).
74. G. F. King, J. C. Dutton, and R. P. Lucht, *Phys. Fluids* **11**, 403–416 (1999).
75. T. Parr and D. Hanson-Parr, L. DeLuca, E. W. Price, and M. Summerfield, eds., *Nonsteady Burning and Combustion Stability of Solid Propellants*, Progress in Aeronautics and Astronautics, American Institute of Aeronautics and Astronautics, vol. 143, Washington, DC, 1992, pp. 261–323.
76. J. M. Donbar, J. F. Driscoll, and C. D. Carter, *Combustion and Flame* **122**, 1–19 (2000).
77. C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*, John Wiley & Sons, NY, 1983.
78. E. J. McCartney, *Optics of the Atmosphere: Scattering by Molecules and Particles*, Wiley, NY, 1976.
79. B. Yip, R. L. Schmitt, and M. B. Long, *Opt. Lett.* **13**, 96–98 (1988).
80. D. A. Feikema, D. Everest, and J. F. Driscoll, *AIAA J.* **34**, 2,531–2,538 (1996).
81. M. B. Long, in A. M. K. P. Taylor, ed., *Instrumentation for Flows with Combustion*, Academic, London, 1993, pp. 467–508.
82. R. B. Miles and W. R. Lempert, *Ann. Rev. Fluid Mech.* **29**, 285–326 (1997).
83. M. Namazian, J. T. Kelly, and R. W. Schefer, *Twenty-Second Symposium (Int.) on Combustion*, The Combustion Institute, Pittsburgh, 1988, pp. 627–634.
84. M. Namazian et al., *Exp. Fluids* **8**, 216–228 (1989).
85. J. B. Kelman, A. R. Masri, S. H. Starner, and R. W. Bilger, *Twenty-Fifth Symposium (Int.) on Combustion, The Combustion Institute*, Pittsburgh, 1994, pp. 1,141–1,147.
86. D. Dabiri and M. Gharib, *Exp. Fluids* **11**, 77–86 (1991).
87. I. Kimura et al., in B. Khalighi, M. J. Braun, and C. J. Freitas, eds., *Flow Visualization*, vol. 85, ASME FED, 1989, pp. 69–76.
88. M. Ozawa, U. Müller, I. Kimura, and T. Takamori, *Exp. Fluids* **12**, 213–222 (1992).
89. M. C. J. Coolen, R. N. Kieft, C. C. M. Rindt, and A. A. van Steenhoven, *Exp. Fluids* **27**, 420–426 (1999).
90. J. Coppeta and C. Rogers, *Exp. Fluids* **25**, 1–15 (1998).
91. A. Alahyari and E. K. Longmire, *Exp. Fluids* **17**, 434–440 (1994).
92. D. C. Fourguette, R. M. Zurn, and M. B. Long, *Combustion Sci. Technol.* **44**, 307–317 (1986).

93. D. A. Everest, J. F. Driscoll, W. J. A. Dahm, and D. A. Feikema, *Combustion and Flame* **101**, 58–68 (1995).
94. J. N. Forkey, W. R. Lempert, and R. B. Miles, *Exp. Fluids* **24**, 151–162 (1998).
95. G. S. Elliott, N. Glumac, C. D. Carter, and A. S. Nejad, *Combustion Sci. Technol.* **125**, 351–369 (1997).
96. R. G. Seasholtz, A. E. Buggele, and M. F. Reeder, *Opt. Lasers Eng.* **27**, 543–570 (1997).
97. B. K. McMillan, J. L. Palmer, and R. K. Hanson, *Appl. Opt.* **32**, 7,532–7,545 (1993).
98. J. L. Palmer, B. K. McMillin, and R. K. Hanson, *Appl. Phys. B* **63**, 167–178 (1996).
99. E. R. Lachney and N. T. Clemens, *Exp. Fluids* **24**, 354–363 (1998).
100. T. Ni-Imi, T. Fujimoto, and N. Shimizu, *Opt. Lett.* **15**, 918–920 (1990).
101. J. M. Seitzman and R. K. Hanson, *Appl. Phys. B* **57**, 385–391 (1993).
102. M. C. Thurber, F. Grisch, and R. K. Hanson, *Opt. Lett.* **22**, 251–253 (1997).
103. M. C. Thurber et al., *Appl. Opt.* **37**, 4,963–4,978 (1998).
104. B. Hiller and R. K. Hanson, *Appl. Opt.* **27**, 33–48 (1988).
105. R. J. Hartfield Jr., S. D. Hollo, and J. C. McDaniel, *AIAA J.* **31**, 483–490 (1993).
106. P. Cassady and S. Lieberg, *AIAA Paper No. 92-2962*, 1992.
107. P. Buchhave, in L. Lading, G. Wigley, and P. Buchhave, eds., *Optical Diagnostics for Flow Processes*, Plenum, NY, 1994, pp. 247–269.
108. A. Melling, *Meas. Sci. Technol.* **8**, 1,406–1,416 (1997).
109. J. E. Rehm, *Ph.D. Dissertation, The University of Texas at Austin*, 1999.
110. J. Westerweel, D. Dabiri, and M. Gharib, *Exp. Fluids* **23**, 20–28 (1997).
111. R. D. Keane, R. J. Adrian, and Y. Zhang, *Meas. Sci. Technol.* **6**, 754–768 (1995).
112. E. A. Cowen and S. G. Monismith, *Exp. Fluids* **22**, 199–211 (1997).
113. J. E. Rehm and N. T. Clemens, *Exp. Fluids* **26**, 497–504 (1999).
114. M. Samimy and M. P. Wernet, *AIAA J.* **38**, 553–574 (2000).
115. J. C. McDaniel, B. Hiller, and R. K. Hanson, *Opt. Lett.* **8**, 51–53 (1983).
116. P. H. Paul, M. P. Lee, and R. K. Hanson, *Opt. Lett.* **14**, 417–419 (1989).
117. M. Allen et al., *AIAA J.* **32**, 1,676–1,682 (1994).
118. G. Kychakoff, R. D. Howe, and R. K. Hanson, *Appl. Opt.* **23**, 704–712 (1984).
119. L. Hesselink, in W. -J. Yang, ed., *Handbook of Flow Visualization*, Hemisphere, NY, 1987.
120. D. W. Watt and C. M. Vest, *Exp. Fluids* **8**, 301–311 (1990).
121. M. Yoda, L. Hesselink, and M. G. Mungal, *J. Fluid Mech.* **279**, 313–350 (1994).
122. B. Yip, R. L. Schmitt, and M. B. Long, *Opt. Lett.* **13**, 96–98 (1988).
123. T. C. Island, B. J. Patrie, and R. K. Hanson, *Exp. Fluids* **20**, 249–256 (1996).
124. H. Meng and F. Hussain, *Fluid Dynamics Res.* **8**, 33–52 (1991).
125. D. H. Barnhart, R. J. Adrian, and G. C. Papen, *Appl. Opt.* **33**, 7,159–7,170 (1994).
126. J. O. Scherer and L. P. Bernal, *Appl. Opt.* **36**, 9,309–9,318 (1997).
127. J. Zhang, B. Tao, and J. Katz, *Exp. Fluids* **23**, 373–381 (1997).

FORCE IMAGING

KIM DE ROY
RSscan INTERNATIONAL
Belgium

L. PEERAER
University of Leuven
Belgium, FLOK
University Hospitals Leuven
CERM

INTRODUCTION

The study of human locomotion has generated a substantial number of publications. Starting from evolutionary history, the authors try to give some insight into the transition from quadripedal to bipedal locomotion (1,2). Evolving to a semiaquatic mode of life, it may be assumed that the human foot developed from its original versatile function for swimming (3), support, and gripping to a more specialized instrument that can keep the body in an upright position. This change of function enables all of the movements that are specific to humans such as walking and running. Data deduced from the literature show that the effects of walking speed on stride length and frequency are similar in bonobos, common chimpanzees, and humans. This suggests that within extant Hominidae, spatiotemporal gait characteristics are highly comparable (4) (Fig. 1). Despite these similarities, the upright position and erect walking is accepted as one of the main characteristics that differentiate humans from animals. No wonder that throughout the search for human evolution the imprints of the feet of the first humanoid creatures that were found are studied and discussed as much as their skulls. Whatever the causes for this evolution to the erect position, the fact remains that static and dynamic equilibrium must be achieved during bipedal activities, thus dramatically reducing the supporting area with respect to the quadripedal condition where this total area is formed by more than two feet.

The available area of support during bipedal activities is thus restricted to that determined by one or both feet. The anatomical structure of the human foot, as well as the neuromuscular and circulatory control must have evolved to a multijoint dynamic mechanism that determines the complex interaction between the lower limb and the ground during locomotion (5). Consequently, besides gravitation, the external forces acting on the human body act on the plantar surface of the foot and generate movement according to Newton's laws. Thus, studying the latter, called the ground reaction force (GRF), is essential to our understanding of human normal and pathological locomotion. The GRF may, however, vary in point of application, magnitude, and orientation, necessitating the

Color Figures

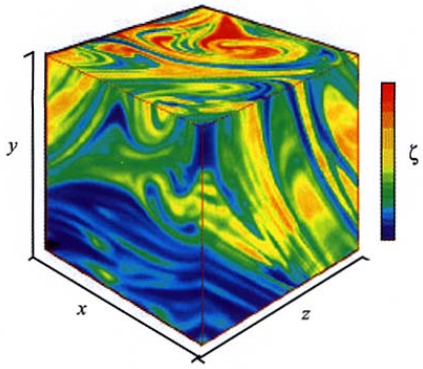

Flow Imaging, Figure 1. Three-dimensional rendering of the conserved scalar (ζ) field measured in a turbulent water jet using laser-induced fluorescence of a fluorescent dye seeded into the jet fluid. The cube is approximately 27 mm on each side, and the data resolve the finest scalar and vorticity scales in the flow. (Reprinted with permission from Quantitative Flow Visualization via Fully-Resolved Four-Dimensional Spatio-Temporal Imaging by W. J. A. Dahm and K. B. Southerland, in *Flow Visualization: Techniques and Examples,* A. J. Smits and T. T. Lim, eds., Imperial College Press, London, 2000.)

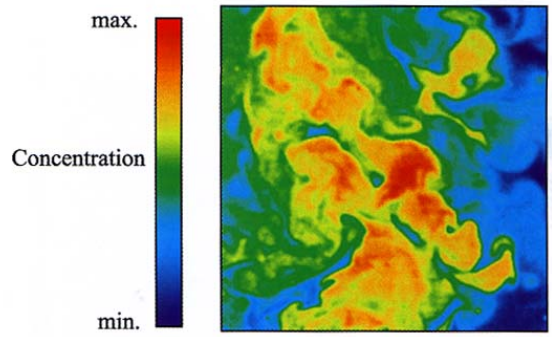

Flow Imaging, Figure 16. Example of a planar Rayleigh scattering image of a turbulent propane/acetone jet. The jet issued into a slow co-flow of filtered air, the field-of-view was 35 × 35 mm, and the local Reynolds number at the measuring station was 5600. The signal is proportional to the concentration of jet fluid. (Reprinted with permission from Planar Measurements of the Full three-Dimensional Scalar Dissipation Rate in Gas-Phase Turbulent Flows by L. K. Su and N. T. Clemens, *Experiments in Fluids,* **27**, 507–521, copyright 1999 Springer-Verlag.)

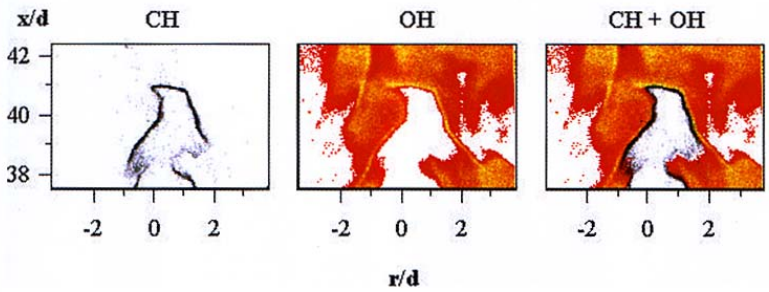

Flow Imaging, Figure 15. Sample of simultaneously acquired CH/OH PLIF images in a turbulent methane–oxygen jet flame. The CH field is shown at left, the OH field at center, and the superposition of the two at the right. The coordinates x and r refer to axial and radial distances, respectively, and d is the diameter of the jet nozzle. (Reprinted with permission of Elsevier Science from Reaction Zone Structure in Turbulent Nonpremixed Jet Flames — From CH-OH PLIF Images by J. M. Donbar, J. F. Driscoll and C. D. Carter, *Combustion and Flame,* **122**, 1–19, copyright 2000 Combustion Institute.)

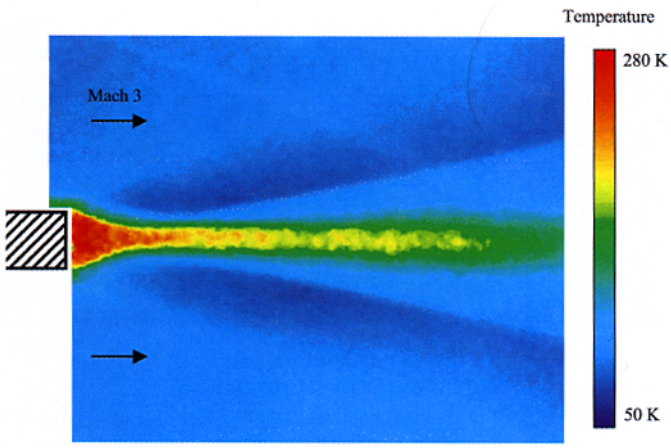

Flow Imaging, Figure 18. The mean temperature field of a supersonic bluff-body wake derived from two-line NO PLIF imaging. The field of view is 63 mm wide by 45 mm high. (Reprinted with permission from PLIF Imaging of Mean Temperature and Pressure in a Supersonic Bluff Wake by E. R. Lachney and N. T. Clemens, *Experiments in Fluids,* **24**, 354–363, copyright 1998 Springer-Verlag.)

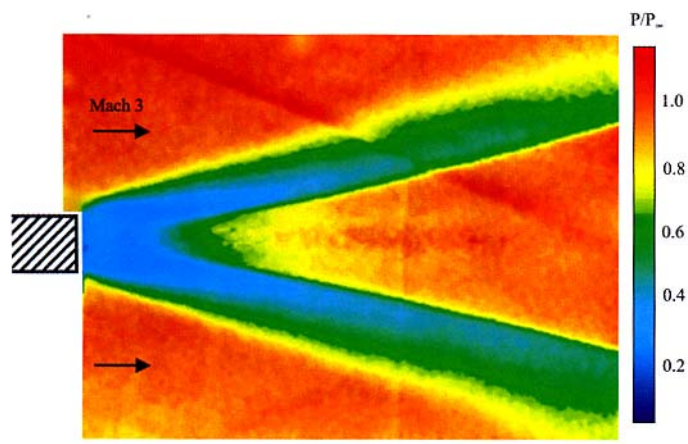

Flow Imaging, Figure 19. The mean pressure field of a supersonic bluff-body wake derived from NO PLIF imaging. The field of view is 63 mm wide by 45 mm high. (Reprinted with permission from PLIF Imaging of Mean Temperature and Pressure in a Supersonic Bluff Wake by E. R. Lachney and N. T. Clemens, *Experiments in Fluids,* **24**, 354–363, copyright 1998 Springer-Verlag.)