

Optimizing Precision and Accuracy of Quantitative PLIF of Acetone as a Tracer for Hydrogen Fuel

Victor M. Salazar and Sebastian A. Kaiser
Sandia National Laboratories

Fabien Halter
Université d'Orléans, France

This is the work of the U.S. Government.

ABSTRACT

Quantitative planar laser-induced fluorescence (PLIF) of gaseous acetone as a fuel-tracer has been used in an optically accessible engine, fueled by direct hydrogen injection. The purpose of this article is to assess the accuracy and precision of the measurement and the associated data reduction procedures. A detailed description of the acetone seeding system is given as well. The key features of the experiment are a high-pressure bubbler saturating the hydrogen fuel with acetone vapor, direct injection into an optical engine, excitation of acetone fluorescence with an Nd:YAG laser at 266 nm, and detection of the resulting fluorescence by an unintensified camera. Key steps in the quantification of the single-shot imaging data are an in-situ calibration and a correction for the effect of local temperature on the fluorescence measurement. We assess the accuracy of the measurement in terms of drift in acetone-vapor concentration, linearity of fluorescence with laser energy, absorption of the beam within the probe volume, spatial inhomogeneity in the calibration measurements, and uncertainties in the temperature correction. The precision is impacted by camera read-out noise, shot (quantum) noise, shot-to-shot variations in total laser-pulse energy and the transverse energy profile, and "beam steering" by thermal gradients. Procedures to quantify and if possible minimize all of the above factors are described. Among the factors investigated, the single greatest impact on accuracy and precision has uncertainty in the calibration of global equivalence ratio and refractive beam steering, respectively.

INTRODUCTION

For vehicle use, advanced hydrogen-fueled internal combustion engines (H2ICEs) can be highly efficient power plants meeting present and future emissions regulations at reasonable cost [1, 2]. A key element of such advanced H2ICEs is direct injection (DI) of the hydrogen fuel into the cylinder [3, 4]. However, as with hydrocarbon-fueled engines, the design of DI engines requires knowledge about injection event, in-cylinder mixing, and resulting pre-combustion mixture distribution. Much can be learned from careful parametric studies on single-cylinder research engines [5]. However, ideally these metal-engine studies are complemented by laser-based measurements in optically accessible engines and numerical simulations to gain physical insight into these processes. As in optical gasoline and diesel engines, planar laser-induced fluorescence (PLIF) of a fuel-tracer can be used in DI-H2ICEs to make two-dimensional measurements of the instantaneous fuel distribution.

Using optical diagnostics, the Hydrogen Engine Lab at Sandia's Combustion Research Facility aims to provide some of the physical understanding of in-cylinder processes that will be necessary to design future DI-H2ICEs. Previous studies on this engine included OH* chemiluminescence imaging [6], and measurements of velocity and fuel distribution [7, 8]. The PLIF imaging technique used to measure the in-cylinder fuel distribution has undergone constant improvement. The purpose of this paper is to document the implementation

of a new high-pressure seeding system and the accuracy and precision of the optical equivalence-ratio measurement. The measurement precision will be assessed assuming propagation of the various errors via the classic quadrature sum. The accuracy, *i.e.*, the extent of the various systematic biases, will be characterized for the worst- and a typical-case scenario.

Although the seeding system is specialized for the high-pressure gaseous fuel, the results presented here can also serve as a benchmark for comparisons to fuel imaging in liquid-fueled engines. Indeed, Herold and Ghandhi [9] performed an in-depth analysis of the precision of fuel-tracer PLIF in very weakly stratified mixtures in a gasoline HCCI engine. Additional motivation for a thorough characterization of the measurements comes from the fact that the experimental results will be used to validate computational modeling of the engine.

ENGINE AND ENGINE OPERATION

The engine under consideration is a passenger-car sized single-cylinder research engine with a flat-topped piston, pent-roof, and four valves. Optical access is of the Bowditch type, *i.e.*, through parts of the liner and pentroof as well as through the piston. Hydrogen fuel is delivered directly into the cylinder through an injector developed by Westport Inc. with injection pressures up to 100 bar. If not noted otherwise, the injector used for the current study was an electro-magnetic injector with 6 orifices at an included cone angle of 90°. The orifice diameter is 0.54 mm. The injector is located towards the side between the intake valves at an angle of 40° with respect to the fire deck. Essential engine and testing parameters are summarized in Table 1.

Bore	92 mm
Stroke	85 mm
Displacement	560 cm ³
Compression ratio	11
Speed	0 (static), 1200, 1500 rev/min
Intake pressure	1 bar

Table 1: Engine specifications and experimental conditions.

To separate various error sources, four different modes of operation were used: static, fully homogeneous, pre-IVC DI, and post-IVC DI with various injection timings (IVC: intake valve closure). In static tests, a few drops of acetone were released into the combustion chamber of the stopped engine. After the acetone has evaporated, experiments that do not depend on the engine actually running could be performed in a very clean and repeatable environment. PLIF images were used to verify that the acetone had mixed with the air in the

cylinder to the point where no inhomogeneities could be detected anymore. In the motored engine, fully homogeneous operation was achieved by mixing the bulk gas with a hydrogen/acetone mixture from a gas bottle upstream of the intake plenum. Pre-IVC injection from -300 to -240° CA yields a nearly homogeneous mixture. (The crank angle is given as -360° CA to 360° CA, with 0° CA corresponding to top dead center of the compression stroke.) In using this fueling scheme to create reference conditions, any lack of homogeneity will cause corresponding measurement errors, an inaccuracy that will be assessed below. Finally, retarded, post-IVC injection is the actual target mode of operation, about which we wish to learn through quantitative imaging measurements.

To familiarize the reader with the underlying problem of fuel-distribution measurements in the DI-H2ICE, typical equivalence-ratio fields from an early post-IVC injection are shown in Figure 1. A detailed analysis of direct injection, mixing, fuel distribution, etc., is not subject of this paper. Some general observations include that the single shot shows details of turbulent mixing of hydrogen jets that have hit the cylinder liner opposite from the injector (located at the top of the image). Resolution and noise of the images seem reasonably good, though there is clear evidence of shot noise in Figure 1a. In the mean image of Figure 1b the location of the fuel clouds is in the bottom corners, towards the cylinder liner. The rms image of Figure 1c shows that intense mixing with the lean center is restricted to a narrow layer at mean equivalence ratios of about 0.5.

Since the data acquisition rate in the experiments is much slower than the in-cylinder processes in the running engine, statistical information about different planes or points in time has to be obtained from additional sets of measurements. For the present purpose of investigating the details of the experimental method itself, planes were chosen as appropriate for a particular aspect and the location of the plane is indicated when relevant.

Since these PLIF measurements are aimed at the pre-combustion mixture preparation only, the engine was motored with nitrogen as the bulk gas. Bulk flows were metered through sonic orifices. The fuel injector was calibrated off-line by injecting from a holder heated to engine temperature into a pressure vessel. The external surface temperature of the injector was 72°C. Temperature and pressure of the vessel were monitored and the amount of hydrogen injected per cycle for a given injection duration was calculated from this information. Neither the accuracy of the bulk-gas metering nor that of the injector calibration has been assessed here. Bias from these calibrations should be on the order of a few percent. Any bias in the nitrogen or hydrogen flow will add to the biases discussed below.




Figure 1: Examples of equivalence-ratio images from acetone PLIF in the DI-H2ICE. The field of view covers the central 65 mm of the 92 mm bore. (a) a single shot, (b) mean, and (c) rms. Intake pressure 1 bar, 1500 rpm, injection pressure 82 bar, injection timing 140°CA to -122°CA yielding a global equivalence ratio of $\phi = 0.25$. Imaging at -100°CA in a plane 11.5 mm below the fire deck. Injection is from a 6-hole injector located roughly at the top of the field of view [8].

TRACER AND SEEDING SYSTEM

CHOICE OF TRACER

Although most commercial liquid hydrocarbon fuels exhibit strong fluorescence upon ultraviolet excitation, for quantitative measurements it is common practice to substitute “real” fuel with a non-fluorescing model fuel which is mixed with a small amount (maybe 1 – 5%) of a fluorescing tracer. In this way the spectroscopic properties of only one single compound have to be taken into account, which is already not an easy task at engine temperatures and pressures.

In the case of hydrogen the pure fuel is completely non-fluorescing at any practically useful excitation wavelength. Both Rayleigh scattering [10] and Raman scattering [11-13] have been used to measure hydrogen concentrations in flows, but either technique is impractical for two-dimensional measurements in an engine, mainly due to weak signal strength. PLIF of a tracer added to the hydrogen fuel is therefore the only viable technique to obtain quantitative images of the in-cylinder fuel distribution. However, the choice of a tracer for hydrogen is not obvious. In general, the tracer substance should be as “similar” as possible to the fuel it is meant to represent, but of course also have good fluorescence properties, preferably with well-known photophysics. In the case of hydrogen for DI-H2ICE application, the most restrictive requirement is that the tracer must have a high enough vapor pressure to be present in the high-pressure fuel upstream of the injector at a concentration that allows adequate signal generation for the in-cylinder imaging after injection. Since both tracer and fuel are already in the vapor phase before injection, there is no concern about fuel-tracer co-evaporation, as there is for liquid fuel-tracer systems. However, to quantify the PLIF measurement, some

knowledge of the dependence of the fluorescence signal on temperature, pressure, and bath gas is required. Unless significant time is spent in preliminary investigations, it is therefore preferable to use a tracer with readily available spectroscopic data, and this restricts the choices severely.

The only tracer (other than acetone in our own work) for which application in a DI-H2ICE is reported in the literature is triethylamine (TEA). Using this tracer, Kirchweyer *et al.* [14] obtained ensemble-averaged PLIF images of the in-cylinder fuel distribution of hydrogen following direct injection. Previously, TEA had been used as a tracer for visualizing hydrogen after port-injection [13] and together with benzene for Exciplex PLIF tracing isooctane in an SI engine [15]. Rubas *et al.* [39] report qualitative PLIF measurements in the jets from a high-pressure methane injector in a large-bore diesel engine. As in the present work, they seeded the fuel with acetone using a bubbler.

For liquid fuels, the literature reports on the use of many different tracers, but the two most commonly employed ones are 3-pentanone and toluene. The photophysics of these two tracers, and also those of acetone, the simplest ketone, have been investigated extensively for engine-relevant temperatures and pressures. An excellent review on this topic is given by Schulz and Sick [16]. This review also contains an overview of alternative, less-used tracers such as biacetyl, acetaldehyde, and tetrahydrothiophen.

As in previous work [7, 8] we have chosen acetone as a gaseous fuel tracer. Table 2 shows a few key physical properties of acetone as compared to toluene and TEA. Also indicated are the approximate wavelength ranges of absorption and emission spectra.

	Acetone	Toluene	TEA
Mol. Mass [g/mol]	58.1	92.1	101.1
Boiling point at 1 bar [K]	329.4	383.8	361.9
Saturation concentration in H ₂ at 100 bar, 300K [% Vol.]	0.33	0.042	0.10
Diffusion Coefficient [cm ² /s]	0.11	0.079	0.072
Excitation wavelengths [nm]	230 - 320	240 - 270	Max. at 240
Emission wavelengths [nm]	325 - 500	270 - 320	270 - 330

Table 2: Physical properties of potential tracers for hydrogen. The diffusion coefficient is given for binary diffusion into nitrogen at standard temperature and pressure.

All three tracers can be excited with a frequency-quadrupled Nd:YAG laser at 266 nm, the laser available for the present work. The fluorescence of acetone is to a large fraction in the visible, which means inexpensive, large-aperture camera lenses can be used. Also, the quantum efficiency of most detectors is highest for visible light. If we exclude the control-intensive option of heating the hydrogen fuel, acetone can be seeded into the fuel at higher concentrations than both TEA and toluene, due to its higher vapor pressure.

For a given concentration and laser energy, the fluorescence signal depends on the tracer's absorption cross-section and quantum yield. The absorption cross-section of acetone is about 10 times smaller than those of toluene and TEA [16, 17], but applying higher laser energy can compensate for this in part. Also, if excessive attenuation of the laser sheet throughout the combustion chamber is to be avoided, laser absorption can become a limiting factor and hence a greater absorption cross-section may not confer any advantage. At room temperature and in the absence of oxygen, the quantum yield for acetone is about three orders of magnitude higher than that of toluene. Increasing temperature decreases the quantum yield for both tracers, but much more so for toluene [18, 19]. The dominant mechanism by which the excited state is depopulated in acetone is intersystem crossing, a process that is influenced very little by the presence of oxygen. Toluene, on the other hand, is strongly quenched by oxygen. Since the mixing measurements to be made here will be with the engine motored with nitrogen as the bulk gas (instead of air), a stronger signal might be obtainable with toluene, but it would be more sensitive to the necessary temperature corrections.

For TEA, little photophysical data are available at engine-relevant conditions. Measurements in a high-

pressure cell and a motored engine imply a moderate dependence of the fluorescence on temperature, and no pressure dependence [13]. The lack of detailed data on the temperature dependence of the fluorescence process prevents application of a local temperature-correction scheme. Both studies that used TEA in a hydrogen-fueled engine relied on a global calibration method, which can be inaccurate for local mixture conditions.

The diffusion coefficient of all three tracers is about one order of magnitude less than that of hydrogen (0.73 cm²/s into nitrogen). Hence there is concern that the tracer may not adequately follow the fuel that it is supposed to trace. This question of the significance of such "differential diffusion" is difficult to address, and the answer depends on what it is that we are actually trying to measure. On one end of the spectrum is the mean field of the scalar (here, of the hydrogen concentration), such as the image in Figure 1b, with length scales so large that molecular diffusion is unlikely to be relevant. On the other end are instantaneous realizations of the scalar field, as captured by the single shot in Figure 1a, with gradients occurring down to the resolution of the imaging technique. At the smallest scales of turbulent flows molecular diffusion is always important. If the imaging technique can resolve those scales, differential diffusion will certainly impact the measurement. For example, gradients in the actual hydrogen concentration will be smaller than those measured by PLIF of the less diffusive tracer. The spatial distribution of the fluctuations in the fuel distribution, as given by the ensemble rms in Figure 1c, is in between these two extremes. It is useful to consider the relevance of differential diffusion between fuel and tracer for this quantity.

The mixing action of turbulence can be quantified as a "turbulent diffusivity"

$$\nu_t = u^* l^* \quad (1)$$

Since this quantity is of central importance in turbulence modeling, many sophisticated approaches exist for its mathematical description. Here very crude estimates of the velocity scale u^* and the length scale l^* will have to suffice. Taking the velocity fluctuations to be correlated to the scalar fluctuations, a reasonable estimate for u^* would be the ensemble rms of the velocity. In previous work [8] the rms of the magnitude of the two velocity components measured was found to average about 6 m/s towards the end of the compression stroke, with hydrogen DI and operating conditions similar to those targeted in our current work. The corresponding length scale l^* is the integral scale of the flow. Heywood [40] suggests a value of 2 mm for the end of the compression stroke in automotive-sized engines. This yields a turbulent diffusivity of $\nu_t = 120 \text{ cm}^2/\text{s}$, much greater than the molecular diffusion coefficients of both hydrogen into nitrogen and acetone into nitrogen. Therefore for the

measurement of fluctuations in the fuel concentration the tracer can be considered to follow the hydrogen fuel well. However, for derived quantities with much smaller length scales this might not be the case, and a separate evaluation of the importance of differential diffusion would be required. Also, the above estimates for u^* and l^* were for the top of the compression stroke, with significant time passed since the direct-injection event. During the injection and close to the hydrogen jet, length scales can be expected to be significantly smaller and velocity fluctuations to be higher, again necessitating renewed estimation of v_T . Finally, Blotevogel *et. al* [13] have compared fluorescence of TEA in hydrogen to direct, simultaneous hydrogen-concentration measurements via Raman scattering and concluded that effects of differential diffusion were not observable.

In conclusion, given the lack of data on tracers other than toluene and acetone, the latter appears to be a reasonable choice when investigating the mixture formation in a DI-H2ICE. However, we have not directly compared different tracers, and it is well possible that other tracers could perform better.

SEEDING SYSTEM

To orient the reader in terms of experimentally useful tracer concentrations, Table 3 presents the acetone saturation concentration for a few pressures that are expected to be relevant for current and future DI-H2ICE development. The table also shows that at room temperature relatively small temperature differences result in significant changes in the acetone saturation-concentration.

Mixture Pressure [bar]	Mixture Temperature [K]	Acetone fraction [% Volume]	Acetone fraction [% Mass]
50	300	0.66	19.1
100	300	0.33	9.6
150	300	0.22	6.4
200	300	0.17	4.8
100	295	0.27	7.7
100	305	0.41	11.8

Table 3: Acetone volume and mass fraction in hydrogen at the evaporative equilibrium for selected temperatures and pressures.

In previous experiments [7, 8] a pre-mixture of hydrogen and 0.5% acetone (by volume), stored in a standard gas bottle at 54 bar, was used to supply the engine with fuel. Allowing for a reasonable duration of the experiment, the injection pressure was limited to 25 bar. While pre-mixing yields a precisely defined tracer concentration, this method of seeding is unsuitable for higher injection pressures, because the tracer is already present in the

fuel delivery system where pressure is the highest. Instead, by seeding directly upstream of the injector, the gaseous tracer concentration is limited by the lowest pressure in the system.

The acetone seeder employed here is a high-pressure vessel with hydrogen bubbling through the acetone. The functionally essential parts of the seeding system are shown schematically in Figure 2. Auxiliary components like pressure gages, vents, and safety feature have been omitted for clarity. Apart from the bubbler itself, the seeder contains a bypass that together with a check valve prevents spilling liquid acetone into the output or input lines when pressurizing or de-pressurizing the system. The seeder is refilled through a transparent riser that also serves as a level gauge. A thermocouple measures the temperature of the liquid acetone, another one that of the gas stream leaving the seeder.




Figure 2: Schematic of the high-pressure acetone seeder.

The main advantage of this bubbler-type seeder, as compared to injector-based alternatives (for example [13]) is its relative simplicity and reliability. The major disadvantage is that the acetone flow rate cannot be controlled directly. Instead it depends on the system temperature and the mass transfer inside of the hydrogen bubbles. In tests it was found that the hydrogen does not necessarily saturate with acetone vapor in the bubbler. Figure 3 shows the spatially integrated fluorescence signal from four separate engine runs, spaced over the course of a day. The engine was operated with pre-IVC DI. Each trace has been corrected for variations in the average laser power. As will be shown below, for pre-IVC injection timing the cycle-to-cycle variability of the PLIF signal is mostly due to shot-to-shot fluctuations of the laser energy. The data are also normalized by the saturation vapor-pressure at the temperature of the seeder for each experiment (varying between 296 and 300 K). The temperature used for this correction is that of the liquid acetone. The gas-stream temperature was found to follow the liquid temperature to within 1 K, and both remained constant within 1 K during each experiment.




Figure 3: Spatially integrated PLIF signal for four different engine runs under nominally identical conditions over the course of a day. Average laser power and acetone temperature differ slightly for each run and have been corrected for.

Figure 3 shows that for the length of time for each PLIF experiment (3 – 6 minutes) the signal level, and therefore the acetone concentration, stays constant. The exception is the first experiment, which was the initial run after filling and pressurizing the seeder. Successive experiments show decreasing acetone concentration, probably due to a lesser degree of vapor saturation in the lowering liquid column.

An on-line response calibration can account for the variation in acetone concentration: an experiment consisted of a loop over a set of cycles, with each set containing a cycle with a nearly homogeneous charge prepared by pre-IVC DI (for calibration and response correction), an unfueled cycle, (for background correction), and the data cycle (with typically inhomogeneous charge). Also, as usual in optical engine experiments, the engine was “skip-fired”, or in this case, “skip-fueled”.

LASER AND CAMERA SYSTEM

EXPERIMENTAL ARRANGEMENT

Figure 4 shows a schematic of the laser and camera system. Light from a frequency-quadrupled Nd:YAG laser is formed into a thin, horizontal sheet, exciting acetone fluorescence in the cylinder, which is imaged through the bottom of the Bowditch-type piston onto a CCD camera. The field of view extends over the central 65 mm of the 92 mm bore. Integrated into this typical PLIF arrangement is an electro-mechanical shutter, a passive single-cavity beam stretcher, a variable attenuator, and a laser-energy meter.

The Nd:YAG laser provides pulses with about 100 mJ of energy and 8 ns duration at 10 Hz repetition rate. The mechanical shutter is used to “downsample” this repetition rate to the 2.5 Hz at which the CCD can be read out. A passive-cavity beam stretcher lengthens the effective pulse duration and lowers the peak power by

creating a short pulse train from a single input pulse. Pulse stretching is more commonly employed in Raman line-measurements [13, 20-24] to prevent laser-induced plasma breakdown in the probe volume. In this experiment, the lower peak power allows higher laser energies without damaging the engine windows. The longer pulse duration also avoids depleting the ground level of the fluorescence transition, an issue that will be investigated in more detail below. The current single-cavity stretcher lowers the peak power by about a factor of 2.5, while preserving the total energy per pulse.




Figure 4: Schematic of the optical arrangement. Focal lengths are in millimeters.

An attenuator consisting of a half-wave plate and a thin-film polarizer allows for adjusting the laser energy without changing the beam’s energy-profile. Finally, the energy of each laser shot is measured and recorded by a laser-energy meter (Coherent LabMax TOP with J25-MUV-193 sensor head). For this measurement an aperture spatially filters the back-reflection of the final focusing lens such that the sensor head receives approximately the same transverse section of the beam that illuminates the probe volume.

DETECTION SYSTEM

In previous PLIF experiments [8] an intensified CCD (Princeton Instruments PI-MAX 512, “super-blue” Gen II intensifier) detected the fluorescence signal. The current experiment employs an unintensified back-illuminated CCD camera (Princeton Instruments PIXIS 1024B, UV-enhanced). For all results presented here the chip is binned 2 x 2 prior to read-out. Projected into the measurement plane, the pixel area is 132 x 132 μm^2 . The actual in-plane resolution of the imaging system, including piston window and viewing mirror, was measured using a scanning knife-edge technique [25]. The resulting line spread function (LSF) is very well approximated by a Gaussian function with a full width at half-max of 270 μm . The corresponding resolution in the

out-of-plane dimension, given by the thickness of the laser sheet, is estimated to be 500 μm .

The unintensified camera's spectral response is not optimized for the visible part of the acetone fluorescence spectrum. Nevertheless, its quantum efficiency of about 62% compares very favorably to the 22% of the photocathode on the intensified camera. In addition, the

resolution of the detector is only limited by the pixel pitch, and not degraded by the intensifier. Since increased resolution can always be traded off with shot noise via filtering, this translates into a better signal/noise ratio at a given resolution. Figure 5 illustrates how this trade-off can be used to compare the detection noise of the currently used CCD to that of the previously used ICCD in an approximate fashion.




Figure 5: Visual comparison of images from the intensified vs. the unintensified CCD. A little over a quarter of the entire field of view is shown for better resolution in print. All images were taken under identical conditions for injection, engine, and imaging. Top row: *Three different* single shots from the intensified CCD at native resolution. Bottom row: *One single* image from the unintensified CCD with different amounts of filtering (given as the kernel width of the Gaussian filter) applied.

The images in Figure 5 were all taken under the same late-injection condition, which results in significant fuel-concentration gradients in the field of view. Visual comparison of the three single shots from the ICCD to the unfiltered image from the CCD shows that the former much less than the latter is able to capture the fine structure associated with the high gradients in the mixing zone. That is to say, the CCD has better resolution than the ICCD. For a proper comparison, the resolution has to be degraded to match that of the ICCD. This is done here by Gaussian filtering with different kernel sizes. Judged by the maximum gradients in the images, a kernel with $\sigma_{\text{gsm}} = 1.5$ pixels yields roughly comparable resolution. Images with a homogenous concentration field (not shown) are then evaluated in terms of single-shot detection noise, smoothing the images from the CCD with $\sigma_{\text{gsm}} = 1.5$ pixels while leaving those from the ICCD at native resolution. The resulting signal/noise ratios (SNRs) are

- SNR (ICCD) = 15
- SNR (CCD) = 54.

At comparable resolution, the CCD's detection noise is therefore 3.6 times less than that of the ICCD. If somewhat less or more smoothing is deemed the "proper" amount of filtering in this approximate comparison ($\sigma_{\text{gsm}} = 1.0$ pixels and 2.0 pixels), the noise advantage ranges from a factor of 3 to 4.

It should be stressed that it is not the intention of the previous paragraph to deliver an exact comparison of the two camera systems, but rather to give a rough idea of how the detectors compare so that a decision on which one to use can be made. Rothamer and Ghandhi [38] have performed a more complete comparison of different camera types with respect to optical engine measurements. Indeed, unintensified CCDs have in general the major disadvantage that they rely on a (slow) mechanical shutter to block unwanted light, and therefore the shortest exposure times are several orders of magnitude longer than for ICCDs, which routinely are gated for less than 100 ns. However, in the motored experiments targeted here, there is no need to discriminate against light from a combustion event. Hence, for this limited purpose, the unintensified CCD is

the better choice. The noise characteristics of this detector will be assessed in more detail in a section below.

DATA REDUCTION

Except where explicitly noted, standard background and response corrections were applied to the PLIF images. This correction can be written as

$$I_i^{Corr} = \frac{I_i - \langle I_i^{BG} \rangle}{\langle I^{Resp} \rangle - \langle I^{BG} \rangle} \quad (2)$$

where I is the PLIF intensity, i is the image number, and the superscripts BG , $Resp$ and $Corr$ refer to background, response ("flat field"), and corrected images, respectively. Following the nomenclature of Herold and Gandhi [9], $\langle \rangle$ denotes ensemble averaging over a set of images. The over-bar $\bar{}$ will be used to indicate spatial averaging within one image or image sub-region.

In the PLIF measurements, the background images are acquired by firing the laser in a cycle without hydrogen injection. The response of the optical system, often imprecisely called "flat field", is recorded by averaging cycles in which pre-IVC injection has created a nearly homogenous mixture at the time of imaging. Both background and response correction are not perfect, but have limits in their accuracy, which will need to be assessed.

MEASUREMENT PRECISION

The precision of the PLIF measurement is taken here to be any deviation of a given single-shot image from the ensemble average that is not caused by an actual variation of the hydrogen concentration in the measurement plane. The measurement precision has no influence on the ensemble (multi-cycle) mean, but can instead be characterized by the ensemble root-mean-square (rms) of the deviation of the instantaneous measurements from their physical value:

$$rms = \sqrt{\left\langle \left(I_i^{corr} - \langle I_i^{corr} \rangle \right)^2 \right\rangle} \quad (3)$$

This is a useful measure of precision when for a particular experiment the true, physical value of any instantaneous measurement can be approximated by the ensemble mean over a set of such measurements (as for time-invariant systems or periodic ones with phase-locked data acquisition).

Precision of the PLIF measurement is assessed in terms of detection noise, laser energy variation, transverse laser-mode variation, "beam steering", and variation in concentration over the course of the measurement. We

will seek to quantify each of these potential sources of single-shot error by its ensemble rms. Depending on the error assessed, the rms is evaluated over the whole image or a suitable sub-region.

DETECTION NOISE

In a previous section a rough comparison between intensified and unintensified camera was made. The better signal/noise characteristics motivated the choice of the unintensified CCD for the targeted motored experiments. To assess the precision of this detector in more detail, a series of image sets in the static engine was taken. For each set, a small amount of acetone was placed with a pipette into the combustion chamber of the engine, yielding a uniform air/acetone mixture upon evaporation. Different amounts of acetone result in different signal strengths. After background subtraction, the signal intensity was evaluated in a region in the center of the image. The image region was then "flattened" by the standard corrections and ray-mean normalization [9] such that the remaining pixel noise could be evaluated independently of shot-to-shot variations in the laser sheet. Figure 6 shows the spatial rms over the evaluated region versus the fluorescence signal from that region. The data point at zero signal corresponds to an image with the lens covered, *i.e.*, complete darkness.




Figure 6: PLIF rms versus square root of the PLIF signal. Also indicated are the theoretically expected shot noise and the combined shot noise and read-out noise. In the part of the graph designated as "shot-noise limited", the read-out noise is less than 10% of the total detector noise.

The straight line through the origin corresponds to the theoretically expected noise photon (quantum) noise. In the data acquisition software, a nominal conversion factor of 1 photoelectron to 1 camera count was chosen, which would yield a slope of unity in this plot. However, as was also the case in the investigation of Herold and Gandhi [9], it is apparent that the manufacturer-reported gain factor is slightly in error. Based on the highest-intensity data point, the slope was adjusted to 1.05. Since dark noise is usually negligible for cooled