

Quantitative Laser Diagnostics for Combustion



Benjamin A. O. Williams

Merton College.

DPhil, Trinity 2009

Quantitative Laser Diagnostics for Combustion

Benjamin A. O. Williams, Merton College. DPhil, Trinity 2009.

Abstract

Quantitative Planar Laser Induced Fluorescence (QPLIF) is developed as a diagnostic technique and then applied to a prototype Jaguar optical internal combustion engine. QPLIF derives quantitative, two-dimensional, spatially-resolved measurements of fuel concentration. This work reports the first demonstration of a fully-fractionated surrogate fuel which exhibits all the characteristics of a typical gasoline. This ‘pseudo’ fuel, developed in association with Shell UK, is blended to accept a fluorescent tracer which may track one of the light, middle or heavy fractions of the fuel, each of different volatility. The traditional weaknesses of PLIF for quantitative measurements are addressed by use of a fired *in-situ* calibration method, which maps the quantum efficiency of the tracer and concurrently corrects for window fouling and exhaust gas residuals (EGR). Fuel distributions are presented with an estimated super-pixel accuracy of 10% at different operating conditions, and then compared to the computational fluid dynamics (CFD) predictions of an in-house Jaguar model.

Fuel/Air Ratios by Laser Induced thermal Gratings (FARLIG) is developed theoretically, and results of validation experiments conducted in a laboratory setting are reported. FARLIG conceptually enables the measurement of fuel concentration, oxygen concentration and temperature within a spatially-localised probe volume. Uniquely, the technique exploits the dominant influence of molecular oxygen on non-radiative quenching processes in an aromatic tracer molecule. The changing character of a model quenching mechanism potentially allows the oxygen concentration in the measurement volume to be derived. Absolute signal strength is used to determine fuel concentration, while the oscillation period of the signal provides a precise measurement of temperature ($\sim 0.3\%$ uncertainty), with accuracy limited by knowledge of the gas composition.

Acknowledgements

I would like to take this opportunity to thank those without whom the last four years of research - and this thesis - would not have been possible. In common with most things, there is much more to it than meets the eye - so whilst my name may stand alone on the title page, I owe much to excellent supervision from Prof. Paul Ewart and Prof. Richard Stone. I am very grateful to them both for their superb counsel and supportive natures. I also thank Prof. Ian Walmsley for his presentation about the “Lasers and QIP” option, without which I might not have chosen *the* correct fourth year option. I also thank him for taking me in as a fourth-year MPhys project student, and then warmly welcoming me to the Atomic and Laser Physics sub-department as a graduate student.

I would not be where I currently am without the valuable friendship and inspiring presence of Dr. Rob Stevens, who has probably taught me more than he realises. Indeed, I have never met anyone before or since who could quite so effortlessly embarrass my software knowledge. For re-shaping my brain on a regular basis, Rob, thank you.

My other colleagues on the Physics side of things - Dr. Yoshihiko Arita and Dr. Katherine Richard - have both been highly influential, and I thank them for their friendship. In particular, Yoshi for those late-evening discussions, Kate for demonstrating the power of an iron will, and them both for taking me to the hospital following a rather unplanned trip over the handlebars of my bicycle!

I also extend my warmest thanks to my fellow students in the Engineering Science department - to Dr. Phil Price, Dr. Xiaowei Wang, Ben Twiney (of Ben² fame), Steve Marshall, Mike Braisher and Longfei Chen. Thank you all not just for your aid and friendship, but also for enduring a (probably seemingly endless) series of UnICEG and CCSPV presentations about QPLIF!

A special mention goes to the support staff of the Clarendon Laboratory. Rob Harris in the Research Workshop is a master of his domain, and I am honoured to have received his tuition at the helm of a variety of machines. George, Terry and Alan from Maintenance deserve many thanks for their help, as do Alan and Mohammed in Stores.

Finally, I would like to thank my family and friends for understanding whilst I have been otherwise engaged with work. I am indebted to you all for the time away from Oxford we have spent together - even if it may have only been for a couple of hours, it makes all the difference.

Ben Williams, 20th April 2010.

Contributions to this research

The work presented in this thesis is my own, except where identified in the main text or indicated with a reference, and with due consideration to the text of the acknowledgements.

The following is a list of the significant areas where I have used the results of (or otherwise incorporated) work done by others:

- Design and testing of the multi-component fuel used in QPLIF tests
- Initial build of `HDAQRed.m`, used to analyse HDAQ data from the single cylinder optical engine
- CoBRA software and burn rate analysis correlation against QPLIF
- CFD modelling, and generation of CFD images to compare to QPLIF distributions
- Initial build of LITGS fitting code

Ben Williams, 20th April 2010.

Contents

| | |
|---------------------------------------------------------------------|----------|
| 1. Introduction | 1 |
| 1.1. Motivation | 1 |
| 1.2. Thesis Structure | 3 |
| 2. Literature Review | 5 |
| 2.1. Chapter overview | 5 |
| 2.2. Terminology | 5 |
| 2.2.1. Engine terminology | 5 |
| 2.2.2. Combustion terminology | 8 |
| 2.3. Relevant parameters | 9 |
| 2.4. The optical advantage | 10 |
| 2.5. Optical techniques summary | 11 |
| 2.6. Linear optical | 12 |
| 2.6.1. Absorption Spectroscopy | 14 |
| 2.6.2. Photon emission techniques | 18 |
| 2.6.3. Scattering Techniques | 24 |
| 2.7. Non-linear optical | 27 |
| 2.7.1. Coherent Anti-Stokes Raman Scattering (CARS) | 28 |
| 2.7.2. Degenerate Four Wave Mixing (DFWM) | 33 |
| 2.7.3. Laser Induced Thermal Grating Spectroscopy (LITGS) | 34 |
| 2.8. Technique selection and closest rivals | 37 |

| | |
|------------------------------------------------------------------|-----------|
| 2.8.1. QPLIF | 38 |
| 2.8.2. FARLIG | 40 |
| 3. Theory of LIF, LITGS and Tracer photophysics | 43 |
| 3.1. Chapter overview | 43 |
| 3.2. LIF | 43 |
| 3.2.1. LIF signal production | 43 |
| 3.2.2. The impact of quenching on LIF | 46 |
| 3.2.3. Measurements with LIF | 48 |
| 3.2.4. Pros and Cons of LIF | 49 |
| 3.3. Thermal Gratings | 50 |
| 3.3.1. Phenomenological discussion of LITGS | 50 |
| 3.3.2. LITGS theory | 53 |
| 3.3.3. Deriving parameters from LITGS signals | 59 |
| 3.3.4. Pros and Cons of LITGS | 61 |
| 3.4. Tracer Photophysics | 63 |
| 3.4.1. Toluene (aromatic) | 63 |
| 3.4.2. Acetone (ketone) | 69 |
| 3.4.3. Motivations for tracer selection | 72 |
| 4. Quantitative Planar Laser Induced Fluorescence - QPLIF | 73 |
| 4.1. Chapter overview | 73 |
| 4.2. Technique overview | 73 |
| 4.2.1. Measuring fuel distributions with QPLIF | 73 |
| 4.2.2. 2-level model and QPLIF | 74 |
| 4.2.3. CCSPV Consortium | 75 |
| 4.3. Experimental details | 76 |
| 4.3.1. Multi-component fuel | 76 |
| 4.3.2. Engine details | 80 |

| | | |
|-----------|-------------------------------------------------------------------|------------|
| 4.3.3. | Optical details | 83 |
| 4.4. | Calibrating QPLIF | 91 |
| 4.4.1. | Calibration loop | 91 |
| 4.4.2. | The EGR effect | 97 |
| 4.4.3. | 'Bookend' calibration | 101 |
| 4.4.4. | Calibration summary | 105 |
| 4.5. | Analysis system | 107 |
| 4.6. | Testing QPLIF | 113 |
| 4.6.1. | Quantifying PLIF signals | 113 |
| 4.6.2. | Verification of QPLIF | 114 |
| 4.6.3. | Accuracy and precision of QPLIF | 116 |
| 4.7. | Applying QPLIF | 119 |
| 4.7.1. | QPLIF imaging | 119 |
| 4.7.2. | Average distributions for Late-DI | 120 |
| 4.7.3. | Application of CoV to QPLIF | 124 |
| 4.7.4. | Correlations with burn-rate analysis | 126 |
| 4.7.5. | Comparison with Jaguar CFD | 127 |
| 4.8. | Summary | 130 |
| 4.8.1. | Achievements | 130 |
| 4.8.2. | Prerequisites for Quantitative PLIF | 131 |
| 5. | Fuel/Air Ratios by Laser Induced thermal Gratings - FARLIG | 133 |
| 5.1. | Chapter overview | 133 |
| 5.2. | Technique overview | 133 |
| 5.3. | Measuring fuel/air ratios with LITGS | 135 |
| 5.4. | Experimental details | 137 |
| 5.4.1. | Lasers and optics | 137 |
| 5.4.2. | Timing and data acquisition | 142 |

| | | |
|-----------|---------------------------------------------------|------------|
| 5.4.3. | Gas cell considerations | 143 |
| 5.4.4. | Detector characterisation | 147 |
| 5.4.5. | Alignment procedure | 150 |
| 5.5. | Data analysis | 153 |
| 5.5.1. | Modelling theory | 153 |
| 5.5.2. | Modelling a LITGS signal with MATLAB | 155 |
| 5.5.3. | Fitting to a LITGS signal with MATLAB | 157 |
| 5.6. | Results | 162 |
| 5.6.1. | Fuel concentration | 162 |
| 5.6.2. | Oxygen concentration by quenching model | 164 |
| 5.6.3. | Oxygen concentration by signal contrast | 182 |
| 5.6.4. | Temperature | 184 |
| 5.7. | Summary | 186 |
| 6. | Conclusions | 188 |
| 6.1. | QPLIF | 188 |
| 6.2. | FARLIG | 189 |
| 6.3. | Comparisons | 190 |
| 6.3.1. | QPLIF vs. FARLIG | 190 |
| 6.3.2. | Novelty in QPLIF | 191 |
| 6.3.3. | Novelty in FARLIG | 192 |
| 6.4. | Summary | 192 |
| 7. | Future Prospects | 194 |
| 7.1. | QPLIF | 194 |
| 7.2. | FARLIG | 195 |
| A. | QPLIF supplement | 213 |
| A.1. | QPLIF preparation procedure | 213 |

| | |
|------------------------------------------|------------|
| A.2. QPLIF schedule | 217 |
| B. Semi-QPLIF measurements | 219 |
| B.1. Shot-wise distributions | 219 |
| B.2. IVO PLIF | 220 |
| B.3. Imaging sprays | 221 |
| C. FARLIG supplement | 223 |
| C.1. Gas cell refill procedure | 223 |
| C.2. Alignment technique | 224 |
| C.3. Schedule | 225 |
| D. FARLIG fitdata | 229 |

1. Introduction

1.1. Motivation

The roots of the modern internal combustion (IC) engine can be traced back to the end of the 19th century. In the last hundred or so years, many improvements have been made - for example, in terms of efficiency, specific power output and so on. However, it has become increasingly clear in the past decade or two that the effects of IC engine emissions on global warming is an issue that cannot be ignored. In particular, carbon dioxide emissions are an important contributor to this effect. Government data from 2008 indicates that road transport is responsible for 22% of carbon dioxide emissions in the UK¹. Following the failure of the *voluntary* ACEA agreement whereby all European car manufacturers had agreed to reduce CO₂ emissions to 140gkm⁻¹ by 2008[†], the prospect of EU regulation of CO₂ emissions seems increasingly likely. (Of course, carbon dioxide is not the sole emission from internal combustion engines, but pollutants such as NO_x are already covered by the EURO standards².) As a significant source of man-made greenhouse gas emissions and pollutants, and in view of the reality that internal combustion engines will continue to be in widespread use³, the internal combustion engine bears further attention to improve its characteristics.

Optical techniques for measuring combustion parameters have clear advantages over their traditional counterparts (e.g. the thermocouple), since they offer the opportunity

[†]Actually, Fiat and Peugeot did manage it, with fleet averages of 133.7gkm⁻¹ and 138.1gkm⁻¹ respectively.

to take measurements without perturbing the phenomenon under investigation. In this thesis, two optical diagnostic techniques will be investigated with respect to their ability to measure parameters of interest to the combustion process in an engine.

Mixture preparation - i.e. the distribution of fuel, temperature and so on within a cylinder - just prior to ignition is of great importance to the emissions and efficiency of the burn. Careful manipulation of this distribution is a major focus of modern engine research, and in this thesis two techniques are presented which allow for important insights to be made in this area. In the case of spark-ignited (SI) gasoline engines (such as the spray-guided direct injection SI prototype used in this thesis), the nature of the early phase of the burn will be strongly influenced by the environment in the region immediately surrounding the spark plug. In the case of compression ignition engines (i.e. diesel or homogeneous charge compression ignition - HCCI - gasoline), ignition might be distributed across the cylinder. In this study, the optical access afforded by a window in the side of the pent-roof region of the cylinder allows measurements to be taken in the vicinity of the spark plug.

Owing to the complex nature of the in-cylinder interactions which lead up to the mixture distribution at the time of ignition, some variability is expected between the behaviour of sequential engine cycles. One poorly-prepared cycle may also affect subsequent cycles, as some residual burnt gases are generally retained from one cycle to the next. A well-calibrated, smooth-running engine is hard to achieve in the presence of high cyclic variability, and so this project also seeks to measure such fluctuations, and investigate how they may affect combustion.

One strategy adopted in modern IC engines at low to part-load is that of stratification - where the fuel distribution is manipulated by careful combination of such mechanisms as in-cylinder flow and injection duration and timing, in order to generate a slightly rich mixture in close proximity to the spark plug, which is lean overall. This kind of fuel distribution ensures ignition will occur, but also comes with the benefits of reducing fuel consumption and improving emissions. It is a very challenging problem

to ensure that the correct fuel distribution is created cycle after cycle.

The development cycle for this kind of engine would be significantly accelerated if significant levels of prototyping could be completed in a computer-based modelling environment. Naturally, this is only possible if the computer-based model is considered to be an accurate predictor of reality. In this respect, comparisons between computer models and experimental measurements are becoming increasingly important for modern and future engine development⁴. As an example of one such comparison, quantitative measurements of fuel distributions are compared to their computer-modelled counterparts towards the end of chapter 4.

1.2. Thesis Structure

This thesis is globally structured to progress from general discussion of a topic to specific details of implementation. Accordingly, the thesis begins with a general introduction to internal combustion engines and the role optical diagnostics plays in their development.

This introductory chapter is followed by chapter 2, which explores some of the technical background and subject-specific terminology that will be important in this thesis. A literature review forms the bulk of the chapter, and describes - in general terms - the key characteristics of a range of relevant optical diagnostic techniques.

Continuing in the progression from the general to the specific, chapter 3 contains a description of the theory of Planar Laser-Induced Fluorescence (PLIF), Laser Induced Thermal Grating Spectroscopy (LITGS), and the photophysics of commonly-used tracer molecules. This understanding forms a strong basis for the description of the two techniques implemented in this thesis which are described in the subsequent two chapters.

Chapter 4 is an account of the development and application of the QPLIF (Quantitative-PLIF) technique. The chapter includes a discussion of the fuel and tracer system and

the methodology required for quantitative measurements and analysis. Experimentally-measured fuel distributions are presented, and compared to their computer-modelled counterparts.

Chapter 5 details the development and application of FARLIG (Fuel/Air Ratios by Laser Induced thermal Gratings). FARLIG is an extension of LITGS, and the chapter describes how the FARLIG technique derives a fuel/air ratio, along with details of the experimental apparatus and methodology. The results of the validation trials are analysed and presented at the end of the chapter.

The presented work is then drawn together in chapter 6, where QPLIF and FARLIG are compared to one another and against previously-published work.

The final topic of discussion in this thesis regards future extensions to this work, and forms the content of chapter 7.

2. Literature Review

2.1. Chapter overview

This chapter begins by explaining some useful subject-specific terminology and important parameters, before selecting a subset of optically-measurable parameters relevant to this research. The motivations for optical diagnostics will then be briefly covered, before reviewing a number of techniques that attempt to measure the chosen parameters. The final element of this chapter explains the motivation for the choice of PLIF and LITGS, before describing in some detail the closest competitors to the two techniques developed in this thesis. This theme will be revisited later for a comparison in chapter 6, where the innovations in QPLIF and FARLIG will be compared to their rivals.

2.2. Terminology

In common with all technical subjects, the topic of internal combustion engine research involves many acronyms which may not already be familiar to the reader. This section attempts to provide sufficient grounding to understand the contents of this thesis.

2.2.1. Engine terminology

The following is a description of a simplified 4-stroke, spark ignition, direct-injection, gasoline engine. An example in-cylinder pressure trace is presented in figure 2.1.

An internal combustion (IC) engine operating on a four-stroke cycle repeats its actions with a period defined by two full rotations of the crankshaft (this rotation will eventually cause the driven wheels to spin). The progression of an engine cycle is tracked by reference to these 720 crank angle degrees (CAD). Unsurprisingly, the four-stroke cycle splits these 720 CAD into quarters. By convention, the cycle is said to begin with the *induction stroke*. During these initial 180 CAD, the piston, which is sealed to the bore of the cylinder, travels downwards. This motion pulls intake gas into the cylinder through the intake valves which are situated on the upper face of the cylinder. In the case of a modern engine featuring direct injection, the gas pulled into the cylinder at the start of the cycle is simply air. A direct injection (DI) engine simulates the character of older designs - which inducted a mixture of air and fuel into the cylinder - by injecting fuel into the engine at some point during the intake stroke[†]. Once the piston reaches the bottom of its stroke - referred to as ‘bottom dead centre’ (BDC), the *compression stroke* begins. With the intake valves closed, the piston travels back up its stroke towards the cylinder head. This reduces the volume of the mixture of fuel and air by a factor referred to as the compression ratio, which is typically approximately 10:1. The compression process is polytropic ($pV^n = \text{const}$), becoming adiabatic (i.e. no heat transfer) with rising engine speeds. This compression process happens on a short timescale[‡], and causes both pressure and temperature to rise in the mixture - though not so far that it self-ignites. The spark plug which is often centrally-located in the roof of the cylinder will tend to be timed to produce a discharge arc some time before the piston reaches the top of its stroke at ‘top dead centre’ (TDC). Following ignition of the mixture, the piston is pressed back down the bore of the cylinder during the *expansion stroke* by the rapidly-expanding hot gases produced by the burning of the air/fuel mixture, which releases heat. It is during these 180 CAD that this cylinder will do work and aid in turning the crankshaft. Following

[†]Since so-called ‘early’ injection will allow maximum time for the mixture to homogenise.

[‡]The UK motorway speed limit corresponds to an engine speed of 3250 rpm in a typical gasoline SI engine car - compressing the in-cylinder gases in just 10ms.

the next BDC event, the exhaust valve(s) open in the cylinder roof, and as the piston travels once more towards TDC the exhaust gases are expelled from the cylinder during the *exhaust stroke*. As the piston reaches TDC, the exhaust valves close, and the intake valves open, beginning the next cycle...

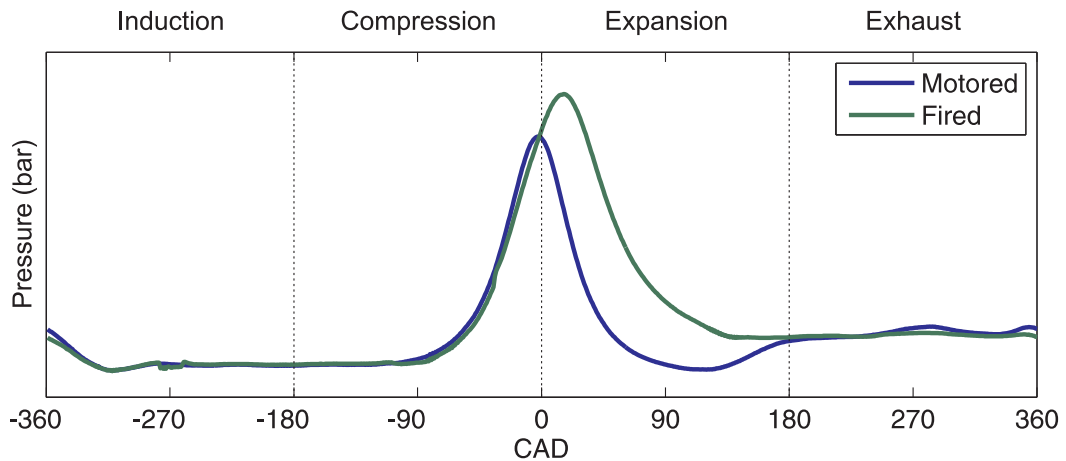


Figure 2.1.: Typical pressure traces recorded on the optical engine employed in this work, showing convention adopted for labelling events by CAD. The difference between fired and motored operation is explained in the main text.

There is no universal convention, but this thesis shall take zero CAD to be the TDC event at the end of the compression stroke (when the combustion event is occurring, see figure 2.1). This means that injection and ignition events occur at some negative angle, or, equivalently at some positive angle ‘BTDC’ - simply, Before Top Dead Centre.

Events are often labelled by acronyms:

- | | |
|-----------------------------|-----------------------------|
| IVO - Intake Valve Opening | IVC - Intake Valve Closing |
| EVO - Exhaust Valve Opening | EVC - Exhaust Valve Closing |
| SOI - Start Of Injection | EOI - End Of Injection |

Modern engines tend to make much use of EGR (exhaust gas recirculation/residuals). This is a process whereby either some of the intake air is substituted by exhaust gas from the previous cycle by use of pipework external to the cylinder (external-EGR), or where the valve events are chosen so that EVC happens some time before TDC, trapping some exhaust gas in the cylinder, which will take part in the next cycle

(internal-EGR). This selection of valve timing whereby EVC occurs before IVO is known as NVO - or negative valve overlap. In the case of complete combustion, the residual gases are largely inert, and take a passive role in a combustion event - merely acting as some kind of ‘filler’. As such, when the engine is under low load conditions (e.g. a light touch on the throttle) where a small quantity of fuel is injected, fuel economy can be improved without running ‘lean’ (which would otherwise increase pollutant levels, e.g. NO_x).

The practical implementation of IC engines in laboratories often sees the crankshaft (or flywheel) connected to a dynamometer. A dynamometer (‘dyno’ for short) is a device which is normally able to measure engine speed and torque, dissipating the power produced by the engine which is driving it. More advanced dynamometer designs employ electric motors and can essentially run in reverse - in this case, driving the piston up and down the bore etc. without combustion occurring. This is known as ‘motored’ operation (no ignition, engine driven externally), and can be useful when the combustion event itself is not the focus of a project. When the fuel/air mixture that is admitted into the engine is burnt, the engine is said to be running in ‘fired’ mode. Figure 2.1 shows pressure traces for both motored and firing engine conditions.

2.2.2. Combustion terminology

In a gasoline engine, the issue of stoichiometry is very important. For the complete combustion of the mixture which is inducted into the cylinder, the ratio of oxygen (i.e. air) to fuel must be kept at a particular value. This balance is described by the following equation, solved for the reference fuel iso-octane which has the chemical formula C_8H_{18} :

$$1 \text{ part fuel} + \alpha \text{ parts air} \Rightarrow \beta \text{ parts water} + \gamma \text{ parts carbon dioxide} \quad (2.1)$$

$$\text{C}_8\text{H}_{18} + \alpha(0.21 \times \text{O}_2 + 0.79 \times \text{N}_2) \Rightarrow \beta \times \text{H}_2\text{O} + \gamma \times \text{CO}_2 + \alpha \times 0.79 \times \text{N}_2 \quad (2.2)$$

Comparing coefficients in carbon, C, hydrogen, H, and oxygen, O:

$$8 = \gamma \quad (2.3)$$

$$18 = 2\beta \quad (2.4)$$

$$0.21 \times 2\alpha = \beta + 2\gamma \quad (2.5)$$

Which leads to the result that a stoichiometric mixture of iso-octane and air being mixed - *by number* - in the ratio of $\sim 1 : 59$. If the air/fuel mixture contains either surplus air or insufficient fuel for stoichiometry, the mixture is said to be lean. In the opposite case, the mixture is rich. The level to which the mixture is lean or rich is quantified by reference to the equivalence ratio. If twice as much fuel is present as the stoichiometric case, the mixture has an equivalence ratio, ϕ of 2. This may also be referred to as a ' λ ' of 0.5[‡].

2.3. Relevant parameters

This thesis is primarily concerned with the measurement of fuel concentration ('[fuel]') or air/fuel ratio (AFR), because it is of primary importance in a combustion event. However, it is recognised that temperature (T), pressure (p) and the presence of EGR (see section 2.2 for explanation of this acronym) all have important roles to play in the combustion process. However, to the best of the author's knowledge, in-cylinder experiments generally rely on the use of a pressure transducer, with no apparent drawbacks. Therefore, this chapter will only review techniques which have applications in measuring the presence of fuel, EGR and the value of T . Some of the listed techniques will have other applications, and where possible these will be acknowledged, but not considered in detail.

[‡]Lambda sensors measure oxygen, so higher lambda corresponds to greater dilution

2.4. The optical advantage

There are a multitude of devices targeted at taking measurements in IC engines. For the purposes of this discussion, they are categorised into two groups - the ‘traditional’ diagnostics tools, and the optical diagnostics tools.

‘Traditional’ measurement devices would include anything from thermocouples and pressure transducers through to flame ionisation detectors⁵ (FIDs) and 4-gas analysers. In general, these measurement devices have a major drawback regarding *in-situ* deployment, which is that in order to make a measurement, a probe must be introduced into the measurement location. This invariably has the effect of altering the parameter under investigation. To give a simple example, the introduction of a thermocouple into a flame to measure temperature will introduce a heat loss mechanism by conduction back along the length of the thermocouple itself, thus rendering the measurement inaccurate.

By contrast, an optical measurement technique avoids the requirement to introduce anything into the measurement region physically. Providing there is sufficient optical access, one or more laser beams can achieve the same result with no perturbation to the local environment. Optical measurements also provide the potential for highly spatially- and temporally-resolved measurements. For example, an imaging technique may be able to determine fuel concentration in $\sim 150,000$ locations simultaneously at a repetition rate of 12kHz⁶ - i.e. measure a distribution every crank-angle degree.

It is said that ‘there’s no such thing as a free lunch’ and that statement holds true here for the non-intrusive and non-perturbative nature of optical techniques is usually associated with increased cost and/or complexity. For example, one may manually measure a temperature with a K-type thermocouple and a suitable digital multimeter at the cost of approximately £50, to an accuracy and precision in the region of $\pm 3^\circ\text{C}$ [†].

[†]To re-iterate, the measured temperature is probably different from the temperature at the measurement location in the absence of the thermocouple - and so extra inaccuracy is likely to be introduced.

At the other end of the scale, the application of a technique such as LITGS requires the purchase of optics, a trio of fast detectors, a pair of lasers and an oscilloscope - and all the associated overheads, for a cost in the region of £100,000. However, the measurement *can* now be both accurate (owing to the non-invasive nature) and precise, producing a measurement with a precision in the region of 0.16%⁷.

Besides their cost in equipment, optical diagnostics clearly require optical access to the phenomenon under investigation. In the case of an engine, this necessitates the installation of windows into the cylinder head and/or cylinder walls. Windows are typically made from crown glass or fused silica, which has low thermal conductivity relative to metal. In order that these windows are not excessively scuffed during operation, the piston rings usually have to be moved lower on the piston, which increases crevice volumes. These and other considerations mean that the environment of an optical engine is not entirely representative of the production model, with possible differences in - for example - compression ratio, wall temperature and heat transfer.

2.5. Optical techniques summary

Despite the increased financial overhead, the ‘traditional’ measurement instruments and techniques will now largely be ignored due to their invasive and perturbative nature when used in-cylinder. (It should be emphasised that their use in environments where perturbations are less important - e.g. in the exhaust, as a lambda-sensor, or in the cylinder-wall, as a pressure transducer - is nevertheless still very valuable.) In concentrating on the optical measurement techniques, two classes are defined. The distinction between such ‘linear’ and ‘non-linear’ techniques is derived from the way the incident radiation interacts with the medium. In the general case, the induced polarisation, P , for a given incident E field is given by:

$$P(\omega) = \epsilon_0(\chi^{(1)}E(\omega) + \chi^{(2)}E(\omega)^2 + \chi^{(3)}E(\omega)^3 + \dots) \quad (2.6)$$

Where $(\mu_0\epsilon_0)^{-\frac{1}{2}} = c$, ϵ_0 being the permittivity of free space, μ_0 the permeability of free space, c is the speed of light in vacuum, and $\chi^{(n)}$ is the n^{th} order in the expansion of the susceptibility of the medium.

For the most part, daily life is only concerned with the linear response term, proportional to $\chi^{(1)}$, the linear term in the expansion of the susceptibility of the medium. However, when a sufficiently high driving force is applied, the response of the medium can no longer accurately follow the driver - leading to non-linear behaviour. As a result, at high intensities, the higher orders of the electric susceptibility (which have progressively smaller values) become important, and effects such as harmonic generation - where radiation is produced by the induced polarisation at a higher frequency than the input - become important. In this sense, the higher-order susceptibilities can be seen as coupling coefficients allowing multiple oscillating fields to interact with one another.

The properties of available optical techniques for diagnosing air/fuel ratio (or fuel concentration), temperature and EGR fraction are summarised in table 2.1.

Each optical technique will be reviewed independently in the following pages, firstly by considering the underlying concepts of operation, and secondly by describing applications from the literature. Techniques are split by class - linear in section 2.6 and non-linear in section 2.7.

2.6. Linear optical

A number of important, widely-used techniques lie in the class of linear optical diagnostics. It is instructive to consider the behaviour of a simple optical system in order to begin to see how these techniques are related to one another. Figure 2.2 is a simple graphical description of a system composed of two discrete electronic energy levels. Under equilibrium conditions, a group of such systems would lie in their lowest energy configurations - i.e. in the lower state, labelled '1'. When this system is excited by the

| Technique name | [fuel] or AFR | T | EGR | Measured on a 2D plane? | Class |
|-------------------------|------------------|-----|-----|----------------------------|-------------|
| LITGS | + | + | o | o | non-linear |
| CARS | + | + | + | o | non-linear |
| DFWM | o | + | o | + | non-linear |
| Absorption spectroscopy | + | + | o | o | linear |
| PLIF/PLIEF | + | + | + | + | linear |
| Phosphor doping | - | + | - | + | linear |
| Chemiluminescence | o | o | - | + | linear |
| Raman scattering | + | + | + | + | linear |
| Rayleigh scattering | + | + | - | + | linear |
| Thermocouple | - | + | - | - | traditional |
| Lambda-sensor | + | - | - | - | traditional |
| FID | + | - | - | - | traditional |

Table 2.1.: Summary of properties for optical techniques which can measure fuel concentration (‘[fuel]’) or AFR, T and EGR fraction. ‘+’ denotes that the technique has demonstrated this property. ‘o’ signifies the case of partial success, or at least possibility. ‘-’ denotes that a technique does not have a particular attribute. Technique acronyms such as ‘LIEF’ will be expanded in sections 2.6 and 2.7.

application of a (coherent) laser pulse, it is possible that population can be transferred from one level to the other in a similarly coherent fashion - e.g. Rabi oscillation - as described by the optical Bloch equations. However, it is generally the case that under the conditions which are prevalent in this thesis, collisions and interactions with other molecules will rapidly remove the strict phase relationships that are required for this behaviour to be exhibited, and so the state of the molecule can be well-described by the use of rate equations.

Figure 2.2 is a graphical representation of the rate equation approach in the low intensity limit (where stimulated emission can be disregarded), with arrowed lines indicating the pathways by which population can pass around the system. B_{12} is the Einstein coefficient for absorption, and $\rho(\nu)$ is the energy density per unit bandwidth of exciting radiation so that the rate of absorption is the product $B_{12}\rho(\nu)$. A_{21} is the Einstein coefficient for spontaneous emission, and Q_{21} is a similarly-defined rate at which quenching processes occur. Values N_1 and N_2 are used to represent the number of




Figure 2.2.: A representation of a two level system. Transition rates include photon absorption $B_{12}\rho(\nu)$, spontaneous emission A_{21} producing a photon (wavy line), and a quenching process Q_{21} arising from molecular collisions.

molecules in each state. The application of a laser pulse to a collection of these systems transfers population into the upper state - and so causes N_2 to rise. The system will then seek to return to its low-energy configuration. There are two possible avenues - that of a return to the ground state via the process of spontaneous emission (at rate A_{21}), or the ‘quenching’ of the excited state energy via non-radiative de-excitation (at rate Q_{21}). The former process results in the emission of a photon at some wavelength which is characteristic of the energy spacing of the levels via $\Delta E = hc/\lambda$ (h is Planck’s constant, and ΔE the energy difference between the levels). Spontaneous lifetimes are typically on the order of 10ns. The quenching process does not emit a photon, it is often the result of a collision with another molecule, and provides a mechanism to convert the electronic energy which was stored in the system into translational kinetic energy - i.e. heat[†].

2.6.1. Absorption Spectroscopy

In absorption spectroscopy, one or many wavelengths are chosen with which to probe a sample. The Beer-Lambert law specifies the following relationship between the intensity incident on a sample, I_0 , and the intensity, I , after propagating a distance l

[†]In general, quenching may also release energy stored in vibrational and rotational states, though these levels are not represented in figure 2.2.

through the medium:

$$I(l, \lambda) = I_0 e^{-n\sigma(\lambda, T, p, \dots)l} \quad (2.7)$$

In this relationship, n is the number density of the absorbing species and σ is the absorption cross-section (which is proportional to B_{12}). The exponent in this equation, $n\sigma l$, is often termed the ‘optical thickness’ of a medium. Measuring the absolute intensity of a transmitted probe beam as the absorber concentration is increased will result in a reduction in signal, as straightforwardly-indicated by equation 2.7. It should be noted that σ is a function of the internal state of the molecules of the sample - and so its environment - and the wavelength, λ , at which the probe is applied. In particular, the distribution of population amongst levels inside a molecule is strongly affected by temperature, as given by the Boltzmann distribution:

$$N_2/N_1 = e^{-\frac{E_2 - E_1}{kT}} \quad (2.8)$$

Where N_2 and N_1 are the population densities in two levels of energy E_2 and E_1 , k is the Boltzmann constant, and T represents temperature in degrees Kelvin. By measuring the intensity transmitted through a sample as the temperature of the sample is changed, the relative strengths of a handful of probe wavelengths tuned to different transitions would be seen to vary relative to one another as the result of this thermal redistribution of population. The way this ‘fingerprint’ varies is therefore an indicator of temperature, T . Equation 2.7 which underlies all absorption spectroscopy indicates that these techniques all share a ‘line-of-sight’ integration characteristic, which means that the spatial selectivity of the measurement is defined by the cross-sectional shape and propagation vector of the probe beam. That is to say, the effect of an absorption spectroscopy measurement is to determine the mean value of a parameter along the interaction path length.

Absorption spectroscopy is a robust, extensible technique that has seen wide adop-

tion. Indeed, it is even sufficiently robust to find application in industrial sensors.

The first example is that of the measurement of atmospheric methane. Edwards and Burch⁸ used the $3.39\mu\text{m}$ line of a Helium-Neon laser to measure absorption in the C-H stretch band of the hydrocarbon. The authors note a relatively strong absorption coefficient, which is of use in the next example.

The second example system is that of the ‘Laser Spark Plug’ sold by LaVision GmbH⁹. Their system is a replacement for the standard-issue spark plug, and incorporates a small measurement ‘cage’ next to the spark plug ground electrode. This cage has open sides, a mirror at the base, and two optical fibres terminating on the upper face. Infra-red radiation is coupled from one fibre to the other after reflecting from the mirror. The size of the cage determines the path length along which absorption can occur. In the standard system, a tungsten filament lamp with a spectral filter is used as the source of probe radiation. Suitable choice of bandpass filters over two or three detectors allows different parts of the spectrum to be detected simultaneously, along with soot luminance. The system is said to be able to discern the presence of fuel, water and EGR. Advantages of this system include the high sampling frequency (up to 30kHz - more than sufficient for a sample per crank angle degree), the ability to deploy the system easily on any test engine without modification for optical access, and, as an industrial tool, a ‘plug-and-go’ nature. Disadvantages include the dependence of σ , on T , making quantitative measurements non-trivial, the constrained locality of the measurement region, and a worry that the cage next to the spark plug ground strap may cause interference with in-cylinder flows. Similar devices have also been reported, see, for example Tomita et al.¹⁰.

Instead of using a range of relatively spectrally wide filters to measure the overall absorption strength of a band of absorption lines, as just described, absorption spectroscopy can be applied to measuring the strength of particular transitions. As previously mentioned, measuring the ratio of absorption strengths for a handful of transitions allows the derivation of parameters related to the local environment of the

probed species (e.g. T and p). Such multi-line techniques can either be performed by using a narrow bandwidth broadly-tunable laser¹¹, a multiplexed set of narrowband lasers¹², or even leveraging the potential for multi-mode operation of a laser to create a ‘comb’ of frequencies to probe the absorption features of the sample¹³. In all cases, the techniques allow for the measurement of temperature and pressure - given sufficient spectral knowledge, and an understanding of the broadening mechanisms in effect (e.g. pressure broadening, Doppler broadening). These methods can also be used to derive species concentration.

A number of researchers have augmented absorption spectroscopy in order to attempt to perform spatially-resolved measurements. In the simplest approach, and for a sample of known symmetry - e.g. a cylindrically-symmetric flame - ‘raster-scanning’ across the flame is enough to reveal spatial information. That is to say, taking an absorption measurement along the y axis for each (x_i, z_i) path, where z represents the axis of cylindrical symmetry, and x lies perpendicular. This set of measurements is sufficient because the known symmetry allows the problem to be addressed like ‘unpeeling an onion’, where each successive step towards the centre-line reveals information about the next layer towards the axis of symmetry. Understandably, this approach compounds errors, as each successive layer is calculated. (This approach relies upon the Abel transform, and is mentioned in Chen and Goulard¹⁴.)

Naturally, a good many problems do not exhibit suitable symmetry - and engine research is no exception. A method for dealing with arbitrary distributions was presented by Chen and Goulard¹⁴, refined and applied by Santoro et al.¹⁵. These techniques employ a Fourier transform to derive spatial information. In the interests of not straying far off-course, the details are not considered here. Suffice to say that even modern approaches to this problem have trouble resolving detail. Terzija et al.¹⁶ found that objects smaller than $\sim 20\%$ of the diameter of the measurement volume were poorly reconstructed when using 27 measurement beams. In the context of the optical engine, that means that structures smaller than $\sim 18\text{mm}$ in diameter would remain unresolved.

2.6.2. Photon emission techniques

This group of linear optical techniques is related by its reliance on the measurement of emitted photons, generated by processes such as spontaneous emission (which was represented by the green line in figure 2.2). Each technique will now be considered.

Laser Induced Fluorescence

Laser Induced Fluorescence (LIF) is a conceptually simple technique, yet it offers a wealth of information about the probed volume. In a typical LIF experiment, some tracer species is doped into the region to be tested. Following exposure to a laser beam at a wavelength corresponding to an absorption feature in the tracer molecule, a number of fluorescence photons will be spontaneously emitted (i.e. the blue and green parts of figure 2.2). Spontaneous emission is an incoherent process - see figure 2.3 - illustrating that the directional laser beam is converted into fluorescence photons which are radiated into all 4π of solid angle.

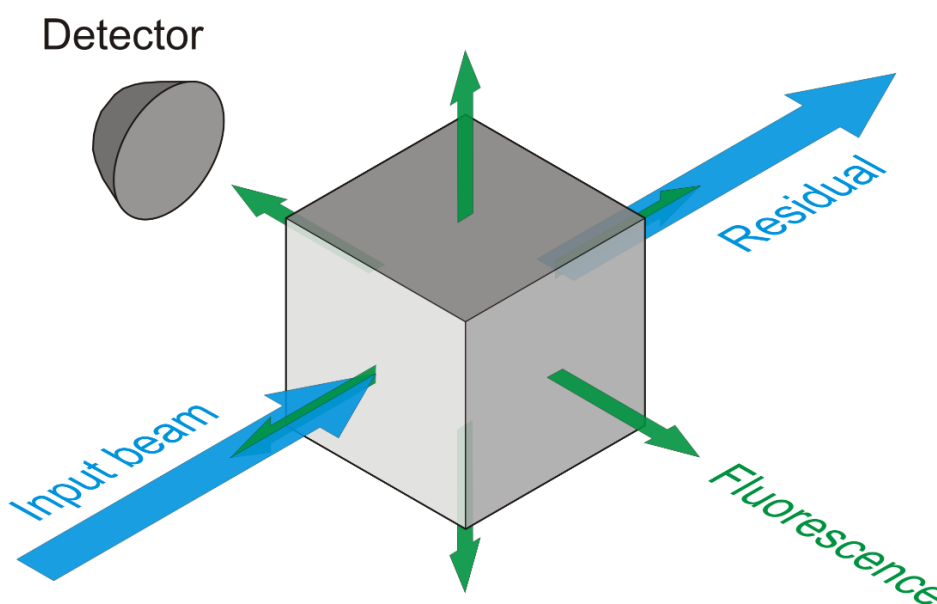


Figure 2.3.: A graphical representation of the incoherent LIF process.

Trivial extension of the LIF technique to shaping the incoming laser beam into a thin sheet allows many locations to be simultaneously excited. The use of a camera

with a suitable lens positioned at right angles to this sheet allows some fraction of the emitted photons to be collected, and a spatial distribution to be measured from the image. This is particularly valuable, because, as may be expected, LIF signal strength increases with increasing dopant concentration. In fact, in the low-signal regime where the incoming laser intensity does not saturate the transition inside the molecule, this relationship between fluorescence signal strength and number density *can* be linear.

Spatially-resolved measurements across a plane can provide great insights into a problem, and LIF-based techniques can be extended to measure other parameters of interest, such as temperature. LIF has its drawbacks, however. As with any technique that relies upon the measurement of intensity, beam steering and/or window fouling can cause incorrect interpretation of signal strength. Quantitative measurement of concentration has often been claimed, but in practice is a very challenging problem. A brief description of the requirements for quantitative PLIF will be given in this section, but the full discussion along with the system employed by QPLIF is presented in chapter 4. For more details of LIF theory, refer to section 3.2.

The value of the *in-situ* spatially-resolved distributions produced by PLIF has been widely recognised. Accordingly, many examples of the application of PLIF to IC-engine research are outlined in the recent review paper by Schulz and Sick¹⁷. (Note that in general, the emission from LIF is spectrally quite broadband - corresponding to a wide array of energy levels in a tracer molecule - but most imaging experiments are insensitive to this, simply integrating the spectrum.)

PLIF was one of the techniques chosen by the researchers at Mitsubishi in order to characterise and develop the first commercial gasoline direct injection (GDI) engine¹⁸. Though the distributions that they measured were only qualitative in nature, the information was nonetheless very valuable in correctly shaping the in-cylinder fuel distribution.

Quantitative measurement of fuel distributions with PLIF has seen a handful of

approaches. Reboux et al.¹⁹ is the first report on the use of the FARLIF (Fuel/Air Ratio by Laser Induced Fluorescence) technique, in which the strength of toluene fluorescence is taken to be linearly proportional to the fuel/air ratio, based on bench experiments with toluene vapour. This assumption has later been proven to be incorrect by Koban et al.²⁰ - in fact, the quantum efficiency of toluene (an aromatic hydrocarbon) varies strongly at engine-relevant temperatures and pressures. Indeed, it appears to be generally true that fluorescent tracers exhibit a strong, complex dependence on environmental conditions which is usually excitation-wavelength dependent^{21,22,23,24,25}. Quantitative measurement therefore relies upon accurate determination of the quantum efficiency of the fluorescent tracer being used.

Two options are available to quantify this dependence. In the first solution, researchers have chosen to model (or, at least parameterise) the quantum efficiency of the tracer under study as a function of T and p , based on measurements made in the controlled environment of a test cell. This approach then relies upon modelled in-cylinder T values to apply the quantum efficiency model to experimental conditions^{26,27,28}. In the second solution, researchers have employed *in-situ* calibration to characterise the behaviour of the tracer under study as a function of crank angle, during operation at controlled conditions. In the work presented by de Sercey et al.²⁹, a recirculation loop which feeds the exhaust from the cylinder back to the intake is used under motored conditions to produce a homogenised mixture in the cylinder. Measuring the fluorescence from this uniform mixture of air and fuel as the cycle progresses produces a map of fluorescence signal strength as a function of compression, which can later be used in reverse to convert inhomogeneous distributions to fuel concentrations. Smith and Sick³⁰ reports on a method which uses *in-situ* motored calibration with early injection for homogeneous fuel distributions to derive a model, which is applied to motored engine operation. Blotevogel et al.³¹ describes the quantitative measurement of fuel distributions for a hydrogen-powered IC engine. The approach employs a homogeniser that mixes helium and the liquid tracer triethylamine (TEA)

for collecting homogeneous *in-situ* calibration images. The authors indicate that TEA exhibits no pressure dependence, but calibration is required for temperature dependence. Owing to safety concerns, no fired operation is undertaken. The authors also note that TEA decomposes above $\sim 250^\circ\text{C}$, potentially limiting its application. Finally, Han and Steeper³² include quantitative measurements in a fired gasoline engine with two fuel components, employing a calibration scheme which is very briefly described, but involves calibration with both homogeneous mixtures and fired operation. This final example bears the most similarity with QPLIF, and will be re-visited in more detail in the final section (2.8) of this chapter.

PLIF has been extended by the use of multiple-colour excitation and/or simultaneous detection in multiple spectral bands to enable the measurement of temperature distributions^{33,34,35}. Exhaust gas residual (EGR) distributions have also been measured, by adding a fluorescent tracer to the EGR³⁶, or its complement - by adding fluorescing tracer to inducted air and fuel vapour³⁷. Dual-tracer PLIF has also been used to quantitatively image the distribution of oxygen in an engine, taking advantage of the high sensitivity of toluene fluorescence to O_2 concentration relative to 3-pentanone^{38,39}.

Amongst other applications, of relevance in an engine is that LIF has also been used to diagnose the thickness of a film of liquid deposited on a surface - e.g. to characterise the impingement pattern of a fuel spray onto the piston crown in a direct-injection engine⁴⁰.

Laser Induced Exciplex Fluorescence

Laser Induced Exciplex Fluorescence, LIEF, is an extension of the LIF technique to the use of a tracer system which exhibits different behaviour depending upon the density of the tracer. As originally demonstrated by Melton and Verdick⁴¹ and more recently Wieske et al.⁴², this system makes it possible to study the concentration of both liquid and vapour phases simultaneously. With the industry moving towards universal deployment of direct injection, the evolution of the spray is becoming an increasingly

important topic. However, this thesis is concerned with measurements in the vapour phase only, to judge mixture preparation prior to ignition, so an exciplex tracer system is an unnecessary complication.

Phosphorescent tracers

These techniques are strongly related to the LIF techniques explained above. However, instead of fluorescence photons from a relatively short-lived ($\sim 10\text{ns}$) upper state being collected, phosphorescence photons produced on much longer timescales are collected.




Figure 2.4.: Illustration of the phosphorescence process. Heavy lines denote the lowest energy level of a manifold. Thin lines are vibrational sub-levels.

Figure 2.4 describes the energy levels in a more complex example system than that of figure 2.2. Three electronic configurations are illustrated (the three heavy lines), each of which has a family of sub-levels, corresponding to extra quanta of energy being stored in vibrational and rotational modes of the molecule. The lowest energy configuration - the ground state - is referred to as 'S₀'. Absorption of a photon promotes population from the ground state to the first singlet excited state 'S₁'. The non-radiative process known as 'inter-system crossing' (ISC) allows population to transfer from S₁ to T₁. T₁ is the first excited triplet state, which lies at slightly lower energy than S₁. The distinction between these two states is determined by the alignment of the electron spins in the wavefunction of the system. If the total spin angular momentum quantum

number, s , is one, then the system is in the triplet state (three spin configurations have $s = 1$). If $s = 0$, then the system is in the singlet state.

Following absorption into S_1 and conversion to T_1 via ISC, the molecule will rapidly lose rotational and vibrational energy, eventually reaching the lowest level of the manifold in T_1 . The selection rules for electric dipole radiation prevent spontaneous emission from the triplet state, requiring the use of a higher moment - e.g. electric quadrupole - to decay back to the ground state singlet radiatively. This process is typically about 1000 times slower than electric dipole radiation, so the triplet state can produce photons over a relatively long timescale by phosphorescence.

Temperature measurements within engines have been conducted by use of phosphorescent tracer compounds. Hasegawa et al.⁴³ demonstrates the use of Dy:YAG seeded into the in-cylinder gas to measure temperature. In this case, the ratio of two spectral regions is used to monitor temperature, with a precision of $< 5\%$.

Bruebach et al.⁴⁴ presents a scheme which determines surface temperature by monitoring the lifetime of the phosphorescence emitted from a phosphor ($Mg_4FGeO_6:Mn$) which is ‘painted’ onto the surface of interest, and excited by the third harmonic of an Nd:YAG laser. This technique is shown to resolve temperature fluctuations on the order of 3K.

Chemiluminescence

A final method for the production of photons will now be discussed, which is of relevance in IC engine research. Where LIF and phosphorescence-based techniques rely on absorption from a pump laser, the excited state population is produced in chemiluminescence by a chemical reaction[†]. For example, following/during the chemical reactions which drive combustion, excited OH, CH and C_2 molecules are formed. (The excited state denoted by the asterisk, e.g. OH^* .) In common with the other techniques described in this section, photons will be emitted from these excited molecules at wave-

[†]It might be reasonable, therefore, to call this ‘CIF’ for Chemically Induced Fluorescence

lengths which characterise the spacing of the energy levels in the molecule. Emission from OH* is detected at 308nm, CH* at 430nm, and C₂* at 516nm⁴⁵.

Amongst others, Panoutsos et al.⁴⁶ have developed a technique which relies upon measuring the ratio of OH* and CH* chemiluminescence as a diagnostic of equivalence ratio. Chemiluminescence techniques are often linked to low signal strengths, so in this study, a relatively large-area Cassegrain-telescope-style optic is employed to maximise the photon collection efficiency. Since chemiluminescence measurements rely upon flame chemistry to create the excited state population, these techniques are of limited value in measuring pre-ignition fuel distributions.

2.6.3. Scattering Techniques

In this subsection, the two relevant scattering techniques will be reviewed.

Rayleigh Scattering

Rayleigh scattering is the process whereby light is elastically (i.e no energy transfer) scattered from molecules. The cross-section for this process increases as λ^{-4} , so that light in the UV has a significantly higher chance of being elastically-scattered[†]. Rayleigh scattering is dominant when the molecular diameter, d , is much smaller than the wavelength of the radiation falling upon it. The amount of scattered light is linearly dependent upon the number density of the scattering species, as follows⁴⁷:

$$I_R = CI_0N\sum_{i=1}^jX_i\sigma_{r,i} \quad (2.9)$$

Where C is a calibration constant to account for efficiency of collection optics, I_0 is the incident intensity, N is the number density of the sample, X_i is the molar fraction of each species present and $\sigma_{r,i}$ is the Rayleigh cross section of each of these species.

[†]This wavelength dependence is the reason the sky appears to be blue - without scattering, the sky would be dark. With scattering, some small wavelength-dependent fraction of the Sun's light which would otherwise pass by through the atmosphere is redirected towards the planet's surface.

This highlights that Rayleigh scattering does not produce a species-specific signal, but the strength of the scattered signal can be used in combustion diagnostics. Zhao and Hiroyasu⁴⁷ considers the application of Rayleigh scattering to measurements of mole fraction, density and temperature. The suitability of Rayleigh scattering to making these measurements depends on the ability to make simplifying assumptions about the system under investigation - e.g. in the intake stroke of an IC engine assume there is small variation in T and that a pressure transducer is present, in order to determine X_f , the mole fraction of fuel.

In all cases, the application of Rayleigh scattering to combustion diagnostics is highly dependent on the minimisation of interference from background sources. The presence of even a small amount of particulate matter is likely to swamp the Rayleigh signal, as would any contributions from other sources such as stray reflections.

Despite these drawbacks, Espey et al.⁴⁸ report on the application of Rayleigh scattering to the quantitative measurement of fuel concentration in a diesel spray. Certain modelling assumptions regarding the mixing process are required to analyse the recorded data, and measurements are performed only in the vapour-phase region of the jet, where interference resulting from fuel droplet Mie scattering can be avoided.

Raman Scattering

Spontaneous Raman scattering is an inelastic scattering process, where there is some energy transfer between the molecule and the photon. In the case where energy is transferred from the photon to the molecule, the emitted photon is at a longer wavelength, and is said to be Stokes shifted. (In the opposite case, where the photon gains energy, it is anti-Stokes shifted.) The creation of Stokes and anti-Stokes shifted photons is illustrated by considering the interaction of an incident field $E = E_0 \cos \omega_0 t$ with a molecule. This E field induces a polarisation, P :

$$P = \alpha \epsilon_0 E \tag{2.10}$$

Where α is the polarisability of the molecule, and ϵ_0 is the permittivity of free space. α can be expanded in first order as:

$$\alpha = \alpha_0 + \frac{\partial\alpha}{\partial Q}Q \quad (2.11)$$

Where Q is some inter-nuclear coordinate, and α_0 is the value of the polarisability when at equilibrium. Suppose a natural mode of vibration for the molecule causes $Q = Q_0 \cos \omega_v t$. This will give rise to the following polarisation:

$$P = \left(\alpha_0 + \frac{\partial\alpha}{\partial Q}Q_0 \cos \omega_v t \right) \times \epsilon_0 E_0 \cos \omega_o t \quad (2.12)$$

$$= \epsilon_0 E_0 \left(\alpha_0 \cos \omega_o t + \frac{\partial\alpha}{\partial Q}Q_0 \cos \omega_v t \cos \omega_o t \right) \quad (2.13)$$

$$= \epsilon_0 E_0 \left(\alpha_0 \cos \omega_o t + \frac{\partial\alpha}{\partial Q}Q_0 \frac{\cos(\omega_o - \omega_v)t + \cos(\omega_o + \omega_v)t}{2} \right) \quad (2.14)$$

Clearly, the induced polarisation now contains components shifted up and down in frequency by ω_v , the vibrational frequency of the molecule, leading to the formation of the first-order Stokes and anti-Stokes frequencies. By allowing higher terms in the expansion of equation 2.11, direct coupling to higher orders can be achieved, but each successive order of the expansion is much weaker (i.e. $\frac{\partial\alpha}{\partial Q} \gg \frac{\partial^2\alpha}{\partial Q^2}$), so higher-order Raman shifts are generally not seen in spontaneous Raman scattering.

This is a non-resonant process which scales again as λ^{-4} . The cross section for spontaneous Raman scattering is typically weaker than Rayleigh scattering by a factor of $\times 10^3$, however⁴⁹. In contrast to Rayleigh scattering, Raman scattering does provide species-specific signals. Every molecule has a different vibrational (and rotational[†]) character, and the transitions between these levels produces a ‘fingerprint’ which characterises the molecule.

The signal from a Raman scattering experiment must be detected with some dis-

[†]Rotational shifts are very small, and thus challenging to separate spectrally from the excitation wavelength

persive element - e.g. a spectrometer. As previously mentioned, each molecule is associated with one or more characteristic vibrational modes. The Raman shift of the hydrogen molecule, H_2 , is known to be $4155\text{cm}^{-1\dagger}$. The effect of Raman scattering would be to produce a spectrum consisting of a series of peaks centred at the excitation wavelength and spaced by this shift, with each peak having a strength which is representative of the population density in that particular vibrational energy level (vibrational energy levels exist as a ladder of sub-levels, as illustrated in figure 2.4). Since energy levels are in general populated according to equation 2.8, the relative strengths of these spectral lines delivers information about the temperature of the molecule.

An impressive facility has been assembled by the researchers at Sandia National Labs, which combines Rayleigh and Raman Scattering with LIF and has been reported on a number of occasions, e.g. Barlow and Carter⁵⁰, Masri et al.⁵¹, Barlow et al.⁵² and Cabra et al.⁵³. This apparatus can quantify the concentration of N_2 , O_2 , H_2O , CH_4 and H_2 by Raman and Rayleigh scattering, together with OH , CO and NO by LIF. This system is understandably complex, and has not been applied to an IC engine to the best of the author's knowledge - one major reason for this is that high laser intensity (in this case doubled Nd:YAG pulses of $\sim 2\text{J}$) is required to obtain sufficient signal from spontaneous Raman scattering, which is generally incompatible with the use of windows due to damage threshold limitations.

2.7. Non-linear optical

Recall from equation 2.6 that for sufficiently high field strength, non-linear relationships between the induced polarisation and the applied field become observable. For a medium with an inversion centre (i.e. an origin can be chosen such that the properties of an arbitrary location (x, y, z) cannot be distinguished from those at $(-x, -y, -z)$), it can be shown⁵⁴ that the second order susceptibility is zero. This is the case in a

[†]wavenumbers, $\bar{\nu}$, measured in cm^{-1} and given by $1/\lambda$ are converted to energy by $E = hc\bar{\nu}$.

gas, and so the lowest-order non-linear effects that are observable involve $\chi^{(3)}$ - the so-called four-wave mixing processes. This section will include discussion of two such processes: DFWM and CARS.

Another relevant technique is Polarisation Spectroscopy (PS). This technique does not see any application in IC engines, however, owing to small (and variable) stress-induced birefringence in the window material causing additional unwanted polarisation rotation. This technique will therefore not be considered.

2.7.1. Coherent Anti-Stokes Raman Scattering (CARS)

CARS is a four-wave mixing process which in its most commonly-used form employs three pump beams of two different colours to produce a fourth anti-Stokes shifted output beam. Algebraically, the frequencies of the lasers are related by $2\omega_1 - \omega_2 = \omega_3$. Figure 2.5 indicates that when the frequency separation between ω_1 (pump) and ω_2 (Stokes) is chosen to coincide with a Raman resonance, the anti-Stokes shifted beam at ω_3 is created. This is a non-resonant process, and so the diagram in figure 2.5 makes use of virtual levels in order to enable the transitions.




Figure 2.5.: Diagram of energy levels involved in CARS process. Levels labelled with n values correspond to electronic resonances. Increasing values of v denote increasing quanta of stored vibrational energy.

The spectrum produced by CARS following scanning of the Stokes wavelength, or dispersed detection of a broadband Stokes laser is not dissimilar to that produced by spontaneous Raman scattering. However, the CARS signal strength is related to the square of the susceptibility, so interference effects can alter the detail of the recorded structure. Providing such effects can be sufficiently-well modelled, CARS can be used for temperature measurements by fitting a model to the measured spectrum. In certain cases, spectral information can also provide concentration[†], or, more generally, signal strength can be used to infer number density.

In order to address multiple species simultaneously, two or three spectrally narrow ('narrowband') Stokes lasers may be simultaneously employed, or alternatively, a Stokes beam that is itself sufficiently broadband to encompass the range of Raman shifts can be used. In this context, Raman shifts for relevant molecules vary in the range of 500-4000cm⁻¹, and typical broadband laser sources may be on the order of 200-300cm⁻¹ - enabling only a small number of molecules to be simultaneously probed[‡].

Non-linear, frequency-mixing processes are generally associated with phase-matching conditions, which ensure the constructive generation of the output beam as the pump beams propagate through the interaction region. In the case of a gas, which exhibits very little to no dispersion, the phase-matching condition is simply the momentum conservation equation:

$$\underline{k}_a + \underline{k}_b = \underline{k}_c + \underline{k}_d \quad (2.15)$$

Where the underline indicates vector properties, $k_a = k_b$, and $k = 2\pi/\lambda = \omega/c$. Referring to figure 2.5, 'a' and 'b' are associated with the upwards green arrows, 'c' is associated with the orange Stokes arrow, and 'd' is represented by the blue downward anti-Stokes arrow. A variety of solutions are available to the phase-matching condition of equation 2.15, including a trivial collinear one. The collinear geometry results in a

[†]By means of the interaction between the resonant CARS process described here, and the non-resonant background processes. For more detail, see Eckbreth⁴⁹.

[‡]Rotational-CARS produces shifts up to $\sim 400\text{cm}^{-1}$, enabling better multi-species detection - providing the nearby excitation wavelength can be adequately suppressed

poorly-defined interaction region (much like Absorption Spectroscopy, in fact), and for this reason the ‘BOXCARS’ configurations⁵⁵ illustrated in figure 2.6 are preferred. The ‘folded BOXCARS’ configuration affords extra separation between the signal beam and the pump beams, improving the discrimination against background in the detection process.

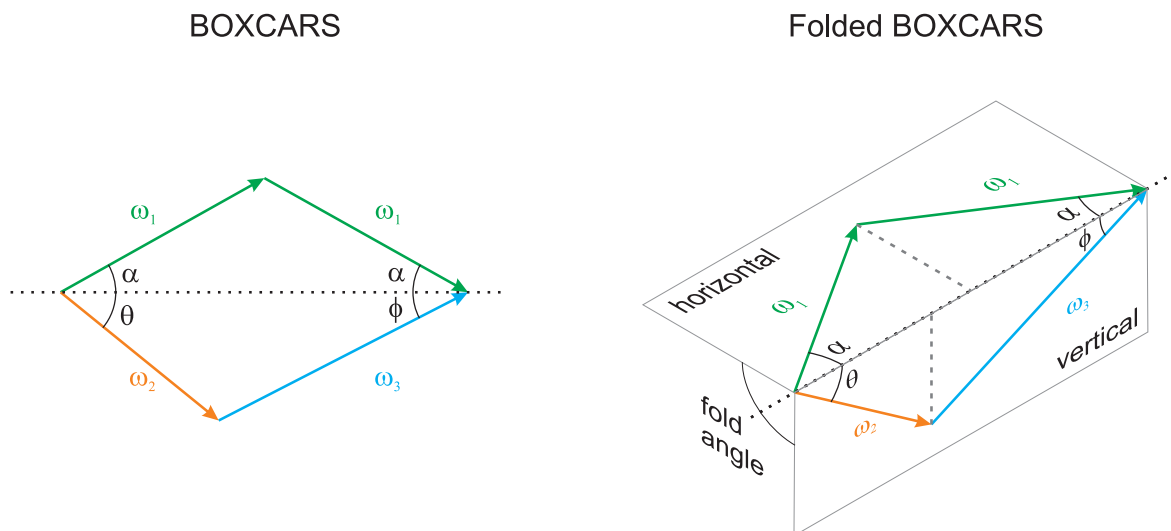


Figure 2.6.: The phase matching condition requires particular interaction geometries. The ‘BOXCARS’ and ‘folded BOXCARS’ arrangements yield the most tightly localised probe volume, and are therefore popular. Colours are consistent with those of figure 2.5.

CARS is a very widely adopted technique, so complete coverage here is impossible (a full review of CARS-based techniques can be found in Eckbreth⁴⁹). Some relevant applications taken from the literature will now be briefly covered:

Multi-species CARS was demonstrated by Eckbreth and Anderson⁵⁶, employing two broadband lasers and a third narrowband laser. Species that could be probed with this arrangement included O_2 , N_2 , CO_2 , NO , H_2O and the C-H stretch vibrational mode (i.e. hydrocarbons).

Temperature measurements based on N_2 CARS or O_2 CARS became popular, but the limited computing power at the time led to similarly limited accuracy. Dreier et al.⁵⁷ reported the use of a simplified fitting system which allowed for timely fitting

of spectra recorded in a high temperature furnace, at up to 2000K. The accuracy of the fit was estimated at between 40-80K, and verified against a thermocouple which had corrections for gas flow, convection, heat conduction and radiation applied to its readings. At this time, the ‘CARSFIT’ computer code developed and maintained by SANDIA National Labs became widely adopted⁵⁸.

Rotational CARS also quickly became a strong research interest, due to the improved multi-species probing character. The folded BOXCARS geometry is invaluable for rotational-CARS in isolating the signal beam from the much stronger pump and probe lasers. Alden et al.⁵⁹ reports on a comparison of vibrational and rotational CARS for temperature measurement. It is found that rotational-CARS is better around room temperature, and vibrational CARS is the method of choice around flame temperatures, owing to population distribution at the respective temperatures. Temperatures were said to be accurate to about 10%, and a single spectrum could be computer generated in $\sim 60\text{s}$ [†]. Further comparisons between rotational and vibrational CARS are reported, for example, by Seeger and Leipertz⁶⁰, where accuracy is now improved to between 10-40K, and precision estimated from a population standard deviation at 10-60K, depending on technique and temperature range[‡]. Measurement of O₂ concentration is also reported in this paper, and is said to be accurate to within 1%.

CARS has been demonstrated in sooting flames⁶¹, where other techniques like spontaneous Raman scattering suffer from contamination due to the soot luminosity. Beyrau et al.⁶² reports on the application of CARS to the measurement of temperature and air/fuel ratio in a sooting ethene-air flame. Temperature is said to be accurate to 2-5%, and the concentrations of oxygen and ethene are evaluated to a standard deviation of 6.8% and 5.5%, respectively.

In a recent publication, Weikl et al.⁶³ report on the application of rotational-CARS

[†]The paper says the PC used was a bit faster than a DEC VAX 8200, which had dual 5MHz CPUs

[‡]By this time, the Intel Pentium would have been widely available in desktop PCs, running at 100MHz - which probably explains why the paper makes no reference to a lengthly wait for a modelled spectrum!

to the measurement of temperature and EGR fraction in an HCCI engine. The authors made some modifications to the engine for line-of-sight optical access, and used a mixture of iso-octane and n-heptane to minimise the formation of combustion deposits. Since rotational-CARS measures many species simultaneously, both EGR fraction (reported, and calculated by comparison of O₂ and N₂ concentrations) and air/fuel ratio (not reported, instead the authors used a lambda sensor to determine oxygen levels in EGR gas) are measurable.

In summary, the CARS technique offers a number of advantages:

1. A signal *beam* (in common with other non-linear techniques), aiding background rejection. This means it can be applied even in sooting environments, unlike scattering techniques
2. Detection at shorter wavelengths than the pump beams, which helps to avoid possible LIF interference
3. Measurement of temperature (and concentration) based on spectral shape

CARS also suffers from some disadvantages:

1. Typically limited to major species, in common with spontaneous Raman scattering
2. Temperature and concentration derivation depends on accurate modelling of CARS spectra, which can be complex and time-consuming
3. While imaging is possible, phase-matching makes it impractical to achieve

Note that amongst the many variations of the CARS technique is ERE-CARS, or Electronic Resonant Enhancement-CARS, in which electronic resonance is used to boost the CARS signal strength⁶⁴. The implementation of this technique is challenging, because in order to maximise the sensitivity of the technique, all three excitation laser

beams (probably at UV-wavelengths) need to be tuned into resonance. A simpler technique, DFWM, achieves similar sensitivity and will be described in the next section.

2.7.2. Degenerate Four Wave Mixing (DFWM)

DFWM is, as the title suggests, a four-wave mixing process, which employs three degenerate (i.e. equal frequency) pump beams. When the frequency of these beams is tuned to match an electronic molecular resonance, the fields are coupled via the third order susceptibility, and a fourth signal beam is generated at the same frequency. A scanned spectrum can therefore be used as with CARS or Raman scattering for the determination of temperature from relative peak heights/spectral shape and concentration from signal magnitude.

The DFWM technique is, as a result of the frequency-degenerate pump beams, much easier to phase-match than CARS. Although at first the spectral degeneracy of the signal may seem a drawback, background interference is rarely an issue thanks to the highly angularly-selective detection possible with a signal beam. Owing to the resonant nature of DFWM, the signal is generally some orders of magnitude stronger than the equivalent CARS signal. This allows detection of minor combustion species such as OH, CO and NO down to ppm levels. The creation of a DFWM signal is also free from the quenching interferences of LIF, and thanks to the relaxed phase-matching requirements, DFWM can be used in imaging applications.

Since matching an electronic resonance in the probed molecule is required to generate a DFWM signal, the laser source must either lie in the ultra-violet or infra-red parts of the spectrum, and yet retain tunability across tens of wavenumbers to resolve enough spectral lines. Sum- and difference-frequency generation and harmonic generation are therefore likely to be required, and tuning the non-linear crystals required to perform this task is likely to impose a limit.

Full reviews of DFWM and its applications can be found in Farrow and Rakestraw⁶⁵

and Kohse-Hoeinghaus and Jeffries⁶⁶. Some example applications of DFWM from the literature include initial demonstrations of multiplex DFWM measurements to a sodium-seeded flame⁶⁷, then adapted to OH⁶⁸ enabling the measurement of temperature⁶⁹. 2D imaging using DFWM was first demonstrated by Ewart et al.⁷⁰, again in a sodium-seeded flame. Later work combined these concepts to produce spatially-resolved temperature measurement from OH DFWM⁷¹. Reichardt et al.⁷² is a more recent report considering the application of DFWM to the measurement of quantitative OH concentrations in flames of varying stoichiometry. Grant et al.⁷³ presents the first application of DFWM to measurement of NO concentration in a firing SI engine.

2.7.3. Laser Induced Thermal Grating Spectroscopy (LITGS)

This section will present only a brief overview of the details of the LITGS technique. For an in-depth discussion, refer to section 3.3. A LITGS signal is formed by the crossing of two pulsed pump beams at some small intersection angle. In the crossing region, a spatial density modulation is generated in the gas by the heating resulting from quenching processes acting on molecules which have been driven to the excited state. By Bragg-scattering a probe beam from this three-dimensional density grating, a signal similar to that shown in figure 2.7 is produced. The formation and readout geometry is typically very similar to a BOXCAR arrangement[†], as shown in figure 2.6.

The temporal evolution of the shape of such a LITGS signal relates to the fluid dynamic properties of the medium that the grating is written into. Again, no justification is given here (instead, see section 3.3), but the high-frequency oscillation provides information which leads to temperature, the decay lifetime leads to pressure, and the overall signal strength can be used as a measure of concentration.

At high laser intensities, another process may become visible. This is the process

[†]Which is why DFWM experiments sometimes see LITGS signal interference, particularly at higher pressures - as reported in Danehy et al.⁷⁴




Figure 2.7.: A simulated LITGS signal, showing high frequency oscillation and exponential decay characteristics.

of electrostriction, whereby molecules are moved towards maxima of the electric field by application of a force to the molecule's (induced) dipole moment. This process is not dependent on absorption in an electronic resonance in the target molecules, so any species can be affected. However, electrostrictive gratings are generally much weaker than thermal gratings.

Electrostrictive signals can often be distinguished by their prompt appearance, in contrast to thermal gratings, which are only generated following the transfer of stored electronic energy to heat (which may take ~ 10 ns). Electrostrictive gratings also oscillate at twice the frequency of their thermal counterparts since they only have an acoustic component.

The laser intensity employed in this work was found to be insufficient to produce a measurable electrostrictive grating, and so their contribution will now be ignored. It should be noted that electrostrictive grating signals encode much the same information

as thermal gratings, allowing, for example, the retrieval of temperature, as described by Stampanoni-Panariello et al.⁷⁵ and Hubschmid et al.⁷⁶.

A recent review paper by Stampanoni-Panariello et al.⁷⁷ discusses experiments conducted at the Paul-Scherrer Institute in Switzerland. Some specific examples of LITGS that can be found in the literature will now be reviewed:

The LITGS technique was first demonstrated in the gas phase by Cummings⁷⁸. Shortly afterwards, Paul et al.⁷⁹ published a paper describing the theory underlying the formation and evolution of a LITGS signal.

Hart et al.⁸⁰ reported on the use of LITGS in measuring temperature in atmospheric pressure air between 300-650K, finding accuracy and precision of 1%, even over a period of days. Other authors applied the LITGS technique to the measurement of T and p , including the extension of the technique to the measurement of these parameters along a line (rather than simply in the small crossing volume) by use of sheet-shaped laser beams⁸¹.

The investigation of composition by reference to the measurable (though small) change in oscillation frequency has been reported by Hemmerling et al.⁸². Overall signal strength has also been used to measure species concentration⁸³. The former approach has a sensitivity which is dependent on the difference in molar mass between the species of interest and the bulk gas. The latter approach, having a reliance upon intensity measurements, suffers from the same drawbacks as LIF.

LITGS has been used to measure velocity. The detail will not be described in this thesis, but the following references describe the method: Walker et al.⁸⁴, Hart et al.⁸⁵ and Kozlov⁸⁶.

Recently, Seeger et al.⁸⁷ reports on a method for using LITGS to measure equivalence ratio. (Also, the same technique is compared with Raman scattering in Kiefer et al.⁸⁸.) This technique employs both electrostriction (non-resonant laser induced grating formation, see section 3.3) and on-resonant LITGS. The technique is calibrated

in a binary mixture of nitrogen and propane (C_3H_8), and then applied to injection of propane into a test cell containing air. Since the propane is weakly absorbing at the pump wavelength, a small fraction of the incident energy will be converted into a thermal grating. The strength of the thermal grating is proportional to the partial pressure of propane present in the crossing region. The electrostrictive signal strength is proportional to total pressure, and so the ratio describes equivalence ratio.

The authors found that the LITGS signal became overwhelming beyond a 1.6% propane concentration, preventing the measurement of the electrostrictive peak height. Therefore, above this limit the characterisation of propane concentration by its effect on the oscillation frequency is used to derive concentration information. As the closest competitor to the FARLIG technique, this work will be considered more fully in the next section.

2.8. Technique selection and closest rivals

In this section, the choice of techniques developed in this thesis are justified. Each technique is then considered against the most similar work previously reported in the literature.

The choice of a technique for making imaging-mode fuel concentration measurements in the optical engine was relatively straightforward. While Raman and Rayleigh scattering do offer such spatially-resolved measurements across a plane, the laser power required to do so is not friendly to the windows. Attempting a non-linear method like DFWM was also ruled out on the grounds of the very restricted optical access. The most interesting region in the cylinder - near the spark plug - is only accessible throughout the cycle via the window in the side of the cylinder head, and through the piston crown window (see section 4.3.2 for details). This clearly motivates the adoption of a PLIF-based technique to measure spatial fuel distributions.

Although not applicable throughout the engine cycle with the current optical access,

the further development of a LITGS-based technique in a gas cell was motivated by its robust temperature measuring ability, which has both higher accuracy[†] and precision than CARS, and is also much more readily derived from the temporal signature of the signal. Since the LITGS technique makes no use of spectral information, the experimental configuration is also simpler and cheaper, with no bulky spectrometer or expensive CCD-array required to image the spectrum. In terms of alignment and signal separation, the two techniques are comparable.

2.8.1. QPLIF

The most direct competitor for the Quantitative-PLIF (QPLIF) technique has been identified as that described by Han and Steeper³². Their paper will now be reviewed in detail:

Han and Steeper³² discusses the application of PLIF to the quantitative measurement of equivalence ratio in an optical research engine. The quantum efficiency of the tracer employed in the experiments is calibrated *in-situ* by means of a premix generator which produces a homogeneous mixture of known composition. This mixture is then imaged at 30 BTDC, and is later used to convert LIF intensity to equivalence ratio, and deal with any laser profile inhomogeneity. Shot-to-shot energy fluctuations are also corrected by monitoring the beam with a joule meter.

Two fuel/tracer combinations are employed in the system described by Han and Steeper³². Both are tested by measuring the fluorescence of the vapour evolved from the evaporating mixture, and comparing the signal strength against the lost mass. The tracer composition is adjusted if required to obtain the best agreement in this test. The first fuel/tracer combination simulates a two-component fuel, based on iso-octane and cyclo-pentane, with a tracer comprising a mixture of 2-butanone, 3-pentanone and 3-hexanone. The second fuel/tracer combination is a single component fuel, and uses cyclohexane and 3-pentanone as base fuel and tracer, respectively. Good agreement is

[†]Given sufficient knowledge of gas composition.

observed in the evaporation experiment in both cases.

The engine investigated by Han and Steeper³² employs a centrally-mounted swirl injector and two spark plugs, one mounted either side of the top of the cylinder. The cylinder itself is built from a water-cooled metal liner and a 1-inch high fused silica annulus. The laser sheet which is ~ 0.7 mm thick enters and exits through the annulus, approximately 11mm below the spark plugs. The cylinder head and cylinder liner are held at 50°C, and the engine is run in a skipped burst-fired mode, such that a PLIF image is recorded on the last cycle of a sequence of 10 fired cycles. SOI is either 280 BTDC or 180 BTDC, and all images are recorded at 30 BTDC, which is 1-5 CAD before ignition.

PLIF images recorded using DI are calibrated with the reference images already described, and then the pixels are binned into 4x4 super-pixels, presumably to improve noise characteristics without losing much spatial information. Each super-pixel is said to have $\pm 20\%$ error in its reported equivalence ratio, with the authors taking care to attempt to quantify the potential impact of temperature fluctuations by measuring fluorescence strength at a range of CAD.

The authors report that the single component fuel produces a mixture of similar stratification to the two-component fuel when SOI = 280 BTDC, but that the agreement is not nearly so good at SOI = 180 BTDC. They conclude that single component fuels - even those chosen with a similar volatility (as iso-octane is often chosen to model gasoline) - may not appropriately model mixing or combustion of a multi-component fuel.

Some areas can be selected where the technique described by Han and Steeper³² could be improved. These include:

1. Pump gasoline has more than just two fractions - can this be accounted for?
2. Skipped engine operation introduces extra undesirable variability

3. Improve the calibration system, if possible
4. The reported measurements collect data in a plane which is 11mm distant from the spark plugs
5. Tracing fuel around a cycle to monitor distribution evolution could yield further insight
6. Different fuel fractions may exhibit different behaviour which should be captured
7. Cycle-resolved fuel distribution measurements could allow the relationships between subsequent cycles to be investigated

Chapter 6 will revisit this topic, to see whether the QPLIF technique can address some of these points.

2.8.2. FARLIG

The most direct competition to the FARLIG (Fuel/Air Ratio by Laser Induced thermal Gratings) technique is that detailed by Seeger et al.⁸⁷, which will now be described:

In Seeger et al.⁸⁷, two forms of grating are employed to derive the equivalence ratio of a propane/air mixture in a test chamber. In this technique, the gratings are formed by the fundamental wavelength of an Nd:YAG laser, at 1.064 μm . Propane is weakly absorbing at this wavelength, allowing the production of a similarly weak thermal grating. The pump laser beams are also focussed down into the crossing region, and as a result of the high intensity that this produces, an electrostrictive grating is formed. The cw argon-ion laser used as a probe beam is scattered by both gratings, producing a prompt grating formed by electrostriction just before the thermal grating. Since the magnitude of the density modulation produced by electrostriction is related to the total gas density (little species sensitivity), the electrostrictive peak scales with total pressure. The strength of the contribution from the thermal grating depends on the

concentration of propane. Taking the ratio of the initial peaks of the two gratings therefore derives a value proportional to equivalence ratio.

The authors report that this ratio is largely independent of pressure, at least in the measured range of 1-4 bar. Calibration tests conducted in a known concentration of propane mixed with nitrogen at 298K demonstrate that the thermal grating becomes dominant above a concentration of 1.6% propane, motivating a second approach to concentration measurement. The second approach that is used in the paper is to derive propane concentration from the oscillation period of the grating signal. This can be achieved because the introduction of more propane increases the mean molecular mass of the gas, which in turn reduces the sound speed and increases the inter-fringe transit time.

The uncertainty in concentration using the ratio method was evaluated by using a population of 100 single shot measurements, and found to be $\sim 14\%$ at 0.8% propane. The second approach based on oscillation period is found to be less repeatable, with a lower detection limit at 100% uncertainty (i.e. a measurement of $x \pm x$) estimated at 3.2% propane. Typical uncertainty, measured with a further 100 shots at 6% propane was evaluated to be 39%. It is suggested that the intermediate range of concentration - between 1.6% and 3.2% - requires a combination of the two approaches to yield a solution, employing the concentration mole fraction, $x\%$, as a fitting variable in a model including both electrostriction and thermal grating production.

Calibration complete, the technique was applied to equivalence ratio measurement in a test chamber fitted with a direct-injection system from a gasoline engine. Propane at 13 bar and 330K was admitted into the chamber which was itself filled with 3 bar of air and held at 345K. (Change in temperature between the calibration and measurement when employing the ratio method is estimated to cause less than 8% error, and changing from nitrogen to air a further $<3\%$ error.) By varying the delay between the injection event and the measurement, the arrival and departure of the propane jet could be measured. The jet itself introduced further uncertainty which

was estimated to be on the order of 28% shot-to-shot.

It can be easily calculated from equation 2.2 that a stoichiometric mixture of propane and air is achieved at 4% mole fraction propane. Propane is flammable between 2.12% and 9.35% in air⁸⁹, so the ratio mechanism for measurement is actually less useful than the period approach for this particular fuel. The reliance upon measurement of the period to determine fuel concentration is unfortunate, since this precludes using this feature of the signal for temperature measurement. Upon examination, the example signal shown in this paper from a 6% propane measurement also suffers from digitisation artifacts. This may be the result of a weak signal, which would likely be due to the use of a 70mW argon-ion laser as a grating probe.

The FARLIG technique should aim to improve on the $\sim 40\%$ measurement accuracy of this approach, whilst retaining the ability to measure temperature. Improved signal strength may aid in the reduction of uncertainty. One improvement in the FARLIG technique is that the beams are not focussed at the crossing point, which reduces the chance of gas breakdown and allows more fringes to be formed. Following the discussion of the FARLIG technique in chapter 5, this topic will be revisited in the conclusions to determine what other advances the FARLIG technique has introduced.

3. Theory of LIF, LITGS and Tracer photophysics

3.1. Chapter overview

This chapter will serve as the basis for further discussion in chapters 4 and 5 regarding the adaptation of the fundamental techniques into QPLIF and FARLIG. Section 3.2 deals with LIF, and section 3.3 covers LITGS. The chapter then closes in section 3.4 by considering the interactions of the tracers toluene and acetone with the pump laser and the environment.

3.2. LIF

In this section, the theory of LIF signal production will be reviewed, followed by a detailed discussion of quenching. Consideration is then given to common approaches for making LIF measurements, before finally discussing the inherent advantages and disadvantages of LIF-based techniques.

3.2.1. LIF signal production

The fundamental concepts of Laser Induced Fluorescence (LIF) are not hard to capture. A molecule has a series of energy states which are addressable by absorption of photons of suitable energy. Once such a photon has been absorbed, the molecule will, after some

time, relax toward its ground state configuration and will probably emit one or more photons in the process, which will be at a wavelength equal to or longer (i.e. lower energy) than the incident photon.

In a typical LIF experiment, a tracer molecule is added to the system which marks the location of the (non-fluorescing) species of interest. Taking a simplified model case, and ignoring any coherent excitation phenomena, the tracer species that is being excited by the laser beam can be represented by a two-level system of a lower ground state and upper excited state (labelled by subscript ‘1’ and ‘2’, respectively - refer to figure 2.2). Furthermore, it is assumed that a low laser intensity is used such that stimulated emission is negligible. The evolution of population over time can then be described with the following rate equations:

$$\frac{dn_1(\nu)}{dt} = -n_1(\nu)B_{12}\rho(\nu) + n_2(\nu)A_{21} + n_2(\nu)Q_{21} \quad (3.1)$$

$$\frac{dn_2(\nu)}{dt} = n_1(\nu)B_{12}\rho(\nu) - n_2(\nu)A_{21} - n_2(\nu)Q_{21} \quad (3.2)$$

Where $n_i(\nu)$ represents the number density of the species addressed by laser radiation of frequency ν , B_{12} is the Einstein coefficient for absorption on the transition, A_{21} is the Einstein coefficient for spontaneous emission, Q_{21} is the rate of decay of the upper state by non-radiative processes, and $\rho(\nu)$ is the energy density per unit bandwidth of the incident laser radiation.

Taking a steady state solution, where $\frac{dn_i(\nu)}{dt} = 0$ (and dropping frequency dependencies),

$$n_2A_{21} + n_2Q_{21} = n_1B_{12}\rho \quad (3.3)$$

The first term of the equation above is the rate of production of fluorescence photons by the process. Denoting this term by R_{fluo} :

$$R_{\text{fluo}} = n_1B_{12}\rho - n_2Q_{21} \quad (3.4)$$

Equation 3.3 gives an expression for n_2 which can be substituted into equation 3.4 to find:

$$R_{\text{fluor}} = n_1 B_{12} \rho - \frac{n_1 B_{12} \rho}{A_{21} + Q_{21}} Q_{21} \quad (3.5)$$

$$= n_1 B_{12} \rho \frac{A_{21}}{A_{21} + Q_{21}} \quad (3.6)$$

Since this process is incoherent, these fluorescence photons will be emitted into all 4π steradians. Naturally, detectors do not enclose the entire volume under measurement, and a multiplicative factor of $\frac{\Omega}{4\pi}$, where Ω is the solid angle subtended by the collection optics at the source, is employed.

Multiplying the rate of production of fluorescence photons (equation 3.6) by the energy of the output photon, $h\nu$, gives the power output of the entire process.

Even in a simple system of two electronic states, the upper and lower states will in reality be composed of a number of sub-levels, that are at different energies owing to the rotational or vibrational excitation of the molecule. These individual levels are also not perfectly defined in energy, due to processes like pressure broadening, which *inter alia* modify the effective lifetime of the state.

Given that the laser is sufficiently spectrally broad to address all molecules in the sample - e.g. the frequency change caused by Doppler broadening is, at most, below the laser's spectral linewidth - and if the laser is also of a uniform intensity over this width, the factor $n_1(\nu) \times \rho(\nu)$ integrates across frequency to $N_1^0 \times \frac{I_\nu}{c}$. Here, I_ν is the laser intensity at the centre of its lineshape, and N_1^0 is the initial number of species in the lower state. In this way, it is possible to relate a measured intensity to the number density of the tracer species.

The loss of addressable population is rarely a problem in LIF - barring tracer molecule decomposition. Conceivably, very high temperatures could promote molecular population to higher electronic states which could make the excitation photon energy non-resonant - but this is assumed not to be the case in the current analysis.

Rolling these assumptions together, and denoting fluorescence power collected by F , the following relationship is derived⁴⁹:

$$F = N_1^0 \times \frac{B_{12}I_\nu}{c} \times \frac{A_{21}}{A_{21} + Q_{21}} \times \frac{\Omega}{4\pi} \times lA \times h\nu \quad (3.7)$$

Where A is the cross-sectional beam area and l is the path length of the absorption. Implicit here is a further assumption, that the laser intensity is not significantly affected as it propagates. That is to say that scattering losses are negligible, and the medium is optically thin (see discussion of absorption spectroscopy in section 2.6).

Note the presence of the factor $\frac{A_{21}}{A_{21} + Q_{21}}$. This is termed the ‘quantum efficiency’, and is a measure of the probability that an absorption event will lead to the emission of a fluorescence photon. In many cases, the quenching rate can be significantly higher than the spontaneous emission rate.

3.2.2. The impact of quenching on LIF

At this point, it is valuable to use some simple arguments to motivate the investigation of non-radiative de-excitation pathways - i.e. quenching processes - which have until now been simply denoted by the rate Q_{21} . The details of the processes at work inside a tracer molecule will be covered in section 3.4. For now, suffice to say that non-radiative de-excitation involves the loss of energy by means other than photon emission. This process is dominated by the contributions from interactions (mainly collisions) with other molecules.

Kinetic theory describes the concept of a mean free path ‘ λ ’ between collisions, treating gas molecules as hard spheres:

$$\lambda = \frac{RT}{\sqrt{2}\pi d^2 N_A p} \quad (3.8)$$

where R is the molar gas constant, T is the gas temperature in Kelvin, d is the molecular

diameter, N_A is Avogadro's number and p is the ambient pressure. Taking standard temperature and pressure with an estimated molecular diameter of 3 Angstroms, the mean free path (i.e. the distance travelled between collisions) is $\sim 100\text{nm}$. Travelling at 350ms^{-1} , the result is a collision every 300ps, or a collision rate of $\sim \times 10^{10}\text{s}^{-1}$.

Another result of kinetic theory indicates that the inter-molecular collision rate is linearly dependent on the gas pressure. The nature of the collision itself is somewhat affected by the gas temperature, since this increases the average molecular speed and causes more energetic impacts. Note that it is not actually required for two molecules impact upon one another - since Van der Waals forces, dipole moments or H-bonds can allow longer-range interactions. This means that the rate at which quenching may have an impact on the state of a molecule could be substantially higher than the simple calculation above - potentially causing changes on a timescale as short as $\sim 10\text{ps}$.

Different collision partners will also have different cross sections for de-excitation, since the presence of similar energy levels increases the chance of perturbing the excited electronic state when the molecules are brought together.

Energy is stored in a molecule in three principal forms. Displacements of the atoms in the centre of mass frame lead to rotational and vibrational motions, and energy can also be held in the excitation of electrons to higher electronic states. A molecule will rapidly 'fall down' the ladder of rotational energy states following interactions with other gas molecules, converting the energy stored in these internal motions into kinetic energy of the 'collider' molecule (picosecond timescales). Vibrational and electronic states, with their progressively larger energy gaps decay more slowly[†]. In this way, the kinetic energy of the gas is increased by the decay of excited state energy on different timescales. That is to say, a real molecular system is composed of many sub-levels, each with associated quenching rates - leading to $Q_{i,i-j}$ -type terms.

The reality of the situation is that the simple model presented in figure 2.2 does

[†]This is the 'energy gap' rule - smaller energy transitions typically occur more quickly. Note that this is the opposite to spontaneous emission, where larger energy gaps cause faster decay!

not account for this diversity of processes, where each may have inherent temperature, pressure and collision partner dependencies. Naturally, this means that simply writing a quenching rate as ‘ Q_{21} ’ could be quite misleading - though it serves a purpose in introducing the concepts. Note that an IC engine is a perfect example of an environment where a simple fixed quenching rate is a bad assumption, owing to the rapidly-changing environmental conditions during an engine cycle.

3.2.3. Measurements with LIF

Some applications of LIF were discussed in section 2.6.2. In the sphere of engine-based research, PLIF is normally undertaken by combining a base fuel which has no inherent fluorescent properties with a fluorescent tracer molecule. The tracer is chosen to mimic the behaviour of the fuel accurately, such that it serves as a good ‘marker’ of fuel location. Tracers which may be used include (but are by no means limited to) ketones such as 3-pentanone or acetone and aromatics such as toluene. The fluorescence from the tracer molecule is spectrally integrated when performing measurements of concentration, but by choosing a pair of spectral windows and ratioing signal strengths, the temperature sensitivity of the spectral emission can be harnessed to measure temperature. In both cases, accurate interpretation of the recorded signals is totally reliant upon determining the fluorescent behaviour of the tracer molecule as a function of the ambient conditions - that is, T , p and the composition of the surrounding gas will generally all have an effect.

For imaging purposes, LIF is employed in a planar geometry, PLIF (Planar-LIF), by shaping the exciting laser beam into a sheet (for example, with a cylindrical lens). The sheet is usually $\lesssim 1\text{mm}$ in thickness, and the fluorescence from the molecules situated in the illuminated plane is collected by a spherical lens sited at right angles to the plane, producing an image on a camera. The spatial resolution of the system is usually limited by either the acuity of the lens or the size of the pixels in the imaging array.

3.2.4. Pros and Cons of LIF

LIF, and particularly PLIF, has clear advantages. The technique is easily deployed for spatially-resolved qualitative measurement of distributions in optically-accessible environments. The amount of information acquired in a single shot can provide a wealth of insight, for example, to diagnose the cause of a persistent misfire as being caused by a lean mixture at the time of ignition. The comparison of these experimental results against the predictions of computer simulations is also clearly of high relevance to future engine developments. PLIF is also a very sensitive technique, and is equally well suited to measuring distributions of minor species such as OH down to concentrations of parts-per-million.

On the other hand, some fundamental limitations make getting the most from a LIF-based technique very challenging. A common drawback of all techniques that depend on intensity measurements is the reliance on the collection efficiency of the optics being constant with respect to time. In an engine environment, where combustion by-products may be deposited on the windows used for optical access, this cannot be assured. This is naturally particularly concerning when attempting to derive quantitative measurements. Worse still for quantitative analysis is the issue of the tracer's inherent quantum efficiency. As has been briefly described, the probability of any given absorption event producing a photon is affected by the local environment that the tracer molecule finds itself in. In order to attribute a measured intensity to a particular concentration, this variation cannot be ignored.

Chapter 4 of this thesis will present the development and application of a robust method for mitigating the effects of these problems and obtaining quantitative measurements in an IC engine.

3.3. Thermal Gratings

This section opens with a brief overview of the LITGS technique. The technique is then explored phenomenologically, illustrating the important concepts and some practical considerations, before the theory behind LITGS (as described by Paul et al.⁷⁹) is presented. This section closes with a consideration of how a LITGS signal is used to perform measurements, and a discussion of the advantages and disadvantages of the technique in general.

Laser Induced Thermal Grating Spectroscopy (LITGS) is a non-linear optical diagnostic technique. The thermal grating is produced by quenching energy absorbed from laser-generated fringes. The strength of this grating is then monitored as a function of time using a Bragg-scattered probe beam. The temporal character of the grating's behaviour is governed by the fluid dynamics of the medium into which the grating is written, allowing the derivation of, for example, temperature and pressure.

3.3.1. Phenomenological discussion of LITGS

The formation, evolution and readout of a thermal grating is illustrated in the three sub-diagrams of figure 3.1. (For a more rigorous account of the theory behind the LITGS technique, see the following section.) In the first step towards the production of a LITGS signal, an intensity fringe pattern is overlaid onto the medium by the interfering pump beams - step **a**. The spacing of these fringes, Λ , is given by

$$\Lambda = \frac{\lambda_p}{2 \sin(\theta_x/2)} \quad (3.9)$$

where θ_x is the crossing-angle of the pump beams and λ_p is the pump wavelength.

The laser wavelength is tuned to an electronic resonance of a constituent in the gas, which causes energy to be stored in excited states of the molecules in the high-intensity regions of the fringe pattern following absorption. This energy is subsequently released

by quenching processes leading to localised heating, as represented in step **b**. Since adding heat to a particular parcel of gas causes its temperature and pressure to rise, this mechanism allows a periodic three-dimensional density modulation to be written into the crossing zone - which is the thermal grating.




Figure 3.1.: A schematic illustrating formation and readout of a thermal grating as a three-step process. a) Formation of a fringe pattern by the interference of two coherent pump beams. b) Thermalisation of energy absorbed from the pump beams causing localised heating. c) Monitoring the visibility of the grating by use of a probe beam.

Step **c** shows the application of a third (‘probe’) beam, usually incident at the first-order Bragg-angle defined by

$$\lambda = 2\Lambda \sin \theta \quad (3.10)$$

where λ is the probe wavelength, Λ is the grating plane spacing and θ is the angle of incidence. The intensity of this coherently-scattered probe beam is dependent on the modulation depth of the induced grating.

Since the rapidly-deposited heat causes both temperature and pressure to change, the thermal grating is said to be formed from two components. The first component to consider is that due to the temperature modulation - the *temperature grating*. Where the local temperature has been raised, the local density will be reduced, such that the density change from temperature modulation follows the shape of the intensity fringe pattern. The strength (i.e. visibility) of this component of the overall thermal grating

will reduce over time as diffusion returns the system to equilibrium, leading to a LITGS signal which monotonically decays over time.

The second component is known as the *acoustic grating*. With the *swift* deposition of heat into the gas as the upper state population is quenched, there is an associated pressure increase. The response of the gas to this imbalance is to create shockwaves, which only propagate constructively in the direction through the grating planes. With one acoustic wave travelling in each direction, a standing wave is formed. The strength of this standing wave is co-sinusoidally modulated with time, and also decays as time passes due to viscous effects damping the motion.

The overall reflected signal is therefore composed of a part which simply decays (*temperature*), and a part which not only decays with time, but also oscillates in strength (*acoustic*). The relative strengths of these processes is related to the rate at which the quenching occurs. If heat is only slowly released into the gas, then no pressure imbalance will be generated, and the acoustic grating component will not form. The reflected probe beam - i.e. the LITGS signal beam - describes how these two components have affected the grating visibility in combination.

Three approaches to grating readout are generally adopted by researchers in this field. The most common approach is that of the use of a continuous-wave source, such as a 1W argon-ion laser^{86,90,91}. Others have used a short-pulse laser, and scanned the time delay between the readout pulse and the pump pulses⁹², or used a cavity to prolong the lifetime of a short pulse⁹³. An approach that is rarely used is that of a long pulsed laser^{7,80}. A long-pulse laser such as the one employed in the experiments reported in this thesis can be capable of delivering approximately 100mJ per μs , which leads to an improvement in the amount of scattered light available for detection which is in the region of 5 orders of magnitude above that of a 1W cw source. This has obvious repercussions on improved detectability against sources of background interference, and allows the use of faster (i.e. higher bandwidth), less sensitive detectors.

Naturally, scanned short-pulse techniques are inherently unable to perform single-shot measurements, with a full grating decay event requiring many shots to capture as the pump/probe delay is scanned.

For a typical experiment, the reflected beam will contain oscillations at approximately 50MHz (i.e. a reflectivity peak every 20ns). The decay processes will limit grating lifetimes to approximately $1\mu\text{s}$, during which time acoustic waves will travel about 0.3mm. Since the width of the crossing area of the pump beams is normally several millimetres, effects arising from waves propagating out of the crossing region can be ignored.

Owing to this oscillatory nature, the detector used to acquire the signal must be able to respond to high frequency signals without distortion. However, given that a LITGS signal may also be in the region of $1\mu\text{s}$ in length, faithful capture of the entire decay requires good lower-frequency response also.

3.3.2. LITGS theory

The analysis presented in this section draws heavily on Paul et al.⁷⁹, and provides the theoretical basis of the preceding section.

In the process of Laser Induced Thermal Grating Spectroscopy (LITGS), pump beams are crossed to produce fringes. Population is quenched, as described previously, leading to heating at the intensity maxima, separated by typically 5-50 μm . To describe this process in detail, the effect this heating has on the gas dynamics has to be addressed. Considering the scale of the problem, local thermal equilibrium is assured in experiments since the mean free path is much shorter than the grating spacing. In addition, the mean time between inter-molecular collisions is small compared to the laser pulse duration and the decay time of the induced grating. It is therefore acceptable to use the linearised hydrodynamic equations⁹⁴ for the evolution of the density, ρ . Equations 3.11 to 3.13 are the continuity equation, and the two linearised hydrody-

dynamic equations (in one dimension, taken as moving through the grating planes), in a form that suits this calculation.

$$\frac{\partial \rho'}{\partial \zeta} + \frac{\partial u'}{\partial \xi} = 0 \quad (3.11)$$

$$\frac{\partial u'}{\partial \zeta} + \frac{1}{\gamma} \frac{\partial P'}{\partial \xi} - \frac{4}{3\text{Re}} \frac{\partial^2 u'}{\partial \xi^2} = 0 \quad (3.12)$$

$$\frac{1}{\gamma} \frac{\partial P'}{\partial \zeta} - \frac{\partial \rho'}{\partial \zeta} - \frac{1}{\text{RePr}} \left(\frac{\partial^2 P'}{\partial \xi^2} - \frac{\partial^2 \rho'}{\partial \xi^2} \right) = \phi_{\text{ext}} = \frac{\gamma - 1}{\gamma} \frac{\epsilon h \nu_a Q \Lambda}{c_0} \frac{N^*(\zeta, \xi)}{P_0} \quad (3.13)$$

Here, the equations have been written for normalised perturbations (e.g. $\rho'(x, t) = \Delta\rho(x, t)/\rho_0$, where ρ_0 is the unperturbed density value and $\Delta\rho$ is the change). In these equations, p represents pressure and u flow speed through the grating planes. ϵ is the fraction of the energy stored internally in the molecule that is released as heat, after absorbing a photon of energy $h\nu_a$. $\zeta = c_0 t/\Lambda$ and $\xi = x/\Lambda$ are normalised time and distance parameters, that render the set of equations dimensionless. γ is the ratio of specific heat capacities (c_p/c_v) and c_0 is the isentropic speed of sound. N^* is the excited state population number density, and varies as a function of x and t - or in the dimensionless formulation, ξ and ζ . Q is the quenching rate, and Λ is the fringe spacing. ϕ_{ext} is the external driving term which is perturbing the system.

The linearised ideal gas law $P' = \rho' + T'$ and the equation of state for the enthalpy of an ideal compressible medium are used in casting the above equations as written.

The Reynolds and Prandtl numbers are defined as: $\text{Re} = c_0 \Lambda \rho_0 / \mu$ and $\text{Pr} = c_p \mu / K$. μ is the dynamic viscosity and K is the thermal conductivity. The Reynolds and Prandtl numbers are non-dimensional, and formed from characteristic values of the problem to be analysed. The Reynolds number is concerned with the relative magnitudes of the inertia of the flow and the viscosity. High Reynolds number is therefore related to turbulent flow fields, where the inertia of the flow dominates, and a low Reynolds number is related to a laminar ('smooth') flow field, where the fluid viscosity is dominant. Similarly, the Prandtl number relates viscous diffusion and thermal

diffusion. This means that a small Prandtl number indicates a flow regime where the thermal diffusion rate is high, leading to thermal boundary layers which are thick relative to velocity boundary layers (as in flows with liquid metals).

Equation 3.11 is the transformed continuity equation, and the fact that it is set to zero means that there is no loss or creation of mass. Equation 3.12 is the momentum transport equation, and is set to zero because there is no net force acting on the molecules, since the experimental fields are too low to cause electrostriction. The rightmost term in equation 3.13, ϕ_{ext} , drives the energy transport equation, and is the result of the transfer of internal molecular energy to heating. It can be seen that the driving term has a dependence on position, which may be understood as the result of the spatially-varying fringe pattern. It also varies in time, as the excited state population decays by non-radiative processes.

As an irrotational problem with a source function that is separable in space and time, the system of equations is reduced to one wave equation which is solved by Fourier transforms in space and Laplace transforms in time. Finding the solution for a spatially infinite grating - which is fair if the acoustic waves don't propagate significantly outside the crossing zone in the grating decay time - the normalised density perturbation is:

$$\rho' \propto \cos(2\pi\xi)Z(\zeta) \quad (3.14)$$

$$Z(\zeta) \propto (W(\zeta) * G(\zeta) * M(\zeta)) \quad (3.15)$$

$$W(\zeta) = a_1 \exp(s_1\zeta) - 2a_2 \exp(-u\zeta) \cos(v\zeta + \phi) \quad (3.16)$$

The constants a_1 , s_1 , a_2 , u , v and ϕ are found as part of the solution.

$$s_1\zeta = -\frac{(2\pi)^2}{\text{RePr}}\zeta = -\frac{4\pi^2 K}{\Lambda^2 \rho_0 c_p} t \quad (3.17)$$

$$u\zeta = \frac{(2\pi)^2 (\gamma - 1 + 4\text{Pr}/3)}{\text{RePr}} \zeta = -\frac{4\pi^2 K}{\Lambda^2 \rho_0 c_p} \left[\frac{\gamma - 1 + \frac{4c_p \mu}{3K}}{2} \right] t \quad (3.18)$$

$$v\zeta = 2\pi\zeta = 2\pi\frac{c_0}{\Lambda}t \quad (3.19)$$

The fraction of the probe beam scattered from a spatially-modulated grating such as that described here is given by^{95,96}:

$$\eta = \left(\frac{\pi\Delta nd}{\lambda}\right)^2 + \left(\frac{\Delta\alpha d}{4}\right)^2 \quad (3.20)$$

Where d is the grating thickness, n is the refractive index, α is the absorption coefficient and λ is the probe wavelength. Assuming there is no electronic resonance at the wavelength of the readout beam, $\Delta\alpha = 0$. This is the case of a ‘phase grating’ - where only the refractive index modulation amplitude, Δn , is non-zero. In general, the refractive index of a medium is a function of a parameter X in the following sense:

$$\Delta n = \frac{\partial n}{\partial X}\Delta X \quad (3.21)$$

The coefficient $\frac{\partial n}{\partial \rho}$ is non-zero[†], and so this mechanism is responsible for the reflectivity of the grating structure. Equation 3.14 describes the density modulation that is produced, and so its form is of utmost interest. Also of note is the convolution of functions shown in equation 3.15. This equation indicates that the medium’s response is a convolution of the temporal shape of the laser pulse, G , the decay shape of the quenching process, M , and the function W , which describes the grating impulse response function.

Equation 3.16 gives the form of $W(\zeta)$. The first term in this equation is identified as the stationary temperature grating and its decay. The second term is associated with a standing pressure wave, which is composed of two counter-propagating acoustic waves, and their associated decay constant. The identification of each part of equation 3.16 is aided by studying the factors that combine to make the constants s_1 and u .

s_1 is a grouping that is dependent on K , the thermal conductivity, and describes

[†]and $\frac{\partial n}{\partial T} \ll \frac{\partial n}{\partial \rho}$, so the influence of temperature is negligible.

the transport of heat energy away from its source. u is directly proportional to the classical acoustic wave attenuation coefficient. u is also seen to differ from s_1 simply by the factor in square brackets. Evaluating this factor for commonly-used gases and at realistic pressures finds a range of values from 0.67 to 0.83. This indicates that the standing acoustic wave decays at almost the same rate as the static temperature contribution. The presence of Λ underlines the dependence of these constants on the geometry of the experiment.

Figure 3.2 shows how the two components described by equation 3.16 evolve over time. In order to present the concepts clearly, the grating has been forced to decay more quickly than would be expected in a real situation. Four times have been selected for interrogation, denoted by the vertical dashed lines on the upper two traces, and labelled **a**, **b**, **c** and **d**. The spatial density modulation in the direction through the grating planes is then presented in the four traces in the lower section of the figure. Note that in the first instant, the density modulation is zero, as the temperature and acoustic parts cancel one another. At later times, the acoustic and temperature components interfere with one another to determine the overall grating reflectivity.

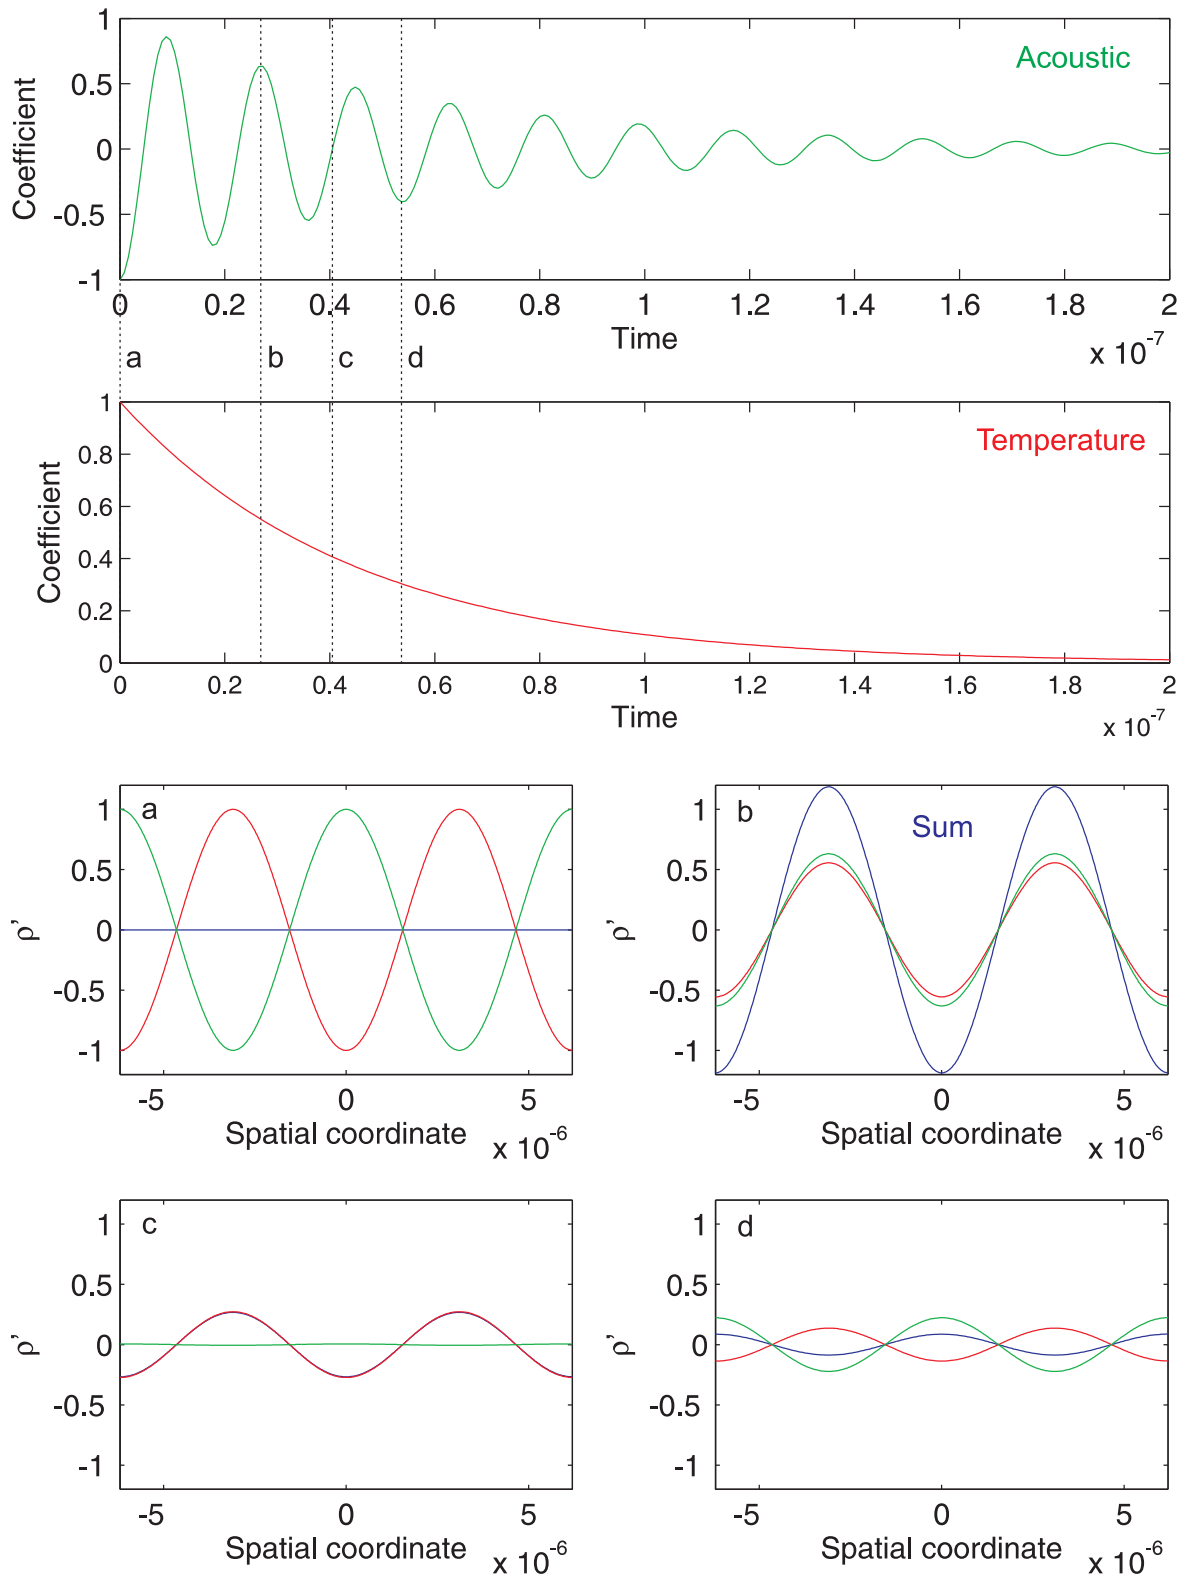


Figure 3.2.: Evolution of the density perturbation at the heart of a LITGS signal. Modulation arising from pressure waves in green, from temperature in red, and the sum in blue. Temporal coefficients are plotted in the upper section, with four interrogation times marked. These four times are those used in plotting the spatial modulation in the lower section of the figure.

3.3.3. Deriving parameters from LITGS signals

The temporal character of a LITGS signal encodes much information which is of interest in the context of an internal combustion engine. Though it may not prove possible to measure all of the following simultaneously, a LITGS signal has been used to measure temperature, pressure, composition/concentration and velocity. The review paper of Stampanoni-Panariello et al.⁷⁷ discusses a range of experiments - also including the use of LITGS in imaging - at the Paul-Scherrer Institute in Switzerland. The concepts behind measuring T , p and composition will now be discussed. (For velocity, refer to e.g. Walker et al.⁸⁴, Hart et al.⁸⁵ and Kozlov⁸⁶).

Temperature measurement

By measuring the frequency of oscillation in the signal, the time taken for an acoustic wave to travel a fringe spacing, τ , can be recovered. Assuming the fringe spacing Λ is already known, a speed can easily be calculated. It was seen previously that the travelling waves in the acoustic grating acquire phase as $2\pi\zeta$ (equation 3.19). The phase speed of a wave is given by ω/k . In this case, $\omega = \frac{2\pi c_0}{\Lambda}$, and the wavelength is clearly Λ . The calculation thus yields c_0 - the sound speed in the gas, showing that acoustic waves travel at the sound speed. A familiar formula, $c_0 = \sqrt{\frac{\gamma RT}{m}}$ demonstrates the relationship between sound speed and gas temperature. Note that this analysis relies on foreknowledge of γ/m , so knowledge of the gas composition is required for accurate derivation of T .

Pressure measurement

With reference to equations 3.17 and 3.18, the decay characteristics of the components of the grating are expected to exhibit a strong dependence on the gas density, ρ_0 . It can be expected that fitting a modelled LITGS signal to a measurement will lead to this information. However, the derivation of pressure requires accurate estimates of

gas kinetic parameters to be made. To outline a *possible* method for the simultaneous measurement of temperature and pressure, the fitting method employed in this thesis is used as an example. In the course of determining the best fit of the model to the experimental data, the routines that have been developed select optimal values for the Reynolds number, Re , and the inter-fringe transit time, τ . In a control experiment, where data are recorded at a known temperature and composition, the oscillation frequency of the acoustic part of the grating may be used to determine the grating spacing, Λ , accurately using the following relationship:

$$T = \frac{\Lambda^2 m}{\tau^2 \gamma k_B} \quad (3.22)$$

Where k_B is the Boltzmann constant, and the oscillation frequency is simply $1/\tau$. Substituting the above formulation of temperature, and the definition of the Reynolds number into the ideal gas law:

$$p = \frac{\rho_0 k_B T}{m} = \frac{\mu Re}{c_0 \Lambda} \frac{k_B c_0^2 m}{m \gamma k_B} = \frac{\mu Re c_0}{\Lambda \gamma} \quad (3.23)$$

It is seen that the accurate determination of pressure, p , depends upon foreknowledge of μ , the dynamic viscosity, and γ , the ratio of specific heats, for the particular gas mixture under investigation. (Recall that $c_0 = \Lambda/\tau$.) Determination of gas kinetic properties has been previously considered by Stevens and Ewart⁷. Deriving pressure from the decay timescale of the grating requires that the decay lifetime is not limited by the size of the grating.

This thesis is not concerned with the measurement of pressure, but as can be seen in the results section of chapter 5, the values of Λ , τ , μ , γ etc have conspired to cause:

$$Re \approx 100 \times (\text{pressure in bar})$$

Which is easily verified with the following approximate values for air at room tem-

perature: $\gamma = 1.38$, $\rho_0 = 1.184\text{kgm}^{-3}$, $c_0 = 345\text{ms}^{-1}$, $\Lambda = 6.2 \times 10^{-6}\text{m}$ and $\mu = 1.983 \times 10^{-5}\text{kgm}^{-1}\text{s}^{-1}$ - giving $\text{Re} = 128$.

Composition/concentration measurement

There are two main approaches to the problem of composition or concentration measurements with LITGS. The first is easily anticipated from the previous two measurement types, and relies upon measuring the small change in grating transit time, τ , as a result of the change in average molar mass affecting the sound speed⁸². The uncertainty in this measurement is affected by the difference in molar mass between the diluent and the absorbing species. Owing to the variety of molecules present in an IC engine, this technique may prove challenging to implement (in contrast to, e.g. measuring the composition of a binary mixture, where $\sim 1\%$ accuracy is achieved on CH_4/N_2 and H_2/N_2 mixtures⁹⁷).

An alternative measurement solution is to determine the concentration of the absorbing species by the intensity of the scattered probe beam⁸³. This method relies upon increased number density of the absorbing species causing greater heat deposition and therefore stronger density modulation. It should be noted that this method has the downside of relying on intensity measurement, and is thus vulnerable to the same window fouling problems as the LIF techniques.

Other approaches to concentration measurement (particularly, of oxygen) will be investigated in chapter 5 of this thesis.

3.3.4. Pros and Cons of LITGS

In contrast to the spatially-resolved distributions that may be produced by a technique like PLIF, the LITGS technique is generally used to perform localised measurements in a small volume. (The technique can be extended to measurements along a line, with sheet-like laser beams⁸¹.) The LITGS technique is more demanding to deploy than

PLIF, since it typically employs two laser sources and three detectors[†]. On the other hand, there is no requirement for an expensive camera or lens.

The reduction in spatial information gathered by LITGS is compensated by high precision measurements of gas parameters (T , p) which do not rely on intensity measurement. The coherent nature of the signal generation also aids in background suppression, since the signal propagates from the interaction region in a well-defined beam which may be detected far from the probed volume with little or no loss in signal strength.

This thesis will present an extension to the LITGS concept, allowing for the measurement of fuel/air ratios in chapter 5. A competing technique described by Seeger et al.⁸⁷ exists, which was described in section 2.8 and will be used as a benchmark.

[†]One each for measuring the pump and probe pulses, and a third for the LITGS pulse itself, in order to later decouple the influences of the pump and probe pulse shapes from the signal.

3.4. Tracer Photophysics

In this section, the photophysical properties of tracers are explored. Toluene - an aromatic hydrocarbon - will be considered in the first instance, with the second subsection explaining the differences in behaviour seen with a ketonic tracer like acetone. Much of the information in this section is derived from the review paper of Schulz and Sick¹⁷.

3.4.1. Toluene (aromatic)

Toluene is an aromatic hydrocarbon, often used in PLIF experiments as a fuel tracer. It has chemical formula $C_6H_5CH_3$ - i.e. a benzene ring with a methyl group attached to it - leading to the alternative name 'methyl-benzene'. It is this benzene ring that leads to the distinctive smell of an aromatic hydrocarbon. Toluene has a molar mass of 92.14g/mol, a density of 0.8669g/ml and a boiling point of 110.6°C.

Toluene exhibits a broad absorption spectrum, capturing photons below ~ 280 nm in wavelength. The shape of the absorption spectrum is temperature-dependent, and the peak at ~ 265 nm red-shifts to longer wavelengths at increasing temperature. Subsequent fluorescent emission from toluene occurs at longer wavelengths than excitation. For excitation at 248nm, the emission spectrum covers 265-320nm, with a peak at ~ 285 nm.

The typical quantum efficiency of toluene is approximately 10-30% at room temperature and pressure in a nitrogen bath gas, but this value is strongly affected by the presence of oxygen and the ambient temperature. In fact, the measured behaviour indicates \sim exponential decay with increasing T or O_2 concentration. As a result of this and the potential for variation in absorption cross-section, the fluorescent yield of toluene may vary dramatically over the range of conditions found inside an IC engine.

A brief review of available papers relating to the behaviour of toluene at different T , p and oxygen concentration will now be conducted.

Koban et al.⁹⁸ reports on the absorption and emission spectra of toluene, considering the change in toluene absorption cross section at 248nm and 266nm as a function of temperature. (Note, these two wavelengths are of particular importance, being respectively the fundamental emission of the KrF excimer laser and the fourth-harmonic of Nd:YAG.) The authors discovered that the absorption cross-section at 248nm is roughly constant from 300-900K, encompassing the compression stroke in a GDI engine. Fluorescence quantum yield was found to decay approximately exponentially with increasing temperature in a nitrogen bath gas, with the value at 900K about three orders of magnitude weaker than at 300K (slightly different functional dependence is found at 248nm and 266nm). The authors also note the potential for temperature measurement by ratio-ing two parts of the emission spectrum.

The effect of increasing temperature and oxygen concentration on toluene fluorescence is detailed in Koban et al.²³. Oxygen is demonstrated to cause a slight red-shift in the emission spectrum. Stern-Volmer diagrams, plotting $\frac{\text{Signal strength without oxygen}}{\text{Signal with oxygen}}$ for a range of T and oxygen concentration are included for excitation at 266nm and 248nm. These plots - together with information from the literature - motivate the development of a model for toluene. This model forms most of the basis for the explanation of toluene's photophysics that will shortly be introduced. This model is then applied in Koban et al.²⁷ to describe toluene fluorescence in an IC engine, with reasonable agreement in evidence.

Zimmermann et al.⁹⁹ present measurements of the fluorescence lifetime of toluene utilising excitation at 248nm and varying temperature and oxygen concentration. The best fit to the experimental data is achieved by using a bi-exponential decay model, consisting of a faster and slower decay process. As temperature and oxygen concentration are increased, the two decay lifetimes are observed to shorten. Increasing temperature from 300-650K in a nitrogen atmosphere has the same effect on the long lifetime as increasing the oxygen concentration from zero to 6% at 300K, causing the long lifetime to drop from 45ns to \sim 2ns in both experiments. The effect on the short lifetime com-

ponent is a little less pronounced, causing a decrease from $\sim 1.8\text{ns}$ to $\sim 0.5\text{ns}$ in both cases. It is important to note that these data are only able to determine information about the behaviour of the S_1 state, since the phosphorescence from the excited triplet state will not be included in measurements on these short timescales of $< 200\text{ns}$.

The full description of all the processes and rates which are active in the photophysics of toluene is beyond the scope of this thesis, but it will prove instructive to consider a model of its behaviour which encompasses most of the important concepts. Figure 3.3 is a Jablonski diagram which describes the processes acting in toluene. This illustration and the accompanying explanation is largely derived from a combination of information from Schulz and Sick¹⁷, Koban et al.²³ and Koch¹⁰⁰.

Discussion of figure 3.3 begins by describing the fundamental features of the diagram, continues to consider the general features, and then discusses the unusual aspects of toluene.

The diagram consists of a series of lines, each corresponding to a particular molecular configuration, where the energy of any particular configuration increases vertically. The levels are grouped into two columns which represent two multiplicities - the left column those states with a multiplicity of one and the right column those states with a multiplicity of three. Multiplicity is determined by the number of combinations of electron spin that are possible for a given configuration. The left hand column is further split into two, with the lower-lying levels the ground state and the upper levels the first excited state. The groups of levels are labelled with a letter - where 'S' is singlet (multiplicity 1), T is triplet (multiplicity 3) - and a subscript number where '0' is the ground state, '1' the first excited state *etc.*

Note that T_1 lies slightly lower than S_1 in energy. This occurs because the spin wavefunction of the electrons in the triplet state is symmetric. Since the overall wavefunction must be anti-symmetric, the spatial part is required to be anti-symmetric. This leads to reduced overlap between the electrons and so the Coulomb repulsion

term is reduced in magnitude, lowering the energy slightly. In the case of toluene, this energy change is known to be $\sim 1.3\text{eV}^\dagger$.

Each group ('manifold') of levels consists of a level where no energy is stored in the vibrational modes of the molecule - this is illustrated at the base of a manifold with a thick line. Above this relaxed level lie many sub-levels, represented by thin lines, where increasing energy denotes ever more quanta of energy being stored in the vibrational modes of the molecule. Note that rotational modes also exist but are more finely spaced than their vibrational counterparts, and not represented here.




Figure 3.3.: A Jablonski diagram illustrating the processes active in the toluene molecule. Dashed lines indicate processes enhanced by the presence of oxygen. Dotted vertical lines within a manifold indicate vibrational redistribution. Wavy lines are used for processes which absorb (up) or emit (down) photons.

The distribution of population between energy levels when the system is in thermal equilibrium is described by the Boltzmann distribution, as given in equation 2.8. Under the application of a laser pulse, such as that shown by the upward-pointing light blue wavy lines in figure 3.3, this distribution may be reproduced in the first excited singlet state at some non-equilibrium offset above the relaxed state. This energy will be quickly dissipated by interactions with other surrounding gas molecules, allowing the population to fall down the vibrational ladder. This process is illustrated by the

$^\dagger 1\text{eV}$ is $1.6 \times 10^{-19}\text{J}$

downward-pointing red dotted lines in each manifold, and inevitably leads to heating of the bulk gas.

Following the absorption event, the molecule may fluoresce. This process produces a photon, and is illustrated by the downward-pointing green wavy lines. The fluorescence is likely to occur from the lower-lying levels of the upper manifold, but the rate of vibrational relaxation is important in determining this. In the gas phase, vibrational relaxation may not be significantly faster than the fluorescence lifetime. Competing with fluorescence is the process known as inter-system crossing (ISC). ISC allows the system to change multiplicity by altering the spin of an electron. This process does not alter the total energy stored in the molecule, and since the triplet relaxed state lies at a lower energy than the first excited singlet, the excess energy is held in vibrational modes of the molecule. This process is illustrated by purple horizontal lines in the diagram. The rate at which the ISC process occurs is known to increase with vibrational excitation, due to increased interaction between the electronic and vibrational parts of the wavefunction.

Once a molecule has transferred to the triplet state and vibrationally relaxed, it is likely to undergo further ISC to the S_0 state¹⁰¹. The competing radiative process of phosphorescence is comparatively extremely slow, because spontaneous emission by the dipole transition moment is forbidden from this level due to the change in multiplicity. The molecule must instead use a higher-order moment (which has suitable selection rules) in order to emit a photon whilst undergoing this transition to the ground state - for example, the electric quadrupole operator. Transitions requiring the use of higher-order moments may take up to seconds or even minutes depending on how strongly forbidden the transition is.

When toluene is excited in the presence of a nitrogen bath gas and at a wavelength of 266nm, the processes described above are a fair representation of the mechanisms available. The fluorescence lifetime of S_1 is measured to be on the order of 100ns, with ISC to T_1 occurring with an equivalent lifetime of order 10ns.

This picture is modified somewhat in the presence of molecular oxygen. O₂ is unusual, in that it exists in a triplet ground state. Kikuchi et al.¹⁰² reports that the collisions between excited singlet molecules and oxygen causes the ISC rate to be highly increased. This process occurs through an intermediate encounter complex, and leaves oxygen in its singlet state (which is metastable), because the triplet-singlet energy difference in oxygen is 0.98eV. It is reported that virtually every collision between oxygen and the excited singlet state will experience this mechanism, resulting in *extremely* rapid removal of population from fluorescing states, and a dramatic reduction in fluorescence efficiency. Kawaoka et al.¹⁰³ further describes that molecular oxygen is able to quench triplet states of toluene with near unity efficiency (again producing singlet oxygen) - i.e. limited by gas collision rate, and that oxygen also provides an enhancement to the rate of ISC from T₁ to S₀. They quantify the ISC process as typically occurring 100-1000 times more infrequently than triplet state quenching leading to O₂ singlet production. Usui et al.¹⁰⁴ summarises the available pathways for oxygen quenching of aromatics, and Schweitzer and Schmidt¹⁰⁵ is a comprehensive review of the generation and deactivation mechanisms of singlet oxygen.

A further modification to the standard picture is required at high excitation energies - e.g. 248nm radiation coupled with high temperatures, or simply under high temperature conditions. Burton and Noyes¹⁰⁶ first noticed a marked increase in non-radiative de-excitation with high excitation in the S₁ state. Jacon et al.¹⁰⁷ extended this work and proposed the mechanism could be internal conversion (IC - non-radiative decay without a change in multiplicity) back to the S₀ state. They suggest that this event is triggered by the strong excitation of particular vibrational modes of the molecule causing a strong overlap with states in the S₀ manifold. This IC mechanism therefore only applies to particular highly-excited sub-levels of the S₁ manifold, and is said to occur on a very short - perhaps picosecond - timescale.

To summarise, the model contains absorption from S₀ to S₁, occurring in $\sim \times 10^{-15}$ s, spontaneous emission back to S₀ at $\sim \times 10^{-7}$ s, ISC from S₁ to T₁ at $\sim \times 10^{-8}$ s, phos-

phorescence from T_1 at $\sim \times 10^{-6} - \times 10^{-3}$ s. Vibrational relaxation occurs at a rate which depends on the gas pressure, since higher density results in more perturbation to the oscillation - but is $\sim \times 10^{-9}$ s. Some states are vulnerable to IC at $\sim \times 10^{-12}$ s. The presence of oxygen will dramatically increase the rates of ISC and quenching, making phosphorescence completely negligible, and reducing fluorescence strongly.

3.4.2. Acetone (ketone)

Acetone is a simple ketone with chemical formula CH_3OCCH_3 - a double-bonded carbon-oxygen group and two methyl groups. Acetone is in wide use as a solvent, and can be found in, for example, paint thinner and nail polish remover. Acetone has a molar mass of 58.08g/mol, a density of 0.79g/ml and a boiling point of 56.5°C.

Acetone has a spectrally broad absorption feature which covers $\sim 240\text{nm} - 320\text{nm}$, with a peak at $\sim 280\text{nm}$. As temperature increases, the peak gains a slight red-shift and the cross-section generally rises. Acetone fluoresces in the range $\sim 320\text{nm} - 500\text{nm}$ with a peak at 400nm (slightly dependent on excitation wavelength). This enables the detection of acetone fluorescence with consumer-grade camera lenses, in contrast to aromatic fluorescence which requires special UV-transmitting optics, typically manufactured using fused silica elements.

Very briefly, another, heavier member of the ketone group is 3-pentanone. This molecule exhibits very similar photophysical behaviour to acetone. 3-pentanone has the chemical formula $\text{C}_5\text{H}_{10}\text{O}$, i.e. the same structure as pentane, with a carbon-oxygen double bond replacing one of the hydrogens on the central carbon in the chain. 3-pentanone has a molar mass of 86.03g/mol, a density of 0.815g/ml and a boiling point of 101°C.

A brief review of papers that describe the temperature and pressure dependencies of ketone tracers is now given:

Results of measurements on acetone and 3-pentanone fluorescence as a function of

pressure at 383K and 248nm excitation are presented by Grossmann et al.²¹. They also plot the temperature dependence of 3-pentanone at 248nm, 277nm and 312nm. The authors are quick to recognise the potential for temperature measurement, and assess the potential of a ratioing technique.

Ossler and Alden¹⁰⁸ reports on the application of quadrupled Nd:YAG mode-locked picosecond pulses to acetone and 3-pentanone fluorescence. The paper presents more results from 3-pentanone than acetone, but includes emission spectra and effective upper state lifetime measurements for a range of temperatures and pressures.

Thurber et al.³³ covers the variation of the absorption cross section of acetone at specific excitation wavelengths (rather than a spectrum) as a function of temperature. These data are used to demonstrate the temperature dependence of the absorption profile. The authors develop a model for ketone fluorescence and discuss application in temperature measurement using the ratio method.

The behaviour of acetone fluorescence at 248nm, 266nm and 308nm excitation with change in pressure and bath gas is mapped out in Thurber and Hanson²². The authors spend some time testing and discussing the impact of adding oxygen to the bath gas.

The effect of changing temperature and excitation wavelength on 3-pentanone fluorescence is reported by Koch and Hanson²⁴. The fluorescence spectrum is presented, and absorption is tested at 248nm, 266nm and 308nm. Finally, the fluorescence signal strength is measured as a function of temperature at these same wavelengths in a bath gas of nitrogen.

Rayleigh scattering is used by Koch et al.¹⁰⁹ to calibrate measurements of acetone and 3-pentanone fluorescence spectra. Acetone fluorescence is investigated as a function of pressure and excitation wavelength.

Braeuer et al.¹¹⁰ introduces some improvements to the photophysical model introduced by Thurber et al.³³, particularly in the case of high excitation energy - i.e. the use of a 248nm laser. Acetone and 3-pentanone are investigated using air and CO₂ bath gases with a range of pressures and temperatures.

Further measurements of 3-pentanone fluorescence with changes in temperature and pressure are given in Modica et al.¹¹¹. Their results are compared against the Thurber et al.³³ model, as modified for 3-pentanone.

Koch et al.¹¹² treats measurements of 3-pentanone and acetone absorption spectra when shock-heated up to 1100K.

The low saturation intensity of 3-pentanone is highlighted by the measurements performed by Petersen et al.¹¹³. Whilst the effective lifetime of ketones is short, this is only due to the high ISC rate - which does not return population to the ground state sufficiently quickly to allow the use of high laser intensities without risk of saturation.

The model presented in the previous section (3.4.1) for toluene is largely accurate for acetone. The differences can be summarised as follows.

The energy difference between the singlet S_1 and triplet T_1 states is much reduced in ketones, at only $\sim 0.2\text{eV}$. Recalling the ‘energy gap’ rule, this results in much faster ISC between the singlet and triplet states. However, since this energy is no longer sufficient to promote oxygen from its ground state triplet to its excited singlet state, the dramatic enhancement of the ISC rate which is seen in toluene is not present here¹¹⁴. The addition of oxygen therefore does not have such a marked impact on ketone fluorescence (though it is not negligible in many cases).

In the ketone molecule, the radiative transition to S_0 from S_1 is actually symmetry- and orbital-forbidden. This results in a long spontaneous emission lifetime. However, owing to the very rapid ISC rate, the effective lifetime of the upper state still appears to be short, and is typically measured to be on the order of 5ns.

Finally, there is no evidence reported in the literature for an IC process with a high activation energy as found in toluene.

3.4.3. Motivations for tracer selection

As discussed in section 3.4.1, the internal relaxation processes of the aromatic tracers are quite complex. However, the higher fluorescence yield compared to ketones (section 3.4.2) and thus stronger signal is a compelling argument for their use. The polar nature of oxygenated hydrocarbons also tends to cause problems with achieving ideal co-evaporation from multi-component mixtures. Ultimately, when this knowledge is combined with the very low temperature sensitivity of the absorption cross-section at the experimentally-accessible wavelength of 248nm, the tracer for the medium fuel fraction was chosen to be toluene rather than 3-pentanone. Tri-methyl-benzene ('TMB') was a natural choice for tracing the heavy fraction of the fuel.

The choice of the light fraction is largely motivated by the available species. Acetone does not exhibit ideal co-evaporation with its fuel fraction, as will be seen in section 4.3.1, but lighter aromatics simply are not available. By exciting acetone at 277nm, the temperature dependence of the fluorescence yield is noticeably reduced relative to 248nm³³.

4. Quantitative Planar Laser Induced Fluorescence - QPLIF

4.1. Chapter overview

This chapter explores the development and application of the Quantitative PLIF technique. The chapter will begin by giving a brief description of the technique and project. The next section begins discussing the details, covering the multi-component fuel, the optical engine and the optics used in these tests. Of central importance to the application of QPLIF is the issue of calibration, which will be treated separately in section 4.4. Another important building block in QPLIF is the analysis system which translates the images into fuel concentrations, so this is considered just before the presentation of experimental results in sections 4.6 and 4.7. This chapter closes with a summary, which includes a discussion of the requirements for a quantitative PLIF implementation.

4.2. Technique overview

4.2.1. Measuring fuel distributions with QPLIF

Quantitative Planar Laser Induced Fluorescence is a linear incoherent optical diagnostic technique which employs a sheet-shaped laser beam to excite fluorescence from tracer molecules doped into a non-fluorescing base fuel. By imaging the ‘glow’ from the

tracer, the distribution of the fuel inside the engine can be determined. By means of *in-situ* calibration, the images of fluorescent intensity can be converted into quantitative fuel distributions.

Unique to this work is the application of Quantitative PLIF to the case of a multi-component fuel with three co-evaporating tracers. The base fuel (designed in collaboration with Shell UK) comprises six components, paired into three fractions, and is designed to have similar properties to a standard gasoline. One fuel fraction can then be traced by addition of the corresponding tracer to the base fuel. For more details, see section 4.3.

The calibration system that is employed in this work was independently developed, but shares some features with that reportedly previously by Han and Steeper³². The approach presented in section 4.4 does have some additional improvements, however.

This work also clearly reveals the effect of exhaust gas residuals, which are not accounted for in some calibration schemes, e.g. that reported by de Sercey et al.¹¹⁵. This will be shown to have a significant impact in section 4.4.3.

4.2.2. 2-level model and QPLIF

The simple 2-level system referred to in section 2.6 is reproduced in figure 4.1. This diagram presents a basic model system, and the pathways that are available within it - absorption, spontaneous emission and quenching (in the low-intensity regime where stimulated emission is negligible). To recap, the PLIF technique is reliant upon capturing the photons produced by the spontaneous emission process.

The upper state population is effectively competed over by spontaneous emission and non-radiative quenching processes. This leads to the concept of a ‘quantum efficiency’, a factor of the form $\frac{A_{21}}{A_{21}+Q_{21}}$ which is essentially the probability of a single absorption event producing a fluorescent photon. Since the engine environment causes Q_{21} to vary substantially, the quantum efficiency must be known in order to make sense of




Figure 4.1.: A representation of a two level system. Transition rates include photon absorption $B_{12}\rho(\nu)$, spontaneous emission A_{21} producing a photon (wavy line), and a quenching process Q_{21} arising from molecular collisions.

the fluorescent signal strength. The implementation of this calibration constitutes a significant fraction of this chapter.

4.2.3. CCSPV Consortium

The QPLIF technique was developed as part of the contribution from Oxford University towards the CCSPV consortium. The Combustion Concepts for Sustainable Premium Vehicles (CCSPV) project was formed with researchers from the universities of Oxford, Leeds, Loughborough and UCL, together with industrial partners Jaguar Cars UK and Shell UK. The project was funded by the EPSRC, Shell and Jaguar, with all fuels supplied by Shell.

The project was convened to investigate the sources of cyclic variability. PLIF was selected in order to determine what role the fuel distribution might have to play. The reduction of cyclic variability is motivated by a desire to achieve the best possible optimisation of an engine, which is naturally limited by how reproducible an engine cycle is. Other universities investigated different aspects of the engine's behaviour - Leeds investigated the combustion event itself, Loughborough the airflow structures inside the engine, and UCL the characteristics of the direct injection fuel spray.

Loughborough, UCL and Oxford were each supplied with the same model of prototype single-cylinder Jaguar optical engine. This engine features a multi-hole spray-

guided direct injector, a centrally mounted spark plug, and modifications providing optical access, enabling techniques such as PLIF and PIV to be applied.

The production version of this engine has recently been released onto the market, titled the ‘V8 AJ133 Gen-III’, and will see use in models such as the Jaguar XF and XK. The normally-aspirated version of this engine improves on the old design by increasing peak power and low engine speed torque by 27% and 34%, respectively. Brake-specific fuel consumption has also improved by 5% over the old model at 1500rpm and 2.62bar BMEP, which is indicated to be class-leading performance compared with other manufacturers¹¹⁶.

4.3. Experimental details

In this section, the experimental arrangements will be detailed. Discussion begins with the development and composition of the multi-component fuel, followed by a description of the optical engine. The final subsection regards the optical system, covering a range of topics from the source laser to equipment modifications to improve optical collection efficiency.

4.3.1. Multi-component fuel

The issue of selecting a suitable fuel and tracer for investigating fuel distributions in IC engines has seen several different approaches. Typically, the base fuel is chosen to be iso-octane, on account of its similar volatility to pump gasoline. Iso-octane does not fluoresce in itself, so a tracer is normally added at around 10% by volume. In order to ensure the tracer tracks the fuel, the boiling point and density are chosen to be as similar as possible. Two tracers are commonly used with iso-octane - toluene (also known as methyl-benzene) and 3-pentanone. Recent work has identified that mixtures of iso-octane and 3-pentanone form what is known as an ‘azeotrope’, where preferential evaporation of one or other of the tracer or base fuel is caused if the correct

(temperature-dependent) mixture is not chosen¹¹⁷.

Since pump gasoline actually contains many hydrocarbon components, a single-component fuel may not accurately represent behaviour in an engine. To this end, the work of Han and Steeper³² presents the use of a two-component fuel. Whilst that is a valuable step towards modelling pump gasoline, this thesis presents the first use of a multi-component fuel that can more closely model the behaviour of the gasoline that one would buy in a petrol station.

| Fraction name | Component | Volume fraction | BP (°C) | Tracer species ($\lambda_{\text{excitation}}$, doping) | BP (°C) |
|---------------|-------------|-----------------|---------|----------------------------------------------------------|---------|
| Light | n-butane | 14.4% | -0.45 | Acetone (277nm, 5%) | 56.2 |
| | iso-pentane | 16.7% | 30 | | |
| Medium | iso-octane | 45.4% | 99.2 | Toluene (methyl-benzene) (248nm, 2%) | 110.6 |
| | n-octane | 11.2% | 126 | | |
| Heavy | n-decane | 2.2% | 174.1 | 1,2,4 Tri-methyl-benzene (248nm, 2%) | 169 |
| | Isododecane | 10.1% | 176 | | |

Table 4.1.: The Shell component fuel, illustrating tracers chosen and relative proportions. Despite apparent differences in boiling points (BPs), the tracers and fuel components have been shown to co-evaporate.

The multi-component fuel was developed in partnership between researchers at Oxford University and Shell UK prior to the author’s involvement in this project. The multi-component (‘pseudo’) fuel has the composition listed in table 4.1. The fuel can be seen to contain six components, paired into three fractions. The composition of each fraction is tailored to ensure that the fuel fraction co-evaporates with the listed tracer. The overall composition of the fuel is also very similar to commercial products, having a relatively large contribution from octanes (i.e. 8 carbon atoms).

This base fuel comprises solely of saturated hydrocarbons (*i.e.* no double-bonded carbons), and exhibits no absorption and no fluorescence when excited at 248nm. This enables the addition of a controlled amount of tracer, allowing a balance to be struck between optical thickness and fluorescent signal. The optimal concentrations are also given in table 4.1, and are chosen such that negligible absorption occurs as the laser beam passes through a uniform air/fuel mixture of equivalence ratio $\phi = 3$. This

condition is very important in producing analysable images, otherwise beam absorption would need to be accounted for¹¹⁸. The requirement for minimal dopant levels in the fuel imposes a severe limitation on the achievable signal strengths in quantitative PLIF experiments.

The characteristics of the component fuel were investigated by Shell UK, using a validated in-house Vapour-Liquid Equilibrium (VLE) modelling program. This represents quite a strict test of the co-evaporative behaviour of fuel and tracer, where the evaporation process occurs on a relatively long timescale. The evaporation process in the engine is likely to occur more quickly, owing to the fine atomisation of the fuel in the spray and the prevailing high-speed turbulent in-cylinder airflow characteristics. Such ‘diffusion controlled’ evaporation is predicted to have even better agreement between liquid and vapour compositions than that shown in figure 4.2¹¹⁹.

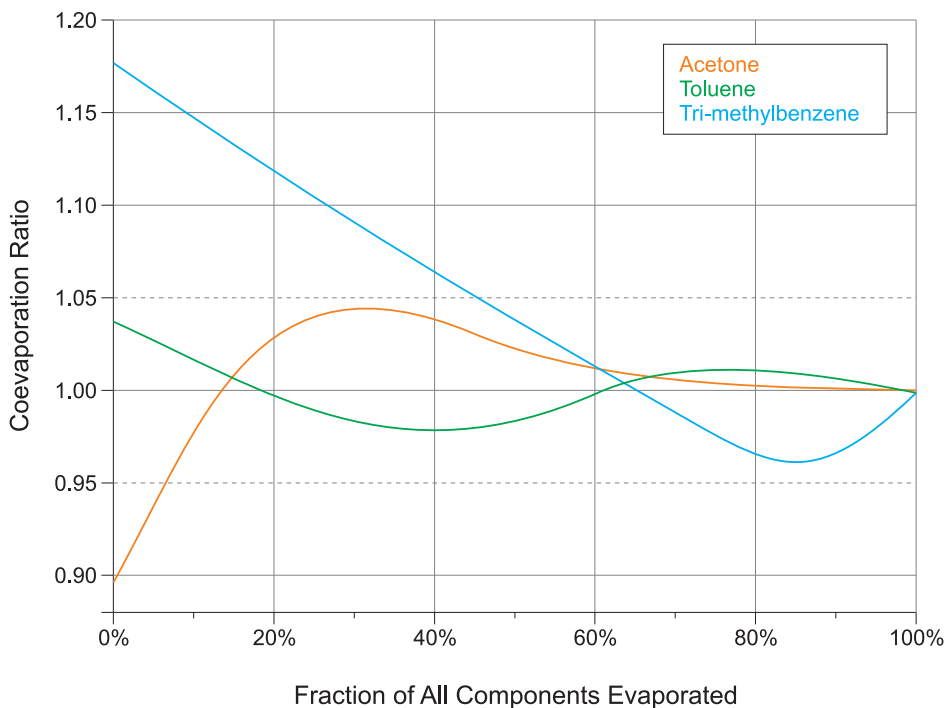


Figure 4.2.: A plot of the Shell ‘pseudo’ fuel co-evaporation characteristics for the three different tracers used in this experiment.

Figure 4.2 shows the modelled co-evaporation characteristics of the Shell component

From left to right, figure 4.3 illustrates that the behaviour of pump gasoline is quite different to iso-octane, but very similar to that of the pseudo fuel. This test provides important validation that the base fuel used in these experiments is representative of a real fuel, and is a marked improvement on using a simple single-component fuel as a surrogate.

4.3.2. Engine details

The single-cylinder optical research engine used in these experiments has been provided by Jaguar. It is effectively a single cylinder from a prototype of the recently-released ‘V8 AJ133 Gen-III’. The engine has a swept volume of 562cc, and a compression ratio of 11:1. The engine features a multi-hole spray-guided direction injector that is centrally mounted, next to the spark plug (figure 4.4 c). The fuel injector has six nozzles - four of which point downwards into the cylinder, and two that spray fuel towards the spark plug, one plume to either side. The engine has two intake valves and two exhaust valves. The engine is fitted with a piezo-electric pressure transducer which is always able to measure the pressure of the in-cylinder gases. This pressure transducer is referenced against a barrel-mounted transducer which is exposed either to ambient air or to cylinder contents, depending on the timing during the engine cycle. Provision is also made for fuel injection in the intake plenum, so the engine can be fired on homogeneous mixtures.

Optical access to the cylinder contents is provided by a 70% bore diameter fused silica window in the crown of the piston, which can be viewed from below by use of a 45° mirror mounted inside the extended piston skirt, to which the piston itself is attached. The cylinder is composed of a metal section through which temperature-controlled water may flow (the ‘barrel’), with a 25mm high annulus sandwiched between the metal barrel and the cylinder head. The piston rings are sited so as to avoid scuffing the optical annulus, but cause increased crevice volume relative to the standard engine.

The final element of optical access is provided by a roughly triangular window situated in the pent-roof region, mounted in the cylinder head. The fused silica annulus and pent-roof window are illustrated in figure 4.4 a.




Figure 4.4.: Images of the engine, showing: **a** - Optical access via pent-roof window and annulus, and rear of the cylinder pressure transducer. **b** - Integrated light from many cycles of fired operation. **c** - Cylinder head close-up with injector, spark plug, pent-roof window, valves, and beam termination stripe. Exhaust valves are on the left hand side of the image.

The piston is sealed to the bore of the cylinder with piston rings made from Torlon[®], a wear-resistant polymer manufactured by Solvay Advanced Polymers which contains PTFE and can thus be run without oil lubrication. This absence of oil in the cylinder is certainly advantageous for optical diagnostics. The engine is also connected to a dynamometer, allowing for both fired and motored operation (i.e. with and without ignition).

| | | | |
|-----------------|-------------------|----------------------|----------|
| Engine speed | Fuel pressure | Air flow rate | SOI |
| 1500 rpm | 150 bar | 1.21Ls ⁻¹ | 280 BTDC |
| Ignition timing | Manifold pressure | Intake air | Lambda |
| 35 BTDC | -0.479 bar gauge | 30°C | 1.0 |
| IVO | IVC | EVO | EVC |
| 24 ATDC | 274 ATDC | 244 BTDC | -6 BTDC |

Table 4.2.: WWMP definition, as supplied by Jaguar. Acronyms and items of terminology, such as ‘IVO’ and ‘Lambda’ are covered in section 2.2.

The Jaguar engine was operated at the World Wide Mapping Point, defined by the combination of a number of parameters, summarised in table 4.2. The standard injection timing for WWMP operation is 280 BTDC. It was found that the latest injection timing for CoV IMEP[†] < 8% was 140 BTDC. The use of these two injection timings is the only difference between the ‘early’ and ‘late’ operating conditions.

ETCS - the Engine Timing Control System is responsible for overseeing the operation of the engine. For example, it issues the commands to trigger the injection at the requested timing during the cycle. ETCS also interfaces with the data acquisition (DAQ) systems which log analogue signals from a number of sources. The low speed DAQ (LDAQ) system takes temperatures and other low frequency data once per engine cycle, whilst the high speed DAQ (HDAQ) acquires a sample from all the connected sources once every CAD. Table 4.3 lists the signal recorded on each of the high and low speed DAQ channels. ETCS and the DAQ systems are all implemented in National Instruments LabView, utilising National Instruments DAQ cards PCI-MIO16E for high speed and PCI-6024E for low speed data logging.

Following the delivery and installation, two major problems were revealed related to QPLIF. The first apparent issue was that the valve stem seals did not prevent oil seeping down from the cylinder head into the engine. This was diagnosed by the presence of a very fine fluorescent mist and extremely strong fluorescent signal from the

[†]IMEP is a measure of engine performance. CoV, or Coefficient of Variation is defined as the standard deviation of a population divided by the population mean, and is therefore a normalised measure of variability.

| Channel No. | HDAQ signal | LDAQ signal |
|-------------|-------------------|-------------------|
| 0 | Unused | Exhaust temp. |
| 1 | Cylinder Pressure | Coolant In temp. |
| 2 | Barrel Pressure | Coolant Out temp. |
| 3 | Exhaust Pressure | Barrel Out temp. |
| 4 | Inlet Pressure | Air temp. |
| 5 | Laser trigger TTL | Oil temp. |
| 6 | Airflow | Plenum temp. |
| 7 | Lambda sensor | Lambda sensor |
| 8 | Fuel pressure | - |
| 9-11 | - | - |
| 12 | - | Throttle position |
| 13 | - | Airflow |
| 14 | Camshaft flag | Engine speed |
| 15 | Long TDC flag | Engine torque |

Table 4.3.: Listing of HDAQ and LDAQ channels used in the final QPLIF experiments.

rear of the valves when open. This problem was resolved by adopting roller bearings and grease lubrication in the cylinder head. The second major issue related to the material the cylinder head is made from. The fabrication process of prototype parts involves the use of a low-pressure cast. Since this may result in a porous metal, prototype parts are sealed with a resin (in this case, methyl methacrylate is used). Unfortunately, this resin fluoresces very brightly when illuminated at 248nm. Figure 4.4 c shows the fillet of aluminium solder which was embedded into the cylinder head in order to suppress this source of background fluorescence.

The engine was maintained and operated during the reported experiments with the aid of staff and students in the Department of Engineering Science.

4.3.3. Optical details

Laser system

Tracer excitation in the QPLIF experiment is achieved by use of a Lambda Physik LPX120i. The gas used for these experiments is a krypton-fluorine mixture, which lases at 248nm. This laser was built in 1991, and has required some refurbishment

work to keep it running (for example, replacement electrodes built following advice from Ernst¹²¹, replacement of seals etc). This laser is currently capable of producing up to 150mJ per pulse, in a pulse length of $\sim 15\text{-}20\text{ns}$. Maximum repetition rate is on the order of 200Hz.

For the excitation of the ketonic tracer, acetone, a longer excitation wavelength is known to reduce the temperature sensitivity of the fluorescence yield. For this reason, the excimer laser is optionally focussed into a 70cm-long gas cell that is pressurised to 30 bar with hydrogen. At the focus of the beam, the intensity is sufficiently high to induce Stimulated Raman Scattering (SRS) - a coherent version of spontaneous Raman scattering, as discussed in section 2.6.3. This process converts $\sim 30\%$ of the radiation at 248nm into radiation at 277nm. The fundamental wavelength is also converted into other Raman orders (e.g. 312nm, 359nm Stokes and 225nm anti-Stokes, all spaced in energy by 4155cm^{-1}), although with much lower efficiency. The single Raman order of interest is selected by dispersing the output of the gas cell with a large fused silica prism, and using a mask, as illustrated in figure 4.5. (Originally, this was achieved with Schott glass filters, but the losses to absorption in the glass were quite high - as much as $\sim 70\%$ - and there was a constant risk of thermal damage to the filters from absorbing the energy in the radiation at 248nm.)

One concern for the quantitative analysis of PLIF images is that the fluorescent signal is linearly proportional to laser intensity provided the laser intensity lies below the saturation intensity of the molecule. Saturation occurs when the rate of photon absorption is greater than the rate of the processes that return the population back to the ground state. Saturating behaviour is often described by a factor $\frac{1}{1+I/I_{\text{sat}}}$, where I is the incident intensity and I_{sat} is the saturation intensity. In order to avoid the effects of saturation, the laser intensity should be kept to $\lesssim 0.5 \times I_{\text{sat}}$.

For acetone, the lifetime of the upper state is typically quoted to be approximately 5ns³³. This would correspond to a saturation intensity which is much higher than can be achieved with a laser sheet from this source. However, recalling discussion




Figure 4.5.: A schematic of the optical arrangements employed to convey the excimer beam to the engine. Upper section: Michelson-style pulse splitter used for aromatic tracers. Lower section: Raman shifter for acetone (light fraction) tracer. Dashed lines are ‘flippable’ mirrors. Orange line denotes 50% beamsplitter. Red lines illustrate mask on dispersed Raman orders. Dumbbell-like shape is the 30-bar hydrogen cell.

from section 3.4.2, this is merely an *effective* lifetime - the rate at which population is returned to the ground state in a ketone is much lower, causing a similarly reduced I_{sat}^{113} . However, it is determined experimentally that this is not a problem in this system because of the relatively low intensity produced by the SRS process in the hydrogen cell.

The saturation of the aromatic tracers is more concerning, however, since their upper state lifetimes are $\sim 100\text{ns}$ and there is a relatively large amount of energy in the 248nm laser pulse. Saturation was avoided in this case by splitting the source pulse from the excimer laser into two sub- I_{sat} pulses using an optical arrangement bearing some similarities to the Michelson interferometer. The creation of multiple low-energy pulses is preferred, since it increases the amount of laser energy delivered to the engine, and so increases the fluorescence signal strength as compared to a single sub- I_{sat} pulse. The chosen optical arrangement incorporates a 50% beamsplitter at 45° and two normal incidence high-reflecting mirrors. The short ‘arm’ of the Michelson-

like setup adds a negligible path length, whilst the long ‘arm’ is sufficiently far ($\sim 6\text{m}$) to delay the second pulse long enough to produce an output consisting of two well-resolved copies of the input pulse. This arrangement has the disadvantage of causing half of the beam energy to be reflected in the wrong direction, but this drawback is compensated for by the reduced alignment complexity relative to other optical setups featuring ‘loop’ beampaths.

Imaging system

Recall from equation 3.7 that the signal from a LIF process is proportional not only to number density but also to laser intensity. This relationship has the potential to introduce ambiguity into the interpretation of an image - a strong signal may either be the result of a high number density, or the result of a high intensity. In order to remove this uncertainty, a method is required to measure the intensity profile of the laser sheet used to excite the fluorescence.

Figure 4.6 is a schematic diagram of the experimental apparatus used to capture the profile of the laser sheet in the QPLIF experiment. The system works by splitting off a small fraction of the beam, and passing it into a fused silica cell which contains a low concentration of laser dye in methanol. The path distance between the low-reflectivity beamsplitter (simply an uncoated window) and the dye cell is chosen to be the same as the path length between the beamsplitter and the imaging region inside the engine, to account for any beam evolution that may occur over this distance.

The fluorescence produced at visible wavelengths from the dye in the cell is then imaged onto a ThorLabs LC1 USB CCD line-array using a 50mm f/1.4 Nikon lens set to f/8[†]. A transverse profile is saved to the computer for every shot of the experiment, allowing any shot-wise beam shape variation to be accounted for. Integrating the trace from the LC1 camera also enables the total laser pulse energy to be monitored and corrected on a shot-wise basis - a method which is clearly superior to the use of a

[†]Best edge contrast is recorded in the image by using a decreased aperture, such as this.




Figure 4.6.: Schematic of line-camera beam profile monitoring system

photodiode for the same purpose. Note that commercial devices, such as that used by Pfadler et al.¹²² are available, which can be used to ‘homogenise’ the laser sheet, obviating the need for laser sheet profile correction.

Figure 4.6 includes two lenses. The first (cylindrical) lens that the laser beam passes through has a focal length of approximately 1m, and is used to focus the beam down slowly to a thin sheet of 0.7-1mm thickness at the cylinder head in the engine. The second (also cylindrical) lens is used to ‘fan’ the beam out, in order to generate a laser sheet which is of equal width to the field of view of the camera/lens combination used to record the PLIF image ($\sim 60\text{mm}$).

Figure 4.7 illustrates the admission of the laser sheet into the engine via the 45° mirror in the extended Bowditch piston and the window in the piston crown. The fluorescence photons are then imaged into a Princeton Instruments ICCD-576G/RB with a Coastal Optics $f/4.5$ 105mm UV lens. The Intensified Charge Coupled Device (ICCD) camera consists of a thermo-electrically cooled CCD chip (576x384 pixels) with a close-coupled micro-channel plate (MCP) intensifier. The front face of the intensifier acts as a photocathode - a material with low work function which produces photoelectrons when illuminated. The photoelectrons are then accelerated down thin channels by a potential, before striking a phosphorescent layer (of short lifetime). This

system enables the detection of very faint signals.




Figure 4.7.: Diagram indicating optical arrangement used to image LIF plane onto ICCD chip

Nevertheless, it was rapidly determined that the images acquired by the camera were generally limited by the shot noise generated when running the intensifier at maximum gain. In order to alleviate these symptoms, the Coastal Optics lens was fitted with an adapter which allowed an anti-reflection coated fused silica lens to be placed just in front of the objective lens element. This modification resulted in a focal length reduction, and a decrease from $f/4.5$ to $f/3$. Concerns over image distortion were dismissed after conducting imaging trials using a sheet of squared paper placed in the laser sheet position.

The spatial resolution of the system was measured by imaging a back-illuminated $200\mu\text{m}$ pinhole. Figure 4.8 shows the comparison between the measurement and a model which performs a 2-D convolution of a circular step function with an Airy disc of variable width. The spatial resolution of the system, measured from the character of the Airy disc, was determined to be $350\mu\text{m}$ in the object plane. This spatial resolution is consistent with the specification of the camera - and is limited by the MCP intensifier. This spatial resolution is approximately equivalent to three pixels, and is a solid basis for employing a 3×3 median filter in the image processing stages.

This section on optical considerations closes with figure 4.9, an image of acetone




Figure 4.8.: Spatial resolution measurement, comparing measured and modelled intensity distributions at the CCD sensor.

fluorescence recorded by the author’s digital camera, which shows the location of the laser sheet inside the cylinder. Comparison against figure 4.4 a, particularly with reference to the stick-on temperature gauge on the metal barrel may aid in understanding the orientation of this image. Note also that the pent roof window is visible in the upper-right of figure 4.9.

Timing scheme

Figure 4.10 illustrates the timing scheme in use for the QPLIF experiments. When the timing control software (ETCS) determines that the engine is in the correct part of the cycle, a TTL pulse is created which triggers the system. Firstly, the pulse is amplified in order to trigger the line camera. The trigger pulse is then delayed by $5\mu\text{s}$, in order to allow the line camera sufficient time to prepare. The TTL signal then triggers the excimer mini-controller, which communicates with the laser via fibre-optics. There is some jitter between the application of the trigger to the mini-controller and the production of a laser pulse. This jitter has been minimised by using a signal derived from the high voltage thyatron trigger inside the excimer. This signal triggers



Figure 4.9.: Illustrating the in-cylinder location of the laser sheet with acetone fluorescence.

a Berkeley Nucleonics 555 pulse generator, which is able to control the delay precisely (resolution of 1ns) before the camera gate is opened. The TTL pulse from the BNC-555 triggers a Princeton Instruments PG200 which generates a 100ns long high-voltage pulse, activating the intensifier inside the camera.

The repetition rate of the system is constrained by the speed of the ICCD readout to ~ 1.5 Hz. When the engine is running at 1500rpm, as at the WWMP conditions, 12.5 engine cycles occur per second. This is clearly too fast for the camera, so a skip-triggering mode in the ETCS software is employed, whereby a laser trigger is produced for 1 in 9 engine cycles[†]. In addition to reducing background, the use of a short integration time in this experiment is of practical convenience, since it allows the experiment to be conducted with the room lighting turned on.

[†]Note that the engine is always run in ‘equilibrium’, such that the ignition and/or injection event occurs every cycle, in an attempt to reduce variability.

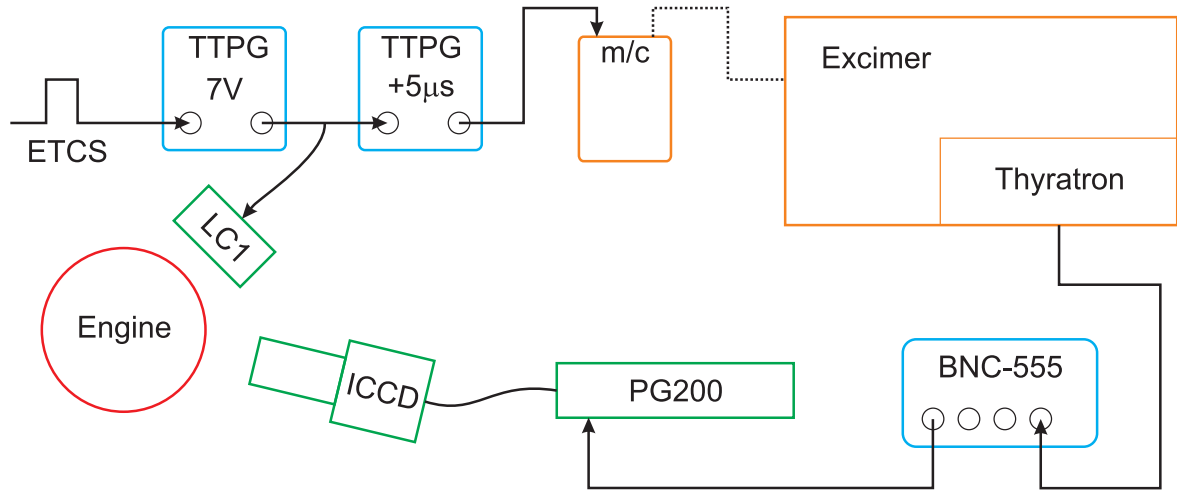


Figure 4.10.: Schematic of QPLIF timing system to synchronise line camera and ICCD. TTPG = Thurlby Thandar Pulse Generator, LC1 = ThorLabs line camera, m/c = mini controller for excimer laser, BNC-555 = Berkeley Nucleonics 555 pulse generator, PG200 = Princeton Instruments ICCD gate generator.

4.4. Calibrating QPLIF

Accurate calibration is pivotal to the correct interpretation of PLIF images when deriving quantitative fuel distributions. This work features an *in-situ* calibration method which uses a typical engine cycle to derive the behaviour of the quantum efficiency of the tracer as a function of crank angle during the compression stroke. This section opens by presenting a method relying upon use of a recirculation loop and motored calibration (similar to de Sercey et al.²⁹), followed by discussion of why this appears to be inadequate in section 4.4.2. The improved ‘bookend’ calibration method (independently developed, yet similar to Han and Steeper³²), which addresses the shortcomings of motored calibration is then described.

4.4.1. Calibration loop

Conceptually, the *in-situ* calibration method has a number of advantages over the alternative approach, which is to develop an exhaustive model of both the tracer and the engine. Modelling requires not only that the absorption cross-section and fluores-

cence quantum efficiency are accurately modelled as functions of T , p and composition, but also that the in-cylinder environment can be predicted. Whilst pressure is easily measured using a pressure transducer, temperature is much less straightforward. In addition, a model must also account for mass loss as the result of blow-by (where cylinder gases escape past the piston rings), which is typically more significant in an optical engine than its all-metal equivalent due to different ring material and sealing characteristics.

The first method that was investigated for *in-situ* calibration was that of a motored recirculation loop, conceptually very similar to that presented by de Sercey et al.²⁹. Figure 4.11 illustrates the hardware which is used in this calibration scheme.




Figure 4.11.: A simplified diagram illustrating the arrangements for conducting a calibration on the engine. fFID = fast FID.

In this motored (i.e. *without ignition*) loop scheme, a section of pipework carries the mixture which is expelled from the engine back around to the intake side. A bleed valve (not illustrated) and large intake plenum allows the air pressure to be stabilised, and any fuel losses are compensated for by use of the fuel injector in the plenum, which is connected to a tank of Shell pseudo fuel with added tracer. The recirculation

pipework is held at a constant temperature, in order to prevent fuel condensation. As the mixture is continually pumped around the loop, it becomes homogenised. The fuel concentration in the loop is monitored by a fast-response Flame Ionisation Detector (FID) - a device which, loosely-speaking, produces a signal proportional to the amount of carbon it burns⁵. Bursting discs are fitted to the cylinder and the recirculation pipework in the unlikely event that the gases self-ignite, allowing the pressure to be harmlessly released into the laboratory exhaust system. A filter is also fitted to the loop to catch any small particulate matter (e.g. fine dust from piston ring material) that is carried around the loop.

A full set of calibration data is collected by running the loop with concentrations of fuel that represent a range in equivalence ratio of approximately 0.3 to 3. These tests are conducted three times, each time using a different tracer in the fuel, so that the individual fluorescence yield characteristics of each tracer may be ascertained. 32 PLIF images are recorded every 10 CAD during the compression stroke, from 90 BTDC to TDC. Images are also recorded with no fuel or tracer present, in order to correct for the influence of any stray reflections and similar phenomena inside the engine.

Once the background frame has been subtracted from the fuelled image and shot-to-shot laser pulse-energy fluctuations have been accounted for, a mean fluorescence value is determined by averaging over as many of the available pixels as possible. These measurements are then plotted in 3-D, with co-ordinate axes of camera signal, equivalence ratio and in-cylinder pressure. The collection efficiency of the optical system is assumed to remain constant for all tests, and laser pulse energy is monitored before collecting a dataset with a Scientech-365 volume-absorbing power meter.

A routine written in MATLAB[®] is then used to perform a least-squares fit to these data points using the following equation - which is, itself, based on equation 3.7:

$$S = a \times (\phi + b) \times p^c \frac{A_{21}}{A_{21} + d \times p^e + f \times p^g \times \phi} \quad (4.2)$$

where $a - g$ are fitting variables which are determined by the non-linear least-squares fitting routine, and p , ϕ and S are the measured values of pressure, equivalence ratio and camera signal. A_{21} is the Einstein coefficient for spontaneous emission for the tracer. Coefficient a is an overall scaling factor. b allows for an offset in ϕ , but is always near zero. c , e and g allow the pressure and temperature dependencies of these three terms to be rolled together into one exponent. d is a coefficient related to the rate of pressure-dependent quenching, and f is a similar coefficient which handles any ‘self-quenching’, whereby fuel-fuel collisions result in non-radiative de-excitation.

The three datasets and their associated fitting surfaces can be seen in figure 4.12. It is straightforward to see how these calibration surfaces would be used; one simply uses the surface as a ‘look-up table’ to convert a value of camera signal and in-cylinder pressure into a value of equivalence ratio.

The details of the application of this form of calibration rely on highly methodical operator behaviour. In order to scale the signal axis of the 3-D surface correctly, the laser sheet cross-section and pulse energy must be known for calibration and measurement. Any change in the transmission efficiency of the optics will have a negative impact on experimental accuracy. Any relocation of the camera - for example, to view a smaller area in more detail - must be accompanied by a controlled experiment to monitor how the collection efficiency is altered.

Figure 4.12 presents PLIF signal strength as a function of ‘equivalence ratio’ and ‘pressure’. Note that a value ϕ of ‘equivalence ratio’ means that the fuel concentration was set to ϕ times that calculated for stoichiometry, as measured on a Combustion fast-FID HFR 400 (featuring *in-situ* zero and span gas calibration). Note also that whilst the ‘pressure’ axis does correspond to the pressure measured by the in-cylinder pressure transducer, the in-cylinder temperature would also be changing along this axis. This is the very essence of the *in-situ* method, and, providing the calibration cycle is representative of the measurement cycle, all is well.

Notice the shape of the calibration surface is markedly different between the aro-

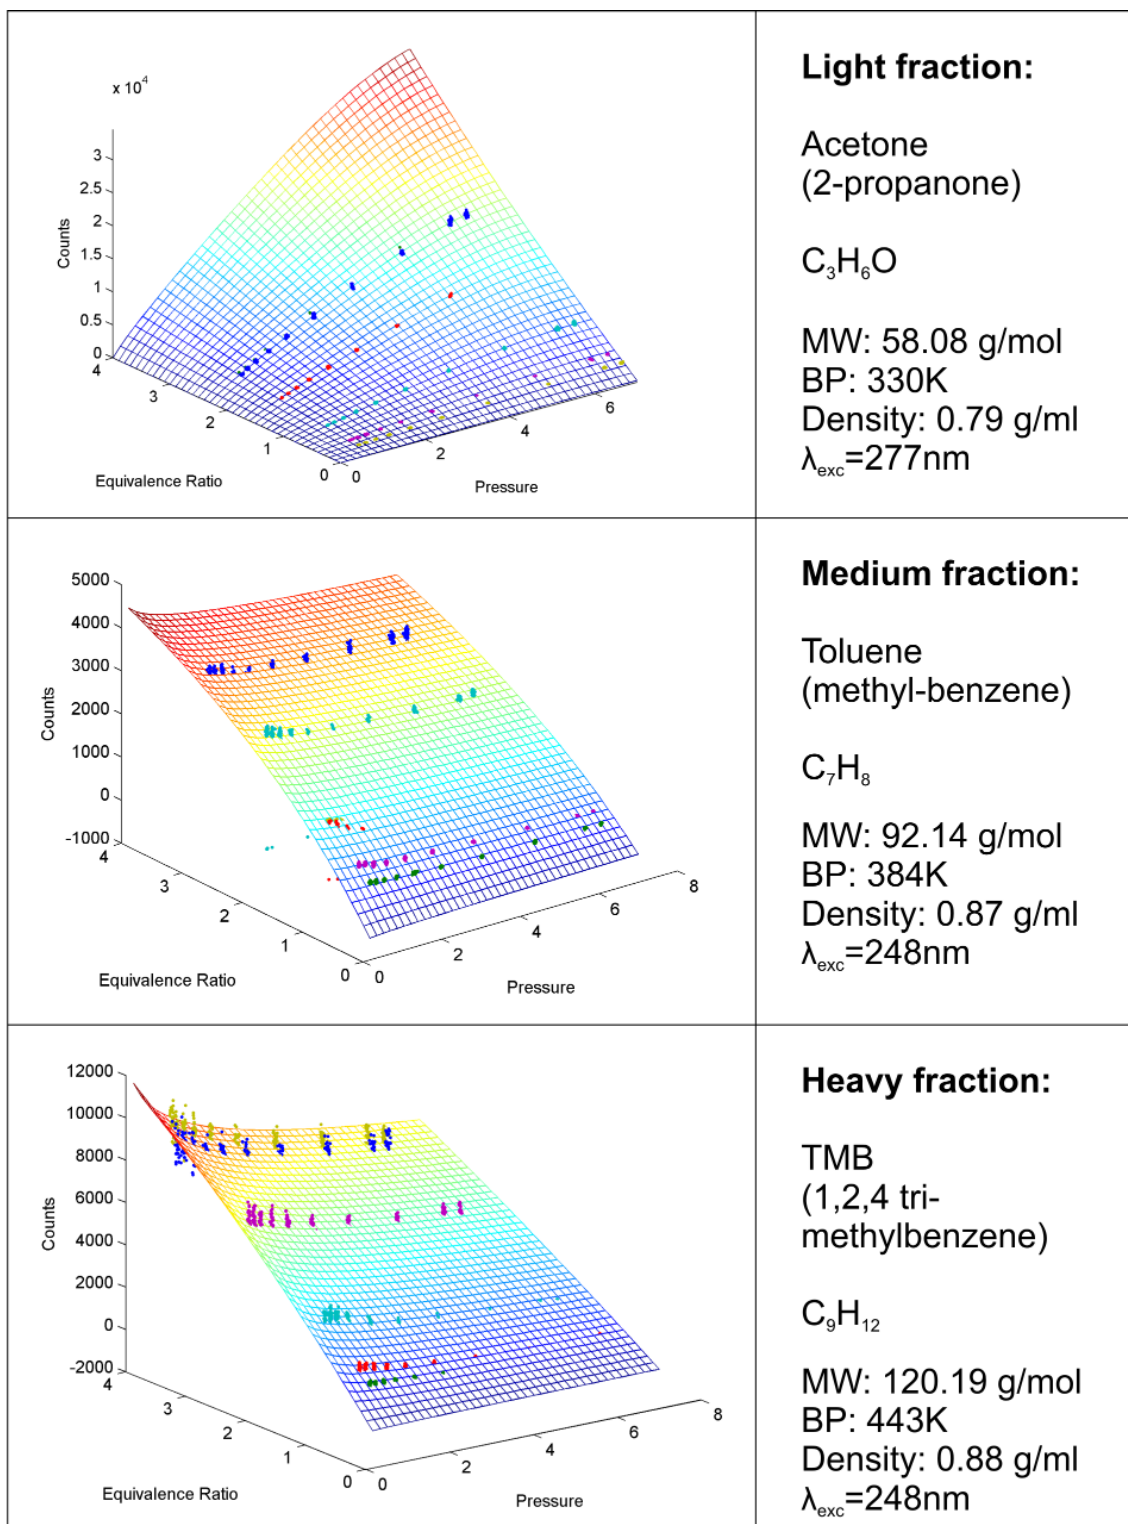


Figure 4.12.: The three calibration surfaces produced from the motored calibration loop data. Each data point corresponds to the mean PLIF signal recorded in a single image.

matic and ketonic tracers. Where acetone generally increases in signal strength with increasing pressure, the signal recorded from toluene and TMB is liable to decrease. This indicates that the increased number density at higher pressure is more than cancelled out by the negative effect of increased temperature on quantum efficiency when aromatic tracers are used.

Attempts to use these data to quantify the images produced by PLIF measurements of firing mixtures using direct injection in the engine without the recirculation loop produced unexpected results. Despite fastidious efforts to monitor laser intensity and optical efficiency, unaccountably large discrepancies arose. The mismatch was particularly profound in the case of early direct injection, where the mixture should have been homogeneous and roughly stoichiometric at the time of ignition - instead, the calibrated measurements were clearly systematically low, often reporting a fuel concentration approximately half the expected value.

The following sub-section discusses reasons for this behaviour, before a solution is presented in section 4.4.3.

4.4.2. The EGR effect

The disparity between the calibration data and the measurements was investigated by performing some comparisons between motored and fired operation. While early direct injection is able to yield a relatively uniform fuel distribution at the time of ignition, the fuel injector fitted to the intake plenum on this engine is clearly an even better candidate for achieving a homogeneous in-cylinder mixture, owing to the time available for mixing to occur.

Direct comparisons between the various types of operation were then performed with toluene - notionally the most important tracer, since it represents the behaviour of the largest fraction of the base fuel. Figure 4.13 shows how signal levels changed during these trial runs.

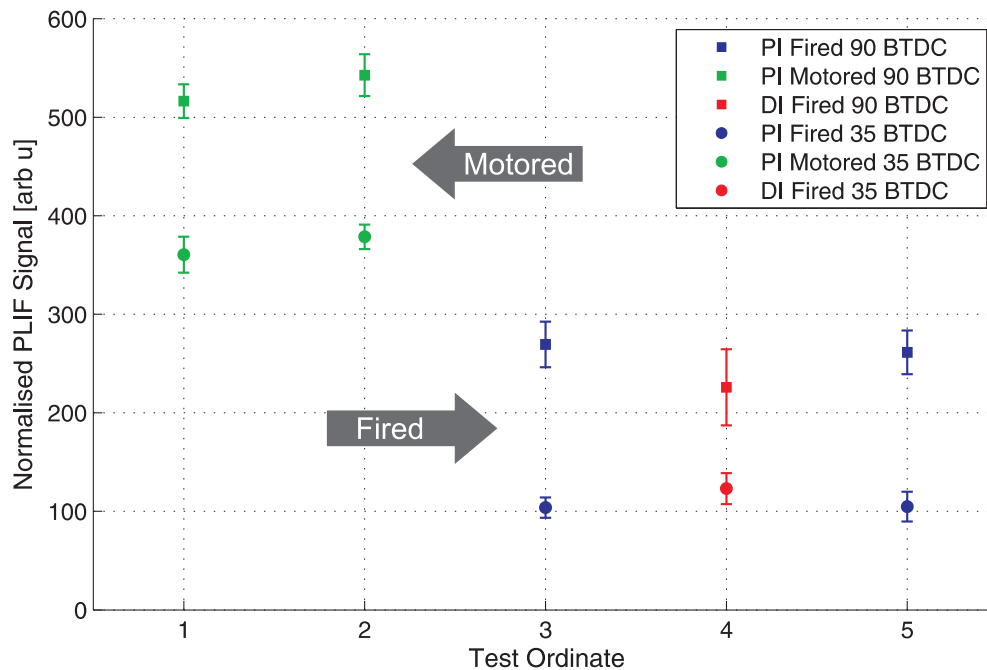


Figure 4.13.: The marked effect of fired operation on toluene fluorescence. DI = direct injection, PI = plenum injection.

Figure 4.13 plots the 'normalised signal' - that is, the fluorescence signal divided by the airflow and multiplied by the engine lambda. The engine was operated with the

same injection duration, air flow rate, intake air temperature and so on for motored and fired operation. When fired, a comparison was drawn between the signal measured with direct injection and plenum injection.

For each data point on the graph, 60 images were taken. A large region of interest in the pent-roof section of the image was selected and averaged to produce a mean signal strength for each image. This sample of 60 values was then averaged to yield the data point, and had its standard deviation calculated to produce the error bar.

The most striking result from figure 4.13 comes from comparing test numbers 2 and 3. The only difference between these two tests is that there was no ignition in test 2, and yet the signal strength can be seen to have declined by between 50 and 72% depending on the timing of the image (90 BTDC or 35 BTDC, respectively). Tests 1 and 2 and tests 3 and 5 are repeats of one another, in order to gauge the stability of this measurement. Comparing the size of the error bars against the magnitude of the effect, there is clearly a systematic trend here. Tests 3-5 are all fired, but test 4 employs early direct injection. It can be seen that the sample taken at 90 BTDC has higher fluctuation and a lower mean than those recorded with plenum injection. This is probably due to a less homogeneous mixture and a slightly non-uniform distribution. However, at 35 BTDC, the direct injection (DI) data is very similar in magnitude and variability to the plenum injection (PI) data, indicating that the mixture is now both more uniformly distributed and more repeatable.

Similar tests were conducted for the other two tracers, in order to investigate their behaviour under these conditions. Figure 4.14 shows this comparison.

Note that the dashed lines represent motored plenum injection (PI M) tests, and the solid lines denoted fired plenum injection tests (PI F). Each data point and error bar is produced in the same way as the data presented in figure 4.13. No direct injection tests were explicitly conducted with these tracers, since the difference in behaviour between DI and PI appears to be clear from the previously-presented toluene data.

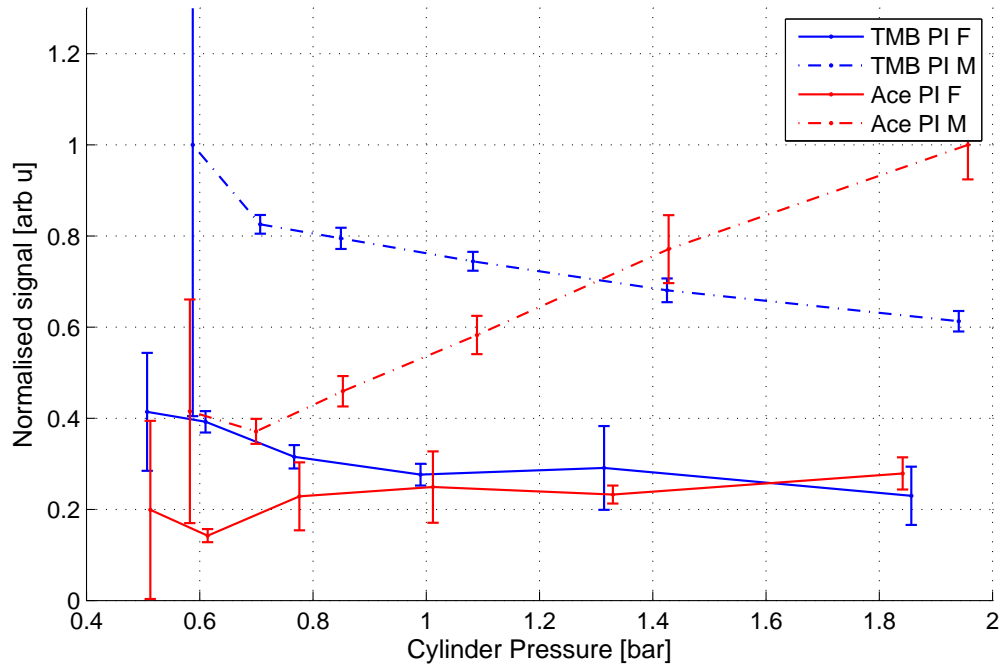


Figure 4.14.: Figure showing the effect of fired operation on signal levels for acetone (Ace) and tri-methyl-benzene (TMB). PI stands for ‘plenum injection’, F for fired operation, and M for motored operation.

Figure 4.14 indicates a similar story to that which was reported in figure 4.13 for toluene. The blue lines on the figure illustrate the behaviour of TMB, and show a signal drop of $\sim 50 - 60\%$. The red lines describe the behaviour of acetone, and show a drop of $\sim 55 - 70\%$, a gap which increases with increasing compression. Note that the trends shown in the motored data displayed in this figure are consistent with the shapes of the calibration surfaces that were presented in figure 4.12.

A number of hypotheses are put forwards to explain this behaviour. The major difference between fired and motored engine operation is the presence of EGR - exhaust gas residuals - in the former. That is to say, due to the valve timing on this engine, some residual burnt gas which would otherwise be expelled is kept between cycles.

The absorption properties of EGR were tested using a xenon flashlamp, a USB spectrometer (Ocean Optics USB4000) and a band-pass Schott glass filter. The trans-

mission spectrum of the cylinder contents was taken with and without the presence of exhaust gas residuals in the range 300-390nm, covering a significant fraction of the emission spectrum of toluene fluorescence. No absorption features were detected.

Review of the literature suggests that hot CO₂ and/or hot H₂O *can* cause significant absorption in the ultra-violet^{123,124}. However, the papers suggest that the temperatures achieved in the present experiments are insufficient to pose a problem at 248nm or 277nm.

Modelling and measurements previously undertaken in the Department of Engineering Science¹²⁵ have indicated that at the WWMP condition, 15% of the cylinder contents is composed of residuals from the previous cycle. Given that the residual gases from a fired cycle will not contain any tracer, whereas the residuals from motored operation will, this is an immediate source of a 15% reduction in signal strength by reduced number density. This will be called the *dilution effect*.

An important characteristic of EGR gases from fired cycles is that they possess a high temperature. At the point of IVC - approximately 90 BTDC, the mixture has not yet undergone any compression. The in-cylinder contents will thus be composed of air at $\sim 300\text{K}$ and EGR gases. Assuming complete mixing, an EGR temperature of 900K and an EGR fraction of 15%, the final temperature of the mixture would be $0.15 \times 900\text{K} + 0.85 \times 300\text{K} = 390\text{K}$. (This analysis also assumes the heat capacity at constant volume, C_v , is the same for both in-cylinder compositions.) A very recently published communication by Devillers et al.¹²⁶ includes, in figure 6, a plot of integrated toluene emission against temperature. Using these data, the change in temperature that has just been calculated produces a reduction in signal of 40%. Taken with the 15% of the dilution effect, and allowing for the uncertain estimate of temperature rise, the signal drop in the case of toluene is entirely accounted for. Note that as the compression proceeds, the initial temperature difference is 'magnified', so the drop in signal strength is expected to increase, a trend which is demonstrated in figure 4.14.

With the dilution and temperature effects accounting for the signal drop in the

case of toluene, attention was turned to acetone (no literature exists for TMB, but its behaviour is not dissimilar to toluene). The same dilution effect is in operation for acetone fluorescence, leaving 40% signal drop yet to be explained at 90 BTDC. Similar temperature data for acetone is presented by Thurber et al.³³. Figure 3 of this paper indicates that for a nitrogen atmosphere, one must increase from 300K to 550K to get only a 5% drop in fluorescence yield. Section 3.4 discusses why ketone fluorescence does not have the same dependence on oxygen concentration as found in aromatics - and a paper by Braeuer et al.¹¹⁰ supports this, reporting that -

“...nitrogen and air test gases below 2-3 atm can be treated as equivalent from the standpoint of temperature dependences, as the fluorescence quenching effect of the oxygen in air remains small enough to cause little change in the relative temperature dependence data...”

Therefore, the major cause of the decrease in acetone fluorescence remains to be determined.

Regardless of the cause of the disparity between the motored and fired operation of the engine, a fired *in-situ* calibration method would incorporate this phenomenon and allow quantitative measurements to be made. The next section describes the solution which was developed and deployed in these experiments.

4.4.3. ‘Bookend’ calibration

In-situ calibration relies upon the use of a representative engine cycle under controlled conditions to calibrate the response of the tracer under test conditions. The previous section illustrates how a motored cycle is not sufficiently representative of a fired cycle, and thus a fired calibration method is desirable. Once again, although independently conceived, the method presented here appears to be similar to that reported by Han and Steeper³².

One potential drawback of a fired calibration system relates to the limited range of

equivalence ratios over which the engine will run. However, the calibration data from the motored calibration loop (figure 4.12) can be re-used, creating plots such as that of figure 4.15.

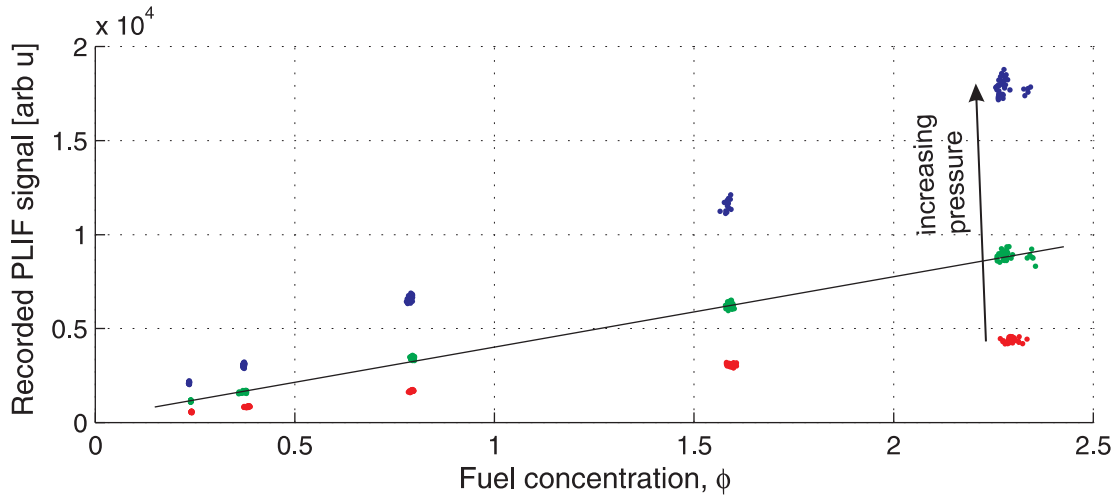


Figure 4.15.: A plot of PLIF signal as a function of equivalence ratio in the calibration loop, demonstrating linearity. Data presented for light fraction tracer, acetone, but all tracers exhibit this behaviour. The colours represent different pressures in the engine.

Figure 4.15 indicates that the PLIF signal intensity is highly linear with respect to change in fuel concentration, meaning that self-quenching i.e. fuel-fuel collisions are of negligible importance. The PLIF model, equation 4.2 can therefore be simplified to the following:

$$S = a \times \phi \times p^c \frac{A_{21}}{A_{21} + d \times p^e} \quad (4.3)$$

This equation shows that a calibration at a single fuel concentration is sufficient to describe PLIF behaviour. Self-quenching would probably be of more importance at higher fuel partial pressures - but at an equivalence ratio of 1, less than 2% fuel is present in the vapour.

The ‘bookend’ calibration scheme takes advantage of this linear behaviour to perform a fluorescence yield calibration in a firing engine with a single, stoichiometric, air/fuel mixture. The fuel injector mounted in the intake plenum is active during these calibration tests, in order to introduce a mixture of maximum homogeneity into

the engine. Any slight deviations in the air/fuel ratio of the mixture - i.e. drifts and other fluctuations from stoichiometry - are monitored throughout the calibration procedure, and are taken into account.

The ‘bookend’ character of the scheme relates to the use of two calibration sequences - one at the start of an experimental run, and another at the end. The combined data from these two calibration runs taken with stoichiometric fired plenum injection is used both to calibrate the fluorescence yield and to determine how the transmission of the windows has changed during the experimental run.

The fluorescence yield is calibrated using a method which is analogous to that of the motored calibration - images at a sequence of timings during the cycle are used to build a new calibration surface relating camera signal, fuel concentration and cylinder pressure by fitting to equation 4.3 in MATLAB[®].

Compensation for window transmission is an important feature delivered by the ‘bookend’ system, since the absolute intensity of the PLIF signal must be accurately measured to obtain meaningful results. From experience in running this engine, combustion deposits appear to come in three major forms. The first is combustion-generated soot, which has proven to be relatively minimal at the WWMP operating point. The second source of combustion deposits relates to fine particulate matter generated by the wearing of the piston rings (and, to a lesser extent, the edges of the sealing gaskets). These first two types of combustion deposits may be effectively removed from optical elements by use of a solvent such as acetone. More subtle and thus more concerning is the third class of combustion deposit. It is observed that over a period of time, a faint orange/yellow tint will accumulate on the optical surfaces that are exposed to the combustion event. It is not clear what the cause is, though some correlation has been seen between increased yellowing and extended motored operation with injection followed by firing. This coating has the unfortunate property of being strongly absorbing at the laser wavelengths, and so piston crown windows in particular have a limited lifespan.

The construction of the engine makes replacement of the pent-roof window an extremely time-consuming task, and while many combustion deposits can be cleaned off, the orange/yellow tint has not demonstrated a weakness to any solvent, requiring mechanical polishing to be removed. The optical annulus is more easily replaced, however, so the section of the image which views the engine through the annulus is used to reference the change in window transmission during a test. Figure 4.16 outlines the experimental sequence adopted in the ‘bookend’ calibration, and shows the initial and final derived window transmission images.

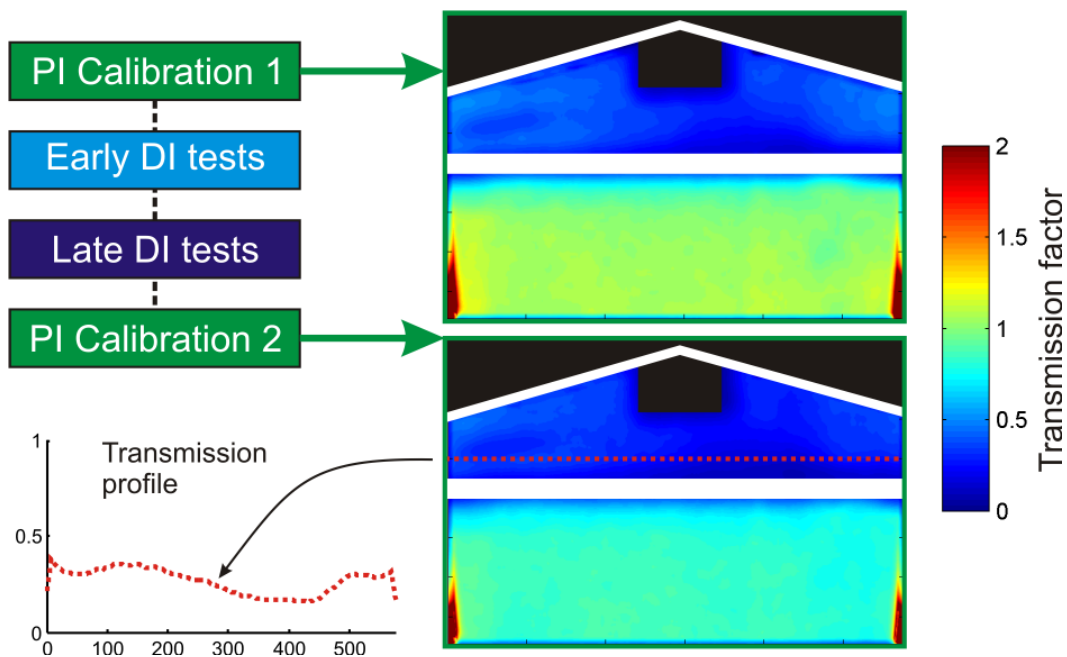


Figure 4.16.: Left hand side: ‘Bookend’ calibration method sequence. Right hand side: Window transmission correction images, initial and final, with colour bar. Annulus region of initial mean value = 1, final has mean value ~ 0.8 . Pent-roof window region has relative transmission of just ~ 0.4 owing to stubborn combustion deposits (see text). Bottom left figure is a plot of window transmission along the red dotted line crossing the second bookend correction image.

Figure 4.16 shows typical data, indicating that the pent-roof window is not only of reduced transmission relative to the clean annulus, but also that the combustion deposits are non-uniformly distributed upon its surface. The figure also indicates that the overall transmission of the windows is reduced by $\sim 20\%$ over the course of the

tests between the first and last calibration sequences, owing to the build-up of the various deposits. (Typically, 3-4 test conditions amounting to $\sim 10,000$ fired engine cycles are investigated between the bookends.)

In detail, the ‘bookend’ calibration records 32 images at 10 CAD steps in the compression stroke between 90 BTDC and TDC - covering, as usual, the entire range of relevant pressures. The engine is always operated at an equivalence ratio of 1, as indicated by the lambda sensor positioned in the engine exhaust. The fuel is supplied to the engine by use of an injector which is sited in the (heated) intake plenum, allowing maximum time for the mixture to homogenise. As with any other PLIF test, a background dataset is also collected at the same timings, in order to remove any counts recorded on the camera which are not due to tracer fluorescence. The recorded data are plotted in 3-D (pressure, equivalence ratio and signal), just like the motored calibration data, and the simplified equation 4.3 is used to find a best-fit solution with MATLAB[®]. This procedure essentially generates a new fired calibration surface, which looks qualitatively like those presented in figure 4.12, but differs in quantitative details from experiment to experiment as tracers and optical conditions vary.

One example of the ‘bookend’ calibration is shown in figure 4.17 for a toluene tracer test. Data from both bookend calibration runs are used, once the signal from the second dataset has been scaled to match the first dataset by comparison of the average signal levels in the annulus region of the recorded PLIF images. Note that a linearly-interpolated version of the window fouling correction image is used to scale the intermediate tests, before the calibration surface is used to convert the recorded PLIF intensity into fuel concentration.

4.4.4. Calibration summary

The *in-situ* calibration concept, employing homogeneous mixtures for uniform PLIF signal production provides an effective method of allowing for the variation of fluores-

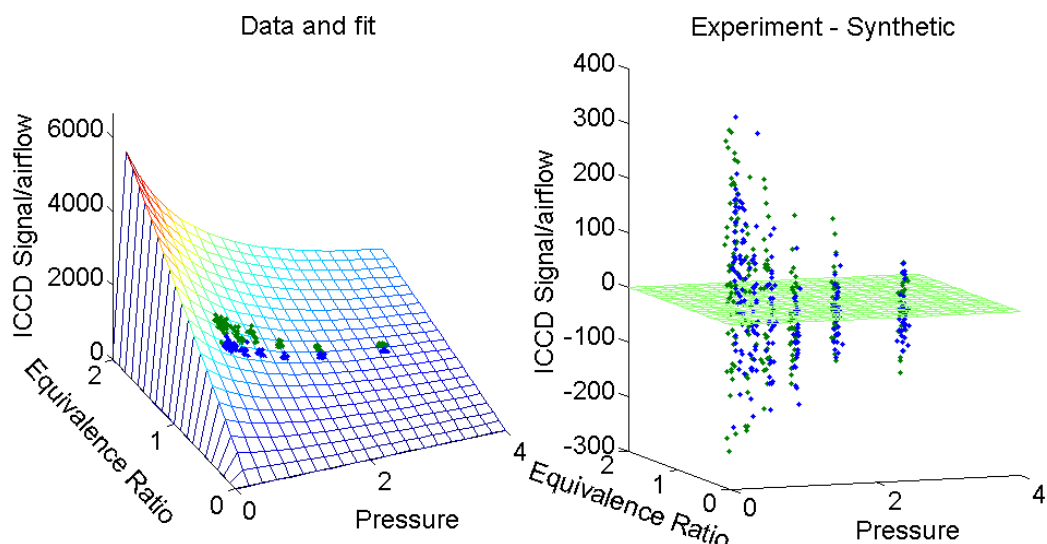


Figure 4.17.: An example of the data acquired during bookend calibration runs, and the associated fitting surface. Blue and green dots represent the normalised fluorescence signal strength taken from the two plenum injection tests conducted at the start and end of an experimental sequence. The right hand plot shows the deviation of the experimental data (points) from the model (green plane), demonstrating even scatter about the fitted surface.

cence yield of a tracer. However, implicit in the calibration is a particular temperature history. If this temperature is not uniformly achieved across the imaged region, then its effect will remain uncorrected, and shall be construed as a change in number density. There is evidence in the literature that the cyclic variability in temperature is relatively low (Einecke et al. ²⁸ suggests $\lesssim 20\text{K}$, including measurement uncertainty), and it will be assumed to be a negligible effect in this work.

The bookend calibration system not only accounts for the effect of fired operation on the overall fluorescence yield of the tracer under study, but also allows for the degree of window fouling to be determined. This is believed to be an important improvement over the technique presented by Han and Steeper ³². The assumptions made in using this calibration method are that the PLIF signal is linear with ϕ for constant T , p and composition (see figure 4.15), and that the annulus begins the test in good condition (any combustion deposits should uniformly coat the surface, if at all).

Note that the overall effect of fired operation is, approximately, to halve the magni-

tude of the recorded signal. This naturally has a deleterious effect on the signal-to-noise ratio achieved. Figure 4.12, presenting the calibration surfaces from motored conditions shows that even a very rich mixture of acetone only causes $\sim 30,000$ counts to be registered on the 16-bit intensified camera at maximum gain, using half its dynamic range. With the effect of EGR, and fuel concentrations closer to stoichiometry, the signal could be substantially lower. Whilst this signal reduction might be ameliorated by an increase in tracer doping levels, such an increase would inevitably lead to increased absorption, potentially complicating analysis.

The true beauty of the bookend calibration scheme lies not only in its ability to account for the effects of fired operation on PLIF signal strength, nor its window fouling correction, but also the fact that the system has very low requirements of the operator. The camera may be freely moved between tests, the absolute laser intensity does not have to be monitored with high precision[†] - and even slight changes in tracer doping levels do not matter. All of these things are automatically accounted for by this *in-situ*, online calibration method.

4.5. Analysis system

The analysis routines which convert the raw data from the line camera, data acquisition systems and intensified camera are implemented in a mixture of scripts and functions written for MATLAB[®] (a product of The MathWorks, Inc.) and an open-source java-based image processing tool called 'ImageJ', maintained largely by the National Institutes of Health (NIH) in the USA¹²⁷.

The great strength of MATLAB is the library of built-in functions which are well documented in the online help feature. Combined with the powerful integrated development environment for code generation, many data analysis problems can be handled. However, for simple repetitive tasks, the processing speed of ImageJ comes to the fore.

[†]Particularly if the intensity is well below I_{sat} .

Though ImageJ has no development environment to speak of, it is ideal for conducting simple tasks - like background subtraction, measuring mean and standard deviation in a region of interest and applying scaling factors or false-colour maps - much more quickly than MATLAB can achieve, particularly on large files.

Since a typical experimental dataset will consist of approximately 2GB of data of mixed origin, improvements in processing speed are often worthwhile. In the current iteration of the analysis system, a full dataset can be processed in around four hours. The analysis system leaves a kind of ‘accounting trail’ of data files and diagnostic information which consumes two to three times the original dataset size. While potentially wasteful on disk space, this allows the user to verify each stage has progressed correctly, or investigate where an unexpected result may have arisen.

Figure 4.18 provides an overview of the data analysis suite. Each block is labelled and colour coded, such that orange refers to raw source data, blue to scripts and functions written in MATLAB, and green to macros written for ImageJ. The process is roughly split into two columns, read from top-to-bottom and joined by the ‘IJ 1’ processing module. The directionality of the arrows indicates the ‘flow’ of data between blocks, and allows the resolution of the dependencies that any given block possesses.

An exhaustive description of all the elements in the analysis system may be of little benefit to the reader, and is beyond the scope of this thesis, so instead a brief description of each block will be given. These short descriptions will be followed by a more detailed discussion of the system which processes the signals recorded on the beam-shape monitoring line camera system.

The analysis procedure begins with the ‘HDAQ_red’ processing block, which consists primarily of the MATLAB script called `Red.HDAQ_LC1.PLIFIMG.m`. This script loads the high speed (‘HDAQ’) and low speed (‘LDAQ’) data acquisition files, together with the line camera laser beam profiles (‘LC1’) for each shot. The code determines which engine cycles were used to record PLIF images, and derives all the necessary values from the logged data (e.g. engine lambda, in-cylinder pressure). The system also

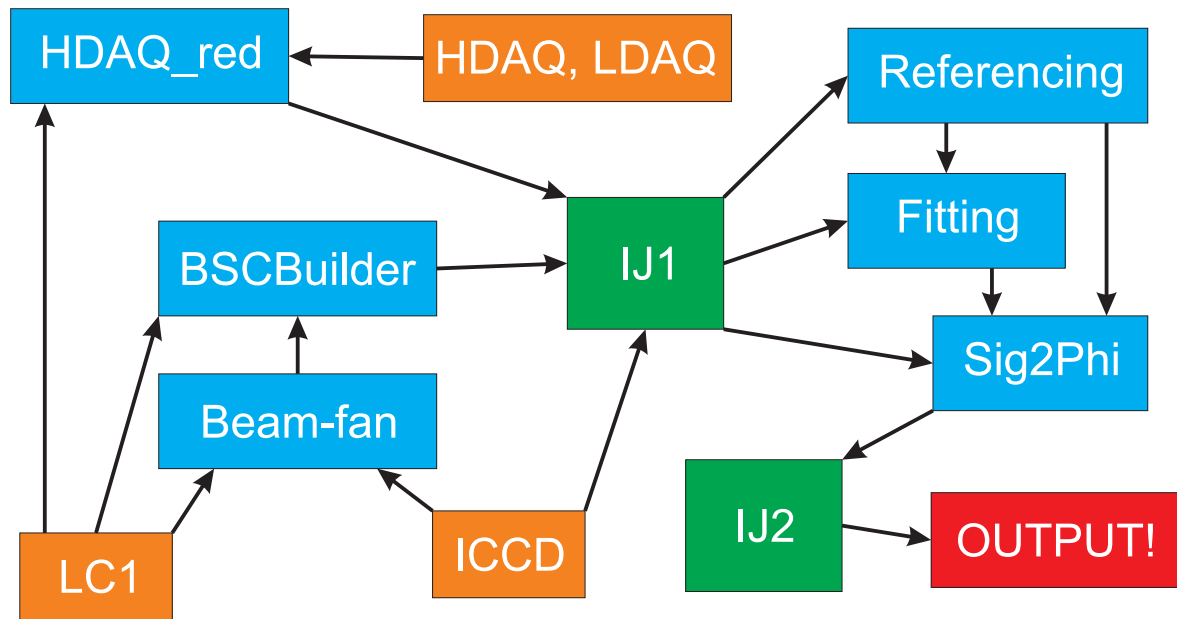


Figure 4.18.: Flow diagram, indicating the sequence of steps undertaken in order to analyse raw ICCD, LC1 and DAQ data acquired during an experimental run. Orange = data source, blue = MATLAB scripts or functions, green = ImageJ macros.

integrates the line camera ('LC1') data and calculates a normalisation value that will correct for laser pulse energy fluctuation. Any images recorded where the laser shot energy was outside of preset limits is automatically discarded. The normalisation values are stored in ASCII text files and passed on to the 'IJ1' block.

Laser profile correction images are produced by the combination of 'Beam-fan' (Minimiser_4.m) and 'BSCBuilder' (BSCBuilder.m). These two modules respectively a) determine the propagation of the laser sheet inside the engine from a reference image and then b) use the line camera data to create a custom correction image for each frame recorded by the intensified camera. The function of these two blocks will be considered in further detail presently.

Following the creation of the normalisation files and the beam-shape correction files, the 'IJ1' module is run. This block consists of four distinct steps, all executed in ImageJ. Firstly, the ICCD images recorded without fuel or tracer are averaged to produce background images ('BG'), such that there is a background image for every

imaged CAD position. The background images are then used to correct the signal ('SG') files by subtraction. In the third step, the laser sheet profile correction images are applied. In the final step, average fluorescence images are created for each visited CAD position. During these processing steps, one of the first operations performed on the raw data is the application of a 3×3 median filter (the size of which is justified by the discussion in section 4.3.3), which acts to remove some of the high frequency spatial noise on the PLIF images with minimal loss in resolution[†].

The processed images which now contain only fluorescent signal and have been corrected for laser pulse energy and beamshape fluctuations are then fed into three MATLAB components. The 'Referencing' block (`Referencing52.m`) is used to derive the window fouling correction images from the data, by referencing the initial and final bookend calibration tests against one another. The 'Fitting' block (`FitSurfPF2.m`) is used to take measurements from the two bookend calibration sequences and derive a best fit solution of equation 4.3. The 'Sig2Phi' block (`Sig2Phi_avg_Mk64.m` and/or `Sig2Phi_shot_Mk64b.m`) then applies this fitted equation to the PLIF data to produce quantitative fuel distributions.

The final block, 'IJ2', running in ImageJ, simply applies a false-colour map to the fuel distributions and effects a handful of changes to improve the aesthetics of the output images.

Having discussed the overall operation of the analysis system, a more detailed discussion is presented of the beam-shape correction sub-system, as denoted by the 'Beamfan' and 'BSCBuilder' functional blocks. Figure 4.19 illustrates the operation of the beam-shape correction system.

Firstly, the MATLAB script `Minimiser_4b.m` loads a reference image from the ICCD and a reference beamshape from the line camera. This reference test is conducted every time the experiment is performed, in order to account for any shift in the positions

[†]A median filter is used in preference to - for example - a moving average, because median filtering tends to retain edges, where other filters effectively blur detail away

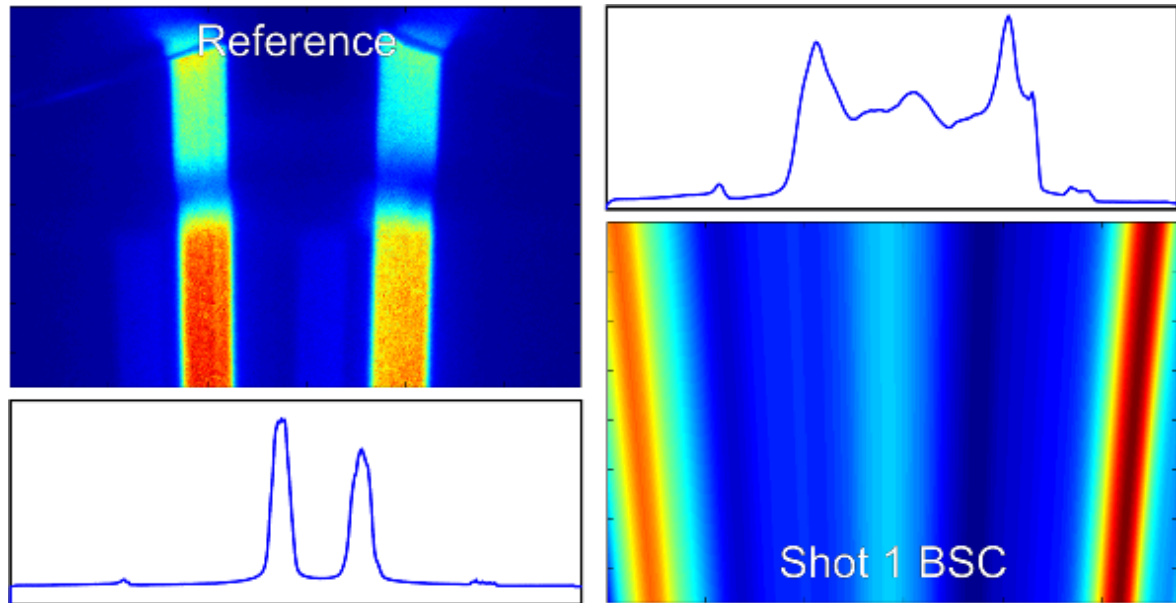


Figure 4.19.: Left hand side: reference image and beamshape used to determine the propagation of the beam in the engine. Right hand side: the correction frame produced by applying the propagation information to the single-shot beamshape. (BSC = beam shape correction.) ‘Hotter’ colours indicate higher values. Note, the images are presented on different scales.

of any optical components. These reference images are shown in the left hand side of figure 4.19, and are created by placing a mask with two slits in the path of the laser beam, and sealing the combustion chamber with a small pot of acetone placed inside. The acetone evaporates, forming a saturated vapour, which fluoresces strongly when excited by the laser beam. Note that the gentle divergence of the two stripes is caused by the second cylindrical lens drawn in figure 4.6, which is used to ensure the beam is wide enough to fill the field of view of the camera. The code in `Minimiser_4b.m` then uses the functions `localMaximum.m`[†] and `deterp30.m` to determine how the centres of the two peaks in the LC1 data spread under a range of stretch and offset factors. The function `LHA.m` is then used to find the location of the centre of the two stripes in the ICCD frame as a function of height. By comparing these two datasets, the optimum stretch and offset of the line camera profile can be determined for each pixel row in the image. This process is not without noise, so a further function `makelinparms.m` is used

[†]Downloaded from the [MathWorks website](#), developed by Yonathan Nativ.

to make a linear fit to the stretch/offset parameters. The resulting linear fits describe the propagation of the laser sheet inside the engine, allowing the line camera profile to be matched against the ICCD image.

In the next step, the MATLAB script `BSCBuilder.m` is used to produce the beam shape correction images, as shown in the right hand side of figure 4.19. This process employs functions `BeamImg.m` to create the image, row by row - with a stretched and offset image created according to the linear fit created by `makelinparms.m` using the function `deterp3.m`.

The overall result of these steps is to create a batch of multi-frame TIFF[‡] files, where each TIFF frame includes a beam shape correction (BSC) image that is generated by using the relevant measurement from the line camera data. These BSC images are employed in module ‘IJ1’, to divide each image by the spatial variation of laser intensity.

For the sake of completeness, it should be mentioned that the analysis software can run in an alternative mode, using `plenumBS_4.m` in place of `BSCBuilder.m`. This module emulates the line-camera laser sheet profile by using the data about beam propagation from `makelinparms.m` to derive a profile from the homogeneous-mixture plenum injection bookend tests, rather than the shot-wise data recorded by the line camera. `plenumBS_4.m` then interpolates these two ‘fake’ line camera laser sheet profiles to approximate the laser beamshape in the intervening tests, and creates a set of BSC images. For a stable excimer beamshape, this method can produce slightly higher quality images. This is potentially because the accuracy of the LC1 system is hampered if the Bowditch mirror and/or the piston crown window are not uniformly transmitting. In this second mode of operation, the LC1 data is still used for pulse energy fluctuation corrections - and is far preferred for doing so than sampling the beam with a small-area photodiode.

[‡]TIFF (Tagged Image File Format) files store image data in an extensible, flexible format. Commonly supported options include the use of either no compression or lossless LZW compression, with anything up to 32-bit precision per pixel.

4.6. Testing QPLIF

This section is dedicated to presenting results which test and characterise the QPLIF technique, as applied with bookend calibration and a multi-component fuel to the Jaguar optical engine. The first topic to be covered in this section is that of how to present quantitative data. The second sub-section is concerned with assessing the performance of QPLIF against expectations. The third and final section describes a method for estimating the accuracy of the QPLIF technique as applied.

4.6.1. Quantifying PLIF signals

The experimental system, incorporating bookend calibration and analysis routines is able to convert a PLIF image into a quantitative map of fuel distribution. It is important to consider how to interpret the values the system generates.

Since the quantitative calibration is conducted relative to a reference point at stoichiometric operation, and the PLIF signal has been calibrated to derive tracer concentration linearly, it is tempting to simply consider the derived values to be the equivalence ratio (ϕ) of the mixture at any given location. However, since only one tracer is investigated at any one time, only one of the three fuel fractions has a known concentration. Since the stoichiometry of the mixture is dependent on the mean composition of all the present hydrocarbons (see equation 2.2), this value is not the mixture's equivalence ratio. It is, in some sense, the equivalence ratio of that particular fraction - but this terminology can lead to confusion. Therefore, this thesis chooses to use the term 'fuel fraction concentration' (FFC). The meaning of a particular value of FFC is simple: if the FFC is evaluated to be 1.5 in a particular volume for a particular fraction, then the concentration of that fuel fraction is 1.5 times greater than the concentration that fuel fraction had in the stoichiometric mixture.

Note that if the FFC values of all three fractions were to coincide at some value x , then the mixture would have the same composition as the stoichiometric mixture in the

reference test, and therefore the equivalence ratio would indeed simply be x . However, in the general case, the overall equivalence ratio of the fuel can only be determined by combining the FFC values of each fraction.

4.6.2. Verification of QPLIF

The QPLIF method is verified by confirming that the behaviour of early direct injection (early-DI or EDI) is as expected. For early-DI, the fuel injection event occurs at 280 BTDC. This provides 27 ms for the fuel to be mixed with the in-cylinder air by the turbulent flow field which exists inside the engine. It is therefore expected that a homogeneous distribution will be in evidence at the time of ignition.

This expectation is tested in figure 4.20 and table 4.4. Figure 4.20 is a plot of the mean fuel fraction concentration (FFC) in a large region of interest[†] (ROI) in the quantitative PLIF images against the value recorded on the lambda sensor. A green line denotes ideal behaviour, where the two measurements are in perfect agreement. The upper and lower red lines indicate the bounds of $\pm 10\%$ agreement, which could be considered as acceptable limits.

Figure 4.20 shows the data points from the two bookend calibration tests clustered tightly around the line of equality. This behaviour is expected, since the calibration will cause this to occur. Note how the calibration mechanism has compensated for the variation in engine lambda between *and during* calibration tests. Of greater interest is the longer trail of data points which corresponds to the EDI data. The image timing information is lost in this form of presentation, but the approach of the EDI data points towards the line of equality is strongly tied to the ongoing compression and mixing processes (i.e. in progression to lower CAD BTDC). This confirms that the QPLIF technique is behaving as expected, showing that homogeneous stoichiometric mixtures (all fractions exhibit this behaviour) are produced by early-DI.

[†]This ROI is illustrated on the CFD distributions shown in figure 4.27 as a rectangle with a thin black outline.

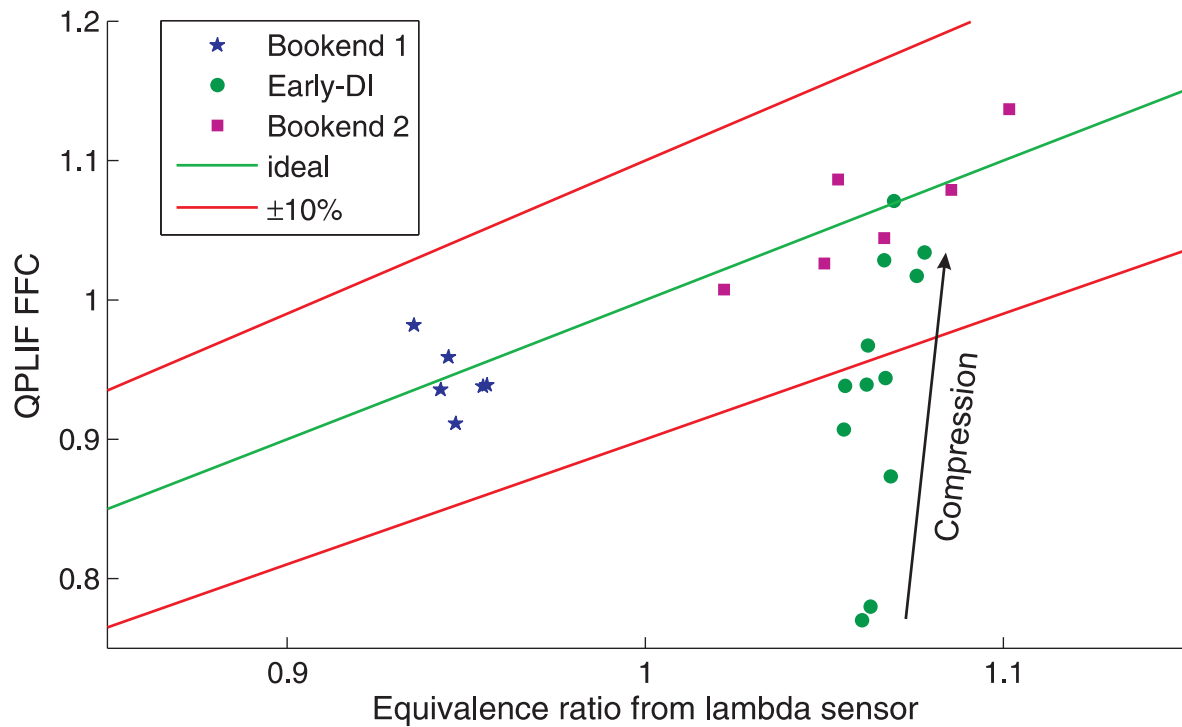


Figure 4.20.: FFC vs lambda sensor for toluene tracer. Plot shows data from the two bookend calibration measurements, and an intervening early-DI test. The green line denotes equality, and the red lines show the $\pm 10\%$ limits.

The mean FFC values in a large ROI are presented in table 4.4. Each successive row is a step closer to TDC-compression. A column is given for each fuel fraction, and the two values in each cell correspond to the mean and standard deviation of the sample distribution of the average value in the ROI in the 32 QPLIF images (that is to say, the values are derived in the same way as the data points in figures 4.13 and 4.14).

Table 4.4 lends further justification to the assertion that the QPLIF technique works. Each fraction is seen to produce a homogeneous stoichiometric mixture just prior to ignition. Note that the double horizontal line in the table represents ignition timing. Following ignition, there is a dramatic rise in pressure and temperature, and this part of the cycle is not included in the modelled calibration surface. At this time, the combustion event will consume the tracer, as can clearly be seen for the acetone measurements, where the FFC value dips below unity. This same behaviour is not seen for the aromatic tracers toluene and TMB. In these two cases, the FFC value

| CAD BTDC | Acetone | | Toluene | | TMB | |
|-------------|----------|-----------|----------|-----------|----------|-----------|
| | mean FFC | std. dev. | mean FFC | std. dev. | mean FFC | std. dev. |
| 60 | 1.05 | 0.18 | 0.91 | 0.17 | 1.17 | 0.15 |
| 50 | 1.12 | 0.19 | 0.98 | 0.18 | 1.06 | 0.15 |
| 40 | 1.15 | 0.18 | 1.02 | 0.19 | 1.17 | 0.20 |
| 30 | 1.19 | 0.17 | 1.05 | 0.24 | 1.28 | 0.24 |
| 20 | 1.10 | 0.19 | 1.12 | 0.26 | 1.33 | 0.27 |
| 10 | 0.74 | 0.26 | 1.26 | 0.35 | 1.47 | 0.30 |
| 0 | 0.93 | 0.33 | 2.68 | 0.62 | 2.04 | 0.41 |

Table 4.4.: FFC values for all three fractions and early-DI. Ignition occurs at 35 BTDC. Lambda sensor readings remained within 5% of stoichiometry during tests.

appears to rise, which is entirely non-physical. The explanation for this phenomenon probably lies in the increased production of poly-cyclic aromatic hydrocarbons (PAH) with these tracers present. PAH molecules fluoresce strongly, and are linked with soot formation. Note also that the standard deviation of FFC increases post-ignition, as the cyclically-variable flame introduces scatter.

4.6.3. Accuracy and precision of QPLIF

The imaging of homogeneous distributions of known equivalence ratio is an ideal opportunity to attempt to assess the precision and accuracy of the QPLIF method. Figure 4.20 includes data points derived from homogeneous PI mixtures. These tests indicate that the QPLIF technique should not suffer from any systematic offsets degrading its accuracy. The distribution of the average images about the line of equality suggests that $\lesssim 5\%$ scatter remains in the dataset, which must be related to effects such as residual uncorrectable scatter in laser pulse energy, affecting all pixels in the entire frame equally.

Figure 4.21 illustrates the justification for the accuracy estimate given here. Accuracy estimation begins using a plenum injection fuel distribution, where nominally the same fuel concentration should be present at every location in the field of view. (Left of figure 4.21.)

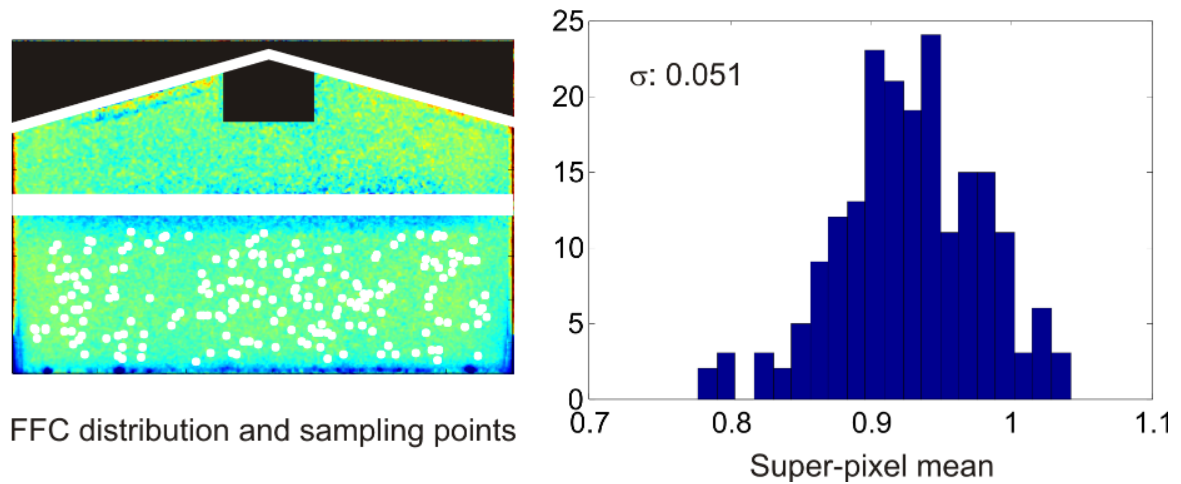


Figure 4.21.: Accuracy estimation, employing 200 randomly-chosen super-pixels. Left hand side: example PI fuel distribution and sampling points in the annulus region. Right hand side: distribution of super-pixel mean values

Next, a large number of sampling locations are randomly chosen, as demonstrated in the overlay on the left hand side of figure 4.21. A ‘super-pixel’ is defined as a 3×3 pixel matrix centred at each sampling point. Each super-pixel has an associated mean value across its 9 pixels, and the group of super-pixel mean values forms a distribution, such as that plotted in the right hand side of figure 4.21. Assuming this distribution to be normal, a standard deviation can be calculated - 0.051 in the case of the figure.

By running this analysis approximately 5 times each on a range of crank angles with PI fuel distributions (causing a new set of sampling points to be chosen at random each time), the standard deviation of the mean super-pixel value is consistently found to be in the range of 4-6%. To the extent to which each super-pixel represents an independent measurement of the same fuel concentration, this value is taken to be a good estimate of super-pixel accuracy.

Each super-pixel is formed from 9 pixels. The standard deviation of the pixel-wise FFC value can be calculated across these 9 pixels, yielding a super-pixel standard deviation. Doing this for all sampled super-pixels individually and then averaging produces a mean super-pixel standard deviation. Dividing this mean super-pixel standard deviation by the square root of the number of pixels per super-pixel (3), derives the

standard error. This produces a precision estimate on the mean FFC value for any given super-pixel of 1%.

Combining this estimation with that derived from figure 4.20, the overall accuracy on a super-pixel in an average fuel distribution can be estimated at 10%.

Note that since normal distributions have been assumed throughout, that this 10% error estimate means that 68% of the time, the super-pixel measurement of FFC will lie within $\pm 10\%$ of the true value.

4.7. Applying QPLIF

In this section, the QPLIF technique is applied to more challenging problems involving imaging fuel distributions produced by direct injection. Some consideration will also be made regarding how the quantitative data may be used. This section closes with a direct comparison between measured fuel distributions from QPLIF and Jaguar-calculated CFD fuel distribution predictions at the same operating point.

4.7.1. QPLIF imaging

The ETCS system governing the synchronisation of data acquisition and the timing of various engine events is able to schedule the capturing of QPLIF images with high flexibility. A typical arrangement is illustrated in figure 4.22. This figure shows how a full set of PLIF images is collected by capturing sets of 32 images at 10 different points during the compression stroke. (Note that for each such fired test, a motored test exists which is used for background subtraction.) Injection and ignition timing are also indicated on this diagram.

Figure 4.22 illustrates the case of ‘90:10:0’ with 32 ‘snaps’, where 32 images are taken every 10 CAD from 90 BTDC to 0 BTDC, for a total of 320 images. An example test schedule for a QPLIF experiment is included in appendix A.2. (Note that appendix A.1 describes the preparatory steps for a typical QPLIF experiment in detail.)

Figure 4.22 highlights that average QPLIF distributions are the most commonly used output of the QPLIF analysis. This is because individual frames are recorded at very low signal levels (~ 2000 counts on a camera which digitises up to 65,536 counts - 16-bits), and are strongly affected by shot noise as a result. As previously discussed, the adoption of a 3×3 median filter helps to overcome this issue without losing any spatial resolution, but the use of averaging is also very valuable. In this thesis, average images are those created by averaging across all shot-wise images taken in one experimental run at a particular crank angle.

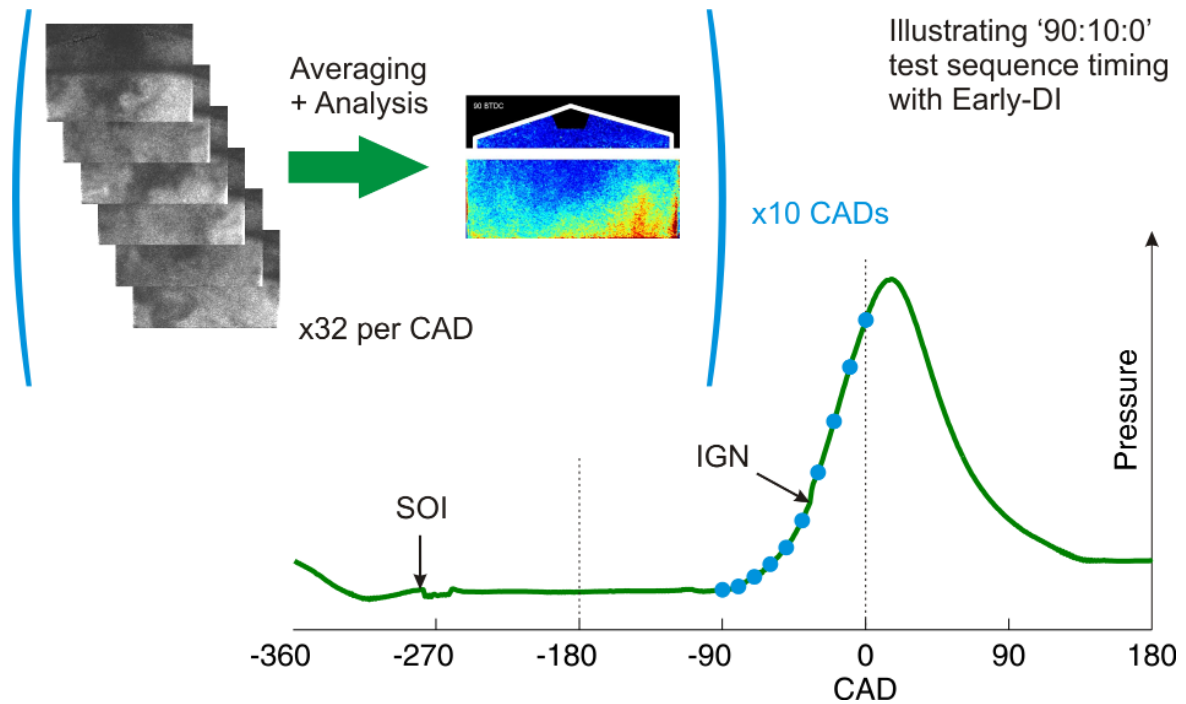


Figure 4.22.: Illustration of image acquisition timing in a QPLIF experiment. Acronyms are explained in section 2.2.

4.7.2. Average distributions for Late-DI

Following the successful testing of QPLIF in the previous section, the system is now applied to the measurement of a more challenging condition - WWMP with late DI (i.e. injection 140 BTDC). This operation condition is not suggested by Jaguar, nor will it be used in their production engines. However, it is expected that much higher stratification will be visible at the time of ignition, so any differences between tracers can be more easily studied. This late injection timing represents the most retarded timing possible, whilst retaining less than 8% CoV IMEP - i.e. this is the latest that injection can be delayed before cyclic variability becomes overwhelming, ultimately leading to the production of many misfiring cycles.

Figures 4.23 and 4.24 present a matrix of images from the LDI condition. The leftmost column indicates the angle that the distribution is recorded at. There are then three columns - one for each tracer. Each distribution is an average. The image representing each distribution is false-coloured and has had an overlay applied which

outlines the edges of the pent-roof window in white (roughly triangular region at the top), and includes a horizontal thicker white line which represents the gasket which seals the fused-silica annulus from the cylinder head. A black mask is applied over the location of the spark plug, and the colour bar which relates image colour to FFC value is given at the bottom of the figure. Note that the black mask on the spark plug is important to cut down on the high levels of scatter and fluorescence in this area owing to the laser beam terminating on the cylinder head very close-by. Close scrutiny of figures 4.4 and 4.9 will reveal the opaque square on the pent-roof window which blocks this light from entering the camera lens.

When interpreting these distributions, it is important to bear in mind that the measurement only provides data on fuel concentration in a two-dimensional plane, ~ 1 mm thick. Nothing is known about what may be occurring outside this region, other than that similar stratification is expected elsewhere.

Figures 4.23 and 4.24 demonstrate the transport of fuel within the cylinder. The fuel spray (which started at 140BTDC and lasted for less than 10 CAD) is seen to pass through the field of view, probably impinging on the piston crown. As the piston rises, the combination of the flow created by induction through the intake valves and the momentum of the spray jet induces an anticlockwise ‘wall-guided’ flow, which carries fuel vapour up the right hand side of the cylinder, and over the exhaust valves towards the spark plug. A roughly stoichiometric mixture is seen to be present at the spark plug at around the time of ignition, allowing the combustion process to occur. Note that the rising cut-off in the images recorded from 50 BTDC to TDC is caused by the piston encroaching upon the field of view of the camera, eventually reaching the top of the annulus at ~ 10 BTDC.

Interestingly, the three fuel fractions exhibit very similar behaviour, despite their clearly different boiling points and molar weights. This discovery indicates that the vaporisation process which is in effect in the injected fuel spray is relatively insensitive to differences in these physical properties. Note also that in this case, where the

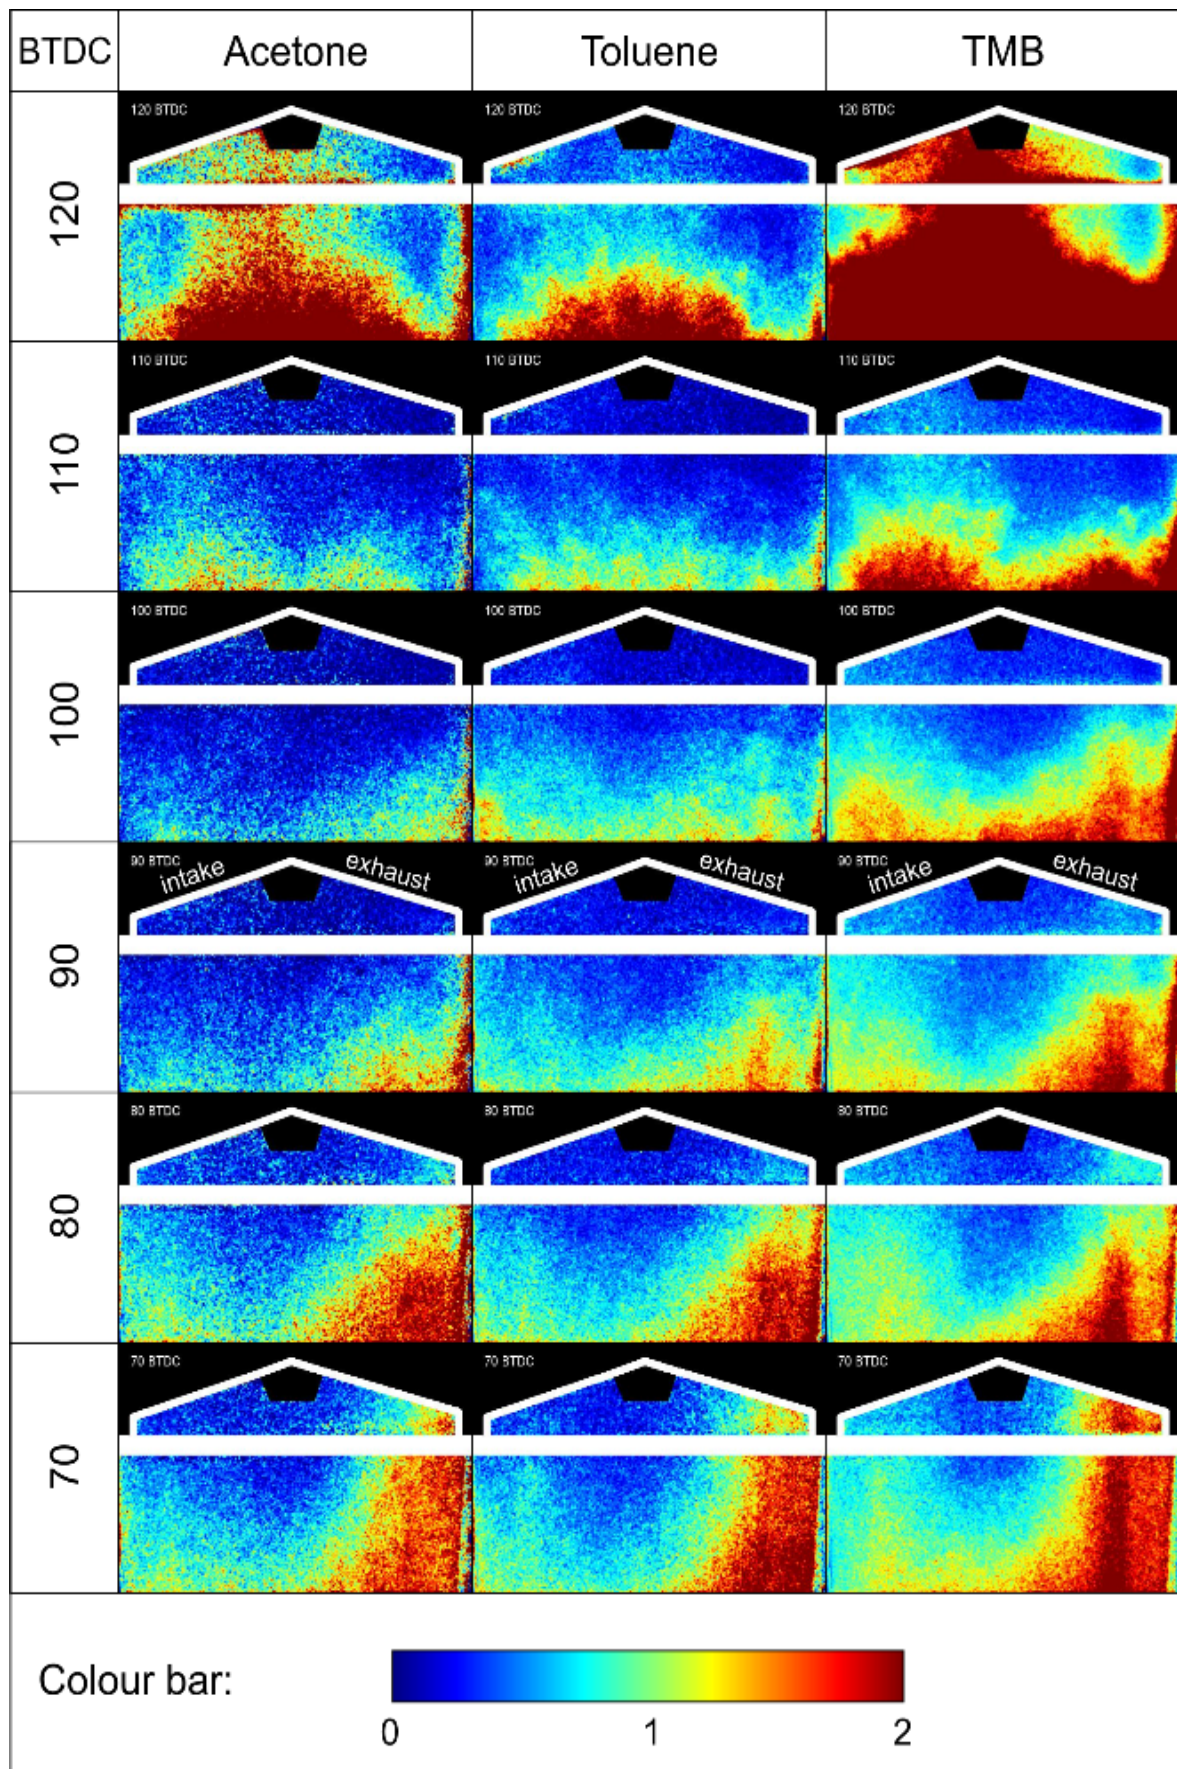


Figure 4.23.: QPLIF Late-DI distributions. Colour bar in Fuel Fraction Concentration (FFC). (1/2)

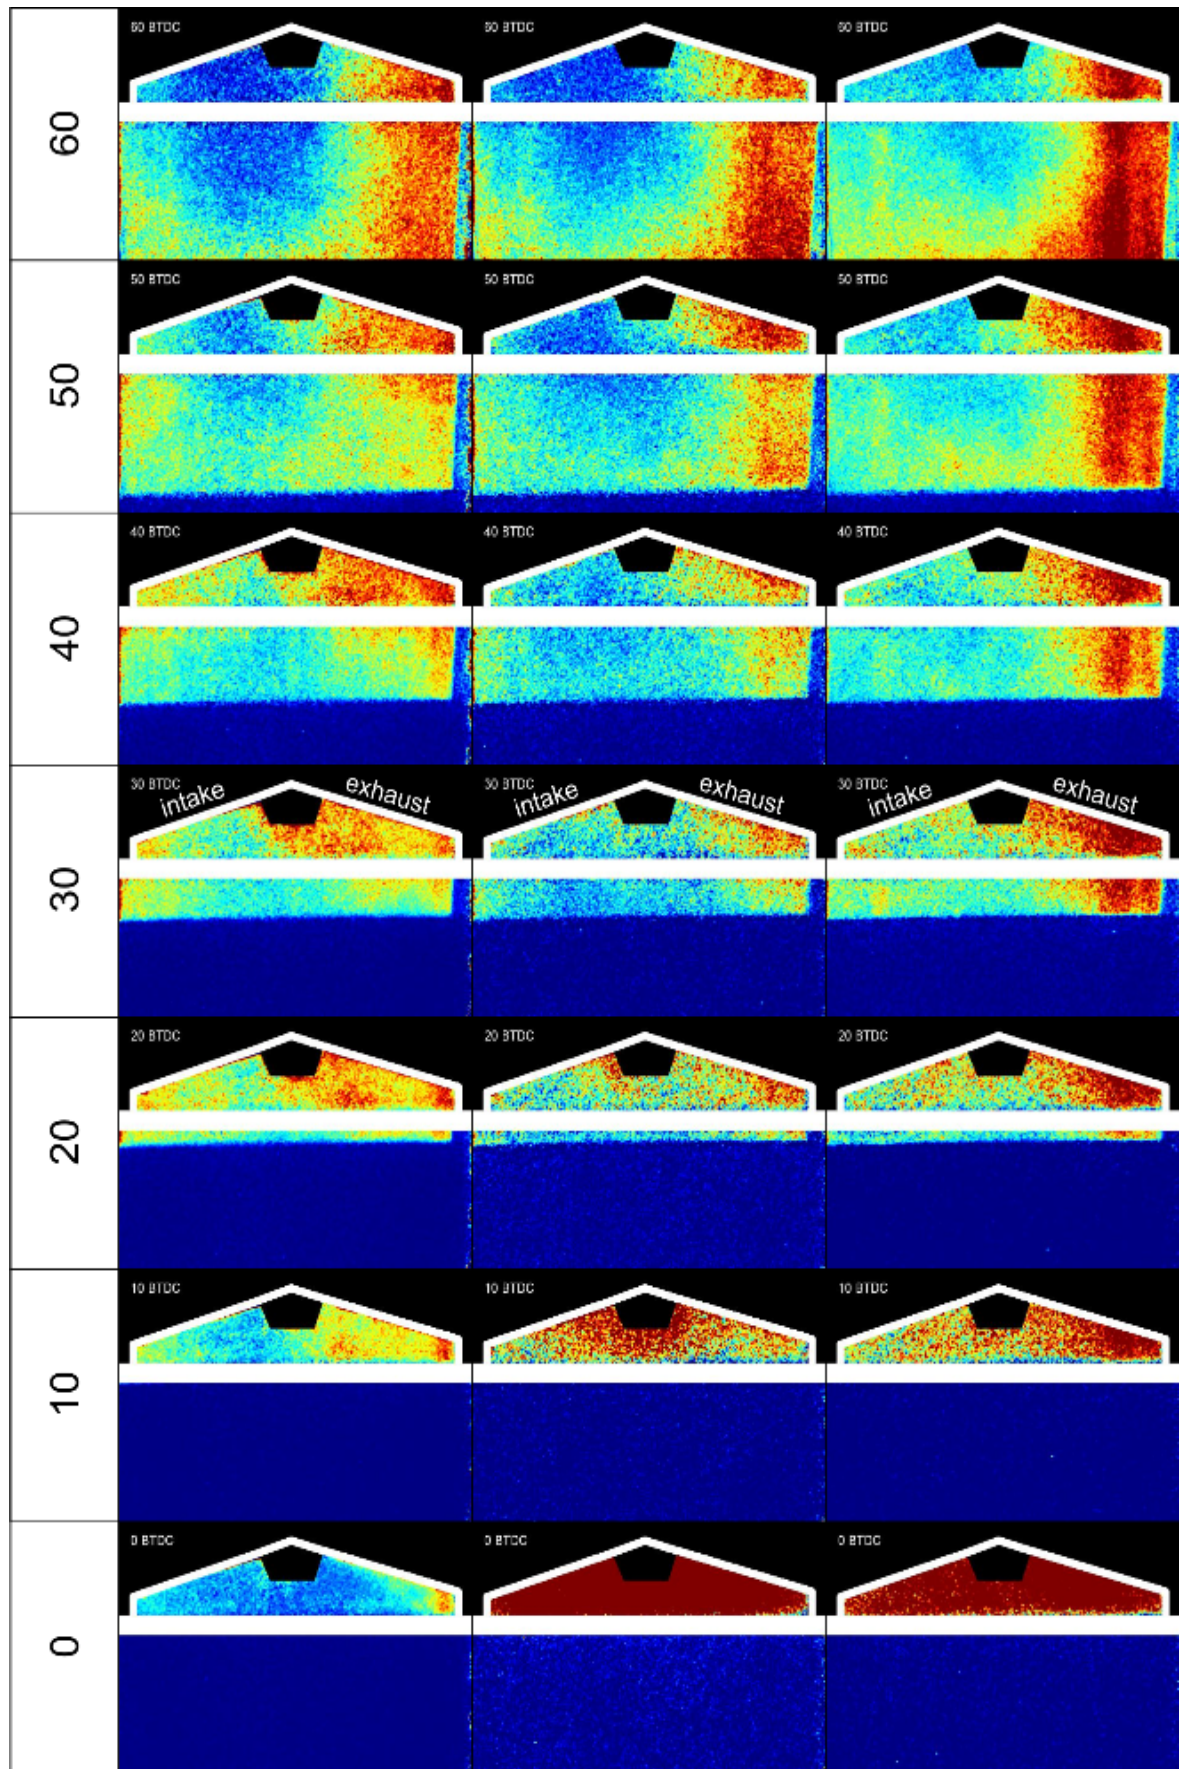


Figure 4.24.: Matrix of QPLIF Late-DI distributions. (2/2)

FFC value of each fuel component is virtually identical, the FFC value becomes the equivalence ratio, ϕ , of the mixture at that location.

4.7.3. Application of CoV to QPLIF

The low fluorescent signal levels which are characteristic of quantitative PLIF experiments (as a result of sub-saturation intensity operation and low doping levels) do not immediately render the investigation of cyclic variability impossible. (Appendix B includes some examples which indicate just how noisy these images are.) Even in the case where signal levels are bountiful, a single shot PLIF image is of little use by itself. The answers to questions such as: ‘is this a repeatable distribution?’ and ‘where is the distribution most variable/stable?’ can only be answered by considering the behaviour of an ensemble of such shots. This naturally leads to the formation of an average, which is where QPLIF is most usable anyway.

Cyclic investigations with QPLIF are aided if the spatial resolution requirements are relaxed. In the distributions presented in figure 4.25, a moving average filter of 6.5 pixels diameter is employed, leading to a reduced spatial resolution of approximately 1.5mm. Note that this is approximately the same as the width of the spark plug ground electrode, so the resulting distribution is still of use.

Having formed the average fuel distribution by taking the pixel-wise mean of all images taken at a particular crank angle, an analogous standard deviation image can be retrieved by calculating the pixel-wise standard deviation from the same dataset. This standard deviation will be a measure of both fuel delivery fluctuations and noise introduced by the measurement method and analysis system. However, the scatter induced by the measurement will be uniform across the field of view, unlike that caused by fuel delivery, allowing the two to be separated, to some degree.

The ‘coefficient of variation’ or CoV is defined as the standard deviation of a population divided by its mean, and is therefore a normalised measure of variability. If

the standard deviation image is divided on a pixel-wise basis by the average image, the CoV image is calculated. The CoV image records in each location the percentage variability of the distribution. Areas in the fuel distribution which are consistently delivered the same quantity of fuel will have a low CoV (e.g. a value of 0.1, or 10% fluctuation in FFC). Areas with CoV values as high as 0.6 would be locations where fuel delivery is highly variable, with a shot-to-shot fluctuation of 60% of the mean.




Figure 4.25.: CoV images of QPLIF fuel distributions, produced by performing pixel-wise calculation of standard deviation divided by mean.

With this description in mind, figure 4.25 presents the CoV distributions for the middle fuel fraction using plenum injection and early-DI (injection at 280 BTDC) at 60 BTDC, alongside late-DI (injection at 140 BTDC) images at 60 BTDC and 80 BTDC. The colour bar in the figure encodes the CoV value from zero to one.

Figure 4.25 indicates that the inherent variability of the early-DI condition is all but identical to that of plenum injection - with a slight advantage to PI. This probably indicates that the measurement technique is becoming the limiting factor in determining the scatter in the shot-to-shot measurement reproducibility. By comparison of the late DI images at 80 BTDC and 60 BTDC, the degree of cyclic variability is seen to reduce

as the mixture becomes more uniformly mixed. In general, the regions of highest variability are seen to be those on the edges of high average FFC areas (refer to figures 4.23 and 4.23 for average late-DI distributions). It should be noted that the division process responsible for producing CoV from standard deviation and mean must be interpreted with caution in areas where the mean value is close to zero, since this will tend to amplify noise levels and produce spurious results. The size of the smoothing mean filter applied to this data has been chosen to attempt to minimise such problems, but the data in the pent-roof window in the 80 BTDC late-DI image remains less reliable. Appendix B includes presentation of a similar analysis to qualitative images of sprays taken at much higher signal levels, showing much less noise.

4.7.4. Correlations with burn-rate analysis

The analysis reported in this subsection was performed by Xiaowei Wang, a DPhil graduate in Engineering Science from Professor Richard Stone's Internal Combustion Engines research group. The combustion burn rate analysis software (CoBRA) described in this section was developed by Hongrui Ma (another DPhil graduate from Professor Stone's group). The QPLIF data were collected and analysed into FFC distributions by the author.

This analysis proceeds by measuring the mean FFC in a region of interest which surrounds the spark plug like an arc, with inner radius 6mm and outer radius 10mm (i.e. within the pent-roof window). The pressure trace measured on the high speed DAQ system is then supplied to CoBRA, the burn rate analysis software, which determines the rate at which the air/fuel mixture is consumed by combustion.

The sub-plots in figure 4.26 demonstrate the relationship between the mean FFC value in the region immediately surrounding the spark plug, and the early burn rate of the flame. The y axis on these plots, 0-10% mass fraction burned, is the time taken for 10% of the combustion event to occur in crank angle degrees.

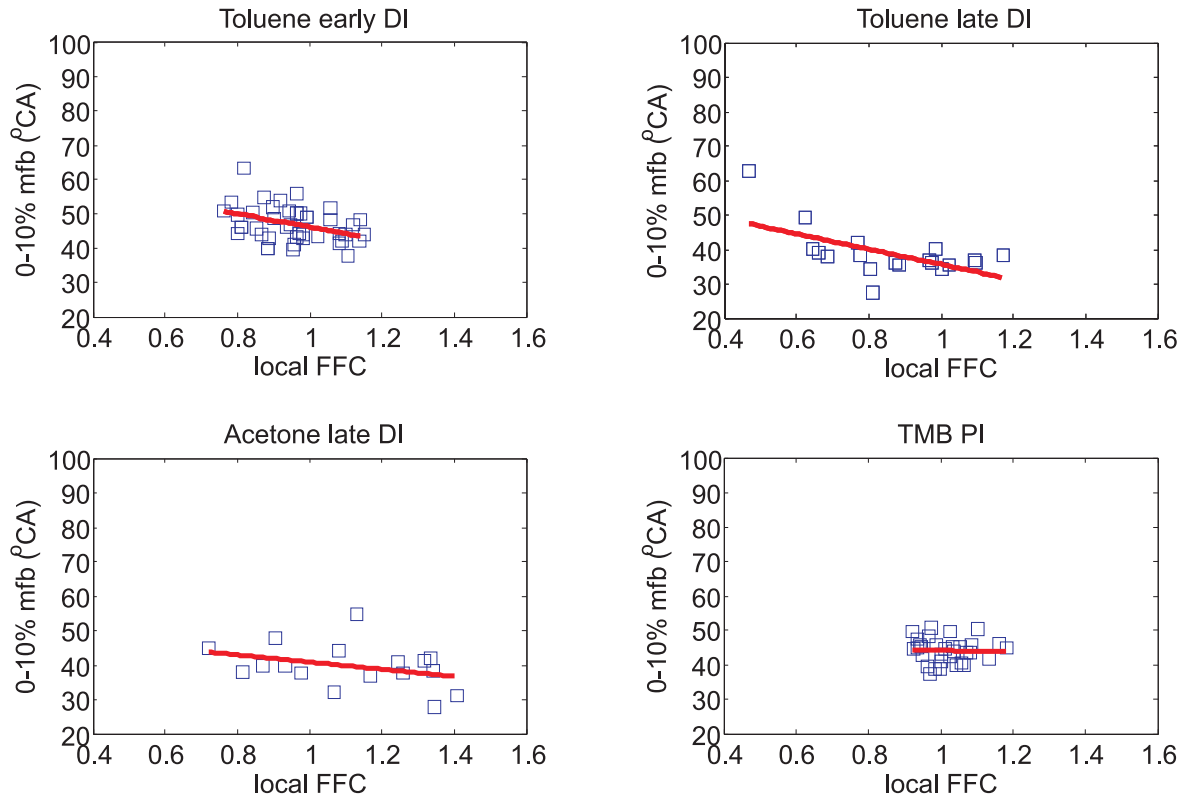


Figure 4.26.: Plots illustrating the correlations between local fuel concentration at ignition and initial burn rate.

Considering firstly the relative widths of the distributions in local FFC, the repeatability of plenum injection (PI) is seen to be better than early-DI, which in turn is better than late-DI. The three DI tests also present a weak correlation, where a richer mixture strength causes a faster burn. The lack of spread in the plenum injection data makes the resolution of any such trend challenging. Note that all fractions of the fuel exhibit similar behaviour - i.e. having indistinguishable influence upon early combustion rate - as illustrated by the upper right and lower left plots.

4.7.5. Comparison with Jaguar CFD

In this subsection, measured QPLIF distributions are compared to predicted distributions produced by CFD (computational fluid dynamics) calculations conducted at Jaguar Cars UK by Dr Huiyu Fu. The interpretation of the raw three-dimensional

mesh of CFD data to produce planar fuel distribution images was jointly undertaken by the author and a summer student in the Physics department, Thomas Harty.

Dr H. Fu provides the following description of the Jaguar CFD modelling system:

The CFD simulation was carried at Jaguar Land Rover Powertrain Analysis. The Ford in-house CFD code MESIM (**M**ultidimensional **E**ngine **S**imulation) was used to simulate three-dimensional, transient and compressible two-phase engine flow. MESIM is an extension of KIVA-3V*[†] with enhanced dynamic mesh algorithm[‡] and spray models[§]. The mesh movement methodology was developed for a generic pent-roof chamber, 4-valve and bowl in piston engine geometry. The enhanced spray models consist of the detailed physics including fuel injection, spray atomization, droplet vaporization, spray/wall impingement and turbulent fuel/air mixing. The models have been validated with optical engine experiments. Turbulence was modelled by the RNG k- ϵ turbulence model. Single component fuel, iso-octane, was used in the simulation to represent gasoline used in engine test. This may lead to somewhat inadequate representation of the detail structure of spray due to the difference in vaporization characteristics between the two fuels. Computation was started 20° before intake TDC and terminated 20° after compression TDC.

Having conducted the computer modelling, Dr Huiyu Fu of Jaguar Cars UK very kindly supplied a dataset describing his predictions for the late-DI WWMP operating condition. This dataset included a predicted fuel distribution at 20 CAD intervals during the cycle, from injection to TDC. For the purposes of clarity, only the distributions

*A. A. Amsden, “KIVA-3V: A Block-Structured KIVA Program for Engines with Vertical or Canted Valves”, Los Alamos National Laboratory Report LA-13313-MS, 1997

†A. A. Amsden, “KIVA-3V, Release 2: Improvements to KIVA-3V”, Los Alamos National Laboratory Report LA-13608-MS, 1999

‡J. Yi and Z. Han, “Methodology of Rapid Mesh Generation and Dynamic Mesh Management with Moving Valves for Engine CFD”, Ford Research Report, SRR-2001-0268

§Z. Han, Z. Xu, S. Wooldridge, J. Yi and G. Lavoie, “Modelling of DISI Engine Sprays with Comparison to Experimental In-Cylinder Spray Images”, SAE 2001-01-3667.

at 80 BTDC, 60 BTDC and 40 BTDC are displayed in figure 4.27.

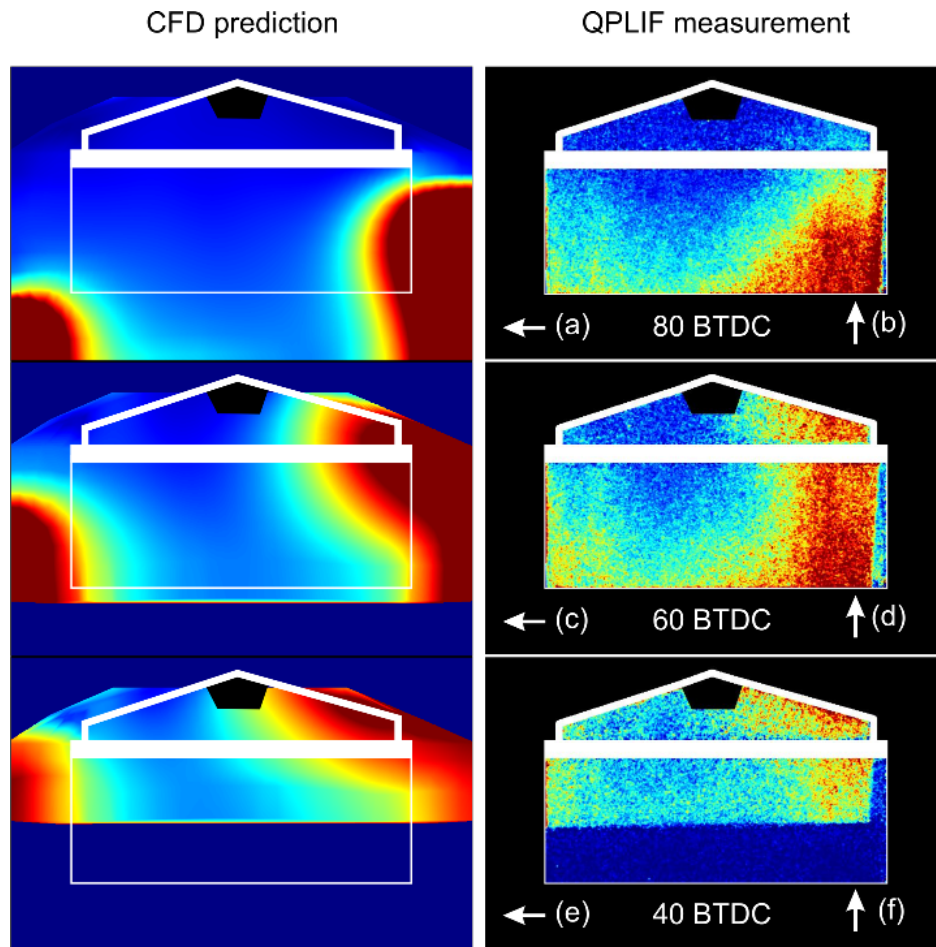


Figure 4.27.: Comparing QPLIF measurements with the predictions of Jaguar's CFD model.

Figure 4.27 shows CFD and QPLIF fuel distributions presented on the same false colour scale as has been used throughout this chapter - see, for example, figure 4.23. The distributions are also presented here at nominally the same scale. Whereas the CFD model is able to predict fuel concentrations across the entire width of the cylinder, the QPLIF technique is restricted by the available optical access. The limits of the QPLIF image are therefore reproduced as an overlay on the CFD distribution images.

It can clearly be seen that the two methods agree on the global trends, with the counter-clockwise fuel transport mechanism previously described also evident in the CFD model. The use of the same colour map allows more direct comparison of fuel

concentrations to be assessed, and acceptable agreement is observed. The computational challenges implicit in generating such fuel distribution predictions must not be understated, requiring accurate modelling of spray formation, evolution and evaporation, in-cylinder gas flow properties and so on. Therefore, this level of agreement is very encouraging in view of the fact that engine development is bound to become increasingly dependent on the availability of accurate computer models.

4.8. Summary

Two topics will be handled in this section. Firstly, the achievements detailed in this chapter will be summarised. Secondly, a description is given of the key requirements to achieve quantitative imaging of fuel distributions with PLIF.

4.8.1. Achievements

In this chapter, the development of a quantitative technique for fuel distribution measurements has been detailed. Particular attention has been paid to the properties of the base fuel, which has been demonstrated to have gasoline-like behaviour. The co-evaporation of tracer and the corresponding fuel fraction has also been investigated.

The *in-situ* calibration mechanism has been addressed in some detail, describing the evolution from a motored ‘calibration loop’ system to the fired ‘bookend’ calibration system which accounts for window fouling and the presence of internal EGR gases trapped between cycles by negative valve overlap (NVO).

The ‘bookend’ calibration system was then tested in the early-DI environment, where engine behaviour is most easily predicted. Following the success of the QPLIF technique in revealing data trends which conform to these expectations, the QPLIF technique was applied to the more challenging case of late-DI.

The late-DI test condition produces significant structure in the fuel distribution measured in the engine. Comparing the evolution of the measured distributions for

each tracer as a function of crank angle during an engine cycle has illustrated that there are no significant differences in behaviour between the three fuel fractions. This discovery allows for future experiments to be conducted with a single tracer - however, a multi-component base fuel must still be used to get meaningful results that can be compared to real-world gasolines.

Despite the low signal levels which are an inherent property of quantitative PLIF techniques, some information about cyclic variability has been derived. The data have also been compared to combustion burn rate, discovering a correlation between local FFC and early burn rate.

A comparison between QPLIF measurements and Jaguar CFD predictions has also been drawn, showing very encouraging agreement. The validation of computational models is an important application for QPLIF, since accurate computer models of engine behaviour are becoming increasingly important tools for engine development.

4.8.2. Prerequisites for Quantitative PLIF

The implementation of Quantitative-PLIF is subject to fulfilling a number of prerequisites. These will now be briefly reviewed.

Base fuel properties

The base fuel to which the fluorescent tracer is added must not itself fluoresce - otherwise the measurement will be contaminated. This can be achieved by using only saturated hydrocarbon components in the fuel mixture. Leading on from this point, the base fuel must also be non-absorbing at both the laser excitation wavelength and at the fluorescence wavelengths.

For the experimental results to have any validity, the base fuel must also have the same behaviour as standard pump gasoline. This can be achieved by including fuel components with a range of volatilities.

Tracer knowledge

Tracers must be selected (and/or the fuel designed) such that they co-evaporate with the base fuel, for maximum validity. Care should be taken that the tracer is not affected by, for example, thermal decomposition or chemical reactions.

The tracer doping level must be very carefully selected. The doping level must be kept to a minimum in order that the mixture does not become optically thick (i.e. that the laser is not attenuated as it propagates through the mixture). The doping level should also be kept low to minimise perturbation of the base fuel's properties. However, a low doping level will obviously cause a low signal level, which may have a significant negative impact on the quality of the images obtained.

The quantum yield (the combination of absorption cross-section and fluorescence efficiency) of the tracer must be known at all environmental conditions which will be present in the engine. This knowledge can either be determined by a large campaign of bench tests at varying temperature, pressure and composition, or by employing an *in-situ* method. The *in-situ* method has the advantage of calibrating quantum yield for a 'typical' engine cycle. The modelling approach requires that one independently determines temperature, composition (notably EGR) and other effects such as blow-by for the engine cycle under investigation.

Laser intensity

Care should be taken to avoid saturating the transition in the molecule. If the laser intensity is allowed to increase too far, then the assumption of a linear signal dependence on laser intensity (common in PLIF analysis) will be invalidated, and the analysis of PLIF images becomes challenging or even impossible. (That is, unless the entire sheet is at such high intensity that it is all saturated - but even so, the laser pulse is practically-speaking not a delta-function, so saturation cannot be totally assured.)

5. Fuel/Air Ratios by Laser Induced thermal Gratings - FARLIG

5.1. Chapter overview

In this chapter, the details of the FARLIG (Fuel/Air Ratio by Laser Induced thermal Gratings) technique will be presented. The opening sections give a brief overview of the technique, and an explanation of how FARLIG builds upon the standard thermal grating technique described in section 3.3.

The experimental details are then covered in section 5.4, including information about the apparatus and the procedures which are used in the acquisition of FARLIG signals.

The analysis system which has been developed for use in this project is then described, detailing the fitting system and the LITGS model in section 5.5.

This chapter then includes a section on results, where the application of FARLIG to fuel concentration (5.6.1), oxygen concentration (5.6.2, 5.6.3) and temperature measurements (5.6.4) is presented and discussed.

The chapter concludes with a summary section.

5.2. Technique overview

FARLIG is a non-linear optical diagnostic technique which in principle is capable of simultaneously determining fuel/air ratio and gas temperature in a probe volume.

Figure 5.1 shows the simple two-level model of a molecule again. With reference to the processes indicated in this system, it is appealing to note that the two techniques described in this thesis are complementary. LIF/QPLIF probes the spontaneous decay of the upper state via the Einstein A coefficient, and LITGS/FARLIG measure signals due to the quenching decay channel. As a result, it is conceivable that simultaneous LIF and LITGS measurements on a sample would provide a wealth of information about the decay of the excited species.




Figure 5.1.: A representation of a two level system. Transition rates include photon absorption $B_{12}\rho(\nu)$, spontaneous emission A_{21} producing a photon (wavy line), and a quenching process Q_{21} arising from molecular collisions.

The FARLIG technique extends standard LITGS in two ways - firstly, by deriving fuel concentration from integrated signal strength. Secondly, oxygen concentration can be derived by discovering how changes in the quenching function affect the shape of a LITGS signal. These two measurements, taken together, define a fuel/air ratio. FARLIG includes the measurement of temperature, as usual in a LITGS technique from the oscillation frequency of the grating visibility. (Pressure could also be measured from the decay time of the grating signal, but has not been performed in this work.)

In this thesis, experimental FARLIG signals recorded under a range of controlled conditions are presented. The test matrix is designed to validate the concepts that underlie the technique, and demonstrate its ability to measure fuel/air ratio and temperature simultaneously.

5.3. Measuring fuel/air ratios with LITGS

The formation and readout of a LITGS signal was described in detail in section 3.3. Figure 5.2 serves as a reminder of the important steps involved. First, a pair of pump laser beams are crossed at a small angle, to create a set of intensity fringes. The laser frequency is chosen to coincide with an electronic molecular resonance in a tracer species, and so population is driven to the upper state. Quenching processes convert this stored energy into heat over time. This localised heating creates a three-dimensional density modulation - a grating. The depth of this density modulation is then monitored as a function of time by Bragg-scattering a probe beam from the grating. The temporal shape of the signal recorded by detecting the scattered portion of the probe beam encodes information about the gas in the crossing region - e.g. T and p .

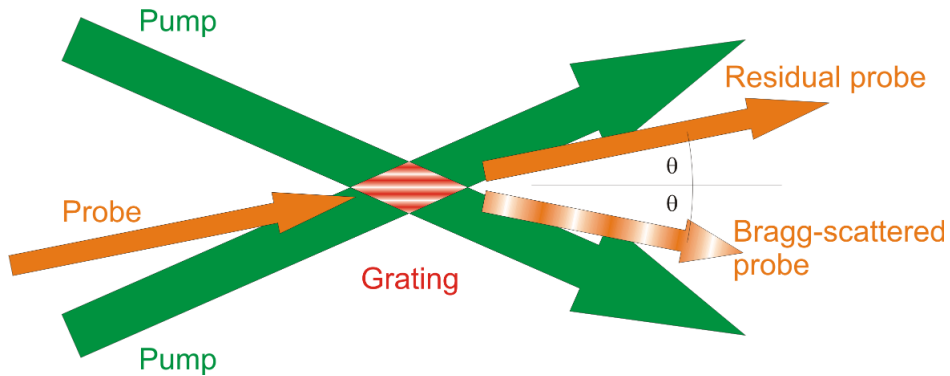


Figure 5.2.: Schematic of LITGS geometry, illustrating the key steps which contribute to the formation and readout of a thermal grating.

One of the key results of the theory described in section 3.3 is that the scattered signal is related to the convolution of three source terms:

$$S(t) \propto L(t) * Q(t) * G(t) \tag{5.1}$$

Where S is the signal, L is the laser pulse which excites the grating, Q is the quenching function which converts stored laser energy into heat, and G is the grating impulse

function, which describes the response of the medium to being heated.

In the FARLIG tests reported in this thesis, toluene is used as the absorbing species. Section 3.4.1 describes how such aromatic molecules are very strongly influenced by the presence of oxygen - almost every oxygen collision leads to quenching. Such a marked effect must be reflected in a change in the function $Q(t)$. In this thesis, two methods for determining oxygen concentration will be presented.

In the first method (section 5.6.2), the LITGS signal is fitted using a model quenching function. The model comprises two variable quenching decay timescales (fast and slow) and a variable branching ratio factor which alters the relative importance of the two channels. This model is not intended to represent the myriad internal processes of the excited molecule accurately, but should provide a quantitative measure of oxygen concentration following examination of how the model's three parameters vary across the matrix of validation tests.

In the second method (section 5.6.3), the heights of the peaks and troughs on the LITGS signal are measured directly, and the signal contrast calculated. This method reveals some sensitivity to oxygen concentration because a slow deposition of heat into the medium will result in a weak acoustic grating. Increasing oxygen concentration causes a faster quenching rate, which should be revealed as a signal with increased contrast.

The amount of heat deposited into the gas following the excitation event is dependent upon the density of the absorbing species. Greater heat deposition leads to greater density modulation, causing a stronger signal. The integrated strength of a LITGS signal is therefore directly related to fuel concentration (section 5.6.1).

It should be noted that the increased quenching rates in effect at increased oxygen concentrations are liable to cause a slight increase in LITGS signal, as energy which would otherwise be carried away by fluorescence photons is converted to heat. The magnitude of this signal increase is limited by the quantum efficiency of the tracer with no oxygen present, which at 266nm is approximately 30%¹⁰⁶.

By considering air to be 21% oxygen, and knowing both oxygen and fuel concentrations, a fuel/air ratio can be derived, which is the major new feature of the FARLIG technique. Further to this, the oscillation period of the grating signal can be readily used to derive temperature (section 5.6.4). It must not be overlooked that the accuracy of the temperature measurement is dependent on good knowledge of gas composition - which may be informed by the fuel/air ratio measurement.

By application of de-convolution methods, the quenching function itself may be recovered - providing the pump laser pulse shape and grating impulse response functions are well known.

Section 5.5 will describe how the analysis which generates the fuel/air ratio from experimental data is performed.

5.4. Experimental details

In this section, the details of the implementation of FARLIG are discussed. The subsections deal with the lasers and optical bench layout, the timing and data collection systems, the details of the gas cell, a consideration of detector characteristics, and finally, how the system is aligned.

5.4.1. Lasers and optics

A Continuum Surelite I-10 provides the pump pulses for FARLIG. This is a flashlamp-pumped multi-mode Nd:YAG laser specified to produce $\sim 200\text{mJ}$ at 532nm at a repetition rate of 10Hz , with a pulse length of about 5ns . The grating readout functionality is provided by use of a long-pulse flashlamp-pumped dye laser (FLPDL), which delivers up to 1J per pulse with a pulse length of approximately $1\mu\text{s}$ at around 600nm when a dye solution of Rhodamine-6G is used. Alignment duties are performed by a 10mW helium-neon (HeNe) laser.

The FLPDL is a modified Candela LFDL-8. The development of the current FLPDL

system is described in Dr. R. Stevens' DPhil thesis¹²⁸. For the experiments reported here, the laser was narrowed to $\sim 1\text{nm}$ FWHM with an intra-cavity interference filter, at 610nm. An intra-cavity adjustable aperture is used to confine the laser operation to the TEM₀₀ cavity mode. This limits laser pulse energy to approximately 50mJ, with a pulse length of approximately 600ns. The dye mixture used consists of 250mg of Rhodamine-6G in 10 litres of deionised water, for a molar concentration of approximately 5×10^{-5} . A splash[†] of Ammonyx[®] LO - an effective dye molecule triplet state quencher, which also aids in the mixing process - is also added to the dye solution.

The long-pulsed output of this laser is preferable to a continuous-wave source, such as a 1W argon-ion laser, simply because of the increase in intensity that is used to read the thermal grating's temporal character. Delivering 100mJ in 1 μs corresponds to an intensity increase of five orders of magnitude compared to the CW source. This would naturally be expected to improve the signal-to-noise ratio dramatically, and in some cases even allows for the signal spot to be seen with the naked eye (under normal ambient illumination levels, as in a typical office environment). The increased intensity also allows for the use of higher bandwidth detectors, which typically have lower sensitivity.

The optical bench layout in the FARLIG experiment is given in figure 5.3. Colour-coding is used to discriminate between the different laser wavelengths.

The output laser pulse from the Surelite is split at a dichroic mirror, reflecting only the second harmonic at 532nm, and transmitting the fundamental at 1064nm into a beam dump. The green beam is then passed into a Glan-Taylor (G.T.) prism to ensure its polarisation. The beam is then incident on a non-linear, frequency-doubling crystal made from KD*P (potassium dideuterium phosphate, highly birefringent), which converts a fraction of the 532nm radiation into laser light at 266nm (the fourth harmonic of the laser's fundamental emission). The fourth harmonic is produced in the orthogonal polarisation to the input, so a second G.T. prism is used to remove the majority

[†]About 20ml.




Figure 5.3.: Schematic of FARLIG optical bench setup, for details see text.

of the green from the output of the KD*P crystal.

The UV light at 266nm passes through a lens which focusses the beam at the back focal plane of the crossing lens. The beam is then split into two at a 50% beamsplitter, and two parallel beams of equal intensity are created, ready to be crossed to form fringes and excite toluene.

High quality dielectric mirrors are used to reflect the 532 and 266nm laser light. The radiation from the FLPDL and HeNe is manipulated across the table using broadband metallic mirrors. Note that the FLPDL beam also has a lens which produces a focus on the back focal plane of the crossing lens.

The crossing lens, placed just before the optical rail begins, is a 2" diameter fused silica lens, with a focal length at 266nm of approximately 235mm. This lens is responsible for crossing the laser beams involved in the FARLIG technique at a point, defining the measurement volume and producing a signal.

Note that the FARLIG experiment employs the defocussed crossing geometry, whereby a crossing lens is used to bring the beams together in the interaction region. Each beam has a lens introduced into its path which causes the beam to come to a focus in the back focal plane of the crossing lens. The combination of these two lenses can act as a telescope, to alter the diameter of the laser beams - but critically, the beams leaving the crossing lens are collimated. In focussed geometries, a 're-collimating' lens is usually required to collimate the output from the interaction region and create a signal beam. Experience suggests that avoiding the use of a 're-collimating' lens allows for the production of a cleaner signal, since less scatter is collected from cell windows.

As part of the alignment procedure, a pair of masks are placed into the mask holders indicated by the yellow circles at the ends of the optical rail. The technique used to align the beams to these masks is described in figure 5.9 and the accompanying text.

Figure 5.4 illustrates the shape of the masks used in the FARLIG experiment. The spacing of the holes on the mask is informed by considering firstly the desired grating fringe spacing, and secondly the angle of incidence for the Bragg scattering condition

to be satisfied at that fringe spacing.

The grating fringe spacing, Λ is given by equation 3.9. The beam crossing angle θ_x is straightforwardly determined from the combination of lens focal length and the spacing of the holes on the alignment mask. In the experiments reported here, the pump beam alignment hole spacing is 9.5mm. For a lens focal length of 235mm, this would produce a fringe spacing of approximately $6.5\mu\text{m}$.

Equation 3.10 describes the Bragg condition which must be satisfied for the probe beam to be scattered from the thermal grating. This equation is used to calculate the spacing of the probe and signal-tracing holes on the mask from the centre-line. For the masks in this experiment, these holes are found 10.90mm either side of the vertical centre line of the mask. Additionally, a vertical offset is introduced of 10mm above and below the horizontal centre line, in order to aid spatial discrimination between beams.




Figure 5.4.: Illustration of beam crossing geometry in FARLIG experiment. HeNe beam traces the path of the Bragg-scattered probe (FLPDL) beam.

Figure 5.4 shows how the masks are used to define the crossing geometry of the beams. Note that figure 5.4 could be a little misleading - obviously a mask with holes in it cannot by itself cause the depicted deviation of the beams! For the details of the alignment method, see section 5.4.5.

The HeNe beam (the red line in all figures) is used to trace the path that the Bragg-

scattered signal from the grating will follow, allowing the detector leg to be aligned. The LITGS signal which will follow this path is either collected by a photodiode (PD), a photo-multiplier tube (PMT), or a combination of a PD and a PMT. When both detectors are used concurrently, there is a beamsplitter in the detector box which delivers approximately 5% of the signal beam to the PMT. Note that the FLPDL and LITGS detectors are screened from detecting stray signals by use of Schott glass longpass filters, with strong blocking below 600nm. The LITGS detectors are also enclosed in a box which blocks stray light.

5.4.2. Timing and data acquisition

Once alignment is performed, and the experiment is otherwise prepared, all timing is controlled from a Berkeley Nucleonics Corporation 555 (BNC-555) timing and pulse generator box. A summary of the settings can be found in table 5.1.

| Channel # | 1 | 2 | 3 | 4 |
|-----------|------------|-----------------|-----------------------------|----------------------|
| Delay | 0 | 180-270 μ s | Q-SW + \sim 2.760 μ s | Q-SW + 10.35 μ s |
| Mode | Continuous | 1-in-10 | 1-in-10 | 1-in-10 |
| Purpose | YAG Lamps | Q-Switch | FLPDL Trigger | Oscilloscope |

Table 5.1.: Summary of settings on BNC-555 pulse generator.

The BNC-555 is run with an internally-generated 10Hz clock, to simulate the Surelite's internal 10Hz rate, and maintain any equilibrium (e.g. thermal lens) achieved during the alignment procedure. This internal clock directly triggers the firing of the flashlamps inside the Surelite. The experiment is constrained to run at 1Hz by the repeat rate of the flashlamp pumped dye laser (FLPDL). The Q-switch input on the Surelite is therefore also used, to produce a laser pulse from 1 in 10 flashlamp pulses. The output from the two lasers is synchronised by use of an offset of approximately 2.760 μ s. The oscilloscope is then triggered some 10 μ s later.

The oscilloscope used in these experiments is a LeCroy WaveRunner WR6050A. It records data on all four channels at 5×10^9 samples per second, with a bandwidth of

500MHz. All channels are terminated with 50 ohm impedance, to minimise the presence of cable reflections. The datafiles are saved to a hard disk inside the oscilloscope, and later copied to a PC for analysis. The oscilloscope is always connected to two photodiodes - one monitoring the pulse shape of the 266nm light, and one monitoring the un-scattered probe beam. As previously mentioned, the LITGS signal is collected by use of either a PD, a PMT, or a combination of both on the remaining channels.

5.4.3. Gas cell considerations

The high pressure gas cell used in the FARLIG experiments is a modified version of the high pressure cell used for SRS in QPLIF (figure 4.5).




Figure 5.5.: Gas handling considerations. Upper - part **a**: gas cell setup. Lower - part **b**: gas manifold.

Figure 5.5 shows the modifications to the standard SRS high pressure gas cell. The diagram illustrates the connection to the gas manifold, allowing the introduction of nitrogen and oxygen. A K-type thermocouple is fed into the cell via a plug, and is situated as depicted. On the opposite side of the cell, and offset slightly is a second welded-on fitting. A ball valve is connected to this fitting, and capped by a standard

Swagelok[®] nut. A PTFE disc with a 1mm diameter hole drilled through the centre is captured between the nut and the ball valve. The hole is just large enough to accept the outer diameter of the needle, which can be used to introduce a measured amount of toluene into the cell in conjunction with a syringe of 10 microlitre graduations. Note, the filling procedure is detailed in section C.1. The temperature of the cell may be adjusted by use of heater tape which is wrapped around the body of the cell and connected to an autotransformer, allowing variable voltage to be used to warm the gas.

A pressure transducer is connected to the gas manifold, as can be seen towards the lower right hand side of figure 5.5, which allows the overall pressure of the mixture in the cell to be monitored (using the associated digital display). A set of needle valves are present, one of which allows the gas cell's contents to be isolated. The cell contents may be purged by use of a vacuum pump. A three-way valve is also included which may either be closed, or expose the manifold to regulated gas bottles containing oxygen free nitrogen and oxygen.

The thermal grating is formed by absorption in toluene molecules. In the initial FARLIG experiments, a saturated vapour of toluene was employed in a lower-volume gas cell. This arrangement proved to be less than ideal, owing to the tendency of the toluene to condense onto the windows - causing window damage as a light 'etching' when the laser was incident. Indeed, this damage, though minor to the eye, was sufficient to eliminate the LITGS signal.

In the current evolution of the apparatus, the gas cell has a much larger volume, which aids in rejecting scattered light from the windows. The introduction of the ball valve and PTFE disc allows for a more carefully metered volume of toluene (or toluene-doped middle fraction component fuel) to be admitted to the cell.

The location of the crossing point inside the cell determines the optimum toluene concentration. A relatively high concentration is preferred as it means more laser light is absorbed in the crossing region, increasing the visibility of the density modulations forming the grating. However, this must be balanced against the absorption which

occurs in the gas prior to reaching the crossing point.

The energy absorbed in the crossing region can be described by the following analysis:

$$\Delta E(x) = \frac{dI}{dx} \Delta x A \tau \quad (5.2)$$

where ΔE is the energy absorbed, $I(x)$ is the intensity of the laser beam as a function of propagation distance, x through the vapour, and Δx is the length of the crossing region - for a suitably short crossing region. A is the cross-sectional area of the crossing region, and τ the laser pulse length. The Beer-Lambert law, equation 2.7, describes how I varies, leading to:

$$\Delta E(x) \propto -I_0 n \sigma e^{(-n\sigma x)} \quad (5.3)$$

where all symbols are as defined before. Differentiating this relationship, the following is found -

$$\frac{d\Delta E(x)}{dn} \propto -\sigma e^{(-n\sigma x)} - n\sigma e^{(-n\sigma x)} \times (-\sigma x) \quad (5.4)$$

which may be set to zero, to look for a maximum -

$$\sigma e^{(-n\sigma x)} = \sigma x \times n\sigma e^{(-n\sigma x)} \quad (5.5)$$

$$n = 1/(\sigma x) \quad (5.6)$$

Equation 5.6 predicts the optimal concentration as that when the beam decays to $1/e$ (i.e. 37%) in the distance before interaction. For an absorption cross-section of $2 \times 10^{-19} \text{cm}^{-2}$, and 15cm propagation to reach the crossing point, the optimal n is $3.3 \times 10^{17} \text{cm}^{-3}$, or approximately 1.2% concentration, by number.

The test cell is 34mm in internal diameter, and 30cm long, for a volume of 1090cm^3 . To create the optimal concentration, approximately $65 \mu\text{l}$ of toluene is required. A saturated vapour of toluene requires $208 \mu\text{l}$ - a concentration of 4% - which is known to lead to toluene condensing on the windows. It is experimentally determined that a

single droplet from the syringe and needle, of approximately 20 μ l strikes an adequate balance between limiting window condensation and maintaining grating strength. (The analysis above suggests that changing from 65 μ l to 20 μ l decreases the energy deposited in the interaction region to about 60% of the maximum value - an acceptable loss, given that the cell windows should remain clear for longer.)

In the case of quantitative fuel concentration measurements, the optimisation above is not applicable. This is because any change in fuel concentration is liable to have a marked effect on the optical thickness of the mixture. In order for the laser intensity at the crossing point to be consistent - whatever the toluene concentration in the cell - the mixture *must* be optically thin. Indeed, note that the doping level in the QPLIF experiments - where an optically thin toluene mixture was employed in the engine - corresponds to roughly 1.7 μ l of toluene in this test cell. Alternative liquid mixtures for use in the cell, comprising toluene doped into the medium fraction of the Shell component fuel (this fuel was described in section 4.3.1), will therefore be used in section 5.6.1.

Following extended running, a 'mist' is observed in the test cell, which appears to scatter the laser beams quite strongly. Figure 5.6 shows a visible HeNe beam as it passes through this highly scattering atmosphere following \sim 200 laser shots. The onset of this problem can be observed as the LITGS signal gains a significant contribution from the incoherent scatter of the FLPDL as it passes through the cell. This unwanted scatter into the signal path can be later corrected in the analysis stage, but is better avoided from the outset and eventually becomes very strong.

This phenomenon is a further reason to reduce the toluene concentration where possible, since a lower toluene vapour concentration does appear to slightly increase running time. A gas flush following the procedure described in section C.1 is required to reverse this behaviour - the particulate matter is not observed to separate out, even after hours of inactivity.

It is hypothesised that this phenomenon may be some form of photochemical 'smog'.




Figure 5.6.: Strong scattering visible following 200 FARLIG shots.

Note that this effect would not be problematic in an IC engine application, where a fresh charge is admitted to the engine on every cycle.

5.4.4. Detector characterisation

The FARLIG experiment employs a combination of photodiodes (PDs) and a photomultiplier tube (PMT). The PDs are ThorLabs DET210 reverse-biased PIN silicon detectors, with a specified rise and fall time of 1ns, and a spectral response covering 200-1100nm. The PMT in use in these experiments is a Thorn EMI 9783B, driven by a high-voltage power supply, producing up to 900V. This tube is specified to have a typical rise time of 2ns, a transit time of 20ns, and a photocathode which is sensitive from 160-680nm.

LITGS experiments reported previously in the literature are usually conducted with PMTs, owing to the low intensity of the Bragg-scattered beam. The FLPDL allows

these experiments to employ a less sensitive PD. The detectors have been characterised by exposing them to 50fs laser pulses, which are short enough to appear as delta functions to the detectors. Figure 5.7 shows how the PD and PMT compare in this test. 50-Ohm cabling and termination is used in order to minimise cable reflections which may otherwise contaminate the results.

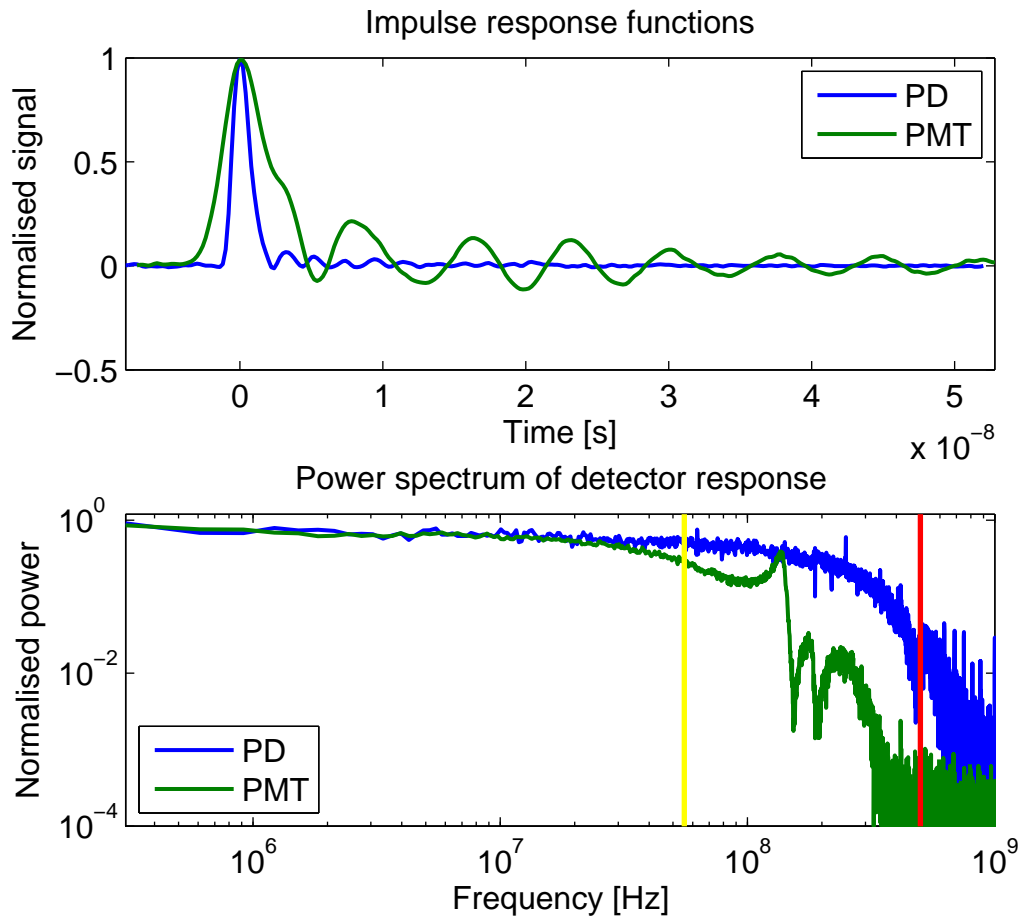


Figure 5.7.: Illustration of the impulse response functions of the detectors used in FARLIG (averaged over ~ 200 single shot measurements). The three PDs were independently measured, but all produced similar results.

Figure 5.7 demonstrates quite a marked difference in behaviour between the photodiode and photomultiplier used to collect the LITGS signal. The upper part of the figure shows the detectors' temporal response to the very short incident laser pulse. The PD conforms well to its specification in this test, with a FWHM of 1.1ns. The

PMT fares less well, with a FWHM of 3.6ns. More worrying is the oscillatory behaviour that is present on the PMT trace. This was confirmed not to be symptomatic of cable reflections, and may instead be an artifact of the detector and/or power supply. The lower half of figure 5.7 shows the power spectrum (i.e. frequency response) of the two detectors. The vertical yellow line denotes the typical oscillation frequency of a LITGS signal, and the vertical red line is the bandwidth of the oscilloscope. Both detectors do a good job up to the LITGS oscillation frequency, but at around 150MHz the response of the PMT becomes erratic, comparing poorly to the smooth, monotonic roll-off of the PD. Indeed, the PD appears to be well matched to the bandwidth of the oscilloscope. (Note that despite a delta function containing equal contributions from all frequencies, ideal behaviour in this test would be little different to the trace registered by the PD - some frequency roll-off is expected due to the bandwidth limit on the oscilloscope.)

This test suggests that the photodiode is more able to capture the detail in a FARLIG signal accurately - a valuable improvement.

Another important aspect of dealing with the detection of a FARLIG signal relates to the signal level generated by the detector. For both the photodiodes and photomultiplier, the detector response becomes non-linear above a certain threshold, as illustrated in figure 5.8.

Figure 5.8 shows LITGS signals recorded at four different intensities on the photodiode used for capturing the FARLIG signal. The signals are scaled by eye such that a ‘best match’ is achieved between all four in the lower voltage tail, from around 100ns onwards on the x axis. The signals which were recorded with higher peak voltage output - as given in the key in the upper right of the figure - demonstrate far lower modulation than those recorded with lower intensity beams. As the signals become lower in magnitude, they tend asymptotically to that collected with the minimum intensity. A practical compromise between ideal signal reproduction and high signal to noise ratio is selected as the 150mV signal level. Neutral density filters can be employed to drop the signal strength if necessary to achieve this target. A similar test




Figure 5.8.: Illustration of the impact of detector saturation on the recorded shape of a LITGS signal, comparing the average of 4 single shot LITGS signals at each signal strength.

was performed with the photomultiplier (and a 700V setting on the associated power supply), yielding a saturation output voltage of $\sim 700\text{mV}$.

5.4.5. Alignment procedure

A painstaking alignment procedure is required for the generation of a LITGS signal. The experiment begins by ensuring that the lasers are all operating optimally - i.e. the FLPDL is running on the lowest-order cavity mode (tested by inspection, and pulse shape), and the Surelite is warmed up, with good beam quality and maximised second harmonic conversion. The KD*P crystal is then confirmed to be in the correct orientation, before the major alignment procedure, which is illustrated in figure 5.9.

To summarise the alignment method (a full description is given in section C.2), the beams are aligned using the masks and a pinhole centered at the crossing point, in

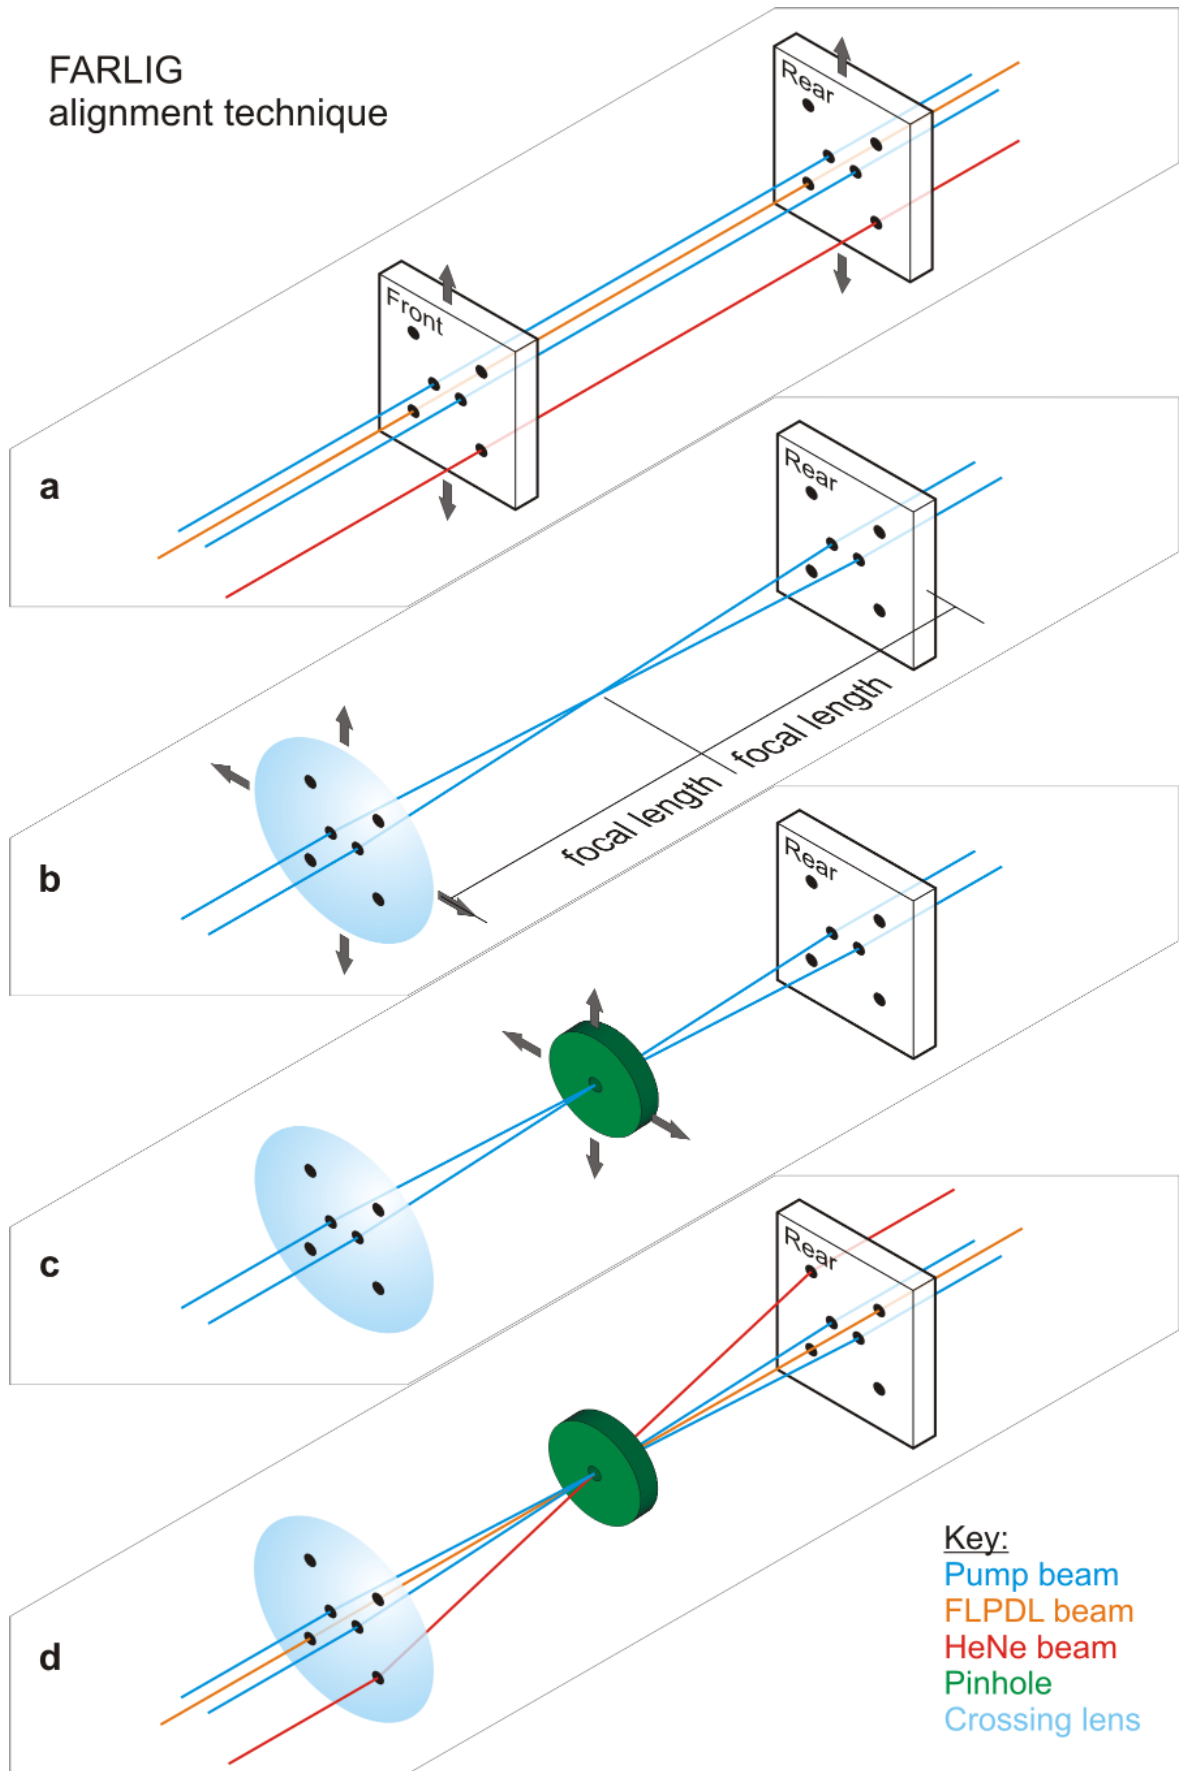


Figure 5.9.: FARLIG alignment technique. See section C.2 for detailed method.

conjunction with the adjustments on the crossing lens.

Note that the masks are coated with a fluorescing film (highlighter pen), in order that the fourth harmonic light is converted into visible fluorescence, aiding alignment. In the alignment method of figure 5.9, the fourth harmonic beams define, in some sense, the ‘optical axis’ about which all other beams are aligned. The use of a pinhole which is spaced one focal length from the rear mask allows the chromatic aberration of the crossing lens to be corrected. The cell windows will act to delay the focus slightly, but the effect is small and the alignment procedure described in figure 5.9 has been successfully applied repeatedly.

In order to achieve the optimal alignment of the detection leg, a short length of 1/2” diameter tube is placed around the crossing point of the beams - i.e. a little larger than the 266nm pump beams. This tube is slightly bowed, and a few drops of toluene are dripped into the dip. As the toluene evaporates, it creates a strong vapour concentration, which generates a significantly stronger LITGS grating than normal. Providing that the alignment steps above have resulted in the alignment being close to ideal, all that will be required is some tweaking to maximise the detected LITGS signal. (The focussing lens in the FLPDL beam is often a good candidate for adjustment.) Note that this operating condition is able to create a LITGS signal which can be seen by eye when it hits a target[†].

One final experimental detail relates to the synchronisation of the detectors with respect to one another. It is crucial to be able to time-align the various signals during the later processing of the data. The various time offsets are determined by removing any coloured glass filters from the detectors, and allowing the 266nm pulse to propagate to each detector from the crossing point, by introducing a diffuser at this location. Naturally, this test accounts for not only optical delay in propagating the laser pulses across the table, but also cable delay, as the signals pass from detector to oscilloscope.

[†]With the lab lighting at normal levels, as in a standard office environment.

5.5. Data analysis

5.5.1. Modelling theory

Equation 5.1 indicates that the LITGS signal is composed of the convolution of three source functions. The three functions are: $L(t)$ which describes the laser pump pulse, $Q(t)$ which describes the quenching of stored laser energy into heat, and $G(t)$ which describes the response of the medium to impulse heating - i.e. the grating itself. The convolution theorem states that

$$F.T.(f * g) \propto F.T.(f) \times F.T.(g) \quad (5.7)$$

where $*$ denotes the operator for a convolution, and $F.T.(f)$ is the Fourier transform of the function f . Since the Fourier transform of a function of time is its equivalent in frequency space, equation 5.1 may be re-written in the following form:

$$z(\omega) \propto l(\omega) \times q(\omega) \times g(\omega) \quad (5.8)$$

where ω is angular frequency, $l(\omega)$ is the laser pump pulse in the frequency domain, $q(\omega)$ is the quenching function and $g(\omega)$ is the grating impulse response function - which follows from the theory of Paul et al.⁷⁹, described in section 3.3. The function $q(\omega)$ is represented in this thesis by a model which incorporates two lifetimes and a branching ratio - and will now be discussed:

Modelling quenching

The Fourier transform of a decaying exponential multiplied by the Heaviside step function (i.e. $f(t) = 0$ $t < 0$ and $f(t) = e^{-\alpha t}$ $t \geq 0$) is given by the following:

$$F.T.(e^{-\alpha t}u(t)) = \frac{1}{\sqrt{2\pi}} \frac{1}{\alpha + i\omega} \quad (5.9)$$

$$F.T.(e^{-t/Q_i}u(t)) = \frac{1}{\sqrt{2\pi}} \frac{Q_i}{1 + i\omega Q_i} \quad (5.10)$$

where α is a proportionality constant, Q_i is the decay lifetime of a particular pathway, and $u(t)$ is the unit step function.

Calculations are actually performed in the developed fitting routines in dimensionless time: $B_i = Q_i/\tau$, where τ is the inter-fringe transit time. In this representation, the quenching is modelled with two timescales: B_1 is the faster decay, and the B_2 the slower. The two exponentials are mixed with a split ratio, so that:

$$q(\omega) = \text{split} \times \frac{1}{\sqrt{2\pi}} \frac{B_2}{1 + i\omega B_2} + (1 - \text{split}) \times \frac{1}{\sqrt{2\pi}} \frac{B_1}{1 + i\omega B_1} \quad (5.11)$$

where i is $\sqrt{-1}$, and ‘split’ is the parameter controlling the splitting ratio between the two pathways ($0 \leq \text{split} \leq 1$).

Physically, it is difficult to ascribe these two timescales to any particular processes in the toluene molecule. The discussion in section 3.4.1 indicates that it is highly unlikely that internal conversion is active at this excitation wavelength (266nm). Heating is expected to occur by very rapid vibrational and rotational relaxation of the molecule (sub-nanosecond) whenever a non-equilibrium population distribution is formed. In the absence of oxygen, Burton and Noyes¹⁰⁶ measures a fluorescence quantum yield of 30% at 25°C in a low pressure toluene vapour. The fluorescence lifetime is calculated to be 142ns, and the inter-system crossing ‘lifetime’ to be 51ns. In the presence of oxygen, the rate of intersystem crossing is known to be dramatically enhanced, both from the excited singlet state (S_1) to the triplet state (T_1), and from the triplet state (T_1) to the ground electronic state (S_0). Kawaoka et al.¹⁰³ explains that oxygen-enhanced triplet quenching results in the larger aromatic molecule carrying away the excess energy above the triplet-singlet energy gap in oxygen. This is because the Franck-Condon factor is very small for transitions to higher vibrational states in singlet oxygen[†] - particularly

[†]Perhaps this production of reactive singlet oxygen is related to the formation of the ‘smog’ in the test cell over extended periods - since the singlet state of oxygen is metastable and reactive...

relative to the same energy jump to vibrationally-excited levels in the aromatic.

Grossmann et al.²¹ indicates that the effective fluorescence lifetime of toluene excited at 248nm reduces from ~ 50 ns to ~ 2 ns as oxygen concentration is increased to just 6%. This behaviour is consistent with possessing an oxygen-enhanced rate of inter-system crossing, as described above.

5.5.2. Modelling a LITGS signal with MATLAB

The modelling and fitting code employed in this thesis was previously described by Stevens¹²⁸. The fitting system has been developed using the functions and built-in development environment provided by MATLAB[®], a product of The Mathworks, Inc. The core of the fitting system is a module which conducts a non-linear least squares fit to a measured LITGS signal by adjusting the parameters of a model. In order to improve the speed of fitting, the modelling is conducted in frequency space - where the primary advantage is to swap the time-consuming convolution operations of equation 5.1 for products, as in equation 5.8. The implementation and features of the MATLAB model of a LITGS signal will now be discussed.

The modelling of a LITGS signal is conducted by the MATLAB function `LMin_Func3.m`, which calls the function `TEMPORAL4.m`. The model formed by these two functions has five parameters, and can vary the Reynolds number, Re , the inter-fringe transit time, τ , the fast quenching timescale, Q_1 , the splitting ratio between quenching ratios, 'split', and the slow quenching timescale, Q_2 . The following discussion describes how the fitting system generates a model LITGS signal from these five parameters and the measured laser pulses.

`LMin_Func3.m` almost immediately calls `TEMPORAL4.m`. The model of `TEMPORAL4.m` assumes that the grating is spatially sinusoidal (i.e. not suffering from saturation effects), such that the only source of variable grating reflectivity is the temporal part of the density perturbation, i.e. that given in equation 5.1. As described previ-

ously (section 5.5.1), this equation can be written in frequency space as:

$$z(\omega) \propto l(\omega) \times q(\omega) \times g(\omega) \quad (5.12)$$

The experiment measures $L(t)$, so a fast Fourier transform is used to generate $l(\omega)$. The Fourier-domain representation of the quenching function, $q(\omega)$, is given by equation 5.11, and was incorporated into the modelling code by the author. The frequency-domain representation of the grating impulse heating response function, $g(\omega)$, is given by the following equation⁷⁹:

$$g(\omega) = ((i\omega)^3 + c_1 \times (i\omega)^2 + c_2 \times (i\omega) + c_3)^{-1} \quad (5.13)$$

where

$$c_1 = \left[\frac{4(2\pi)^2}{3\text{Re}} + \frac{\gamma(2\pi)^2}{\text{RePr}} \right] \quad (5.14)$$

$$c_2 = \left[(2\pi)^2 + \frac{4(2\pi)^4\gamma}{3\text{Pr}(\text{Re})^2} \right] \quad (5.15)$$

$$c_3 = \frac{(2\pi)^4}{\text{RePr}} \quad (5.16)$$

and γ is the ratio of specific heats for the gas composition at the crossing point. Note that all equations are given in normalised time - i.e. the frequency ω has been normalised by multiplication with the inter-fringe transit time.

Once $z(\omega)$ is known, the inverse-fast-Fourier transform may be taken to recover $Z(t)$. In order to obtain the grating's intensity reflection coefficient, the modulus-square is calculated:

$$n(t) = Z(t) \times Z(t)^* \quad (5.17)$$

where the superscript asterisk denotes the complex conjugation operation.

This result is returned to `LMin_Func3.m`, where the modelled grating reflectivity $n(t)$ is multiplied by the measurement of the probe pulse (i.e. `FLPDL(t)`), to produce a

modelled LITGS signal.

The author has also included at this point a step which convolves this modelled LITGS signal with the measured instrument response function of the FARLIG photodiode. This step results in a slight improvement in the agreement between the model and the measurement.

5.5.3. Fitting to a LITGS signal with MATLAB

In this sub-section, the adoption of the model described above into a MATLAB fitting code is described.

Fitting system

The operation of the fitting system will now be summarised by describing the function `LIGSfitAuto4.m`. Significant modifications the author has introduced will be highlighted as the routine is discussed, though it should be noted firstly that the majority of the code was developed prior to the author's involvement in this project by Dr. R Stevens, and secondly that small adjustments which have been made throughout the code will remain unaddressed.

The function takes 4 input arguments and produces 3 outputs. The input arguments are used to inform the routine which shot from which dataset is to be analysed. The output arguments contain the best-fit parameter set, the comparison of the model and the experiment, and a fit error.

`LIGSfitAuto4.m` begins by locating the data file specified by the input arguments, and extracting the relevant traces that were recorded on the oscilloscope (modified for this thesis to accept the new file format). The fitting parameters and a variety of constants are then configured by the `SET_LIMITS.m` function. Importantly, the detector synchronisation time offsets are specified at this point (note that these offsets will change, should any detectors be moved or coaxial cables replaced).

The code then subtracts the zero baseline level from each channel, and makes an estimate of the noise on the baseline. The time offsets are then used to synchronise the pump, probe and signal detectors with a spline-based interpolation, and truncated.

A module developed by the author is then incorporated (`pumpdeconv.m`), which allows the measured impulse response function of the signal photodiode to be de-convolved from the measurement of the pump laser pulse. This de-convolution process (using fast Fourier transforms) reveals a little extra longitudinal mode beating in the pulse, and results in a reduction of the measured FWHM by about 0.5ns.

The `LIGSfitAuto4.m` routine continues by rewriting the time vector to align $t=0$ to the peak of the laser pump pulse, and fits a Gaussian to the laser pump pulse to estimate its FWHM.

The code then determines which samples contain data worth fitting to, by locating the rising edges of the probe and signal traces, and where both fall to a level where they are masked by the noise. Assuming sufficient samples exist between the rise of the probe and the inception of the signal, the signal trace can be examined for probe leakage that is the result of incoherent scattering into the detector leg. If the signal trace is not zero in this region, a scaled version of the probe pulse is subtracted from the signal such that it is. The data is once again truncated, to include only the samples worth fitting to.

By fast-Fourier transforming the signal trace and searching for the strongest peak in frequency space, the oscillation frequency of the grating signal may be determined accurately. This prediction of inter-fringe transit time is used to update the value configured by `SET_LIMITS.m`.

The final step before minimisation is to plot the initial guess against the source data. The agreement is usually relatively poor at this stage, since the predicted parameter values from `SET_LIMITS.m` are quite general.

It is at this stage that the MATLAB built-in function `lsqcurvefit` - the non-linear least squares curve fitting function - is called upon to minimise the fit error by adjusting the

fit parameters. Experimentation suggests that the best fit is achieved most quickly by first fitting Re , Q_1 and split, then allowing just split and Q_2 to vary, before fitting all parameters simultaneously.

This fitting mechanism calls upon the `LMin_Func3.m` and `TEMPORAL4.m` functions (that were described in the previous sub-section) to model a LITGS signal at many different combinations of parameters in order to determine the best fit. Naturally, `lsqcurvefit` does not perform an exhaustive and time-consuming search of every possible combination of parameters - rather, it approaches the minimum using the 'Trust-region interior-reflective Newton optimisation' algorithm.

Once the best fit has been achieved, the model and experiment are compared visually for the user, and a fit error residual is calculated. These, and similar, useful, pieces of information are returned by `LIGSfitAuto4.m` in its output arguments.

Other fitting considerations

The viability and repeatability of the FARLIG technique may only be assessed following the application of the fitting system to as large a group of test signals as possible. A wrapper routine (`LIGSBatch.m`) was therefore also developed by the author, to track the combination of fitting parameters that were chosen for each LITGS signal against the test conditions (e.g. cell gas composition and pressure). With the saturation characteristics of the detectors (described previously in section 5.4.4) in mind, the dataset is pre-conditioned by selecting only those signals with a peak signal level below the saturation voltage limit.

The dataset of signals recorded with the photodiode and within acceptable voltage bounds (50-150mV peak value) is composed of approximately 1500 individual shots. The entire dataset may be processed by the automated fitting routines overnight - in approximately 8 hours. (Each shot takes approximately 20 seconds to find an optimal solution, which is perfectly acceptable when developing improvements and so on.)

As an aside, it should be noted that the least-squares fitting routines applied in this

analysis (and, indeed, generally) do not guarantee to determine the global minimum in the fitting parameter space. Experience suggests that fitting routines work best when presented with a series of ‘centred and scaled’ parameters - i.e. that the upper bounds, lower bounds and estimates for each of the fitting parameters are all scaled by choice of appropriate factors to values on the order of unity. It is also important that the chosen fitting parameters are orthogonal. The orthogonality of fitting parameters may be verified by conducting a ‘brute force’ trawl across parameter space in proximity to the fit-determined minimum, and determining the error on the fit at each combination of these nearby fitting parameters. For each pair of parameters chosen from the set a 2D plot may be created, which represents the ‘height’ of the ‘error surface’ (i.e. the fit error at any particular combination of parameters) using a false-colour mapping. In general, the shape of the fit minimum is expected to be an ellipse - with increasing error as the parameters depart from the minimum. If the axes of the ellipse do not lie parallel to the axes of the plot, then the parameters are not orthogonal, and varying one parameter has an impact on the best choice for the other. For example, two parameters are quite strongly related if the ellipse’s major axis is at 45° to the horizontal axis, and the non-linear least squares fitting system will not behave optimally.

The non-linear fitting routines incorporated into MATLAB usually allow the Jacobian at the calculated minimum to be returned as an output argument. The Jacobian, J , is a matrix populated by elements which are calculated by taking the partial derivative of the function to be fitted with respect to the fitting parameters. The Jacobian takes this form:

$$J = \begin{pmatrix} \frac{\partial F_1}{\partial x_1} & \frac{\partial F_1}{\partial x_2} & \frac{\partial F_1}{\partial x_3} & \dots & \frac{\partial F_1}{\partial x_j} \\ \frac{\partial F_2}{\partial x_1} & \frac{\partial F_2}{\partial x_2} & \frac{\partial F_2}{\partial x_3} & \dots & \frac{\partial F_2}{\partial x_j} \\ \frac{\partial F_3}{\partial x_1} & \frac{\partial F_3}{\partial x_2} & \frac{\partial F_3}{\partial x_3} & \dots & \frac{\partial F_3}{\partial x_j} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial F_i}{\partial x_1} & \frac{\partial F_i}{\partial x_2} & \frac{\partial F_i}{\partial x_3} & \dots & \frac{\partial F_i}{\partial x_j} \end{pmatrix} \quad (5.18)$$

Where F_i represents the i^{th} element in time of the function to be minimised, and x_j represents the j^{th} minimisation parameter. This matrix, the values of the fitting parameters at the minimum, and the residual (i.e. difference between the model and experiment) may then be supplied together to the function `nlparci.m` to determine 95% confidence intervals on each parameter. Alternatively, a rough measure of relative sensitivity to each fitting parameter may be obtained by investigating the values of the partial derivatives in the matrix. Note that this Jacobian matrix is numerically calculated by Matlab during the course of the fitting operation by determining the change in the fitting function with small changes in the fit parameters. A further, user-led approach to investigating the sensitivity of the fit is discussed in the text accompanying figure 5.20, in the following section.

5.6. Results

This section will include a presentation of the results and discussions regarding the FARLIG technique. The components of a fuel/air ratio measurement are addressed in turn - considering first the measurement of fuel concentration with LITGS, and then the derivation of oxygen concentration. This section will close by presenting the measurement of temperature, which is the third parameter covered by the FARLIG technique.

5.6.1. Fuel concentration

The fuel concentration measurement component of FARLIG is achieved by integrating the recorded LITGS signal. As discussed in section 5.4.3, the ‘standard’ toluene concentration adopted for use in these experiments is unsuitable for quantitative fuel concentration measurements, since it is optically thick.

Therefore, these tests continue to employ single droplets of fuel from the syringe and needle to create the fuel vapour in the gas cell, but instead of neat toluene, mixtures of toluene and the middle fraction of the Shell component fuel are used. Two blends are created, chosen to reduce the toluene density by a factor of 20 - 40 versus the standard case. The two dilutions are created by adding 0.5ml of toluene to a small, empty measuring cylinder, and topping it up to 10ml with component fuel middle fraction. Half of this mixture is then removed and stored for use as the ‘1 in 20’ mix. The remaining half is diluted by doubling the quantity of fluid remaining in the measuring cylinder by adding more middle fraction. This more dilute mixture is then at a ‘1 in 40’ concentration level. At low doping levels such as these, the thermal grating that is written into the medium becomes too weak to be reliably detected with the LITGS photodiode, and so the LITGS photomultiplier is used instead. The gas cell was pressurised to 4 bar absolute with dry compressed air for both trials.

Figure 5.10 shows a comparison of normalised LITGS signal strengths for the two

concentrations. Approximately 200 shots were recorded for each concentration, and the PMT PSU was set to 700V throughout. The integrated signal on the LITGS PMT is divided by the product of the integrated signals on the fourth harmonic and probe beam detectors, in order to perform the normalisation. (Further normalisation is applied in order to achieve a mean of unity for the higher concentration, purely for the purposes of presentation here.) Detector saturation is not a concern; the signal levels were sufficiently low.

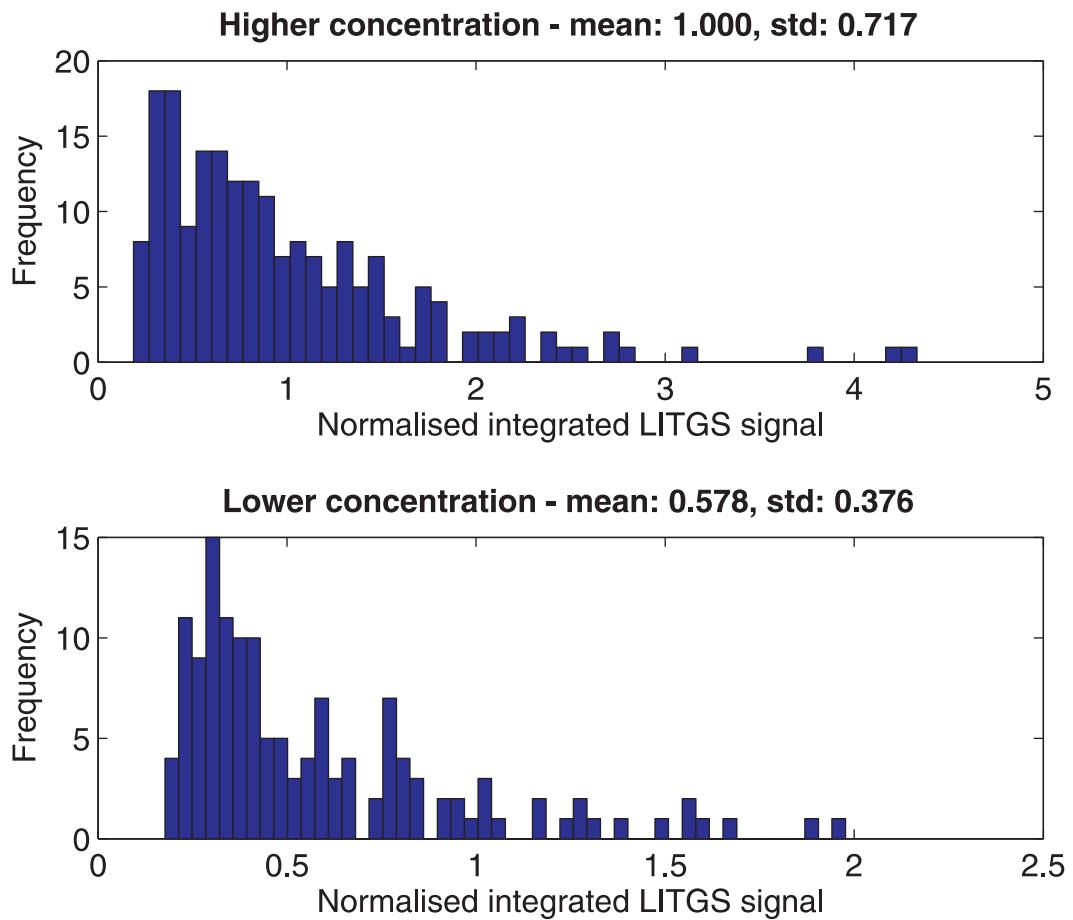


Figure 5.10.: Comparison of normalised FARLIG signal strengths for two fuel concentrations.

The spread in probe pulse energy (i.e. the standard deviation of the distribution) during these tests was approximately 15%, and the spread in the fourth harmonic energy was less than 2.5%. As can be seen in figure 5.10, the spread in the normalised

LITGS signal strength is significantly larger - in the region of 70%.

However, the two distributions in figure 5.10 demonstrate very similar shapes (the lower edge of the low concentration distribution may be artificially cut, due to the elimination of signals below a threshold voltage value). Indeed, the mean of the low concentration distribution is measured to be 0.58, which matches well with the expected value of 0.5, particularly in view of the fact that the toluene number density in the lower concentration case will be subject to some variation about the 50% ideal. Compositional uncertainty is introduced in the mixing stage and when injecting a single droplet of fuel. The repeatability of the volume of the fuel droplet is not known, but any variation is taken to be small.

Clearly, the limiting factor in applying the FARLIG technique to the measurement of fuel concentration using this method is the shot-to-shot fluctuation in the signal strength. Unfortunately, this shot-wise instability renders the assessment of a small increase in signal strength due to the addition of oxygen (as discussed in section 5.3) highly challenging. Nevertheless, this demonstration indicates that the technique holds good promise for making single-shot fuel concentration measurements, provided a reproducible signal strength can be attained.

The source of the shot-wise instability is currently unknown, but future experiments will attempt to resolve this issue.

5.6.2. Oxygen concentration by quenching model

The determination of oxygen concentration in the FARLIG technique depends upon knowledge of how the quenching model's parameters vary in the presence of oxygen. Assuming that this behaviour is found to be reproducible, and that the gradient of some fit parameter (or combination of fit parameters) with the addition of oxygen is sufficiently high, a sensitive measurement could be performed.

Test conditions

In order to validate the FARLIG technique for the measurement of oxygen concentration, a matrix of test conditions at which signals would be measured was drawn up, as shown in table 5.2.

| O ₂ pressure | 1 bar total | 2 bar total | 4 bar total | 8 bar total |
|-------------------------|-------------|-------------|-------------|-------------|
| 0 mb | | | | |
| 100mb | | | | |
| 200mb | | | | |
| 400mb | | | | |
| 600mb | | | | |
| 800mb | | | | |
| 1000mb | | | | |
| 2000mb | | | | |
| 4000mb | | | | |
| Air (21%) | | | | |

Table 5.2.: Matrix of gas cell pressures and compositions tested, in order to validate the FARLIG technique. Blacked-out cells denote compositions that were not tested.

Table 5.2 shows the testing matrix that the FARLIG experiments covers. The test conditions encompass both changing the composition of the gas by varying the fraction of oxygen present, and varying the total pressure in the gas cell. Note that the test matrix is built to cover - as a minimum - the range of environments found in an IC engine during a typical cycle - with pressure from 1 to 8 bar, and oxygen content varying between 20% (i.e. air) and 0% (i.e. burned gas).

Each gas composition is generated by mixing the gas from a cylinder of oxygen and a cylinder of nitrogen, with the correct ratio obtained by using partial pressures. A further test using bottled air is performed, to verify that the composition created from the partial pressure mixing procedure is accurately generated.

Example fits

As discussed in section 5.5, the many single-shot LITGS measurements taken at each of the conditions indicated in table 5.2 are processed by the fitting system as a large batch, enabling the variation of the 5 fitting parameters (Reynolds number, inter-fringe transit time, fast quenching time, splitting and slow quenching time) to be monitored as composition and pressure are varied.

Figures 5.11 and 5.12 show single shot signals recorded in a nitrogen atmosphere at 1, 2, 4 and 8 bar, with their associated fits. The following two examples describe how to decode the plot titles: The plot title ‘090225 0-1 #7, FitErr=1.0877%’ means a signal recorded on the 25th February 2009 with 0mb oxygen in 1 bar total pressure. It was shot number 7 in the sequence recorded, and the fitting system achieved a 1.0877% residual fit error. The plot title ‘090225 400-4 #21, FitErr=0.90414%’ means a signal recorded on the 25th February 2009 with 400mb oxygen in 4 bar total pressure (balance nitrogen). It was shot number 21 in the sequence recorded, and the fitting system achieved a 0.90414% residual fit error. The ‘residual fit error’ is produced by calculating the standard deviation of the residual vector and dividing by the magnitude of the recorded LITGS signal. This value thus measures ‘goodness of fit’.

Each plot consists of two panels: the lower panel shows the experimental data as blue dots, overlaid by the model as a red line. The FLPDL readout pulse is shown with green dots, and the magenta trace illustrates the pump pulse. The upper panel is the difference between the model and the experiment.

The values of the fitting parameters themselves are also incorporated into each plot in a text box. Note that these values are given to an arbitrary number of significant figures, and the actual error/variability in each parameter should be assessed by considering the size of the error bars in figures 5.14-5.16, and/or looking at the relevant data in appendix D[†]

[†]A numerical method for determining the confidence interval limits on the fitted parameter estimates, taking advantage of the Jacobian matrix was also discussed in section 5.5.3.

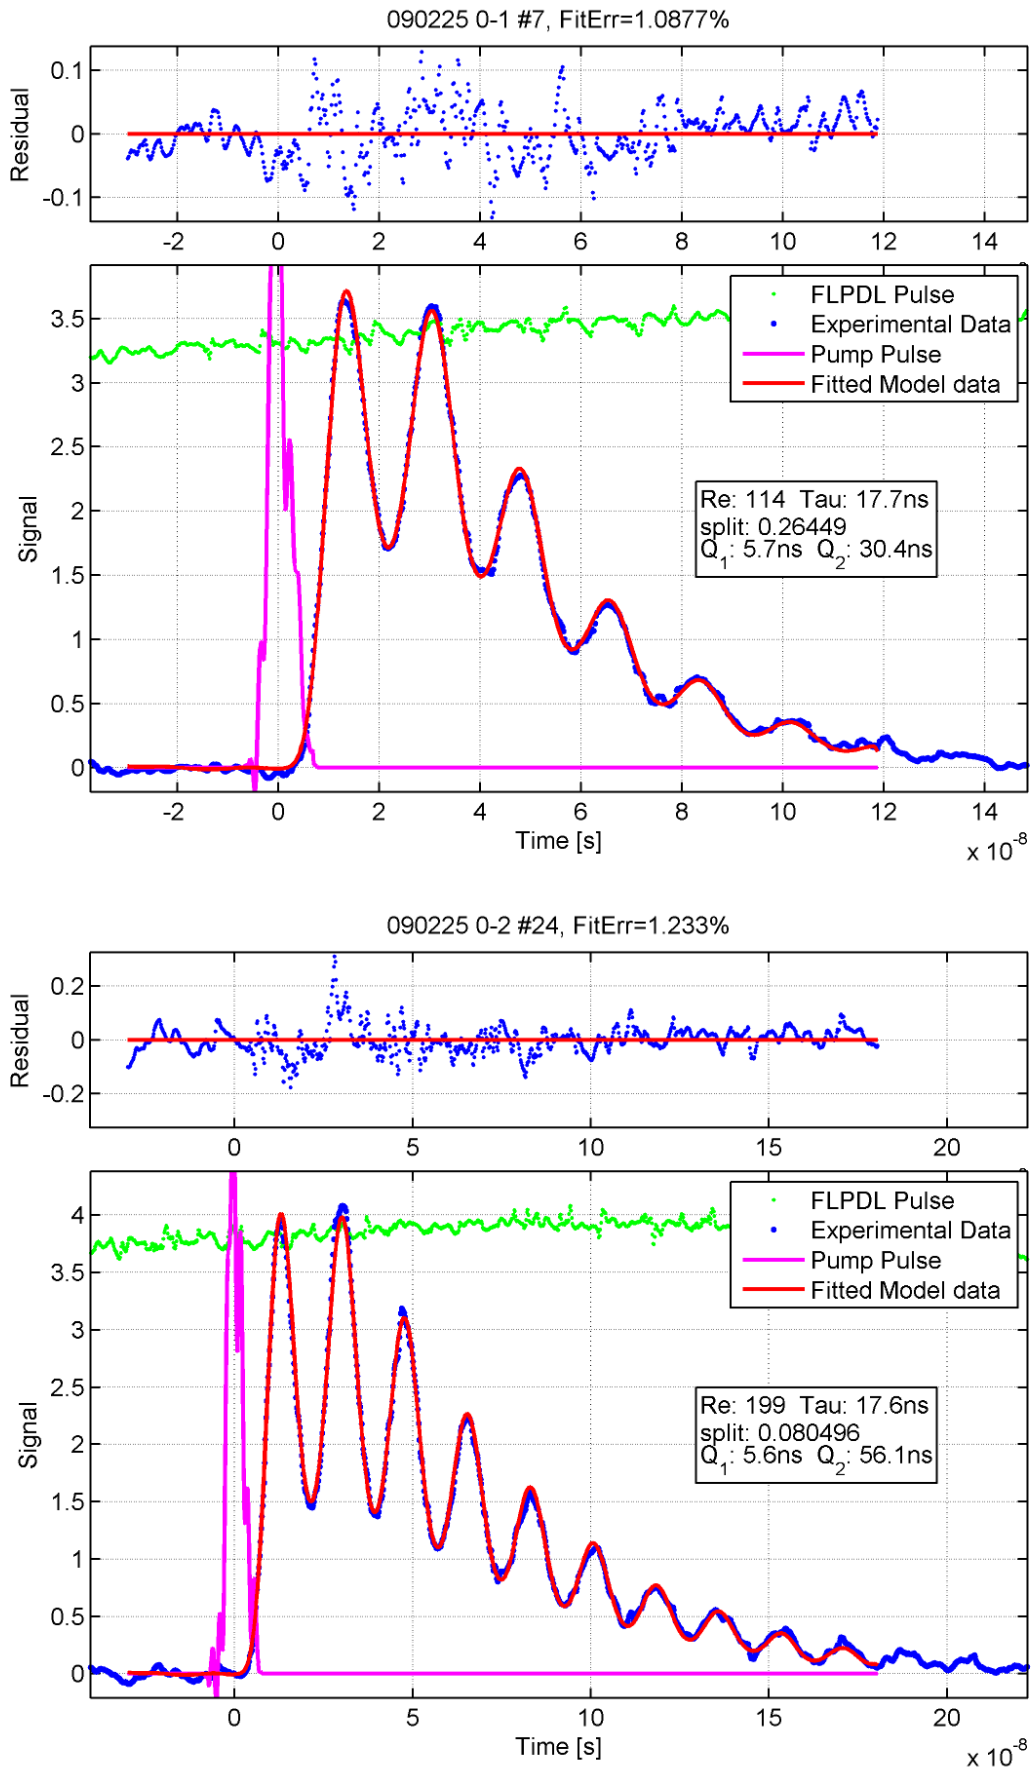


Figure 5.11.: FARLIG shots recorded in nitrogen at 1 bar and 2 bar total pressure.

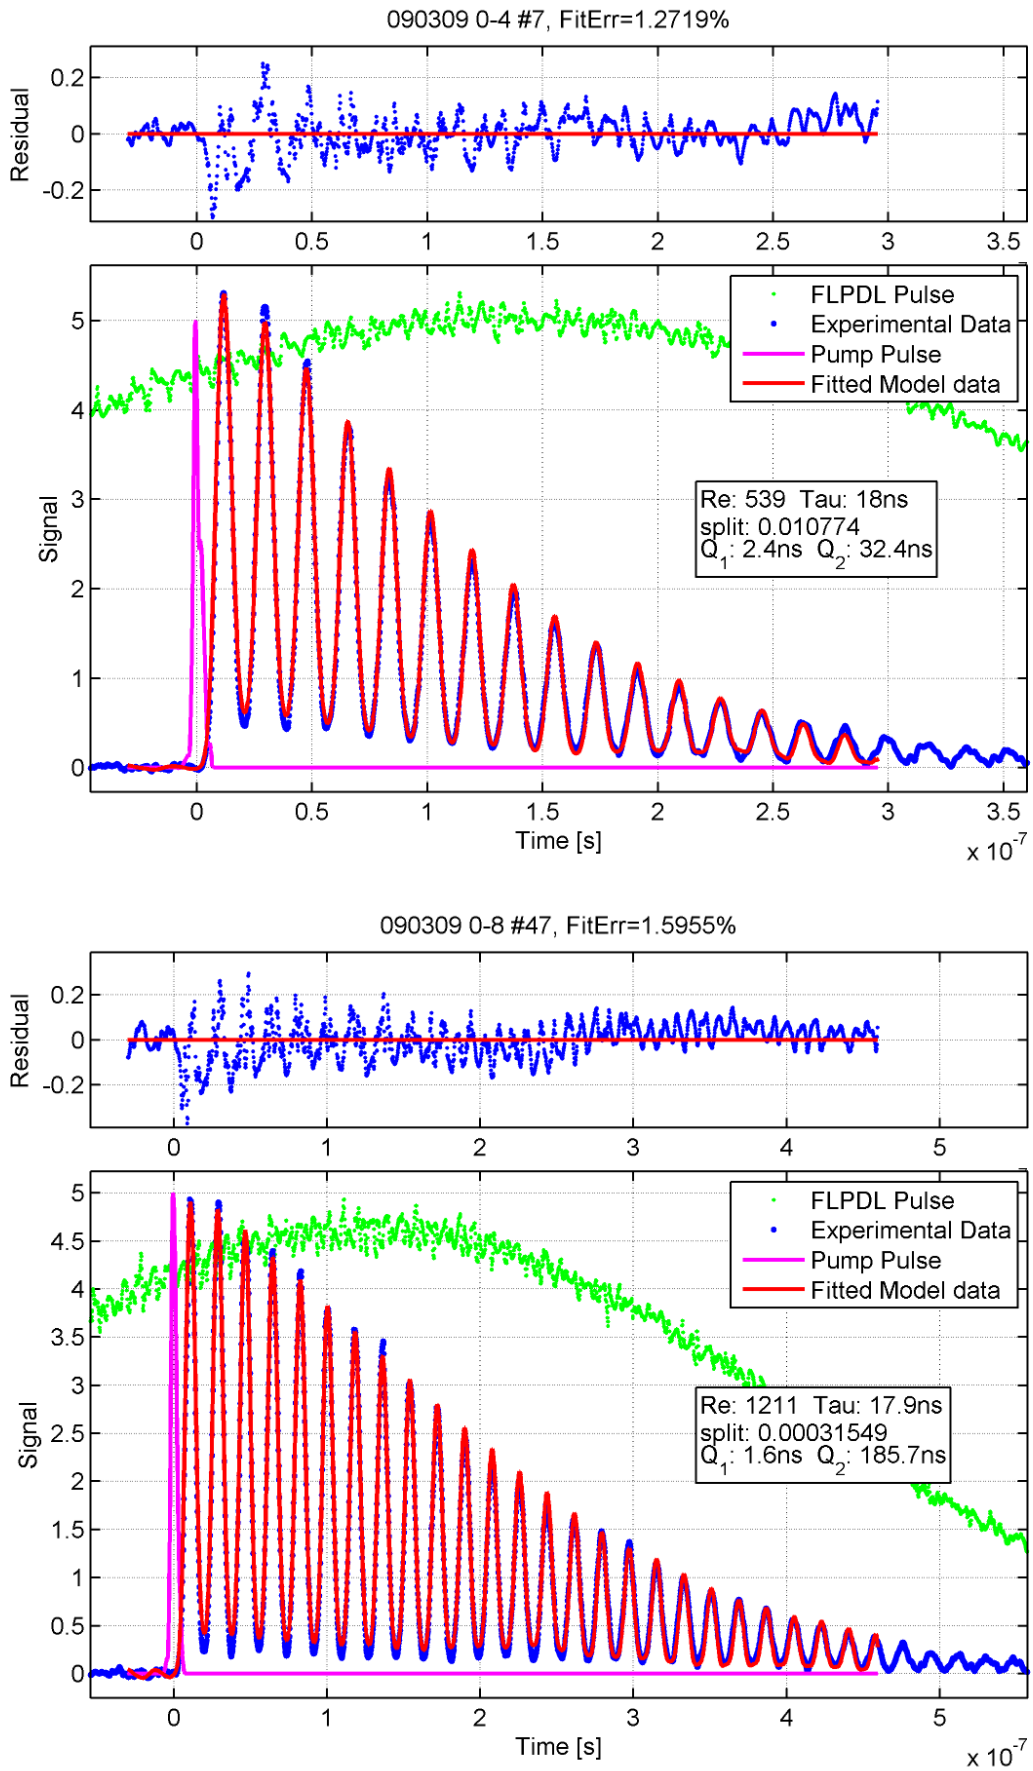


Figure 5.12.: FARLIG shots recorded in nitrogen at 4 bar and 8 bar total pressure.

Figures 5.11 and 5.12 demonstrate that the fitting system operating with the five-parameter model is capable of producing very good likenesses of the experimental data. This behaviour lends confidence to the ability of the fitting system and model to resolve any differences which may be present between the various test conditions of table 5.2.

Presenting trends in fitting parameters

The output of the batched fitting system will now be examined, in order to attempt to discover trends which may be used to allow FARLIG to measure oxygen concentration. A full table of the outputs of the batched fitting process is given in appendix D, and is the resource from which the figures in this sub-section are produced. Note that the datasets are post-processed by preferentially selecting all shots which have a fit error of less than 2%. If this would result in removing more than half of the fitted shots, then the top 50% best-fitting shots are kept, and the rest discarded.

The analysis of trends begins with figure 5.13, which shows how the inter-fringe transit time, τ varies with the changing composition of the gas in the test cell. As a reminder, τ is given by Λ/c , and c is dependent on the mean molecular mass m and ratio of specific heats γ , as will be shown later in equation 5.20. The amount of toluene present in each test is thought to be very similar, so changes in m and γ are largely due to the varying ratios of nitrogen and oxygen.

Figure 5.13 shows a series of data points, where each data point is calculated by finding the mean of the shot-wise best-fit values. The determination of τ is quite stable, with error bars which would be hard to discern. Two distinct groupings are seen in the plot - one for each of the two days on which the majority of the FARLIG data was recorded. Two best-fit lines are included, one for each day. It is likely that the difference between the best fit lines is caused by a slightly different experimental alignment on the two days, causing a minor change in grating plane spacing (Λ). The scatter of the data points about the best fit lines indicates that the gas compositions (from table 5.2) had been created as intended.




Figure 5.13.: Inter-fringe transit time dependence on composition

Note that a simple model is used to derive the best fit line - where the mean molecular mass m is calculated as a function of composition as $m = \frac{f \times 32 + (1-f) \times 28}{100}$. In this equation, f is the percentage of the mixture which is oxygen. The optimum proportionality constant (α) is then fitted in order to relate τ and m via $\tau = \alpha \sqrt{m}$.

With figure 5.13 showing good agreement with expectations of the gas composition, the variation of the quenching parameters as oxygen partial pressure changes can be presented. Figures 5.14, 5.15 and 5.16 each show the variation of one of the quenching model fit parameters against oxygen partial pressure.

Figure 5.14 shows the relationship between the fast quenching timescale and the oxygen partial pressure for all the test conditions visited during the data collection exercise. Note that the colouring used in this figure is consistent with figure 5.13, and figures 5.15 and 5.16. Each data point in the figure represents the mean of the best-fit




Figure 5.14.: Fast quenching time as a function of oxygen partial pressure

values generated by the fitting engine. In addition, the data from the 4 bar test on the 25th February 2009 (i.e. ‘090225 4 bar’) is plotted with error bars, which are the same length above and below the data point as the standard deviation of the shot-wise best-fit values at each condition.

This figure highlights some inconsistency between the repeat datasets at 2 bar and 4 bar which were recorded on both the 3rd March and the 25th February. It is not clear why this discrepancy has arisen.

The trend in the data is rather noisy, but does tend to indicate that the fast quenching time probably becomes shorter as both oxygen is added, and as pressure is increased. However, with reference to the error bars on the ‘090225 4 bar’ dataset, the suitability of this parameter for determining oxygen concentration appears limited.

Note that the ‘090225 4 bar’ conditions each consist of approximately 30 shots, from which the means and standard deviations are derived.

Figure 5.15 shows how the split parameter varies with oxygen partial pressure.




Figure 5.15.: Split parameter as a function of oxygen partial pressure

Figure 5.15 consists again of mean data points, with error bars plotted on the ‘090225 4 bar’ dataset, just like figure 5.14. In this figure, a marked decrease in the split parameter with pressure can be seen. The agreement between the data captured on the two different days is acceptable, and a gradual reduction of the split parameter with increasing oxygen concentration appears to be present. Owing to the size of the error bars on the ‘090225 4 bar’ dataset, this parameter may be more suitable for use in oxygen concentration measurements.

Figure 5.16 illustrates the evolution of the slow quenching timescale parameter as a function of oxygen partial pressure, and is plotted following the conventions of the previous two figures. Somewhat erratic behaviour is illustrated in this figure. The agreement between the data collected on the two days is passable only, and the ‘090225 4 bar’ dataset and associated error bars indicates that this parameter is probably not




Figure 5.16.: Slow quenching time as a function of oxygen partial pressure

well-suited to determining oxygen concentration. With reference to figure 5.15, it may be seen that the fraction of the quenching function which is composed of the long timescale becomes small once pressure rises. This probably makes the fitting of the exact length of the long decay timescale, Q_2 , somewhat unstable, leading to the relatively tall error bars. That is to say, once only 2% of the decay is in the slow channel, the fit is probably quite insensitive to the difference between a long quenching timescale of 100ns as compared to 150ns.

Figure 5.17 is a log-lin plot, which shows the split parameter on a logarithmic scale, as plotted against total gas pressure. Consistent with the previously-presented data, figure 5.17 consists of data points which are mean values. The almost linear trend through the clusters of data points appears to indicate that the split parameter varies roughly as the inverse square of the total pressure. The groupings of data points also




Figure 5.17.: Split parameter as a function of total pressure

generally have the data recorded at higher oxygen partial pressures towards the bottom of the cluster, as may be verified from figure 5.15. Interestingly, this indicates that the split parameter may be a convenient indicator of total pressure.

The trends in the quenching parameters that have been presented here are, sadly, not sufficiently compelling to pronounce FARLIG a success. Some discussion follows.

Oxygen addition

Figures 5.18 and 5.19 show the effect of adding 20% oxygen at 1 bar total pressure and 10% oxygen at 4 bar total pressure, respectively. The difference is subtle in both cases, but is generally related to the relative positioning of the peaks and troughs of the first two or three oscillations. The two figures also serve to demonstrate the accuracy of the fitting procedure, which generally achieves $< 2\%$ fit error ('FitErr'). It is somewhat surprising that despite the quality of these fits, the overall trends are poorly resolved.




Figure 5.18.: FARLIG shots at 1 bar total pressure, in pure nitrogen and 20% oxygen.

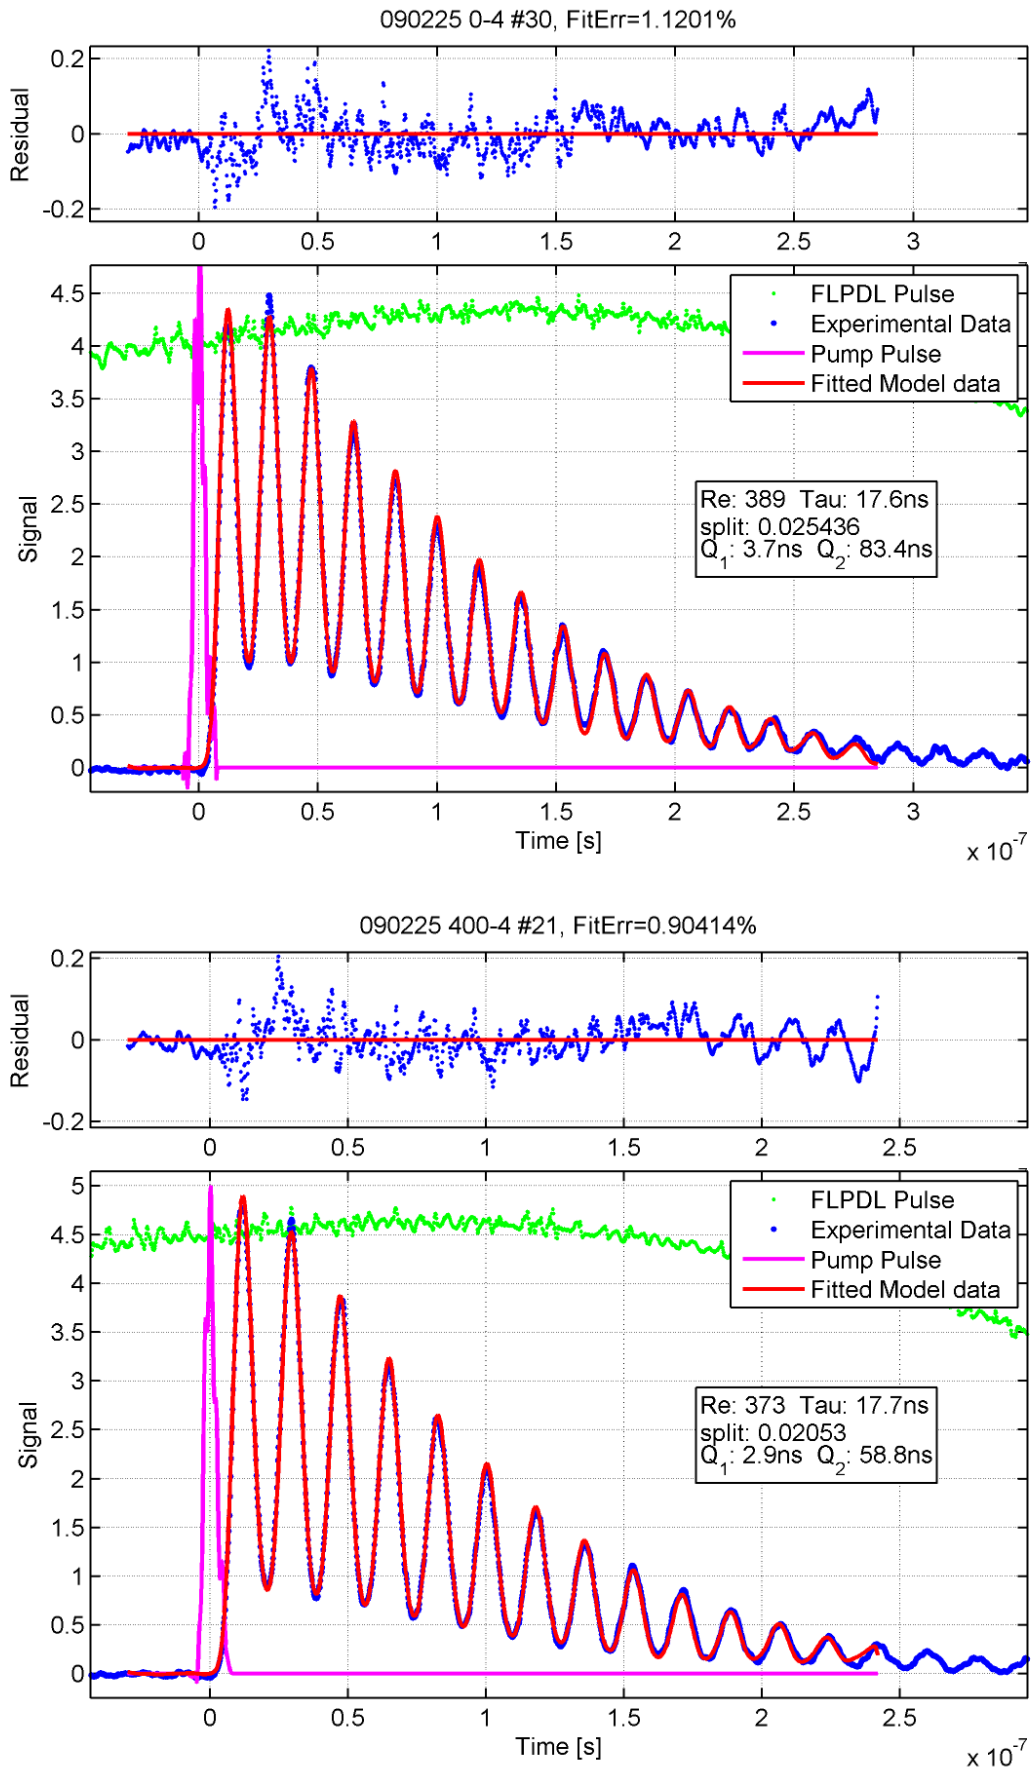


Figure 5.19.: FARLIG shots at 4 bar total pressure, in pure nitrogen and 10% oxygen.

Sensitivity of the fitting and model

The sensitivity of the fitting system may be further demonstrated by manually affecting the fit outcome. Figure 5.20 demonstrates what happens when the split factor is forced to zero, with all other parameters left at their original values.

The upper part of figure 5.20 shows the fit as produced by the automated process. The lower plot indicates the manually-altered split value, assigned the smallest value which MATLAB accepts (i.e. effectively zero)[†]. The effect of removing this 2% long decay pathway from the quenching function has quite a noticeable impact, illustrating that the fit and model are highly sensitive to changing the fit parameters.

Discussion of fitdata trends

Figures 5.14, 5.15, 5.16 and 5.17 illustrated the following: Firstly, that the magnitude of the fast quenching time exhibits a weak dependence on oxygen concentration with moderate scatter, and secondly that the long quenching timescale exhibits no clear behaviour as a function of oxygen concentration due to high variability. Thirdly, the split parameter offers stronger oxygen concentration dependence than Q_1 , the fast quenching timescale, together with reduced scatter. It is also noted that the split parameter is a strong function of pressure.

Despite being the most promising candidate, the data presented here regarding the split parameter is probably insufficient to build an oxygen concentration mapping for FARLIG.

However, if one of the other quenching parameters could also be used, then the ambiguity may be slightly relieved. Of the two, the shorter timescale appears to show the most promise. However, bearing in mind that the pulse width of the pump laser is approximately 4.5ns FWHM, almost all the data points presented in figure 5.14 are actually poorly resolved by the experimental system - since the width of the pump

[†]Note that the fit error in the lower plot was manually set to infinity, and was not calculated by the system to have this value!

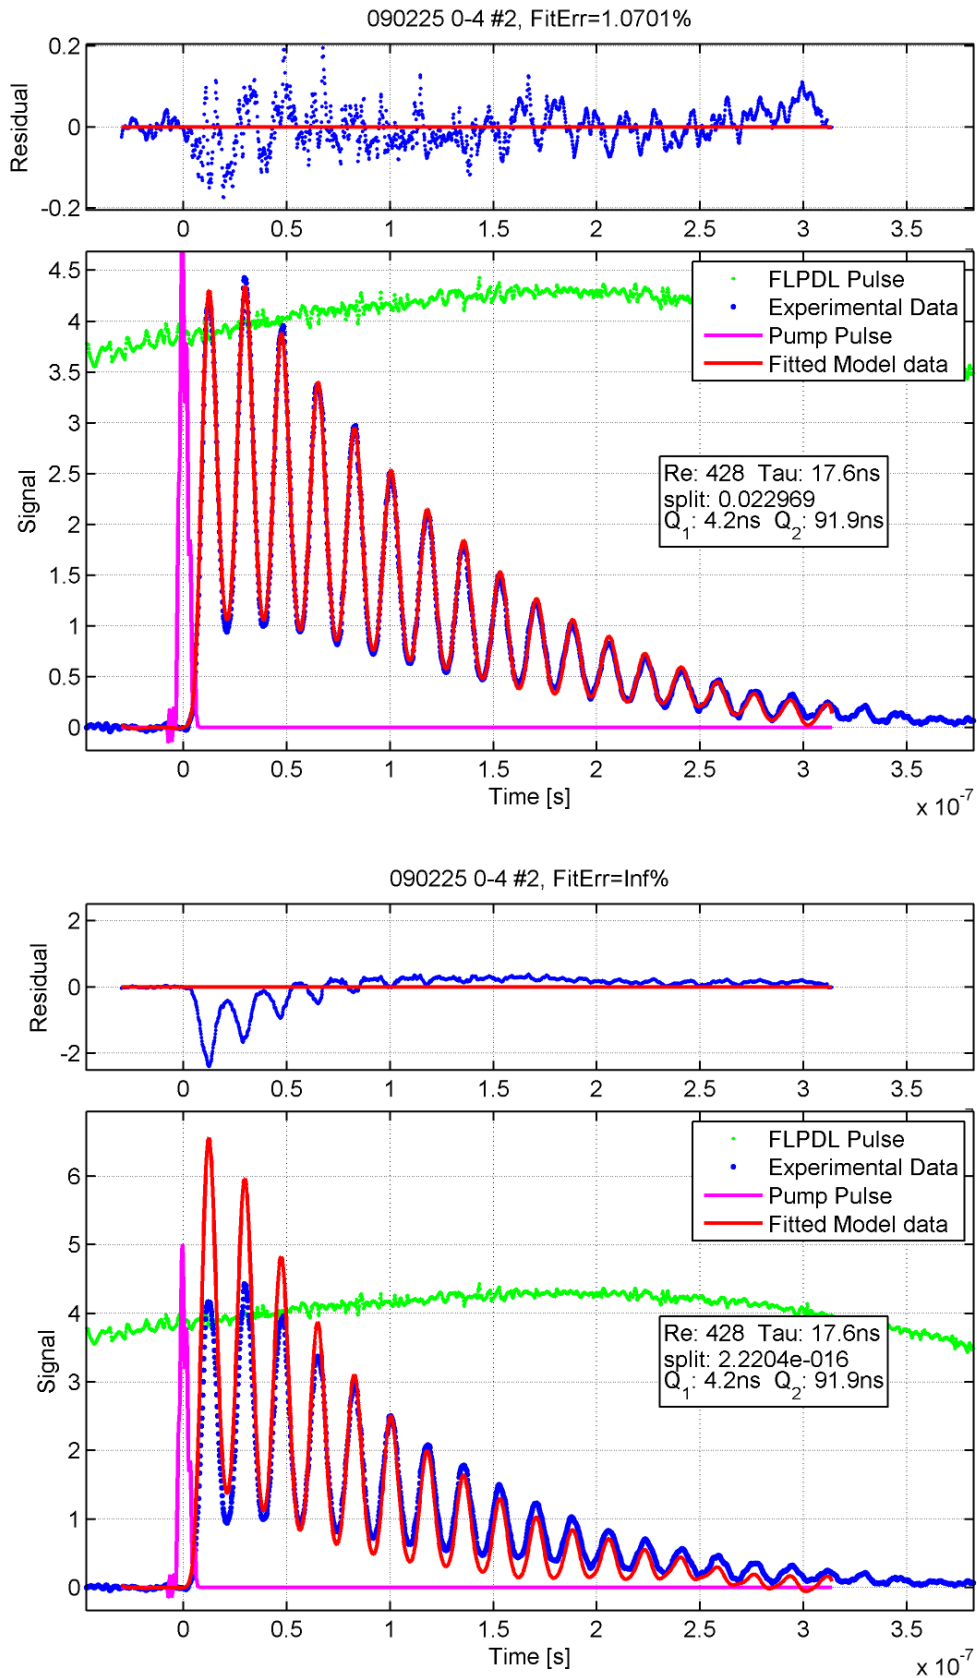


Figure 5.20.: FARLIG shots at 4 bar total pressure, in pure nitrogen. Upper plot is the best fit, lower plot has the 2.3% slow quenching time component removed, producing a noticeably poorer fit.

pulse will act to ‘smear’ these shorter timescales out, and affect the accuracy with which they may be measured.

Indeed, recalling earlier discussions in this thesis regarding tracer photophysics (e.g. section 3.4), it is expected that some processes will occur on even faster timescales - surpassing the limits of both the photodiode (IRF FWHM $\sim 1\text{ns}$) and the oscilloscope (500MHz). A simple estimate of the time between collisions of gas molecules at STP was given in section 3.2.2, giving a value of 300ps. Working from this baseline value at 1 bar absolute, the time between collisions with oxygen molecules may be estimated - and such values are given in table 5.3.

| Pressure | 20% O ₂ | 10% O ₂ | 5% O ₂ | Nitrogen |
|----------|--------------------|--------------------|-------------------|----------|
| 1 bar | 1.5ns | 3ns | 6ns | 300ps |
| 2 bar | 0.8ns | 1.5ns | 3ns | 150ps |
| 4 bar | 0.4ns | 0.8ns | 1.5ns | 75ps |
| 8 bar | 0.2ns | 0.4ns | 0.8ns | 38ps |

Table 5.3.: Estimation of time between collisions with oxygen molecules for a series of compositions and total pressures. The rightmost column estimates the time between collisions with nitrogen molecules in a pure nitrogen atmosphere at the given pressures - following the method suggested in section 3.2.2

The time between collisions with oxygen sets an important timescale in the system, since it is known that each collision with an oxygen molecule is likely to cause inter-system crossing (ISC) from the excited singlet state, S_1 to the triplet state, T_1 . Oxygen collisions also enhance the rate of ISC from T_1 to highly vibrationally-excited levels of S_0 .

These timescales are important because one of the easiest paths for the conversion of stored energy inside the molecule to heat is by conversion of vibrational energy into translational energy following interactions or collisions with other molecules. Absorption of a 266nm photon - as employed here - does not produce a molecule in a high vibrational level in S_1 , so an ISC event to T_1 must occur before energy may be released.

One may attempt to use this picture of the internal behaviour of the toluene molecule to attempt to associate particular processes with the fast and slow quenching decay

timescales. Clearly, the rate of energy release from vibrational de-excitation is extremely fast, and so this must be part of the fast timescale. When in isolation, the ‘lifetime’ of ISC from S_1 in toluene has been estimated in the range of 80ns. In the presence of collisions with other molecules - and particularly oxygen, this rate will probably fall somewhat, perhaps causing it to first be part of the long timescale, and then contribute to the short timescale as total pressure and oxygen partial pressure are increased. Again, in isolation, the ISC lifetime from the T_1 state is measured to be $\sim 50\mu\text{s}$. This is also bound to reduce in the presence of collisions - again, particularly oxygen. Note also that the oxygen-enhanced ISC from S_1 to T_1 often produces the metastable singlet oxygen state, which is long lived (approximately 80ms^{105} at STP) - eventually, though, this energy will be released as heating.

Ultimately, the attribution of particular processes to the long and short timescales is not important to the creation of a measurement technique, so long as some variation is discernible. The author suggests that the temporal resolution of this experimental system is insufficient to resolve how the fast timescale is changing, due to the limits of the pump pulse, oscilloscope and detector (ordered from most to least problematic). A future experiment may address this claim by addressing as many of these issues as possible.

1-bar quenching parameters review

Bearing the timescales of table 5.3 in mind, it seems that the only condition at which the experimental system in its current form might resolve quenching parameter information is at 1 bar total pressure. Therefore, figure 5.21 reviews the fitted quenching parameters at this condition.

Note that due to the reduced and consistently erratic signal magnitude, that each data point and associated error bar is formed from a set of only ~ 8 individual shots. This low sample size makes the evaluation of a trend very challenging - so even in this regime, where all timescales are at their longest, the data does not tell a clear story.




Figure 5.21.: Quenching model parameters for 1-bar data. Thin black line indicates rough trend in data.

Figure 5.21 could nonetheless indicate a poorly resolved and weak trend in both Q_1 and Q_2 . The split parameter, on the other hand, would be compelling, if only the error bars were not so tall. The author suggests that with a larger dataset and/or a more consistent signal strength, that this plot may well demonstrate sufficiently convincing behaviour.

5.6.3. Oxygen concentration by signal contrast

The previous sub-sections discuss the use of the parameters of the quenching function in order to discern the impact of an unknown concentration of oxygen on a LITGS signal. A further method has been explored to derive this information in a more straightforward way, which is now presented.

The visibility or contrast of a LITGS signal, is defined quite simply as:

$$\text{Contrast} = \frac{A - B}{A + B} \quad (5.19)$$

where the height of the first LITGS peak has value A and the first trough is recorded with a signal level B . Figure 5.22 shows how this contrast characteristic varies as a function of oxygen concentration.




Figure 5.22.: LITGS signal contrast as a function of oxygen partial pressure

Figure 5.22 plots mean values as data points, and standard deviations as error bars

(on the ‘090225’ data), with the same colour coding as all previous plots. Least-squares linear fits through the ‘090225’ datasets are also plotted on the graph. Note that the x axis is fractional oxygen content (expressed as a percentage), easily allowing the behaviour of air - at 21% oxygen - to be determined and compared to pure nitrogen at 0%.

The data and fits indicate that the contrast in the LITGS signal increases with total pressure, but also, more importantly, with increasing oxygen partial pressure[†]. It is not known why the trend is less pronounced in the data from the 9th March.

The physics underlying this behaviour may be explained by considering the following: The oscillations on a LITGS signal are caused by the interference of a static grating and an acoustic grating. If the two gratings are not created with equal magnitudes in density, then the signal will not be fully modulated as the acoustic standing wave evolves. As gas pressure is increased, the medium becomes increasingly ‘stiff’, facilitating the production of a pressure wave formed by a (temporally-) short, sharp deposition of heat. Furthermore, the inclusion of increasing levels of oxygen in the gas mixture shortens the timescale on which the quenching occurs (to timescales which are not resolved by the previous methods), which is also beneficial to the instigation of the acoustic shock wave. Ergo, increasing total pressure and oxygen partial pressure cause increased contrast.

This method may be developed by using more than just the first peak and trough to derive the contrast - instead, perhaps, encompassing the first five oscillations of the LITGS signal. Such an analysis has the potential to reduce the magnitude of the error bars, and thus increase the sensitivity of the method.

Before this method may be applied with confidence it must be investigated more thoroughly, to confirm that the variation is reproducible and that the physics of the underlying mechanism are well-understood.

[†]The reader may confirm this by examining the previously-plotted figures 5.18 and 5.19.

5.6.4. Temperature

The final aspect of a FARLIG measurement that will be addressed in this thesis is that of concurrently measuring temperature. (The technique could additionally be extended to measure pressure.) Temperature measurement is achieved, as has been described previously in sections 3.3.3 and 5.3, by use of the relationship between the inter-fringe transit time and gas temperature:

$$\tau = \sqrt{\frac{m}{\gamma k T}} \Lambda \quad (5.20)$$

where τ is the inter-fringe transit time, m is the mean molecular mass, γ is the ratio of specific heats of the gas, k is Boltzmann's constant, T is the temperature in Kelvin, and Λ is the grating plane spacing. Note that τ is easily measured by observing the oscillation frequency of a LITGS signal.

A simple demonstration of the application of LITGS to temperature measurement is now presented, whereby the temperature of the gas in the test cell is increased by heating the cell using heater tape with voltage/current varied by an autotransformer. A thermocouple is present inside the cell (see figure 5.5), which provides a reference signal. Note that the gas is not replaced in the cell as temperature is varied, such that the composition does not change. For these experiments, the gas cell was pressurised to 4 bar absolute with dry compressed air, and the 'standard' dose of toluene was admitted to the cell as a single droplet from the syringe/needle. The system was given some time to reach a thermal equilibrium following increased autotransformer voltage, with the intention of generating a uniform gas temperature inside the cell. (Equilibrium is confirmed by monitoring the thermocouple readout.)

Under these controlled circumstances, equation 5.20 can be used at some reference point of known temperature to derive a proportionality constant. Exactly this is performed in figure 5.23, where the point at 38°C is used to calibrate the LITGS measurement to match the thermocouple.

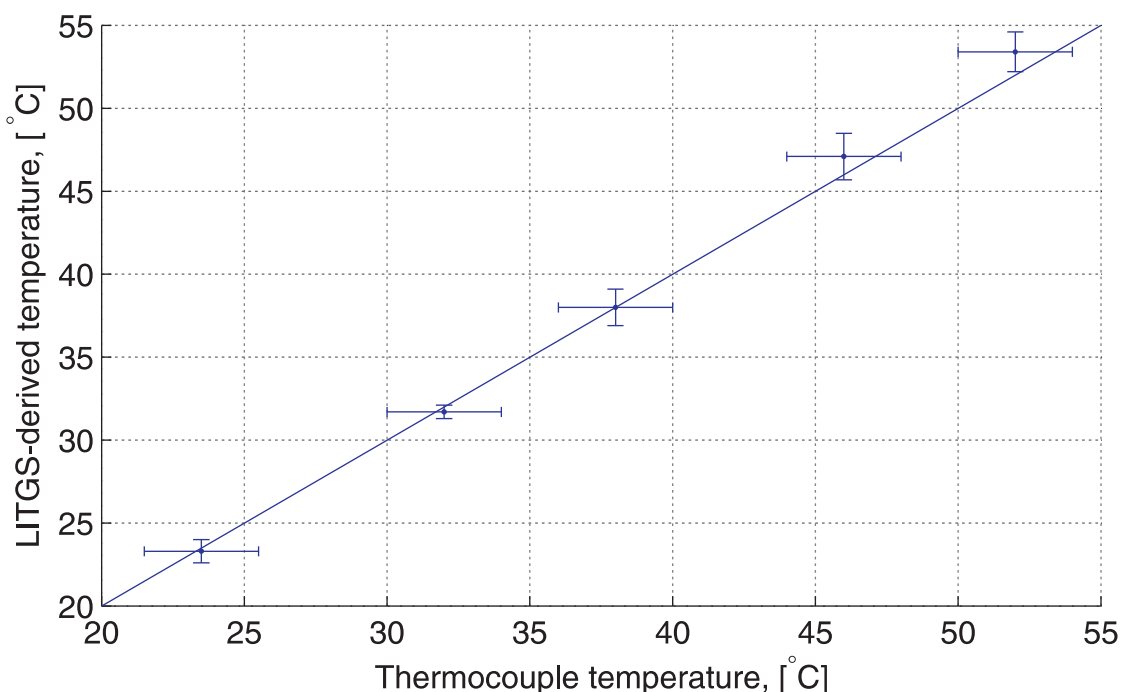


Figure 5.23.: Comparison of temperature measured by LITGS against that measured with a thermocouple.

Figure 5.23 shows very good agreement between the temperature derived by LITGS and the measurement on the thermocouple. The error bars[†] in the plot are determined in the case of the thermocouple by estimation at $\pm 2^\circ\text{C}$, and in the case of LITGS, by considering the spread of inter-fringe transit times produced by the `LIGSBatch.m` fitting routine (hence of variable height). The temperature monitoring capabilities of the LITGS technique are clearly very good, with a precision (and probable accuracy) of approximately 1 degree in 300 - or $\sim 0.3\%$.

However, it must be stressed that whilst the precision of LITGS is always good (and will improve at higher pressures, when more oscillations are visible) - the accuracy of the derived measurement is wholly dependent on good knowledge of the composition of the gas at the crossing point, without which the values of γ and m - which are required for equation 5.20 - must be guessed. (Note that Λ is determined with a test at known T and composition in preference to using inaccurate geometric methods.)

[†]Dual errorbar plotting code 'errorbare.m' developed by Yu Heng, and downloaded from <http://www.mathworks.com/matlabcentral/fileexchange/23465>

5.7. Summary

This chapter has demonstrated that FARLIG is capable of performing fuel concentration measurements, and accurate temperature measurements.

Although the analysis presented in the sections on fuel and oxygen concentrations is not totally compelling - a clear trend is yet to be derived from the oxygen concentration dataset, and the shot-to-shot fluctuations are too strong for fuel concentration measurements - both aspects of the technique are close to providing convincing results. In this respect, it is worth considering how the final FARLIG technique may operate. The author proposes that the FARLIG technique, when presented with a signal recorded at arbitrary conditions would:

1. Determine oxygen concentration, using one of two methods - either a ‘look-up’ table based on the quenching parameters Q_1 , split and Q_2 , or a similarly-calibrated version of the ‘contrast method’ of section 5.6.3.
2. Determine fuel concentration, potentially correcting the measurement based on the change in fluorescence quantum yield with oxygen concentration (since a side-effect of oxygen causing the quenching to occur more quickly is to cause energy which would otherwise be carried away in photons to be deposited as heat).
3. Determine temperature, using the information derived in steps 1 and 2 to inform the technique about the composition at the crossing point, thus limiting the compositional uncertainty.

The next steps for FARLIG are centered about three topics. The first important aspect which should be addressed is that of the shot-wise instability of the signal strength. A series of diagnostic experiments to determine the cause of the fluctuations would begin by attempting to isolate the problem to either the production of the grating (i.e. the pump laser) or the readout of the grating (i.e. the probe laser).

Potential sources of error include problems relating to beam-steering (of pump and/or probe) and unstable pump beam interference.

The second important investigation is to repeat the experiments presented here with a shorter pump pulse (ideally, on the order of 100ps) and a suitably faster detector/oscilloscope combination (e.g. minimum 1GHz bandwidth). It is expected that this repeated experiment would be more able to resolve the temporal detail which is encoded in the beginning of a FARLIG signal which is key to recovering oxygen concentration via the modelling method.

The third topic of future work relates to the ‘contrast method’, which was described in section 5.6.3. This method has the advantage of not requiring such fast detection equipment as the model-derived oxygen concentration method - provided the oscillations of a LITGS signal can be accurately resolved (something the current apparatus clearly *can* do), then the signal visibility/contrast may be determined. This method needs further investigation - firstly to verify that changes in signal contrast are repeatable, and secondly, to confirm that the physics of the underlying mechanism are correctly understood.

Providing these investigations are successful, the FARLIG technique would provide accurate, spatially-resolved, single shot measurements of fuel/air ratio and temperature. Such a technique could prove an asset to future engine development.

6. Conclusions

In this chapter, the QPLIF and FARLIG techniques' strengths and weaknesses will be discussed. The techniques are then compared to each other, and to their respective closest competitors. This chapter closes by briefly summarising the thesis.

6.1. QPLIF

PLIF (Planar Laser Induced Fluorescence) has an innate strength - which is its ability to create an image, simultaneously probing many locations and revealing spatial relationships. This strength is countered by the reliance of quantitative interpretation on absolute intensity measurements, which are prone to error resulting from any change in optical collection efficiency. The quantum efficiency of the fluorescent tracer must also be carefully calibrated to generate quantitative distributions, since the environment the molecule finds itself in can have a pronounced effect on the number of photons that are emitted - and any uncalibrated behaviour can easily lead to mis-interpretation of the recorded images.

The *in-situ* online calibration method that has been developed in this thesis, which forms a key element of the Quantitative-PLIF (QPLIF) technique provides a means for generating a calibration map of quantum efficiency and to monitor the transmission of the window elements. The fired nature of the 'bookend' calibration scheme also allows for the effects of EGR (exhaust gas residuals) to be incorporated into the calibration it produces. The window fouling correction is seen to account for some 20% signal

reduction over the course of a test, while the presence of EGR typically cuts the fluorescent signal strength in half.

This point leads to an important limit inherent to the QPLIF technique, which is that the low tracer doping levels combined with fired operation conspire to produce diminished fluorescent intensity. This unavoidably small signal manifests itself in the high noise that is present in the measured fuel distributions.

Whilst the calibration method employed in the QPLIF technique is able to account for the average change in temperature which occurs during an engine cycle, any stratification or non-uniformity in the temperature field would be incorrectly interpreted as differing fuel concentration. (Clearly, this limitation also applies to how well-mixed the EGR gas is with the fresh charge.) The literature indicates that such fluctuations can vary in magnitude, from little²⁸ to significant³⁷. The degree of inhomogeneity appears to depend largely upon the operating mode of the engine and the timing of the imaging.

Note that the methods employed in the QPLIF technique are adaptable to any IC engine, provided a homogeneous mixture may be formed and used in a fired cycle to calibrate the system.

6.2. FARLIG

LITGS (Laser Induced Thermal Grating Spectroscopy) offers measurements of relevant combustion parameters in a well-defined localised probe volume, defined by the intersection of its pump laser beams. This measurement can be extended along a line. LITGS - and, by extension FARLIG (Fuel/Air Ratio by Laser Induced thermal Gratings) - provides a very precise measurement of temperature, which does not rely upon the measurement of intensity, and is thus immune to window fouling errors. The power of this important attribute is somewhat tempered by the accuracy of the measurement relying upon foreknowledge of the composition of the probed gas. Fortunately, fuel

vapour normally comprises a small fraction of the overall mixture - but it remains to be determined just how large an impact its relatively large molecular mass and ratio of specific heats may have under typical mixing conditions. Helpfully, LITGS techniques require only line-of-sight access to the probe volume.

Unfortunately, the work presented in this thesis has been unable to demonstrate the viability of the FARLIG extension to LITGS conclusively. However, assuming that forthcoming experiments yield satisfactory results, the technique would derive temperature and oxygen concentration from the temporal shape of the signal, and only the fuel concentration measurement would rely upon measuring absolute intensities accurately. This is a step in the right direction from a PLIF-based technique.

Note that both methods suggested in this thesis for oxygen concentration measurements probably have an effective limit on the pressure up to which they may be employed. For the use of the quenching model, the limit is imposed by the rate at which quenching occurs, and seems likely to be in the region of 4-8 bar. For the use of the 'contrast method', the limit will be defined by the pressure at which the contrast is already high, before the addition of oxygen - and could be around 8 bar.

6.3. Comparisons

6.3.1. QPLIF vs. FARLIG

When choosing a technique to deploy for a measurement campaign, the choice between QPLIF and FARLIG is largely motivated by whether one needs information across a plane, or if the evolution of a single location is of more relevance. For example, if the transport of fuel around the cylinder is of primary interest, then PLIF has a clear advantage. If an investigation of a localised region - for example, near the spark plug - is called for, then monitoring both fuel/air ratio (not just fuel concentration) *and* having a precise measure of temperature may be of greater utility (i.e. using FARLIG).

The two techniques do also complement each other somewhat. With relaxed requirements on quantitative analysis, the doping concentration may be increased for PLIF experiments, resulting in clear - albeit qualitative - images. A working FARLIG technique could then be deployed alongside to monitor the fuel concentration in the most important region inside the engine (perhaps at a higher repetition rate than is easily achieved with PLIF), and be used to ‘peg’ the PLIF image to a rough value, whilst also providing temperature and FAR history.

6.3.2. Novelty in QPLIF

Section 2.8.1 identified the technique described by Han and Steeper³² to be the most relevant previously-published work. The QPLIF technique reported in this thesis utilises what appears to be a similar calibration method, but a number of improvements have been made. To address the points raised in section 2.8.1:

Firstly, the QPLIF technique covered in this thesis uses a fully-fractionated base fuel, which has been shown to have all the properties of a pump gasoline - making the results and conclusions of this work highly relevant to the development of the engine.

Secondly, the engine is always operated in the steady-state in these experiments, such that no variability caused by skipped operation is introduced.

Third, the ‘bookend’ calibration independently developed in this thesis includes quantum efficiency calibration on a fired engine cycle, and provides a window fouling correction. This thesis also provides - to the author’s best knowledge - the first report of the typical impact of EGR on PLIF signal strength.

To address the fourth point, the plane of the measurements lies in the same plane as the spark plug, enabling the most important region in the engine to be investigated.

Taking points 5 and 6 together - images have been recorded across a range of crank-angles around the engine cycle, indicating that the three fuel fractions exhibit very similar behaviour.

To address point 7 - the rate of image capture is limited in these experiments by the ICCD. Lack of a camera with faster readout rate is all that prevents the QPLIF technique from being applied to capture multiple images within a single engine cycle.

6.3.3. Novelty in FARLIG

Section 2.8.2 selected Seeger et al.⁸⁷ as the most similar previously-published research. The FARLIG technique offers the following relative advantages and disadvantages:

FARLIG can conceptually measure oxygen concentration from the signal shape, and has been demonstrated to be able to measure temperature from the oscillation frequency (not possible in Seeger et al.⁸⁷, where oscillation frequency is used for composition). FARLIG has the advantage of not requiring the focussed geometry to generate a signal, which avoids problems such as gas breakdown.

The techniques share a dependence on an electronic resonance, though FARLIG is tuned onto it, whereas the compared work is tuned far from resonance.

While the FARLIG technique does appear to have a stronger signal than that reported by Seeger et al.⁸⁷ - which is probably due to the use of a long-pulsed dye laser as readout beam - the $\sim 40\%$ accuracy of the fuel concentration measurement reported by Seeger et al.⁸⁷ could not be improved upon in the present work.

Indeed, the full capabilities of FARLIG remain to be identified, since the experiments reported here were unable to demonstrate a wholly convincing measurement of oxygen concentration, and signal stability limits the accuracy of the measurement of fuel concentration.

6.4. Summary

In summary, therefore, two optical diagnostic techniques have been discussed in detail and applied in this thesis. The following serves as a précis:

The QPLIF technique - incorporating for the first time a 6-component base fuel acting as a surrogate for pump gasoline, mated with three fluorescent tracers of differing volatility - has been applied to the quantitative measurement of fuel distributions. The average fuel distribution was reported as a function of crank angle with 10% error on a super-pixel. Encouraging agreement was also demonstrated, when these measured distributions were compared to Jaguar's in-house CFD model of the same engine.

The three fuel fractions were independently tracked by the co-evaporating tracers during an engine cycle. No significant differences were found between the behaviour of the fuel fractions. This observation suggests that whilst the base fuel must be representative of a real gasoline in order to generate the correct spray pattern, it is not important to consider the behaviour of each fraction individually.

The 'bookend' calibration mechanism has been shown to account for the varying quantum efficiency of the three tracers which have been employed, and also provides a correction for the deposition of combustion products during an experimental run. The use of a fired calibration cycle has allowed the effect of EGR to be incorporated. The magnitude of the effect of EGR has been found to be very significant, roughly halving the fluorescent intensity.

The FARLIG technique has been demonstrated able to measure temperature precisely (just 0.3% uncertainty), and good grounds have been provided to indicate accurate measurement of fuel concentration should be achievable.

The presented results are not sufficient to demonstrate the viability of either of the two proposed methods (quenching model or contrast method) to the measurement of oxygen concentration. However, both concepts show promise and will be the subject of future investigation.

A fully-functional FARLIG technique would enable a spatially- and temporally-resolved single-shot measurement of fuel concentration, oxygen concentration (hence FAR) and temperature.

7. Future Prospects

In this chapter, future topics of investigation for the QPLIF and FARLIG techniques are discussed.

7.1. QPLIF

The *in-situ* calibration technique at the core of QPLIF is able to correct for many variables in its calibration map. However, the technique is still vulnerable to inhomogeneous temperature fields. Previously published work indicates that temperature inhomogeneity and cyclic variability may be significant, depending on engine operating conditions. It is also well known that temperature has a pronounced effect on the magnitude of a PLIF signal, so a means to monitor its change might prove beneficial. Temperature field measurements with PLIF have already been demonstrated - for example, by Luong et al.³⁵. Integrating a similar technique into QPLIF would prove challenging on the detector - since it would involve splitting the signal spectrally - but the extra information may outweigh this drawback.

Indeed, future experiments would benefit from a more sensitive/lower noise intensified camera - should one be available.

Future experiments should also aim to determine the cause of the signal drop in the presence of EGR for acetone - since it does not appear to be explained on the grounds of quantum efficiency reduction resulting from temperature rise. Further investigation of absorption and emission spectra may help to resolve this issue.

Finally, the QPLIF technique should be applied at other operating conditions, in order to generate further points of reference between experiments and computer models, and to aid in the development of new strategies for reducing engine emissions and improving efficiency.

7.2. FARLIG

One important area of investigation for future FARLIG work is to determine the cause of the shot-wise instability of the LITGS signal in these experiments - be it related to the creation or the read-out process - and, if possible, apply a solution. The fuel concentration measurement component of FARLIG could then be applied on a shot-wise basis.

Early results from a path-length compensated pair of probe beams has shown that matching the path length of the pump pulses to less than the coherence length of the pump laser causes better shot-to-shot reproducibility and much higher signal strength. This has important consequences for FARLIG, in terms of allowing the measurement of fuel concentration, the adoption of weaker probe lasers and/or the measurement of weaker transitions. Physically, this improvement is due to increased contrast in the laser fringes when the two compensated pump beams overlap in the test volume, which leads to increased density (and thus refractive index) modulation.

The accuracy of the temperature measurement in all LITGS-based techniques is limited by knowledge of the gas composition in the measurement volume. It would therefore be useful to perform a 'sensitivity analysis', estimating and/or practically determining the effect of varying the fuel concentration on the inter-fringe transit time, τ .

The 'contrast method' for determining oxygen concentration in FARLIG warrants further investigation.

The determination of oxygen concentration by use of a fit to a quenching model,

which has been a major focus of the work reported here, appears to currently be limited by the duration of the laser pump pulse. As mentioned in the conclusions section, if this were reduced to ~ 100 ps in length, and the detector and oscilloscope had matching bandwidth, then it is hoped that some variation in the Q_1 parameter could be identified. Coupled with the ‘split’ parameter - which already appears to exhibit reasonable sensitivity - the technique would then be able to determine oxygen concentration.

The final topic of future FARLIG research would be to apply the technique to the measurement of temperature in an environment where charge cooling occurs. Charge cooling is an effect important in direct-injection engines, where the vaporisation of fuel droplets causes a reduction in gas temperature since the latent heat of vaporisation must be supplied to perform the phase change. The change in temperature is thought to be only approximately 10K, which is challenging for most techniques to resolve, but should be readily achievable for LITGS. Measurement of fuel concentration, coupled with the knowledge generated by performing the ‘sensitivity analysis’ described above would probably enhance the accuracy of the temperatures measured using this method.

Bibliography

- [1] DEFRA. e-Digest Statistics about: Climate Change, 2008.
- [2] EC. Regulation 715/2007 - Euro 5 and Euro 6 standards, 2007.
- [3] D. K. Manley, A. McIlroy, and C. A. Taatjes. Research needs for future internal combustion engines. *Physics Today*, 61(11):47–52, 2008.
- [4] M. C. Drake and D. C. Haworth. Advanced gasoline engine development using optical diagnostics and numerical modeling. *Proceedings of the Combustion Institute*, 31:99–124, 2007.
- [5] W. K. Cheng, T. Summers, and N. Collings. The fast-response flame ionization detector. *Progress in Energy and Combustion Science*, 24(2):89–124, 1998.
- [6] J. D. Smith and V. Sick. Crank-angle resolved imaging of biacetyl laser-induced fluorescence in an optical internal combustion engine. *Applied Physics B-Lasers and Optics*, 81(5):579–584, 2005.
- [7] R. Stevens and P. Ewart. Single-shot measurement of temperature and pressure using laser-induced thermal gratings with a long probe pulse. *Applied Physics B-Lasers and Optics*, 78(1):111–117, 2004.
- [8] B. N. Edwards and D. E. Burch. Absorption of 3.39-micron Helium-Neon laser emission by methane in atmosphere. *Journal of the Optical Society of America*, 55(2):174–&, 1965.

- [9] T Berg, V Beushausen, O Thiele, and H Voges. Fibre Optics Spark Plug Sensor for the Optimisation of Engine Combustion Processes. *Motortechnische Zeitschrift (MTZ)*, 67, 2006.
- [10] E Tomita, N Kawahara, M Shigenaga, A Nishiyama, and RW Dibble. In situ measurement of hydrocarbon fuel concentration near a spark plug in an engine cylinder using the 3.392 μm infrared laser absorption method: discussion of applicability with a homogeneous methane-air mixture. *Measurement Science & Technology*, 14(8):1350–1356, 2003.
- [11] R. K. Hanson and P. K. Falcone. Temperature measurement technique for high-temperature gases using a tunable diode-laser. *Applied Optics*, 17(16):2477–2480, 1978.
- [12] D. S. Baer, R. K. Hanson, M. E. Newfield, and Nkjm Gopaul. Multiplexed diode-laser sensor system for simultaneous H₂O, O₂, and temperature measurements. *Optics Letters*, 19(22):1900–1902, 1994.
- [13] Y. Arita, R. Stevens, and P. Ewart. Multi-mode absorption spectroscopy of oxygen for measurement of concentration, temperature and pressure. *Applied Physics B-Lasers and Optics*, 90(2):205–211, 2008.
- [14] F. P. Chen and R. Goulard. Retrieval of arbitrary concentration and temperature fields by multiangular scanning techniques. *Journal of Quantitative Spectroscopy & Radiative Transfer*, 16(10):819–827, 1976.
- [15] R. J. Santoro, H. G. Semerjian, P. J. Emmerman, and R. Goulard. Optical Tomography for Flow Field Diagnostics. *International Journal of Heat and Mass Transfer*, 24(7):1139–1150, 1981.
- [16] N Terzija, JL Davidson, CA Garcia-Stewart, P Wright, KB Ozanyan, S Pegrum, TJ Litt, and H McCann. Image optimization for chemical species tomography

- with an irregular and sparse beam array. *Measurement Science & Technology*, 19(9):–, 2008.
- [17] C. Schulz and V. Sick. Tracer-LIF diagnostics: quantitative measurement of fuel concentration, temperature and fuel/air ratio in practical combustion systems. *Progress in Energy and Combustion Science*, 31(1):75–121, 2005.
- [18] K. Kuwahara and H. Ando. Diagnostics of in-cylinder flow, mixing and combustion in gasoline engines. *Measurement Science & Technology*, 11(6):R95–R111, 2000.
- [19] J. Reboux, D. Puechberty, and F. Dionnet. A New Approach of Planar Laser Induced Fluorescence Applied to Fuel/Air Ratio Measurement in the Compression Stroke of An Optical S.I. Engine. *Society of Automotive Engineers*, 941988, 1994.
- [20] W. Koban, J. D. Koch, R. K. Hanson, and C. Schulz. Toluene LIF at elevated temperatures: implications for fuel-air ratio measurements. *Applied Physics B-Lasers and Optics*, 80(2):147–150, 2005.
- [21] F. Grossmann, P. B. Monkhouse, M. Ridder, V. Sick, and J. Wolfrum. Temperature and pressure dependences of the laser-induced fluorescence of gas-phase acetone and 3-pentanone. *Applied Physics B-Lasers and Optics*, 62(3):249–253, 1996.
- [22] M. C. Thurber and R. K. Hanson. Pressure and composition dependences of acetone laser-induced fluorescence with excitation at 248, 266, and 308 nm. *Applied Physics B-Lasers and Optics*, 69(3):229–240, 1999.
- [23] W. Koban, J. D. Koch, R. K. Hanson, and C. Schulz. Oxygen quenching of toluene fluorescence at elevated temperatures. *Applied Physics B-Lasers and Optics*, 80(6):777–784, 2005.

- [24] J. D. Koch and R. K. Hanson. Temperature and excitation wavelength dependencies of 3-pentanone absorption and fluorescence for PLIF applications. *Applied Physics B-Lasers and Optics*, 76(3):319–324, 2003.
- [25] P. Guibert, V. Modica, and C. Morin. Influence of pressure, temperature and gas phase composition on biacetyl laser-induced fluorescence. *Experiments in Fluids*, 40(2):245–256, 2006.
- [26] M Berckmueller, N Tait, and D Greenhalgh. The Influence of Local Fuel Concentration on the Cyclic Variability of a Lean Burn Stratified Charge Engine. *Society of Automotive Engineers*, 970826, 1997.
- [27] W. Koban, J. D. Koch, V. Sick, N. Wermuth, R. K. Hanson, and C. Schulz. Predicting LIF signal strength for toluene and 3-pentanone under engine-related temperature and pressure conditions. *Proceedings of the Combustion Institute*, 30:1545–1553, 2005.
- [28] S Einecke, C Schulz, and V Sick. Measurement of temperature, fuel concentration and equivalence ratio fields using tracer LIF in IC engine combustion. *Applied Physics B - Lasers and Optics*, 71(5):717–723, 2000.
- [29] G. de Sercey, M. Heikal, M. Gold, S. Begg, R. Wood, G. Awcock, and O. Laguitton. On the use of laser-induced fluorescence for the measurement of in-cylinder air-fuel ratios. *Proceedings of the Institution of Mechanical Engineers Part C-Journal of Mechanical Engineering Science*, 216(10):1017–1029, 2002.
- [30] J. D. Smith and V. Sick. Quantitative, dynamic fuel distribution measurements in combustion-related devices using laser-induced fluorescence imaging of biacetyl in iso-octane. *Proceedings of the Combustion Institute*, 31:747–755, 2007.
- [31] T. Blotevogel, M. Hartmann, H. Rottengruber, and A. Leipertz. Tracer-based laser-induced fluorescence measurement technique for quantitative fuel/air-ratio

- measurements in a hydrogen internal combustion engine. *Applied Optics*, 47(35): 6488–6496, 2008.
- [32] D. Han and R. R. Steeper. An LIF equivalence ratio imaging technique for multicomponent fuels in an IC engine. *Proceedings of the Combustion Institute*, 29:727–734, 2003.
- [33] M. C. Thurber, F. Grisch, B. J. Kirby, M. Votsmeier, and R. K. Hanson. Measurements and modeling of acetone laser-induced fluorescence with implications for temperature-imaging diagnostics. *Applied Optics*, 37(21):4963–4978, 1998.
- [34] M. C. Thurber and R. K. Hanson. Simultaneous imaging of temperature and mole fraction using acetone planar laser-induced fluorescence. *Experiments in Fluids*, 30(1):93–101, 2001.
- [35] M. Luong, R. Zhang, C. Schulz, and V. Sick. Toluene laser-induced fluorescence for in-cylinder temperature imaging in internal combustion engines. *Applied Physics B-Lasers and Optics*, 91(3-4):669–675, 2008.
- [36] C. Morin, V. Modica, and P. Guibert. Measurement of exhaust gas recirculation rate by laser-induced fluorescence in engine. *Measurement Science & Technology*, 19(10):–, 2008.
- [37] DA Rothamer, JA Snyder, RK Hanson, RR Steeper, and RP Fitzgerald. Simultaneous imaging of exhaust gas residuals and temperature during HCCI combustion. *Proceedings of the Combustion Institute*, 32:2869–2876, 2009.
- [38] W. Koban, J. Schorr, and C. Schulz. Oxygen-distribution imaging with a novel two-tracer laser-induced fluorescence technique. *Applied Physics B-Lasers and Optics*, 74(1):111–114, 2002.
- [39] D. Frieden, V. Sick, J. Gronki, and C. Schulz. Quantitative oxygen imaging in an engine. *Applied Physics B-Lasers and Optics*, 75(1):137–141, 2002.

- [40] B. D. Stojkovic and V. Sick. Evolution and impingement of an automotive fuel spray investigated with simultaneous Mie/LIF techniques. *Applied Physics B-Lasers and Optics*, 73(1):75–83, 2001.
- [41] Lynn A. Melton and James F. Verdieck. Vapor/Liquid Visualization for Fuel Sprays. *Combustion Science and Technology*, 42(3 & 4):217–222, 1985.
- [42] P Wieske, S Wissel, G Grunefeld, and S Pischinger. Improvement of LIEF by wavelength-resolved acquisition of multiple images using a single CCD detector - Simultaneous 2D measurement of air/fuel ratio, temperature distribution of the liquid phase and qualitative distribution of the liquid phase with the Multi-2D technique. *Applied Physics B - Lasers and Optics*, 83(2):323–329, 2006.
- [43] R Hasegawa, I Sakata, H Yanagihara, B Johansson, A Omrane, and M Alden. Two-dimensional gas-phase temperature measurements using phosphor thermometry. *Applied Physics B - Lasers and Optics*, 88(2):291–296, 2007.
- [44] J. Bruebach, E. van Veen, and A. Dreizler. Combined phosphor and CARS thermometry at the wall-gas interface of impinging flame and jet systems. *Experiments in Fluids*, 44(6):897–904, 2008.
- [45] J Kojima, Y Ikeda, and T Nakajima. Basic aspects of OH(A), CH(A), and C-2(d) chemiluminescence in the reaction zone of laminar methane-air premixed flames. *Combustion and Flame*, 140(1-2):34–45, 2005.
- [46] CS Panoutsos, Y Hardalupas, and AMKP Taylor. Numerical evaluation of equivalence ratio measurement using OH* and CH* chemiluminescence in premixed and non-premixed methane-air flames. *Combustion and Flame*, 156(2):273–291, 2009.
- [47] F. Q. Zhao and H. Hiroyasu. The applications of laser Rayleigh-scattering to

- combustion diagnostics. *Progress in Energy and Combustion Science*, 19(6):447–485, 1993.
- [48] C. Espey, J. E. Dec, T. A. Litzinger, and D. A. Santavicca. Planar laser Rayleigh scattering for quantitative vapor-fuel imaging in a diesel jet. *Combustion and Flame*, 109(1-2):65–86, 1997.
- [49] Alan C. Eckbreth. *Laser Diagnostics for Combustion Temperature and Species*. Combustion Science and Technology. Gordon and Breach Publishers, second edition edition, 1996.
- [50] R. S. Barlow and C. D. Carter. Raman-Rayleigh-LIF measurements of nitric-oxide formation in turbulent hydrogen jet flames. *Combustion and Flame*, 97(3-4):261–280, 1994.
- [51] A. R. Masri, R. W. Dibble, and R. S. Barlow. The structure of turbulent non-premixed flames revealed by Raman-Rayleigh-LIF measurements. *Progress in Energy and Combustion Science*, 22(4):307–362, 1996.
- [52] R. S. Barlow, A. N. Karpetsis, J. H. Frank, and J. Y. Chen. Scalar profiles and NO formation in laminar opposed-flow partially premixed methane/air flames. *Combustion and Flame*, 127(3):2102–2118, 2001.
- [53] R. Cabra, T. Myhrvold, J. Y. Chen, R. W. Dibble, A. N. Karpetsis, and R. S. Barlow. Simultaneous laser Raman-Rayleigh-LIF measurements and numerical modeling results of a lifted turbulent H-2/N-2 jet flame in a vitiated coflow. In *29th International Combustion Symposium*, pages 1881–1888, Sapporo, Japan, 2002. Combustion Inst.
- [54] Robert W. Boyd. *Nonlinear Optics*. Academic Press, 3 edition, 2008.
- [55] A. C. Eckbreth. BOXCARS - crossed-beam phase-matched CARS generation in gases. *Applied Physics Letters*, 32(7):421–423, 1978.

- [56] A. C. Eckbreth and T. J. Anderson. Dual broad-band CARS for simultaneous, multiple species measurements. *Applied Optics*, 24(16):2731–2736, 1985.
- [57] T. Dreier, B. Lange, J. Wolfrum, and M. Zahn. Determination of temperature and concentration of molecular nitrogen, oxygen and methane with Coherent-Anti-Stokes-Raman-Scattering. *Applied Physics B-Photophysics and Laser Chemistry*, 45(3):183–190, 1988.
- [58] R. E. Palmer. The CARSFT Computer Code for Calculating Coherent Anti-Stokes Raman Spectra: Use Guide and Programmer Information. *SANDIA National Labs. Reports*, SAND89-8206, 1989.
- [59] M. Alden, P. E. Bengtsson, H. Edner, S. Kroll, and D. Nilsson. Rotational CARS - A comparison of different techniques with emphasis on accuracy in temperature determination. *Applied Optics*, 28(15):3206–3219, 1989.
- [60] T. Seeger and A. Leipertz. Experimental comparison of single-shot broadband vibrational and dual-broadband pure rotational coherent anti-Stokes Raman scattering in hot air. *Applied Optics*, 35(15):2665–2671, 1996.
- [61] C. Brackmann, J. Bood, P. E. Bengtsson, T. Seeger, M. Schenk, and A. Leipertz. Simultaneous vibrational and pure rotational coherent anti-Stokes Raman spectroscopy for temperature and multispecies concentration measurements demonstrated in sooting flames. *Applied Optics*, 41(3):564–572, 2002.
- [62] F. Beyrau, T. Seeger, A. Malarski, and A. Leipertz. Determination of temperatures and fuel/air ratios in an ethene-air flame by dual-pump CARS. In *European Conference on Nonlinear Optical Spectroscopy (ECONOS)*, pages 946–951, Besancon, France, 2003. John Wiley & Sons Ltd.
- [63] M. C. Weikl, F. Beyrau, and A. Leipertz. Simultaneous temperature and exhaust-gas recirculation-measurements in a homogeneous charge-compression ignition

- engine by use of pure rotational coherent anti-Stokes Raman spectroscopy. *Applied Optics*, 45(15):3646–3651, 2006.
- [64] T. Doerk, M. Hertl, B. Pfelzer, S. Hadrich, P. Jauernik, and J. Uhlenbusch. Resonance enhanced coherent anti-stokes Raman scattering and Laser induced fluorescence applied to CH radicals: A comparative study. *Applied Physics B-Lasers and Optics*, 64(1):111–118, 1997.
- [65] R. L. Farrow and D. J. Rakestraw. Detection of trace molecular species using degenerate 4-wave-mixing. *Science*, 257(5078):1894–1900, 1992.
- [66] Katharina Kohse-Hoeinghaus and Jay B. Jeffries. *Applied Combustion Diagnostics*. Combustion: An International Series. Taylor & Francis, London, 1 edition, 2002.
- [67] P. Ewart and P. Snowdon. Multiplex degenerate 4-wave-mixing in a flame. *Optics Letters*, 15(23):1403–1405, 1990.
- [68] B. Yip, P. M. Danehy, and R. K. Hanson. Degenerate 4-wave-mixing temperature measurements in a flame. *Optics Letters*, 17(10):751–753, 1992.
- [69] T. Dreier and D. J. Rakestraw. Measurement of OH rotational temperatures in a flame using degenerate 4-wave mixing. *Optics Letters*, 15(1):72–74, 1990.
- [70] P. Ewart, P. Snowdon, and I. Magnusson. Two-dimensional phase-conjugate imaging of atomic distributions in flames by degenerate 4-wave mixing. *Optics Letters*, 14(11):563–565, 1989.
- [71] P. Ewart and M. Kaczmarek. 2-dimensional mapping of temperature in a flame by degenerate 4-wave-mixing in OH. *Applied Optics*, 30(27):3996–4000, 1991.
- [72] T. A. Reichardt, W. C. Giancola, C. M. Shappert, and R. P. Lucht. Experimental investigation of saturated degenerate four-wave mixing for quantitative concentration measurements. *Applied Optics*, 38(33):6951–6961, 1999.

- [73] A. J. Grant, P. Ewart, and C. R. Stone. Detection of NO in a spark-ignition research engine using degenerate four-wave mixing. *Applied Physics B-Lasers and Optics*, 74(1):105–110, 2002.
- [74] P. M. Danehy, P. H. Paul, and R. L. Farrow. Thermal grating contributions to degenerate 4-wave-mixing in nitric-oxide. *Journal of the Optical Society of America B-Optical Physics*, 12(9):1564–1576, 1995.
- [75] A. Stampanoni-Panariello, B. Hemmerling, and W. Hubschmid. Temperature measurements in gases using laser-induced electrostrictive gratings. *Applied Physics B-Lasers and Optics*, 67(1):125–130, 1998.
- [76] W. Hubschmid, R. Bombach, B. Hemmerling, and A. StampanoniPanariello. Sound-velocity measurements in gases by laser-induced electrostrictive gratings. *Applied Physics B-Lasers and Optics*, 62(1):103–107, 1996.
- [77] A. Stampanoni-Panariello, D. N. Kozlov, P. P. Radi, and B. Hemmerling. Gas-phase diagnostics by laser-induced gratings II. Experiments. *Applied Physics B-Lasers and Optics*, 81(1):113–129, 2005.
- [78] E. B. Cummings. Laser-Induced Thermal Acoustics - simple accurate gas measurements. *Optics Letters*, 19(17):1361–1363, 1994.
- [79] P. H. Paul, R. L. Farrow, and P. M. Danehy. Gas-phase thermal-grating contributions to 4-wave mixing. *Journal of the Optical Society of America B-Optical Physics*, 12(3):384–392, 1995.
- [80] R. C. Hart, R. J. Balla, and G. C. Herring. Nonresonant referenced laser-induced thermal acoustics thermometry in air. *Applied Optics*, 38(3):577–584, 1999.
- [81] R. Stevens and P. Ewart. Simultaneous single-shot measurement of temperature and pressure along a one-dimensional line by use of laser-induced thermal grating spectroscopy. *Optics Letters*, 31(8):1055–1057, 2006.

- [82] B. Hemmerling, P. Radi, A. Stampanoni-Panariello, A. Kouzov, and D. Kozlov. Novel non-linear optical techniques for diagnostics: laser-induced gratings and two-colour four-wave mixing. *Comptes Rendus De L Academie Des Sciences Serie Iv Physique Astrophysique*, 2(7):1001–1012, 2001.
- [83] B. Hemmerling, D. N. Kozlov, O. M. Stel'makh, and B. Attal-Tretout. Diagnostics of water-containing gas mixtures using thermal laser-induced gratings. *Chemical Physics*, 320(2-3):103–117, 2006.
- [84] DJW Walker, RB Williams, and P Ewart. Thermal grating velocimetry. *Optics Letters*, 23(16):1316–1318, 1998.
- [85] R. C. Hart, G. C. Herring, and R. J. Balla. Pressure measurement in supersonic air flow by differential absorptive laser-induced thermal acoustics. *Optics Letters*, 32(12):1689–1691, 2007.
- [86] D. N. Kozlov. Simultaneous characterization of flow velocity and temperature fields in a gas jet by use of electrostrictive laser-induced gratings. *Applied Physics B-Lasers and Optics*, 80(3):377–387, 2005.
- [87] T. Seeger, J. Kiefer, M. C. Weikl, A. Leipertz, and D. N. Kozlov. Time-resolved measurement of the local equivalence ratio in a gaseous propane injection process using laser-induced gratings. *Optics Express*, 14(26):12994–13000, 2006.
- [88] J. Kiefer, D. N. Kozlov, T. Seeger, and A. Leipertz. Local fuel concentration measurements for mixture formation diagnostics using diffraction by laser-induced gratings in comparison to spontaneous Raman scattering. In *European Conference on Nonlinear Optical Spectroscopy*, pages 711–721, St Petersburg, RUSSIA, 2007. John Wiley & Sons Ltd.
- [89] Robert C. Weast. *CRC Handbook of Chemistry and Physics*, 1980.

- [90] E. B. Cummings, H. G. Hornung, M. S. Brown, and P. A. Debarber. Measurement of gas-phase sound speed and thermal diffusivity over a broad pressure range using Laser Induced Thermal Acoustics. *Optics Letters*, 20(14):1577–1579, 1995.
- [91] H. Latzel, A. Dreizler, T. Dreier, J. Heinze, M. Dillmann, W. Stricker, G. M. Lloyd, and P. Ewart. Thermal grating and broadband degenerate four-wave mixing spectroscopy of OH in high-pressure flames. *Applied Physics B-Lasers and Optics*, 67(5):667–673, 1998.
- [92] W. Hubschmid, B. Hemmerling, and A. Stampanonipaniariello. Rayleigh and Brillouin modes in electrostrictive gratings. *Journal of the Optical Society of America B-Optical Physics*, 12(10):1850–1854, 1995.
- [93] B. Hemmerling and D. N. Kozlov. Generation and temporally resolved detection of laser-induced gratings by a single, pulsed Nd : YAG laser. *Applied Optics*, 38(6):1001–1007, 1999.
- [94] Walter G. Vincenti and Charles H. Kruger. *Introduction to Physical Gas Dynamics*. Wiley, New York, 1965.
- [95] H Eichler, P Guenter, and D Pohl. *Laser-Induced Gratings*, volume 50 of *Springer Series in Optical Science*. Springer, Berlin, 1986.
- [96] A. Stampanoni-Panariello, D. N. Kozlov, P. P. Radi, and B. Hemmerling. Gas phase diagnostics by laser-induced gratings I. theory. *Applied Physics B-Lasers and Optics*, 81(1):101–111, 2005.
- [97] B. Hemmerling, W. Hubschmid, and A. Stampanoni-Panariello. Temperature and mixture fraction measurements in gases by laser-induced electrostrictive gratings. In A. R. Burgess and F. L. Dryer, editors, *27th International Symposium on Combustion*, pages 69–75, Boulder, Co, 1998. Combustion Institute.

- [98] W. Koban, J. D. Koch, R. K. Hanson, and C. Schulz. Absorption and fluorescence of toluene vapor at elevated temperatures. *Physical Chemistry Chemical Physics*, 6(11):2940–2945, 2004.
- [99] F. P. Zimmermann, W. Koban, C. M. Roth, D. P. Herten, and C. Schulz. Fluorescence lifetime of gas-phase toluene at elevated temperatures. *Chemical Physics Letters*, 426(4-6):248–251, 2006.
- [100] Jon Koch. *Fuel tracer photophysics for quantitative Planar Laser-Induced Fluorescence*. PhD thesis, 2005.
- [101] T. G. Dietz, M. A. Duncan, and R. E. Smalley. Time evolution studies of triplet toluene by 2-color photo-ionization. *Journal of Chemical Physics*, 76(3):1227–1232, 1982.
- [102] K. Kikuchi, C. Sato, M. Watabe, H. Ikeda, Y. Takahashi, and T. Miyashi. New aspects on fluorescence quenching by molecular oxygen. *Journal of the American Chemical Society*, 115(12):5180–5184, 1993.
- [103] K. Kawaoka, A. U. Khan, and D. R. Kearns. Role of singlet excited states of molecular oxygen in quenching of organic triplet states. *Journal of Chemical Physics*, 46(5):1842–&, 1967.
- [104] Y. Usui, N. Shimizu, and S. Mori. Mechanism on the efficient formation of singlet oxygen by energy-transfer from excited singlet and triplet-states of aromatic hydrocarbons. *Bulletin of the Chemical Society of Japan*, 65(3):897–902, 1992.
- [105] C. Schweitzer and R. Schmidt. Physical mechanisms of generation and deactivation of singlet oxygen. *Chemical Reviews*, 103(5):1685–1757, 2003.
- [106] C. S. Burton and W. A. Noyes. Electronic energy relaxation in toluene vapor. *Journal of Chemical Physics*, 49(4):1705–&, 1968.

- [107] M. Jacon, C. Lardeux, R. Lopezdelgado, and A. Tramer. 3rd decay channel and vibrational redistribution problems in benzene-derivatives. *Chemical Physics*, 24(2):145–157, 1977.
- [108] F. Ossler and M. Alden. Measurements of picosecond laser induced fluorescence from gas phase 3-pentanone and acetone: Implications to combustion diagnostics. *Applied Physics B-Lasers and Optics*, 64(4):493–502, 1997.
- [109] J. D. Koch, R. K. Hanson, W. Koban, and C. Schulz. Rayleigh-calibrated fluorescence quantum yield measurements of acetone and 3-pentanone. *Applied Optics*, 43(31):5901–5910, 2004.
- [110] A. Braeuer, F. Beyrau, and A. Leipertz. Laser-induced fluorescence of ketones at elevated temperatures for pressures up to 20 bars by using a 248 nm excitation laser wavelength: experiments and model improvements. *Applied Optics*, 45(20):4982–4989, 2006.
- [111] V. Modica, C. Morin, and P. Guibert. 3-Pentanone LIF at elevated temperatures and pressures: measurements and modeling. *Applied Physics B-Lasers and Optics*, 87(1):193–204, 2007.
- [112] J. D. Koch, J. Gronki, and R. K. Hanson. Measurements of near-UV absorption spectra of acetone and 3-pentanone at high temperatures. *Journal of Quantitative Spectroscopy & Radiative Transfer*, 109(11):2037–2044, 2008.
- [113] B. R. Petersen, J. B. Ghandhi, and J. D. Koch. Fluorescence saturation measurements of 3-pentanone. *Applied Physics B-Lasers and Optics*, 93(2-3):639–644, 2008.
- [114] W. M. Nau and J. C. Scaiano. Oxygen quenching of excited aliphatic ketones and diketones. *Journal of Physical Chemistry*, 100(27):11360–11367, 1996.

- [115] G. de Sercey, G. Awcock, and M. Heikal. Development of a calibrated technique for in situ investigation of air-fuel mixing in engines. *Imaging Science Journal*, 54(1):33–45, 2006.
- [116] Malcolm Sandford, Graham Page, and Paul Crawford. The All New AJV8. *Society of Automotive Engineers*, 2009-01-1060, 2009.
- [117] M. Davy, P. Williams, D. Han, and R. Steeper. Evaporation characteristics of the 3-pentanone-isooctane binary system. *Experiments in Fluids*, 35(1):92–99, 2003.
- [118] V. Sick and B. Stojkovic. Attenuation effects on imaging diagnostics of hollow-cone sprays. *Applied Optics*, 40(15):2435–2442, 2001.
- [119] R. E. Stevens, H. Ma, C. R. Stone, H. L. Walmsley, and R. Cracknell. On planar laser-induced fluorescence with multi-component fuel and tracer design for quantitative determination of fuel concentration in internal combustion engines. *Proceedings of the Institution of Mechanical Engineers Part D-Journal of Automobile Engineering*, 221(D6):713–724, 2007.
- [120] Z. van Romunde and P. G. Aleiferis. Effect of Fuel Properties on Spray Development from a Multi-Hole DISI Engine Injector. *Society of Automotive Engineers*, 2007-01-4032, 2007.
- [121] G. J. Ernst. Uniform-field electrodes with minimum width. *Optics Communications*, 49(4):275–277, 1984.
- [122] S Pfadler, F Beyrau, M Loffler, and A Leipertz. Application of a beam homogenizer to planar laser diagnostics. *Optics Express*, 14(22):10171–10180, 2006.
- [123] C. Schulz, J. B. Jeffries, D. F. Davidson, J. D. Koch, J. Wolfrum, and R. K. Hanson. Impact of UV absorption by CO₂ and H₂O on NO LIF in high-pressure

- combustion applications. *Proceedings of the Combustion Institute*, 29:2735–2742, 2003.
- [124] J. B. Jeffries, C. Schulz, D. W. Mattison, M. A. Oehlschlaeger, W. G. Bessler, T. Lee, D. F. Davidson, and R. K. Hanson. UV absorption of CO₂ for temperature diagnostics of hydrocarbon combustion applications. In *30th International Symposium on Combustion*, pages 1591–1599, Chicago, IL, 2004. Combustion Inst.
- [125] Xiaowei Wang. *Instantaneous In-cylinder Heat Transfer and Combustion Analysis in Spark Ignition Engines*. PhD thesis, 2008.
- [126] R. Devillers, G. Bruneaux, and C. Schulz. Investigation of toluene LIF at high pressure and high temperature in an optical engine. *Applied Physics B-Lasers and Optics*, 96(4):735–739, 2009.
- [127] W. S. Rasband. ImageJ, 1997-2009.
- [128] R. E. Stevens. *Laser-Induced Grating Techniques for Combustion Diagnostics*. PhD thesis, 2004.

A. QPLIF supplement

A.1. QPLIF preparation procedure

QPLIF experiments proceed by preparing the engine, laser, intensified camera, line camera, data acquisition systems, ETCS, fuel and tracer mixing etc.

Engine preparation includes cleaning all optical surfaces which are prone to contamination, such as the piston crown window (pcw), pent-roof window (prw), Bowditch 45° mirror and fused silica annulus. Once all items are clean, the annulus can be fitted onto the barrel and raised into position using compressed air from the laboratory feed. The pressure should be increased slowly and controllably, such that the annulus and pent-roof window do not experience any sudden shocks. With the barrel raised to the cylinder head, the annulus should be verified to be central in the locating plate. The gasket on the upper surface of the annulus should also be verified to still be concentrically located. The engine can now be manually rotated by use of a suitable bar or ratchet on the rear driveshaft from the dynamometer. A few full cycles of rotation listening for unusual noises or unexpected friction is wise. Once the normal rotation of the engine is ascertained, the barrel may be raised to full pressure by switching to bottled nitrogen regulated to 20 bar. Should the pneumatic pressure holding the barrel in place fall below this value during a test, a deadlock will be tripped, stopping the engine. Fuel pressure is set by adjusting the regulator on the fuel pump. The high pressure transducer on the fuel line should read 150 bar.

Line camera preparation is a matter of starting the micropump which circulates

the dye solution, and confirming the collection of laser beam profiles which are below ~ 1500 counts in peak magnitude so as not to saturate the detector response.

Intensified camera preparation is begun by allowing a slow flow of dry nitrogen to purge the detector head. This is important to prevent condensation of water vapour in the air onto the delicate components inside. Once the nitrogen is flowing (confirm by, for example, testing the flow by feeling for a cooling sensation on the lips when the tube is presented), the cooling water flow may be started. This flow should also be relatively slow, to allow heat to be carried from the thermoelectric cooler. The next step enables the Peltier cooling element (turn on the camera controller, ST-138), which over the following 20-30 minutes will reduce the temperature of the CCD to approximately minus 40-50°C. The PG200 gate generator is then turned on. The camera may be operated in two modes - 'gated', where the intensifier is only operative upon the application of an external high voltage to the BNC on the camera - or 'shutter', where the ICCD camera is itself in control of the gating process (and may be controlled by software). The use of an external gate unit allows a gate down to ~ 100 ns to be employed. It is good practice to leave the camera in 'gated' mode, where the intensifier is normally disarmed. Extreme care must be taken when increasing the gain control on the camera, that the chip is not saturated. The ICCD camera will emit a warning tone when saturation is detected, which is a sign to the user to pay more attention! When reading out the entire 576×384 array, the maximum frame rate is approximately 1.5Hz. A background signal level of 350 counts is typical, and related to the setting of the ADC converter stage in the ST138 controller. This does make for a slight reduction in dynamic range, but also allows any drift in baseline counts to not go unnoticed. If there has been a change of wavelength, the Schott glass filter on the ICCD should be changed to account for the altered laser frequency.

Data acquisition systems should always be checked with a short test before taking experimental data, to confirm that all analogue channels are recording 'sensible' data values. Equally, the operation of the ETCS computer is verified at this point, e.g. by

switching between firing modes and triggering the DAQ systems.

Fuel and tracer preparation proceeds by measuring out $\sim 500\text{ml}$ of Shell Component Fuel. The requisite amount of tracer is added to this fuel. (2% for toluene and TMB, 5% for acetone.) In the case of ‘bookend’ calibration being in effect, half of this fuel is poured into the plenum injection tank, and half into the DI tank. If there has been a change of tracer, both tanks must be emptied and the fuel lines purged by running in fired mode on pure iso-octane.

Laser preparation includes adjusting the flipper mirrors to select either 248nm or 277nm operation. The normal propagation of the laser beam through the first set of optical cabinets should be checked before moving into the engine room. Adjustments are now made to bring the laser sheet into the dye cell for line-camera monitoring. A small crucible of acetone should be introduced into the engine, and the barrel closed, with the valves shut. With the experiment frequency controlled by the first TTPG box of figure 4.10, the image on the intensified CCD camera is used to check the laser sheet is positioned optimally and of the correct width inside the engine. Care must also be taken at this stage that the laser sheet continues to terminate upon the aluminium solder stripe on the cylinder head. The energy read-out on the excimer laser mini-controller should be checked, to confirm that the laser is producing an acceptable pulse energy ($\gtrsim 100\text{mJ}$). The pulse energy should be expected to decay during the test, so a marginal result implies the need to replace the KrF gas mixture. The mini-controller has a function to perform this operation, but the various gas valves and regulators must be operated manually. Should any doors be opened or the extraction system be disarmed during a test, the laser interlock will immediately trip.

In the test schedule illustrated in figure A.1, these steps[†] are summarised. The page is dominated by a table which describes the sequence of tests to be performed. Each row describes another test condition, and each test condition consists of a motored

[†]and a few more - like checking the plenum heater tape is plugged in - it is impractical to write an exhaustive test schedule, since no two tests are 100% identical.

‘BG’ test (no injection or ignition, only laser light) and a fired ‘SG’ test. The images recorded during the ‘BG’ test are used as background reference images for the ‘SG’ tests. As each set of images is collected, the corresponding box in the table ought to be ticked, allowing progress to be monitored.

Note also the presence of the beam-fan reference tests (here referred to as ‘linparms’), the application of which is described in more detail in section 4.5. In practical terms, two sets of reference images are collected. Firstly, a background with no laser beam, and secondly a test where the laser beam is masked to produce two ‘stripes’.

The engine is slowly increased in speed with the oil pressure interlock overridden - whilst listening for unusual noises - to approximately 700rpm, when the button can be released. The engine should be increased in steps to its final operating speed (1500 rpm for WWMP), allowing time for the piston rings to settle, the airflow to stabilise and the coolant to come up to temperature and so on.

For each firing run, the engine should be allowed to stabilise before taking measurements. This may take 30-60 seconds of firing, during which time the engine surfaces will increase in temperature. During a PLIF run, the ETCS system is unable to stop the ignition or injection, so the user should be aware of a method to disable one or both, should the torque readout become very negative, or any temperatures go too high. ‘Normal’ motoring torque is in the region of -25Nm, and ‘normal’ firing torque around -15Nm. A torque higher than -30Nm whilst firing can rapidly become highly concerning. Each run lasts for ~ 320 images, with a skipped laser trigger recording one frame for every 9 engine cycles. This corresponds to $\sim 3,000$ cycles or about two minutes. Including, e.g. warm-up time, the engine probably runs for about 3-4 minutes of continuous fired operation for any given set of images.

A full QPLIF experimental run as described here is likely to take in the region of 3-4 hours when performed by a practiced user. The experiment will produce in the region of 2GB of data spread over three different computers, which should be collated, ready for analysis.

A.2. QPLIF schedule

An example schedule from a QPLIF experiment is given in figure A.1. This particular experiment consists of the two bookend tests to calibrate the dataset, two tests on early-DI conditions and a further test with late-DI. Note that tests are kept to the same overall length by reducing the number of snaps taken at each crank angle position. This is done to avoid running the optical engine in fired mode for an extended period of time. Excessive heating is liable to cause the ring material to melt, which can result in anything from mild panic to weeks of downtime for repair!

System checks

- Line camera alignment – is the shape on a flat background, and maximised?
- Is the laser sheet filling the field of view?
- Is there a beam block above the 2nd cylindrical lens?
- Cleaning checklist:
 - Piston crown window
 - Pent-roof & window
 - Bowditch mirror
- Laser energy sufficiently high?
- Fuel mixed and ready to go?
 - If changing tracer, remember to purge with iso-octane.
- Low ICCD counts during BG1? Should be ~350, especially outside the prw.
- **SET H/C and O/C ratios in the lambda sensor (ECM AFRMeter)**

PLIF test sequence – trial runs

| TYPE | CAD range | TRIGS | SNAPS | BG | SG |
|------------|------------|-------|-------|----|----|
| Plenum inj | 90:10:0 | 320 | 32 | | |
| DI early 1 | 280:10:160 | 338 | 26 | | |
| DI early 2 | 150:10:30 | 338 | 26 | | |
| DI late | 140:10:30 | 336 | 28 | | |
| Plenum inj | 90:10:0 | 320 | 32 | | |

- Set engine to an air flow rate of 1.21Ls^{-1}
- Exhaust gas analyser armed
- Cooled barrel
- Heated plenum to keep fuel in vapour state

Post-experiment sequence

- background: LC1 & ICCD background with no beam present (but lights on)
- linparms: mask static engine, acetone vapour

Figure A.1.: QPLIF experimental schedule

B. Semi-QPLIF measurements

In this appendix, some extra results from the application of PLIF to the Jaguar optical engine are presented. The results in this section are either less quantitative in nature than those presented in the main text, or simply deemed surplus to the main theme.

B.1. Shot-wise distributions

Figure B.1 shows a representative set of four single-shot fluorescence images recorded with the tracer acetone added to the component fuel, at 40 BTDC - just prior to ignition. The images have been corrected for the removal of background signal sources, but have not been converted into FFC by the analysis system. They are presented on the same scale, however, so comparison between the four images is valid.

The four images of figure B.1 highlight the variability that is present from shot-to-shot in the late-DI WWMP operating point. These images also illustrate typical single-shot noise levels after the application of a 3×3 median filter to the raw images. This low signal level is in some sense the Achilles' heel of QPLIF, but it is an unavoidable side-effect of desiring quantitative measurements. If either absorption or laser energy are too high, then quantitative analysis will become highly challenging, if not impossible[†].

[†]i.e. higher signal levels are not easily achieved.




Figure B.1.: Shot-wise qualitative measurements of light tracer fluorescence

B.2. IVO PLIF

The valve event timings given in table 4.2 indicate that the IVC event - intake valve closure - occurs at 274 ATDC - i.e. 86 BTDC. This means that as the piston is travelling up the bore of the cylinder during the compression, the intake valve is open until about half way. This means that some of the contents of the cylinder are bound to get pushed back out and into the intake runner. The images presented in figure B.2 are average fuel distributions measured around the time of IVO, which occurs at 24 ATDC (336 BTDC). These images *have* been processed by the analysis system, and so are presented as quantitative fuel distributions.

These images are presented on the same scale as commonly used in this thesis, with the MATLAB JET colour map quantifying FFC in the range from 0 to 2. The tracer used was toluene. The sequence of six average images shows no immediate change at the time of intake valve opening. However, the profile of the valve is visible in the distributions measured at 310-280 BTDC, and a mixture of fuel and air is seen to be drawn into the engine. At 290 BTDC, the mixture just below the intake valve has an




Figure B.2.: Images of toluene fluorescence, showing fuel entering the cylinder at IVO. The number in the lower left of each image is the image timing, in CAD BTDC.

FFC of approximately 0.5. This mixture becomes diluted as the stroke proceeds.

Notably, this behaviour is largely absent when late-DI is used, because the fuel distribution is relatively lean under the intake valves until about 50 BTDC, as can be verified from figure 4.24. This measurement probably does not betray any unexpected behaviour - but does illustrate how QPLIF can be a useful tool, providing information throughout the engine cycle.

B.3. Imaging sprays

Section 4.7 included some discussion of the use of CoV for investigating the repeatability/cyclic variability of measured fuel distributions. Figure B.3 shows a matrix of fluorescence images, examining the behaviour of the fuel spray. These images have been analysed, but owing to the presence of liquid fuel the results cannot be relied upon, so no particular calibration is applied to the first two rows in the table. However, as a relative measure, the CoV is still perfectly valid, and is presented on the scale of zero to two, as given at the bottom of the figure.

The images in figure B.3 demonstrate that the edges of the spray plumes exhibit

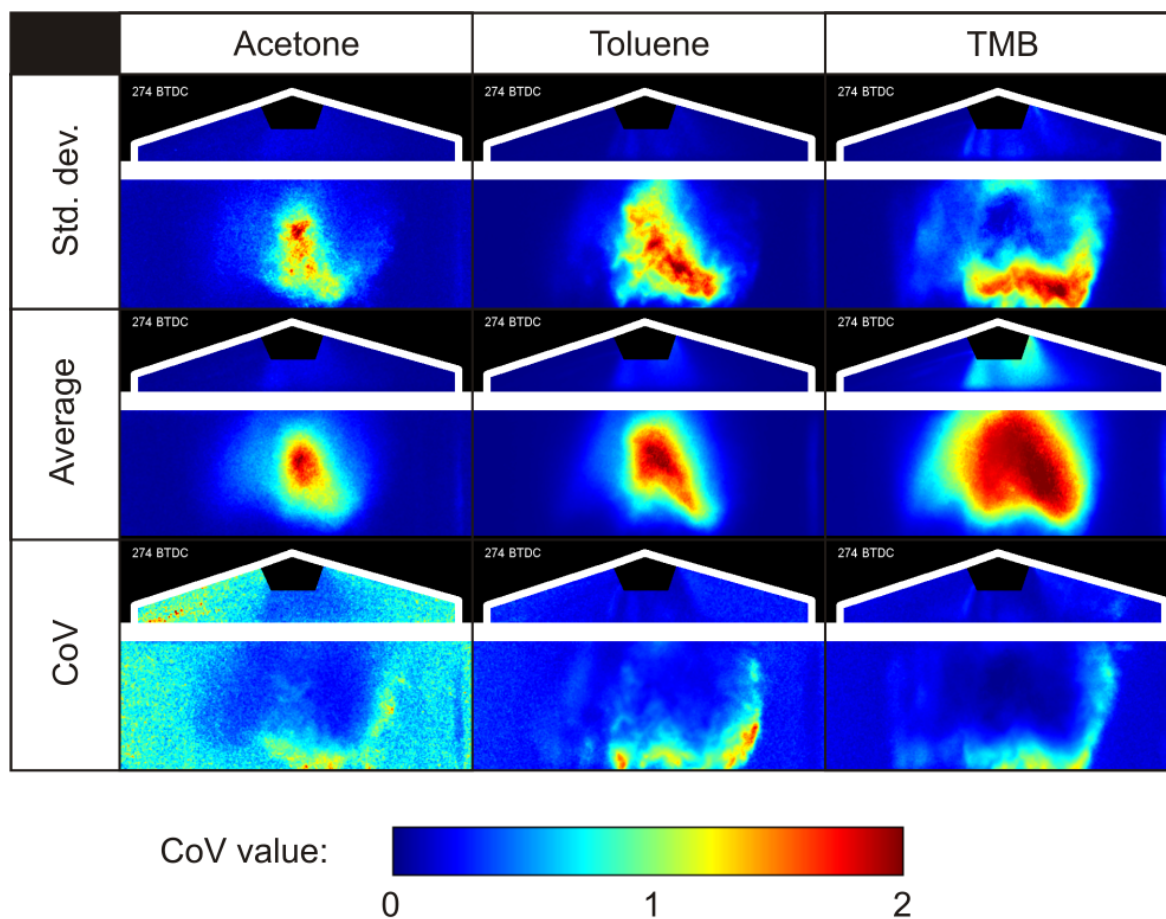


Figure B.3.: Characterisation of spray by consideration of CoV. Most variability is seen at the edges of the spray plumes.

the greatest cyclic variability, with anything up to and exceeding 100% variation in the cycle-to-cycle quantity of fuel delivered. Note firstly the unexpected dip in the TMB standard deviation image, which is the result of mild ICCD saturation which clipped the distribution reducing variability in this region. Note secondly that the acetone CoV image is different to the others. This is because the quantum efficiency of acetone is markedly lower at these low pressures than the aromatic tracers, and so the intrinsic camera noise becomes more of a problem, introducing a uniform variability to the image. Nevertheless, the corona of higher CoV values is still present, though less well defined. Finally, note that the CoV is relatively low at the tops of the jets where the fuel delivery is very consistent - as would be expected.

C. FARLIG supplement

C.1. Gas cell refill procedure

The gas cell is refilled by adopting the following procedure, where cell window heating is conducted by using a source of hot air directed at the windows, alternating between the two ends of the cell every 30 seconds.

1. Slowly evacuate the cell. Monitor thermocouple reading to ensure temperature does not drop more than $\sim 2^{\circ}\text{C}$. Whilst the cell contents are evacuated, the cell windows should be kept heated to prevent condensation.
2. Following the experimental schedule (see section C.3), fill the cell to the specified partial pressure of oxygen (or maximum 1 bar absolute).
3. Top up cell pressure to 1 bar absolute with nitrogen (if oxygen partial pressure < 1 bar abs.).
4. Keeping the ball valve closed, insert the needle into the PTFE disc
5. Open the ball valve, and guide the needle into the cell
6. Allow one drop of toluene to fall from the end of the needle, by depressing the plunger on the syringe
7. Withdraw the needle, applying gentle pressure to the ball valve mechanism. The resistance will reduce as the needle clears the ball valve, at which point the valve can be closed as the needle is fully withdrawn

8. Keep the windows heated, and allow a minute or two for the toluene to vaporise (the small ‘pool’ of toluene can be seen by eye to evaporate)
9. Proceed slowly to test pressure by adding nitrogen or oxygen, heating windows. Monitor thermocouple temperature to ensure no large deviations in temperature are caused by the addition of gas.
10. If the cell is being heated to a temperature, ensure thermal equilibrium is reached.
11. Acquire FARLIG signals

Keeping the cell windows clear of condensed toluene is crucial to maintaining the quality of the FARLIG signal - the LITGS signal will rapidly degrade if toluene is allowed to condense out.

C.2. Alignment technique

Figure 5.9 illustrates the sequence of steps undertaken to align the FARLIG experiment. The steps are described as follows:

1. Firstly, centre the aperture immediately following the polarisers and KD*P on the fourth harmonic beam, and then use this iris to limit the beam diameter to $\sim 1.5\text{mm}$.
2. Now align the two fourth harmonic beams through the masks, using the adjustments on the beamsplitter and mirrors in the horizontal plane. The mask heights may be adjusted. (Figure 5.9 **a**.)
3. The HeNe alignment beam and the FLPDL beam are also aligned through the masks at this time. (Figure 5.9 **a**.)
4. The front mask is now removed, and the crossing lens placed into its holder. (Figure 5.9 **b**.)

5. The transverse adjustment screws on the crossing lens are now used to place the fourth harmonic beams on the rear mask correctly. (Figure 5.9 **b.**)
6. The pinhole is then introduced about the crossing point. The pinhole is $\sim 100\mu\text{m}$ in diameter, and may be accurately placed with a system of translation stages. (Figure 5.9 **c.**)
7. The HeNe and FLPDL are now adjusted to travel through the pinhole and out onto the correct hole on the rear mask. This step corrects for any chromatic aberration in the crossing lens. (Figure 5.9 **d.**)
8. The pinhole is now removed, and the focussing lenses which draw the HeNe and FLPDL beams to focus in the rear focal plane of the crossing lens are now put into place, and adjusted such that the beams lie once more in the correct location on the rear mask.
9. The rear mask is now removed, and the HeNe alignment beam is used to adjust the mirrors in the detector leg to carry the HeNe beam (acting as a signal beam tracer) to the detectors which will be used.
10. Finally, the aperture in the detector leg is put in place, to block any stray light from reaching the detectors. The aperture used to narrow the width of the fourth harmonic is also fully opened, prior to installing the gas cell and beginning experiments.

C.3. Schedule

The schedule for a FARLIG experimental sequence is included in this appendix, as figure C.1.

This figure shows the document which is used to plan an experiment. Some important factors to consider during the experiment are noted at the top of the sheet,

followed by a matrix of test conditions to be evaluated. A brief experimental procedure is then outlined. The end of the page gives a space for any extra notes, e.g. if the experiment also includes measurements at different temperatures, it is noted here.

A FARLIG experiment proceeds by turning on all the apparatus and verifying normal operation, before the alignment procedure which is explained in sections 5.4.5 and C.2, and then the various test conditions as described in figure C.1.

Also included in this section is the start-up sequence for the flashlamp-pumped dye laser (FLPDL), as figure C.2. Note that the chiller and nitrogen purge are not strictly required; the laser may be operated with just cooling water flowing through the laser, which is itself cooled at the heat exchanger. (Cooled dye at 4°C should result in a more stable output beam, however.)

LIGS Experiments

THINGS TO BEAR IN MIND

- **Keep cell windows heated to avoid condensation**
- **Add and remove gas very slowly! Particularly avoid lowering T**
- Approx. 200 shots before scattering becomes an issue.
- Variac at 50V – cell temperature ~27°C

EXPERIMENTAL CONDITIONS TO TRY

| O ₂ pressure | 1 bar total | 2 bar total | 4 bar total | 8 bar total |
|-------------------------|-------------|-------------|-------------|-------------|
| 0mb | | | | |
| 100mb | | | | |
| 200mb | | | | |
| 400mb | | | | |
| 600mb | | | | |
| 800mb | | | | |
| 1000mb | | | | |
| 2000mb | | | | |
| 4000mb | | | | |
| AIR | | | | |

- Clean and fit windows
- Vacuum out cell
- Heat windows, 2 mins each end
- Fill to specified partial pressure of oxygen (or 1 bar max)
- Top up with nitrogen to 1 bar if required
- Insert needle and admit 1 drop of toluene
- Keep windows heated, 1 min each end
- Allow 5 mins for toluene to evaporate (look at pool disappearing)

- Continue to test pressure
- Acquire signals
- Tick table, include temperature reading
- Heat windows
- Advance to next pressure
- Acquire signals
- Tick table, include temperature reading
- Heat windows
- ...

Once O₂ tests are complete, turn variac up and use a 4 bar mix to get signals at increasing temperatures.

Notes here:

PMT bias voltage:
Experiment date:

(use table if variable)

Figure C.1.: FARLIG experimental schedule

Dye laser start up sequence

=====

Preflight A:

1. Check the chiller reservoir
2. Start chiller and nitrogen purge

Leave system for ~1hour to get down to temperature

Preflight B:

1. Turn on the water for the heat exchanger (dark blue tap, low-ish flow)
2. Turn on the water pump (passes deionised water around lamps and heat exchanger)
3. Turn on the laser electronics
Interlock lamp comes on while the thyatron warms up
4. Turn on dye circulator pump

[check thyatron heater is at 6.3V every now and then]

When the interlock lamp goes out, the system is ready to go.

Firing:

1. Set trigger to manual
2. Wind HV up to ~12kV (any higher may cause damage to lamps)
3. Simmer current to maximum
4. Set to bypass flashlamp check
5. Tap PTM to arc the flashlamps, and quickly reduce simmer current to ~ 2A
6. Increase HV to get a nice output
Play around with iris etc, to get TEM00 operation
7. Triggering can be performed using the internal trigger, which is set with the dial at the top left of the control box.

Turning off:

1. Wind down HV and simmer
2. Turn off control units, pumps and power supplies.

Figure C.2.: FLPDL startup and shutdown procedures

D. FARLIG fitdata

The entire listing of fitdata is given in this appendix as figures D.1 and D.2.

| date | partial o2 | %O2 | totalp | Re | std(Re) | tau | std(tau) | Q1 | std(Q1) |
|------|------------|-----------|--------|--------|---------|----------|----------|----------|----------|
| 9 | 0 | 0 | 2 | 248.47 | 15.85 | 1.80E-08 | 4.28E-11 | 2.71E-09 | 4.22E-10 |
| 9 | 400 | 20 | 2 | 213.15 | 23.16 | 1.81E-08 | 4.04E-11 | 2.25E-09 | 4.38E-10 |
| 9 | 420 | 21 | 2 | 202.89 | 18.098 | 1.82E-08 | 4.60E-11 | 2.49E-09 | 3.10E-10 |
| 9 | 1000 | 50 | 2 | 198.41 | 17.768 | 1.86E-08 | 4.58E-11 | 2.56E-09 | 4.56E-10 |
| 9 | 0 | 0 | 4 | 474.89 | 49.284 | 1.80E-08 | 3.54E-11 | 2.47E-09 | 2.56E-10 |
| 9 | 400 | 10 | 4 | 443.36 | 37.61 | 1.80E-08 | 3.92E-11 | 2.02E-09 | 2.97E-10 |
| 9 | 800 | 20 | 4 | 415.64 | 29.654 | 1.82E-08 | 3.21E-11 | 2.35E-09 | 4.11E-10 |
| 9 | 840 | 21 | 4 | 412.17 | 22.992 | 1.82E-08 | 3.35E-11 | 2.41E-09 | 4.05E-10 |
| 9 | 1000 | 25 | 4 | 413.54 | 36.454 | 1.83E-08 | 3.46E-11 | 2.32E-09 | 4.23E-10 |
| 9 | 0 | 0 | 8 | 1046.9 | 150.59 | 1.79E-08 | 2.87E-11 | 1.89E-09 | 2.10E-10 |
| 9 | 400 | 5 | 8 | 795.49 | 87.112 | 1.80E-08 | 3.41E-11 | 1.67E-09 | 3.32E-10 |
| 9 | 800 | 10 | 8 | 760.09 | 90.996 | 1.80E-08 | 3.28E-11 | 1.66E-09 | 3.65E-10 |
| 9 | 1000 | 12.5 | 8 | 749.48 | 72.873 | 1.81E-08 | 3.43E-11 | 1.69E-09 | 3.04E-10 |
| 9 | 1680 | 21 | 8 | 734.71 | 72.625 | 1.82E-08 | 3.54E-11 | 1.47E-09 | 3.10E-10 |
| 25 | 0 | 0 | 1 | 112.88 | 4.8796 | 1.76E-08 | 1.13E-10 | 4.65E-09 | 9.08E-10 |
| 25 | 200 | 20 | 1 | 105.65 | 20.212 | 1.78E-08 | 7.79E-11 | 2.06E-09 | 1.40E-09 |
| 25 | 400 | 40 | 1 | 96.473 | 16.459 | 1.80E-08 | 7.85E-11 | 3.89E-09 | 2.57E-09 |
| 25 | 600 | 60 | 1 | 111.43 | 13.912 | 1.84E-08 | 1.75E-10 | 5.35E-09 | 2.91E-09 |
| 25 | 800 | 80 | 1 | 95.359 | 7.8846 | 1.85E-08 | 1.31E-10 | 4.70E-09 | 1.27E-09 |
| 25 | 1000 | 100 | 1 | 96.81 | 26.133 | 1.87E-08 | 2.53E-10 | 4.15E-09 | 1.11E-09 |
| 25 | 0 | 0 | 2 | 217.25 | 19.344 | 1.76E-08 | 5.28E-11 | 5.51E-09 | 4.72E-10 |
| 25 | 200 | 10 | 2 | 192.04 | 21.456 | 1.76E-08 | 5.35E-11 | 4.08E-09 | 1.13E-09 |
| 25 | 400 | 20 | 2 | 183.4 | 13.827 | 1.78E-08 | 1.10E-10 | 4.39E-09 | 1.09E-09 |
| 25 | 1000 | 50 | 2 | 175.03 | 13.659 | 1.82E-08 | 9.77E-11 | 4.07E-09 | 9.63E-10 |
| 25 | 0 | 0 | 4 | 432.67 | 36.663 | 1.76E-08 | 4.60E-11 | 4.02E-09 | 6.89E-10 |
| 25 | 200 | 5 | 4 | 347.67 | 38 | 1.77E-08 | 4.20E-11 | 3.44E-09 | 7.99E-10 |
| 25 | 400 | 10 | 4 | 368.36 | 45.09 | 1.77E-08 | 4.54E-11 | 3.54E-09 | 6.92E-10 |
| 25 | 600 | 15 | 4 | 374.5 | 37.556 | 1.78E-08 | 3.92E-11 | 3.40E-09 | 5.77E-10 |
| 25 | 800 | 20 | 4 | 367.5 | 32.367 | 1.79E-08 | 3.68E-11 | 3.31E-09 | 6.15E-10 |
| 25 | 1000 | 25 | 4 | 381.12 | 30.727 | 1.79E-08 | 3.89E-11 | 2.99E-09 | 6.64E-10 |
| 27 | 210 | 21 | 1 | 92.82 | 5.0264 | 1.82E-08 | 4.85E-11 | 1.10E-09 | 2.86E-10 |

Figure D.1.: Full table of fitdata parameters, part 1. (Left hand side of full table.)

| split | std(split) | Q2 | std(Q2) | Err | std(Err) | numshots | contrast | std(contrast) |
|---------|------------|----------|----------|--------|----------|----------|----------|---------------|
| 0.0434 | 0.00806 | 2.88E-08 | 7.77E-09 | 1.5553 | 0.2299 | 36 | 0.725139 | 0.018168142 |
| 0.05138 | 0.00985 | 3.03E-08 | 8.65E-09 | 1.4031 | 0.2229 | 21 | 0.726906 | 0.016367572 |
| 0.05378 | 0.00636 | 3.19E-08 | 8.44E-09 | 1.4357 | 0.2171 | 19 | 0.719689 | 0.011783391 |
| 0.04159 | 0.00451 | 3.42E-08 | 7.12E-09 | 1.2966 | 0.1923 | 24 | 0.751827 | 0.015820622 |
| 0.00909 | 0.00132 | 7.79E-08 | 3.62E-08 | 1.6554 | 0.1705 | 13 | 0.79351 | 0.010630223 |
| 0.01591 | 0.00321 | 4.77E-08 | 1.67E-08 | 1.4453 | 0.2348 | 35 | 0.769973 | 0.131306152 |
| 0.01294 | 0.00221 | 5.56E-08 | 1.81E-08 | 1.3806 | 0.2392 | 41 | 0.796853 | 0.017479981 |
| 0.01259 | 0.00241 | 5.36E-08 | 1.36E-08 | 1.4617 | 0.2644 | 16 | 0.797154 | 0.014498006 |
| 0.01163 | 0.0014 | 6.46E-08 | 1.76E-08 | 1.4189 | 0.2204 | 24 | 0.803781 | 0.015614987 |
| 0.00122 | 0.00094 | 2.44E-07 | 1.29E-07 | 1.8439 | 0.1873 | 23 | 0.837493 | 0.010574316 |
| 0.00634 | 0.00145 | 9.29E-08 | 3.76E-08 | 1.6451 | 0.1838 | 40 | 0.829024 | 0.011927021 |
| 0.0046 | 0.00102 | 1.21E-07 | 5.32E-08 | 1.548 | 0.2087 | 33 | 0.837538 | 0.011712756 |
| 0.00354 | 0.00075 | 1.53E-07 | 6.11E-08 | 1.5844 | 0.1818 | 26 | Inf | NaN |
| 0.00287 | 0.0007 | 1.51E-07 | 6.35E-08 | 1.6019 | 0.2055 | 35 | 0.859733 | 0.008051643 |
| 0.26158 | 0.01244 | 2.93E-08 | 1.85E-09 | 1.4944 | 0.304 | 7 | 0.406219 | 0.025029928 |
| 0.2026 | 0.09762 | 2.12E-08 | 6.47E-09 | 1.4661 | 0.241 | 8 | 0.445578 | 0.014688881 |
| 0.21024 | 0.09873 | 9.12E-08 | 1.80E-07 | 1.5554 | 0.2487 | 7 | 0.476698 | 0.03185023 |
| 0.14322 | 0.0995 | 9.56E-08 | 1.78E-07 | 1.4824 | 0.3187 | 7 | 0.491049 | 0.035565665 |
| 0.18795 | 0.04066 | 2.67E-08 | 4.75E-09 | 1.9479 | 0.5379 | 6 | 0.518871 | 0.020837015 |
| 0.18104 | 0.09042 | 8.35E-08 | 1.62E-07 | 2.658 | 0.5736 | 8 | 0.550838 | 0.041828575 |
| 0.07378 | 0.0067 | 5.28E-08 | 9.42E-09 | 1.4468 | 0.2425 | 16 | 0.480895 | 0.019443624 |
| 0.09885 | 0.01569 | 3.54E-08 | 1.06E-08 | 1.5989 | 0.2391 | 8 | 0.522536 | 0.046382338 |
| 0.06236 | 0.01371 | 4.62E-08 | 9.56E-09 | 1.3959 | 0.3276 | 14 | 0.567629 | 0.043657588 |
| 0.04834 | 0.00655 | 5.16E-08 | 1.23E-08 | 1.5749 | 0.2104 | 9 | 0.631602 | 0.047396298 |
| 0.02274 | 0.00217 | 7.77E-08 | 1.49E-08 | 1.4551 | 0.2491 | 25 | 0.619085 | 0.035587304 |
| 0.02979 | 0.00361 | 6.85E-08 | 1.33E-08 | 1.5157 | 0.2643 | 35 | 0.632059 | 0.036257128 |
| 0.0196 | 0.00506 | 7.45E-08 | 1.58E-08 | 1.2567 | 0.2612 | 41 | 0.655721 | 0.035808298 |
| 0.01501 | 0.0041 | 7.82E-08 | 1.71E-08 | 1.1871 | 0.2642 | 20 | 0.69019 | 0.03667194 |
| 0.01435 | 0.00306 | 8.32E-08 | 1.69E-08 | 1.2124 | 0.2026 | 19 | 0.699867 | 0.037956708 |
| 0.01093 | 0.00217 | 8.55E-08 | 3.03E-08 | 1.4691 | 0.2775 | 13 | 0.729227 | 0.031137099 |
| 0.06822 | 0.01575 | 3.02E-08 | 1.86E-09 | 1.2963 | 0.2418 | 30 | 0.639759 | 0.021803444 |

Figure D.2.: Full table of fitdata parameters, part 2. (Right hand side of full table.)