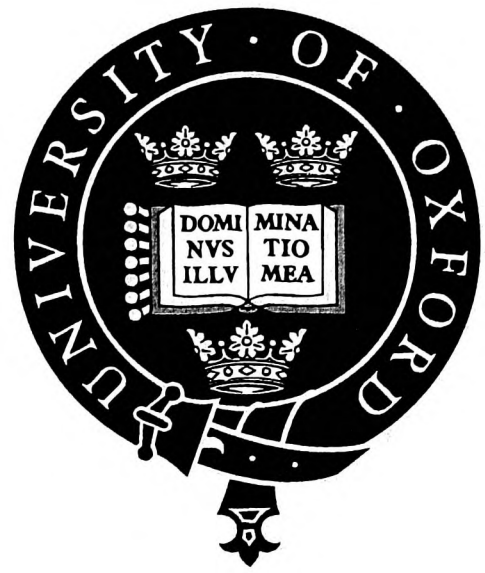


Optical Diagnostics and Combustion Analysis in a Gasoline Direct Injection Engine

Hongrui Ma

Somerville College

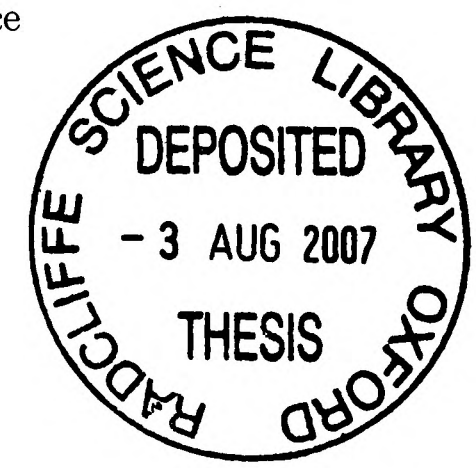
University of Oxford

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Department of Engineering Science

University of Oxford

Michaelmas Term 2006



Optical Diagnostics and Combustion Analysis in a Gasoline Direct Injection Engine

Hongrui Ma
Somerville College, Oxford

Doctor of Philosophy
Department of Engineering Science

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford, Michaelmas Term 2006

Abstract

Gasoline Direct Injection (GDI) engines work with stratified charge at part load and burn with lean mixtures in order to save fuel, whilst at full load, the fuel and air mix homogeneously for maximum power output. The higher compression ratio and the absence of throttling are two of the most significant benefits of GDI engines. The key issues facing GDI combustion include in-cylinder mixture preparation and post-combustion soot formation. This work was intended to investigate these aspects and was undertaken on a dedicated Jaguar single-cylinder optical GDI engine with a spray-guided combustion system. The spray-guided concept does not rely as much on charge motion or piston design, and can avoid wall-wetting effects so as to reduce engine emissions. Relevant engine control hardware and data acquisition equipment were commissioned. Data/image processing software was also developed to suit the measurements.

A data-processing case study with data from a small two-stroke glow ignition engine has been conducted to develop a method to combine the burn rate and heat release analyses in the study of engines with premixed charge but compression ignition. Difficulties such as unknown ignition timing and polytropic index have been addressed. Results for all operating conditions have shown good correlations between the two methods.

The technique of quantitative planar laser-induced fluorescence is useful for measuring 2-D fuel distribution in GDI engines. The relevant physics and literature were reviewed in depth. A multi-component fuel was designed to give reasonable co-evaporation characteristics with tracers matching different fuel fractions. The absorption and fluorescence features of each fuel component and tracer were characterised. Optimisation of hardware and signal-to-noise ratio was performed. A recirculating loop was set up for the calibration of the technique.

The technique of colour-ratio pyrometry (CRP) for estimating the temperature and loading of soot was applied on the GDI engine. Critical features of the candidate CCD colour camera including its spectral response and noise behaviours were fully studied. Validation tests with reference sources together with an error analysis suggested an accuracy of ± 50K within the combustion temperature range. Engine combustion images were then taken under various operating conditions. Temperature estimates were shown to be insensitive to the concentration of soot. Simulation with a thermodynamic modelling package, ISIS, was introduced for comparison with the experimental data. With careful tuning, ISIS gave outputs comparable to the CRP and proved to be a cost-effective tool to study GDI engines. High-speed combustion imaging was carried out using a CMOS camera, allowing the study of flame properties as well as crank-angle resolved CRP. By using a lens in the piston crown to give full bore optical access and appropriate image processing, the flame front could be detected reliably throughout the main combustion process.

Acknowledgements

Great gratitude must be drawn to my family, had it not been for their endless support, I would not have made the progress this far.

I would like to show special thanks to my supervisor, Prof. Richard Stone, for his patient supervision, cheerful encouragement and generousness in sharing his fantastic knowledge. Not only does he constantly point me to the right directions in academic research, he has also impacted a great deal on my personal life, so that as an overseas student I never feel any kind of unease or discomfort.

I would also like to thank my colleagues Dr Rob Stevens, Dr Kenneth Kar, Philip Price, Xiaowei Wang and Steve Marshall, for their warm introduction to the living, working and research circumstances, their cooperation, and their precious advice on improving my work. Prof. Paul Ewart and Prof. Robert Raine deserve as many thanks, for their kind help, inspiration and collaboration with my work.

I am grateful to the sponsors of this project, EPSRC, Jaguar Cars and Shell Global Solutions, as well as those academic colleagues in Oxford and other universities who have been kindly supportive.

Finally, thanks to all the dear friends that I have met during my stay in Oxford — you have made this otherwise-solitary journey one of the sweetest and most liberating experiences in my life.

Contents

List of Figures ix

List of Tables xiv

Nomenclature xv

Chapter 1 Background 1

1.1 Introduction 1

1.2 Concept of direct injection spark ignition 2

1.2.1 Developmental history 2

1.2.2 Recent applications 4

1.3 Advantages of the DISI concept 6

1.3.1 Fuel economy 6

1.3.2 Charge cooling 7

1.3.3 Fuelling control 8

1.4 Characteristics of DISI engines 8

1.4.1 Air intake and charge motion 8

1.4.2 Mixture preparation 10

1.4.3 Spray guiding concepts 11

1.4.4 Fuel injection technology 13

1.4.5 Combustion and emissions 16

1.4.6 Fuel efficiency strategies 17

1.5 Diagnostics and instrumentation for DISI research 17
1.5.1 Optical diagnostics 17
1.5.2 Combustion imaging and analysis 19
1.5.3 Laser diagnostics 20
1.5.4 Emissions measurement 22
1.6 Modelling DISI engines 24
1.7 Jaguar Optical DISI Engine 25
1.8 Summary 25

Chapter 2 Data Acquisition and Processing **27**

2.1 Introduction 27
2.2 Instrumentation system 28
2.3 Engine data analysis 31
2.3.1 Introduction of CoBRA 31
2.3.2 Pressure data analysis 32
2.3.3 Mass fraction burned (MFB) 34
2.3.4 Completeness of combustion 38
2.3.5 Heat release analysis 39
2.3.6 Implementing CoBRA in MATLAB 42
2.3.7 Graphical user interface of CoBRA 46
2.4 Image acquisition and processing 49
2.4.1 Charge coupled device (CCD) 49
2.4.2 PLIF imaging equipment 50
2.4.3 CCD combustion imaging equipment 52
2.4.4 High-speed combustion imaging equipment 55
2.4.5 Image processing 55
2.4.6 Image processing software 61
2.5 Summary 65

Chapter 3 Two-Stroke Engine Combustion Analysis 67

3.1 Introduction 67

3.2 Experimental aspects 68

3.3 Analysis issues 69

3.4 Results and discussion 72

3.4.1 Correlation between burn rate and heat release analyses 72

3.4.2 Effects of fuel blend and glow plugs 75

3.4.3 Cycle-by-cycle variations 79

3.5 Summary 82

Chapter 4 Development of Quantitative PLIF for AFR Measurement 84

4.1 Introduction 84

4.2 Principles of PLIF 85

4.3 Interpretation of PLIF images 93

4.4 Equipment and device 98

4.5 Design of component fuel 101

4.5.1 Selection criteria for fuel and tracers 101

4.5.2 Fuel/tracer design 103

4.5.3 Characterisation of fuel components and tracers 108

4.6 Optimisation of PLIF signal 115

4.6.1 Systematic noise correction 116

4.6.2 Optimization of signal-to-noise ratio 118

4.7 PLIF calibration 122

4.7.1 Calibration loop 122

4.7.2 PLIF vs. AFR calibration 124

4.8 Summary 127

Chapter 5 Soot Imaging and Pyrometry 129

5.1 Introduction 129

5.2 Principles of soot pyrometry 132
5.2.1 Radiation of glowing soot 132
5.2.2 Colour-ratio pyrometry 133
5.3 CCD camera calibration 137
5.3.1 Colour spectral response 137
5.3.2 Camera response study 143
5.3.3 Calibration image processing 148
5.4 Validation of CRP 150
5.4.1 Black body imaging 150
5.4.2 Tungsten filament imaging 152
5.4.3 Tungsten ribbon imaging 154
5.4.4 Concluding remarks on the temperature calibration 157
5.4.5 Error analysis 157
5.5 DISI engine results 160
5.6 Summary 166

Chapter 6 High-Speed Combustion Imaging **169**

6.1 Introduction 169
6.2 Experimental aspects 170
6.3 High-speed image sequence vs. burn rate 173
6.4 Flame growth and propagation 176
6.4.1 Flame front detection 178
6.4.2 Results with Perspex piston 184
6.5 High-speed soot pyrometry 191
6.6 Summary 195

Chapter 7 DISI Engine Modelling **197**

7.1 Introduction 197
7.2 Combustion modelling with ISIS 198

7.2.1 Graphical user interfaces 198
7.2.2 Tuning ISIS 200
7.2.3 Engine discussions 201
7.3 Summary 204

Chapter 8 Conclusions and Future Work **205**

8.1 Conclusions 205
8.1.1 Optical diagnostics 205
8.1.2 Combustion analysis 211
8.2 Suggestions for future work 212

References **217**

Appendix A Incidental Work — Lumenera Camera Software **224**

List of Figures

1.1 Targets for NEDC drive-cycle fleet average CO_2 emissions in the ACEA agreement 2

1.2 Development of gasoline engines 3

1.3 Swirl and tumble in-cylinder flows 9

1.4 Two injection modes of DISI engine 10

1.5 Three GDI combustion concepts 11

1.6 Spray characteristics of a GDI injector 13

1.7 Principle of operation of the high-pressure swirl nozzle 14

1.8 A model Bosch piezoelectric injector 15

1.9 Typical arrangement for an engine with optical access 18

2.1 Engine instrumentation system and structure of data flow 29

2.2 Experimental layout on the Jaguar DISI engine 31

2.3 Measured pressure change as a sum of piston motion and combustion 35

2.4 An example of heat release analysis 42

2.5 Graphical user interface of CoBRA 47

2.6 Simple model of charge coupled device 49

2.7 Schematic of the intensified charge coupled device (ICCD) detector 51

2.8 Schematic of the CCD imaging apparatus 52

2.9 Pulnix TMC-9700 CCD colour video camera 53

2.10 Schematic of image registration 58

2.11 Sliding neighbourhood blocks 61

2.12 The graphical user interface of the image processing software 62

2.13 An example of white image masking 63

2.14 An overview of the control point selection tool 64

2.15 An example of image registration 64

2.16 Example images from image processing 65

3.1 Correlation between 'nominal' spark timing and 1% MFB timing 70

3.2 Sample results of cycle-averaged Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$ 71

3.3 Power density spectrum of the pressure signal 72

3.4 Correlation between 0–1% MFB and 0–10% MFB 73

3.5 Correlation between 1% MFB and Q_{nmin} 74

3.6 Correlation between Q_{nmin}–10% MFB and 10–90% MFB 75

3.7 Q_{nmin}–10% MFB for various settings 76

3.8 The effect of nitromethane on Q_{nmin}–10% MFB and 10–90% MFB 77

3.9 Effect of nitromethane content, fuel needle setting and glow plug type 78

3.10 Cycle-by-cycle correlation between IMEP and Q_{nmin} 80

3.11 Correlation between the CoC of the n^{th} cycle and IMEP of the $(\mathrm{n+1})^{th}$ cycle 81

3.12 Scatter plot of cycle-by-cycle data 82

4.1 Three-state energy level model for electronic transitions 86

4.2 Schematic of laser-induced scattering and fluorescence 90

4.3 The time-resolved correlation between fluorescence yield and laser signal 92

4.4 The time-resolved fluorescence signal 97

4.5 Absorption spectra of acetone varying with temperature 100

4.6 Setup of the Raman cell 101

4.7 Coevaporation of (a) acetone/isooctane, (b) 3-pentanone/isooctane mixtures at 303 K [58] 106

4.8 Coevaporation ratios at 303K according to equilibrium VLE calculations 108

4.9 Absorption spectra of first-batch components and tracers 109

4.10 Absorption spectra of second-batch components 110

4.11 System for measuring the fluorescence spectra of fuels and tracers 111

4.12 Absorbance and fluorescence spectra of tracers and component fuel 114

4.13 Prediction of signal-to-noise ratios for tracers 115

4.14 Excimer laser power curve for standard neon fill 119

4.15 Estimation of the SRS efficiency 120

4.16 Example of estimation of the saturation intensity for toluene vapour 121

4.17 A schematic of the pulse stretching 122

4.18 Schematic of the PLIF calibration loop (plan view) 123

4.19 Schematic of the optical access in the pentroof region and a sample raw image 125

4.20 Sample images of background and PLIF signal 125

4.21 PLIF calibration results against crank angle with different lambda 126

4.22 PLIF calibration results against mixture strength at different crank angles 127

5.1 Procedures of the two strategies for colour-ratio pyrometry 134

5.2 Look-up tables calculated with the mapping with a variable KL strategy 135

5.3 Variation of soot emissivity with wavelength and KL 136

5.4 Expected R/G and R/B ratios of CCD camera response with different KL values 136

5.5 CCD camera calibration setup 138

5.6 Transmission characteristics of filters 139

5.7 Images at different wavelengths as seen by the camera 140

5.8 Colour filter array (CFA) of the camera 140

5.9 Effect of the re-demosaicing process 142

5.10 The camera's response test for linearity with aperture 144

5.11 Dark noise with respect to gain 145

5.12 Dark noise with respect to exposure 146

5.13 The camera's response test for linearity with sensitivity 147

5.14 The camera's response test for linearity with exposure 148

5.15 Photodiode sensitivity and signal from the tungsten bulb 149

5.16 Normalized spectral response of the camera 149

5.17 Sample images from the validation experiments 151

5.18 Correlation between measured and inferred temperatures of the black body 151

5.19 Tungsten emissivity data used in the tungsten filament test 152

5.20 Pyrometry results from filament lamp experiment 153

5.21 Pyrometry results from ribbon lamp experiment 156

5.22 The effect of 1 bit digitization error on the estimated temperature 158

5.23 The effect of random noise and global averaging on temperature estimation 160

5.24 Set-up for measurements on the optical DISI engine with CRP 161

5.25 Sample image from the Jaguar DISI engine 162

5.26 Temperature and KL distribution for an sample image 163

5.27 Relation between temperature and crank angle with a fixed KL 164

5.28 Relation between temperature and mixture strength with a fixed KL 164

5.29 Relations between variable-KL results and crank angle 165

5.30 Relations between variable-KL results and mixture strength 166

6.1 The control pulse trains for the high-speed imaging 171

6.2 Sample combustion images with different fuels 172

6.3 Image sequence vs. cylinder pressure and MFB for combustion event A 175

6.4 Image sequence vs. cylinder pressure and MFB for combustion event B 176

6.5 Image sequence vs. cylinder pressure and MFB for combustion event C 177

6.6 Effect of smoothing an image using non-linear diffusion filtering 180

6.7 Detected flame fronts from sample images 181

6.8 Sample results of critical point detection 183

6.9 Sample results of temporal contour interpolation 184

6.10 The Perspex piston made at Oxford 185

6.11 Distortion correction for the Perspex piston 187

6.12 Example of distortion correction on a combustion image 187

6.13 Sample combustion images with distortion correction 188
6.14 Example of flame fronts with Perspex piston 188
6.15 Analysis of flame properties 189
6.16 Example of flame front propagation with Perspex piston 190
6.17 Example results of high-speed CRP (1) 193
6.18 Example results of high-speed CRP (2) 194
6.19 Example results of high-speed CRP (3) 195
6.20 Example results of high-speed CRP (4) 196

7.1 Graphical user interface ISISxp 199
7.2 Graphical user interface ISIS2MAT 200
7.3 Tuning ISIS with data from CoBRA 202
7.4 Comparison between ISIS and CRP (1) 202
7.5 Comparison between ISIS and CRP (2) 203

A.1 The user interface of LuCamICE at start-up 226
A.2 The user interface of LuCamICE when the camera is connected 227
A.3 Pages of the image and video properties 232

List of Tables

2.1 An example of the statistical output data from CoBRA 46

2.2 A list of common transformations 59

3.1 Effect of γ on the phasing of Q_n and $dQ_n/d\theta$ 71

4.1 Properties of component-fuel components 107

4.2 Selected common tracers and their properties 107

5.1 Summary of pyrometry results for a sample image 162

6.1 Engine operating conditions and statistics of the three cines 174

6.2 Combustion analysis results of the three combustion events 174

Nomenclature

S.I. units and standard notations for chemical elements, formulae, and abbreviations are used in this thesis. Other abbreviations are listed below.

Acronyms

ABDC	After bottom dead centre
ACEA	European automobile manufacturers association
ADC	Analogue to digital converter
AFR	Air/fuel ratio
ATDC	After top dead centre
BBDC	Before bottom dead centre
BDC	Bottom dead centre
BMEP	Brake mean effective pressure
BNC	Bayonet Neill-Concelman connector
BTDC	Before top dead centre
CA	Crank angle
CAR	Continuously adjustable resolution
CARS	Coherent anti-Stokes Raman spectroscopy
CCD, ICCD	Charge-coupled device and intensified charge-coupled device
CDF	Cumulative distribution function
CFA	Colour filter array
CFD	Computational fluid dynamics
CMOS	Complementary metal oxide semiconductor
CoBRA	Combustion burn rate analysis
CoC	Completeness of combustion
CoV	Coefficient of variation
CRP	Colour-ratio pyrometry
DAQ	Data acquisition
DC	Direct current
DFWM	Degenerate four-wave mixing
DISI	Direct injection spark ignition
DSP	Digital signal processing
DVM	Digital voltmeter
EDR	Extreme dynamic range
EGR	Exhaust gas recirculation
EPO	Exhaust port opening
EOC	End of combustion
ETCS	Engine timing control software

EVC, EVO	Exhaust valve closure and exhaust valve opening
FID	Flame ionisation detection
fps	Frames per second
FVVT	Fully variable valve train
GDI	Gasoline direct injection
GUI	Graphical user interface
HC, UBHC	Unburned hydrocarbons
HCCI	Homogeneous charge compression ignition
HPLC	High performance/pressure liquid chromatography
IC	Internal combustion
IMEP	Indicated mean effective pressure
IPC	Inlet port closure
ISIS	Integrated spark-ignition engine simulation
IVC, IVO	Inlet valve closure and inlet valve opening
KLSA	Knock limited spark advance
LabVIEW	Laboratory virtual instrument engineering workbench, a graphically programmed computer language for real-time instrumentation
LIEF	Laser-induced exciplex fluorescence
LII, LIS	Laser-induced incandescence and laser-induced scattering
LNT	Lean NO_x trap
LNV	Lowest normalised value
LWM	Local weighted mean
MATLAB	Matrix laboratory, a high level programming language widely used in engineering and scientific computing
MBT	Minimum ignition advance for best torque
MCP	Microchannel plate
MFB, x	Mass fraction burned
MON	Motor octane number
MPI	Multipoint port injection
NEDC	New European drive cycle
PIV	Particle image velocimetry
PLIF	Planar laser-induced fluorescence
PMT	Photomultiplier tube
RGB	Red, green and blue
RIPE	Raw image processing for Engines
RON	Research octane number
RPM	Revolutions per minute
SBS, SRS	Stimulated Brillouin scattering and stimulated Raman scattering
SMD	Sauter mean diameter
SNR, NSR	Signal-to-noise/noise-to-signal ratio
SOC	Start of combustion
TDC	Top dead centre
TMB	Trimethylbenzene
TTL	Transistor-Transistor logic
UV	Ultraviolet
VLE	Vapour liquid equilibrium

Symbols

$a, 1-a$	Weighting coefficients
A	Absorbance or aperture area
c	Specific heat capacity, kJ/(kg·K) or speed of light or calibration factor
C	Ratio or constant
E	Characteristic function for the variable-KL strategy
h	Planck's constant, $6.6260755 \times 10^{-34}$ J·s
i, j	i^{th}, j^{th} sample of the digitized data
I	Intensity
k	Rate (of spontaneous emission *etc.*), ns
K	Absorption coefficient per unit flame thickness
L	Length of an optical path
m	Mass
n	Polytropic index or number of subdivisions or order of a digital filter
N	Number of samples or number density or f number
P	Pressure
Q	Heat
rev	Revolution
R	Specific gas constant, kJ/(kg·K) or ratio
R, G, B	Red, green and blue pixel values
S	Discrete cumulative distribution
T	Temperature, °C or K
t	Time, s
U	Internal energy, kJ
V	Volume
W	Work, kJ
x, y, z	Cartesian coordinates

Greek letters

Δ	Difference
Ψ	Total normalised pressure change due to combustion
α	Constant in the Hottel and Broughton correlation, 1.39 for visible wavelengths
α, β	Look-up table entries
γ	Ratio of specific heats
θ	Crank angle, °
ε	Emissivity of a radiating medium or truncation error
η	Efficiency
λ	Wavelength, nm or reciprocal of equivalence ratio $= \frac{\text{actual air/fuel ratio}}{\text{stoichiometric air/fuel ratio}}$
ρ	Density, kg/m^3
σ	Cross section
τ	Time, s

χ	Completeness of combustion
ω	Cutoff frequency of a digital filter

Subscripts

∞	Infinity
0	Initial/reference value
1, 2	Values at state 1, 2
10, 50, 90	The states at which 10, 50 and 90% of the fuel has evaporated during the distillation process
abs	Absorption
b	Black body
c	Combustion
cal, C	Calibration
flu, PLIF	Fluorescence
hr	Heat release
ht	Heat transfer
i, j	i^{th}, j^{th} sample or species i, j
in	Inlet manifold
m, max	Maximum
min	Minimum
n	Order of a digital filter
net	Net
p	Pressure
q	Quenching
QE	Quantum efficiency
t	Time
S	Swept volume
sat	Saturation
t	Tungsten
tran	Transmission
v	Volume or piston motion or laser intensity
vap	Vapour phase
λ	Spectrum

Superscripts

$*$(asterisk)	Normalised value or decay time constant or excited state
$\cdot$(dot)	Quantity per unit time
i, j	i^{th}, j^{th} sample
R, G, B	Pixels with the true colour information

Chapter 1

Background

Genius without education is like silver in the mine.

— Benjamin Franklin (1706–1790)

1.1 Introduction

For over a century in the history of internal combustion (IC) engines, two concepts have played a dominant role during the process of development and production. The spark ignition or gasoline engine shows the distinctive feature of high specific power and large speed range due to its good level of air utilisation; whilst the compression ignition or Diesel engine, offers better thermal efficiency from the use of a higher compression ratio, and the absence of throttling. Combining the positive aspects of these two engine types has always been a goal of engine research and development in the automobile industry, and the Direct Injection Spark Ignition (DISI, or Gasoline Direct Injection) engine is seen as one of the most promising ways to achieve this aim.

The concept of DISI engines has gained a resurgence of interest in recent years as a way of reducing both fuel consumption and emissions to meet the current strict legislation as well

as the future targets for new European drive-cycle (NEDC) fleet average CO_2 emissions (Figure 1.1), owing to the improvements in fuel injection technology as well as advances in the understanding of the processes involved in mixture preparation and stratified charge combustion.

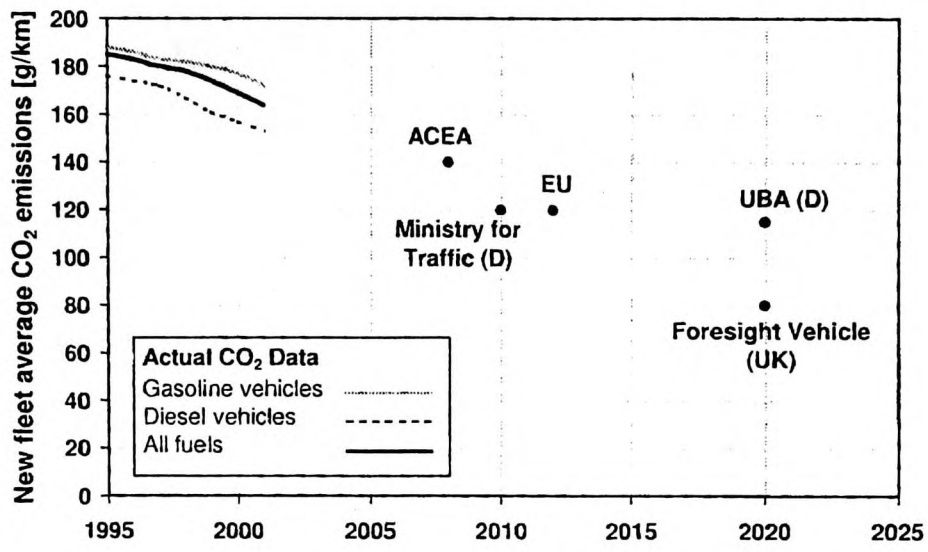


Figure 1.1: Targets for NEDC drive-cycle fleet average CO_2 emissions in the ACEA agreement, adapted from Lake *et al.* [40]

1.2 Concept of direct injection spark ignition

1.2.1 Developmental history

Direct injection systems were first adopted on highly boosted airplane gasoline engines in the 1930s. In 1937, an engine with a Bosch mechanical gasoline direct injection system first took off in an aircraft. Ricardo developed their 2-stroke sleeve valve engine, which was the basis of the V12 Rolls-Royce 'Crecy' GDI engine (Hiett and Robson [25], cited by Jackson *et al.* [31]). The Crecy achieved a much higher output than other types of engine at that time, but was not developed further because of the advent of gas turbines.

The first passenger car with a series-produced GDI engine, a Gutbrod, hit the road in 1952. In 1954, Mercedes-Benz 300SL introduced a direct injection system (cited by Iwamoto *et al.* [30]) to solve the inherent performance problems of carburettor systems.

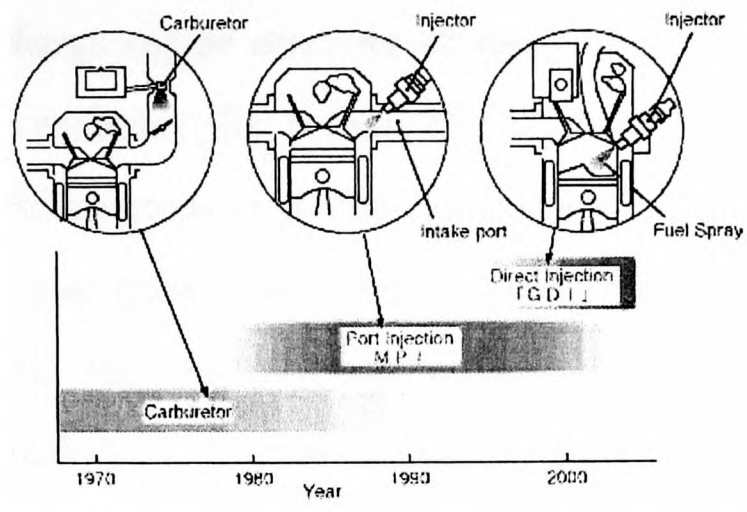


Figure 1.2: Development of gasoline engines (from www.mitsubishi-motors.co.jp)

Carburetors were replaced by inlet port fuel injection systems, which were more robust and provided better mixture without excessive charge motion. Low pressure fuel injection then gradually dominated the development of automotive gasoline engines, initially using a single injector in the inlet manifold, and then moving to multipoint port injection (MPI) into the back of each inlet valve (see Figure 1.2). Gasoline direct injection engines initially adopted an early injection strategy, that is, they injected fuel into the cylinder during the intake stroke for a homogeneous mixture close to stoichiometric. Therefore the objective was to improve full load output rather than fuel economy.

The energy crisis of the early 1970s emphasised the need to improve the fuel efficiency of internal combustion engines. Spark ignition engines, which operate unthrottled in an extremely lean condition by distinctively stratifying the charge and provide rich air-fuel mixture in the proximity of the spark plug, were of particular interest. Several stratified charge combustion concepts were proposed, most notably, prechamber injection (Pischinger and Klocker [47]) with which the charge stratification occurs in a divided combustion chamber, and the conventional undivided combustion chamber. A slightly rich mixture is formed by injecting fuel into a prechamber, where it is then ignited and the flame propagates to the lean mixture in the main combustion chamber. Alternatively, spray impingement and air motion is employed to achieve stratification in a conventional chamber.

These early stratified charge engine concepts all used injector technology derived from the Diesel system. For example, the Texaco TCCS system (Alperstein *et al.* [4]) and VW's research engine (Schapertons *et al.* [53]), used a narrow spacing configuration with the spark plug located near diesel type injectors, and considerably higher compression ratios were achieved. But the relatively poor atomization and narrow injection timing window still called for high level of intake air motion to improve fuel vaporization. The Ford PROCO engine (Scussel *et al.* [55] cited by Anderson *et al.* [5]) used an outwardly opening pintle atomizer with vibration to enhance fuel atomization, with low fuel pressure and early injection during compression stroke.

The MAN-FM (cited by Jackson *et al.* [31]) and OSKA (Kato *et al.* [34]) combustion systems used fuel spray wall-impingement to help break up the spray to promote evaporation and mixing. A relatively high level of swirl was needed to achieve stratification and vaporization. A major problem with these systems was the lack of variability of injection characteristics between part load and full load, when significant wall impingement of large amounts of fuel would occur. Research activities on GDI engines had reduced by the end of this stage.

1.2.2 Recent applications

The great majority of the recent systems use a combination of homogeneous charge for increased full-load performance and stratified charge to take advantage of the part-load fuel economy. The main organisations involved and the prototype and production engines developed in this stage are listed briefly below.

Mitsubishi first used its 'Gasoline Direct Injection' system for production, which was introduced into the Japanese market in 1995, then the European market in 1997. It is an in-line 4-cylinder engine with 4 valves per cylinder and a total swept volume of 1864 cm^2. The main concept is to use upright inlet ports and reverse tumble air motion (see

Section 1.4.1) to facilitate a stratified mixture formation.

Ricardo patented a research engine with a much smaller swept volume than the Mitsubishi GDI which uses the reverse tumble system. Toyota brought their engine onto the Japanese market in 1996, which use a "swirl control valve" to vary in-cylinder air motion according to the engine operating point. The Nissan and AVL prototype systems use a similar swirl control valve to the Toyota system, but with a piston crown shape similar to the Mitsubishi GDI.

Mercedes, Ford *et al.* have used a centrally mounted injector, close to the spark plug, and use the spray shape to influence stratified charge mixture formation. The Orbital Engine Company's combustion system is based on a pressurised-air-assisted injector, which they have recently applied to automotive four stroke engines (Houston and Cathcart [29]).

Peugeot Citroen (Georjon *et al.* [20]) and Delphi (Befrui *et al.* [8]) have utilised similar closely-arranged spray-guided (see Section 1.4.3) combustion systems. The former conducted experimental validation of numerical modelling and prediction, whilst the latter investigated injector characteristics.

As engine downsizing is becoming more attractive to auto-manufacturers in Europe because of the cost-effectiveness and benefits on fuel consumption introduced by this concept, Ricardo developed turbocharging systems for downsized DISI engines [40] (detailed in Section 1.4.6). More recently, Lotus, together with other consortium members, have been investigating a spray-guided homogeneous charge DISI engine equipped with a fully variable valve train (FVVT) system; this strategy is believed to reduce the engine cost and circumvent the need for NO_x traps, hence providing global fuel compatibility with high-sulphur fuels.

1.3 Advantages of the DISI concept

1.3.1 Fuel economy

When the engine works under low load, the Diesel compression ignition engines inject less fuel into the cylinder, as self-ignition occurs spontaneously at points within the combustion chamber where the mixture reaches the appropriate conditions. Thus combustion is completed for any amount of excess air. While in SI engines, the ignition source is fixed spatially and temporally, that is, an ignitable mixture must be available at the fixed proximity of the spark plug. To this end, carburettor and port-injection systems are designed to provide a homogeneous fixed air/fuel ratio (AFR), which is likely to release much extra energy at low load within the whole combustion chamber. If the mixture is leaner than stoichiometric, the speed of combustion will decrease and the combustion will be incomplete. Thus throttling is used to reduce the mass of fuel while still maintaining a mixture close to stoichiometric. But throttling reduces the intake manifold pressure and therefore the density of the air charge; hence it will inherently decrease the volumetric efficiency and increase the work dissipated in the gas exchange process.

With the elimination or reduction of throttling, the gasoline direct injection concept attempts to achieve an ignitable mixture near the spark plug, but to have the charge stratified so as to reduce the total amount of fuel burned. Hence a remarkable fuel economy at part load could be gained, because the pumping loss is significantly reduced and the volumetric efficiency is increased. In addition, a higher thermodynamic efficiency can be achieved with homogeneous lean charge through higher ratio of specific heats.

The low fuel consumption of GDI also provides the capability of reducing CO_2 emissions which is of great concern for the greenhouse effect.

1.3.2 Charge cooling

In a Direct Injection engine, the fuel is injected directly into the cylinder as a well atomized spray of fuel droplets, of which the vaporisation absorbs thermal energy from the air, reducing the charge temperature in a process known as charge cooling.

If this decrease in charge temperature occurs during the intake stroke, before inlet valve closure (IVC), the density of the intake charge will increase. Since the volume of air inhaled depends on the geometry of the engine, a greater mass of air will be present in the cylinder after intake. Therefore the volumetric efficiency will increase. The KIVA (CFD) code analysis used by Han *et al.* [23] showed that the charge temperature had decreased by 15 °C by IVC, and that the trapped mass increased by 2% for early (before IVC) injection. Wyszynski *et al.* [74] undertook research on the single-cylinder Mitsubishi GDI spark ignition engine to quantify the extent of charge cooling, and it was shown that the volumetric efficiency depends both on the fuel and its time and place of injection; for stoichiometric operation with unleaded gasoline, changing from port injection to direct injection led to a 9% increase in volumetric efficiency, which was improved by a further 3% when 10% methanol was blended with the gasoline.

If the in-cylinder vaporisation occurs after IVC, the mass of air inhaled into the cylinder cannot be influenced, so the volumetric efficiency remains unchanged. However, for the same compression ratio, the charge cooling reduces the in-cylinder temperature and hence improves the knock tolerance, or allows earlier ignition timing for a given probability of knock, thereby advancing the knock limited spark advance (KLSA), which will give greater flexibility in setting ignition timing to give optimum torque output, known as minimum ignition advance for best torque (MBT) timing. On the other hand, for the same level of knock susceptibility, a higher compression ratio is realisable so that the thermal efficiency of the engine will be increased according to the ideal Otto analysis.

1.3.3 Fuelling control

Fuelling control for every combustion cycle is very important so as to keep performance and emissions within acceptable margins. During transient operation with port injection, such as a sudden open or closure of the throttle, the fuel film will have a hysteresis of fuel vaporisation and mixture delivery from the injector to the cylinder.

The implementation of direct injection allows easier control of mixture preparation and will speed up the response, for instance, providing the capability of immediate fuel cut-off during deceleration. In addition, because of the elimination of the inlet condition-dependent fuel film, the requirement of rich mixture operation for cold starting can be released, reducing unburned HC emissions compared to the port injection engine (Han *et al.* [23]).

1.4 Characteristics of DISI engines

1.4.1 Air intake and charge motion

The rotational component of the in-cylinder air motion with its axis parallel to the cylinder axis is denoted as swirl, whilst the component perpendicular to this is known as tumble (see Figure 1.3). The tumble motion is generally realised by an upright straight intake port, in contrast with the transverse intake port of a swirl system. The turbulence is necessary to help the air/fuel mixing process, whilst the mean flow helps to transport the mixture to the spark vicinity and to stabilise the stratified mixture plume. Homogeneous SI combustion needs a high turbulence intensity and a low mean flow velocity to combust the mixture effectively. Whilst for stratified charge combustion, to ensure a stable and stratified mixture, a greater mean flow velocity and lower level of turbulence is required.

The cylindrical geometry is more favourable to maintaining swirling flows with less viscous

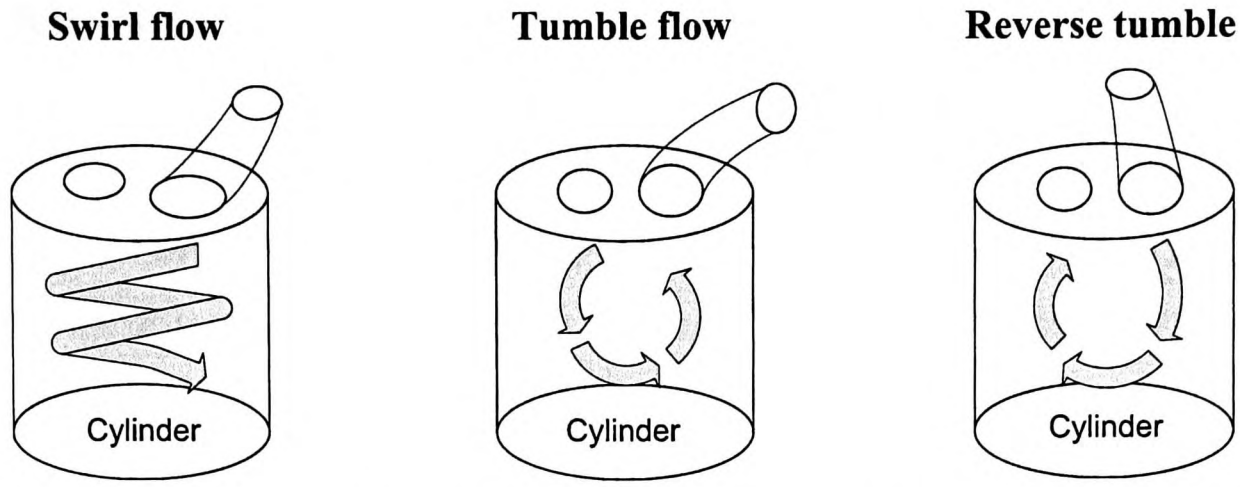


Figure 1.3: Swirl and tumble in-cylinder flows

dissipation. Tumble flows decay relatively quickly into secondary vortices during the compression stroke, especially for a flat piston crown. Specific piston crown shape designs can however attenuate this decay to give higher mean flow rates.

Numerical investigations (Han *et al.* [23]) showed that the spray direction and penetration distance are heavily influenced by the intake air flow, and conversely that the spray-induced gas motion affects the large scale flow structures, and can increase the mean velocities of the gases.

All of the early stratified charge combustion systems (Section 1.2.1) used high levels of swirl to achieve stratification, which limited torque output at high loads, as the flow velocities generated by the inlet ports scaled proportionally to engine speed, introducing excessive flow losses. Also, high swirl rates can centrifuge larger fuel droplets onto the cylinder wall.

The problem was addressed differently by numerous workers, including Toyota, Nissan and AVL. These engines used a control valve in one of the inlet ports to vary the levels of swirl in the cylinder according to the load required and the injection timing strategy.

Another problem is that, the momentum of charge motion is proportional to the engine speed, while the momentum of fuel spray does not depend on the engine speed. Therefore, the range of engine speeds in which the adequate air/fuel mixing is realised is limited.

1.4.2 Mixture preparation

For a homogeneous charge, enough time is needed for fuel/air mixing and evaporation before ignition, so injection usually occurs during the intake stroke, known as early injection mode (see Figure 1.4). Charge cooling would occur as the air flowed into the cylinder, increasing the volumetric efficiency. It is important to minimise impingement of the fuel spray onto the walls of the cylinder or onto the piston, as then thermal energy from the combustion chamber surface will be taken to vaporise the fuel in preference to thermal energy from the charge air, thereby reducing the charge cooling effect.

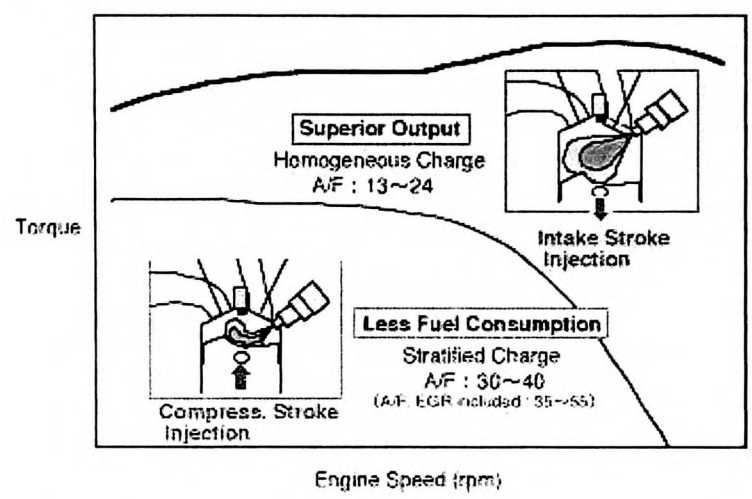


Figure 1.4: Two injection modes of DISI engine (from www.mitsubishi-motors.co.jp)

The mixture formation of a stratified charge is more complex and is dependent on the interaction of the fuel spray, the in-cylinder air motion, and the changing combustion chamber surface in the form of the moving piston. To avoid the excessive mixing which would result in a reduction of stratification and a more homogeneous mixture, the time between injection and ignition must be reduced. Thus the injection usually occurs during the compression stroke, as the piston is moving towards TDC and ignition, known as late injection mode (see Figure 1.4). In this case, charge cooling cannot affect the volumetric efficiency but can still limit knock. However, it is generally good to avoid spray impingement onto the walls (wall-wetting) so as to limit UHC and soot emissions.

1.4.3 Spray guiding concepts

There are three GDI-combustion concepts, described as wall-guided, air-guided and spray-guided combustion systems (see Figure 1.5). The distinction between these different concepts is the method used to guide the spray and mixture to the spark plug.

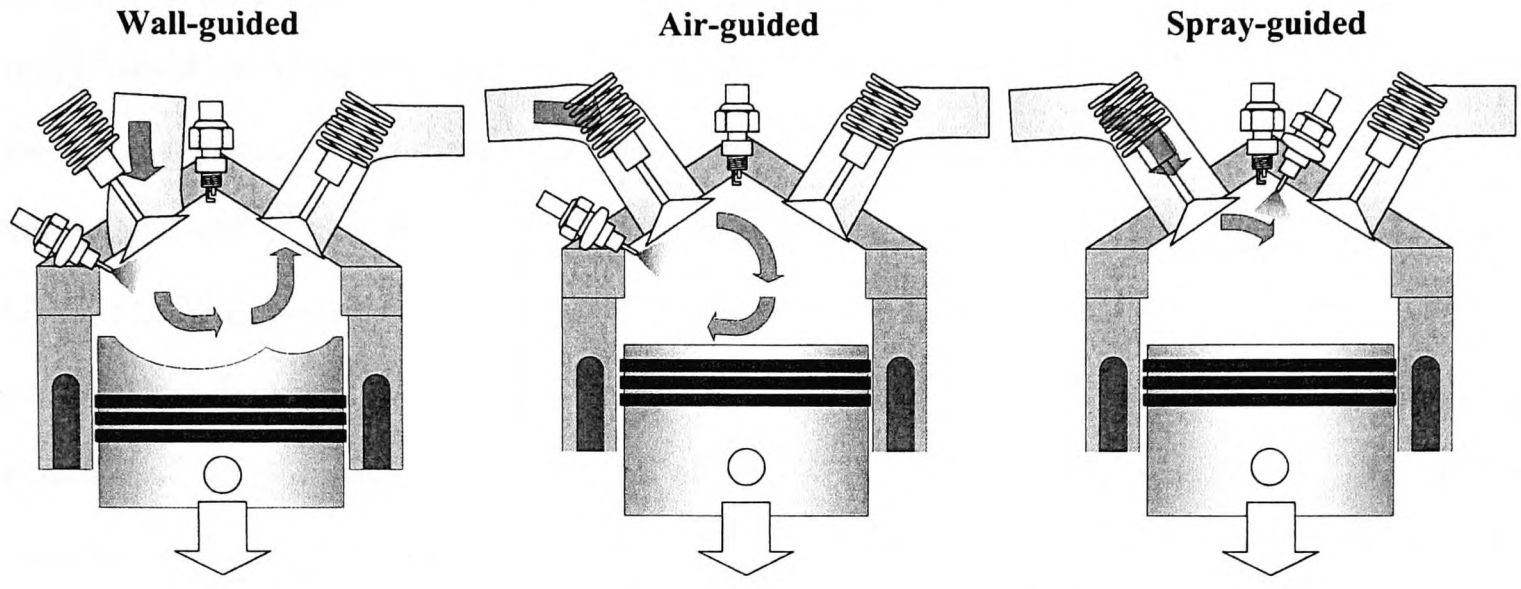


Figure 1.5: Three GDI combustion concepts

In the wall-guided concept, the injected fuel is guided to the spark plug by an interaction with the combustion chamber wall, generally a well defined piston cavity. In-cylinder flow including tumble, reverse tumble or swirl (see Section 1.4.1) is needed for the desired mixture formation. Since the air/fuel mixing is supported by the piston shape and motion, the fuel spray inevitably impinges onto the piston. This, as well as the unintentional impingement during homogeneous charge, can form fuel films on the piston and walls, which is known as wall-wetting. Combustion of such films causes diffusion flames which are quite likely to increase soot generation, HC emissions and cause deposits. Moreover, the increased surface-to-volume ratio due to the irregular piston shape leads to more heat transfer.

The air-guided concept can be described as air assisted mixture transport. The evaporated fuel is led towards the spark plug by a well defined in-cylinder flow. The interaction between the injected spray and the charge motion has to be appropriate under a large

part of the engine operating regime. As in Figure 1.5, the air flow is now clockwise in order to prevent the fuel from impinging onto the piston.

Referred to as the second generation of DISI combustion concept, spray-guided concept has a significantly different approach for realising mixture preparation for stratified and homogeneous operation, that is, it directs the fuel spray straight to the spark plug. Normally based on a narrow spacing configuration, which means the injector and spark plug are closely arranged, the mixture preparation and combustion do not rely as much on charge motion or piston design as they do on injector characteristics. The appropriate AFR at different engine loads (a different amount of injected fuel) must be realised by the penetration depth of the injected spray, and is therefore mainly determined by the spray dynamics. A short mixture formation time until ignition is available due to the narrow spacing, and the engine requires late injection to maintain desirable stratification. According to the small extension of the ignitable mixture zone at the periphery of the spray, this system is more sensitive to spray and installation tolerance and calls for stringent requirements on injector parameters, such as fine atomization, perfectly axisymmetric spray plume and a compromise between high pressure (for atomization) and short penetration (for avoidance of wall-impingement). In the case of a narrow spacing, the liquid fuel spray impinged onto the spark plug as well as the over-rich mixture located around the spark plug can lead to soot and deposit accumulation, which will result in spark plug fouling. If a high energy ignition system is adopted to reduce spark plug fouling, then the durability of the spark plug could decrease.

Nevertheless, the spray-guided concept has shown greater potential for ultra-lean burning than the two wide-spaced concepts, representing a promising solution, and it thus gains notably more interest in very recent research activities. Compared to a wall-guided DI system, Ford and Bosch reported a 5–6% average NEDC fuel economy improvement with a spray-guided DI engine, a 40% reduction in unburned HC and more rapid catalyst light-off times; extension to the stratified operating window was achieved as well, leading to a wider

overall operating regime (0.5–5bar BMEP) preferable for real-world DISI engines [71].

1.4.4 Fuel injection technology

The fuel system requirements for GDI injection are entirely different and more rigorous than those for Diesel or port injection systems, primarily in terms of minimum droplet size. MPI injectors produce droplets of Sauter Mean Diameter (SMD) of 120 to 200μm, whilst DISI investigations quote maximum SMDs of 30μm or so. Thus for all operating conditions a well-atomised fuel spray is needed, and for stratification the spray geometry needs to be compact and stable (Kume *et al.* [39], cited by Zhao *et al.* [76]). For late injection the fuel is injected at high rates during the compression stroke at cylinder pressures up to 2MPa. At full load, a well-dispersed spray is needed even at the higher fuel quantities (see Figure 1.6). The injector must be resistant to deposit formation and to higher body and tip temperatures than a PI injector. The SMD of the spray droplets must be less than 25μm, with the DV90 not exceeding 45μm.

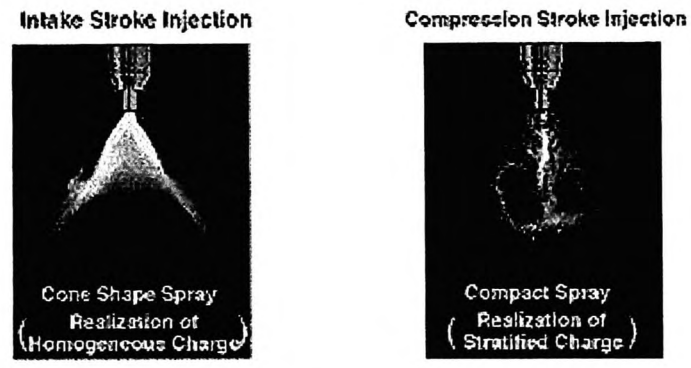


Figure 1.6: Spray characteristics of a GDI injector (from www.mitsubishi-motors.co.jp)

Three solenoid-actuated injector configurations are presently under development for spray-guided systems, namely, high-pressure swirl, outward-opening and multi-hole injectors. The high-pressure swirl injector (see Figure 1.7) is the type most often seen in DI applications. It can be considered as a multi-hole nozzle with an infinite number of holes. It is noted by Tomoda *et al.* [67] that the swirl nozzle achieves smaller fuel spray droplets for the same pressure than the multi-hole nozzle, allowing the use of a lower pressure

system whilst maintaining vaporisation rates. The swirl nozzle also has a smaller axial penetration distance, helping to limit wall-wetting by impingement of the spray onto the surface of the combustion chamber. Laser diffraction measurements showed a different distribution of droplet size for the swirl and multi-hole nozzle, with similar Sauter Mean Diameter, but significantly more large diameter droplets in the asymmetric size distribution seen for the multi-hole nozzle. This means that the DV90 may be a better measure of the atomisation performance of a given injector, which is related closely to the unburned HC emissions.

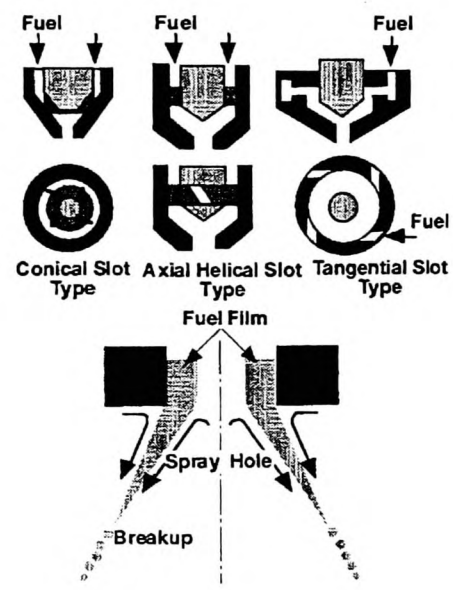


Figure 1.7: Principle of operation of the high-pressure swirl nozzle, Iwamoto *et al.* [30]

To avoid wall-impingement and leakage, it is necessary to reduce the design pressure in high pressure swirl injector systems, but this competes with the requirement for smaller droplet diameters achieved with increases in operating pressure.

The outward-opening configuration features a needle valve and a circular cross section flow path. Such injectors are more robust against deposits than swirl ones but lack the flexibility of spray pattern.

In contrast, the multi-hole configuration is flexible through varying the number and distribution of the holes, the orientation of the injector and the injection angle of every single stream. Ortmann *et al.* [45] argued that, compared to a swirl injector, a multi-hole injector with its enhanced air entrainment in the spray and stability against ambient

pressure changes, can lead to a smaller cycle-by-cycle variation of IMEP and less soot emissions for wall-guided combustion system, provide wider ignition windows under high injection pressures for air-guided system, and allow a stable engine operation at all loads for spray-guided systems. Wirth *et al.* [71] also favoured a multi-hole configuration with an injection pressure up to 20MPa on the basis of robustness against deposit formation on the injector, cost and soot emission level.

In the other category of DI injector configuration are piezoelectric injectors (Figure 1.8). Such injectors are fitted with several hundred thin crystals within the body. The expansion of the piezo wafers in the presence of an electric field is transferred to the nozzle needle in less than 0.1ms. Therefore, they can cycle twice as often as magnetically-triggered injectors and are well suited for desired high injection pressure. Such short injection duration can be used to generate mixture flows easy to ignite. Additionally, the injector's circular side-opening creates a cohesive and reproducible cone-shaped stream flowing from the top center of the combustion chamber and can be timed to produce a charge of varying density. Its inline design means very little fuel return flow is needed, which reduces the delivery rate and energy required from the high-pressure fuel pump. The impediment to the take-off of such injectors is their larger size and expense than solenoid-controlled ones.

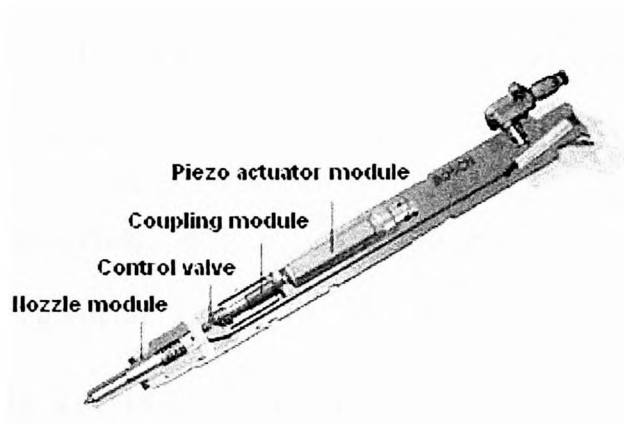


Figure 1.8: A model Bosch piezoelectric injector

The design consideration of GDI Injector location is of great importance as well. As described in Section 1.4.3, close-coupling of the spark plug and injector is likely to cause

spark plug fouling. For a wide-spacing, the inlet side of the cylinder is more favourable for injector position due to packaging and thermal reasons, and this also allows easier interaction between the inlet air flows and the fuel spray for mixture formation. However, injector location should not compromise the valve sizes.

1.4.5 Combustion and emissions

Because of the increased compression ratio and overall lean environment with most DISI applications, NO_x emissions can increase over conventional PFI engines. Many investigations have found that the DISI engine can tolerate much higher levels of EGR, which can help to reduce these emissions. At partial load, there exist areas of local high air/fuel ratio, where most of the NO_x is produced. Higher levels of in-cylinder flow will disperse the fuel mixture plume, leading to slow combustion, especially with EGR. Future strict legislation will require NO_x emissions that are unrealisable by the DISI engine, and the tendency for overall lean combustion will cause difficulty for exhaust aftertreatment, because of the extreme difficulty of reducing NO_x with excess oxygen, as well as the sensitivity of catalysts to the sulphur content and temperature of the exhaust gas. The future of the DISI concept is thus tightly tied to the progress of lean-NO_x reduction technology, such as lean NO_x traps (LNT) with compromises, however, in the economy benefits due to catalyst cost and regeneration as well as the requirement of high-quality fuels.

A significant amount of unburned HC is also generated at partial load due to several reasons: 1) the probability of quenching when the flame propagates to the relatively lean mixture region; 2) relatively low in-cylinder temperatures close to the end of combustion at high air/fuel ratios, in which case the HC may not be oxidised sufficiently; 3) wall-wetting for the wall-guided combustion system, which can cause liquid fuel film burning and therefore an increase in unburned HC emissions and particulates.

Soot generation in DISI engines is much higher than PFI engines owing to locally rich

mixtures and diffusion combustion where there is insufficient vaporisation, but still lower than Diesel engines by at least an order of magnitude (Eichlseder *et al.* [17]).

1.4.6 Fuel efficiency strategies

For the measurement of fuel efficiency, the standard drive cycles used both for numerical simulation and experimental evaluation of engine operation do not represent a large part of the operating requirement within the stratified charge region of the engine map, where the fuel economy benefits are potentially the greatest.

Air-boosting a GDI engine, either by turbo- or super-charging has been investigated recently. In general a fuel efficiency benefit is seen for boosted engines, when for a given power requirement a smaller engine can be used. This is known as downsizing by turbocharging. The main benefits would occur at lower load levels where there is a large air requirement during stratified charge operation.

One of the effects of pressure charging a Spark-Ignition engine is to increase the knock sensitivity by raising the peak temperatures seen during the compression stroke. Therefore there is a limit on the so-called admissible boost pressure, which also depends on the compression ratio.

1.5 Diagnostics and instrumentation for DISI research

1.5.1 Optical diagnostics

Engines with optical access are invariably used in studies of flow, mixture preparation, and the combustion process. There are basically two types of arrangement for the optical access, namely, quartz windows through the piston which provide a vertical view of the combustion chamber, and full or partial annular quartz windows which allow a side view

of the combustion chamber. A typical arrangement for an optical engine is illustrated in Figure 1.9.

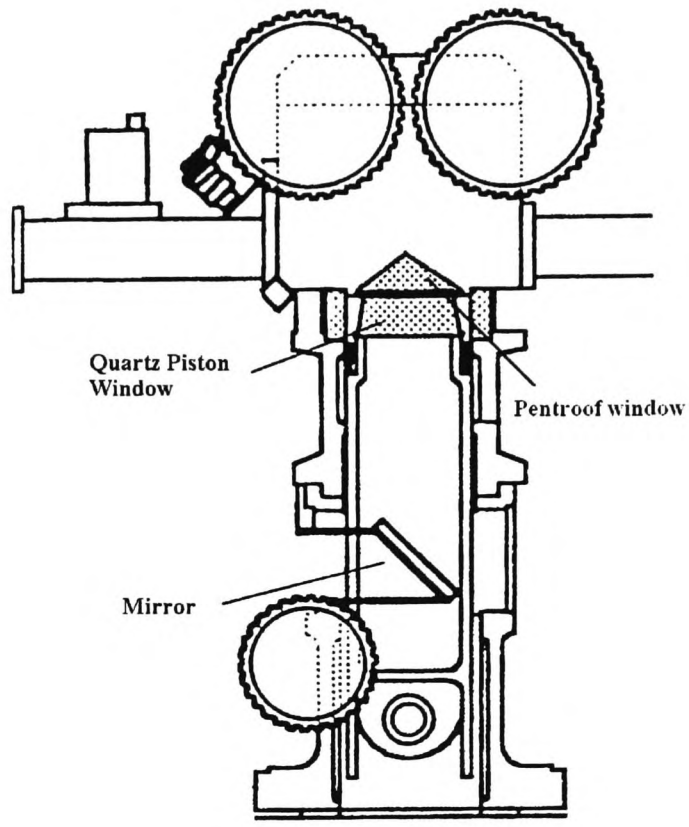


Figure 1.9: Typical arrangement for optical access through the piston, and an annulus below the cylinder head

A slotted tubular extended piston, which supports the optical window in its piston crown, is mounted on the conventional piston crown. The piston slots enable an elliptical mirror to be fixed at a certain angle, so that an image of the combustion chamber can be viewed through the opposite slot. The piston crown window is typically 75% of the bore diameter. The piston rings in the combustion chamber are often based on graphite and fibre-reinforced polymers, so as to avoid any lubrication that can foul the optical path. Such rings can run on cast-iron cylinder liners or indeed an annular window. The lower piston will have a conventional oil control ring to prevent oil contamination of the mirror and piston window.

The windows are usually made of fused silica (also referred to as quartz) due to its high strength and low coefficient of thermal expansion, which leads to a good resistance to thermal shock. However, quartz is very fragile, so great care needs to be taken in the

installation and servicing of the windows. The choice of window material is also governed by the radiation wavelengths that are of interest, *e.g.* sapphire (Al_2O_3) may be used when ultraviolet radiation is being studied, while germanium or silicon may be used when infrared radiation needs to be transmitted.

Several problems arise during research with optical engines. When the optical access is the full length of the stroke, operation is usually confined to non-firing or skip-firing tests. An annular window as well as a concave piston window would both cause substantial optical distortion, but this can be calibrated and corrected for in a particular test. The alternative to correcting for optical distortion is to eliminate it, either by the use of plane windows set into the bore, or by means of an engine that has a rectangular piston.

1.5.2 Combustion imaging and analysis

A variety of visualization techniques have been widely reported for combustion imaging and analysis. These techniques, as well as image processing, make it feasible to investigate mixture formation, combustion characterization, spray and flame propagation, and in-cylinder soot concentration and temperature in DISI engines in either a qualitative or quantitative way.

One such technique is endoscopy, which has historically been performed on Diesel engines for spray and combustion investigations; whilst nowadays it is often used in combination with laser techniques to minimize engine modifications or to gain access to complex geometries. With an intensified CCD camera and endoscopic access, Schanzlin *et al.* [52] observed the flame luminosity in a research DISI engine. Schanzlin also carried out parallel research with ion current sensing probes and two-colour pyrometry, using the former technique to sense the ionization in the flame front, and using the latter method to extract information on soot formation and concentration as well as estimation of the flame temperature. Extensive comparison between homogeneous and stratified operation was made

using results from the three techniques, and this showed that, in the stratified mode there was highly intensive luminosity originating from soot produced in diffusion flames or in premixed flames with rich reaction zones, whilst in the homogeneous mode cycle-by-cycle variations were least noticeable.

Simonini *et al.* [57] employed colour-ratio pyrometry and a CCD camera imaging system to study soot temperature and concentration measurements in a single-cylinder optical version of the port-injected Rover K16 SI engine. The colour-ratio method is an evolution of the widely used two-colour pyrometry method and has the following advantages: 1) only one CCD colour camera is needed; 2) once the spectral response of the camera is known, no further calibration is required; 3) no absolute calibration of light intensity is needed as only the ratios between the colour components are used, making the method easy to adapt to different situations with different flame brightnesses.

A similar pyrometry method and CCD camera setup, and a high speed cine camera was adopted by Wyszynski *et al.* [73] on the single-cylinder Mitsubishi GDI SI engine for combustion imaging and analysis. Discussions were detailed about soot temperature and formation in a single engine cycle as well as cycle-by-cycle variations.

1.5.3 Laser diagnostics

With the increasing capability and reliability of lasers, laser-based diagnostic techniques for research on IC engines, particularly combustion, have been extensively explored. Eckbreth [16] gave a very detailed survey and description of various laser diagnostic techniques that are potentially suitable for temperature and species concentration measurements in combustion, including: Rayleigh scattering for total density determinations (and thus temperature under isobaric conditions); spontaneous Raman scattering and coherent anti-Stokes Raman spectroscopy (CARS) for major species concentrations (*i.e.* $\sim$1% or greater) and temperature measurements; and laser-induced fluorescence spectroscopy

(LIF) and degenerate four-wave mixing (DFWM) for thermometry and detection of reactive intermediates at low levels (ppm). In the infrared, DFWM is also applicable to the detection of major product species.

Grant [21] employed DFWM, CARS and planar LIF to approach combustion from three different viewpoints. From a pollution perspective, DFWM was applied for the detection of nitric oxide in an IC engine. Secondly, spark-ignited hydrogen-air combustion in an ignition cell was studied from a safety point of view, with CARS adopted to characterise the cooling effect of the spark electrodes on the surrounding gas. Finally from an efficiency perspective, self-ignition in an IC engine was investigated using double pulse planar LIF imaging.

There are a myriad of lasers, but for diagnostics the choices can be limited to a few basic types which emit in the visible and ultraviolet (UV), because of the signal strength scaling of the Raman and Rayleigh cross sections, the spectral location of electronic resonances for fluorescence and DFWM, and the peak sensitivity range of most optical detectors (Eckbreth [16]). Pulsed lasers are widely used for light scattering and LIF measurements, because with the detection gated, their high intensities can significantly increase the signal-to-noise ratios. Two of the most important and widely used pulsed lasers are the Nd:YAG (neodymium in yttrium aluminium garnet) and Excimer lasers.

As an economical option, Nd:YAG lasers can output three wavelengths (532nm, 355nm and 266nm) in the visible and UV region by frequency mixing, typically with energies of 1J and repetition rates of 30Hz. Care must be taken when triggering the laser externally, because to obtain optimized performance, the frequency of the flashlamp should be approximately the same as the nominal working frequency, *e.g.* 30Hz. For example, if a 10Hz YAG laser is used in an engine experiment, the available operating speed is restricted to 1200RPM in order to generate one single laser pulse per cycle at a specified crank angle.

By selecting different rare gases (*e.g.* XeCl for 308nm, KrF for 248nm *etc.*) as the lasing medium, an Excimer laser can provide desired wavelengths separately. Typical energy levels and repetition rates are a few hundred mJ and several hundred Hz, respectively. Despite the fact that they tend to be large, heavy, expensive and difficult to maintain and service, Excimer lasers are becoming more popular in PLIF imaging due to their high pulse intensity in the UV region. Unlike YAG lasers, the operating frequency of Excimer lasers is only limited by their maximum repetition rates. Furthermore, they typically have rectangular beams (a feature due to the mechanism to drive them) which make them well suited to provide uniform sheets for PLIF.

1.5.4 Emissions measurement

Exhaust emissions have a major impact on engine development strategy because of legislation, and also because of the insights the measurements provide into engine performance. While for the DISI concept, it has been shown, both theoretically and experimentally, that emission levels are relatively low, based on the engine output. However, emissions legislation is focused on vehicle emissions, which means the vehicle-level output rather than the engine-level output is crucial.

The major problems due to emissions encountered in a GDI engine were detailed in Section 1.4.5. The demand to fully understand the mechanisms of emission formation and to seek effective elimination calls for delicate and sophisticated emissions analysis systems, which can provide accurate and instantaneous (fast-response) measurements. As far as a GDI engine is concerned, HC, NO_x and soot particulates are of most interest.

Flame ionisation detection (FID) is widely used to detect hydrocarbons. When hydrocarbons are burned, electrons and positive ions are formed. If the unburned hydrocarbons (sampled from the engine emissions) are burned in air and an electric field, the current flow corresponds very closely to the number of carbon atoms present, hence quantifying

the HC measurement. A fast-response flame ionisation detector (FFID) has a response time of the magnitude of 1ms, since the flame ionisation detector has been miniaturised and can be sited remotely from the analyser, and adjacent to the sample source. This system is capable of making cycle-resolved measurements, including: 1) in-cylinder sampling to look at cycle-by-cycle variations in unburned HC before and after combustion; 2) sampling of the combustion chamber gases in a lean-burn or stratified-charged engine to establish the air/fuel ratio distribution; 3) measurement of the unburned HC in the exhaust port of a gasoline engine.

NO_x (nitrogen oxides) emissions are measured using a chemiluminescence technique, that is, detection of light emissions from reactions such as:

$$NO + O_3 \rightarrow NO_2^* + O_2 \rightarrow NO_2 + O_2 + \text{photon}$$

where NO_2^* is nitrogen dioxide in an activated state.

Ferguson [19] shows that if: a) the reaction chamber is sufficiently large; b) the ozone flow rate is steady and high compared with a steady sample flow; c) the reaction chamber is at a fixed temperature, then the emitted light is proportional to the concentration of nitric oxide in the sample stream. Both NO and NO_2 can exist in the exhaust of an engine. NO_2 can be measured by passing the sample through an oven that converts NO_2 to NO. As with an FFID, a fast-response chemiluminescence NO analyser (with a response of a few ms) can be constructed by miniaturising the reaction chamber and installing it close to the sampling point.

The size and concentration of exhausted soot particulates can be measured by the established principle of electrical mobility classification. The sample from engine exhaust is first passed through an electrical charger which places a prescribed charge on each particle. The charged sample then passes the classifier column where the particles are exposed to a strong electrical field, causing them to drift across the flow to an array of detectors. Due to the difference in aerodynamic drag, different particle sizes are collected in different locations along this array. The signals from the detectors are processed in real time to

give the particle size spectrum and any desired total emission parameters.

1.6 Modelling DISI engines

The aim with modelling internal combustion engines is twofold [60]:

- to predict engine performance without having to conduct tests;
- to deduce the performance of parameters that can be difficult to measure in engine tests.

Engine modelling helps to save development costs as well as provide guidance for engine structural and mechanical design. However, engine models rely heavily on experimental data and empirical correlations, for example for predicting processes such as heat transfer, ignition delay and burn rate.

Although sophisticated computational fluid dynamics (CFD) codes have been developed and widely applied to engine simulations in recent years, zero-dimensional models, *i.e.* no fluid flow modelling, remain a useful and yet comparatively straightforward tool for calculations based on engine thermodynamics and combustion chemistry. Ferguson [19] listed cycle calculation models for spark ignition and compression ignition engines in FORTRAN, which were later extensively modified and expanded by Raine *et al.* [48] as Integrated Spark Ignition engine Simulation (ISIS), with the emphasis on the prediction of the generation of NO_x. In particular, ISIS allows for the possibility of non-homogeneous charge and hence the modelling of DISI engines in stratified mode. Other major additional function of ISIS include:

- incorporation of a NO kinetic model
- multi-zone analysis
- options for burn rate and different heat transfer routines

- incomplete combustion
- oxidant other than air
- hydrogen as a fuel

Simulation results from ISIS have been validated against data from experiments on a number of engines [6] [50] with good agreement between the trends in the predicted and experimental results.

1.7 Jaguar Optical DISI Engine

The engine used in this project was a Jaguar single cylinder research engine with a spray-guided DISI combustion system and a Bowditch piston arrangement. Optical access could be achieved in five ways:

- a full quartz cylinder liner
- a short quartz annulus which goes on top of a metal liner
- a similar short stainless steel annulus with four side windows
- a flat quartz window (60mm diameter) fixed in the piston crown and a 45° mirror underneath
- a triangular quartz window embedded in the cylinder head alongside the pentroof area of the combustion chamber

The engine has a bore of 89mm, a stroke of 90mm and a compression ratio of 11.

1.8 Summary

Direct injection spark ignition engines (DISI) have been of great interest to automotive researchers and manufacturers because they possess the potential of improvements to both

fuel economy and emissions reduction. In this concept where throttling is eliminated, the pumping losses and engine response lag in transient operation can be reduced. By injecting fuel late in the compression stroke, stratification of the fuel/air mixture can be achieved resulting in overall lean burning which helps with fuel efficiency. Additionally, injecting fuel early in the induction stroke enhances the charge cooling effect which is beneficial for loosening the restrictions on engine compression ratio design posed by combustion knock.

As the new generation DISI spray-guiding concept, the spray-guided combustion system is distinguished from the air- or wall-guided concepts by how the fuel/air mixture preparation is controlled. Sitting close to the spark plug, the injector squirts fuel towards the spark plug gap so that desirable mixture stratification can form in the proximity of the spark plug and fuel impingement on the piston and chamber walls can be minimised. A spray-guided combustion system requires careful considerations of the design and control of the injector whilst easing the design of the piston crown.

Various engine diagnostic techniques can be applied on DISI engines, most notably optical diagnostics for their non-invasive features. Natural in-cylinder illumination *i.e.* combustion incandescence and chemiluminescence offers the opportunity of using digital cameras for flame imaging as a cost-effective 2-D solution with high temporal and spatial resolution. Further image analysis allows for temperature and soot concentration estimation with colour pyrometry. Laser-based measurements provide even better temporal/spatial resolution and accuracy usually at higher cost and with a more delicate setup.

In another direction, simulation models can be used to predict engine performance and comparisons with experimental data. As the new generation DISI concept is still under development, engine modelling can prove a useful and yet more straightforward tool to give guidance on engine design and optimisation.

Chapter 2

Data Acquisition and Processing

Experts often possess more data than judgment.

— Colin Powell (1937–)

2.1 Introduction

Data acquisition and subsequent processing is the key to extracting information from experimental data and understanding the fundamentals beyond the observations. This chapter discusses the instrumentation system and the data acquisition and processing tools used in this project.

The Jaguar single cylinder optical research GDI engine has been integrated into the engine instrumentation system based on the previous Oxford Rover K4 engine test bed. Facilities for both GDI and PFI fuelling are included and standard measurements such as cylinder pressure and flow rates measurements can be carried out.

The hardware and software for the engine control and data acquisition system are described. For the engine to run effectively, timing signals *e.g.* injection and ignition must

be controlled promptly and precisely. This is undertaken with an engine timing control software developed in LabVIEW for this engine system. The National Instruments data acquisition system has been used extensively and is capable of handling both high-speed and low-speed signals from the many channels in use. The integration of the timing control and data acquisition systems allows for customised experiments to be conducted.

In particular, three camera systems are introduced for the optical research including planar laser induce fluorescence, soot pyrometry and high-speed combustion imaging. The working principles as well as the details of the camera specs and operating procedures are given in Section 2.4.

Emphasis is also drawn to the data/image processing tools developed in MATLAB. Combustion analysis with experimental data is performed by a MATLAB package CoBRA that implements the Rassweiler and Withrow method for burn rates analysis, estimates the completeness of combustion based on statistics and calculates the heat release from the cylinder pressure data. The original version of CoBRA has also been extended substantially to have functions for diesel engines and 2-stroke combustion systems. The powerful image processing toolbox in MATLAB provides a platform on which a graphical user interface has been created to analyse images of different kinds from the optical work. Examples of image processing are listed at the end of Section 2.4.5 to illustrate the capability of the software.

2.2 Instrumentation system

The engine instrumentation system was based on the previous Oxford Rover K4 engine test facility [72] [74], including a DC dynamometer system with a Control Techniques four-quadrant controller, a Heypac high pressure fuelling system for DISI operations, a PFI system for both vapour and liquid fuels, a temperature measurement and control system, an inlet air flow measurement and control system, an AFR recorder *etc.* A block

diagram of the structure of the engine instrumentation system and data flow is shown in Figure 2.1.

To measure the in-cylinder pressure, a water-cooled Kistler 6041A piezoelectric pressure transducer was installed in the cylinder head. A water-cooled Kistler 4075A10 piezoresistive pressure transducer was installed in the cylinder barrel to correct for thermal drift and set a datum for the cylinder transducer (see details in Section 2.3.2). Both transducers and their corresponding amplifiers were calibrated using a dead weight tester, showing good linearity. The gauge inlet manifold pressure was measured by a DRUCK PDCR A20-0800 transducer.

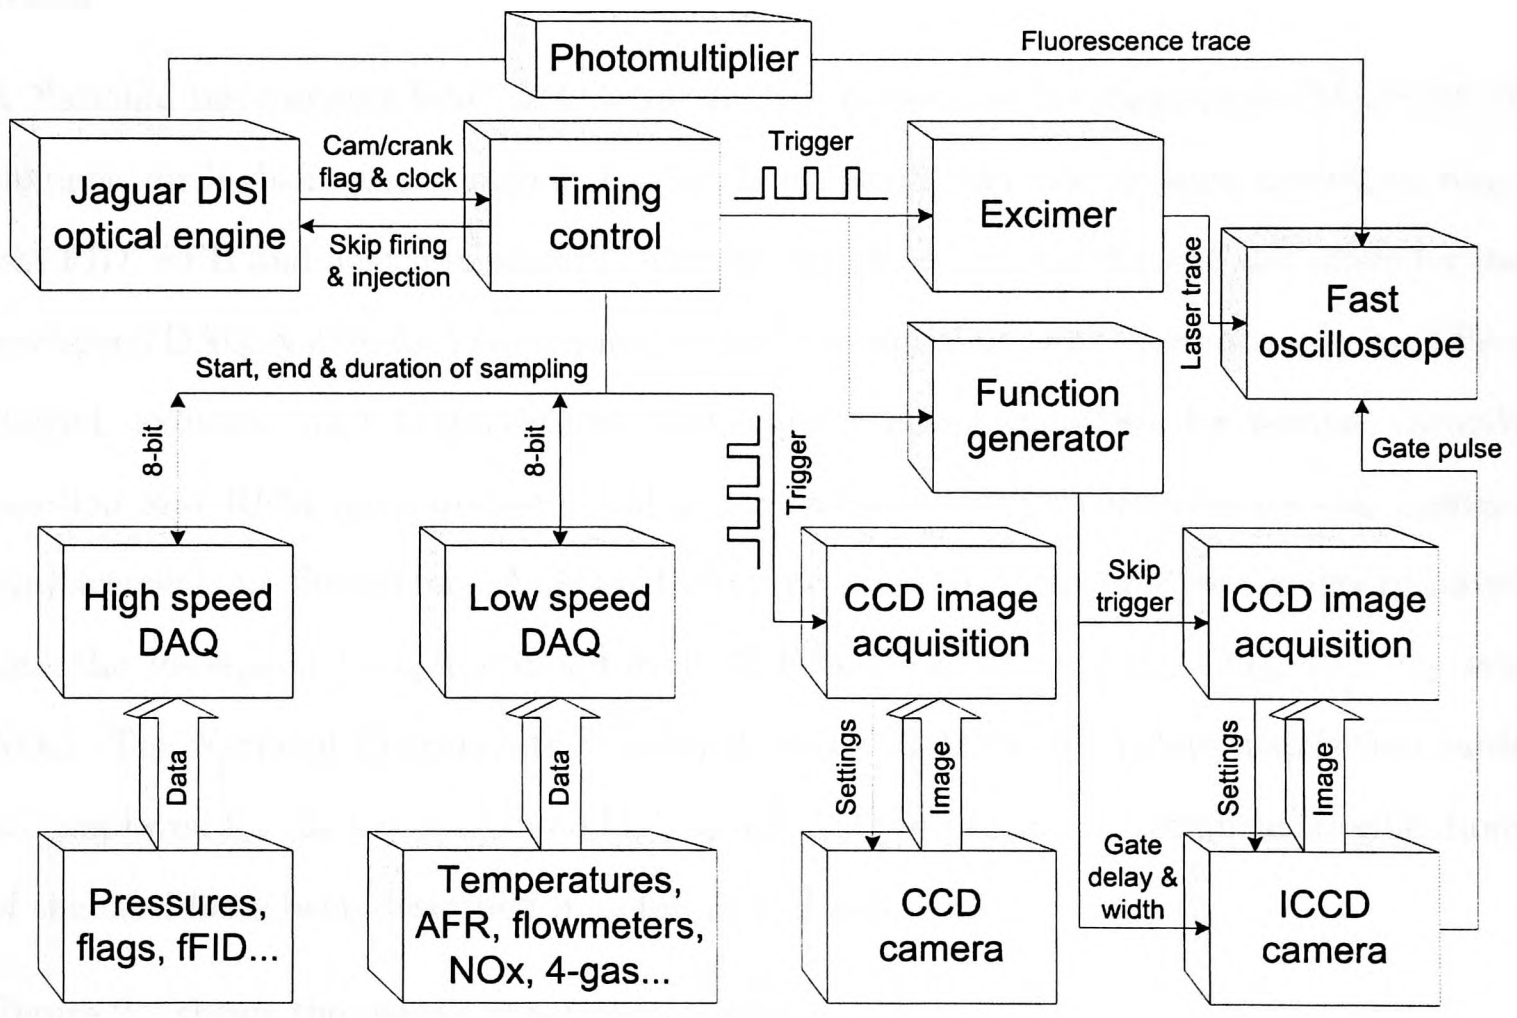


Figure 2.1: Engine instrumentation system and structure of data flow

In order to control the engine experiments and log the ensuing data, three software modules were developed, namely, an engine timing control software (ETCS) [75], a high-speed data acquisition (DAQ) system and a low-speed DAQ system.

Trigger (flag) and clock signals from the engine shaft encoder (crankshaft or camshaft) are fed into the ETCS program, by counting the number of clock signals using a National Instruments counter card, parameters such as pulse delay and duration can be controlled to precision as good as one clock signal interval(*e.g.* 1° crank angle if the encoder clock is set to 360 pulses/rev). The program sends out TTL pulses (injection, ignition *etc.*) to an interface box according to the settings on the front panel. The box, with an opto-isolator protection mechanism, then outputs appropriate signals to the ignition and injection drives. Meanwhile, the ETCS program communicates with the high-speed and low-speed DAQ systems via RS-232 serial links to synchronise the start and end of the data-logging events.

A National Instruments BNC-2040 breakout box is used for the high-speed DAQ with 16 channels for high-frequency signals *e.g.* cylinder, barrel and inlet plenum pressures, flags, fast FID, AFR and heat flux sensors. Another interface box was devised and made for the low-speed DAQ system for low-frequency signals comprising eight thermocouple amplifiers (barrel, exhaust, inlet temperatures *etc.*), three voltage amplifiers for torque, throttle position and RPM measurements and a frequency-to-voltage converter for the positive displacement air flowmeter. A Sykes-Pickavant portable 5-gas analyser is incorporated into the low-speed DAQ for monitoring AFR and emissions (CO, CO_2, HC, O_2 and NO_x). The National Instruments PCI-6024E and PCI-MIO-16E-1 data acquisition cards are employed for the low-speed and high-speed DAQ respectively. Technical specifications of the card have been described in detail in [32] and [72].

Figure 2.2 shows the overall experimental setup.

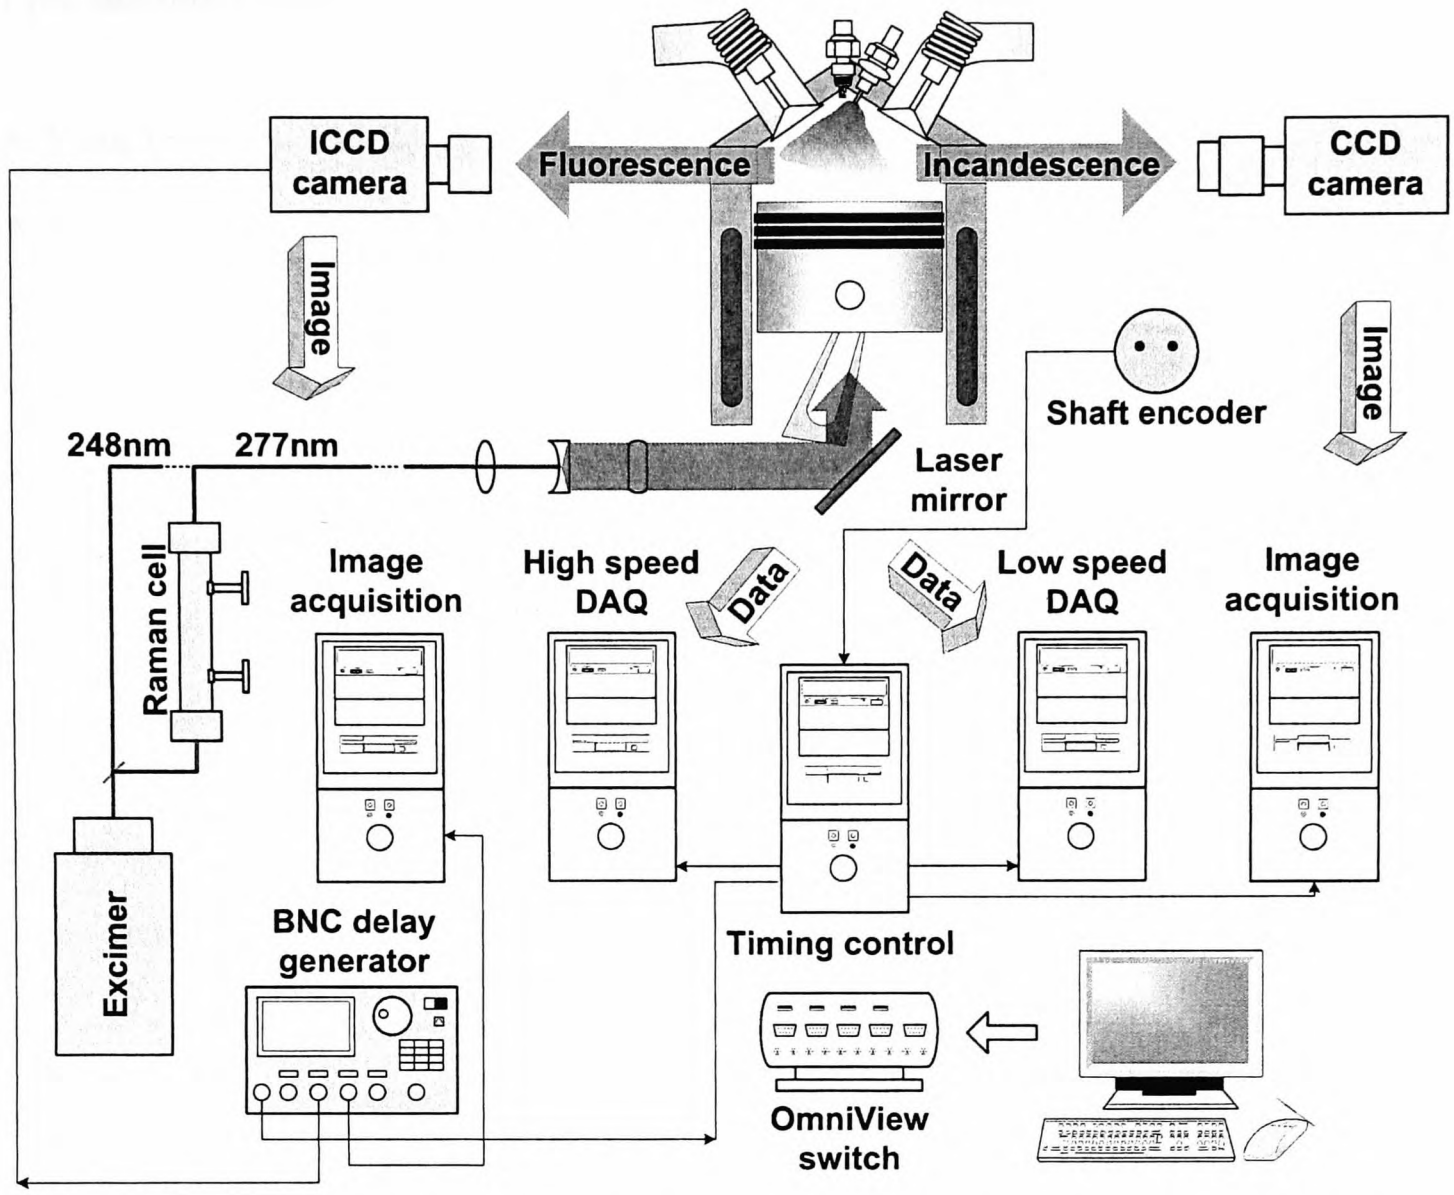


Figure 2.2: Experimental layout on the Jaguar DISI engine including measuring devices and timing control and DAQ system

2.3 Engine data analysis

2.3.1 Introduction of CoBRA

CoBRA (**Co**mbustion **B**urn **R**ate **A**nalysis) is a computer program initially written in FORTRAN to perform combustion analysis of cylinder pressure data from an internal combustion engine. Later on, CoBRA was rewritten in MATLAB, mainly to take advantage of the vast mathematical function library available, the convenience of matrix arithmetic and the graphical user interface (GUI), and also to add particular routine for optical engine study. The present MATLAB code version covers the analysis of combus-

tion parameters below:

- Peak pressure and the crank angle at which peak pressure occurs (p_{max}, θ_{pmax})
- Rate of pressure rise ($\mathrm{d}p/\mathrm{d}\theta$)
- Net and gross IMEP
- Mass fraction burned (MFB)
- Normalised pressure rise due to combustion (ψ)
- Heat release analysis (net heat release Q_{n} and release rate $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$)
- Completeness of combustion (χ)
- CCD image grabbing sequence

Two versions of CoBRA are available for the time being, the text version and the GUI version, with an almost identical kernel but slightly different functions. In the latest text version of CoBRA, version 6.7, a list of data files and their operating points can be input via a spreadsheet file instead of from individual MATLAB scripts saved as text files. This saves time when processing a large number of data files, and helps to keep track of operating points. Also, two MATLAB routines have been written which take in the same input spreadsheet file, and extract the required information to output spreadsheet files from the .mat files created during the analysis.

The GUI package of CoBRA, CoBRAxp (see Figure 2.5), provides a variety of options in case a Diesel engine is investigated, as well as a plug-in graphical output display. Both versions can deal with an internal software trigger and 2 or 4-stroke engine data.

2.3.2 Pressure data analysis

The data file logged from the data acquisition system typically comprises 2–4 columns representing corresponding channels. At least 2 channels, the cylinder pressure and the flag (1/rev crank flag or 1/cycle cam flag), are needed for data processing.

In order to analyse the pressure signal, it is necessary to reference it to a pressure datum to know its absolute value, because of the effect of thermal drift on the pressure transducer and the integration of the transducer signal by the charge amplifier. The datum can be fixed by assuming that the cylinder pressure is equal to atmospheric pressure (or another user-specified value) at either exhaust top dead centre (ETDC) or inlet bottom dead centre (IBDC). By the end of the exhaust stroke the pressures in the cylinder and in the exhaust manifold have equalised, and the pressure in the exhaust manifold can be measured. Similarly, the cylinder pressure at IBDC can be assumed equal to the pressure in the inlet manifold. In these two cases if the engine is turbocharged, the gauge pressure (relative to atmospheric pressure) of the turbine inlet or the compressor outlet, respectively, needs to be added to the atmospheric pressure as a datum pressure, *i.e.* the exhaust or inlet manifold pressure.

To avoid the approximation in these approaches, data from a barrel pressure transducer can be used. The barrel transducer is in a position such that it measures in-cylinder pressure when the piston is near BDC but atmospheric pressure when the piston rings are above it, near TDC. Hence, knowing atmospheric pressure allows the datum to be set for the barrel transducer around TDC which ideally provides an identical pressure with the cylinder pressure transducer around BDC. The datum for the cylinder transducer can then be set. In practice, two ranges of crank angles, around BDC and TDC respectively, can be specified and the averaged pressure data are used in the correction.

Alternatively, a range of crank angles can be specified (*e.g.* $\pm 2°$ around ETDC). The sampled pressures within this range are then averaged. The pressure reference is set by assuming this mean pressure to be equal to a datum pressure. This option is particularly useful to deal with noisy data such as data from extremely fast sampling.

Occasionally the cylinder transducer is mounted somewhere else in the cylinder rather than the default site, *e.g.* near or even in the spark plug. In such cases, the acquired pressure data are bound to be different from that by a default transducer. For example,

when a small pressure transducer is put in the spark plug, it has less effective cooling and poorer sensitivity and hence is more susceptible to effects such as thermal drift than a normal-size pressure transducer located in the cylinder barrel. Therefore, a dynamic calibration is required and can be conducted using the following steps:

1. Mount both pressure transducers and obtain a set of data from both.
2. Take out the mean difference between the two sets of pressure data and save it as a profile, noting its tendency to vary with operating conditions.
3. Whenever the pressure transducer is moved or changed, the dynamic correction profile is loaded and added to the pressure data as compensation.

In order to know the combustion chamber volume at each pressure sample point, the flag is used to find the crank angle at each sample point. With knowledge of the engine swept volume and the nominal compression ratio, the combustion chamber volume can be easily calculated from the crank angle. Then the calculation of $p_{\max}$, θ_{pmax}, $\mathrm{d}p/\mathrm{d}\theta$ and IMEP is straightforward.

2.3.3 Mass fraction burned (MFB)

The technique used here for computing the mass fraction burned (MFB) is the method first developed by G. M. Rassweiler and L. Withrow in 1938 (Rassweiler and Withrow [51]). The combustion process was approximated by considering a series of small crank angle intervals ($\Delta\theta$), during each of which the change in pressure (Δp) is the sum of two parts: the pressure rise due to piston motion (Δp_{v}) and the pressure rise due to combustion (Δp_{c}).

The concept is shown graphically in Figure 2.3. First, the charge is compressed polytropically[1] by the piston, then a small amount of charge burns causing a pressure rise.

[1] A process in which there is heat transfer, with pV^n = constant, known as a polytropic process.

In the next crank angle interval, the pressure again changes polytropically and another bit of charge burns, and this continues throughout the combustion period. Rassweiler and Withrow illustrated the method by using a 'spread sheet' calculation procedure as an example, but this was not expounded mathematically.

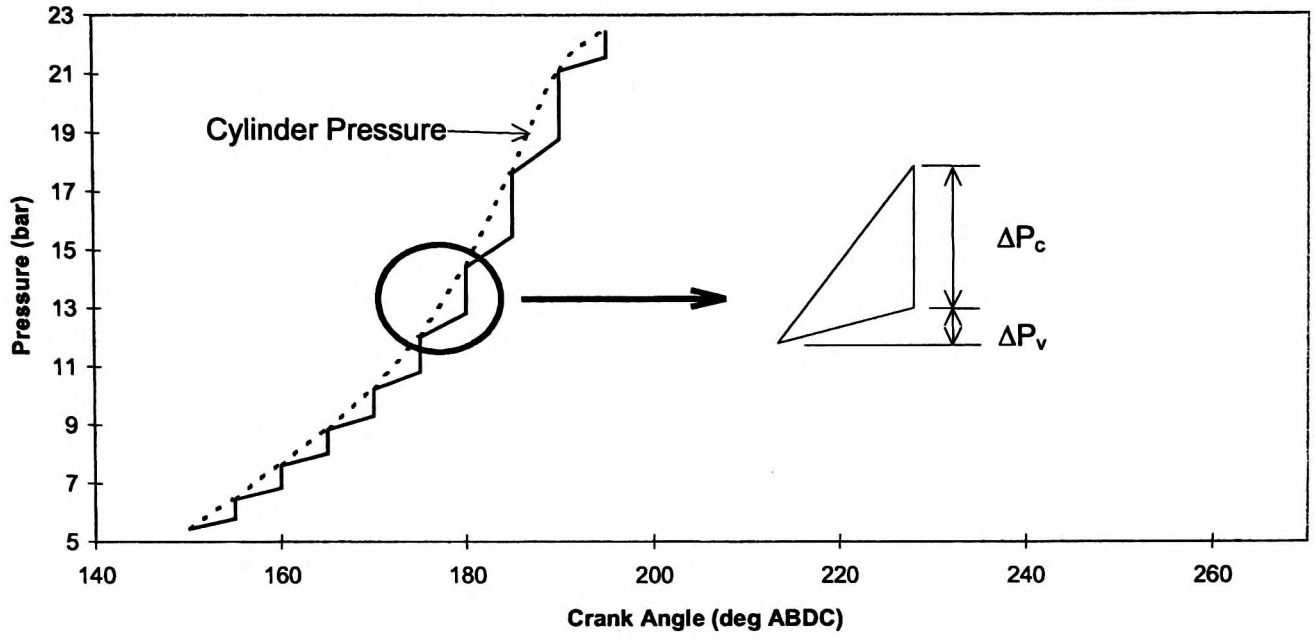


Figure 2.3: Measured pressure change as a sum of piston motion and combustion

A number of texts, *e.g.* Stone [60] gives a mathematical derivation of the method which will henceforth be referred to as the 'standard' method. Since the total pressure rise over one crank interval is the sum of the pressure rise due to piston motion (Δp_{v}) and the pressure rise due to combustion (Δp_{c}),

$$\Delta p = \Delta p_{\mathrm{c}} + \Delta p_{\mathrm{v}} \tag{2.1}$$

As the crank angle θ_i increments to θ_{i+1}, the volume changes from V_i to V_{i+1} and the pressure from p_i to p_{i+1}. It is assumed that the pressure change due to piston motion at time step i, ($\Delta p_{\mathrm{v,i}}$), is approximated by a polytropic process with an index n, hence:

$$p_i V_i^n = (p_i + \Delta p_{\mathrm{v},i}) V_{i+1}^n \tag{2.2}$$

$$\Rightarrow \Delta p_{\mathrm{v},i} = p_i \left[\left(\frac{V_i}{V_{i+1}} \right)^n - 1 \right] \tag{2.3}$$

Substituting Eq.(2.3) into Eq.(2.1) produces

$$p_{i+1} - p_i = \Delta p_{\mathrm{c},i} + p_i \left[\left(\frac{V_i}{V_{i+1}} \right)^n - 1 \right] \tag{2.4}$$

from which the pressure rise due to combustion can be calculated

$$\Delta p_{\mathrm{c},i} = p_{i+1} - p_i \left(\frac{V_i}{V_{i+1}} \right)^n \tag{2.5}$$

Since combustion is not taking place at constant volume, Eq.(2.5) must be normalised to some reference volume, so that the resulting parameter is proportional to the heat released by the charge burned during each time step. This reference volume is arbitrary (Rassweiler and Withrow used the volume at ignition); here the cylinder volume at TDC was used. This normalised pressure change is denoted by $(\Delta p_{\mathrm{c}}^*)$

$$\Delta p_{\mathrm{c}}^* = \Delta p_{\mathrm{c},i} \left(\frac{V_i}{V_{\mathrm{TDC}}} \right) \tag{2.6}$$

The point at which $(\Delta p_{\mathrm{c}}^*)$ is no longer positive indicates the end-of-combustion (EOC), because then no more heat is being released by charge burn. Assuming that the cumulative normalised pressure rise due to combustion is proportional to the mass fraction burned x, then for N total increments

$$x = \frac{\sum_0^i \Delta p_{\mathrm{c}}^*}{\sum_0^N \Delta p_{\mathrm{c}}^*} \quad (\textit{standard method}) \tag{2.7}$$

The total normalised pressure change due to combustion is known as $\Psi = \sum_0^N \Delta p_{\mathrm{c}}^*$.

Several key observations were pointed out in the original work regarding this method:

1. Since the piston motion and combustion are considered to progress in a stepwise fashion, small step sizes should be used to minimise errors.
2. The effects of heat transfer are accounted for only in the choice of the polytropic index n. The accuracy of the index thus influences the accuracy of the calculations.
3. Blowby will affect the results, by making the compression and expansion process deviate from being polytropic.
4. In order to normalise the pressure rises due to combustion, it was assumed that the pressure rise was inversely proportional to total volume.
5. In order to calculate the pressure rise due to combustion, the polytropic indices for the burned and unburned regions were assumed to be equal.

Further assumptions are also inherent in this method, Heywood [24], Stone [60]:

1. it assumes that the referenced pressure rise is proportional to the mass of fuel burned rather than the amount of chemical energy released, and
2. there is no allowance for heat transfer, dissociation, or change in composition beyond the fact that the polytropic index differs from the ratio of specific heats

Despite these limitations, comparisons of mass fraction burned calculated using the Rassweiler and Withrow method with the results of other more complex thermodynamic models show remarkably close agreement, Brunt and Emtage [9], Stone and Green-Armytage [61].

As mentioned above, the sensitivity of the method to the calculation and selection of the polytropic index is very important. If the compression index is used throughout the calculation, it results in a characteristic fall in the mass fraction burned curve after the end of combustion. This is due to heat transfer and the presence of combustion products resulting in a different polytropic index after the combustion event than before it. An accurate index would produce a curve that rose to 1.0 and remained there through to the end of the cycle. Rassweiler and Withrow evaluated the index during compression

and expansion, and used an appropriately averaged value during combustion. Other researchers, detailed by Ball [6] propose switching from one index to the other at TDC, and also various methods for determining EOC for the purposes of calculating the expansion index.

Calculation of the compression index must also be done with care since the pressure rise at the beginning of the stroke is small and discretisation errors (ADC steps) from the data acquisition system are more significant, Stone [60]. This can be avoided by ignoring the initial part of compression, for example, by calculating the index from halfway between IVC and ignition to ignition. With regard to the concerns about the calculation and selection of the polytropic index, the calculation procedure adopts the following definitions and methods:

1. The compression index is evaluated from halfway between IVC and ignition to ignition.
2. EOC is defined as the point at which the cumulative sum of normalised pressure changes due to combustion is a maximum.
3. The compression index is applied throughout the calculation. The MFB results after EOC do not remain at unity as they should, because of heat transfer and leakage as well as the change in composition and temperature.

2.3.4 Completeness of combustion

The completeness of combustion χ is defined as the fraction of the inducted fuel which is usefully burned during the cycle (Sztenderowicz and Heywood, [66] [65]). As mentioned above, the term $\Psi = \sum_{0}^{N} \Delta p_c^*$, for the ideal conditions which are the basis of the Rassweiler and Withrow analysis, is expected to yield the adiabatic combustion pressure rise at the reference volume. If this term is evaluated for any combustion cycle and compared with some 'maximum' of Ψ, referred as Ψ_m, which represents 'complete combustion', then an

estimation of the relative completeness of combustion can be obtained by [7]:

$$\chi = \frac{\Psi}{\Psi_{\mathrm{m}}} \tag{2.8}$$

The frequency distribution of Ψ departs from a normal distribution, so the Kaplan-Meier approach, Cox *et al.* [11] is adopted to estimate the empirical cumulative distribution function (CDF). According to this approach, the calculated data array of Ψ is firstly sorted in ascending order, forming a new array Ψ'. Then the interval of cumulative distribution, [0 100%], is divided into several subdivisions by the total number of cycles N, for instance, [50%–55%], [95%–100%] *etc.* when $N = 20$. The corresponding values of the cumulative distribution S_n (discrete) to Ψ' are given by:

$$S_n = 1 - \prod_{j=0}^{n-1} \left(1 - \frac{1}{N-j}\right) = \frac{n}{N} \tag{2.9}$$

where n is the number of subdivisions involved.

For given results of Ψ, interpolation is used to obtain the consecutive CDF. A significant point at which S_n=97.5% is used to define Ψ_{m}. This is considered to be a satisfactory measurement of complete combustion, Davis *et al.* [12], that is, 97.5% of the Ψ data calculated from the engine data are below Ψ_{m}. For those 2.5% above Ψ_{m}, truncation is made such that they are all set to Ψ_{m}, in order to treat such cycles as 'complete' ones.

2.3.5 Heat release analysis

A heat release analysis is normally applied to Diesel engines, but can also be used on spark ignition engines. The heat release analyses compute how much heat would have to be added to the cylinder contents, in order to produce the observed pressure variations. It is usually assumed that the products and reactants are fully mixed. Applying the 1st Law of Thermodynamics to a control volume in which there is no mass transfer gives the

heat released by combustion (δQ_{hr}), see for example Heywood [24]:

$$\delta Q_{\mathrm{hr}} = \mathrm{d}U + \delta W + \delta Q_{\mathrm{ht}} \tag{2.10}$$

where δQ_{ht}= the heat transfer with the chamber walls.

Another implication of Eq.(2.10) is that there is no allowance for differences in the properties of the reactants and products and that there is a uniform temperature. Each of the terms in Eq.(2.10) has to be evaluated:

$$\mathrm{d}U = mc_{\mathrm{v}}\mathrm{d}T \tag{2.11}$$

From the equation of state ($pV = mRT$)

$$m\mathrm{d}T = (p\mathrm{d}V + V\mathrm{d}p)/R \tag{2.12}$$

Eq.(2.11) together with Eq.(2.12) gives

$$\mathrm{d}U = c_{\mathrm{v}}\,(p\mathrm{d}V + V\mathrm{d}p)/R \tag{2.13}$$

Substituting Eq.(2.13) into Eq.(2.10) and noting that $\delta W = p\mathrm{d}V$ gives on an incremental angle basis:

$$\frac{\mathrm{d}Q_{\mathrm{hr}}}{\mathrm{d}\theta} = \frac{c_{\mathrm{v}}}{R}\left(p\frac{\mathrm{d}V}{\mathrm{d}\theta} + V\frac{\mathrm{d}p}{\mathrm{d}\theta}\right) + p\frac{\mathrm{d}V}{\mathrm{d}\theta} + \frac{\mathrm{d}Q_{\mathrm{ht}}}{\mathrm{d}\theta} \tag{2.14}$$

The $\mathrm{d}V/\mathrm{d}\theta$ term is defined from the geometry of the engine and the term $\mathrm{d}p/\mathrm{d}\theta$ can be obtained from the engine data. If semi-perfect gas behaviour is assumed (such that $c_{\mathrm{p}}/c_{\mathrm{v}} = \gamma$ and $R = c_{\mathrm{p}} - c_{\mathrm{v}}$), then Eq.(2.14) can be written as

$$\frac{\mathrm{d}Q_{\mathrm{n}}}{\mathrm{d}\theta} = \frac{\mathrm{d}Q_{\mathrm{hr}}}{\mathrm{d}\theta} - \frac{\mathrm{d}Q_{\mathrm{ht}}}{\mathrm{d}\theta} = \frac{\gamma}{\gamma - 1}p\frac{\mathrm{d}V}{\mathrm{d}\theta} + \frac{1}{\gamma - 1}V\frac{\mathrm{d}p}{\mathrm{d}\theta} \tag{2.15}$$

where $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$= the net heat release rate. Then the cumulative net heat release Q_{n} is an

integration of the results of the net heat release rate.

To be more precise, the heat capacity c_v should be referred to as a function of temperature. Therefore Eq.(2.13) becomes:

$$dU = c_v\,(p dV + V dp)/R + pV dc_v/R \tag{2.16}$$

And hence Eq.(2.15) evolves to the following:

$$\frac{dQ_n}{d\theta} = \frac{dQ_{hr}}{d\theta} - \frac{dQ_{ht}}{d\theta} = \frac{\gamma}{\gamma-1} p \frac{dV}{d\theta} + \frac{1}{\gamma-1} V \frac{dp}{d\theta} + \frac{pV}{R}\frac{dc_v}{d\theta} \tag{2.17}$$

c_v is given by an empirical approximation developed by AVL (Kahn [33]):

$$c_v = 700 + T(0.155 + A) \tag{2.18}$$

where T is the temperature at the given increment and A is given by:

$A = 0.1$ for SI engines

$A = \frac{\sum Q \cdot 56}{V_S p_{in}}$ for CI engines

where $\sum Q$ = released fuel heat (kJ)

V_S = swept volume (cm^3)

p_{in} = inlet manifold pressure

The same definition as it remains, γ changes with temperature in this case.

Since the mass is assumed to be constant, calculation of the net heat release is only valid between IVC and EVO.

Figure 2.4 shows an example of heat release analysis. The negative net heat release rate implies that there is heat transfer to the cylinder surfaces. Ignition should be close to the minimum net heat release rate because once combustion commences, the heat release rate will add to the negative heat transfer. Near the end of combustion, the heat transfer

to the combustion chamber becomes greater than the heat release from the combustion process. Thus, the net heat release rate becomes negative, in other words the cumulative heat release starts to fall, and this can be used to estimate the end of combustion.

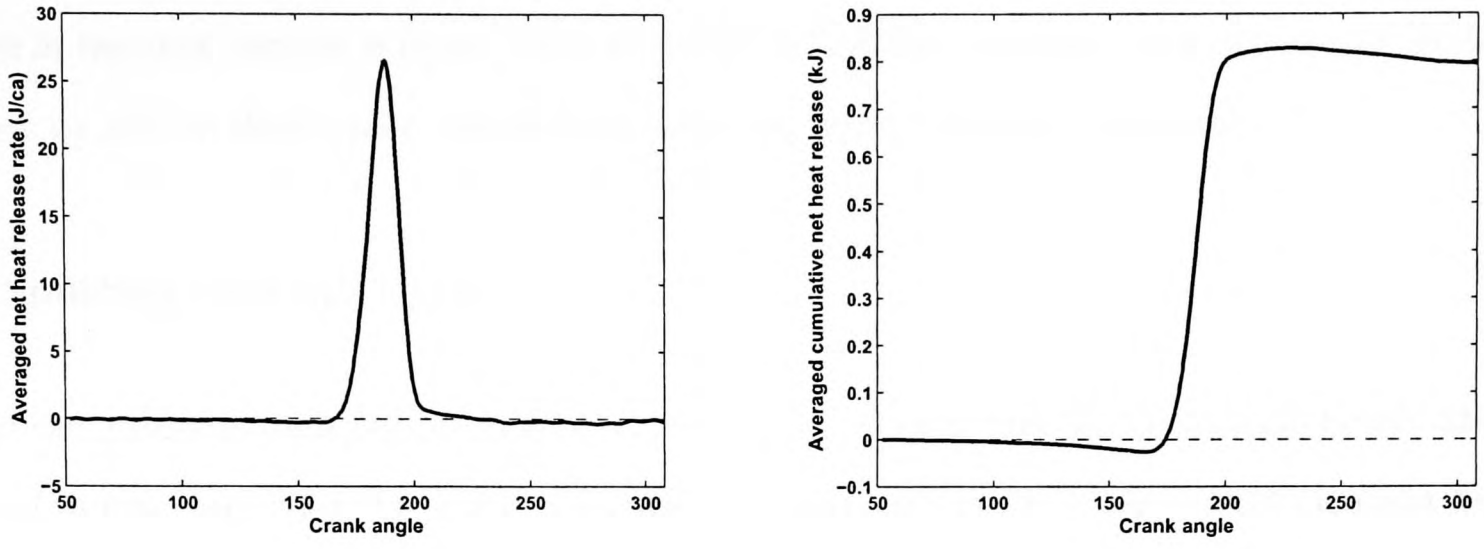


Figure 2.4: An example of heat release analysis — the net heat release rate (left) and the cumulative net heat release (right)

2.3.6 Implementing CoBRA in MATLAB

This section details the implementation of the CoBRA code, including how to process data from the data file, how to fulfill the functions, and how to detect and deal with errors. All the calculation is done within a single .mat file, stepping through as follows:

1) Reading in the data file

To read in the data file, CoBRA assigns each column of the data file to a variable, *e.g.* a two-channel data file results in two variables, cylinder pressure data and flag data respectively. Additional channels can be either a barrel pressure transducer or a CCD image grabbing sequence VINIT, or both.

2) Filtering pressure data

A digital filter is used to filter the noise in the pressure data. An order n lowpass digital Butterworth filter with normalized cutoff frequency ω_n is selected by default. The filter type in the text version is fixed, while the GUI of CoBRA provides more options for filter types as well as their order, cutoff frequency and other relevant parameters.

3) Splitting data into cycles

The raw pressure data fed in CoBRA is in a single column, and all the sample points are listed in time sequence. In order to analyse the data on a cycle-by-cycle basis, it must be split into cyclic columns. This is done by firstly picking the flag pulse, which represents a known crank angle during one revolution, *e.g.* BDC. The flag pulse can be picked up by either the rising edge or the falling edge, depending on the shaft encoder set-up. For example, with the Rover-Oxford K16 engine this is set to the rising edge. Because a certain threshold (1V in the code) of the voltage jump is used for sifting out those sample points which stand for high voltage, it should be noted that extra false pulses may be picked up. Such points are then deleted from the flag array.

If the external trigger[2] is used for sampling, that is, the nominal sampling rate is a multiple of 180/rev. Take 360/rev for example in the following description. Normally in this case, 360 points should be sampled per revolution, but due to interference in the circuits and signal paths, there could be more or less than 360 sampled points. If more than half of the revolutions have such error pulses, there must be a serious problem in the data acquisition system or the shaft encoder, and the calculations (especially the statistics) become meaningless. Thus the code will stop running when detecting such problems in the data. For cases when fewer than half of the revolutions have extra points, the code

[2]The external trigger is from the shaft encoder. For engines in Oxford, sampling rates of 180/rev, 360/rev and 720/rev are available. Therefore, the code needs to be changed when encoders with different sampling rates are used.

truncates them to 360/rev; for cases when fewer than half of the revolutions have missing points, the code pads points to the end of the revolution which are the same as the last point. The idea to use truncation or padding rather than interpolation and extrapolation here is because normally there are only a few revolutions that have more than 1 extra or missing points.

When the internal trigger (clock) from the data acquisition card is used in sampling, the pulse width between each two sample points is fixed, not varying with engine operating conditions such as engine speed. So this can be useful in many ways, *e.g.* detecting the variation of engine speed. In this case, the code calculates the mean value of the number of sample points per revolution, and gives the CoV[3] of the sampling rate as points/rev so that the user can tell whether the internal trigger is reliable and whether the engine operating condition is stable. For revolutions that have more points than the mean rate, the code does interpolation according to the mean rate; for revolutions that have fewer points than the mean rate, the code does extrapolation according to the mean rate.

Once all the revolutions have a consistent sampling rate, the flag is used to determine the points of BDC and TDC. The flag is not necessarily the BDC point, but is worked out from the *flag* specification in the engine spec file or in the GUI front panel. Then the data are split into cyclic columns starting with the BDC of the compression stroke, with each column containing data from one cycle.

4) Implementing functions

The code then implements all the functions one by one, namely, the calculation of the sample point resolved cylinder volume, finding the image grabbing points if the data relates to CCD imaging, the calculation of gross, net IMEP and maximum pressure, the analysis of heat release and the calculation of mass fraction burned and completeness of combustion. The results of MFB calculation are given as sample point resolved MFBs,

[3]CoV = Coefficient of Variation (%) = standard deviation / mean.

and the burn durations of MFB percentage which are of interest. For instance, the 10%, 50% and 80% burn rates are chosen particularly for the Rover-Oxford K4 optical engine, and this can also be specified by the user.

5) Detecting motored and misfired cycles

During the calculation of mass fraction burned, the motored and misfired cycles are detected. When the engine is motored by the dynamometer, the gross IMEP (IMEP during the compression and expansion strokes) is negative. Misfired cycles are defined as partially burned cycles, detected by a threshold parameter named *burn cutoff*. Misfire detection is carried out by determining which cycles have $\Psi < cutoff * \Psi_{max}$, where Ψ_{max} is the maximum total normalised pressure change due to combustion over all cycles. Cycles that have been determined to be misfired can be excluded in the statistical analysis. Therefore more cycles can be added to the statistical calculations by reducing the value of *burn cutoff* unless all the cycles are motored, in which case statistics are not available.

6) Writing output files

A statistical analysis is made across the cycles, with the mean, standard deviation, CoV and LNV[4] of the calculated parameters also in the output data. An example of this data is presented in Table 2.1.

The CoBRA output information is contained on the GUI panel, as well as within the MATLAB variables *summ* and *cyc*, in structure format. Then all the variables can be saved to separate binary data files *.mat for each operating condition. These files can be read from the MATLAB command window.

[4]LNV = Lowest Normalised Value (%) = minimum value / mean.

Mean values	Std. deviations	Coeffs. Of Variation	Lowest Norm. Values
» summ.mean ans = IMEPn: 5.6752 IMEPg: 5.9262 pMax: 35.0428 apMax: 194.6942 psi: 34.5643 aEOC: 203.7961 iComp: 1.2447 CoC: 0.9719	» summ.std ans = IMEPn: 0.0579 IMEPg: 0.0588 pMax: 1.6073 apMax: 1.4215 psi: 0.4686 aEOC: 4.6376 iComp: 0.0209 CoC: 0.0129	» summ.COV ans = IMEPn: 1.0196 IMEPg: 0.9922 pMax: 4.5865 apMax: 0.7301 psi: 1.3558 aEOC: 2.2756 iComp: 1.6801 CoC: 1.3269	» summ.LNV ans = IMEPn: 98.0069 IMEPg: 98.0714 pMax: 86.3494 apMax: 98.6162 psi: 97.4660 aEOC: 96.4199 iComp: 95.5566 CoC: 97.4779

Table 2.1: An example of the statistical output data from CoBRA for a set of data from one operating condition

2.3.7 Graphical user interface of CoBRA

The GUI panel is divided into six sections (see Figure 2.5):

Menus — Unlike the text version, the GUI menu *File* → *Open data file* allows the user to load data files through a pop-up dialogue window, so that the files are not necessarily in the same folder as the code. Once the calculation is complete, the user can save the *mat* file which contains cyclic data, engine specification data and summary data (if available). The menu provides two options; with *File* → *Save MAT file* saving as the default name (same as the input data file) and *File* → *Save as MAT file* saving as any user specified name.

Reference — this section is for reference only. The information displayed here is not calculated directly from the engine data, but input by the user.

Engine Specification — this section lets the user specify the basic configuration data for the engine, including geometry, valve timing, compression ratio and flag. Three engines, the Rover-Oxford K16 MPI, the Rover-Oxford K4 optical GDI, and the Jaguar optical GDI are predefined for the user's convenience. More options — user defined SI engine and user defined Diesel engine (2-stroke or 4-stroke) — are provided for further investigations on any other engines.

Input — in this section, various inputs for the code are presented for the entire calcu-

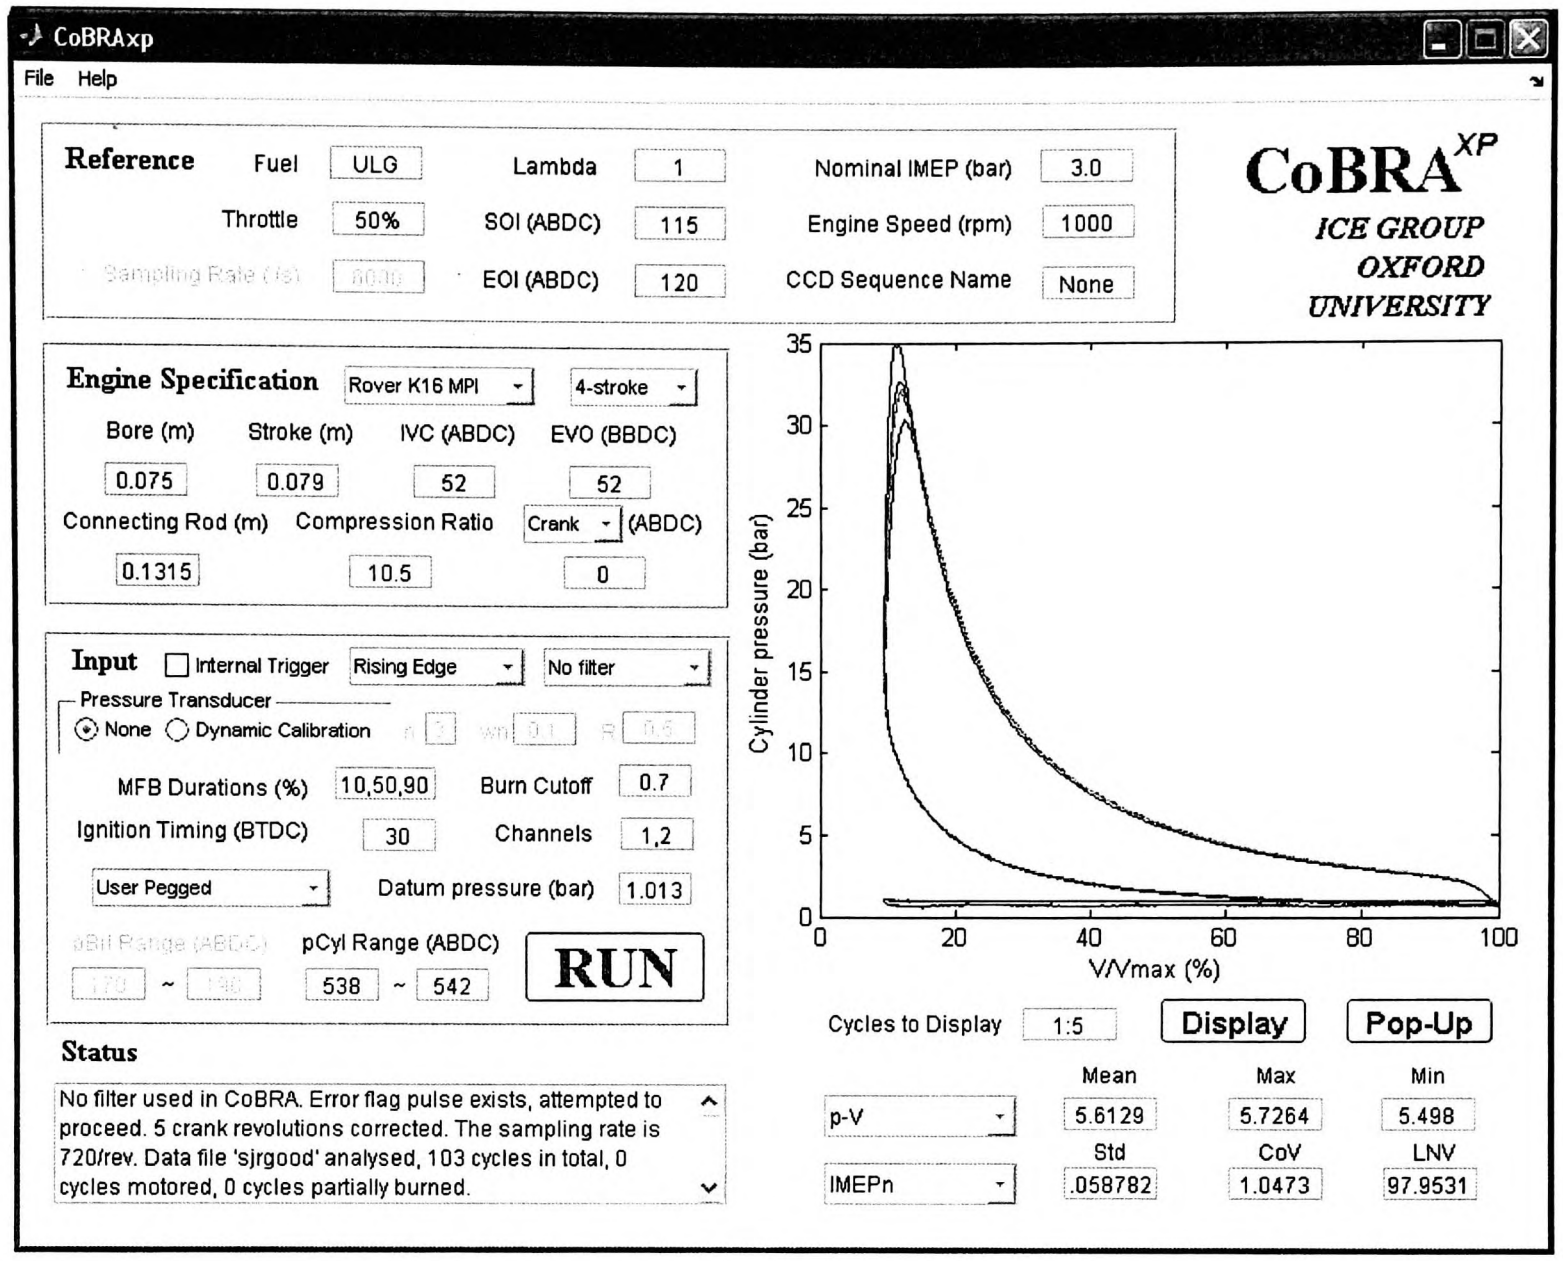


Figure 2.5: Graphical user interface of CoBRA

lation. As detailed in Section 2.3.2, four options are available for the correction of in-cylinder pressure data. When the *ETDC* or *IBDC pressure* method is adopted, the *Datum pressure* textbox is normally set to atmospheric pressure as the engine is naturally aspirated. But if the engine is turbocharged, the *Datum pressure* can be estimated as the sum of the atmospheric pressure and the turbine gauge pressure. If *barrel transducer* is selected, two settings of crank angle range are enabled. The first one, *pBrl Range*, is used for setting the datum of the barrel pressure transducer, whilst the other, *pCyl Range*, is used for setting the datum of the cylinder pressure transducer. If pressure averaging (*User Pegged*) is employed, only *pCyl Range* is used for setting the datum of the cylinder pressure transducer.

If a dynamic calibration for the pressure transducer has been carried out, click on the *Dynamic Calibration* option button, then a dialog will pop up asking which profile needs to be input; by default, the *None* button is activated.

To sample the flag, either the rising edge or the falling edge of the flag signal is used. This is specified during the synchronization of the shaft encoder and needs to be input here for data processing. Owing to the noisy nature of the pressure data, three filters, Butterworth, Chebyshev I and Elliptic, can be adopted with the order, cutoff frequency and value of ripple defined by the user. Alternatively, the user can choose not to apply any filter on the raw pressure (or pre-processed data) data by selecting *No filter*.

A checkbox is available for sampling with an internal trigger, *i.e.* with a fixed sampling rate. In this case, the number of sample points from each engine cycle is not necessarily equal to the clock from the shaft encoder, *e.g.* 360 samples/rev, but is calculated from the given engine data. Also presented here are the inputs of *ignition timing* for IC engines, *injection timing* for Diesel engines, MFB percentages for burn duration calculations, *burn cutoff* for misfire detection and *channels* for specification of the desired columns in the data file.

Output — this section provides a simple interface to display the results. Extensive graphs and diagrams can be displayed, and for some cyclic results such as the p-V diagram, the indices of those cycles which are of interest can be specified. Instead of ticking the checkbox to save and recall, the user can check the statistical summary readily here.

Status — this is a dedicated textbox to show the 'status' of the running code. Major information and errors are displayed here immediately after the code has finished or suspended.

2.4 Image acquisition and processing

This section discusses the experimental equipment and procedure for obtaining combustion and fuel distribution images from the single-cylinder optical engine, as well as the subsequent image processing strategies and approaches developed to extract information from the images.

2.4.1 Charge coupled device (CCD)

A CCD (Charge Coupled Device) consists of MOS (Metal-Oxide-Semiconductor) capacitors arranged in a rectangular array. Each MOS capacitor performs three functions connected with handling electrical charges: 1) photoelectric conversion (photo sensor); 2) accumulation of electrical charges; 3) transmission of electrical charge. The generated charges are fed into the vertical shift registers all at once. The charges are then transferred from the vertical transmission section to the horizontal shift registers successively and finally reach the output amplifier to be read out successively. A simplified model of a charge coupled device is shown in Figure 2.6.

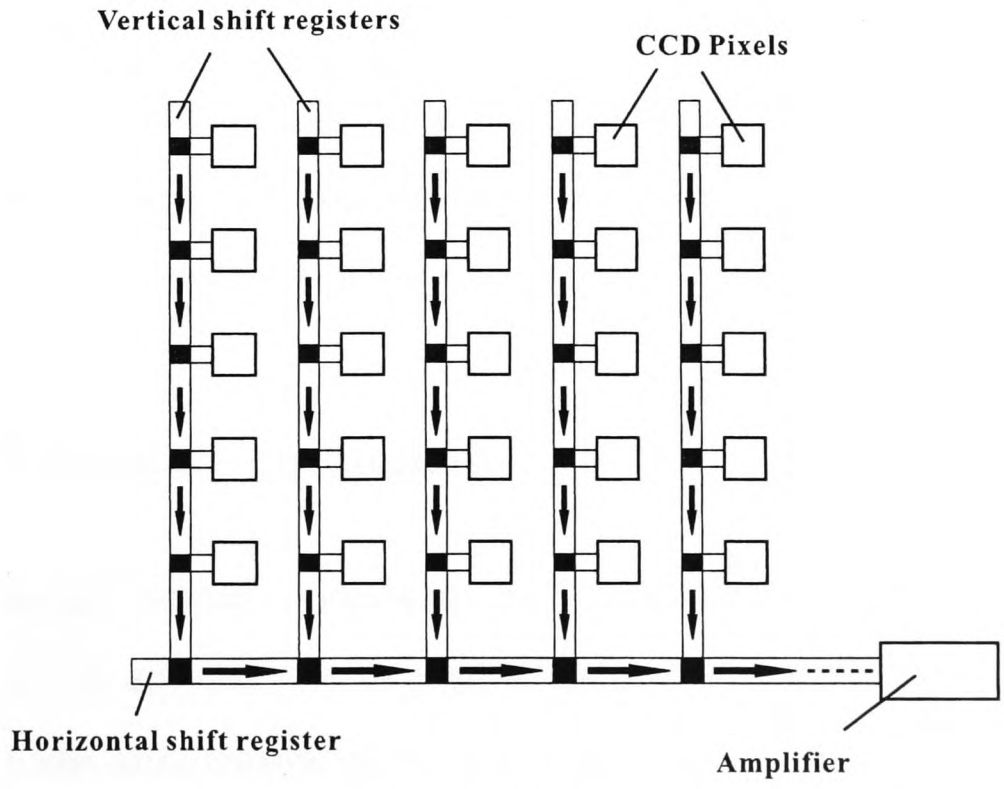


Figure 2.6: Simple model of charge coupled device

The length of time during which pixels are allowed to integrate the incoming light is known as the exposure. The choice of exposure strategy needs to be worked out, based on the estimated signal strength and the consequent noise. Two strategies for exposure setting are commonly used, namely, on chip accumulation and post averaging.

For signal levels low enough to be comparable with read out noise, a longer exposure time for the detector, *i.e.* signal accumulation in the CCD (followed by a single readout scan), improves the signal-to-noise performance approximately linearly with the length of exposure time. The maximum time limit for on-chip accumulation is determined by either the saturation of the CCD by the signal, or the loss of dynamic range due to the dark charge build-up in the pixels. Dark charge is the thermally induced build-up of charge in the CCD which varies between individual CCDs and depends on the temperature of the detector.

When the signals are strong, on-chip accumulation will cause saturation, thus the scan ought to be taken each time the CCD is charged with a short exposure time. Then all the readouts can be averaged to reduce the noise from the CCD. This is what is meant by post averaging.

If the imaging plane is not infinitely thin, the light from different depths is all captured by the camera at the same time, and gets integrated on the CCD chips to give one single image. In this case, information from the third dimension can not be discriminated. This effect is known as line-of-sight integration.

2.4.2 PLIF imaging equipment

To detect fluorescence against a luminous background and intense scattering from the excitation source, it is necessary to implement significant discrimination. This is particularly true when the total fluorescent energy emitted is low. Discrimination is achieved using spectral filters and fast gating of the detector (coupled with sudden excitation) in

the time domain. Intensified CCD cameras are ideal tools for these purposes.

An intensified CCD (ICCD) series detector consists of a CCD optically interfaced to a microchannel plate (MCP) image intensifier through an optical-fibre coupler (see Figure 2.7).

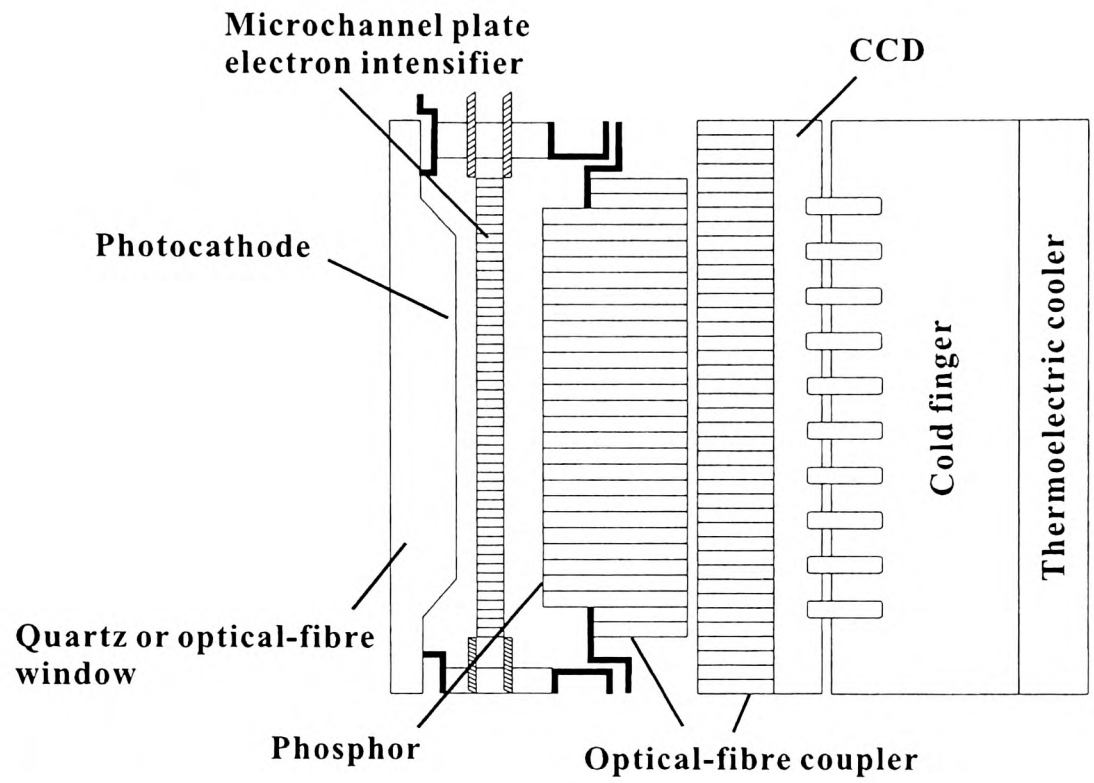


Figure 2.7: Schematic of the intensified charge coupled device (ICCD) detector

The MCP comprises hundreds of individual miniature electron multipliers with an excellent input to output spatial geometric accuracy. When incident light goes through the front window, photons hit the photocathode and photoelectrons come out. These electrons are then accelerated by the high voltage across the MCP electron intensifier. When hit by the electrons, the phosphor screen generates photons, which then pass the optical-fibre coupler and are finally recorded by the CCD. The intensifier gain can be set by adjusting the voltage across the MCP. The combined high gain of the MCP image intensifier and the low readout noise of the CCD array result in a very sensitive detector capable of responding to a single photoelectron.

2.4.3 CCD combustion imaging equipment

Pictures of combustion were taken using a digital system (Figure 2.8), which included the following components:

- Digital CCD camera
- Frame grabber
- Computer and software

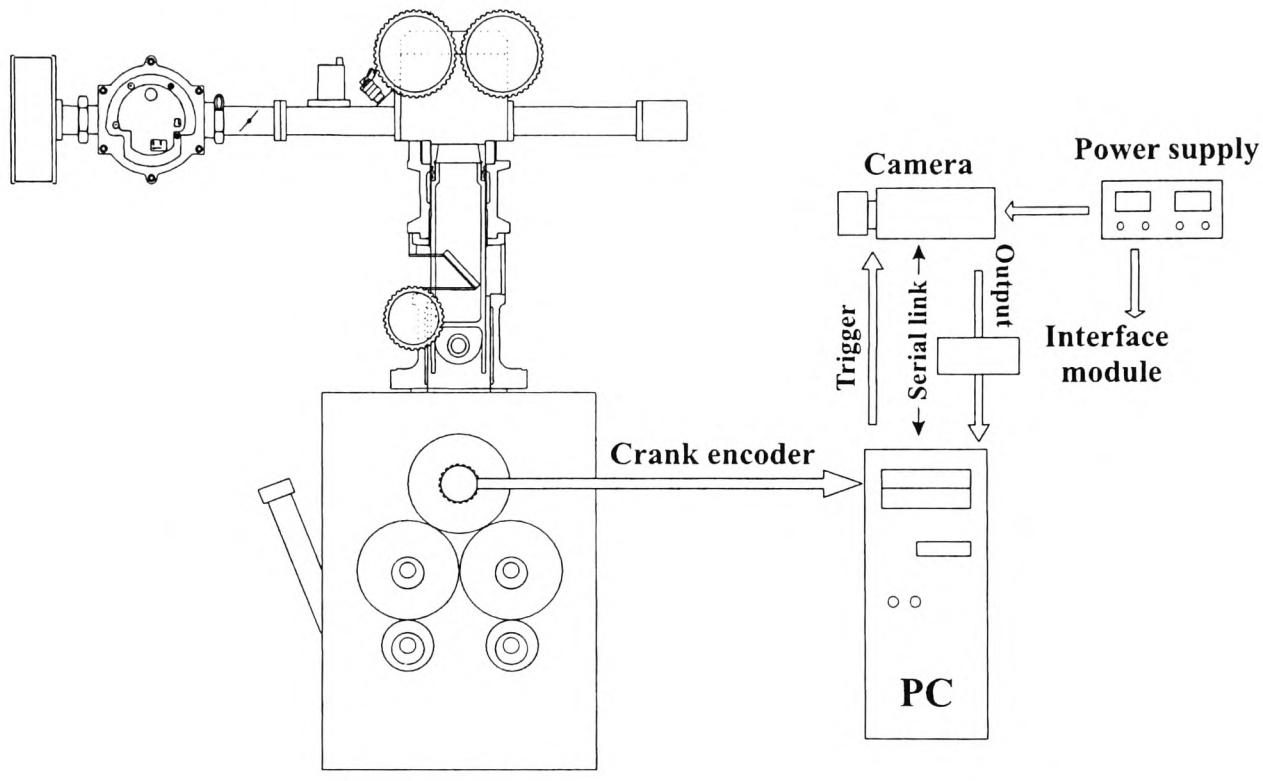


Figure 2.8: Schematic of the CCD imaging apparatus

The Pulnix TMC-9700 (Figure 2.9) is a digital process and output CCD colour video camera which uses a 2/3 inch high-resolution progressive scan interline transfer CCD (8.8×6.6 mm). It scans all lines sequentially from top to bottom at a frame rate of up to 30Hz. Like a non-interlace computer screen, it generates a stable image without alternating lines and provides full TV resolution of 768×484 pixels. The camera's pixels are not square but 13.6×11.6 μm thus the images from the camera have an aspect ratio of 1.17:1.

Because each pixel in the CCD array is only sensitive to overall intensity and does not differentiate colour, the TMC-9700 uses RGB primary colour filters to generate colour

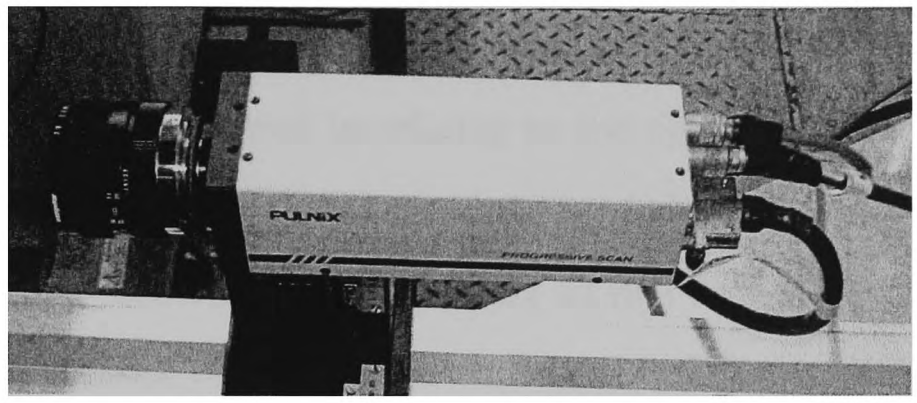


Figure 2.9: Pulnix TMC-9700 CCD colour video camera

reproduction. This CFA (Colour Filter Array) is a 3G and R/B staggered structure which mimics the human colour vision response. The pattern of the 4 by 2 pixel RGB filter is shown in Figure 5.8.

There are more green than red or blue pixels because the human eye is more sensitive to green. The camera has internal DSP (Digital Signal Processing) functions which assign RGB components to every pixel using dithering and neighbouring pixel values. Tests with monochromatic light showed that this would introduce cross talk between colour channels and was corrected for in the image processing described in Section 5.3.1.

The digital camera handled the whole process digitally so as to avoid D/A and A/D conversion errors. Moreover, an analogue signal would carry a ± 1 pixel uncertainty along a scan line. The camera framing rate could be set up to a maximum of 30 frames per second (fps), and the shutter speed could be set as low as 1/16000 of a second for filming very bright events. The working parameters, such as colour balance and shutter speed, could be set manually using the panel at the back of the camera, or using dedicated software on a PC via an RS-232C link. The operating parameters on the camera were set using the TMC9700P software provided by Pulnix. The camera's RS232C connector was linked to the computer's COM communications port, which enabled the exchange of data between the two. The interface module (CV-100) converted the digital output from the camera from TTL to RS-422 differential video. It was powered from a separate 12V power supply.

The frame grabber was a Mutech MV-1000 for the PCI bus, fitted with the optional MV-1100 Digital Module, which allowed interfacing to the digital camera. The grabber read video with 8 bits resolution (for each of the three fundamental colours), storing the data in the on-board 1Mbyte frame buffer and allowing for real time display as well as real time image capture. The grabber was configured and controlled through its own software, and triggering could be achieved both through software or an external TTL pulse (the same pulse that triggered the camera). The frame grabber's functions were activated through the MV1000 Capture Sequence software supplied by MuTech.

For combustion imaging, the camera was set to 'asynchronous reset mode' via the serial link. In this mode a 5V pulse signal (called VINIT) was applied to the external trigger input; this input was accessed through the power cable. The camera reset at every negative going pulse edge and captured the frame regardless of the shutter speed. It should be noted that in the progressive mode, the camera kept capturing pictures at a frequency of 30Hz when VINIT was low, therefore a 'narrow' low pulse of VINIT must be maintained so that the camera only captured once during every period of VINIT. This could be ensured by using crank angle flag signal when engine combustion was investigated. The video was kept and output continuously during VINIT high. The VINIT signal could also be recorded on the data acquisition system.

The ETCS program provided the negative going edge and a specified delay needed to trigger the CCD camera at the required crank angle timing. Even with no input, the output was high, thereby not triggering off the camera. With this setup, all of the images from an image capture sequence were from the same crank angle timing. Increments between sub-sequences could also be added so that after taking a certain number of images at one crank angle setting, the image timing would be changed according to the increment setting and the same number of images would be taken at the new crank angle setting.

2.4.4 High-speed combustion imaging equipment

A Phantom v7.1 high-speed CMOS (complementary metal oxide semiconductor) colour camera was used to carry out crank-angle-resolved combustion imaging. In general, CMOS imagers can achieve faster frame rates and higher system integration at lower cost than CCD imagers but suffer from greater dark noise and lower sensitivity. The high speed and sensitivity of the Phantom 7 series makes such a camera ideal for recording extremely fast events, particularly combustion and flame in an IC engine, tasks previously undertaken by rotating prism film cameras.

The Phantom v7.1 camera has a resolution of 800 by 600 pixels and a dynamic range of 8 bit. It offers a recording speed of 4800 fps for full frame pictures and a maximum frame rate of 160000 fps with reduced resolution. Using the CAR (Continuously Adjustable Resolution) feature, the active image area is adjustable in 16×8 pixel increments to maximise recording speeds and recording times allowing the user to tailor the field of view to the test subject's shape and dimensions.

Its other advanced features include the on-chip circuitry that provides true 'snap shot' shutter capabilities (down to 2μs), auto exposure, and pixel level Extreme Dynamic Range (EDR) control to extend picture quality in harsh lighting conditions.

A PC with the Phantom software package installed was used for communication with the camera, such as setup, control, viewing images, analysis, and download of image files, via a gigabit ethernet interface.

2.4.5 Image processing

To obtain quantitative results from the imaging experiments, the raw images need to be post-processed. This section introduces the general concepts and strategies in image processing. Procedures on how to extract information from images according to theories

for combustion analysis and laser-based techniques will be detailed in the subsequent sections.

The first process involves the removal of systematic errors present in the imaging system. Normally systematic errors include noise from the imaging equipment (cameras, photo-multipliers *etc.*), background signals and distortion due to optical paths. Instrument noise and background signals can be subtracted from the raw images straightforwardly, albeit most of the time they are highly dependent on the experimental setup and conditions. Optical distortion can be corrected by image registration (re-mapping). Other systematic errors in particular applications need specialised calibration.

1) Image storage and format

Digital images are constantly stored as two-dimensional arrays, *i.e.* matrices, in which each element of the matrix corresponds to a single pixel in the displayed image. A pixel is derived from picture element and usually denotes a single dot on a computer display. For example, an image composed of 200 rows and 300 columns of different coloured dots would be stored as a 200-by-300 matrix. Some images, such as RGB, require a three-dimensional array, where the three planes in the three dimensions represent the red, green and blue pixel intensities, respectively. Apart from the RGB format, images can also be stored as indexed (colourmap), intensity (greyscale) or binary (black-and-white) format.

These principles of image storage make image arithmetic (addition, subtraction *etc.*) quite straightforward — with matrices. More complicated image processing which is commonly used in engine research includes image registration and filtering.

2) Image registration

Image registration is the process of aligning two or more images of the same scene. Typically, one image, called the base image, is considered the reference to which the other

images, called input images, are compared. The object of image registration is to bring the input image into alignment with the base image, by applying a spatial transformation to the input image. A spatial transformation maps locations in one image to new locations in another image. Determining the parameters of the spatial transformation needed to bring the images into alignment is the key to the image registration process. Image registration is necessary in several common applications:

- To correct for the viewpoint shift of a camera due to system vibration
- To bring the viewpoints of a multi-camera system into consistence
- To correct for optical distortion present in the image acquisition system
- To convert image pixel coordinates to real-world coordinates
- To remove the effect of video interlacing

If only one camera is employed, the major viewpoint shift is caused by vibration of the test rig. For experiments which need two or more nominally identical cameras to take images from the same event, the fields of view from different cameras cannot be practically identical because the cameras cannot be in exactly the same place. In addition, the noise levels differ between cameras. Optical paths where there are lenses or non-flat optical windows can cause distortion, for example, to make a circle appear elliptical or even more warped in the images. Thirdly, to extract a quantity such as velocity from images calls for the conversion from pixel coordinates to real-world coordinates, even though this is not necessary for the extraction of a relative quantity such as position.

All these five tasks can be done by image registration, usually by recording a set of rectangular grids or concentric circles. Firstly, grids are recorded from cameras in their different positions. Optical distortion means that the recorded shapes are no longer regular. And with spacing or radius exactly known, the calculation of real-world coordinates and the correction for interlacing become feasible.

Point mapping is widely used to determine the parameters of the transformation required to bring an image into alignment with another image. Point mapping uses particular points in a pair of images, namely base image and input image, that identify the same feature or landmark in the images (Figure 2.10). Then, spatial mappings for the coordinates are inferred from the positions of these control points.

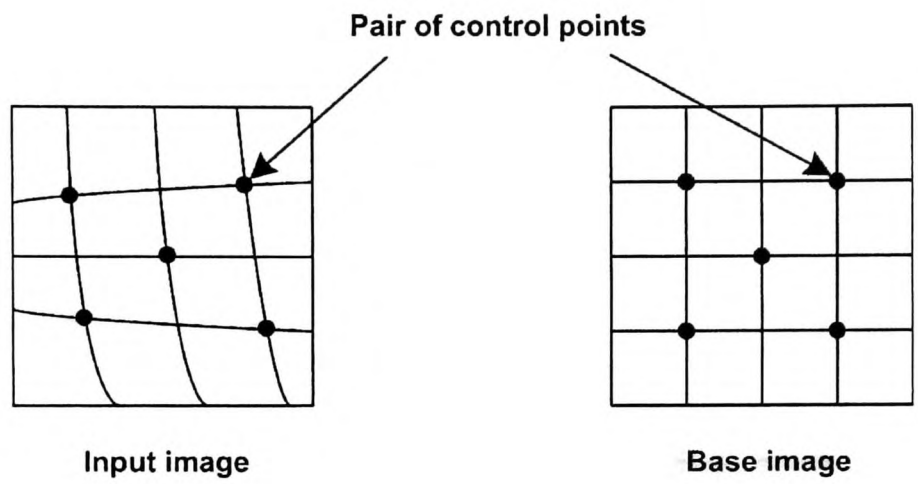


Figure 2.10: Schematic of image registration

Spatial mapping can be in the form of either global transformation or local transformation. In global transformations, a single mathematical expression (2 mappings for x and y coordinates, respectively) applies to an entire image. When using high-order polynomials, it is difficult to make the global mappings numerically stable. Furthermore, distortion cannot always be satisfactorily corrected by global transformation due to the great spatial variation and high computing expense. Therefore local transformations are more commonly applied. In such cases, different mathematical expressions apply to different regions within an image, which means a significantly greater number of mappings are required, yet in much simpler forms.

Cross-correlation can be used to ensure that the control points are accurately selected. A number of regions around each manually pre-selected control point can be defined in the input image and around the matching control point in the base image, and then the correlation between the values at each pixel in the region are calculated. Next, the position with the highest correlation value is picked and used as the optimal position of

the control point.

The features and applicability of the most commonly applied transformations are listed in Table 2.2. It should be noted that, if more than the minimum control points are selected, a least squares solution will be employed.

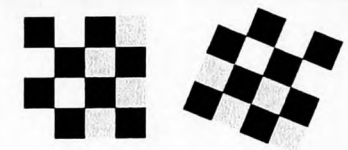
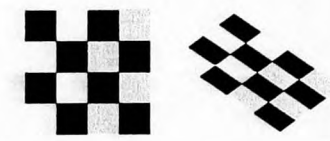
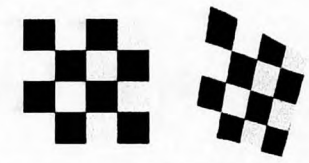
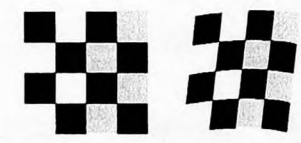
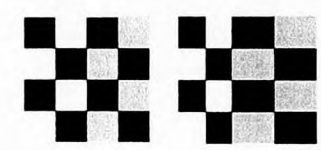
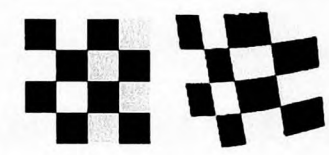
Type, applicability & description	Min. control points	Example
Linear conformal: use when shapes in the input image are unchanged, but distorted by some combination of translation, rotation, and scaling. Straight and parallel lines remain straight and parallel.	2 pairs	
Affine: use when shapes in the input image exhibit shearing. Straight and parallel lines remain, but rectangles become parallelograms.	3 pairs	
Projective: use when the scene appears 'tilted'. Straight lines remain straight, but parallel lines converge.	4 pairs	
Polynomial: use when objects in the image are curved. The higher the order of the polynomial, the better the fit, but the result can contain more curves than the base image.	6 pairs (2^{nd} order), 10 pairs (3^{rd} order), 15 pairs (4^{th} order)	
Piecewise linear: local affine mapping on triangular regions	4 pairs	
LWM: local 2^{nd} polynomial fitting	6 pairs	

Table 2.2: A list of common transformations

3) Image filtering

Filtering is a technique for modifying or enhancing an image. For example, you can filter an image to emphasize certain features or remove other features. Filtering is a neighbourhood operation, in which the value of any given pixel in the output image is

determined by applying some algorithm to the values of the pixels in the neighbourhood of the corresponding input pixel. A pixel's neighbourhood is a set of pixels, defined by their locations relative to that pixel, which is called the centre pixel. The neighbourhood is a rectangular block, and as the pixel moves from one element to the next in an image matrix, the neighbourhood block slides in the same direction. Figure 2.9 shows the neighbourhood blocks for some of the elements in a 6-by-5 matrix with 2-by-3 sliding blocks. The centre pixel for each neighbourhood is marked with a dot.

To perform a sliding neighbourhood operation:

- Select a single pixel
- Determine the pixel's neighbourhood
- Apply a function, which must return a scalar, to the values of the pixels in the neighbourhood
- Find the pixel in the output image whose position corresponds to that of the centre pixel in the input image and set this output pixel to the value returned by the function
- Repeat the above 4 steps for each pixel in the input image

For instance, suppose Figure 2.11 represents an averaging operation, with a function summing the values of the six neighbourhood pixels and then dividing the sum by 6. The result is passed to the value of the output pixel. For more demanding linear-filtering applications, such as smoothing, sharpening and edge enhancement, the function can be a discrete convolution — a neighbourhood operation in which each output pixel is a weighted sum of the neighbouring input pixels.

More sophisticatedly, dedicated filters can be designed according to their response in the frequency domain. In such cases, image processing is always implemented in combination with signal processing tools.

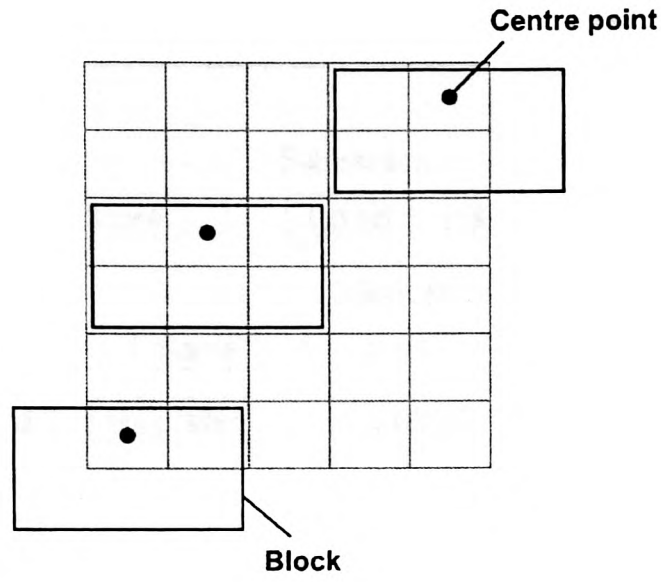


Figure 2.11: Sliding neighbourhood blocks

4) Image masking/overlay

When capturing images, some solid parts of the object are inevitably invisible because there is almost no light signal generated or scattered from them. For example, the inlet and exhaust valves from a vertical viewpoint (through the piston), the valves, the cylinder head and the moving piston from a side viewpoint (parallel to the piston crown), as well as the walls of the combustion chamber in all cases. Hence to make the images clearly represent all the features (especially geometric) of the field of view, a mask/overlay image, usually artificially made, needs to be added to the raw images, resulting in a constantly shadowed 'background'. This process is known as image masking/overlay.

2.4.6 Image processing software

Based on the image processing strategies detailed in Section 2.4.5, as well as the principles of image analysis for PLIF and colour-ratio pyrometry in subsequent sections, a MATLAB graphical user interface, **R**aw **I**mage **P**rocessing for **E**ngines (RIPE, see Figure 2.12) has been built to perform the following functions:

- Image loading and saving

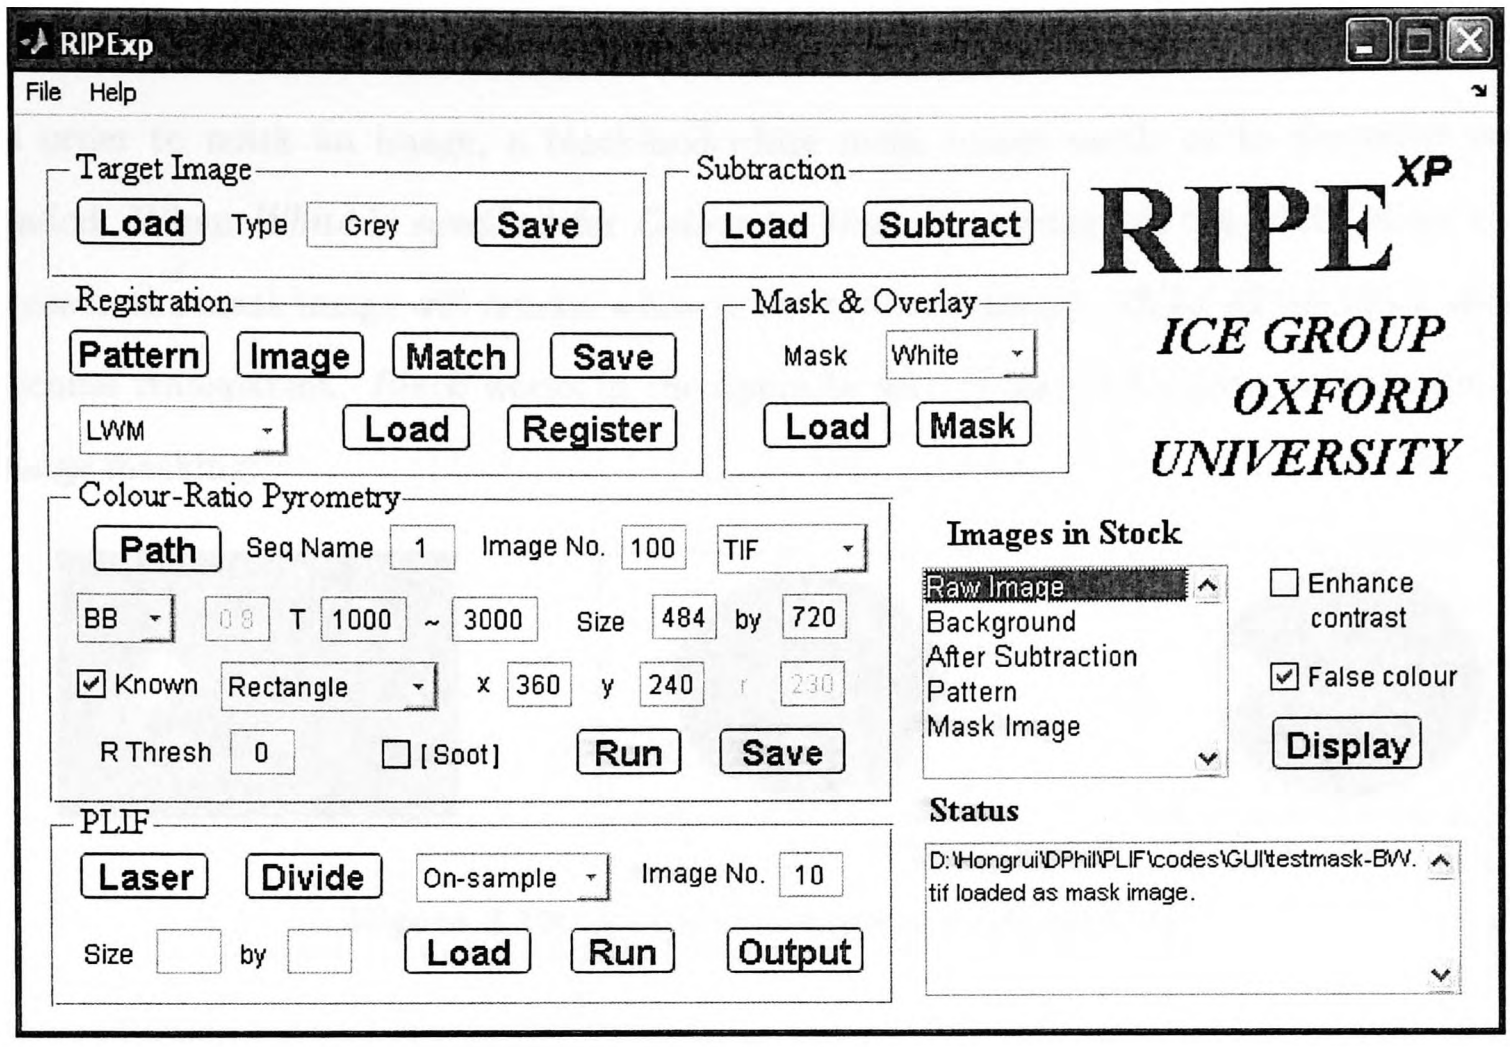


Figure 2.12: The graphical user interface of the image processing software

- Subtraction of background noise
- Image masking and overlay
- Image registration (matching control points, generating, loading and saving transformations, *etc.*)
- Calibration for the non-uniformity in laser sheet intensities
- Colour-ratio pyrometry analysis
- Statistics for images
- Multiple image display, contrast enhancement and false colour for non-RGB images
- Log text-file to record activities

Raw images can be loaded from either menu *File* → *Load image...* or the *Load* button in *Target Image*. As soon as an image is loaded, its format is displayed. Background noise comes in the form of background images, which will then be subtracted from the target

image.

In order to mask an image, a black-and-white mask image needs to be prepared and loaded. When *White* is specified for *Colour* in *Mask & Overlay*, all the white edges and areas in the mask image will remain white in the resultant image, whilst all the black areas become transparent. *Black* works in the opposite way. Figure 2.13 gives an example of image masking.

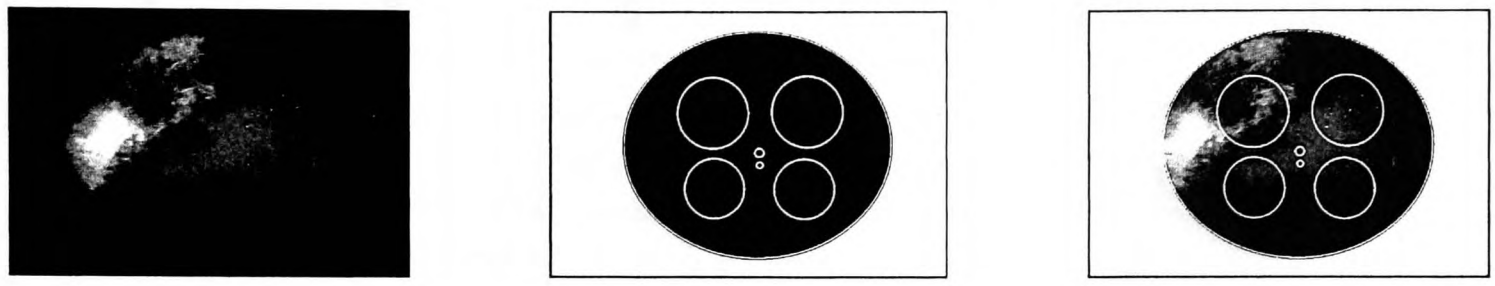


Figure 2.13: An example of white image masking

For registration, the software provides six types of transformation, namely, linear conformal, affine, projective, polynomial, piecewise linear and local weighted mean (LWM), with the complexity and applicability increasing accordingly, as detailed in Section 2.4.5. Once both the base image (pattern) and the input image (image of the pattern from the camera) have been loaded, the Match button is used to pop-up the control point selection window (see Figure 2.14). The selection of control points can also be carried out using non-interactive code lines, which is particularly useful when there are many control points to select and they are distributed in a regular manner (*e.g.* a set of circles with only radial distortion).

Sufficient control points need to be selected from both the base image on the right side and the input image on the left side. Then a transformation (according to which type is specified in the pop-up menu) will be generated and saved via the *Save* button, which can be loaded via the *Load* button later.

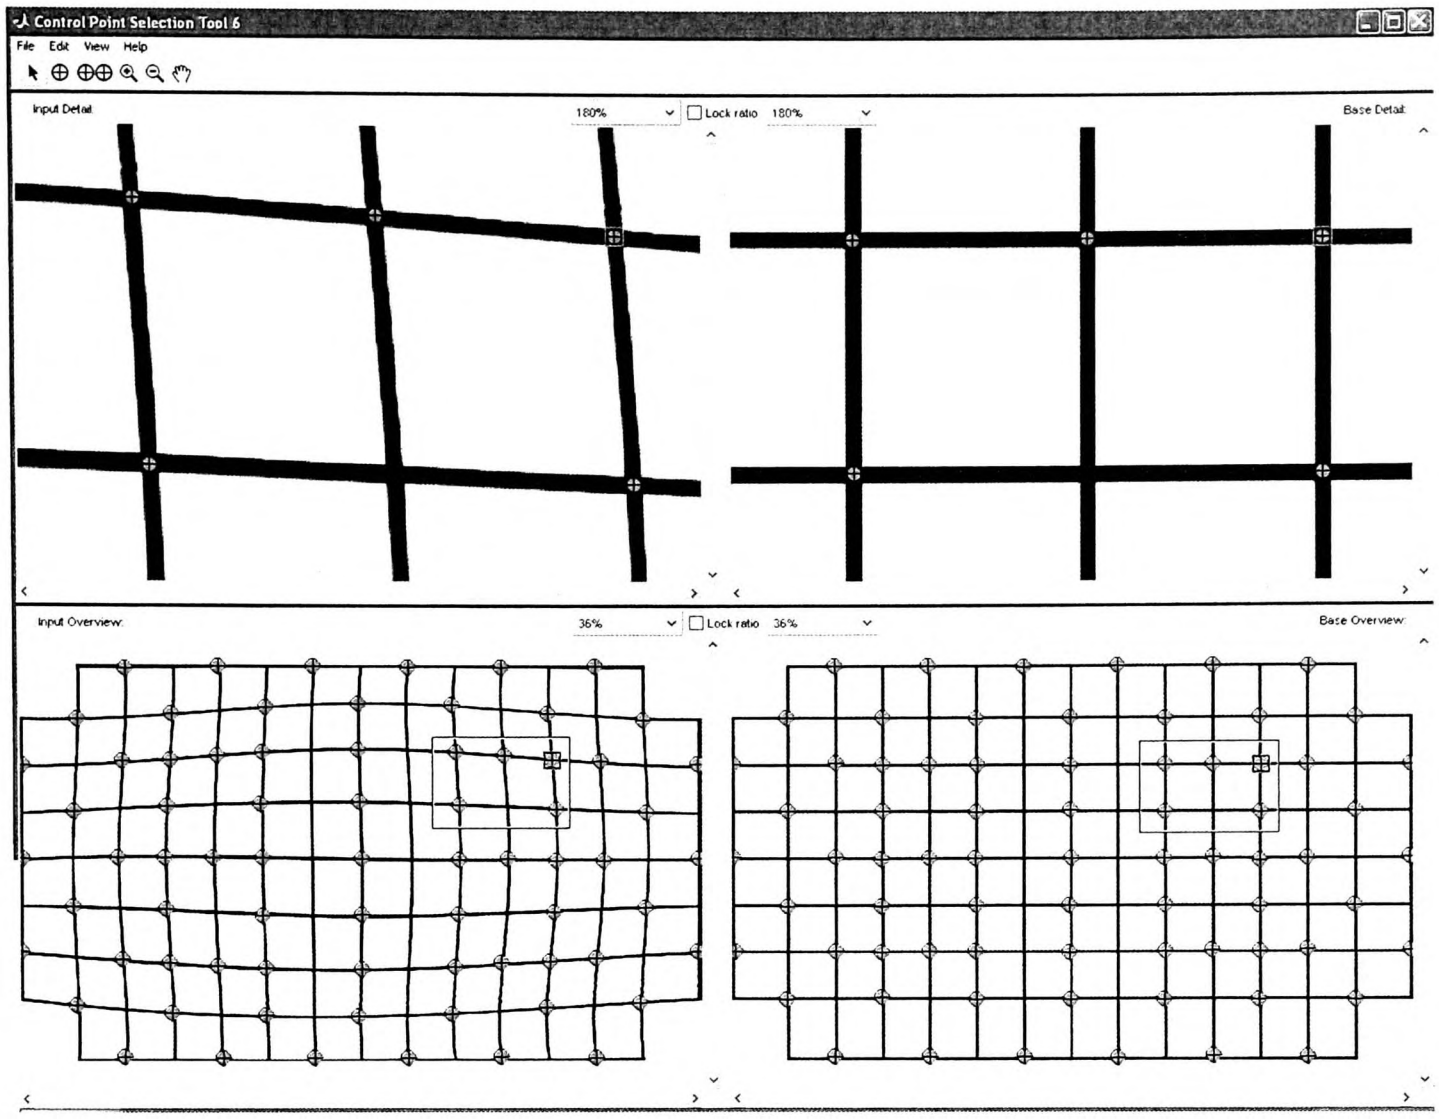


Figure 2.14: An overview of the control point selection tool

Figure 2.15 gives an example of image registration using the LWM transformation. The registered pattern image (left) shows moderate restoration from the previous distortion. Since only 25 out of 221 pairs of control points were selected, a better restoration can be achieved if more points are included.

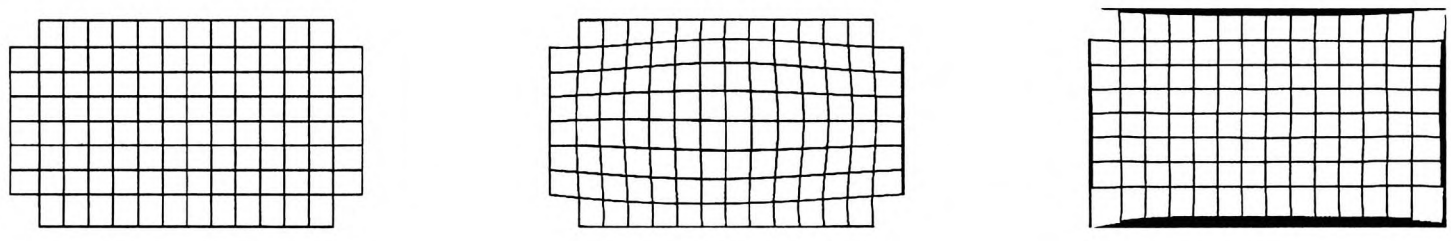


Figure 2.15: An example of image registration

The statistical calculation depends on which strategy is chosen for image acquisition, *e.g.* on-chip accumulation or off-line averaging (on-sample averaging), and hence determines

how much information about cycle-by-cycle variation can be extracted from the averaged images.

The colour-ratio pyrometry analysis is based on the theory described in Chapter 5. Various results *e.g.* pixel distribution with respect to object temperature for red, green and blue in the images from the CCD camera as well as statistics are calculated and saved in this section.

Images in Stock lists all the images before processing, in process and after processing. This is also where images are selected for processing, display, saving *etc.* Information about the activities is shown in the *Status* textbox, and can also be saved in a step-by-step log file for reference later.

The software also provides options to improve the quality of the image display, such as contrast enhancement and false colour for grey-scale images, as shown in Figure 2.16.



Figure 2.16: Example images — (a) raw grey-scale image (Davy [13]); (b) contrast enhancement; (c) contrast enhancement and false colour

2.5 Summary

Details of the engine instrumentation system and the data acquisition and processing tools have been presented. The engine instrumentation system was based on the previous

Oxford Rover K4 engine test bed. The facilities have been established for both GDI and PFI operations, both conventional ignition coils and a coil-on-plug ignition system, and measurements of pressures and flow rates, temperatures, emissions *etc.*

To control the engine timing effectively, an engine timing control software developed in LabVIEW has been used for this particular engine system. A National Instruments data acquisition system has been used to handle both high-speed and low-speed signals from the instrumentation system.

Three camera systems are described for the use with planar laser induce fluorescence, soot pyrometry and high-speed combustion imaging, respectively. Their working principles as well as the specs and operating procedures have been given in detail. The suitability of each camera for its corresponding applications has also been discussed.

The data/image processing tools have been introduced in depth. Various combustion analyses with the experimental data have been incorporated into a MATLAB code package CoBRA. These include burn rates analysis, estimation of completeness of combustion, pressure-data-based heat release analysis *etc.* The original version of CoBRA has also been extended substantially to have functions for diesel engines and 2-stroke combustion systems.

The image processing software package developed in MATLAB provides extensive functions to analyse images from different optical work. Finally, examples of image processing have been given as demonstration of the capability of the software.

Chapter 3

Two-Stroke Engine Combustion Analysis

Learning without thought is labor lost.

— Confucius (551–479 B.C.)

3.1 Introduction

A case study with data from a small homogeneous charge, two-stroke glow ignition engine has been conducted. One of the main objectives was to develop a method to combine burn rate and heat release analyses, in order to investigate the combustion process of engines with a premixed charge, but ignited by compression. The findings also include the combustion and performance characteristics of different nitromethane/methanol fuel blends, and the characteristics of different glow plugs. Difficulties such as an unknown ignition timing and polytropic index have been addressed by combining both heat release and mass fraction burn analyses. Results for all operating conditions have shown a good correlation between the two methods.

3.2 Experimental aspects

Combustion analysis has been conducted on a small two-stroke glow ignition engine of the type used in model aircraft [41]. The experiments in this engine were carried out by Prof. Robert Raine and his colleagues at The University of Auckland. The subsequent data analysis was carried out in this work.

Traditionally, burn rate analysis is widely used for spark ignition engines, while the heat release analysis is more suitable for compression ignition engines. This type of miniature engine does not fall into either category. In spark ignition engines, the mixture is premixed and the combustion is initiated by a spark, so the time of start of combustion is known. In compression ignition engines, combustion starts from auto ignition of the fuel. The start of combustion is unknown, but it has to be later than the fuel injection timing. This miniature engine, with a pre-mixed mixture and glow plug assisted autoignition is a special case. It is likely that the compression and ignition processes would be most similar to homogeneous charge compression ignition (HCCI) engines. In this section, both analyses of burn rate are considered together to establish correlations between them. The aim is to gain an understanding of the applicability of both analyses to this type of engine. Since the combustion process is similar to HCCI engines, it is anticipated that the study here would also be relevant to those engines.

In the experiments a K&B.45 Sportster model engine was used. The engine is a single cylinder, air cooled, two-stroke engine with dimensions: bore 21.6mm, stroke 20.3mm, swept volume 7.4cm^3 and geometric compression ratio of 7.6. Maximum rated power is 0.62kW at a speed of 11500RPM. In the experiments, engine speed was restricted to approximately 7500RPM. The engine was fuelled with a mixture of methanol, nitromethane and lubrication oil, using a glow plug and piston compression to assist the ignition of the charge. The supplied fuel mass flow rate was varied (not precisely quantitatively) by setting the needle valve position. The detailed experimental equipment, arrangement and

results from measured engine power have been previously reported [49].

With an external sampling rate of 200 points/cycle, 200 cycles of data at 27 different engine operating conditions (all at wide open throttle) were measured by varying three parameters:

- 3 different fuel compositions (nitromethane 0%, 10%, 20%)
- 3 different glow plugs ('Hot', 'Medium', 'Cold')
- 3 different fuel-supply needle settings ('2', '3', '4'); the greater the number is, the higher the fuel flow rate is.

3.3 Analysis issues

For convenience, the names of the data files are used instead of listing their corresponding operating conditions in the following discussion. For example, '0-4-Cold' represents the condition of 0% nitromethane concentration, needle setting '4', cold plug. This naming system will be used hereafter.

In this glow ignition engine, there are issues with unknown timing for the start of combustion, choice of polytropic index for heat release analysis, and sampling rate. These will be addressed as follows.

The burn rate analysis in CoBRA requires the spark timing to calculate the compression polytropic index prior to combustion. This is chosen to be the polytropic index (polynomially) fitted from a point halfway between inlet closure and spark timing to the point of spark timing. Given the fact that this engine has no spark ignition, the effect of varying the 'nominal' spark timing specified in the code on the phase of 1% mass fraction burned was investigated. Figure 3.1 shows the result of this investigation for the data set for mean values for case '0-4-Cold'. This data set was chosen as it is one showing extreme variations in cycle-by-cycle pressure data, so that it was felt this would give the most

rigorous test of the effects being studied. As shown in Figure 3.1, a variation of 30°CA in the 'nominal' spark timing resulted in a variation of only 3°CA in the phase of 1% MFB. Therefore, 145°CA ABDC was selected as the 'nominal' spark timing in the subsequent analysis.

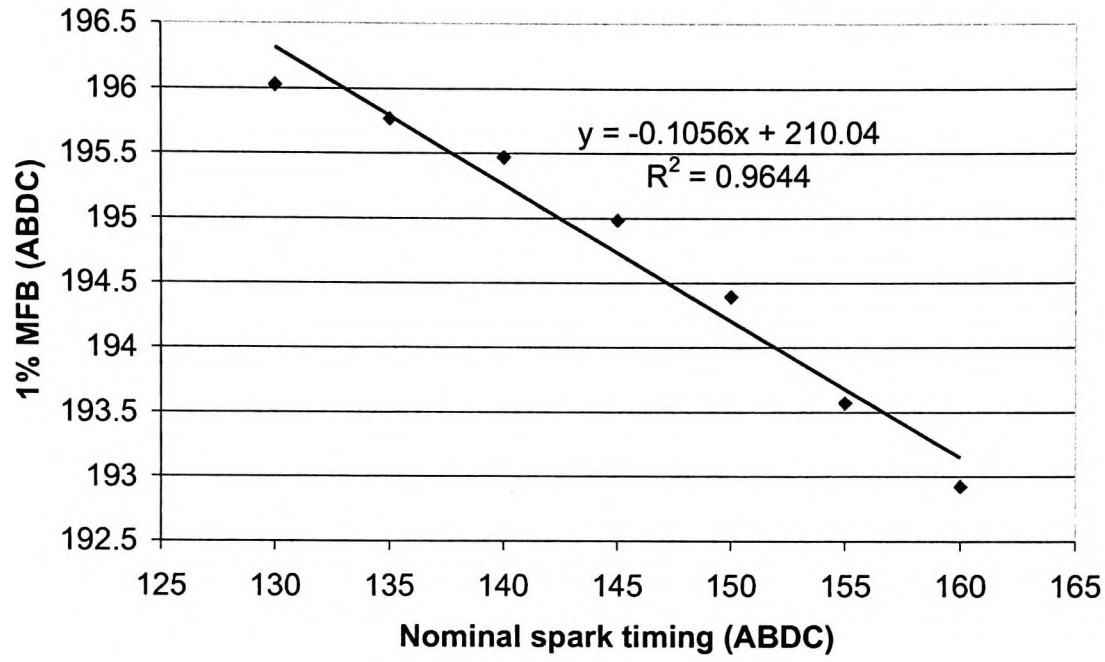


Figure 3.1: Correlation between 'nominal' spark timing and 1% MFB timing, '0-4-Cold'

In heat release analysis, γ is a function of temperature and air/fuel ratio and has an influence on the magnitude of Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$ according to Eq.(2.15). In order to investigate the effects of γ, results for calculated Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$ from '0-4-Cold' are shown in Table 3.1 and Figure 3.2. Table 3.1 shows that, with γ varying between 1.2–1.4, the changes in the phasing of Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$ were only within 1.8°CA (one sample point). It was also noted that CA of Q_{nmin} and CA of $(\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta)\,|_0$ were not exactly the same. This difference was probably because of the numerical error from the computation in MATLAB. However, it is the phase of the start of combustion (SOC) that is of most interest, and hence it was assumed that the effect of changes in γ within a typical range on the phasing of Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$ was negligible. A value of $\gamma = 1.4$ was therefore selected in the subsequent analysis.

The cylinder pressure signal was measured at 200 samples/cycle, which is lower than what

γ	1.2	1.3	1.4
CA of Q_{nmin}	176.4	176.4	176.4
CA of Q_{nmax}	280.8	280.8	282.6
CA of $(\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta)_{\mathrm{min}}$	167.4	165.6	163.8
CA of $(\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta)_{\mathrm{max}}$	217.8	217.8	219.6
CA of $(\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta)\,\vert_0$	176.4	177.3	177.8

Table 3.1: Effect of γ on the phasing of Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$, '0-4-Cold'

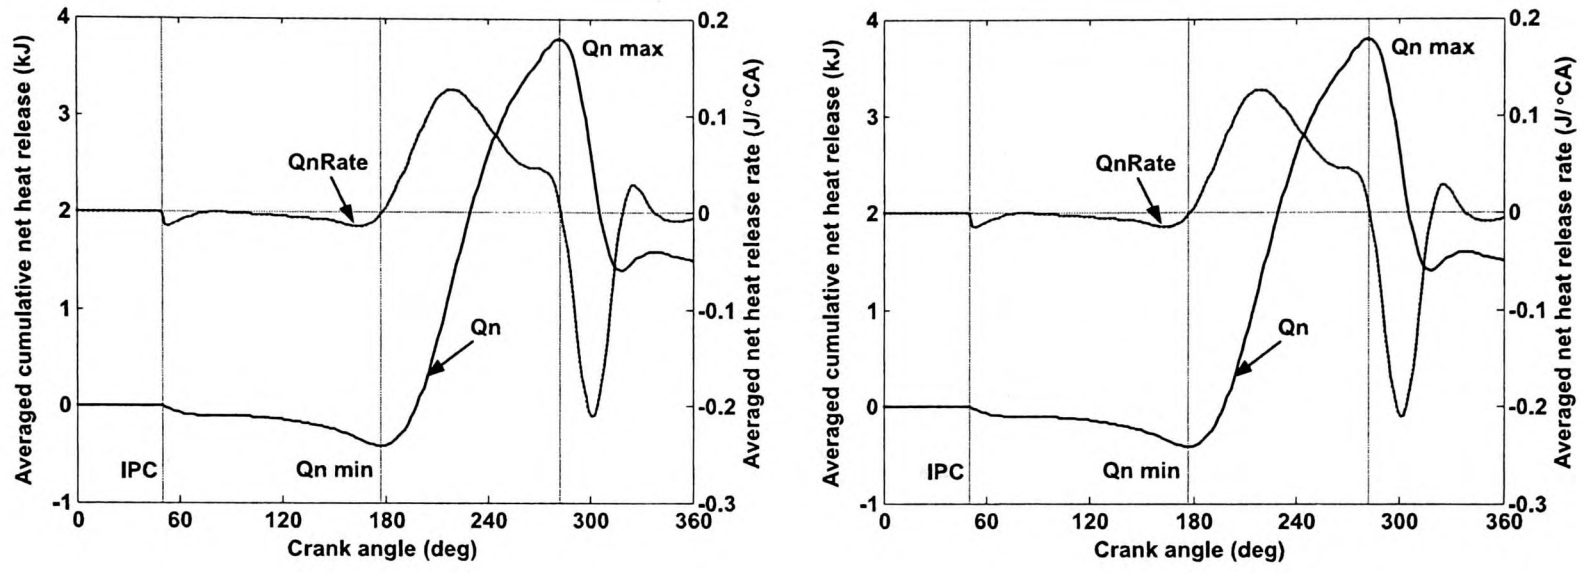


Figure 3.2: Sample results of cycle-averaged Q_{n} and $\mathrm{d}Q_{\mathrm{n}}/\mathrm{d}\theta$, '0-4-Cold' (left) and '0-4-Hot' (right), calculated with $\gamma = 1.4$

is typically used for heat release analysis in the literature [60]. However, since heat release analysis is usually applied in compression ignition engines with much higher compression ratio, and so much faster pressure rise, it is difficult to judge whether the current sampling rate is sufficient without further analysis. Too low a sampling rate will introduce aliasing in the signal and truncation errors when evaluating $\mathrm{d}p/\mathrm{d}\theta$. Figure 3.3 shows the spectrum of the cylinder pressure signal in one data set. There are distinct peaks at the fundamental frequency and higher harmonics, indicating the cyclic nature of the signal. The spectrum flattens off above $0.08°\mathrm{CA}^{-1}$, indicating there is little or no power above that frequency. This assures that the sampling rate is fast enough to capture most information contained in the signal.

Errors also arise when approximating the derivative of a continuous signal by finite difference. In CoBRA, $\mathrm{d}p/\mathrm{d}\theta$ is calculated by the forward difference scheme, where the

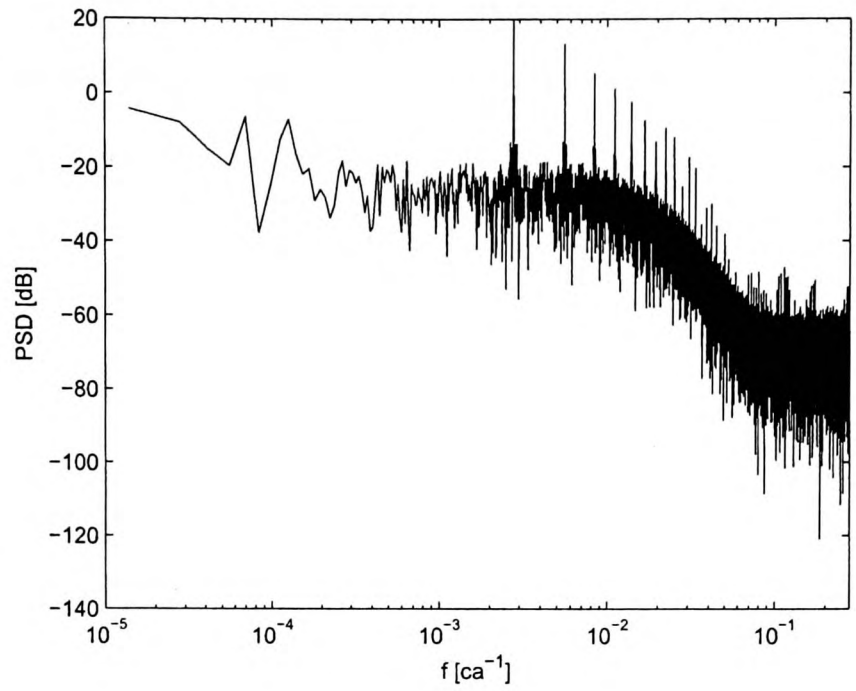


Figure 3.3: Power density spectrum of the pressure signal in '0-4-Cold', normalized to the DC value

truncation error (ε) is in the order of a crank angle step (k), *i.e.* $\varepsilon = O(k)$. 3-point and 5-point central difference schemes have been considered, which have ε in the order of k^{-2} and k^{-4} respectively. The three schemes have been compared against each other. The percentage differences between them are typically $\pm 0.1\%$. When $dp/d\theta$ is close to zero, because of the quantization and electrical noise, the difference can go up to $\pm 1\%$. This is however irrelevant because heat release analysis is only valid between IPC and EPO. This is when $dp/d\theta$ are high. Therefore, it is concluded that the truncation error is insignificant compared to random noise in $dp/d\theta$, and that the sampling rate is sufficient.

3.4 Results and discussion

3.4.1 Correlation between burn rate and heat release analyses

Historical research on engine thermodynamics has shown that in conventional spark ignition engines, the 0–10% MFB is well correlated to the initial period of combustion and the 10–90% is representative of the main combustion period during which most of the

fuel burns. On the other hand, the minimum of Q_n (where the net heat release rate becomes positive, as shown in Figure 3.2) and the maximum of Q_n (where the net heat release rate becomes negative) are commonly defined as the start and end of combustion for compression ignition engines, respectively. As the combustion process in this engine is not the same as either spark ignition engines or compression ignition engines, the initial period and main combustion need to be defined when employing heat release analysis.

Figure 3.4 shows the highly linear correlation between the 0–1% MFB and 0–10% MFB durations based on data for all engine operating conditions. In this figure, the time for 0% MFB refers to the 'nominal' spark timing of 145°CA ABDC. This correlation implies that features of the initial combustion can be described by either of the two durations. However, in order to obtain an absolute value for the rate of combustion, the SOC needs to be defined.

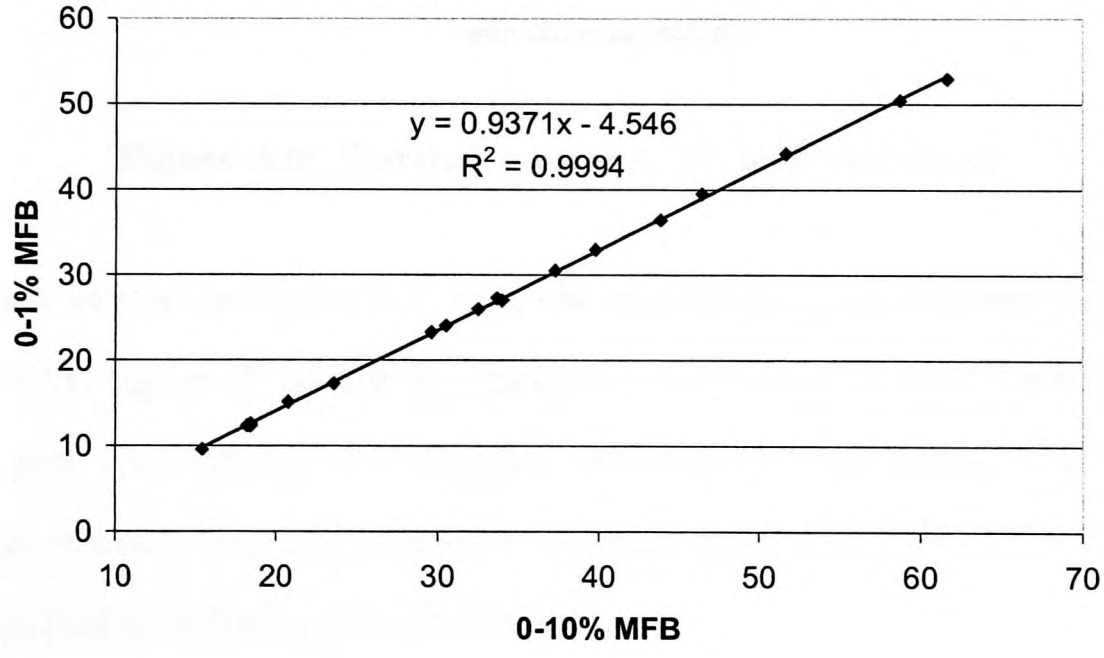


Figure 3.4: Correlation between 0–1% MFB and 0–10% MFB

Figure 3.5 shows that the MFB analysis and heat release analysis have close correlation except for 2 anomalies (represented by hollow symbols in Figure 3.5). Note that both the minimum of Q_n and 1% MFB are indicators of SOC. Since the minimum of Q_n is an absolute timing and it correlates well with 1% MFB, it will be used to define SOC

to complement the burn rate analysis. The anomalies correspond to '0-4-Cold' and '0-4-Medium', both of which contain several misfire cycles. These two sets of data also have the greatest cycle-by-cycle variation and extremely slow combustion. This was established from the IMEP and 0-90% MFB results to be discussed later. For these two cases, the rate of combustion will be less certain because SOC is not well defined. It has been noted

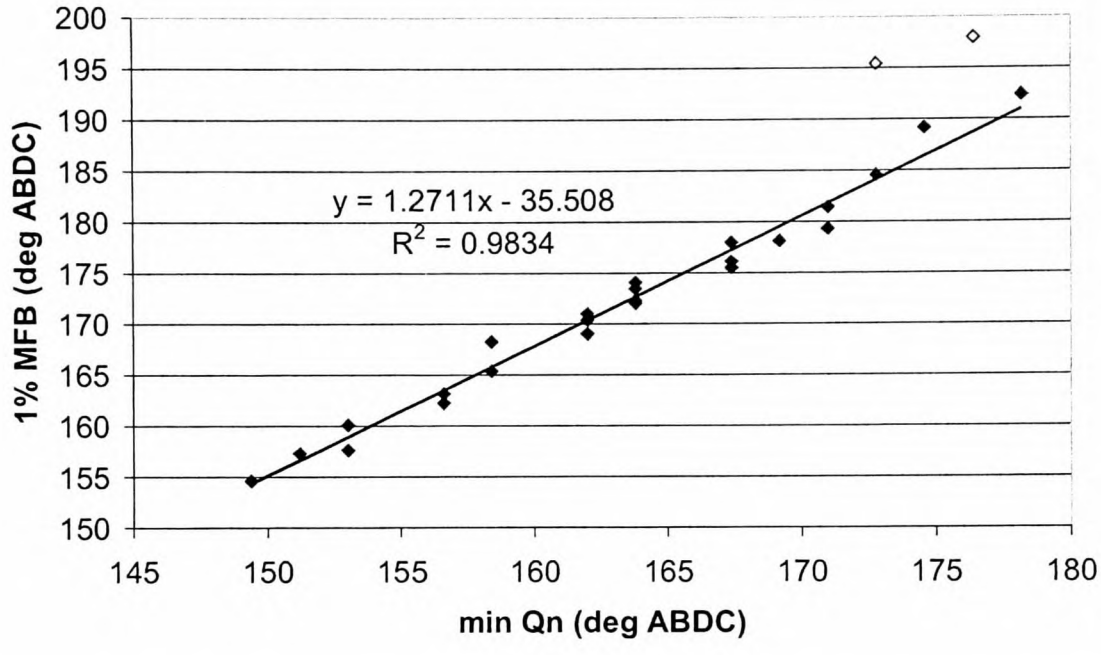


Figure 3.5: Correlation between 1% MFB and Q_{nmin}

from data such as that in Figure 3.2, that the maximum of Q_{n} remains almost constant (280.8–282.6°CA) under all operating conditions and hence we infer that this estimates the exhaust port open timing. In summary, therefore, for this engine, the initial period of combustion is defined as the difference between Q_{nmin} and 10% MFB, and the main combustion period is defined as 10–90% MFB.

Correlation between Q_{nmin}–10% MFB and 10–90% MFB is shown in Figure 3.6. As expected, a quicker initial period of combustion tends to be followed by a quicker main combustion. Again, the correlation cannot apply well to the data from '0-4-Cold' and '0-4-Medium' (represented by hollow symbols in Figure 3.6) because both of them present abnormally slow initial and main periods of combustion.

This analysis, making use of both 'mass fraction burned' and 'heat release' methodologies,

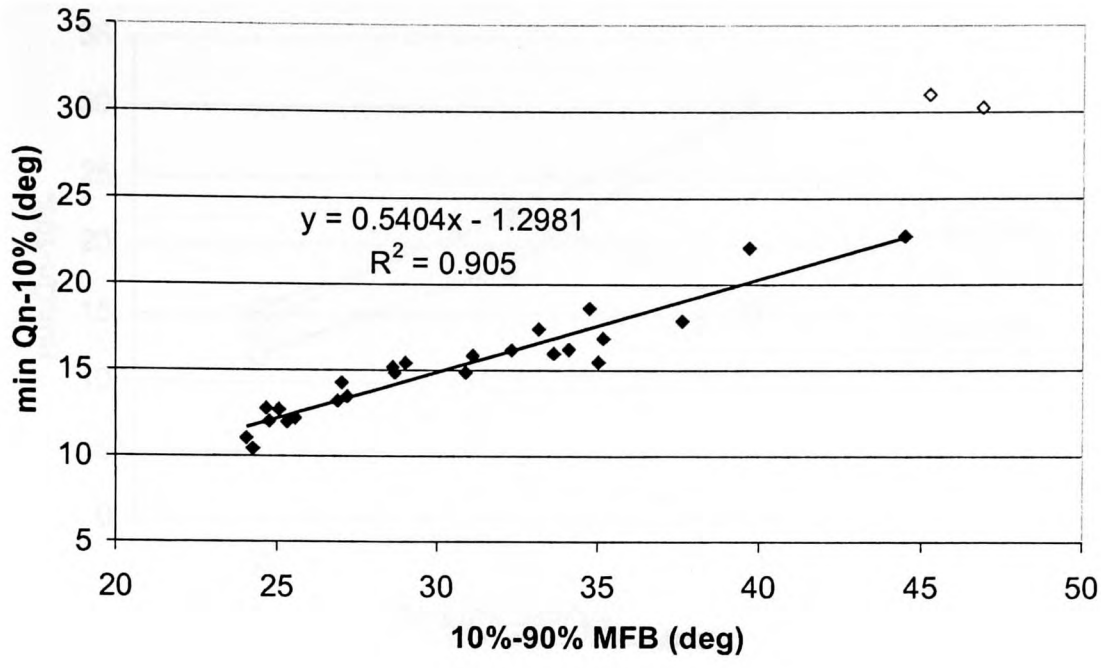


Figure 3.6: Correlation between Q_{nmin}–10% MFB and 10–90% MFB

has established that both give useful results for the homogeneous, compression assisted engine under study. Comparisons between results from the two methods show good agreement in those parameters which can be compared.

3.4.2 Effects of fuel blend and glow plugs

The air-fuel ratio was not monitored or controlled in the experiment, other than keeping fixed needle settings for the fuel supply. This makes the interpretation of the results difficult, particularly when the nitromethane content changes, since both fuel and mixture strength effects are present. The problem can be first isolated by looking at the results of pure methanol as shown in Figure 3.7.

The initial period of combustion, as indicated by Q_{nmin}–10% MFB, is a strong function of laminar burning velocity in spark ignition engines, and more generally a function of chemical kinetics [24]. As the needle setting is reduced, less fuel is supplied, so relative air/fuel ratio (λ) is increased. Figure 3.7 shows that the initial combustion period consistently reduces as the needle setting is reduced for the 'cold' and 'medium' plugs. Experiments [60] show that the laminar burning velocity of methanol peaks at a relative λ of about 0.9,

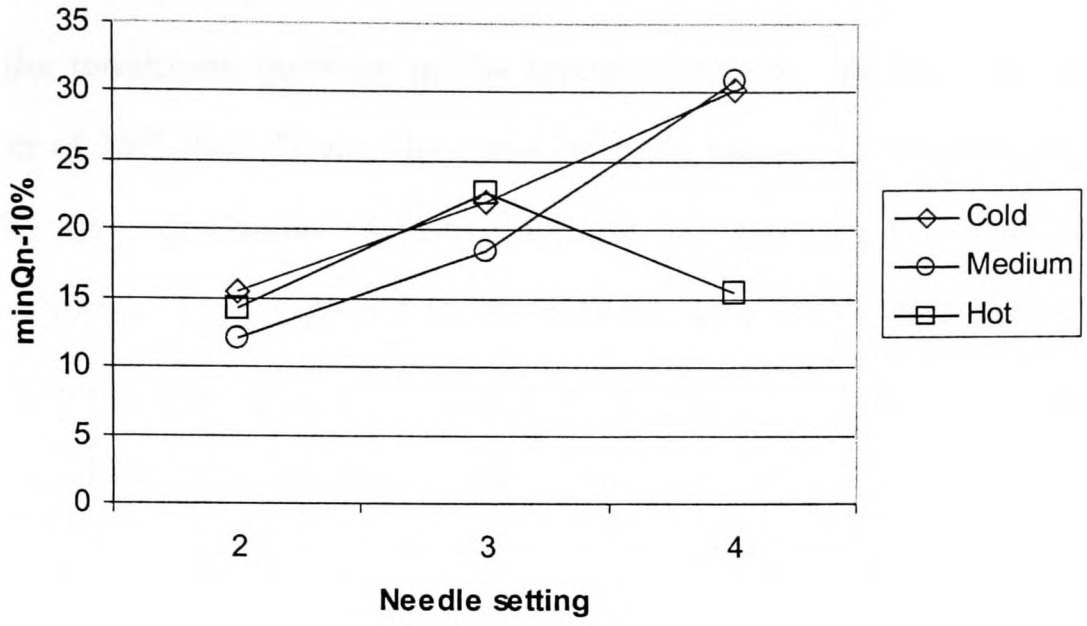


Figure 3.7: Q_{nmin}–10% MFB for various needle settings and plug types using pure methanol

and it decreases monotonically either side of that ratio. For λ to increase, and at the same time increase the laminar velocity, the mixture must be richer than stoichiometry. Furthermore, the reduction in Q_{nmin}–10% MFB is almost linear. This can only happen for very rich mixtures, as the closer it is to $\lambda = 0.9$, the smaller the change in the laminar velocity. The fact that '0-4-Cold' and '0-4-Medium' have very poor combustion also indicates that the mixture is near the rich burning limit, hence λ is much smaller than 0.9. '0-4-Hot' does not follow the general trend. With needle setting 4 and 0% nitromethane, λ will be the lowest amongst all tested conditions. It is likely that the hot plug has facilitated the combustion by supplying enough heat to vaporize the methanol, while the cold and medium plugs cannot. In other words, a high-temperature plug effectively extends the rich burning limit of the fuel.

When 20% nitromethane is added to the methanol, the initial period of combustion changes is shown by the filled bars in Figure 3.8 (left), relative to the values for methanol only. The figure also shows limit bars for the cycle-by-cycle variations. This figure indicates that in 5 out of 9 cases, the initial period of combustion is shorter with nitromethane added than without. As discussed in [41], adding 20% nitromethane to the fuel will lean the mixture by 19%. Hence the shorter initial period of combustion could be caused by

the leaner mixture and/or by the nitromethane. However, for methanol and other hydrocarbon fuels, the maximum increase in the laminar burning velocity for 19% increase in λ is of the order of 15% [60]. Given the cycle-by-cycle variations (limit bars) indicated in Figure 3.8 (left), five conditions ('3-Cold, Medium and Hot', and '4-Cold and Medium') definitely have shorter Q_{nmin}–10% MFB. All of these have about 40% reduction in Q_{nmin}–10% MFB. Assuming that Q_{nmin}–10% MFB reduces by the same order as the laminar burning velocity, then nitromethane certainly has speeded up the combustion. For the other four conditions (all needle setting '2' and '4-Hot'), the initial combustion period has changed by an insignificant amount given that the cycle-by-cycle variation is large compared with the changes in the mean values. Interestingly these four conditions also have the shortest initial combustion periods (~15°CA), as shown in Figure 3.7. For '0-4-Hot' (rich limit), the hot plug not only helps initiating the combustion, it also speeds up the combustion. In contrast, needle setting 2 may be close to or leaner than the stoichiometric condition, such that the benefit of nitromethane is offset by the slower combustion of the weak mixture.

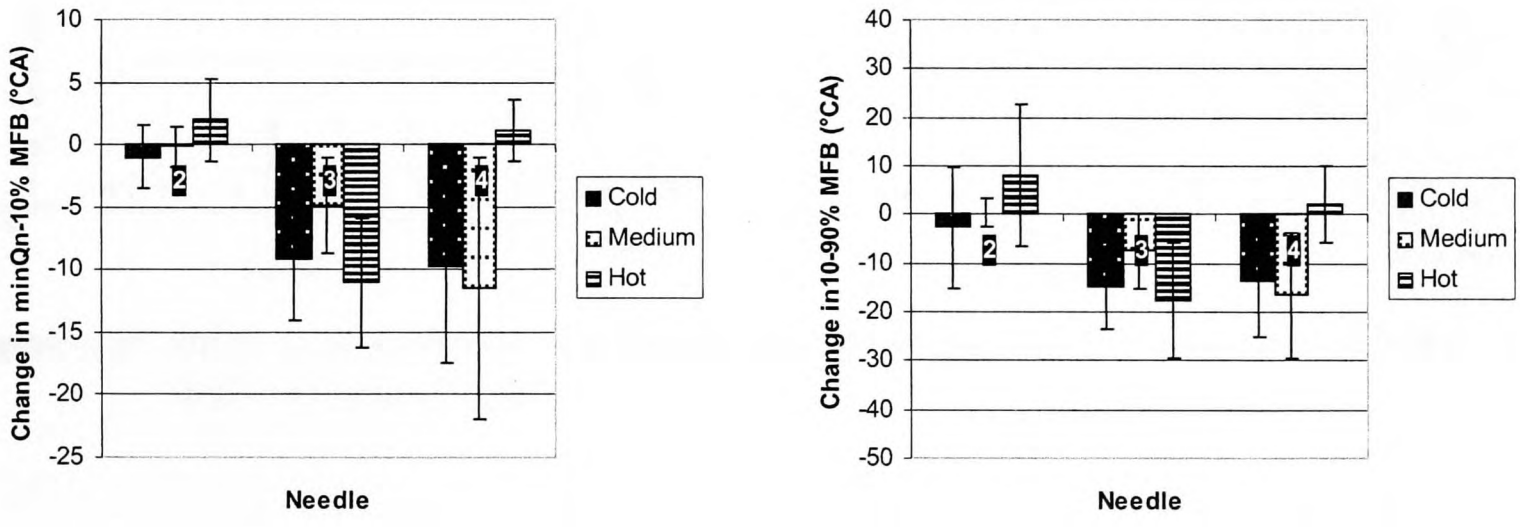


Figure 3.8: The effect of nitromethane on Q_{nmin}–10% MFB (left) and 10–90% MFB (right) for various needle settings and plug types; change relative to 0% nitromethane fuel

There is a similar picture for the main combustion period (Figure 3.8 (right)). For those conditions which have shorter initial combustion periods, the main combustion period

is also shorter. Therefore, nitromethane can also increase the rate of main combustion. The cycle-by-cycle variations in the main combustion period are larger than for the initial period. This is expected as other factors such as turbulence and the flame development also affect the main combustion period.

Two 3-D graphs of cycle-averaged data from all the engine operating conditions are shown in Figure 3.9, giving information about the effect of nitromethane content, fuel supply needle setting and glow plug type on IMEP and Q_{nmin} (SOC). In each graph, the x-axis (projecting into the page) presents the nitromethane concentration of the fuel. The y-axis (horizontal) presents the needle setting and the z-axis (vertical) gives the output parameter from the analysis of combustion data. Different symbols are used to represent the three different glow plug types.

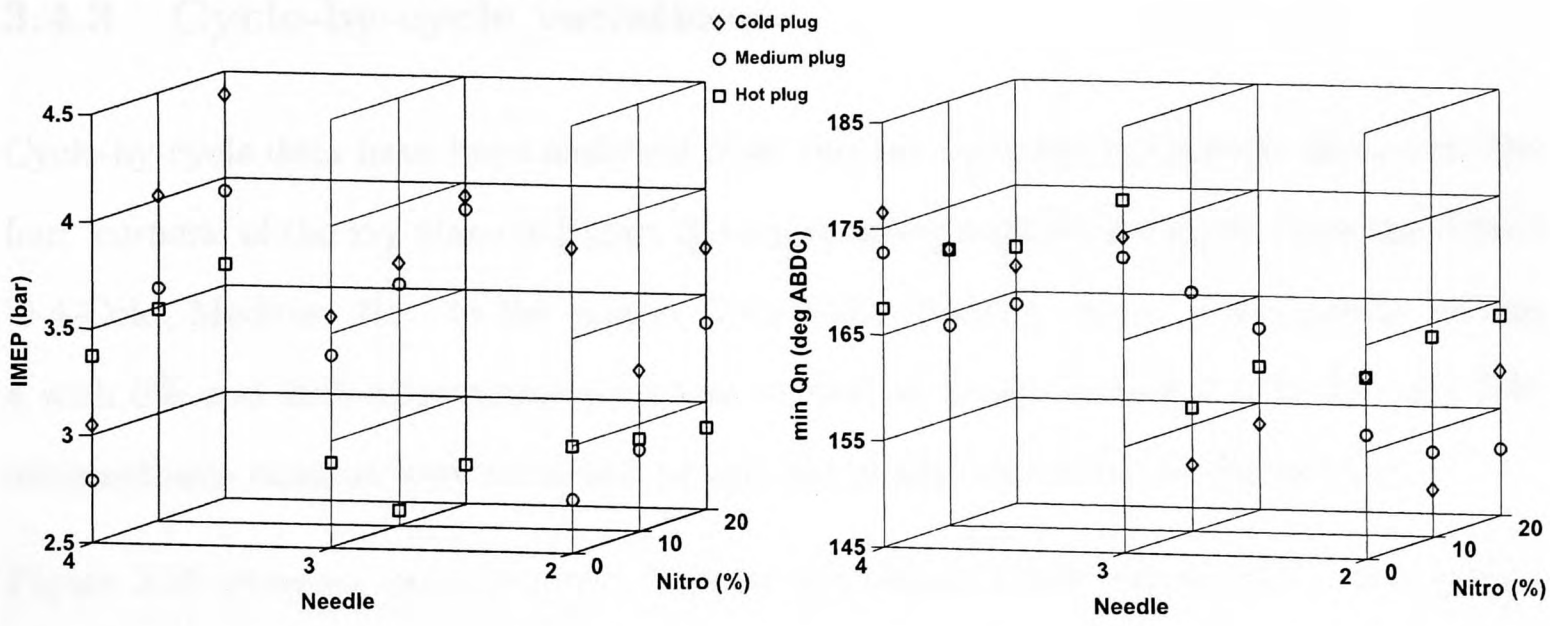


Figure 3.9: Effect of nitromethane content, fuel needle setting and glow plug type on IMEP (left) and Q_{nmin} (right)

Results presented in Figure 3.9 (left) suggest that IMEP generally increased with increased nitromethane content. This is likely to be the results of faster and more efficient combustion in the case of higher nitromethane content. Figures 3.8 have shown the faster combustion results. The combustion is more efficient because all operating conditions are rich relative to the stoichiometry, so burning at closer to stoichiometric point would have fewer incomplete combustion products and less unburnt fuel.

The results for the hot plug do not follow the trend well. The IMEP do not change much as the nitromethane content increases, especially for needle settings 2 and 3. The burn rate analysis has shown the hot plug either has little effect or has enhanced the combustion rate. The answer lies in the combustion phasing. The start of combustion (SOC or Q_{nmin}) in Figure 3.9 (right) shows that using cold and medium plugs, SOC is close to TDC (~170°CA) and occurs earlier when the nitromethane content increases, so the ignition time is closer to the MBT timing. However, when the hot plug is used, the SOC are generally early (~160°CA) before any nitromethane is added. The SOC stays approximately the same once nitromethane content increases. Hence there is no improvement in IMEP due to better ignition timing.

3.4.3 Cycle-by-cycle variations

Cycle-by-cycle data have been analysed from the four extreme fuel/needle data sets (the four 'corners' of the x-y plane in Figure 3.9 representing mixture strengths from the richest '0-4-Cold, Medium, Hot' to the leanest '20-2-Cold, Medium, Hot'). Thus, needle setting 4 with 0% and 20% nitromethane content as well as needle setting 2 with 0% and 20% nitromethane content were investigated and the results are shown in Figure 3.10.

Figure 3.10 presents cycle-by-cycle data for calculated IMEP versus SOC. The graphs show a characteristic of maximum IMEP against SOC, analogous to the MBT timing study commonly applied to spark ignition engines. Each graph consists of three operating conditions using different plugs. Particularly in the '0-4-Cold, Medium, Hot' graph, an operating envelope has been defined/sketched such that all the outlier points represent poor engine operating conditions. In this case, both cold and medium plugs show a much wider scatter of points than the hot plug. This clearly means large cycle-by-cycle variation in both indicated power and start of combustion. Apart from several misfire cycles showing very low IMEP, a noticeably high percentage of points in these two cases

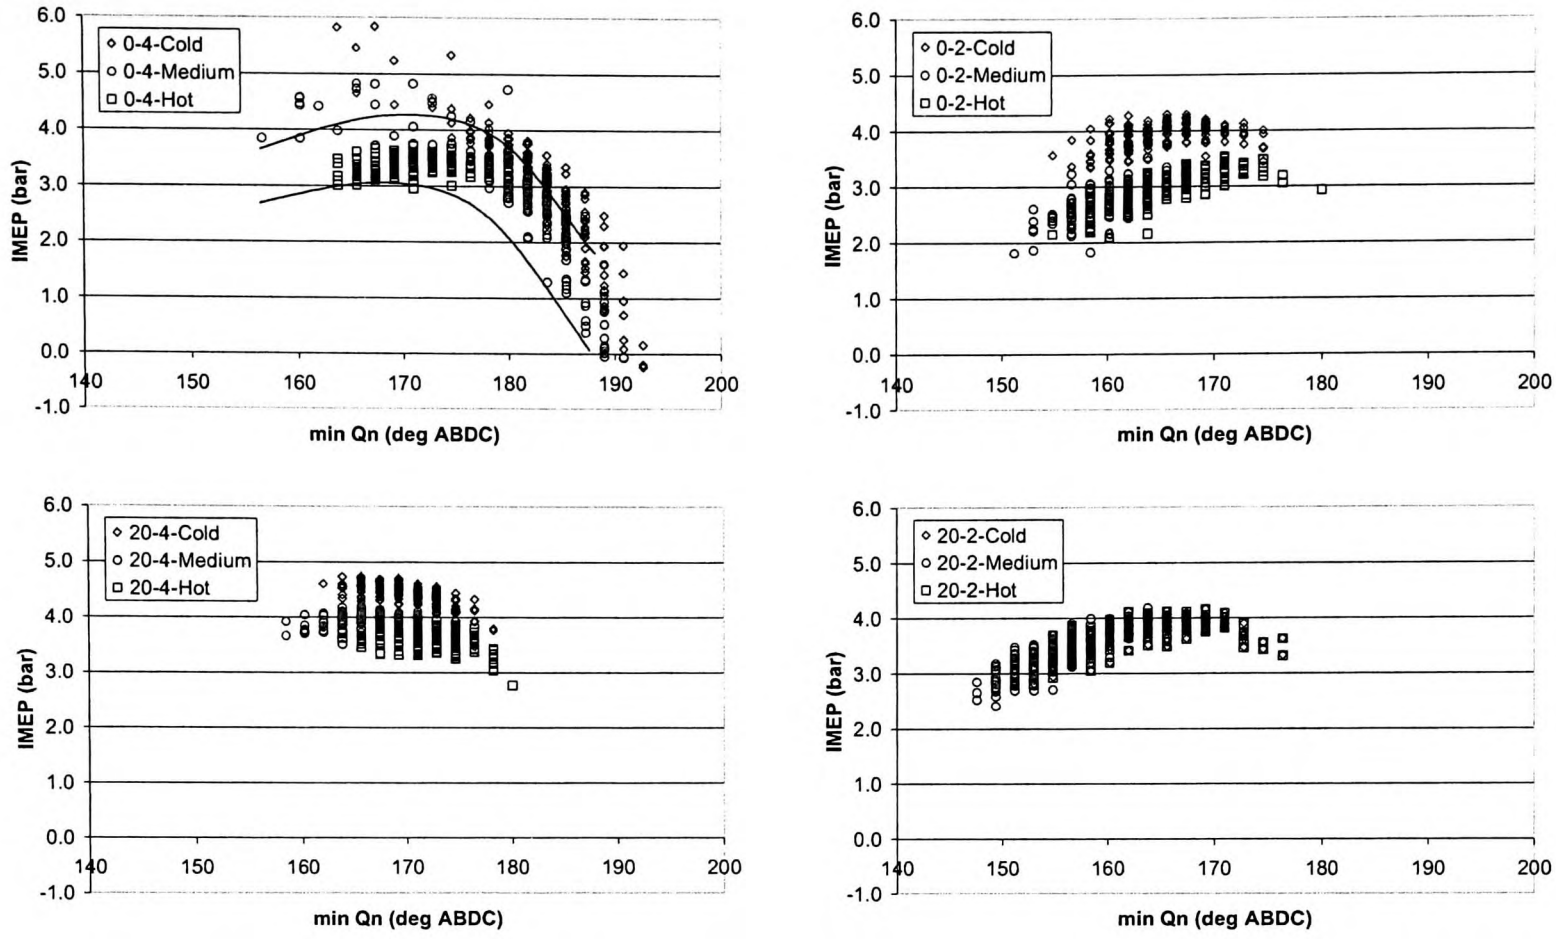


Figure 3.10: Cycle-by-cycle correlation between IMEP and Q_{nmin} from four fuel/needle data sets — needle setting 2 and 4 with 0% and 20% of nitromethane content

lie in the right bottom corner of the plot, confirming that partial and late combustion leads to low power.

Generally speaking, 20% nitromethane (the lower pair of Figure 3.10) gives a higher IMEP, less variation and earlier start of combustion than 0% nitromethane. Furthermore, the hot plug presents a relatively more stable range of SOC than the other two plugs. Therefore, to optimize the operating regime of this engine at wide open throttle, 20% nitromethane and hot plug would be selected as a best compromise.

Inspection of original pressure versus crank angle traces raise questions about how, in this type of combustion system, one cycle might affect the subsequent cycle. That is, if the n^{th} cycle has late or incomplete combustion, does this affect the $(n+1)^{th}$ cycle? In order to address this question, comparisons in terms of completeness of combustion (CoC) were made between two cases: '0-4-Cold' and '0-4-Hot', with CoVs of IMEP of 31.3% and 4.5% respectively.

Figure 3.11 gives the cycle-by-cycle correlation between the CoC and IMEP, with the x-axis presenting the CoC of the n^{th} cycle and the y-axis presenting the IMEP of the $(n+1)^{th}$ cycle. The left graph in Figure 3.11, which is for the case '0-4-Cold, with very large CoV of IMEP, clearly shows the trend that a partially burning (n^{th}) cycle (or even misfire cycles, when CoC = 0) is always followed by a cycle $(n+1)^{th}$ with higher than average IMEP (3.05bar). This is likely to be because for a partially burning cycle (probably slow burning as well) the trapped in-cylinder residuals tend to add to the air/fuel ratio and initial temperature of the following cycle which is then likely to perform a more complete combustion (*i.e.* better scavenging). The right graph gives a dense cluster of points randomly positioned on the right top corner of the plot. This means that the combustion of these cycles is relatively stable and complete as expected due to the low cycle-by-cycle variation.

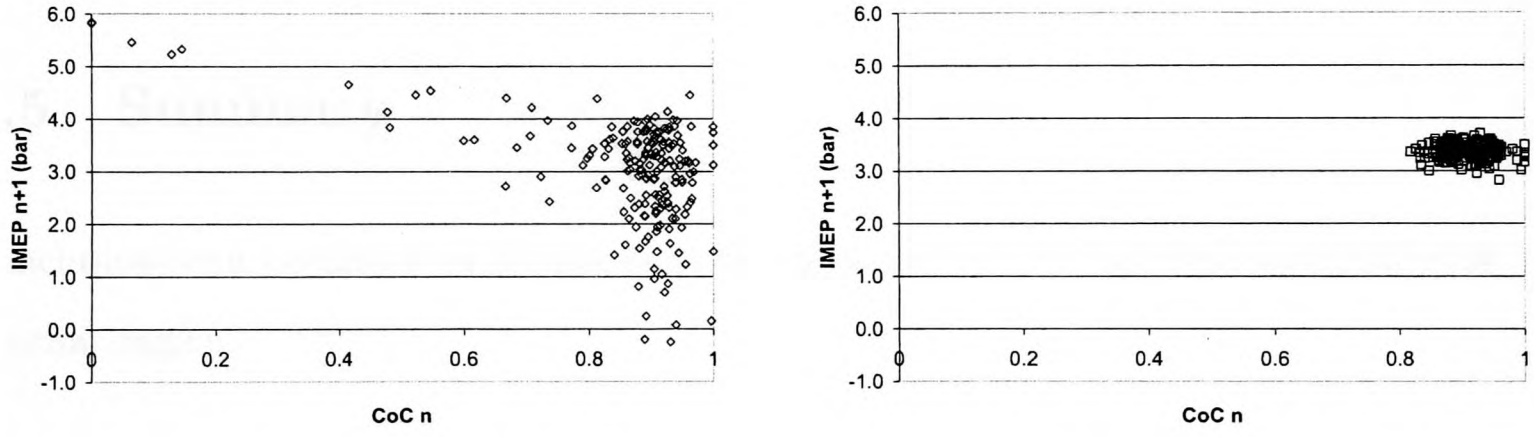


Figure 3.11: Correlation between the completeness-of-combustion of the n^{th} cycle and IMEP of the $(n+1)^{th}$ cycle — '0-4-Cold' (left) and '0-4-Hot' (right)

Figure 3.12 presents plots of cycle-by-cycle crank angle of Q_{nmin} and crank angle of 90% MFB from the above two sets of data (excluding misfire cycles in '0-4-Cold'). With Q_{nmin} and 90% MFB representing the SOC and end of main combustion, the left graph agrees with the perception that cycles with early ignition tend to finish combustion quickly as well. The right graph shows data for crank angle of 90% MFB for cycle n against crank angle of Q_{nmin} for cycle (n+1) and implies that a slow burning cycle (the right bottom wing of '0-4-Cold', marked with an ellipse) is more likely to trigger a following cycle with

early ignition. Data from '0-4-Hot' in the right graph supports the observation in the previous paragraph that stable combustion features a random and tight distribution of data points.

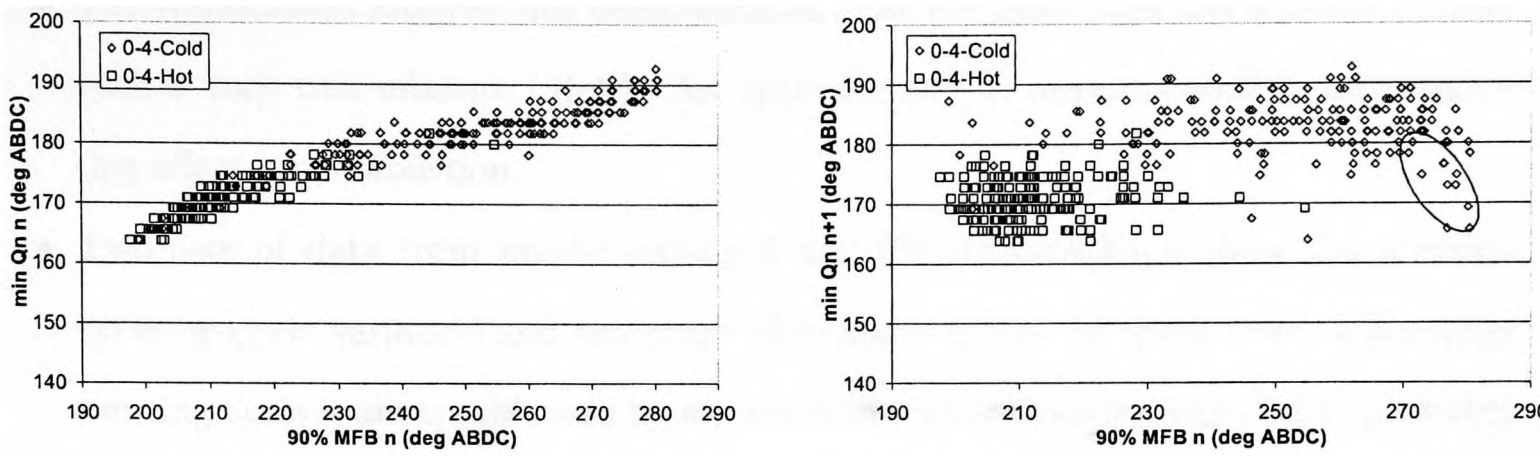


Figure 3.12: Scatter plot of cycle-by-cycle data — Q_{nmin} vs. 90% MFB of the same cycle (left) and Q_{nmin} of the (n+1)th cycle vs. 90% MFB of the nth cycle (right)

3.5 Summary

Conclusions can be drawn as follows from the work carried out on this two-stroke glow ignition engine:

- The burn rate and heat release analyses have produced useful results from the measured pressure data.
- The 'nominal' spark timing and ratio of specific heats specified in the program has no significant effect on the calculated results on the phasing of heat release.
- Nitromethane improves both the initial and main combustion burn rate of pure methanol, beyond the uncertainty created by the unknown mixture strength and cycle-by-cycle variations.
- IMEP generally increases with increased nitromethane content, only if the start of combustion occurs earlier so that the combustion phasing is approaching the optimum.

- Comparison with other engines shows that the combustion period of the engine under study is remarkably similar to traditional research engines, even though the engine design and operating regimes are quite different.
- The combustion analysis has demonstrated that the glow plug has a marked effect with a very rich mixture ('0-4'). An optical study is recommended to investigate this effect on combustion.
- Two sets of data from needle setting 4 and 0% nitromethane show the greatest cycle-by-cycle variation and extremely slow combustion. In these cases, a partially burning cycle is always followed by a cycle with higher than average IMEP, probably because for such a cycle the trapped in-cylinder residuals and initial temperature enhance combustion in the following cycle.

These results may provide insight of the performance and combustion process of such engines and facilitate more efficient applications.

Chapter 4

Development of Quantitative PLIF for AFR Measurement

What do you know and how do you know it?

— Ayn Rand (1905–1982)

4.1 Introduction

It has been shown that gasoline direct injection offers the potential to combine the high performance of a conventional port-injected gasoline engine but with greater fuel economy. However, the realisation of this potential requires greater control of the in-cylinder fuel distribution and mixture preparation processes than that required by a conventional port-injected engine.

Laser diagnostic techniques are widely implemented in fuel distribution measurement owing to their non-invasive nature, that is, they require very few modifications to the engine test rig and hence minimize the performance change of the engine. There is particular interest in planar laser-induced fluorescence (PLIF), because it produces a signal with

high spatial and temporal resolution which may be spectrally separated from scattered laser light. Numerous studies have been carried out using PLIF for visualising species of combustion and reacting flows [16] [54] [78]. However, few results have been justified to be fully quantitative in IC engine environments. The work in this chapter aims to develop a quantitative and reliable PLIF technique for the air/fuel ratio (AFR) measurement on the optical DISI engine.

The chapter starts with an exposition of the theory relevant to quantitative PLIF measurements in an engine. A dedicated multi-component fuel as well as the corresponding tracers have been designed to meet the requirements of the DISI engine experiments. The fuel distillation characteristics, construction and key features are then discussed. Such a fuel is shown to exceed the specifications of accepted single-component fuel/tracer set. Attention has also been drawn to the optimisation of both hardware and signal-to-noise ratio. Finally, a recirculating loop has been set up for the calibration of PLIF signal against AFR. Example calibration results from the first stage of tests are shown.

4.2 Principles of PLIF

Basic physics

In general terms, the laser-induced fluorescence (LIF) process is best viewed as a photon absorption, followed after some finite period of time, by spontaneous emission (fluorescence) from the excited manifold (Eckbreth [16]). The basic concepts of the technique can be shown by a simple three-state model as illustrated in Figure 4.1.

As illustrated in Figure 4.1, the ground state (S_0) and excited states (S_1 — singlet, T_1 — triplet) of a molecule are all represented by a system of discrete energy levels which are schematically shown here as a series of lines (only the vibrational levels are shown here, with the finer rotational levels ignored for simplicity). As the position of the levels rises,

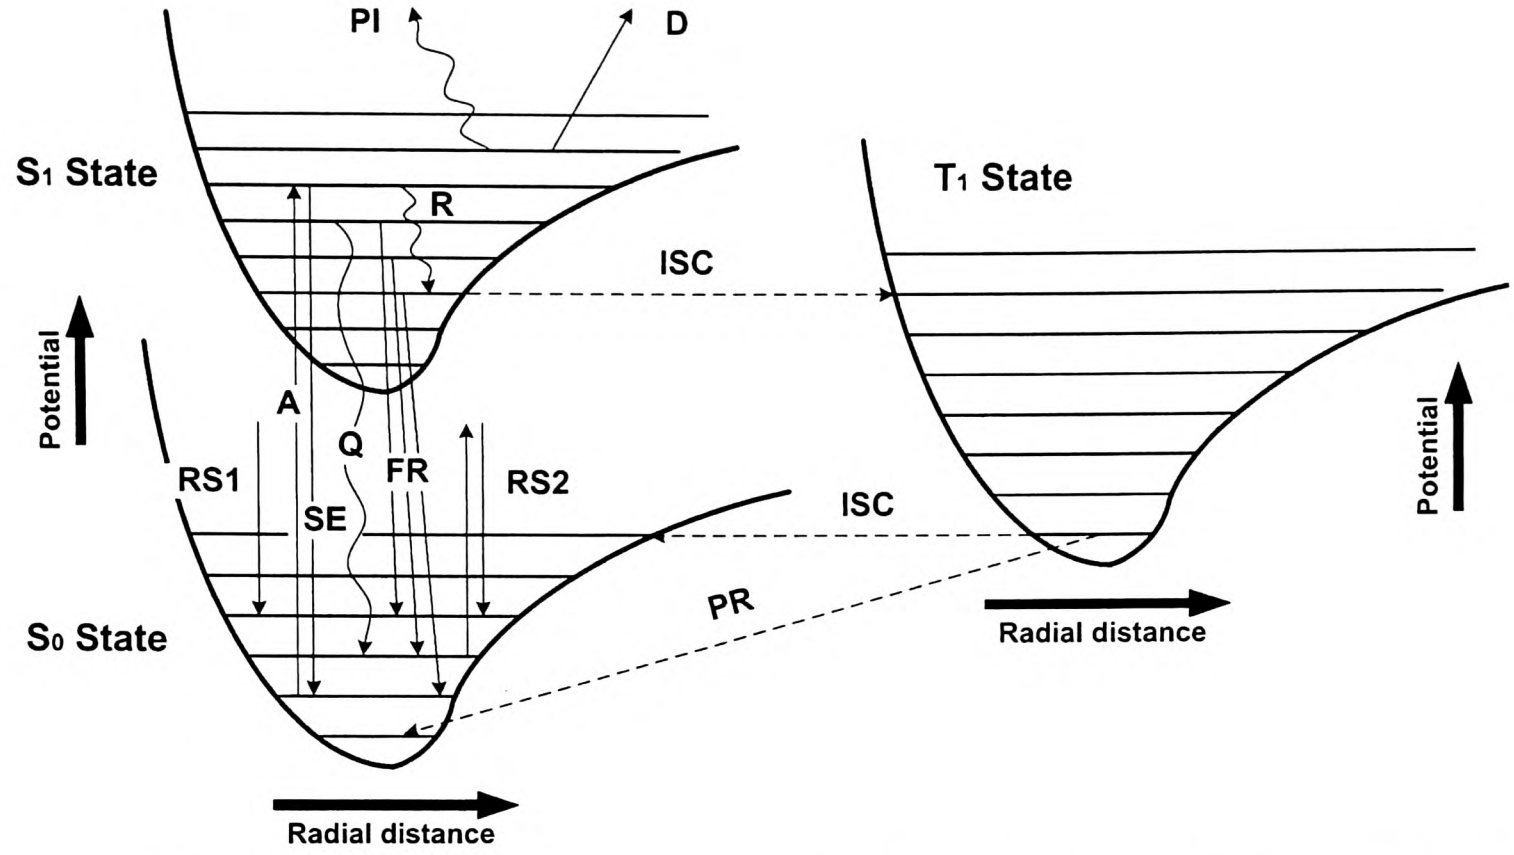


Figure 4.1: Three-state energy level model for electronic transitions: S_0 — ground state; S_1 — excited singlet state; T_1 — excited triplet state; ISC — intersystem crossing; RS_1 — Rayleigh scattering; RS_2 — Raman scattering; A — absorption; SE — stimulated emission; Q — quenching; FR — fluorescence (spontaneous emission); PR — phosphorescence; R — relaxation (collisional or vibrational); PI — photoionization; D — dissociation. Straight lines represent radiative transitions, whilst oscillatory lines represent non-radiative transitions.

the potential energy of individual levels gets greater. In the horizontal direction within each potential curve, the extension of the lines from left to right stands for an increasingly larger radial distance between the atom nuclei of the molecule.

In a process known as absorption, molecules in the ground state (S_0) are first pumped to the excited singlet state (S_1) by laser irradiance at a wavelength corresponding to the energy difference between the levels. Note that molecules cannot be excited directly to the T_1 state. All the triplet states must be formed from singlet states through intersystem crossing, because singlet-triplet transition is (almost) forbidden, because to a first-order approximation, a photon (absorbed or emitted) does not act on the angular momentum of an electron. Therefore, an electron moving between a singlet and a triplet state does not result in (direct) conservation of angular momentum. The excited molecules then return to the ground state via various energy loss processes *i.e.* de-excitation pathways.

Absorption

As the laser beam interacts with molecules, part of the energy, in the form of photons, will be absorbed. The molecules that have absorbed a photon are then pumped to an excited state. Absorption is the cause of LIF and tracer absorption is necessary for generating the signal; however, the inevitable attenuation of laser intensity throughout the medium under investigation due to absorption can lead to non-uniformity in terms of fluorescence production. Furthermore, absorption by other species in the medium , such as fuel components and alien constituents introduced during the process of transporting the fuel through pumps, pipes, injectors *etc.*, will add noise to the signal. In order to account for these aspects, it is essential to understand and quantify the spatial effect of absorption on the light sheet. More preferably, absorption from all sources should be minimised.

The concept of absorption cross section needs to be introduced to describe absorption quantitatively. In physics, the cross section, usually denoted with σ, is defined as the equivalent cross section of a particle as a solid object onto which other particles impinge. In the case of interaction between molecules and light, the two particle objects are molecules and photons. Absorption cross section may be used to define the probability of light passing through the medium that absorbs; it depends on the wavelength of the light. The relationships between absorption cross-sections and probabilities of absorption is instructively described by Hilborn [26] [27].

When the absorption level is low (unsaturated), Beer's Law (Eq. 4.1) can be applied:

$$I(x) = I_{x=0} \cdot \mathrm{e}^{-\sigma(\lambda)\int_0^L N_0(x,T)\mathrm{d}x} \tag{4.1}$$

where I is the intensity of light transmitted, $I_{x=0}$ is the input intensity, λ is the wavelength of the incoming light, L is the length of the optical path and N_0 the number density of

the medium. Usually the term absorbance is used to simplify calculation:

$$A = \log_{10} \frac{I_{x=0}}{I} = \sigma(\lambda) \int_0^L N_0(x,T)\mathrm{d}x \tag{4.2}$$

In a homogeneous medium ($N_0(x) = const.$, $A = \sigma N_0 L$), $I(x)$ features a simple exponential decay. In the general case however, $I(x)$ is a non-uniform function that requires knowledge of the species distribution. The distribution may be inferred from a PLIF image with known absorption. Therefore, it is necessary to seek a self-consistent value for N_0 that predicts the correct level of absorption and number distribution predicted by the PLIF images, with the assumption that there are no other effects on the magnitude of the fluorescence or absorption. When absorption is caused by multiple species and/or non-fluorescing species, this analysis is not valid. Furthermore, scattering can affect absorption by attenuating the laser beam along a given line. Significant scattering or reflection at a single point *e.g.* a droplet can direct laser light into various directions and complicate the local intensity contributing to absorption. Again in conclusion, minimisation of the total effect of absorption and scattering on a laser beam should be achieved wherever possible.

As in Eq. 4.1 the process of absorption is fundamentally temperature dependent. Temperature affects the average energy stored by a molecule, both as kinetic and internal energy. At a high temperature the distribution of electrons in states is spread over a larger number of levels than at a low temperature. So depending on the energy of the absorbing state, N_0 may increase or decrease with temperature. It is therefore important to excite tracer species at wavelengths which cause absorption from an energy level with a weak temperature dependence (further described in Section 4.4).

Finally, laser intensities that are too strong will saturate the absorption by the tracers and totally change the way of understanding the PLIF results. The absorption of a photon by a molecule causes the transfer of an electron from a lower state to an upper state and then the population of electrons in the lower state (N_0) has been reduced by one. This effect

will reduce the ability of the medium to absorb until the electron returns to the original state *e.g.* by fluorescence after some time t^*. Here t^* is the signal decay time constant, *i.e.* the time for the fluorescence to decay to 1/e of its maximum value (also referred to as fluorescence lifetime), which can be obtained experimentally through proper calibration. When there are so few electrons in the ground state that the absorption no longer follows Beer's Law, saturation occurs. The saturation intensity is given by:

$$I_{\text{sat}} = \frac{hc}{\lambda\sigma(\lambda)t^*} \tag{4.3}$$

where hc/λ is the energy of a photon (h is Planck's constant, c is the speed of light). The saturation intensity is a limit on the application of PLIF in the unsaturated or linear regime (detailed in Section 4.3).

Fluorescence

After absorption of laser irradiance, some of the excited molecules then return to the ground state via the spontaneous emission of photons, a process known as fluorescence. In order to measure the concentration of a specific species, the laser is tuned — normally in the UV spectrum — so that its wavelength corresponds to a known absorption band of the target species. Photon absorption from the laser beam only raises the energy state of the target species. The resulting fluorescence emission, which is normally Stokes shifted (red-shifted) from the excitation wavelength, is directly proportional to the mole fraction of the species and therefore proportional to the concentration. It should be noted that fluorescence does not necessarily occur directly from the pumped level, but that the molecules may relax collisionally or internally to a lower level prior to photon emission (transitions among vibrational or rotational levels, and hence Stokes shift occurs). A schematic of the relations among laser-induced scatterings and fluorescence is shown in Figure 4.2, where λ is the wavelength.

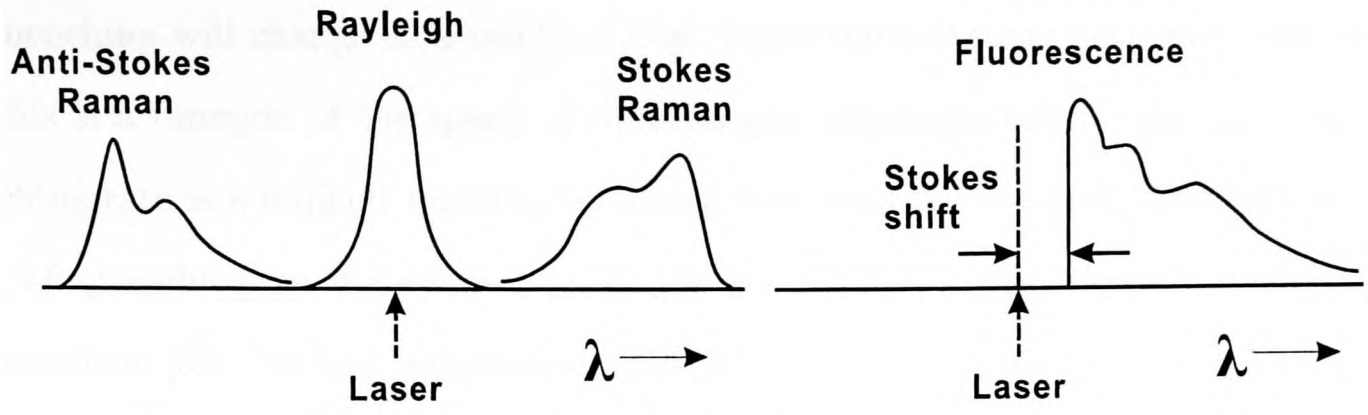


Figure 4.2: Schematic of laser-induced scattering and fluorescence

As shown in Figure 4.1, fluorescence is just one of the several energy loss processes available for the excited electrons. In addition to the absorption process, the laser excitation may stimulate a relaxation of some excited state molecules to the ground state. This is known as stimulated emission. Alternatively, an excited molecule may dissociate prior to the emission of a photon or, (if the laser irradiation is sufficiently high for it to absorb multiple times) it may become ionised. These processes are known as predissociation and photoionization respectively. Electronic transition from the T_1 state to the ground state is another radiative process termed phosphorescence, for which the lifetime is significantly longer than fluorescence and hence the signal intensity is relatively weak.

Collisions with molecules from the surrounding molecular bath may also cause an energy loss (radiationless, as heat) from the excited state. According to historical research, it is this final process, known as collisional quenching, that has most significantly limited the accuracy of quantitative measurements using the LIF technique (Eckbreth [16]). Since quantitative LIF is usually applied in the linear regime (when the absorption is weak, to be detailed in Section 4.6), *i.e.* the fluorescence intensity is linearly proportional to the input laser irradiance, quenching can reduce fluorescence intensities, especially by oxygen at elevated pressures, in a way that depends on the local pressure, temperature and composition, therefore making quantitative measurements in an engine much more difficult than qualitative fluorescence imaging. The quenching rate increases with pressure (as $\sim P^2$) owing to the change in the density of the gas; and with temperature (as $\sim T^{1/2}$ owing to the the change in the root mean square speed of the gas molecules. Furthermore,

the quenching will change as a result of how 'hard' the collisions between molecules are and this is a function of the speed of collision and the temperature. In conclusion, the quenching rate is a difficult quantity to model and predict. Practical illustrations of the effect of quenching on fluorescence yield are given for toluene with respect to species concentration [36] [38] and temperature [35] [37].

The fluorescence quantum efficiency (also termed fluorescence quantum yield, as denoted in curly brackets in Eq. 4.4) is defined as the ratio of fluorescence rate over the absorption rate including rates of all intramolecular de-excitation pathways and quenching among all species. Non-radiative processes of de-excitation of excited molecules hence make the fluorescence quantum efficiency less than unity.

At a certain excitation wavelength, the fluorescence signal intensity is given by

$$I_{\mathrm{PLIF}} = c \cdot I_{\lambda} \cdot N \cdot V \cdot B_{\mathrm{abs}}^{0\rightarrow 1} \cdot \left\{ \frac{k_{\mathrm{flu}}}{k_{\mathrm{total}} + \sum_{i} k_{\mathrm{q}}\left(p, T, N_i\right)} \right\} \tag{4.4}$$

where N is the number density of the fluorescing species in the observation volume V in which the local laser intensity I_{λ} interacts with molecules with a probability of $B_{\mathrm{abs}}^{0\rightarrow 1}$ causing absorption to occur in those molecules, and c is the factor concerning geometrical arrangements and properties of the detection system; k_{flu}, k_{total}, $\sum_{i} k_{\mathrm{q}}$ denote the rate of spontaneous emission, the rate for the sum of all intramolecular de-excitation pathways, and the overall rate of quenching (intermolecular) from surrounding species, respectively, p, T, N_i the pressure, temperature and number density of the mixture composition. It should be noted that this expression is only valid for sufficiently weak laser intensities, that is, when PLIF is applied in the linear regime.

Figure 4.3 shows the time-resolved correlation between the fluorescence yield and the laser signal. If the laser pulse is infinitely short (impulse) and assume the fluorescence only

occurs between two energy levels, the fluorescence yield is given by:

$$I_{\text{PLIF}} = CI_{\lambda}\mathrm{e}^{-t/t^*} \tag{4.5}$$

where t is time, and C is a physical constant derived from Eq. 4.4.

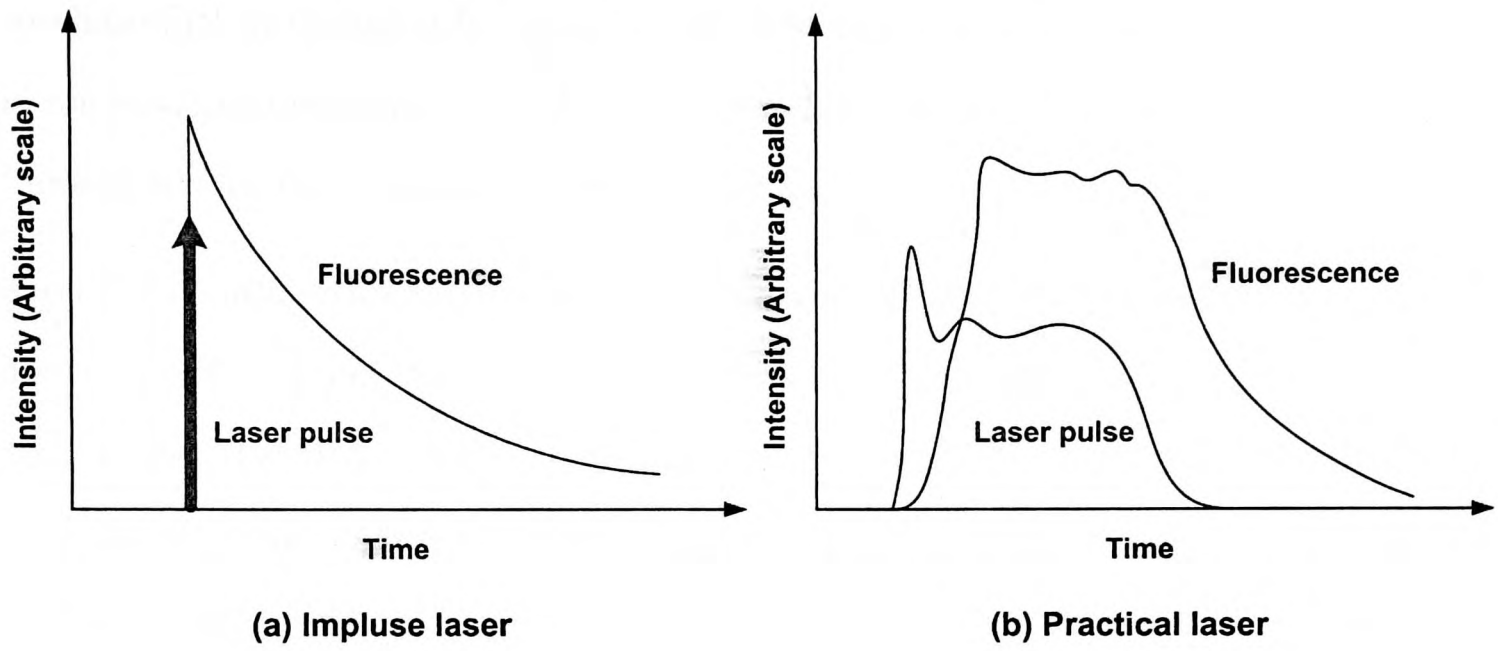


Figure 4.3: The time-resolved correlation between fluorescence yield and laser signal

When employing practical laser pulses, the fluorescence yield is the convolution of the laser pulse and the exponential term in Eq. (4.5), as the following:

$$I_{\text{PLIF}} = \int_0^{\infty} CI_{\lambda}(t-\tau)\mathrm{e}^{-\tau/t^*}\mathrm{d}\tau \tag{4.6}$$

where $I_{\lambda}(t)$ is the laser signal as a function of time.

Planar laser-induced fluorescence (PLIF) uses a single laser beam to form a sheet by a combination of spherical and cylindrical lenses (see Figure 2.2). This allows fluorescence measurement from a nearly 2-D sheet and thus the PLIF images imaged normal to the laser sheet are spatially resolved and do not suffer from line-of-sight integration problems. However, any undesirable noise along the line of sight will still be recorded by the camera.

4.3 Interpretation of PLIF images

PLIF strategies for GDI engines

In a GDI engine, the direct injection of liquid fuel into the cylinder inevitably results in a two-phase fuel distribution for much of the induction and compression strokes. Many of the conventional measurement techniques are either unable to distinguish between phases, or unsuitable for use in the presence of liquid fuel droplets.

Felton [18] identifies three strategies for PLIF imaging: natural fluorescence from the fuel, fluorescence from a molecular tracer (also referred to as dopant or fluor), and fluorescence from exciplex forming tracers. The first two strategies are used for fuels existing in either liquid form or gaseous form, while the third strategy, known as Laser-Induced Exciplex Fluorescence (LIEF), was proposed for fuel phase discrimination. Each of the three strategies has its own advantages and limitations.

Ideally one would want to undertake experiments directly on commercial fuels under production engine operating conditions. But the natural fluorescence from a commercial fuel such as gasoline is due to some uncharacterized components (usually the heavier and less volatile constituents) within the fuel and thus cannot represent the whole fuel. Another drawback of using gasoline as a fluor is its high absorption of UV light mainly due to its aromatic components.

Adding molecular tracers to an otherwise optically transparent fuel, known as tracer-PLIF, correlates the fuel distribution to the characteristics and concentration of the tracers only, if the tracer concentration in the fuel is known and uniform. In order to obtain fuel distribution, the tracer and fuel have to co-evaporate. However, more problems arise in the quantitative interpretation of the PLIF images. Meticulous matching must be made between the fuel and the tracers. To avoid natural fluorescence from commercial gasoline and maintain representativity, a 'model' fuel with close characteristics to gasoline

needs to be modelled. Tracers such as 3-pentanone or acetaldehyde will fluoresce with almost identical spectra in both the liquid and vapour phases, therefore preventing fuel distribution measurement in both liquid phase and vapour phase simultaneously.

In the LIEF case, a fluorescent excited molecule (M^*, such as a fluorobenzene) can react with a second molecule (G, such as diethyl-methyl-amine) to form a second emitting molecule M-G^*, the exciplex. The exciplex has an emission that is red-shifted with respect to the fluorescence wavelength of M^*, and the relative concentrations of the tracers can be adjusted to make the exciplex M-G^* the dominant emitter in the liquid phase, and the monomer (M^*) the dominant emitter in the vapour phase (Felton [18]). The liquid and vapour fluorescence emissions may be viewed separately using appropriate filters, allowing semi-quantitative interpretation of the fuel concentration in each phase. Quantification of the LIEF signal is nonetheless limited by several factors. The equilibrium of exciplex formation as well as the fluorescence yield is temperature dependent. In addition, the emission increment from the excited monomer may lead to a spectral overlap of fluorescence from the liquid and vapour phase. Furthermore, fluorescence can be severely quenched by oxygen and self-quenched, which is particularly serious for the vapour phase.

Quantification of PLIF images

For many tracer species at pressures found in an engine, the quenching rate is much faster than the fluorescence rate, causing the fluorescence quantum efficiency to be strongly pressure (and partly temperature) sensitive in a given medium. It is essential to have a value for the fluorescence quantum efficiency for quantitative PLIF. In the case of LIF, it is possible in some cases to measure the time dependence of the LIF signal, which will decay with this rate, $k_{\text{total}} + \sum_i k_{\text{q}}\left(p, T, N_i\right)$, exponentially. However, for PLIF it is not practically possible, with modern camera sensors, to record the time dependence of the signal in two dimensions. The commonly applied alternative is to apply a method to 'calibrate'

for the overall quantum efficiency of a tracer species in a medium of known composition at a given temperature and pressure. The calibration process generally involves recreating the range of conditions that will be seen experimentally in a comparable controlled environment with known temperature and pressure [14]. For each temperature, pressure and tracer concentration to be measured, a controlled measurement must be taken that determines the fluorescence quantum efficiency in that situation. These values may then be used to understand experimental data. For IC engines, the major unknown in this analysis is cylinder temperature. In many PLIF experiments the temperature dependence is simply ignored because it is difficult to measure. Therefore, it is important to select tracer species for which the temperature dependence of the fluorescence quantum efficiency is small.

The first calibration process of a PLIF image is the correction for spatial and temporal non-uniformity in the laser sheet intensity. The spatial non-uniformity is three dimensional because the laser sheet is not perfectly thin, with laser intensity varying through the thickness as well as along the width and length of the sheet. Calibration can be done by normalising the spatial laser intensity of the reference images (taken from a perfectly uniform mixture) to the pulse-averaged intensity value and then dividing the noise-corrected PLIF images by the 3-D normalised intensity profile. The temporal non-uniformity is due to the varying pulse energy of the laser, which can be corrected by recording the energy of all the pulses during which the PLIF signal is detected and recorded.

Secondly, depending on the term I_v/I_{sat} in Eq. (4.4), two regimes for signal interpretation can be used in practice, namely the linear regime and the saturated LIF regime (detailed in Section 4.3).

In the linear regime when $I_v << I_{\text{sat}}$, the fluorescence signal is linearly proportional to the input laser intensity. The signal must then be calibrated against ambient pressure and temperature in order to obtain the fuel concentration quantitatively (assuming the tracer and fuel have the same vaporising and transport properties), because high pressures

and temperatures facilitate the quenching process and temperature affects the absorption cross section as well.

For IC engines, oxygen is a major species and its collisional cross section is much greater than other colliders. If it can be assumed that: 1) oxygen quenching is the dominant de-excitation pathway at all temperatures and the quenching rate of the in-cylinder charge is proportional to the oxygen concentration at constant temperature; 2) the temperature-dependent variables in Eq. (4.4), *e.g.* absorption cross section and fluorescence quantum efficiency, can be experimentally characaterised and used to interpret the signal, then Eq. (4.4) becomes

$$I_{\mathrm{PLIF}} = C_{\mathrm{cal}} \cdot I_v \cdot \frac{[\mathrm{Tracer}]}{[\mathrm{O_2}]} \sim \frac{[\mathrm{Fuel}]}{[\mathrm{O_2}]} \sim \frac{1}{\mathrm{AFR}} \tag{4.7}$$

where C_{cal} is a calibration factor as a function of system setup, tracer/fuel composition and ambient conditions. This technique for direct fuel/air ratio (FAR) measurements is termed FAR-LIF and favours toluene as the tracer. However, Koban *et al.* [37] argued that for most of the pressure and temperature in the compression stroke of engines, assumptions for toluene FAR-LIF do not hold and the measurement case shifts between signal$\sim$FAR and signal$\sim N_{\mathrm{Fuel}}$ depending on temperature and excitation wavelength, that is, for practical engine applications, oxygen quenching is neither dominant nor negligible, and thus the effects of temperature and oxygen cannot be decoupled.

Another way to account for quenching in the linear regime is to measure the overall quenching decay rate (sometimes in the form of effective fluorescence lifetime) and include it in the equation for the fluorescence yield (Figure 4.4). Such experiments are susceptible to uncertainties of system setup and ambient conditions (p, T, composition) as well as limited by difficulties in signal gating and elimination of scattering *etc.* Nevertheless, literature reports effective fluorescence lifetimes of $\sim$10ns for liquid phase (heavily quenched) and $\sim$100ns for vapour phase (less quenched). This indicates the potential of simultaneous measurement of both liquid and vapour phase fluorescence separated with careful settings of the detection system (camera gating *etc.*) and post data analysis (dou-

ble exponential curve fitting *etc.*).

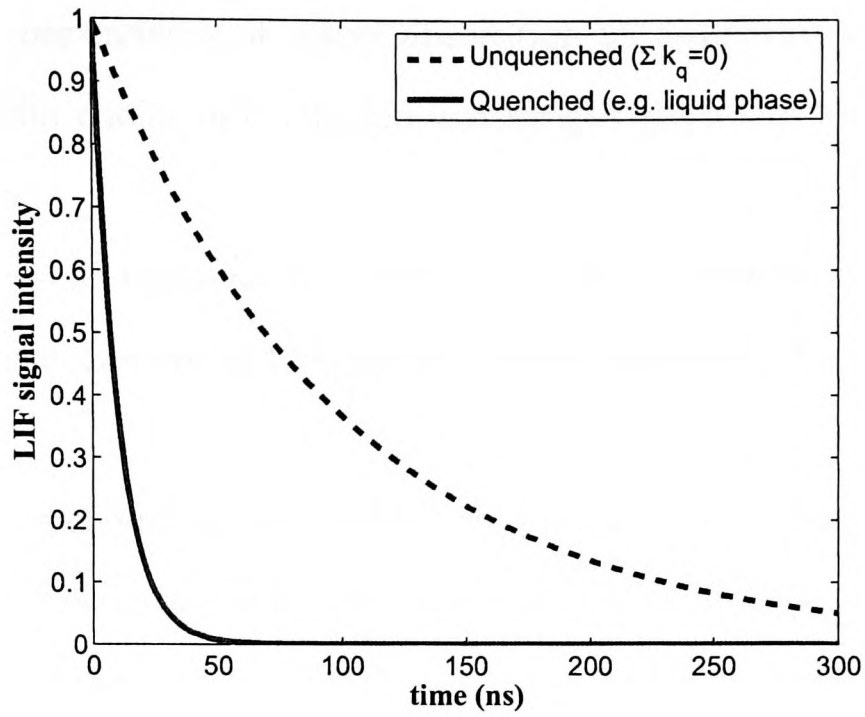


Figure 4.4: The time-resolved fluorescence signal, way to measure quenching decay rate and discriminate liquid and vapour phase

When saturated LIF regime is applied, *i.e.* $I_v >> I_{\text{sat}}$, the laser intensity is greatly increased such that the stimulated emission and absorption processes dominate over quenching. Here the fluorescence signal is independent of the laser intensity, thus the linear regime no longer exists. For a given tracer concentration, saturated LIF gives the strongest signal and therefore the highest detection sensitivity. Saturated LIF can be attempted (see discussion in [16] [58]) and can eliminate problems with non-uniform excitation across the laser path. However, it is not always possible to saturate the excitation over the limited duration of the laser pulse and at the edges of the laser beam. Also, the saturation intensity is often too large to be practically attainable with available laser sources. Even with sufficiently powerful lasers, the interpretation of the signal may be formidable when a highly uncertain combination of saturation LIF and linear LIF comes into the experiments, due to either the temporal intensity decay of the laser pulse or the pulse-to-pulse variation. In general, for a quantitative PLIF study one must ensure that the laser intensity is either far stronger or far weaker than the saturation intensity and apply analysis that is appropriate to the chosen regime.

Further steps are required to derive a 2-D fuel distribution out of the PLIF signal:

⇒ minimise the dependence of tracer absorption on temperature — this can only be done via careful choice of excitation wavelength given the limited flexibility in selecting tracers;

⇒ neglect the spatial temperature variation in the dependence of quenching, and assume a uniform pressure and temperature field because the point values are difficult to measure;

⇒ create a homogeneous fuel/air mixture with known temperature, pressure and composition (only based on AFR, otherwise impossible to identify), sample the corresponding fluorescence signal within wide yet realisable ranges and make look-up tables for subsequent local measurements of inhomogeneous AFR;

⇒ estimate the experimental errors in the quantification process *e.g.* fluctuations in the temperature, pressure field, and uncertainties of the detection system.

4.4 Equipment and device

This section briefs the key equipment and devices for the PLIF experiments in this project.

Intensified CCD camera

The fluorescence of many organic tracers is found in the UV region of the spectrum and is typically ~100nm wide. The spectral regions of intense fluorescence are therefore well separated from the thermal emissions of combustion, and may be isolated using conventional quartz filters. Furthermore, for these experiments it is necessary to employ quartz optics that can transmit the UV fluorescence emitted by organic tracers in order to maximise the detection efficiency.

PLIF measurements in this project used the Princeton Instruments ICCD camera (Thomp-

son UV-extended chip, 576×384 pixels) with a 135mm quartz objective (UV lens) from Coastal Optical. The ICCD camera is cooled by a Peltier pump to minimise dark noise and is gated using a PG200 high voltage pulse generator with a 1ns rise time and a gate width down to 4ns, thereby minimising signal interference from other optical effects. Short gate durations can also ensure that the images are captured during the time when the fluorescence is much more intense than the background light such as chemiluminescence and incandescence in the combustion environment being imaged. Besides the water cooling circuit needed for taking heat away from the Peltier pump, pure nitrogen was fed to the intensifier to avoid moisture condensation on the micro-channel plate (MCP).

As shown in Figure 4.3, there is a certain amount of temporal overlap between the fluorescence yield and the laser pulse. The fluorescence signal decays fast as well. The aim to both avoid the laser-induced scattering noise and maximize the available signal intensity necessitates a proper gate setting of the ICCD camera. Moreover, the gain and the gate setting (gate width, shape and timing) of the ICCD camera have important effects on both the signal and the noise level. Therefore, it is sensible to maximise the PLIF signal by a process of optimising these settings.

Excimer Laser

A Lambda Physik LPX-100 KrF Excimer laser was refurbished and installed in Oxford for this project. It is capable of producing 300mJ×10ns pulses at 248nm and at up to 300Hz, sufficient to perform saturated LIF.

This specification is ideal for exciting most organic tracers, and has sufficient energy for work demanding wavelength shifting and stimulated Brillouin scattering (SBS, for pulse cutting). Safety issues including fluorine F_2 leaking and laser shielding are of great concern for this system.

As mentioned in Section 4.5, absorption features of tracers like the ketone group vary

with temperature. Figure 4.5 depicts this phenomenon with acetone.

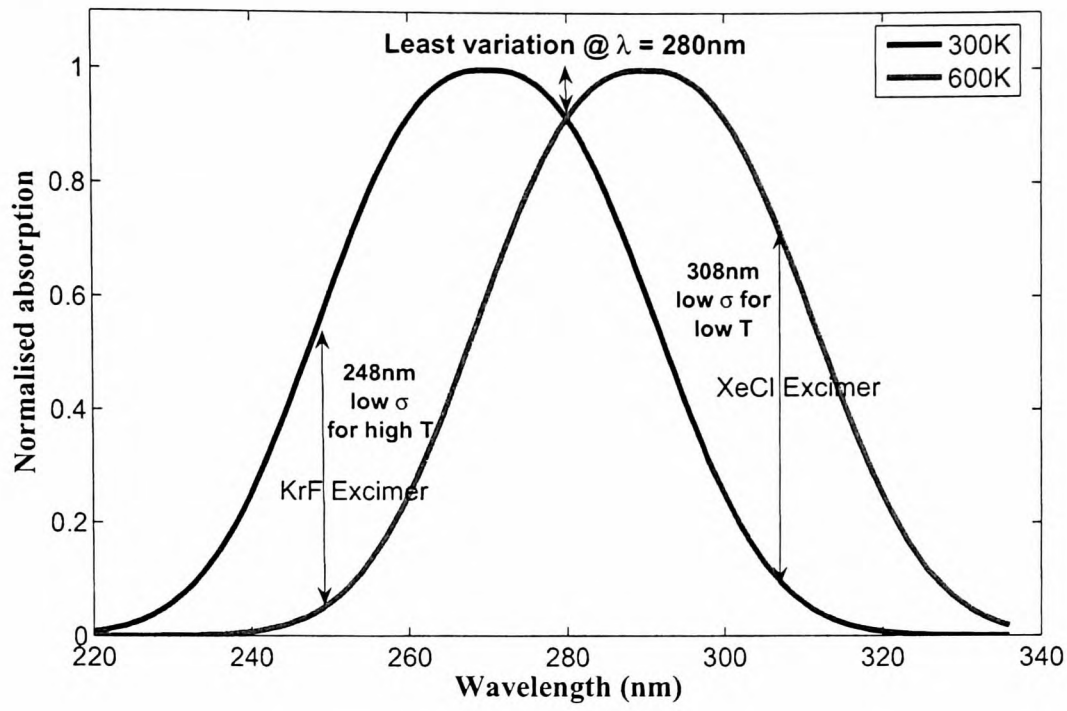


Figure 4.5: Absorption spectra of acetone varying with temperature

To minimise the change in the absorption spectrum, an excitation wavelength of 277nm was selected. This was obtained by feeding and focusing the Excimer laser light (248nm) into a Raman cell (Figure 4.6, designed and constructed by colleagues Robert Stevens and Yoshihiko Arita from the Department of Physics at Oxford) and it was shifted to 277nm via a process known as Raman shifting (stimulated Raman scattering, SRS). The cell was filled with pressurised H_2 of 30bar. Coherent photons interact with the hydrogen molecules and a portion of the energy gets dumped in the form of heat so that the wavelength is red-shifted. A series of discrete wavelengths (Stokes and anti-stokes shifted as well as the fundamental) emerge and a prism and masks can be used to select the appropriate frequency.

It should be noted that the absorption cross section of toluene is shown to be insensitive to temperature at 248nm but temperature-dependant above 260nm [54]. Moreover, toluene fluorescence peaks around 280nm. It is not sensible to excite toluene at 277nm because it would be impractical to separate the fluorescence signal from the laser wavelength. The spectral characteristics of more complex tracers such as trimethylbenzene have not been

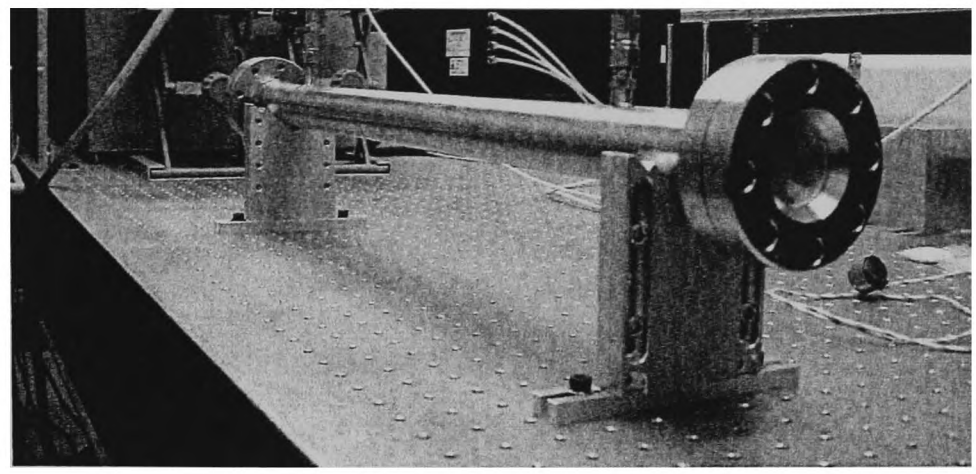


Figure 4.6: Setup of the Raman cell

well understood. For these reasons, The Raman Cell is not used for tracers other than acetone.

Spectrophotometer

A Perkin-Elmer Lambda 9 spectrophotometer (usable range 185–3200nm) was used to measure the absorption characteristics of the fuel components and tracers. The spectrometer inside produces two identical light beams, one passing through the sample and the other acting as a reference. These two beams then enter separate apertures and are collected by a comparator, which outputs the absorbance of the sample.

4.5 Design of component fuel

4.5.1 Selection criteria for fuel and tracers

Quantitative tracer-PLIF measurement depends most crucially on the assumption that the tracer behaves and moves as if it were the fuel itself, in constant proportion to the fuel, and without perturbing the fuel properties.

Real fuels *e.g.* gasoline contain fluorescent species, signal from which cannot be identified and interpreted quantitatively. Such fuels comprise components of different volatility —

each of which must be tracked independently (as separate fuels) to understand the in-cylinder mixture formation and combustion behaviour. Therefore replacement of the real fuel with non-fluorescing substances of known composition is required. This is termed 'component fuel' or 'base fuel'. Appropriate tracers can then be added to the component fuel each fraction of which co-evaporates with a corresponding tracer, and thus each tracer will identify the spatial distribution of that fraction in the environment of interest. This fuel must also have features as close to the reference commercial fuel (the Shell Dutch Pura for this project) as possible in terms of distillation, octane rating (RON and MON), calorific value, density *etc.* In addition, the fuel should be optically transparent, *i.e.* it should not absorb much laser light in the UV range.

For organic molecules, fluorescence is considerably stronger where π bonds are present: the π-bonding electrons facilitate the polarisation of the molecule. Aromatic (containing benzene rings) compounds are therefore strong fluors. So the base fuel components should be organic molecules without any π bonds or benzene rings *e.g.* alkanes, and the tracers should be molecules with π bonds *e.g.* alkenes, aldehydes, ketones and aromatics.

Ideally a base fuel should be matched with a large number of tracers each closely lined up with a distinct fraction of the fuel. However, not many naturally occurring compounds exist for much of the gasoline distillation curve up to normal boiling points of 350K. Above 350K many aromatic tracers exist and can be paired with heavy fuel component for co-evaporation, and this could be applicable to diesel fuels. The more volatile tracers are mostly oxygenates and the polar nature of their bonding means that they do not co-evaporate well with the non-polar hydrocarbons.

The fluorescence quantum efficiency is directly related to the fluorescence lifetime and both depend on the quenching rate. Ketone tracers have fast fluorescence rates (k_{flu} ~few ns) and are therefore affected little by quenching. Aromatic tracers tend to have much slower fluorescence rates (~100ns), and are therefore affected much more strongly by changes in the quenching rate. On the other hand, tracers with longer fluorescence lifetimes produce

fluorescence that is readily separated from scattered light (Rayleigh and Raman) which occurs instantaneously.

The main desirable properties of tracers are listed below:

- absorbing the available laser wavelength
- high fluorescence yield
- co-evaporating with the corresponding components in the component fuel
- non-toxic, stable and soluble in fuel
- sufficiently red-shifted to facilitate separation of elastically scattered laser light from the PLIF signal
- a fluorescence lifetime of 10ns–10μs for convenient detection
- different fluorescence wavelengths and/or lifetimes across the selected set of tracers for simultaneous detection

4.5.2 Fuel/tracer design

Pairs of single-component fuels and co-evaporating tracers have been in use for some time, even when little justification was shown to prove the tracer behaved reasonably closely to the fuel. 'Co-evaporate' in this sense means that the ratio of the tracer concentration to that of its fuel fraction should remain, as far as possible, the same in both liquid and vapour phases throughout the evaporation process, to within the accuracy of the PLIF measurement. The closeness of co-evaporation may be described by a 'co-evaporation ratio', C_{vap}, defined as the vapour-phase tracer to fuel (or fuel-fraction) ratio normalised by the tracer to fuel-fraction molar ratio overall (see Eq. 4.8).

$$C_{\mathrm{vap}} = \frac{[\mathrm{tracer}_{\mathrm{vap}}]}{[\mathrm{fuel-fraction}_{\mathrm{vap}}]} \cdot \frac{[\mathrm{fuel-fraction}]}{[\mathrm{tracer}]} \tag{4.8}$$

For systematic effects in the evaporation to be smaller than the typical precision of an engine PLIF measurement, the co-evaporation ratio should lie between 0.85 and 1.15 whenever more than 5% of a fuel fraction is in the vapour phase.

The component fuel used in this project was supplied by Shell Global Solutions and consists of 3 pairs of components. Stevens *et al.* [58] describe the process of modelling the component fuel and the tracer/fuel-component co-evaporation. Appropriate tracers may be added to the multi-component fuel, so as to co-evaporate with different fractions of the fuel such that each tracer will identify the spatial distribution of that fraction in the engine. The separate identification of different fuel fractions is essential for understanding the progressive vaporisation of the fuel and the way in which the variation in local air/fuel ratio develops. Issues of co-evaporation have been identified recently by a number of authors, in particular the co-evaporation of pure iso-octane reference fuel with ketonic tracer [22].

Early work to examine the effect of different fuel fraction led to the addition of tracers of different volatility to a single fuel, *e.g.* [64]. In these cases, quantitative interpretation was not possible given that the co-evaporation of the fuel fraction and tracer is not provable. Two-component exciplex tracer fuels have been modelled and used for quantitative PLIF diagnostics [62]. A three-component fuel has been demonstrated with ketonic/aromatic tracers [22]. This fuel is most similar to the fuel used in this project, yet this fuel/tracer set has been shown not to meet the criteria for co-evaporation as defined earlier in this section.

T_{10}, T_{50} and T_{90}, the temperatures at which 10, 50 and 90% of the fuel has evaporated, were nominated as the typical boiling points for the volatility classes of prime interest: *i.e.* a light (volatile), a medium and a heavy fraction. The temperatures of these classes for Dutch Pura (from an ASTM curve) are given as ~60°C, 100°C and 170°C. This provided a basic scope for selecting the tracers.

A comprehensive review of tracer-LIF diagnostics has been presented by Schulz *et al.* [54] including extensive discussions on the properties of organic tracers. The two most frequently reported tracers are acetone and 3-pentanone. 3-pentanone gives a fluorescence yield a few times stronger than acetone and has a boiling point (102°C) close to iso-octane (99.2°C, most representative compound of gasoline), however, it only co-evaporates with iso-octane when in a mixture of 20% 3-pentanone [58] and it degrades in the mixture. The high vapour pressure of acetone (boiling point 56°C) makes it an ideal tracer for the light fraction. A feature of the ketone family is that the absorption spectrum of acetone shifts red-wise as the temperature gets higher; this effect has been accounted for and discussed in Section 4.4. Sercey *et al.* [14] used gas chromatography to study the fuel breakdown phenomenon and found that iso-octane degraded with acetone present in the mixture in a motoring engine environment. Detection of acetone fluorescence also suffers from its low quantum efficiency and short lifetime. Nevertheless, the start of the distillation curve of a typical gasoline (in the region up to 80°C) is only covered by a small set of aldehydes and ketones, the least inappropriate fluor of which is acetone in terms of overall benefits. Additionally, to co-evaporate with the light fraction, acetone is the best choice. Figure 4.7 describes the evaporation of iso-octane with either acetone or 3-pentanone and shows considerable deviation from co-evaporation. The addition of the ketones to the the mixture increases the overall volatility of the mixture by around 10% at ~303K [58].

Toluene is another widely studied tracer species for combustion research. It is present in gasoline and hence was selected to match T_{50}. For the heavy fraction, there is a large group of suitable aromatics, so 1,2,4-trimethylbenzene (TMB) was chosen simply for a high fluorescence yield and a close match of boiling point to T_{90}.

To find a fuel blend that gives satisfactory co-evaporation, some preliminary vapour liquid equilibrium (VLE) calculations were performed by colleagues from Shell Global Solutions with an in-house Shell program that uses a modified Redlich-Kwong equation of state and liquid activity coefficients derived from the UNIFAC group contribution method [58].

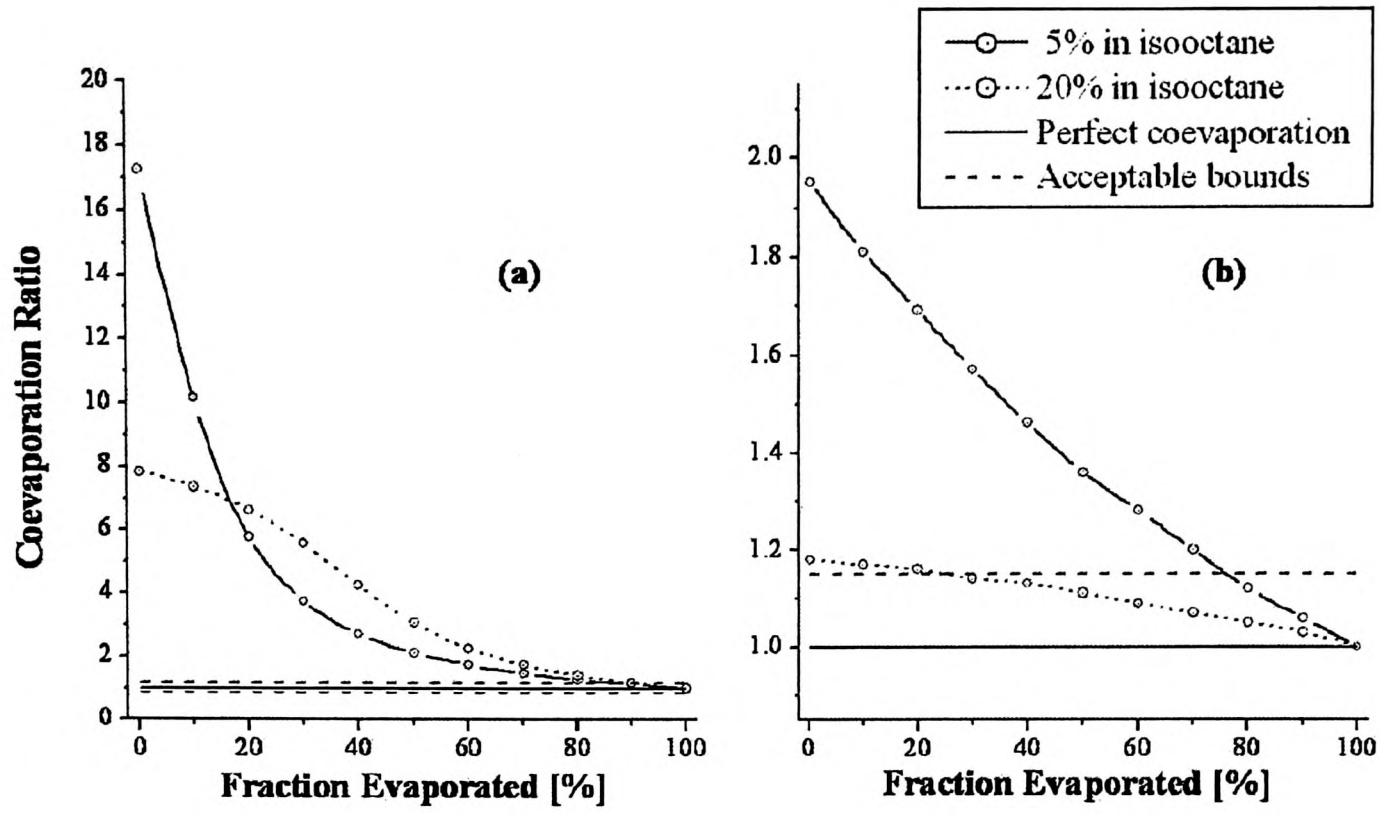


Figure 4.7: Coevaporation of (a) acetone/isooctane, (b) 3-pentanone/isooctane mixtures at 303 K [58]

With the chosen three tracers, a three-compound blend was initially identified by the VLE calculations such that each compound co-evaporated acceptably with one of the tracers. Unfortunately, these compounds could not be sourced economically in practical quantities. An iterative design procedure was then taken with a six-compound multi-component fuel blend, using three pairs of readily available compounds such that one compound from each of three pairs was more volatile, and the other less volatile, than the corresponding tracer species. By selecting suitable pairs and adjusting the proportions, a composition was found for which, according to the VLE calculations, each pair evaporated in the same proportion to the tracer to within the required tolerance.

Table 4.1 and 4.2 summarise the properties of the component-fuel components and selected tracers. On blending the six-component mixture, it was found that n-butane was most readily available as a mixture with other C_3 and C_4 compounds. So a final small adjustment was made to the blend for this light fraction mix.

Figure 4.8 shows the co-evaporation ratios for the multi-component fuel. The co-evaporation

Species		$T_{boiling}$(°C)	Density (g/cm^3)	Volume fraction (%)
Light	butane (mix)	-1	599 (25°C)	11.7
fraction	iso-pentane	30	626.2 (20°C)	18.7
Medium	iso-octane	99	691.9 (20°C)	44.1
fraction	n-octane	127	703 (25°C)	10.9
Heavy	n-decane	174	726.3 (25°C)	2.2
fraction	iso-dodecane	~180	754.6 (25°C)	12.4

Table 4.1: Properties of component-fuel components

Species	Family	$T_{boiling}$(°C)	Density (g/cm^3)	LIF score	LIF band (nm)	LIF lifetime (ns)
acetone	ketones	56	789.9 (20°C)	1	300–600 (liquid) 350–550 (vapour)	4.0
toluene	aromatics	111	866.9 (20°C)	~700	270–330	177
TMB	aromatics	168	880 (20°C)	~900	270–350	77

Table 4.2: Selected common tracers and their properties taken from the literature. Note: spectral properties vary considerably with the environment and phase — the values also vary considerably in the literature. Therefore, these values should be taken as a guide only.

ratios (with respect to each fuel component pair) of the tracers across the range of fractions evaporated compares favourably with the chosen criteria and the accepted pure component fuel co-evaporating pairs. The co-evaporation ratio for the TMB tracer is relatively high for small fractions evaporated (>1.15 at <10% vapour). This could be corrected by using more iso-dodecane in the blend but this would be of little benefit as less than 5% of the heavy fraction is vaporised until the overall fraction evaporated exceeds about 50%. The region where more than 5% of the heavy material is evaporated is marked by the bold segment of the line in Figure 4.8. The stated blend gives the best calibration in this region.

It has been shown [58] that there is little variation in the overall fuel volatility when 5% tracers are added to the component fuel.

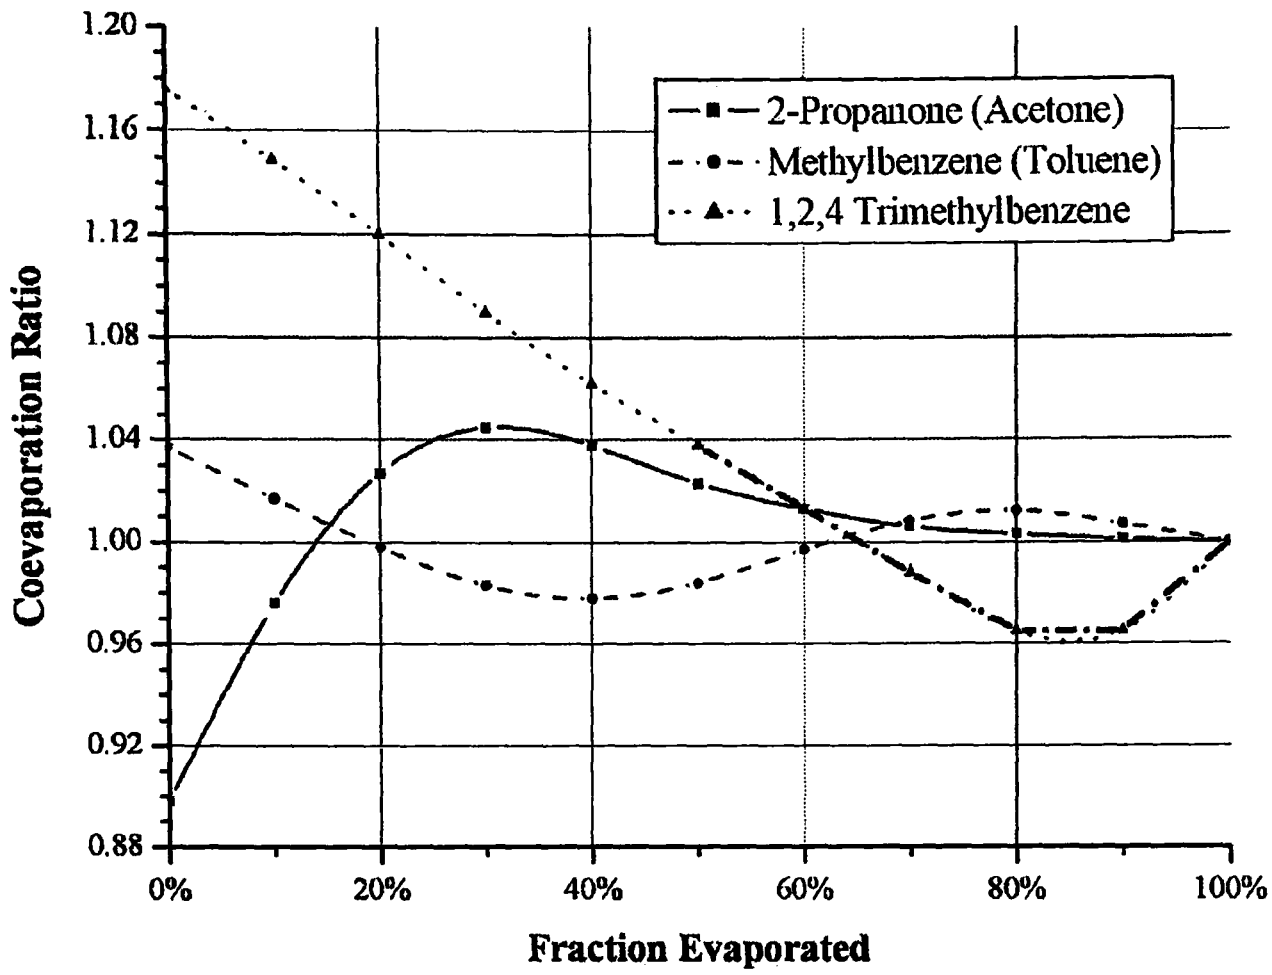


Figure 4.8: Coevaporation ratios at 303K according to equilibrium VLE calculations. The bold region of the trimethylbenzene curve shows the region where greater than 5% of the heavy fuel fraction has evaporated [58]

4.5.3 Characterisation of fuel components and tracers

Preliminary LIF tests on sample component fuel showed that it generated noticeable LIF signals, which were considered to be from contaminants not removed in the distillation process. Therefore, two sets of experiments, namely absorption and LIF spectroscopy with a sample cell, were conducted on all the liquid tracers and fuel components in their liquid phase. N-butane is gaseous under atmospheric conditions and is less likely to contain contaminants that fluoresce, since it is the higher boiling point hydrocarbons that are most likely to fluoresce.

In the absorption experiments, a 10mm thick cell was used and in order not to saturate the spectrophotometer, samples of pure fuel components and solutions of 1% acetone, 0.1% toluene and TMB in spectroscopically pure iso-octane (HPLC grade) were used.

Results from the first batch of component fuel and the components supplied by Shell are presented in Figure 4.9. The 'Shell component fuel' plot is from a sample of the actual shipped batch; the 'Oxford fuel' one is from a sample mixed up at Oxford with the five shipped fuel components according to the specified composition, and the 'Calculated fuel' is derived from the weighted sum of the five components using Eq. 4.2 to give

$$\text{abs}(\%) = 100\,(1 - I/I_0)\,\% = 100\left(1 - 10^{-\sum_i c_i A_i}\right)\% \tag{4.9}$$

where c is the relative quantity of component i.

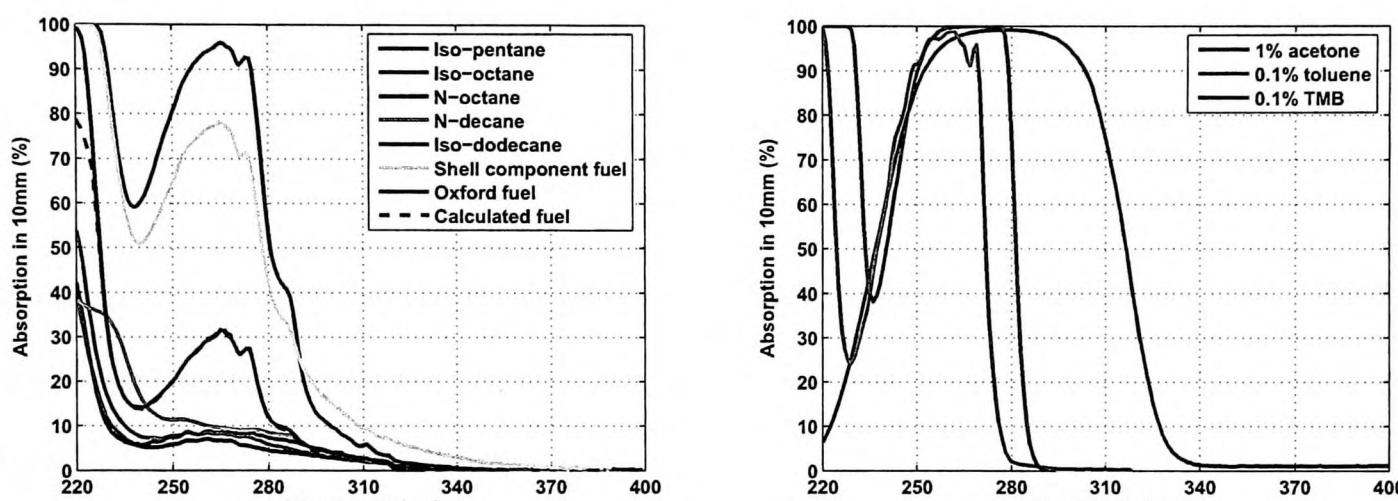


Figure 4.9: Absorption spectra of first-batch components and tracers

It is apparent in the left figure that iso-dodecane and n-decane both have considerable contaminants of some sort. This together with the composition of the component fuel makes the contribution from iso-octane and iso-dodecane more significant than the others, and hence a request for a 'purer' version of these two components was made. The 'Calculated' spectrum matches the one from the Oxford mix both in shape and in magnitude but only in shape with the Shell mix, which indicates: 1) the actual composition may differ from the nominal recipe or 2) more contaminants of a similar sort were introduced in the mixing process at Shell.

Figure 4.10 shows the absorption spectra of the second batch of iso-octane and iso-dodecane. With the quality of the new iso-dodecane much improved but that of the

new iso-octane even worse, the old iso-octane and new iso-dodecane were used to predict the spectrum of the new component fuel, which is seen to be reasonably 'transparent' (left hand side figure, ~10% absorption around 270nm).

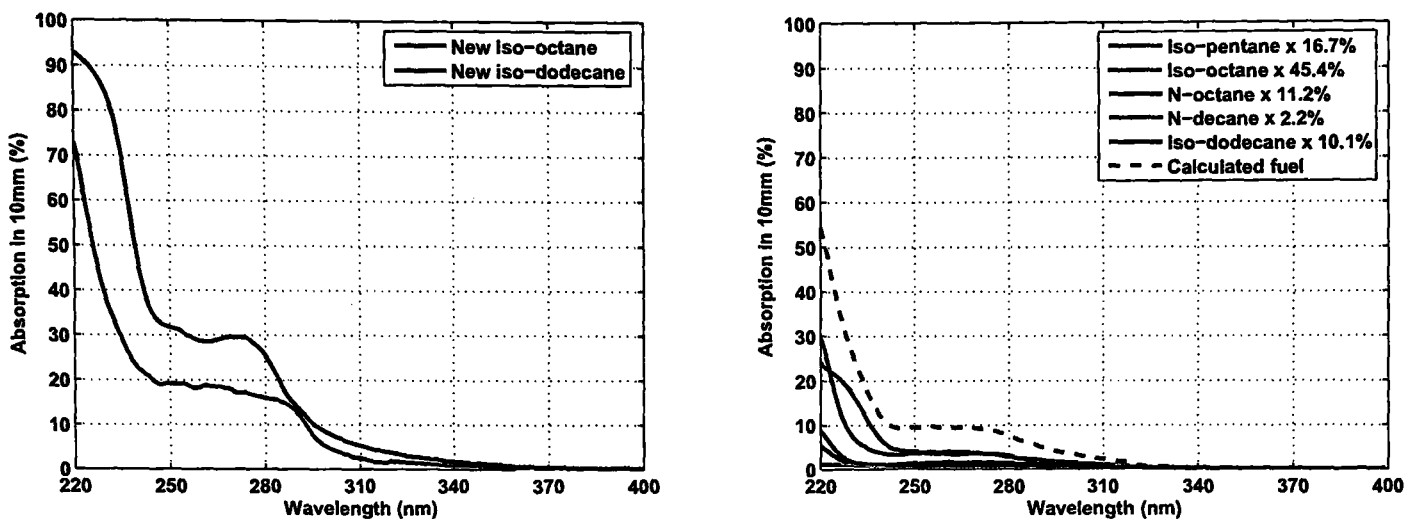


Figure 4.10: Absorption spectra of second-batch components

Knowing the absorption features of the fuel components and tracers, it was advisable to investigate how much fluorescence the component fuel might produce and to estimate the signal-to-noise ratio (SNR) with respect to the LIF signals from the tracers. Figure 4.11 depicts the layout for the LIF experiment.

The beam from the KrF excimer laser was focused into a cell filled with the sample in question, the same as was used in the absorption tests. The ensuing fluorescence was led through a spectrometer for the spectroscopic study. Monochromatic light was then detected by the photomultiplier tube (PMT). The motor on the spectrometer was set to run at 25nm/min so as to automate the spectra acquisition. Due to variations in laser intensity, a photodiode (ThorLabs DET210) was used to collect a portion of the laser light for data normalisation. The boxcar integrator acted as a sample-and-hold device and took both the laser pulse and fluorescence signal over a short predefined gate, longer than the fluorescence lifetime. Data of certain cycles within the hold time (~100ms) were averaged and logged by the DAQ system.

Data with the cell empty was subtracted from all other data sets as background and

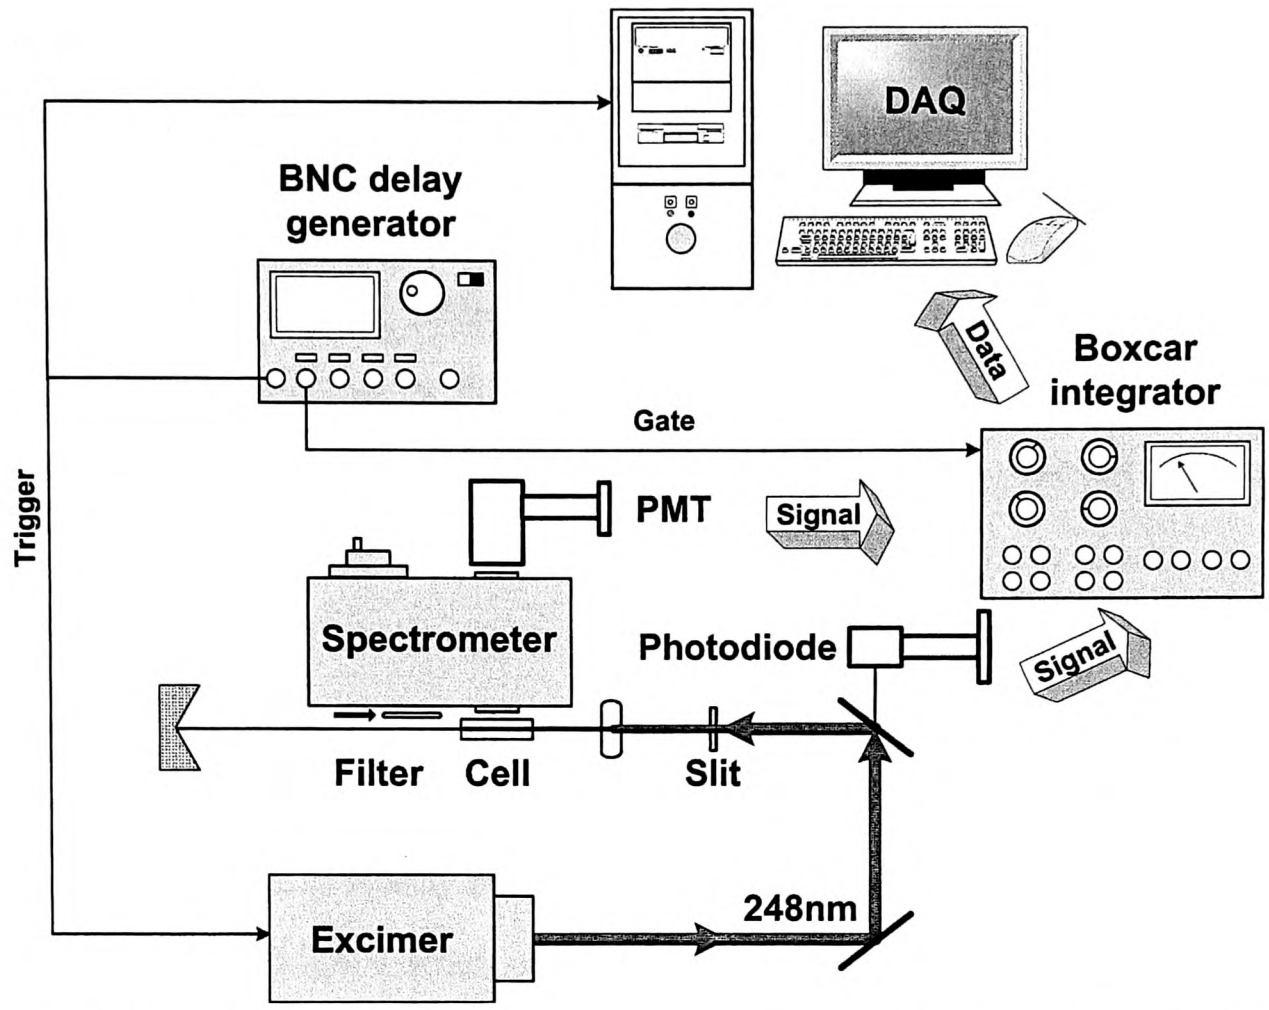


Figure 4.11: System for measuring the fluorescence spectra of fuels and tracers with a spectrometer and photomultiplier tube (PMT)

data from the photodiode were used to normalise those from the PMT. Systematic errors such as laser scatter within the spectrometer and second-order diffraction effects from the grating were observed as small local peaks on the spectra. Data points within these regions were discarded.

The first few figures in Figure 4.12 show the fluorescence and absorbance spectra of the tracers, components and component fuel, with the last being the weighted sum of the five components. Both signals are proportional to molecule number density. It can be seen that most of the fuel components gave fluorescence of similar level to acetone because the contaminants in them are more efficient at fluorescing. This is particularly true for iso-pentane and n-octane and has made the predicted fluorescence from the component fuel significant. However, as the fluorescence spectrum of the component fuel is located sufficiently away from that of acetone, optical filters can be used. The substantial LIF signal from the original iso-dodecane proved that use of the new batch was essential. The

two iso-octane samples showed similar fluorescence characteristics, despite the difference in their absorption properties. The same was true for TMB and toluene.

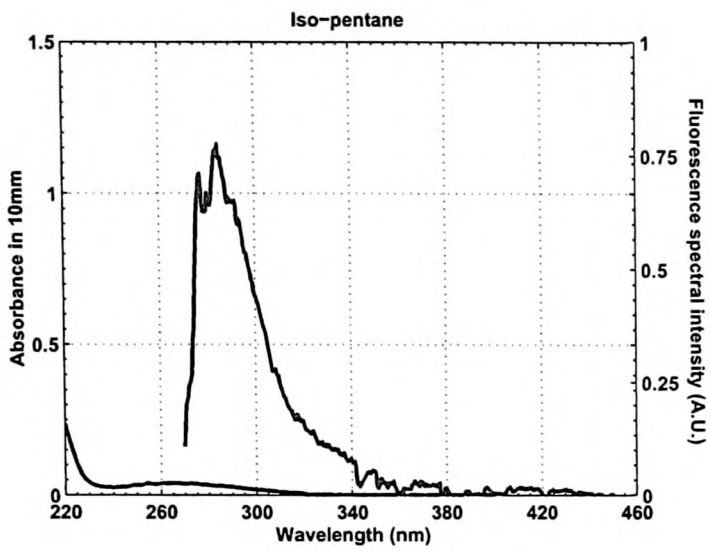
To further understand the signal-to-noise ratio based on the fluorescence from the component fuel and to be able to estimate the desirable concentrations for the tracers, the aforementioned LIF experiments were repeated with two 'long-pass' optical filters, namely WG280 and WG360 (UQG Optics, 50×50×2mm) which have nominal cut-off wavelengths of 280nm and 360nm respectively. Their responses from the spectrophotometer are plotted in Figure 4.12. The fluorescence responses from these tests are also shown in Figure 4.12. By putting one of these filters between the cell and the spectrometer, most systematic errors such as laser scatter can be removed. With the WG360 filter the fluorescence from the component fuel is mostly filtered out and the SNR for acetone can be maximised, likewise for toluene and TMB when using the WG280 filter.

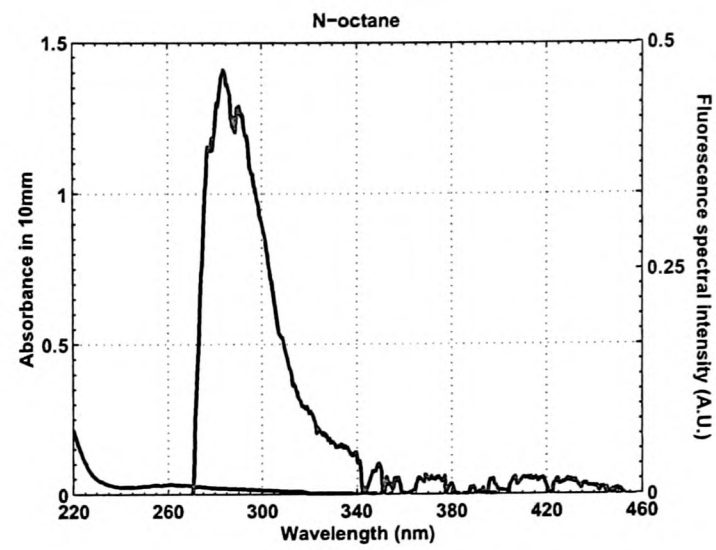
The SNR was calculated as the ratio of the spectrally integrated fluorescence spectra of the tracer and component fuel:

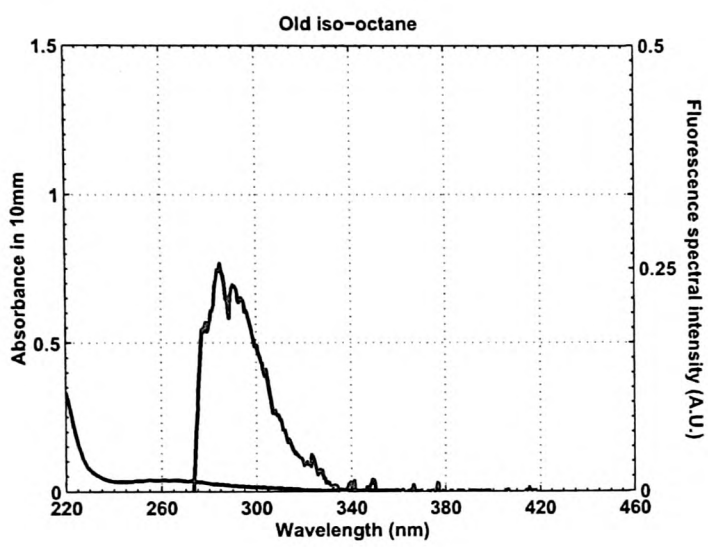
$$\mathrm{SNR} = \frac{\int_{\lambda_1}^{\lambda_2} I_{\mathrm{LIF,tracer}}(\lambda)\mathrm{d}\lambda}{\int_{\lambda_1}^{\lambda_2} I_{\mathrm{LIF,fuel}}(\lambda)\mathrm{d}\lambda} \tag{4.10}$$

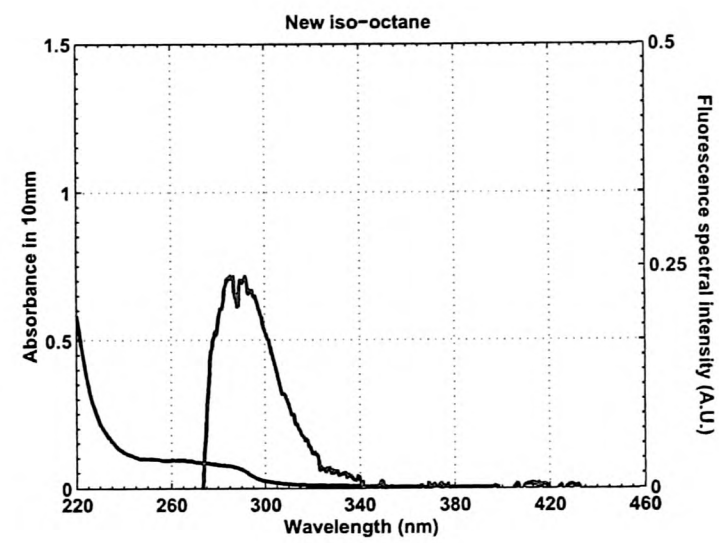
Data for other tracer volume fraction were simply derived by scaling the experimental data. Figure 4.13 presents the SNR results for fuel-tracer mixtures in both the liquid and vapour phase. The SNR limit was obtained when there was no less than 80% transmission through the vapour-phase mixture in the cylinder (89mm bore) at a particular pressure and temperature. The following assumptions were made in the vapour phase analysis:

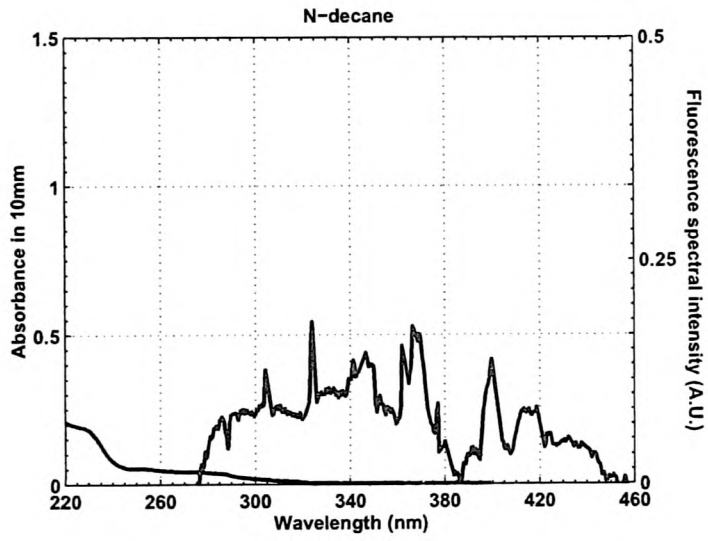
- absorption of the component fuel is neglected
- the mixture is fully vaporised *i.e.* no droplets are present in the cylinder
- the mixture vapour behaves as an ideal gas
- the absorption cross section of the liquid and vapour phase is the same at 248nm

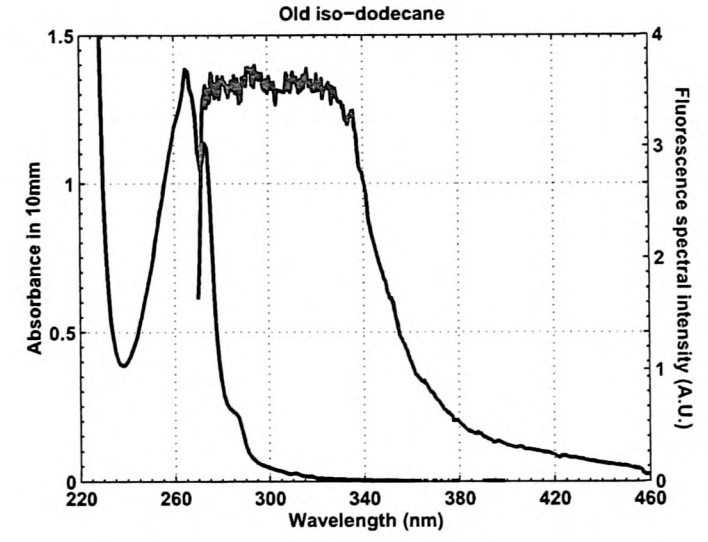

Iso-pentane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)

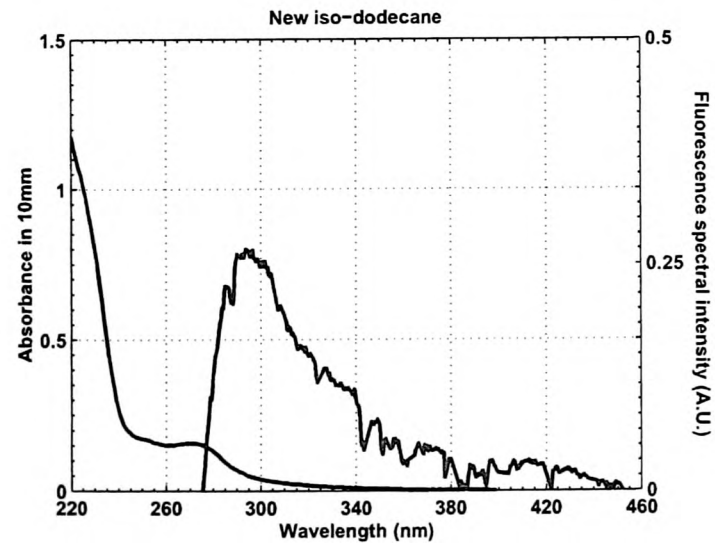

N-octane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)

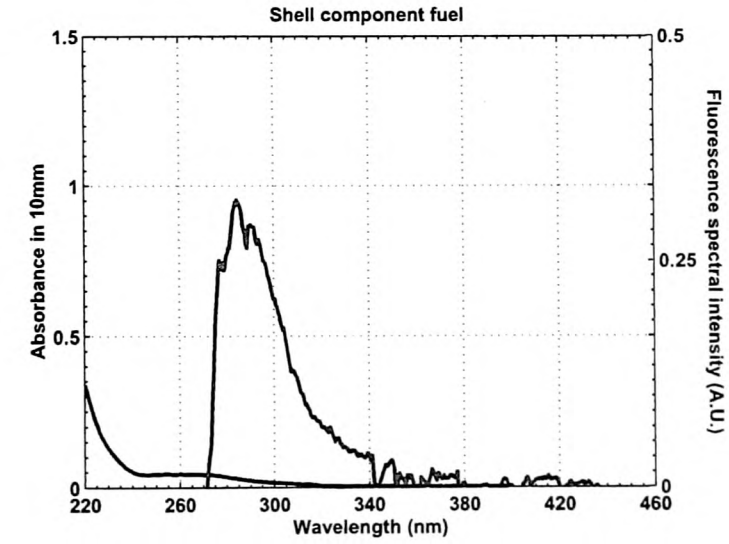

Old iso-octane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


New iso-octane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


N-decane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


Old iso-dodecane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


New iso-dodecane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


Shell component fuel
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)

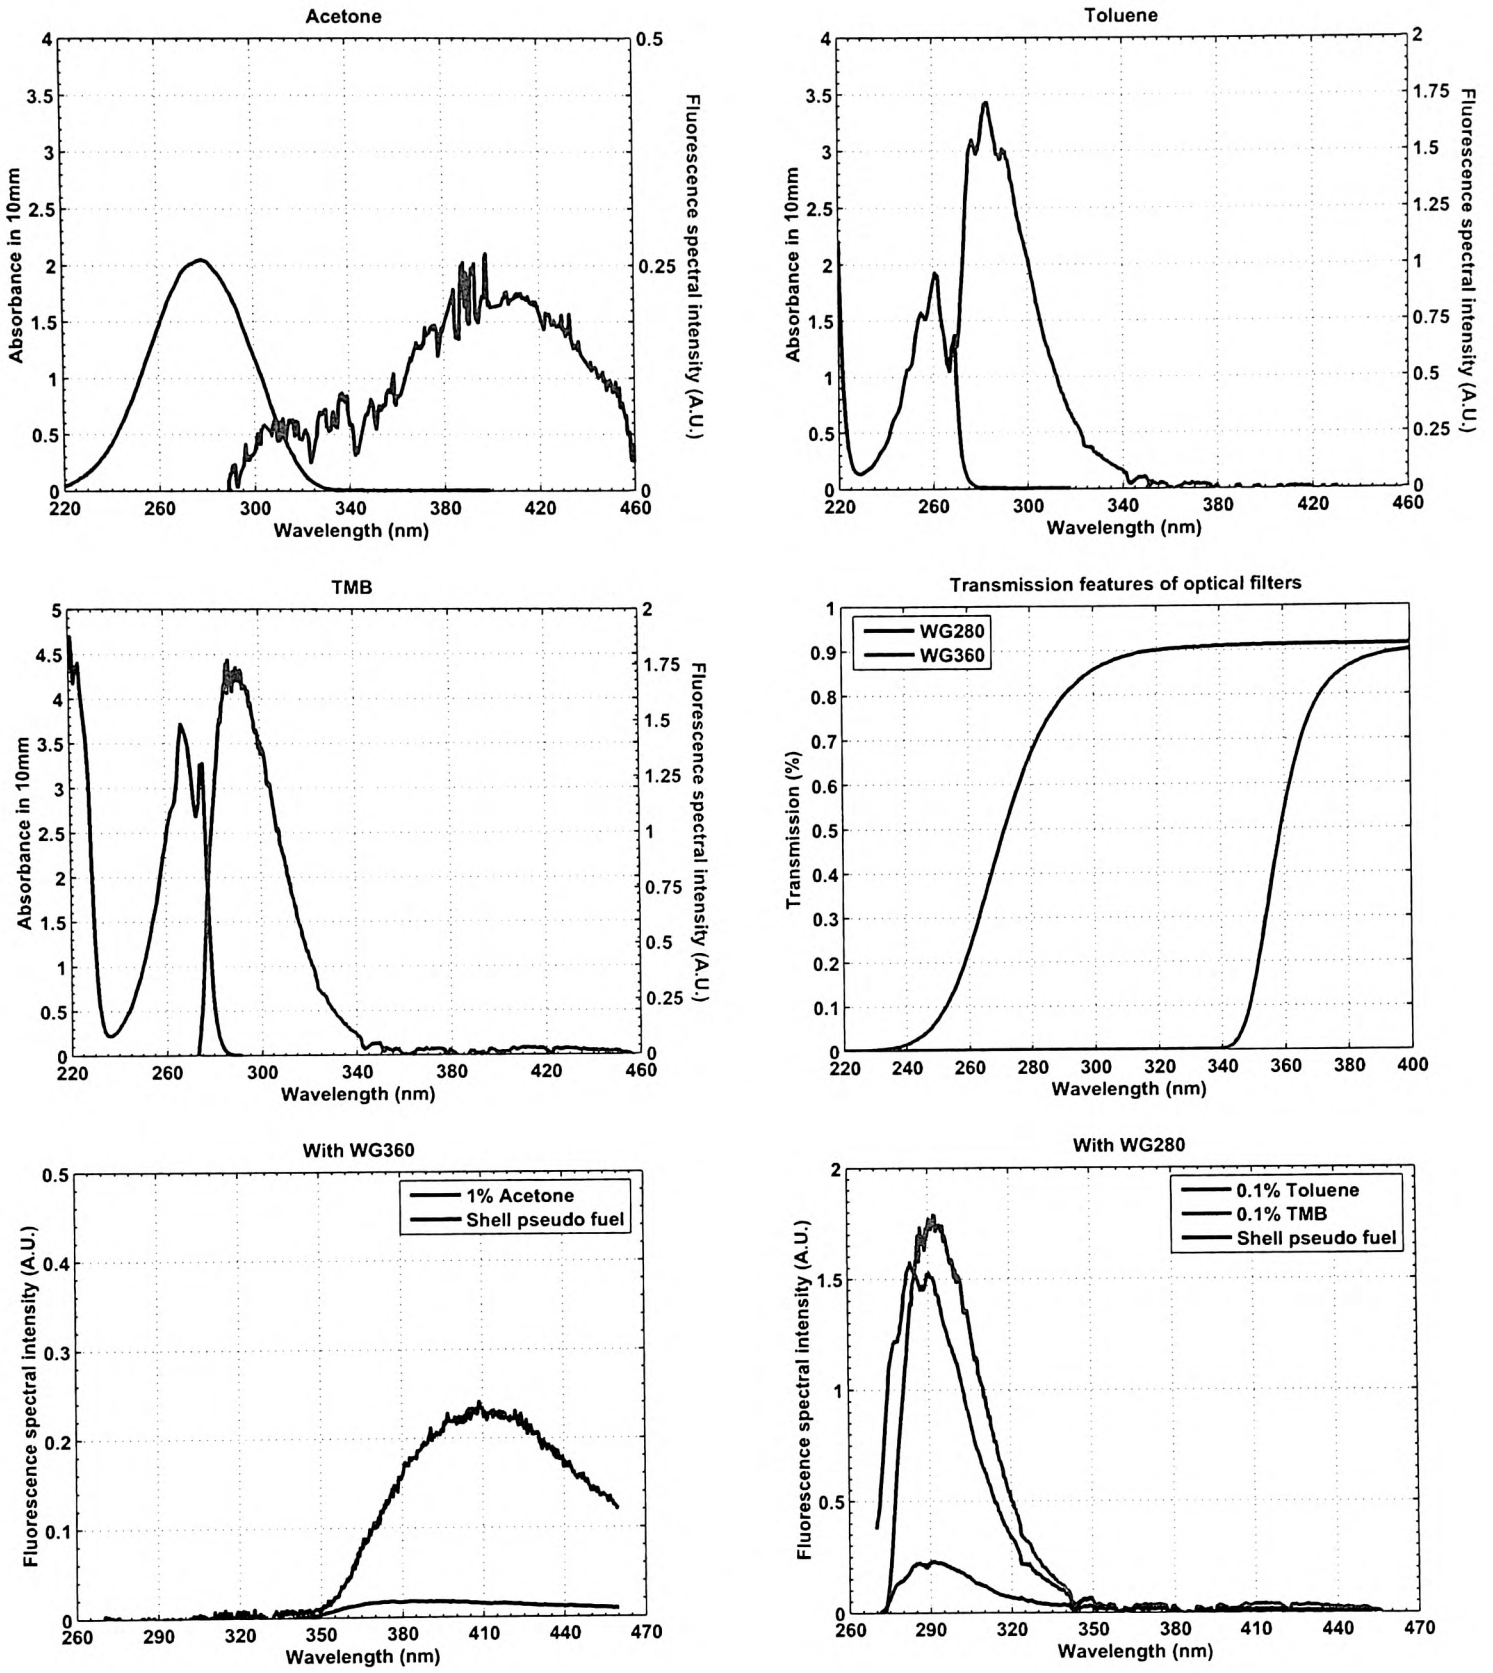


Figure 4.12: Absorbance and fluorescence spectra of tracers and component fuel, with and without optical filters

The SNR for acetone is generally better than the other two tracers although its absolute LIF signal is much weaker. This is because the WG360 filter removes most of the noise from the component fuel while the cut-off wavelength of the WG280 filter sits close to the peak of the noise. This has led to the fact that a relative noise of ~50% is to be

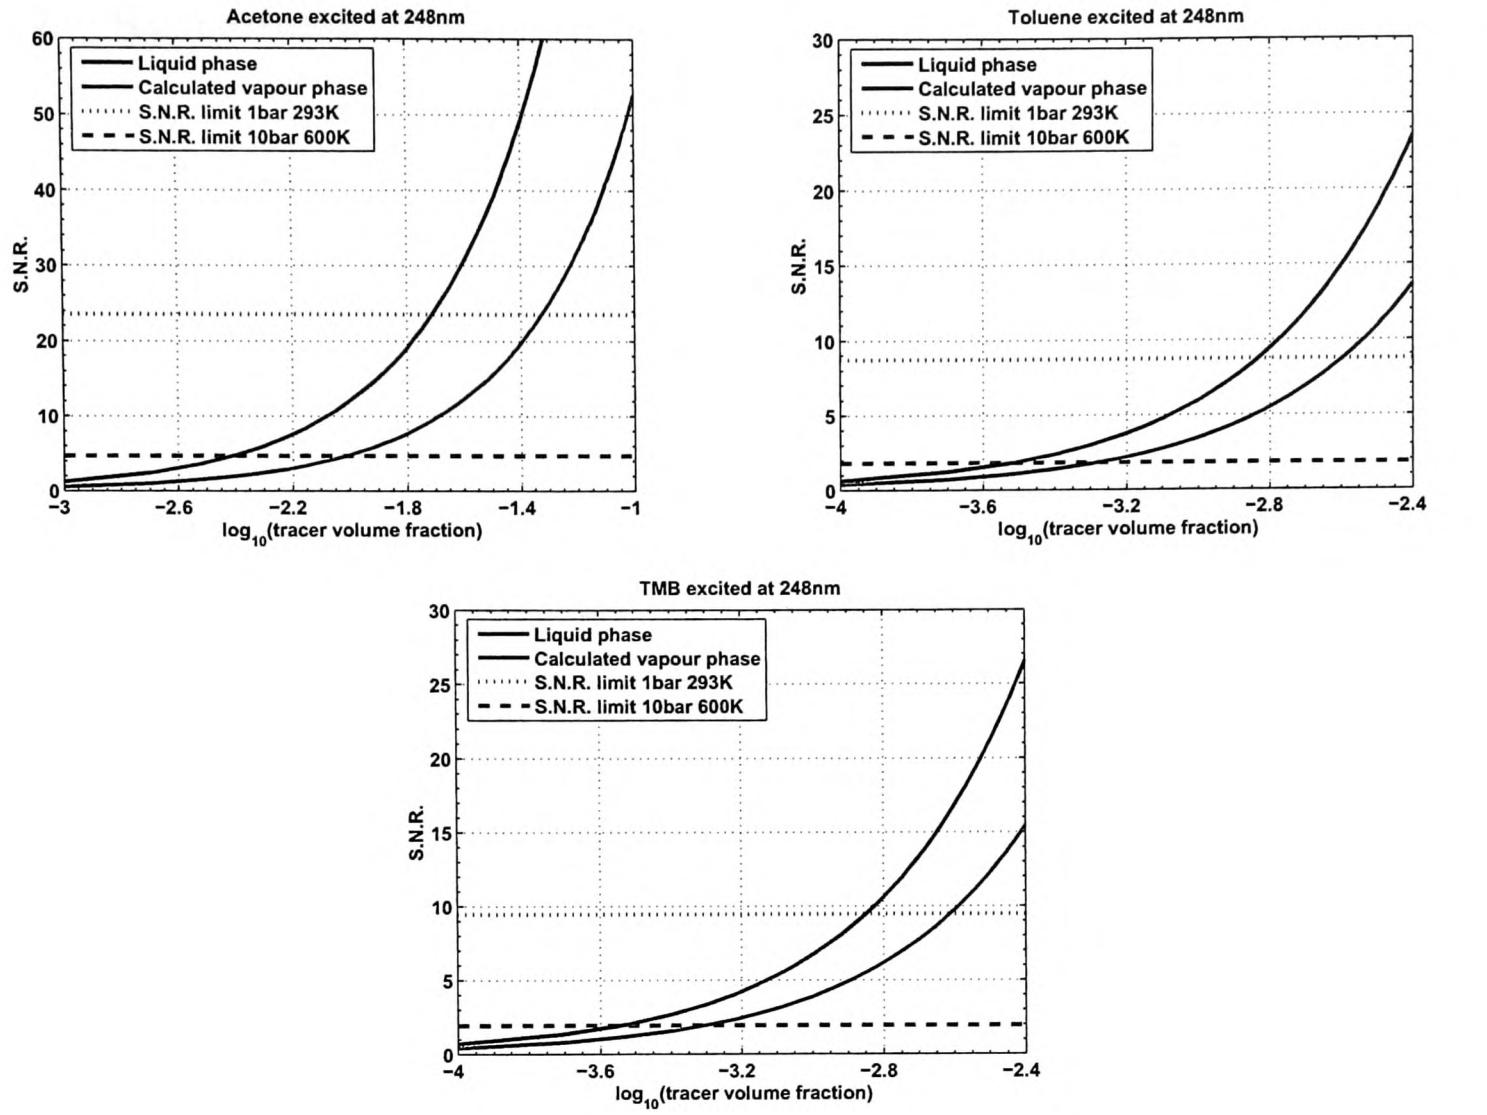


Figure 4.13: Prediction of signal-to-noise ratios for tracers

expected prior to ignition in an IC engine (10bar, 600K) when using toluene/TMB and one WG280 filter. When this is the case during the PLIF experiments, it is advisable to use two filters or even more together.

4.6 Optimisation of PLIF signal

In order to obtain quantitative results, the raw PLIF images must be post-processed for calibration of the PLIF system. This section will discuss details for noise correction and PLIF calibration. The relevant concepts of image processing have been introduced in Section 2.4.5.

4.6.1 Systematic noise correction

The first process is to remove systematic errors present in the imaging system, including:

- background noise
- built-in errors of the intensified CCD camera
- distortion effects
- Laser-induced incandescence (LII)
- phosphorescence and scattering

The background noise consists of two parts, dark noise and scattering noise. Due to the imperfect features of the intensified CCD camera system, even if the camera lens is completely covered, there will be a noise response from the CCD, *e.g.* the effect of the way in which the MCP (microchannel plate) is switched on and off, electrons drifting from the photocathode under high voltage, and deadtime noise. Deadtime noise is the charge that builds up on the CCD when there is no incident light signal and the MCP is switched off. Normally a CCD would be 'cleaned up' in this time to remove the build-up of noise, but in particular applications *e.g.* doing on-chip accumulation, such noise cannot be removed automatically by the camera. Unlike the dark noise, the background scattering noise is dependent upon the level of the gain and the setting of the intensifier gate. Therefore, to remove background noise, it is advisable to measure the background response, that is, to take reference images with constant illumination under different settings of the gate and the gain of the ICCD camera and to make subsequent corrections (background subtraction).

Apart from the imperfections mentioned above, other built-in features of the ICCD camera can also cause significant errors, such as:

- The spatial and temporal non-uniformity in the response of the CCD detectors (pixels);

- The incomplete zeroing after each pixel is moved forward to the shift register (the pixel data can only be read out when it is moved into the shift register).

Such errors can be corrected by some pixel-based transform and filtering methods (*e.g.* spatial Fourier transform).

Both the optical windows in the engine and the camera lens act as lenses and introduce optical distortion errors into the imaging. Images of a deliberately designed pattern can be taken, for instance, rectangular grids (a set of intersecting lines) for the side viewpoint (an imaging plane perpendicular to the piston crown) or concentric circles (with radial lines) for the vertical viewpoint (an imaging plane parallel to the piston crown). It should also be noted that in the vertical view case, with an optical window mounted in the piston, the piston motion can cause changes in the distortion level. This effect ought to be calibrated for as well. When such reference images are available, image registration can be done between the original pattern and the images from the camera to correct distortion effect.

Noise due to laser-induced incandescence (LII) from the fuel droplets and soot is also possible. It adds a noise signal to the fluorescence signal, as well as causing a loss of information behind the soot particles and fuel droplets because they absorb most of the laser light so that the area behind them becomes 'dark'. One feasible way to deal with LII noise is to use optical filters, because generally the wavelength spectra of LII signals and PLIF signals do not overlap.

As mentioned in Section 4.2, apart from the LIF signal there are other laser-induced light signals existing in the experimental process, namely phosphorescence and scattering. Because the lifetime of the phosphorescence signals is significantly longer than fluorescence, the laser-induced phosphorescence can cause blurring of the images (*e.g.* capturing moving objects). Laser-induced scattering processes (typically Mie scattering, Rayleigh scattering and Raman scattering), which almost happen simultaneously with the laser excitation process, are apparently potential sources of signal interference in PLIF experiments. By

gating the ICCD camera shortly after the laser pulse has ended, the scattering noise is expected to be avoided. Another feasible way of dealing with both phosphorescence and scattering noise is to place specialized optical filters in front of the image acquisition system.

Eq. (4.4) implies that the local laser intensity varies spatially as the laser gets transmitted by the media. Also present is the shot-to-shot variation in the input laser intensity. These two effects cause temporal and spatial non-uniformity and can be accounted for by using an averaged image of a homogeneous mixture of air and the base fuel to divide images from PLIF experiments.

Other systematic errors, including steering of the laser beam by density gradients, laser-produced perturbations, partial saturation and radiative trapping (absorption of the fluorescence within the combustion gases), are often avoidable or negligible, albeit needing to be considered in some particular applications (Zhao [78]).

4.6.2 Optimization of signal-to-noise ratio

This section describes the work carried out to optimise the SNR. Any leak of oil and coolant from the cylinder head and the engine block would significantly absorb the laser irradiance and hence reduce the fluorescence yield from the tracer. Furthermore, such contaminants would fluoresce and upset the PLIF signal. Therefore, great care has gone into recommissioning the engine to eliminate such leaks.

The gain and the gate setting (gate width, shape and timing) of the camera intensifier were studied and optimised. The gain was selected to give maximum available signal detection and yet avoid overloading the camera under rich conditions. The excess time the camera gate was open around the laser pulse was minimised. This was realised by interfacing with the laser triggering electronics and seeking an exact indication of the arrival of the light pulse — the signal from the thyratron in the Excimer. A thyratron is a type of gas

filled tube serving as a high energy electrical switch, basically a controlled gas rectifier. The signal from the thyratron was amplified and used to time the camera after the laser was triggered. This then minimised the effect of jitter in the laser pulse. Although careful gating of the camera helped to remove some of the background noise (ambient light, scatter, unwanted background fluorescence from contaminants *etc.*), a spectral filter that suited the acetone fluorescence spectrum was put in front of the camera.

The relation between the discharge potential on the Excimer and the laser pulse energy was investigated (Figure 4.14). Energy saturation occurred when the discharge potential went above 21kV, beyond which the discharge became less effective. The subsequent discharge potential was selected within the upper and lower bounds. The composition of the gas mixture used for the Excimer laser was also optimised to give an even discharge and improved spatial beam uniformity. An unstable resonator was designed and fitted in the Excimer laser cavity to reduce the laser beam divergence and improve the beam quality; this was also necessary for optimising the output beam from the Raman cell.

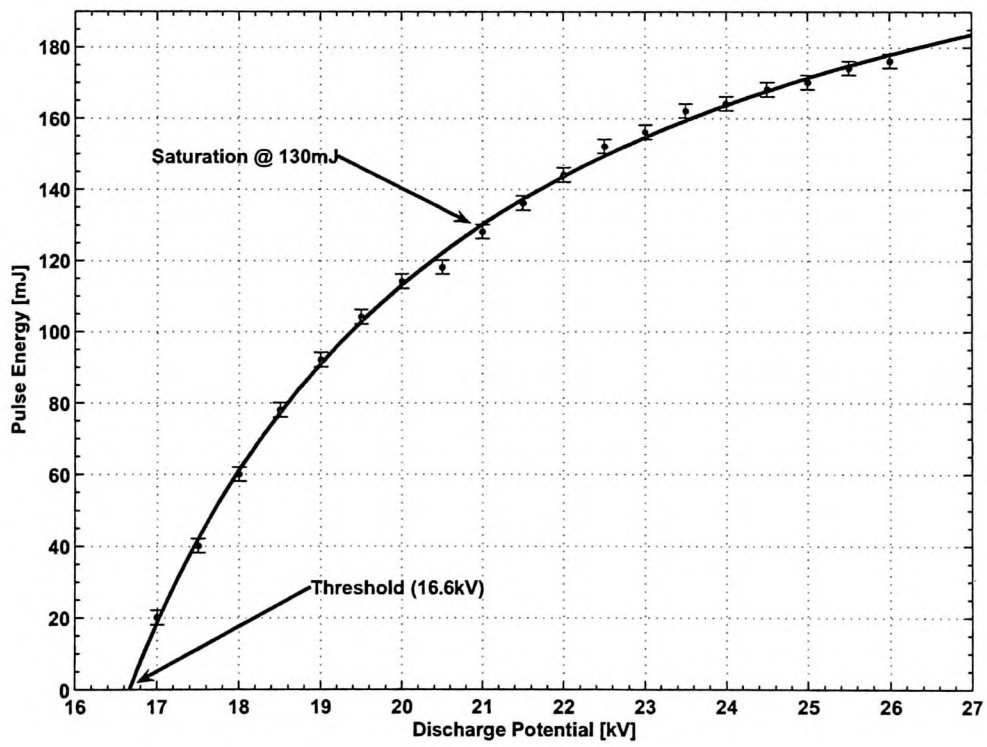


Figure 4.14: Excimer laser power curve for standard neon fill

Furthermore, acetone has a relatively low fluorescence yield, so the maximum laser in-

tensity and a good beam quality is critical to ensure a good SNR. The output from the Raman cell is proportional to the input laser energy until the input becomes so low that beam non-uniformities start to form (Figure 4.15). Hence, the selected input laser intensity should be in the linear range. When the input laser was focused with a 550mm lens, the efficiency of the Raman shifting process was estimated to be ~16% and the threshold of input laser intensity for the linear relation to hold was ~76mJ.

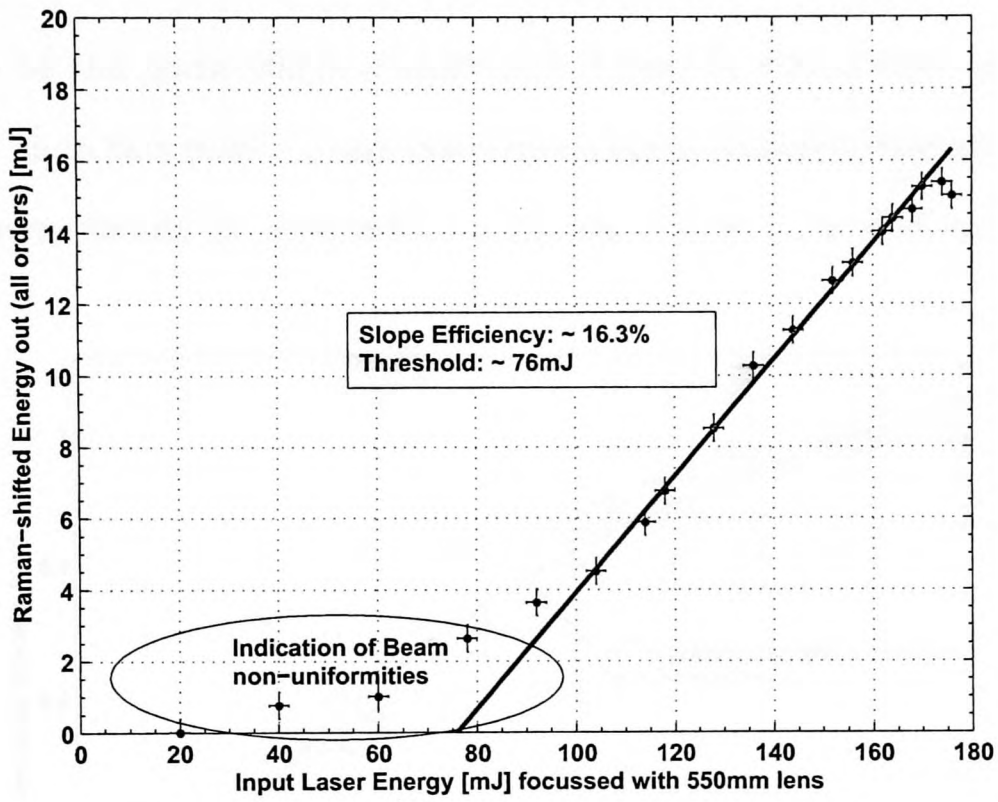


Figure 4.15: Estimation of the SRS efficiency

The SRS process gives an output of discrete wavelengths, consisting of Stokes and anti-Stokes shifted as well as the scattered fundamental. The Stokes shifting is of many orders, and the higher the order, the lower the intensity. In practice, up to five orders of Stokes shifting could be observed. For the output beam to suit experiments with acetone, effective filtering must be performed after the cell to cut out the original laser and high orders of Stokes shifting, other than the first order (277nm), from the Raman cell output. This was realised by using a prism and masks after the cell.

Before it entered the cylinder, the laser sheet was masked by a mask with a slit under the

piston window. In so doing, the thickness of the laser sheet was constrained, minimising the uncertainties due to line-of-sight effect; any unwanted scatter from the optics could also be blocked.

The realistic saturation intensities for the tracers were experimentally measured to obtain the maximum usable laser intensities for maximum signal in detecting the tracers (see Figure 4.16). The measurement looked at the fluorescence intensities from a particular tracer under different input laser intensities, then the saturation intensity could be estimated to be the point where a reasonable linear fit was always possible for input intensities lower than this point. Consequently, an input intensity slightly lower than this saturation intensity should be chosen for maximum fluorescence yield by this tracer.

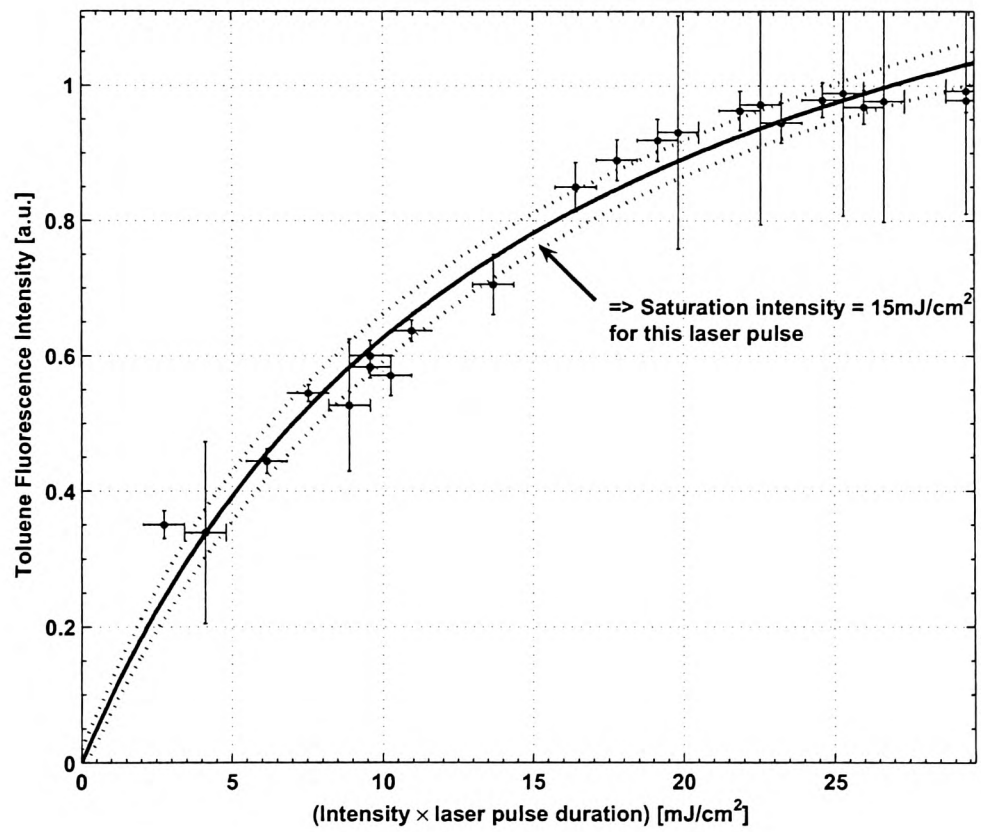


Figure 4.16: Example of estimation of the saturation intensity for toluene vapour

In the linear LIF regime, a single intense laser pulse can strengthen the PLIF signal output, but in practice its peak intensity will probably saturate the tracer absorption. So sometimes, especially in toluene PLIF measurements, the laser pulse needed to be stretched into a series of pulses all of which were below the saturation intensity for the

tracer at that particular wavelength. The schematic of this pulse stretching process and the corresponding optics setup is shown in Figure 4.17. The input beam was separated at the beam splitter, and the main beam carried on along the optical path as the first pulse in the pulse train. The diverted beam was then reflected by the four mirrors and hence became the delayed second pulse. The iterating circuit continued to give the third, fourth...pulses until the trapped energy dissipated. As a result of this setup, the exposure time of the ICCD camera needed to be prolonged to collect all the ensuing fluorescence signal.

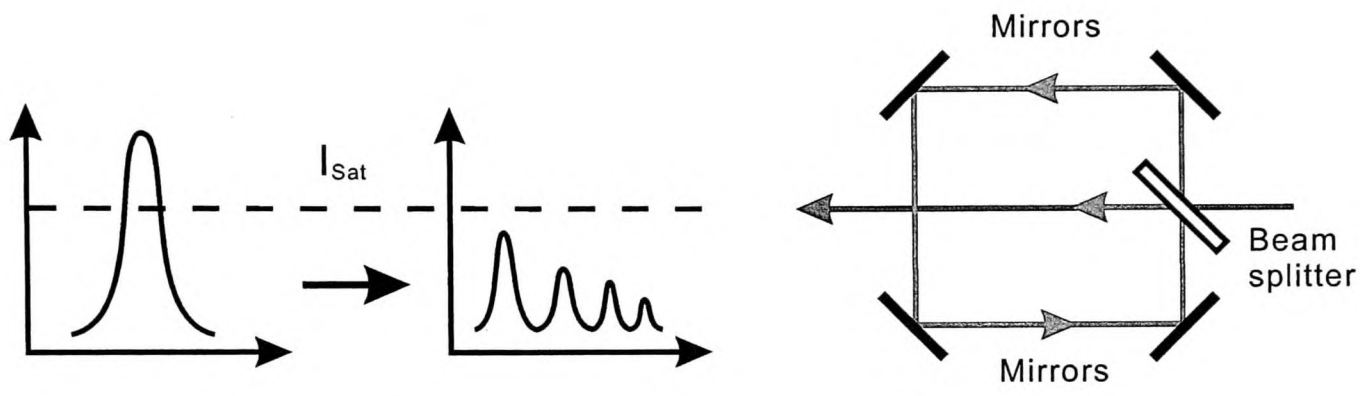


Figure 4.17: A schematic of the pulse stretching for toluene PLIF — pulse stretching (left) and its optics setup (right)

4.7 PLIF calibration

4.7.1 Calibration loop

The dependance of the PLIF signal on pressure and temperature can be calibrated in a pressure vessel where both pressure and temperature can be controlled precisely. However, difficulties will arise in reproducing the same optical and geometrical setup as that in the engine experiment. Also, it is not trivial to measure the instantaneous temperature in an IC engine. Therefore, a calibration based on a running engine is more representative and robust.

Sercey *et al.* [14] introduced a 'closed-loop' system to provide a homogeneous calibration

mixture in their DISI engine cylinder. A similar calibration loop was implemented in Oxford which directed the exhaust back to the inlet. By closely controlling the intake pressure and temperature as well as the injection duration, the fuel and air would be homogeneously mixed under known conditions after several engine cycles. The acquired PLIF images could then be calibrated against air/fuel ratio at various crank angles and hence various cylinder pressures and temperatures. The calibration images were typically taken during the compression stroke when the fuel (with one tracer corresponding to the fuel fraction under study) would have fully vaporised. The concentration of hydrocarbons was measured by a fast FID to derive the AFR in the loop. Cylinder pressure and temperature traces were measured throughout the engine cycles so a calibration 'sweep' of pressures and temperatures was effectively performed within a single test. It should be noted that given the maximum available sampling rate of the ICCD camera, images had to be captured through many engine cycles to build up a full calibration map *i.e.* database. Figure 4.18 shows a schematic of the PLIF calibration loop.

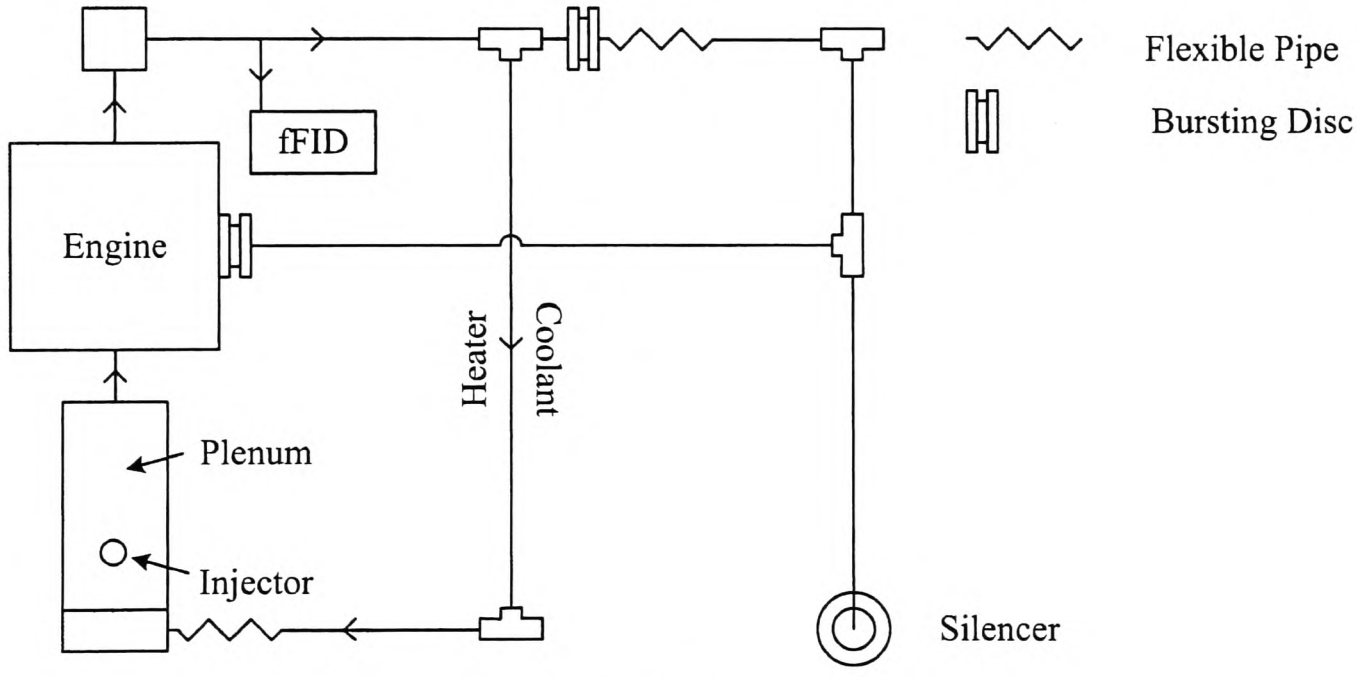


Figure 4.18: Schematic of the PLIF calibration loop (plan view)

Two bursting discs were used as one-time safety valves, with one side of them facing the pressure being monitored, and the other side connected with the exhaust extraction facilities. They would burst when the pressure being monitored had exceeded a certain

threshold. The bigger disc was put between the loop and the exhaust, and the smaller one with a higher pressure rating was installed in the aluminium annulus on top of the cylinder liner. The loop outside the engine was wrapped with a heating element and coolant pipes in order to maintain the required temperature for fuel vaporisation. The engine coolant was heated as well to above 40°C; however, the cylinder barrel was often running warmer, up to 70°C which was necessary to keep all the fuel fractions and tracers vaporised at all conditions down to $\lambda = 0.3$.

Between two measurements the mixture in the loop was purged by motoring the engine with a diaphragm pump; fresh air was drawn in from a bleed valve. The fFID was used to monitor the residual HCs during this process. Since the response time of all instrumentation was fast enough to resolve engine cycles, the blowby effect was accounted for despite the slight change in AFR over time.

The laser sheet entered the engine via the 45° mirror underneath the piston and was set to be in the vertical plane intersecting the cylinder axis. The ICCD camera viewed the fluorescence through the pentroof window. A schematic of the optical access in the pentroof region as well as a sample raw image are shown in Figure 4.19. In the sample image, black lines have been overlaid to show the field of view *i.e.* the pentroof window. The undesirable fluorescence from oil can be seen around the Inlet valve and the edge of the pentroof window.

4.7.2 PLIF vs. AFR calibration

This section presents results from the first set of PLIF calibration tests, carried out at 1500RPM, -0.48bar gauge inlet pressure with acetone as the tracer. The loop temperature and barrel temperature were 50°C and 60°C respectively. A volume fraction of 10% was used for acetone. Too low a tracer concentration would give a poor SNR while on the other hand, too high a concentration would lead to heavy absorption along the laser sheet

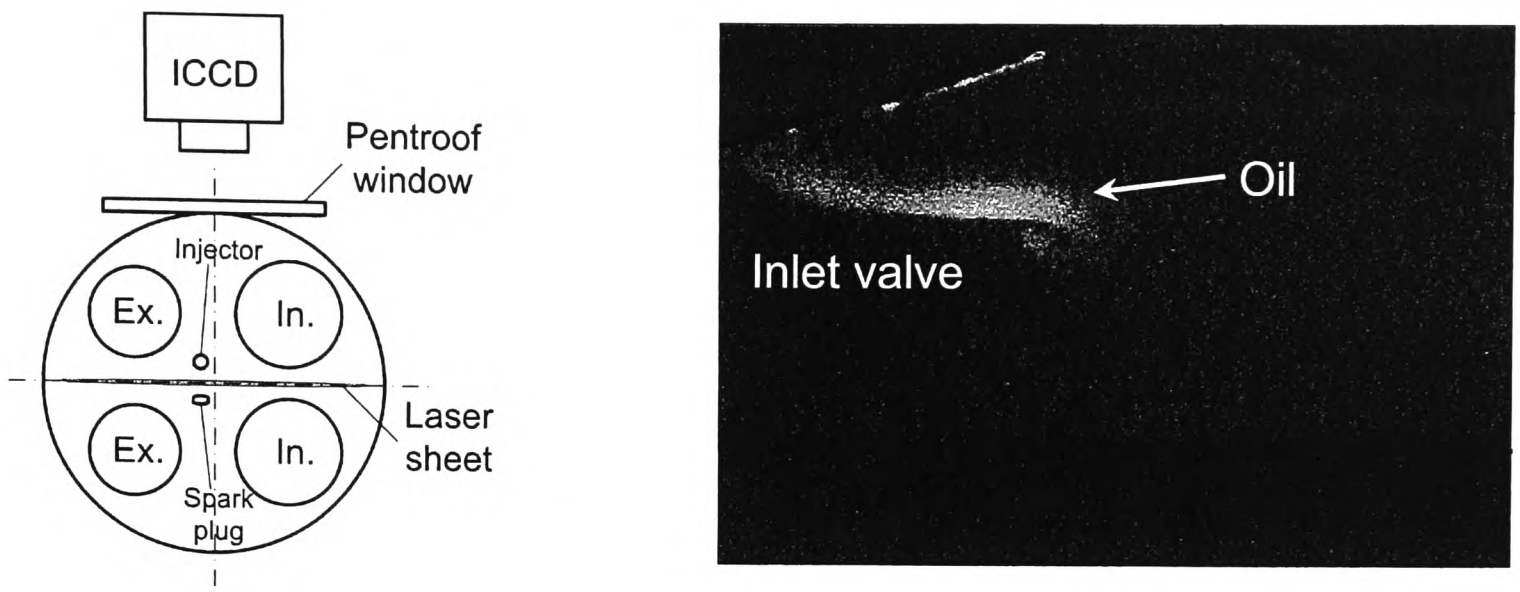


Figure 4.19: Schematic of the optical access in the pentroof region from a plan view (left) and a sample raw image (right)

and spatial non-uniformity in the PLIF images.

The ETCS program sent a signal to trigger the laser, while the ICCD camera was timed by an amplified signal from the thyratron after the laser was triggered. With the camera set to a reduced frame size, 576 by 191 pixels, its sampling rate could match the laser repetition rate of 3.125Hz, that is, skip pulsing once every four engine cycles. The laser energy was recorded by an energy meter to account for the temporal fluctuation in the laser pulse.

Figure 4.20 shows a sample image of the background light (with laser but no fuel) and a PLIF image after the background subtraction, with a lambda of 1.0 at TDC. The latter clearly outlined the laser sheet.

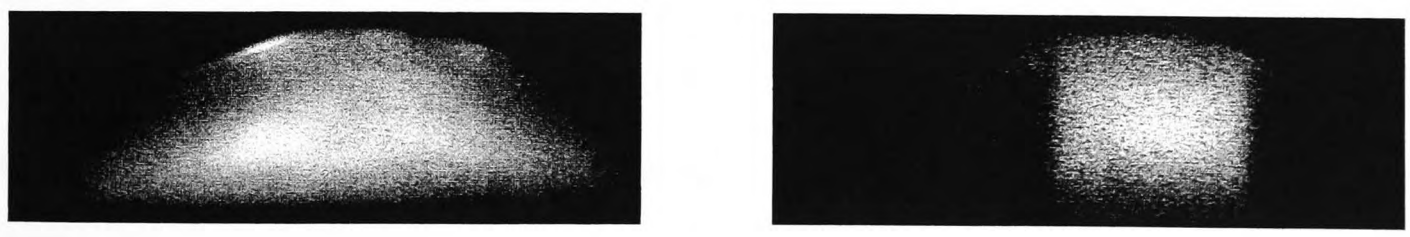


Figure 4.20: Sample images of background (left) and PLIF signal after background subtraction (right), $\lambda = 1.0$ with acetone @ TDC

The camera response to fluorescence yield was calibrated against crank position and hence

the cylinder pressure and temperature at that particular crank angle (Figure 4.21). The camera counts were normalised after taking into account the chosen camera gain and exposure settings. It can be seen that with a lambda of 1.5, the signal level was low and in some cases below zero; the dependance of the signal on p and T was difficult to extract. This was due to overwhelming background noise and a poor SNR. The standard deviation was the least for stoichiometric operation.

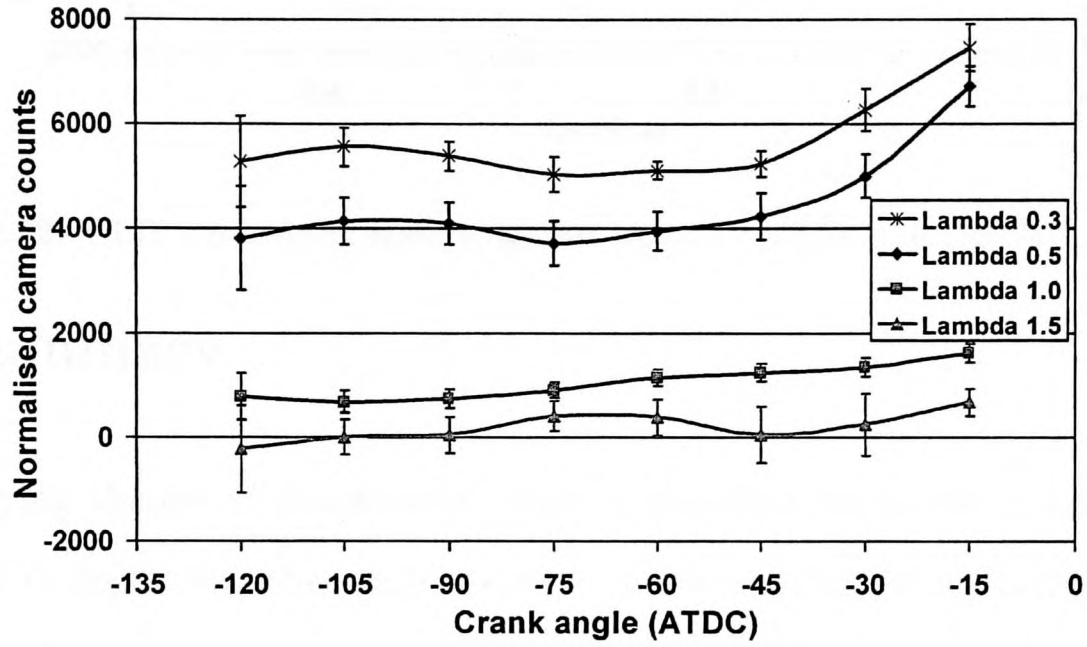


Figure 4.21: PLIF calibration results against crank angle with different lambda

Figure 4.22 shows the calibration results against lambda at a series of crank angles. Again, the SNR was very poor with $\lambda = 1.5$ and above. For stoichiometric and rich operation, the relation between camera response and mixture strength was reasonably linear.

At present, the engine is being overhauled and measures to improve the SNR are being implemented. More PLIF calibration results with other tracers as well as AFR measurements from normal firing engine tests are to follow in the next stage of the ongoing project.

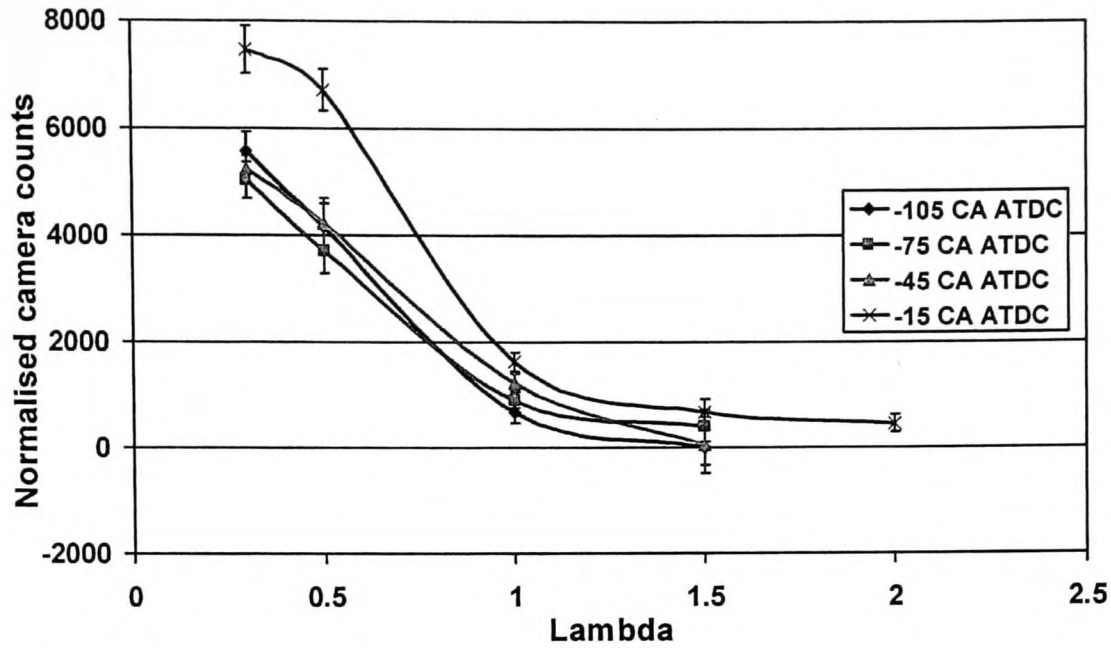


Figure 4.22: PLIF calibration results against mixture strength at different crank angles

4.8 Summary

The underlying theory of quantitative PLIF measurements in an IC engine has been summarised in depth, together with a review of others' relevant work to this date. The methodology for selecting a fuel/tracer set for tracer-PLIF has been discussed. A dedicated six-component three-fraction fuel was designed, with three corresponding tracers to match the light, medium and heavy fraction of the fuel respectively. PLIF images with this fuel/tracer set can reveal quantitatively the distribution of different fuel fractions in an engine. The absorption and fluorescence characteristics of each fuel component and tracer were studied to give an estimate of the bounds of signal-to-noise ratio.

Great care was taken to optimise both hardware and the SNR, including spectral filtering, time gating of the camera, laser pulse stretching for toluene PLIF, study of stimulated Raman scattering for acetone PLIF and absorption saturation *etc.*

A recirculating loop was constructed and set up for the calibration of PLIF signal against AFR. This experimental layout has the advantages of *in situ* optics and a real engine environment. Example calibration results with acetone as the tracer were shown, suggesting

that maximising SNR was the key to a practical calibration for quantitative PLIF.

Chapter 5

Soot Imaging and Pyrometry

There are two kinds of light — the glow that illuminates, and the glare that obscures.

— James Thurber (1894–1961)

5.1 Introduction

CCD imaging and colour pyrometry are among the most commonly used techniques for characterizing flames and combustion [69] [77]. Colour pyrometry relates the temperature of glowing soot and the intensity of the thermal radiation it emits. In conventional two-colour pyrometry, for example [3], the light sensing device is calibrated against a black body to determine the apparent temperatures at two specific wavelengths.

The flame temperature can be determined from the ratio of the light intensities at two wavelengths by eliminating an unknown coefficient in the extended version of Planck's radiation law. Substitution of the flame temperature into equations involving the absolute intensities yields an estimate of the KL factor, which is proportional to the soot concentration. However, the necessity to have an absolute calibration of the light-sensing

apparatus for the two colour method makes the result sensitive to factors that alter the visible intensity, such as deposits on optical surfaces or wall reflections/glare.

In contrast, the colour-ratio pyrometry (CRP) technique [57] [73] uses the ratios of three colours to give two independent measures for determining the temperature and KL. The conventional approach normally uses narrow band filters and two light sensors/cameras with beam splitting optics, but it is also possible to use a single colour CCD camera [63] to avoid complexities such as the optical setup and image registration.

The advantage of CRP, that absolute intensities are not required, means that it is insensitive to the multitude of problems associated with quantitative intensity measurement. Such problems include calibration errors, variations in optical path and window fouling. Ratio measurements mean that the technique is more rapidly and robustly applicable to engine environments. However, sufficient equipment dynamic range is needed for good accuracy.

The work reported here employs a colour CCD camera with three channels (RGB) and uses the ratios of the pixel intensities. The three channels can be reduced to two ratios between red and green (R/G), and red and blue (R/B). In so doing the absolute intensity response is not used and an absolute temperature calibration is not needed. Instead, look-up tables of R/G and R/B ratios are created for the appropriate range of temperature and KL values.

Before describing the current work in detail, it will be useful to review some of the practical problems associated with pyrometry. The ideal form of calibration is with a black body, but it is difficult to operate such objects at combustion temperatures. The usual approach is to use a tungsten ribbon lamp, but the emissivity of the filament is a function of wavelength, temperature, age, surface finish and temperature-history of the lamp; consequently there are wide variations in the published data [68]. This issue is side-stepped as it is common for the tungsten ribbon lamp only to have a calibration in

terms of an apparent (equivalent black body) temperature at a particular wavelength. This can be one of the wavelengths used in traditional two-colour pyrometry (if it does not coincide with a wavelength that is characteristic of chemiluminescence), but there is still the issue of obtaining an absolute calibration at a second wavelength. If absolute measurements are being used, then strictly speaking an *in situ* calibration is needed. One approach is to install a calibrated source in the engine, or a tungsten ribbon lamp can be placed in an equivalent optical path [10].

It must also be remembered that the measurement is not from a single plane, but it is an integration along the optical path. Murphy and Shaddix [44] conclude that the stratification of temperature along an optical path biases temperature measurements towards the higher temperatures along the path, and that the concentration of the hottest soot will bias the KL value. This work suggests that when there is a variation in temperature of 600K with a mean temperature of 1500K, then errors in the measured temperature can readily exceed 100K. There is also the issue of how to model the emissivity of the soot, and this is discussed in detail by di Stasio and Massoli [15]. However, it would appear that this source of error is much less significant than the effect of variations in temperature or KL along the optical path being measured.

With all of these possible sources of error, then why is soot pyrometry a popular technique? It is comparatively cheap and easy to implement, and can be used to give a 2-D plot of the temperature and KL. This is very useful for comparisons between, for example, different fuels or different injection systems.

Colour-ratio Pyrometry has not been extensively applied in the literature; therefore this chapter describes a generic approach to calibrating a CCD camera for use with CRP. The validation of the calibration method and the temperature prediction with bench tests is also given here. Finally, the colour-ratio technique has been applied to the Jaguar optical access DISI engine with a spray-guided combustion system. Results from measurements over a range of engine operating conditions are reported.

5.2 Principles of soot pyrometry

5.2.1 Radiation of glowing soot

Analysing flame images relies on the relationship between the temperature of a glowing particle and the intensity of the thermal radiation it emits. The spectral intensity emitted by a black body, $I_{\lambda b}$, is a function of wavelength and temperature alone, and is given by Planck's Radiation law:

$$I_{\lambda b} = \frac{2C_1}{\lambda^5 \left[\exp\left(C_2/\lambda T\right) - 1\right]} \tag{5.1}$$

where C_1 and C_2 are known as Planck's radiation constants, having values of 1.871×10^{-16} W$\cdot$ m^2 and 1.439×10^{-2} m$\cdot$K respectively.

Real objects such as soot particles are not perfect emitters, and radiate less energy at each wavelength than the equivalent black body value. The spectral emissivity ε_λ is defined as the ratio of the actual intensity emitted by a body at a given wavelength to the black body intensity, so that

$$I_\lambda = \varepsilon_\lambda I_{\lambda b}(\lambda, T) = \varepsilon_\lambda \frac{2C_1}{\lambda^5 \left[\exp\left(C_2/\lambda T\right) - 1\right]} \tag{5.2}$$

The emissivity of the flame has to be estimated and this can be done by the widely used empirical correlation established by Hottel and Broughton [28]

$$\varepsilon_\lambda = 1 - \exp\left(-KL/\lambda^\alpha\right) \tag{5.3}$$

where K is the absorption coefficient per unit flame thickness, and L is the geometric flame thickness along the light axis of the flame detection system (optical path length). The absorption coefficient is a function of the soot particle concentration in the flame. The value of α depends on the physical and optical properties of the soot in the flame and, at visible wavelengths, was found by Hottel and Broughton [28] to have a value of

1.39. Typical temperature values for glowing soot range from about 1900K to 2600K, with the KL factor varying between 0.2 and 1.4 (Wyszynski [74]).

5.2.2 Colour-ratio pyrometry

The most commonly used technique for measuring flame temperatures is two-colour optical pyrometry, which was originally developed by Hottel and Broughton [28]. It is necessary to calibrate the light-sensing apparatus for the two colour method so that the signals can be related to an apparent flame temperature. This can be done using a source of black body radiation, to produce calibration curves of output against black body temperature for each wavelength.

Whereas the two-colour method uses a temperature-calibrated optical system at two specific wavelengths or wavelength bands, a colour camera has the three colour (RGB) filters with overlapping wide spectral responses (see Figure 5.16 in Section 5.3.1). The three channels can be reduced to two ratios between red and green (R/G), and red and blue (R/B). In this way the absolute response is not used and absolute temperature calibration is not needed. Using equations (5.2) and (5.3), and the camera spectral response from Figure 5.16, look-up tables of R/G and R/B ratios can be calculated for typical temperature and KL values.

The sensitivity of the system could be enhanced by using a filter that would attenuate the longer wavelengths, so that at a mid-range temperature each of the three colour channels would have the same signal level.

Soot pyrometry estimates directly the temperature of the soot, but due to the almost instantaneous attainment of thermal equilibrium (in 10^{-6}–10^{-5} s, Schanzlin *et al.* [52]), it can be assumed that the measurement of the soot temperature also represents the burnt gas temperature, *i.e.* the temperature of the surrounding gas.

Two strategies for colour-ratio pyrometry can be employed, namely interpolation with a

fixed value of KL and mapping with a variable KL. The schematics of these procedures are shown in Figure 5.1.

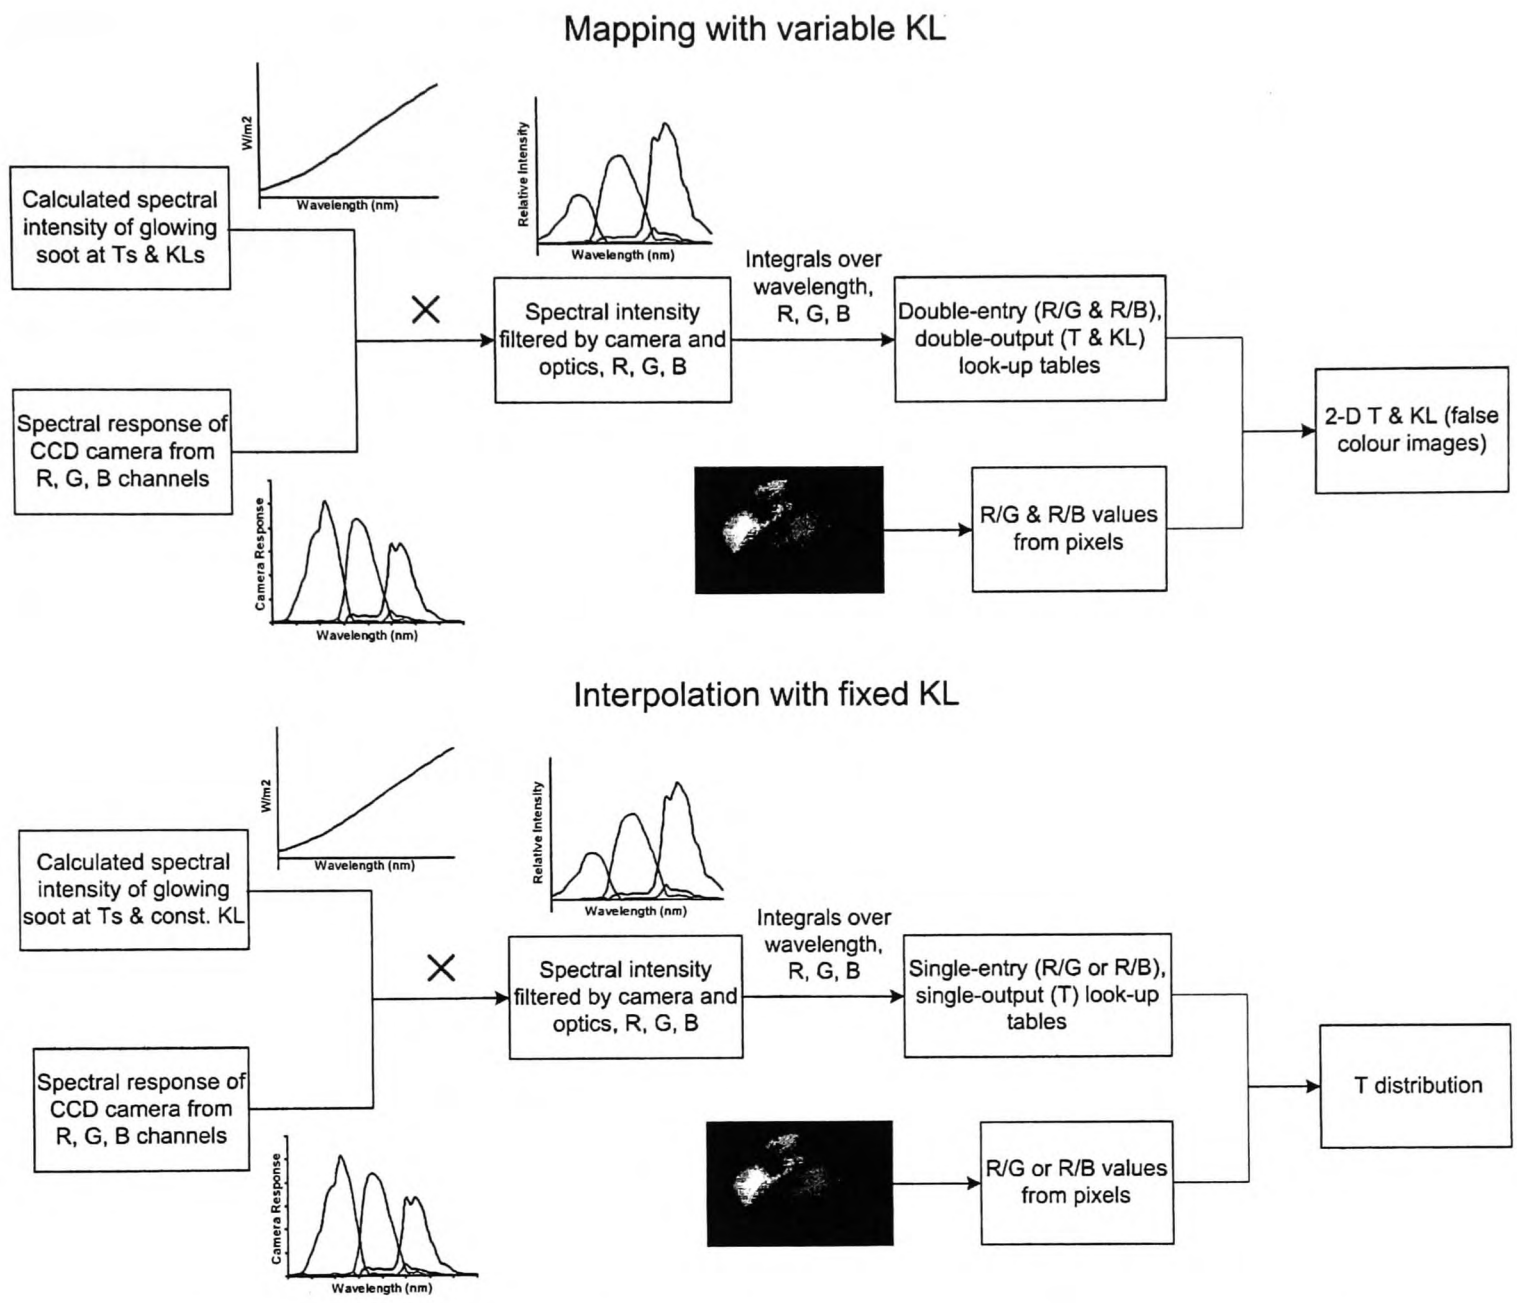


Figure 5.1: Procedures of the two strategies for colour-ratio pyrometry

By means of integrals over the full wavelength range, the colour-ratio technique is less sensitive to the wavelength selections and emissivity variation usually faced in classical two-colour pyrometry. In addition, because the solutions are derived from pre-created look-up tables via interpolation rather than iteratively solving non-linear equations, less computational time is required and the processing is significantly simplified.

When the mapping with a variable KL strategy is applied, both the R/G and R/B ratio values from the image are used and the mapping for temperature and KL is performed

to minimise the following characteristic function:

$$E = \left[(R/G)_{T,KL} - \alpha_{i,j}\right]^2 + \left[(R/B)_{T,KL} - \beta_{i,j}\right]^2 \tag{5.4}$$

where $(R/G)_{T,KL}$, $(R/B)_{T,KL}$ are the inputs from a particular pixel and $\alpha_{i,j}$, $\beta_{i,j}$ are the entries in the look-up tables. Figure 5.2 shows the surface plots of the colour ratios in the two look-up tables created with the 2-D mapping strategy.

Despite the potential of spatially resolving both temperature and KL, previous work [56] [63] has shown that the KL solution from soot pyrometry is often ill-conditioned in some cases, because:

- for a given pair of colour ratios there can be multiple solutions of temperature and KL which all map closely to the measured colour ratios;
- Uncertainties due to noise in the blue channel of a conventional CCD are high compared with the precision required to determine KL accurately;
- KL tends to be dominated by the environment closest to the detector.

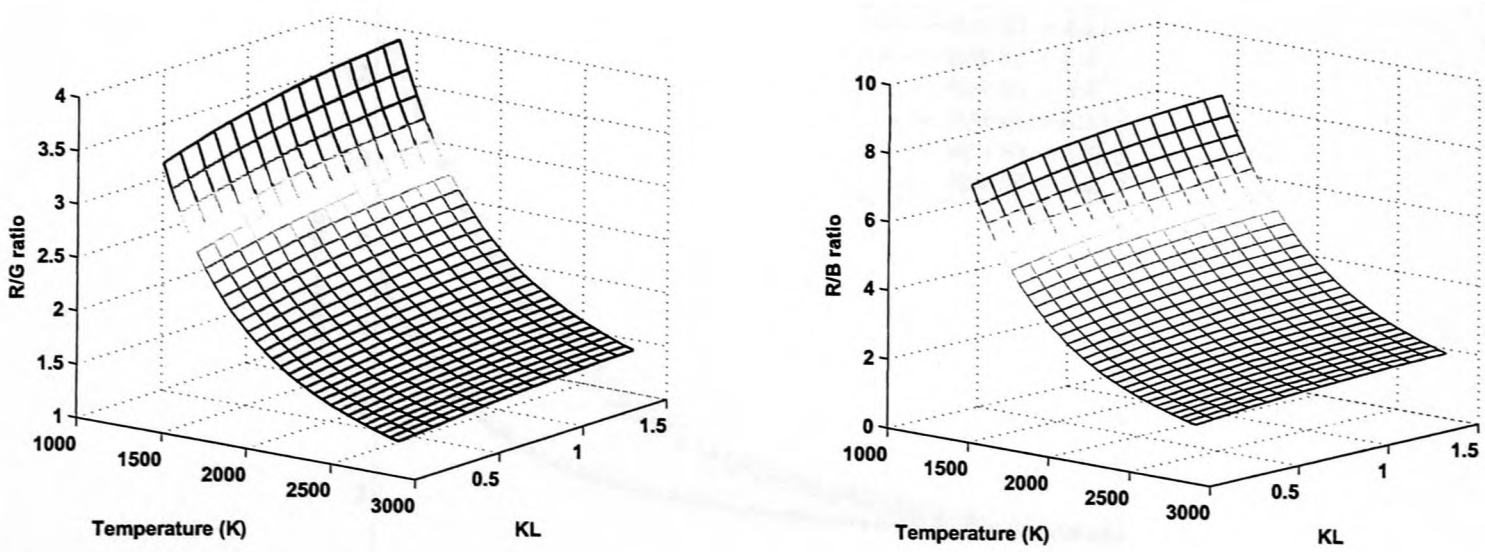


Figure 5.2: Look-up tables calculated with the mapping with a variable KL strategy for the temperature range 1000–3000K in steps of 10K, and for the KL range 0.1–1.5 in steps of 0.1

The variation of emissivity with wavelength at different KLs [28] is shown in Figure 5.3. Assuming a constant KL eliminates the complication of accounting for soot emissivity

at the expense of not being able to detect soot concentration. Figure 5.4 shows the R/G and R/B ratios in the look-up tables created with the interpolation method, assuming a constant KL. The temperature of each pixel can then be derived by interpolation of the R/G or R/B ratio.

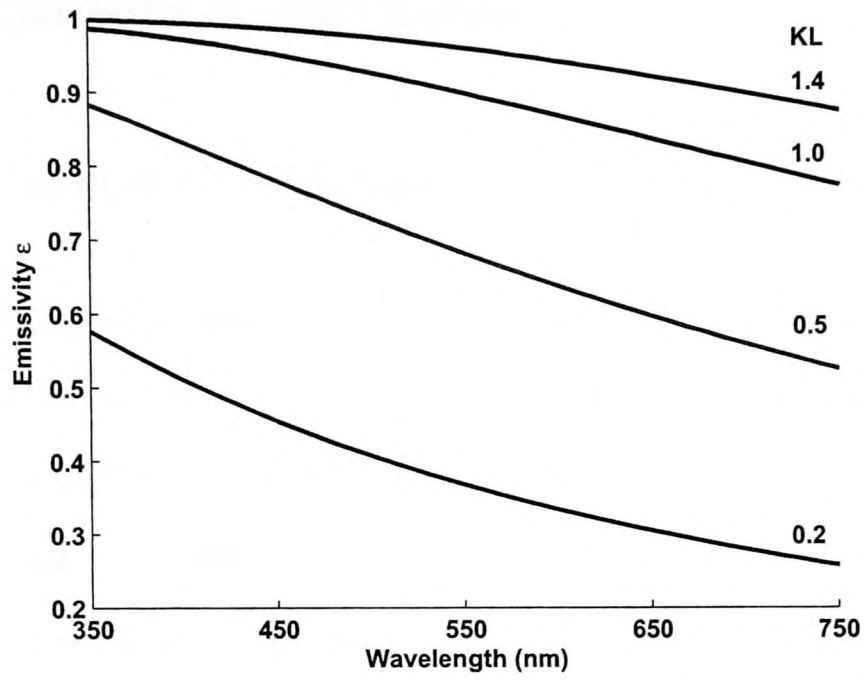


Figure 5.3: Variation of soot emissivity with wavelength and KL [28]

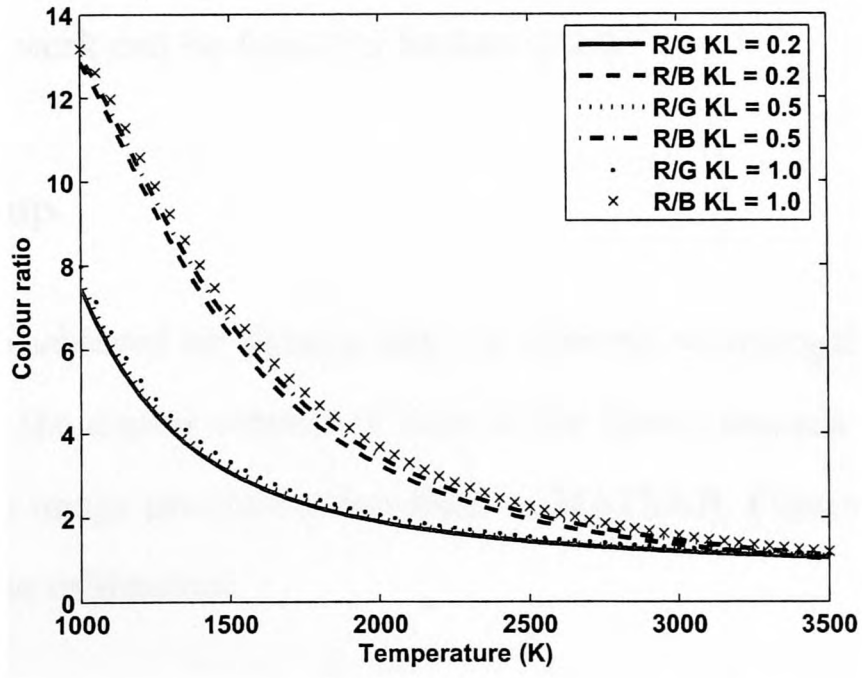


Figure 5.4: Expected R/G and R/B ratios of CCD camera response with different KL values

It is worth noting that the blue green (B/G) ratio is not mathematically independent of

the R/G and R/B ratios. Additionally, because of decreasing light intensity towards the blue end of the emission spectrum, the B/G is the most susceptible of the three ratios to noise and digitization error. For these reasons the B/G ratio has not been used.

5.3 CCD camera calibration

5.3.1 Colour spectral response

Since the soot and tungsten filaments do not have the same emissivity, then a direct calibration is not possible. If the soot was assumed to be a grey body (with constant emissivity), then a black body could be used if one was available at a high enough temperature, however these are difficult to manufacture. Instead it is necessary to determine the spectral calibration of the CCD camera, so as to obtain the camera's colour spectral response. This response (which is seen later in Figure 5.16) is a measure of the camera's sensitivity to different wavelengths for the three colour channels (red, green and blue), and is needed for the CRP. The technical specifications for the Pulnix TMC-9700 video camera used in this work can be found in Section 2.4.3.

Experimental setup

The calibration was achieved by shining light at different wavelengths onto the camera's CCD sensors. Then the digital outputs of each of the three channels were extracted from the images using the image processing functions in MATLAB. Figure 5.5 shows the setup that was used for the calibration.

Light from a tungsten lamp (13.5V 5A) was led through the spectrometer (HILGER & WATTS). It was focused onto grating g (1200 lines/mm, blazed @300nm) and then focused onto the camera CCD chip (no objective lens) or photodiode. Both slits s_1 and

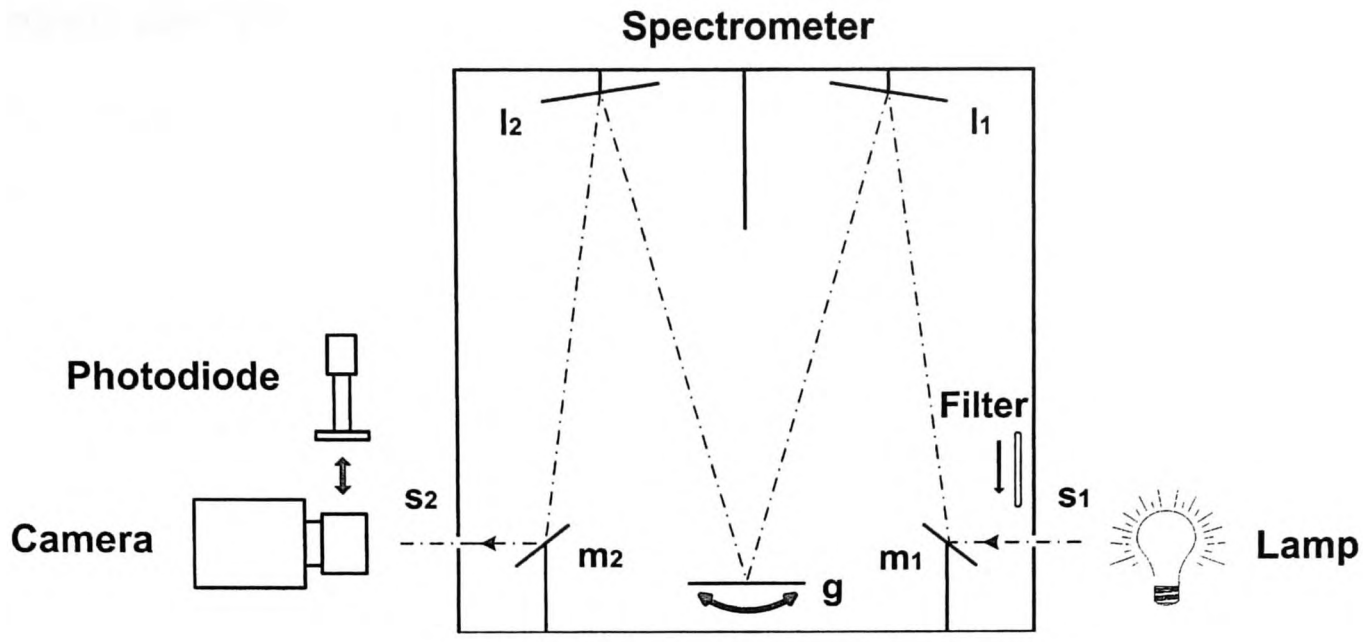


Figure 5.5: CCD camera calibration setup

s_2 were set to 2mm, giving an output bandwidth of 8nm. The output wavelength was varied by rotating the grating around a vertical axis (perpendicular to this sheet). To calibrate the emission spectrum of the lamp and the efficiency of the spectrometer, a photodiode of known spectral sensitivity (ThorLabs DET210) was used to measure the relative intensities when extracting the spectral response curves.

To avoid saturation of images at certain wavelengths yet make full use of the dynamic range of the camera, different shutter speeds were used.

The spectrometer could potentially introduce the following sources of error:

- Because the bulb emitted much less in the blue range than red, light scatter of the red constituents could bias the measurements in the blue range;
- The second order diffraction effect of the grating could cause some blue radiation (300–400nm) to appear in the signal in the red range (600–800nm).

To attenuate these sources of noise, four filters were selected and put behind slit s_1 according to the wavelength range being measured. When the Blue filter was used, most of the radiation above 500nm would have been heavily attenuated, likewise for the LEE 026 filter, hence eliminating the cross talk between wavelength ranges. The transmission char-

acteristics of these filters were obtained from a spectrophotometer (Perkin-Elmer Lambda 9, usable range 185–3200nm) and are shown in Figure 5.6. Also presented in Figure 5.6 are three other filters used when testing the camera's response linearity with sensitivity and exposure (detailed in Section 5.3.2).

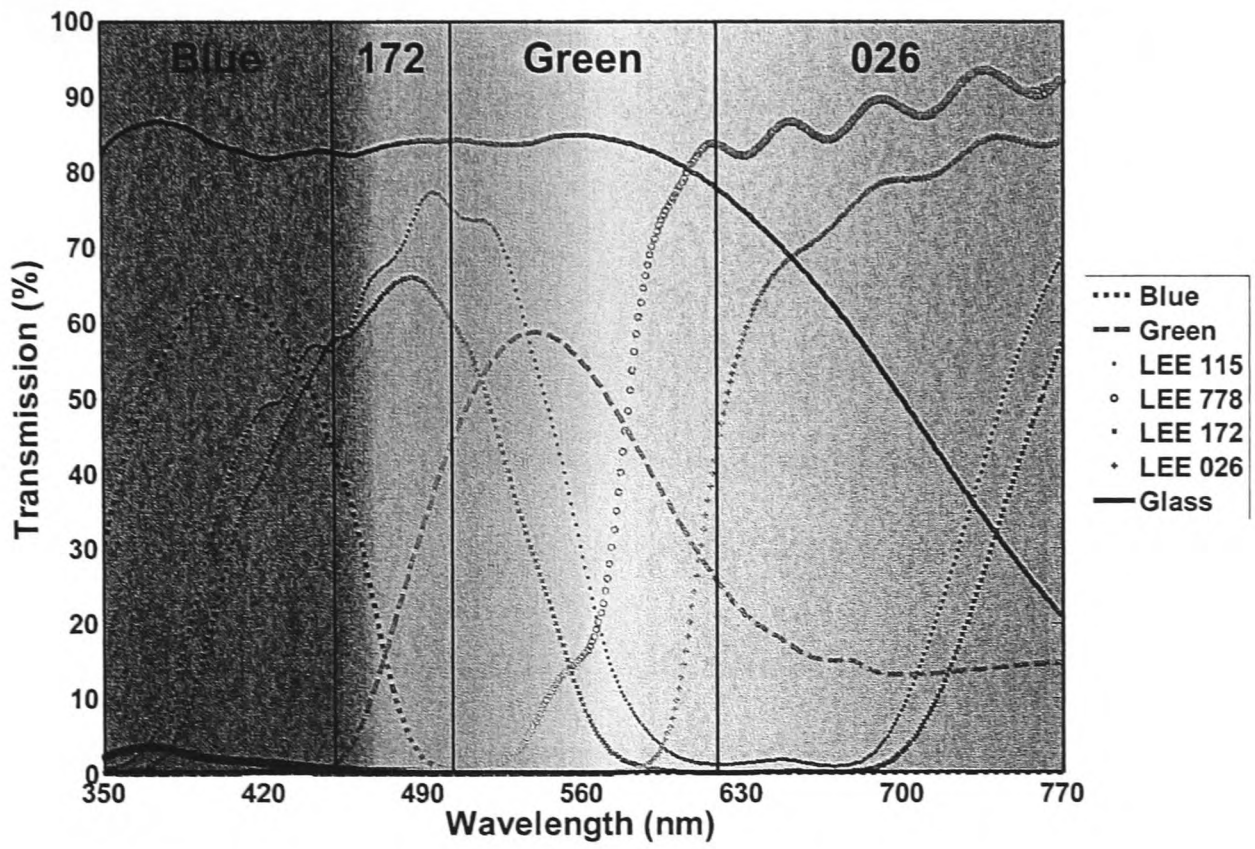


Figure 5.6: Transmission characteristics of filters

To obtain a sufficiently detailed and full calibration curve, the wavelength range of 350–770nm (in steps of 10nm) was chosen and 100 images at each wavelength were taken. Images of the background (with the lamp on and slit s_1 closed) were taken as well for noise subtraction. Figure 5.7 shows images at 470, 540 and 610nm respectively.

General image processing

As described in Section 2.4.3, The Pulnix TMC-9700 camera is a one-chip CCD camera with a colour filter array (CFA) to provide artificial colour information. The staggered pattern of the CFA is shown in Figure 5.8.

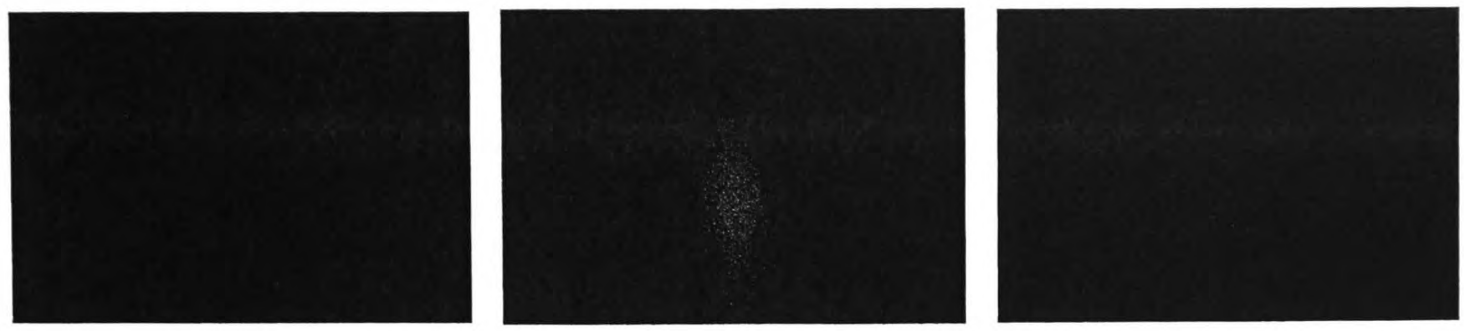


Figure 5.7: Images at different wavelengths as seen by the camera (470, 540 and 610nm left to right)

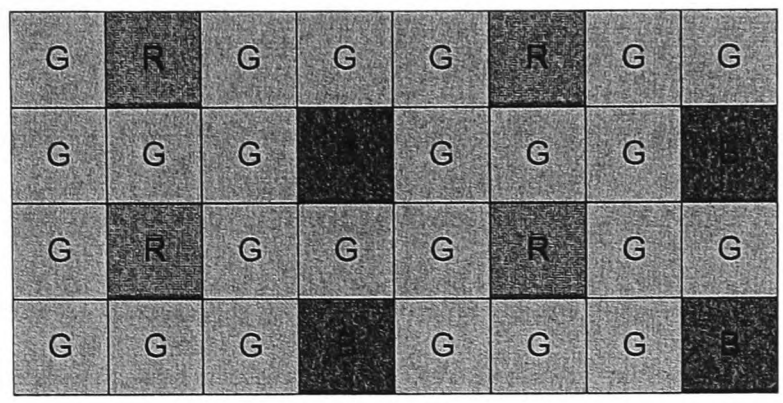


Figure 5.8: Colour filter array (CFA) of the camera

With the CFA present, each pixel only detects the intensity of a part of the whole radiation wavelength range, and the corresponding colour is referred to as its true colour. As the camera outputs RGB values for every pixel, then values of the other two colour channels are derived from on-chip digital signal processing (DSP) This can introduce cross talk between the true-colour channel and the other artificial-colour ones depending on the algorithm that has been implemented to provide the colour reproduction (whether in RGB or HSV). This effect was seen in preliminary work and led to a systematic error in the CRP. Unfortunately, this DSP process is performed before any image is downloaded into the computer and hence cannot be accessed and corrected for by the user. The same is true for most commercial colour cameras, and it should be noted that they may have different CFAs (the most widely used of which is the Bayer pattern [42]). Other on-chip DSP processes can be optionally performed via the camera software including contrast enhancement, edge sharpening and white balance correction, all of which are likely to bias the genuine information in the raw images by introducing artifacts. To minimize the cross talk, raw images from the camera were first processed as follows:

- only the true colour channel for each pixel was kept and the other two channels were discarded;
- then pixel values of the other two colour channels were reproduced through colour demosaicing.

Bilinear interpolation [42], modified for the particular CFA of this camera, was selected as the demosaicing algorithm for two reasons:

- it is straightforward to implement and requires little computational resource;
- interpolation of the artificial colour pixel is only realized within the colour channel, *e.g.* to derive the red value for a true green pixel only its neighboring true red pixels are involved hence avoiding inter-channel cross talk.

The bilinear colour interpolation was implemented differently for the green channel and the red channel while in the same way for the red and blue channels. The interpolation of green value for true red/blue pixels was derived by linear combination of the eight neighboring true green pixels. Take a true red pixel for example:

$$\begin{aligned} G^R_{i,j} = 0.16 \cdot \left(G^G_{i,j-1} + G^G_{i,j+1}\right) + 0.14 \cdot \left(G^G_{i-1,j} + G^G_{i+1,j}\right) \\ +0.1 \cdot \left(G^G_{i-1,j-1} + G^G_{i+1,j-1} + G^G_{i-1,j+1} + G^G_{i+1,j+1}\right) \end{aligned} \tag{5.5}$$

where the superscript denotes the true colour of that pixel, and i and j mean the pixel lies in the i^{th} row and j^{th} column. The coefficients 0.16, 0.14 and 0.1 were derived so that the two closest (in the same row as $R^R_{i,j}$) pixels were weighted more than the two less close ones (in the same column as $R^R_{i,j}$) and the four furthest ones (on the corners of the 3 by 3 region) were weighted the least. This is because the camera CCD sensor is not a perfect square but a rectangle with an aspect ratio of 1.17 and hence a diagonal-to-side ratio of $\sim$1.5. The green value for a true blue pixel was derived in the same way.

The interpolation of the red values for true green pixels was calculated in two steps. Those true green pixels which have two neighbouring true red pixels above and underneath them

were processed first as following:

$$R^G_{i,j} = 0.5 \cdot \left(R^R_{i-1,j} + R^R_{i+1,j}\right) \tag{5.6}$$

Then for the remaining undetermined non-true red pixels (of which there are three in a row between two pixels with known red values), linear interpolation was applied. The interpolation of the blue values for the true green pixels was carried out in the same way. Another benefit of this re-demosaicing process is that it smoothes each individual image (see Figure 5.9).



Figure 5.9: Effect of the re-demosaicing process

The thermal radiation from glowing tungsten or soot at a temperature lower than 3000K is red-dominant whilst the radiation from chemiluminescence is blue-dominant, so pixel saturation is more likely to happen on the red and blue channels than the green one. When pixel saturation occurs on a red or blue channel, pixels within the 3 by 7 array, the center of which is the saturated red or blue pixel, will be biased by the saturated pixel in the on-chip demosaicing process. Therefore, such pixels are assigned 255 at the end of the re-demosaicing code so that they can be readily excluded at a later stage.

5.3.2 Camera response study

Flat field correction

Owing to geometry and artifacts of manufacture, the sensitivity of different areas of the CCD detector may vary. This can be calibrated by exposing the detector to a uniformly illuminated field, namely a flat field. In practice, it is almost impossible to produce a perfectly uniform illumination, so keeping the variation of the field below the readout noise level is necessary, which is likely to be 1 bit for this camera. Hence for common applications where only half of the dynamic range is employed, it is advisable to make sure the flat field has less than 1% variation in intensity with light sensing devices such as a photodiode.

The same lamp as in Section 5.3.1 was put ~2m away from the camera to provide quasi-uniform illumination. The lamp was set to run at maximum current and voltage to generate similar levels of red, green and blue constituents *i.e.* grey light so that all three channels could be corrected for at the same time to a reasonable accuracy. Firstly the bare CCD chip was exposed to the field and a 50mm lens was used to account for the lens effect afterwards. With the lens, two pieces of diffuser paper were put between the camera and the lamp to avoid focusing on the bulb. In both cases, a photodiode was used to verify the flatness of the field.

The captured images were processed similarly in MATLAB as described in the previous sections. By dividing each pixel with the global mean of all pixels, a 2D matrix/image of normalised flat field correction coefficients was obtained and could be used to divide the combustion images to better resolve the spatial temperature distribution estimated by the colour-ratio technique.

It should be noted that when CRP is applied, although the camera is a single-chip one, a flat field correction must be carried out for the three colour channels separately, *i.e.*

three flat field matrices, otherwise any correction coefficient would be cancelled out when a ratio is generated; to some extent, the flat field correction for the colour-ratio technique is similar to a weak white balance correction.

With the same flat field setup, a study of the correlation between the aperture and the camera response was carried out as well. The aperture was assumed to be a perfect circle despite the actual shape (a curvy triangle) and different aperture areas were normalised to the one with a f number of 1.8 (maximum aperture for this lens):

$$\frac{A_{f_N}}{A_{f_{1.8}}} = \left(\frac{1.8}{N}\right)^2 \tag{5.7}$$

where A is aperture area and N is f number. Figure 5.10 shows the non-linear camera response with the aperture area. This suggests that care must be taken when performing analysis on images from different aperture settings. Although the aperture setting is unlikely to affect the results derived from colour ratios, it is sensible to keep using the maximum aperture unless image saturation or depth of field is of concern.

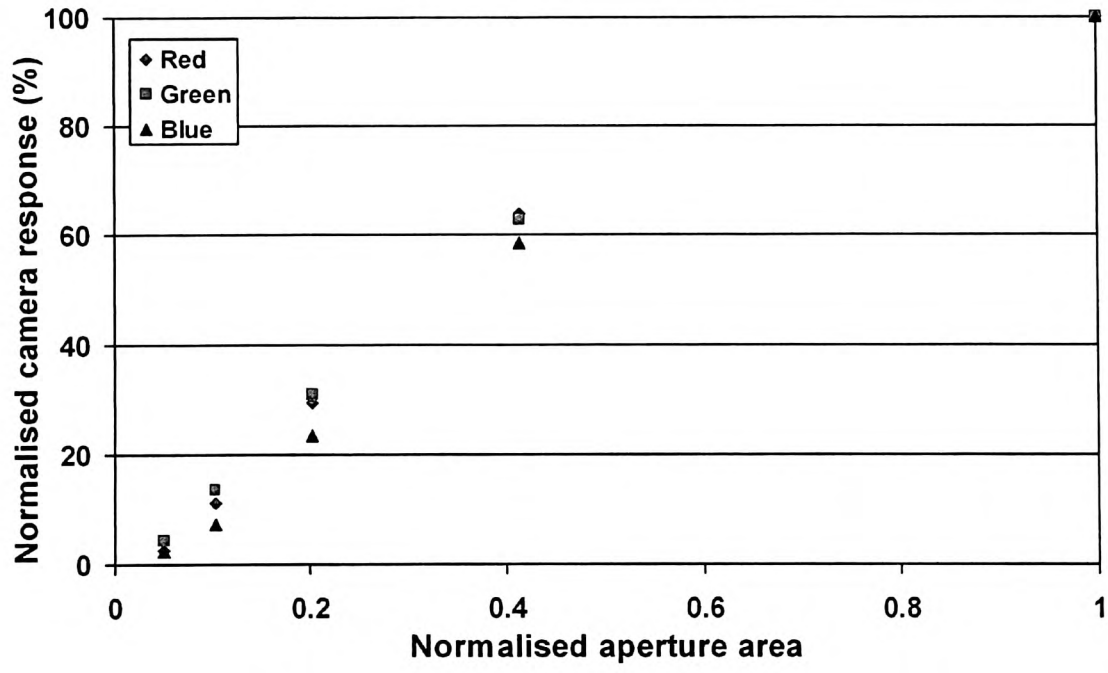


Figure 5.10: The camera's response test for linearity with aperture

Camera response test

Firstly, the camera's dark noise and dark current effect was studied with respect to gain and exposure by completely covering the CCD sensors.

Figure 5.11 (left) shows the correlation between the camera's dark noise and gain with a fixed exposure of 1ms. The camera seems to have a multi-amplifier system and for a gain higher than 190 (maximum 255) the dark noise is minimised both in amplitude and shot-to-shot variation. A higher gain which also minimises dark noise would be ideal in practice but unfortunately this camera features 'sticky bits' with a gain higher than 160, that is, many pixels appear darker than others when imaging a uniformly illuminated object as if they lost 1 or more bits of dynamic range. Therefore, gain 132 was selected at the beginning of this work as a compromise. Figure 5.11 (right) is the average image of 100 dark noise images with gain 132, exaggerated by a factor of 20 to show the pattern of the dark noise.

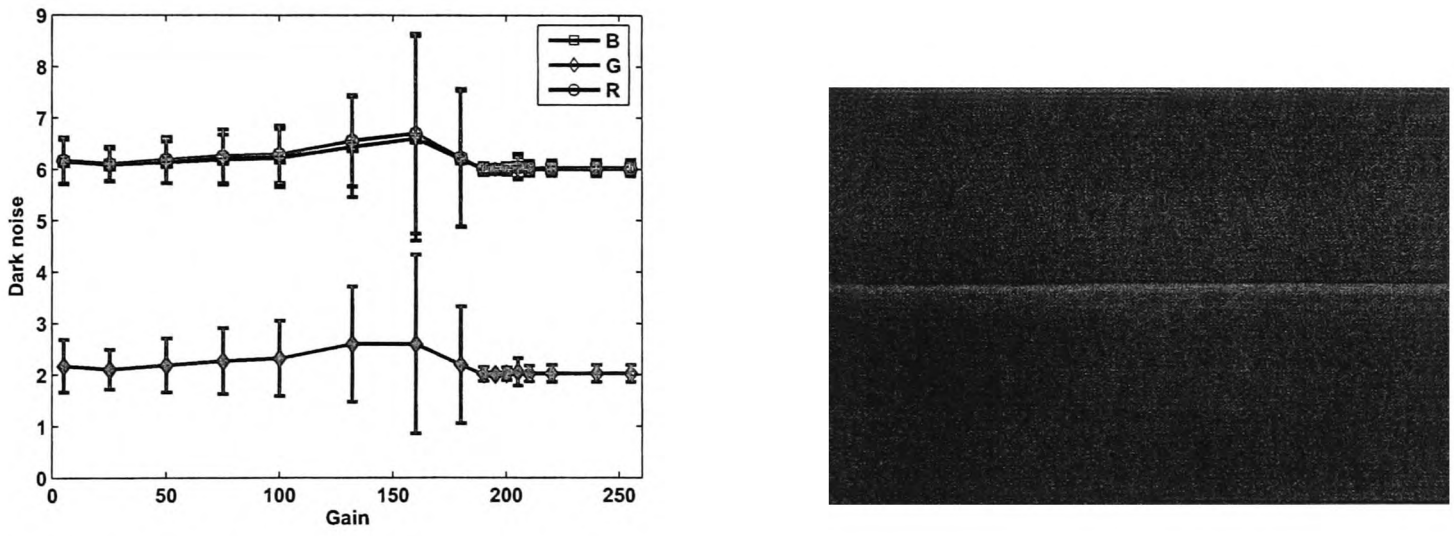


Figure 5.11: Dark noise with respect to gain (left) and amplified dark noise image with gain 132 and exposure 1ms (right)

The dark current effect refers to the fact that more dark noise builds up as exposure time increases. However, this effect is not significant for this camera as shown in Figure 5.12 (left) probably because the camera resolution is low. So the dark current effect is not considered in this research and the dark noise is taken to be constant regardless of

the exposure. Figure 5.12 (right) shows the same pattern as Figure 5.11 (right), also indicating that the dark noise is reproducible and has no correlation with exposure on a statistical basis.

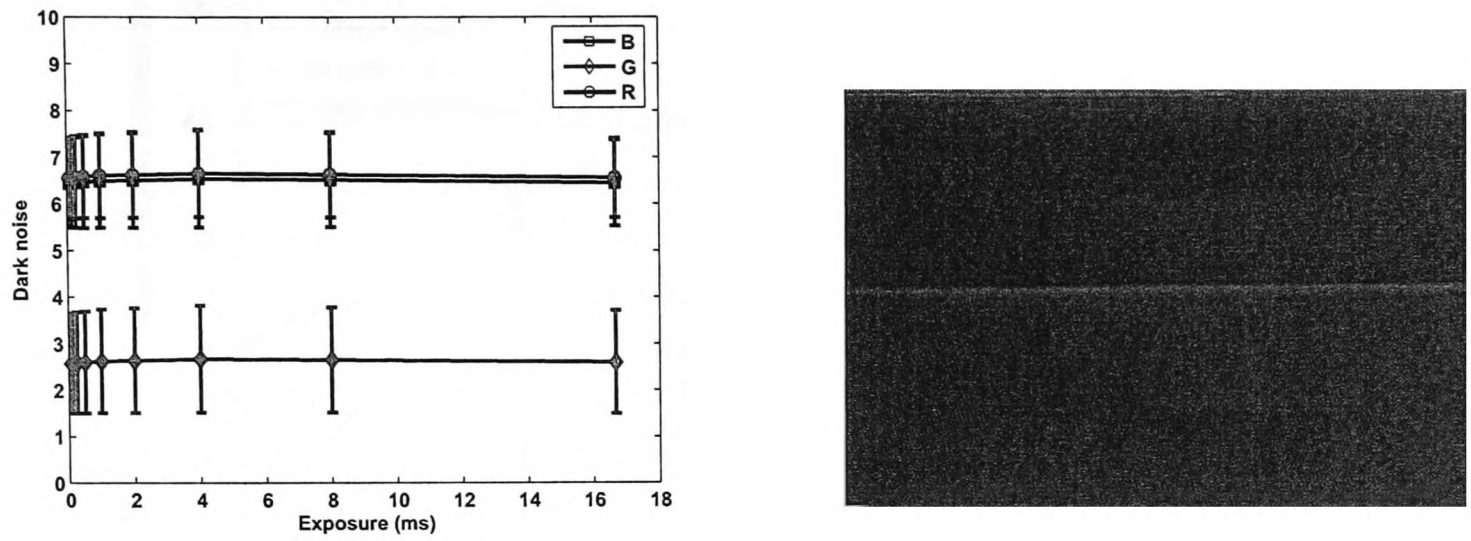


Figure 5.12: Dark noise with respect to exposure (left) and amplified dark noise image with gain 132 and exposure 16.6ms (right)

The camera's sensitivity was tested on the same spectrometer setup with the following approach: images were captured at a series of wavelengths; at each wavelength, a particular shutter setting was first selected without any filter to best utilize the dynamic range. Then different combinations of the seven filters (Figure 5.6) were chosen to reduce the level of radiation into s_1, whilst maintaining the shutter setting for that wavelength. Such images were normalized to the camera's response in the 100% transmission case (no filter). Results regarding the camera's response linearity with sensitivity are shown in Figure 5.13. The responses from the three colour channels are plotted separately. The non-linearity was within 2% over the wavelength range of interest, and hence the camera's response sensitivity was taken to be linear.

The camera's response linearity with exposure was also tested on the same spectrometer setup, so that:

- images were captured at a series of wavelengths;
- at each wavelength, an optimal shutter setting was first selected without any filter;

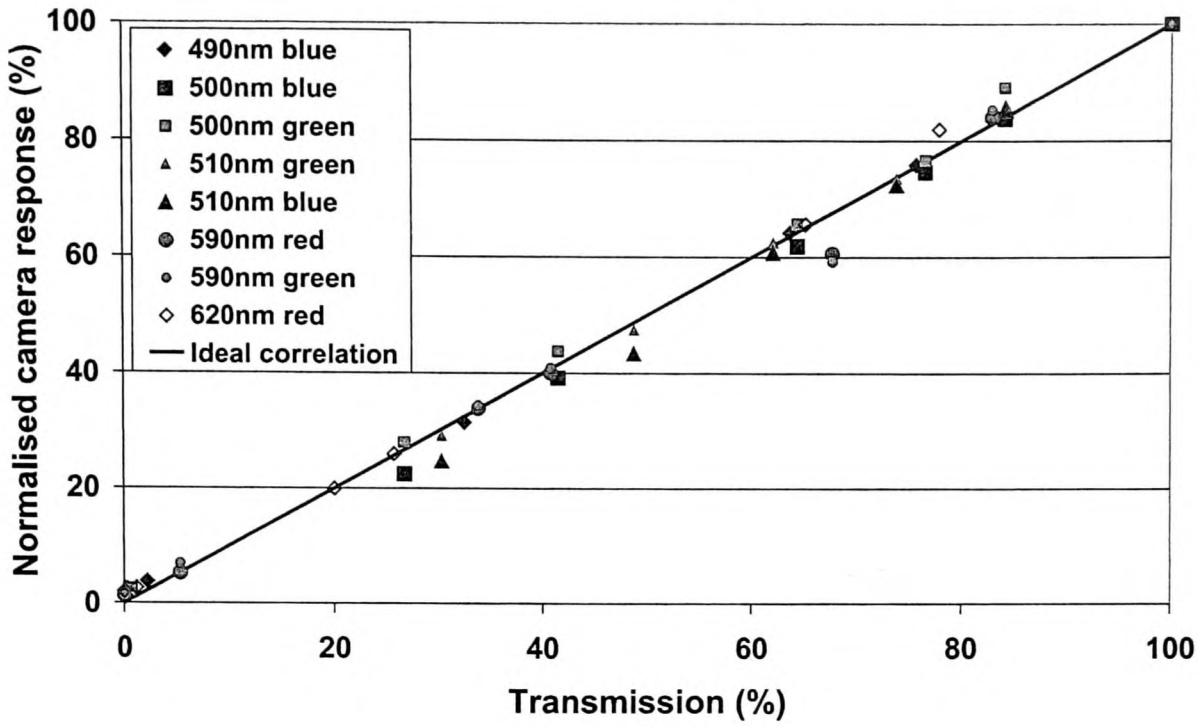


Figure 5.13: The camera's response test for linearity with sensitivity

- then different combinations of the seven filters were chosen to approximately give an each-step-halved transmission so that a doubled exposure time could be used to maintain the camera's response.

Since the camera's sensitivity was proved to be linear, the exposure-response correlation could be derived by

$$R_{i,j} = \frac{I_i}{I_j \cdot \gamma_{\text{tran}}} \tag{5.8}$$

where: $R_{i,j}$ is the response ratio between two shutter settings i and j, I_i and I_j represent measured camera responses with shutter settings i and j, and γ_{tran} is the transmission data involved.

Figure 5.14 presents the camera's response linearity with exposure. The shutter effect was then taken account of by normalizing the spectral calibration images captured using different shutter settings with respect to exposure time.

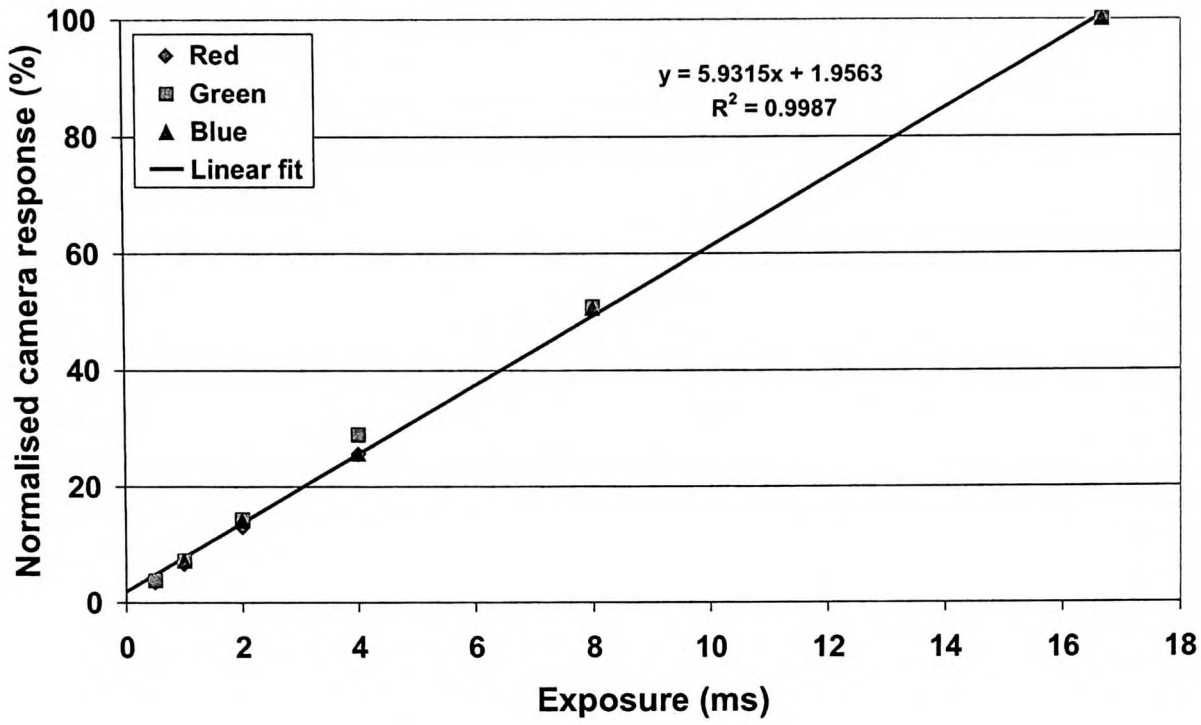


Figure 5.14: The camera's response test for linearity with exposure

5.3.3 Calibration image processing

The diffracted light beam showed an elliptic geometry and its intensity varied across the two dimensions of the image. Therefore only the central part (60 by 50 pixels) of each image was extracted and processed subsequently, and this effectively reduced the bandwidth of the wavelength analyzed to ~1nm.

Because the photodiode has a non-linear sensitivity with respect to wavelength, experiments with lasers of various wavelengths were conducted. Figure 5.15 (left) gave the sensitivity of the photodiode with respect to wavelength. The polynomial fit to these measured points showed good agreement with the datasheet supplied by the manufacturer (the outliers are due to uncertainties and instabilities in laser energy), and hence the manufacturer's calibration was employed in the following steps. This, together with the digital voltmeter (DVM) reading from the photodiode indicated the spectral emission of the tungsten lamp at a particular temperature (Figure 5.15 (right)). The DVM readings were taken when the photodiode detected the bulb radiation through the various filters and hence displayed a multi-stage shape. The 'corrected signal' refers to the photodiode

response after the photodiode sensitivity normalization.

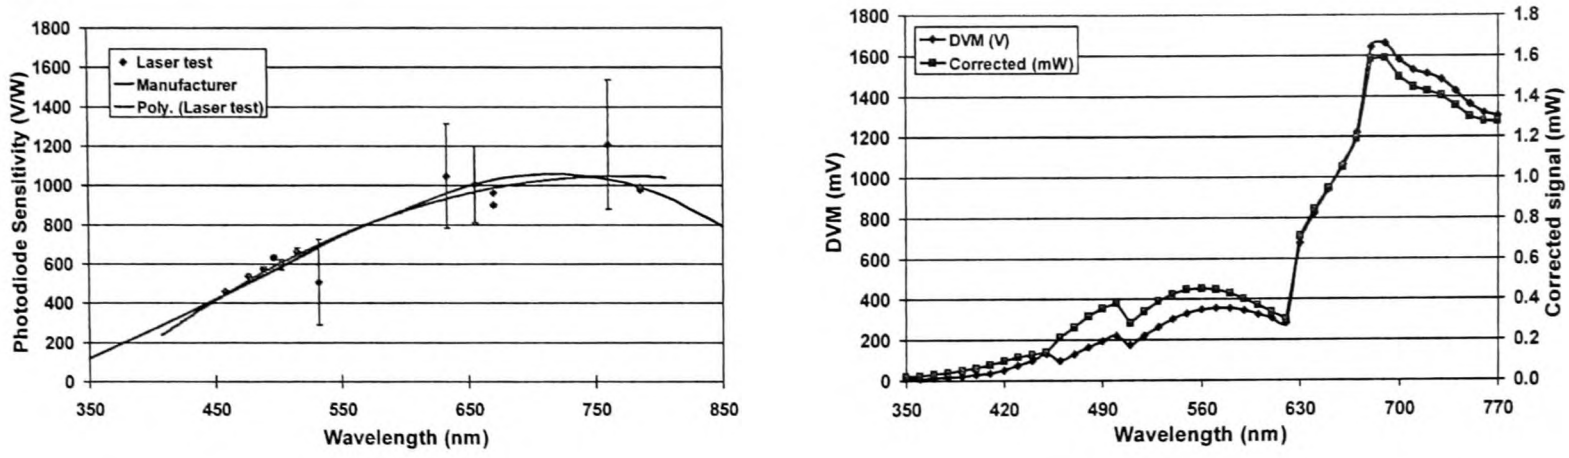


Figure 5.15: Photodiode sensitivity (left) and signal from the tungsten bulb (right) with respect to wavelength

To correct for the intensity variation of the tungsten lamp with wavelength, the RGB components at wavelength λ were normalised to the maximum intensity $I_{\max}$ by multiplying the factor $I_{\max}/I_\lambda$ (yielding Figure 5.16), where I_λ was the intensity at wavelength λ from the photodiode signal (Figure 5.15 (right)).

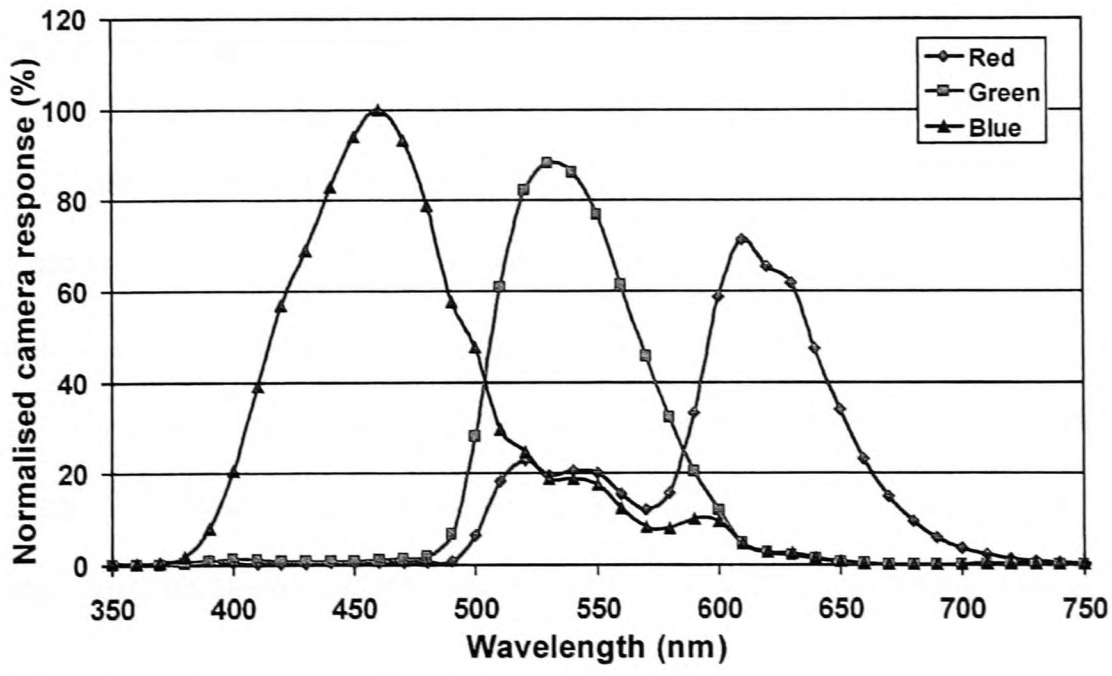


Figure 5.16: Normalized spectral response of the camera

As seen in Figure 5.16, in the range of green (500–600nm) the spectral responses of red and blue follow each other. This suggests that the signal from the green channel has affected the red and blue channels. One possible cause of this cross talk is miss-registration or

misalignment between the CCD chip and the CFA, that is, a signal that falls onto the true-green pixels spreads to the adjacent R or B pixels. As there will be variations between cameras and the published data are only nominal, it is essential to calibrate each camera.

5.4 Validation of CRP

To verify the CRP using the calibrated spectral response of the CCD camera, three sets of experiments were conducted, namely: imaging of a blackbody cavity, a tungsten filament lamp and a tungsten ribbon lamp. A 50mm lens with variable aperture was used for focusing. As only ratios were dealt with in the analysis, the non-linear effect of the aperture setting would not need to be corrected for. The transmission characteristics of the lens and optical pathway were also measured with the spectrophotometer and accounted for in the subsequent analysis.

5.4.1 Black body imaging

The black body used in this experiment was made from a steel bar with a cavity in one end, the diameter of which was an order of magnitude smaller than the length of the tube. The camera was focused to the inside back of the black body through the hole. An aluminum tube was manufactured to cover the camera lens to prevent inclusion of extraneous light. The black body (with a radiation shield) was then heated up by an oxy-acetylene flame, and its temperature was recorded by a thermocouple inserted behind the cavity in the black body.

Temperature readings and images of the black body were taken continuously until the solid end of the black body melted. Figure 5.17 (left) is a typical image of the black body from the experiment.

Comparisons between the measured and inferred temperatures of the black body are

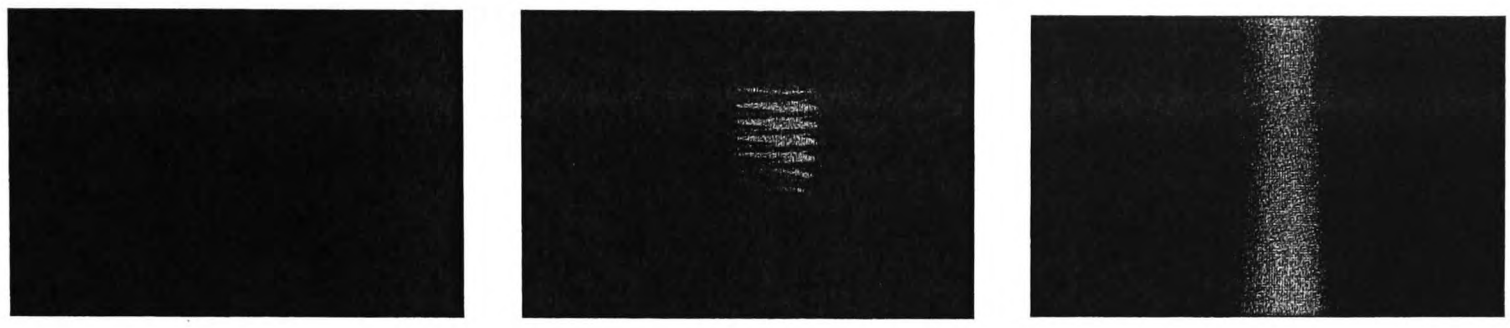


Figure 5.17: Sample images from the validation experiments: black body (left), tungsten filament lamp (middle), and tungsten ribbon lamp (right)

shown in Figure 5.18 together with the corresponding standard deviations. The correlation between results from the R/G ratio and the measured temperatures was better and more linear than that from the R/B ratio. This is because in such a low temperature range the black body would emit much less blue constituents than red and green. Furthermore, the camera's dynamic range is only 8 bits. Therefore, it is difficult to resolve the temperature from the R/B ratio and such results would be associated with large uncertainties. This is also reflected by the standard deviation results. Better precision could be achieved with a 10 or 12-bit camera.

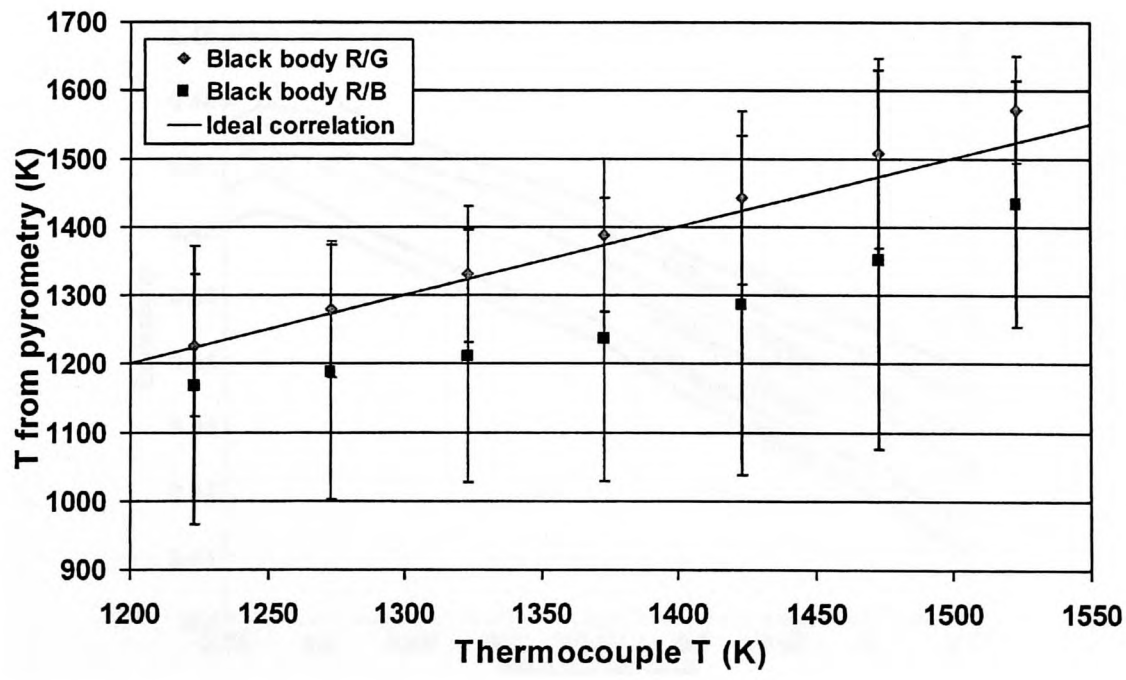


Figure 5.18: Correlation between measured and inferred temperatures of the black body

5.4.2 Tungsten filament imaging

The same tungsten lamp used in the camera spectral calibration process was utilised in this experiment. Figure 5.17 (Middle) is a typical image of the filament from the experiment. Instead of running the lamp under nominal conditions, the voltage and current on the lamp were varied to give different intensities and wavelength distributions of the light, and the voltage and current were recorded to estimate the filament's resistance as a separate temperature indicator. Temperatures from the CRP were then compared with the ones derived from the measured filament resistance (Weast & Astle [70]). The same necessary steps of image processing as above were carried out.

Initially the lamp filament was assumed to be a grey body (constant emissivity). Measured emissivity data for tungsten ribbon lamps at different temperatures and wavelengths (Touloukian & DeWitt [68]) were then used subsequently for improved accuracy. The emissivity data used are plotted in Figure 5.19, but it must be remembered that there are wide variations among the published data.

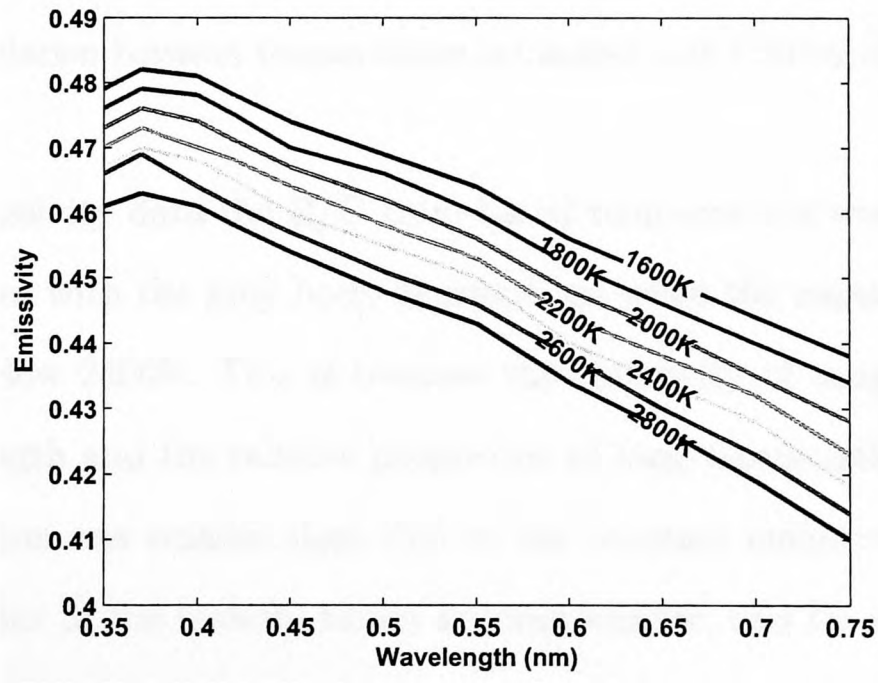


Figure 5.19: Tungsten emissivity data used in the tungsten filament test, data selected from [68]

Figure 5.20 shows the correlation between temperatures predicted by CRP and filament

resistance together with the corresponding coefficient of variation (CoV) for both cases. The CoV presented here was defined as the average of the spatial CoV of each image from the image sequence (100 images under the same condition), hence reflecting the temperature variation across the selected part of the filament.

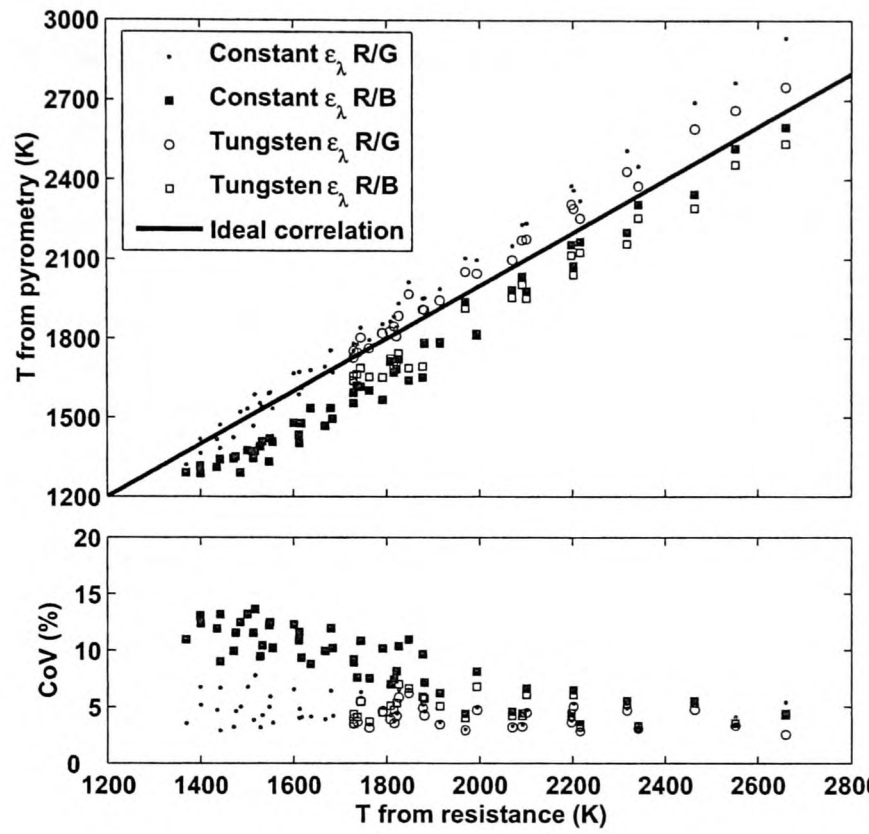


Figure 5.20: Relation between temperatures estimated with CRP and filament resistance

With realistic emissivity data the R/G ratio-based temperatures were generally < 100K lower than the ones with the grey body assumptions when the resistance indicated temperatures were below 2400K. This is because the emissivity of tungsten decreases with increasing wavelength and the relative proportion of long-wavelength radiation to short-wavelength radiation was smaller than that in the constant emissivity case. So the calculated colour ratios in the look-up tables become smaller, and for a given R/G ratio of a pixel, the estimated temperature moves towards the lower values. This effect is not so distinct for results based on the R/B ratio; the R/B data points with realistic emissivity values became lower than the ones with constant emissivity at around 2000K resistance temperature and above. Again this is because the lamp did not emit as much blue light as

green and red at low temperatures and larger uncertainties were likely to be introduced. However, the difference introduced by using experimental emissivity is insignificant within the temperature range for engine combustion. Moreover, the database of measured emissivity is coarse, incomplete, strongly dependant on the surface condition and, to a large extent, self-inconsistent for the wavelength range of interest here [68]. In conclusion, the temperature calibration against a tungsten filament should not be treated as an absolute reference because of uncertainties in the emissivity and uncertainty in the temperature deduced from the filament resistance.

In summary, the temperatures derived from the resistance of the filament agreed with the ones from the colour-ratio method better for R/G ratios at low temperatures (< 2200K) and better for R/B ratios at high temperatures (> 2200K). Again, temperatures derived from the R/B ratios are better resolved at high temperatures, indicating the potential of using R/B results as a supplement for checking the R/G results in the high temperature range.

The discrepancies from the unity line in Figure 5.20 may partly be due to uncertainties in the inference of the temperature from resistance, *e.g.* the non-uniform temperature distribution across the filament, the different resistance-temperature behavior of other parts of the lamp *etc.* The overall CoVs of this experiment scatter within 3–7% in the range of 1600–2600K except those low-temperature R/B points with constant emissivity.

5.4.3 Tungsten ribbon imaging

A high-stability tungsten ribbon lamp (POLARON 10/V) was employed for its capability of calibrating temperature measuring devices. A typical image of the ribbon lamp from the experiment is shown in Figure 5.17 (right).

The lamp was calibrated in terms of radiance temperature by reference to a National Physics Laboratory (NPL) standard tungsten ribbon lamp which has been calibrated

using a precision radiation thermometer, which has itself been calibrated with reference to a blackbody source at the freezing point of gold (~1064°C). At each calibration current the radiance temperature of the test lamp is determined by comparison with the calibrated standard lamp using a comparator.

The radiance temperature corresponds to the temperature of a black body at the calibration wavelength (660nm for this lamp). The radiance temperature is not the true ribbon temperature, but the temperature that a black body would have to be at, in order to emit the same intensity of radiation at the specified wavelength. Therefore, conversion between the archived radiance temperature of the lamp and the actual ribbon operating temperature must be done based on Planck's Radiation law and an estimate of the ribbon emissivity. Rewriting Eq. 5.1 and 5.2

$$I_{\lambda b} = \frac{2C_1}{\lambda^5 \left(e^{C_2/\lambda T_b} - 1\right)} \tag{5.9}$$

$$I_{\lambda t} = \varepsilon_\lambda I_{\lambda b}(\lambda, T_b) = \varepsilon_\lambda \frac{2C_1}{\lambda^5 \left(e^{C_2/\lambda T_t} - 1\right)} \tag{5.10}$$

with subscripts b and t representing black body and tungsten respectively. Letting $I_{\lambda t}$ equal $I_{\lambda b}$ and substituting Eq. 5.9 into Eq. 5.10, the correlation between the temperature of the tungsten ribbon T_t and its corresponding blackbody temperature T_b (also the radiance temperature) can be written as:

$$T_t = \frac{C_2}{\lambda_c \ln\left(\varepsilon_\lambda e^{C_2/\lambda T_b} + 1 - \varepsilon_\lambda\right)} \tag{5.11}$$

As shown in Figure 5.21, a fairly linear correlation existed between the estimated temperatures with CRP and the deduced lamp operating temperatures, it being more linear for the R/G results than the R/B results. The R/B points seemed to match with the lamp operating temperature better than the R/G ones, however, they were more scattered

and the difference between the two series decreased with increasing lamp temperature. Again there is uncertainty in the emissivity of the tungsten ribbon. This manifests itself as uncertainty in its true temperature and as uncertainty in the spectral intensity of the radiation so the ribbon lamp does not provide an absolute reference.

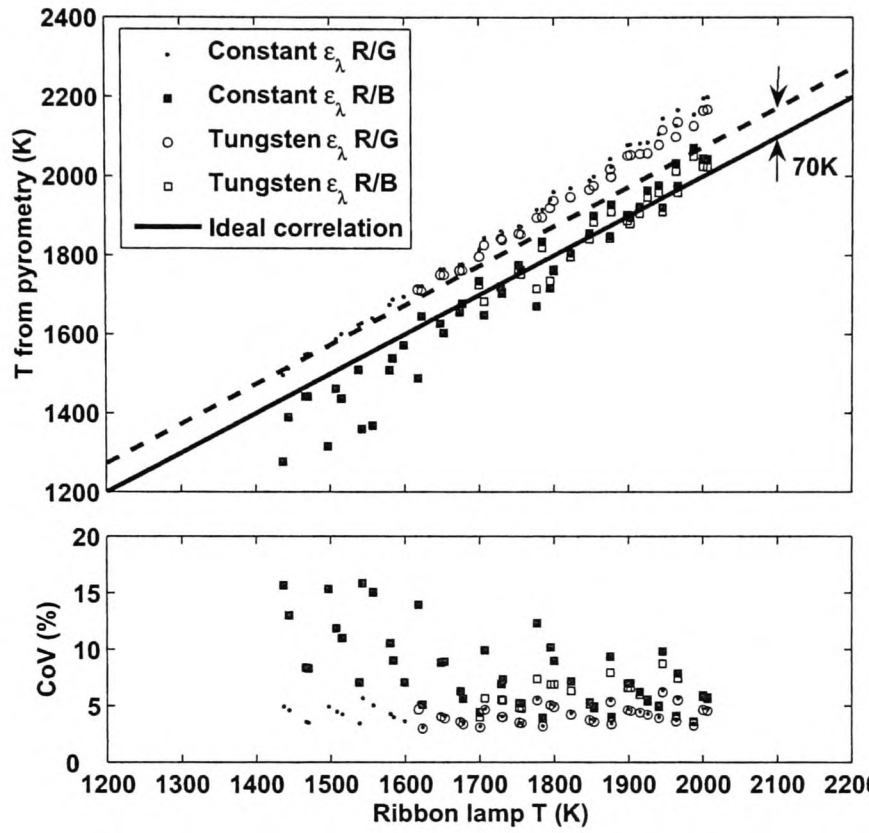


Figure 5.21: Relation between temperatures estimated with CRP and the estimate of the ribbon lamp operating temperature

An examination of the spectral output of the tungsten ribbon lamp in the range 640–910nm at each temperature suggests that there might be a systematic error in the black-body calibration of the lamp temperature; this could be an aging phenomenon. In Figure 5.21 a broken line has been added to account for this deduced temperature offset of 70K.

The overall CoVs of this experiment fall within 3–10% for the temperature range of 1600–2000K except those low-temperature R/B points analysed assuming a constant emissivity.

5.4.4 Concluding remarks on the temperature calibration

In comparisons between the temperatures derived from CRP and the reference temperatures provided by the three validation experiments, the R/G colour ratio results showed excellent linearity; the R/B results were more scattered but were useful as supplementary information, particularly for the high temperature range — this may be improved using a camera of higher dynamic range. To minimize the possible error due to all the noise sources and best utilize both the R/G and R/B colour ratios, a weighted-mean temperature T_c was adopted as the output from the CRP:

$$\begin{aligned} T_C &= a \cdot T_{R/G} + (1-a) \cdot T_{R/B} \\ a &= 1.5 - T_{R/B}/2000 \\ |T_{R/G} - T_{R/B}| &< 200\text{K} \end{aligned} \tag{5.12}$$

where a and $1-a$ are the weighting coefficients with an applicable range of 1000–3000K. A premise of this calibration is that the difference between the temperatures from R/G and R/B estimates must be below 200K. Any image not compliant with this premise is likely to be due to chemiluminescent flames and hence was not taken as valid.

Throughout the range of 1200–2800K, the weighted temperature from the blackbody and tungsten filament experiments showed a deviation of less than ±50K from the ideal correlation and the one from the tungsten ribbon experiment is always 50–100K higher than the reference temperature. Taking account of the 70K systematic offset mentioned in the previous section, the deviation in the tungsten ribbon experiment is also less than ±50K.

5.4.5 Error analysis

Firstly, the ±1 bit digitization error caused by the A/D conversion process of the camera was investigated. Such an error in the colour values (intensity) can lead to an error in the

colour ratios, which results in an intrinsic error on the temperature estimate. Quantifying this error helps to understand the accuracy of the colour-ratio method and set a reasonable step size in the look-up tables. The digitization error can occur on any one, two or three colour channels at the same time, therefore the worst case, in which there was +1 bit error in the red channel and -1 bit error in the green and blue channel was chosen for this study.

For a range of temperatures, a relative intensity of 127 was assigned to the red component, which was stronger than signals from the other two channels, and the green and blue values were calculated by dividing the red value by the ratios in the look-up tables. Then 1 was added to the red values and subtracted from the green and blue values to generate new colour ratios. These new ratios were mapped back using the original look-up tables and the error-added temperatures were returned. The error from the R/G and R/B ratio calculation is presented in Figure 5.22.

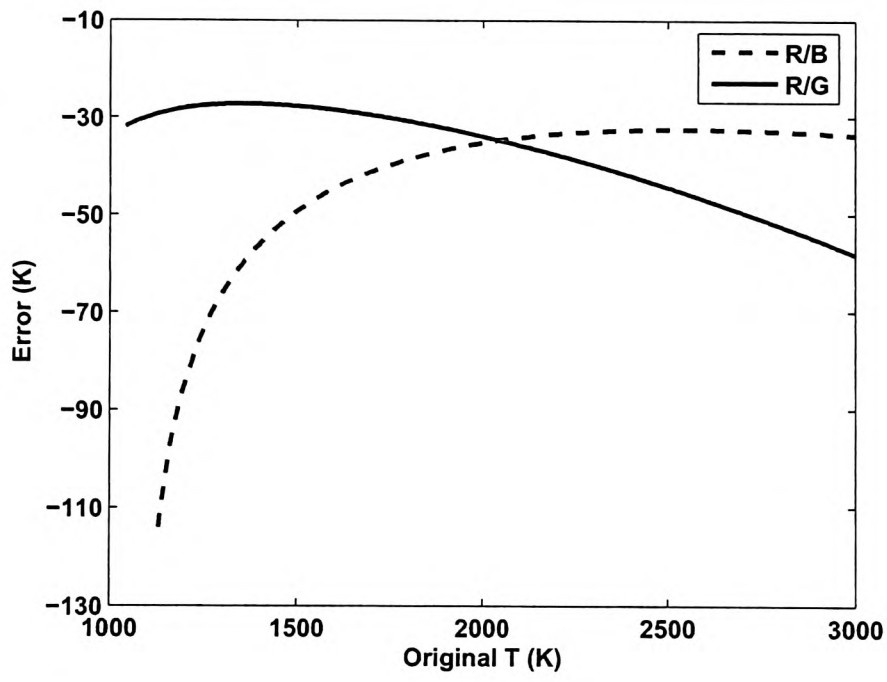


Figure 5.22: The effect of 1 bit digitization error on the estimated temperature

The step size of the look-up tables used for Figure 5.22 was 10K and was found to give reasonable performance. The error presented is always negative because this case referred to ± 1 bit error introduced to the numerator and denominator of the colour ratios

respectively. The error curves would be approximately inverted for the opposite case.

Since the 1 bit digitization error may not happen on all three channels at the same time, the case presented here indicates the worst-case error. However, the relative intensity of 127 may not be achieved and the utilisation of the dynamic range may be poor in practice. Taking these two factors into consideration, error in the temperature estimation due to the ± 1 bit digitization error can be predicted to be within $\pm$50K for the temperature range of engine combustion (1500–2600K). It should also be noted that although Figure 5.13 indicates linearity uncertainties of $\sim$2% (2–3 counts for a camera response of 127 counts), the source of such uncertainties are most likely to be a combination of both shot-to-shot noise and digitization error which cannot be decoupled from one another. Therefore, the ± 1 bit digitization error is still representative of most cases encountered in practice.

The ± 1 bit digitization error analysis above was applicable to a single pixel, whilst the effect of random noise on a global averaging basis was studied as well since spatial averaging was applied in the validation experiments. In this study, standard artificial images of uniform colour ratios were produced according to the look-up tables for a series of temperatures. Then different levels of random noise was added to the three colour channels of these images, namely, noise-to-signal ratio (NSR) of 0.05–0.25 in steps of 0.05, in which case signal means the colour value of a particular colour channel in the images. These new images were processed using the original look-up tables (step size 10K), and the returned error-added temperature field was spatially averaged. Figure 5.23 shows the error from the R/G and R/B ratio calculation.

The resultant error displayed a feature with respect to the original temperature as opposed to a random distribution that one would expect. However, this is a genuine fact in particular for the colour-ratio technique presented here. Firstly, the algorithm for CRP was interpolation with fixed KL in which the boundaries of the look-up tables for interpolation had to be set before any calculation was carried out; in this study the boundaries were 1000 and 3000K. Therefore, as the original temperature approached either boundary,

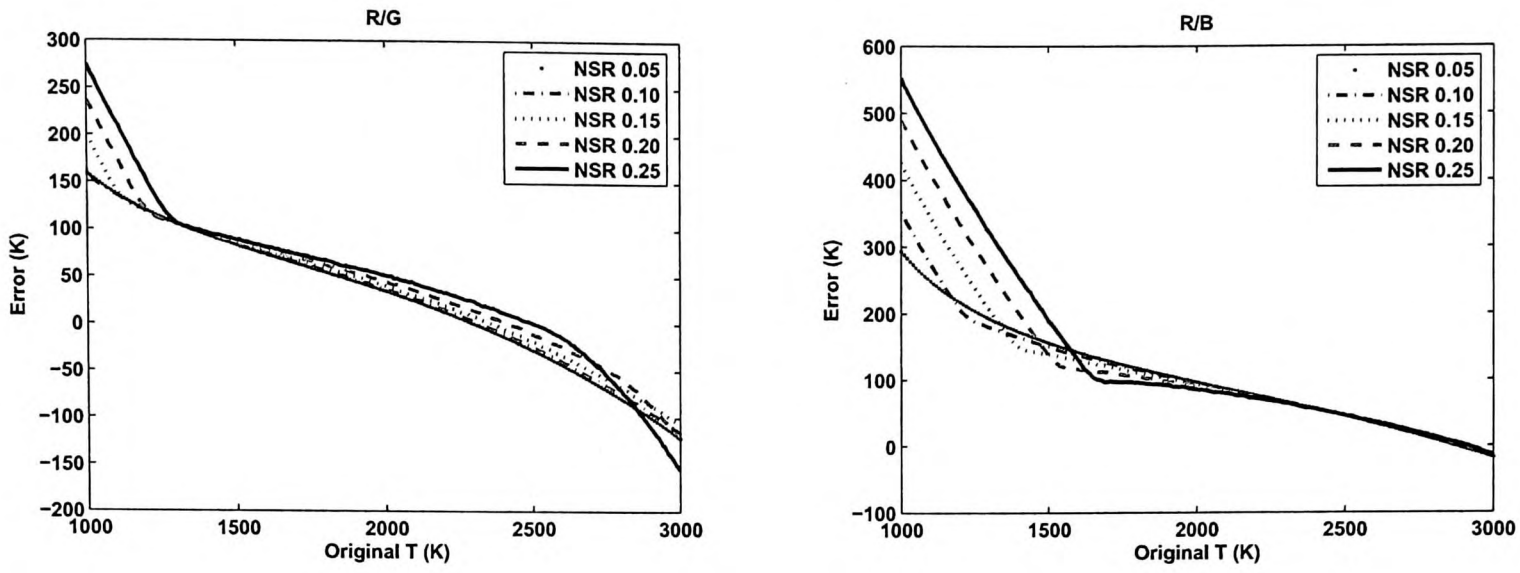


Figure 5.23: The effect of random noise and global averaging on temperature estimation

the random noise distribution added to the original temperature field was cut off on one side and became distorted, resulting in the increasing deviation from zero on both ends of the x axis in 5.23. One reason why the error behaved differently at low temperature and high temperature may well be the shape of the colour ratio curves in 5.4, which monotonically decrease with decreasing gradient. For instance, adding the same amount of numerical noise (NSR) to the standard image at 1000K and 3000K would result in a different amount of temperature error due to the different scales of interpolation. By this means, the temperature error can be estimated to be 50K for the temperature range of 1800–2600K using R/G results and 2400–3000K using R/B results.

Recalling the deviation of the experimental results from the 'unity' line in the above sections, this deviation is likely to be the combined error of the ± 1 bit digitization error of individual pixels, the spatially averaged random error and the uncertainties in the experiments themselves.

5.5 DISI engine results

A typical arrangement for the pyrometry measurement on the Jaguar engine is shown in Figure 5.24. An image sequence of 100 images was taken for each engine operating condi-

For systematic effects in the evaporation to be smaller than the typical precision of an engine PLIF measurement, the co-evaporation ratio should lie between 0.85 and 1.15 whenever more than 5% of a fuel fraction is in the vapour phase.

The component fuel used in this project was supplied by Shell Global Solutions and consists of 3 pairs of components. Stevens *et al.* [58] describe the process of modelling the component fuel and the tracer/fuel-component co-evaporation. Appropriate tracers may be added to the multi-component fuel, so as to co-evaporate with different fractions of the fuel such that each tracer will identify the spatial distribution of that fraction in the engine. The separate identification of different fuel fractions is essential for understanding the progressive vaporisation of the fuel and the way in which the variation in local air/fuel ratio develops. Issues of co-evaporation have been identified recently by a number of authors, in particular the co-evaporation of pure iso-octane reference fuel with ketonic tracer [22].

Early work to examine the effect of different fuel fraction led to the addition of tracers of different volatility to a single fuel, *e.g.* [64]. In these cases, quantitative interpretation was not possible given that the co-evaporation of the fuel fraction and tracer is not provable. Two-component exciplex tracer fuels have been modelled and used for quantitative PLIF diagnostics [62]. A three-component fuel has been demonstrated with ketonic/aromatic tracers [22]. This fuel is most similar to the fuel used in this project, yet this fuel/tracer set has been shown not to meet the criteria for co-evaporation as defined earlier in this section.

T_{10}, T_{50} and T_{90}, the temperatures at which 10, 50 and 90% of the fuel has evaporated, were nominated as the typical boiling points for the volatility classes of prime interest: *i.e.* a light (volatile), a medium and a heavy fraction. The temperatures of these classes for Dutch Pura (from an ASTM curve) are given as ~60°C, 100°C and 170°C. This provided a basic scope for selecting the tracers.

A comprehensive review of tracer-LIF diagnostics has been presented by Schulz *et al.* [54] including extensive discussions on the properties of organic tracers. The two most frequently reported tracers are acetone and 3-pentanone. 3-pentanone gives a fluorescence yield a few times stronger than acetone and has a boiling point (102°C) close to iso-octane (99.2°C, most representative compound of gasoline), however, it only co-evaporates with iso-octane when in a mixture of 20% 3-pentanone [58] and it degrades in the mixture. The high vapour pressure of acetone (boiling point 56°C) makes it an ideal tracer for the light fraction. A feature of the ketone family is that the absorption spectrum of acetone shifts red-wise as the temperature gets higher; this effect has been accounted for and discussed in Section 4.4. Sercey *et al.* [14] used gas chromatography to study the fuel breakdown phenomenon and found that iso-octane degraded with acetone present in the mixture in a motoring engine environment. Detection of acetone fluorescence also suffers from its low quantum efficiency and short lifetime. Nevertheless, the start of the distillation curve of a typical gasoline (in the region up to 80°C) is only covered by a small set of aldehydes and ketones, the least inappropriate fluor of which is acetone in terms of overall benefits. Additionally, to co-evaporate with the light fraction, acetone is the best choice. Figure 4.7 describes the evaporation of iso-octane with either acetone or 3-pentanone and shows considerable deviation from co-evaporation. The addition of the ketones to the the mixture increases the overall volatility of the mixture by around 10% at ~303K [58].

Toluene is another widely studied tracer species for combustion research. It is present in gasoline and hence was selected to match T_{50}. For the heavy fraction, there is a large group of suitable aromatics, so 1,2,4-trimethylbenzene (TMB) was chosen simply for a high fluorescence yield and a close match of boiling point to T_{90}.

To find a fuel blend that gives satisfactory co-evaporation, some preliminary vapour liquid equilibrium (VLE) calculations were performed by colleagues from Shell Global Solutions with an in-house Shell program that uses a modified Redlich-Kwong equation of state and liquid activity coefficients derived from the UNIFAC group contribution method [58].

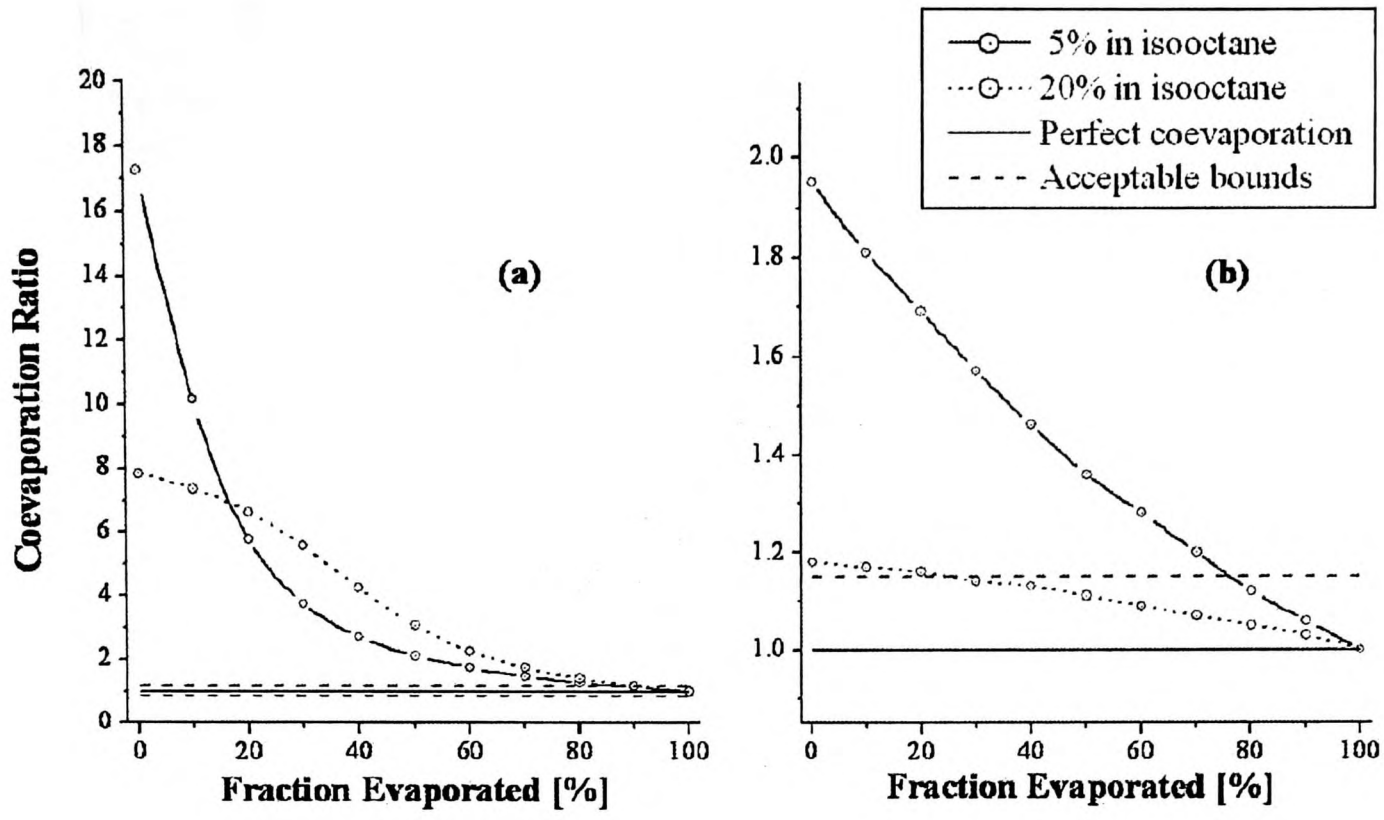


Figure 4.7: Coevaporation of (a) acetone/isooctane, (b) 3-pentanone/isooctane mixtures at 303 K [58]

With the chosen three tracers, a three-compound blend was initially identified by the VLE calculations such that each compound co-evaporated acceptably with one of the tracers. Unfortunately, these compounds could not be sourced economically in practical quantities. An iterative design procedure was then taken with a six-compound multi-component fuel blend, using three pairs of readily available compounds such that one compound from each of three pairs was more volatile, and the other less volatile, than the corresponding tracer species. By selecting suitable pairs and adjusting the proportions, a composition was found for which, according to the VLE calculations, each pair evaporated in the same proportion to the tracer to within the required tolerance.

Table 4.1 and 4.2 summarise the properties of the component-fuel components and selected tracers. On blending the six-component mixture, it was found that n-butane was most readily available as a mixture with other C_3 and C_4 compounds. So a final small adjustment was made to the blend for this light fraction mix.

Figure 4.8 shows the co-evaporation ratios for the multi-component fuel. The co-evaporation

Species		T_{boiling}(°C)	Density (g/cm^3)	Volume fraction (%)
Light	butane (mix)	-1	599 (25°C)	11.7
fraction	iso-pentane	30	626.2 (20°C)	18.7
Medium	iso-octane	99	691.9 (20°C)	44.1
fraction	n-octane	127	703 (25°C)	10.9
Heavy	n-decane	174	726.3 (25°C)	2.2
fraction	iso-dodecane	~180	754.6 (25°C)	12.4

Table 4.1: Properties of component-fuel components

Species	Family	T_{boiling}(°C)	Density (g/cm^3)	LIF score	LIF band (nm)	LIF lifetime (ns)
acetone	ketones	56	789.9 (20°C)	1	300–600 (liquid) 350–550 (vapour)	4.0
toluene	aromatics	111	866.9 (20°C)	~700	270–330	177
TMB	aromatics	168	880 (20°C)	~900	270–350	77

Table 4.2: Selected common tracers and their properties taken from the literature. Note: spectral properties vary considerably with the environment and phase — the values also vary considerably in the literature. Therefore, these values should be taken as a guide only.

ratios (with respect to each fuel component pair) of the tracers across the range of fractions evaporated compares favourably with the chosen criteria and the accepted pure component fuel co-evaporating pairs. The co-evaporation ratio for the TMB tracer is relatively high for small fractions evaporated (>1.15 at <10% vapour). This could be corrected by using more iso-dodecane in the blend but this would be of little benefit as less than 5% of the heavy fraction is vaporised until the overall fraction evaporated exceeds about 50%. The region where more than 5% of the heavy material is evaporated is marked by the bold segment of the line in Figure 4.8. The stated blend gives the best calibration in this region.

It has been shown [58] that there is little variation in the overall fuel volatility when 5% tracers are added to the component fuel.

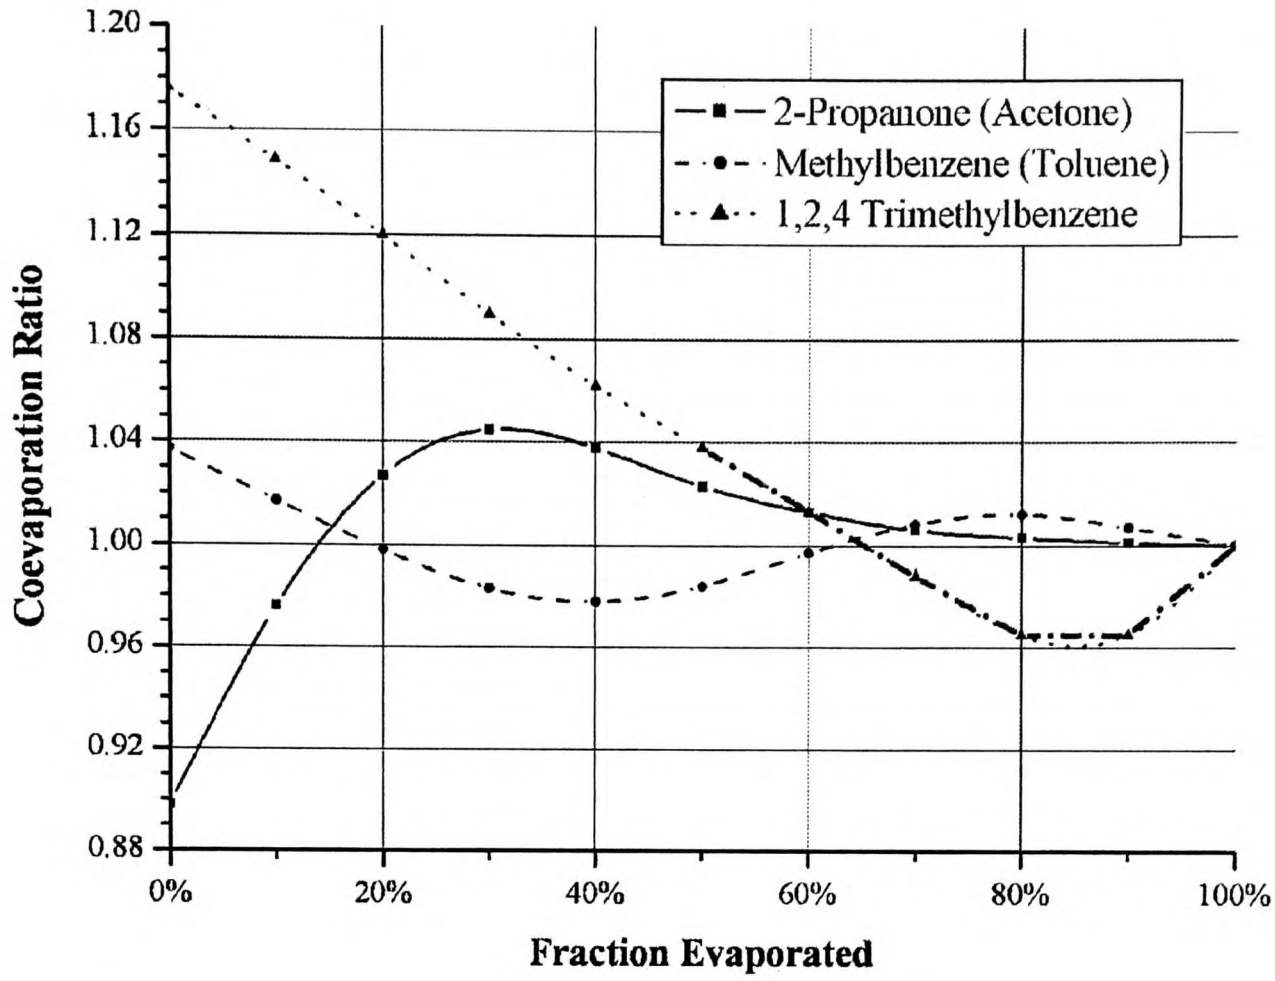


Figure 4.8: Coevaporation ratios at 303K according to equilibrium VLE calculations. The bold region of the trimethylbenzene curve shows the region where greater than 5% of the heavy fuel fraction has evaporated [58]

4.5.3 Characterisation of fuel components and tracers

Preliminary LIF tests on sample component fuel showed that it generated noticeable LIF signals, which were considered to be from contaminants not removed in the distillation process. Therefore, two sets of experiments, namely absorption and LIF spectroscopy with a sample cell, were conducted on all the liquid tracers and fuel components in their liquid phase. N-butane is gaseous under atmospheric conditions and is less likely to contain contaminants that fluoresce, since it is the higher boiling point hydrocarbons that are most likely to fluoresce.

In the absorption experiments, a 10mm thick cell was used and in order not to saturate the spectrophotometer, samples of pure fuel components and solutions of 1% acetone, 0.1% toluene and TMB in spectroscopically pure iso-octane (HPLC grade) were used.

Results from the first batch of component fuel and the components supplied by Shell are presented in Figure 4.9. The 'Shell component fuel' plot is from a sample of the actual shipped batch; the 'Oxford fuel' one is from a sample mixed up at Oxford with the five shipped fuel components according to the specified composition, and the 'Calculated fuel' is derived from the weighted sum of the five components using Eq. 4.2 to give

$$\text{abs}(\%) = 100\,(1 - I/I_0)\,\% = 100\left(1 - 10^{-\sum_i c_i A_i}\right)\% \tag{4.9}$$

where c is the relative quantity of component i.

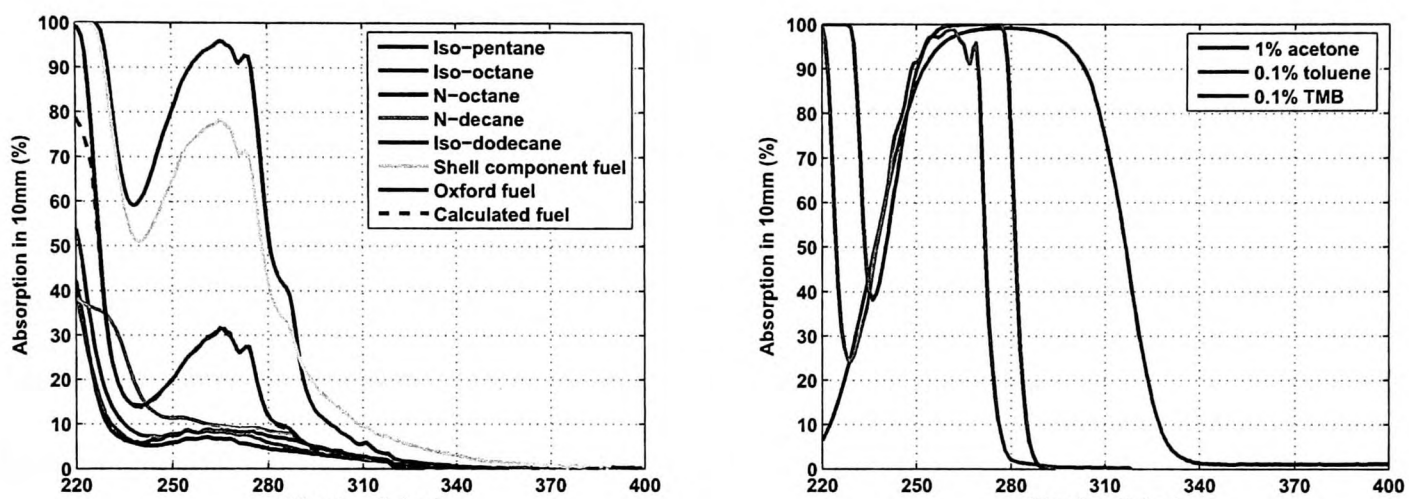


Figure 4.9: Absorption spectra of first-batch components and tracers

It is apparent in the left figure that iso-dodecane and n-decane both have considerable contaminants of some sort. This together with the composition of the component fuel makes the contribution from iso-octane and iso-dodecane more significant than the others, and hence a request for a 'purer' version of these two components was made. The 'Calculated' spectrum matches the one from the Oxford mix both in shape and in magnitude but only in shape with the Shell mix, which indicates: 1) the actual composition may differ from the nominal recipe or 2) more contaminants of a similar sort were introduced in the mixing process at Shell.

Figure 4.10 shows the absorption spectra of the second batch of iso-octane and iso-dodecane. With the quality of the new iso-dodecane much improved but that of the

new iso-octane even worse, the old iso-octane and new iso-dodecane were used to predict the spectrum of the new component fuel, which is seen to be reasonably 'transparent' (left hand side figure, ~10% absorption around 270nm).

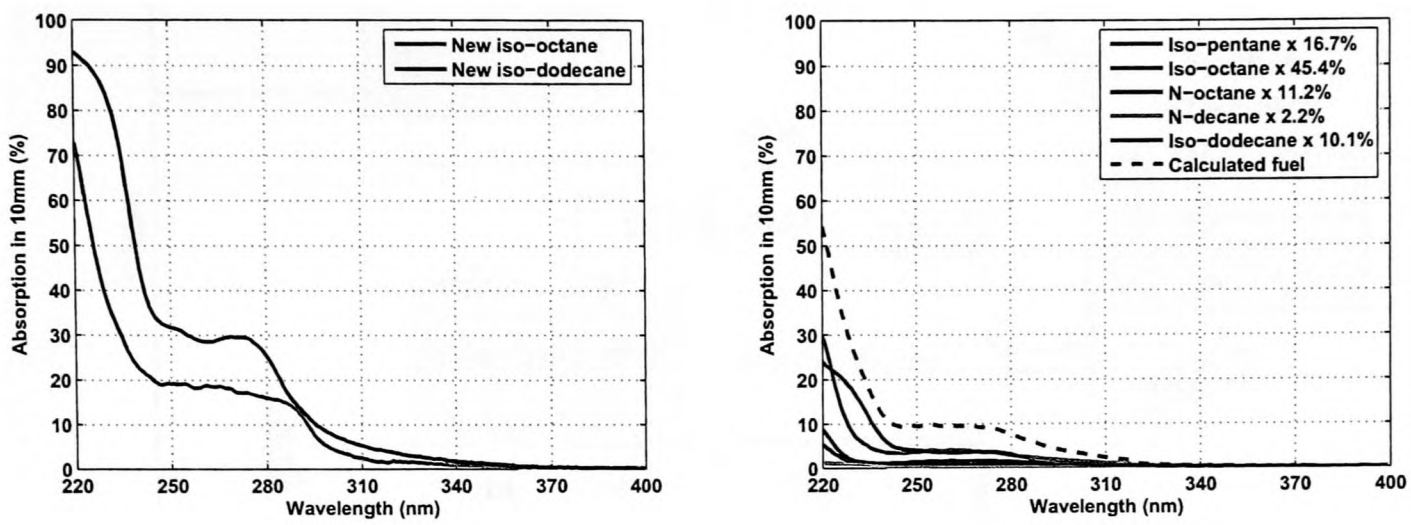


Figure 4.10: Absorption spectra of second-batch components

Knowing the absorption features of the fuel components and tracers, it was advisable to investigate how much fluorescence the component fuel might produce and to estimate the signal-to-noise ratio (SNR) with respect to the LIF signals from the tracers. Figure 4.11 depicts the layout for the LIF experiment.

The beam from the KrF excimer laser was focused into a cell filled with the sample in question, the same as was used in the absorption tests. The ensuing fluorescence was led through a spectrometer for the spectroscopic study. Monochromatic light was then detected by the photomultiplier tube (PMT). The motor on the spectrometer was set to run at 25nm/min so as to automate the spectra acquisition. Due to variations in laser intensity, a photodiode (ThorLabs DET210) was used to collect a portion of the laser light for data normalisation. The boxcar integrator acted as a sample-and-hold device and took both the laser pulse and fluorescence signal over a short predefined gate, longer than the fluorescence lifetime. Data of certain cycles within the hold time (~100ms) were averaged and logged by the DAQ system.

Data with the cell empty was subtracted from all other data sets as background and

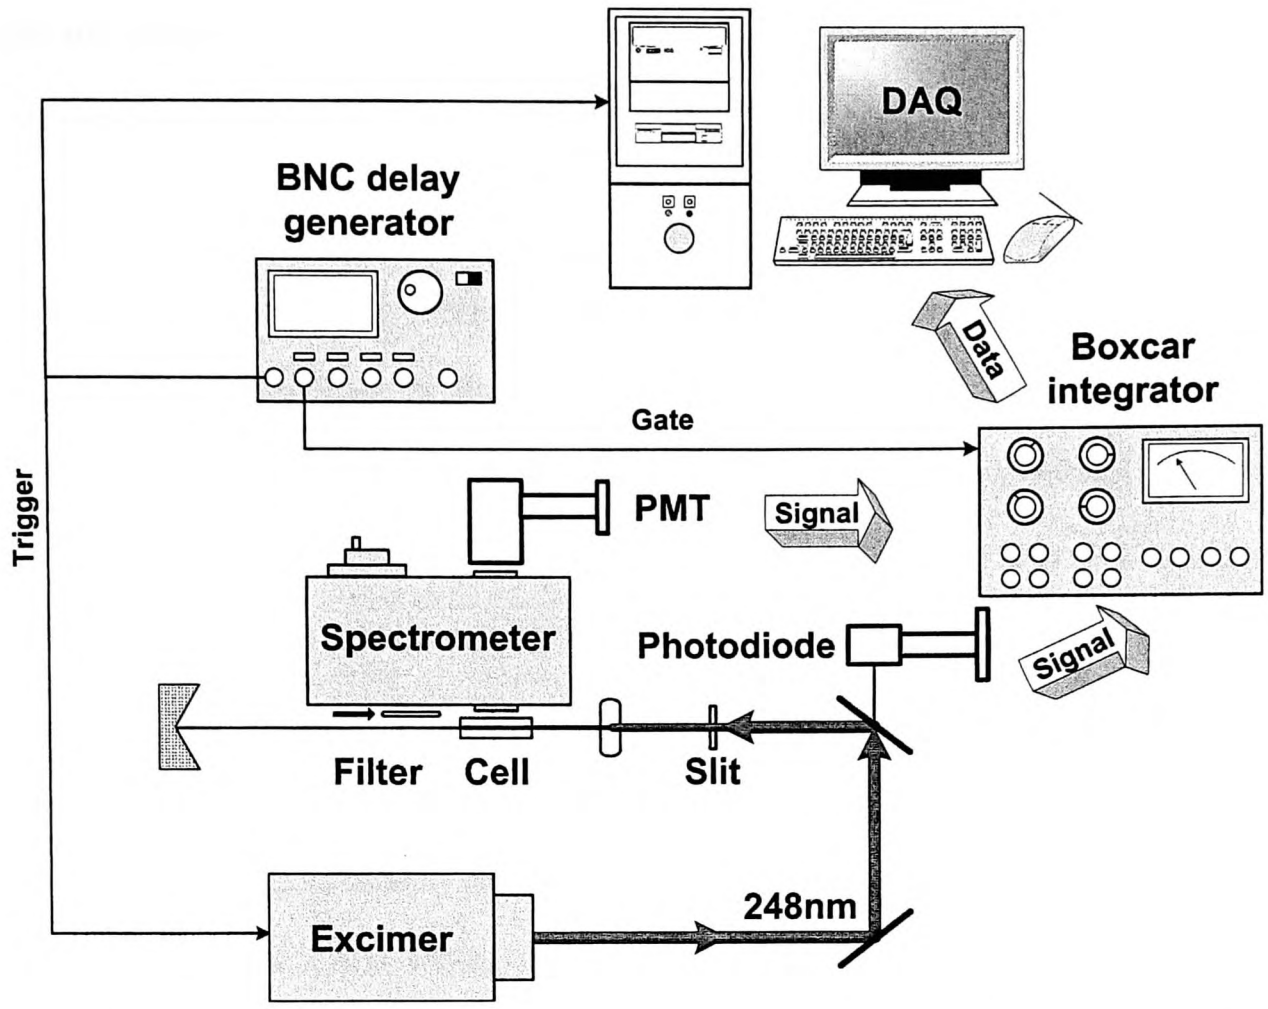


Figure 4.11: System for measuring the fluorescence spectra of fuels and tracers with a spectrometer and photomultiplier tube (PMT)

data from the photodiode were used to normalise those from the PMT. Systematic errors such as laser scatter within the spectrometer and second-order diffraction effects from the grating were observed as small local peaks on the spectra. Data points within these regions were discarded.

The first few figures in Figure 4.12 show the fluorescence and absorbance spectra of the tracers, components and component fuel, with the last being the weighted sum of the five components. Both signals are proportional to molecule number density. It can be seen that most of the fuel components gave fluorescence of similar level to acetone because the contaminants in them are more efficient at fluorescing. This is particularly true for iso-pentane and n-octane and has made the predicted fluorescence from the component fuel significant. However, as the fluorescence spectrum of the component fuel is located sufficiently away from that of acetone, optical filters can be used. The substantial LIF signal from the original iso-dodecane proved that use of the new batch was essential. The

two iso-octane samples showed similar fluorescence characteristics, despite the difference in their absorption properties. The same was true for TMB and toluene.

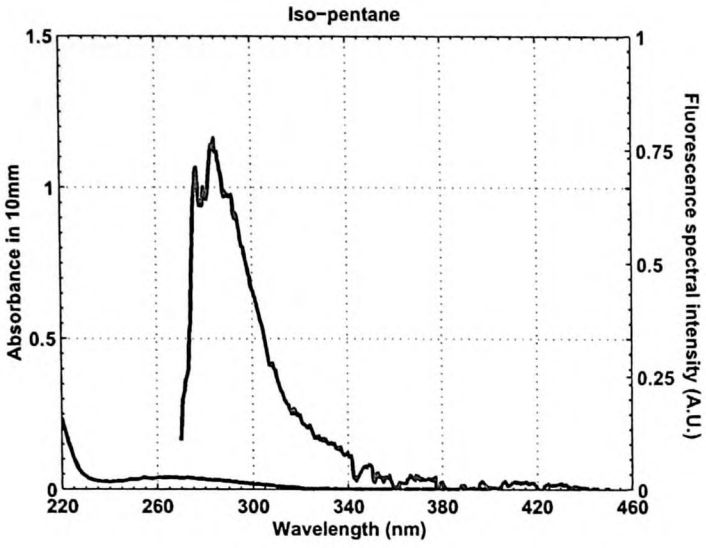
To further understand the signal-to-noise ratio based on the fluorescence from the component fuel and to be able to estimate the desirable concentrations for the tracers, the aforementioned LIF experiments were repeated with two 'long-pass' optical filters, namely WG280 and WG360 (UQG Optics, 50×50×2mm) which have nominal cut-off wavelengths of 280nm and 360nm respectively. Their responses from the spectrophotometer are plotted in Figure 4.12. The fluorescence responses from these tests are also shown in Figure 4.12. By putting one of these filters between the cell and the spectrometer, most systematic errors such as laser scatter can be removed. With the WG360 filter the fluorescence from the component fuel is mostly filtered out and the SNR for acetone can be maximised, likewise for toluene and TMB when using the WG280 filter.

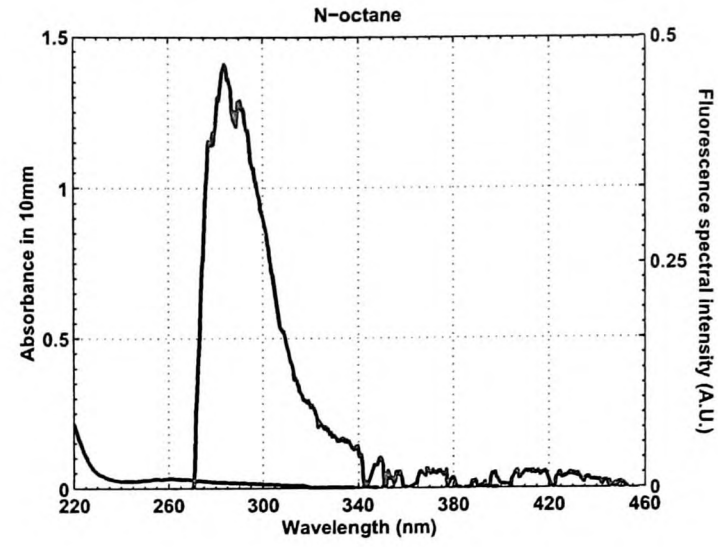
The SNR was calculated as the ratio of the spectrally integrated fluorescence spectra of the tracer and component fuel:

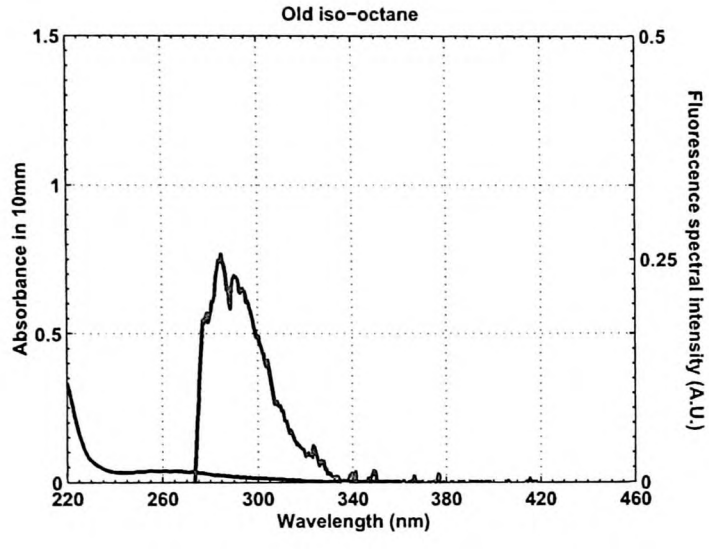
$$\mathrm{SNR} = \frac{\int_{\lambda_1}^{\lambda_2} I_{\mathrm{LIF,tracer}}(\lambda)\mathrm{d}\lambda}{\int_{\lambda_1}^{\lambda_2} I_{\mathrm{LIF,fuel}}(\lambda)\mathrm{d}\lambda} \tag{4.10}$$

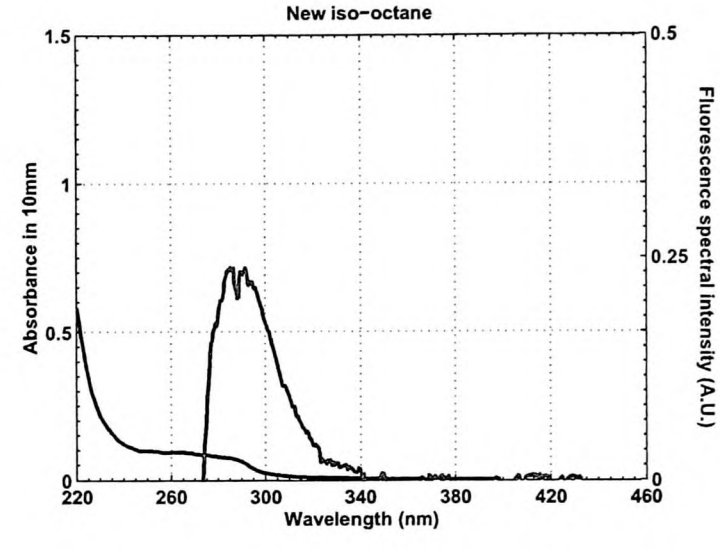
Data for other tracer volume fraction were simply derived by scaling the experimental data. Figure 4.13 presents the SNR results for fuel-tracer mixtures in both the liquid and vapour phase. The SNR limit was obtained when there was no less than 80% transmission through the vapour-phase mixture in the cylinder (89mm bore) at a particular pressure and temperature. The following assumptions were made in the vapour phase analysis:

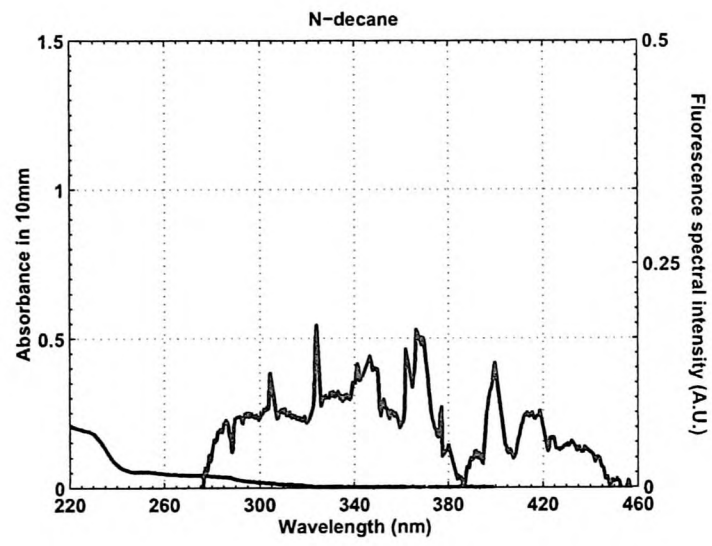
- absorption of the component fuel is neglected
- the mixture is fully vaporised *i.e.* no droplets are present in the cylinder
- the mixture vapour behaves as an ideal gas
- the absorption cross section of the liquid and vapour phase is the same at 248nm

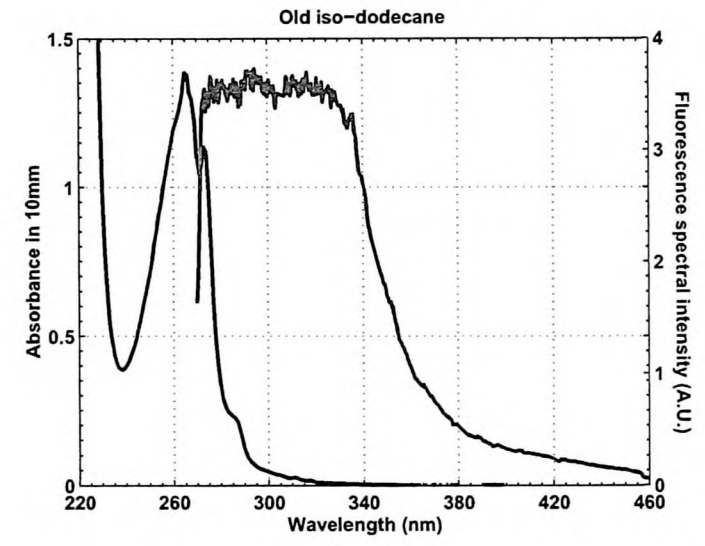

Iso-pentane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)

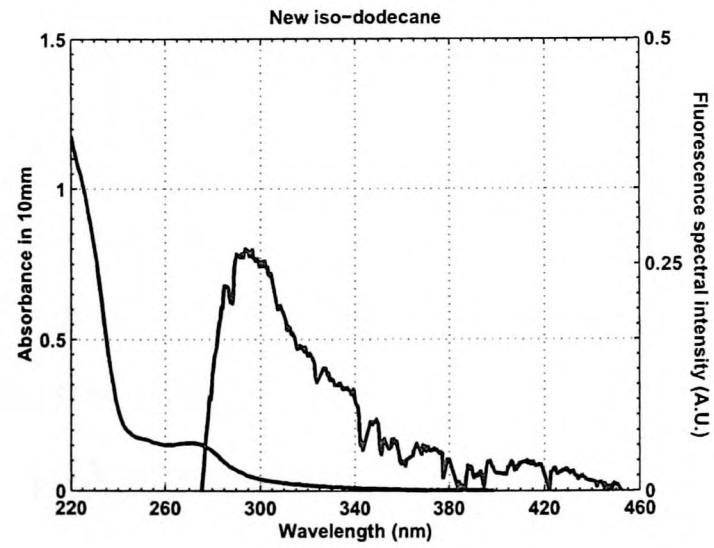

N-octane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)

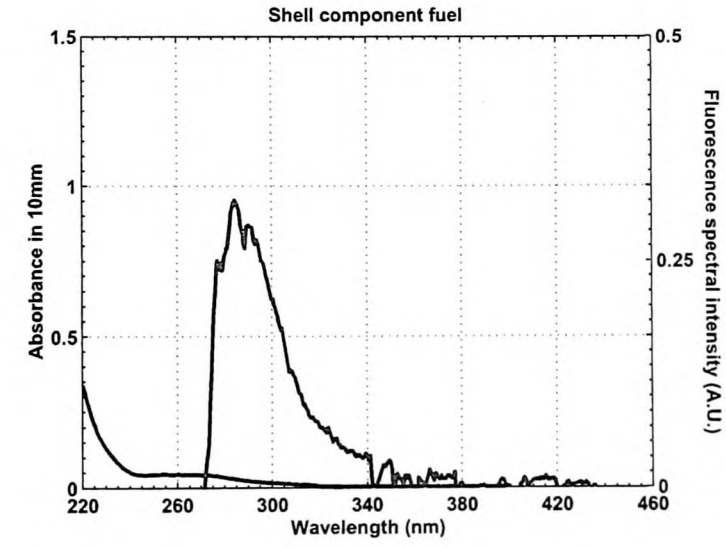

Old iso-octane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


New iso-octane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


N-decane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


Old iso-dodecane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


New iso-dodecane
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)


Shell component fuel
Absorbance in 10mm
Fluorescence spectral intensity (A.U.)
Wavelength (nm)

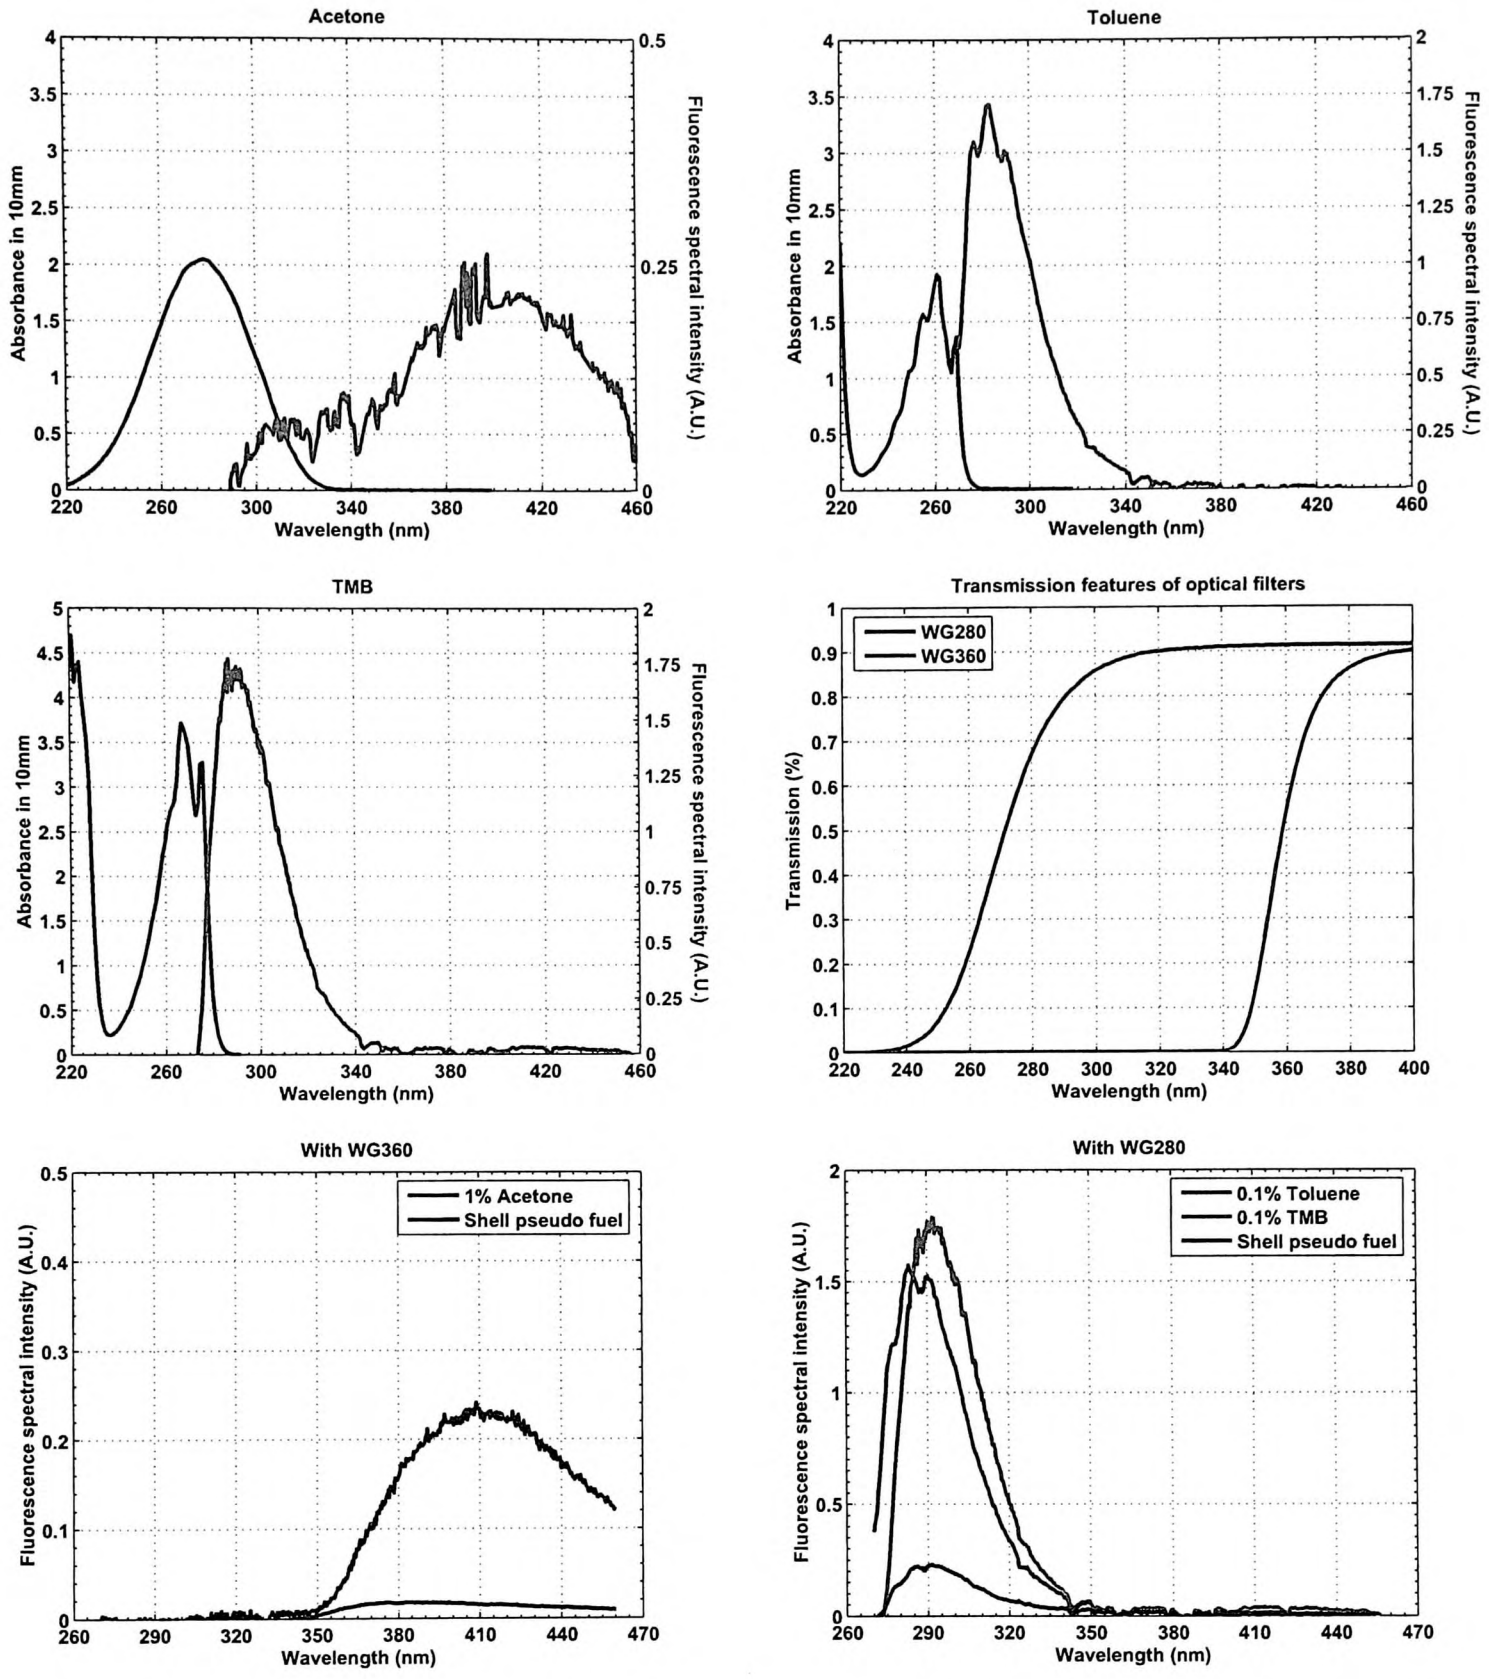


Figure 4.12: Absorbance and fluorescence spectra of tracers and component fuel, with and without optical filters

The SNR for acetone is generally better than the other two tracers although its absolute LIF signal is much weaker. This is because the WG360 filter removes most of the noise from the component fuel while the cut-off wavelength of the WG280 filter sits close to the peak of the noise. This has led to the fact that a relative noise of ~50% is to be

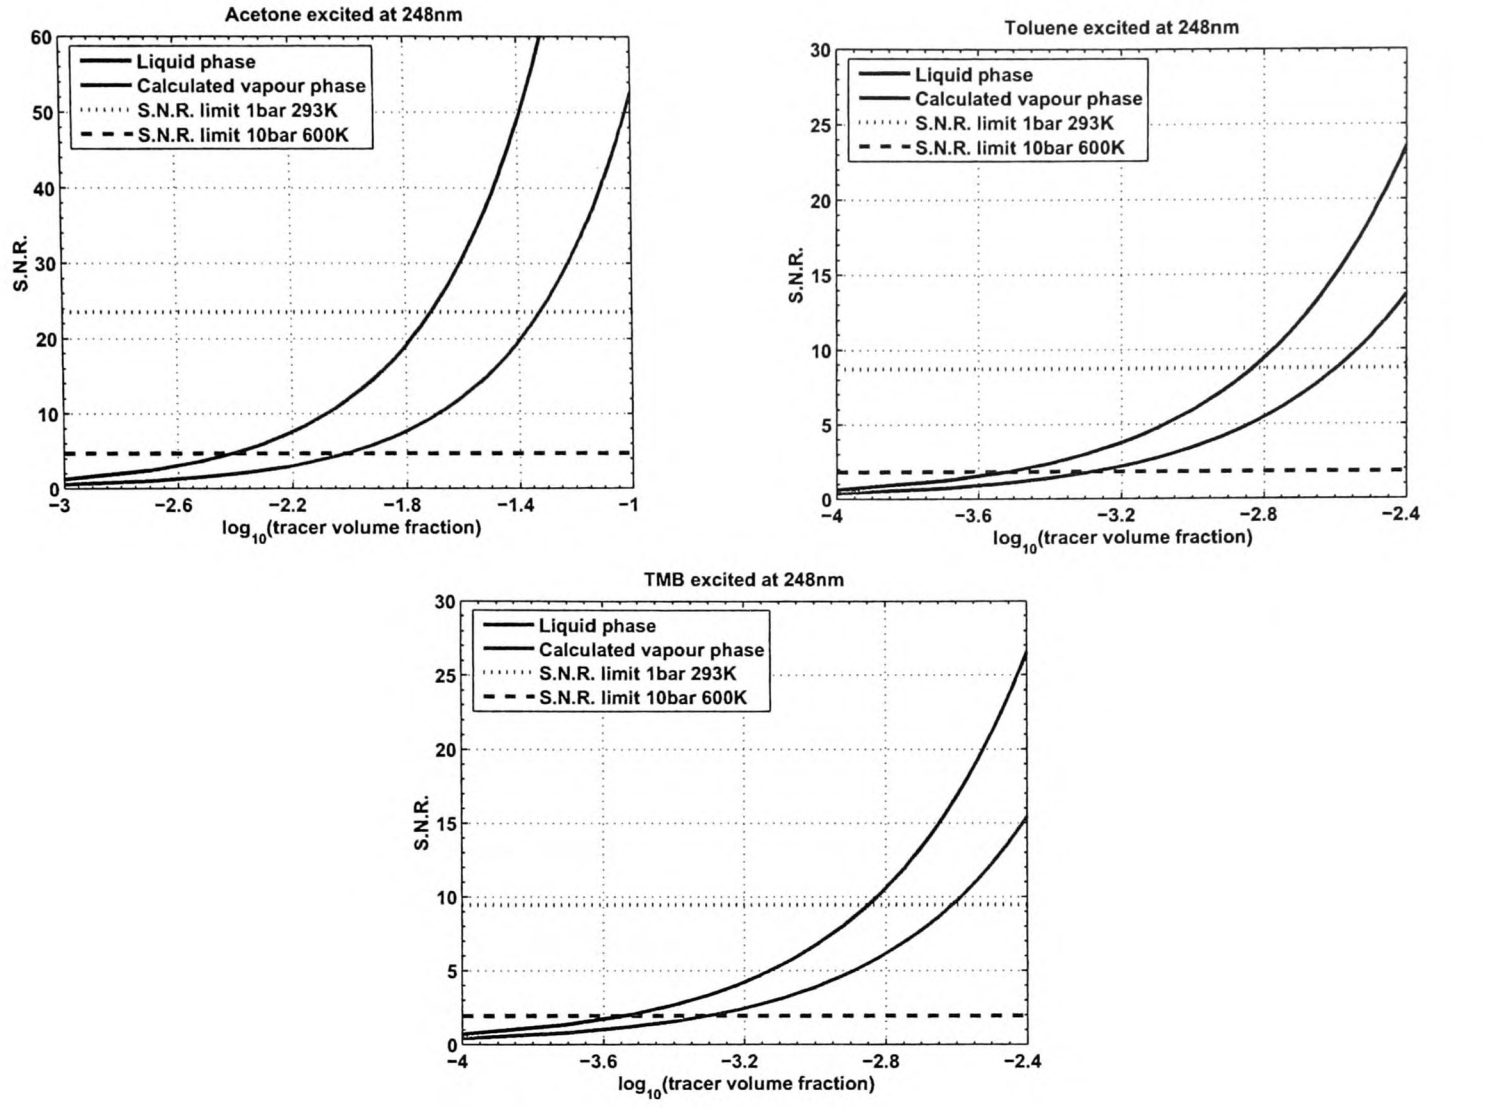


Figure 4.13: Prediction of signal-to-noise ratios for tracers

expected prior to ignition in an IC engine (10bar, 600K) when using toluene/TMB and one WG280 filter. When this is the case during the PLIF experiments, it is advisable to use two filters or even more together.

4.6 Optimisation of PLIF signal

In order to obtain quantitative results, the raw PLIF images must be post-processed for calibration of the PLIF system. This section will discuss details for noise correction and PLIF calibration. The relevant concepts of image processing have been introduced in Section 2.4.5.

4.6.1 Systematic noise correction

The first process is to remove systematic errors present in the imaging system, including:

- background noise
- built-in errors of the intensified CCD camera
- distortion effects
- Laser-induced incandescence (LII)
- phosphorescence and scattering

The background noise consists of two parts, dark noise and scattering noise. Due to the imperfect features of the intensified CCD camera system, even if the camera lens is completely covered, there will be a noise response from the CCD, *e.g.* the effect of the way in which the MCP (microchannel plate) is switched on and off, electrons drifting from the photocathode under high voltage, and deadtime noise. Deadtime noise is the charge that builds up on the CCD when there is no incident light signal and the MCP is switched off. Normally a CCD would be 'cleaned up' in this time to remove the build-up of noise, but in particular applications *e.g.* doing on-chip accumulation, such noise cannot be removed automatically by the camera. Unlike the dark noise, the background scattering noise is dependent upon the level of the gain and the setting of the intensifier gate. Therefore, to remove background noise, it is advisable to measure the background response, that is, to take reference images with constant illumination under different settings of the gate and the gain of the ICCD camera and to make subsequent corrections (background subtraction).

Apart from the imperfections mentioned above, other built-in features of the ICCD camera can also cause significant errors, such as:

- The spatial and temporal non-uniformity in the response of the CCD detectors (pixels);

- The incomplete zeroing after each pixel is moved forward to the shift register (the pixel data can only be read out when it is moved into the shift register).

Such errors can be corrected by some pixel-based transform and filtering methods (*e.g.* spatial Fourier transform).

Both the optical windows in the engine and the camera lens act as lenses and introduce optical distortion errors into the imaging. Images of a deliberately designed pattern can be taken, for instance, rectangular grids (a set of intersecting lines) for the side viewpoint (an imaging plane perpendicular to the piston crown) or concentric circles (with radial lines) for the vertical viewpoint (an imaging plane parallel to the piston crown). It should also be noted that in the vertical view case, with an optical window mounted in the piston, the piston motion can cause changes in the distortion level. This effect ought to be calibrated for as well. When such reference images are available, image registration can be done between the original pattern and the images from the camera to correct distortion effect.

Noise due to laser-induced incandescence (LII) from the fuel droplets and soot is also possible. It adds a noise signal to the fluorescence signal, as well as causing a loss of information behind the soot particles and fuel droplets because they absorb most of the laser light so that the area behind them becomes 'dark'. One feasible way to deal with LII noise is to use optical filters, because generally the wavelength spectra of LII signals and PLIF signals do not overlap.

As mentioned in Section 4.2, apart from the LIF signal there are other laser-induced light signals existing in the experimental process, namely phosphorescence and scattering. Because the lifetime of the phosphorescence signals is significantly longer than fluorescence, the laser-induced phosphorescence can cause blurring of the images (*e.g.* capturing moving objects). Laser-induced scattering processes (typically Mie scattering, Rayleigh scattering and Raman scattering), which almost happen simultaneously with the laser excitation process, are apparently potential sources of signal interference in PLIF experiments. By

gating the ICCD camera shortly after the laser pulse has ended, the scattering noise is expected to be avoided. Another feasible way of dealing with both phosphorescence and scattering noise is to place specialized optical filters in front of the image acquisition system.

Eq. (4.4) implies that the local laser intensity varies spatially as the laser gets transmitted by the media. Also present is the shot-to-shot variation in the input laser intensity. These two effects cause temporal and spatial non-uniformity and can be accounted for by using an averaged image of a homogeneous mixture of air and the base fuel to divide images from PLIF experiments.

Other systematic errors, including steering of the laser beam by density gradients, laser-produced perturbations, partial saturation and radiative trapping (absorption of the fluorescence within the combustion gases), are often avoidable or negligible, albeit needing to be considered in some particular applications (Zhao [78]).

4.6.2 Optimization of signal-to-noise ratio

This section describes the work carried out to optimise the SNR. Any leak of oil and coolant from the cylinder head and the engine block would significantly absorb the laser irradiance and hence reduce the fluorescence yield from the tracer. Furthermore, such contaminants would fluoresce and upset the PLIF signal. Therefore, great care has gone into recommissioning the engine to eliminate such leaks.

The gain and the gate setting (gate width, shape and timing) of the camera intensifier were studied and optimised. The gain was selected to give maximum available signal detection and yet avoid overloading the camera under rich conditions. The excess time the camera gate was open around the laser pulse was minimised. This was realised by interfacing with the laser triggering electronics and seeking an exact indication of the arrival of the light pulse — the signal from the thyratron in the Excimer. A thyratron is a type of gas

filled tube serving as a high energy electrical switch, basically a controlled gas rectifier. The signal from the thyratron was amplified and used to time the camera after the laser was triggered. This then minimised the effect of jitter in the laser pulse. Although careful gating of the camera helped to remove some of the background noise (ambient light, scatter, unwanted background fluorescence from contaminants *etc.*), a spectral filter that suited the acetone fluorescence spectrum was put in front of the camera.

The relation between the discharge potential on the Excimer and the laser pulse energy was investigated (Figure 4.14). Energy saturation occurred when the discharge potential went above 21kV, beyond which the discharge became less effective. The subsequent discharge potential was selected within the upper and lower bounds. The composition of the gas mixture used for the Excimer laser was also optimised to give an even discharge and improved spatial beam uniformity. An unstable resonator was designed and fitted in the Excimer laser cavity to reduce the laser beam divergence and improve the beam quality; this was also necessary for optimising the output beam from the Raman cell.

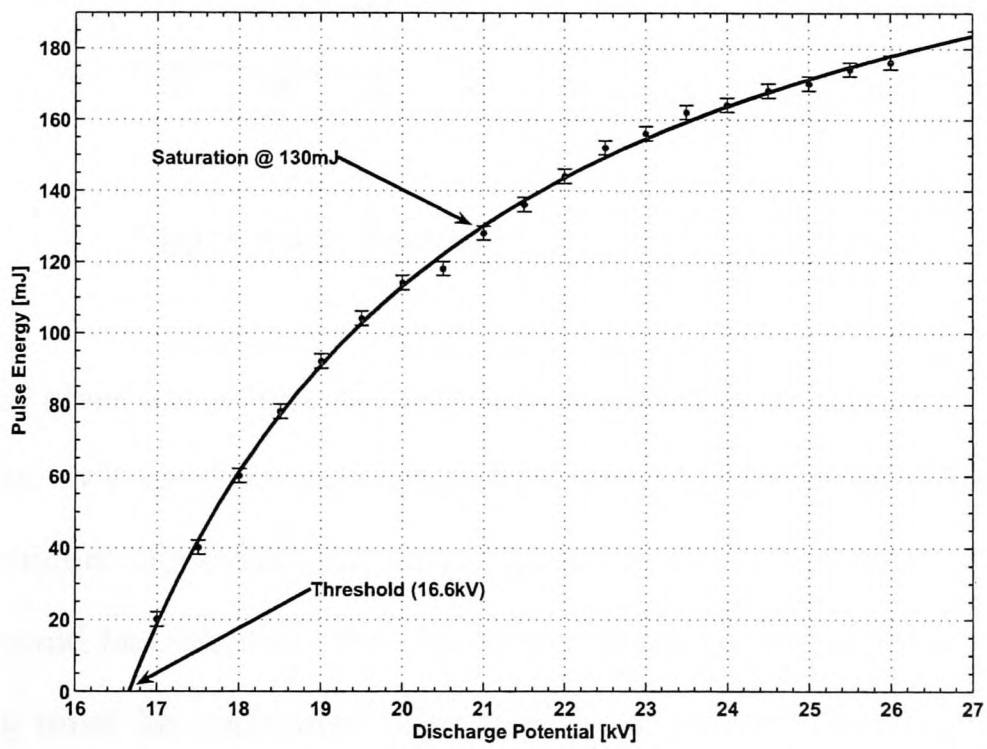


Figure 4.14: Excimer laser power curve for standard neon fill

Furthermore, acetone has a relatively low fluorescence yield, so the maximum laser in-

tensity and a good beam quality is critical to ensure a good SNR. The output from the Raman cell is proportional to the input laser energy until the input becomes so low that beam non-uniformities start to form (Figure 4.15). Hence, the selected input laser intensity should be in the linear range. When the input laser was focused with a 550mm lens, the efficiency of the Raman shifting process was estimated to be ~16% and the threshold of input laser intensity for the linear relation to hold was ~76mJ.

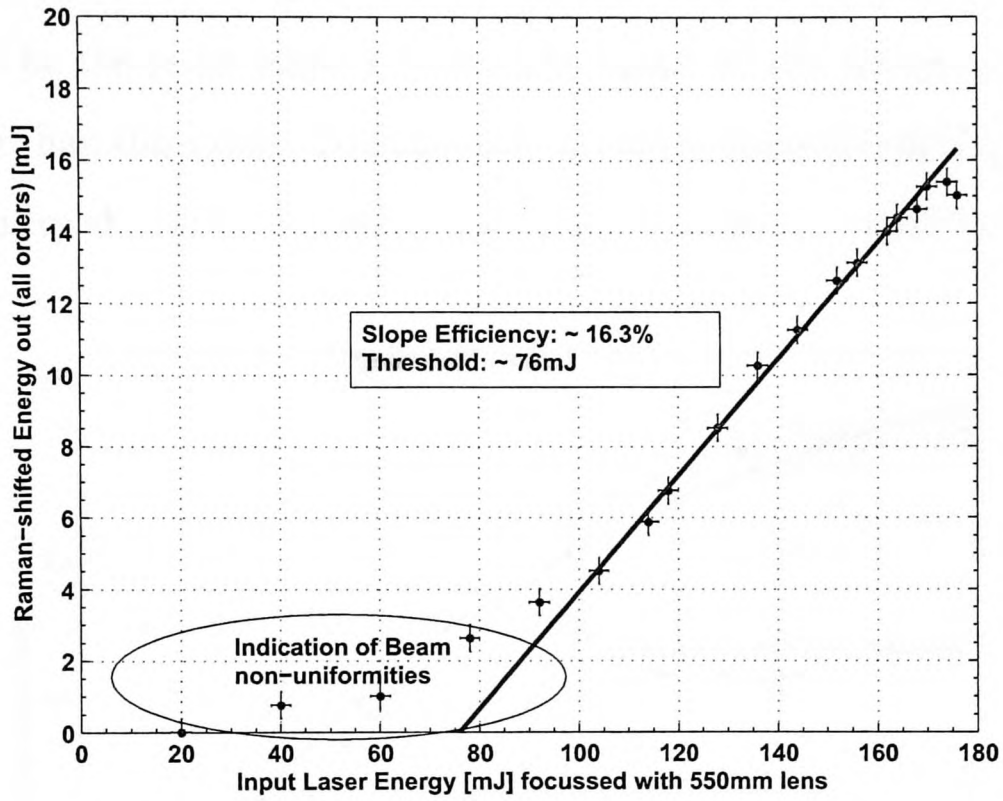


Figure 4.15: Estimation of the SRS efficiency

The SRS process gives an output of discrete wavelengths, consisting of Stokes and anti-Stokes shifted as well as the scattered fundamental. The Stokes shifting is of many orders, and the higher the order, the lower the intensity. In practice, up to five orders of Stokes shifting could be observed. For the output beam to suit experiments with acetone, effective filtering must be performed after the cell to cut out the original laser and high orders of Stokes shifting, other than the first order (277nm), from the Raman cell output. This was realised by using a prism and masks after the cell.

Before it entered the cylinder, the laser sheet was masked by a mask with a slit under the

piston window. In so doing, the thickness of the laser sheet was constrained, minimising the uncertainties due to line-of-sight effect; any unwanted scatter from the optics could also be blocked.

The realistic saturation intensities for the tracers were experimentally measured to obtain the maximum usable laser intensities for maximum signal in detecting the tracers (see Figure 4.16). The measurement looked at the fluorescence intensities from a particular tracer under different input laser intensities, then the saturation intensity could be estimated to be the point where a reasonable linear fit was always possible for input intensities lower than this point. Consequently, an input intensity slightly lower than this saturation intensity should be chosen for maximum fluorescence yield by this tracer.

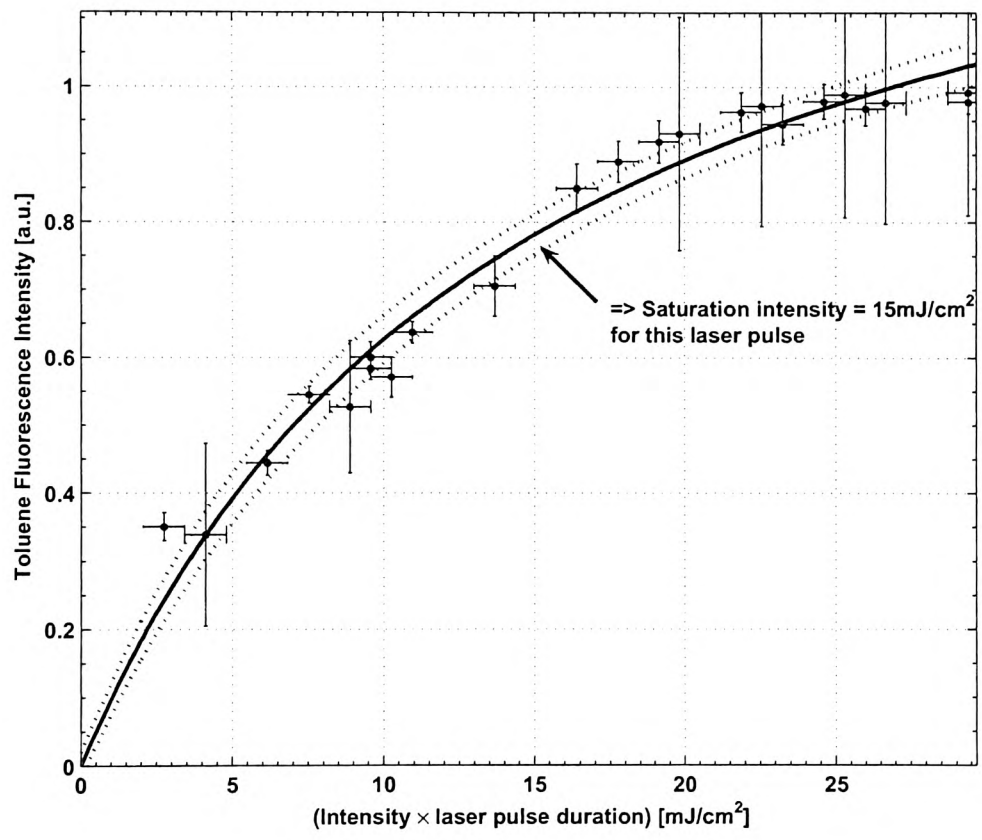


Figure 4.16: Example of estimation of the saturation intensity for toluene vapour

In the linear LIF regime, a single intense laser pulse can strengthen the PLIF signal output, but in practice its peak intensity will probably saturate the tracer absorption. So sometimes, especially in toluene PLIF measurements, the laser pulse needed to be stretched into a series of pulses all of which were below the saturation intensity for the

tracer at that particular wavelength. The schematic of this pulse stretching process and the corresponding optics setup is shown in Figure 4.17. The input beam was separated at the beam splitter, and the main beam carried on along the optical path as the first pulse in the pulse train. The diverted beam was then reflected by the four mirrors and hence became the delayed second pulse. The iterating circuit continued to give the third, fourth...pulses until the trapped energy dissipated. As a result of this setup, the exposure time of the ICCD camera needed to be prolonged to collect all the ensuing fluorescence signal.

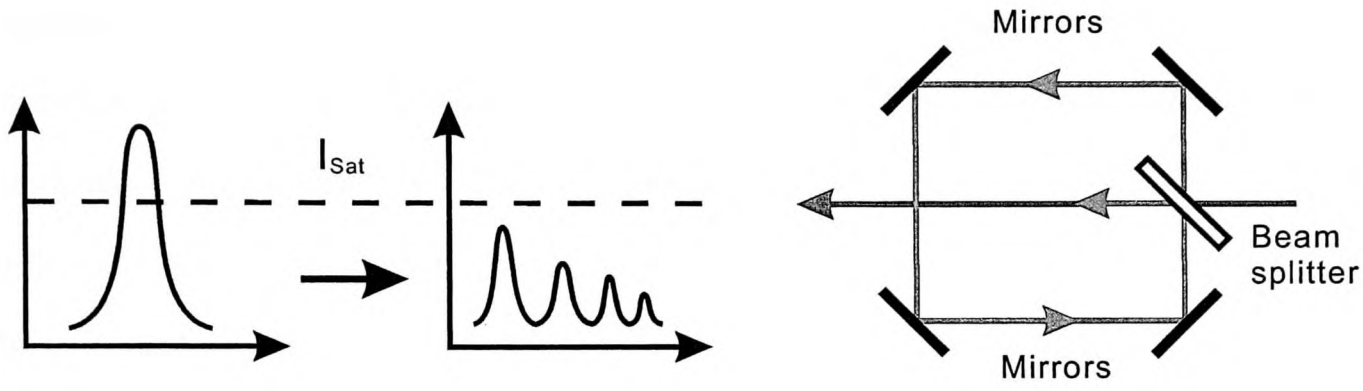


Figure 4.17: A schematic of the pulse stretching for toluene PLIF — pulse stretching (left) and its optics setup (right)

4.7 PLIF calibration

4.7.1 Calibration loop

The dependance of the PLIF signal on pressure and temperature can be calibrated in a pressure vessel where both pressure and temperature can be controlled precisely. However, difficulties will arise in reproducing the same optical and geometrical setup as that in the engine experiment. Also, it is not trivial to measure the instantaneous temperature in an IC engine. Therefore, a calibration based on a running engine is more representative and robust.

Sercey *et al.* [14] introduced a 'closed-loop' system to provide a homogeneous calibration

mixture in their DISI engine cylinder. A similar calibration loop was implemented in Oxford which directed the exhaust back to the inlet. By closely controlling the intake pressure and temperature as well as the injection duration, the fuel and air would be homogeneously mixed under known conditions after several engine cycles. The acquired PLIF images could then be calibrated against air/fuel ratio at various crank angles and hence various cylinder pressures and temperatures. The calibration images were typically taken during the compression stroke when the fuel (with one tracer corresponding to the fuel fraction under study) would have fully vaporised. The concentration of hydrocarbons was measured by a fast FID to derive the AFR in the loop. Cylinder pressure and temperature traces were measured throughout the engine cycles so a calibration 'sweep' of pressures and temperatures was effectively performed within a single test. It should be noted that given the maximum available sampling rate of the ICCD camera, images had to be captured through many engine cycles to build up a full calibration map *i.e.* database. Figure 4.18 shows a schematic of the PLIF calibration loop.

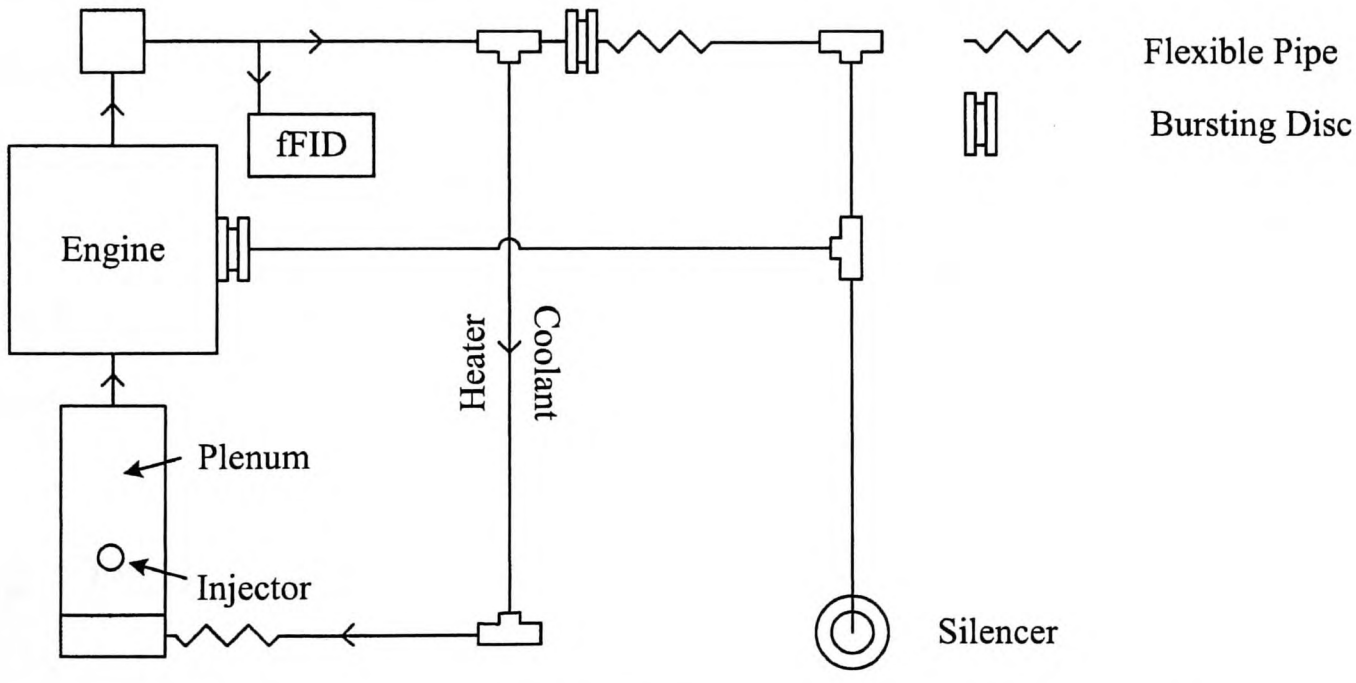


Figure 4.18: Schematic of the PLIF calibration loop (plan view)

Two bursting discs were used as one-time safety valves, with one side of them facing the pressure being monitored, and the other side connected with the exhaust extraction facilities. They would burst when the pressure being monitored had exceeded a certain

threshold. The bigger disc was put between the loop and the exhaust, and the smaller one with a higher pressure rating was installed in the aluminium annulus on top of the cylinder liner. The loop outside the engine was wrapped with a heating element and coolant pipes in order to maintain the required temperature for fuel vaporisation. The engine coolant was heated as well to above 40°C; however, the cylinder barrel was often running warmer, up to 70°C which was necessary to keep all the fuel fractions and tracers vaporised at all conditions down to $\lambda = 0.3$.

Between two measurements the mixture in the loop was purged by motoring the engine with a diaphragm pump; fresh air was drawn in from a bleed valve. The fFID was used to monitor the residual HCs during this process. Since the response time of all instrumentation was fast enough to resolve engine cycles, the blowby effect was accounted for despite the slight change in AFR over time.

The laser sheet entered the engine via the 45° mirror underneath the piston and was set to be in the vertical plane intersecting the cylinder axis. The ICCD camera viewed the fluorescence through the pentroof window. A schematic of the optical access in the pentroof region as well as a sample raw image are shown in Figure 4.19. In the sample image, black lines have been overlaid to show the field of view *i.e.* the pentroof window. The undesirable fluorescence from oil can be seen around the Inlet valve and the edge of the pentroof window.

4.7.2 PLIF vs. AFR calibration

This section presents results from the first set of PLIF calibration tests, carried out at 1500RPM, -0.48bar gauge inlet pressure with acetone as the tracer. The loop temperature and barrel temperature were 50°C and 60°C respectively. A volume fraction of 10% was used for acetone. Too low a tracer concentration would give a poor SNR while on the other hand, too high a concentration would lead to heavy absorption along the laser sheet

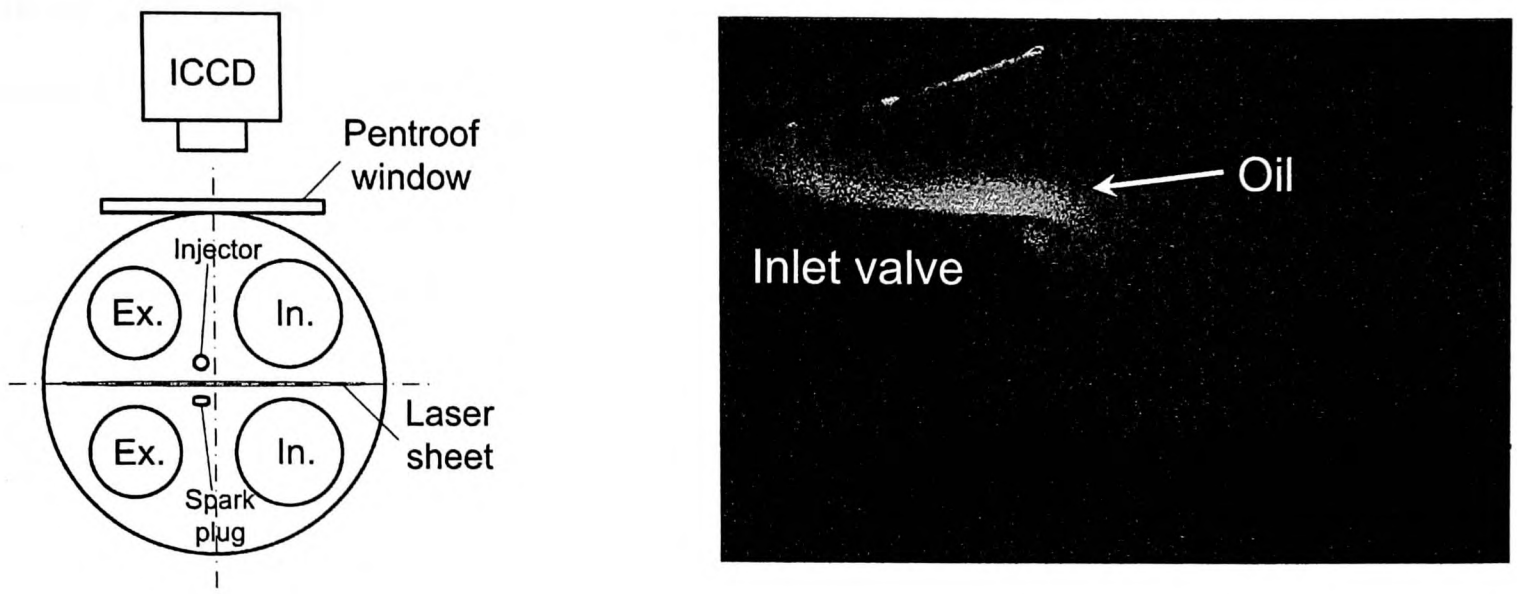


Figure 4.19: Schematic of the optical access in the pentroof region from a plan view (left) and a sample raw image (right)

and spatial non-uniformity in the PLIF images.

The ETCS program sent a signal to trigger the laser, while the ICCD camera was timed by an amplified signal from the thyratron after the laser was triggered. With the camera set to a reduced frame size, 576 by 191 pixels, its sampling rate could match the laser repetition rate of 3.125Hz, that is, skip pulsing once every four engine cycles. The laser energy was recorded by an energy meter to account for the temporal fluctuation in the laser pulse.

Figure 4.20 shows a sample image of the background light (with laser but no fuel) and a PLIF image after the background subtraction, with a lambda of 1.0 at TDC. The latter clearly outlined the laser sheet.

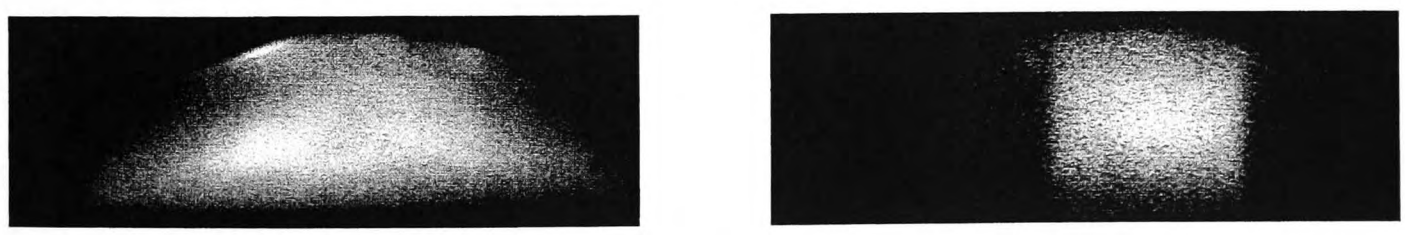


Figure 4.20: Sample images of background (left) and PLIF signal after background subtraction (right), $\lambda = 1.0$ with acetone @ TDC

The camera response to fluorescence yield was calibrated against crank position and hence

the cylinder pressure and temperature at that particular crank angle (Figure 4.21). The camera counts were normalised after taking into account the chosen camera gain and exposure settings. It can be seen that with a lambda of 1.5, the signal level was low and in some cases below zero; the dependance of the signal on p and T was difficult to extract. This was due to overwhelming background noise and a poor SNR. The standard deviation was the least for stoichiometric operation.

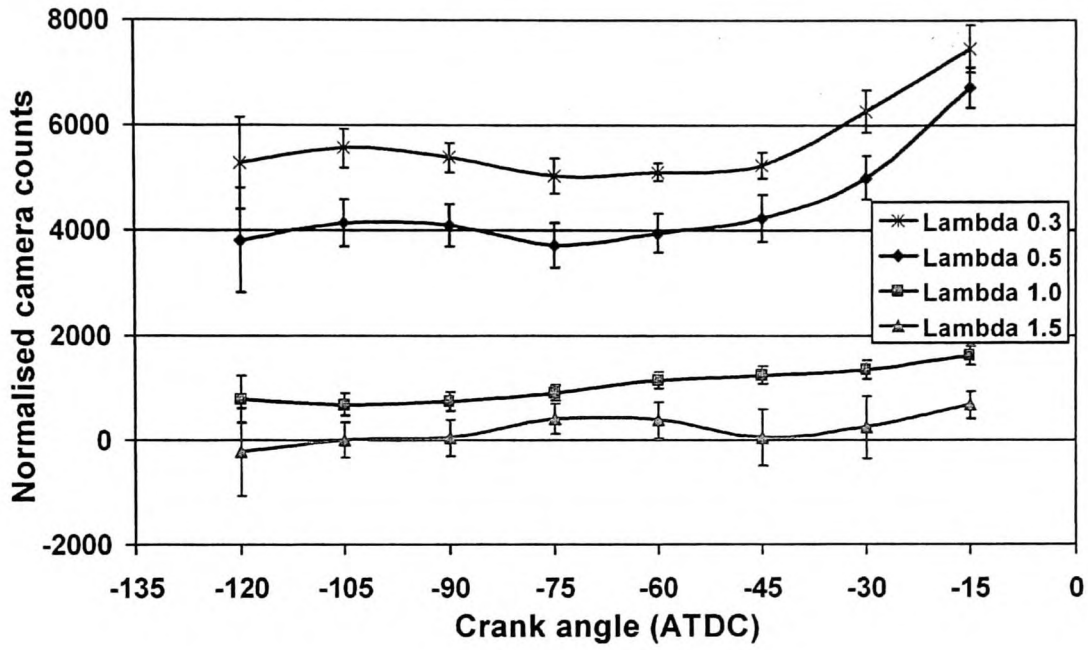


Figure 4.21: PLIF calibration results against crank angle with different lambda

Figure 4.22 shows the calibration results against lambda at a series of crank angles. Again, the SNR was very poor with $\lambda = 1.5$ and above. For stoichiometric and rich operation, the relation between camera response and mixture strength was reasonably linear.

At present, the engine is being overhauled and measures to improve the SNR are being implemented. More PLIF calibration results with other tracers as well as AFR measurements from normal firing engine tests are to follow in the next stage of the ongoing project.

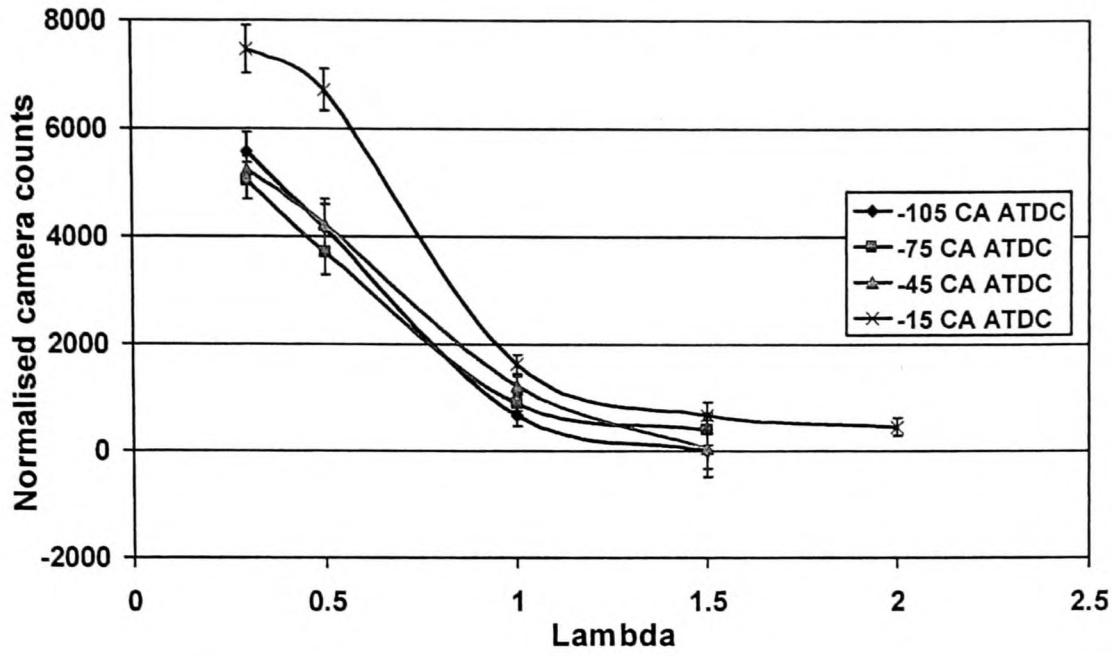


Figure 4.22: PLIF calibration results against mixture strength at different crank angles

4.8 Summary

The underlying theory of quantitative PLIF measurements in an IC engine has been summarised in depth, together with a review of others' relevant work to this date. The methodology for selecting a fuel/tracer set for tracer-PLIF has been discussed. A dedicated six-component three-fraction fuel was designed, with three corresponding tracers to match the light, medium and heavy fraction of the fuel respectively. PLIF images with this fuel/tracer set can reveal quantitatively the distribution of different fuel fractions in an engine. The absorption and fluorescence characteristics of each fuel component and tracer were studied to give an estimate of the bounds of signal-to-noise ratio.

Great care was taken to optimise both hardware and the SNR, including spectral filtering, time gating of the camera, laser pulse stretching for toluene PLIF, study of stimulated Raman scattering for acetone PLIF and absorption saturation *etc.*

A recirculating loop was constructed and set up for the calibration of PLIF signal against AFR. This experimental layout has the advantages of *in situ* optics and a real engine environment. Example calibration results with acetone as the tracer were shown, suggesting

that maximising SNR was the key to a practical calibration for quantitative PLIF.

Chapter 5

Soot Imaging and Pyrometry

There are two kinds of light — the glow that illuminates, and the glare that obscures.

— James Thurber (1894–1961)

5.1 Introduction

CCD imaging and colour pyrometry are among the most commonly used techniques for characterizing flames and combustion [69] [77]. Colour pyrometry relates the temperature of glowing soot and the intensity of the thermal radiation it emits. In conventional two-colour pyrometry, for example [3], the light sensing device is calibrated against a black body to determine the apparent temperatures at two specific wavelengths.

The flame temperature can be determined from the ratio of the light intensities at two wavelengths by eliminating an unknown coefficient in the extended version of Planck's radiation law. Substitution of the flame temperature into equations involving the absolute intensities yields an estimate of the KL factor, which is proportional to the soot concentration. However, the necessity to have an absolute calibration of the light-sensing

apparatus for the two colour method makes the result sensitive to factors that alter the visible intensity, such as deposits on optical surfaces or wall reflections/glare.

In contrast, the colour-ratio pyrometry (CRP) technique [57] [73] uses the ratios of three colours to give two independent measures for determining the temperature and KL. The conventional approach normally uses narrow band filters and two light sensors/cameras with beam splitting optics, but it is also possible to use a single colour CCD camera [63] to avoid complexities such as the optical setup and image registration.

The advantage of CRP, that absolute intensities are not required, means that it is insensitive to the multitude of problems associated with quantitative intensity measurement. Such problems include calibration errors, variations in optical path and window fouling. Ratio measurements mean that the technique is more rapidly and robustly applicable to engine environments. However, sufficient equipment dynamic range is needed for good accuracy.

The work reported here employs a colour CCD camera with three channels (RGB) and uses the ratios of the pixel intensities. The three channels can be reduced to two ratios between red and green (R/G), and red and blue (R/B). In so doing the absolute intensity response is not used and an absolute temperature calibration is not needed. Instead, look-up tables of R/G and R/B ratios are created for the appropriate range of temperature and KL values.

Before describing the current work in detail, it will be useful to review some of the practical problems associated with pyrometry. The ideal form of calibration is with a black body, but it is difficult to operate such objects at combustion temperatures. The usual approach is to use a tungsten ribbon lamp, but the emissivity of the filament is a function of wavelength, temperature, age, surface finish and temperature-history of the lamp; consequently there are wide variations in the published data [68]. This issue is side-stepped as it is common for the tungsten ribbon lamp only to have a calibration in

terms of an apparent (equivalent black body) temperature at a particular wavelength. This can be one of the wavelengths used in traditional two-colour pyrometry (if it does not coincide with a wavelength that is characteristic of chemiluminescence), but there is still the issue of obtaining an absolute calibration at a second wavelength. If absolute measurements are being used, then strictly speaking an *in situ* calibration is needed. One approach is to install a calibrated source in the engine, or a tungsten ribbon lamp can be placed in an equivalent optical path [10].

It must also be remembered that the measurement is not from a single plane, but it is an integration along the optical path. Murphy and Shaddix [44] conclude that the stratification of temperature along an optical path biases temperature measurements towards the higher temperatures along the path, and that the concentration of the hottest soot will bias the KL value. This work suggests that when there is a variation in temperature of 600K with a mean temperature of 1500K, then errors in the measured temperature can readily exceed 100K. There is also the issue of how to model the emissivity of the soot, and this is discussed in detail by di Stasio and Massoli [15]. However, it would appear that this source of error is much less significant than the effect of variations in temperature or KL along the optical path being measured.

With all of these possible sources of error, then why is soot pyrometry a popular technique? It is comparatively cheap and easy to implement, and can be used to give a 2-D plot of the temperature and KL. This is very useful for comparisons between, for example, different fuels or different injection systems.

Colour-ratio Pyrometry has not been extensively applied in the literature; therefore this chapter describes a generic approach to calibrating a CCD camera for use with CRP. The validation of the calibration method and the temperature prediction with bench tests is also given here. Finally, the colour-ratio technique has been applied to the Jaguar optical access DISI engine with a spray-guided combustion system. Results from measurements over a range of engine operating conditions are reported.

5.2 Principles of soot pyrometry

5.2.1 Radiation of glowing soot

Analysing flame images relies on the relationship between the temperature of a glowing particle and the intensity of the thermal radiation it emits. The spectral intensity emitted by a black body, $I_{\lambda b}$, is a function of wavelength and temperature alone, and is given by Planck's Radiation law:

$$I_{\lambda b} = \frac{2C_1}{\lambda^5 \left[\exp\left(C_2/\lambda T\right) - 1\right]} \tag{5.1}$$

where C_1 and C_2 are known as Planck's radiation constants, having values of 1.871×10^{-16} W$\cdot$ m^2 and 1.439×10^{-2} m$\cdot$K respectively.

Real objects such as soot particles are not perfect emitters, and radiate less energy at each wavelength than the equivalent black body value. The spectral emissivity ε_λ is defined as the ratio of the actual intensity emitted by a body at a given wavelength to the black body intensity, so that

$$I_\lambda = \varepsilon_\lambda I_{\lambda b}(\lambda, T) = \varepsilon_\lambda \frac{2C_1}{\lambda^5 \left[\exp\left(C_2/\lambda T\right) - 1\right]} \tag{5.2}$$

The emissivity of the flame has to be estimated and this can be done by the widely used empirical correlation established by Hottel and Broughton [28]

$$\varepsilon_\lambda = 1 - \exp\left(-KL/\lambda^\alpha\right) \tag{5.3}$$

where K is the absorption coefficient per unit flame thickness, and L is the geometric flame thickness along the light axis of the flame detection system (optical path length). The absorption coefficient is a function of the soot particle concentration in the flame. The value of α depends on the physical and optical properties of the soot in the flame and, at visible wavelengths, was found by Hottel and Broughton [28] to have a value of

1.39. Typical temperature values for glowing soot range from about 1900K to 2600K, with the KL factor varying between 0.2 and 1.4 (Wyszynski [74]).

5.2.2 Colour-ratio pyrometry

The most commonly used technique for measuring flame temperatures is two-colour optical pyrometry, which was originally developed by Hottel and Broughton [28]. It is necessary to calibrate the light-sensing apparatus for the two colour method so that the signals can be related to an apparent flame temperature. This can be done using a source of black body radiation, to produce calibration curves of output against black body temperature for each wavelength.

Whereas the two-colour method uses a temperature-calibrated optical system at two specific wavelengths or wavelength bands, a colour camera has the three colour (RGB) filters with overlapping wide spectral responses (see Figure 5.16 in Section 5.3.1). The three channels can be reduced to two ratios between red and green (R/G), and red and blue (R/B). In this way the absolute response is not used and absolute temperature calibration is not needed. Using equations (5.2) and (5.3), and the camera spectral response from Figure 5.16, look-up tables of R/G and R/B ratios can be calculated for typical temperature and KL values.

The sensitivity of the system could be enhanced by using a filter that would attenuate the longer wavelengths, so that at a mid-range temperature each of the three colour channels would have the same signal level.

Soot pyrometry estimates directly the temperature of the soot, but due to the almost instantaneous attainment of thermal equilibrium (in 10^{-6}–10^{-5} s, Schanzlin *et al.* [52]), it can be assumed that the measurement of the soot temperature also represents the burnt gas temperature, *i.e.* the temperature of the surrounding gas.

Two strategies for colour-ratio pyrometry can be employed, namely interpolation with a

fixed value of KL and mapping with a variable KL. The schematics of these procedures are shown in Figure 5.1.

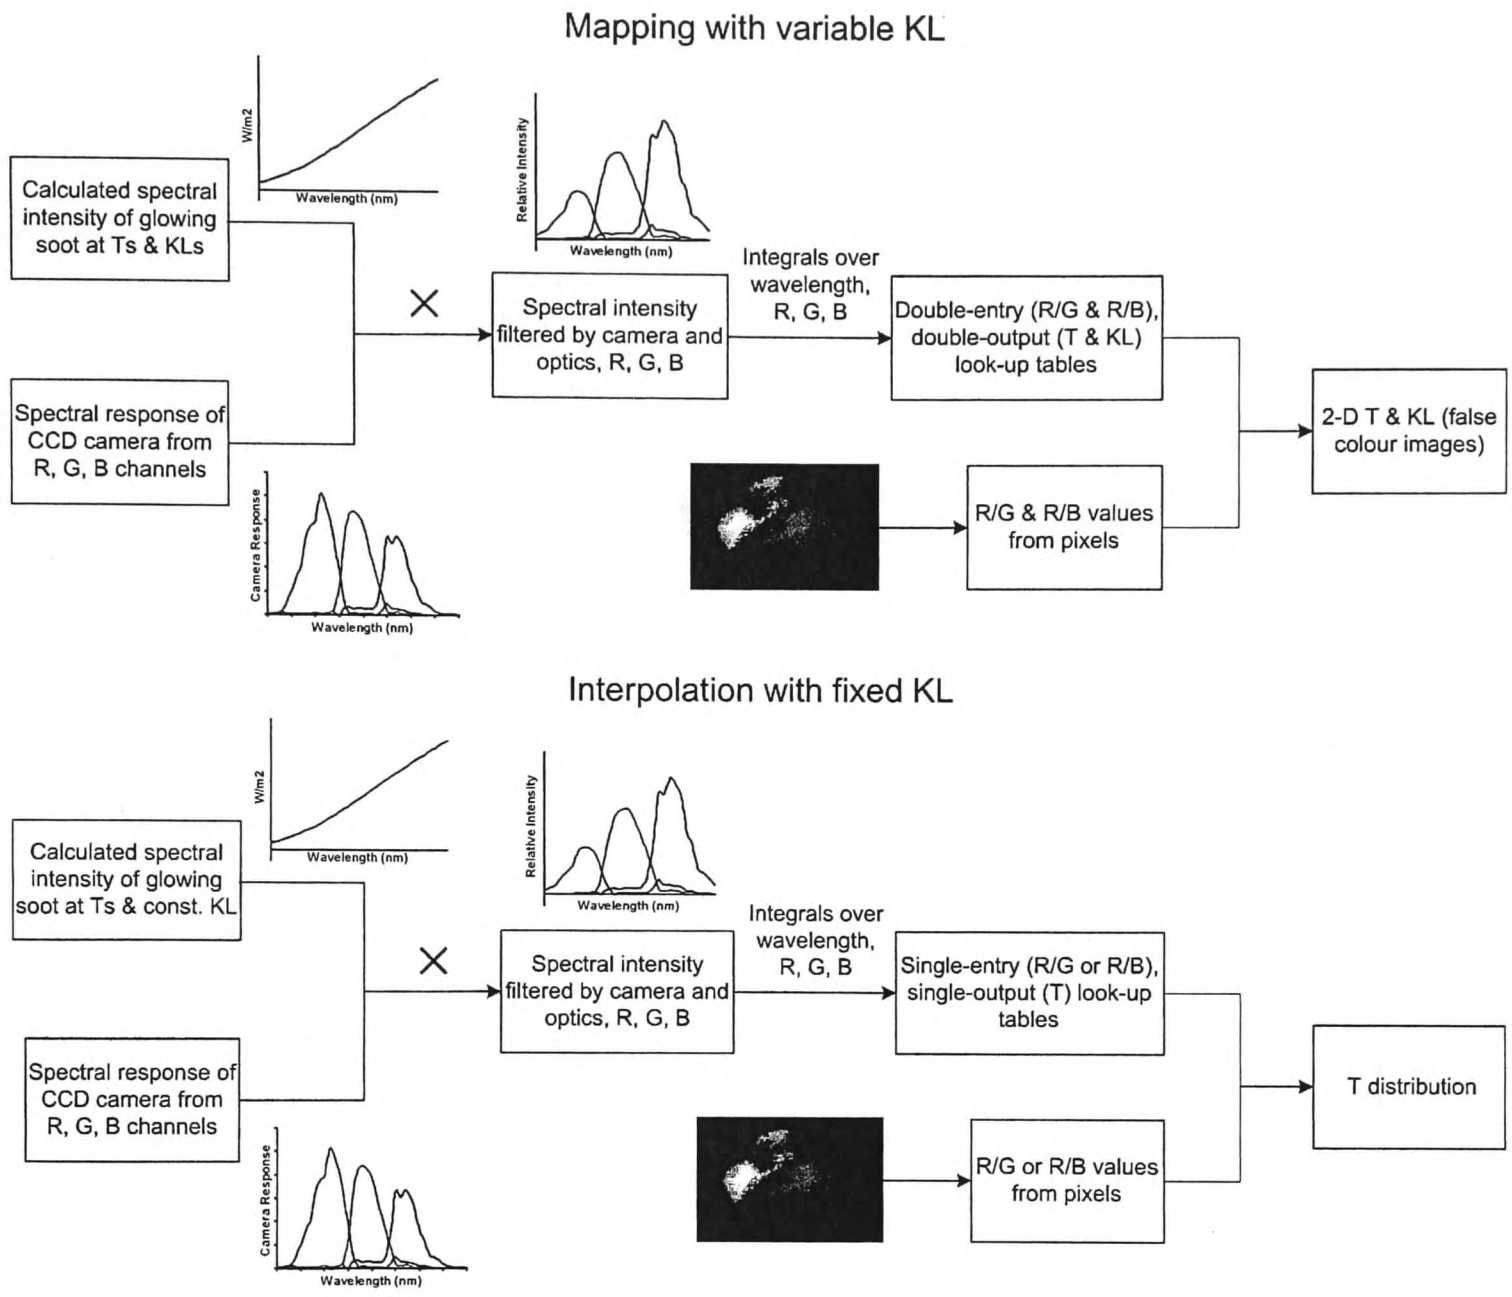


Figure 5.1: Procedures of the two strategies for colour-ratio pyrometry

By means of integrals over the full wavelength range, the colour-ratio technique is less sensitive to the wavelength selections and emissivity variation usually faced in classical two-colour pyrometry. In addition, because the solutions are derived from pre-created look-up tables via interpolation rather than iteratively solving non-linear equations, less computational time is required and the processing is significantly simplified.

When the mapping with a variable KL strategy is applied, both the R/G and R/B ratio values from the image are used and the mapping for temperature and KL is performed

to minimise the following characteristic function:

$$E = \left[(R/G)_{T,KL} - \alpha_{i,j}\right]^2 + \left[(R/B)_{T,KL} - \beta_{i,j}\right]^2 \tag{5.4}$$

where $(R/G)_{T,KL}$, $(R/B)_{T,KL}$ are the inputs from a particular pixel and $\alpha_{i,j}$, $\beta_{i,j}$ are the entries in the look-up tables. Figure 5.2 shows the surface plots of the colour ratios in the two look-up tables created with the 2-D mapping strategy.

Despite the potential of spatially resolving both temperature and KL, previous work [56] [63] has shown that the KL solution from soot pyrometry is often ill-conditioned in some cases, because:

- for a given pair of colour ratios there can be multiple solutions of temperature and KL which all map closely to the measured colour ratios;
- Uncertainties due to noise in the blue channel of a conventional CCD are high compared with the precision required to determine KL accurately;
- KL tends to be dominated by the environment closest to the detector.

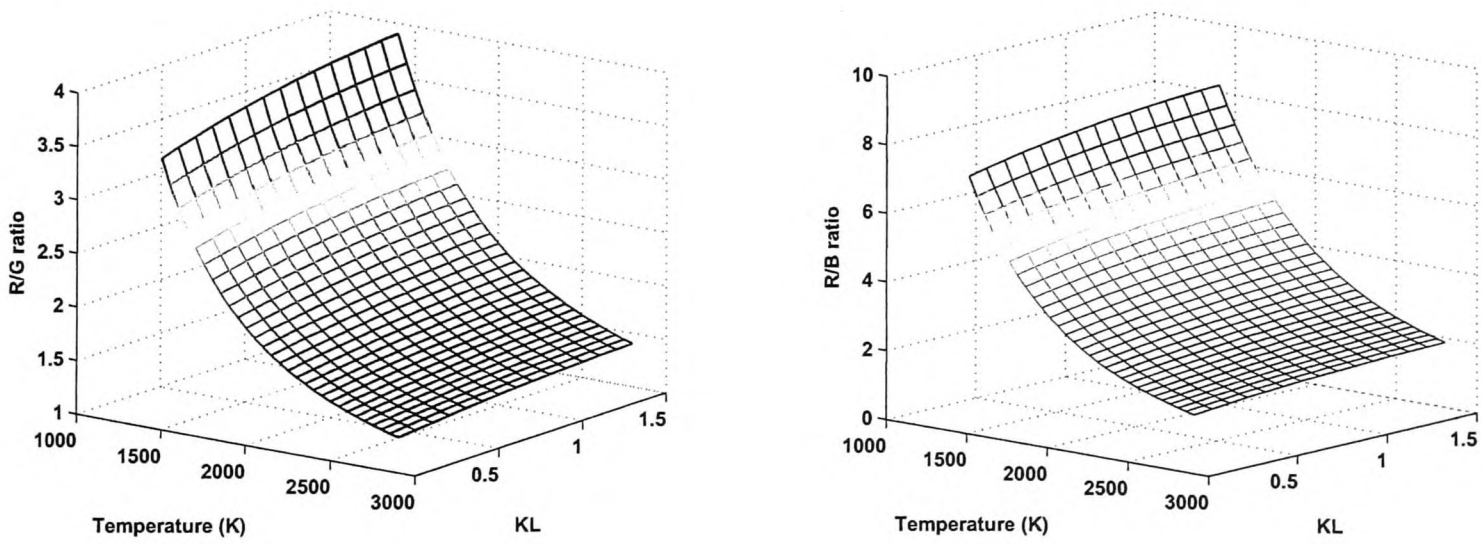


Figure 5.2: Look-up tables calculated with the mapping with a variable KL strategy for the temperature range 1000–3000K in steps of 10K, and for the KL range 0.1–1.5 in steps of 0.1

The variation of emissivity with wavelength at different KLs [28] is shown in Figure 5.3. Assuming a constant KL eliminates the complication of accounting for soot emissivity

at the expense of not being able to detect soot concentration. Figure 5.4 shows the R/G and R/B ratios in the look-up tables created with the interpolation method, assuming a constant KL. The temperature of each pixel can then be derived by interpolation of the R/G or R/B ratio.

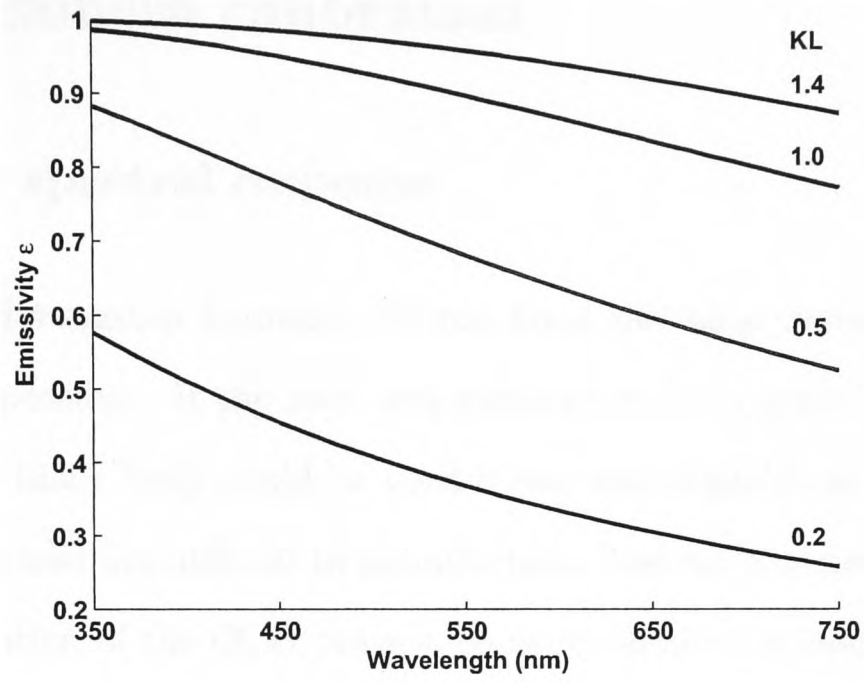


Figure 5.3: Variation of soot emissivity with wavelength and KL [28]

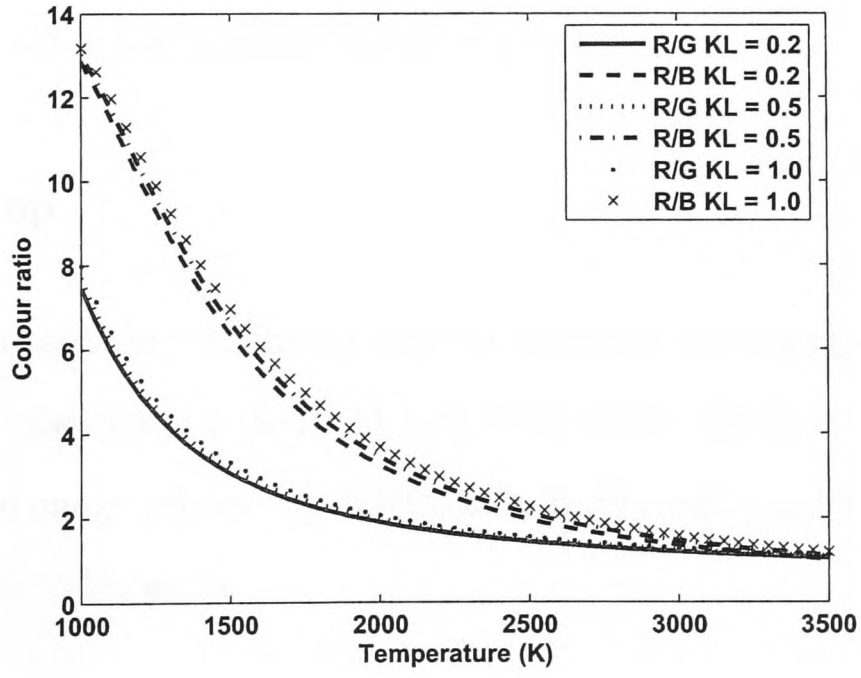


Figure 5.4: Expected R/G and R/B ratios of CCD camera response with different KL values

It is worth noting that the blue green (B/G) ratio is not mathematically independent of

the R/G and R/B ratios. Additionally, because of decreasing light intensity towards the blue end of the emission spectrum, the B/G is the most susceptible of the three ratios to noise and digitization error. For these reasons the B/G ratio has not been used.

5.3 CCD camera calibration

5.3.1 Colour spectral response

Since the soot and tungsten filaments do not have the same emissivity, then a direct calibration is not possible. If the soot was assumed to be a grey body (with constant emissivity), then a black body could be used if one was available at a high enough temperature, however these are difficult to manufacture. Instead it is necessary to determine the spectral calibration of the CCD camera, so as to obtain the camera's colour spectral response. This response (which is seen later in Figure 5.16) is a measure of the camera's sensitivity to different wavelengths for the three colour channels (red, green and blue), and is needed for the CRP. The technical specifications for the Pulnix TMC-9700 video camera used in this work can be found in Section 2.4.3.

Experimental setup

The calibration was achieved by shining light at different wavelengths onto the camera's CCD sensors. Then the digital outputs of each of the three channels were extracted from the images using the image processing functions in MATLAB. Figure 5.5 shows the setup that was used for the calibration.

Light from a tungsten lamp (13.5V 5A) was led through the spectrometer (HILGER & WATTS). It was focused onto grating g (1200 lines/mm, blazed @300nm) and then focused onto the camera CCD chip (no objective lens) or photodiode. Both slits s_1 and

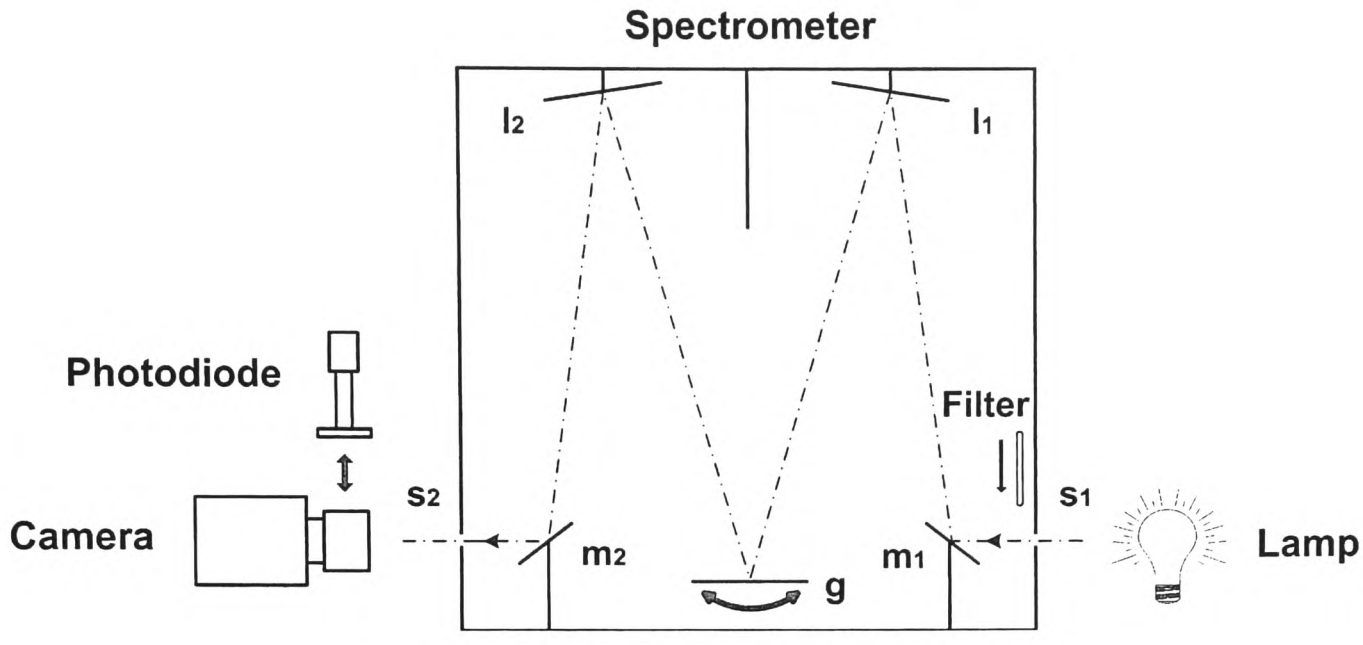


Figure 5.5: CCD camera calibration setup

s_2 were set to 2mm, giving an output bandwidth of 8nm. The output wavelength was varied by rotating the grating around a vertical axis (perpendicular to this sheet). To calibrate the emission spectrum of the lamp and the efficiency of the spectrometer, a photodiode of known spectral sensitivity (ThorLabs DET210) was used to measure the relative intensities when extracting the spectral response curves.

To avoid saturation of images at certain wavelengths yet make full use of the dynamic range of the camera, different shutter speeds were used.

The spectrometer could potentially introduce the following sources of error:

- Because the bulb emitted much less in the blue range than red, light scatter of the red constituents could bias the measurements in the blue range;
- The second order diffraction effect of the grating could cause some blue radiation (300–400nm) to appear in the signal in the red range (600–800nm).

To attenuate these sources of noise, four filters were selected and put behind slit s_1 according to the wavelength range being measured. When the Blue filter was used, most of the radiation above 500nm would have been heavily attenuated, likewise for the LEE 026 filter, hence eliminating the cross talk between wavelength ranges. The transmission char-

acteristics of these filters were obtained from a spectrophotometer (Perkin-Elmer Lambda 9, usable range 185–3200nm) and are shown in Figure 5.6. Also presented in Figure 5.6 are three other filters used when testing the camera's response linearity with sensitivity and exposure (detailed in Section 5.3.2).

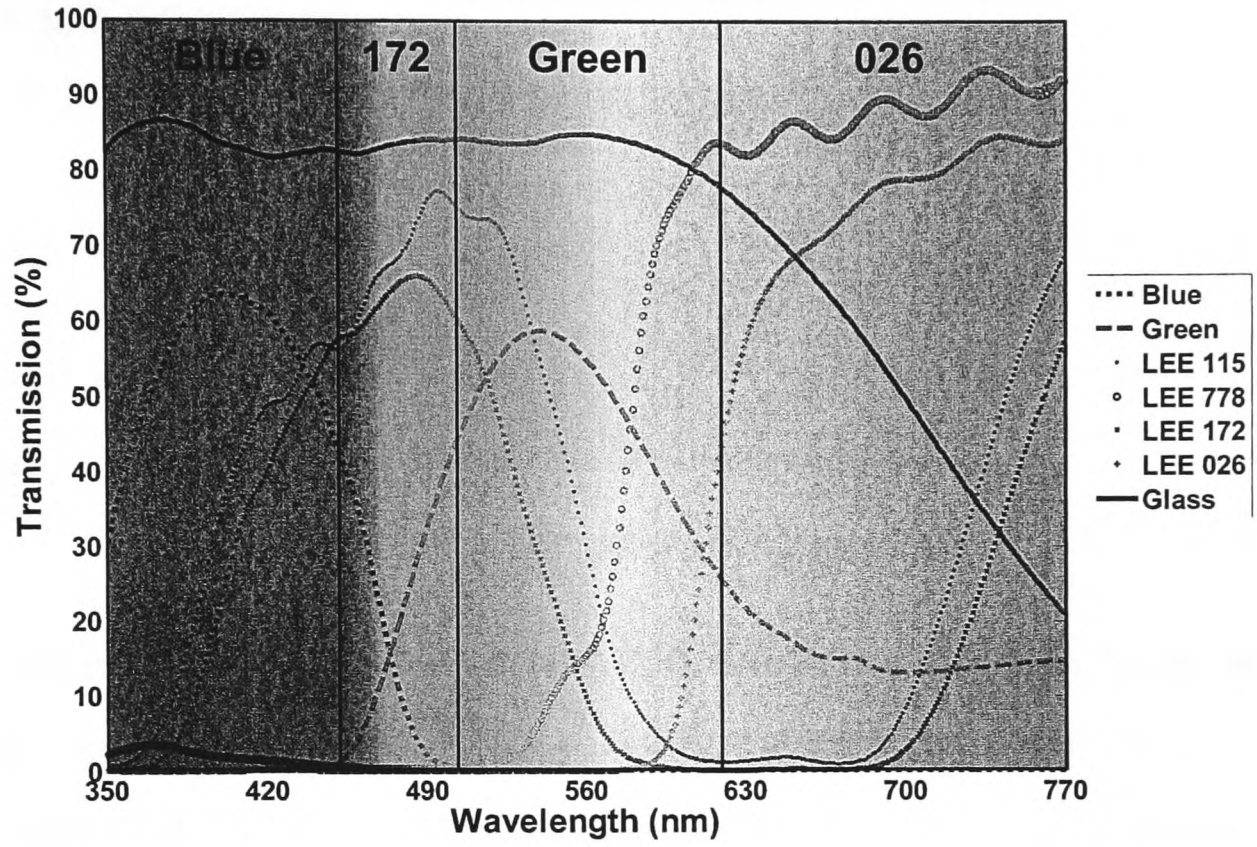


Figure 5.6: Transmission characteristics of filters

To obtain a sufficiently detailed and full calibration curve, the wavelength range of 350–770nm (in steps of 10nm) was chosen and 100 images at each wavelength were taken. Images of the background (with the lamp on and slit s_1 closed) were taken as well for noise subtraction. Figure 5.7 shows images at 470, 540 and 610nm respectively.

General image processing

As described in Section 2.4.3, The Pulnix TMC-9700 camera is a one-chip CCD camera with a colour filter array (CFA) to provide artificial colour information. The staggered pattern of the CFA is shown in Figure 5.8.



Figure 5.7: Images at different wavelengths as seen by the camera (470, 540 and 610nm left to right)

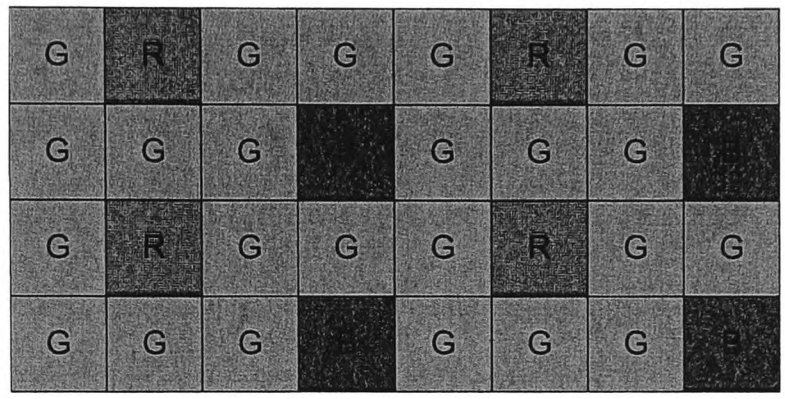


Figure 5.8: Colour filter array (CFA) of the camera

With the CFA present, each pixel only detects the intensity of a part of the whole radiation wavelength range, and the corresponding colour is referred to as its true colour. As the camera outputs RGB values for every pixel, then values of the other two colour channels are derived from on-chip digital signal processing (DSP) This can introduce cross talk between the true-colour channel and the other artificial-colour ones depending on the algorithm that has been implemented to provide the colour reproduction (whether in RGB or HSV). This effect was seen in preliminary work and led to a systematic error in the CRP. Unfortunately, this DSP process is performed before any image is downloaded into the computer and hence cannot be accessed and corrected for by the user. The same is true for most commercial colour cameras, and it should be noted that they may have different CFAs (the most widely used of which is the Bayer pattern [42]). Other on-chip DSP processes can be optionally performed via the camera software including contrast enhancement, edge sharpening and white balance correction, all of which are likely to bias the genuine information in the raw images by introducing artifacts. To minimize the cross talk, raw images from the camera were first processed as follows:

- only the true colour channel for each pixel was kept and the other two channels were discarded;
- then pixel values of the other two colour channels were reproduced through colour demosaicing.

Bilinear interpolation [42], modified for the particular CFA of this camera, was selected as the demosaicing algorithm for two reasons:

- it is straightforward to implement and requires little computational resource;
- interpolation of the artificial colour pixel is only realized within the colour channel, *e.g.* to derive the red value for a true green pixel only its neighboring true red pixels are involved hence avoiding inter-channel cross talk.

The bilinear colour interpolation was implemented differently for the green channel and the red channel while in the same way for the red and blue channels. The interpolation of green value for true red/blue pixels was derived by linear combination of the eight neighboring true green pixels. Take a true red pixel for example:

$$\begin{aligned} G^R_{i,j} = 0.16 \cdot \left(G^G_{i,j-1} + G^G_{i,j+1}\right) + 0.14 \cdot \left(G^G_{i-1,j} + G^G_{i+1,j}\right) \\ +0.1 \cdot \left(G^G_{i-1,j-1} + G^G_{i+1,j-1} + G^G_{i-1,j+1} + G^G_{i+1,j+1}\right) \end{aligned} \tag{5.5}$$

where the superscript denotes the true colour of that pixel, and i and j mean the pixel lies in the i^{th} row and j^{th} column. The coefficients 0.16, 0.14 and 0.1 were derived so that the two closest (in the same row as $R^R_{i,j}$) pixels were weighted more than the two less close ones (in the same column as $R^R_{i,j}$) and the four furthest ones (on the corners of the 3 by 3 region) were weighted the least. This is because the camera CCD sensor is not a perfect square but a rectangle with an aspect ratio of 1.17 and hence a diagonal-to-side ratio of $\sim$1.5. The green value for a true blue pixel was derived in the same way.

The interpolation of the red values for true green pixels was calculated in two steps. Those true green pixels which have two neighbouring true red pixels above and underneath them

were processed first as following:

$$R^G_{i,j} = 0.5 \cdot \left(R^R_{i-1,j} + R^R_{i+1,j}\right) \tag{5.6}$$

Then for the remaining undetermined non-true red pixels (of which there are three in a row between two pixels with known red values), linear interpolation was applied. The interpolation of the blue values for the true green pixels was carried out in the same way. Another benefit of this re-demosaicing process is that it smoothes each individual image (see Figure 5.9).



Figure 5.9: Effect of the re-demosaicing process

The thermal radiation from glowing tungsten or soot at a temperature lower than 3000K is red-dominant whilst the radiation from chemiluminescence is blue-dominant, so pixel saturation is more likely to happen on the red and blue channels than the green one. When pixel saturation occurs on a red or blue channel, pixels within the 3 by 7 array, the center of which is the saturated red or blue pixel, will be biased by the saturated pixel in the on-chip demosaicing process. Therefore, such pixels are assigned 255 at the end of the re-demosaicing code so that they can be readily excluded at a later stage.

5.3.2 Camera response study

Flat field correction

Owing to geometry and artifacts of manufacture, the sensitivity of different areas of the CCD detector may vary. This can be calibrated by exposing the detector to a uniformly illuminated field, namely a flat field. In practice, it is almost impossible to produce a perfectly uniform illumination, so keeping the variation of the field below the readout noise level is necessary, which is likely to be 1 bit for this camera. Hence for common applications where only half of the dynamic range is employed, it is advisable to make sure the flat field has less than 1% variation in intensity with light sensing devices such as a photodiode.

The same lamp as in Section 5.3.1 was put ~2m away from the camera to provide quasi-uniform illumination. The lamp was set to run at maximum current and voltage to generate similar levels of red, green and blue constituents *i.e.* grey light so that all three channels could be corrected for at the same time to a reasonable accuracy. Firstly the bare CCD chip was exposed to the field and a 50mm lens was used to account for the lens effect afterwards. With the lens, two pieces of diffuser paper were put between the camera and the lamp to avoid focusing on the bulb. In both cases, a photodiode was used to verify the flatness of the field.

The captured images were processed similarly in MATLAB as described in the previous sections. By dividing each pixel with the global mean of all pixels, a 2D matrix/image of normalised flat field correction coefficients was obtained and could be used to divide the combustion images to better resolve the spatial temperature distribution estimated by the colour-ratio technique.

It should be noted that when CRP is applied, although the camera is a single-chip one, a flat field correction must be carried out for the three colour channels separately, *i.e.*

three flat field matrices, otherwise any correction coefficient would be cancelled out when a ratio is generated; to some extent, the flat field correction for the colour-ratio technique is similar to a weak white balance correction.

With the same flat field setup, a study of the correlation between the aperture and the camera response was carried out as well. The aperture was assumed to be a perfect circle despite the actual shape (a curvy triangle) and different aperture areas were normalised to the one with a f number of 1.8 (maximum aperture for this lens):

$$\frac{A_{f_N}}{A_{f_{1.8}}} = \left(\frac{1.8}{N}\right)^2 \tag{5.7}$$

where A is aperture area and N is f number. Figure 5.10 shows the non-linear camera response with the aperture area. This suggests that care must be taken when performing analysis on images from different aperture settings. Although the aperture setting is unlikely to affect the results derived from colour ratios, it is sensible to keep using the maximum aperture unless image saturation or depth of field is of concern.

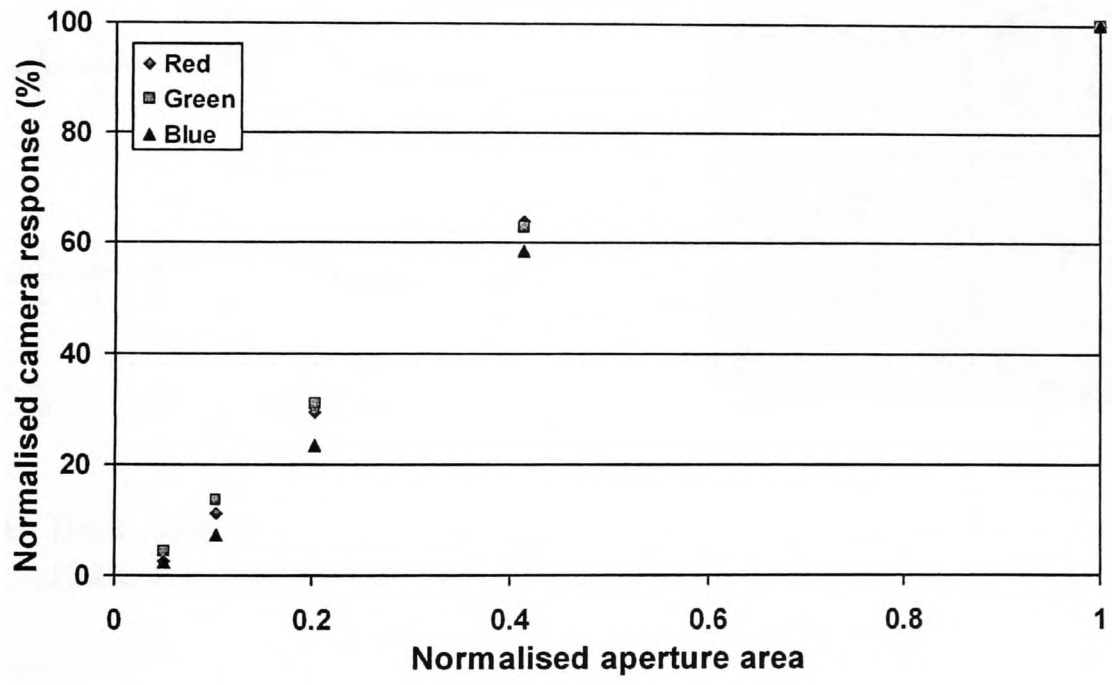


Figure 5.10: The camera's response test for linearity with aperture

Camera response test

Firstly, the camera's dark noise and dark current effect was studied with respect to gain and exposure by completely covering the CCD sensors.

Figure 5.11 (left) shows the correlation between the camera's dark noise and gain with a fixed exposure of 1ms. The camera seems to have a multi-amplifier system and for a gain higher than 190 (maximum 255) the dark noise is minimised both in amplitude and shot-to-shot variation. A higher gain which also minimises dark noise would be ideal in practice but unfortunately this camera features 'sticky bits' with a gain higher than 160, that is, many pixels appear darker than others when imaging a uniformly illuminated object as if they lost 1 or more bits of dynamic range. Therefore, gain 132 was selected at the beginning of this work as a compromise. Figure 5.11 (right) is the average image of 100 dark noise images with gain 132, exaggerated by a factor of 20 to show the pattern of the dark noise.

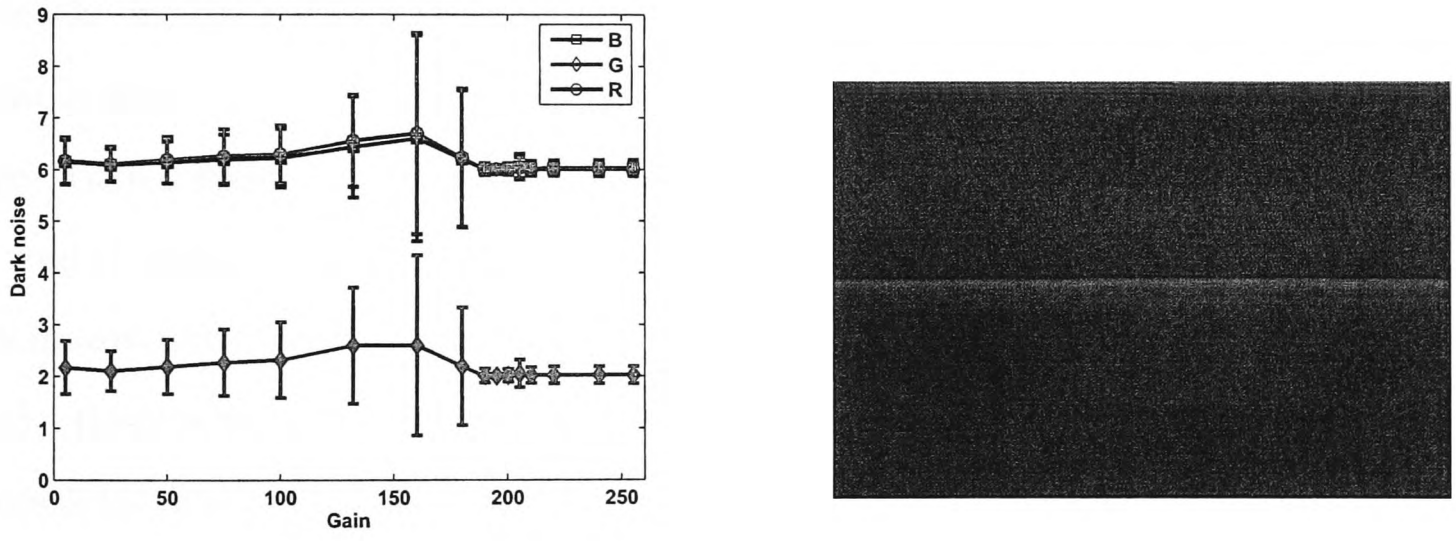


Figure 5.11: Dark noise with respect to gain (left) and amplified dark noise image with gain 132 and exposure 1ms (right)

The dark current effect refers to the fact that more dark noise builds up as exposure time increases. However, this effect is not significant for this camera as shown in Figure 5.12 (left) probably because the camera resolution is low. So the dark current effect is not considered in this research and the dark noise is taken to be constant regardless of

the exposure. Figure 5.12 (right) shows the same pattern as Figure 5.11 (right), also indicating that the dark noise is reproducible and has no correlation with exposure on a statistical basis.

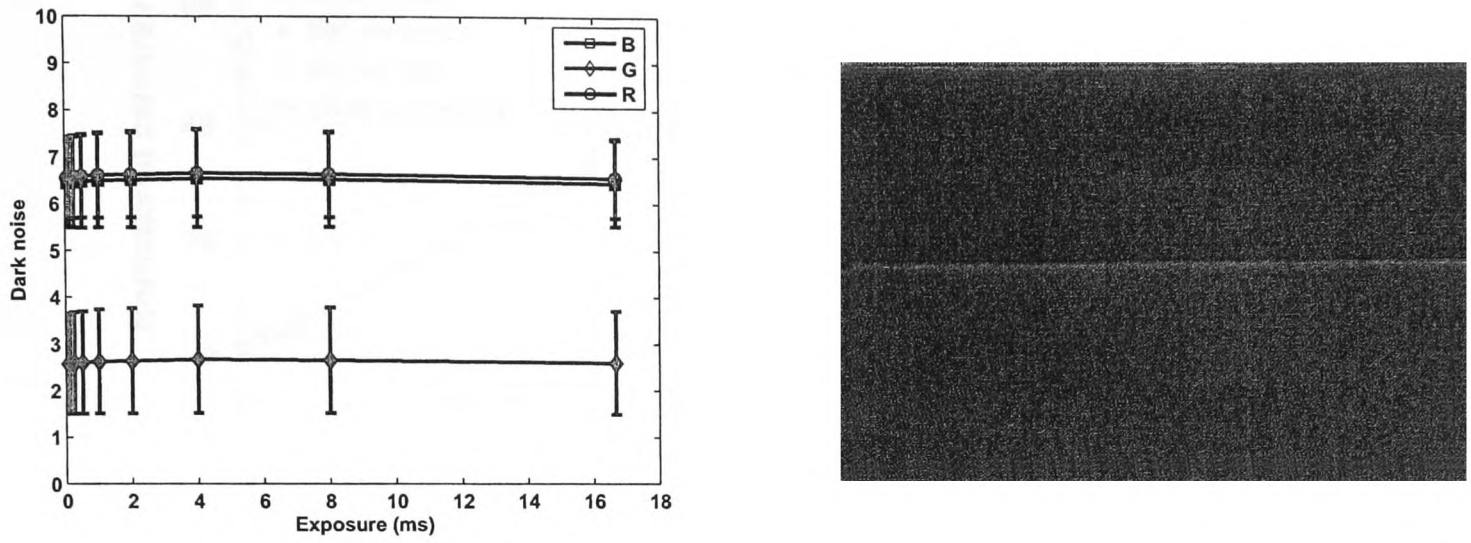


Figure 5.12: Dark noise with respect to exposure (left) and amplified dark noise image with gain 132 and exposure 16.6ms (right)

The camera's sensitivity was tested on the same spectrometer setup with the following approach: images were captured at a series of wavelengths; at each wavelength, a particular shutter setting was first selected without any filter to best utilize the dynamic range. Then different combinations of the seven filters (Figure 5.6) were chosen to reduce the level of radiation into s_1, whilst maintaining the shutter setting for that wavelength. Such images were normalized to the camera's response in the 100% transmission case (no filter). Results regarding the camera's response linearity with sensitivity are shown in Figure 5.13. The responses from the three colour channels are plotted separately. The non-linearity was within 2% over the wavelength range of interest, and hence the camera's response sensitivity was taken to be linear.

The camera's response linearity with exposure was also tested on the same spectrometer setup, so that:

- images were captured at a series of wavelengths;
- at each wavelength, an optimal shutter setting was first selected without any filter;

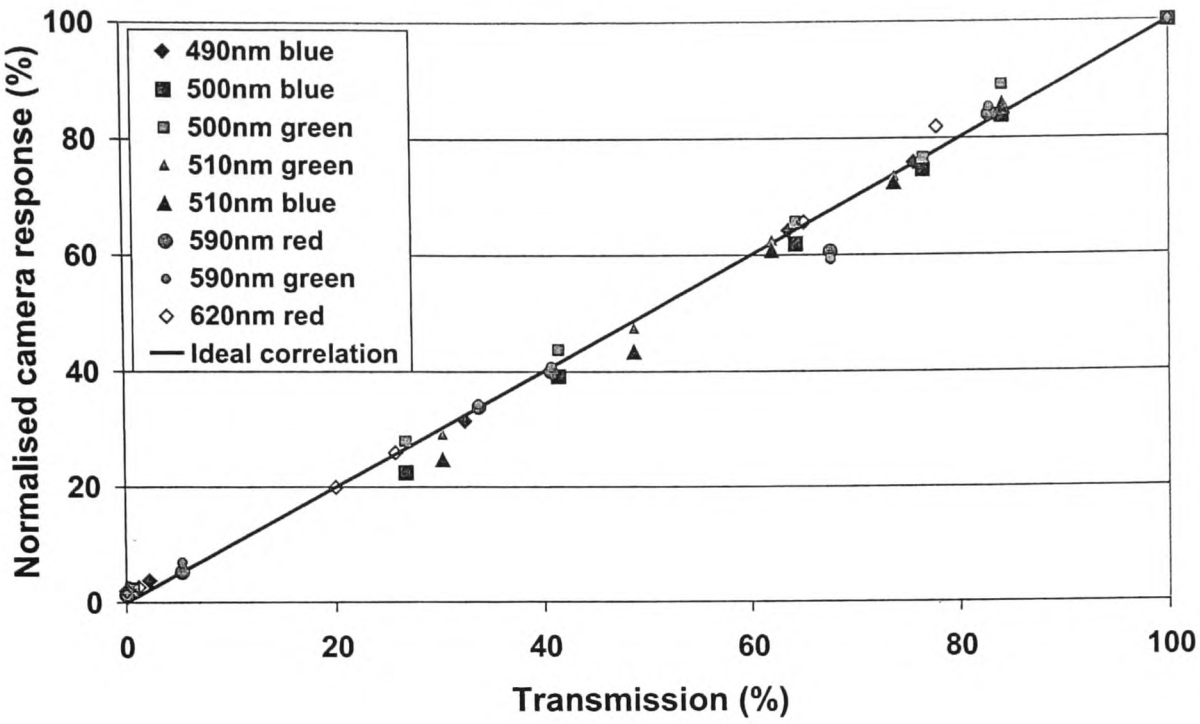


Figure 5.13: The camera's response test for linearity with sensitivity

- then different combinations of the seven filters were chosen to approximately give an each-step-halved transmission so that a doubled exposure time could be used to maintain the camera's response.

Since the camera's sensitivity was proved to be linear, the exposure-response correlation could be derived by

$$R_{i,j} = \frac{I_i}{I_j \cdot \gamma_{\text{tran}}} \tag{5.8}$$

where: $R_{i,j}$ is the response ratio between two shutter settings i and j, I_i and I_j represent measured camera responses with shutter settings i and j, and γ_{tran} is the transmission data involved.

Figure 5.14 presents the camera's response linearity with exposure. The shutter effect was then taken account of by normalizing the spectral calibration images captured using different shutter settings with respect to exposure time.

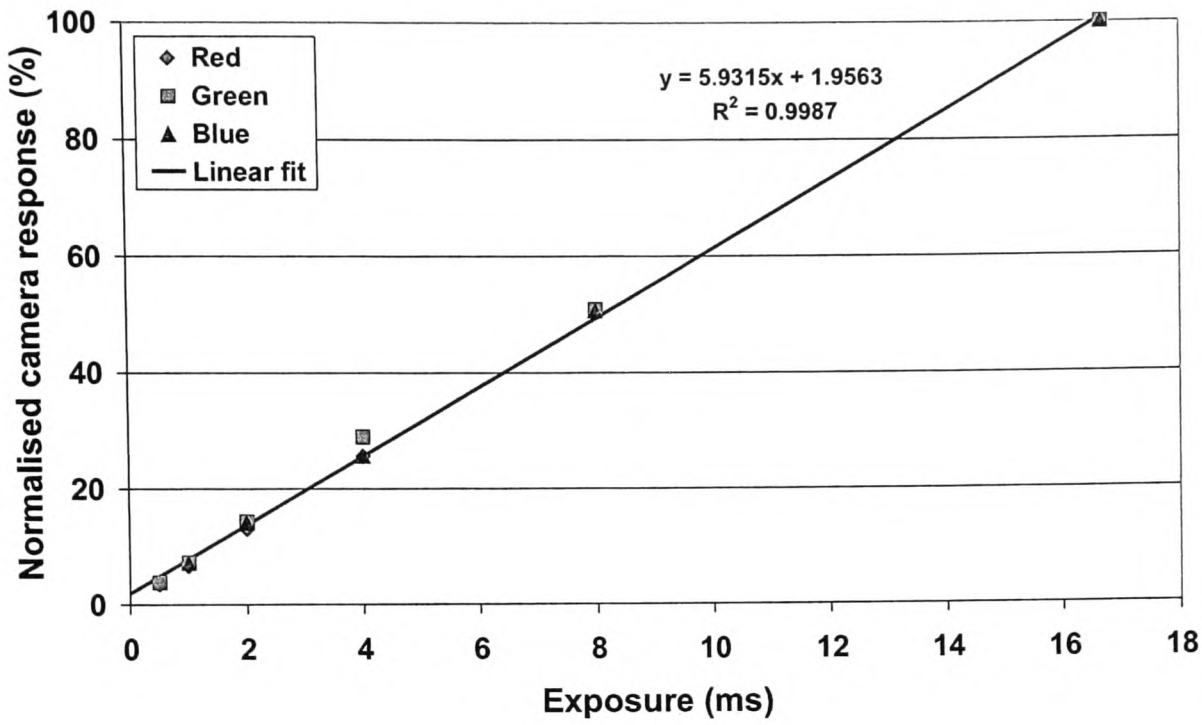


Figure 5.14: The camera's response test for linearity with exposure

5.3.3 Calibration image processing

The diffracted light beam showed an elliptic geometry and its intensity varied across the two dimensions of the image. Therefore only the central part (60 by 50 pixels) of each image was extracted and processed subsequently, and this effectively reduced the bandwidth of the wavelength analyzed to ~1nm.

Because the photodiode has a non-linear sensitivity with respect to wavelength, experiments with lasers of various wavelengths were conducted. Figure 5.15 (left) gave the sensitivity of the photodiode with respect to wavelength. The polynomial fit to these measured points showed good agreement with the datasheet supplied by the manufacturer (the outliers are due to uncertainties and instabilities in laser energy), and hence the manufacturer's calibration was employed in the following steps. This, together with the digital voltmeter (DVM) reading from the photodiode indicated the spectral emission of the tungsten lamp at a particular temperature (Figure 5.15 (right)). The DVM readings were taken when the photodiode detected the bulb radiation through the various filters and hence displayed a multi-stage shape. The 'corrected signal' refers to the photodiode

response after the photodiode sensitivity normalization.

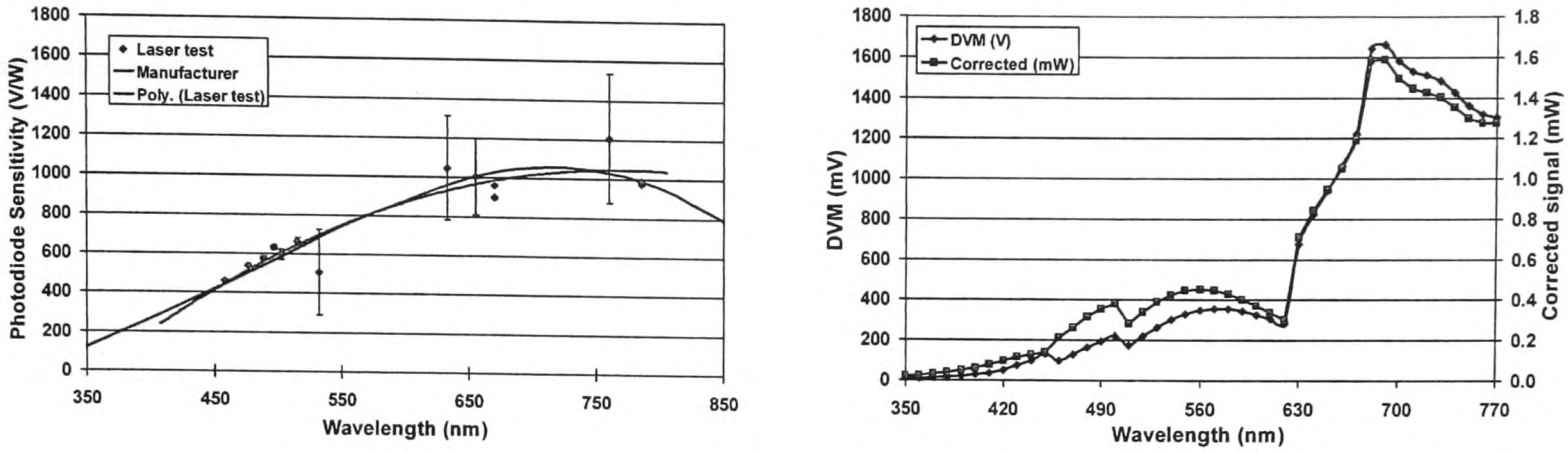


Figure 5.15: Photodiode sensitivity (left) and signal from the tungsten bulb (right) with respect to wavelength

To correct for the intensity variation of the tungsten lamp with wavelength, the RGB components at wavelength λ were normalised to the maximum intensity I_{max} by multiplying the factor I_{max}/I_λ (yielding Figure 5.16), where I_λ was the intensity at wavelength λ from the photodiode signal (Figure 5.15 (right)).

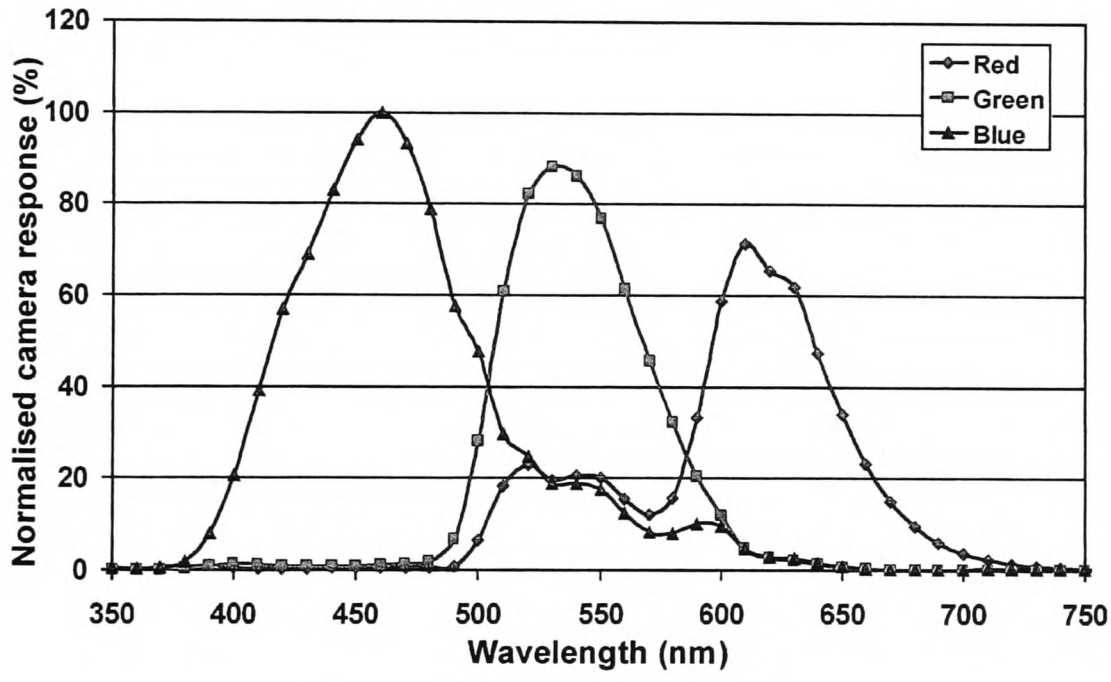


Figure 5.16: Normalized spectral response of the camera

As seen in Figure 5.16, in the range of green (500–600nm) the spectral responses of red and blue follow each other. This suggests that the signal from the green channel has affected the red and blue channels. One possible cause of this cross talk is miss-registration or

misalignment between the CCD chip and the CFA, that is, a signal that falls onto the true-green pixels spreads to the adjacent R or B pixels. As there will be variations between cameras and the published data are only nominal, it is essential to calibrate each camera.

5.4 Validation of CRP

To verify the CRP using the calibrated spectral response of the CCD camera, three sets of experiments were conducted, namely: imaging of a blackbody cavity, a tungsten filament lamp and a tungsten ribbon lamp. A 50mm lens with variable aperture was used for focusing. As only ratios were dealt with in the analysis, the non-linear effect of the aperture setting would not need to be corrected for. The transmission characteristics of the lens and optical pathway were also measured with the spectrophotometer and accounted for in the subsequent analysis.

5.4.1 Black body imaging

The black body used in this experiment was made from a steel bar with a cavity in one end, the diameter of which was an order of magnitude smaller than the length of the tube. The camera was focused to the inside back of the black body through the hole. An aluminum tube was manufactured to cover the camera lens to prevent inclusion of extraneous light. The black body (with a radiation shield) was then heated up by an oxy-acetylene flame, and its temperature was recorded by a thermocouple inserted behind the cavity in the black body.

Temperature readings and images of the black body were taken continuously until the solid end of the black body melted. Figure 5.17 (left) is a typical image of the black body from the experiment.

Comparisons between the measured and inferred temperatures of the black body are



Figure 5.17: Sample images from the validation experiments: black body (left), tungsten filament lamp (middle), and tungsten ribbon lamp (right)

shown in Figure 5.18 together with the corresponding standard deviations. The correlation between results from the R/G ratio and the measured temperatures was better and more linear than that from the R/B ratio. This is because in such a low temperature range the black body would emit much less blue constituents than red and green. Furthermore, the camera's dynamic range is only 8 bits. Therefore, it is difficult to resolve the temperature from the R/B ratio and such results would be associated with large uncertainties. This is also reflected by the standard deviation results. Better precision could be achieved with a 10 or 12-bit camera.

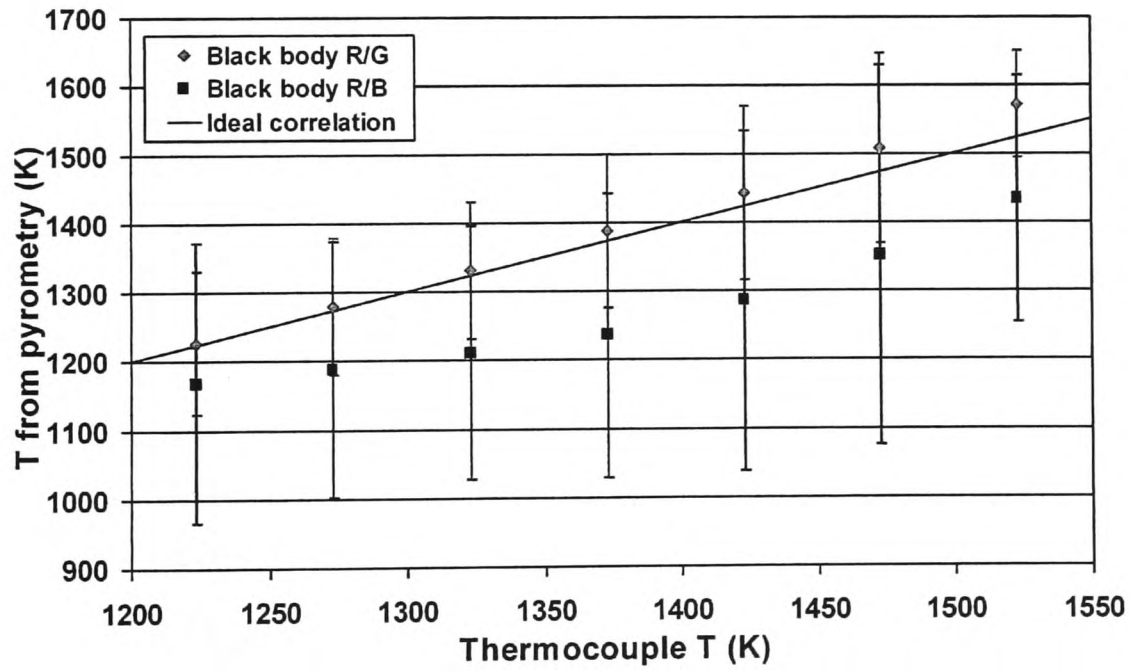


Figure 5.18: Correlation between measured and inferred temperatures of the black body

5.4.2 Tungsten filament imaging

The same tungsten lamp used in the camera spectral calibration process was utilised in this experiment. Figure 5.17 (Middle) is a typical image of the filament from the experiment. Instead of running the lamp under nominal conditions, the voltage and current on the lamp were varied to give different intensities and wavelength distributions of the light, and the voltage and current were recorded to estimate the filament's resistance as a separate temperature indicator. Temperatures from the CRP were then compared with the ones derived from the measured filament resistance (Weast & Astle [70]). The same necessary steps of image processing as above were carried out.

Initially the lamp filament was assumed to be a grey body (constant emissivity). Measured emissivity data for tungsten ribbon lamps at different temperatures and wavelengths (Touloukian & DeWitt [68]) were then used subsequently for improved accuracy. The emissivity data used are plotted in Figure 5.19, but it must be remembered that there are wide variations among the published data.

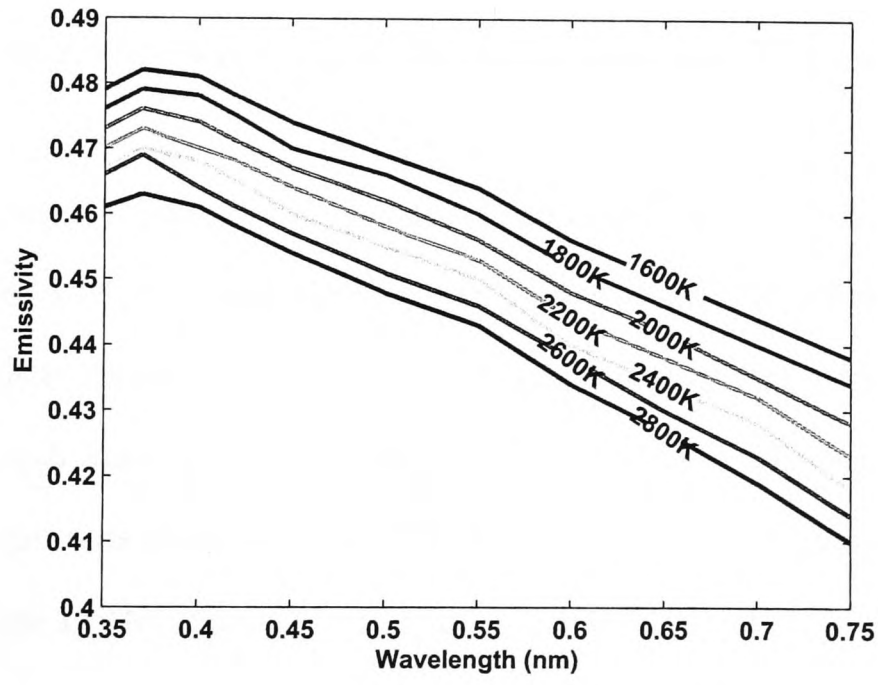


Figure 5.19: Tungsten emissivity data used in the tungsten filament test, data selected from [68]

Figure 5.20 shows the correlation between temperatures predicted by CRP and filament

resistance together with the corresponding coefficient of variation (CoV) for both cases. The CoV presented here was defined as the average of the spatial CoV of each image from the image sequence (100 images under the same condition), hence reflecting the temperature variation across the selected part of the filament.

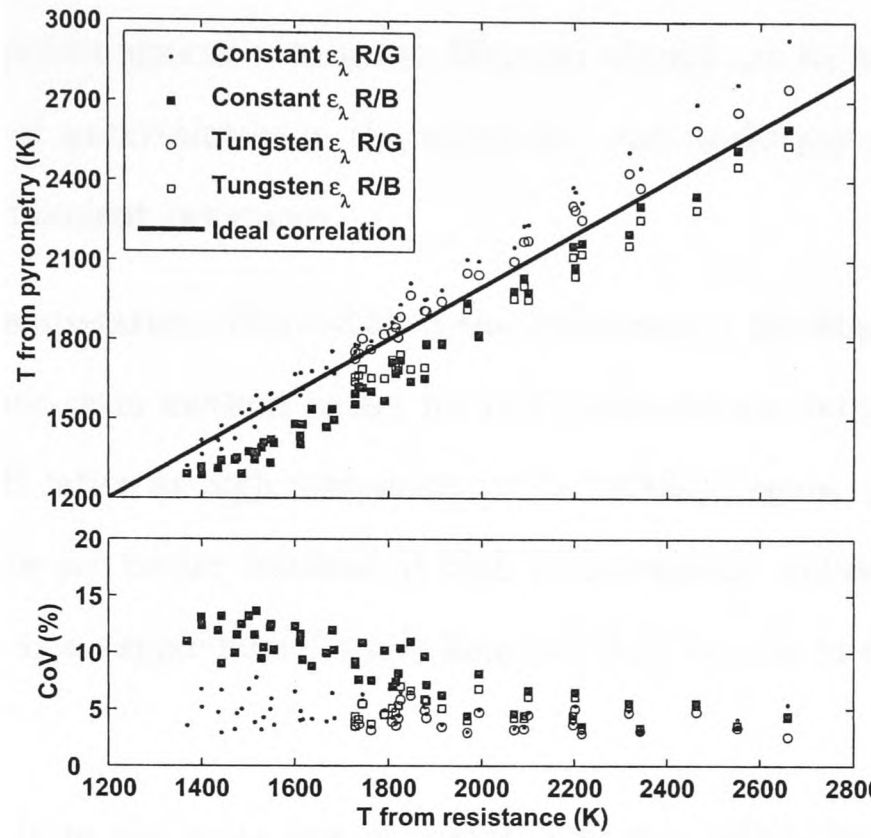


Figure 5.20: Relation between temperatures estimated with CRP and filament resistance

With realistic emissivity data the R/G ratio-based temperatures were generally < 100K lower than the ones with the grey body assumptions when the resistance indicated temperatures were below 2400K. This is because the emissivity of tungsten decreases with increasing wavelength and the relative proportion of long-wavelength radiation to short-wavelength radiation was smaller than that in the constant emissivity case. So the calculated colour ratios in the look-up tables become smaller, and for a given R/G ratio of a pixel, the estimated temperature moves towards the lower values. This effect is not so distinct for results based on the R/B ratio; the R/B data points with realistic emissivity values became lower than the ones with constant emissivity at around 2000K resistance temperature and above. Again this is because the lamp did not emit as much blue light as

green and red at low temperatures and larger uncertainties were likely to be introduced. However, the difference introduced by using experimental emissivity is insignificant within the temperature range for engine combustion. Moreover, the database of measured emissivity is coarse, incomplete, strongly dependant on the surface condition and, to a large extent, self-inconsistent for the wavelength range of interest here [68]. In conclusion, the temperature calibration against a tungsten filament should not be treated as an absolute reference because of uncertainties in the emissivity and uncertainty in the temperature deduced from the filament resistance.

In summary, the temperatures derived from the resistance of the filament agreed with the ones from the colour-ratio method better for R/G ratios at low temperatures (< 2200K) and better for R/B ratios at high temperatures (> 2200K). Again, temperatures derived from the R/B ratios are better resolved at high temperatures, indicating the potential of using R/B results as a supplement for checking the R/G results in the high temperature range.

The discrepancies from the unity line in Figure 5.20 may partly be due to uncertainties in the inference of the temperature from resistance, *e.g.* the non-uniform temperature distribution across the filament, the different resistance-temperature behavior of other parts of the lamp *etc.* The overall CoVs of this experiment scatter within 3–7% in the range of 1600–2600K except those low-temperature R/B points with constant emissivity.

5.4.3 Tungsten ribbon imaging

A high-stability tungsten ribbon lamp (POLARON 10/V) was employed for its capability of calibrating temperature measuring devices. A typical image of the ribbon lamp from the experiment is shown in Figure 5.17 (right).

The lamp was calibrated in terms of radiance temperature by reference to a National Physics Laboratory (NPL) standard tungsten ribbon lamp which has been calibrated

using a precision radiation thermometer, which has itself been calibrated with reference to a blackbody source at the freezing point of gold (~1064°C). At each calibration current the radiance temperature of the test lamp is determined by comparison with the calibrated standard lamp using a comparator.

The radiance temperature corresponds to the temperature of a black body at the calibration wavelength (660nm for this lamp). The radiance temperature is not the true ribbon temperature, but the temperature that a black body would have to be at, in order to emit the same intensity of radiation at the specified wavelength. Therefore, conversion between the archived radiance temperature of the lamp and the actual ribbon operating temperature must be done based on Planck's Radiation law and an estimate of the ribbon emissivity. Rewriting Eq. 5.1 and 5.2

$$I_{\lambda b} = \frac{2C_1}{\lambda^5 \left(e^{C_2/\lambda T_b} - 1\right)} \tag{5.9}$$

$$I_{\lambda t} = \varepsilon_\lambda I_{\lambda b}(\lambda, T_b) = \varepsilon_\lambda \frac{2C_1}{\lambda^5 \left(e^{C_2/\lambda T_t} - 1\right)} \tag{5.10}$$

with subscripts b and t representing black body and tungsten respectively. Letting $I_{\lambda t}$ equal $I_{\lambda b}$ and substituting Eq. 5.9 into Eq. 5.10, the correlation between the temperature of the tungsten ribbon T_t and its corresponding blackbody temperature T_b (also the radiance temperature) can be written as:

$$T_t = \frac{C_2}{\lambda_c \ln\left(\varepsilon_\lambda e^{C_2/\lambda T_b} + 1 - \varepsilon_\lambda\right)} \tag{5.11}$$

As shown in Figure 5.21, a fairly linear correlation existed between the estimated temperatures with CRP and the deduced lamp operating temperatures, it being more linear for the R/G results than the R/B results. The R/B points seemed to match with the lamp operating temperature better than the R/G ones, however, they were more scattered

and the difference between the two series decreased with increasing lamp temperature. Again there is uncertainty in the emissivity of the tungsten ribbon. This manifests itself as uncertainty in its true temperature and as uncertainty in the spectral intensity of the radiation so the ribbon lamp does not provide an absolute reference.

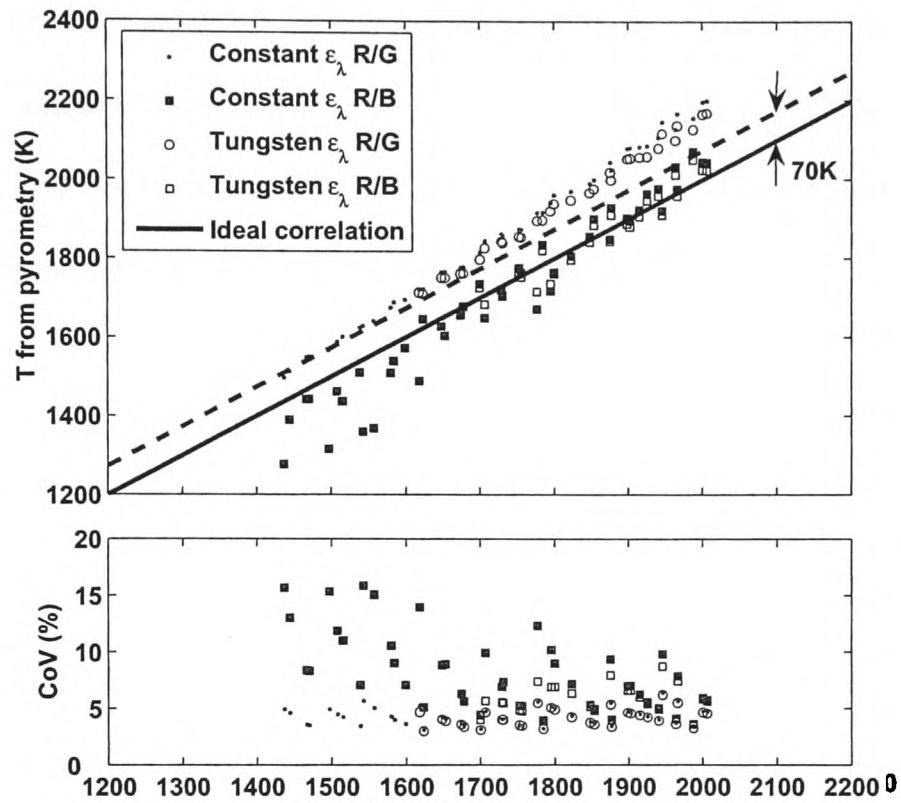


Figure 5.21: Relation between temperatures estimated with CRP and the estimate of the ribbon lamp operating temperature

An examination of the spectral output of the tungsten ribbon lamp in the range 640–910nm at each temperature suggests that there might be a systematic error in the black-body calibration of the lamp temperature; this could be an aging phenomenon. In Figure 5.21 a broken line has been added to account for this deduced temperature offset of 70K.

The overall CoVs of this experiment fall within 3–10% for the temperature range of 1600–2000K except those low-temperature R/B points analysed assuming a constant emissivity.

5.4.4 Concluding remarks on the temperature calibration

In comparisons between the temperatures derived from CRP and the reference temperatures provided by the three validation experiments, the R/G colour ratio results showed excellent linearity; the R/B results were more scattered but were useful as supplementary information, particularly for the high temperature range — this may be improved using a camera of higher dynamic range. To minimize the possible error due to all the noise sources and best utilize both the R/G and R/B colour ratios, a weighted-mean temperature T_c was adopted as the output from the CRP:

$$\begin{aligned} T_C &= a \cdot T_{R/G} + (1-a) \cdot T_{R/B} \\ a &= 1.5 - T_{R/B}/2000 \\ |T_{R/G} - T_{R/B}| &< 200\text{K} \end{aligned} \tag{5.12}$$

where a and $1-a$ are the weighting coefficients with an applicable range of 1000–3000K. A premise of this calibration is that the difference between the temperatures from R/G and R/B estimates must be below 200K. Any image not compliant with this premise is likely to be due to chemiluminescent flames and hence was not taken as valid.

Throughout the range of 1200–2800K, the weighted temperature from the blackbody and tungsten filament experiments showed a deviation of less than ±50K from the ideal correlation and the one from the tungsten ribbon experiment is always 50–100K higher than the reference temperature. Taking account of the 70K systematic offset mentioned in the previous section, the deviation in the tungsten ribbon experiment is also less than ±50K.

5.4.5 Error analysis

Firstly, the ±1 bit digitization error caused by the A/D conversion process of the camera was investigated. Such an error in the colour values (intensity) can lead to an error in the

colour ratios, which results in an intrinsic error on the temperature estimate. Quantifying this error helps to understand the accuracy of the colour-ratio method and set a reasonable step size in the look-up tables. The digitization error can occur on any one, two or three colour channels at the same time, therefore the worst case, in which there was +1 bit error in the red channel and -1 bit error in the green and blue channel was chosen for this study.

For a range of temperatures, a relative intensity of 127 was assigned to the red component, which was stronger than signals from the other two channels, and the green and blue values were calculated by dividing the red value by the ratios in the look-up tables. Then 1 was added to the red values and subtracted from the green and blue values to generate new colour ratios. These new ratios were mapped back using the original look-up tables and the error-added temperatures were returned. The error from the R/G and R/B ratio calculation is presented in Figure 5.22.

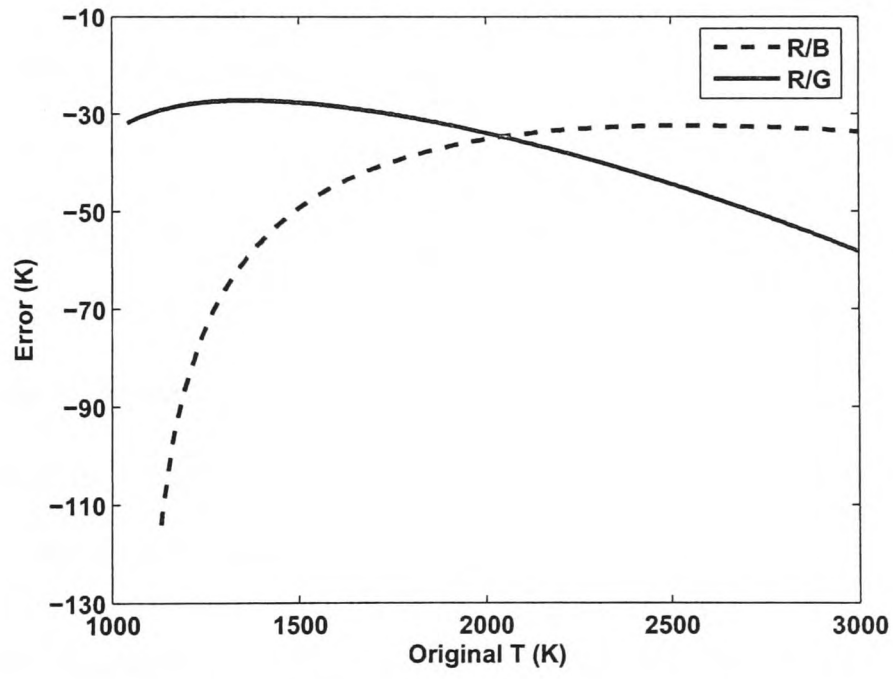


Figure 5.22: The effect of 1 bit digitization error on the estimated temperature

The step size of the look-up tables used for Figure 5.22 was 10K and was found to give reasonable performance. The error presented is always negative because this case referred to ± 1 bit error introduced to the numerator and denominator of the colour ratios

respectively. The error curves would be approximately inverted for the opposite case.

Since the 1 bit digitization error may not happen on all three channels at the same time, the case presented here indicates the worst-case error. However, the relative intensity of 127 may not be achieved and the utilisation of the dynamic range may be poor in practice. Taking these two factors into consideration, error in the temperature estimation due to the ± 1 bit digitization error can be predicted to be within $\pm$50K for the temperature range of engine combustion (1500–2600K). It should also be noted that although Figure 5.13 indicates linearity uncertainties of $\sim$2% (2–3 counts for a camera response of 127 counts), the source of such uncertainties are most likely to be a combination of both shot-to-shot noise and digitization error which cannot be decoupled from one another. Therefore, the ± 1 bit digitization error is still representative of most cases encountered in practice.

The ± 1 bit digitization error analysis above was applicable to a single pixel, whilst the effect of random noise on a global averaging basis was studied as well since spatial averaging was applied in the validation experiments. In this study, standard artificial images of uniform colour ratios were produced according to the look-up tables for a series of temperatures. Then different levels of random noise was added to the three colour channels of these images, namely, noise-to-signal ratio (NSR) of 0.05–0.25 in steps of 0.05, in which case signal means the colour value of a particular colour channel in the images. These new images were processed using the original look-up tables (step size 10K), and the returned error-added temperature field was spatially averaged. Figure 5.23 shows the error from the R/G and R/B ratio calculation.

The resultant error displayed a feature with respect to the original temperature as opposed to a random distribution that one would expect. However, this is a genuine fact in particular for the colour-ratio technique presented here. Firstly, the algorithm for CRP was interpolation with fixed KL in which the boundaries of the look-up tables for interpolation had to be set before any calculation was carried out; in this study the boundaries were 1000 and 3000K. Therefore, as the original temperature approached either boundary,

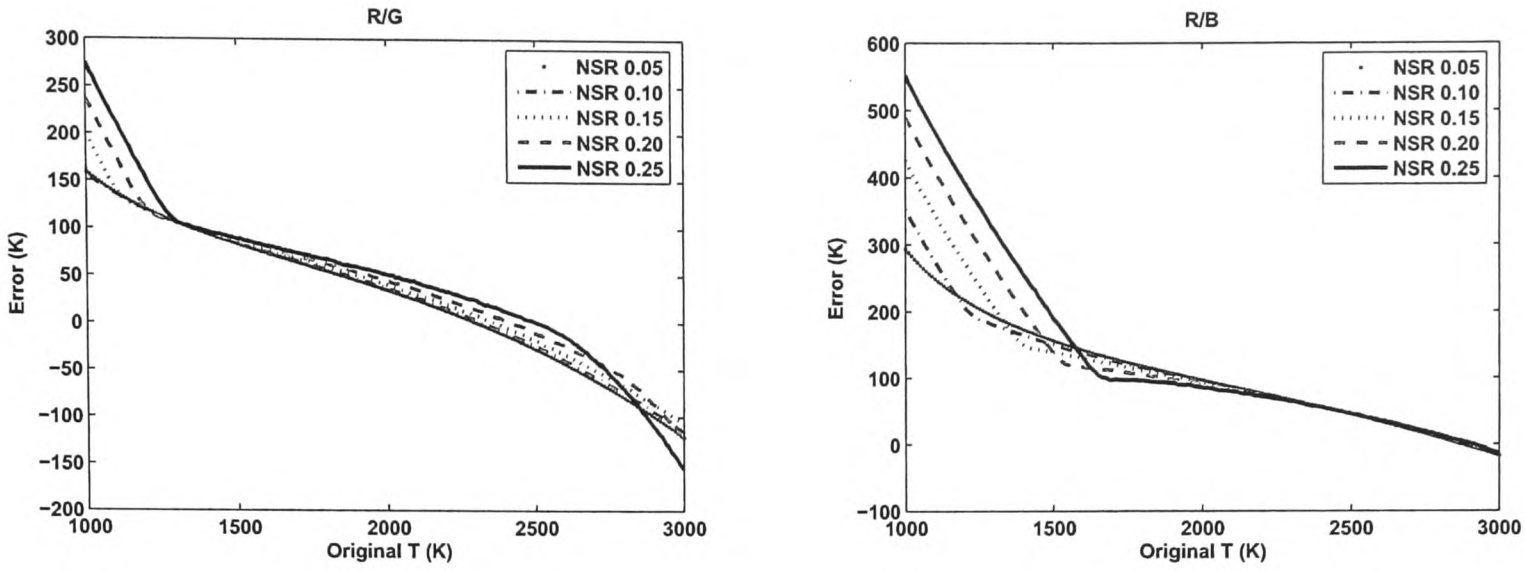


Figure 5.23: The effect of random noise and global averaging on temperature estimation

the random noise distribution added to the original temperature field was cut off on one side and became distorted, resulting in the increasing deviation from zero on both ends of the x axis in 5.23. One reason why the error behaved differently at low temperature and high temperature may well be the shape of the colour ratio curves in 5.4, which monotonically decrease with decreasing gradient. For instance, adding the same amount of numerical noise (NSR) to the standard image at 1000K and 3000K would result in a different amount of temperature error due to the different scales of interpolation. By this means, the temperature error can be estimated to be 50K for the temperature range of 1800–2600K using R/G results and 2400–3000K using R/B results.

Recalling the deviation of the experimental results from the 'unity' line in the above sections, this deviation is likely to be the combined error of the ± 1 bit digitization error of individual pixels, the spatially averaged random error and the uncertainties in the experiments themselves.

5.5 DISI engine results

A typical arrangement for the pyrometry measurement on the Jaguar engine is shown in Figure 5.24. An image sequence of 100 images was taken for each engine operating condi-

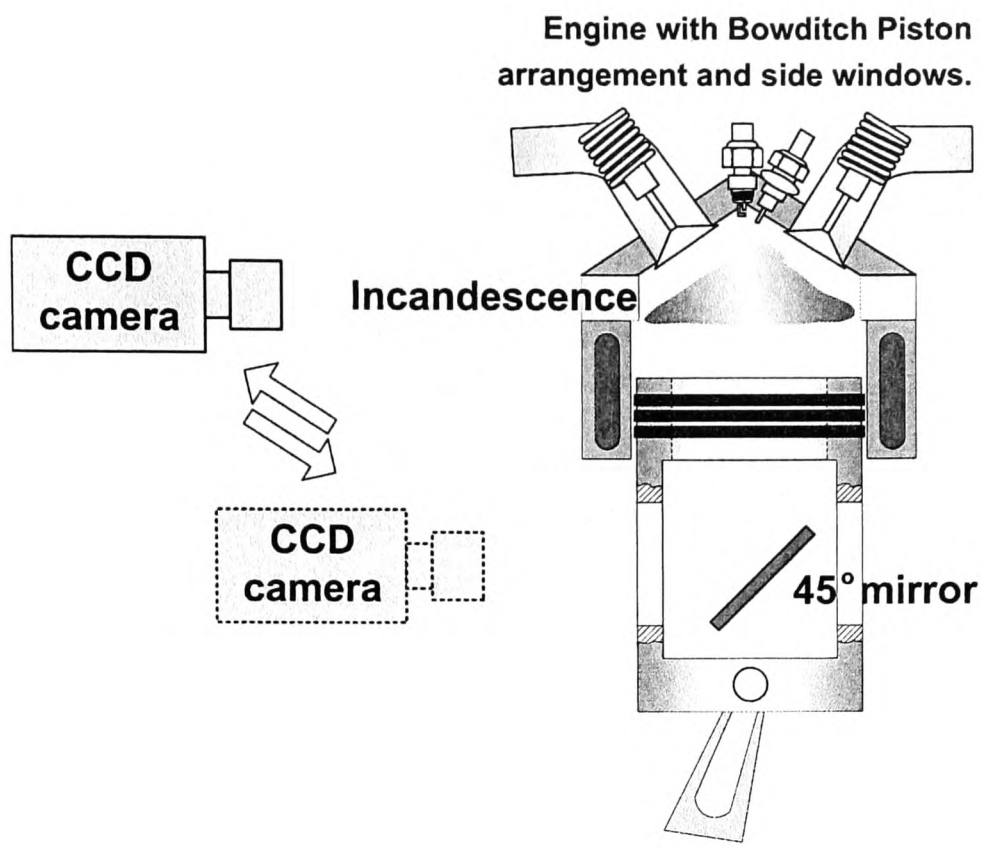


Figure 5.24: Set-up for measurements on the optical DISI engine with CRP

tion with an angle increment of 5° after taking 10 images at each crank angle setting. This gave an imaging duration of 50° after ignition. The exposure time was adjusted accordingly to avoid major image saturation but fixed within each image sequence. All image results presented in this section were flat-field corrected and were for engine operating conditions of:

1500RPM, ignition timing 35° before TDC, injection timing 280° before TDC and inlet plenum pressure -0.48bar (gauge).

A sample image from the Pulnix TMC-9700 CCD camera through the pentroof window (marked with white lines) is shown in Figure 5.25, with the engine operating in a rich homogeneous mode, with toluene and a lambda value of 0.8.

The pyrometry results for this sample image have been analyzed with different methods, and the results are summarized in Table 5.1. With variable KL, both the R/G and R/B colour ratios were input to the map created beforehand with temperature limits of 1400–3000K and KL limits of 0.1–1.4 representing typical soot conditions. This resulted in a single temperature output as well as KL.

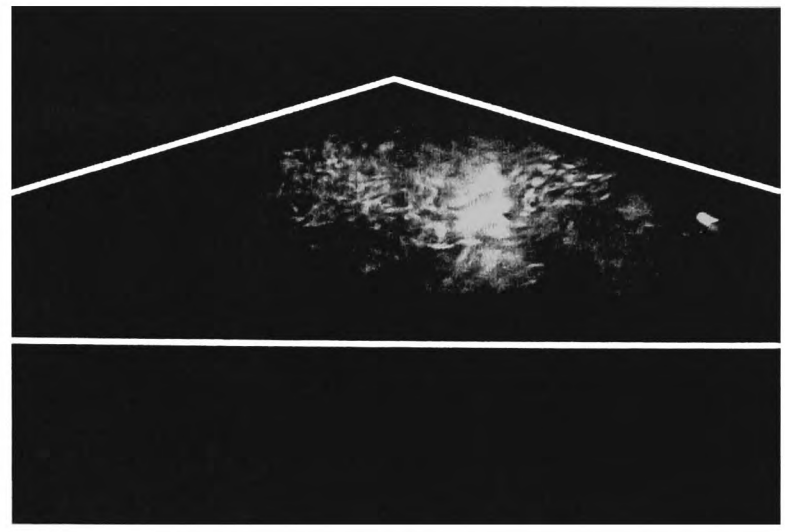


Figure 5.25: Sample image from the engine in homogeneous mode, with toluene, $\lambda = 0.8$, ignition timing 35° before TDC, injection timing 280° before TDC, image timing -10° after TDC

	KL				
	Variable	**0.2**	**0.5**	**1.0**	**Grey body**
mean$_{R/G}$	2397	2502	2550	2601	2649
mean$_{R/B}$		2422	2471	2523	2572
st.d.$_{R/G}$	241	271	265	260	255
st.d.$_{R/B}$		283	277	271	266

Table 5.1: Summary of pyrometry results for the sample image in Figure 5.25 with different KL settings

The predicted temperature with fixed KL increases with the prescribed KL and was the highest when the grey-body assumption was made. This is because the emissivity decreases with KL, more with a smaller KL, according to Hottel and Broughton's formula (Figure 5.3) and hence the longer wavelength range was weighted less in the integrating process when a prescribed KL was introduced.

There was a small difference in temperature between the variable KL and fixed KL cases. An unfortunate corollary of this is that the calculations of KL will be ill-conditioned. The standard deviations of the spatial temperature distribution in all cases are comparable and led to a CoV of ~10%.

Figure 5.26 shows the false colour images of in-cylinder temperature and KL distribution from the sample image (Figure 5.25) using the variable KL strategy. Invalid areas such as outside of the pentroof window, the spark plug and the injector tip were overlaid as

a black layer. Algorithms were incorporated for detecting saturated pixels and potential chemiluminescence pixels (R/B < threshold). Such pixels are obvious in Figure 5.25 and shown in Figure 5.26 as the higher limit values but not included when the spatial average was calculated.

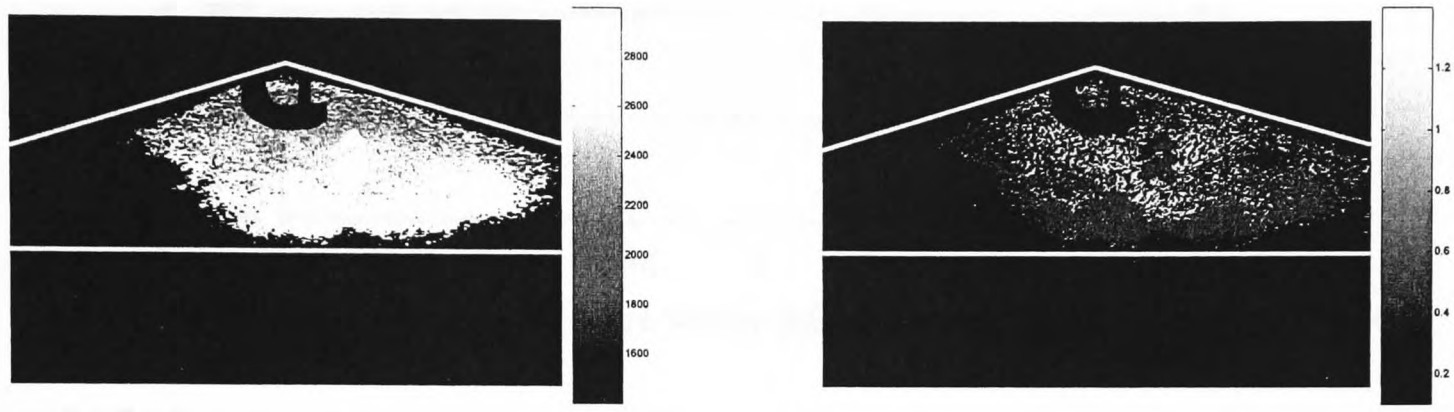


Figure 5.26: Temperature (left) and KL (right) distribution for the sample image of Figure 5.25 with variable KL

Figure 5.27 presents the results for the spatially averaged temperature T_c and the crank angle timing from the image sequence of $\lambda = 0.8$. The first detectable flame occurred 15° after ignition and the predicted temperature seemed to peak around 10° after TDC. This coincided with the time when ~50% mass fraction burnt and when the maximum heat release rate occurred. Because the intensity of the soot emission varied dramatically over time, the image timing range was limited by the relatively narrow dynamic range of the camera system.

The relation between the spatially averaged temperature T_c and the mixture strength is shown in Figure 5.28 with an image timing of 5° after TDC. The data points were sparse and more scattered at lean operations because

- the combustion chemiluminescence became dominant in most of the images which must be excluded from the analysis;
- the engine instabilities increased when operating lean.

It should be noted that he polynomial fit in Figure 5.28 may be optimistic, and in the

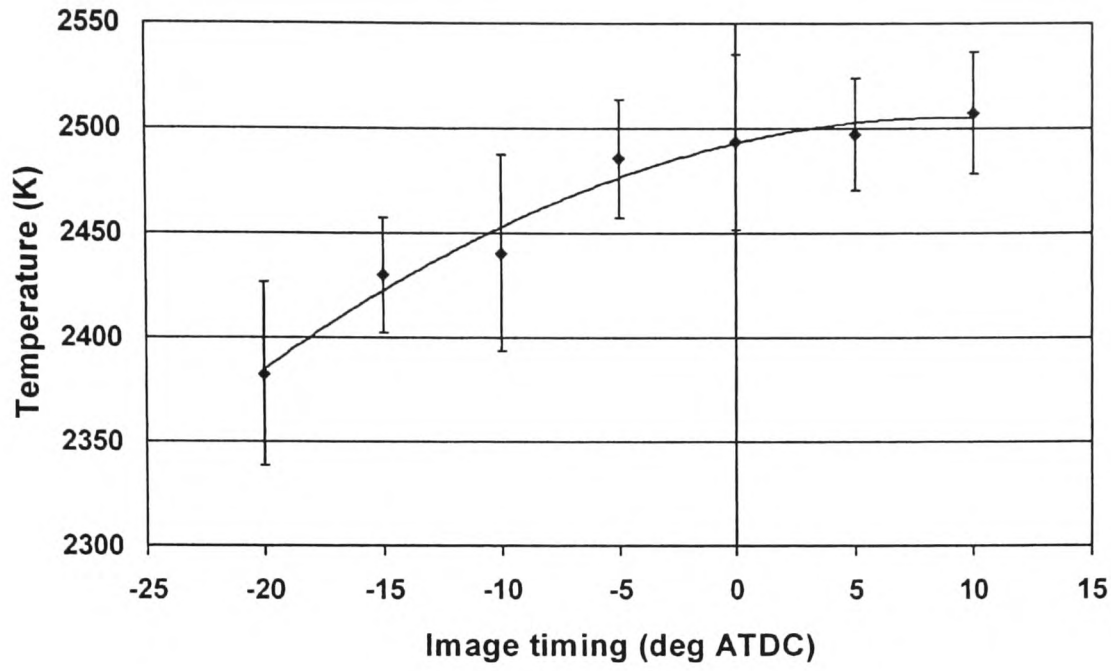


Figure 5.27: Relation between temperature from CRP and crank angle with $KL = 0.5$, $\lambda = 0.8$

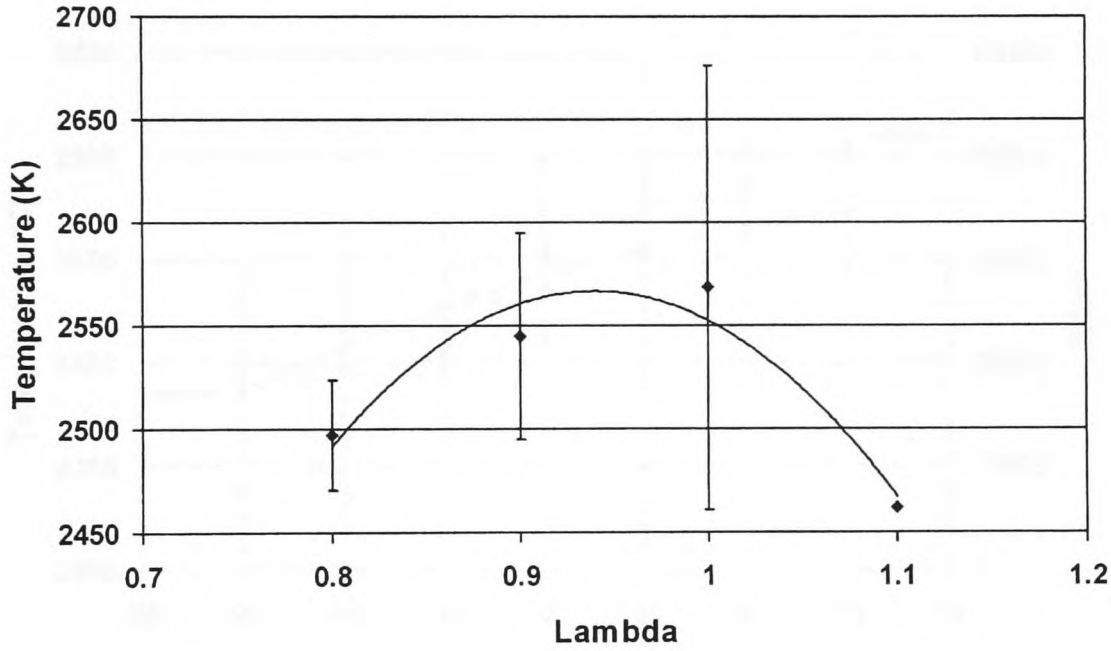


Figure 5.28: Relation between temperature from CRP and mixture strength with $KL = 0.5$; image timing = 5° after TDC. Each data point comes from a separate image

case of $\lambda = 1.1$ no error bars are shown because under such conditions only one image was analysable. The above statement applies to Figure 5.30 as well.

Due to the large associated uncertainties, caution is needed when interpreting images from mixtures leaner than stoichiometric with CRP. However, as would be expected, the polynomial-fit curve shows a temperature peak occurring slightly richer than stoichiometric.

Despite the limited KL resolution when adopting the variable KL strategy, a parameter, $N * KL$, can be used to quantify the loading of soot:

$$N * KL = \sum_i KL_i \cdot N_{KL_i} \tag{5.13}$$

where: N stands for the number of pixels, representing the area within which the pixels give a certain KL value. Therefore, $N*KL$ serves as an indicator of the total soot number within the enflamed volume.

Any KL value above 1.3 was likely to be a result of image saturation or chemiluminescence, and hence it was excluded from the calculation. Results with variable KL from the same image sets as in Figure 5.27 and Figure 5.28 are shown in Figure 5.29 and Figure 5.30.

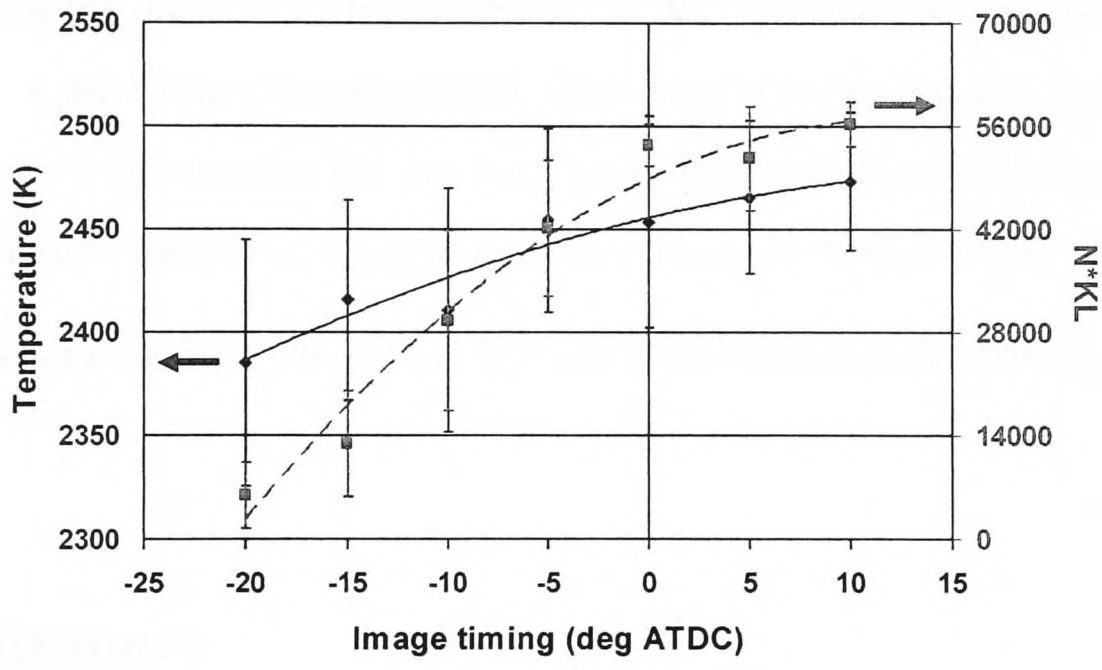


Figure 5.29: Relations between variable-KL results from CRP and crank angle, the other conditions the same as in Figure 5.27

The temperature curve in Figure 5.29 looks largely the same as Figure 5.27, again indicating that the temperature estimate is insensitive to the concentration of soot. As expected, the loading of soot increased as combustion proceeded and reached the highest level at the latest imaging crank angle.

The relation between $N * KL$ and mixture strength in Figure 5.30 shows that the soot

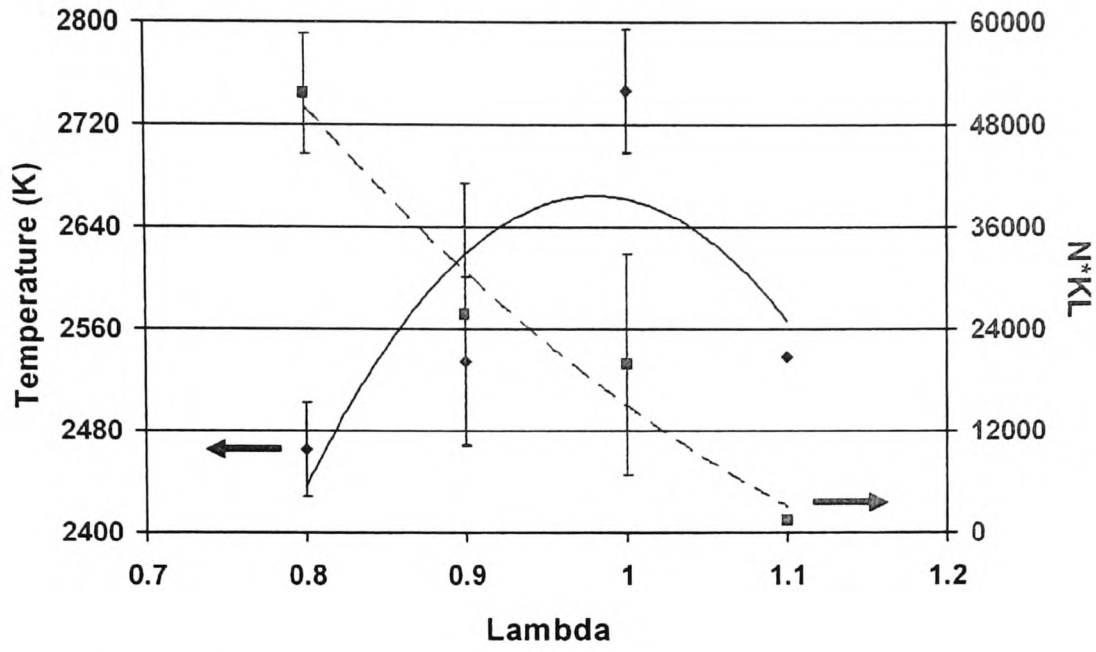


Figure 5.30: Relations between variable-KL results from CRP and mixture strength, the other conditions the same as in Figure 5.28

production was the highest at lambda 0.8 and quickly diminished to extremely low levels as the mixture went lean. The significant discrepancies in temperature between Figure 5.28 and Figure 5.30 when the mixture went lean could be because that the variable-KL algorithm is more sensitive to chemiluminescence than the fixed-KL one.

Comparisons between results from the CRP and engine simulations will be given in Chapter 7.

5.6 Summary

A spectral calibration was carried out to identify the spectral response of a colour CCD camera, with which colour-ratio pyrometry was applied to images from three validation experiments, namely: imaging of a blackbody cavity, a tungsten filament lamp and a tungsten ribbon lamp. The algorithm and implementation of CRP and the required image processing have been discussed.

Results from the bench experiments showed satisfactory agreement between the temperature derived from CRP and the 'reference' temperature. The sources of uncertainty in the

temperature and emissivity of tungsten filaments have been discussed. The R/G results were better at lower temperatures and the R/B results were better at higher temperatures, revealing the potential of an overall temperature calibration for the colour-ratio technique to best utilize information obtained from both the R/G and R/B ratios. The deviation of experimental data from a linear fit was within ±50K in the temperature range of 1200–3000, hence rendering the CRP a practical technique for estimating in-cylinder temperatures of internal combustion engines.

Error analysis was performed to understand the accuracy of the colour-ratio method, including the effect of camera digitization error and spatial averaging of random noise on the temperature estimation. Taking the worst case of the ±1 bit digitization error into consideration, the error in the temperature estimate was still within ±50K for the temperature range of engine combustion (1500–2600K). The possible error caused by spatially averaged random noise was estimated to be within ±50K for the temperature range of 1800–2600K. It should be noted that the aforementioned accuracy was obtained in an engineering application context with careful calibration and an 8-bit colour camera. Research *e.g.* [15] has shown that the accuracy of soot pyrometry is sensitive to soot properties such as size distribution and complex refractive index, and largely depends on the theoretical assumptions introduced. So the accuracy of CRP is unlikely to exceed ±20K even if an ideal digital colour camera could become available. The uncertainty in the KL solution would be inherently higher than that of temperature.

Combustion images from a spray-guided DISI engine were then interpreted using the CRP. A comparison of results from different prescribed KL values indicated that the predicted temperature was insensitive to the KL setting (whether variable or fixed). The general trend was that the temperature increased with increasing prescribed KL, and reached the highest when the soot was assumed to be a grey body. False colour images of spatial temperature and KL distribution have been demonstrated; however, the KL solutions were ill-conditioned. Using a camera with a better dynamic range would help increase

the reliability of the KL measurements.

Finally relations between the spatially averaged temperature and the crank angle timing as well as mixture strength have been presented. The camera's dynamic range posed constraints on the selectable range of image timing. Similarly, results with variable KL were reported and the parameter $N * KL$ was introduced to extract information on soot formation. Due to the presence of combustion chemiluminescence, additional care is necessary when dealing with images from lean engine operating conditions.

Chapter 6

High-Speed Combustion Imaging

Speed provides the one great modern pleasure.

— Aldous Huxley (1894–1963)

6.1 Introduction

The pyrometry results shown in Chapter 5 were not crank-angle resolved simply because of the frame rate limit posed by the CCD camera. In order to gain insights into the engine performance during an individual combustion event, photographic facilities with an imaging speed sufficient to resolve sub-five degrees crank angle is required, more preferably, in line with the cylinder pressure measurements. High-speed digital cameras with advanced CMOS technologies nowadays are suitable for such tasks which formerly had to be carried out by cine film equipment, *e.g.* rotating prism cine cameras, in earlier research work.

This chapter furthers the discussions of soot pyrometry on the DISI engine in a crank-angle resolved context. In addition, qualitative comparisons are made between the high-speed image sequences and burn rate data from combustion analysis. It will be seen that visual perceptions of flames do not necessarily coincide with the mass fraction burned at a

particular stage of the combustion process. Advanced image processing is then introduced for detecting the flame front and studying its growth. The developed techniques also provide the potential for deriving the velocity field of the flame front propagation with planar optical diagnostics.

Results for image sequences with full bore optical access are presented as well.

6.2 Experimental aspects

The high-speed camera used in this work was a Phantom v7.1 CMOS camera, the specification of which has been described in Section 2.4.4. Using a 55mm Nikon objective lens and the 45° mirror optical access, the camera was set up to look at the inner cylinder head. The centre of the field of view was set to be the gap between the spark plug and the DI injector. In all images, the orientation of the combustion chamber was such that the injector was above the spark plug.

The ETCS program sent out a 'fsync' signal every engine cycle to initiate the image sequence for each cycle and the camera took the encoder clock pulse (360/rev) to trigger image capturing. A timing control box was built with the function to select the image timing interval in degrees CA and to ignore the clock pulses when there was no combustion event [43]. For example, when an interval of 5°CA was set on the box and an imaging duration of 100°CA was set on the ETCS panel, 20 images would be taken after the spark timing. The imaging interval needed to be longer than the exposure time to give the camera enough time for shutter opening and closing, data downloading *etc.* For bright yellow sooty flames, a short exposure time must be used to avoid pixel saturation so a faster frame rate (*e.g.* 1 image/2°CA) was then available for a better temporal resolution. On the other hand, when combustion chemiluminescence was dominant in the image, the exposure had to be prolonged to get sufficient contrast against the background, thus a bigger image interval *e.g.* 5°CA must be set via the box.

A camera resolution of 512 by 512 pixels was chosen to allow for an appropriate frame rate, typically 6000fps. The exposure time was varied between 50 and 500μs to accommodate the change of the radiation intensity. The imaging intervals used were 2°CA and 5°CA for rich and lean operations respectively. The imaging duration was 100°CA throughout the whole experiment. Cylinder pressure and image sequences were recorded in a sychronised fashion for three fuels, namely the Shell component fuel, iso-octane and toluene.

Because the major component of the Shell component fuel was iso-octane and both the component fuel and iso-octane did not contain aromatics, the flame radiation was chemiluminescence-dominated even under rich operating conditions. So the images taken in these situations were not suitable for pyrometry analysis. However, the images taken with toluene were distinct and best suited for demonstrating the analyses and techniques discussed in this chapter. Therefore, discussions were concentrated on the results of toluene in the following sections. Further details on the camera setup and the timing control box can be found in [43].

The phasing of the pulse trains for controlling the Phantom camera together with the cylinder pressure is given in Figure 6.1. The zoomed-in plot on the right hand side clearly shows the number of the imaging pulses.

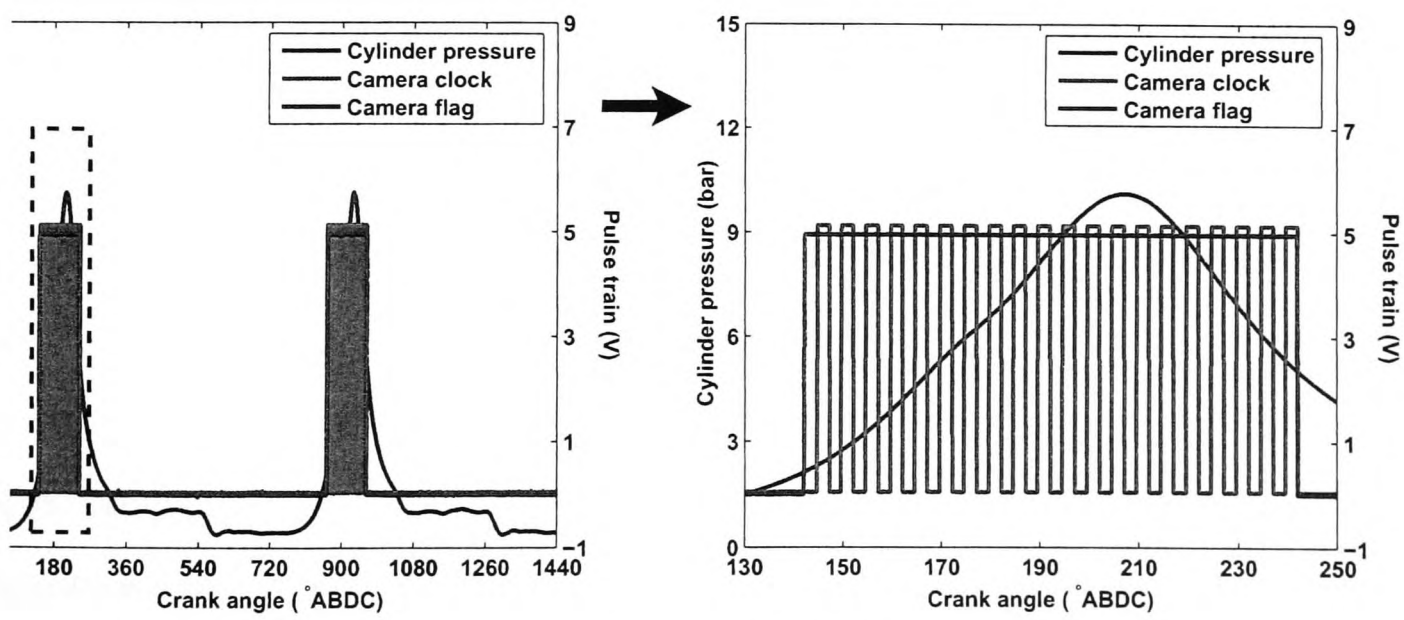


Figure 6.1: The control pulse trains for the high-speed imaging

All the data were recorded at an engine speed of 1500RPM, an ignition timing of 35° before TDC, an injection timing of 280° before TDC and an inlet plenum pressure of -0.48bar (gauge). Figure 6.2 lists some sample combustion images from the three fuels under rich and stoichiometric operating conditions.

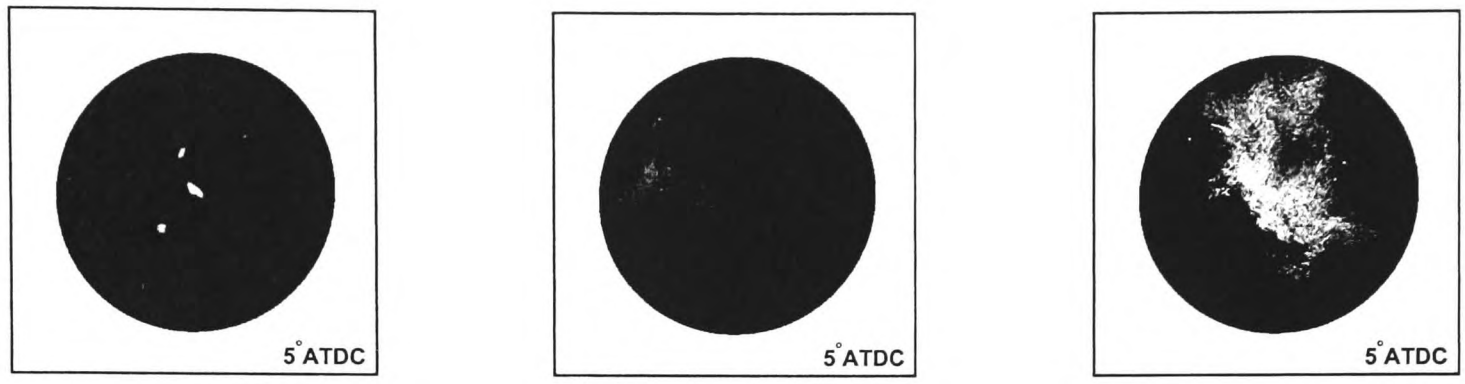


Figure 6.2: Sample combustion images with different fuels under rich ($\lambda = 0.8$, top row) and stoichiometric (bottom row) operating conditions — iso-octane (left), Shell component fuel (middle) and toluene (right)

It should be noted that when the flame radiation intensity was high, a short camera exposure time was used, but this tended to emphasise sooty flame rather than chemiluminescence. For instance, in the top right image, the yellow flame was so intense that hardly any blue chemiluminescent flame could be seen. In general, iso-octane showed well pre-mixed combustion even when the mixture was rich. However, using a lambda of 0.8 with iso-octane gave more evidence of droplet burning *i.e.* mixture stratification. Images with the Shell component fuel were generally the darkest taking account of the exposure effect. Since the component fuel comprised six components of various volatilities,

the air/fuel mixing was not as good as with iso-octane. Toluene showed a considerable amount of diffusion combustion at stoichiometry; nevertheless, the intensity level of the diffusion flame was similar to that from the surrounding chemiluminescence in this case.

6.3 High-speed image sequence vs. burn rate

Three particular engine cycles (combustion event A, B and C) were extracted from three cine sets (cine 1, 2 and 3) to illustrate the comparisons between the high-speed image sequences and the corresponding MFB. The engine operating conditions, the statistics of the three cines from CoBRA and the combustion analysis results for the three combustion events are shown in Table 6.1 and 6.2.

In Table 6.1, cine no. 1 and 2 were both stoichiometric operation, the former with standard early injection (homogeneous mode) and the latter with late injection in the compression stroke. It can be seen that cine 2 presented lower indicated power and $p_{\max}$, a slower burn (especially a longer initial period of combustion *i.e.* 0–10% MFB), and larger cycle-by-cycle variations. Because the orientation of the spark plug (hence the gap) was not optimised for stratified operation, operation with late injection could suffer from fuel impingement onto the back of the spark plug earth electrode. This could lead to undesirable fuel/air mixing and unstable combustion. Cine 3 was rich operation ($\lambda = 0.8$) with standard injection timing and this generally showed the highest indicated power and $p_{\max}$, fastest burn and least cycle-by-cycle variations.

It should be noted that 10%–80% MFB was selected in this analysis for this particular engine instead of the widely used 10%–90% MFB, because the optical engine had larger crevices (polymeric piston rings, no lubricant oil, differences in thermal expansion between the engine block and piston *etc.*) than normal production engines and there could be more trapped fuel in the crevices which might be burned quite late during the expansion stroke. Therefore, 10%–90% MFB would not be representative of the main combustion period.

Cine no.	λ	Injection (°BTDC)	Statistics	IMEP (bar)	p_{max} (bar)	0–10% MFB (°CA)	10%–80% MFB (°CA)
1	1.0	280	mean	2.19	9.82	46.4	43.6
2	1.0	100		1.72	8.92	51.2	44.3
3	0.8	280		2.93	12.05	40.6	24.8
1	1.0	280	CoV (%)	6.5	20.5	12.4	18.9
2	1.0	100		22.0	29.3	25.1	23.9
3	0.8	280		3.5	14.9	7.5	11.6

Table 6.1: Engine operating conditions and statistics of the three cines

Table 6.2 together with Figure 6.3, 6.4 and 6.5 show the image sequences and combustion analysis results from the three combustion events, A, B and C. One common feature for these three examples was that they all showed a higher p_{max} than the average of their cines.

Combustion event	Cine no.	Cycle no.	IMEP (bar)	p_{max} (bar)	0–10% MFB (°CA)	10%–80% MFB (°CA)
A	1	37	2.18	14.82	33.6	35.4
B	2	21	1.80	15.12	31.6	21.8
C	3	25	2.28	14.67	37.2	63.4

Table 6.2: Combustion analysis results of the three combustion events

In Figure 6.3, because the mixture was stoichiometric and well premixed, the flame radiation was chemiluminescence-dominated except that at the very beginning of the combustion period, a glaring trace followed the spark discharge in the vicinity of the plug. The fact that the blue flames were not visible any more after 70% MFB was probably either because the radiation intensity was too low or the flames went out of the field of view, not that the combustion became unnoticeably weak. Combustion chemiluminescence due to chemical reactions is arguably a mark of the flame front. The flame front propagation of this sequence, from the animation, was regular in shape and followed the radial direction as opposed to the irregular swirling of sooty flames observed in event C.

Combustion event B was particularly a fast burning cycle with exceptionally high p_{max}

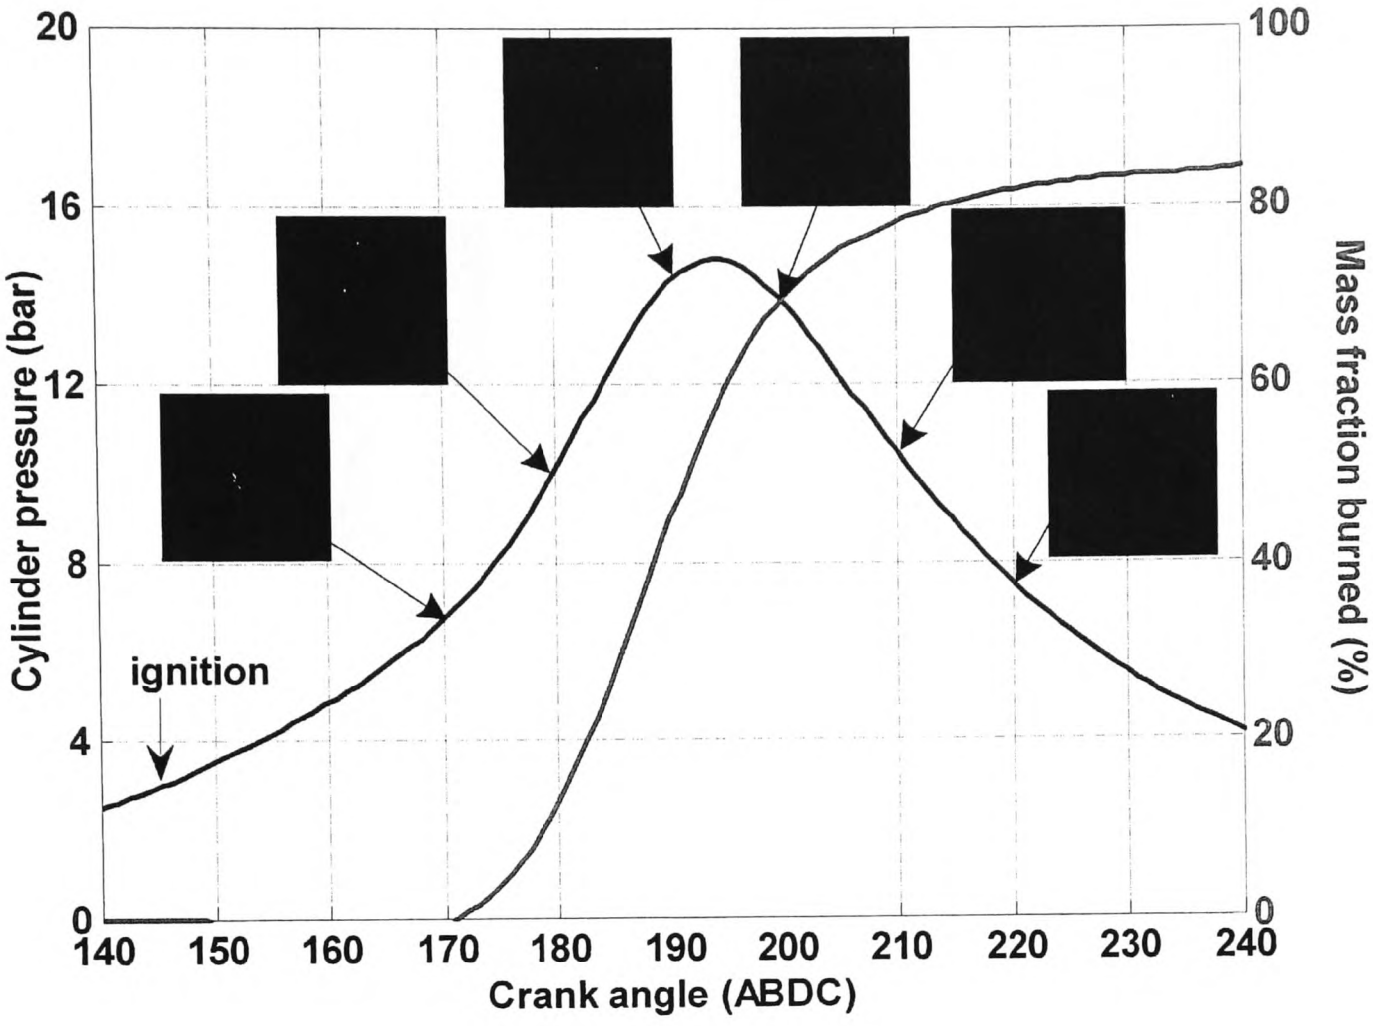


Figure 6.3: Image sequence vs. cylinder pressure and MFB for combustion event A

among the cycles in cine 2. The strength of incandescence quickly developed in the early stage of combustion and diffusion flames (sooty) remained dominant in the field of view in terms of visible inflamed area. Although this was a stoichiometric run, significantly more glowing soot was generated during the combustion period than in combustion event A due to the late injection which might have led to mixture inhomogeneity and droplet burning to some extent. Burning of 'blobs' near the injector in the late stage of combustion (after 210°ABDC, ~82% MFB) was observed, which might be the secondary combustion of a small amount of mixture previously too rich locally to ignite. With the 80% MFB defined as the end of the main combustion, it is worth noting that visible combustion could still last beyond this point which could be from evaporation of the fuel film formed earlier in the cycle.

Event C was the slowest burning cycle of all the three examples, although cine 3 on average

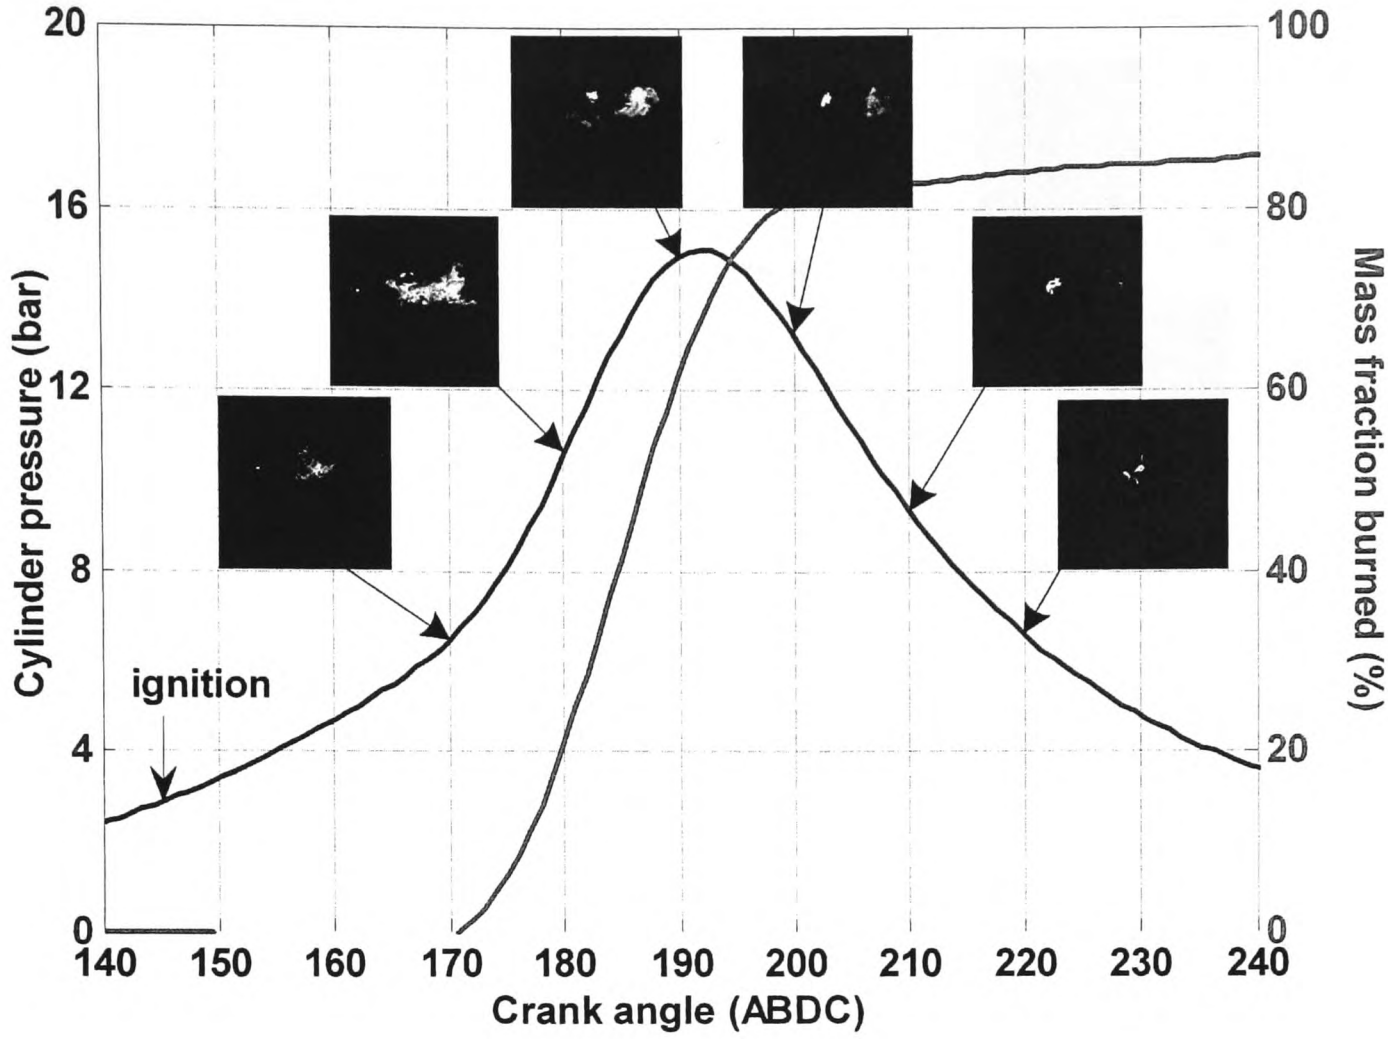


Figure 6.4: Image sequence vs. cylinder pressure and MFB for combustion event B

was the fastest burning one. In particular, the 10%–80% MFB duration of event C was so long that it went beyond the imaging duration for this cycle. However, the flames were invisible after 210°ABDC. This could either mean heavy reduction in the radiation intensity or the flames moving outside the field of view (the field of view from the piston window covered about 50% of the combustion chamber area). In any case, care must be taken when making correlations between visible inflamed area and the corresponding mass fraction burned data.

6.4 Flame growth and propagation

Extensive research has been undertaken in the past on the behaviour of turbulent flames with various techniques *e.g.* OH-PLIF, CFD and chemical kinetics simulations. Under-

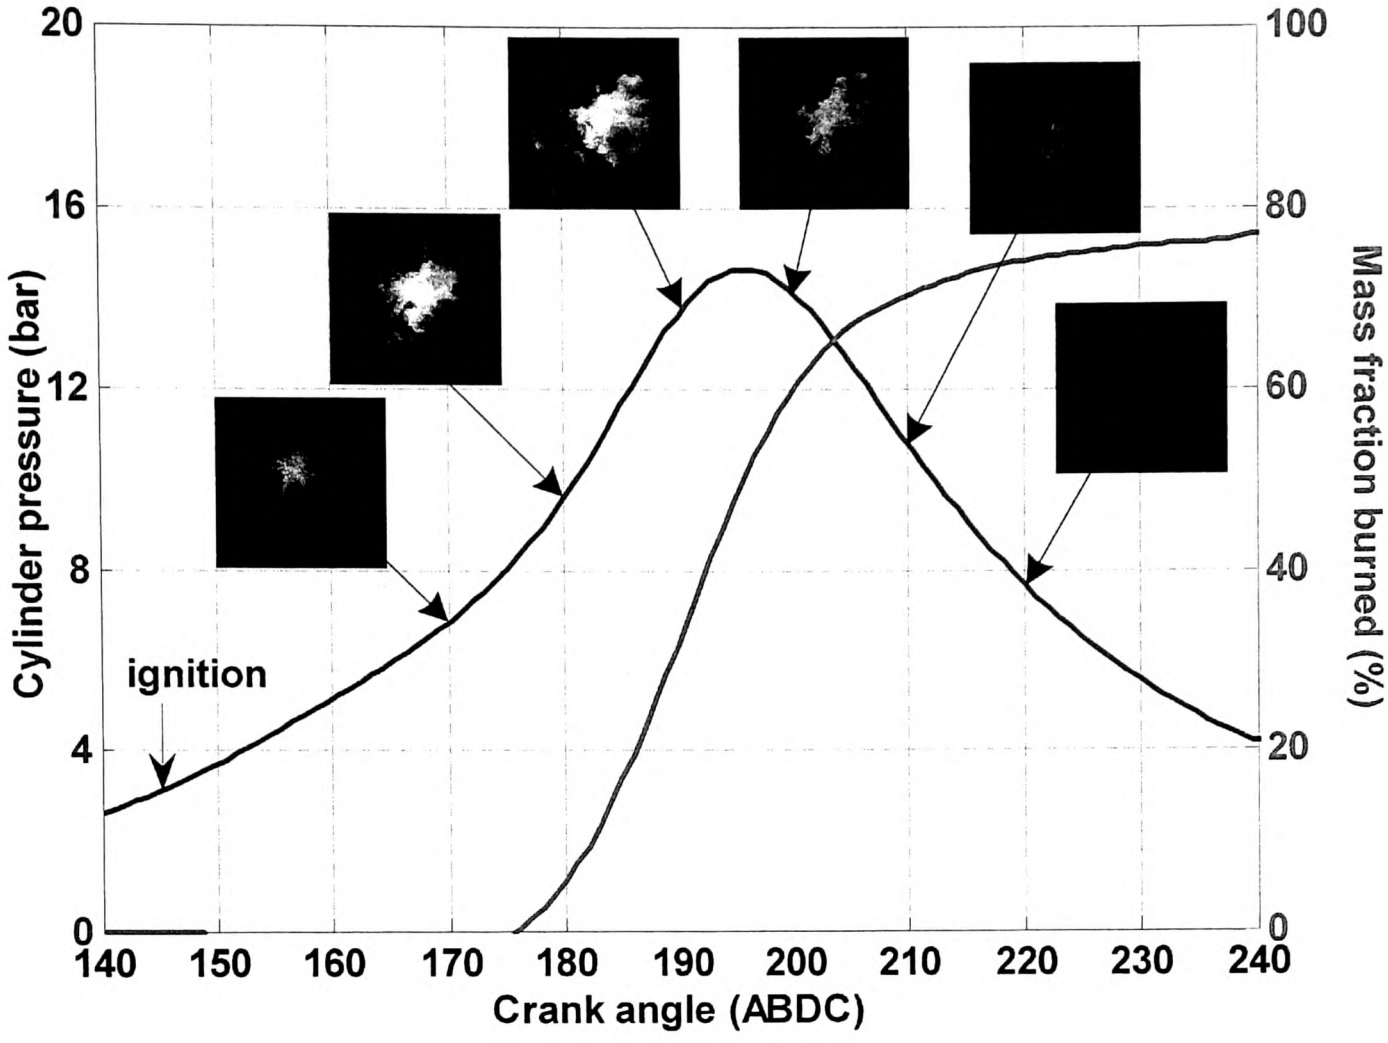


Figure 6.5: Image sequence vs. cylinder pressure and MFB for combustion event C

standing the stability of such flames is of both fundamental and practical significance for IC engine workers. However, because of the inherent time and environment-dependent features of turbulence as well as the extreme complexity of the experiments and subsequent data interpretation associated with IC engines, time correlated measurement and advanced data evaluation must be set up and implemented. Despite the line-of-sight effect, direct high-speed combustion imaging can still prove to be a straightforward and cost-effective approach for flame studies.

Abu-Gharbieh *et al.* [1] [2] employed advanced image analysis for OH-PLIF images to track the flame front and derive its velocity field. This provided a framework for the study carried out in this section. The example images used in the following illustration were from the engine operation with toluene and a lambda of 1.2, the other conditions as described in Section 6.2.

6.4.1 Flame front detection

Non-linear diffusion filtering

To improve the SNR in the flame image sequences and avoid picking up redundant or spurious flame edges, the images were smoothed using nonlinear diffusion filtering before applying an edge detecting algorithm.

Non linear diffusion filtering is a very powerful image processing tool for image noise reduction and image simplification. The method is based on the original approach established by Perona and Malik [46]. The principle is to smooth out noise locally by diffusive flow while at the same time prevent flow across object boundaries (edges).

The idea can be expressed with an analogy of physical diffusion processes such as the temperature diffusion in a metal bar or the diffusion between two fluids mixed together. These physical diffusion processes are modelled by the following differential equation

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \mathrm{div}\,(D \cdot \nabla I) \tag{6.1}$$

where I is the concentration (temperature) and D is the diffusivity (thermal diffusivity). Assuming constant diffusivity, which is true for the hot metal bar example and many others, the diffusion will be faster where the concentration gradient is higher. This phenomenon is known as linear diffusion.

The solution of the linear diffusion equation with a scalar diffusivity $D = d$ turns out to be the same operation as convolving the concentration image I with a Gaussian kernel of width $\sqrt{2t}$. However, for many image processing applications, it is required to preserve image borders (which are, by definition, high gradient regions). Using a linear diffusion would quickly damp out the gradient and destroy the borders. To solve this problem, nonlinear diffusion was proposed to exchange the scalar diffusion constant d with a scalar valued function g of the concentration gradient of the grey levels in the image, with g

decreasing for high gradients. The diffusion equation then becomes

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \mathrm{div}\left(g\left(|\nabla I|\right)\nabla I\right) \tag{6.2}$$

This new formula allows for image smoothing while preserving the borders. In this work, the imaged data were filtered using the equation

$$\frac{\mathrm{d}I}{\mathrm{d}t} = \mathrm{div}\left(g\left(|\nabla\left(G_\sigma * I\right)|\right)\nabla I\right) \tag{6.3}$$

where I represents the intensity of the image under consideration, and $g\left(|\nabla\left(G_\sigma * I\right)|\right)$ represents a locally adaptive diffusive strength. The latter is made inversely proportional to the gradient ∇I in the image itself after smoothing with a Gaussian kernel G_σ of width σ, which is done for stability reasons [2]. The diffusivity function used in this work was

$$g\left(s\right) = \begin{cases} 1 - \exp\left(\frac{-C_m}{(s/\lambda)^m}\right) & g > 0 \\ 1 & g \leq 0 \end{cases} \tag{6.4}$$

where m is a positive constant defining the speed of the diffusivity change; $\lambda > 0$ acts as a contrast parameter separating regions of forward diffusion from regions of backward diffusion; the positive constant C_m determines the direction of the flux function and is calculated to make the flux $g\left(|\nabla I|\right)$ ascending for $s < \lambda$ and descending for $s \geq \lambda$.

By a proper choice of the diffusion kernel G_σ, object boundaries may be enhanced and physical gradients sharpened hence simplifying subsequent location of object boundaries. Figure 6.6 shows the effect of non-linear smoothing on an example flame image.

Edge detection

The MATLAB edge detecting functions were applied on the filtered flame images to track the flame front. The routine took intensity images as input and returned binary images

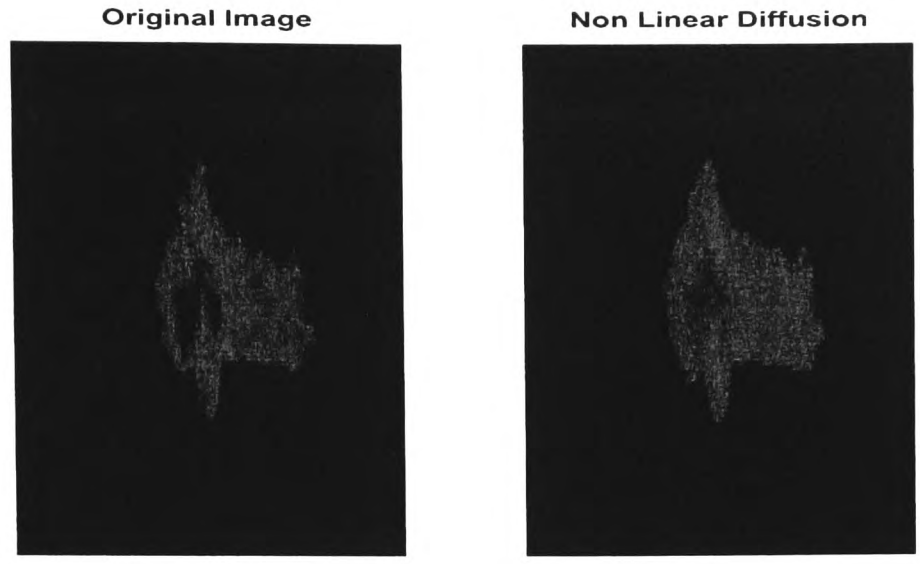


Figure 6.6: Effect of smoothing an image using non-linear diffusion filtering

of the same size as the input images, with 1's where the function found edges and 0's elsewhere. The Sobel approximation (Sobel edge detector) was adopted for the derivative calculation which returned edges at those points where the gradient of I was a maximum.

Figure 6.7 presents the results of flame front detection from four sample images along with the overlaid boundary posed by the field of view. These images were from the same combustion event under lean homogeneous conditions and thus represented turbulent premixed flames. Also, the blue chemiluminescence indicated where the flame front occurred because it coincided with the process of chemical reactions where the flame had propagated. Occasionally, the edge detection routine generated 'holes' within the detected boundary or 'spots' outside it. These were due to extremely bright pixel clusters such as a cloud of soot. Median filtering and region filling were then applied in MATLAB to remove these false features and make the boundary a simple closed contour. The median filtering process also helped to smooth the edge and eliminate redundant shape features.

Temporal contour interpolation

The detected flame fronts were a series of 2-D contours. The change in their shape could reflect the flame propagation in time. However, the time interval between two images was

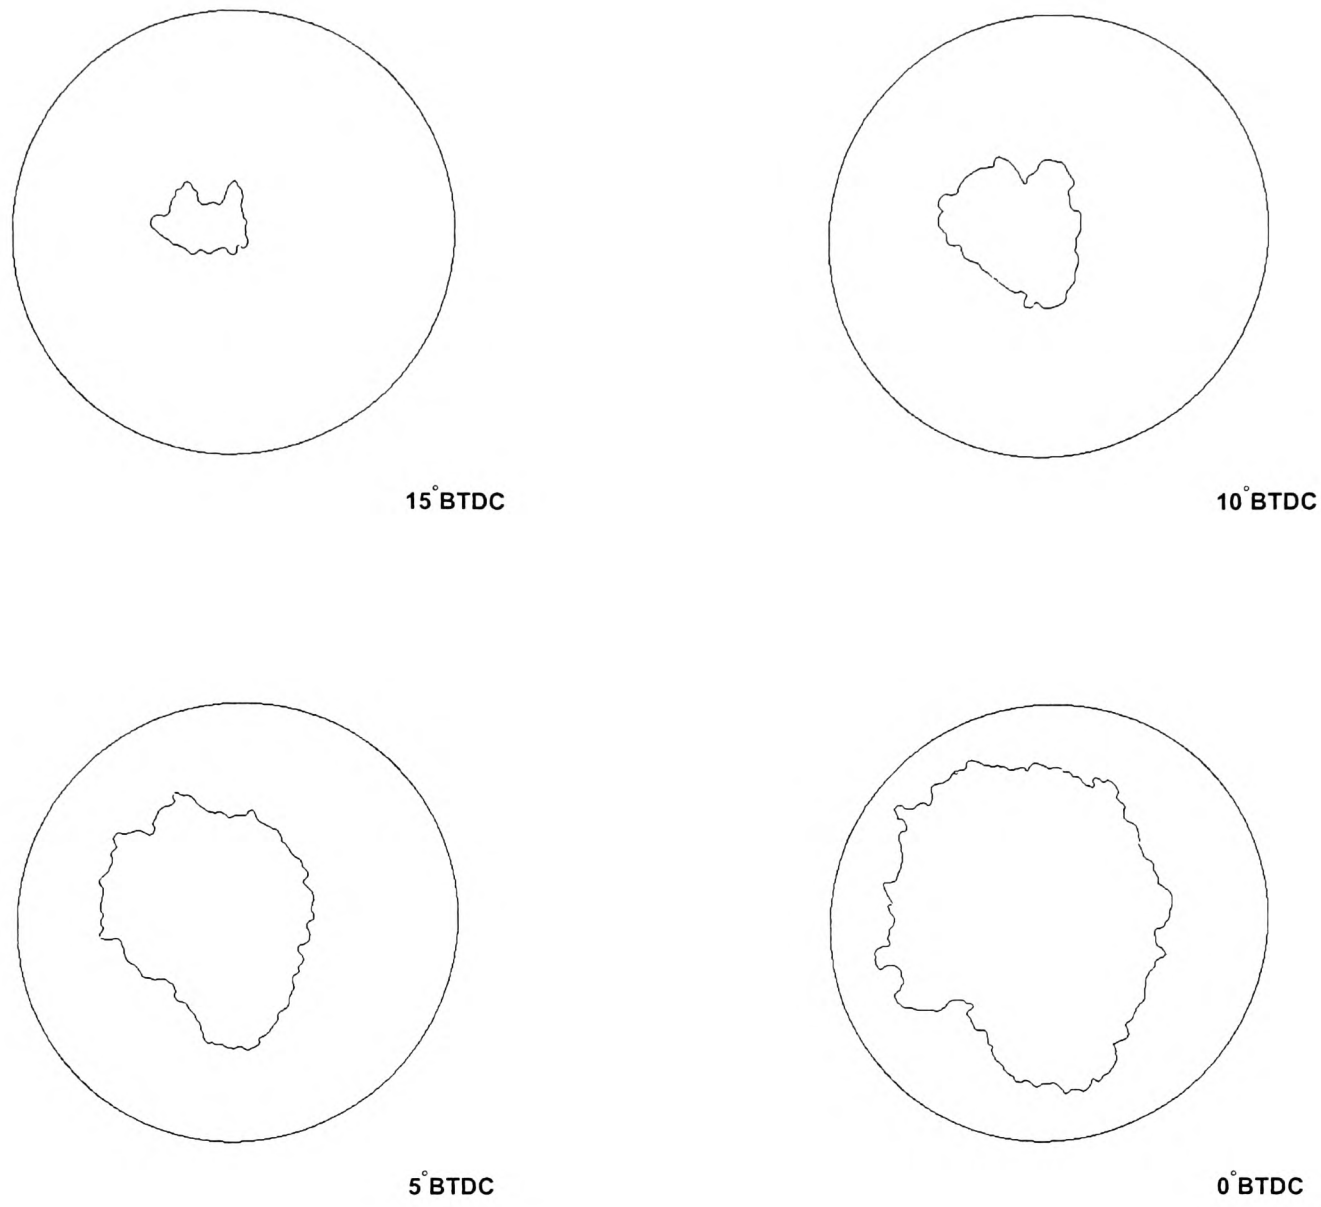


Figure 6.7: Detected flame fronts from sample images

5°CA and hardly sufficient for resolving the temporal information *e.g.* the local flame front velocity. Therefore, morphological interpolation must be undertaken among the contours.

Interpolation in the frequency domain with the 1-D discrete cosine transform (DCT) was employed for the following reasons [1]:

- it is straightforward to implement in MATLAB
- it produces real coefficients
- it has excellent energy compaction properties

- it is the coefficients that correspond with each other (frequency component correspondence), not the coordinates of the pixels (point-to-point correspondence), so no node-to-node match between different contours was required

For each frame j there was a flame front contour $v(n,j) = (x(n,j), y(n,j)), n = 1, 2, \ldots, N$ for $j = 1, 2, \ldots, F$ where F is the number of frames ($F = 4$ in this demonstration). The 1-D DCT of the sequence of $x(n,j)$ contour coordinates (and similarly for the $y(n,j)$ coordinates), $n = 1, 2, \ldots, N$, was defined as the following [1]:

$$X(k,j) = \omega(k) \sum_{n=1}^{N} x(n,j) \cos \frac{\pi (2n-1)(k-1)}{2N} \tag{6.5}$$

where

$$\omega(k) = \begin{cases} \sqrt{\frac{1}{N}} k = 1 \\ \sqrt{\frac{2}{N}} 2 \leq k \leq N \end{cases} \tag{6.6}$$

and $k = 1, 2, \ldots, N$. Then direct interpolation between corresponding frequency coefficients could be performed. Finally the inverse discrete cosine transform was used to transform the interpolated components back into the spatial domain

$$x(n,j') = \omega(k) \sum_{k=1}^{N} X(k,j') \cos \frac{\pi (2n-1)(k-1)}{2N} \tag{6.7}$$

where $n = 1, 2, \ldots, N$ and j' contained the interpolated frames including the original ones.

In practice, the detected flame front contours did not necessarily have the same number of points. This would cause difficulty with the interpolation programming in MATLAB. To ensure that all contours had the same number of points for the interpolation, first the number of points on the smallest contour was set to be the target for the other contours. A routine based on a critical point detection (CPD) algorithm for a simple 2-D closed contour [79] was then used to reduce the number of boundary points whilst maintaining

the contour shape as far as possible. This CPD algorithm does not require curvature approximation or Gaussian filtering (contour smoothing). Instead, a set of pseudocritical points which represent the major features of the contour shape is initially selected from the boundary. This is carried out in the polar coordinate system where ρ and θ have local minima or maxima. The critical level of each point in this pseudocritical point set, defined by the area of the triangle formed by this pseudocritical point and its two neighbours, can be calculated. The points with the lowest critical level are then deleted. Critical levels of only the neighbouring points to the deleted points need to be updated. The routine was run recursively until the number of points of the contour reached the target. Example results of the CPD are shown in Figure 6.8 with the original contour being a solid line and the detected critical points being the red dots.

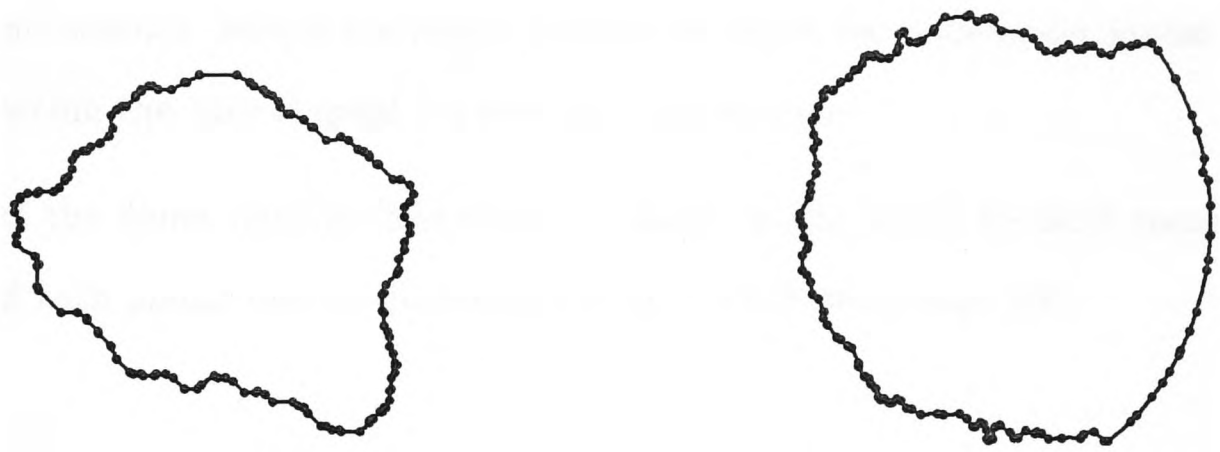


Figure 6.8: Sample results of critical point detection — original contour (solid line), detected critical points (red dots)

Figure 6.9 presents the results of applying the temporal contour interpolation on the flame fronts shown in Figure 6.7.

Flame front velocity estimation

The temporally interpolated contours could be used to estimate the local flame front velocities. The direction of the velocity vector at each node on a particular contour was

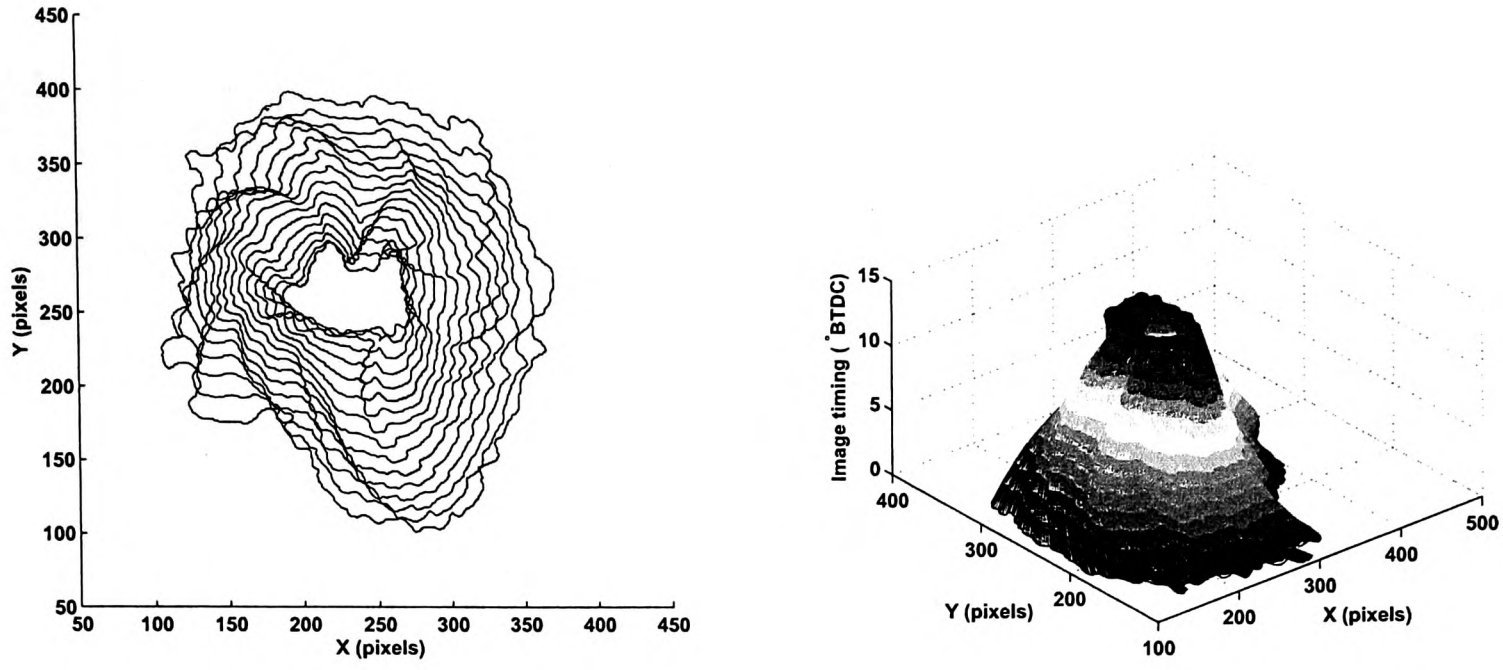


Figure 6.9: Sample results of temporal contour interpolation

locally normal to the contour curve. The intersection point of this normal vector with the subsequent contour defined the target location to which the node under investigation had moved within the time elapsed between the two contours.

Estimation of the flame velocity field based on these results would be more meaningful when applied with planar optical diagnostics such as PLIF involving OH*.

6.4.2 Results with Perspex piston

Images presented in the previous sections had a restricted field of view caused by the quartz piston window. This has introduced uncertainties to the results and discussion in Section 6.3 as well as limitation on the application of flame front detection. Therefore, a Perspex (tradename for polymethyl methacrylate) piston (Figure 6.10) was designed and made at Oxford to allow for a wider imaging field and further flame analysis. The inner surface of the piston was made with a spherical concave surface so as to give an enlarged field of view fully covering the whole combustion chamber.

The design of the concave surface is illustrated in the schematic in Figure 6.10. The

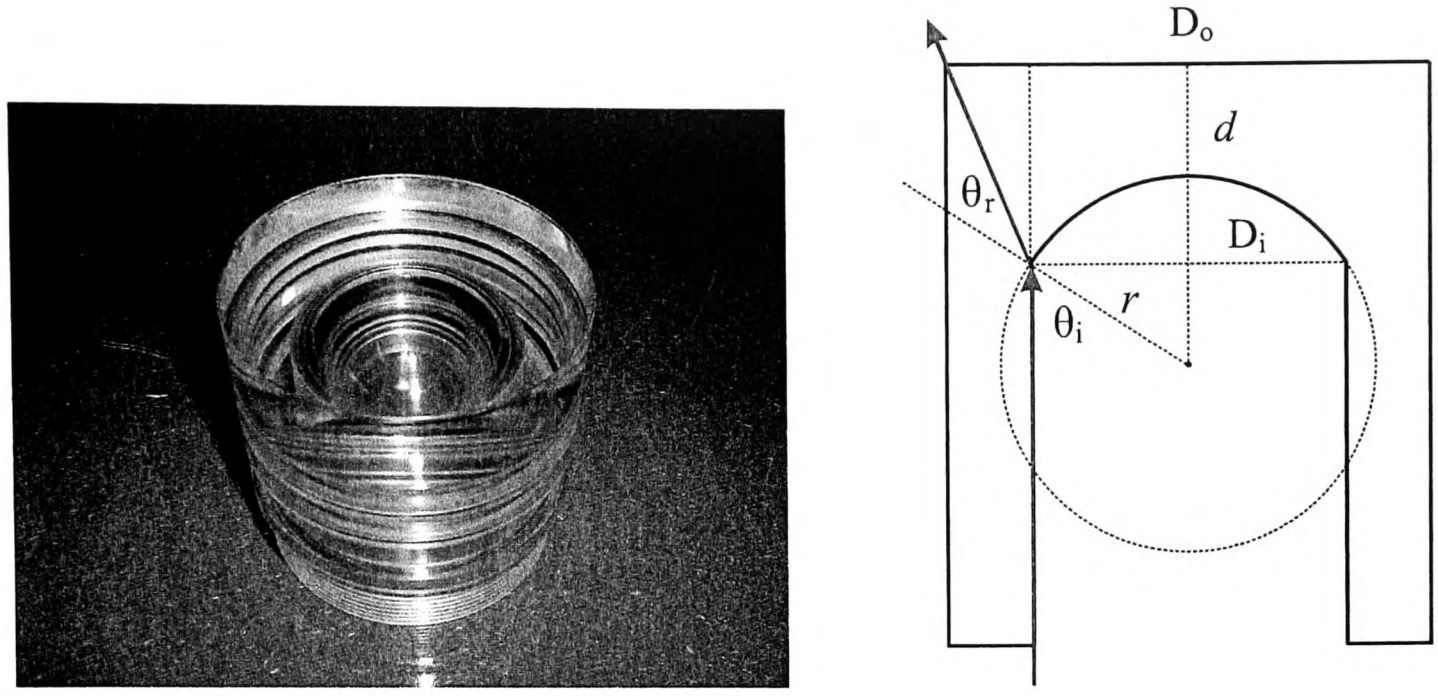


Figure 6.10: The Perspex piston made at Oxford (left) and a schematic of its design (right)

incident and refracted light (red arrows) represents the extreme case when a camera underneath with a collimating lens can see the whole combustion chamber through the piston. From the geometry of the piston's cross section:

$$\tan(\theta_i - \theta_r) = \frac{D_o - D_i}{2\left[r\left(1 - \cos\theta_i\right) + d\right]} \tag{6.8}$$

$$\sin\theta_i = \frac{D_i}{2r} \tag{6.9}$$

where r is the radius of the spherical concave, and θ_i and θ_r are the incident and refracted angles, respectively. D_i and D_o are the inner and outer diameters of the Perspex piston, respectively, and d is the distance between the top and bottom surfaces along the piston axis. According to Snell's Law of Refraction:

$$\sin\theta_i = n\sin\theta_r \tag{6.10}$$

where n is the ratio of refractive indices of Perspex and air. Here $n = \frac{1.492}{1.005}$.

The design values of D_i, D_o and d were 58mm, 88.4mm and 20mm respectively. Eq. 6.8,

6.9 and 6.10 were then solved to give the value $r = 34.4$mm.

Tests under similar operating conditions to Section 6.2 were undertaken with the Perspex piston except that only ~30 engine cycles were recorded in each test due to the piston's vulnerability to thermal damage and sooting. Extreme care was taken so that the piston crown was cleaned when heavily sooted and polished or machined after melting.

Before any of the image analysis described in Section 6.4 could be applied, images taken using the Perspex piston needed to be corrected for optical distortion because of the concave inner surface. Thus a grid of evenly distributed circles, with the largest one having the same diameter as the engine bore, was fixed in the plane of the cylinder head base and imaged before and after the piston was put in place. Since the degree of distortion also changed with crank position, images were taken at 0, 30, 60, 90 and 120°CA ATDC to cover the whole combustion period. The correction was carried out using control point selection and LWM spatial transformation (see Section 2.4.5 for details). Photos of the original, distorted and corrected grid are shown in Figure 6.11. All corrected grid images showed an accuracy in circle diameter of ~1%. Because the distance between the grid and the piston crown changed in a sinusoidal manner *i.e.* the piston accelerated as it moved away from TDC, cubic interpolation rather than linear was used on the coordinates of the selected control points at known crank angles to derive the transformations for other crank positions within 0–120°CA. The largest circle in the corrected image was then used as the overlay of the combustion chamber on the flame images. To enclose the whole restored grid, the image border was extended to 713 by 713 pixels from 512 by 512.

It should be noted that the Perspex piston acted as a lens. Since the focal plane (camera plus piston) was set to be the grid plane when the piston was at TDC, the grid image became increasingly out of focus as the piston moved away from TDC. This effect could not be quantitatively corrected for because in practice the focusing of the camera was fixed during an engine cycle. However, the main combustion period is typically 30°CA each side of TDC, within which the piston was close enough to TDC and hence the focusing

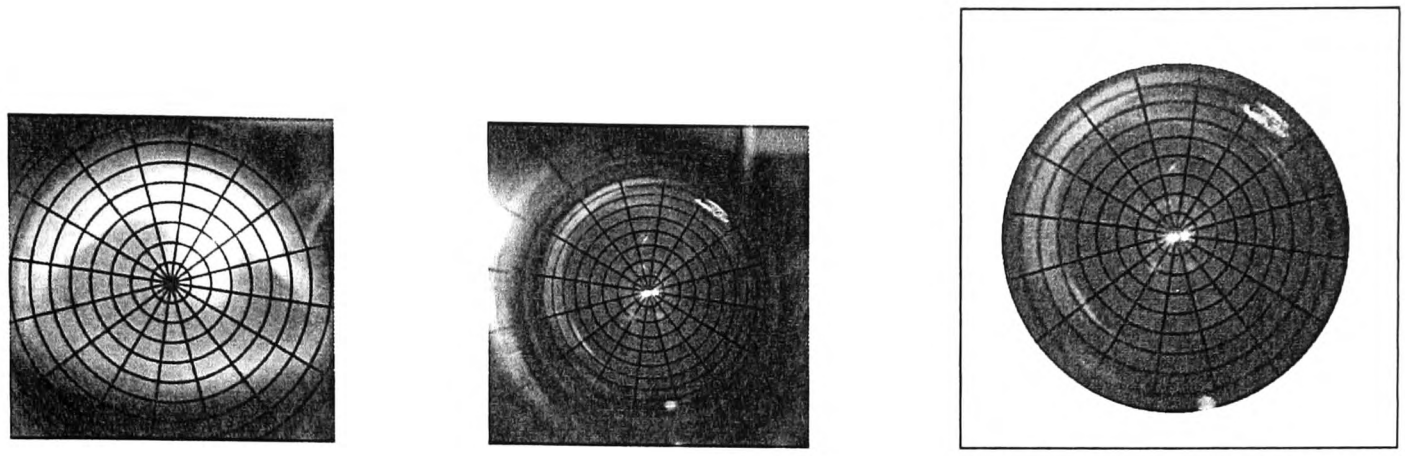


Figure 6.11: Distortion correction for the Perspex piston at 30°ATDC — original grid (left), distorted grid (middle), distortion corrected (right); all images amplified by a factor of 3 for better visualisation

was acceptable.

Figure 6.12 illustrates the effect of the distortion correction on a combustion image taken at TDC. More sample combustion images after distortion correction are shown in Figure 6.13 taken with different mixture strength. The images with a lambda of 1.2 showed a significant amount of droplet burning, which probably resulted from the engine not being fully warmed up, because the Perspex piston could not have lasted until a steady state.



Figure 6.12: Distortion correction on a combustion image taken at TDC with toluene and a lambda of 1.2 — original image(left), distortion corrected (middle), combustion chamber overlaid (right)

Following the image analysis procedure described in Section 6.4, flame images taken using the Perspex piston were processed. Illustrative results shown here were with a lambda of 1.2 to emphasise the chemiluminescent flame which is representative of the

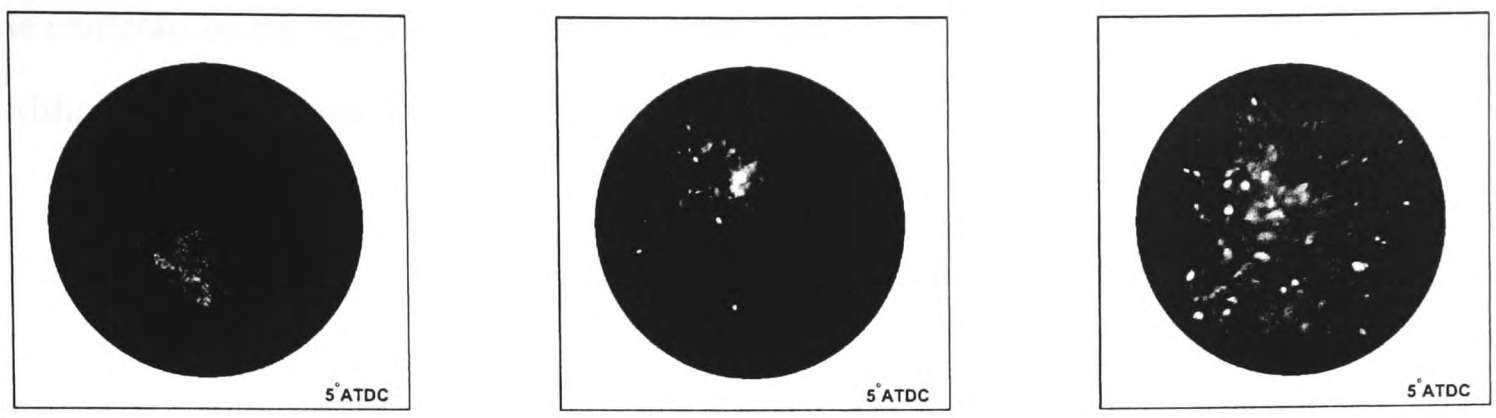


Figure 6.13: Sample combustion images with distortion correction taken at 5°ATDC with toluene and different mixture strength — $\lambda = 0.8$ (left), $\lambda = 1.0$ (middle), $\lambda = 1.2$ (right)

flame front. Figure 6.14 presents all the original flame fronts detected in one engine cycle from 25°BTDC to 15°ATDC using the image processing procedure introduced above.

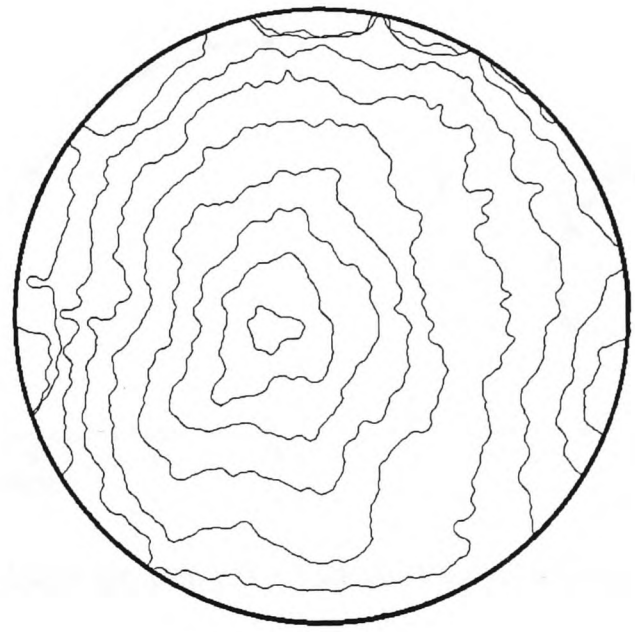


Figure 6.14: Flame fronts from 25°BTDC to 15°ATDC in steps of 5°CA

With such a set of contours, some flame properties can be studied, such as the mean flame radius and the flame centroid locus. The mean flame radius is defined as the equivalent radius of a circle that has the same area as the enflamed area of a flame image. The change of the mean flame radius in time indicates the mean growth rate of the flame. The

flame centroid locus, on the other hand, shows how the flame kernel moved in space. The coordinates of the flame centroid are determined by:

$$X_c = \frac{M_{10}}{M_{00}} \qquad Y_c = \frac{M_{01}}{M_{00}} \tag{6.11}$$

where

$$M_{ij} = \sum_{\text{column}} \sum_{\text{row}} X^i Y^j I(X,Y) \tag{6.12}$$

Here $I(X,Y)$ is pixel intensity of the flame image, in particular, the blue channel. Figure 6.15 gives results of these flame properties from the above image sequence. In the right figure, the flame kernel started left of the spark plug probably because the spark plug gap faced left in the picture. It then moved anti-clockwise to the top right of the spark plug.

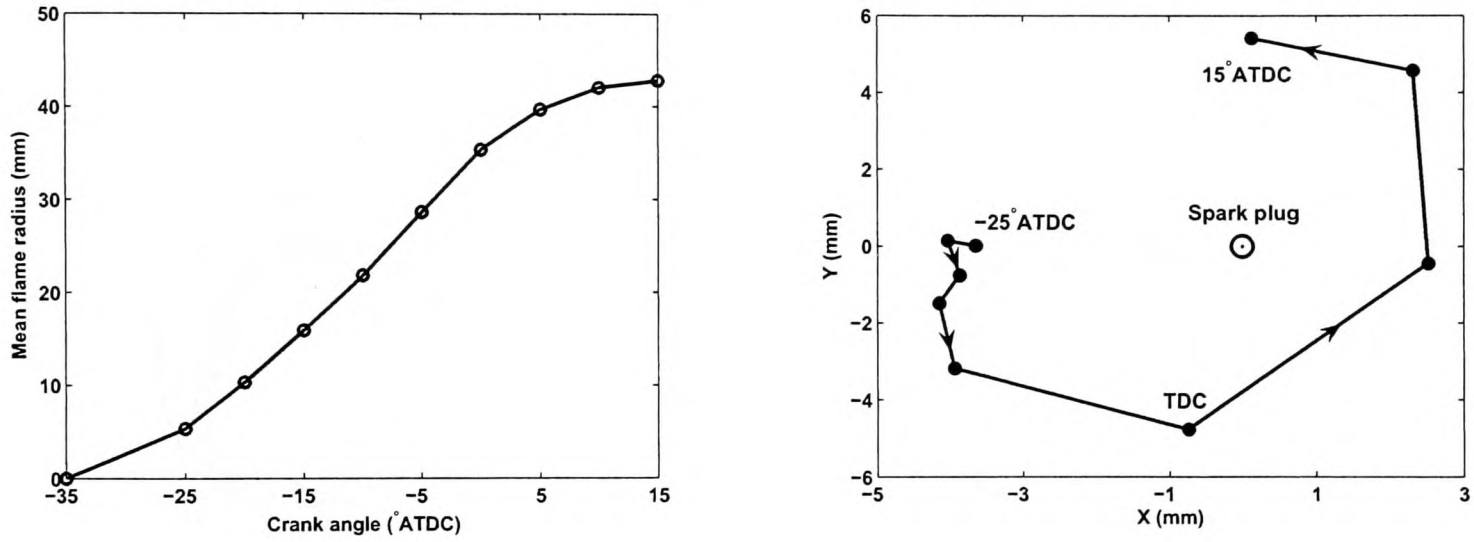


Figure 6.15: Analysis of flame properties — mean flame radius (left) and flame centroid locus (right)

Figure 6.16 lists the results of flame front propagation with the interpolated contours (in blue, every 1°CA) between each pair of consecutive original ones (in red, 5°CA apart). The crank angle range was from 25°BTDC to 15°ATDC from left to right, top to bottom. It can be seen that during the early stage of this engine cycle, the flame growth was strongly affected by the in-cylinder swirl (clockwise); the flame propagated faster to the bottom left than to the other directions.

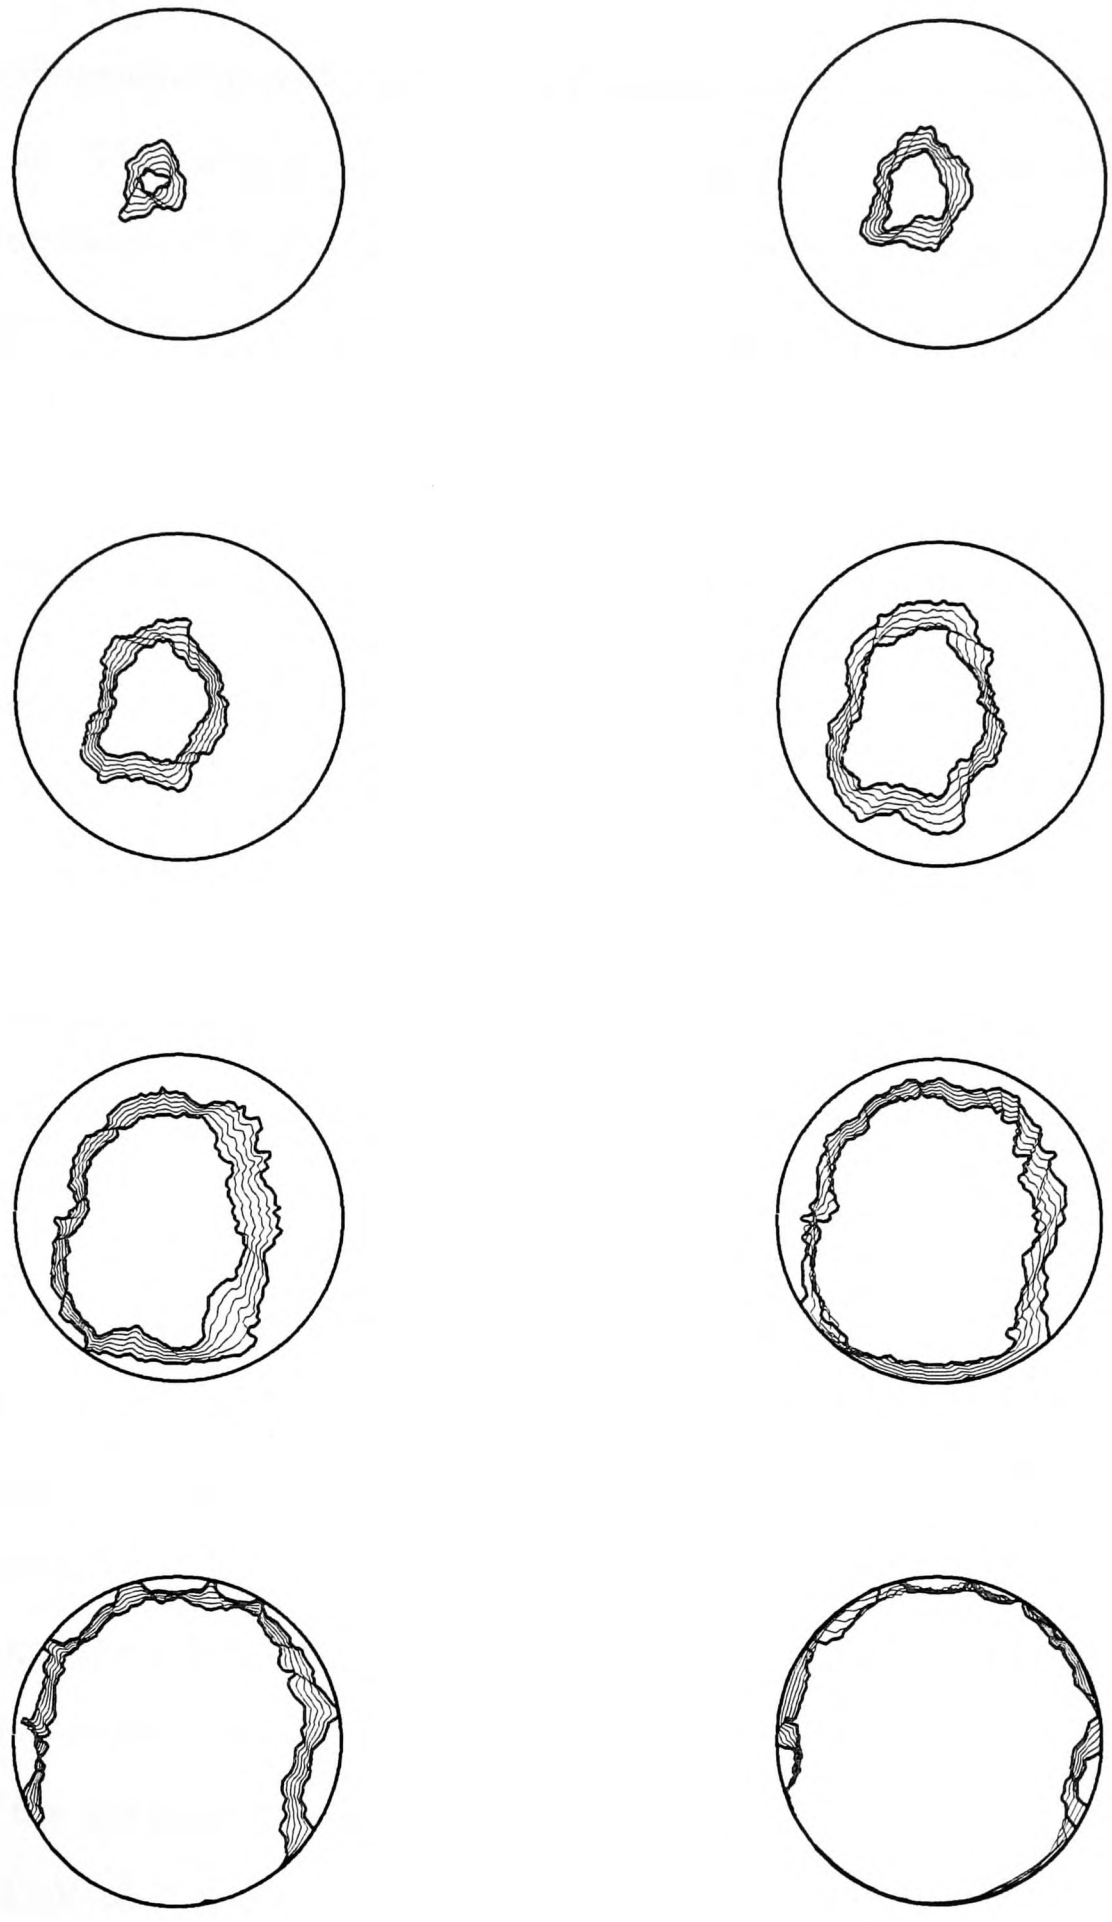


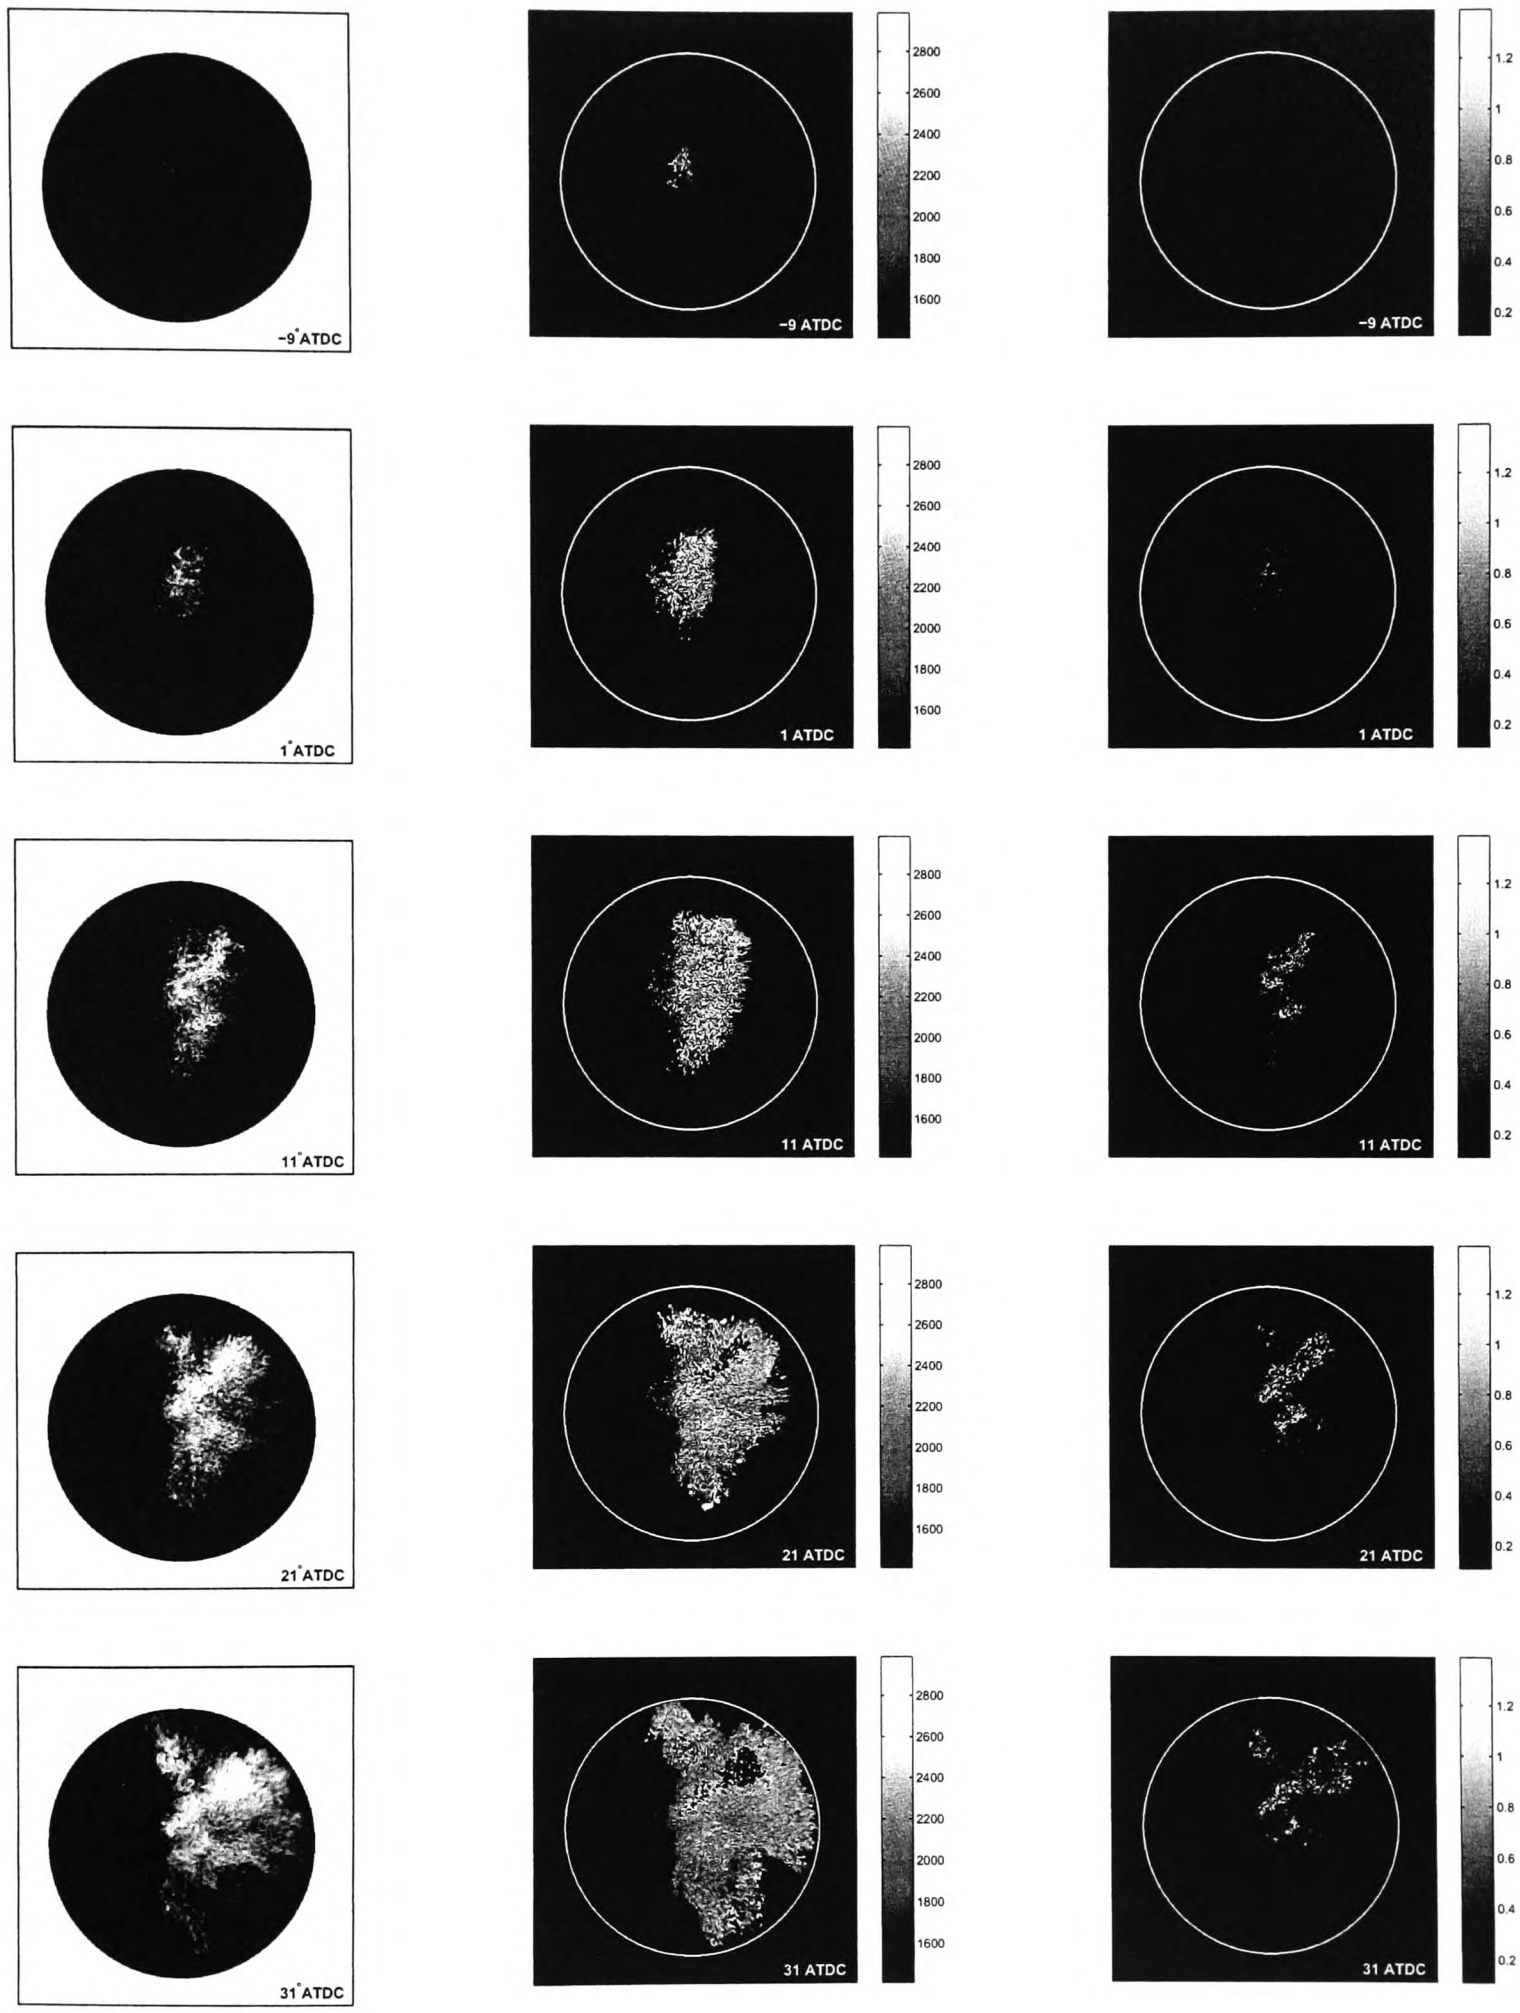
Figure 6.16: Flame front propagation — interpolated flame fronts (in blue) every 1°CA between two consecutive original ones (in red), 89mm bore diameter

6.5 High-speed soot pyrometry

A spectral calibration was performed on the Phantom camera in the same way as described in Chapter 5. The variable-KL strategy was adopted here to give both crank angle resolved temperature and KL results. Figure 6.17 presents the high-speed CRP results from a sample combustion sequence taken with toluene and a lambda of 0.9. The images on the left were corrected in the following sequence: colour re-demosaicing, background subtraction, flat field correction, correction for optical distortion, selection and exclusion of invalid pixels. The invalid pixels contained:

- saturated pixels (the dark 'holes' in the middle of the enflamed area)
- pixels that had a red or blue value below the specified threshold, because such pixels are prone to noise and tend to show great error (for details on error analysis, refer to Section 5.4.5); a dynamic threshold — a quarter of the mean of all red/blue pixel values between 0 and 255 exclusively — was employed so as to account for the different intensity levels of the images taken at different crank angles
- pixels that gave out-of-range R/G or R/B ratios; the T and KL ranges used here were 1400–3000K and 0.1–1.4 respectively

Figure 6.18 shows the averaged crank-angle resolved temperature and $N * KL$ results of the whole image sequence shown in Figure 6.17. First the temperature and soot loading did not peak at the same crank angle but 11°ATDC and 33°ATDC respectively. This suggested that the oxidation of soot did not exceed soot formation until several tens of degrees CA after the maximum soot temperature was achieved. This lag is likely to be a function of the in-cylinder temperature. The CoV of T was ~5% for most crank angles whilst the CoV of $N * KL$ was poor, 41% at the minimum (@ 29°ATDC). It should be kept in mind that the calculation of KL is inherently ill-conditioned, so the CoV of $N * KL$ here did not necessarily reflect the actual cycle-by-cycle variation. The rise of CoV at early and late crank angles was probably due to the low flame intensity and hence


-9°ATDC
-9 ATDC
2800
2600
2400
2200
2000
1800
1600
-9 ATDC
1.2
1
0.8
0.6
0.4
0.2
1°ATDC
1 ATDC
1 ATDC
11°ATDC
11 ATDC
11 ATDC
21°ATDC
21 ATDC
21 ATDC
31°ATDC
31 ATDC
31 ATDC

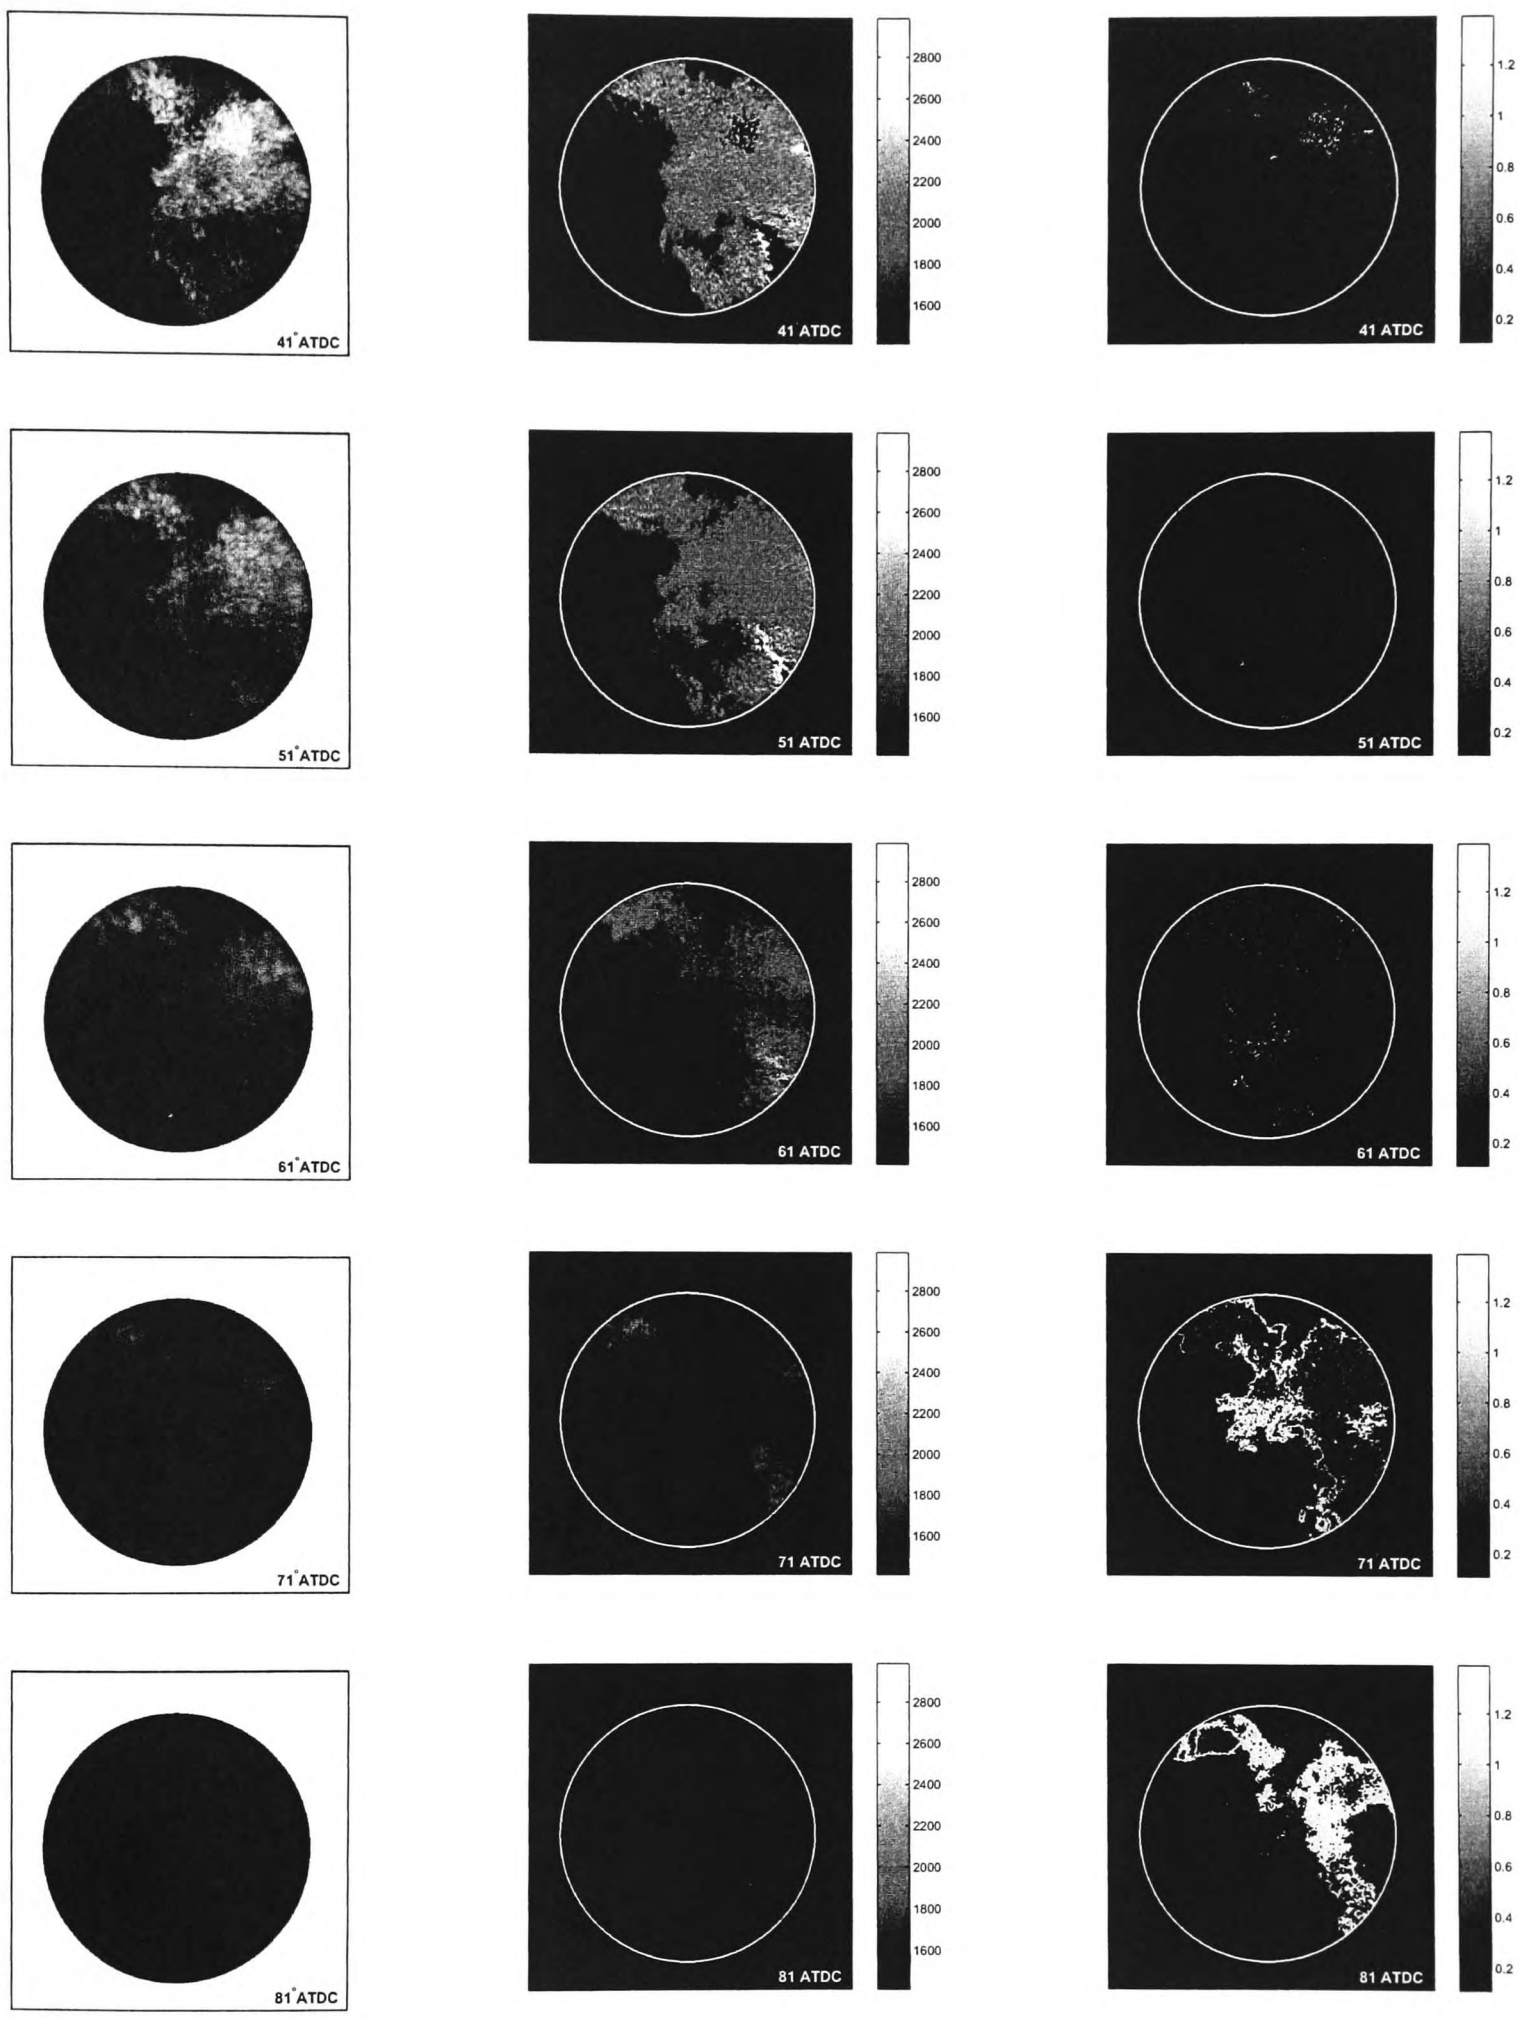


Figure 6.17: Example results of high-speed CRP from a combustion sequence taken with toluene and a lambda of 0.9 — corrected flame image (left), temperature (middle) and KL distribution (right), 89mm bore diameter

could be improved by using a camera with a better dynamic range. The new Phantom v7 series cameras offer options for 36 and 48-bit colour images and would be better suited for high-speed CRP.

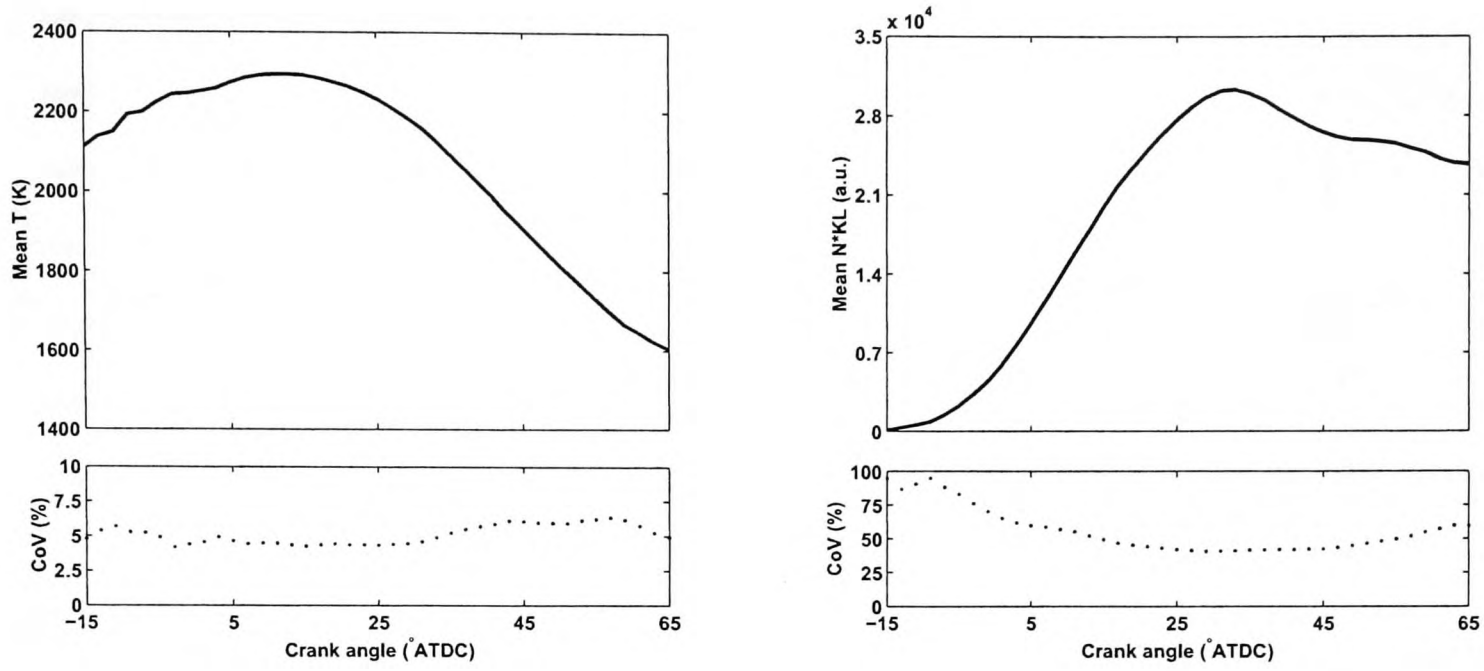


Figure 6.18: Averaged crank-angle resolved results of T (left) and $N * KL$ (right) from an image sequence with toluene and a lambda of 0.9

Since there were more saturated pixels in the middle of the combustion process (typically 0–30°ATDC) than in the later stages, excluding such pixels would bias the mean T and $N * KL$ in the middle CA range as they would have given high T and KL values instead of 'null' if properly captured with a camera of better dynamic range. As a result, $N * KL$ was occasionally abnormally high at late crank angles (comparable to that in the middle range), whilst in fact the soot formed earlier during combustion was expected to be further oxidised and reduced in number concentration. The $N * KL$ solution suffered from low flame intensity as well. Again such errors could be suppressed by using a better camera.

It should be noted that in each of the experiments with the Perspex piston, engine firing was switched on for only 2–3s and hence the engine was in warm-up mode all the time. The cycle-by-cycle variation in this case represented the extreme i.e. transient operation. Such cycle-by-cycle results are shown in Figure 6.19. This could also be a reason why the Phantom T results were noticeably lower than the Pulnix ones shown in Chapter 5 and

the ISIS simulation results in Chapter 7. However, it has been shown in Chapter 5 that the relative accuracy of CRP is no better than 5–10%, therefore such a difference between engine tests and cameras is within tolerance.

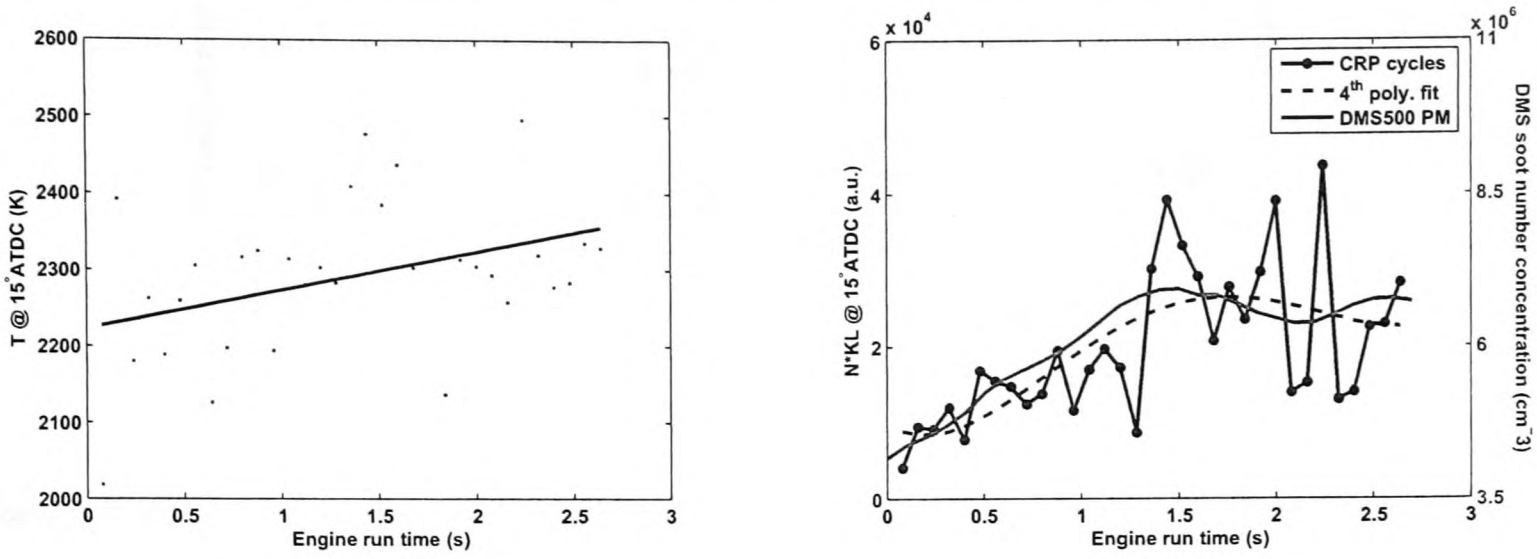


Figure 6.19: Cycle-by-cycle results of T (left) and $N * KL$ (right) from the same image sequence as in Figure 6.18 (toluene, *lambda* = 0.9) with supplementary results from particulate measurement in the exhaust using the DMS500

Some supplementary results from particulate measurements using the DMS500 (sampling rate 10Hz) are also included in Figure 6.19 to illustrate the transient effect. Although the two ordinates have different physical meanings, both give similar trends in terms of soot quantity.

Relations between pyrometry results using variable KL (@ 15°ATDC) and mixture strength is given in Figure 6.20. Lean operation results were not shown here considering the noise due to chemiluminescence.

Sample animations of the high-speed CRP results can be found on www.eng.ox.ac.uk/ice.

6.6 Summary

High-speed combustion imaging was carried out on the DISI engine using a Phantom v7.1 CMOS camera and provided a crank-angle resolved investigation of flame growth and

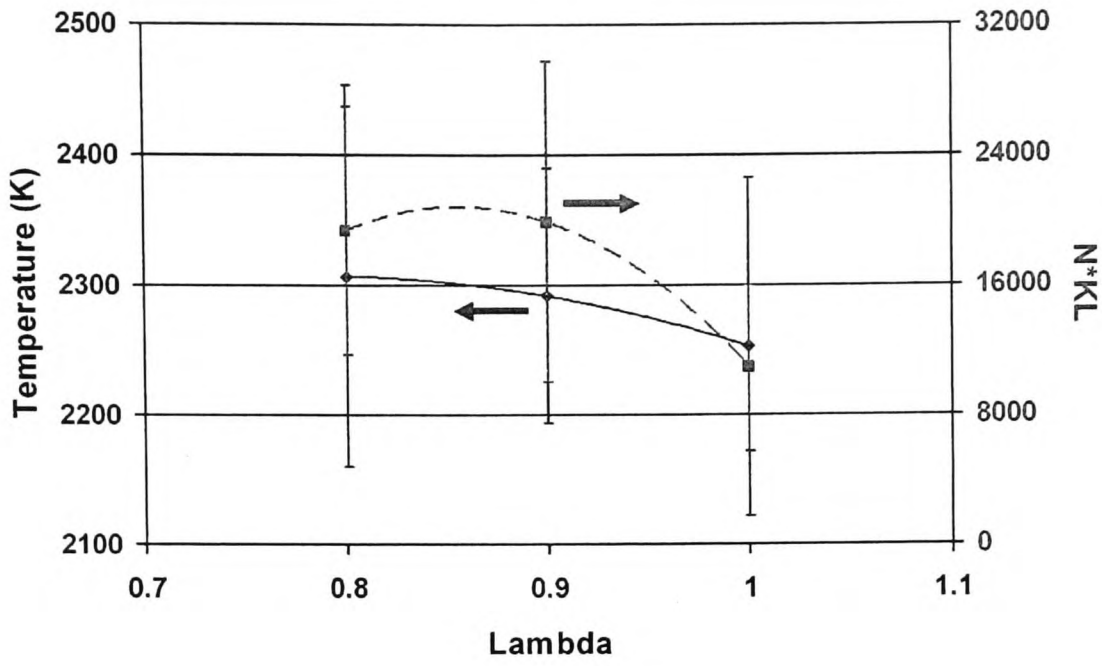


Figure 6.20: Relations between variable-KL results from high-speed CRP and mixture strength @ 15°ATDC

colour-ratio pyrometry.

A quartz window embedded in the metal piston was first used in a series of experiments and qualitative comparisons were made between the combustion sequences and the burn rate data from CoBRA. With such image sequences available, advanced image processing techniques for detecting the flame front and deriving its propagation were developed and demonstrated.

It was found that such a piston window gave restricted field of view beyond which the study of flame behaviour was not possible. A Perspex piston with a concave inner surface was then designed and made to allow a view of the whole combustion chamber. As a results, the analysable crank angle range for flame study was extended and the pyrometry results became more spatially representative. However, the Perspex piston was vulnerable to heating and soot fouling and hence demanded a very short firing period, typically less than 50 engine cycles. Therefore, measurements with this piston were from transient operation.

Chapter 7

DISI Engine Modelling

Positive anything is better than negative nothing.

— Elbert Hubbard (1856–1915)

7.1 Introduction

Thermodynamic modelling can be used as a substitute for experimental engine testing with considerable environmental, safety and cost benefits. Additionally, it allows parameters to be investigated that are beyond the reach of normal engine measurements.

This chapter details the utilisation of the ISIS engine modelling package, followed by discussions on the comparisons between ISIS results and experimental data such as cylinder pressure, burn rates from CoBRA, emissions and more specifically, temperature estimation from soot pyrometry. For convenience of use, two graphical user interfaces have been developed in MATLAB as pre- and post-processing tools respectively. Control parameters can then be preset by the user without having to deal with the FORTRAN code and DOS prompt commands.

7.2 Combustion modelling with ISIS

7.2.1 Graphical user interfaces

To take advantage of the data format and graphical functions in MATLAB, two graphical user interface (GUI) have been built up for ISIS, namely, ISISxp as an input front panel and ISIS2MAT as a post-processing package for conversion of data files to MATLAB format and plotting. A full list of the parameters present on the GUIs and their physical meanings are given in [48].

To run the executable ISIS code, the following files from FORTRAN are required in the same directory:

- Isis411.exe — the main executable
- ISISAUX.* *etc.* — file containing tolerances and NO reaction rate data, *e.g.* .HEY for data from Heywood [24]
- isoctane.dat *etc.* — fuel data files
- *.MAP — a MAP file that contains all the input information and can be ported into the main executable

One of the greatest advantages of ISIS is that, the user can input a selection of parameters as well as the ranges of these parameters at one run. For example, if the influence of both the air/fuel ratio and the engine speed on the combustion performance is of interest, two items can be included in the MAP file: [0.8 1.3 0.1] for Air/fuel ratio and [1000 2000 200] for engine speed, with the first and second element in the array representing the lower and upper range of the parameter, and the third giving the step. As a result, the code will output data files for 6x6=36 condition points.

One problem of the MAP file is, as the user varies some of the input parameters, the sequence of items in the file may change accordingly. This sometimes makes the customi-

sation of the MAP file not straightforward. Besides, manually inputting every parameter in the DOS prompt can be laborious. A graphical input panel, ISISxp, was hence created for the user's convenience (Figure 7.1). Two code files, ISISxp.m and ISISxp.fig, need to be put into the same directory with the files of the FORTRAN version.

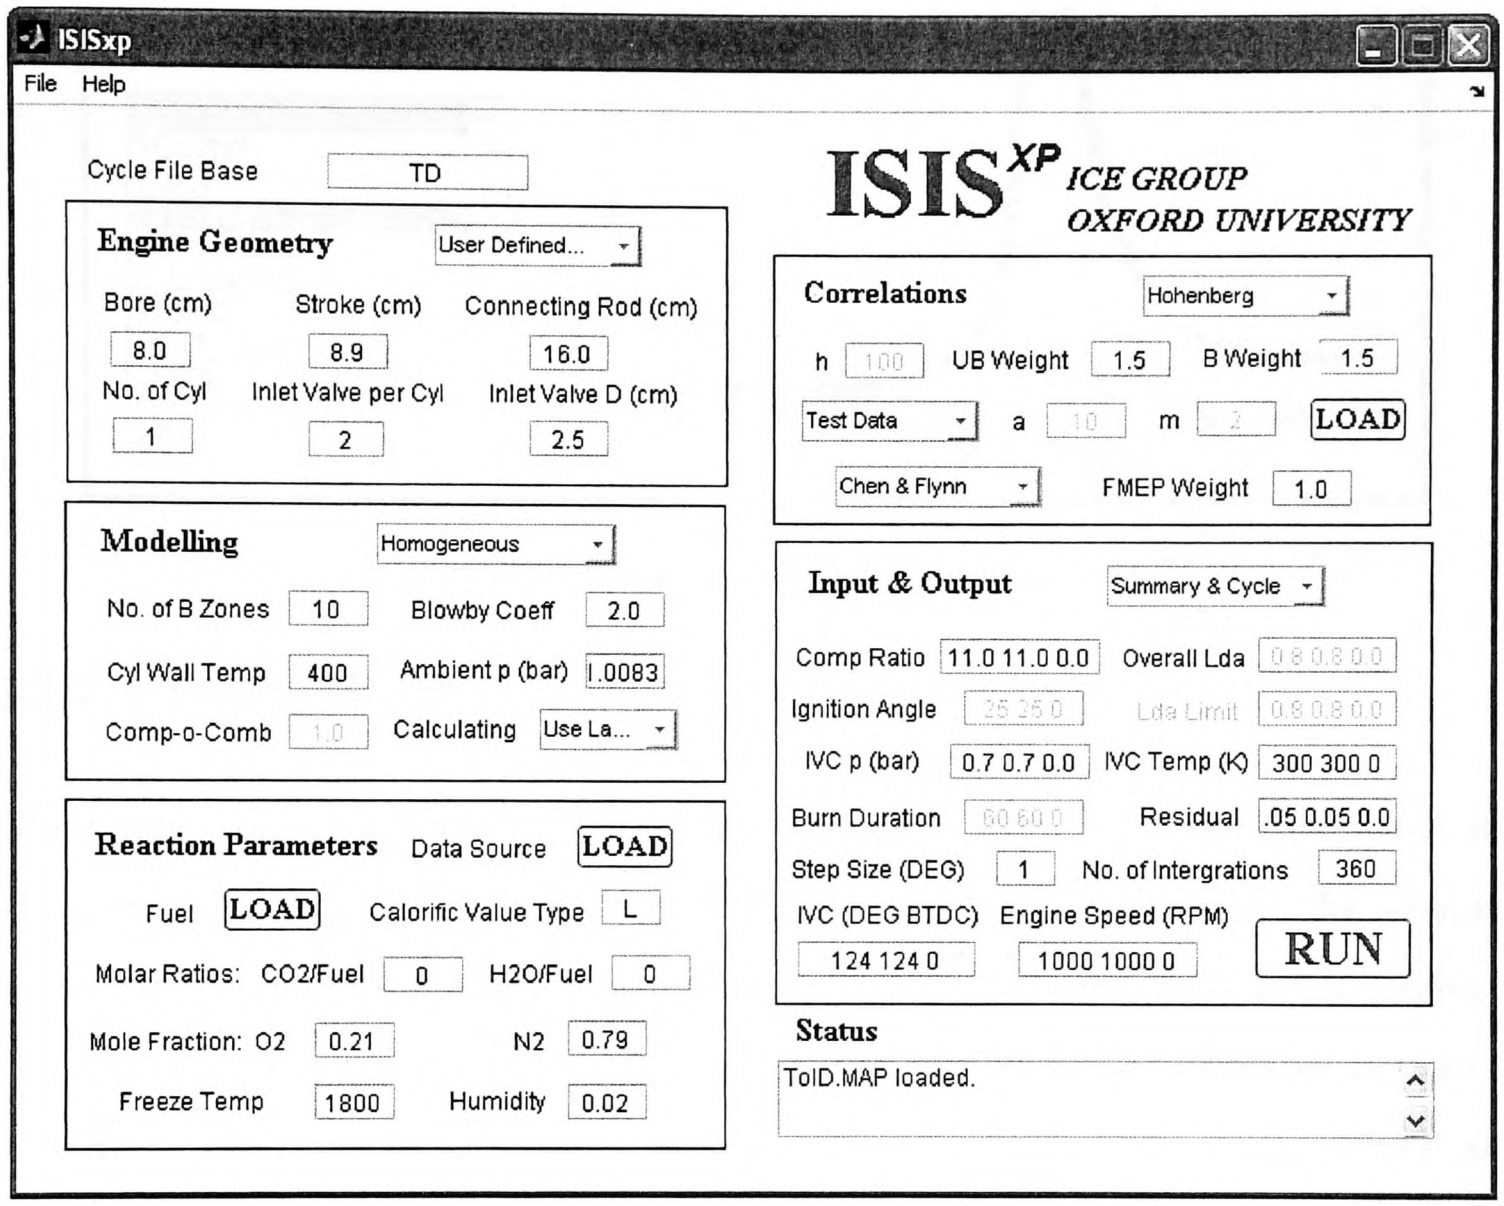


Figure 7.1: Graphical user interface ISISxp

Data files from ISIS were originally all text files which would cause too much burden when the users wanted to extract the data into a spreadsheet such as Excel and make plots. Another GUI, ISIS2MAT, was therefore built to perform this post processing and make graphical display user-friendly. Similarly, two code files, ISIS2MAT.m and ISIS2MAT.fig, are required for this GUI, but do not have to be in the same directory as the data files that are to be processed. An example of running ISIS2MAT is shown in Figure 7.2.

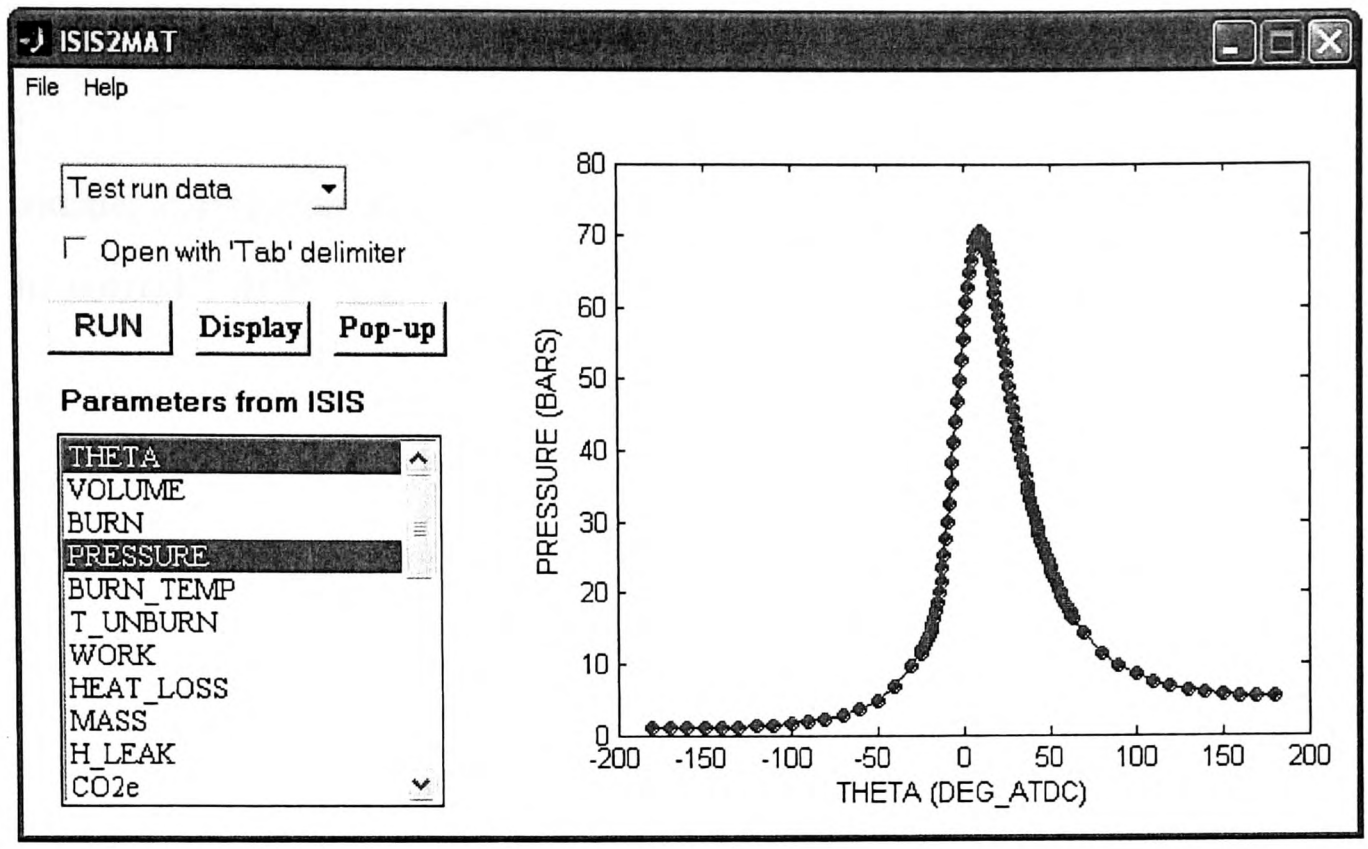


Figure 7.2: Graphical user interface ISIS2MAT

7.2.2 Tuning ISIS

ISIS ensures conservation of mass and energy, which is also true of all engines, but not all experimental data from engines. It is hence important to be confident in the accuracy of any experimental data before making conclusions on the matching between results of ISIS and the engine testing. ISIS is then required to be tuned to give a match as close to the experimental data as possible, particularly on $p_{\max}$, crank angle of $p_{\max}$, IMEP and mass fraction burned.

The most critical step in tuning the model is to ensure that the correct mass of fuel is being trapped in the engine. The quantity of trapped fuel in ISIS depends on:

- the pressure at IVC
- the volume at IVC
- the temperature at IVC
- and the trapped residual fraction

The blowby coefficient should be adjusted to give the correct mass leakage, but this is usually only of second order significance. The most important tuning parameter is the heat transfer correlation weighting factor for burned zones. This should be adjusted to give the correct IMEP, $p_{\max}$, and $\theta_{\max}$.

The burn rate and completeness of combustion data from CoBRA should be used to tune the burn rate model in ISIS. For example, with the 10% and 80% MFB data, the Wiebe function [19] can be solved to give a nominal pair of empirical constants a and m, which can be further tuned until the Wiebe function gives a MFB curve close enough to the one from CoBRA.

Furthermore, a motored engine cycle can be modelled in ISIS by setting AFR to an extreme value *e.g.* 100 and setting the ignition timing to a very late one. However, in this case, blowby and heat transfer will have an effect of the same order of magnitude on the simulation which makes it difficult to decouple these two factors in the tuning process.

7.2.3 Engine discussions

An example of tuning ISIS with data from CoBRA is shown in Figure 7.3. The engine operating conditions in this example are the same as in Figure 5.25, with toluene and a lambda value of 0.8. Particular care was taken to match those data points around TDC because the temperature in this range was of the most interest for comparison with the pyrometry results given in Chapter 5. The noticeable difference between the two MFB curves is mainly because the ISIS MFB data has taken account of the completeness of combustion, which is about 90% in this case, whilst the MFB data from CoBRA was scaled to 100%.

A series of ISIS runs were undertaken to simulate the operations of the Jaguar engine described in Chapter 5, each tuned accordingly. The results are shown in Figure 7.4 and 7.5.

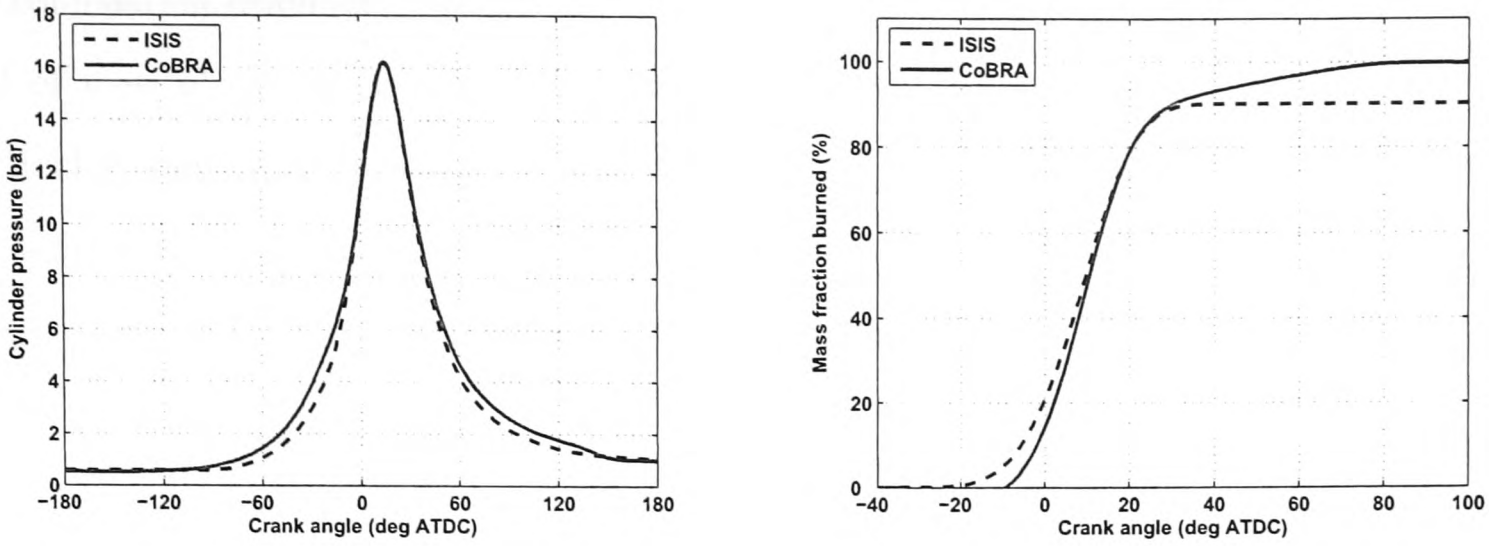


Figure 7.3: ISIS vs. CoBRA: cylinder pressure (left) and mass fraction burned (right)

In Figure 7.4, the step size used in ISIS was 1°CA. The best match between ISIS and pyrometry results occurred around TDC and there was increasing deviation on either side of up to $\sim \pm 200$K.

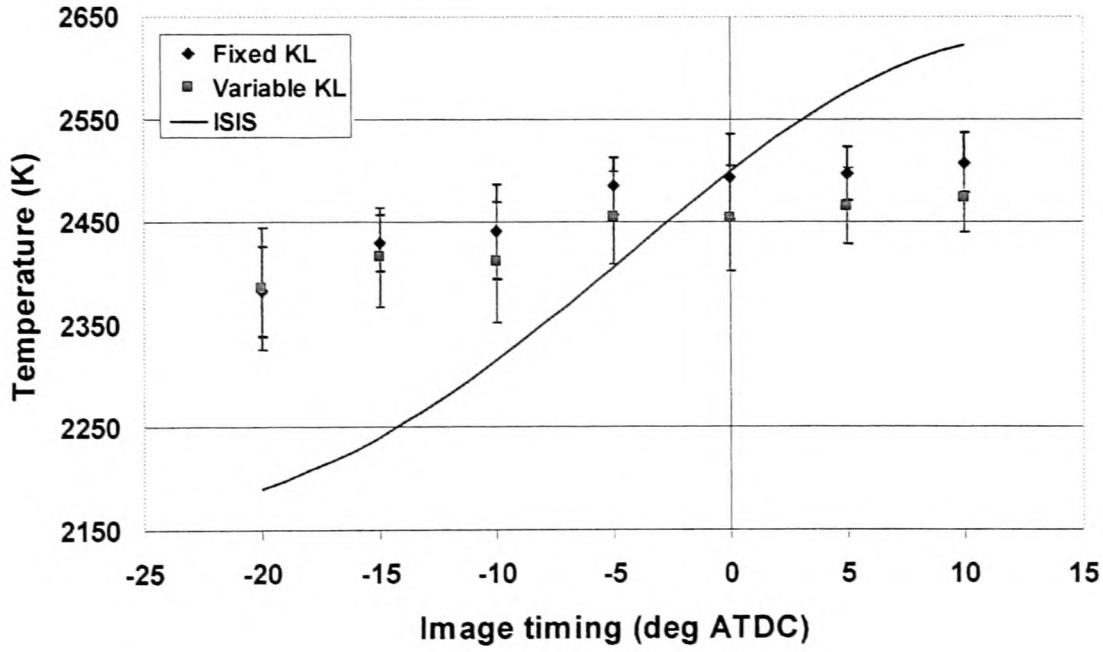


Figure 7.4: Comparison between ISIS and CRP: Temperature vs. image timing @ $\lambda = 0.8$

It should be kept in mind that ISIS calculates multi-zone burned gas temperatures based on zero-dimensional thermodynamic models, whilst soot pyrometry works on glowing soot from combustion. Therefore, a fundamental difference in the two temperatures may exist especially in the early stage of combustion when soot is not broadly distributed across

the combustion chamber. Besides, in cases where the optical access did not give a full field of view of the whole combustion chamber, flames occurring elsewhere other than the field of view would be ignored by soot pyrometry, but would still be accounted for in the simulation. Other possible reasons why there was such a difference in the pyrometry temperature prediction from the ISIS output may include:

- trade-offs in the tuning of parameters in ISIS
- line-of-sight errors in the optical diagnostics
- insensitivity of the pyrometry temperature prediction to the actual soot properties
- ISIS's incapability of simulating cycle-by-cycle variations

However, given the aforementioned overall error of the CRP, the ISIS simulation results agree remarkably well with the CRP estimation, particularly within the main combustion period (-5° to 10°CA ATDC, corresponding with 10–50% MFB).

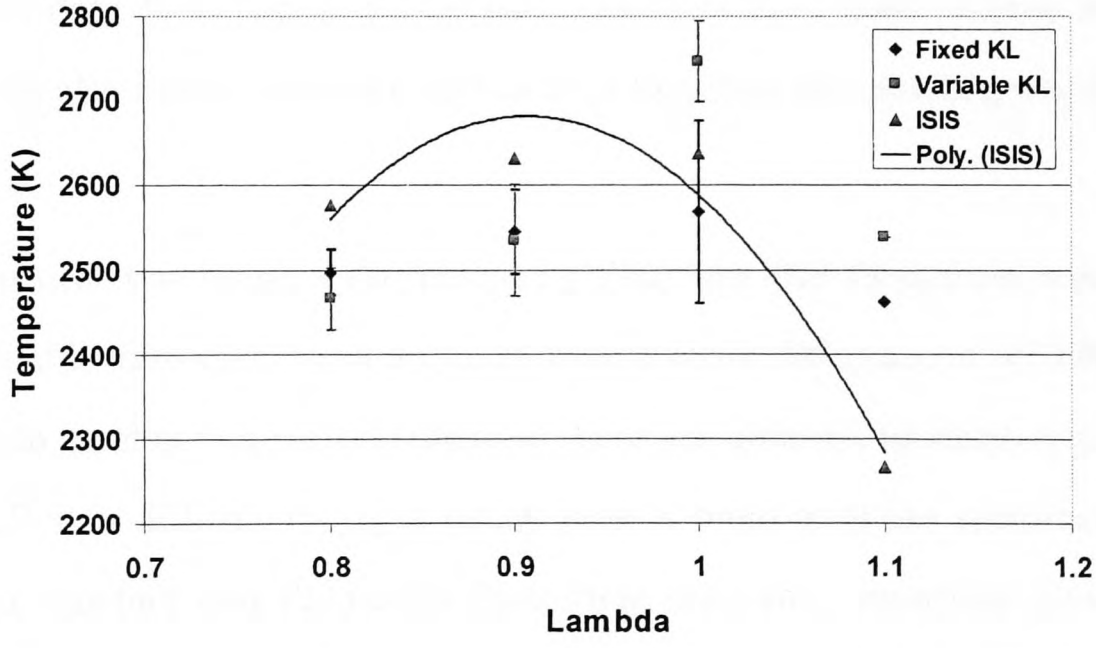


Figure 7.5: Comparison between ISIS and CRP: Temperature vs. mixture strength @ 5° after TDC

Figure 7.5 shows how the mixture strength affected the temperature from both ISIS and CRP. As expected, ISIS gave a peak temperature slightly rich of stoichiometry, similar to CRP. Again, there is insignificant difference between the two temperatures, ~100K in

most cases. The ISIS data differed the most from CRP under lean conditions. This was probably mainly due to the error in CRP introduced by combustion chemiluminescence.

7.3 Summary

To further justify the accuracy of colour-ratio pyrometry, a computer engine model, ISIS, has been tuned to simulate various operations of the Jaguar DISI engine. The tuning process was not always straightforward — it sometimes required many iterations before a reasonable match could be made between the ISIS output and the experimental data; furthermore, a compromise might need to be made when considering the overall effect of all tuning parameters.

Originally, ISIS was developed in FORTRAN and needed to be run in DOS environments. To save the user from fiddling with each input parameter and extracting data from the load of output data files, two graphical user interfaces have been created in MATLAB. The user would also benefit from the ease of data handling and plotting when calling ISIS in MATLAB.

Compared with the temperature estimation by CRP, the ISIS simulation results generally showed reasonable agreement with a typical temperature difference of $\sim$100K when looking at the main combustion period. Because there are different intrinsic error sources for ISIS and CRP, it is difficult to argue which gives a more accurate temperature calculation. However, the fact that the results from them both are comparable greatly supports the applicability of CRP. And with careful tuning, ISIS proved to be a powerful and yet cost-effective tool to study DISI engines.

Chapter 8

Conclusions and Future Work

Knowledge is not wisdom, unless used wisely.

— J. D. Anderson

8.1 Conclusions

A single-cylinder spray-guided optical DISI research engine was commissioned. Necessary control and data acquisition instrumentation was set up and relevant data/image processing software was developed in preparation for the engine measurements.

8.1.1 Optical diagnostics

Two optical techniques were studied and applied on the optical DISI engine, namely, flame imaging and quantitative planar laser-induced fluorescence (Q-PLIF). Soot pyrometry was developed to analyse images from combustion incandescence and flame front analyses were carried out with high-speed combustion imaging facilities.

Quantitative planar laser-induced fluorescence

The principles of Q-PLIF as it applies in a linear regime to the measurement of fuel concentration were outlined in the context of IC engine research. The underlying theory as well as the existing literature were reviewed to show the mathematical background, a prerequisite for justification of quantitative measurements. It has been shown that without careful measurement and understanding of the parameters that contribute to absorption and fluorescence quantum efficiency, Q-PLIF measurements cannot be justified.

Secondly, the design and characteristics of a six-component three-fraction fuel was described. This multi-component fuel was used with fluorescent tracers to determine quantitatively the distribution of different fuel fractions in an engine. The three tracers were selected to match the light, medium and heavy fraction of the fuel respectively. A definition of co-evaporation of the fuel components and the corresponding tracer was proposed. With such criteria, this multi-phase fuel and its tracers were shown to co-evaporate at least as well as existing, accepted, single-component fuel and tracer combinations with negligible effect on overall volatility.

The absorption and fluorescence characteristics of each fuel component and tracer were studied to give an estimate of the bounds of signal-to-noise ratio. The first delivered batch of component fuel was studied spectroscopically and shown to generate considerable LIF signals due to contaminants not removed in the distillation process. Fuel components with better purity were then sourced for the blend of a new batch. This second-batch component fuel was proved to be of good quality for the subsequent PLIF experiments; however, these fuel components were not 100% non-fluorescing compounds. Therefore, samples of the six fuel components and three tracers were tested in a LIF set-up to inverstigate their fluorescence yield individually and hence estimate how much fluorescence the mixture might produce as a whole. The fluorescent spectra from these compounds were recorded, processed and weighted to give an estimation of the signal-to-noise ratio

expected in a real engine test. With the spectral filters in use, acetone generally gave a better SNR than the other two tracers although its absolute LIF signal was weaker.

Great care was then taken to optimise both hardware and SNR, including:

- background light minimisation (removal of contaminating signals from the combustion chamber, oil and coolant leaks *etc.*) and subtraction
- correction for optical distortion, temporal and spatial fluctuation in the laser sheet and so on
- time gating of the ICCD camera and optimisation of the camera settings
- spectral filtering in the fluorescence wavelength range of the tracers
- study of stimulated Raman scattering for acetone PLIF
- laser pulse stretching for toluene PLIF
- measurements of the laser intensities for absorption saturation of the tracers

A recirculating loop was constructed and set up for the calibration of PLIF signal against AFR. The conditions in the loop (HC concentration of the fill, pressure and temperature) were monitored closely and recorded to produce a database for the calibration, which could be used later on to map the PLIF signal against crank position and AFR in normal engine experiments. This experimental layout has the advantages of *in situ* optics and a real engine environment. Example calibration results with acetone as the tracer were shown, suggesting that maximising SNR was the key to a practical calibration for quantitative PLIF.

Soot pyrometry

Soot pyrometry relates the temperature of glowing soot and the intensity of the thermal radiation it emits. Colour-ratio pyrometry (CRP) uses the ratios of three colours to give two independent measures for determining the temperature and loading of soot, thus avoiding measurements of absolute intensities. The technique is more rapidly and robustly

applicable to engine environments. However, it has been shown that sufficient equipment dynamic range is needed for good accuracy.

A systematic calibration was carried out to identify the properties of the Pulnix colour CCD camera used in this work. Such properties included its spectral response, linearity in terms of gain, sensitivity and exposure time, dark noise and flat field performance. With these data CRP was applied to images from three validation experiments, namely: imaging of a blackbody cavity, a tungsten filament lamp and a tungsten ribbon lamp. Results from the bench experiments showed satisfactory agreement between the temperature derived from CRP and the reference source temperature. The sources of uncertainty in the temperature and emissivity of tungsten filaments were discussed. The R/G results were better at lower temperatures and the R/B results were better at higher temperatures, suggesting an overall temperature calibration for the CRP to best estimate the temperature. Error analysis was performed to understand the accuracy of the colour-ratio method, including the effect of camera digitization error and spatial averaging of random noise on the temperature estimation. Taking the worst case of the 1 bit digitization error into consideration, the error in the temperature estimate was still within ±50K for the temperature range of engine combustion (1500–2600K). The possible error caused by spatially averaged random noise was estimated to be within 50K for the temperature range of 1800–2600K. The overall accuracy of CRP was about 5–7% throughout normal combustion temperature range.

Combustion images from the DISI engine were then interpreted using CRP. A comparison of results from different prescribed KL values indicated that the predicted temperature was insensitive to the KL setting (whether variable or fixed). The general trend was that the temperature increased with increasing prescribed KL, and reached the highest when the soot was assumed to be a grey body. False colour images of spatial temperature and KL distribution were generated; however, the KL solutions were ill-conditioned. Using a camera with a better dynamic range would help increase the reliability of the

KL measurements. Relations between the spatially averaged temperature and the crank angle timing as well as mixture strength have been presented. The camera's dynamic range posed constraints on the selectable range of image timing. Similarly, results with variable KL were reported and the parameter $N * KL$ was introduced to extract information on soot formation. Due to the presence of combustion chemiluminescence, additional care is necessary when dealing with images from lean engine operating conditions.

High-speed combustion imaging

High-speed CMOS imaging facilities were employed to further the discussions of soot pyrometry on the DISI engine in a crank-angle resolved context. A timing control box was built with the function to select the image timing interval in degrees CA and to ignore the clock pulses when there was no combustion event. The imaging interval and camera exposure time were varied for different engine operating conditions. Due to the limited camera dynamic range and highly variable flame radiation intensity, compromise had to be made sometimes between temporal resolution and signal strength.

Experiments were undertaken with three fuels, namely, the Shell component fuel, iso-octane and toluene. However, the images taken with toluene were distinct and best suited for demonstrating the soot pyrometry analyses and flame front detection techniques. Therefore, discussions were concentrated on the results of toluene. Three particular engine cycles were extracted from three combustion cines to illustrate the comparisons between the high-speed image sequences and the corresponding mass fraction burned data.

It was seen that, with a standard quartz window embedded in the metal piston, visual perceptions of flames did not necessarily coincide with the mass fraction burned at a particular stage of the combustion process. Such a piston window also gave a restricted field of view beyond which the study of flame behaviour was not possible. This led to the

design and use of a Perspex piston which allowed full bore optical access. This piston had a flat top surface and a concave bottom inner surface to enlarge the field of view, covering the whole combustion chamber. As a result, the analysable crank angle range for flame study was extended and the pyrometry results became more spatially representative. However, the Perspex piston was vulnerable to heating and soot fouling and hence demanded a very short firing period, typically less than 50 engine cycles. Therefore, measurements with this piston were of transient operation. Optical distortion occurred due to this design and was corrected for before any further image analyses were carried out, using a local weighted mean spatial transformation.

Advanced image processing techniques for detecting the flame front and deriving its propagation were developed and demonstrated with chemiluminescent flame images taken with toluene under lean operating conditions. These techniques include:

- non-linear diffusion filtering for image smoothing and effective flame edge detection
- gradient-based edge detecting routines
- a critical point detection routine for a simple 2-D closed contour to reduce the number of boundary points while still maintaining the contour shape as far as possible
- temporal contour interpolation with the 1-D discrete cosine transform, for the derivation of flame propagation where the imaging interval was not sufficiently fine; with this algorithm, no node-to-node match between different contours was required

With the contour set from the above image processing routines, flame properties such as the mean flame radius and the flame centroid locus were studied. The change of the mean flame radius in time indicates the mean growth rate of the flame, and the flame centroid locus shows how the flame kernel moved in space.

Finally, crank-angle resolved CRP results were shown with different lambda. The CoV of T was around 5% for most crank angles; the cycle-by-cycle variation in soot loading was larger mainly due to ill-conditioned KL solutions. All in all, a camera with a better

dynamic range will greatly improve the accuracy of high-speed CRP. CRP's capability of resolving transient measurements was verified by comparison with simultaneous particulate measurements using a DMS500.

8.1.2 Combustion analysis

Data from a small homogeneous charge, two-stroke glow ignition engine was processed by the combustion analysis software CoBRA with the objective of developing a method to combine the burn rate and heat release analyses in order to investigate the combustion process of engines with premixed charge, but compression ignition. The combustion and performance characteristics of different nitromethane/methanol fuel blends, and the characteristic of different glow plugs were also investigated. Difficulties such as unknown ignition timing and the polytropic index were addressed by combining both heat release and mass fraction burn analyses. Results for all operating conditions showed a good correlation between the two methods. The 'nominal' spark timing and ratio of specific heats specified in the software had no significant effect on the calculated results on the phasing of heat release. Nitromethane improved both the initial and main combustion burn rate of pure methanol, beyond the uncertainty created by the unknown mixture strength and cycle-by-cycle variations. IMEP generally increased with increased nitromethane content, only if the start of combustion occurred earlier so that the combustion phasing was approaching the optimum. Comparison with other engines showed that the combustion period of the engine under study was remarkably similar to traditional research engines, even though the engine design and operating regimes were quite different. The combustion analysis has demonstrated that the glow plug had a marked effect with a very rich mixture. An optical study would be desirable to investigate this effect on combustion. In some cases, a partially burning cycle was always followed by a cycle with higher than average IMEP, probably because for such a cycle the trapped in-cylinder residuals and initial temperature enhanced combustion in the following cycle.

A thermodynamic phenomenological engine modelling package, ISIS, was introduced to simulate the combustion of this DISI engine. For convenience of use, two graphical user interfaces were developed in MATLAB as pre- and post-processing tools respectively. Following careful tuning, discussions were undertaken on the comparisons between ISIS results and engine experimental data such as cylinder pressure, burn rates from CoBRA, emissions and more specifically, temperature estimation from soot pyrometry. Compared with the temperature estimation by CRP, the ISIS simulation results generally showed reasonable agreement with a typical temperature difference of ~100K when looking at the main combustion period (-5° to 10°CA ATDC, corresponding with 10–50% mass fraction burned). In examination of the effect of mixture strength on the cylinder temperature, ISIS gave a peak temperature slightly rich of stoichiometric, similar to CRP. The ISIS data differed the most from CRP under lean conditions. This was probably mainly due to the error in CRP introduced by combustion chemiluminescence.

8.2 Suggestions for future work

The automotive market has seen several successful commercial models equipped with DISI engines, most of which possess wall-guided combustion systems. One of the most fundamental and yet challenging issues in order for the next generation of spray-guided DISI engines to be successfully marketed, remains their need to meet increasingly stringent emissions legislations, particularly particulate emissions. Although particulate matter has not been legislated for gasoline engines, environmental concerns have been calling for thorough understanding of the soot formation in DISI engines that perform stratified charge operations, because fuel stratification can lead to high level emissions of unburned HCs and soot. The Jaguar engine in this project has not been able to fire in stratified mode successfully because the current arrangement of the spark plug and injector is not optimal. It is therefore sensible to realise stratified charge operation in this engine in the future

and to allow for insights into the combustion phenomena related to fuel stratification. This can also better validate the developed Q-PLIF technique given the higher levels of fuel inhomogeneity.

The techniques discussed in the following sections have been shown to be very useful in studying fuel stratification and soot formation, and hence will be the focus of the suggested future work.

Combustion imaging

Intensive investigation of soot pyrometry has been carried out and reported at length in this thesis with an emphasis on the colour-ratio technique. However, the validation tests for this technique were indirect — the blackbody test did not cover a wide enough temperature range and the tungsten tests were based on scarce (perhaps even inaccurate) information from archived emissivity data. Furthermore, the calibration of a tungsten ribbon lamp was pointwise rather than spectroscopic and will become less valid over time due to material aging. Ideally, a black body calibration source that spans the whole temperature range seen in a combustion environment needs to be constructed. This may be a ceramic object with a cavity heated up with oxy-acetylene flames, or a furnace. Another proposal is the compact Ocean Optics LS-1-CAL unit, which is a spectrometer with a calibrated light source. It can calibrate the absolute spectral response of a radiometric system. This will greatly simplify the spectral calibration of a camera used for CRP.

Secondly, a colour camera (CCD for low-speed use and CMOS for high-speed applications) with more bits is desirable for an improved signal-to-noise ratio in CRP, in particular, the solution of KL. A Lumenera USB2.0 Lu165C colour CCD camera has been purchased and commissioned at Oxford. This 12-bit camera offers 1.4 Megapixel resolution and a frame rate of 15fps at full frame and up to 60fps at reduced frame resolution. Cycle-by-cycle variations of cylinder T and soot loading can be studied on the existing engine test

bed. The new Phantom v7 series offers 36- or 48-bit high-speed colour imaging, and is ideal for crank-angle resolved CRP as well as flame analyses based on chemiluminescence imaging. A variant of such a high-speed system may be a multi-sensor colour camera with internal light-splitting and relaying device. By exposing the sensors in series, an image sequence could be captured with an interframe time resolution as low as 10ns.

Stojkovic *et al.* [59] applied conventional two-colour pyrometry using an intensified three-camera high-speed imaging system on a stratified-charge wall-guided DISI engine and compared the results with OH* chemiluminescence imaging. A similar experimental layout may be implemented on the Jaguar DISI engine using two high-speed cameras, a colour one for CRP and an intensified monochrome one for OH* chemiluminescence detection. OH-PLIF may also be readily incorporated in this layout as a measure of studying flame front propagation and flame velocity field (planar). By moving the plane of the laser sheet a quasi-3D flame front and its velocity field may be identified. Appropriate spectral filtering will be the key issue for effective discrimination of signals between sub imaging systems. With such diversity of experiments, together with the simultaneous pressure traces, cycle-by-cycle variation can be investigated more closely and impartially.

Alternatively, three inexpensive monochrome cameras can be combined with three corresponding spectral filters to give red, green and blue signals respectively. With proper optics (light splitting *etc.*), CRP can be applied with further improvement in SNR due to the absence of colour de-mosaicing and possible cross talk between colour channels. The calibration procedure of the cameras will also be simplified.

Abu-Gharbieh *et al.* [2] proposed advanced image processing techniques different from the ones implemented in this work for flame front analyses. The techniques they advocated included active contour models for contour segmentation and a curve matching algorithm that incorporates level sets and geodesic path computation to track the propagation of curves representing successive flame contours within a sequence. These techniques are generally more mathematically complex than the ones used in this thesis and have not

been validated on natural flame images. There is considerable scope for investigating the use of these techniques.

Last but not least, particle image velocimetry (PIV) has been applied extensively on engines to study various flow phenomena. Its mathematical foundation of using cross correlation to derive the most probable particle motions between consecutive frames provides the potential of tracking incandescent soot particles (or little soot clouds) in combustion images. Therefore, analysing combustion images using the standard image processing code for PIV may reveal useful information of droplet burning and turbulence during combustion.

Air/fuel ratio measurements

As the next stage of the ongoing project, the calibration of the PLIF signal against air/fuel ratio under various pressure and temperature conditions needs to be completed. With the current component fuel and three tracers, a typical lambda range of 0.3 to 3 and crank angle range of 120 to 15 BTDC will be covered.

The effect of tracer concentration on the fluorescence yield needs to be understood so that an optimal tracer concentration can be used for each of the tracers. Compromise may have to be made between SNR and the dynamic range of the AFR measurements *i.e.* detectable lambda range. It has been shown in [14] that although a 2% concentration of acetone in iso-octane gives the strongest fluorescence intensity across the concentration range studied, the linear relationship between fluorescence yield and equivalence ratio $(1/\lambda)$ did not hold throughout the AFR range of interest.

The degradation of fuel/tracer mixtures may also be investigated. This can be performed by sampling the fill in the calibration loop occasionally during the calibration test and analysing the samples using a gas composition analyser such a gas chromatograph or a Fourier transform infrared spectrometer.

AFR measurements in a firing engine can then be undertaken using Q-PLIF. Previous combustion imaging experiments have shown evidence of fuel inhomogeneity in the Jaguar engine when running in early injection mode. However, the spatial variation of AFR may be small and not enough to validate the Q-PLIF technique throughout a wide range of AFR. Hence it will be interesting to run the engine in stratified mode and examine the results from Q-PLIF.

Laser-induced soot imaging

The techniques of laser-induced incandescence (LII) and laser-induced scattering (LIS) are useful tools to obtain spatially resolved (no line-of-sight integration), two-dimensional soot distribution and information on soot formation and morphology. When a cloud of soot is irradiated with an intense laser sheet, the soot particles absorb the laser energy and are hence heated to a temperature well above the surrounding gas temperature. The hot soot particles will then emit radiation that can be characterised by the elevated soot temperature provided the carbon ingredients in the soot particles are not vaporised at that temperature. By solving the relevant continuity, energy and mass conservation equations and Planck's radiation equation with certain assumptions [77], the temperature and diameter of soot can be determined. The LII signal can be calibrated to give soot volume fraction.

With the existing setup of laser and optics, LII and LIS may be implemented to study soot in the DISI engine. This can provide a different and complementary viewpoint to the understanding of soot emissions by DISI engines aside from soot pyrometry.

References

[1] R. Abu-Gharbieh, G. Hamarneh, T. Gustavsson, and C. F. Kaminski. *Flame Front Tracking by Laser Induced Fluorescence Spectroscopy and Advanced Image Analysis.* Opt. Exp., 8:278–287, 2000.

[2] R. Abu-Gharbieh, C. Kaminski, T. Gustavsson, and G. Hamarneh. *Level Set Curve Matching and Particle Image Velocimetry for Resolving Chemistry and Turbulence Interactions in Propagating Flames.* Journal of Mathematical Imaging and Vision 19:199–218, 2003.

[3] S. K. Ahn, Kamimoto T. Matsui, Y. and, and S. Matsuoka. *Measurement of Flame Temperature Distribution in a D. I. Diesel Engine By Means of Image Analysis of Nega-Color Photographs.* SAE Transactions, Vol. 90, SAE Paper 810183, 1981.

[4] M. Alperstein, G. H. Schafer, and I. Villforth. *Texaco's Stratified Charge Engine — Multifuel, Efficient, Clean and Practical.* SAE Technical Paper No. 740563, 1974.

[5] R. W. Anderson, J. Yang, D. D. Brehob, J. K. Vallance, and R. M. Whiteaker. *Understanding the Thermodynamics of Direct Injection Spark Ignition (DISI) Combustion Systems: An Analytical and Experimental Investigation.* SAE Technical Paper No. 962018, 1996.

[6] J. K. Ball. *Cycle-by-Cycle Variation in Spark Ignition Internal Combustion Engines.* PhD thesis, University of Oxford, 1998.

[7] J. K. Ball, C. R. Stone, and Raine R. R. *A Technique for Estimating Completeness of Combustion and Its Use in Modeling Cycle-by-Cycle Variations in Combustion.* SAE Technical Paper No. 2000–01–0953, 2000.

[8] B. Befrui, R. Kneer, S. Breuer, W. Reckers, D. Robart, H. Wanlin, and C. Weiten. *Investigation of a DISI Fuel Injector for a Close-Arranged Spray-Guided Combustion System.* SAE Technical Paper No. 2002–01–1133, 2002.

[9] M. F. J. Brunt and A. L. Emtage. *Evaluation of Burn Rate Routines and Analysis Errors.* SAE Technical Paper No. 970037, 1997.

[10] F. Cignoli, S. De Iuliis, V. Manta, and G. Zizak. *Two-Dimensional Two-Wavelength Emission Technique for Soot Diagnostics.* Applied Optics, Vol. 40, No. 30, P5370–5378, 2001.

[11] D.R. Cox and Oakes D. *Analysis of Survival Data.* Chapman & Hall, London, 1984.

[12] C. S. Davis and M. A. Stephens. *Approximate Percentage Points using Pearson Curves.* Applied Statistics, Royal Statistical Society, 1983.

[13] M. H. Davy. *Two-Phase Fuel Visualisation in a Direct-Injection Gasoline Engine.* PhD thesis, University College London, 2000.

[14] G. de Sercey, M. Heikal, M. Gold, S. Begg, R. Wood, and G. Awcock. *On the Use of Laser-Induced Fluorescence for the Measurement of In-Cylinder Air-Fuel Ratios.* Proc. IMechE, Part C: J. Mechanical Engineering Science, Vol. 216(10), P1017–1029, 2002.

[15] S. di Stasio and P. Massoli. *Influence of the Soot Property Uncertainties in Temperature and Volume-Fraction Measurements by Two-Color Pyrometry.* Measurement Science and Technology, Vol. 5, No. 12, P1453–1465, 1994.

[16] A. C. Eckbreth. *Laser Diagnostics for Combustion Temperature and Species.* Gordon and Breach Publishers, Netherlands, 1996.

[17] H. Eichlseder, E. Baumann, Muller P., and Rubbert S. *Gasoline Direct InjectionVA Promising Engine Concept for Future Demands.* SAE Technical Paper No. 2000–01–0248, 2000.

[18] P. G. Felton. *Laser Diagnostics for Gasoline Direct-Injection Engines.* In Proc. IMechE Seminar on Lean Burn Combustion Engines, 1996. 3–4 Dec, London, MEP.

[19] C. R. Ferguson. *Internal Combustion Engines — Applied Thermosciences.* Wiley, New York, 1986.

[20] T. Georjon, E. Bourguignon, T. Duverger, B. Delhaye, and Voisard P. *Characteristics of Mixture Formation and Combustion in a Spray-Guided Concept Gasoline Direct Injection Engine: An Experimental and Numerical Approach.* SAE Technical Paper No. 2000–01–0534, 2000.

[21] A. J. Grant. *Laser Diagnostics of Spark-Ignition Combustion Systems.* PhD thesis, University of Oxford, 2000.

[22] D. Han and R. R. Steeper. *Examination of Iso-Octane/Ketone Mixtures for Quantitative LIF Measurements in a DISI Engine.* SAE Technical Paper No. 2002–01–0837, 2002.

[23] Z. Han, R. D. Reitz, J. Yang, and R. W. Anderson. *Effects of Injection Timing on Air-Fuel Mixing in a DISI Engine.* SAE Technical Paper No. 970625, 1997.

[24] J. B. Heywood. *Internal Combustion Engine Fundamentals.* McGraw-Hill, 1988.

[25] G. F. Hiett and J. V. B. Robson. *A High Power Two-Cycle Sleeve-Valve Engine for Aircraft.* Aircraft Engineering, 1950.

[26] R. C. Hilborn. *Einstein Coefficients, Cross-sections, f Values, Dipole-Moments, and All That.* Am. J. Phys. 50(11), P982–986, 1982.

[27] R. C. Hilborn. *Einstein Coefficients, Cross-sections, f Values, Dipole-Moments, and All That, Correction.* Am. J. Phys. 51(5), P471–471, 1983.

[28] Hottel and Broughton. *Determination of True Temperature and Total Radiation from Luminous Gas Flames.* Industrial Engineering and Chemistry, Analytical Edition, Vol. 4, 1932.

[29] R. Houston and G. Cathcart. *Combustion and Emissions Characteristics of Orbital's Combustion Process Applied to Multi-Cylinder Automotive Direct-Injected 4-stroke Engines.* SAE Technical Paper No. 980153, 1998.

[30] Y. Iwamoto, K. Noma, O. Nakayama, T. Yamauchi, and H. Ando. *Development of Gasoline Direct Injection Engine.* SAE Technical Paper No. 970541, 1997.

[31] N. S. Jackson, J. Stokes, P. A. Whitaker, and T. H. Lake. *Stratified and Homogeneous Charge Operation for the DISI Engine — High Power with Low Fuel Consumption and Emissions.* SAE Technical Paper No. 970543, 1997.

[32] Kar K. *Three-Thermocouple Technique for Fluctuating Temperature Measurement.* PhD thesis, University of Oxford, 2004.

[33] P. Kahn. *Development of a Combustion Analysis System for Internal Combustion Engines.* Final Year Project Report, Brunel University, 1988.

[34] S. Kato, S. Onishi, and S. Kobayashi. *DI Stratified Charge Engine by Impingement of Fuel Jet (OSKA) — Performance and Combustion Characteristics.* SAE Technical Paper No. 900608, 1990.

[35] W. Koban, J. D. Koch, R. K. Hanson, and C. Schulz. *Absorption and Fluorescence of Toluene Vapor at Elevated Temperatures.* Phys. Chem. Chem. Phys., 6(11), P2940–2945, 2004.

[36] 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), P777–784, 2005.

[37] 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), P147–150, 2005.

[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), P111–114, 2002.

[39] T. Kume, Y. Iwamoto, K. Iido, M. Murakami, K. Akishino, and H. Ando. *Combustion Control Technologies for Direct Injection SI Engine.* SAE Technical Paper No. 960600, 1996.

[40] T. Lake, J. Stokes, R. Murphy, R. Osborne, and A. Schamel. *Turbocharging Concepts for Downsized DI Gasoline Engines.* SAE Technical Paper No. 2004–01–0036, 2004.

[41] H. Ma, K. Kar, C. R. Stone, R. R. Raine, and H. Thorwarth. *Analysis of Combustion in a Small Homogeneous Charge Compression Assisted Ignition Engine.* International Journal of Engine Research, Vol. 7, No. 3, P237–253, 2006.

[42] H. S. Malvar, L. He, and R. Cutler. *High-Quality Linear Interpolation for Demosaicing of Bayer-Patterned Color Images.* Proceedings of IEEE ICASPP, 2004.

[43] S. Marshall. *Combustion Analysis in a Direct Injection Gasoline Engine.* Final Year Project Report, University of Oxford, 2006.

[44] J. J. Murphy and C. R. Shaddix. *Influence of Scattering and Probe-Volume Heterogeneity on Soot Measurements Using Optical Pyrometry.* Combustion and Flame, 143, P1–10, 2005.

[45] R. Ortmann, S. Arndt, J. Raimann, R. Grzeszik, and G. Wurfel. *Methods and Analysis of Fuel Injection, Mixture Preparation and Charge Stratification in Different Direct Injected SI Engines.* SAE Technical Paper No. 2001–01–0970, 2001.

[46] P. Perona and J. Malik. *Scale-space and edge detection using anisotropic diffusion.* IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 12, No. 7, P629–639, 1990.

[47] F. F. Pischinger and K. J. Klocker. *Single-Cylinder Study of Stratified Charge Process with Prechamber Injection.* SAE Technical Paper No. 741162, 1974.

[48] R. R. Raine. *Computer Modelling of Nitric Oxide Formation in Spark Ignition Engines.* Technical Report, 2004.

[49] R. R. Raine and H. Thorwarth. *Performance and Combustion Characteristics of a Glow-Ignition Two-Stroke Engine.* SAE Technical Paper No. 2004–01–1407, 2004.

[50] R. R. Raine, L. P. Wyszynski, and C. R. Stone. *Modelling of NO Emissions from Homogeneous and Stratified Charge Spark Ignition Engines.* Proc. I.Mech.E, J of Auto. Eng. 216 Part D, 403–412, 2002.

[51] G. M. Rassweiler and L. Withrow. *Motion Pictures of Engine Flames Correlated with Pressure Cards.* SAE Technical Paper No. 800131, 1980.

[52] K. Schanzlin, T. Koch, and K. Boulouchos. *Characterization of the Combustion in a Direct injection Spark Ignition Engine.* SAE Technical Paper No. 2002–01–0834, 2002.

[53] H. Schapertons, K.-D. Emmenthal, H.-J. Grabe, and W. Oppermann. *VW's Gasoline Direct Injection (GDI) Research Engine.* SAE Technical Paper No. 910054, 1991.

[54] 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 P75–121, 2005.

[55] A. J. Scussel, A. O. Simko, and W. R. Wade. *The Ford PROCO Engine Update.* SAE Technical Paper No. 780699, 1978.

[56] S. Simonini. *Combustion Analysis in IC Engines Using a CCD Camera and Imaging Processing.* Master's thesis, University of Oxford, 1998.

[57] S. Simonini, S. J. Elston, and Stone C. R. *Soot Temperature and Concentration Measurements from Colour Charge Coupled Device Camera Images Using a Three-Colour Method.* In Proc IMechE, Vol. 215, 2001. Pt C.

[58] R. Stevens, H. Ma, Stone C. R., Walmsley H., and Cracknall R. *On Planar Laser-Induced Fluorescence with Multi-Component Fuel and Tracer Design for Quantitative Determination of Fuel Concentration in IC Engines.* Proc. IMechE, Vol. 220, Part D: J. Automobile Engineering, 2006.

[59] B. D. Stojkovic, T. D. Fansler, Drake M. C., and Volker S. *High-Speed Imaging of OH^* and Soot Temperature and Concentration in a Stratified-Charge Direct-Injection Gasoline Engine.* Proceedings of the Combustion Institute, Vol. 30, P2657–2665, 2005.

[60] C. R. Stone. *Introduction to Internal Combustion Engines.* Macmillan Press, 1999.

[61] C. R. Stone and D. I. Green-Armytage. *Comparison of Methods for the Calculation of Mass Fraction Burnt from Engine Pressure-Time Diagrams.* In Proceedings IMechE, Vol. 201(D1), P61–67, 1987.

[62] J. P. Styron, P. L. Kelly-Zion, C. F. Lee, J. E. Peters, R. A. White, and R. P. Lucht. *Multicomponent Liquid and Vapor Fuel Distribution Measurements in the Cylinder of a Port-Injected, Spark-Ignition Engine.* SAE Technical Paper, No. 2000–01–0243, 2000.

[63] K. I. Svensson, A. J. Mackrory, M. J. Richards, and D. R. Tree. *Calibration of an RGB, CCD Camera and Interpretation of its Two-Color images for KL and Temperature.* SAE Technical Paper, No. 2005–01–0648, 2005.

[64] J. C. Swindal, D. P. Dragonetti, R. Hahn, P. A. Furman, and W. P. Acker. *In-Cylinder Homogeneity during Cold-Start Studied with Fluorescent Tracers Simulating Different Fuel Distillation Temperatures.* SAE Technical Paper, No. 950106, 1995.

[65] M. L. Sztenderowicz and J. B. Heywood. *Cycle-to-Cycle IMEP Fluctuations in a Stoichiometrically-Fueled SI Engine at Low Speed and Load.* SAE Technical Paper No. 902143, 1990.

[66] M. L. Sztenderowicz and J. B. Heywood. *Mixture Nonuniformity Effects on S.I. Engine Combustion Variability.* SAE Technical Paper No. 902142, 1990.

[67] T. Tomoda, S. Sasaki, D. Sawada, A. Saito, and H. Sami. *Development of Direct Injection Gasoline Engine — Study of Stratified Mixture Formation.* SAE Technical Paper No. 970539, 1997.

[68] Y. S. Touloukian and D. P. DeWitt. *Thermal Radiative Properties — Metallic Elements and Alloys, Thermophysical Properties of Matter (Vol. 7).* IFI/Plenum, New York-Washington, 1970.

[69] J. Vattulainen, V. Nummela, R. Hernberg, and J. Kytölä. *A System for Quantitative Imaging Diagnostics and Its Application to Pyrometric In-Cylinder Flame-Temperature Measurements in Large Diesel Engines.* Measurement Science and Technology, Vol. 11, No. 2, P103–119, 2000.

[70] R. C. Weast and M. J. Astle. *CRC Handbook of Chemistry and Physics.* CRC Press, Inc., 1981.

[71] M. Wirth, D. Zimmermann, R. Friedfeldt, J. Caine, M. Schamel, A. Davies, G. Peirce, A. Storch, K. Ries-Müller, K.-P. Gansert, G. Pilgram, R. Ortmann, G. Würfel, and J. Gerhardt. *A Cost Optimised Gasoline Spray Guided Direct Injection System for Improved Fuel Economy.* In Proceedings IMechE — Fuel Economy and Engine Downsizing, 2004.

[72] L. P. Wyszynski. *Investigation of Combustion Processes in a Direct Injection Spark Ignition Engine.* PhD thesis, University of Oxford, 2002.

[73] L. P. Wyszynski, R. Aboagye, C. R. Stone, and G. T. Kalghatgi. *Combustion Imaging and Analysis in a Gasoline Direct Injection Engine.* SAE Technical Paper No. 2004–01–0045, 2004.

[74] L. P. Wyszynski, C. R. Stone, and G. T. Kalghatgi. *The Volumetric Efficiency of Direct and Port Injection Gasoline Engines with Different Fuels.* SAE Technical Paper No. 2002–01–0839, 2002.

[75] Wang X. *Instantaneous In-Cylinder Heat Transfer and Combustion Analysis in Spark Ignition Engine.* Transfer report, University of Oxford, 2005.

[76] F.-Q. Zhao, M.-C. Lai, and D. L. Harrington. *Review of Mixture Preparation and Combustion Control Strategies for DISI Engines.* SAE Technical Paper No. 970627, 1997.

[77] H. Zhao and N. Ladommatos. *Optical Diagnostics for Soot and Temperature Measurement in Diesel Engines.* Progress in Energy and Combustion Science, Vol. 24, No. 3, P221–255, 1998.

[78] H. Zhao and N. Ladommatos. *Engine Combustion Instrumentation and Diagnostics.* SAE Publication Group, 2001.

[79] P. Zhu and P. M. Chirlian. *On Critical Point Detection of Digital Shapes.* IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 17(8), P737–748, 1995.

Appendix A

Incidental Work — Lumenera Camera Software

For better accuracy in combustion analysis and soot pyrometry, a Lumenera USB2.0 Lu165C camera was commissioned and equipped at Oxford. It is a 12-bit 1.4 Megapixel (1392 by 1040) colour camera with a frame rate of 15fps at full frame and up to 60fps at reduced frame resolution. It also has four I/O channels for control of peripherals and imaging synchronization (*e.g.* strobe lighting, external triggering). Other features include high quality CCD sensor, high sensitivity, high dynamic range, Bayer-pattern colour filter array, video streaming and so on. To accommodate the use of this camera, control software LucamICE was developed in-house.

The LucamICE software was built in Microsoft Visual Basic (VB) 6.0 based on the Lumenera USB Camera API (Application Programming Interface), which provides a comprehensive set of functions allowing the control and operation of any Lumenera USB camera or camera module.

Straightforward applications of this VB software include: setting the camera properties, retrieving, displaying and capturing image data etc. Advanced functions allow for very

powerful features such as: continuously capturing multiple frames, external camera triggering and strobe triggering, streaming video capture to a file, video overlay, synchronous image capture from multiple cameras.

This document serves as a guide to users about the camera operation with the LucamICE software. Each of the software functions is described in detail followed by examples with a visual aid. The last section briefly discusses the structure of the software in VB with the Software Developer's Kit (SDK) provided by Lumenera. This SDK package contains many examples on how to use the API functions in real-world applications. Those users who wish to debug, modify or extend LuCamICE are advised to read the PDF documents 'Lumenera USB Camera API Reference Manual' and 'Lumenera USB Camera User's Manual', as well as the Microsoft VB 6.0 Documentation for relevant guidance on programming. The PDF documents can be found in the directory where the LuCam Capture Software is installed, and the VB 6.0 documentation is referred to as Microsoft Developer Network 6.0 (MSDN).

User Interface of LuCamICE

Before running the LuCamICE software, make sure at least one Lumenera USB camera is plugged into the PC and the camera driver has been properly installed. For details about the installation of the camera driver as well as the LuCam Capture software, please refer to 'Lumenera USB Camera User's Manual'.

The LuCamICE software has been compiled to a self-executable (EXE) file, LuCamICE.exe. Double click on the EXE to run the software. Figure A.1 shows the user interface of LuCamICE when loaded.

As no camera is connected at this stage, most of the functions (command buttons, text boxes etc.) on the panel are not available (in grey colour) as shown in Figure A.1. The

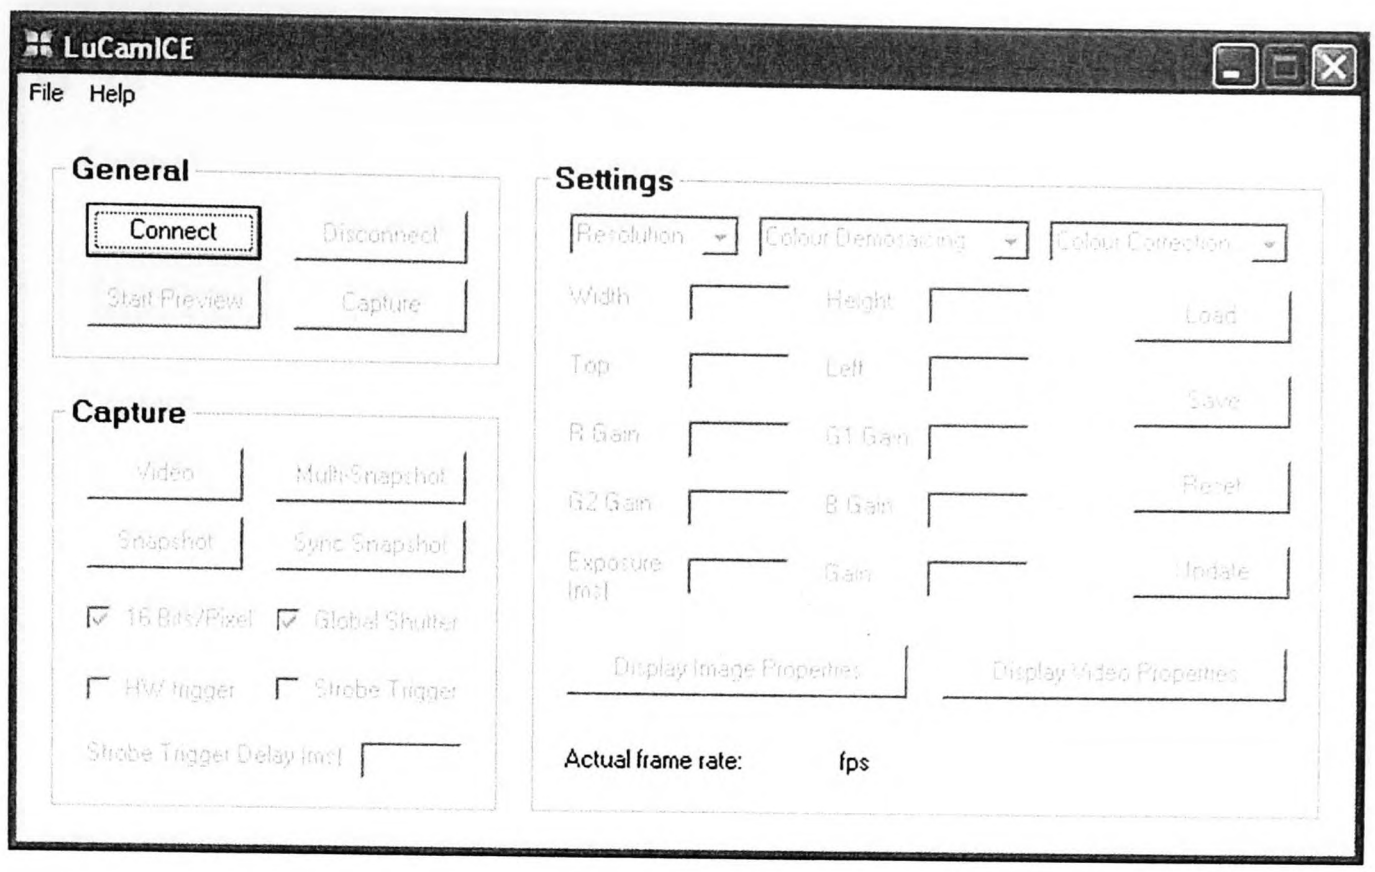


Figure A.1: The user interface of LuCamICE at start-up

interface consists of three parts: *General*, *Capture* and *Settings*, each of which is a frame containing several VB objects.

General camera operation

In order to become functional and controllable, a Lumenera USB camera must be 'connected' by calling the API function *LucamEnableFastFrames()*. In fact, almost every camera operation with the user interface is associated with one API function (sometimes more). All the API functions are described with examples in 'Lumenera USB Camera API Reference Manual'. Therefore, they are not covered in this manual to avoid repetition.

To connect the camera, simply click on the command button *Connect*. Then most software functions will be enabled (See Figure A.2). Similarly, click on *Disconnect* to disconnect the camera.

Also in the *General* frame is the preview function. Clicking *Start Preview* will initiate a preview video window streaming at the fastest available frame rate in video mode. This

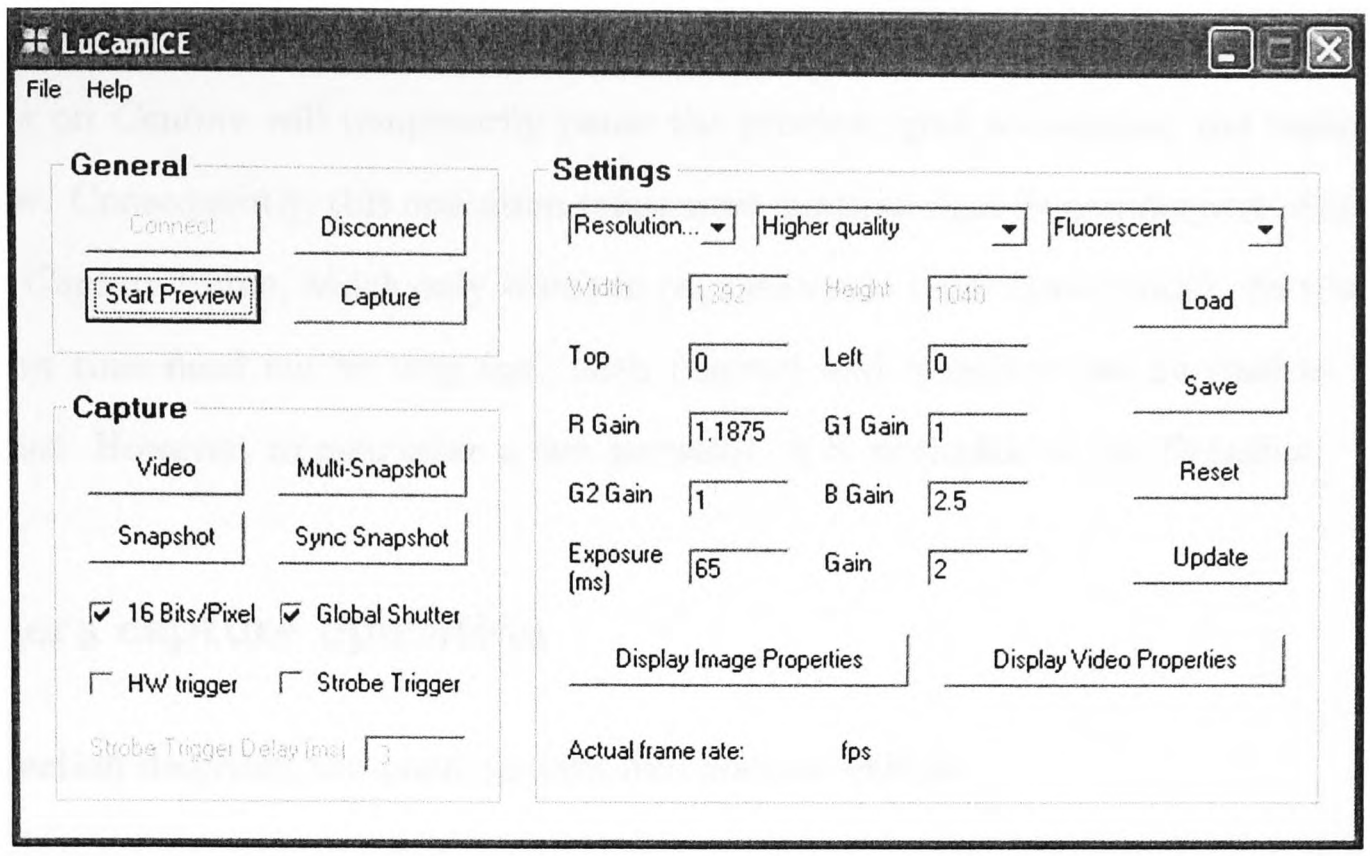


Figure A.2: The user interface of LuCamICE when the camera is connected

function allows the user to check whether the camera settings are appropriate for the current use. No data can be saved in this mode. When the preview window is displayed, the *Start Preview* button becomes *Stop Preview*. Click on it to stop the video preview.

The *Capture* button under this category allows the user to take a single snapshot. When a snapshot frame is obtained and displayed on the screen, the user can examine the image before saving it. When trying to close the window (clicking the 'X' button on the top right corner of the window), the program will ask if the user wants to save the frame. Click 'Yes' to save the image in one of the four supported formats: TIFF (*.tif), Bitmap (*.bmp), JPEG (*.jpg) and raw data (*.raw). Be aware that only the TIFF and raw data formats support 16-bit-per-pixel mode. When the *16 Bits/Pixel* check box in the *Capture* frame is ticked, the captured frame will not be displayed properly on the screen because Windows cannot handle 48-bit Bitmap colour display. Use other image viewers or processors such as MATLAB to display the saved 48-bit colour image. In particular, when reading 8-bit images (LZW-compressed TIFF) saved with LuCamICE into MATLAB, the user must have a version of R14SP2 or above.

Capture works whether the preview mode is on or not. If the preview video is streaming, clicking on *Capture* will temporarily pause the preview, grab a snapshot and resume the preview. Consequently, this operation takes more runtime than its counterpart, *Snapshot*, in the *Capture* frame, which only works in capture mode (non-video mode). So when the read-out time need not be very fast, both *Capture* and *Snapshot* can be used to take a snapshot. However, to guarantee a fast snapshot, it is advisable to use *Snapshot*.

Camera capture operation

This section describes the main capture functions as follows:

Video — Clicking this button will allow the user to capture and save a video.

Firstly, the program pops up an input box asking the user to specify the video record time in seconds. Then there will be another input box asking the user to specify the desirable frame rate to capture the video. Note that a fastest available frame rate exists depending on the camera resolution. If too fast a frame rate is input, the program will report an error message and return to the main interface. If the exact specified frame rate is not available, the nearest available one will be used.

When the video capture is finished, a file I/O dialog will show up asking where the user wants to save the raw video data. This raw data will be used to generate standard 8-bit AVI video file later but this format cannot be recognised by a common Windows media player. Finally, another file I/O dialog pops up asking where the user wants to save the AVI file.

The actual frame rate of the video will be displayed at the bottom of *Settings*.

Snapshot — This button performs the same function as the *Capture* button under *General* except that *Snapshot* gives a faster read-out time but does not work when preview is on.

Multi-Snapshot — This function calls the same set of API functions with a VB 'For'

loop.

The user can specify the number of frames to grab and the imaging idle time interval (only available in software triggering mode) on input boxes after click on *Multi-Snapshot*. The imaging idle time interval is the time the program has to wait until each 'Next' step is taken in the 'For' loop. The actual capture time of each image will depend on the exposure time set in the *Settings* frame and cannot be controlled precisely. Besides, the user may sometimes want to take images continuously at a slow frame rate with short exposure time. So the actual frame rate will be the total time used to take the number of frames specified by the user, i.e. the image capture time + the imaging idle time, divided by the number of frames.

When all the frames have been obtained, a file I/O dialog will pop up asking the user where to save the images. Here the user only needs to type in the prefix of the image names. All the images will be saved in sequence such as ***1.tif, ***2.tif...

The actual frame rate of the image sequence will be displayed at the bottom of *Settings*. No real-time display is available because of the need to have the read-out time as fast as possible.

Sync Snapshot — This button allows the user to grab simultaneous images with multiple Lumenera USB cameras. As the Internal Combustion Engine Group at Oxford only possess one such camera at the moment, this function has not been fully tested. In case of error, future users of this function are advised to look into the VB source code.

16 Bits/Pixel — Ticking this check box enables the 16-bit pixel format.

The Lumenera Lu165C camera at Oxford supports 12-bit pixel bit depth (a dynamic range of 72dB). However, all images are saved in either 8-bit (24-bit colour) or 16-bit (48-bit colour) file format. In 16-bit mode, the unused 4 bits are filled with 0s. It should be noted that when this check box is ticked, the captured colour frame by *Capture* and *Snapshot* will not be displayed properly because Windows systems

cannot deal with 48-bit bitmap colour display. Images saved in the 16-bit format can be opened with image viewers and processors such as MATLAB. In particular, when reading 8-bit TIFF images (LZW-compressed TIFF) saved with LuCamICE into MATLAB, the user must have a version of R14SP2 or above.

Global Shutter — Ticking this check box enables the global shutter of the connected camera (only when it supports this option). Global shutter has advantages over the other types of shutters such as rolling shutter. Please refer to 'Lumenera USB Camera User's Manual' for detailed discussion.

HW trigger — Ticking this check box enables the hardware trigger (rising edge, minimum width 0.5μs) fed to the camera I/O interface from an external source (shaft encoder, function generator, ETCS program etc.).

When HW triggering is used, the camera awaits the trigger signal to start the capture and no imaging idle time interval needs to be specified.

Strobe trigger — Ticking this check box enables the output of a strobe trigger signal that triggers a strobe or flashing device.

When strobe triggering is used, the *Strobe Trigger Delay* text box will become available for the user to specify the time in milliseconds between the start of exposure and the rising edge of the strobe trigger pulse. If 0 delay is set, the strobe trigger signal can also indicate when the exposure starts and hence tell the actual imaging interval when multiple snapshots are taken.

Camera control and settings

The *Settings* frame contains the major camera parameters for image capture, apart from the ones already set under the *Capture* category, including image resolution (*Width* and *Height*), position of sub-window (*Top* and *Left* relative to the top left corner of a full resolution frame), exposure time, camera gain, individual gains for white balance correction

(R, G1, G2, B gain) as well as methods for colour demosaicing and colour correction. Please refer to 'Lumenera USB Camera User's Manual' and 'Lumenera USB Camera API Reference Manual' for further description of colour demosaicing and colour correction.

NOTE:

1. The image resolution must be changed by a multiple of 8 pixels (subwindowing);
2. with different image resolution, the longest available exposure time may vary, therefore a maximum exposure time apply when setting it, i.e., when the user specifies an unavailable exposure time the program will correct for it;
3. when the user specifies an unavailable gain the program will correct for it.

The four command buttons on the right are for quick setting. Whenever the user changes any parameter on the panel, the *Update* button needs to be pressed to make the change effective. Once the user is satisfied with how the camera behaves, (s)he can save the listed camera setting by clicking on the *Save* button. The program then generates a text file containing values of these parameters, which can be loaded at a later stage by clicking the *Load* button. This feature is particularly useful when many sets of experiment are to be carried out on different occasions with identical camera settings. Additionally, the *Reset* button can be used to retrieve the original camera setting when it was powered on.

If the user wants to see all the image/video properties, either press *Display Image Properties* or *Display Video Properties* to pop up dedicated dialogs (Figure A.3). However, these are merely used for monitoring purpose, changing the parameter values on them will not take effect (the program may be extended to have this function if needed).

On the bottom of *Settings*, there is a text box displaying the actual frame rate per second of the video capture or multi-snapshot.

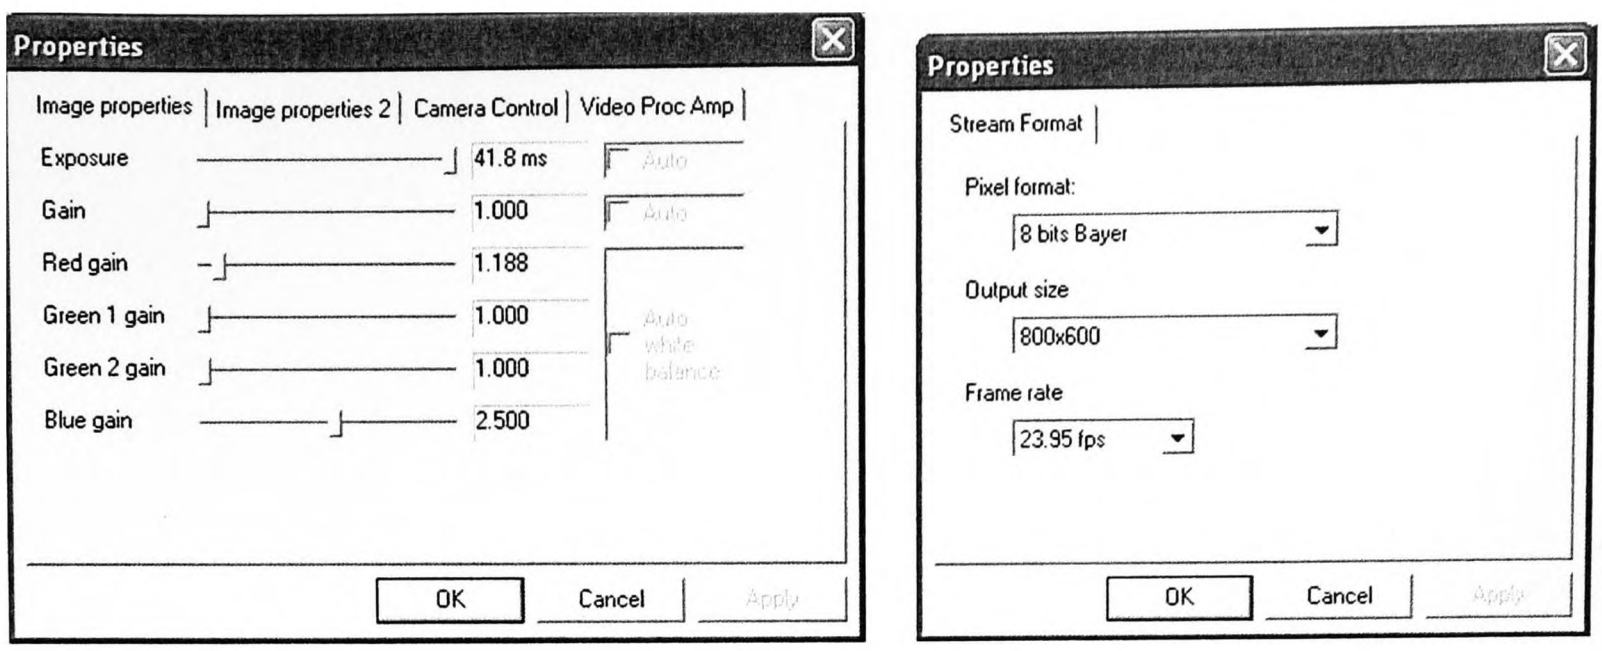


Figure A.3: Pages of the image properties (left) and video properties (right)

Programming with Lumenera SDK

The Lumenera SDK is purchased separately from the camera as a package. It is a built-in part of the LuCam Capture software but in order to install it the user needs a password (Lumenera SDK Key) during the installation of the software. For the Oxford Lu165C camera, the SDK product key is (case-sensitive)

Sierra07

With the SDK properly installed, the user can extend the existing software function by introducing more advanced API functions, for example, to save 48-bit colour images or the raw Bayer data. There are also various examples in the SDK package for the developer to get a taste of the programming. The official LuCam Capture software was developed in Visual C++ 6.0 whilst the LuCamICE was written in Visual Basic 6.0. These are quite different computer languages but the user can use either to achieve much the same result. In case the user wants to upgrade the software with Microsoft Visual Studio .NET, load the code files of version 6.0 in the .NET environment, Microsoft will do the automatic upgrade, converting the code to a readable state for the .NET platform.