

**Progress Towards a Demonstration of
Multi-Pulse Laser Wakefield Acceleration
and Implementation of a Single-Shot
Wakefield Diagnostic**

Stephen John David Dann
Lincoln College, Oxford

Thesis submitted in partial fulfilment of the requirements for the
degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2015

Abstract

An ongoing experiment is described to demonstrate the principle of multi-pulse laser wakefield acceleration, in which a plasma wakefield is resonantly excited by a train of laser pulses, spaced by the plasma wavelength. Particle-in-cell simulations of the initial single-pulse experimental setup are presented, in order to calculate the expected signal. Preliminary results are presented and future plans, based on work done so far, are discussed.

Part of this work involves the implementation of a single-shot wakefield diagnostic — frequency-domain holography, which records the phase shift caused by passage of a probe pulse through the plasma. This implementation is described in detail, along with the associated analysis procedure. Practical difficulties encountered while implementing the diagnostic are discussed, along with possible ways of mitigating them in the future. A method is presented by which the noise level in the resulting phase measurements can be predicted, much more accurately than any previously published method for this technique.

Methods of generating pulse trains for use in future multi-pulse laser wakefield acceleration experiments are presented. These include techniques proposed for use in this demonstration experiment, as well as one intended for use in a dedicated high-efficiency, high repetition-rate, multi-pulse driver laser. This last method, based on programmable pulse shaping using a spatial light modulator, requires a suitable mask to be computed based on the parameters of the required pulse train; an algorithm is described to perform this computation.

Acknowledgements

I would first like to thank my supervisors, Laura Corner and Roman Walczak, for their advice, guidance, and support throughout my studies. Not only did they make this whole project possible, they also provided an excellent working environment, relaxed yet professional. I would also like to thank the other members of our research group, both for their advice and support on my project, and for our occasional but much-appreciated social activities. Whether we were celebrating a successful day in lab, cheering ourselves up when things weren't going so well, or simply bored, I could always count on them to provide both moral and liquid support.

Special thanks should go to those who helped me during my experimental work. First this was James Cowley, whose many long hours of assistance were rendered with a good-humoured cynicism and ready wit that made even the most difficult experiment seem a little less bleak. Later, his boots were ably filled by Peter Tudor, whose steely determination in the face of an uncooperative experiment was truly inspirational. And thanks also to Nicolas Bourgeois for his help in keeping a rather temperamental laser in approximately working order.

Thanks also to my fellow students in the sub-department of Particle Physics, whose regular lunches and coffee breaks helped to break up the tedium of the working day. Particular thanks go to Neven Blaskovic Kraljevic, whose tireless enthusiasm and endless good spirits are positively infectious.

I would also like to thank the members, past and present, official and unofficial, of the Oxford University Society of Change Ringers, who are far too numerous to be listed here. Over my many years in Oxford, the OUSCR has provided diversion, relaxation, distraction, entertainment, and copious quantities of alcohol, but most of all friendship and a sense of community. Throughout the course of my studies, this eclectic band gave me the valuable opportunity to socialise away from the Department of Physics, and even to get away from Oxford once in a while.

It is hard to imagine my time here without the OUSCR, its presence both constant and ever-changing, and few of its present members can remember a time before I arrived. On which note, my commiserations to those members who, with the completion of this thesis, are officially becoming badgers. It is the end of an era.

Finally, I would like to thank my family for supporting me during my research, for staying in touch with me rather better than I've stayed in touch with them, and for the fact that they never once told me to stop messing around and get a proper job.

Contents

Acknowledgements	i
Table of Contents	ii
List of Figures	vi
List of Tables	ix
1 Introduction	1
1.1 Laser wakefield acceleration	2
1.1.1 History	4
1.1.2 EM waves in a plasma	5
1.1.3 Ponderomotive force	7
1.1.4 Linear wakefields	9
1.1.5 Non-linear wakefields	11
1.1.6 Wakefields in the blow-out régime	11
1.1.7 Applications to light sources	14
1.1.8 Limitations	15
1.2 Multi-pulse laser wakefield acceleration	17
1.3 Outline of project	20
2 Frequency-domain holography	23
2.1 Overview	24
2.2 Mathematical description	31
2.2.1 Spectral interferometry	31
2.2.2 Temporal phase reconstruction	33
2.3 Experimental setup	34
2.3.1 Probe chirp measurement	37
2.3.2 Reference chirp measurement	38
2.3.3 Proposed reference chirp measurement	39
2.3.4 Reference spectrum measurement	39
2.3.5 Practical problems	40
2.4 Analysis steps	43
2.4.1 Data input	44
2.4.2 Isolating the side-band	45
2.4.3 Spectral phase map	47
2.4.4 Reference spectrum division	51

2.4.5	Reference phase	52
2.4.6	Temporal phase	52
2.4.7	Reference alignment	53
2.5	Sensitivity	55
2.5.1	Noise model	56
2.5.2	Analytical prediction	58
2.5.3	Experimental value	60
2.5.4	Numerical prediction	65
2.5.5	Averaging	67
2.6	Practical difficulties	68
2.6.1	Sign of probe chirp	69
2.6.2	Spectral ripples	71
2.6.3	Imaging the probe correctly	73
2.6.4	Spectral variations	73
2.6.5	Systematic effects	75
2.6.6	Measuring references	76
2.7	Direct spectral recovery	77
2.7.1	Isolating the spectra	78
2.7.2	Simplifying the problem	80
2.7.3	Using a target spectrum	81
2.7.4	Spectrometer resolution	81
2.7.5	Summary	83
2.8	Other wakefield diagnostics	84
2.8.1	Spectral side-bands	85
2.8.2	Chirped-pulse spectral modulation	86
2.8.3	Comparison to FDH	89
3	Experimental set-up	91
3.1	Laser system	92
3.1.1	Spectrum	93
3.1.2	Autocorrelation measurements	94
3.1.3	Beam profile	98
3.2	FDH optics	99
3.2.1	Probe compressor	99
3.2.2	Pulse pair generation	101
3.2.3	Pulse spectra	102
3.2.4	Probe/reference chirp	105
3.3	Gas and vacuum equipment	107
3.4	Gaussian beam model	109
3.4.1	Modelling a beam as a Gaussian	111
3.5	Old interaction region	113
3.5.1	Layout	114
3.5.2	Transmitted energy	114
3.5.3	Pump focal profile	116
3.5.4	Gas cell	119
3.6	New interaction region	120
3.6.1	Design process	120

3.6.2	Layout	122
3.6.3	Transmitted energy	123
3.6.4	Pump focal profile	123
3.6.5	Pump-probe alignment	127
3.6.6	Gas cell	129
3.7	Post-IP diagnostics	130
3.7.1	Spectrometer	132
3.7.2	Imaging system and focussing	132
3.8	Data acquisition and analysis software	133
4	Simulations	136
4.1	Particle-in-cell simulations	137
4.1.1	EPOCH	138
4.2	Summary of laser parameters	139
4.3	Gas density	140
4.4	2D EPOCH simulations	143
4.4.1	Old IP	144
4.4.2	New IP as designed	145
4.4.3	New IP as measured	148
4.5	Fluid simulation	151
4.5.1	Different contributions to the wakefield	153
4.5.2	Finding the phase shift	154
4.5.3	Comparison with EPOCH simulations	154
4.6	3D EPOCH simulations	157
4.6.1	Electron density snapshots	158
4.6.2	Phase shift	159
4.7	Alignment sensitivity	161
4.7.1	Translational misalignment	162
4.7.2	Angular misalignment	163
4.7.3	Alignment summary	163
4.8	Frequency-domain holography simulation	164
4.8.1	Simulation parameters	165
4.8.2	Simulation with Gaussian spectrum	166
4.8.3	Simulation with experimental spectrum	168
4.9	Summary	170
5	Results	172
5.1	Old interaction region	172
5.2	New interaction region	175
5.2.1	Timing scan	177
5.2.2	Transverse phase map	179
5.3	Search for a wakefield	181
5.4	Summary	187

6	Pulse train generation	188
6.1	Multiple Michelson pulse splitter	189
6.1.1	Dispersion and non-linear effects	191
6.1.2	Pulse splitting before compression	194
6.2	Single-crystal pulse splitter	197
6.2.1	Simulations	199
6.3	Pulse train generation by pulse shaping	204
6.3.1	Brief history	205
6.3.2	Simple model	206
6.3.3	Optimisation	209
6.3.4	Spatial chirp	211
6.3.5	Finite beam size	216
6.3.6	Non-uniform pulse spacings	218
6.3.7	Interaction with amplification	219
6.4	Summary	221
7	Conclusions and future work	222
7.1	Conclusions	222
7.1.1	Experimental work	222
7.1.2	Pulse train generation	224
7.2	Future work	225
7.2.1	Further MP-LWFA experiments	225
7.2.2	FDH development	227
7.2.3	Laser development and pulse train generation	228
	Bibliography	229

List of Figures

1.1	Laser wakefield in the linear régime	10
1.2	Laser wakefield in the non-linear régime	12
1.3	Laser wakefield in the bubble régime	13
1.4	Laser wakefield excited by multiple pulses	18
2.1	The co-ordinate system used to describe FDH	25
2.2	Schematic of ionisation front	29
2.3	Schematic design of an FDH experiment	30
2.4	Side-bands corresponding to fringe pattern	32
2.5	Generation of probe and reference pulses	36
2.6	FDH data-taking — experimental setup	36
2.7	Probe chirp measurement — experimental setup	38
2.8	Reference chirp measurement — experimental setup	39
2.9	Proposed reference chirp measurement	40
2.10	Reference spectrum measurement — experimental setup	41
2.11	FDH analysis flowchart	42
2.12	Example frequency-domain hologram	45
2.13	FFT of an FDH spectrum	46
2.14	An example spectral phase map	48
2.15	Phase unwrapping in 1D	49
2.16	Phase unwrapping in 2D	50
2.17	An example temporal phase map	53
2.18	Definition of actual and predicted noise	60
2.19	Measured signal:noise ratio	62
2.20	Predicted phase error	63
2.21	Actual spectral phase error	63
2.22	Line-outs of spectral phase error	64
2.23	Actual temporal phase error	64
2.24	Simulated spectral phase error	66
2.25	Simulated temporal phase error	66
2.26	Ripples in spectral phase	72
2.27	Ambiguity when spectra cross	79
2.28	Simulated spectral modulation data	88
3.1	Spectrum of the pump pulse	93
3.2	Top-down view of autocorrelator	94
3.3	Side-on view of autocorrelator	94
3.4	Autocorrelation of the pump pulse	95

3.5	Estimated temporal profile	97
3.6	Pump beam profile	98
3.7	Probe compressor	100
3.8	Autocorrelation of the probe pulse	100
3.9	Pump and probe beam-lines	101
3.10	Spectra of the probe and reference pulses	102
3.11	Stability of the reference spectrum	103
3.12	Effects of compressor and SHG crystal on spectrum	104
3.13	Comparison of the pump SHG and reference spectra	105
3.14	Measurements of probe GDD	108
3.15	Measurements of reference GDD	108
3.16	Residuals from chirp fit	115
3.17	Original layout of interaction region	115
3.18	Profiles of the pump pulse near the focus	116
3.19	Beam sizes and M^2 fit	117
3.20	Gaussian beam model — fluence comparison	118
3.21	Original design of gas cell	119
3.22	Updated layout of interaction region	123
3.23	Profiles of the pump pulse near the focus	124
3.24	Beam sizes and M^2 fit	125
3.25	Gaussian beam model — fluence comparison	126
3.26	Gaussian beam model fit to fluence data	126
3.27	Pump-probe overlap	128
3.28	New gas cell	131
3.29	New gas cell photograph	131
4.1	The steps in a particle-in-cell simulation	138
4.2	Circular Gaussian beam model — fluence comparison	141
4.3	Snapshots of laser pulse in old IP	144
4.4	Snapshot of electron density — new IP as designed	145
4.5	Simulated beam sizes	146
4.6	Map of estimated phase shift	147
4.7	Estimated phase shift along y axis	147
4.8	Snapshot of electron density — new IP as measured	148
4.9	Simulated beam sizes	149
4.10	Map of estimated phase shift	149
4.11	Map of estimated phase shift — detail	150
4.12	Estimated phase shift along y axis	151
4.13	x-z snapshot from 3D EPOCH simulation	159
4.14	x-y snapshots from 3D EPOCH simulation	160
4.15	Sensitivity to translational misalignment	162
4.16	Sensitivity to angular misalignment	164
4.17	FFTs of simulated FDH data — Gaussian spectrum	167
4.18	Temporal phase recovered from simulated FDH data	168
4.19	FFTs of simulated FDH data — experimental spectrum	169
4.20	Temporal phase recovered from simulated FDH data	170

5.1	Old IP — Snapshot of phase shift	173
5.2	Old IP — Phase shift as a function of pressure	174
5.3	New IP — Snapshot of phase shift	176
5.4	New IP — Phase shift as a function of pressure	177
5.5	Results of timing scan	178
5.6	Map of effective interaction length	180
5.7	Possible wakefield signal at 38 mbar	182
5.8	Data at 40 mbar	183
5.9	Fourier transform of possible wakefield signal	184
5.10	Data at 21 mbar	184
5.11	Data at 25 mbar	185
5.12	Fourier transform of possible wakefield signal	186
5.13	Line-outs of candidate wakefield signals	186
6.1	Michelson interferometer used to generate pulse pair	189
6.2	Multiple Michelson interferometer	190
6.3	Pulse duration dependence on beamsplitter thickness	195
6.4	Single crystal method of pulse train generation	198
6.5	Simulated pulse train from single crystal method	201
6.6	Pulse spacing from simulated pulse train	201
6.7	Simulated pulse train with 20 nm bandwidth	202
6.8	Pulse spacing from 20 nm simulated pulse train	203
6.9	Schematic of a $4f$ line	205
6.10	Example SLM phase filter	211
6.11	Example 7-pulse train	212
6.12	Important parameters in a $4f$ line	212
6.13	The effect of SLM spatial chirp	215
6.14	The effect of SLM spatial chirp on pulse chirp	215
6.15	The effect of correcting for SLM spatial chirp	216
6.16	The effect of further optimisations	217
6.17	A pulse train with non-uniform spacing	219

List of Tables

3.1	Sellmeier parameters for N-BK 7 glass	106
3.2	Gaussian beam parameters for old pump focus	117
3.3	Gaussian beam parameters for new pump focus	124
3.4	Gaussian beam parameters from fluence fit	125
4.1	Summary of laser parameters	140
4.2	Summary of beam parameters for old IP	140
4.3	Summary of beam parameters for new IP	141
4.4	Summary of optimal densities	142
4.5	Summary of fluid simulation results	154
4.6	Summary of modified fluid simulation results	157
6.1	Dispersion values used for single crystal simulation	202
6.2	Parameters of modelled pulse shaper	214

1. Introduction

Ever since the discovery of the X-ray in 1895 [1] and the electron in 1897 [2], particle accelerators have been at the forefront of scientific discovery. Today, accelerators are valuable research tools in fields as diverse as high-energy physics [3], biochemistry [4], and even archaeology [5]. Beyond research, accelerators have many medical and industrial applications such as radiotherapy, radioisotope production, and electron beam machining. Not too long ago, nearly every home in the Western world contained its own accelerator, in the form of a cathode ray tube television set. Accelerators are an integral, but largely invisible, part of the modern world.

One major use of particle accelerators in research is as synchrotron light sources. When a charged particle follows a curved path due to the presence of an external magnetic field, it emits electromagnetic radiation, called synchrotron radiation. A modern (3rd generation) synchrotron light source includes undulators and wigglers, which consist of an alternating series of magnetic poles. These cause the accelerated electrons to “zig-zag” down the beam line, emitting synchrotron radiation as they do. The wavelength of the emitted radiation can range from the terahertz to hard X-ray régimes, depending on the energy of the electrons and on the strength and period of the magnetic structure. Most importantly, the emitted radiation can achieve extremely high brightness¹; this is particularly important in the X-ray régime, where it enables a number of useful experimental techniques that would not be practical with other X-ray sources.

¹Defined as number of photons per second, per mrad² of angular divergence, per mm² of beam area, per 0.1% bandwidth

However, these capabilities are only available at a limited number of facilities, which generally receive more requests for beam time than they can fulfil. More widespread availability of synchrotron light sources would provide a boost for many areas of research, but is limited by the significant cost of constructing and running them, largely due to the required size of the facilities. One way to reduce the size, and potentially the cost, of synchrotron light sources would be to use stronger electric fields to accelerate the electrons. However, the radiofrequency (RF) cavities used to provide this acceleration experience electrical breakdown at high field strengths, with the maximum field being limited to about 100 MV m^{-1} [6, p. 1230].

1.1 Laser wakefield acceleration

One way to overcome the limit of electrical breakdown is create an accelerating structure from a material that has already broken down — that is, a plasma. Such a structure can take the form of an electron plasma oscillation, sometimes called a Langmuir wave, in which the electrons oscillate due to the electrostatic restoring force provided by the quasi-static background of much more massive ions. In the case of a ‘warm’ plasma, the thermal motion of the electrons acts as an additional restoring force, and the result is a propagating wave governed by the Bohm-Gross dispersion relation [7, p. 494].

However, the case considered here will be the low-temperature limit, in which the thermal motion of the electrons is ignored. In this case the oscillation occurs at a characteristic plasma frequency ω_p regardless of the value of the wave vector $\mathbf{k}$. This causes the group velocity of the oscillation to be zero (given the usual definition $v_g = \frac{d\omega}{dk}$), so the oscillation does not propagate and cannot transport any energy. Instead, some kind of travelling driver can excite a plasma oscillation in place, where the phase velocity of the oscillation matches the velocity of the driver. This oscillation continues until it is damped, transferring its energy to the thermal motion of the plasma electrons.

The plasma oscillation contains a electric field component, providing the restoring force to the electrons. An initially stationary electron experiencing this force will sim-

ply move back and forth, participating in the plasma oscillation. However, an electron which is already moving at the phase velocity of the oscillation will always experience the same electric field, causing it to be accelerated. A plasma oscillation with a phase velocity close to c can therefore accelerate relativistic electrons over a considerable distance. This accelerating field can be orders of magnitude greater than the electrical breakdown limit of RF cavities; a plasma with density 10^{18}cm^{-3} can sustain an accelerating field of 100 GV m^{-1} [6, p. 1230].

Since the plasma oscillation must have a phase velocity close to c , the driver must also be highly relativistic. One option is to use a bunch of charged particles, accelerated using RF technology, which excite the plasma oscillation by their electrostatic field. This approach, known as plasma wakefield acceleration (PWFA), has been shown to allow a doubling of the original electron energy [8] and could be used as a kind of “plasma afterburner”, efficiently providing a significant boost to the energy of a conventional (RF) accelerator. However, the technique as demonstrated so far does not permit the orders-of-magnitude reduction in accelerator length envisaged earlier.

Another option for the driver is to use an intense laser pulse, since they too propagate at close to the speed of light². As the laser propagates through the plasma, its main effect on the free electrons is to cause them to “quiver” in the rapidly oscillating laser field. The interaction between this quiver motion and the electric and magnetic fields of the laser pulse gives rise to the ponderomotive force, which repels electrons from regions of high laser intensity. Conceptually, there is an energy associated with the quiver motion — the ponderomotive potential — so there is an associated force pushing electrons towards regions of lower laser intensity where this potential is reduced. It is this ponderomotive force which excites the plasma oscillation. Since the plasma oscillation trails the laser pulse, it is called a laser wakefield, by analogy with the wake of a boat, and the electrons are said to “surf” on it.

²In fact, the presence of the plasma causes the laser pulse to propagate at slightly less than c . As a result, the accelerated electrons will eventually out-run the accelerating phase of the wakefield and begin to be decelerated. This phenomenon, known as dephasing, is one of the factors limiting the acceleration length

1.1.1 History

The concept of laser wakefield acceleration was first proposed in 1979 by Tajima and Dawson [9]. However, laser technology at the time was not able to generate sufficiently short (less than the plasma period) pulses with enough intensity to drive a significant wakefield. Therefore, Tajima and Dawson suggested the use of a longer laser pulse, modulated at the plasma frequency, to resonantly excite a wakefield. This modulated pulse was generated by beating together two pulses of different frequencies, where the frequency difference $\Delta\omega = \omega_2 - \omega_1 = \omega_p$, and the scheme was thus called plasma beat wave acceleration (PBWA). Another approach was the self-modulated laser wakefield accelerator (SMLWFA) [10, 11], in which a single long laser pulse becomes modulated at the plasma frequency due to the self-modulation instability (a kind of Raman instability). Although electron acceleration was demonstrated under both schemes [11, 12], the accelerating field was limited by resonant detuning — as the wakefield amplitude increased, the plasma period increased due to non-linearities, so that the modulated driver was no longer resonant with the wakefield. In addition, they were susceptible to a number of laser-plasma instabilities due to the long duration of the pulses.

The development of chirped-pulse amplification (CPA) in 1985 made it possible to generate high-intensity ultrashort laser pulses, paving the way for the realisation of the laser wakefield accelerator (LWFA). Although early experiments demonstrated gradients in excess of 100 GV m^{-1} [13, 14], the energy distribution of the accelerated bunches was typically exponential, with only a minority of electrons being accelerated to high energy. It wasn't until the “dream beam” experiments of 2004 [15–17] that bunches of electrons were accelerated to $\sim 100 \text{ MeV}$ with low energy spread. With the 2006 development of plasma channels to guide the driving laser pulse over a longer distance, acceleration to 1 GeV was demonstrated [18].

More recent work on LWFA continues to push the energy frontier, with acceleration to over 3 GeV being demonstrated in 2013 [19] using separate injector and accel-

erator stages. More recently, acceleration to 4.25 GeV in a plasma channel was demonstrated by Leemans *et al.* using the BELLA laser [20]. Meanwhile, work is continuing to improve the quality of the accelerated bunches and control over their parameters via various methods of controlled injection [21–25].

1.1.2 EM waves in a plasma

Consider a plane electromagnetic wave propagating through a plasma, with frequency ω and wave vector $\mathbf{k}$. The electric and magnetic fields, as well as the current density due to the free electrons, can be defined in the standard way:

$$\mathbf{E} = \mathbf{E}_0 \exp i(\omega t - \mathbf{k} \cdot \mathbf{r}) \quad (1.1a)$$

$$\mathbf{B} = \mathbf{B}_0 \exp i(\omega t - \mathbf{k} \cdot \mathbf{r}) \quad (1.1b)$$

$$\mathbf{J} = \mathbf{J}_0 \exp i(\omega t - \mathbf{k} \cdot \mathbf{r}) \quad (1.1c)$$

Neglecting the effect of the magnetic field (valid in the non-relativistic limit), the current density and electric field can be related:

$$\frac{1}{en_e} \frac{\partial \mathbf{J}}{\partial t} = \frac{e}{m_e} \mathbf{E} \quad (1.2)$$

$$\mathbf{J} = -i \frac{e^2 n_e}{\omega m_e} \mathbf{E} \quad (1.3)$$

Making the appropriate substitutions into Maxwell's equations gives:

$$\nabla \cdot \mathbf{B} = 0 \quad \mathbf{k} \cdot \mathbf{B} = 0 \quad (1.4)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad \mathbf{k} \times \mathbf{E} = \omega \mathbf{B} \quad (1.5)$$

$$c^2 \nabla \times \mathbf{B} = \frac{\mathbf{J}}{\epsilon_0} + \frac{\partial \mathbf{E}}{\partial t} \quad c^2 \mathbf{k} \times \mathbf{B} = \left(\frac{e^2 n_e}{\epsilon_0 \omega m_e} - \omega \right) \mathbf{E} \quad (1.6)$$

Defining the plasma frequency ω_p :

$$\omega_p^2 = \frac{e^2 n_e}{\epsilon_0 m_e} \quad (1.7)$$

and combining Equations 1.5 and 1.6 gives:

$$c^2 \mathbf{k} \times (\mathbf{k} \times \mathbf{E}) = (\omega_p^2 - \omega^2) \mathbf{E} \quad (1.8)$$

In the case where $\omega \neq \omega_p$, this gives the dispersion relation for a transverse³ EM wave propagating through a cold plasma:

$$\omega^2 = \omega_p^2 + c^2 k^2 \quad (1.9)$$

Such a wave can only propagate if $\omega > \omega_p$. This is equivalent to the requirement that $n_e < n_c$, where $n_c = \frac{\omega_p^2 \epsilon_0 m_e}{e^2}$ is called the critical density. The refractive index η of the plasma is given by:

$$\begin{aligned} \eta &= \frac{c}{v_p} = \frac{ck}{\omega} \\ &= \sqrt{1 - \frac{\omega_p^2}{\omega^2}} \\ &= \sqrt{1 - \frac{n_e}{n_c}} \end{aligned} \quad (1.10)$$

If $\omega = \omega_p$, Equations 1.4 and 1.6 together imply that $\mathbf{B} = 0$. It can then be seen from Equation 1.5 that $\mathbf{k}$ and $\mathbf{E}$ are parallel. This longitudinal electrostatic field forms part of the plasma oscillation discussed earlier; the other component is an electron density oscillation. For small amplitudes, it is straightforward to relate the electron

³Equations 1.5 and 1.6 show that $\mathbf{k}$, $\mathbf{E}$, and $\mathbf{B}$ are mutually orthogonal.

density oscillation δn to the electric field $\mathbf{E}$ through the continuity equation:

$$\nabla \cdot \mathbf{J} = -\frac{\partial \rho}{\partial t} = e \frac{\partial n_e}{\partial t} \quad (1.11)$$

$$-i \mathbf{k} \cdot \mathbf{J} = i \omega_p e \delta n \quad (1.12)$$

$$k E = -i \frac{e}{\epsilon_0} \delta n \quad (1.13)$$

Setting $\delta n = n_e$ and $k = \frac{\omega_p}{c}$ gives the cold non-relativistic wave breaking field, which indicates how large an accelerating field a given plasma might be expected to sustain:

$$E_0 = \frac{c m_e \omega_p}{e} \quad (1.14)$$

1.1.3 Ponderomotive force

In a laser wakefield accelerator, a plasma oscillation is driven by the ponderomotive force of a short, intense laser pulse. The following derivation of the ponderomotive force approximately follows that given by Esarey *et al.* in [6, p. 1233], although I use SI units throughout.

Consider the motion of an electron in an electromagnetic field, described by:

$$\frac{d\mathbf{p}}{dt} = -e(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (1.15)$$

$$\mathbf{E} = -\frac{\partial \mathbf{A}}{\partial t} \quad \mathbf{B} = \nabla \times \mathbf{A} \quad (1.16)$$

where $\mathbf{A}$ is the vector electromagnetic potential. The contribution to $\mathbf{E}$ from the scalar potential ϕ has been omitted, since it is zero in the case of an electromagnetic wave. It is often convenient to work with the (dimensionless) normalised vector potential $\mathbf{a} = \frac{e\mathbf{A}}{m_e c}$.

Although the kinetic momentum $\mathbf{p}$ of this electron is not conserved, under certain conditions the canonical momentum $\mathbf{P} = \mathbf{p} - m_e c \mathbf{a}$ will be conserved. As a concrete example, a homogeneous but time-varying vector potential $\mathbf{a}$ gives rise to a homoge-

neous and variable electric field, in which the canonical momentum $\mathbf{P}$ is conserved. This is an approximation to a slowly-varying laser field, provided the magnetic field component of the laser can be neglected ($v \ll c$).

If the electron is initially stationary ($\mathbf{p} = 0$) before the laser field turns on (when $\mathbf{E} = \mathbf{A} = 0$), its canonical momentum $\mathbf{P}$ will be 0. Since the canonical momentum is conserved in this case, it will remain 0 after the laser field is turned on. At that point, the kinetic momentum of the electron is given by:

$$\mathbf{p} = m_e c \mathbf{a} + \mathbf{P} = m_e c \mathbf{a} \quad (1.17)$$

This is termed the quiver momentum.

When the laser field varies in space as well as time, the canonical momentum is no longer conserved. In the non-relativistic limit $a \ll 1$ (so-called because the quiver motion of the electron is non-relativistic), the quiver momentum $\mathbf{p}_q = m_e c \mathbf{a}$ remains the leading-order term in the kinetic momentum. But the changing canonical momentum gives rise to another term:

$$\mathbf{p} = m_e c \mathbf{a} + \mathbf{P} = \mathbf{p}_q + \delta\mathbf{p} \quad (1.18)$$

The additional term $\delta\mathbf{p}$ is equal to the total canonical momentum $\mathbf{P}$; the rate of change of $\delta\mathbf{p}$, which is identified as the ponderomotive force, is also the rate of change of the canonical momentum:

$$\mathbf{F}_p = \frac{d\delta\mathbf{p}}{dt} = \frac{d\mathbf{P}}{dt} \quad (1.19)$$

Instead of considering a single electron, we now consider the electrons in a plasma as a cold fluid. This involves expanding $\frac{d}{dt}$ in terms of partial derivatives:

$$\frac{d}{dt} = \frac{\partial}{\partial t} + (\mathbf{v} \cdot \nabla) \quad (1.20)$$

By assuming that $\delta p \ll p_q$, we can solve for the leading-order term of the pondero-

motive force $\mathbf{F}_p$:

$$\begin{aligned}\frac{d\mathbf{p}}{dt} &= \frac{d\mathbf{p}_q}{dt} + \frac{d\delta\mathbf{p}}{dt} \\ &= \frac{\partial\mathbf{p}_q}{\partial t} + (\mathbf{v} \cdot \nabla)\mathbf{p}_q + \frac{d\delta\mathbf{p}}{dt} \\ &= -e\mathbf{E} - e\mathbf{v} \times \mathbf{B}\end{aligned}\tag{1.21}$$

$$\frac{d\delta\mathbf{p}}{dt} = \mathbf{F}_p = -(\mathbf{v} \cdot \nabla)\mathbf{p}_q - e\mathbf{v} \times \mathbf{B}\tag{1.22}$$

$$\mathbf{F}_p = -m_e c^2 [(\mathbf{a} \cdot \nabla)\mathbf{a} + \mathbf{a} \times (\nabla \times \mathbf{a})] = -\frac{1}{2}m_e c^2 \nabla a^2\tag{1.23}$$

where the last step uses the approximation $\mathbf{v} \approx \frac{\mathbf{p}_q}{m_e} = c \mathbf{a}$. This is the 3D ponderomotive force in the non-relativistic limit, which takes the form of the gradient of a ponderomotive potential.

1.1.4 Linear wakefields

The laser wakefield accelerator is based on the combination of these two concepts: the ponderomotive force of a propagating laser pulse displaces electrons from equilibrium, exciting a plasma oscillation, referred to as a wakefield in this context. The longitudinal electric field of the wakefield is used to accelerate relativistic electrons, and is able to do so efficiently as the phase velocity of the wakefield is close to the speed of light.

Esarey *et al.* [6, p. 1233] present this more concretely, with the electron density oscillation $\delta n = n_e - n_0$ varying in the linear régime according to:

$$\left(\frac{\partial^2}{\partial t^2} + \omega_p^2 \right) \frac{\delta n}{n_0} = \frac{c^2}{2} \nabla^2 a^2\tag{1.24}$$

Note that the second term on the left-hand side represents the restoring action of the plasma oscillation, while the right-hand side represents the action of the ponderomo-




FIGURE 1.1 – A schematic of a laser wakefield in the linear (non-relativistic) régime. **Top:** the electron density behind a laser pulse propagating from left to right. The electric fields associated with this structure are also shown. **Bottom:** the strength of the accelerating and focussing forces felt by the electrons, as functions of the distance behind the laser pulse. The shaded regions indicate where these are greater than zero, so that the electrons are accelerated and focussed.

tive force driving the oscillation. This can be solved to give:

$$\frac{\delta n}{n_0} = \frac{c^2}{2\omega_p} \int_0^t dt' \sin[\omega_p(t - t')] \nabla^2 a^2(\mathbf{r}, t') \quad (1.25)$$

If the temporal extent of the laser pulse is finite, it is possible to select a t_{max} such that $a^2 = 0$ for all $t' > t_{max}$. Then for $t > t_{max}$ this results in a simple sinusoidal oscillation whose amplitude depends on the Fourier transform of the temporal profile of the driver a^2 [26]. The transverse profile of the driver also affects the shape of the wakefield, resulting in transverse electric fields in addition to the longitudinal accelerating fields; these transverse fields can be used to focus the accelerated electrons, resulting in the same betatron motion seen in conventional accelerators.

The structure of a linear wakefield is shown in Figure 1.1. This shows the oscillating electron density behind a laser pulse, which takes the form of a sine wave. The direction of the electric field associated with this structure is also shown. In certain regions, this electric field points out and backwards, so that electrons are accelerated forwards and inwards. This simultaneous acceleration and focussing takes place only over a quarter of each plasma wavelength, as shown by the bottom half of the same figure.

1.1.5 Non-linear wakefields

As the amplitude of the wakefield increases, it becomes non-linear, due to the relativistic increase in the inertial mass of the moving electrons. This distorts the shape of the electron density profile, causing the peaks to steepen and narrow. It also causes the period of the oscillation to increase with amplitude. This is the cause of the resonant detuning that affects plasma beatwave acceleration — as the wakefield amplitude increases, its frequency decreases below the beat frequency that was driving it.

Another effect of non-linearity is shown in Figure 1.2. Since the intensity of the laser pulse decreases with increasing radius, the amplitude of the wakefield does too. In the non-linear régime, this causes the plasma wavelength to be longer on-axis than off-axis, causing the wavefronts to become increasingly curved as time passes. In the plot, this manifests as wavefronts further behind the laser pulse being more curved, as they have had longer to oscillate since they were excited by the laser pulse. Another contribution to this effect is the radial expulsion of electrons by the ponderomotive force, which increases the average electron density off-axis, increasing the plasma frequency.

1.1.6 Wakefields in the blow-out régime

A single laser pulse of sufficient intensity is able to completely expel all nearby plasma electrons, resulting in a region of zero electron density propagating behind. First ob-




FIGURE 1.2 – A laser wakefield in the non-linear régime. Data from a 2D PIC simulation (see § 4.1) with $a_0 = 1$. Note the increasing curvature of the wavefronts, due to the dependence of plasma wavelength on radial position.




FIGURE 1.3 – A laser wakefield in the bubble régime. Data from a 2D PIC simulation (see § 4.1) with $a_0 = 8$. A number of electrons have undergone self-injection, and have formed a bunch (indicated by an arrow) which is being accelerated and focussed inside the bubble.

served in simulations by Pukhov and Meyer-ter-Vehn [27], this is referred to as the blow-out or bubble régime (depending on whether or not a plasma wave is present behind the bubble), and has several beneficial properties for electron acceleration. The accelerating electric field is approximately constant with radius, while the focussing field is linear; this is the case (ideally) with an RF accelerator, and it ensures that the normalised emittance of the bunch is conserved during acceleration.

In addition, a wakefield in the blow-out régime is able to undergo self-injection, in which ambient plasma electrons are injected into the accelerating phase of the wakefield. This can be seen in Figure 1.2. This obviates the need for a separate injection mechanism, but the non-linear nature of the process causes significant shot-to-shot fluctuation.

1.1.7 Applications to light sources

There is a great deal of interest [28–33] in the application of LWFA to the generation of high-brightness EM radiation, principally X-rays, due to the compactness of the system compared with a typical synchrotron light source. In fact, LWFA is well suited to use as a light source, since the duration and source size of the X-ray pulses are determined by the small scale of the wakefield; this results in a high peak brightness.

The simplest method of generating radiation using LWFA is simply to allow the electrons to execute betatron oscillations in the focussing electric field present inside the wakefield. This betatron radiation has already been used to image biological specimens through phase contrast imaging, made possible by the micrometre-scale source size of the radiation [30]. Coherence of the radiation might be enhanced by injecting electrons off-axis, resulting in collective betatron oscillations [34]. In fact, whereas for most applications it is desirable for the accelerated bunch of electrons to have a small transverse size and to remain on-axis, in this case it is necessary to use either a large on-axis bunch or a small off-axis one in order to emit betatron radiation.

Another compact method of generating X-rays is to use inverse Compton scattering [35–37], which is related to Thomson or Compton scattering. In Compton scattering, a photon is scattered by a free electron at rest (or nearly so), transferring part of the photon’s energy to the electron. Thomson scattering is the same effect, but without a significant transfer of energy. Inverse Compton scattering is simply Compton scattering viewed from a frame of reference in which the kinetic energy of the electron is initially much greater than the energy of the photon, so that energy is transferred from the electron to the photon.

This is typically realised by using a laser to generate visible-wavelength photons, which are scattered from the accelerated electrons to become X-rays travelling in the same direction as the electrons. In the case of low laser intensity ($a_0 \ll 1$) and a head-on laser-electron collision, the wavelength of the scattered photons is given by

[37]:

$$\lambda_x = \frac{\lambda_L}{4\gamma^2} \quad (1.26)$$

where λ_L is the laser wavelength, γ is the relativistic factor of the electrons, and λ_x is the wavelength of the scattered photons. The small value of λ_L (hundreds of nm) allows relatively modest energies to be used while still reaching X-ray wavelengths. This technique has been used in conjunction with conventional accelerator technology to construct compact X-ray sources [38]. The approach has also been shown to work with LWFA [32], resulting in an even more compact, all-laser X-ray source.

Alternatively, the accelerated electrons can be passed through an undulator or wiggler, as in a conventional synchrotron light source. This has already been used to generate soft X-rays [29]. In principle, this method could even be used to build a free-electron laser (FEL), in which microbunching of the accelerated electrons allows the emitted radiation to be amplified coherently, resulting in significant enhancements to peak brightness. A table-top FEL based on LWFA appears feasible [28], but the required energy spread of the accelerated electrons is difficult to achieve. This problem might be overcome by using a transverse-gradient undulator [31], in which each energy slice is exposed to a different magnetic field within the undulator, so that each slice amplifies at the same wavelength.

1.1.8 Limitations

Despite the many advantages of LWFA over conventional accelerators, there are a number of challenges still to be overcome before it can be truly competitive. In particular, shot-to-shot stability and the energy spread are both typically poor by the standards of RF accelerators. These problems are gradually being overcome, with finer control over plasma parameters to improve stability and the use of controlled injection to reduce energy spread.

But the underlying laser technology poses problems of its own. Most LWFA experiments use titanium-sapphire (Ti:sapph) lasers for their ability to produce ultrashort

pulses with very high peak powers. However, these lasers typically suffer from both low wall-plug efficiency (typically less than 1%) and low repetition rate (often less than 10 Hz) [33]. The main cause of these problems is the build up of heat in the laser gain medium, wasting energy and requiring the average power to be limited so that the system can cool sufficiently between pulses. The low repetition rate is a particular impediment to the use of LWFA in a practical light source, as it reduces the average photon flux reaching the users, requiring more experimental time to collect a given quantity of data.

Practical LWFA seems to require a different laser technology to increase the average power available. This can be achieved through a different geometry of the gain medium, improving cooling. Examples include the thin-disk laser, in which the gain medium comprises a thin disk of material mounted on a heat sink, and the fibre laser, in which the doped core of a single-mode optical fibre is used as the gain medium. Both these technologies allow heat to be removed from the gain medium efficiently through their large surface area to volume ratios.

High power fibre lasers [39, 40] typically use ytterbium-doped glass (Yb:glass) as the gain medium, which allows for high power efficiency (around 50%) and low heat load due to its small quantum defect⁴. However, the peak power generated by fibre lasers is limited by the small mode area within the fibre, so that relatively small peak powers produce high intensities, resulting in damage or detrimental non-linear effects. Photonic crystal fibres (PCFs) allow for larger mode areas while remaining single mode [41–43] and coherent combination permits several fibre amplifiers to be used in parallel, effectively increasing the mode area further. [44–47] Nevertheless, it would be useful if a significant laser wakefield could be excited by a driver with a more modest peak power.

⁴The difference between pump and laser photon energies, which generally ends up as heat.

1.2 Multi-pulse laser wakefield acceleration

One possible method of generating a significant laser wakefield with lower peak power is to resonantly excite the wakefield using multiple pulses, spaced by the plasma period, as shown in Figure 1.4. This is the same basic concept as plasma beatwave acceleration, but the use of multiple pulses allows for greater control over pulse spacing to compensate for resonant detuning⁵. The concept of multi-pulse laser wakefield acceleration (MP-LWFA) was first studied theoretically during the 1990s [48–52], when it was shown to significantly increase the amplitude of a non-linear wakefield, compared to a single laser pulse of the same total energy [53]. However, the effect was never demonstrated experimentally and interest waned over the following years. More recently, interest in MP-LWFA has been renewed due to the possibility of using it, in conjunction with fibre laser technology, to drive a compact radiation source with kHz repetition rates [33].

One key advantage of MP-LWFA over conventional single-pulse LWFA is the lower peak laser power incident on various optics. This reduces the risk of optical damage, which is a major problem when working with high-intensity lasers. Alternatively, the size of the optics can be reduced whilst keeping the peak intensity the same. This is especially useful in the case of the final focussing optics, since the smaller size allows a shorter focal length to be used when coupling into the plasma; this might dramatically reduce the separation between acceleration stages in a multi-stage accelerator [54] which would otherwise significantly reduce the average accelerating gradient. In addition, the use of an optimised pulse train has been shown to allow significantly stronger wakefields to be driven with the same total laser energy in the non-linear régime [53].

The use of multiple pulses also provides more control over the acceleration process. For example, a second pulse train, out of phase with the first, would remove energy from the wakefield; this would reduce the energy dissipated into the plasma as heat.

⁵The plasma frequency decreases as amplitude increases, so that a perfectly periodic driver moves out of resonance.




FIGURE 1.4 – A schematic of a laser wakefield resonantly excited by four pulses. **Top:** the electron density within and behind a short pulse train propagating from left to right. **Bottom:** the strength of the accelerating and focussing forces felt by the electrons. Note how they increase from zero to their maximum values over the length of the pulse train.

Such a scheme would be needed to reduce the thermal load on the plasma channel in many future high repetition rate applications. As an extreme example, the 1 TeV electron-positron collider discussed in [54] would have a repetition rate of 15 kHz with 32 J laser pulses, resulting in a heat load on the order of 150 kW; this is in a 2 m long capillary. The same article mentions the possibility of reducing this load by using a second anti-resonant laser pulse.

Another possible modification is to use a single intense pulse following the driver pulse train to perform ionisation injection [55], as well as one at the beginning to pre-ionise the plasma. In conjunction with a programmable pulse train, this would allow controlled injection into any period of the wake, or multiple periods, and easy control over the wakefield intensity and final electron energy.

Various considerations apply to MP-LWFA that do not apply to single-pulse LWFA. For example, the wakefield generated by earlier pulses affects the path taken by subsequent pulses (due to the transverse variation in electron density), which may lead to a “hosing” instability. Tolerances on the pulse separation must be considered, and the increase in required pulse spacing with wakefield amplitude must be accounted for. The resonance condition makes the wakefield amplitude sensitive to the plasma density, especially in the case of long pulse trains, so additional care must be taken to control the plasma density. The much longer time-scales involved in MP-LWFA compared with LWFA allow significant ion motion to occur; this has a serious effect on the structure of the wakefield, and limits the number of pulses that can be used. These problems, among others, are the subject of ongoing study [33].

The work described in this thesis forms part of an ongoing project which aims to demonstrate the MP-LWFA concept experimentally. This project itself forms part of a larger body of work to develop the technology and theoretical knowledge needed to construct a practical radiation source based around a multi-pulse laser wakefield accelerator.

1.3 Outline of project

There are two challenges to be overcome before work on a practical, compact accelerator based on MP-LWFA can begin. Firstly, the MP-LWFA concept itself must be demonstrated — despite the existing body of theoretical work, it has not been the subject of any experiments. Secondly, the technology needed for a dedicated, high efficiency, multi-pulse driver laser must be developed. My work focussed mainly on the first of these challenges.

Since a dedicated driver laser is not yet available, an interim solution for pulse train generation was needed for the MP-LWFA demonstration experiment. This was to be achieved by taking the single-pulse output of a conventional TW-class Ti:sapphire laser and splitting it up. This pulse train was then to be focussed to a small size at an interaction point (IP), positioned inside a helium-filled gas cell; here it would ionise the helium to form a plasma, in which it would then resonantly excite a wakefield.

Most LWFA experiments measure the generated wakefield indirectly, by detecting and analysing the accelerated electrons using an electron spectrometer, but this requires some method of injecting electrons into the wake. The usual method for this is self-injection in the bubble régime [56], where electrons from the background plasma are injected into the back of the bubble and accelerated. However, MP-LWFA is designed to operate in the linear or quasi-linear régimes, preventing the use of self-injection. Other controlled injection methods exist, such as ionisation injection [57–60], colliding-pulse injection [61, 62], and density gradient injection [63]; however, these were considered to be too complex for this early stage in the development of MP-LWFA, so another method was needed.

In addition, future accelerators based on MP-LWFA are likely to need some kind of diagnostic to monitor the build-up of the wakefield over multiple pulses, in order to optimise the pulse train and identify problems. Therefore, the realisation and use of such a diagnostic was another aim of this project. The method selected was frequency-domain holography (FDH), due to its single-shot measurement capability and ability to

resolve individual periods of the plasma wave. Although transverse plasma shadowgraphy [64, 65] also has these capabilities, the very short laser pulses required were not available. Early calculations suggested that FDH would be capable of measuring the small phase shifts expected. However, later work would show that these calculations underestimated the level of noise.

FDH was originally developed in the late 1990s by Le Blanc *et al.* [66], with later work demonstrating that it could be used to image laser wakefields in single-shot mode [67, 68]. A probe pulse co-propagates with the wakefield and accumulates a phase shift due to the refractive index of the plasma. This phase shift varies across the pulse, both transversely and longitudinally, according to the electron density of the portion of the wakefield at that position (in the co-moving frame). A separate reference pulse propagates in front of the wakefield, such that it experiences no significant phase shift. These pulses are then directed into an imaging spectrometer, where they interfere to produce fringes in the frequency domain. Analysis of this fringe pattern, together with the appropriate calibration data, allows recovery of the phase shift that was imprinted on the probe pulse, and hence an electron density map of the wakefield.

As far as the author is aware, despite the achievements in imaging laser wakefields by the original developers of FDH, and a number of more advanced variants [69, 70], no other group has implemented the technique. As well as the difficulties involved in dealing with a fairly novel experimental method, the wakefields I was planning to measure were significantly weaker than those measured previously with FDH, due to the use of a laser with an order of magnitude lower peak power, pushing the limits of the technology. The application of FDH to this project was a major focus of my work, requiring the construction of experimental apparatus, development of software and procedures needed for analysis, and an understanding of the capabilities and limitations of the technique. This work forms the subject of Chapter 2, which also includes a more complete description of FDH.

Once a wakefield had been driven by a pulse train and measured using FDH, the

aim was to verify that it behaved as predicted. Such tests might include varying the pulse spacing within the train while leaving the plasma density fixed, or vice versa, which should result in the amplitude of the resulting wakefield changing as the pulse spacing and plasma period move in and out of resonance.

I began work on this project using a TW-class laser at the University of Oxford, with the first aim being to generate a wakefield using a single pulse and image it using FDH. This process took longer than anticipated, due to the complexities of reimplementing FDH and the requirement to change the design of the IP. Due to the fact that the available laser energy and sensitivity of FDH were both less than expected, I was unable to observe definitive evidence of a wakefield. It was therefore decided to continue the project using the higher-energy ASTRA laser at the Rutherford Appleton Laboratory, with the results of my work informing the design of the experiment.

The setup of the initial experiment is described in Chapter 3. I performed various simulations based on the parameters of that experiment, which are described in Chapter 4, and the results I obtained are presented in Chapter 5. In Chapter 6 I discuss possible methods of pulse train generation, including methods which were considered for use in the proof-of-concept experiment and a proposed method that might be used in a dedicated MP-LWFA driver laser. Finally, I present conclusions and ideas for future work in Chapter 7.

2. Frequency-domain holography

The frequency-domain holography (FDH) technique is a descendent of the earlier frequency-domain interferometry (FDI), in which the phase difference between a probe pulse and a reference pulse is measured by detecting the shift in the spectral fringes produced when the two pulses interfere. This technique had been used to image a snap-shot of a solid-density plasma in a single shot [71] and was later used to probe a laser wakefield [26, 72] by performing multiple measurements at different times behind the pump pulse.

FDH was developed at the University of Texas at Austin, with the aim of using a long probe pulse to measure the phase at multiple time delays in a single shot. Early experiments [66] reduced the bandwidth of the probe pulse in order to increase its duration, and combined the probe and reference using a Michelson interferometer. Later experiments which succeeded in imaging laser wakefields [67, 68] introduced a more stable interferometer to generate the probe-reference pair, and applied a linear chirp to the pulses to lengthen them instead of limiting the bandwidth.

The same group developed a number of variations of FDH, to work in different régimes or to provide more resolution. Frequency-domain shadowgraphy [68, 73] measures the reconstructed amplitude of the probe pulse, and is designed to operate in high-density régimes where the probe suffers significant refraction due to the plasma, turning into optical “bullets”. Frequency-domain streak camera [68, 69, 74] is designed to resolve the evolution of the wakefield, instead of averaging over the propagation distance z as FDH does. Finally, frequency-domain tomography [68, 70]

uses multiple frequency-domain streak cameras at different angles to provide a tomographic reconstruction of the 3D structure of the wakefield.

Despite the relatively long history of FDH, the literature does not record any implementations by other groups. Therefore, I believe that this is only the second independent implementation of the technique. This chapter presents a detailed account of how FDH works, how I implemented it, the unforeseen problems I encountered and how I solved them, and ideas for future work to improve the technique. Whereas the improvements described above have focussed on extending the capabilities of FDH, the changes I propose focus on simplifying the experimental requirements, hopefully allowing FDH to be more widely adopted.

2.1 Overview

The principle of FDH is that a probe pulse is propagated behind the pump pulse, overlapping the wakefield generated by the pump. As it propagates, it accumulates a phase shift dependent on the refractive index of the plasma, which depends on the electron density. The key point is that there is very little relative motion between the probe and pump pulses, and hence also between the probe pulse and the wakefield. A given point in the probe pulse can be characterised by (t, x, y) co-ordinates, where t is the delay relative to the centre of the pulse and (x, y) are the axes perpendicular to the direction of propagation. The position of this point in space varies with the propagation distance z , as depicted in Figure 2.1. Since the wakefield and probe pulse are co-moving, each such point experiences the refractive index associated with a particular point in the wakefield, regardless of propagation distance z .

In particular, the accumulated phase shift ϕ depends on the optical path length, so it can be calculated by integrating the refractive index η over the propagation distance z . Both are functions of (t, x, y) , as different parts of the probe pulse experience different electron densities and thus different phase shifts. The refractive index depends on the electron density n_e and the critical density n_c , which is a function of the laser




FIGURE 2.1 – The co-ordinate system used in this description of FDH. The transverse position (perpendicular to the direction of propagation) is described by x and y . The distance that the probe pulse has propagated is z . The temporal profile of the pulse is given in terms of t , which measures longitudinal offset from the centre of the pulse. The longitudinal position of such a point within the laboratory frame varies as the pulse propagates, and can be given by $z' = z - ct$. The probe and reference pulses are separated by a time Δt , or (in vacuum) by a distance $c\Delta t$.

frequency $\omega = 2\pi f$ (see Equation 1.10).

$$\phi = \frac{2\pi}{\lambda} \int \eta \, dz \quad (2.1)$$

$$\eta = \sqrt{1 - n_e/n_c} \quad (2.2)$$

$$n_c = \frac{\epsilon_0 m_e \omega^2}{e^2} \quad (2.3)$$

The phase shift depends on ω as well as n_e . Given that the bandwidth of the laser is small compared to its central frequency, the main dependence of the phase shift will be on n_e . I will therefore approximate ω as a constant, equal to the central frequency of the laser pulse. In addition, the portion of the phase shift due to propagation through vacuum is uninteresting, so it is subtracted to leave the additional phase shift caused by the plasma $\Delta\phi = \phi - kL$ (where k is the wave number and L is the total propagation distance). In the limit of an under-dense plasma, $n_e \ll n_c$ (equivalent to $\omega \gg \omega_p$), this

gives:

$$\eta \approx 1 - \frac{n_e}{2n_c} \quad (2.4)$$

$$\Delta\phi \approx -\frac{\pi}{\lambda n_c} \int n_e dz \quad (2.5)$$

And thus the phase shift is proportional to the electron density within the wakefield, integrated over z . This integration results in any evolution of the wakefield being averaged out. This limits FDH to imaging non-evolving wakefields, although more sophisticated techniques [68–70] provide z -resolution.

In particular, if the probe pulse propagates at a different group velocity than the pump pulse (e.g. because it has a different wavelength) the wakefield will appear to be moving temporally with respect to the probe pulse, resulting in a temporal averaging of the phase shift. This acts to limit the useful interaction length between the probe and the wakefield, since if this temporal averaging occurs over more than half a plasma period it will suppress the wakefield signal. The walk-off distance over the interaction length L is given by:

$$\Delta z = \left(1 - \frac{v'_g}{v_g}\right) L = \left(1 - \frac{\eta'}{\eta}\right) L \quad (2.6)$$

where $v_g = \frac{d\omega}{dk} = c^2 v_p^{-1} = \eta c$ (see the dispersion relation – Equation 1.9) is the group velocity of the pump pulse, and v'_g is that of the probe pulse. Considering the case where $\omega' = 2\omega$ (which was the case in previous FDH experiments) and requiring $\Delta z \lesssim \frac{cT_p}{2}$ (where $T_p = \frac{2\pi}{\omega_p}$ is the plasma period) gives the condition:

$$L \lesssim \frac{4}{3} \frac{n_c}{n_e} cT_p \quad (2.7)$$

Since there is no way of measuring the phase of the probe pulse directly, FDH works by measuring the phase of the probe relative to a reference pulse, by interfering them; this requires that both pulses are generated from the same source, so that they

are coherent (having a definite phase relationship). Part of this phase difference is due to the refractive index of the wakefield. The residual phase difference, in the absence of a wakefield, can be measured separately and subtracted from the data. Typically, the reference pulse propagates along exactly the same beam path as the probe, but is temporally separated, so that it propagates in front of the pump pulse.

The temporal resolution of the measurement is limited by the bandwidth of the reference pulse. Any temporal features of the modulated probe that are smaller than the bandwidth limit will result in new spectral components appearing outside the original spectrum. These cannot interfere with the reference pulse, so they will not provide any useful data. The spectrum of the pulse must therefore be broad enough that a bandwidth-limited pulse would be as small as the smallest feature of interest. Since the probe pulse itself must be several plasma periods long, to measure them all, this implies that the probe and reference pulses are not bandwidth-limited. In fact, they are chirped by propagating them through a length of glass before entering the interaction region, increasing their duration while preserving their bandwidth.

After leaving the interaction region, the pulses are interfered by sending them into an imaging spectrometer. In the same way that two separate frequencies will produce beats or fringes in the time domain, two temporally separated pulses produce fringes in the frequency domain¹. The phase of these fringes, that is the position of the peak of each fringe, depends on the spectral phase difference between the two pulses. This spectrum is called a frequency-domain hologram, which I abbreviate as hologram.

The phase difference can be calculated by analysing the fringe pattern and can be plotted as a function of (ω, y) , where the slit at the entrance to the imaging spectrometer has selected and averaged over a small range of x . It is important to note that the phase shift caused by the wakefield occurs in the time domain, but the phase difference between the two pulses is measured in the frequency domain. A linear chirp is

¹Note that the pulses will interfere in the spectrometer even though they do not temporally overlap outside it. The pulses are stretched by the diffraction grating inside. This is required by the fact that time and frequency are conjugate variables, so as the spectrometer measures frequency accurately it necessarily stretches the pulses.

applied to the probe and reference pulses to ensure there is a strong correlation between the two, but it is not exact. In particular, the temporal phase shift affects the spectral amplitude and phase, so both must be measured if $\phi(t)$ is to be reconstructed. Note that there is assumed to be no absorption within the plasma, so that the temporal amplitude is unchanged.

In addition to the hologram, we require some other data for a complete reconstruction. A null hologram is a hologram taken in the normal way, but without a pump pulse. It therefore encodes a uniformly zero phase shift, and is used to subtract any background phase from the data. Data characterising the probe and/or reference pulses is also needed to calculate the temporal phase given the spectral phase; these are referred to as references.

In addition to the wakefield itself, a feature of interest in much of the data is the ionisation front. This arises due to fact that the pump pulse ionises the helium gas into plasma as it passes. The refractive index of this background plasma applies a phase shift to the probe pulse, but only to that portion propagating behind the pump. Therefore, if the probe pulse overlaps the pump pulse, there will be a sudden jump in phase at the position of the pump pulse, as pictured in Figure 2.2. This phase change is extremely sudden due to the highly non-linear dependence of ionisation rate on laser intensity.

In § 2.2 I cover the mathematics of FDH in more detail, including a rough outline of the analysis method. In § 2.3 I give an overview of the experimental setup needed to record frequency-domain holograms and their associated references. In § 2.4 I cover the analysis method in detail, including a description of the software I wrote to automate it and the algorithms I used. In § 2.5 I compare predictions and measurements of the sensitivity of FDH. In § 2.6 I describe various problems I encountered while implementing FDH, some of which I managed to solve. In § 2.7 I describe a possible extension to FDH, which may make measurements easier. Finally, in § 2.8, I describe alternative methods of measuring a wakefield that I investigated.




FIGURE 2.2 – A schematic diagram of the expected phase shift due to an ionisation front and a wakefield. A 2D slice is taken through the plasma in the longitudinal/vertical plane. The refractive index, measured in the co-moving frame, is integrated over a certain propagation length (about 1 mm here), to give a phase shift. This is plotted as a function of the vertical distance y and the delay t from the centre of the probe pulse to which this phase shift is applied.

The position of the pump pulse is marked at $t = 0$. At $t > 0$ the probe experiences a phase shift due to the plasma, which is ionised to either the He^+ or He^{2+} level depending on the intensity of the pump pulse. On axis, the pump pulse is intense enough to excite a wakefield. At $t < 0$, the probe is propagating through unionised helium gas, so experiences no phase shift relative to the reference.




Image courtesy of James Cowley.

FIGURE 2.3 – Schematic design of an FDH experiment. A pulse arrives from the left. A small amount is split off to act as the probe/reference; the optics at the top are described in more detail in § 2.3. The probe and pump are recombined at the dichroic mirror, before being sent into a vacuum chamber. There they are focussed by a parabolic mirror into a helium gas cell, where the pump ionises the helium and excites a wakefield. The probe is separated from the pump by another dichroic mirror and analysed by an imaging spectrometer.

2.2 Mathematical description

Consider two pulses propagating co-linearly with the pump pulse: a reference pulse in front and a probe pulse behind, separated by a time Δt . The spectra of the two pulses are given by $S_{ref}(\omega)$ and $S_{probe}(\omega)$; in general they are not Gaussian, and may be slightly different due to being generated separately. Both pulses will be chirped, having a spectral phase $\phi_{ch}(\omega)$. Due to the method used to generate the pulse pair, the probe pulse will have a relatively small additional chirp $\Delta\phi_{ch}(\omega)$; this will be explained in more detail in § 2.3,

The phase difference between the two pulses will be dominated by a linear term ($\Delta t \omega$) due to the different group delays. Assuming identical polarisation, the electric fields of these two pulses can be represented in the frequency domain as:

$$\tilde{E}_{ref}(\omega) = \sqrt{S_{ref}(\omega)} \exp i\phi_{ch}(\omega) \quad (2.8)$$

$$\tilde{E}_{probe}(\omega) = \sqrt{S_{probe}(\omega)} \exp i[\phi_{ch}(\omega) + \Delta t \omega + \Delta\phi_{ch}(\omega)] \quad (2.9)$$

2.2.1 Spectral interferometry

If these pulses are analysed by a spectrometer, they will interfere, resulting in a spectrum of the following form:

$$\begin{aligned} I_s(\omega) &= |\tilde{E}_{ref}(\omega) + \tilde{E}_{probe}(\omega)|^2 \\ &= S_{ref} + S_{probe} + 2\sqrt{S_{ref}S_{probe}} \cos[\omega\Delta t + \Delta\phi_{ch}(\omega)] \end{aligned} \quad (2.10)$$

As well as the simple sum of the spectra, there is a fringe pattern of period $\frac{2\pi}{\Delta t}$ with an ω -dependent phase shift. We can recover this phase shift from the fringe pattern




FIGURE 2.4 – Fourier transforming a spectrum featuring fringes, as shown in (a), will result in pair of side-bands either side of the base-band, as shown in (b). A band-pass filter can isolate one of the side-bands, thus separating the fringe pattern in frequency space from the underlying spectrum.

by Fourier transforming $I_s(\omega)$ ($\mathcal{F}\{\cdot\}$ denotes the Fourier transform):

$$\begin{aligned} \mathcal{F}\{I_s(\omega)\} = \tilde{I}_s(t) &= \tilde{S}_{ref} + \tilde{S}_{probe} \\ &+ \mathcal{F}\left\{\sqrt{S_{ref}S_{probe}}e^{i\Delta\phi_{ch}(\omega)}\right\}(t - \Delta t) \\ &+ \mathcal{F}\left\{\sqrt{S_{ref}S_{probe}}e^{-i\Delta\phi_{ch}(\omega)}\right\}(t + \Delta t) \end{aligned} \quad (2.11)$$

The fringe pattern manifests as a pair of side-bands, centred at $t = \pm\Delta t$, as shown in Figure 2.4(b). Provided Δt is sufficiently large that they are well separated from the base-band, one can be isolated, translated to $t = 0$ and inverse Fourier transformed to give:

$$\sqrt{S_{ref}S_{probe}} \exp i\Delta\phi_{ch}(\omega) \quad (2.12)$$

and thus the phase difference $\Delta\phi_{ch}(\omega)$ can be recovered.

Now consider an additional temporal phase shift $\Delta\phi_{pl}(t)$ applied to the probe pulse

by the plasma. This will alter both the spectral phase and spectral amplitude, giving:

$$\tilde{E}'_{probe}(\omega) = \sqrt{S'_{probe}(\omega)} \exp i[\phi_{ch}(\omega) + \Delta t \omega + \Delta\phi_{ch}(\omega) + \Delta\phi_{pl}(\omega)] \quad (2.13)$$

Interfering this with the reference pulse and applying the same analysis method as above gives:

$$\sqrt{S_{ref}S'_{probe}} \exp i[\Delta\phi_{ch}(\omega) + \Delta\phi_{pl}(\omega)] \quad (2.14)$$

At this point, the difference in the phases of Equation 2.12 and Equation 2.14 gives the extra spectral phase caused by the plasma, $\Delta\phi_{pl}(\omega)$. However, the density oscillation of the wakefield affects the temporal phase. Due to the linear chirp of the probe pulse, one might expect that there is a direct correspondence between the spectral and temporal phase, requiring only an appropriate re-scaling of the frequency axis to time. In fact, the temporal phase shift affects both the spectral phase and amplitude, and so both are needed to reconstruct an image of the wakefield.

2.2.2 Temporal phase reconstruction

The temporal phase (phase as a function of time) can be calculated if the complete electric field of the probe, as shown in Equation 2.13, is known. An inverse Fourier transform then gives the temporal electric field, from which the phase can be extracted. Starting from Equation 2.14, the remaining data needed are S_{ref} and $\phi_{ch}(\omega)$. These are the references mentioned previously.

S_{ref} can be measured directly, by altering the experimental setup to eliminate the probe pulse. The complete phase of the reference pulse, $\phi_{ch}(\omega)$, is somewhat more difficult. The chirp intentionally introduced to the pulse by a length of glass (or some other means), can be measured by a simple spectral interferometry experiment, where the glass is placed in one arm of an interferometer. If the two arm lengths are slightly different, this is equivalent to an FDH measurement, and the same mathematics are used to analyse the result. More details on these measurements can be found in § 2.3.

However, this does not account for any chirp present in the probe pulse before it enters the glass, which can only be measured by a technique such as FROG [75, 76] or SPIDER [77]. Since I did not have access to an instrument capable of characterising the probe pulse, I assumed that the residual phase was negligible compared with the chirp introduced deliberately. This is somewhat justified by autocorrelation measurements of the probe pulse, before frequency doubling (see Figure 3.8). These show the effects of third-order dispersion and a central FWHM pulse duration of less than 100 fs. Frequency doubling is expected to reduce this to less than 70 fs, and applying $2.2 \times 10^4 \text{ fs}^2$ of GDD (see Figure 3.14) is expected to increase this to about 900 fs. The autocorrelation is less than half as wide as this, suggesting that the residual third-order dispersion should not have a significant effect on the final pulse.

There is now sufficient information to reconstruct the temporal phase. It will be dominated by a quadratic term due to the linear chirp, which must be subtracted to leave only the modulation caused by the plasma. This can be done by reconstructing the electric field of the unmodulated probe pulse (from Equation 2.12), taking its temporal phase, and subtracting it from the data.

2.3 Experimental setup

An FDH experiment, as shown in Figure 2.3, requires several major components. A laser system generates synchronised pump and probe pulses; this is usually achieved by splitting a small amount of energy from the pump pulse to form the probe. The single probe pulse is used to generate the temporally separated probe and reference pulses used for the FDH measurement. We followed the examples of the previous FDH experiments [67, 68], and used a frequency-doubling method for this. The probe and reference pulses are combined with the pump pulse at a dichroic mirror, which allows light of certain wavelengths to be transmitted, while reflecting certain other wavelengths. A variable delay line (not shown) is used to adjust the relative temporal offset.

The pump is then focussed by a parabolic mirror into an interaction region, where it generates a wakefield. In the case of my experiment, this was a helium-filled gas cell, which was ionised by the passage of the pump pulse itself. The probe and reference pulses are then separated from the pump pulse after the interaction region (by another dichroic mirror) and relay-imaged into an imaging spectrometer. As discussed in § 2.6.3, it is important that this imaging system is adjusted correctly — that the image plane coincides with the slit at the entrance of the spectrometer and that the object plane passes through the interaction region.

Most of this apparatus is standard for a pump-probe experiment, regardless of the diagnostic used, so in this section I will focus on the dedicated FDH optics, pictured in Figure 2.5. In their default configuration, these optics generate the probe/reference pulse pair needed for FDH, but they can also be reconfigured in order to measure the references described in § 2.2.2, i.e. the spectrum of the reference pulse and the chirp of either the reference or probe pulse.

Note that previous experiments [68] measured the chirp of the probe pulse. As described in § 2.4.5, similar results can also be obtained by measuring the chirp of the reference pulse. I will describe the experimental setups I used for both measurements in this section, as well as a proposed setup that would make reference chirp measurements easier.

The method of generating probe and reference pulses for use in normal data-taking is shown in Figure 2.5, with the particular layout I used shown in Figure 2.6. A pair of non-linear crystals, phase-matched for second harmonic generation (SHG), are separated by a length of ordinary glass (for instance, fused silica or BK-7) called the walk-off glass. The residual fundamental pulse after the first SHG crystal propagates more quickly through this glass than the second-harmonic pulse (since it exhibits normal dispersion). The fundamental pulse will generate a second frequency-doubled pulse when it reaches the second crystal, which will be followed by the frequency-doubled pulse generated in the first crystal. This setup is extremely easy to align, resulting




Image courtesy of James Cowley.

FIGURE 2.5 – The basic method by which the probe and reference pulses are generated: (a) a single fundamental pulse is injected; (b) it is frequency-doubled, producing a second-harmonic pulse overlapping the fundamental pulse; (c) the pulses are temporally separated in the walk-off glass; (d) the fundamental pulse is frequency-doubled again, generating another second-harmonic pulse. The frequency-doubled pulses are then chirped by the stretching glass, and then combined with the pump pulse at a dichroic mirror (not pictured). Amplitudes of fundamental and frequency-doubled pulses are not to scale.




FIGURE 2.6 – The experimental setup used for FDH data-taking. (P) and (R) are the SHG crystals that generate the probe and reference pulses respectively. (W-O) is the walk-off glass and (Str) is the stretching glass. The double-headed arrow indicates the translation stage that controls the pump-probe relative timing. The dichroic mirror that combines the probe beam with the pump is not pictured.

in two second-harmonic pulses propagating co-linearly. The temporal separation depends only on the thickness and dispersion of the walk-off glass, so that too is extremely stable.

Although the walk-off glass is only intended to separate the fundamental and second-harmonic pulses by introducing different amounts of group delay to each, it will also introduce different amounts of dispersion to the two pulses. This results in the probe and reference pulses being differently chirped, a fact which the FDH analysis must take into account.

After generating a pair of pulses, it is necessary to stretch them, in order for the probe pulse to cover multiple plasma periods. This is done by passing both pulses through a length of glass (the stretching glass) before they are combined with the pump. Note that previous experiments [67, 68] required less chirp so simply used the substrate of the combining dichroic mirror for this purpose.

2.3.1 Probe chirp measurement

This setup can be adapted as shown in Figure 2.7 to measure the chirp of the probe pulse, which allows the same spectrometer to be used in situ. The probe chirp is measured by interfering the chirped probe pulse with another unchirped, frequency-doubled pulse. This other pulse is generated by using a dichroic mirror to separate the residual fundamental pulse after the first SHG crystal; it is then passed through the other SHG crystal (which is removed from its usual place to suppress the reference pulse) to frequency-double it. Since this other pulse doesn't pass through any glass, it is not chirped. The pulses are then combined at a beam-splitter placed after the stretcher glass. The delay line is then adjusted to match the path lengths and ensure the two pulses interfere and generate spectral fringes. The variable spacing of these fringes can be analysed to give the chirp.

Measurements of the probe chirp via this method varied somewhat from day to day (hence the importance of taking these reference measurements every day), but




FIGURE 2.7 – The experimental setup used for measuring the chirp of the probe pulse. (D) is a short-pass dichroic mirror and (BS) is a thin-film beamsplitter. All other labels are as described in Figure 2.6.

typical values were in the range of $2.25\text{--}2.30 \times 10^4 \text{ fs}^2$. More details will be given in § 3.2.4.

2.3.2 Reference chirp measurement

The same apparatus can be used to measure the chirp of the reference pulse, as shown in Figure 2.8. In this case, the probe pulse must be suppressed; this cannot be done simply by removing the first SHG crystal, since this will affect the remaining fundamental pulse and thus change the reference pulse.

Instead, the two blocks of glass are moved to the alternate beam-line. This ensures that the reference pulse passes through all the same optics in the same order as it does during FDH data-taking, before interfering with an unchirped probe pulse at the beamsplitter.

Again, measurements of the reference chirp varied from day to day, but typical values fell in the range of $1.82\text{--}1.88 \times 10^4 \text{ fs}^2$. More details will be given in § 3.2.4.




FIGURE 2.8 – The experimental setup used for measuring the chirp of the reference pulse. All labels are as described in Figure 2.7.

2.3.3 Proposed reference chirp measurement

This measurement could be made more easily and more accurately if it were not necessary to move so many optics. This might be achieved by replacing the short-pass dichroic mirror with a long-pass one, as shown in Figure 2.9. No optics need to be adjusted from the data-taking setup shown in Figure 2.6, so after performing the chirp measurement the dichroic mirror and beamsplitter can simply be removed, without the need for further re-alignment.

2.3.4 Reference spectrum measurement

The setup used to measure the reference spectrum depends on the setup used to measure the chirp. Either setup to measure the reference chirp (Figure 2.8 or Figure 2.9) can be adapted to measure the reference spectrum simply by inserting a beam block at the point marked (a); this blocks the unchirped probe pulse, leaving only the reference pulse to reach the spectrometer. This significantly reduces the time needed to take the reference measurements, and is an advantage of the reference chirp measurement over




FIGURE 2.10 – The experimental setup used for measuring the spectrum of the reference pulse. All labels are as described in Figure 2.7. Equipment in grey is not used for this measurement.

used for data-taking.

Both these problems could be alleviated by using the reference chirp measurement method shown in Figure 2.9, since no optics need to be moved from their data-taking positions. Although alignment would still be necessary when moving from data-taking to reference measurement, little or none would be needed when subsequently returning to data-taking mode.

Another problem was the slight amount of frequency-doubled light reflected from the short-pass dichroic mirror. This didn't have a significant effect on the chirp measurements, but in the reference spectrum measurements it resulted in a second low energy pulse temporally offset from the main one, similar to the pulse-pair used for data-taking. This gave rise to low-visibility fringes in the recorded data that were not representative of the real spectrum. My solution was to apply a low-pass filter to the data before using it for the analysis. Although this might, in general, affect the spectrum proper, in this case the side-bands are well-separated from the base-band by design, in order for the FDH analysis to work.




FIGURE 2.11 – An overview of the FDH analysis. Some pre-processing and post-processing steps are not shown. The cross-term mentioned is Equation 2.14. FFT and IFFT refer to the fast Fourier transform and its inverse. See text (§ 2.4) for more details.

2.4 Analysis steps

In this section, I will detail the method I used to analyse frequency-domain holograms. Since this is still a very new technique, there is no standard software to do this analysis; I therefore wrote my own in C, using mathematical subroutines from the GNU Scientific Library [78]. The software, when run on a modern desktop or laptop computer, is capable of analysing a single hologram in about a second. This allowed me to build it into a real-time data analysis system, which displays a plot of spectral or temporal phase with only a few seconds' delay. Such a system is extremely useful, showing the effects of parameter changes immediately.

The flowchart in Figure 2.11 illustrates the main steps in an FDH analysis, including which data is required for which step. The first piece of data required is a null hologram – an FDH hologram taken without using a pump pulse. This encodes zero plasma-induced phase shift, and it is needed as a reference against which the data can be compared. This could be called a reference hologram, but I avoid that term to prevent possible confusion with the FDH references, which are the reference pulse spectrum and the probe pulse chirp. This null hologram undergoes a complete FDH analysis to produce spectral and temporal phase data which will be used later in the analysis.

The data hologram is then Fourier transformed and one of the side-bands representing the fringe pattern is isolated and translated to the centre. The Fourier transform is inverted, leaving Equation 2.14. Subtracting $\Delta\phi_{ch}(\omega)$ (from the null hologram analysis) leaves the spectral phase $\Delta\phi_{pl}(\omega)$.

Starting from Equation 2.14 again, the data is divided by the square root of the reference spectrum i.e. $\sqrt{S_{ref}}$; then the phase of the reference pulse $\phi_{ch}(\omega)$ is added on, leaving the spectral electric field $\tilde{E}'_{probe}$, as seen in Equation 2.13.

Inverse Fourier transforming this gives the temporal amplitude and phase, but the phase will be dominated by the chirp that was applied to the pulse. This is subtracted by using the temporal phase value produced by the null hologram, leaving only the difference in phase between the two data sets; that is, the phase due to the plasma.

In the following, I will cover these steps in more detail.

2.4.1 Data input

The main input data required is the hologram, described by:

$$I'_s(\omega) = \left| \tilde{E}_{ref}(\omega) + \tilde{E}'_{probe}(\omega) \right|^2 \quad (2.15)$$

Also needed are the null hologram, described by Equation 2.10, the reference spectrum S_{ref} , and the reference chirp $\phi_{ch}(\omega)$.

In fact, the standard analysis, as described in [68], uses the chirp of the probe pulse ($\phi_{ch} + \Delta\phi_{ch}$). I investigated the possibility of measuring the chirp of the reference pulse and using that in the analysis, because I thought it would improve the results, as described later in this section. Although later analysis showed that both methods worked equally well, the reference chirp is easier to measure, as discussed in § 2.3. Only a small change to the analysis procedure is needed when using the reference chirp, and both the standard and modified versions will be described.

The various spectra are collected by the CCD of an imaging spectrometer, so each spectrum is represented as a matrix of ADC counts (proportional to intensity), each column of the matrix corresponding to a different wavelength λ and each row corresponding to a different vertical position y . Note that the values of λ that are sampled are not equally spaced since the dispersion of a diffraction grating (as used in the spectrometer) varies with wavelength; the corresponding values of ω are not equally spaced either. An example of such a spectrum is shown in Figure 2.12.

Although the data is 2D, the underlying mathematical model in § 2.2 is formulated in only one dimension, that being frequency or time. Therefore most of the analysis is performed on a line-by-line basis, each row being considered separately.

The first step shown in Figure 2.11 is to Fourier transform the recorded hologram. The discrete Fourier transform (DFT) is a good approximation to the (continuous) Fourier transform provided the samples are equally spaced. So first the recorded spec-




FIGURE 2.12 – An example of the spectrum produced by interfering probe and reference pulses – a frequency-domain hologram. Fringes are clearly visible.

trum must be transformed onto a grid where points are equally spaced in frequency. The new grid has the same number of points as the original grid and the same endpoints. The software I wrote performs this transformation by using cubic spline interpolation.

2.4.2 Isolating the side-band

Now that the data is represented by samples equally spaced in ω , the fast Fourier transform (FFT) algorithm can be used to obtain a discrete version of Equation 2.11. The result of such an FFT is shown in Figure 2.13. The data is then translated so that one side-band (generally the positive frequency one) is centred at zero frequency.

The amount of translation, in pixels, is a required input to the analysis; in my implementation, it is found by manual inspection of the FFT of the spectrum. This value depends only on the spectrometer geometry and probe-reference spacing, and as such it can remain stable (within a few pixels) for several months, provided no major experimental changes are made. In any case, an error here leads to a linear




FIGURE 2.13 – The discrete Fourier transform of the spectrum shown in Figure 2.12. **Top:** the whole of the transformed spectrum, shown as an image with a logarithmic colour scale. **Bottom:** a horizontal line-out taken through $y = 0$.

The fringes have manifested as side-bands either side of the base-band. Note: (i) that the side-bands are well separated from the base-band, allowing one of them to be isolated; and (ii) that the time axis indicates the separation of the two pulses, which is just under 5 ps.

phase term in the result, which is subtracted out using the null hologram; therefore, precise measurements are unnecessary.

In order to isolate the side-band, I then multiply the data by a band-pass filter W_j , parametrised by a width w measured in pixels. This width is again an input to the analysis, and is selected by manual inspection of the data shown in Figure 2.13. The value selected must be sufficiently large that the edges of the side-band are not clipped, while being sufficiently small that none of the base-band remains.

Various different filters could be used. I originally used a top-hat function:

$$W_j = \begin{cases} 1 & \text{if } |j| < w/2 \\ 0 & \text{otherwise} \end{cases} \quad (2.16)$$

but after concern about the effect of the sharp edges on the data, I switched to a super-

Gaussian:

$$W_j = \exp \left[- \left(\frac{2j}{w} \right)^8 \right] \quad (2.17)$$

Both produced good results.

After isolating the side-band, the Fourier transform is inverted (an IFFT) resulting in the complex data described by Equation 2.14.

2.4.3 Spectral phase map

Repeating this analysis with the null hologram gives the data described by Equation 2.12. Taking the phase from these two sets of data and subtracting them gives $\Delta\phi_{pl}(\omega, y)$. This is termed a spectral phase map, and an example is shown in Figure 2.14. Although this is not useful by itself for imaging the wakefield, it can indicate whether the analysis is working as expected. Large temporal features, such as an ionisation front, are also visible, giving it utility as an experimental diagnostic. This is also the method used to recover the chirp $\phi_{ch}(\omega)$ from a spectral interferometry measurement, as described in § 2.2.2. The software I wrote is versatile enough to do both.

However, the directly measured phase is wrapped to e.g. $[-\pi, \pi]$. When the true spectral phase extends beyond this range, this results in discontinuities appearing, as shown in Figure 2.15. The phase can be unwrapped by adding an integer multiple of 2π to each pixel. The problem is to choose the values that give the smoothest result.

In the case of 1D data, this is a simple problem, as there is only one solution. However attempting to unwrap each row of the data separately is likely to lead to vertically adjacent pixels being assigned very different phases. I needed a 2D phase unwrapping algorithm that eliminates vertical discontinuities as well as horizontal ones.

2D phase unwrapping is a much more difficult problem, firstly because multiple possible solutions exist, and secondly because there will be some regions where local discontinuities are impossible to avoid. Further details can be found in [79, 80].

There are many well-known 2D phase unwrapping algorithms, but they are generally designed to work on data sets with complex features, spanning many multiples of




FIGURE 2.14 – The spectral phase map reconstructed from the spectrum shown in Figure 2.12 and a corresponding null hologram. The phase shift caused by the plasma can be clearly seen (green-red), but only in those frequencies that propagated behind the pump.

2π . The phase maps produced by FDH analysis generally have no such complex features, and rarely exceed 2π . Given this, I chose to write my own algorithm prioritising simplicity and speed over versatility.

This algorithm unwraps one pixel at a time, selecting the highest-quality pixel out of all wrapped pixels adjacent to an already unwrapped pixel. I first used the amplitude as a measure of quality, but switched to the second derivative of the phase, since it gave better results. This pixel is then unwrapped based on the unwrapped values of adjacent unwrapped pixels. In this way a contiguous “blob” of unwrapped pixels grows from a starting point, growing in high-quality areas first, thus avoiding discontinuities. Provided the set of adjacent pixels is stored in a suitable data structure², this is a fast algorithm, able to unwrap a 146×1024 pixel image in less than 100 ms, using a single core of a common (2010-era) laptop computer. Example results are shown in

²A binary heap is ideal, since inserting a new element or removing the maximum element takes time proportional to the logarithm of the number of elements. The simpler option of an unsorted list allows constant-time insertions, but finding and removing the best element is a linear-time operation, resulting in the algorithm becoming unusably slow.




FIGURE 2.15 – An example phase function (in this case, phase as a function of transverse position for a beam with curved wavefronts), shown unwrapped (bottom) and wrapped to $[0, 2\pi]$ (top). The wrapped form, which would be obtained directly from an interferometric measurement, is difficult to work with. The data must be unwrapped before e.g. fitting a polynomial to it. In 1 dimension there is only one phase unwrapping algorithm, which works by constraining the phase difference between two adjacent pixels to the range $[-\pi, \pi]$.




FIGURE 2.16 – An example phase function is shown: (a) in its original form, in the range $[-\pi, 8\pi]$; (b) wrapped to $[-\pi, \pi]$; (c) unwrapped on a row-by-row basis; (d) unwrapped by my custom algorithm. The area of noise in the middle has a serious effect on the row-by-row unwrapping, without affecting my 2D algorithm. The gradient bars (left and bottom) have not been unwrapped properly, because the correct version features a long discontinuity, which the algorithm avoids.

Figure 2.16.

2.4.4 Reference spectrum division

At this point in the procedure, the value of Equation 2.14 has been calculated, that is:

$$\sqrt{S_{ref}S'_{probe}} \exp i[\Delta\phi_{ch}(\omega) + \Delta\phi_{pl}(\omega)] \quad (2.18)$$

The aim is to calculate the electric field:

$$\tilde{E}'_{probe}(\omega) = \sqrt{S'_{probe}(\omega)} \exp i[\phi_{ch}(\omega) + \Delta t \omega + \Delta\phi_{ch}(\omega) + \Delta\phi_{pl}(\omega)] \quad (2.19)$$

in order to inverse Fourier transform it; this gives the electric field as a function of time, from which the temporal phase can be extracted.

The first step is to divide by the square root of S_{ref} , the reference spectrum. However, before doing this we need to apply the same interpolation that we applied to the original hologram data, for consistency. Failing to do this results in quite noticeable problems, as one of the effects of the interpolation is to shift the spectrum horizontally; unless we interpolate the reference spectrum it won't overlap with the data.

Another problem arises in regions where S_{ref} is smaller than or comparable to the level of random noise. Although the amplitude of the numerator will also be small in these regions, there will be some pixels where a spuriously large value of $\sqrt{S_{ref}S'_{probe}}$ is divided by a correctly small (but non-zero) value of $\sqrt{S_{ref}}$. This gives rise to random spikes in the calculated value of $\sqrt{S'_{probe}}$, which cause serious problems for the inverse Fourier transform.

The solution I chose was simply to zero all sufficiently-small values of $\sqrt{S_{ref}S'_{probe}}$, the cut-off point being chosen to be just above the maximum noise level of the spectrometer. Any data removed by this method is already dominated by noise, so would contribute little to the reconstruction.

2.4.5 Reference phase

The next step is to add the chirp of the reference pulse $\phi_{ch}(\omega)$ to the phase, reconstructing the electric field in the frequency domain, as shown in Equation 2.19. The exact method depends on whether the chirp of the reference pulse or probe pulse was measured. The chirp of the reference pulse, $\phi_{ch}(\omega)$, is added to Equation 2.18 directly. The chirp of the probe pulse is $\phi_{ch} + \Delta\phi_{ch}$, so the spectral phase calculated from the null hologram, i.e. $\Delta\phi_{ch}$ must be subtracted as well as adding the probe chirp.

These two methods give the same result, but using the probe chirp requires an additional step in the analysis. In particular, it involves the use of an additional piece of input data, which will include a certain amount of noise. This led me to investigate whether using the reference chirp improved the quality of the reconstruction. Although I eventually concluded that it did not, using the reference chirp makes the process of measuring the references simpler as described in § 2.3, leading to an improvement over the standard method.

Note that it is important that the sign of the chirp is correct — using a typical interferometric technique it is not uncommon to measure a negative chirp instead of a positive one, which leads to incorrect reconstructions; this will be explained in more detail in § 2.6.1.

2.4.6 Temporal phase

The analysis so far has calculated the electric field of the probe pulse in the frequency domain. Applying an inverse Fourier transform to this data gives the temporal electric field, from which the phase can be extracted. This phase is not usable directly as it is dominated by a large quadratic term, corresponding to the linear chirp of the probe pulse.

$$E'_{probe}(t) = |E'_{probe}| \exp i[\phi_{ch}(t) + \Delta\phi_{pl}(t)] \quad (2.20)$$

The complete FDH analysis on the null hologram gives the electric field of a probe




FIGURE 2.17 – The temporal phase map reconstructed from the spectrum shown in Figure 2.12.

pulse that has not been modulated by the plasma:

$$E_{probe}(t) = |E_{probe}| \exp i \phi_{ch}(t) \quad (2.21)$$

The difference of the two measured phases is then $\Delta\phi_{pl}(t, y)$, the phase imprinted on the probe pulse by the wakefield. This is referred to as a temporal phase map and an example is shown in Figure 2.17. Since temporal phase is proportional to electron density, this gives an indication of electron density as a function of ζ and y – a snapshot of the wakefield.

2.4.7 Reference alignment

The reference spectrum division, spectral chirp addition, and temporal chirp subtraction all require matching up each row of the input data with a corresponding row in the appropriate reference. This is only straightforward when all data sets have the same number of rows and the spectrum is centred on the same row. In the case of the temporal chirp this is expected, since it is calculated from the null hologram, which is

recorded soon after or before the actual data.

This may not be the case for the reference/probe chirp and reference spectrum. The realignment required between measuring these references and measuring data means that they may not be centred on the same row as the data.

One option would be to create an algorithm to automatically align the references vertically with the data. Instead I chose to simply select the best (i.e. lowest noise) rows, average them, and use that average for all rows of the input data, which appears to have been the method adopted for previous experiments [68]. Although this method discards any vertical dependence in the references, it has the advantage that every row of the input data is matched up with high-quality reference data. Using the raw, but correctly aligned references, would mean that rows near the top and bottom of the recorded spectrum are matched with noisy references, adversely affecting the reconstruction. Note that there is no need to correctly normalise the reference spectrum for each row, since a constant factor in Equation 2.19 propagates directly through the inverse Fourier transform to the temporal amplitude and does not affect the reconstructed phase.

In addition, the reference chirp becomes noisy at the edges of the spectrum, which can also adversely affect the reconstruction. So I decided to fit a second-order polynomial to the recorded chirp and extrapolate that, minimising the noise introduced by the chirp addition. Although any higher order chirp, not included in the fit, will affect the output data, it is a systematic effect which will largely be cancelled out by subtracting the temporal phase computed from the null hologram.

This averaging gives good results, although a possibility for future work would be to devise a method that preserves some vertical dependence of the references without sacrificing the signal-to-noise ratio.

2.5 Sensitivity

Designing an experiment requires knowing the smallest density oscillation that can be detected by FDH. This depends on the level of random noise in the phase map and the relation between the electron density and the phase shift, described by Equation 2.5. While planning the experiment, it was assumed that this noise level would be comparable to that which can be achieved with frequency-domain interferometry (FDI) — about 5 milliradians [26, 71].

However, FDH is a significantly more complex technique, requiring a separate analysis to predict the level of noise expected in its results. This has not, to the best of my knowledge, been done previously, with the only published prediction of noise in an FDH measurement using an equation applicable to FDI [66]. In this case, the experimentally realised noise level was ~ 10 mrad, significantly greater than the prediction of 0.1 mrad. This is unsurprising, since an FDI measurement uses all the visible fringes to make a single phase measurement, whereas in FDH each fringe is used to make an independent phase measurement. In this section, I address the problem of how to predict the level of phase noise in an FDH measurement more accurately.

Firstly, I take an analytic approach, propagating errors from the hologram to the spectral phase. Although I do not attempt to extend this method to the considerably more complex analysis needed to estimate the error in the temporal phase, it does provide a prediction which can be compared with experimental data.

Secondly, I examine experimental data. I measure the signal-to-noise ratio and use the analytical model to predict the level of noise in the spectral phase. I then look at the noise level found by taking the standard deviation of experimentally obtained phase maps, and compare them to the prediction. This also provides a measurement of the error in the temporal phase measurements, which can be compared to the error in the spectral phase.

Finally, I use a complete numerical model of the FDH process and analysis, which I wrote in order to make quantitative predictions about the performance of FDH. I

repeat the process used for the experimental data to test my predictions about the noise level in the absence of confounding factors from the experimental setup.

2.5.1 Noise model

The main source of error will be noise in the data recorded by the CCD of the spectrometer. This has two main sources. Firstly, there is a certain level of noise associated with the readout electronics, which is expected to be independent of the signal level. This can be measured by analysing a dark portion of the CCD; it was found to be normally distributed (as expected from the central limit theorem) with a standard deviation in the range 3–5 ADC counts, depending on the spectrometer settings. This will be referred to as the Gaussian noise.

The second source of noise is the fact that the photons arriving at the detector and the electrons being deposited in the pixels are discrete entities. This means that the actual number of electrons measured by the readout electronics is a Poisson-distributed random variable. The fluctuations of this value about its mean constitute the Poisson noise.

The standard deviation of the Poisson noise depends on the expected number of electrons being measured. The readout electronics record the number of electrons N in a pixel and convert that to the ADC count A . Ideally, these are linearly related by a proportionality constant n , such that $N = nA$.

The number of electrons N is a Poisson-distributed random variable with mean μ ; its standard deviation is $\sigma_N = \sqrt{\mu} \approx \sqrt{N}$. The standard deviation of the ADC count is thus:

$$\sigma_A = \frac{\sigma_N}{n} = \frac{\sqrt{N}}{n} = \sqrt{\frac{A}{n}} \quad (2.22)$$

This is the Poisson noise in the spectrometer signal, denoted σ_p . Its standard deviation varies from one pixel to another, depending on the amount of light falling on that pixel. Adding this to the Gaussian noise σ_g , which is caused by the readout electronics

and is the same in every pixel, gives the total standard deviation:

$$\sigma = \sqrt{\sigma_p^2 + \sigma_g^2} \quad (2.23)$$

The value measured by the spectrometer consists of various parts. First there is a fixed pedestal, caused by a zero-offset of the readout electronics. This can easily be measured and subtracted, so it will not be mentioned further. Then there is the signal caused by incident light; this consists of a slowly-varying (DC) part and the quickly-varying fringe pattern, which itself has a slowly-varying amplitude and phase. Since the signal is a Poisson-distributed random variable, it can further be split into its expectation value and an error term, representing the Poisson noise. Finally, there is a normally-distributed error term representing the Gaussian noise associated with the readout electronics.

This can be represented thus:

$$I = \hat{I} + \epsilon_p + \epsilon_g \quad (2.24)$$

where I is the measured signal, $\hat{I}$ is the expectation value of the signal, ϵ_p is a random variable representing the Poisson noise and ϵ_g is a random variable representing the Gaussian noise. Splitting $\hat{I}$ into its DC and fringe-pattern terms and combining ϵ_p and ϵ_g (this assumes that ϵ_p can be modelled as normally-distributed) gives:

$$I = I_D + \frac{I_F}{2} \cos(\omega \Delta t) + \epsilon \quad (2.25)$$

$$\epsilon \sim \mathcal{N}(0, \sigma^2) \quad (2.26)$$

where I_D is the expectation value of the local amplitude of the slowly varying (DC) part and I_F is the expectation value of the local peak-to-peak amplitude of the fringe pattern; ϵ is a normally-distributed random variable with standard deviation σ , accounting for both Poisson and Gaussian sources of noise.

It is useful to define a signal-to-noise ratio; since I_D is not used in the FDH analysis, it is a background, contributing to the noise but not to the signal. Therefore the signal-to-noise ratio is given by the ratio of I_F to σ and its reciprocal is defined as the relative noise level:

$$\sigma_r = \frac{\sigma}{I_F} \quad (2.27)$$

2.5.2 Analytical prediction

Ignoring the DC part of the intensity and introducing a slowly varying phase to the fringe pattern, Equation 2.25 becomes:

$$I = \frac{I_F}{2} \cos(\omega\Delta t + \phi) + \epsilon \quad (2.28)$$

where the j subscripts have been dropped. The value of ϕ is estimated by setting $\epsilon = 0$:

$$\frac{I_F}{2} \cos(\omega\Delta t + \hat{\phi}) = I = \frac{I_F}{2} \cos(\omega\Delta t + \phi) + \epsilon \quad (2.29)$$

where $\hat{\phi}$ is an estimator for the value of ϕ .

The standard error propagation formula allows the standard deviation of $\hat{\phi}$ to be calculated:

$$\sigma_{\hat{\phi}} = \sigma_I \left| \frac{\partial \hat{\phi}}{\partial I} \right| = \sigma \left| \frac{\partial I}{\partial \hat{\phi}} \right|^{-1} \quad (2.30)$$

$$\frac{\partial I}{\partial \hat{\phi}} = -\frac{I_F}{2} \sin(\omega\Delta t + \phi) \quad (2.31)$$

$$\sigma_{\hat{\phi}} = 2\sigma_r \csc(\omega\Delta t + \phi) \quad (2.32)$$

where $\csc$ denotes the cosecant.

The model so far does not accurately represent the FDH analysis, since the real analysis is performed on only one side-band. As such, it incorrectly predicts that the error varies depending on the phase within the fringe pattern, diverging to infinity at

the peaks and troughs. A more complete model can be built by splitting I into positive and negative frequency parts:

$$I_+ = \frac{I_F}{4} e^{i(\omega\Delta t + \phi)} + \epsilon_+ \quad (2.33)$$

$$I_- = \frac{I_F}{4} e^{-i(\omega\Delta t + \phi)} + \epsilon_- \quad (2.34)$$

$$I = I_+ + I_- \quad (2.35)$$

$$\epsilon = \epsilon_+ + \epsilon_- \quad (2.36)$$

Since ϵ represents white noise, it is assumed to be uniformly distributed between independent positive and negative frequency parts ϵ_+ and ϵ_- , with equal standard deviations σ_+ and σ_- . The independence condition gives:

$$\sigma^2 = \sigma_+^2 + \sigma_-^2 = 2\sigma_\pm^2 \quad (2.37)$$

$$\sigma_\pm = \frac{\sigma}{\sqrt{2}} = \frac{\sigma_r I_F}{\sqrt{2}} \quad (2.38)$$

Repeating the earlier analysis:

$$\frac{\partial I_+}{\partial \hat{\phi}} = i \frac{I_F}{4} e^{i(\omega\Delta t + \phi)} \quad (2.39)$$

$$\sigma_{\hat{\phi}} = \sigma_+ \left| \frac{\partial I_+}{\partial \hat{\phi}} \right|^{-1} = 2\sqrt{2} \sigma_r \quad (2.40)$$

This roughly agrees with Equation 2.32, but has no infinities or ϕ -dependence.

Although this provides an estimate of the error in the reconstructed spectral phase, it is important to note that this may not be the same as the error in the temporal phase. Although a more complete analysis might give an estimate of the error in the temporal phase, I have only included this analysis as an example of the way FDH errors might be estimated during experimental planning. I will show that such an analysis underestimates the actual error obtained in an FDH experiment, which will need to be considered when planning future experiments.




FIGURE 2.18 – The process of calculating the actual and predicted values of the spectral phase noise, $\sigma_{\hat{\phi}}$. These values can be compared to check the analytical prediction presented in § 2.5.2.

2.5.3 Experimental value

Checking the analytical predictions of the previous section requires comparing values of $\sigma_{\hat{\phi}}$ computed in two different ways. Firstly, a predicted value can be obtained from the signal-to-noise ratio. This requires knowledge of the noise σ and the fringe amplitude I_F at each pixel of the holograms being used. Secondly, an actual value can be obtained from phase maps, computed from the holograms by the FDH analysis process. This requires computing the standard deviation of the phase at each pixel across many phase maps. This is depicted in Figure 2.18.

To accomplish this, I recorded a large number (60) of null holograms under similar conditions. I analysed them to give spectral and temporal phase maps, then discarded 5 clearly anomalous results³. This provided, for each pixel, 55 values of the random variable $\hat{\phi}$. I used these to calculate the standard deviation of $\hat{\phi}$, giving the actual phase error $\sigma_{\hat{\phi}}$ as a function of (y, ω) or (y, t) .

To obtain predicted values of $\sigma_{\hat{\phi}}$ for comparison, the signal-to-noise ratio σ_r^{-1} of the raw data is needed. The Gaussian readout noise was measured to have a stan-

³These anomalous results had phase shifts of several radians, rather than the < 100 mrad phase shifts of the good data. It is believed that they were caused by equipment errors while recording the data.

standard deviation of 4.6 ADC counts. The Poisson noise varied per pixel according to Equation 2.22. The number n of electrons per ADC count was calculated by analysing the high-frequency portion of the Fourier transform; since there is neither signal nor background in this region this gives, for each row, the sum of the squared errors. This was found to match the model if $n = 4$. The total standard deviation, which varies from pixel to pixel, was calculated using Equation 2.23.

However, this overestimates the noise level. Since both Gaussian and Poisson noise are examples of white noise (different pixels are uncorrelated) the noise is spread across the entire spectrum when the Fourier transform is applied. Isolating the side-band as part of the FDH analysis also filters out a large amount of the noise. About 80% of the spectrum is removed, and each fifth of the spectrum is assumed to contribute equally and independently to the total noise level. Therefore the total noise level σ can be related to the contribution from one fifth of the spectrum σ_1 thus:

$$\sigma^2 = 5\sigma_1^2 \quad (2.41)$$

since summing independent normal distributions results in the variances being summed. Thus removing 80% of the spectrum and going from σ to σ_1 results in reducing the standard deviation of the noise by a factor of $\sqrt{5}$.

Having obtained an estimate of σ , the value of σ_r then depends on the peak-to-peak amplitude of the fringe pattern, which varies across the spectrum. This can be measured by band-pass filtering one of the side-bands, giving a measurement of $\sqrt{S_{ref}S_{probe}}$. Comparing the factors of the cosine terms in Equation 2.10 and Equation 2.25 shows that this must be multiplied by 4 to give I_F . Note that a naïve calculation of the signal-to-noise ratio might use the peak intensity of the spectrum; this will overestimate the signal-to-noise level, due to the limited visibility of the fringes.

A map of σ_r^{-1} as a function of (y, f) is shown in Figure 2.19, with typical values being about 100. Since the amount of Poisson noise is less in the troughs of the original hologram than in the peaks, the fringe pattern manifests in the plot of signal-to-noise




FIGURE 2.19 – The measured signal-to-noise ratio, I_F/σ , where I_F is the peak-to-peak amplitude of the fringe pattern, as defined in Equation 2.25. The dashed line in the 2D data (**top**) shows where the line-out (**bottom**) was taken.

ratio. This is not expected to affect the experimental noise values, since each fringe results in only a single independent phase measurement. Using this, I calculated predicted values of $\sigma_{\hat{\phi}}$ from Equation 2.40, which are plotted as a function of (y, f) in Figure 2.20. Typical values range from 20 to 30 mrad.

The actual phase error is shown in Figure 2.21 (frequency domain) and Figure 2.23 (time domain). Typical values of the spectral phase error range from 30 to 40 mrad, while the temporal phase error is generally about 80 mrad. The actual spectral phase error is slightly larger than the prediction from the analytical model — see the comparison in Figure 2.22. These values are significantly greater than typical errors in FDI measurements; in fact given a signal-to-noise ratio S_n and a detector with N pixels, an FDI measurement should result in a phase error of $\sigma_{\phi} \sim (S_n \sqrt{N})^{-1}$ [66], leading to a value < 1 mrad here. As discussed previously, this difference is due to the fact that in FDH each fringe constitutes a separate measurement, to which only a few detector pixels are able to contribute.




FIGURE 2.20 – The predicted values of $\sigma_{\hat{\phi}}$, the error in the reconstructed phase, based on the measurements shown in Figure 2.19 and Equation 2.40. The dashed line indicates where the line-out (see Figure 2.22) was taken.




FIGURE 2.21 – Actual values of $\sigma_{\hat{\phi}}$ in the frequency domain. The dashed line indicates where the line-out (see Figure 2.22) was taken.




FIGURE 2.22 – Line-outs of the spectral phase error from Figure 2.20 and Figure 2.21. The significant difference at the edges is in the noise-dominated region, where phase unwrapping becomes unreliable, significantly increasing the measured phase error.




FIGURE 2.23 – Actual values of $\sigma_{\hat{\phi}}$ in the time domain. The dashed line in the 2D data (**top**) shows where the line-out (**bottom**) was taken.

However, the temporal phase error is larger than the spectral phase error by a factor of 2. This is due to the more complex parts of the FDH analysis introducing other sources of error, such as noise in the measurement of the reference spectrum. This shows that an additional margin of error will be needed when using the analytical model of § 2.5.2 to predict the noise level of temporal phase data.

2.5.4 Numerical prediction

In order to cross-check the experimental results, I used the FDH simulation code I wrote (described in more detail in § 4.8) to investigate the effect of noise on the FDH analysis process. For this purpose, the code was augmented with a basic simulation of Poisson noise — transforming each pixel into a Poisson-distributed random variable — and readout noise — adding a normally-distributed value to each pixel — applied to all output spectra. No phase shift due to plasma was added, and the spectrum was assumed to be a perfect Gaussian.

After generating 64 rows of simulated data and analysing them, I took the standard deviations of the spectral and temporal phase data to give actual values of $\sigma_{\hat{\phi}}$ as functions of frequency and time respectively. I then followed the same procedure as in § 2.5.3 to estimate the signal-to-noise ratio of the input spectrum, and predict the noise in the spectral phase from that. This includes the division by $\sqrt{5}$ to account for the noise eliminated by Fourier-domain filtering during the FDH analysis process. Since the fringe contrast ratio was much greater than seen in the experimental data, the computed signal-to-noise ratio varied significantly from pixel to pixel; I therefore applied a moving-average filter (covering one fringe) to the predicted noise data to allow it to be compared with the actual noise. The actual and predicted values of the frequency-domain noise are compared in Figure 2.24, which shows them to be in close agreement.

The values of the temporal phase noise are shown in Figure 2.25. Note that the minimum value is significantly greater than the minimum spectral phase error (6 mrad




FIGURE 2.24 – Comparison of actual and predicted values of $\sigma_{\dot{\phi}}$ in the frequency domain, based on simulated data. The significant difference at the edges is in the noise-dominated region, where phase unwrapping becomes unreliable, significantly increasing the measured phase error.




FIGURE 2.25 – Actual values of $\sigma_{\dot{\phi}}$ in the time domain, based on simulated data.

compared to 4 mrad), in agreement with the experimental observation of § 2.5.3. This suggests that the increase in noise when transforming from the frequency domain to the time domain is an intrinsic property of the analysis process, rather than a fault of the experimental setup.

2.5.5 Averaging

In principle, it might be possible to detect and measure a wakefield with amplitude less than the noise level. This could be done by averaging over multiple periods of the wakefield, averaging transversely (since the wakefield is expected to cover more than one pixel), or by taking multiple measurements and averaging over those.

Averaging over multiple periods of the wakefield could be achieved by fitting a sine wave to the temporal phase data, or by taking a Fourier transform. In principle, this would bring FDH into line with FDI, since every fringe on the spectrometer would then contribute to the same phase measurement. There are some practical difficulties with this idea, however: the temporal extent of the wakefield may not be clear, and adding data without a wakefield will worsen rather than improve the signal-to-noise ratio.

Similarly, the signal-to-noise ratio could be improved by averaging over several rows of phase data (averaging along the transverse direction), but this also requires the extent of the wakefield to be known to avoid including rows without a wakefield. Although both these approaches could be used to improve the signal-to-noise ratio of the final measurement, they are of little use if the wakefield signal is too faint to see at all.

Taking multiple measurements and averaging over those is another possibility, albeit one which sacrifices the single-shot nature of FDH. Consequently, it is susceptible to shot-to-shot fluctuations, which might change the size, shape, or position of the wakefield from one measurement to another. This particular approach was not used during my experimental analysis, since jitter in the laser beam position caused the po-

sition and measured amplitude of the wakefield to vary. Averaging over all recorded data would have reduced the signal-to-noise ratio by including many measurements without a wakefield, while selecting only those measurements which appear to include a wakefield would be unsound, leading to a kind of selection bias.

In general, intensive data mining to extract a signal from noisy data runs the risk of generating a completely fictitious signal unless the data is well understood. Although possible in principle, the statistical methods needed to perform this kind of analysis safely are quite complex, and a simpler solution is to ensure that the signal is visible above the noise before performing any averaging.

2.6 Practical difficulties

At the beginning of this experiment, we expected that FDH would be relatively straightforward to implement. However, over the course of the project, many unforeseen problems were discovered.

Firstly, since FDH is a very new technique and we are, to my knowledge, only the second group to implement it, there is not a large volume of detailed literature on the subject. I therefore rediscovered several subtleties in the analysis. I have already discussed some of these, such as the need to ensure correct vertical alignment of the references (see § 2.4.7) and the need to zero small values when dividing by the reference spectrum (see § 2.4.4). I will describe some others in this section.

Secondly, the performance of the laser I used was worse than had been expected when planning the experiment. Particular problems were the lower-than-expected energy reducing the phase shift caused by the wakefield, (see § 3.1) the poor shot-to-shot stability affecting the performance of the FDH analysis, and lower than expected bandwidth of the probe pulse, as well as an unusual spectral shape (see § 3.2.3). These all affected the ability of my FDH analysis to detect a wakefield.

Finally, the need to take separate reference measurements, with the associated frequent re-alignment of the optical setup, appears to be a fundamental limitation of

the FDH technique as presented by its creators [67, 68]. I will therefore also present ideas for possible future work, whereby the need for separate reference measurements could be partially or wholly eliminated.

2.6.1 Sign of probe chirp

First note that although I discuss measurements of the probe chirp in this section, all the same arguments apply to the reference chirp, but the maths is more complex in the case of the reference chirp.

The probe chirp was measured by interfering the chirped probe pulse with an unchirped reference pulse and analysing the result in a spectrometer. This is equivalent to the first part of an FDH measurement i.e. the part needed to measure the spectral phase difference between two temporally separated pulses. Therefore, the two pulses must be temporally separated by a certain amount, to ensure the fringe pattern has a suitable period for analysis. This period can be changed by means of an adjustable delay line in one of the two arms of the interferometer.

The best method of setting the spacing is to adjust the delay line until the pulses are overlapped, which can be seen clearly on the spectrometer read-out as the broadest possible fringe pattern. Then the delay line is moved by a known amount to set the spacing appropriately; for example, a move of 0.75 mm results in a spacing of 5 ps (since the distance is divided by c to find the travel time, and then doubled because the light travels both backwards and forwards along the delay line). A simpler method is to adjust the delay line until the fringe pattern looks suitable.

However, it is important whether the reference pulse or probe pulse arrives first. In the first case the phase of the probe pulse is being measured relative to the reference; in the other case the phase of the reference pulse is being measured relative to the probe, giving a result of the opposite sign. Therefore, unless care is taken with the measurement, the measured probe chirp may have the incorrect sign.

The main effect of this will be seen to be a reversal in the direction of the time

axis in the reconstructed temporal data. However, this is not the only effect. It also distorts the reconstructed data, for reasons that are not necessarily obvious when first considering the matter.

Consider the data needed to perform an FDH analysis:

$$\sqrt{S_{ref}S_{probe}} \exp i[\Delta\phi_{ch} + \Delta\phi_{pl}] \quad (2.42a)$$

$$\sqrt{S_{ref}} \quad (2.42b)$$

$$\Delta\phi_{ch} \quad (\text{from the null hologram}) \quad (2.42c)$$

$$\phi_{ch} + \Delta\phi_{ch} \quad (2.42d)$$

The analysis process takes (a), divides it by (b), subtracts the phase (c) and adds the phase (d), to give the electric field:

$$\tilde{E} = \sqrt{S_{probe}} \exp i[\phi_{ch} + \Delta\phi_{ch} + \Delta\phi_{pl}] \quad (2.43)$$

Now consider the electric field that would be reconstructed if the sign of the measured probe chirp were wrong, so that (d) is subtracted rather than being added:

$$\tilde{\epsilon} = \sqrt{S_{probe}} \exp -i[\phi_{ch} + \Delta\phi_{ch} - \Delta\phi_{pl}] \quad (2.44)$$

where ϵ denotes an incorrectly calculated electric field.

At this point, it would be convenient to assume that changing the sign of $\Delta\phi_{pl}(t)$ will change the electric field to:

$$\tilde{\bar{E}} = \sqrt{S_{probe}} \exp i[\phi_{ch} + \Delta\phi_{ch} - \Delta\phi_{pl}] \quad (2.45)$$

where $\bar{E}$ is the same as E , but with the opposite sign of $\Delta\phi_{pl}$ applied in the time domain.

However, this is not correct, because flipping the sign of $\Delta\phi_{pl}$ in the *time domain*

does not simply result in a sign flip in the *frequency domain*; the effect on the phase will be somewhat more complex and there will also be an effect on the spectrum. However, continuing under this assumption, ε can be related to $\bar{E}$, by noting that $\tilde{\varepsilon} = \bar{E}^*$, where $*$ denotes the complex conjugate.

Behaviour under complex conjugation is one of the basic properties of the inverse Fourier transform (denoted by $\mathcal{F}^{-1}\{\circ\}$ below):

$$\mathcal{F}^{-1}\{A^*\}(t) = \mathcal{F}^{-1}\{A\}(-t)^* \quad (2.46)$$

$$\begin{aligned} \varepsilon(t) &= \bar{E}(-t)^* \\ &= |E(-t)| \exp i[-\phi_{ch}(-t) + \Delta\phi_{pl}(-t)] \end{aligned} \quad (2.47)$$

The result is time-reversed and has the wrong sign of temporal chirp (which will be removed by the last step of the analysis), but the sign of $\Delta\phi_{pl}$ is correct. This explains why the main effect of changing the sign of the probe chirp is a temporal reversal, but the incorrect assumption above means that there will be other effects; it is not clear from the theory what these effects will be, but experimentally vertical stripes in the data were observed (these sometimes appeared with the correct sign of the chirp, but tended to be stronger when the sign was wrong). Using the correct sign of the probe chirp is therefore critical to achieving the best possible FDH reconstruction. This can be achieved simply by changing the sign of the measured chirp if it is wrong — the correct sign can be found by knowing the sign of the group velocity dispersion of the stretching glass.

2.6.2 Spectral ripples

Another effect that was unforeseen, but easily explained in retrospect, was the effect of an ionisation front on a spectral phase map. The effect on temporal phase is simple: a change in phase (I've observed shifts from a few hundred mrad to tens of radians) over a time shorter than the pump duration. The effect on the spectral phase might be expected to be similar, since every wavelength in the chirped pulse is either in front




FIGURE 2.26 – A simple numerical calculation – whereby a temporal phase shift in the form of a step function at $t = 0$ is applied to a linearly chirped Gaussian, then the result is Fourier transformed to the frequency domain – results in a spectral amplitude and phase shift of this form. The amplitude scale is not shown. Note that: (a) the phase shift is not a simple step function, as might be expected; and (b) significant modulation has appeared in the spectral amplitude. The asymmetry arises because the time-dependent phase shift acts as a frequency shift, moving the centroid of the spectrum either to higher or to lower frequencies depending on the sign of $\Delta\phi(t)$; for example, $\Delta\phi(t) = \alpha t$ would simply translate the whole spectrum (breaking the symmetry) but have no other effect.

of the pump, receiving no phase shift, or behind it, receiving a full phase shift.

Observed spectral phase maps showed something quite different – there was a sudden phase shift due to the ionisation front, but there were also oscillations both before and after it. This is not a problem in the FDH reconstruction, as these oscillations also appear in the output of a simple calculation, as shown in Figure 2.26. This is accompanied by a modulation of the spectral amplitude, which was anticipated in the mathematical description of FDH given in § 2.2.

These ripples in the spectral phase are a real effect, not an artefact of the FDH reconstruction. They do, however, limit the usefulness of the spectral phase map, since they make the position of the ionisation front less clear and may obscure any wakefield. In principle the full FDH reconstruction, using the modulated spectral am-

plitude, should eliminate them — the phase and amplitude modulations should cancel out, leaving a sharp ionisation front.

In practice, simulations (similar to those described in § 4.8) suggest this cancellation might be incomplete, leaving small ripples in the time domain. It is therefore important to avoid imaging the ionisation front in order to minimise this effect.

2.6.3 Imaging the probe correctly

The discussion so far has ignored the propagation of the probe between the interaction region and the spectrometer. A particular effect of this will be the evolution of the transverse profile of the probe pulse, which distorts the recorded phase information. Therefore, an imaging system must be used to form an image plane at the entrance to the spectrometer, with the corresponding object plane being inside the interaction region. Although such an imaging system is straightforward to design, care must be taken when setting it up to ensure that the image formed at the spectrometer is in focus.

This is not a simple task, since there is no obvious indication of when the image is in focus. The method I chose relied on placing a sharp edge in the interaction region, in the same plane as the pump focus. This provided an image on the spectrometer that could be brought into focus more easily.

Failure to set up the imaging system correctly may result in out-of-focus FDH images that cannot resolve a small wakefield, although calculating the precise effects will require a much more thorough analysis.

2.6.4 Spectral variations

The FDH technique as described relies extensively on taking multiple consistent measurements and comparing them, *viz.* the reference spectrum must match the spectrum of the reference pulse as used to produce the hologram, and the null hologram must use the same probe/reference pulses as the data hologram. This is straightforward in

theory, but becomes a problem in practice if this consistency requirement cannot be upheld; in particular if the shape of the reference spectrum changes.

There are two slightly distinct forms of this problem that I encountered, with different effects. Firstly, small random variations of the reference spectrum from shot to shot. Secondly, larger systematic changes of the spectrum on a longer time-scale, typically hours. The severity of these problems depends, naturally, on the inherent stability of the laser being used.

Shot-to-shot variations

Shot-to-shot variations are a particular problem for subtraction of the temporal phase measured by the null hologram. The purpose of that step is to subtract the phase that would be measured without the pump, so if changing the spectrum changes that phase, an incorrect background phase will be subtracted.

A particularly good example of this occurred when there were sharp edges in the spectrum. Narrow vertical lines would appear in the same place in the spectral phase map (even after subtracting $\Delta\phi_{ch}$, measured with the null hologram), and would fluctuate from shot to shot. An example of this can be seen in Figure 2.21 at approximately 735 THz; this corresponds to the low-frequency edge of the narrow peak at 735 THz in the corresponding spectrum (which is visible in Figure 2.19), which causes the phase at this frequency to fluctuate significantly more than at other frequencies.

The specific effect seems to be that the edge is sufficiently steep that it shifts the measured phase of the fringes. As the spectrum fluctuates, the edge moves to slightly higher or lower frequencies, changing the measured phase shift. This has a consequent effect on the temporal phase map, where the effect is spread over the entire data set by the Fourier transform.

Ideally, this wouldn't be a problem, as the spectra of the probe and reference pulses would be optimised to lack such sharp edges and gaps, which prevent phase recovery entirely and cause an even more severe effect on the temporal phase map. However,

as seen in § 3.2.3 my attempts to optimise the spectra in this way were unsuccessful.

Longer-term drift

More substantial changes to the shape of the reference spectrum occurred on time-scales of hours. When these changes occurred between recording the references and recording the holograms, the reference spectrum division would no longer result in the correct modulated probe spectrum $\tilde{E}'_{probe}(\omega)$, so inverse Fourier transforming it produced incorrect results.

This was a particular problem when using the live analysis suite I built. This relied on recording a reference spectrum and probe chirp before starting to take data, allowing a temporal phase map to be displayed in near-to-real time. This feedback could be used to optimise the parameters of the experiment and identify problems before the main analysis. However, if the reference spectrum changed for some reason during data taking, the quality and usefulness of the results dropped substantially, requiring the data-taking session to be suspended while new references were taken. This is not only time-consuming by itself, but the realignment required can adversely affect various parameters that had already been optimised.

2.6.5 Systematic effects in the temporal phase

Many of the temporal phase maps I recorded exhibited a particular systematic error: a series of vertical stripes on one or both sides of the ionisation front. An example of this can be seen in Figure 2.17, where the left-hand side of the plot (before the ionisation front) includes such a set of stripes (with approximate period 200 fs and amplitude 200 mrad).

The presence or absence of such stripes was linked to the presence or absence of an ionisation front in the data. In particular, they did not appear when the laser was on, but the ionisation front could not be seen in the temporal phase map (because the pump pulse did not overlap the probe pulse). This is similar to the behaviour of

the ripples in the frequency domain, described in § 2.6.2 and shown in Figure 2.26. Those are a real effect in the frequency domain, that should be removed by a correct transformation to the time domain. The fact that they still appear to be present to some extent shows that there is something wrong with this transformation.

Notably, a worsening of these stripes is the main effect of a change in the reference spectrum, of the kind described in § 2.6.4, showing that their most likely cause is an incorrect measurement of the reference spectrum. This shows that even a short time separation between recording the references and the data is not sufficient by itself to ensure consistency between the two.

2.6.6 Measuring references

The biggest problems I encountered with the reconstructed data had a common root cause — inconsistency between the recorded reference spectrum and the holograms. Even if this consistency problem could be fixed, measuring the references (the reference spectrum and the probe/reference chirp) is quite time consuming and requires a realignment of the optical setup each time. These problems would be worsened at a national laboratory, where time is at a premium, and may make the use of FDH impractical for many groups.

The standard FDH technique would be significantly improved if these references did not need to be taken so often. The probe/reference chirp, being simply the dispersion introduced by two pieces of glass, should be quite stable (see § 3.2.4). As such, it should be possible to measure the probe/reference chirp once at the beginning of the experiment and simply use that value throughout, although periodic re-measurements could be taken to confirm consistency.

The reference spectrum, however, is a property of the laser, and thus its stability will depend on the laser's stability. In order to avoid taking frequent separate measurements of the reference spectrum, while avoiding the problems described above that occur when the recorded reference spectrum does not match the hologram, it is

necessary to record the reference spectrum at the same time as the data.

One possibility would be to introduce a second beam-line leading to a second spectrometer, so that the reference pulse could be split and have its spectrum measured at the same time as it is used to produce a hologram. This significantly complicates the experimental setup, so in the following section, I discuss another possibility: that the reference spectrum might be recovered from the hologram directly, by using the information in the base-band that is discarded by the standard FDH analysis.

2.7 Direct spectral recovery

The standard FDH technique relies on separate measurements of the reference spectrum to reconstruct the temporal phase. This is because the amplitude of the cross-term in Equation 2.10 depends on the product of the spectra of the reference pulse and the modulated probe pulse, so a separate measurement of the reference spectrum is needed to isolate the spectrum of the modulated probe pulse.

However, the other term of that equation — corresponding to the base-band in the Fourier transform — is the sum of the same two spectra. In principle, this provides enough information to isolate the two spectra without the need for a separate measurement. This would not only reduce the time needed to record data, it might also improve the quality of the reconstructed phase by eliminating the problems caused by inconsistent reference spectrum measurements discussed in § 2.6.4.

In this section, I discuss the feasibility of this idea, implementation of which is a possibility for future work.

2.7.1 Isolating the spectra

Equation 2.11 shows that the hologram can be separated into a DC term and a side-band:

$$I_{DC} = S_{ref} + S_{probe} \quad (2.48)$$

$$I_{SB} = \sqrt{S_{ref}S_{probe}} \quad (2.49)$$

This gives the sum of the two spectra as well as their product, which in principle allows them to be separated. This would allow a direct measurement of the modulated probe spectrum, without any need for the separate reference division step. For the remainder of this section, I will dispense with the notation specific to FDH, and discuss the general problem of calculating two (real) functions A and B , given only their sum $S = A + B$ and product $P = A \times B$.

The solution to a straightforward quadratic equation gives:

$$\left. \begin{matrix} A \\ B \end{matrix} \right\} = \frac{S \pm \sqrt{S^2 - 4P}}{2} \quad (2.50)$$

where one of A and B takes on the value of the $(+)$ case, and the other the value of the $(-)$ case. Provided S and P are measured correctly, the solutions of this equation should be real numbers, by the definitions of S and P (above). The possibility of incorrect measurements of S and P leading to complex solutions to this equation will be considered later.

Although the values of A and B can be calculated at every point in the recorded hologram, there's an ambiguity at each pixel as to which is which. Resolving this ambiguity requires some method of choosing the sign of $A - B$ at each pixel, which I will refer to as a selection function.




FIGURE 2.27 – Even assuming continuous spectra, ambiguities arise when the spectra cross. In addition to the global ambiguity, where A and B can be interchanged, it is impossible to decide between case (a), where A and B cross over, and case (b), where the lines touch each other but don't cross. In this case (b) may not look physically plausible, but real data will not be so clear.

2.7.2 Simplifying the problem

The problem of creating a suitable selection function can be significantly simplified by requiring that A and B are both continuous functions. In the simple case of 1 dimension without noise, this gives rise to an ambiguity only when the values of A and B cross over (see for example Figure 2.27). These crossing points divide the spectrum into a number of regions, where the continuity condition requires that the selection function is constant within each region. In the zones between these regions, $A = B$ and so the value of the selection function is irrelevant.

An extra dimension and noise both complicate the problem. Consider the discriminant (argument of the square root) in Equation 2.50:

$$D = S^2 - 4P \quad (2.51)$$

In contiguous regions where the discriminant is large and positive, the value of $|A - B|$ is significant, so again the selection function must be constant within these regions. In regions where the discriminant is small (compared to the noise) and positive or zero, A and B are very similar, so the value of the selection function is not so important. Finally, noise in the data can cause the measured value of the discriminant to become negative, resulting in unphysical complex values of A and B . In these cases it is best to set the discriminant to zero, since it is known that its true value is small.

Thus, a suitable algorithm can generally partition the spectrum into a number of “islands” with a large value of D , surrounded by a “sea” of pixels where D is small relative to the noise. The problem is then reduced to assigning a value to the selection function within each island while ignoring the sea. A complete solution would then consider the pixels in the sea at a later step.

2.7.3 Using a target spectrum

Further progress seems to require knowing the approximate form of either A or B . In the case of FDH, this would be realised by recording the reference spectrum in the typical manner and using that. Since this reference spectrum would only be used to help recover the real reference spectrum used for analysis, slight changes would be less important. In principle, the result of the decomposition could be used to keep the target spectrum up to date, tracking any drift in the reference spectrum, and eliminating the need for frequent measurements.

The selection function can then be computed by a simple optimisation algorithm, where the goal function depends on the difference between the target and computed reference spectra. The target will not necessarily have the same normalisation as the data, i.e. the spectrum recovered from the hologram and the target spectrum recorded earlier may be scaled differently, due to differences in optics and alignment between the two measurements, as well as shot-to-shot energy variations. Normalising the two spectra to unit area before comparing them should be sufficient to account for this.

One way to make this part of the analysis simpler would be to make one pulse more intense than the other. This would reduce the number of points where the two spectra have the same value, resulting in fewer, larger “islands”. In turn, this would reduce the number of free parameters of the selection function, and might reduce the noise in the final result.

2.7.4 Spectrometer resolution

Although this approach appears sound in the ideal case, it is likely to run into problems associated with the limited resolution of the spectrometer. A spectrometer operates by illuminating a diffraction grating with a collimated input beam, so that each spectral component is diffracted to a collimated beam at a different angle. A lens or curved mirror then images these onto a detector plane. The size of the spot formed on the detector by a single spectral component depends on how wide the beam is, limiting

the resolution according to the size of the illuminated portion of the diffraction grating.

The size of the spectral fringes will generally be of a similar order to the resolution, reducing their visibility compared to what would be measured by an ideal spectrometer. This is equivalent to a low-pass filter in Fourier space reducing the intensity of the side-bands representing the fringe pattern. This has no substantial effect on FDH because the absolute intensity of the side-band is not relevant to the analysis. But the above method of direct spectral recovery depends sensitively on the ratio between the intensities of the side-band and base-band, which will be affected by the limited resolution of the spectrometer. This reduction in visibility must be computed or measured accurately so that the effect can be corrected. Since it depends on the geometry of the spectrometer and the fringe spacing, it should be extremely stable.

The visibility reduction could be calculated from first principles, using the known geometry of the spectrometer to calculate the shape and cut-off of the filter applied to the measured spectra. However, the recorded data can also be used to obtain a bound on the value, based on the fact that the true value of the discriminant D cannot be negative.

If the measured intensity of the side-band is reduced by some factor α , the measured value of the product term $P = I_{SB}^2$ will be reduced by a factor α^2 , that is:

$$\hat{P} = \alpha^2 P \quad (2.52)$$

where $\hat{P}$ is the measured value and P is the true value. The true value of the discriminant is then:

$$D = S^2 - 4P = S^2 - \frac{4}{\alpha^2} \hat{P} \quad (2.53)$$

Since the true value of D cannot be negative:

$$S^2 - \frac{4}{\alpha^2} \hat{P} \geq 0 \quad (2.54)$$

$$\alpha^2 S^2 \geq 4 \hat{P} \quad (2.55)$$

$$\alpha \geq \frac{2}{S} \sqrt{\hat{P}} \quad (2.56)$$

The value of this lower bound relative to the true value of α depends on the true value of the discriminant:

$$\begin{aligned} \alpha_{LB} &= \alpha \sqrt{\frac{4}{S^2} P} \\ &= \alpha \sqrt{1 - \frac{D}{S^2}} \end{aligned} \quad (2.57)$$

so that the lower bound α_{LB} takes on its maximum value of α only when the discriminant is zero, that is when $A = B$. Provided at least one such point exists in the spectrum, the greatest lower bound on α is the value of α itself.

Again, some care is needed to account for the effect of noise, although the results of preliminary simulations are encouraging. More generally, the combination of computed lower bounds, calculations from first principles, and comparison of recovered spectra to measured spectra should allow the value of α to be determined accurately.

2.7.5 Summary

In principle, the FDH hologram contains enough information to allow a full reconstruction without a separate measurement of the reference spectrum, eliminating several of the challenges of the current FDH technique. Although the algorithm I have outlined in this section produces acceptable results with simple example data, I have not attempted to use it as part of an FDH analysis. As well as the effect of noise on the results, a significant problem is likely to be the finite resolution of the spectrometer, which reduces the visibility of the spectral fringes by an unknown amount. Successful implementation of this direct spectral recovery method requires that this reduction

factor be known, so that the reduction can be corrected. Possible methods of finding this factor include calculation from first principles, comparison of recovered and directly-measured reference spectra, and the use of lower bounds computed across multiple data sets. However, once measured the value should be extremely stable, potentially resulting in significant improvements in the quality of the results of the FDH analysis.

2.8 Other wakefield diagnostics

FDH is not the only diagnostic technique that can be used to measure a plasma wakefield. There are various techniques that work on the same longitudinal probing method as FDH, where a probe pulse is nearly stationary relative to the wakefield, but use different methods to measure the imprinted phase. Other techniques include plasma shadowgraphy [64, 65], which uses a transverse probe pulse to detect the effects of electron density gradients in the wakefield, although the requirement for an extremely short probe pulse made it impractical for this experiment. Since this experiment was only the second time that FDH has been implemented, I evaluated some of these other diagnostics as possible alternatives or complements to FDH.

In this section, I describe two diagnostic techniques that operate on similar principles to FDH; a single probe pulse is propagated behind the pump, then separated and analysed by a spectrometer. In contrast to FDH, these techniques rely on a periodic phase modulation caused by the wakefield changing the spectrum itself. The first method relies on the generation and detection of spectral side-bands; the second method relies on detecting the periodic modulation of the spectrum of a chirped probe pulse.

In particular, a linear wakefield as described in § 1.1.4 will result in a periodic phase modulation of the probe pulse in accordance with Equation 2.5:

$$\phi(t, x, y) = \phi_0(x, y) \sin(\omega_p t) \quad (2.58)$$

Although this phase shift cannot be measured directly, its effect on the electric field might be detected by measuring the spectrum at the end of the interaction region:

$$E(t, x, y) = E_0(t, x, y) \exp i\phi(t, x, y) \quad (2.59)$$

$$I_s(\omega, x, y) = |\mathcal{F}\{E(t, x, y)\}|^2 \quad (2.60)$$

2.8.1 Spectral side-bands

Any kind of periodic modulation of a carrier wave causes spectral side-bands to be generated. In the case of phase modulation, multiple side-bands are generated at integer multiples of the modulation frequency (here ω_p), and their amplitudes depend on the Bessel functions of the first kind. For instance [81]:

$$\begin{aligned} \sin(\omega_0 t + \phi_0 \sin(\omega_p t)) &= J_0(\phi_0) \sin(\omega_0 t) \\ &+ J_1(\phi_0) \left[\sin((\omega_0 + \omega_p)t) - \sin((\omega_0 - \omega_p)t) \right] \\ &+ \text{other side-bands} \end{aligned} \quad (2.61)$$

Provided ω_p is sufficiently large compared to the width of the spectrum, these side-bands can be detected and used as evidence of a wakefield. Their separation from the base-band indicates the plasma frequency and their amplitude indicates the phase modulation depth (which in turn can be used to calculate the amplitude of the electron density oscillation). In particular, for small modulation depth, the following approximations:

$$J_0(x) \approx 1 + O(x^2) \quad (2.62a)$$

$$J_1(x) \approx \frac{1}{2}x + O(x^3) \quad (2.62b)$$

show that the relative side-band amplitude is $\phi_0/2$ to first order in ϕ_0 , and therefore the relative intensity is $\phi_0^2/4$.

2.8.2 Chirped-pulse spectral modulation

An alternative method is to apply a linear chirp to the pulse, so that the periodic modulation is applied in the frequency domain as well as the time domain. Then, a Fourier transform should result in side-bands appearing in the time domain. The advantage of this method is that the spectral modulation is caused by interference between the side-bands and the base-band, and is thus more intense than the side-bands alone. Naturally, this interference can only occur when the plasma frequency is less than the width of the spectrum so that the side-bands overlap the base-band, the opposite of the condition for the above spectral side-band method.

It is instructive to consider this method analytically, although my derivation below requires a small plasma frequency, relative to the laser bandwidth. Since this does not hold for this experiment, I also briefly present a numerical simulation of the method.

First consider a chirped Gaussian pulse, represented here in the frequency domain:

$$\tilde{E}(\omega) = \tilde{E}_0 \exp \left[-\frac{\Delta\omega^2}{\Omega^2} + i\frac{\beta}{2}\Delta\omega^2 \right] \quad (2.63)$$

where $\Delta\omega = \omega - \omega_0$, ω_0 is the centre frequency, Ω is the bandwidth and $\beta = \phi^{(2)}(\omega_0)$ is the group delay dispersion.

Applying a sinusoidal phase modulation, as described in Equation 2.61, results in an electric field of the form:

$$\begin{aligned} \tilde{E}(\omega) = & J_0(\phi_0) \tilde{E}_0 \exp \left[-\frac{\Delta\omega^2}{\Omega^2} + i\frac{\beta}{2}\Delta\omega^2 \right] \\ & + J_1(\phi_0) \tilde{E}_0 \exp \left[-\frac{(\Delta\omega - \omega_p)^2}{\Omega^2} + i\frac{\beta}{2}(\Delta\omega - \omega_p)^2 \right] \\ & - J_1(\phi_0) \tilde{E}_0 \exp \left[-\frac{(\Delta\omega + \omega_p)^2}{\Omega^2} + i\frac{\beta}{2}(\Delta\omega + \omega_p)^2 \right] \\ & + \text{other side-bands} \end{aligned} \quad (2.64)$$

Since $\omega_p \ll \Omega$ (the small plasma frequency condition mentioned previously), the differences in the $\exp \left[-\frac{\Delta\omega^2}{\Omega^2} \right]$ terms can be neglected, resulting in the following ap-

proximation:

$$\begin{aligned}\tilde{E}(\omega) \approx \tilde{E}_0 \exp \left[-\frac{\Delta\omega^2}{\Omega^2} \right] & \left(J_0(\phi_0) \exp \left[i\frac{\beta}{2}\Delta\omega^2 \right] \right. \\ & + J_1(\phi_0) \exp \left[i\frac{\beta}{2}(\Delta\omega - \omega_p)^2 \right] \\ & \left. - J_1(\phi_0) \exp \left[i\frac{\beta}{2}(\Delta\omega + \omega_p)^2 \right] \right)\end{aligned}\quad (2.65)$$

$$\begin{aligned}\tilde{E}(\omega) \propto J_0(\phi_0) + J_1(\phi_0) \exp \left[i\frac{\beta}{2}(-2\omega_p\Delta\omega + \omega_p^2) \right] \\ - J_1(\phi_0) \exp \left[i\frac{\beta}{2}(2\omega_p\Delta\omega + \omega_p^2) \right]\end{aligned}\quad (2.66)$$

$$\tilde{E}(\omega) \propto J_0(\phi_0) - 2iJ_1(\phi_0) \exp \left[i\frac{\beta}{2}\omega_p^2 \right] \sin(\beta\omega_p\Delta\omega) \quad (2.67)$$

The spectral intensity is proportional to $|\tilde{E}|^2$, so:

$$\begin{aligned}I_s(\omega) \propto J_0(\phi_0)^2 + 4J_0(\phi_0)J_1(\phi_0) \sin \left(\frac{\beta}{2}\omega_p^2 \right) \sin(\beta\omega_p\Delta\omega) \\ + 4J_1(\phi_0)^2 \sin^2(\beta\omega_p\Delta\omega)\end{aligned}\quad (2.68)$$

The second term is the spectral modulation, which has a period

$$\Delta\omega_m = \frac{2\pi}{\beta\omega_p} \quad (2.69)$$

and a modulation depth of order ϕ_0 , which is an improvement on the ϕ_0^2 scale of the side-bands themselves.

To isolate this modulation, the spectrum can be Fourier transformed, resulting in side-bands appearing at:

$$\begin{aligned}t &= \pm 2\pi / \Delta\omega_m \\ &= \pm \beta\omega_p\end{aligned}\quad (2.70)$$

Using β as a scaling factor, it is possible to convert the independent co-ordinate of the Fourier transform from t to ω_p , and plot the signal as a function of plasma frequency.




FIGURE 2.28 – The Fourier transform of a simulated Gaussian pulse, chirped and modulated by a sinusoidal phase. The chirp parameter β was used to re-scale the independent axis to plasma frequency. The FWHM bandwidth of the input pulse and the plasma frequency were both 7 THz.

An example of this is shown in Figure 2.28, which is the result of a numerical simulation using a Gaussian spectrum. In this case, the plasma frequency is comparable to the bandwidth, but a side-band still appears at the correct position, showing that the small plasma frequency condition $\omega_p \ll \Omega$ is not necessary in practice.

However, the plasma frequency is still constrained – if it is much larger than the bandwidth, then no spectral modulation appears, since the side-bands of Equation 2.64 are separated from the base-band and cannot interfere. Conversely, if the plasma frequency is too small, the side-bands in the Fourier transform are very close to the base-band and it may not be possible to identify them.

There is a minimum value of the Fourier transform co-ordinate t , which must be found by taking experimental data and identifying the place where the Fourier transform drops below the noise floor. Then, for a fixed plasma frequency ω_p , there is a

minimum value of the chirp parameter β :

$$\beta > \frac{t_{min}}{\omega_p} \quad (2.71)$$

2.8.3 Comparison to FDH

Both these methods (spectral side-bands and chirped-pulse spectral modulation) are simpler to set up than FDH, requiring only a single probe pulse and no separate reference measurements. Their analysis methods are also simpler than the complex process described in § 2.4. The spectral modulation method offers good sensitivity; a spectrum with a signal-to-noise ratio of 400:1 (which is the value found in § 2.5.3) would be expected to provide sensitivity better than 10 mrad. This signal-to-noise ratio is likely to be an underestimate, since the quoted value is based on the peak-to-peak amplitude of the spectral fringes, whereas the spectral modulation method uses the whole of the spectral amplitude.

Another advantage of these alternative methods is that in FDH, the reference pulse must propagate in front of the pump pulse, and it must also be fairly close to the probe pulse (no more than about 10 ps) to ensure the spectral fringes are wide enough to be resolved. This limits how far behind the pump pulse the probe pulse can be, making FDH unsuitable for probing the wakefield at large times after the pump pulse. Although the reference pulse could be propagated beside the pump, this significantly complicates the setup. The simpler methods described in this section have no such limitation, because they use only a single probe pulse propagating after the pump.

However, a working FDH system provides several advantages of its own. The main one is the ability to resolve and measure non-periodic modulations, whereas the alternative methods described in this section are limited to measuring the amplitude and period of sinusoidal modulations. In particular, FDH would be able to measure the amplitude of each individual period of a wakefield, which is useful in a multiple-pulse experiment, as well as non-periodic features, such as the ionisation front. In

addition, very high laser intensities excite wakefields in the bubble régime [56]; these are non-periodic, and cannot be measured by these simpler alternatives to FDH.

3. Experimental set-up

An experiment was conducted at Oxford University to investigate the feasibility of MP-LWFA. The aim was first to generate a wakefield using a single pump pulse and measure it using FDH, then to split that pump pulse up into a train of pulses to demonstrate excitation of a wakefield by multiple pulses. However, as the experiment progressed, it became clear that the diagnostic was less sensitive than originally predicted and that the laser energy was less than I had expected, making it impossible to observe a wakefield. In this chapter, I describe most of the equipment used for the experiment; the experimental setup for taking FDH data is described in more detail in § 2.3 and possible methods of pulse train generation are covered in Chapter 6.

A TW-class laser was used to generate synchronised pairs of pump and probe pulses. The probe pulse was used to generate an FDH probe/reference pulse pair, which was combined with the pump and passed into a vacuum chamber. Here they were focussed by a paraboloidal mirror (the larger beam size of the pump ensuring that its size at the focus was smaller than the probe) into a gas cell containing helium. The pump ionised the helium and generated a wakefield in the resulting plasma, causing a phase shift in the probe pulse. The probe and reference pulses were separated from the pump by a dichroic mirror, then directed out of the vacuum chamber and analysed by an imaging spectrometer.

In the following sections, I describe these components in more detail.

3.1 Laser system

The laser used for the experiment is a chirped-pulse amplification (CPA) [82] system, consisting of a mode-locked oscillator, a pulse stretcher, a series of amplifiers, and a pulse compressor.

The oscillator is a Spectra-Physics Mai-Tai SP, a mode-locked titanium-sapphire laser producing 6.5 nJ pulses at a repetition rate of 84 MHz. The pulses are specified to have a bandwidth of 55 nm centred at 810 nm, and a full-width at half-maximum (FWHM) duration of 35 fs. After being stretched in a standard two-grating stretcher (see e.g. [83, § II.C]), the repetition rate is reduced to 1 kHz and the remaining pulses are amplified in a regenerative amplifier. This amplifier is a Spectra-Physics Spitfire, which amplifies the pulses to an energy of approximately 1 mJ.

After leaving the regenerative amplifier, the pulses' repetition rate is reduced again to 10 Hz, and they are amplified in a Positive Light multi-pass amplifier. The gain medium is pumped by a Spectra-Physics Quanta-Ray PRO-Series, a frequency-doubled Q-switched Nd:YAG laser producing green (532 nm) pulses with up to 1 J of energy. Four passes through the gain medium amplify the seed pulses to approximately 100 mJ, then they enter the pulse compressor.

The pulse energy on entering the compressor is measured by a Gentec-EO QE25ELP pyroelectric energy meter, with an attached attenuator to reduce the incident intensity and protect the energy meter head; the attenuation factor was measured to be 4.15 ± 0.03 and all energy values quoted have been corrected for this attenuation. The pulse energy was sensitive to the alignment of the multi-pass amplifier; a typical value at the conclusion of the experiment was 98 mJ, with a standard deviation of 5.2 mJ. Peak-to-peak energy variation was as high as 30%. Such instability is typical for this kind of laser system.

Before compression, a small amount (roughly 1%) of the pulse is transmitted through one of the mirrors. This will be compressed separately and then used to form the FDH probe and reference pulses. The remainder (pump pulse) is then expanded in a Galilean




FIGURE 3.1 – The spectrum of the pump pulse.

telescope to a diameter of roughly 40 mm and then passed through a standard double-pass grating compressor [83, § II.C]. This compresses the pulse to close to its transform limit, with a FWHM duration of approximately 50 fs.

Energy transmission through the compressor is significantly less than 100%. The throughput of the compressor itself was not measured, since the output beam was too large to fit on the energy meter. Instead, the transmission from the entrance to the compressor to the interaction point (where the pump pulse is focussed inside the vacuum chamber) was measured, and found to be about 30%.

3.1.1 Spectrum

The spectrum of the pump pulse after exiting the compressor was measured using an Andor Shamrock imaging spectrometer (described in more detail in § 3.7); the results are shown in Figure 3.1. The FWHM bandwidth of 20 nm suggests that the spectrum should support pulses with a FWHM duration as short as 50 fs (using a Gaussian time-




FIGURE 3.2 – A top-down view of the custom-built autocorrelator used to measure the pulse duration.




FIGURE 3.3 – A side-on view of part of the autocorrelator pictured in Figure 3.2.

bandwidth product of 0.44). The fundamental probe pulse (before being frequency-doubled) will have the same spectrum.

3.1.2 Autocorrelation measurements

The duration of the pulse was measured using a custom-built single-shot autocorrelator based on the design presented in [84] that was also able to detect pulse-front tilt in the horizontal plane. The design is shown in Figure 3.2. The pulse to be analysed enters at the bottom-left corner; the two irises are used for alignment. It is split at a pellicle beam-splitter and the two parts are re-combined at the SHG crystal, with a small vertical angle between the two beams. This ensures that in addition to the two second-harmonic pulses generated, there will be a sum-frequency generated pulse propagating between them (shown in Figure 3.3).

This sum-frequency pulse is only generated in those parts of the crystal where both




FIGURE 3.4 – An autocorrelation of the pump pulse. **Top:** the image formed on the CCD; the vertical axis corresponds to autocorrelation delay. Structure from the beam profile is also visible, especially diffraction rings due to the small aperture of the autocorrelator. **Bottom:** an average over many columns of the autocorrelation data, to reduce the effect of structure in the beam profile. A Gaussian fit to the central peak, with FWHM 71.8 fs, is also shown.

pulses overlap; the longer the pulses, the greater the vertical extent of this overlap, so a vertical line-out through the sum-frequency pulse gives the autocorrelation. The residual fundamental pulses are removed by a short-pass filter, and the frequency-doubled pulses are imaged onto a CCD. A pinhole placed in the focal plane of the lens blocks the two SHG pulses.

Because of the different geometries of the two arms, one pulse arrives mirrored horizontally relative to the other. If there is any pulse-front tilt in the horizontal plane, this will cause the delay between the two pulses to vary horizontally. Since delay corresponds to the height of the autocorrelation trace, a slope on the trace indicates pulse-front tilt.

A measured autocorrelation of the pump pulse is shown in Figure 3.4. The delay scale was calibrated by adjusting the length of one of the autocorrelator arms and

measuring how far the autocorrelation trace moved vertically. The autocorrelation has the form of a central Gaussian peak, with a FWHM of 71.8 fs, and a pair of broad wings over 200 fs wide. This suggests that the temporal pulse profile consists of a narrow Gaussian, with a FWHM of 50.8 fs, and one or more pre- or post-pulses (which cause the wings to appear).

Since the wakefield is excited by intensity gradients, the pre- or post-pulse (which appears both less intense and longer in duration than the central pulse) will not excite a significant wakefield. I therefore decided to use the autocorrelation to measure what fraction of the energy was in the central peak, although ideally this would be calculated by using a diagnostic such as FROG [75, 76] or SPIDER [77] to measure the temporal profile of the pulse.

This measurement might be made by looking at the areas of the central peak and wings on the autocorrelation, but it is not clear how to convert that ratio to one applicable to the temporal profile. Instead, I converted the measured autocorrelation to an estimated temporal profile. Unlike FROG or SPIDER, this will not be particularly accurate, but it should preserve any important features that appear in the autocorrelation, including the ratio of the wings to the central peak.

An intensity autocorrelation $A(\tau)$ is simply the (mathematical) autocorrelation of the temporal pulse profile $I(t)$:

$$A(\tau) = \int_{-\infty}^{+\infty} I(t)I(t - \tau)dt \quad (3.1)$$

Fourier transforming this gives:

$$\tilde{A}(\omega) = \left| \tilde{I}(\omega) \right|^2 \quad (3.2)$$

Note that $\tilde{I}(\omega)$ here is *not* the spectral intensity — it is the Fourier transform of the temporal intensity profile.

The autocorrelation signal $A(t)$ allows the magnitude of $\tilde{I}(\omega)$ to be computed, but




FIGURE 3.5 – An estimate of the temporal profile of the pump pulse, made by inverting the autocorrelation shown in Figure 3.4. The shaded area consists of 58% of the total area, suggesting that only 58% of the total energy is found within the central pulse, and that the remainder is contained in pre- or post-pulses.

not its phase. Inverting the Fourier transform then presents a 1D phase retrieval problem [85, 86], which is not solvable in the general case due to the existence of multiple valid solutions, and is still difficult even when assumptions are made about the solution. I therefore chose to assume that $\tilde{I}(\omega)$ has flat phase. This will not reproduce the correct temporal profile; in particular, it will be symmetrical around $t = 0$, whereas the real profile is expected to have either pre-pulses or post-pulses but not both. However, this calculation is only intended to provide an estimate of the fraction of the total energy in the central peak, so the exact shape is not important.

This gives the simple calculation:

$$I(t) \approx \mathcal{F}^{-1} \left\{ \sqrt{|\mathcal{F}\{A(\tau)\}|} \right\} \quad (3.3)$$

The result of this calculation applied to the data in Figure 3.4 is shown in Figure 3.5.




FIGURE 3.6 – A digitally-enhanced photograph of the profile of the pump beam, taken by illuminating a piece of paper with the laser. The line labelled “1 cm” was drawn on the paper, although it has been digitally altered from black to white for better visibility.

The shaded central peak contains 58% of the total area, so I estimate that 58% of the measured laser pulse energy is contained within a Gaussian with 50.8 fs FWHM duration. Given an approximate throughput of 30% from the compressor to the interaction point, this gives a useful energy of 17 mJ and a peak power of 200 GW.

3.1.3 Beam profile

After being expanded in the compressor, the beam was too large to fit on any readily available beam profile monitor. In order to record the profile of the pump pulse, I used it to illuminate a piece of paper and took a photograph using a digital camera. The resulting image, after some digital enhancement, is shown in Figure 3.6.

The beam is clearly elliptical, being wider in the vertical direction than the horizontal. Furthermore, its profile is neither Gaussian nor uniform, but irregular; bright

and dark spots are visible, and faint circular features on the edges show where the beam has clipped the edges of earlier optics due to its large size. These will all affect the quality of the focus.

3.2 FDH optics

The optics used to generate the probe/reference pulse pair needed to perform FDH have already been described in § 2.3, so this section will focus on how those optics fit into the rest of the experiment. I will also include measurements of the chirp and spectrum of the generated pulses.

3.2.1 Probe compressor

Although a probe pulse could be generated by picking off a small amount of the compressed pump pulse, one of the pulse-train generation methods proposed for this experiment requires the pump compressor to be significantly offset from optimal compression. Compressing the probe pulse separately ensures that this will not affect the probe pulse; in particular the SHG process would not be efficient if the probe pulse were not properly compressed.

The probe compressor was a custom design, shown in Figure 3.7. The geometry of the pump compressor was measured, and the amount of dispersion it introduced was calculated. The probe compressor was designed to introduce an identical amount of dispersion, so that the probe pulse would be compressed as well as the pump pulse. Although the group delay dispersion was corrected in this way, there was a certain amount of residual third-order dispersion ($\frac{d^3\phi}{d\omega^3}$) after compression; this affected the autocorrelation of the compressed pulse (shown in Figure 3.8) and the spectrum of the frequency-doubled probe and reference pulses. It would have been possible to re-design the probe compressor to correct for this, but lack of space and equipment made this difficult; I decided that the effect on the frequency-doubled spectrum was not severe enough to warrant such an effort.




FIGURE 3.7 – The probe compressor. The first diffraction grating (a) is mounted on a translation stage to control the inter-grating separation. Both gratings (a) and (b) also have rotation control to allow spatio-temporal couplings to be corrected. (c) is a plane mirror angled slightly downwards so that the returning beam strikes the semi-circular mirror (d), while the incoming beam passes above it.




FIGURE 3.8 – An autocorrelation of the probe pulse before being frequency doubled. See also Figure 3.4. **Top:** the image formed on the CCD. **Bottom:** an average over many columns of the autocorrelation data. The asymmetry is caused by the small beam profile, which acts as an envelope to the autocorrelation signal. The large pre-/post-pulses are a typical symptom of uncompensated third-order dispersion [83, Fig. 24].




FIGURE 3.9 – The separate pump and probe beam-lines, and how they fit together. (P) and (R) are the SHG crystals that generate the probe and reference pulses respectively. (W-O) is the walk-off glass and (Str) is the stretching glass. These are described in more detail in § 2.3 and can be replaced by another of the setups shown in that section to take reference measurements (for which the pump is not needed). (D) is a short-pass dichroic mirror used to combine the pump and probe beams. Pulse train generation optics could be installed in the space between the pump compressor and the dichroic mirror.

3.2.2 Pulse pair generation

The FDH probe/reference pulse pair is generated as described in § 2.3; a pair of 100 μm thick β -barium borate (BBO) crystals, phase-matched for second harmonic generation, are separated by a length of glass to separate the 400 nm and 800 nm pulses temporally. The pulse pair is then sent through an adjustable delay line before being stretched by 160 mm of BK7 glass and combined with the pump beam at a short-pass dichroic mirror. The delay line allows control of the pump-probe relative timing.

A schematic of the whole setup is shown in Figure 3.9. References were taken




FIGURE 3.10 – The spectra of the probe and reference pulses, after being normalised to the same total energy. The line represents the mean of 45/47 samples (probe/reference respectively), and the shaded area either side denotes plus/minus one standard deviation. The small scale fluctuations in the reference spectrum are an artefact of the measurement process, as described in § 2.3.5.

by blocking the pump (before the compressor) and rearranging the FDH optics, as described in § 2.3. This allowed the same spectrometer to be used without needing to move or realign it.

3.2.3 Pulse spectra

The performance of FDH depends on the shape of the probe and reference pulse spectra. The spectra should ideally be smooth, with a large bandwidth and no gaps. The stability of the spectra is also important, since FDH analysis relies on several different measurements, taken at different times. Typical spectra of the reference and probe pulses are shown in Figure 3.10.

Although the two pulses are supposed to be identical, they clearly have different spectra. This is because the shape of the second-harmonic spectrum depends on the chirp of the fundamental pulse and the angle of the crystal. Since two different crys-




FIGURE 3.11 – Three different measurements of the reference spectrum, taken from the data set used to produce Figure 3.10, to show how the shape of the spectrum changes. Note the slight variations in the position of the central peak. Measurements are normalised to correct for differing pulse energy, which has little effect on the shape of the spectrum.

tals were used, their angles may have been slightly different, and the walk-off glass changes the chirp of the fundamental pulse between the two SHG crystals.

The spectra also change measurably from one shot to another, as shown in Figure 3.11. This variation, especially in the position of the central peak, causes artefacts to appear during FDH reconstruction, as described in § 2.6.4.

The spectrum can be optimised by altering the angle of the SHG crystal and the separation of the gratings in the probe compressor. The effects of these changes are shown qualitatively in Figure 3.12. I was unable to improve the spectrum significantly from its initial form.

Comments

Given the zero-to-zero bandwidth of the fundamental spectrum (as seen in Figure 3.1) of 50 nm, the expected zero-to-zero bandwidth of the second-harmonic spectrum is 14 nm (which can be derived from $\frac{\Delta f}{f} = \frac{\Delta \lambda}{\lambda}$ and $\Delta f_{SHG} = \sqrt{2}\Delta f$). The probe and




FIGURE 3.12 – The effects of moderate changes to (a) the grating separation in the compressor and (b) the angle of the SHG crystal on the reference spectrum. The nominal settings were chosen by adjusting the compressor separation to optimise the autocorrelation, and adjusting the crystal angle to optimise the bandwidth.

reference spectra seen in this section are narrower than that. In addition, the spectra are very uneven, with a narrow peak in the centre and significant troughs. These properties negatively affect the performance of FDH: the smaller bandwidth limits the temporal resolution of the measurements, the gaps in the spectrum introduce significant noise into the analysis and the sharp edges of the central peak cause anomalous phase to be measured.

It was unclear whether these problems were due to residual third-order dispersion (and therefore could be corrected by redesigning the probe compressor) or inherent to the laser. To determine this, I frequency-doubled the pump pulse and measured its spectrum, on the basis that this would eliminate any problems caused by third-order dispersion. The result is shown in Figure 3.13, where it is compared with the spectrum of the reference pulse; although the frequency-doubled pump has somewhat more bandwidth than the reference pulse, it exhibits similar structure, suggesting that redesigning the probe compressor will not lead to a significant improvement in the




FIGURE 3.13 – Comparison of the second-harmonic spectrum of the pump pulse with the spectrum of the reference pulse (as seen in Figure 3.10). Translating the reference spectrum (the dashed line) shows that the pump SHG spectrum has more bandwidth.

shape of the spectrum.

3.2.4 Probe/reference chirp

The probe and reference pulses will both have a significant amount of chirp, due to the group delay dispersion (GDD) introduced by the stretching glass and the walk-off glass. There will also be some GDD and higher order dispersion left over after the compressor, although I have not measured this¹.

The amount of GDD introduced by the glass can be calculated using the Sellmeier equation²:

$$n^2 - 1 = \frac{B_1 \lambda^2}{\lambda^2 - C_1} + \frac{B_2 \lambda^2}{\lambda^2 - C_2} + \frac{B_3 \lambda^2}{\lambda^2 - C_3} \quad (3.4)$$

where n is the refractive index, λ is the vacuum wavelength in μm and the B_i and C_i parameters depend on the material. Values for these parameters can usually be

¹An attempt was made to characterise the pulse using a GRENOUILLE [87], but the pulse duration was too long for the particular device which was available.

²This is only one of several different forms in use.

Parameter	Value	Units
B_1	1.040	
C_1	6.001×10^{-3}	μm^2
B_2	0.232	
C_2	2.002×10^{-2}	μm^2
B_3	1.010	
C_3	1.036×10^2	μm^2

TABLE 3.1 – The parameters to the Sellmeier equation (3.4) applicable to N-BK 7 glass [88, p. 13].

obtained from the manufacturer of the glass.

The group velocity dispersion (GVD) can then be calculated:

$$\text{GVD} = \frac{\partial}{\partial \omega} \left(\frac{1}{v_g} \right) = \frac{\partial^2 k}{\partial \omega^2} = \frac{2}{c} \frac{\partial n}{\partial \omega} + \frac{\omega}{c} \frac{\partial^2 n}{\partial \omega^2} \quad (3.5a)$$

$$\omega = \frac{2\pi c}{\lambda} \quad (3.5b)$$

$$\frac{\partial n}{\partial \omega} = \frac{\partial \lambda}{\partial \omega} \frac{\partial n}{\partial \lambda} = -\frac{\lambda^2}{2\pi c} \frac{\partial n}{\partial \lambda} \quad (3.5c)$$

$$\frac{\partial^2 n}{\partial \omega^2} = \frac{\partial \lambda}{\partial \omega} \frac{\partial}{\partial \lambda} \left(\frac{\partial n}{\partial \omega} \right) = \frac{\lambda^3}{(2\pi c)^2} \left(2 \frac{\partial n}{\partial \lambda} + \lambda \frac{\partial^2 n}{\partial \lambda^2} \right) \quad (3.5d)$$

$$\text{GVD} = \frac{\partial^2 k}{\partial \omega^2} = \frac{\lambda^3}{2\pi c^2} \frac{\partial^2 n}{\partial \lambda^2} \quad (3.5e)$$

Note that the wavelength λ is measured in vacuum, so it is related to the angular frequency ω rather than the wave number k . The derivatives of n with respect to λ can be computed from the Sellmeier equation (3.4):

$$n \frac{\partial n}{\partial \lambda} = - \sum_{i=1}^3 \frac{B_i C_i \lambda}{(\lambda^2 - C_i)^2} \quad (3.6a)$$

$$n \frac{\partial^2 n}{\partial \lambda^2} = - \left(\frac{\partial n}{\partial \lambda} \right)^2 - \sum_{i=1}^3 \frac{B_i C_i (3\lambda^2 + C_i)}{(\lambda^2 - C_i)^3} \quad (3.6b)$$

The parameters for the BK 7 glass used in this experiment (shown in Table 3.1) give a GVD of $117 \text{ fs}^2/\text{mm}$ at $\lambda = 410 \text{ nm}$. The total GDD introduced by 160 mm of stretching glass and 32 mm of walk-off glass is therefore $2.25 \times 10^4 \text{ fs}^2$. The reference pulse will suffer somewhat less GDD, because it propagates through the walk-off glass before being frequency-doubled, and the GVD is generally lower at longer wavelengths.

The dispersion introduced by these two pieces of glass was measured many times, since it forms one of the FDH references. This data can be compared to the above prediction, but it can also be used to measure the day-to-day stability of the chirp and check whether significant third-order dispersion is introduced by the glass.

Summaries of some of the probe and reference chirp measurements are shown in Figure 3.14 and Figure 3.15 respectively. Many of the measurements of the GDD of the probe pulse are close to the predicted value of $2.25 \times 10^4 \text{ fs}^2$. These values are the results of fitting a second-order polynomial to the measured phase function. This phase function will contain a certain amount of random noise, which affects both the fitted value and the residuals from the fit.

The effect on the fitted values can be seen in the summary plots, where deviations of up to 200 fs^2 within a single data set are typical. Such small errors, on the 1% level, will have a proportional effect on the scale of the time axis of the reconstructed FDH data, but they are not expected to have any more severe effect. On the other hand, significant deviations are sometimes seen from one day to another. This shows the importance of taking new reference data every day, so that these errors do not affect the FDH reconstruction.

These fits only use second-order polynomials. If there were a significant amount of third-order dispersion, it would be visible in the residuals as a third-order polynomial term. Residuals from one fit are shown in Figure 3.16; noise can be seen, but no obvious cubic dependence that would indicate third-order dispersion.

3.3 Gas and vacuum equipment

Many ultra-fast high intensity lasers need to propagate entirely in vacuum, since the material dispersion and non-linear effects of propagation through the atmosphere would result in the pulses being stretched and distorted. That is not the case for the laser used in this experiment, whose pulses are neither particularly short nor intense by modern standards. Nevertheless, the entire interaction region — the focussing




FIGURE 3.14 – Measurements of the GDD of the probe pulse made over the course of the experiment.




FIGURE 3.15 – Measurements of the GDD of the reference pulse made over the course of the experiment.

paraboloid, gas cell, and pump/probe separator — is enclosed within a vacuum chamber. Crucially, this allows the pressure within the gas cell to be substantially less than atmospheric pressure, which is necessary to ensure that the plasma wavelength is long enough to be resolved. It also eliminates non-linear interactions with ambient air near the focus.

The main vacuum chamber is rectangular, approximately $1450 \text{ mm} \times 650 \text{ mm}$, and 450 mm tall. It can be evacuated by a scroll pump, which can reduce the pressure to 0.6 mbar (measured using an MKS Instruments Series 910 DualTrans transducer, combining Piezo and MicroPirani sensors).

In addition to the entrance and exit windows for the laser, the vacuum chamber features a number of feed-throughs, including an electrical feed-through to operate a camera inside the chamber and a gas feed-through to supply the gas cell with helium.

The gas cell itself acts as the target for the experiment, containing the gas which is ionised by the pump pulse and in which a wakefield is excited. The change in the focal length of the paraboloid partway through the experiment (described in more detail in § 3.5) was accompanied by a change to the design of the gas cell. As such, the gas cells are described in more detail in § 3.5.4 and § 3.6.6. Both of them follow the same principle, as they allow the pump laser to drill entry and exit holes through the thin metal walls of the cell; gas escapes through these holes and is refilled through a gas inlet, connected to a gas supply line. The pressure in this supply line is monitored using a second Series 910 transducer, placed just outside the vacuum chamber. The pressure inside the gas cell is assumed to be equal to the pressure in the supply line.

3.4 Gaussian beam model

An important part of characterising a laser pulse, so that predictions can be made about the wakefield it will excite, is understanding its transverse intensity profile and how that changes throughout the focal region. In principle, this can be accomplished by a single measurement of both amplitude and phase in a single plane (which can be made

directly with e.g. a Shack-Hartmann wavefront sensor [89], or reconstructed from several amplitude-only measurements in different planes via the Gerchberg-Saxton algorithm [90]), since this uniquely determines, via the propagation equation, the profile in every other plane.

However, in practice this is not so useful, since the equations and simulations that might be used to predict the wakefield amplitude are rarely set up to handle arbitrary input data. In particular, analytical models (such as that described in § 4.5) require that the profile have a simple analytical form, and the particle-in-cell code EPOCH (see § 4.1) also requires the laser amplitude and phase to be expressed analytically. Some analytical model must be fit to the experimental data; one widely used approach is to model the beam profile as a Gaussian.

The Gaussian beam model [91, Chap. 6] is based on the following exact solution to the propagation equation:

$$E(r, z) = E_0 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w(z)^2}\right) \exp\left(-i \left[kz + \phi_G + \frac{kr^2}{2R(z)} \right]\right) \quad (3.7)$$

where

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2} \quad (3.8)$$

$$R(z) = z \left[1 + \left(\frac{z_R}{z}\right)^2 \right] \quad (3.9)$$

$$\phi_G = \arctan \frac{z}{z_R} \quad (3.10)$$

$$z_R = \frac{\pi w_0^2}{\lambda} \quad (3.11)$$

where $w(z)$ is the e^{-2} intensity radius, w_0 is the minimum value of w (which occurs at the beam waist or focus, here taken to be at $z = 0$), $R(z)$ is the radius of curvature of the wavefronts, ϕ_G is the Gouy phase shift, and z_R is the Rayleigh range (or Rayleigh length).

The Gaussian beam has the smallest product of focal spot size and divergence angle

of any possible beam profile. Therefore, any real beam will have either a larger spot or a wider divergence angle than that predicted by the above model. It is nevertheless desirable to be able to model such beams simply; the usual method used is the so-called M^2 model, where the intensity profile is assumed to be a Gaussian, M times larger than that predicted by the Gaussian beam model. In the usual formulation, w_0 remains the e^{-2} intensity radius at the focus, but the Rayleigh range is divided by a factor of M^2 . Every real beam will have $M \geq 1$.

This model can be extended to cover various other kinds of non-ideal behaviour by real beams. The simplest such extension is to model the beam profile as an ellipse, by selecting two principal axes and describing the radius along each axis. This allows different values of w_0 and M^2 along each axis, allowing for various kinds of elliptical beams. The position of the waist, z_0 , can also be allowed to differ for each axis, allowing astigmatic beams to be modelled. The full model for the intensity profile is thus:

$$I(x, y, z) = \frac{2P}{\pi w_x(z) w_y(z)} \exp -2 \left(\frac{x^2}{w_x(z)^2} + \frac{y^2}{w_y(z)^2} \right) \quad (3.12)$$

$$w_{x,y}(z) = w_{(x,y)0} \sqrt{1 + \left(\frac{z - z_{(x,y)0}}{z_{(x,y)R}} \right)^2} \quad (3.13)$$

$$z_{(x,y)R} = \frac{\pi w_{(x,y)0}^2}{\lambda M_{x,y}^2} \quad (3.14)$$

where P is the total power and the beam is described by the parameters w_{x0} , w_{y0} , z_{x0} , z_{y0} , M_x^2 , and M_y^2 .

3.4.1 Modelling a beam as a Gaussian

This model can be used to describe the laser pulse when performing the simulations described in Chapter 4. Since the purpose of these simulations is to predict the wake-field amplitude that is expected to occur in the experiment, the model needs to accurately replicate the behaviour of the real laser pulse. However, there are many aspects of the laser behaviour, and the Gaussian beam model will only be able to reproduce

some of them accurately. In many cases, the aim of this modelling is to predict the propagation of the beam over long distances, so this behaviour is the most important aspect to preserve. However, for the purposes of this experiment the aim is to predict the behaviour of the wakefield, so the important behaviour is that near the focus. I therefore aimed to model the beam in such a way that the peak intensity as a function of propagation distance z was correct, since that is the main driver of the wakefield.

Firstly, several different measurements of the intensity profile at various z positions throughout the focal region are needed. Enough measurements must be made to allow accurate determinations of:

- The positions of the two (x and y) beam waists
- The sizes of the two beam waists
- The two divergence angles

A set of (orthogonal) principal axes is then selected by examining the most elliptical profiles measured. The analysis then proceeds in two stages: firstly, each profile is analysed to determine values for w_x and w_y ; secondly, Equation 3.13 is fit to this data, giving the parameters for the Gaussian beam model. This second step is a straightforward least-squares fit, so it is the method used for the first step that is most important.

I tried two methods to assign values to w_x and w_y given an arbitrary beam profile. Both involved taking line-outs of the data along the principal axes and analysing those. The first method (the 13.5% method) marks the points where the measured intensity drops below 13.5% of its maximum value, and taking the distance between these points to be $2w$. The second method (the 2σ method) involves measuring the standard deviation of the intensity profile and setting w to twice that. These both give the correct value for a true Gaussian beam profile. For arbitrary data, I found that the 13.5% method gave slightly better values (reproducing the peak intensity more accurately) but was somewhat less reliable. An alternative method I considered was

fitting a Gaussian to the measured data, but I found that it produced values that were consistently too small, and it was therefore abandoned.

3.5 Old interaction region

Here, the term “interaction region” refers to all the equipment associated with the generation of a wakefield, in particular the focussing paraboloid, the gas cell, the pump-probe separator and the beam dump. Interaction point (IP) refers specifically to the focus of the pump beam within the gas cell — these terms are somewhat interchangeable, insofar as the design of the IP (focal spot size) determines the design of the interaction region (the paraboloid focal length and gas cell dimensions). Over the course of the experiment two designs for the IP were used; this section describes the first.

The first consideration when designing the interaction region is which gas is to be used. Helium was chosen due to its small atomic number (and hence predictable ionisation level) and the ease of working with it compared to hydrogen. The next considerations are the gas pressure (see calculations of optimal pressure in § 4.3), which affects the period and amplitude of the wakefield, and the focal length of the paraboloid, which affects the focal spot size, and hence the peak intensity and the interaction length.

The original design for the interaction region used a long focal length paraboloid ($f = 1625$ mm). The resulting large focal spot size leads to the generation of a low-amplitude wakefield, but the relatively long Rayleigh length (~ 15 mm) allows it to interact with the probe pulse over at least 10 mm. In principle, this should allow the probe pulse to accumulate a measurable phase shift. However, there was a significant phase shift due to the background plasma (on the order of 10 radians). Based on measurements of this phase shift as a function of pressure (see § 5.1) I concluded that the plasma was acting as a diverging lens, causing the probe pulse to refract away from the axis long before it had reached the end of the intended interaction length. This was

later confirmed by the simulation described in § 4.4.1, where the pump pulse is seen to refract due to this ionisation-induced defocussing [92, 93] after a few millimetres.

3.5.1 Layout

The mirror used was an 8-inch diameter, $f = 1625$ mm, on-axis, gold-coated paraboloid. The on-axis nature of the optic led to difficulties devising a suitable optical layout, since the angle between the incoming and outgoing beams had to remain small (about 1.4°). In order to ensure this small angle, the mirror directing the beam onto the paraboloid had to be placed as close as possible to the IP without obstructing either the focussed beam or the gas cell.

After the IP (the gas cell is described later), a long-pass dichroic mirror is used to separate the probe beam from the residual pump, which is blocked by a beam dump. Limitations of space required this optic to be placed relatively close to the IP point, resulting in a rather high incident pump intensity. The probe beam was collimated by an achromatic lens and directed to the post-IP diagnostics.

3.5.2 Transmitted energy

The energy reaching the IP was deduced by measuring the total transmission of the optics between the entrance of the compressor (where the laser energy was typically measured) and the IP. The beam exiting the compressor was too large to fit on any available energy meter, so the energy was measured somewhere between the paraboloidal mirror and the interaction point (where the beam was smaller). This had to be done at low energy to avoid damaging the energy meter with the high intensity of the focussed beam, but since no non-linear optics are involved the result can be assumed to hold at higher energies. The measured transmission was $29\% \pm 1\%$.




FIGURE 3.16 – Residuals from a fit of a second-order polynomial to some of the reference chirp data. No evidence of third-order dispersion can be seen.




FIGURE 3.17 – The original layout of the interaction region. (IP) denotes the interaction point, where the gas cell is positioned.




FIGURE 3.18 – Profiles of the pump pulse near the focus. The central image is the focal spot (insofar as that term makes sense for an astigmatic beam) and the longitudinal positions are measured relative to it.

3.5.3 Pump focal profile

The profile of the pump beam at the focus was measured using a camera in the post-IP diagnostics area. A small amount of the pump pulse was reflected by the post-IP dichroic mirror and then collimated. A second lens formed an image of the focus with approximately $0.5\times$ magnification. A $4\times$ microscope objective was used to image this onto the CCD of a DataRay beam profile monitor. The magnification of the whole system was measured using a 1951 USAF resolution test chart as $2.0\times$.

By recording multiple beam profiles at different z positions, the beam propagation can be analysed and compared against the Gaussian beam model, described in § 3.4. Three such profiles are shown in Figure 3.18; these clearly show that the beam is astigmatic, i.e. that the z -position of the focal plane is different when measured along two different principal axes. This appeared to be a property of the laser rather than the paraboloid, since it was visible in foci formed using lenses too.

These profiles were analysed using the method described in § 3.4.1 — in particular the 13.5% method. The principal axes were found to be rotated by 30° from the laboratory frame — the rotated axes are referred to as horizontal (x) and vertical (y) based on which laboratory-frame axis they are closest to. The measured values of $w_{x,y}$ and the best fits using Equation 3.13 are shown in Figure 3.19. The fit parameters are shown in Table 3.2. The sizes measured along the vertical axis seem to be a good fit, but the




FIGURE 3.19 – Measured beam sizes and fit using M^2 model (see § 3.4.1).

Parameter	Value		Units
	x	y	
w_0	69.4 ± 1.7	61.9 ± 1.2	μm
z_0	58.5 ± 0.6	83.3 ± 0.4	mm
M^2	1	1.20 ± 0.02	

TABLE 3.2 – The parameters for the fits shown in Figure 3.19. An unconstrained fit on the x -axis resulted in an unphysical value of $M^2 < 1$, so the fit was repeated with a fixed $M^2 = 1$, hence the lack of a quoted uncertainty.

fit to the horizontal data is rather less good.

To check the results of this analysis, I used the measured profiles to calculate the peak fluence (intensity integrated over time) at each z position:

$$\hat{\Phi} = \frac{E\hat{N}}{A\sum_i N_i} \quad (3.15)$$

where $\hat{\Phi}$ is the peak fluence, N_i are the measured pixel values, $\hat{N}$ is the largest measured pixel value, E is the total pulse energy, and A is the effective area of each pixel (i.e. the area in the object plane corresponding to one pixel; this depends on the pixel size and the magnification factor).




FIGURE 3.20 – Comparison of peak fluence as a function of z between: (a) values measured using a CCD; and (b) values predicted using the Gaussian beam model with the parameters from Table 3.2. The model fits the data as well as it could be expected to.

I also used the Gaussian beam model, with the parameters shown in Table 3.2, to calculate the peak fluence as a function of z . This is also proportional to the pulse energy E , so for the purposes of this comparison the actual value of E is irrelevant. I used 15 mJ to give physically meaningful results, although the actual measurements were taken at a lower energy.

The two sets of values for $\hat{\Phi}$ are compared in Figure 3.20. Since the beam is not actually Gaussian, the data is not expected to fit the curve (representing the Gaussian beam model) perfectly, even in the absence of measurement errors. In this case, the fit appears good (for contrast, a bad fit will be shown in Figure 3.25), showing that the Gaussian beam model with the parameters in Table 3.2 is a good model of the beam in the experiment.




Image courtesy of James Cowley.

FIGURE 3.21 – An early design of the gas cell used in the old IP. The gas inlet can be seen on the left. The circular plate on the right holds the end-cap in place over the gas cell itself (centre, mostly hidden). The window visible at the top, originally intended to make transverse measurements of plasma density possible, was not included in the final design.

3.5.4 Gas cell

Although LWFA experiments can use pulsed gas jets and capillaries as targets, it was decided to use a simple gas cell for this experiment. This consisted of a cylindrical cavity, 10–20 mm long (adjustable in 5 mm steps), fed by helium gas through a connection on the top. An early version of this design is shown in Figure 3.21. The ends of the cylinder were covered by 200 μm thick aluminium plates, through which the pump laser drilled small holes. The gas escaped through these holes, requiring it to be refilled through the inlet. The fill rate was controlled by a needle valve and determined the pressure inside the gas cell. The pressure was monitored by a transducer, attached to the gas supply system just outside the vacuum chamber (see § 3.3).

Although simple, there were a few drawbacks with this design. The flow of gas out of the cell increased the ambient pressure within the vacuum chamber, requiring the vacuum pump to run continuously, causing unwanted vibrations. Over time, variation in the pointing of the pump beam caused the holes to enlarge. Moreover, the probe beam was only visible when passing through the holes made by the pump beam; signif-

icant realignment often required letting the chamber back up to atmospheric pressure and removing the gas cell.

3.6 New interaction region

After setting up the interaction region described in the previous section, I was able to use FDH to image the ionisation front. However, when I plotted the phase shift due to the background plasma as a function of gas pressure, the gradient was significantly less than expected and, instead of having a purely linear dependence on gas pressure, it appeared to saturate at around 2–3 radians (see § 5.1). This could be explained by the large phase shift caused by the background plasma causing the probe pulse to refract, limiting its interaction length. Increasing the pressure would cause it to refract sooner, producing the saturation effect seen. This meant that the original design, with its dependence on a long interaction length to make up for the small amplitude of the wakefield, was unsuitable. I therefore designed a new interaction region to maximise the measured signal while limiting the total phase shift due to the background plasma.

3.6.1 Design process

In order to design a new interaction region, I used the fluid equations described in § 4.5 to predict the amplitude of the wakefield generated by a Gaussian laser pulse, and the resulting phase shift. At this time, the laser energy was significantly higher than later in the experiment; I had also not yet accounted for the fraction of the energy present in a pre-/post-pulse. Therefore these simulations were performed for a laser energy of 60 mJ, rather than the 16.5 mJ of later simulations.

Using the measured beam parameters of the original interaction region, this simulation predicted that a 10 mm interaction length would result in a background phase shift of nearly 23 radians, but a peak-to-peak phase shift due to the wakefield of only 57 mrad. This background phase shift was far too large, causing the pump and probe pulses to suffer ionisation-induced defocussing (as observed in the simulation

described in § 4.4.1), so I needed to reduce it by reducing the interaction length to about 1 mm. This would also reduce the phase shift due to the wakefield, so the focal spot size must be decreased (increasing the wakefield amplitude) to compensate. Reducing the focal spot size reduces the Rayleigh length of the laser beam, limiting the interaction length.

In the 1D limit, where the focal spot size is much larger than the plasma wavelength, these two effects exactly cancel. However, a useful result of the 2D fluid equations is that the density oscillation due to the wakefield is significantly enhanced as the beam size is reduced below the plasma wavelength. There is no enhancement to the accelerating electric field, so this is not useful for acceleration experiments, but it is useful for an experiment which measures the electron density oscillation directly. This also reduces the optimal electron density [26], further reducing the background phase shift. It is therefore advantageous to make the focal spot size as small as possible.

However, if the spot size is made too small the wakefield will become non-linear; I decided to avoid this régime for a couple of reasons. Firstly, the fluid equations used for the simulation only hold for linear wakefields. Secondly, I expected that a non-linear wakefield would be harder to measure, due to the significant phase shift being concentrated into a very small portion of the phase map. The phase shift due to a linear wakefield will be more spread out, making it harder to miss and allowing more use of averaging.

I therefore calculated the focal spot size and peak electron density oscillation that would be produced by paraboloids of various focal lengths. Given that focal spot size is proportional to focal length (this is a simple result of the Gaussian beam model described in § 3.4 — $w_0 \propto f$ when $z = f \gg z_R$), I determined that a focal length of 150 mm would produce a spot size of $w_0 = 9 \mu\text{m}$. Using the equations of § 4.5, I predicted that this would result in a peak value of $\frac{\delta n_e}{n_e}$ of 0.55. Since similar calculations showed that a significantly shorter focal length would result in values of $\frac{\delta n_e}{n_e}$ greater than 1 and hence a non-linear wakefield, I selected 150 mm as the focal length of the

new interaction region. With a 1 mm interaction length, I predicted that this would result in a background phase shift of 1.2 radians and a peak-to-peak phase shift of over 700 mrad due to the wakefield.

However, in order to use the fluid equations, I had to model the laser as a circular Gaussian beam (the approximation described in § 4.5.2 to represent elliptical or astigmatic beams will not be accurate for small beam sizes, since the transverse gradient has a strong effect on the wakefield amplitude). However, the beam is significantly astigmatic. In order to account for this, I took the value of w where the beam was most circular (near the focus) as w_0 .

However, when this design was realised, several differences were observed. The lower energy of the final system was mentioned above; accounting for this reduced the peak-to-peak phase shift predicted by the fluid equations to about 200 mrad. More significantly, the new focal profile was not as astigmatic as I had expected, resulting in a smaller focal spot size than the design had called for. Although this enhanced the phase shift, it also pushed the wakefield towards the non-linear régime; in particular, the transverse size of the electron density peaks became very small, as discussed throughout Chapter 4, making correct alignment of the diagnostic very difficult.

3.6.2 Layout

The mirror used was an $f = 150$ mm, off-axis, gold-coated paraboloid. The angle between the incoming and outgoing beams was designed to be 90° , making the layout a lot simpler. The shorter focal length made this paraboloid a lot more sensitive to misalignment, which would introduce astigmatism or other aberrations. Therefore, the paraboloid was mounted on a manual 5-axis translation/rotation stage. By measuring the profile of the beam around the focus, I could adjust the paraboloid to minimise the observed astigmatism. This had to be done under atmospheric conditions, but no significant change was observed when the chamber was pumped down to vacuum.

The shorter focal length meant that the dichroic mirror could be placed much closer




FIGURE 3.22 – The updated layout of the interaction region. (IP) denotes the interaction point, where the gas cell is positioned.

to the gas cell (described later). Again, the probe beam was collimated by an achromatic lens, but in this layout a series of additional lenses were required to simultaneously reduce the diameter of the collimated probe beam (to fit on the later optics) and increase the magnification factor from $0.5\times$ in the previous layout to approximately $6\times$, in order to image the smaller interaction region.

3.6.3 Transmitted energy

The energy transmission measurement made for the original interaction region was repeated for the new layout. The result was a transmission of $29.6\% \pm 1.5\%$.

3.6.4 Pump focal profile

An additional part of the re-design of the interaction region was improved diagnostics. In particular, a camera was placed inside the vacuum chamber to image the pump focus directly, using a motorised flip mirror and a long working distance microscope objective. This system was calibrated using a 1951 USAF resolution test chart, giving a scale of $(0.548 \pm 0.004) \mu\text{m}/\text{px}$. A translation stage allowed the camera to sample different transverse planes throughout the focal region.

Again, beam profiles were recorded at different z positions, some of which are shown in Figure 3.23. There is no obvious astigmatism here, but the beam quality still appears quite poor. This suggests that realigning the paraboloid appropriately




FIGURE 3.23 – Profiles of the pump pulse near the focus. The top-right ($z = 0$) image is the focal spot and the longitudinal positions are measured relative to it.

Parameter	Value		Units
	x	y	
w_0	5.94 ± 0.15	5.87 ± 0.03	μm
z_0	109 ± 4	98.7 ± 0.4	μm
M^2	1.38 ± 0.05	2.04 ± 0.01	

TABLE 3.3 – The parameters for the fits shown in Figure 3.24.

was able to cancel out the inherent astigmatism of the laser, at the cost of introducing other aberrations.

These profiles were analysed using the method described in § 3.4.1. The principal axes appear to coincide with the axes of the camera. However, the 13.5% method produced very poor results when applied to this data, particularly in the case of the x -axis. I therefore used the 2σ method instead. The measured values of $w_{x,y}$ and the best fits using Equation 3.13 are shown in Figure 3.24. The fit parameters are shown in Table 3.3. Both fits appear good.

However, when I repeated the fluence analysis described in § 3.5.3, the fit between the experimental data and the prediction from the Gaussian beam model was poor,




FIGURE 3.24 – Beam sizes measured using 2σ method, and fit using M^2 model (see § 3.4.1).

Parameter	Value	Units
w_0	5.15 ± 0.05	μm
z_0	101 ± 2	μm
M^2	1.08 ± 0.04	

TABLE 3.4 – The parameters for the fit shown in Figure 3.26.

with the model clearly giving fluence values that were too small (see Figure 3.25). This suggests that the beam sizes used for the Gaussian beam model, estimated by the 2σ method, are too large.

Having no good method of extracting values of $w_{x,y}$ from the beam profiles, I decided to fit the Gaussian beam model to the fluence data directly, although this required the assumption that the beam was circular. The results of this fit are shown in Figure 3.26 and the fit parameters are shown in Table 3.4. The fit is much better, bearing in mind that the experimental data is not expected to fit the Gaussian beam model exactly.

As a sanity check I then looked back at the beam sizes measured with the 13.5% method, which were not of good enough quality to fit a model to, but had produced




FIGURE 3.25 – Comparison of peak fluence as a function of z between: (a) values measured using a CCD; and (b) values predicted using the Gaussian beam model with the parameters from Table 3.3.




FIGURE 3.26 – The results of fitting a circular Gaussian beam model to the recorded peak fluence data. The fit between the data and model is much better than in Figure 3.25.

good results in the case of the old IP. I found that they were within 30% of the values obtained from the above fluence fit. This shows that the fluence fit is not producing unreasonable beam sizes — 100% deviations would have been concerning. I therefore decided to use this fluence fit to model the beam.

Focal shift from air to vacuum

Since the gas cell is aligned to the laser focus manually, under atmosphere, it is important that the position of the focus does not move significantly when the chamber is pumped down to vacuum. To check this, the position of the focus was measured using the camera in the vacuum chamber while under atmosphere, then again after pumping it down to vacuum. The longitudinal position of the focus changed by $125\ \mu\text{m} \pm 25\ \mu\text{m}$, and the transverse position changed by less than $15\ \mu\text{m}$. Since these are both smaller than the relevant dimensions of the gas cell, this motion should not be a problem.

3.6.5 Pump-probe alignment

An FDH measurement requires the pump and probe beams to be aligned, that is both overlapped and co-linear. This was achieved in two stages. First, both beams were aligned to a pair of apertures positioned before the entrance to the vacuum chamber. This was generally not sufficient to ensure that the beams overlapped at the focus, so the second step used the camera positioned outside the vacuum chamber to detect the refraction of the probe beam that occurred when it passed through the plasma generated at the pump focus. This required about 100 mbar of air in the vacuum chamber, and the probe pulse to arrive after the pump pulse. Images from this process are shown in the bottom row of Figure 3.27. After the pump and probe beams were well overlapped, the same diagnostic could be used to correctly position the probe and reference pulses after and before the pump pulse respectively, as shown in the upper row of the same figure.

This alignment procedure ensures that the pump and probe beams are overlapped




FIGURE 3.27 – Images of the probe and reference pulses captured by the camera in the post-IP diagnostics area. **Top row:** The pump pulse ionises gas to plasma in the focus, which refracts the probe and reference pulses. This is used to set the delay stage so that the probe is behind the pump and the reference in front. **Bottom row:** The same refraction is used to ensure good spatial overlap between the probe and pump beams, with both probe and reference behind the pump pulse.

at the interaction point, but it does not ensure that they are perfectly co-linear. An angular offset between the pump and probe beams would result in the accumulated phase shift being spread transversely across the probe beam, reducing the signal. I checked this angular alignment using the camera inside the vacuum chamber and found that the angular offset between the pump and probe beams was (26 ± 7) mrad under atmosphere and (11 ± 9) mrad under vacuum. These are both consistent with a 20 mrad misalignment, although the actual value is expected to change from day to day. The effect of such a misalignment is discussed in § 4.7.

There did not seem to be any practical way to improve the alignment beyond this level, which I believe was limited by how accurately the probe beam could be aligned to the apertures (determined by the probe beam size and shape). Correcting the alignment based on the measurements described in this section might have been possible, but would have been very time-consuming for only a small gain. The best alignment using this method would have been limited by the uncertainty of these measurements — about 10 mrad; this was dominated by the small Rayleigh range of the pump, which caused its direction of propagation to be uncertain.

3.6.6 Gas cell

The new design of the IP called for a significantly shorter gas cell, to reduce the phase shift caused by the background plasma. The original design could not reasonably be adapted to the required length of 1 mm, principally because the attachment for the gas inlet would be larger than the gas cell. Instead, the newer gas cells consisted of 3.6 mm diameter nickel tubing with 50 μm thick walls. These were crushed to the required thickness of 1 mm, resulting in an approximately rectangular cross-section. The gas inlet was then attached to the circular end of the tube via a Swagelok connector, the other end being sealed with epoxy. The design is shown in Figure 3.28, and an example is shown in Figure 3.29. Despite the different method of manufacture, the gas cell works on the same operating principle of the original (see § 3.5.4), and suffers from

the same problems.

Since the gas is flowed through the cell continuously, there will be streams of gas leaving the holes in the front and back of the cell, through which the laser pulses travel. Although in the previous design these might be neglected, since they are much shorter than the total interaction length, here they might be expected to have a significant effect. In particular, the lower density within these plumes of gas means that any wakefield excited would have a longer period. However, calculations (based on the 2D fluid equations which will be introduced in § 4.5) showed that the laser pulse would only be intense enough to excite a significant wakefield within a few hundred μm of the focus; any wakefield excited within the plumes would have an amplitude about 100 times less than the amplitude at the laser focus.

Due to the method of manufacture, the thickness of the gas cells was not uniform – not only did different gas cells have different thicknesses, but a thickness of a single gas cell varied along its length. Although a better design of gas cell is a priority for a future experiment, there is no reason to believe that this variation had a significant effect on the results of this experiment. Slight variations in the thickness of the gas cell will affect the background phase shift, but not significantly affect the phase shift due to the wakefield, since most of that occurs within a few hundred μm of the focus. In any case, the jets of gas exiting the holes at the front and back of the gas cell are expected to be partially ionised, meaning that the interaction length is not strictly limited to the thickness of the gas cell itself.

3.7 Post-IP diagnostics

After the probe pulse interacts with the plasma and is separated from the remaining pump pulse, it must be imaged onto the entrance slit of an imaging spectrometer for analysis. The probe beam is collimated by the time it exits the vacuum chamber, so an image is formed using a single achromatic lens with $f = 300\text{ mm}$. This image is magnified further by a $4\times$ microscope objective mounted on the end of a lens tube,




FIGURE 3.28 – A diagram of the gas cell design used with the new IP. All dimensions are approximate.




FIGURE 3.29 – Two views of one of the gas cells used with the new IP. The top view is facing in the direction of laser propagation, and the hole drilled by the pulse laser can be seen towards the right.

which in turn is attached to the spectrometer. An insertable mirror, placed after the achromatic lens, allows the image to be formed on a CCD instead. This image of the probe beam at the interaction point can be used for alignment.

3.7.1 Spectrometer

The spectrometer used was an Andor Shamrock SR-163 imaging spectrometer, with an Andor Newton DU-940N CCD attached at the exit port to record the spectrum. The entrance slit was adjustable; data was collected at its narrowest setting, 10 μm . Calculations of the signal-to-noise ratio can be found in § 2.5.3.

The spectrometer follows a standard Czerny-Turner configuration [94]; light from the entrance slit, initially diverging, is collimated by a curved mirror and used to illuminate a diffraction grating. The diffracted light is focussed onto the exit plane by a second curved mirror, where it is recorded by a CCD. The curved mirrors form an image of the entrance slit on the CCD, but the position of this image (in the horizontal direction) depends on the wavelength of the light. Effectively there is one image for each wavelength, but the finite width of the entrance slit means that these images overlap. Hence the width of the entrance slit and the dispersion of the grating determine the resolution of the spectrometer. In order to achieve the best possible resolution, the most dispersive grating available (2400 lines/mm) was selected, resulting in a resolution of better than 0.05 nm.

The vertical scale on the spectrometer, that is the correspondence between row number on the CCD and vertical position at the interaction point, was calibrated using a 1951 USAF resolution test chart, giving a result of $(0.58 \pm 0.08) \mu\text{m}/\text{px}$.

3.7.2 Imaging system and focussing

The lenses between the interaction point and the slit of the spectrometer act as an imaging system. If the image plane is fixed at the entrance slit, then the object plane lies somewhere near the interaction point, and it can be moved by adjusting the po-

sition of one of the lenses. When the system is first constructed, the object plane will not coincide with the interaction point, so an additional alignment step is necessary.

The usual treatment of FDH (see Chapter 2) assumes that the probe beam is collimated as it passes through the interaction region, so that a single part of the probe pulse always experiences a phase shift due to the same part of the wakefield. It also neglects the effect of the accumulated phase shift on the propagation of the probe pulse (by contrast, FDS [73] relies on this effect). The collimation requirement is approximately satisfied in practice, because the much smaller size of the probe beam before being focussed ensures that the probe focus is larger than the pump focus, and has a correspondingly larger Rayleigh range. Provided the imaging system is correctly aligned, the propagation between the interaction region and the spectrometer can be ignored. However, if the imaging system is not aligned correctly, the image of the interaction point will be formed in front of or behind the entrance to the spectrometer. The propagation of the probe over the remaining distance will affect both its amplitude and phase, distorting the phase measured by FDH.

This alignment was performed by placing an object (in the form of a test chart) accurately in the plane of the pump focus, by positioning the camera in the chamber so that it was imaging the pump focus, then positioning the test chart so that its image on the camera was sharp. Then the probe beam was used to form an image of the test chart on the spectrometer, and the imaging lens outside the chamber was positioned so that the image was as sharp as possible.

3.8 Data acquisition and analysis software

Several pieces of experimental apparatus are computer-controlled, including the spectrometer, the beam profile monitors, and some of the translation stages (including the stage controlling the pump-probe delay). In most cases the devices are controlled by stand-alone software supplied by the manufacturer. A notable exception is the delay line translation stage, for which I wrote programs in both Python and LabView that

present a graphical user interface (GUI) to accept commands and control the device via a serial line.

I also wrote a program to perform FDH analysis, although it omits certain pre-processing of the FDH references for which manual intervention is required (to select which region of data to average over). For more detail see § 2.4.

During the course of the experiment, it became clear that it was not possible to interpret the FDH holograms by eye, and that some kind of real-time analysis would be needed to guide adjustments to parameters such as gas pressure and the delay stage setting. To this end, I created a program in LabView which:

- Accepts commands from the user through a GUI
- Commands the spectrometer to save data to a well-known location
- Performs FDH analysis on the saved data by running a separate program
- Displays the resulting phase map to the user
- Optionally, saves the raw data to disk for archiving and later analysis

The program presents the user with options to collect a frame from the spectrometer to use as a null hologram or to grab frames continuously and analyse them as data holograms. The whole process typically takes about 2 seconds per frame. The displayed phase map can be the spectral phase or, if suitable references are supplied, the temporal phase. The program also presents an option to automatically save each frame of data it records, allowing large volumes of experimental data to be recorded quickly and easily.

A separate data acquisition program which I created in LabView made it possible to quickly collect a number of spectra at different values of the pump-probe delay, by controlling the spectrometer and delay stage simultaneously. The delay stage was scanned in user-specified increments between user-specified limits, and at each point a user-specified number of spectra were recorded and saved, with the position of the

delay stage encoded into the file name. This idea could not be extended to scanning the pressure in the gas cell, because the gas supply system was not computer controlled.

4. Simulations

In this chapter, I present predictions of the wakefield signal that I expect to be measured by FDH, based on the laser parameters presented in Chapter 3. I consider three cases: the old interaction point (IP), described in § 3.5 and denoted by OIP; the new IP as originally designed, described in § 3.6 and denoted by NIPD; and the new IP as realised and measured, also described in § 3.6 and denoted by NIPM. For each case, I use both EPOCH, a particle-in-cell code described in more detail in § 4.1, and the simple analytical model presented in [26] to calculate the expected density oscillation. I then integrate this over time according to Equation 2.5 to calculate the phase shift that will be applied to the probe pulse. Due to limited availability of computing resources, I was not able to perform a full 3D EPOCH simulation over the entire interaction length. I therefore performed 2D EPOCH simulations over the full interaction length, which ranged from 1 mm for NIPM to 20 mm for OIP. I supplemented these with a 3D EPOCH simulation over a much more limited interaction length of 200 μm , which I performed only for NIPM (the actual setup of the experiment). This provided a better prediction of the wakefield amplitude, at the expense of not knowing the propagation length over which that amplitude was maintained.

Particle-in-cell simulations require considerable computing resources, to which I did not have access until the end of the experiment. Therefore the EPOCH simulations described in this chapter were not used during experimental planning; they would have been useful at this time, as they reveal several problems with the experimental design, in particular the presence of ionisation-induced defocussing in the old design

for the IP.

I will also describe a numerical model of the FDH measurement, which I used to generate a simulated FDH hologram, based on actual measurements of the probe spectrum and chirp. By analysing this I can obtain the expected FDH signal.

4.1 Particle-in-cell simulations

While many properties and behaviours of a plasma can be studied by treating it as a fluid, sometimes a more explicit treatment is required. One option is the particle-in-cell (PIC) simulation, which consists of two coupled simulations: one of the motion of the charged particles within the plasma, and one of the electromagnetic (EM) fields. The charged particles experience the Lorentz force due to the EM fields, whilst simultaneously acting as the source term in Maxwell's equations.

A typical PIC simulation includes a simple finite-difference time-domain electromagnetic solver, acting on a uniform Yee lattice, with cuboid grid cells. In the absence of a source term representing free charges and currents, this models the behaviour of an EM field in vacuum. In order to model a plasma, the simulation tracks the positions and velocities of charged particles of various species (e.g. electrons and ions). Since a typical simulation might include over 10^{12} individual electrons, the simulation does not model them individually; instead they are grouped into macro-particles, each representing a number of particles (the macro-particle's weight) with the same position and velocity. The motion of these particles under the Lorentz force is simulated with the same step size as the EM solver, so that the two solvers act alternately. The EM field acting on each macro-particle is computed by interpolation from nearby grid points, since the position of the macro-particle can take on any value. Likewise, the charge and current of the macro-particle must be deposited on nearby grid points so that the EM solver can use them as a source term.

These are the four essential parts of a PIC simulation: the EM solver; interpolation of EM fields to particles; the particle pusher, solving the relativistic equations of mo-




FIGURE 4.1 – The four steps in a particle-in-cell simulation. Moving once round the loop advances the simulation by a time Δt .

tion; and the deposition of charge and current onto grid points. These are repeated, in order, as many times as necessary to reach the desired simulation end time. This is depicted in Figure 4.1.

PIC simulations are particularly well-suited to laser-plasma interactions, as the laser can be modelled by an appropriate boundary condition on the EM field. The ponderomotive force, wakefields (including the non-linear and bubble régimes), and electron acceleration all emerge naturally from the simulation.

4.1.1 EPOCH

EPOCH (Extendable PIC Open Collaboration) is an open source relativistic PIC code, developed originally by Chris Brady at the University of Warwick, and then by Keith Bennett and other contributors from multiple institutions through the Collaborative Computational Plasma Physics (CCPP) group [95]. Like most modern PIC codes, it is capable of a high degree of parallelism, spreading operations across multiple CPU cores to reduce the total run-time. One particular feature of EPOCH that I took advantage of is its integrated ionisation model (the implementation of which is described in [95]). This allows the initial simulation to include species that are not completely ionised, and models their subsequent ionisation by the electromagnetic field. The code

models multi-photon ionisation, tunnelling ionisation (using the ADK model), barrier-suppression ionisation, and collisional ionisation.

This was important to my work, because the experiments being simulated did not use a pre-ionised plasma. Instead, the helium gas was ionised by passage of the pump pulse, and it was not clear how complete this ionisation would be or whether it would have any adverse effects. In particular, integrating the EM field solver with the ionisation model allowed me to observe the pump pulse being refracted away from the axis by its own plasma (see § 4.4.1), which would not have been possible otherwise.

4.2 Summary of laser parameters

Before performing any simulations, I need the relevant laser parameters. I will model the temporal profile as a Gaussian, for which I need to know the FWHM duration and the energy reaching the interaction point (not including any energy in pre- or post-pulses). The transverse profile will also be modelled as a Gaussian. While the fluid simulation will use the full M^2 model described in § 3.4, the EPOCH simulations are more limited. The 2D EPOCH simulations only allow the beam size to be specified in one dimension, so they require a circular beam model. The 3D EPOCH simulation would allow a non-circular beam model in principle, but the only good beam model I had available in the NIPM case was circular anyway.

Another limitation of the EPOCH simulations is that they require $M^2 = 1$. This is because the Gaussian beam model with $M^2 > 1$ is, strictly speaking, unphysical insofar as it is not a solution of Maxwell's equations. Since EPOCH propagates the laser pulse by solving Maxwell's equations, any Gaussian beam input will necessarily propagate with $M^2 = 1$. In practice, the M^2 model is used to represent an unspecified superposition of higher-order modes, so representing this in EPOCH would require finding and specifying a suitable superposition. I decided that this was unnecessary given the small measured values of M^2 — there would be little effect on the overall beam propagation, while the detailed shape of the focus would not necessarily match

Parameter	Value	Units	Source
Total energy entering compressor	98	mJ	§ 3.1
Transmission to IP	29	%	§ 3.5.2, § 3.6.3
Fraction of energy in central pulse	58	%	§ 3.1.2
Useful pulse energy reaching IP	16.5	mJ	
FWHM pulse duration	50.8	fs	§ 3.1.2

TABLE 4.1 – Summary of the laser parameters that are independent of the choice of paraboloid.

Parameter	Value		Units
	x	y	
w_0	69.4	61.9	μm
z_0	58.5	83.3	mm
M^2	1.0	1.2	
Effective w_0	81		μm

TABLE 4.2 – Summary of the beam parameters for OIP. See Table 3.2. The effective w_0 was derived by the best fit of a circular Gaussian beam to the peak fluence data, shown in Figure 4.2. The significant astigmatism makes it larger than the individual w_0 values.

experimental observations, since the M^2 value does not fully characterise the beam propagation in reality.

The laser energy and temporal duration are summarised in Table 4.1. The Gaussian beam parameters for OIP are summarised in Table 4.2. Modelling this as a circular Gaussian for the EPOCH simulation presents a problem, as the beam is significantly astigmatic. To solve this, I used the recorded fluence data, presented in Figure 3.20, and fit a curve to it in the manner described in § 3.6.4. The resulting fit is shown in Figure 4.2. The calculated value of w_0 , 81 μm , is similar to the value of w in the original data (see Figure 3.19) at the point where the beam is circular.

The Gaussian beam parameters for NIPD and NIPM are summarised in Table 4.3. These models are already circular, so they can be used in the EPOCH simulations directly.

4.3 Gas density

Having the laser parameters, I then proceeded to select gas densities at which to run the simulations. I first selected the electron density which would maximise the ampli-




FIGURE 4.2 – Comparison of peak fluence as a function of z between: (a) values measured using a CCD; and (b) values predicted using a circular Gaussian beam model, fit to this data ($M^2 = 1$, $w_0 = 81 \mu\text{m}$).

Parameter	Value		Units
	NIPD	NIPM	
w_0	9	5.15	μm
M^2	1	1.08	

TABLE 4.3 – Summary of the beam parameters for the new, shorter focal length paraboloid. See Table 3.4.

Parameters	Density / $\times 10^{17} \text{ cm}^{-3}$		He pressure / mbar
	e ⁻	He	
OIP	20.0	20.0	82.9
NIPD	10.4	5.2	21.6
NIPM	7.9	3.95	16.4

TABLE 4.4 – Summary of the optimal values for the density of electrons and helium atoms, and the pressure of helium gas.

tude¹ of the electron density oscillations given the laser parameters. In the case of a small-amplitude Gaussian pulse in the 1D limit (that is, where the beam width is much larger than the pulse length), this resonant density is given by [26]:

$$n_{res} = \frac{24 \ln 2 \epsilon_0 m_e}{q_e^2 \tau^2} \quad (4.1)$$

where ϵ_0 is the permittivity of free space, m_e and q_e are the mass and charge of the electron respectively, and τ is the FWHM pulse duration. Given the known laser parameters, this gives a resonant electron density of $n_{res} = 2.0 \times 10^{18} \text{ cm}^{-3}$.

In order to account for 2D effects (the beam width being of the same order as the pulse length), which reduce the resonant density [26], I used the analytical model described in more detail in § 4.5. Fixing the pulse parameters at their expected values at the focus, I varied the electron density to numerically maximise the predicted density oscillation.

Having values for the electron density, I then needed to calculate the corresponding values for the gas density, for both the experiment and the EPOCH simulations, which will include ionisation effects. For this purpose, I assumed that OIP would reach a sufficient intensity to ionise helium to the He^+ level, but NIPD and NIPM would be able to ionise it to the He^{2+} level. These assumptions are justified by the 2D EPOCH simulations shown in § 4.4, which include an ionisation model. This allowed me to calculate the required gas density and the corresponding pressure (using $p = nk_B T$ and assuming $T = 298 \text{ K}$). The results are summarised in Table 4.4.

¹The more usual aim of a laser-plasma experiment is to maximise the accelerating electric field, which leads to a different solution.

4.4 2D EPOCH simulations

Using these parameters, I constructed three 2D EPOCH simulations. Each consisted of a Gaussian laser pulse propagating through helium gas with uniform density (except for a short density up-ramp at the beginning, to avoid sudden transitions). The probe pulse was not included, due to concerns about its impact on the required computation time: at the very least, it would have required a reduction in the grid spacing to accommodate its shorter wavelength, which would have required a smaller step size (by the Courant condition). Furthermore, the real probe pulse was significantly larger than the pump pulse, so it would have required a wider simulation domain. Instead I calculated the expected phase shift based on measurements of the electron density. The simulations included the ionisation of the helium atoms due to the electric field of the laser pulse. A moving window was used to follow the propagation of the laser pulse over several millimetres, while using a smaller simulation domain for efficiency.

Each simulation used 10 cells per wavelength of the pump pulse longitudinally (80 nm per cell) and 5 helium macro-particles per cell (so 10 electron macro-particles per cell for the new IP, and 5 for the old IP). The transverse cell size varied between 125 nm and 500 nm, depending on the expected transverse size of features in the simulation and on available computing power. The laser pulse was configured so that it reached a focus of the correct size half-way through the simulation.

Since the simulations were only 2D, the pulse only focusses and defocusses in one dimension. Although the measured beam size behaves according to the Gaussian beam model, the peak intensity scales as $\frac{1}{w}$ instead of $\frac{1}{w^2}$ as it should. The peak intensity in the simulation could only match the experimental value at one point; I chose to make the value correct at the focus, so the intensity far from the focus is unrealistically high. This will result in the measured phase shift being overestimated somewhat.




FIGURE 4.3 – OIP: Snapshots of the laser pulse before entering the gas cell, after propagating 1 mm, and after propagating 3 mm. The pulse is propagating from left to right. The effect of refraction from the plasma is apparent after 1 mm and has completely distorted the shape of the laser pulse after 3 mm (by which point reflections from the walls of the simulation domain are causing an interference pattern to appear at the top and bottom edges).

4.4.1 Old IP (OIP)

This simulation was set to follow the laser propagation over 20 mm, which was the longest interaction length considered during the experimental design. However, the plasma gradient generated as the helium was ionised caused the pump pulse to refract rapidly, expanding significantly after only a few millimetres, as shown in Figure 4.3.

These results suggest that we would never have been able to observe a wakefield with the original design for the IP. As discussed in § 3.6 I had already anticipated that the probe pulse might be refracted, preventing the wakefield from being observed. These results suggest that the same happens to the pump pulse, allowing a wakefield to exist only over a short distance. This further justifies our decision to abandon the original design for the IP.




FIGURE 4.4 – NIPD: A snapshot of the electron density, taken when the laser is at the focus. Separate regions of ionisation to He^+ and He^{2+} are visible. A wakefield is just about visible behind the laser pulse.

4.4.2 New IP as designed (NIPD)

This simulation was set to follow the laser propagation over 2 mm, that being the distance over which I expected most of the phase shift to be accumulated (from preliminary results of the fluid simulation described in § 4.5). A snapshot of the electron density is shown in Figure 4.4.

In order to check whether the beam is propagating as expected (checking both the simulation and refraction from the plasma) I measured the beam size using the 13.5% method (see § 3.4.1) at each recorded snapshot. I then plotted the measured beam size as a function of z and compared it to the prediction from the Gaussian beam model using the parameters from Table 4.3. The results, for simulations both with and without particles, are shown in Figure 4.5. The presence of the plasma is seen to cause the beam to refract as it approaches the focus, resulting in a larger than expected waist




FIGURE 4.5 – NIPD: Measurements of the beam size from the simulation after different propagation times. Data from simulations with and without particles are compared to the Gaussian beam model used to create the simulation.

size.

Since I did not include a probe pulse in the simulation, due to concerns about its impact on the time required by the simulation, I estimated the phase shift using snapshots of the electron density. I took each snapshot, and converted it to a map of refractive index (for $\lambda = 405$ nm). I then used these snapshots to evaluate the integral:

$$\phi = \frac{2\pi}{\lambda} \int (\eta - 1) dz \quad (4.2)$$

over the interaction length (see § 2.1). This resulted in a map of the total accumulated phase, shown in Figure 4.6.

The clear transition between He^+ and He^{2+} , visible in Figure 4.4, is absent, due to averaging over different beam sizes. A wakefield is just visible. The phase shift along the $y = 0$ axis is shown in Figure 4.7. This shows the wakefield more clearly — its period is 110 fs and its peak-to-peak amplitude is 100 mrad.




FIGURE 4.6 – NIPD: A map of the estimated phase shift experienced by the probe, as a function of delay and transverse position.




FIGURE 4.7 – NIPD: A plot of the estimated phase shift experienced by the probe, along the propagation axis and as a function of delay.




FIGURE 4.8 – NIPM: A snapshot of the electron density, taken soon after the laser passed through the focus. Separate regions of ionisation to He^+ and He^{2+} are visible. A wakefield is clearly visible behind the laser pulse.

4.4.3 New IP as measured (NIPM)

Since the Rayleigh length in this simulation is considerably shorter than in the previous one, this simulation was set to follow the laser propagation only over 1 mm. A snapshot of the electron density is shown in Figure 4.8.

In contrast to the previous case, the wakefield is clearly visible. Its shape, with broad regions of low electron density and small regions of high electron density, shows that it is entering the non-linear régime. I repeated the beam size measurements described in the previous section. The results, shown in Figure 4.9, indicate that refraction caused by the plasma is less of a problem than in the previous simulation.

Repeating the analysis described in the previous section, I obtained a map of the expected phase shift of the probe pulse, shown in Figure 4.10. The central portion, containing the wakefield, is shown separately in Figure 4.11.




FIGURE 4.9 – NIPM: Measurements of the beam size from the simulation after different propagation times. Data from simulations with and without particles are compared to the Gaussian beam model used to create the simulation.




FIGURE 4.10 – NIPM: A map of the estimated phase shift experienced by the probe, as a function of delay and transverse position.




FIGURE 4.11 – NIPM: A map of the estimated phase shift experienced by the probe, as a function of delay and transverse position — detail of wakefield.




FIGURE 4.12 – NIPM: A plot of the estimated phase shift experienced by the probe along the propagation axis. The two data sets show data taken directly along the axis and an average taken from $y = -2.5 \mu\text{m}$ to $y = +2.5 \mu\text{m}$.

Although a wakefield is clearly visible in this data, the most significant feature covers a very small area. The effect of this can be seen in Figure 4.12, which shows the phase shift both along the $y = 0$ axis and averaged between $y = \pm 2.5 \mu\text{m}$. Measuring the phase shift along the axis gives a peak-to-peak amplitude of nearly 450 mrad, whereas averaging over only $5 \mu\text{m}$ gives a peak-to-peak amplitude of only 230 mrad. In both cases, the period of the wakefield is 125 fs.

4.5 Fluid simulation

In this section, I take a different approach to predicting the amplitude of the wakefield. The behaviour of the plasma is described by linearised non-relativistic fluid equations. These are solved to give an analytic equation for the density oscillation as a function

of the laser parameters:

$$\begin{aligned} \frac{\delta n_e}{n_e} = & \frac{1}{\pi^2 \sqrt{\pi}} \frac{E_p}{m_e c^2} \left(\frac{\lambda r_e}{\sigma_r^2} \right)^2 e^{-\pi r_e n_e \sigma_z^2} \\ & \times e^{-(r/\sigma_r)^2} \left\{ \pi r_e n_e \sigma_r^2 + \left[1 - \left(\frac{r}{\sigma_r} \right)^2 \right] \right\} \sin(k_p \zeta) \end{aligned} \quad (4.3)$$

This is presented as Equation (37) by Siders *et al.* [26]; however, their derivation is only applicable to pulses with circular polarisation, although this is not made clear in the text. Adapting the equations to the case of linear polarisation requires a simple change. The ponderomotive force, which drives the wakefield, is given by [6, p. 1233]:

$$\mathbf{F}_p = -\frac{1}{2} m_e c^2 \nabla a^2 \quad (4.4)$$

where $\mathbf{a}$ is the normalised vector potential, $\mathbf{a} = \frac{e\mathbf{A}}{m_e c}$. This is the instantaneous vector potential, so in the case of linear polarisation it has the form:

$$a^2 = a_L^2 \sin^2(\omega t - kz) \quad (4.5)$$

where a_L is a slowly-varying envelope. The wakefield is then generated by the time-averaged portion of $\mathbf{F}_p$:

$$\langle \mathbf{F}_p \rangle = -\frac{1}{4} m_e c^2 \nabla a_L^2 \quad (4.6)$$

However, the derivation used by Siders *et al.* does not include this time averaging step, leading to a result which is a factor of 2 too large for the case of linear polarisation, although it is correct for the case of circular polarisation. The version given above as Equation 4.3 is correct for this case.

Siders *et al.* parametrise laser pulses using different quantities than the ones dis-

cussed so far. These parameters are related according to:

$$r_e = \frac{1}{4\pi\epsilon_0} \frac{q_e}{m_e c^2} \quad (4.7)$$

$$\sigma_r = \frac{w}{\sqrt{2}} \quad (4.8)$$

$$\sigma_z = \frac{c\tau}{2\sqrt{\ln 2}} \quad (4.9)$$

where τ is the full-width at half-maximum pulse duration and r_e is the classical electron radius.

4.5.1 Different contributions to the wakefield

An important result of Equation 4.3 is that the amplitude of the density oscillations is significantly enhanced when the laser spot size is smaller than the plasma wavelength. Although this does not lead to an increase in the accelerating electric field, it is a great boon to any experiment measuring phase shift. Ignoring this effect, the increase in amplitude as the focal spot size is reduced is exactly offset by the decrease in the Rayleigh length, so that the total phase shift due to the wakefield is independent of the spot size. The 2D enhancement allows the measurable phase shift to be increased by using a sufficiently small focal spot size.

Electron oscillations along the x , y , and z axes each contribute to the wakefield; these contributions depend on the x , y , and z components of the ponderomotive force, and therefore on the gradient of the laser intensity in those directions. Therefore, while the z component of the wakefield depends only on the pulse intensity and duration, the x and y components also depend on the spot size. (When formulated in terms of energy, the z component depends on spot size, but the x and y components have a stronger dependency.)

Parameters	L / mm	ϕ_{pl} / rad	$2\Delta\phi_{pl}$ / mrad
OIP	10	22.8	15.8
NIPD	2	2.37	206
NIPM	1	0.90	482

TABLE 4.5 – Key results from the fluid simulation. L is the simulated interaction length (based on the Rayleigh length and experimental considerations), ϕ_{pl} is the phase shift due to the background plasma, and $2\Delta\phi_{pl}$ is the peak-to-peak phase shift due to the wakefield.

4.5.2 Finding the phase shift

In order to find the total phase shift due to a wakefield, samples are taken at a number of z positions over the interaction length. At each point, the Gaussian beam propagation formula (Equation 3.13) is used to find the σ_r parameter, then Equation 4.3 is used to find the amplitude of the electron density oscillation. Then δn_e is integrated to find the total phase shift due to the wakefield, according to Equation 2.5. At the same time, the total phase shift due to the background plasma is found.

Both sets of parameters for the new IP (see Table 4.3) assume a circular Gaussian beam, while the parameters for the old IP, shown in Table 4.2, treat the beam as elliptical and astigmatic. However, Equation 4.3 treats the beam as circular, and only allows for a single σ_r parameter. In this case I set:

$$\sigma_r = \frac{\sqrt{w_x w_y}}{\sqrt{2}} \quad (4.10)$$

on the basis that this preserves the peak intensity. Although this will not accurately reproduce the effects of transverse oscillations, these are not expected to be significant in the old IP, where $\sigma_r \gg \sigma_z$.

The results of these simulations are summarised in Table 4.5.

4.5.3 Comparison with EPOCH simulations

The phase shifts derived from fluid simulations on the new IP, both as designed and as measured, can be compared to the results of the 2D EPOCH simulations described in § 4.4. In the case of the new IP as designed, the EPOCH prediction was 100 mrad,

a factor of 2 smaller than the prediction from the fluid simulation. The new IP as measured will be addressed later.

This is a significant discrepancy; the wakefield is clearly in the linear régime, so the fluid equations should hold. I hypothesised that the main reason for the difference was because the fluid equations are based on cylindrical symmetry, whereas the 2D EPOCH simulations describe a plasma and laser pulse which are independent of y . Therefore, the fluid equations include electron oscillations in the x , y , and z directions, but the EPOCH simulation only includes oscillations in the x and z directions. This discrepancy was a major motivation for running a 3D EPOCH simulation, but first I tried to address it from the other side, by making the fluid predictions match those from the 2D EPOCH simulations.

Addressing dimensionality

To test this hypothesis, I isolated the three contributions to Equation 4.3 by repeating the first part of the derivation given in [26] for three cases:

$$a_L^2(\zeta, \mathbf{r}) = a_0^2 f(\zeta) e^{-(r/\sigma_r)^2} \quad (4.11a)$$

$$a_L^2(\zeta, \mathbf{r}) = a_0^2 f(\zeta) e^{-(x/\sigma_x)^2} \quad (4.11b)$$

$$a_L^2(\zeta, \mathbf{r}) = a_0^2 f(\zeta) \quad (4.11c)$$

Here a_L is the envelope of the normalised vector potential of a laser pulse with linear polarisation and $\zeta = z - ct$. Case (a) is the case investigated by Siders *et al.* and will include all electron oscillations. Case (b) matches the 2D EPOCH simulations (by removing y -dependence) and will include only oscillations in the x and z directions. Case (c) will include only oscillations in the z direction.

The starting point of the derivation is:

$$\frac{\delta n_e}{n_e} = \frac{1}{4k_p} \int_{\zeta}^{\infty} \sin[k_p(\zeta - \zeta')] \nabla^2 a_L^2(\zeta', \mathbf{r}) d\zeta' \quad (4.12)$$

which can be derived from e.g. Equation (10) of [6, p. 1233] by making appropriate variable substitutions; the conversion from the instantaneous vector potential a to its envelope a_L is accomplished by dividing by 2, accounting for the time-averaging of the ponderomotive force.

In all three cases given above, an equation of the following form can be derived:

$$\frac{\delta n_e}{n_e} = \frac{a_L^2(\zeta, \mathbf{r})}{4} + \frac{k_p A}{4} \int_{\zeta}^{\infty} \sin[k_p(\zeta - \zeta')] a_L^2(\zeta', \mathbf{r}) d\zeta' \quad (4.13)$$

where the form of A depends on the particular case.

In case (a):

$$A = 1 + \frac{4}{(k_p \sigma_r)^2} \left[1 - \left(\frac{r}{\sigma_r} \right)^2 \right] \quad (4.14a)$$

In case (b):

$$A = 1 + \frac{2}{(k_p \sigma_x)^2} \left[1 - 2 \left(\frac{x}{\sigma_x} \right)^2 \right] \quad (4.14b)$$

And in case (c):

$$A = 1 \quad (4.14c)$$

By comparing these results, I modified Equation 4.3 to remove the effect of oscillations in the y direction. For $r = 0$ this gives:

$$\frac{\delta n_e}{n_e} = \frac{1}{\pi^2 \sqrt{\pi}} \frac{E_p}{m_e c^2} \left(\frac{\lambda r_e}{\sigma_r^2} \right)^2 e^{-\pi r_e n_e \sigma_z^2} \left[\pi r_e n_e \sigma_r^2 + \frac{1}{2} \right] \sin(k_p \zeta) \quad (4.15)$$

This modified equation gives much smaller phase shifts, as shown in Table 4.6. After adjusting the focal spot size to 11 μm to account for plasma refraction (as seen in Figure 4.5), the result is very close to that obtained by EPOCH.

It should be noted that these results are expected to be less accurate than those in Table 4.5; they are included only as evidence that the reason for the discrepancy with the 2D EPOCH simulations is due to the different dimensionality of the two models.

Parameters	L / mm	ϕ_{pl} / rad	$2\Delta\phi_{pl}$ / mrad
NIPD	2	2.37	142
NIPD + refraction	2	2.37	112
NIPM	1	0.90	276

TABLE 4.6 – Key results from the fluid simulation, adapted to ignore oscillations in the y direction. See Table 4.5 for headers. (+refraction) in the second row refers to increasing w_0 in the new IP to $11\text{ }\mu\text{m}$ to account for refraction by the plasma.

New IP as measured

Comparison of the fluid and EPOCH results for the new IP as measured (NIPM) is more difficult, because in this case the wakefield is non-linear; in particular, the fluid equations predict a peak value of $\frac{\delta n_e}{n_e}$ in excess of 1.5. In the EPOCH simulations, this non-linearity can be seen as very sharp density peaks, in contrast to broad and smooth density troughs. Therefore, one approach to comparing the EPOCH and fluid simulations would be to take the transversely averaged phase shift from the EPOCH simulations, which should be less strongly affected by the non-linear wakefield (since it averages over the sharp density peak). This results in a peak-to-peak phase shift of 230 mrad. Applying the same transverse averaging to the modified 2D fluid simulation results in a peak-to-peak phase shift of 210 mrad, in fairly close agreement.

4.6 3D EPOCH simulations

The 2D EPOCH simulations and the fluid simulations disagree because they are modelling slightly different physical processes: the fluid simulations are modelling the somewhat realistic case of a laser pulse with cylindrical symmetry, whereas the 2D EPOCH simulations are modelling the case of a laser pulse which is unvarying along one of the transverse dimensions, and infinite in extent along the same dimension. In § 4.5.3 I adjusted the fluid equations to model this latter case; the agreement I obtained between the fluid equations and EPOCH show that this difference in dimensionality is the main cause of the disagreement.

This is an unsatisfactory solution however, since it involves taking a physically

reasonable model and making it physically unreasonable. Therefore, I also ran a 3D EPOCH simulation of the new IP as measured (NIPM), and compared the results with the original fluid simulation.

Since 3D particle-in-cell simulations are much more computationally intensive than their 2D counterparts, and due to the limited computing power at my disposal, I reduced the scope of the 3D simulation to concentrate on the region around the focus itself; the laser pulse was allowed to propagate only over 200 μm (compared with 1 mm in the 2D simulation), focussing in the middle of the domain. By cutting out the regions where the beam size was large, I could significantly reduce the transverse size of the simulation (from $\pm 60 \mu\text{m}$ to $\pm 14 \mu\text{m}$) without affecting its accuracy. I also removed the ionisation model, starting with a fully ionised plasma. Finally, I reduced the number of electron macro-particles per cell from 10 to 4 — each macro-particle represented 250 electrons, with a charge of 40 aC. The cell size remained the same, at 80 nm longitudinally by 125 nm transversely. The final simulation took just over 2500 CPU-hours to run across 64 cores, taking 39 hours of wall-clock time.

4.6.1 Electron density snapshots

I took a 3D snapshot of the electron number density at a simulation time of 720 fs, after the laser had propagated some distance through the plasma and about three periods of the wakefield were visible. Here, I present various 2D slices through this volume. The co-ordinate system I use to describe these slices has the laser pulse propagating in the $+z$ direction, with its electric field being polarised along the y axis.

A slice in the x - z plane, at $y = 0$, is shown in Figure 4.13. The peak electron density is nearly $8 \times 10^{18} \text{ cm}^{-3}$, significantly greater than the peak electron density of approximately $2.8 \times 10^{18} \text{ cm}^{-3}$ obtained from the equivalent 2D simulation (see Figure 4.8). Various x - y slices at different z positions are shown in Figure 4.14. The slices further behind the laser pulse (at smaller z values) show a clear asymmetry. This is connected to the polarisation direction of the laser, with the major axis of the ellipse




FIGURE 4.13 – An x - z slice through the electron density at $t = 720$ fs. The laser pulse is propagating from left to right. The vertical dashed lines indicate where the x - y slices in Figure 4.14 were taken.

being aligned with the electric field of the laser pulse. Since the asymmetry does not occur until after the first density peak, it may be a result of the tail of the laser pulse interacting with the electrons within that density peak. If this is the case, it may be a simulation artefact: not only is the density peak small compared with the transverse cell size, but the laser field here is unrealistically intense due to light being reflected from the walls of the simulation domain².

4.6.2 Phase shift

Due to the short duration of this simulation, my previous method of predicting the phase shift by integrating the electron density will not work; in particular, although the electron density at the focus is known, the length over which its value will be maintained is not. Increasing the propagation distance to match the 2D simulations

²Since the side walls of the simulation domain are not perfectly absorbing, it acts as a rectangular waveguide; the incident laser pulse couples slightly to the higher-order modes of this waveguide, which propagate more slowly than the main laser pulse and so are able to overlap with the first density peak of the wakefield.




FIGURE 4.14 – x - y slices through the electron density at the z positions marked with dashed lines in Figure 4.13.

(which would require that the transverse size increase also, to house the diverging laser beam) would increase the time required by approximately $80\times$, requiring 4 months of computation across 64 CPUs. However, it is clear that the most significant phase shift will only be measured very close to the axis of propagation, due to the small size ($< 1 \mu\text{m}$) of the electron density peaks. Any slight misalignment in the experimental setup will prevent the full phase shift from being measured. In the best case, this will result in some amount of transverse averaging. In the worst case, the portion of the probe carrying the phase shift will miss the slit of the spectrometer entirely, and no signal will be measured.

Furthermore, the density peaks are very small in the longitudinal direction — about $5 \mu\text{m}$. This results in the peaks in the measured phase shift as a function of time having a very short duration ($\sim 15 \text{ fs}$). This is less than the expected temporal resolution of the diagnostic, further reducing the measured phase shift. In summary, although this 3D simulation shows much higher peak values of electron density than the equivalent 2D simulation, they are unlikely to significantly increase the measured phase shift.

There are thus various ways to predict the peak-to-peak phase shift:

- 482 mrad, from the fluid simulation (§ 4.5.2), but this is not expected to be applicable to the non-linear régime;
- 320 mrad, calculated by re-scaling the above value so that $\frac{\delta n_e}{n_e} = 1$;
- 450 mrad, from the 2D EPOCH simulation on axis (see § 4.4.3), but this is expected to be difficult to measure directly and in any case is an underestimate of the true value;
- 230 mrad, from the 2D EPOCH simulation averaged transversely, although again this may be an underestimate of the true value, since the 3D simulation results in greater variations in electron density.

In summary, given that measuring a true on-axis phase shift appears impractical, realistic values for the peak-to-peak phase shift are likely to be around 300 mrad.

4.7 Alignment sensitivity

An important practical concern when carrying out the experiment was the sensitivity of the measured phase shift to relative misalignments of the pump and probe beams. There can be angular misalignments, where the probe and pump beams do not follow parallel paths. This would cause the accumulated phase shift to be “smeared” transversely over the probe beam, reducing its peak value. There can also be translational misalignments: in practical terms, these occur when the region of the probe beam with maximum phase shift is not centred on the slit of the spectrometer.

To estimate the sensitivity of the phase shift to misalignments, I adapted the fluid simulations described in § 4.5 to account for a non-zero and varying value of r . I did not attempt to use the EPOCH simulations for this, as I have demonstrated that the 2D simulations underestimate the density oscillations (see § 4.5.3) and the 3D simulation only covered a short interaction length.




FIGURE 4.15 – The phase shift predicted by a fluid simulation, as a function of the radial co-ordinate r .

All simulations were performed for the new IP as measured (NIPM).

4.7.1 Translational misalignment

To investigate the effect of a translational misalignment, I used the fluid simulation to calculate the phase shift for various values of the radial co-ordinate r . The results are shown in Figure 4.15.

The phase shift is very sensitive to translational misalignment, with offsets as small as $3\ \mu\text{m}$ reducing the phase shift by over 50%. This value is the distance separating the peak of the wakefield and the vertical/longitudinal plane in which the phase shift is measured. The horizontal position of this plane is controlled by the transverse position of a $4\times$ microscope objective placed in front of the spectrometer; this objective takes a $6\times$ (approx) magnified image of the IP and then images that onto the entrance slit of the spectrometer. Ray diagrams can be used to show that moving the objective by $\frac{4}{5}\ \mu\text{m}$ changes the position of the measurement plane *in the $6\times$ magnified image* by

1 μm , and thus the position of the measurement plane at the IP by $\frac{1}{6}$ μm .

Therefore, 9.6 μm of transverse motion of the microscope objective corresponds to 2 μm of transverse offset at the IP; Figure 4.15 shows that at this point, the phase shift will be reduced by 33%. In principle, 5 μm alignment precision should be possible, giving a transverse offset of about 1 μm and a phase shift reduction of about 10%.

Since the image of the IP on the spectrometer slit is expected to be magnified by a total factor of $24\times$, the slit width of 10 μm should not cause a problem, since it corresponds to transverse offsets of less than 0.5 μm , at which the reduction in phase shift is negligible.

4.7.2 Angular misalignment

To investigate the effect of an angular misalignment, I used the fluid simulation to calculate the phase shift given $r = \theta z$ for various values of the angular misalignment θ . This represents the case of no translational misalignment, where the probe and pump intersect exactly at the focus ($z = 0$). The results are shown in Figure 4.16.

In practical terms, the phase shift is less sensitive to angular misalignment than to translational misalignment. Typical misalignments of 10–30 mrad (see § 3.6.5) cause phase shift reductions of 5–27%.

4.7.3 Alignment summary

Angular misalignment seems unlikely to be a severe problem by itself, given the measured alignment between the pump and the probe. While the positioning accuracy of the microscope objective used to align into the spectrometer seems sufficient to eliminate any significant problems from translational misalignment, there is no indication of whether it is correctly aligned; the background phase shift covers a much larger volume, and its peak cannot be located with enough precision. Instead, the alignment must be adjusted while looking for a wakefield, which would be completely impractical without real-time analysis.




FIGURE 4.16 – The phase shift predicted by a fluid simulation, as a function of the angular misalignment θ .

Given a positioning error of $5\ \mu\text{m}$ at the spectrometer ($1\ \mu\text{m}$ at the focus) leading to a 10% reduction in the measured phase shift, and an angular error of 20 mrad leading to a 12% reduction, a total reduction in measured phase shift of at least 20% seems inevitable. Slightly greater errors of $10\ \mu\text{m}$ and 30 mrad, which are certainly possible, would lead to a total reduction in phase shift of at least 50%. This could have a major effect on the ability to detect a wakefield.

Furthermore, the fluid model is only expected to hold in the linear régime. The 3D EPOCH simulation in § 4.6 shows that the wakefield is expected to be significantly non-linear. This results in a significant decrease in the transverse extent of the electron density peaks, and hence even stricter alignment requirements.

4.8 Frequency-domain holography simulation

In order to check whether my implementation of the FDH analysis process (as described in § 2.4) worked correctly, and also to investigate whether the probe spectrum

had sufficient bandwidth to support the plasma frequency, I constructed a numerical simulation of the FDH process. It takes the reference spectrum as input, which can be obtained experimentally or assumed to be a Gaussian. This is used to construct the electric field of the unchirped reference pulse (by assuming flat spectral phase), to which known amounts of group-delay dispersion can be added to give the electric fields of the chirped probe and reference pulses. A fixed delay is added to one of the pulses, and two spectral interferograms are computed and output: the interference between the unchirped reference and chirped probe pulses, which is used to measure the chirp; and the interference between the chirped reference and probe pulses, which acts as the null hologram.

Then, a phase shift is added to the probe pulse of the form:

$$\phi(t) = \phi_0 \sin(\omega_p t) \quad (4.16)$$

representing the phase shift due to a wakefield. Then the data hologram is computed by interfering the probe and reference pulses. This gives all the information needed to perform an FDH analysis, which is carried out by the same code used to analyse the experimental data. The calculated temporal phase shift can be plotted and compared with the known phase shift shown in Equation 4.16.

4.8.1 Simulation parameters

As well as a measured reference spectrum, the required input parameters are:

- the group-delay dispersion applied to the reference pulse;
- the group-delay dispersion applied to the probe pulse;
- the group delay used to separate the reference and probe pulses;
- the peak-to-peak amplitude of the applied phase shift $2\phi_0$; and
- the period of the applied phase shift $T_p = \frac{2\pi}{\omega_p}$.

Values for the first three were selected based on measurements of the experimental setup, while reasonable values for the other two were selected based on the simulations presented earlier in this chapter.

Based on the measurements shown in Figure 3.15 and Figure 3.14 I selected a value of $1.8 \times 10^4 \text{ fs}^2$ for the reference GDD and $2.3 \times 10^4 \text{ fs}^2$ for the probe GDD. The probe-reference separation can be found by looking at the Fourier transform of any FDH hologram: the data shown in Figure 2.13 suggests a value of 5 ps.

The earlier simulations did not result in a conclusive value for the expected phase shift; I took a peak-to-peak amplitude of 300 mrad to be representative, and a plasma period of 125 fs based on the results of the simulation described in § 4.4.3.

4.8.2 Simulation with Gaussian spectrum

First, I ran a simulation using a Gaussian input spectrum, with a FWHM bandwidth of 5 nm. The Fourier transforms of the three interferograms are shown overlaid in Figure 4.17, showing several interesting features.

The data from the null hologram is simplest: a narrow base-band is accompanied by two side-bands, representing the fringe pattern, which appear at $t = \pm 5 \text{ ps}$ as expected. The data hologram, which includes the effect of a periodic phase shift, is more interesting: the periodic modulation generates an additional series of side-bands, as described in § 2.8.2. These appear around the base-band, but also around each of the fringe-pattern side-bands. In this case, the periodic phase shift can be detected (and in principle, its period and amplitude calculated) by observing these side-bands, without any need for the full FDH analysis³. Finally, the chirp measurement results in much broader side-bands: this is due to the large spread in group delay that arises when the probe pulse is chirped and the reference pulse is not.

The full FDH analysis resulted in the temporal phase shown in Figure 4.18. The period matches the expected value of 125 fs, with 16 whole periods visible in the

³This method, dubbed TESS, is the subject of recent work by Nicholas Matlis, which he is preparing for publication[96].




FIGURE 4.17 – Fourier transforms of the three simulated interferograms. The transforms of the simulated chirp measurement, null hologram, and FDH hologram proper are shown vertically offset on a logarithmic scale.




FIGURE 4.18 – The temporal phase computed by performing the full FDH analysis on the simulated data.

2 ps displayed region. The measured peak-to-peak amplitude (varying between 265–300 mrad) is generally slightly less than the 300 mrad input, which may be due to the limited bandwidth of the reference pulse, or may be an issue with the sample spacing.

4.8.3 Simulation with experimental spectrum

I then ran another simulation, using an experimental spectrum instead of an ideal Gaussian; this spectrum was taken from a measurement of the spectrum of the reference pulse. The Fourier transforms of the three interferograms are shown overlaid in Figure 4.19. However, many of the features which are visible in Figure 4.17 are obscured here by the greater width of the Fourier transforms, which is caused by the presence of smaller-scale features in the real spectrum. In particular, the smaller sidebands which appeared in the previous data set, and were a clear sign of the presence of a periodic phase shift, are no longer visible. Clearly, a full FDH analysis is required to recover a phase shift from this data.

The full FDH analysis resulted in the temporal phase shown in Figure 4.20. Again,




FIGURE 4.19 – Fourier transforms of the three simulated interferograms. The transforms of the simulated chirp measurement, null hologram, and FDH hologram proper are shown vertically offset on a logarithmic scale.




FIGURE 4.20 – The temporal phase computed by performing the full FDH analysis on the simulated data.

the period is correct at 125 fs, with 8 periods visible in the central 1 ps of the graph (the very edges of the graph are less clear). The peak-to-peak amplitude is much less uniform than the data presented in Figure 4.18, varying from 160 mrad to 310 mrad in places. Typical values of the amplitude fall in the range 200–250 mrad, 7–25% less than the previous data, presumably due to the smaller spectral bandwidth limiting the temporal resolution.

4.9 Summary

The simulations presented in this chapter suggest that a correctly aligned diagnostic will be able to measure a peak-to-peak phase shift of at least 200 mrad, which is greater than the measured noise level of 80 mrad presented in § 2.5.3. In principle, it should be possible to detect a wakefield. However, the results of § 4.7.1 suggest that this correct alignment will be difficult to achieve, with the microscope objective in front of the spectrometer needing to be correctly placed to better than $\pm 10 \mu\text{m}$.

From a purely mechanical point of view, it should be possible to achieve this precision. However, the problem is the lack of a suitable reference to align against — there is no use in achieving 10 μm precision if there is no indication of where the diagnostics ought to be positioned. Since the wakefield is expected to be much smaller transversely than the ionised volume, the only way to align the diagnostics to the wakefield is to look for the wakefield itself.

Scanning for the wakefield might be possible, given the presence of a real-time data analysis tool, but several factors make this difficult. These are the small signal-to-noise ratio, the small vertical made more difficult by the relatively small signal-to-noise ratio, size of the expected signal region on the spectrometer (about 10 pixels, since each pixel covers 0.5 μm), and the shot-to-shot energy and angular jitter of the laser. As a result, although the only way to align the system is by observing the wakefield, it is doubtful whether it is possible to observe the wakefield before the system is properly aligned. As such, it may not be possible to align the system well enough to observe a significant phase shift.

5. Results

In this chapter, I present a selection of measurements made using FDH on the experiment described in Chapter 3. Early measurements focus on the ionisation front (see Figure 2.2 on page 29) because it is much easier to detect than the predicted wakefield. If the interaction length is known, the size of this phase shift is sufficient to calculate the plasma density. Alternatively, by assuming a value for the plasma density (based on an assumed ionisation state and the temperature and pressure of the input gas) an effective interaction length can be calculated. In the case of the original design of the interaction region, it will be shown that this results in an unrealistically small value for the interaction length. This justifies the move to a new design, whose measured interaction length is in much closer agreement with the gas cell dimensions.

Although no definitive wakefield signal was observed, various signals were detected which might be caused by a wakefield. A selection of these are presented, together with a discussion of whether they are consistent with being generated by a wakefield.

5.1 Old interaction region

As described in § 3.5, the original design of the experiment used a long focal length paraboloid to focus the pump and probe pulses into a 10 mm long gas cell. The larger size of the pump beam at the paraboloid ensures that the pump focus is significantly smaller than the probe beam, while the relatively long Rayleigh length ensures that the pump pulse maintains its small size over the entire interaction length. The result




FIGURE 5.1 – A snapshot of the temporal phase shift experienced by the probe pulse in the old IP, as a function of time and vertical position. The pump pulse is positioned approximately at $t = 0$, centred vertically. Behind the pump ($t > 0$) a phase shift of about 1.2 radians is seen, due to gas ionised by the pump pulse. Outside this region there is no significant phase shift, although vertical stripes can be seen. This anomalous effect is caused by some kind of systematic error in the data.

is that the FDH phase maps feature a distinct narrow column of phase shift due to the plasma, as seen in Figure 5.1. No pump propagates outside this region, hence no phase shift is observed.

The phase shift within this column is the ionisation phase shift — that is the phase shift experienced by the probe due to propagating through a (uniform) plasma rather than unionised helium gas. This phase shift caused by the plasma is given by Equation 2.5, reproduced here:

$$\Delta\phi \approx -\frac{\pi}{\lambda n_c} \int n_e dz \quad (5.1)$$

This phase shift is proportional to the cross-sectional plasma density $\int n_e dz$. Ideally, the electron density n_e is everywhere proportional to the input gas pressure p_{in} , and




FIGURE 5.2 – Phase shift due to ionisation as a function of pressure, in the old interaction region.

we can define an effective interaction length L^* such that:

$$n_0 L^* = q \frac{p_{in}}{k_B T} L^* = \int n_e dz \quad (5.2)$$

Where n_0 is the predicted electron density, q is the assumed ionisation level (+1 or +2) of the helium ions, and T is the (assumed, calculated, or measured) temperature of the helium gas inside the cell. By comparing the effective interaction length against the dimensions of the gas cell, various problems might be identified, either in the diagnostics or in the design of the interaction region.

By making a number of FDH measurements while varying the pressure of the helium gas supplied to the cell, the phase shift can be measured as a function of pressure. This should be a straight line passing through the origin, its slope indicating the effective interaction length L^* . Such a plot, produced from experimental data, is shown in Figure 5.2.

The slope of a plot of phase shift versus pressure is related to the effective interac-

tion length according to:

$$\frac{d\phi}{dp} = \frac{\pi}{\lambda n_c} \frac{qL^*}{k_B T} \quad (5.3)$$

Assuming room temperature and an effective interaction length of 10 mm (the actual length of the gas cell; see § 3.5.4) gives an expected gradient of 300 mrad/mbar, a factor of 30 greater than the slope seen in Figure 5.2. This clearly shows that the pump and probe are not interacting over the full length of the gas cell, and was the original motivation for the re-design of the IP to use a shorter interaction length, as described in § 3.6. This observation was later corroborated by particle-in-cell simulations of the original IP, described in § 4.4.1, in which the pump pulse was quickly refracted away from the axis by ionisation-induced defocussing.

Apart from the shallow slope, several other features of Figure 5.2 stand out. Firstly, the phase shift approaches a non-zero value as pressure approaches zero. This may be due to ionisation of ambient background gas over a much greater interaction length. The decrease in the slope of the graph above 150 mbar may be caused by a decrease in the effective interaction length, as the ionisation-induced defocussing becomes stronger. Finally, there is a significant spread in measured phase shift at each pressure. This may be caused by shot-to-shot energy fluctuations of the laser, which were as large as 30% peak-to-peak, affecting the onset of ionisation.

5.2 New interaction region

The new interaction region, with its much shorter interaction length, effectively prevented the refraction of the pump and probe beams (as will be seen later). However, even at 1 mm, the gas cell was several Rayleigh lengths long. This resulted in a relatively larger volume of ionisation, as the pump pulse was still intense enough to fully ionise the helium even while diverging. Since the FDH measurement effectively integrates the plasma density along the propagation axis, the region of ionisation is much less clearly defined than in Figure 5.1. As shown in Figure 5.3, the phase shift drops off




FIGURE 5.3 – A snapshot of the temporal phase shift experienced by the probe pulse in the new IP, as a function of time and vertical position. This was previously shown as Figure 2.17. The pump pulse is positioned approximately at $t = 0.1$ ps. Behind the pump pulse (larger t) phase shift is seen due to the pump ionising the gas. The distance over which this occurs varies vertically, as the size of the pump pulse changes over the interaction length. The vertical stripes in front of the pump pulse ($t < 0$) appear to be a systematic error present in some of the data, as discussed previously in § 2.6.5.

gradually transversely, as the longitudinal distance over which the pump is sufficiently intense to ionise the gas decreases.

Another distinctive feature in Figure 5.3 is the presence of vertical stripes in the region prior to the ionisation front (on the left). These may be caused by the presence of the ionisation front in the data, as described in § 2.6.2. As suggested in that section, it might be possible to mitigate this effect by moving the probe pulse so that it is no longer imaging the ionisation front. However, the observed effect does not appear to be present after the ionisation front, where the wakefield is expected, and simulations suggest the effect should be smaller than the expected wakefield signal. Since the non-linear nature of the wakefield might cause it to become distorted many periods behind the pump, it was considered better to keep the ionisation front in view, so that the first few periods of any wakefield would be visible.




FIGURE 5.4 – Phase shift due to ionisation as a function of pressure, in the new interaction region. A linear fit to the intermediate pressure region is shown. The inset graph shows detail near the origin, especially the non-zero phase shift near zero pressure.

Similarly to the original interaction region, the ionisation phase shift can be measured as a function of pressure, giving results shown in Figure 5.4. This graph shares several features in common with Figure 5.2 – a decrease in slope above 150 mbar and a non-zero intercept. A linear fit to the intermediate pressure region of the graph is shown. The slope of this fit gives an effective interaction length of 1.07 mm, which is in much closer agreement with the actual thickness of the gas cell – approximately 1 mm (see § 3.6.6). Interestingly, the intercept of the fit is negative, suggesting a lower pressure inside the gas cell than the pressure at the gas inlet. Nevertheless, this is much closer to the ideal straight line case, and justifies the move to the new interaction region.

5.2.1 Timing scan

In order to check the accuracy of the time scale on the reconstructed temporal phase maps, I performed a number of timing scans, in which the temporal separation be-




FIGURE 5.5 – Results of a timing scan. Timing of the ionisation front, measured using FDH, is plotted against the pump-probe delay inferred from the position of the translation stage in the probe delay line. A linear best fit is shown, with gradient 1.14 ± 0.03 . A fit with gradient fixed to 1 is also shown.

tween pump and probe pulses was varied while the position of the ionisation front was recorded. The pump-probe separation was controlled by varying the position of the translation stage in the probe delay line — the position of the stage could be used to calculate the relative pump-probe delay, although this value was subject to an arbitrary zero-offset. Meanwhile, the temporal phase maps resulting from FDH reconstruction were analysed to find the time at which the ionisation front occurred.

The results of this scan are plotted in Figure 5.5. Theoretically, the points should lie on a line with gradient 1. A linear fit, computed using the least-squares method, is shown; its gradient is 1.14 with a statistical error from the fitting procedure of 0.03. This appears incompatible with the hypothesis that the gradient should be 1, suggesting that the measured chirp is incorrect. On the other hand, a fit with the gradient constrained to 1 (also shown in Figure 5.5) does not appear significantly worse than the best fit.

It is possible that this error is due to residual chirp from the probe compressor,

which would affect the relation between time and frequency but would not appear in the chirp measurements (which only measure the chirp introduced by the glass). It is also possible that residual chirp or errors in the reconstruction cause the relation between true and measured time to be non-linear, in which case the gradient of a linear fit will necessarily be incorrect.

The significant spread in the measured times — about ± 200 fs — is discouraging. It may be due to jitter in the pump-probe separation, caused by variation in path length on the order of 50 μm ; given the size of the experiment, the use of separate compressors, and the lack of vibration isolation for many of the optics, this would not be unexpected. Another contributor might be variation in the total phase shift, caused by pump energy variations, which affect the process used to locate the ionisation front. However, the fact that the points fall close to a line with gradient equal to 1 is encouraging, suggesting that the time axis of the reconstructed data is approximately correct.

5.2.2 Transverse phase map

Since the spectrometer used is an imaging spectrometer, with a narrow vertical slit at the entrance, each shot records a cross-section through the ionisation volume, the plane of this cross-section being parallel to the vertical and longitudinal (y and z respectively) axes. By taking multiple sets of data while scanning the horizontal (x) position of this plane, it is possible to build up a complete 3D map of the ionisation front. In practice, the x position being imaged was changed by adjusting the transverse position of the microscope objective placed before the spectrometer (see § 3.7). Extracting the phase shift due to ionisation from this data and converting that to an effective interaction length L^* , I produced the plots shown in Figure 5.6. For the right-hand data, I adjusted the position of the probe pulse after each movement of the microscope objective, to improve the probe intensity and thus increase the horizontal range over which data could be collected.




FIGURE 5.6 – Effective interaction length (as defined by Equation 5.2) as a function of transverse position. Black regions are those without reliable data, due to insufficient probe intensity. **Left:** data taken with fixed probe position. **Right:** data taken by optimising horizontal probe position for each horizontal measurement position.

The effective interaction length peaks at approximately 0.8 mm. This is somewhat less than the value obtained from the fit to the data in Figure 5.4; this can be explained by the non-zero intercept of the fit, giving lower than expected phase shifts at these pressures (40–50 mbar). The elliptical shape of the data can be explained by comparison with the roughly elliptical shape of the near-field laser beam profile, as shown in Figure 3.6. Due to the short Rayleigh length, the ionised region is strongly affected by how the pump pulse diverges.

In summary, this data suggests that the FDH diagnostic is working properly, insofar as it is able to image the ionisation front, useful information given how new the technique is. Furthermore, it shows that the redesign of the IP was successful, and that the pump pulse is propagating through the gas cell without experiencing significant ionisation-induced defocussing. However, the non-linearities in Figure 5.4 suggest that the gas inlet pressure is not a reliable measure of the pressure inside the gas cell, especially in the region of the resonant pressure (17 mbar). It will be necessary to account

for this when searching for a wakefield, by scanning over a range of inlet pressures.

5.3 Search for a wakefield

As described previously, the original aim of the experiment was to investigate the behaviour of multiple-pulse laser wakefield acceleration (MP-LWFA) by measuring the amplitude of the generated wakefield as various parameters of the driving pulse train were changed. Before this could be done, it was necessary to be able to detect and measure the wakefield resulting from a single drive pulse. This turned out to be very challenging, for a number of reasons.

Firstly, even assuming correct alignment, the peak-to-peak phase shift of 200 mrad predicted by the simulations of Chapter 4 is not much greater than the noise level of 80 mrad (see § 2.5.3). Secondly, this alignment was difficult to achieve, as discussed in § 4.7; not only is the required tolerance very tight (about 2 μm), but the much greater size of the probe beam compared to the pump focus, and of the ionisation region compared to the wakefield, made it very difficult to achieve this alignment. Compounding this problem was the effect of beam jitter, which was observed on a scale comparable to the pump focal size, and was largely caused by incomplete decoupling of vibrations from the vacuum pump.

This last problem is especially significant, as even if a wakefield does appear the signal is likely to be intermittent. This makes it difficult not only to optimise the signal strength, but also to check that the signal is real rather than a statistical fluctuation. Essentially, the random beam jitter makes it impossible to reproduce the experimental configuration that resulted in a particular signal.

Finally, although the real-time data analysis system which I set up proved to be invaluable during the course of the experiment, especially while exploring the available parameter space, it was dependent on the availability of valid references. The validity of these references could quickly become questionable, either due to random changes in the laser behaviour or due to adjustments made by the experimenter; such adjust-




FIGURE 5.7 – FDH temporal phase data, with a helium inlet pressure of 38 mbar. The lower black box, marked “Signal”, contains a candidate wakefield signal. The upper box, marked “Background”, provides something to compare the signal region against. The contents of these regions are Fourier transformed to produce the results seen in Figure 5.9.

ments might include changing the transverse plane being imaged, in order to search for the location of the wakefield. This made locating the wakefield especially difficult, especially given the small signal-to-noise ratio even under ideal conditions.

Given these problems, I was unable to observe a definite wakefield signal. I did observe a number of possible such signals, some of which are presented here. However, these were never very clear, nor were they very repeatable; although the data was consistent with being generated by a wakefield, it was impossible to test this hypothesis rigorously.

One example is shown in Figure 5.7. The signal region indicated contains a faint periodic signal, possibly a wakefield. The extent of this region was chosen by visual inspection of the data, looking for evidence of fringes. Another example at a similar pressure is shown in Figure 5.8. In order to verify the presence of a periodic signal and ascertain its frequency, I performed a Fourier transform of the data within the signal region, averaged vertically, and with a Hamming window [97] applied. To distinguish




FIGURE 5.8 – FDH temporal phase data, with a helium inlet pressure of 40 mbar. Signal and background data regions are marked, to be used in the Fourier transform shown in Figure 5.9.

the signal from artefacts of the reconstruction process, I applied the same analysis to a background region, with the same size as the signal region but vertically offset. The results, shown in Figure 5.9, show a slight excess of signal over background around 10 THz in both cases.

The simulations of Chapter 4 predict that, at resonant density¹, the signal should have a frequency of 8 THz (a period of 125 fs; see § 4.4.3), a peak-to-peak amplitude of 200–400 mrad (see § 4.6.2), and cover a vertical extent of about 4 μm (see § 4.7.1). This signal covers a vertical extent of about 3 μm , has a frequency of about 10 THz (a period of about 100 fs), and a peak-to-peak amplitude of 200–300 mrad (see Figure 5.13), in rough agreement with the results of the simulations. The fact that this signal, despite being of the expected amplitude, is so unclear shows the difficulty of detecting a wakefield in this experiment.

Other examples of possible wakefield signals can be seen in Figure 5.10 and Fig-

¹The unclear relation between gas inlet pressure and plasma density means that the density at which this data was taken is not known accurately. I consider the resonant density case as a representative example.




FIGURE 5.9 – Fourier transforms of the data contained within the signal and background regions marked in Figure 5.7 and Figure 5.8. In both cases, the signal has a slight excess over the background around 10 THz.




FIGURE 5.10 – FDH temporal phase data, with a helium inlet pressure of 21 mbar. Signal and background data regions are marked, to be used in the Fourier transform shown in Figure 5.12.




FIGURE 5.11 – FDH temporal phase data, with a helium inlet pressure of 25 mbar. Signal and background data regions are marked, to be used in the Fourier transform shown in Figure 5.12.

ure 5.11; following the same analysis procedure, their Fourier transforms are shown in Figure 5.12. These data were recorded with a lower helium inlet pressure, and accordingly the signals appear to have a lower frequency of 6–7 THz.

If these signals are due to wakefields, their frequencies can be related to the plasma density, since $\omega_p^2 \propto n_e$. Given estimated frequencies of 10 THz and 6.5 THz, the ratio of plasma densities should be approximately 2.4 (with a significant margin of error due to the uncertainty in the measured frequencies). Judging by the helium inlet pressures, this ratio is too high, but Figure 5.4 suggests a slightly non-linear relation between inlet pressure and plasma density. A more accurate measurement might be obtained by looking at the phase shift due to ionisation, giving a ratio of about 2.

Line-outs of all the candidate wakefield signals shown in this section are plotted in Figure 5.13. This shows the peak-to-peak amplitude of the signals. In the data taken at 38 mbar this is about 200–300 mrad. The phase at $t < 0$ appears extremely noisy; this corresponds to the systematic vertical bands that can be seen in the same region of the 2D phase maps.




FIGURE 5.12 – Fourier transforms of the data contained within the signal and background regions marked in Figure 5.10 and Figure 5.11. In both cases, the signal has a slight excess over the background around 6–7 THz.




FIGURE 5.13 – Line-outs of the candidate wakefield signals shown in this section. In each case, the vertical position of the line-out corresponds to the centre of the marked signal box.

In summary, wakefield-like signals can be seen in some of the collected data. The peak-to-peak amplitude and vertical extent of these signals is similar to the predictions presented in Chapter 4, and the dependence of frequency on plasma density is not inconsistent with the behaviour of a wakefield ($\omega_p^2 \propto n_e$). However, practical considerations, especially beam jitter, made these observations very difficult to reproduce, and the quantity and quality of the data that was obtained is insufficient for a robust analysis. Therefore, these signals are not conclusive evidence of the presence of a wakefield.

5.4 Summary

The FDH technique was used to observe the ionisation front, providing information on the cross-sectional plasma density in the gas cell. This showed that the effective interaction length in the original design for the IP was far shorter than expected, motivating the move to a new design. Data from this new design justifies the move and shows that, when it comes to large scale features, the diagnostic works well.

Due to the small signal-to-noise ratio, combined with poor repeatability caused by beam jitter, no definitive evidence of a wakefield was observed. Although a signal that was not inconsistent with the presence of a wakefield was observed intermittently, this was not sufficient to proceed to the main part of the planned experiment, which would have required accurate measurements of the amplitude of the wakefield under varying conditions.

Nevertheless, this experience has provided valuable lessons for future experimental design. These include constraints on focal spot size, with the original design being too large and the later redesign being too small. As well as the practical experience of implementing FDH, these results indicate how large a signal is required to be measurable, as well as highlighting problems such as beam jitter. These points are covered in more detail in Chapter 7.

6. Pulse train generation

Controlled trains of laser pulses have many possible applications. In addition to their possible use in MP-LWFA, they can be used for quasi-phase matching in high-harmonic generation [98], coherent control of atomic and molecular quantum states [99], and to produce controlled trains of electron bunches by illuminating a photocathode [100]. While repetition rates up to several GHz can be generated by mode-locked oscillators directly, many applications, including MP-LWFA require repetition rates of the order of 1 THz. Possible approaches include combining the output of multiple lasers, such that each contributes a different pulse to the train, or repetition rate multiplication [101]. However, the approaches considered here all involve taking individual pulses from a mode-locked laser and splitting them up to form pulse trains.

For the purposes of this work, these pulse-splitting methods fall into two categories. Firstly, there are those which are simple to implement, since they operate on the full-energy compressed pulses and so require minimal changes to the laser. Although these are useful for many applications, including an MP-LWFA demonstration experiment, they are not suitable for a dedicated MP-LWFA driver laser — the fact that the full energy compressed pulse must be supported negates one of the key advantages of using a pulse train, the reduction in optical damage due to the lower peak power. Secondly, there are more invasive techniques, which split the pulse before amplification. These are more efficient and reduce the peak power present in the laser. Although they require more radical changes to the design of the laser system, this would not be a problem for a dedicated MP-LWFA driver laser.




FIGURE 6.1 – Use of a Michelson interferometer to generate a pair of pulses. **(a)** A standard Michelson interferometer, using two separate beamsplitters for clarity; and **(b)** a Michelson interferometer with polarising beamsplitter and half-wave plate, ensuring all the input energy ends up in a single output beam.

This chapter is split into two main parts. In the first, I consider two pulse train generation methods that were considered for the MP-LWFA demonstration experiment. These are of the first type: non-invasive, but inefficient. In the second part, I consider a possible method of the second type, for use in a possible fibre-based MP-LWFA driver laser. In all cases, I aim for a train of evenly spaced pulses with approximately equal energy, and as little wasted energy as possible. I also consider pulse trains with non-uniform spacing, as these could be used to compensate for resonant detuning or for a proposed “energy recovery” mode, where a second pulse train extracts energy from the wakefield to reduce heating of the plasma in a high repetition rate accelerator (see for instance the mention of a trailing anti-resonant laser pulse in [54]).

6.1 Multiple Michelson pulse splitter

It is straightforward to split one pulse into a pair of temporally-separated pulses by building a Michelson interferometer (or some other kind of interferometer) and offsetting the arms, so that when the beams are recombined, one pulse arrives after the other. The separation between the pulses can be adjusted by changing the length of one of the arms. Such a setup is depicted in Figure 6.1(a).




FIGURE 6.2 – Use of multiple Michelson interferometers to generate a pulse train. The pulse trains present at various points in the system are indicated.

A notable drawback of this approach is that, since the pulses do not interfere on the combining beamsplitter, there are two output beam lines, each containing half the energy of the input pulse. This can be overcome by using a polarising beamsplitter to recombine the beam; a half-wave plate placed into one arm ensures that both pulses end up in the correct output beam line, as shown in Figure 6.1(b).

A number of these devices can be chained together, each step doubling the number of pulses in the train [102, 103]. The resulting pulse train can be made uniform by setting the arm lengths appropriately. This is realised by building multiple copies of device (a), where the combining beamsplitter of one stage doubles as the splitting beamsplitter of the next. This ensures that no energy is lost. The final stage uses design (b), since only a single output beam is wanted¹. An example of such a setup is shown in Figure 6.2.

¹It is impossible to chain multiple copies of design (b), since it requires a polarised input but produces a mixed-polarisation output.

6.1.1 Dispersion and non-linear effects

The use of transmissive optics are usually avoided in the context of ultrafast lasers, due to the problems of dispersion and non-linear effects. The large bandwidth of an ultrafast laser pulse renders it susceptible to chromatic dispersion, which separates the different frequency components of the pulse, lengthening it. The high peak intensity of such a pulse may induce various non-linear effects, which also distort the shape and spectrum of the pulse. However, this design requires the use of transmissive optics for the beamsplitter and half-wave plate. Furthermore, if the apparatus is placed after the compressor, the diameter of the optics will need to be large (typically > 75 mm) to reduce the peak laser intensity and avoid damage; this results in a considerable thickness of material, due to the challenges involved in manufacturing large, thin glass substrates of sufficient flatness.

Since this was one of the methods considered for the MP-LWFA demonstration experiment, it was necessary to investigate these effects, to determine how severe they would be and how thick the beamsplitters could be — thinner optics being preferable but more expensive. I accomplished this by developing a numerical simulation that included two effects: dispersion up to 4th order (the $\frac{d^4\phi}{d\omega^4}$ term); and self-phase modulation. Self-phase modulation is a non-linear effect, in which the $\chi^{(3)}$ term of the non-linear susceptibility leads to an intensity-dependent contribution to the refractive index. These effects can be combined in the formula for the refractive index n :

$$\begin{aligned} n &= n_0(\omega) + n_2 I \\ &= n_0^{(0)} + n_0^{(1)}(\omega - \omega_0) + n_0^{(2)}(\omega - \omega_0)^2 + n_0^{(3)}(\omega - \omega_0)^3 + n_2 I \end{aligned} \tag{6.1}$$

Where n_0 is the linear portion of the refractive index, which is Taylor-expanded up to 3rd order in $(\omega - \omega_0)$. n_2 is the (lowest-order) non-linear portion of the refractive index; although this is in principle also a function of frequency ω , this higher-order effect is dwarfed by the uncertainty in the measurement of n_2 itself. I is the intensity, in units of W m^{-2} , so that n_2 has units $\text{m}^2 \text{W}^{-1}$.

Split-step method

Consider a laser pulse with electric field:

$$\mathbf{E} = \sqrt{\frac{n\epsilon_0 c}{2}} A(z, t - \frac{z}{c}) e^{i(k_0 z - \omega_0 t)} \mathbf{e}_x \quad (6.2)$$

so that $I = |A|^2$. Given a refractive index of the above form, this laser pulse evolves according to the non-linear Schrödinger equation [104, Chap. 2]:

$$i \frac{\partial A}{\partial z} = -i\beta_1 \frac{\partial A}{\partial t} + \frac{\beta_2}{2} \frac{\partial^2 A}{\partial t^2} + i \frac{\beta_3}{6} \frac{\partial^3 A}{\partial t^3} - \frac{\beta_4}{24} \frac{\partial^4 A}{\partial t^4} - \frac{n_2 \omega_0}{c} |A|^2 A \quad (6.3)$$

The right-hand side of this equation can be divided into a dispersive part and a non-linear part, and written in operator notation thus:

$$\frac{\partial A}{\partial z} = (\hat{D} + \hat{N}) A \quad (6.4)$$

$$\hat{D} = -\beta_1 \frac{\partial}{\partial t} - i \frac{\beta_2}{2} \frac{\partial^2}{\partial t^2} + \dots \quad (6.5)$$

$$\hat{N} A = i \frac{n_2 \omega_0}{c} |A|^2 A \quad (6.6)$$

The general solution to this equation is:

$$A(z + h) = \exp[h(\hat{D} + \hat{N})] A(z) \quad (6.7)$$

$$A(z + h) \approx \exp(h\hat{D}) \exp(h\hat{N}) A(z) \quad (6.8)$$

While the first equality is exact, the second is only approximate because the $\hat{D}$ and $\hat{N}$ operators do not commute. However, this approximation is useful, because while there is no direct way of computing the effect of the operator $\exp[h(\hat{D} + \hat{N})]$, the $\exp(h\hat{D})$ and $\exp(h\hat{N})$ operators can be computed exactly for any h ; these calculations are performed in the frequency and time domains respectively, which is the underlying cause of the non-commutativity.

This leads to the split-step method of solving the non-linear Schrödinger equation,

which involves iterated application of Equation 6.8 with a sufficiently small value of h . This particular approach can be considered as the equivalent of the Euler method for solving ODEs: in particular, the final error will be of order h , and higher-order methods exist which are more complex but also more computationally efficient.

Simulation

I created a simple implementation of the split-step method to model the propagation of a laser pulse through a multiple-Michelson pulse splitter. The required parameters were:

- The FWHM duration of the incident laser pulse;
- The peak intensity of the incident laser pulse;
- The material dispersion values β_i for the (assumed) fused silica substrate of the beamsplitters, at 800 nm;
- The non-linear refractive index n_2 of fused silica at 800 nm.

This simulation was performed before I had any measurements of the laser performance, so I assumed a FWHM duration of 50 fs, a uniform profile with diameter 50 mm, and a total pulse energy of 90 mJ. The material dispersion values were calculated based on the Sellmeier equation (see § 3.2.4), the parameters for which are readily available (e.g. [105]). It is harder to find a value for the non-linear refractive index n_2 , because it is more difficult to measure; in particular, there may be no measurements available at the wavelength of interest. Based on the measurements given in [106] I interpolated a value of $n_2 = 2.87 \times 10^{-20} \text{ m}^2 \text{ W}^{-1}$ at $\lambda = 800 \text{ nm}$.

The simulation also requires a sequence of material blocks and beamsplitter surfaces (each of which reduces the incident intensity by a factor of 2). Although there are eight possible paths through the apparatus depicted in Figure 6.2, by correctly orienting the beamsplitters, it is possible to ensure that every pulse passes through the equivalent of exactly three beamsplitters. For the sake of brevity, I am neglecting the

difference between the polarising and non-polarising beamsplitters, and ignoring the effect of the half-wave plate; a simulation including these effects was performed, with similar results. Since the beamsplitters are oriented at 45° to the incident beam, and have a refractive index of approximately 1.45, the pulse propagates at 29° through the beamsplitter, increasing its effective thickness by 14%.

This dispersion results in a substantial increase in the pulse duration, even with beamsplitters only a few mm thick. However, it can mostly be compensated by adding an equal and opposite quantity of GDD (second-order dispersion) to the pulse before it passes through the Michelson. In practice, this would be accomplished by detuning the pulse compressor — by adjusting the separation between the diffraction gratings within the compressor, the amount of dispersion it adds to the pulse can be changed. I incorporated this into the simulation by assuming that the pulse would be perfectly compressed when the grating separation was optimised, and using the appropriate equations [83, eqns. 21–23] to calculate the quantities of second through fourth order dispersion that would be added when this separation was changed.

As shown in Figure 6.3, this significantly reduced the pulse durations predicted for longer beamsplitters. In the linear case, this would completely compensate for the dispersion introduced by the beamsplitters (higher-order dispersion is too small to be a problem). However, the presence of self-phase modulation complicates the situation, preventing complete compensation.

6.1.2 Pulse splitting before compression

The high intensity incident on the first beamsplitter in this design causes a number of problems: the beamsplitters must have a diameter of about 75 mm, but be only a few mm thick, and even then the pulses are broadened slightly due to self-phase modulation. The size of the beamsplitters not only makes them more expensive, but also increases the space required by the apparatus.

A simple method of reducing the intensity, allowing the use of smaller, thicker




FIGURE 6.3 – Dependence of the output pulse duration on the (perpendicular) thickness of a single beamsplitter, both with and without compressor adjustment to compensate for dispersion.

optics, would be to place the pulse splitter immediately before the pulse compressor. A typical chirped-pulse amplification system stretches the pulses from about 100 fs to several hundred ps, reducing the peak power by a factor of several thousand. That would allow the beam diameter to be reduced from 50 mm to 10 mm, allowing use of 25 mm diameter optics, while still reducing the peak intensity by an approximate factor of 100, allowing several mm of beamsplitter thickness (easier to achieve with smaller optics) without a significant effect on pulse duration. This could be improved further by placing the pulse splitter before the main amplifiers — this has recently been demonstrated [107] as part of generating short pulse trains for use in illuminating photocathodes.

A secondary advantage of this method is that it would reduce the peak intensity at the exit of the compressor, since the pulse has already been split into a train by that point. This peak intensity is the main consideration in the output beam diameter, since the intensity must be low enough not to cause any damage to the diffraction gratings.

Generating a pulse train before the compressor will therefore allow the beam size to be reduced, allowing the use of smaller diffraction gratings in the compressor. This could result in significant cost savings, due to the difficulty of creating large, high-quality diffraction gratings.

However, it is important to consider whether this pulse train generation method will work correctly before the compressor, where the pulses are much longer than the intended intra-train pulse spacing, and whether the output will affect the performance of the compressor. Firstly, the use of a polarising beamsplitter and half-wave plate to achieve a single output beam is incompatible with placement before the compressor, because compressor efficiency is highly sensitive to polarisation. The compressor requires the entire input pulse to have a particular (linear) polarisation, which cannot be achieved when combining two beams in a polarising beamsplitter. Therefore, a multiple Michelson *without* a final polarising beamsplitter is needed, which will have two equal-energy output beams. In principle, these could be compressed separately and then combined with a polarising beamsplitter, although this would introduce additional complexity. Otherwise, the energy of the laser would be reduced by 50%.

Secondly, since the pulse spacing is shorter than the pulse duration, there will be interference whenever two beams are combined on a beamsplitter. In the case of a single Michelson, if the input electric field has the form:

$$\tilde{E}(\omega) \exp i\frac{\beta_2}{2}(\omega - \omega_0)^2 \quad (6.9)$$

the output electric field will have the form:

$$\frac{1}{2}\tilde{E}(\omega) \exp \left[i\frac{\beta_2}{2}(\omega - \omega_0)^2 \right] (1 \pm \exp i\omega\Delta t) \quad (6.10)$$

where one of the output beams takes the + case, and the other the – case. This results in the spectrum of each output beam taking on a $\sin^2 \omega\Delta t$ dependence. In particular, there will be particular wavelengths where this value is zero for one beam, and all the

intensity ends up in a single output beam. The strong linear chirp means that this also manifests in the temporal profile; at particular times within the pulse, one output beam has the same intensity as the incident beam. This is unlike the case where the apparatus is placed after the compressor, in which each beamsplitter reduces the peak intensity by a factor of 2;

Nevertheless, the superposition principle ensures that the compressor will work correctly, giving a train of separated compressed pulses. The input to the compressor will be a superposition of identical chirped pulses, where each has a different group delay:

$$\frac{1}{\sqrt{2N}} \tilde{E}(\omega) \exp \left[i \frac{\beta_2}{2} (\omega - \omega_0)^2 \right] (1 + \exp i\omega\Delta t + \exp 2i\omega\Delta t + \cdots) \quad (6.11)$$

The compressor will remove the group-delay dispersion, leaving a superposition of compressed pulses, which will not interfere:

$$\frac{1}{\sqrt{2N}} \tilde{E}(\omega) (1 + \exp i\omega\Delta t + \exp 2i\omega\Delta t + \cdots) \quad (6.12)$$

6.2 Single-crystal pulse splitter

A significant difficulty with the multiple Michelson interferometer approach to pulse splitting is in the alignment of the multiple beam paths. Unless carefully aligned, each pulse will emerge from the interferometer at a different position and propagating in a different direction. An alignment-free method of splitting a pulse into two is to use a birefringent crystal — if the fast axis is oriented at 45° to the polarisation direction of the incident light, the difference in group velocities will result in two pulses emerging, with orthogonal polarisations. Multiple crystals, with the direction of the fast axis changing by 45° each time and with the length doubling each time, can thus produce a pulse train [108, 109].

This does not require any alignment, but the device will have a considerable total




FIGURE 6.4 – An illustration of the single crystal method of pulse train generation. A chirped pulse (a) is passed through a birefringent crystal, which causes the polarisation of the pulse to vary with time (b). Passing the pulse through a polariser then applies a $\sin^2$ modulation to the temporal profile (c). A two-pulse train is shown here, but much longer trains can be made.

length. This is not only inconvenient, but the laser pulses will suffer from chromatic dispersion and possibly non-linear effects while propagating through the material. An alternate method [109–111] uses only a single birefringent crystal to act as a multiple-order waveplate. This introduces a wavelength-dependent polarisation to a chirped pulse, which is then passed through a polariser to convert this to a $\sin^2$ temporal profile.

When a linearly-polarised laser pulse is passed through a birefringent crystal, with its fast axis oriented at 45° to the direction of polarisation, the crystal acts as a multiple-order waveplate. It introduces an optical path length difference between the two polarisation components, which in turn causes a phase shift proportional to ω . This phase shift changes the polarisation state; since a phase shift which is a multiple of 2π will have no effect, the final polarisation state is a periodic function of ω . If the pulse is then passed through a polariser set to transmit the original polarisation, the transmission will also be a periodic function of ω , with a $\sin^2$ form. By applying a linear chirp to the laser pulse before entering the crystal (by offsetting the pulse compressor), this translates to a $\sin^2$ modulation of the temporal profile, effectively giving a pulse train. This technique is illustrated in Figure 6.4.

The number of pulses within a train is determined by the period of the modulation in the frequency domain and the bandwidth of the laser: the former is determined by the optical path length difference, and hence by the thickness of the crystal. Given a

particular number of pulses, the pulse spacing is determined by the duration of the chirped pulse, and hence by the amount of linear chirp applied. Since the train has a $\sin^2$ profile, the sub-pulse duration and spacing have a fixed ratio (e.g. 1:2, if measuring the FWHM duration).

The principal advantage of this method over the multiple Michelson method described in § 6.1 is that it is significantly easier to align: instead of having to overlap N separate beams to produce a train of N pulses, the beam is passed through a single crystal followed by a polariser. Each output beam from the polariser contains half the energy of the input pulse, so that if only one is used the whole system is only 50% efficient. However, it is possible to build the polariser into a Michelson interferometer and recombine the two pulse trains; by adding a delay to one arm, the two pulse trains can be separated when they are recombined, giving in effect a single pulse train twice as long as the original.

The main disadvantage of this method is reduced pulse train quality. Firstly, the whole pulse train exists as a $\sin^2$ modulation of an ideally Gaussian envelope, so the energy is not equally distributed between the pulses. Secondly, the pulse train arises due to interference between the slow and fast polarisation components at the polariser: the birefringent crystal introduces an ω -dependent phase shift between these components, but this is a group delay. This delay means that the two components do not perfectly overlap, limiting the degree of interference at the beginning and end of the pulse train. This manifests as a limited visibility at the edges of the pulse train, which can be expected to reduce the efficiency with which the pulse train excites a wakefield.

6.2.1 Simulations

Much of the work done on this method to date has been related to its use in quasi-phase matched high harmonic generation [109] and its application to MP-LWFA had not previously been considered. To check its suitability, I created a simulation of the process. In particular, this was used to investigate whether it would be possible to

compensate for resonant detuning — this requires a pulse spacing that increases along the train, depending quadratically on pulse number in the low-amplitude limit [112].

The simulation was based on decomposing the incident laser pulse into two polarisation components, aligned with the fast and slow axes of the birefringent crystal, and applying the appropriate phase shift to each:

$$\phi(\omega) = \frac{2\pi}{\lambda} n(\omega) L \quad (6.13)$$

where $n(\omega)$ is the refractive index, computed using the Sellmeier equation, and L is the thickness of the crystal. I investigated both quartz and calcite as the birefringent crystal, using the Sellmeier coefficients given by [113]. By adjusting the values of the second, third and fourth order dispersion of the pulse on entering the crystal, I could control the pulse spacing up to second order in pulse number n . Note that such independent control cannot be achieved simply by offsetting a pulse compressor — some more sophisticated mechanism would be required, typically an acousto-optic programmable dispersive filter [114], available commercially as a *Dazzler*.

By way of example, I generated the pulse train shown in Figure 6.5 using a simulated laser pulse with 40 nm FWHM bandwidth, 1 mm of calcite (I could obtain similar results using 20 mm of quartz), and the dispersion values shown in Table 6.1. The main peaks (those with at least 40% of the peak intensity) were identified, and I tuned the dispersion parameters so that the spacing was $(100 + 0.3n^2)$ fs (where n is the pulse number, starting from 0), as shown in Figure 6.6. The quadratic form of this target pulse spacing is approximately what is required to compensate for resonant detuning in the low amplitude limit [112] — these particular parameters were chosen to produce reasonable pulse spacings in the 100–200 fs range.

Since the bandwidth of 40 nm appeared unnecessarily high for a pulse train with 100 fs spacing, I attempted to produce a similar pulse train with a FWHM bandwidth of only 20 nm. The results are shown in Figure 6.7 and Figure 6.8.

The spacings still appear correct, with the anomaly at $n = 0$ due to the very low




FIGURE 6.5 – A simulated pulse train, generated by the single crystal method. Those pulses with a peak intensity at least 40% of the maximum peak intensity are highlighted, and were used for the spacing measurement.




FIGURE 6.6 – Pulse spacings extracted from the data shown in Figure 6.5, compared to the target spacing of $(100 + 0.3n^2)$ fs.

Parameter	Value	Units
$\phi^{(2)}(\omega)$ (GDD)	1.08×10^4	fs^2
$\phi^{(3)}(\omega)$ (TOD)	-1.82×10^4	fs^3
$\phi^{(4)}(\omega)$ (FOD)	9.0×10^4	fs^4

TABLE 6.1 – Values of the group-delay dispersion, third-order dispersion, and fourth-order dispersion used as inputs to produce the pulse train shown in Figure 6.5.




FIGURE 6.7 – As Figure 6.5, but using a simulated laser pulse with only 20 nm FWHM bandwidth.




FIGURE 6.8 – Pulse spacings extracted from the data shown in Figure 6.7, compared to the target spacing of $(100 + 0.3n^2)$ fs.

contrast ratio of that pulse. There are two key differences when compared to the data in Figure 6.5. Firstly, there are fewer pulses; this is because the periodic modulation introduced by the calcite has a fixed period in frequency space, so with less bandwidth there are fewer periods. Secondly, the contrast ratio is worse; this is because the chirped pulse is shorter, but the calcite introduces the same delay between the two polarisation components, so they overlap less and interfere less strongly. Note that attempting to increase the number of pulses by increasing the thickness of the calcite will make the contrast ratio worse, as the polarisation components are further separated.

In summary, my simulations show that control over the spectral phase up to fourth-order allows the birefringent crystal method to produce pulse trains with tailored spacing, with control up to second-order in pulse number. However, the contrast ratio can be quite poor if the bandwidth of the incident laser pulse is insufficient.

6.3 Pulse train generation by pulse shaping

Perhaps the most flexible way to generate a pulse train is to use one of a number of techniques developed for temporal pulse shaping [115, 116]. This process can be understood in terms of a filter being applied to the pulse in the frequency domain, which results in a change to the pulse temporal profile. Various technologies exist that allow such a filter to be programmed nearly in real time, providing the user with nearly arbitrary control over the temporal shape of the pulse. Using such a system to generate a pulse train would allow far greater control than either of the previous methods, as well as combining 100% throughput with a completely polarised output.

The particular pulse-shaping method that I will discuss here uses a $4f$ line as shown in Figure 6.9, also called a zero dispersion line, to create a plane in the middle of the apparatus where frequency is directly mapped to space — that is, each wavelength component is focussed at a different transverse position on this central plane. Various filters can be placed in this plane; in this case I will consider a liquid crystal spatial light modulator (LC-SLM). This is a programmable phase filter based on the same technology used to make liquid crystal displays (LCDs). Each pixel of an LC-SLM acts as an independently controlled, programmable waveplate. In an LCD, the pixels are square and the addition of integrated polarisers converts the programmable waveplates into programmable transmission elements. In this case, the device is pixelated in one dimension only, so that the pixels are long and thin, and there are no polarisers; by orienting the input polarisation along one of the axes of the waveplate, phase can be controlled pixel-by-pixel without any effect on amplitude or polarisation.

If this method were to be used as part of a dedicated MP-LWFA driver laser, the pulse shaper would need to be placed before the laser amplifiers due to its low damage threshold. It is important to consider possible interactions between the pulse train and the amplification process, such as distortions of the pulse train or less efficient amplification. I consider these in more detail in § 6.3.7, concluding that harmful effects can be minimised by avoiding any changes to the spectral amplitude. This is the same




FIGURE 6.9 – Schematic diagram of a $4f$ line. The apparatus consists of two identical diffraction gratings, two identical lenses, and some kind of spatial filter (acting here as a spectral filter). The space between each pair of components is equal to the focal length of the lenses, so that the whole apparatus is 4 focal lengths long, hence the name.

conclusion reached by Umstadter *et al.* – in their discussion of pulse train generation techniques for MP-LWFA [53, Sec. V] they recommend the use of phase masks in preference to amplitude masks (in the spectral domain) to avoid “nonlinear interference between pulses”.

The aim then is to find a method of constructing a phase-only spectral filter which will result in a train with a particular number of pulses. The pulse train must have a good contrast ratio and ideally uniform pulse intensities. Initially I will consider the case of uniform pulse spacing, before extending the method to non-uniform trains. Additionally, the method should work with a commercially-available LC-SLM and be scalable to many tens of pulses without requiring an especially large $4f$ line.

6.3.1 Brief history

An LC-SLM was first used, as part of a $4f$ line, for programmable femtosecond pulse shaping by Weiner *et al.* in 1989 [117]; grey-level control was demonstrated in that experiment, allowing GDD to be added to compress the pulses. This first device consisted of only 32 pixels, their centres $100\text{ }\mu\text{m}$ apart and with $5\text{ }\mu\text{m}$ gaps between adjacent pixels. Further development allowed the number of pixels on a single device to be

increased, with 512-element modulators available by 2011 [118], making it practical to shape higher-bandwidth pulses.

The use of pulse shaping to create pulse trains also dates back to 1989, with Weiner *et al.* demonstrating the generation of a multi-THz pulse train using a phase-only filter inscribed on a fused silica substrate [119]. The filters used were periodic repetitions of an M (maximal length) sequence, a pseudorandom binary code. Since the M sequence is binary, each pixel has a phase shift of either zero or $\Delta\phi$; the value of $\Delta\phi$ controls the intensity of the central peak. The pulse trains thus obtained were not uniform in intensity, but by using Dammann gratings [120, 121], Weiner *et al.* were able to improve this.

However, Dammann gratings do not use uniform pixel sizes, so they are not suitable for use with a programmable LC-SLM. Additionally, the use of binary rather than grey-level control of the phase ignores a major advantage of the LC-SLM. Other pulse train generation techniques include Multiple Independent Comb Shaping [122], which does not appear to scale to many 10s of pulses.

The advantages of the method presented in this chapter are that it is designed to work with the particular properties of an LC-SLM — uniform (spatial) pixel size, grey-level control of phase — it scales to many pulses, and it allows good intensity uniformity in the resulting pulse train.

6.3.2 Simple model

To begin with, I consider a simplified model of the $4f$ line. In particular, I assume the focal spot size at the LC-SLM to be zero, so that each distinct wavelength passes through exactly one pixel, and I assume that the mapping from frequency to spatial position is linear in frequency, that is $x = \alpha\omega$. These assumptions will enable me to construct a simple model relating the filter to the resulting pulse train, which can then be used for optimisation. This is not, as far as I am aware, a standard derivation; it was developed during the course of this work.

First, consider a spectral filter, $\tilde{F}(\omega)$, applied to an input spectrum $\tilde{E}_{in}(\omega)$, to produce an output spectrum:

$$\tilde{E}_{out}(\omega) = \tilde{F}(\omega) \times \tilde{E}_{in}(\omega) \quad (6.14)$$

By the convolution theorem,

$$E_{out}(t) = F(t) * E_{in}(t) \quad (6.15)$$

Regardless of the form of E_{in} , a uniformly spaced pulse train can be produced by making $F(t)$ a uniformly spaced comb of Dirac delta functions:

$$F(t) = V(t) \times \sum_j \delta(t - j\Delta t) \quad (6.16)$$

where $V(t)$ is an envelope function that determines the complex amplitude of each pulse.

Fourier transforming this gives:

$$\tilde{F}(\omega) = \tilde{V}(\omega) * \sum_j \delta(\omega - j\Delta\omega) \quad (6.17)$$

where $\Delta\omega = \frac{2\pi}{\Delta t}$. This shows that $\tilde{F}(\omega)$ is a periodic function, whose period $\Delta\omega$ is equal to 2π times the repetition rate of the pulse train. The precise form of this periodic function, given by $\tilde{V}(\omega)$, determines the envelope $V(t)$ of the pulse train.

When using an LC-SLM, the form of $\tilde{V}(\omega)$ is constrained, because it must be composed of a finite number of adjacent pixels. Given that the LC-SLM consists of N pixels per period of $\tilde{F}(\omega)$, so that each pixel covers a frequency interval $\delta\omega = \frac{\Delta\omega}{N}$, and assuming that there are no gaps between the pixels and that each is perfectly uniform,

$\tilde{V}(\omega)$ takes the form:

$$\tilde{V}(\omega) = \sum_{k=0}^{N-1} X_k \operatorname{rect}\left(\frac{\omega}{\delta\omega} - k\right) \quad (6.18)$$

$$\operatorname{rect}(x) = \begin{cases} 1 & \text{if } |x| < \frac{1}{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (6.19)$$

where X_k is the complex amplitude of the k -th pixel. This can be re-written as a convolution:

$$\tilde{V}(\omega) = \operatorname{rect}\left(\frac{\omega}{\delta\omega}\right) * \sum_{k=0}^{N-1} X_k \delta(\omega - k\delta\omega) \quad (6.20)$$

Inverse Fourier transforming this gives:

$$V(t) = \operatorname{sinc}\left(\frac{t\delta\omega}{2\pi}\right) \times \sum_{k=0}^{N-1} X_k e^{ik\delta\omega t} \quad (6.21)$$

where $\operatorname{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$ is the normalised sinc function.

Recalling Equation 6.16, it is useful to define P_j as the complex amplitude of the j -th pulse in $F(t)$:

$$P_j = V(j\Delta t) = \operatorname{sinc}(j/N) \times \sum_{k=0}^{N-1} X_k \exp i \frac{2\pi j k}{N} \quad (6.22)$$

Thus P_j depends on $\tilde{X}_j$, the discrete Fourier transform (DFT) of X_k :

$$\tilde{X}_j = \sum_{k=0}^{N-1} X_k \exp i \frac{2\pi j k}{N} \quad (6.23a)$$

$$P_j = \tilde{X}_j \operatorname{sinc}(j/N) \quad (6.23b)$$

Due to the basic properties of the DFT, $\tilde{X}_j$ is periodic, with period N . Therefore, a pulse train centred on $j = 0$ will cause duplicate trains to appear at $j = N, 2N, \dots$. However, since these are centred around zeroes of $\operatorname{sinc}\left(\frac{j}{N}\right)$, they will be much lower

in amplitude. The best-case contrast ratio that can be achieved is limited by these duplicate trains. Take a uniform train of n pulses, where n is odd, with all other P_j between $j = -\frac{N}{2}$ and $j = \frac{N}{2}$ being zero. Then the contrast ratio is the ratio of any pulse in the central train e.g. the last ($j = (n - 1)/2$) to the last pulse of the train centred on $j = -N$ ($j = (n - 1)/2 - N$). Since the $\tilde{X}_j$ values are the same for both pulses (by periodicity), this is merely the ratio of sinc functions, which reduces to:

$$\frac{P_{(n-1)/2}}{P_{(n-1)/2-N}} = \frac{2N}{n-1} - 1 \quad (6.24)$$

This is the ratio of amplitudes. The energy contrast ratio is thus given by:

$$C = \left(\frac{2N}{n-1} - 1 \right)^2 \quad (6.25)$$

For example, a train of 21 pulses generated by a SLM with 100 pixels per period, cannot have a contrast ratio better than 81:1.

6.3.3 Optimisation

In principle, given a particular target pulse train P_j (specified for a single period), Equation 6.23 could be reversed to give the required periodic filter X_k to produce that pulse train. In this case, however, we have the additional constraint that $|X_k| = 1$ for all k ; that is, that the filter only affect the spectral phase. One way to deal with such an over-constrained problem is to use an optimisation algorithm. In this case, with many dimensions and a highly non-linear relationship between input (phase) and output (pulse intensity), a global optimisation algorithm is needed, such as simulated annealing [123, § 10.12] or a genetic algorithm [124]. Such an algorithm needs a set of parameters to vary and an objective function of those parameters to minimise. In this case, the parameters would be a set of N phases φ_k . Evaluating the objective function

$G(\varphi_k)$ would consist of three steps:

$$X_k = \exp i\varphi_k \quad (6.26a)$$

$$P_j = P_j(X_k) \quad (\text{see Equation 6.23}) \quad (6.26b)$$

$$G = G'(P_j) \quad (6.26c)$$

Where $G'(P_j)$ is an objective function for the pulse train itself.

I implemented the objective function $G'(P_j)$ as the sum of two terms: a uniformity term and a contrast ratio term. First, the N pulses centred on $j = 0$, forming the central period of P_j , are classified as either “on” pulses or “off” pulses. The “on” pulses are the central (around $j = 0$) pulses which are expected to form part of the pulse train, whereas the “off” pulses are the others, which are expected to have zero energy. The uniformity term is then calculated by taking the relative standard deviation of the energies of the “on” pulses, and the contrast ratio term is calculated by dividing the maximum energy of any of the “off” pulses by the mean energy of the “on” pulses. Given N pixels per period and a target pulse train length of n pulses, where n is odd, these can be defined thus:

$$E_+ = \left\{ |P_j|^2 \text{ for all } j \text{ where } -\frac{n-1}{2} \leq j \leq \frac{n-1}{2} \right\} \quad (6.27)$$

$$E_0 = \left\{ |P_j|^2 \text{ for all other } j \right\} \quad (6.28)$$

$$U = \frac{\sqrt{\text{Var}[E_+]}}{\langle E_+ \rangle} \quad (6.29)$$

$$C = \frac{\text{Max}[E_0]}{\langle E_+ \rangle} \quad (6.30)$$

$$G'(P_j) = U + C \quad (6.31)$$

Clearly G' is minimised (at 0) for a perfectly uniform pulse train (where all the “on” pulses have the same energy) with the best possible contrast ratio (where all the “off” pulses have zero energy). Any loss of quality will result in an increase in G' . Therefore, the optimisation process will correctly improve the quality of the pulse train.




FIGURE 6.10 – An example of an SLM phase filter, generated by an optimisation algorithm. The SLM in this example is made up of 80 pixels, and the filter itself is 20 pixels long, so that it repeats 4 times.

For the optimisation algorithm itself, I used the simulated annealing subroutines from the GNU Scientific Library [78]. As a simple example, I created a filter, shown in Figure 6.10, that generates a 7-pulse train, shown in Figure 6.11. This can be seen to have very good uniformity, contrast ratio, and visibility, especially compared with the pulse trains generated by the single crystal method (see Figure 6.5 and Figure 6.7) which have very poor uniformity and visibility.

6.3.4 Spatial chirp

This simple model does not account for several aspects of the behaviour of a $4f$ line/LC-SLM combination; the most significant of these is the non-linear mapping between frequency and spatial position. The position to which a given frequency is mapped is determined by the grating equation. Using the parameters shown in Figure 6.12, this




FIGURE 6.11 – The pulse train generated by the filter shown in Figure 6.10.




FIGURE 6.12 – Part of a $4f$ line with important parameters labelled. These are the angle of incidence θ_i , the angle of diffraction for the centre wavelength θ_d , the angle of diffraction for some other wavelength θ'_d , the focal length f , and the spatial offset at the SLM x .

can be written:

$$\sin \theta'_d + \sin \theta_i = \frac{2\pi c}{\omega d} \quad (6.32)$$

$$\theta'_d(\omega) = \arcsin \left(\frac{2\pi c}{\omega d} - \sin \theta_i \right) \quad (6.33)$$

$$\theta_d = \theta'_d(\omega_0) = \arcsin \left(\frac{2\pi c}{\omega_0 d} - \sin \theta_i \right) \quad (6.34)$$

$$\begin{aligned} x(\omega) &= f \tan[\theta_d - \theta'_d] \\ &= f \tan \left[\arcsin \left(\frac{2\pi c}{\omega_0 d} - \sin \theta_i \right) \right. \\ &\quad \left. - \arcsin \left(\frac{2\pi c}{\omega d} - \sin \theta_i \right) \right] \end{aligned} \quad (6.35)$$

where d is the line spacing of the diffraction grating. Although this can be approximated as $x = \alpha\omega$, as above, there is a significant second-order effect. This makes the effective pixel size in frequency space, $\delta\omega \approx \frac{d\omega}{dx}\delta x$, greater at one end of the SLM than the other. This non-linear mapping was observed by Wang *et al.* [118], who took advantage of the effect to reduce the intensity of replica pulses.

However, in this case the effect, which I refer to as SLM spatial chirp, will adversely affect the pulse train. Although the period of the phase mask is a constant number of pixels, its period in frequency $\Delta\omega$ varies with wavelength. This causes the pulse spacing Δt to vary with wavelength also, which has the effect of causing the pulses to become increasingly chirped the further from the centre ($j = 0$) they are. This is a particular problem for driving a wakefield, since the peak intensity is important.

A simulation taking account of these effects is much more detailed, since it must sample many points inside each pixel in order to account for the varying effective pixel size correctly. I created such a model using the parameters shown in Table 6.2; these were chosen to be as close as possible to a hypothetical fibre-based MP-LWFA driver laser. In particular, the centre wavelength is a typical value for a large bandwidth ytterbium-doped fibre laser [125], the geometric parameters of the $4f$ line give a reasonable size of apparatus with easily available gratings, and the SLM parameters

Parameter	Value	Units
Centre wavelength	1030	nm
Laser FWHM bandwidth	15	nm
Focal length	330	mm
Grating line density	1400	lines/mm
Grating incident angle	50	°
SLM pixel width	102	µm
SLM pixel count	320	
Pulse train period	271	fs

TABLE 6.2 – Summary of the parameters of the laser and pulse shaper discussed in this section.

are based on a commercially available device (the JENOPTIK SLM-S320). The selected bandwidth can already be achieved by Yb-doped fibre oscillators [47], and might be achieved by high peak power systems in the future.

The effect of SLM spatial chirp can be seen clearly in Figure 6.13, as the pulses become increasingly distorted as their distance from the centre of the train increases. To verify that the chirps of the pulses vary across the train, I calculated the temporal chirp $\frac{d^2\phi}{dt^2}$ at the centre of each pulse, which is shown in Figure 6.14. The linear dependence of chirp on pulse number near the centre of the train is in accordance with my hypothesis. Towards the edges, the pulses become more distorted and the measurement becomes less reliable.

In order to correct the SLM spatial chirp, it is necessary to ensure that the periodicity of the phase mask in frequency is constant. This means that the periodicity in terms of pixels will vary; clearly, the mask cannot be perfectly periodic, since each period will contain a slightly different number of pixels. However, by adding or removing a certain number of evenly-spaced pixels to/from each period, the number of pixels in each period can be changed while preserving approximate periodicity. Applying this to the mask used in Figure 6.14 produces the train shown in Figure 6.15. The strong linear dependence of chirp on pulse number has disappeared, showing that the varying periodicity has been corrected. However, the change to the mask has adversely affected the uniformity of the pulse train.

This uniformity can be improved by using a second round of optimisation, this




FIGURE 6.13 – An example 41-pulse train, shown both without (**top**) and with (**bottom**) the effect of SLM spatial chirp.




FIGURE 6.14 – The pulse train shown in Figure 6.13, with measurements of the temporal chirp (df/dt) of each pulse.




FIGURE 6.15 – The same pulse train shown in Figure 6.14, after correcting for the SLM spatial chirp.

time using the more complete model. Since the mask is already close to an optimum, a simpler and less computationally intensive optimisation algorithm should be effective. This results in the pulse train shown in Figure 6.16. Note the significant improvement over the uncompensated pulse train shown in Figure 6.14, with quality comparable to the ideal case shown in the top of Figure 6.13.

6.3.5 Finite beam size

So far, the beam size at the SLM has been assumed to be zero, so that each frequency passes through exactly one pixel. Realistically, the beam profile at the SLM can be expected to be a Gaussian, whose size depends on the size of the beam incident on the first diffraction grating and on the focal length f . If the beam at the SLM is sufficiently close to the edge of a pixel, two different phase shifts will be applied, each to a different part of the beam. This spatially-varying phase means that the beam after the SLM is no longer a Gaussian; since each frequency component will be affected in a different way, this introduces a complex spatio-temporal coupling into the final output beam.




FIGURE 6.16 – The same pulse train shown in Figure 6.15, after another round of optimisation to improve uniformity.

Although each frequency component is not a Gaussian, it can be decomposed into Hermite-Gaussian modes. The higher-order modes will diffract more quickly than the zero-order mode (the Gaussian), so they can be filtered out by various apertures in the laser system. This is especially true in the case of a fibre laser, which might be used as a dedicated MP-LWFA driver, since only the zero-order mode will couple efficiently into the fibre. Therefore, it is a good approximation to ignore the higher-order modes, leaving only the Gaussian mode, and removing any spatio-temporal coupling. In this case, the only effect of the finite beam size is to convolve the spectral filter with a Gaussian, representing the beam profile.

This convolution can be accounted for straightforwardly in either of the two models discussed previously in this section. However, the resulting spectral filter will no longer be phase-only. In particular, wherever there is a significant phase difference between adjacent pixels, there will be a low-amplitude region in the spectral filter, the width of which depends on the beam size. One of the original aims of this method was to produce a filter which does not affect the spectral amplitude, due to concerns

about its effect on subsequent amplification. It is therefore important to take steps to minimise this effect. The simplest way of doing this is to minimise the beam size on the SLM, which can be done by increasing the beam size on the diffraction gratings. Simulations with the geometry described in Table 6.2 and a beam size of $w_0 = 10$ mm incident on the diffraction gratings do not have a significant effect on the spectrum.

Another method might be to minimise the phase shifts between adjacent pixels. In principle, this could be incorporated into the optimisation routines discussed earlier. I have not done this, since the beam sizes I expect to be used do not have a severe effect on the spectrum.

6.3.6 Non-uniform pulse spacings

This method can be extended to allow much finer control over pulse spacing, allowing the spacing between any pair of pulses to be altered. As described in § 1.2, this could be used to compensate for resonant detuning or to extract energy from the wakefield after acceleration. The basic principle is to double the period of the phase mask, $\Delta\omega$. This halves the pulse spacing, but since individual control over pulse amplitude is possible, every other pulse can be constrained to zero energy, restoring the original pulse spacing. This is achieved by a suitable redefinition of “on” and “off” pulses. Now, any pulse can be moved by half a period, by setting it to “off” and setting one of its neighbours to “on”. Once a suitable target pulse train has been constructed, the same steps used to optimise a uniform pulse train can be carried out. An example of the resulting pulse train, including the effects of SLM spatial chirp as described in § 6.3.4 is shown in Figure 6.17.

This method allows each pulse to be positioned with a precision of half a period. Greater precision could be achieved by increasing $\Delta\omega$ by a larger factor; for example, it could be increased by a factor of 4, then 3 out of every 4 pulses would be set to “off”, allowing quarter-period precision. By this point, looking at the examples above, one period of the phase mask would cover the entire SLM, so further improvements in




FIGURE 6.17 – A simulated pulse train with non-uniform spacing. This was generated by the same method used to generate the pulse train shown in Figure 6.16, and includes the same effects.

precision seem impossible. This is simply because the smallest time-scale over which the temporal profile can be controlled is related to the bandwidth of the spectral filter. The precision could be improved by using a larger SLM, although that would result in large parts of the SLM having almost no incident laser light.

6.3.7 Interaction with amplification

Since an LC-SLM based pulse shaper would necessarily be placed before the main laser amplifiers, it is important to consider how the pulse train generation would interact with amplification. In particular, it is possible that the process of amplification might affect the pulse train or the presence of a pulse train might affect amplification. I have not had time to study these effects in detail, but here I present an overview of some possible interactions. The amplification is assumed to take place after the pulse stretcher in a CPA system, so that the pulse train is chirped to much longer than its original duration.

As part of the amplification process, the pulse train will propagate a considerable distance through the gain medium. The chromatic dispersion caused by this should not be a problem, since it is a linear process which must be compensated by the compressor even in the case of a single pulse. There will also be various non-linear processes, such as self-phase modulation. These act in the time domain, depending on the temporal intensity profile. However, because the pulse is strongly chirped at this point, there is a direct correspondence between frequency and time; this means that it is possible to treat these non-linear processes as acting in the frequency domain and depending on the spectral amplitude. Provided the pulse train generation process does not affect the spectral amplitude, these non-linear processes should be limited to adding the same spectral phase to the pulse train that they would add to a single pulse. This is essentially a linear process, and will not affect the pulse train.

The amplification process itself may also affect the spectral amplitude, through processes such as gain narrowing. While the overall shape of the spectrum determines the shape of each identical pulse within the train, it does not affect the envelope of the pulse train, so this should not be a concern. Conversely, the amplification process may be affected by the spectral amplitude, especially in a heterogeneous gain medium; this is another reason to ensure that the filter does not change the spectral amplitude. The amplification process should not be sensitive to or affect the spectral phase, since CPA relies on the spectral phase added by the stretcher being preserved during amplification until it is removed by the compressor.

Finally, another concern is that earlier pulses in the train might be amplified more than later ones, as the energy available in the gain medium is depleted as the pulse train passes. However, the strong chirp applied to the pulse train by the stretcher ensures that each pulse will be stretched to a much longer duration than the total length of the pulse train; the constituent pulses of the train will merge together and be amplified simultaneously.

6.4 Summary

I have described three methods which could be used to generate a pulse train. The multiple Michelson method allows a certain degree of control, but is difficult to align and requires particularly thin beamsplitters if it to be used after the compressor. This latter problem might be ameliorated if it is placed before the compressor, which would also allow the use of smaller and thus cheaper diffraction gratings in the compressor itself. The single birefringent crystal method is much simpler to set up, but allows less control over the pulse separation and results in somewhat poorer visibility.

Much greater control, including per-pulse programmable spacing, can be achieved by the use of an SLM-based pulse shaper. I have described a process by which a phase-only spectral filter can be created which should generate a highly uniform pulse train using a realistic pulse shaper configuration. However, such a device would need to be placed before the main amplifiers when used in a laser sufficiently powerful to drive a useful wakefield, leaving open the question of how the pulse train will interact with the amplification process. I have addressed some possible problems, but ideally a more in-depth study is needed at some point in the future.

7. Conclusions and future work

The original aim of this project was to demonstrate the principle of multi-pulse laser wakefield acceleration (MP-LWFA), by exciting a wakefield using a train of laser pulses and measuring the amplitude of the resulting wakefield using frequency-domain holography (FDH). A necessary condition for this was that the diagnostic be sensitive enough to measure the wakefield resulting from a single pulse, before that pulse was split into a train. Although preliminary calculations suggested that this would be possible, later measurements showed that the available laser energy and the sensitivity of the FDH diagnostic were both less than expected. Since I was unable to reliably detect the wakefield due to a single laser pulse, it was impossible to proceed to using a train of pulses. Nevertheless, the reimplementing of FDH (believed to be the first one performed independently of the original developers) was successful, and the lessons learned over the course of the project will be valuable when it is repeated using a higher-energy laser.

In this chapter, I summarise the main conclusions of my project and outline suggestions for further work.

7.1 Conclusions

7.1.1 Experimental work

After setting up the experiment using the original design for the interaction point (IP), I was able to use the FDH diagnostic to image the ionisation front, as described in § 5.1.

However, the expected wakefield was not visible, with later calculations (see Table 4.5) giving an expected phase shift significantly less than the noise level of the FDH measurement. Furthermore, the background phase shift due to ionisation was much less than expected given the length of the gas cell and the input gas pressure. This fact, together with the very large value (> 20 radians) of the predicted background phase shift, suggested that either the pump or probe pulse was suffering from ionisation-induced defocussing, significantly reducing the interaction length. This was later supported by the particle-in-cell (PIC) simulation described in § 4.4.1.

To resolve this problem, I redesigned the IP as described in § 3.6.1, aiming for a much shorter interaction length so that the background phase shift would be smaller, thus reducing the impact of ionisation-induced defocussing. Simply reducing the interaction length would also reduce the phase shift caused by the wakefield, so to compensate I also reduced the laser spot size, increasing the amplitude of the wakefield. The fluid equations [26] I was using for this design process predicted an enhancement to the density oscillation as the transverse size of the laser pulse is reduced below its longitudinal size. This meant that smaller focal spot sizes would result in greater phase shifts due to the wakefield, despite the shorter Rayleigh length reducing the interaction length. I therefore designed the focal spot size to be as small as possible, without making it so small that the wakefield would become significantly non-linear — the fluid equations I was using only held for linear wakefields, and I expected a non-linear wakefield to be more difficult to measure.

After reconstructing the interaction region according to the new IP design, the measured focal spot size was found to be less than the design value. Whilst it had not been possible to correct the astigmatism present in the old IP, it was possible to correct any astigmatism in the new IP; the design had assumed that the level of astigmatism would remain the same, so correcting it reduced the effective focal spot size. This reduced the transverse size of the wakefield and pushed it further into the non-linear régime, with the result that the measured phase shift was expected to be very

sensitive to the alignment of the diagnostics, as described in § 4.7. Nevertheless, FDH measurements of the background phase shift (described in § 5.2) were consistent with the expected plasma density and interaction length, showing that the IP redesign had successfully eliminated the problem of refraction seen in the old IP.

The tight alignment constraints, together with significant beam jitter due to vibrations from the vacuum pump, made it difficult to observe a wakefield reliably. Some signals consistent with the presence of a wakefield were observed, as described in § 5.3, but they were too weak and intermittent to be useful in a demonstration of MP-LWFA. Nevertheless, the reimplementation of FDH was successful and many of the lessons learned from this experiment are applicable to future work towards the same goal, as described in § 7.2.

7.1.2 Pulse train generation

I also studied possible methods of pulse train generation. Firstly, I considered methods which might be used as part of an MP-LWFA demonstration experiment, based on splitting an intense pulse after compression. The first method considered was the multiple Michelson interferometer, described in § 6.1. One concern with this design was the effect of self phase modulation on the resulting pulse train, given the high input intensity and the required transmissive elements. My calculations showed that this could be avoided by detuning the pulse compressor, provided the beamsplitters were not too thick.

The second method considered was the single birefringent crystal, described in § 6.2. This has been demonstrated for uniform pulse trains, but the effect of resonant detuning (the increase of plasma wavelength with wakefield amplitude) may necessitate pulse trains with non-uniform spacing. I created a simulation of the process including the effects of higher-order phase, which demonstrated that such a pulse train could be created by proper control over higher order phase of the input pulse.

Beyond this, I considered how a pulse train might be generated in a dedicated MP-

LWFA driver laser, using an SLM-based pulse shaper (see § 6.3). Since the intention was to place the pulse shaper before the amplification stages in the laser, it was necessary to avoid any changes to the spectral amplitude. I developed an algorithm that can be used to create a phase filter which, applied to an SLM pulse shaper, results in a uniform train containing a chosen number of pulses. I improved this algorithm to account for the dependence of angular dispersion on wavelength, and to produce pulse trains with non-uniform spacing (again, to compensate for resonant detuning). This suggests that such a scheme is feasible, and the simulations I developed have reached the point that they could help to design an experiment to demonstrate it.

7.2 Future work

7.2.1 Further MP-LWFA experiments

Although I was not successful in my original aim to demonstrate MP-LWFA, the difficulties in the experiment could be overcome if more laser energy were available. In particular, this would allow the focal spot size to be increased, reducing the problems of a non-linear wakefield and loosening the alignment tolerances, whilst also increasing the total phase shift caused by the wakefield. Therefore, the obvious next step is to repeat the experiment using a more powerful laser. Such an experiment will benefit from the experience of this project, especially in the design of the interaction point and the diagnostics.

The original design of the IP relied on a long interaction length to make up for a small wakefield amplitude. The later design used a very small focal spot size to increase the total phase shift due to the 2D enhancement to the density oscillation. Both of these turned out to be fundamentally flawed approaches — the long interaction length could not be achieved in practice due to ionisation-induced defocussing, and the small focal spot resulted in a very small wakefield, making the diagnostic very sensitive to alignment. A future experiment of this kind must select a focal spot size that is large

enough to avoid alignment difficulties and remain in the linear régime — this depends on the apparatus used and the laser energy. At the same time, it is necessary to check that this will not result in significant ionisation-induced defocussing, which might be accomplished through PIC simulations (or potentially other, less computationally intensive methods).

Once an appropriate focal spot size has been selected, the phase shift depends on the laser pulse duration and energy. The required phase shift depends, in turn, on the level of noise in the FDH measurement. This was originally significantly underestimated during the planning of the experiment, but an accurate estimate is crucial for selecting the required laser pulse energy. The phase error in the frequency domain can be predicted from knowledge of the signal-to-noise ratio of the spectrometer, as described in § 2.5.2; provided the details of the FDH analysis process are properly accounted for, this leads to accurate values, as shown in § 2.5.3 and § 2.5.4. However, these results also show that the phase error in the time domain is significantly greater than in the frequency domain. Additionally, the apparent wakefield signals seen in § 5.3 are very unclear, despite the measured peak-to-peak phase shift being somewhat above the noise floor. Therefore, a significant safety margin between the predicted noise floor and the predicted phase shift is highly advisable.

A particular difficulty during the course of the experiment was the need to take reference data for use during FDH analysis, especially since the instability of the laser made it necessary to make these measurements often, and the layout of the experiment was such that these measurements required a significant amount of work. It would, therefore, be highly advisable for future experiments to take steps to improve this. The apparatus could be designed in such a way that reference data can be taken simultaneously, or nearly so, with experimental data. It might be possible to recover the necessary reference data directly from the experimental spectrum, as described in § 2.7, although this has not yet been demonstrated. Or the references could simply be omitted, resulting in a kind of “FDH-lite”; in addition to the generation of spectral

phase maps, the Fourier transform of the FDH hologram can, under certain conditions, be used to measure the amplitude and phase of the wakefield, as described in § 4.8. Using the Fourier transform as a diagnostic in this way is the subject of recent work by Nicholas Matlis, which he is preparing for publication [96].

7.2.2 FDH development

Despite the difficulties involved in implementing FDH, it has a lot of potential (together with its more sophisticated variants) as a useful diagnostic in laser wakefield experiments. However, there is some way to go before it can become a practical diagnostic, which could be used routinely on many experiments. The biggest obstacle to its wider adoption is probably the complexity involved in taking references. I have discussed possible solutions to this problem above. Although the best solution would appear to be to design the apparatus to allow references to be taken simultaneously with experimental data, that would not necessarily be very practical. The alternative is to reconstruct the reference data from the experimental data, which trades off reduced complexity of the apparatus for a significantly more complex analysis process. In either case, I believe that one of these approaches must be demonstrated before FDH can see more widespread adoption, the latter approach being more challenging but also potentially more useful.

Another obstacle is the need to completely understand and reimplement the analysis process. Compare this to the process of analysing FROG data [76] which is widely used, despite its complexity. This is perhaps due to the ready availability of “black-box” software to perform the analysis, as well as self-contained commercial devices. Although it would be wildly optimistic to predict that a completely self-contained FDH system could be manufactured, the availability of standard and easy-to-use FDH analysis software would be an important step towards the widespread adoption of FDH. In addition to this, there are many subtleties involved in the recording and analysis of FDH data, some of which are described in § 2.6. The clear understanding and doc-

umentation of these points would be another useful step, as would investigation of other matters, such as the relative importance of the different pieces of input data, which qualities of that input data are most important, and how to recognise common problems.

7.2.3 Laser development and pulse train generation

In the longer term, the aim is to develop MP-LWFA to the point where it can be used to develop a compact light source. This will require a compact, efficient, high repetition rate driver laser. Fibre lasers show a great deal of promise in this respect, being capable of high power efficiency, MHz repetition rates, and average powers of 100's of watts [40]. Although they are not yet capable of the peak powers required to excite a useful wakefield, this limitation could be overcome through coherent combination of the output of multiple amplifiers [47].

A dedicated MP-LWFA driver laser will also require a method of generating a pulse train, such as the one outlined in § 6.3. Complex pulse shaping is not usually performed with high-energy pulses, so further work is needed to show not only that this method can be used to generate a pulse train, but also that such a pulse train can be amplified to high energy without undergoing distortion. This should, eventually, be followed by the construction and testing of a full-scale demonstration driver laser, capable of emitting pulse trains with sufficient energy to accelerate electrons to several hundred MeV at a sufficient repetition rate to construct a useful compact light source.

Bibliography

- [1] W. C. Röntgen. “On a new kind of rays”. *Nature* **53** (1369) 274–276 (Jan. 1896).
- [2] J. J. Thomson. “Cathode rays”.
Philosophical Magazine **44** (269) 293–316 (Oct. 1897).
doi: 10.1080/14786449708621070.
- [3] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”.
Physics Letters B **716** (1) 1–29 (2012).
doi: 10.1016/j.physletb.2012.08.020.
- [4] Goedele N. Maertens, Stephen Hare, and Peter Cherepanov. “The mechanism of retroviral integration from x-ray structures of its key intermediates”.
Nature **468** (7321) 326–329 (Nov. 2010). doi: 10.1038/nature09517.
- [5] Alenka Možir *et al.* “A study of degradation of historic parchment using small-angle x-ray scattering, synchrotron-IR and multivariate data analysis”.
Analytical and Bioanalytical Chemistry **402** (4) 1559–1566 (2012).
doi: 10.1007/s00216-011-5392-6.
- [6] E. Esarey, C. B. Schroeder, and W. P. Leemans.
“Physics of laser-driven plasma-based electron accelerators”.
Rev. Mod. Phys. **81** (3) 1229–1285 (Aug. 2009).
doi: 10.1103/RevModPhys.81.1229.
- [7] J. A. Bittencourt. *Fundamentals of Plasma Physics*. 3rd ed.
Springer Science & Business Media, 2013. ISBN: 1475740301.
- [8] Ian Blumenfeld *et al.* “Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator”. *Nature* **445** (7129) 741–744 (Feb. 2007).
doi: 10.1038/nature05538.
- [9] T. Tajima and J. M. Dawson. “Laser electron accelerator”.
Phys. Rev. Lett. **43** (4) 267–270 (July 1979).
doi: 10.1103/PhysRevLett.43.267.
- [10] C. Joshi *et al.* “Forward Raman instability and electron acceleration”.
Phys. Rev. Lett. **47** (18) 1285–1288 (Nov. 1981).
doi: 10.1103/PhysRevLett.47.1285.
- [11] K. Nakajima *et al.* “Observation of ultrahigh gradient electron acceleration by a self-modulated intense short laser pulse”.
Phys. Rev. Lett. **74** (22) 4428–4431 (May 1995).
doi: 10.1103/PhysRevLett.74.4428.

- [12] C. E. Clayton *et al.* “Ultrahigh-gradient acceleration of injected electrons by laser-excited relativistic electron plasma waves”.
Phys. Rev. Lett. **70** (1) 37–40 (Jan. 1993). DOI: 10.1103/PhysRevLett.70.37.
- [13] A. Modena *et al.*
“Electron acceleration from the breaking of relativistic plasma waves”.
Nature **377** (6550) 606–608 (Oct. 1995). DOI: 10.1038/377606a0.
- [14] V. Malka *et al.* “Electron acceleration by a wake field forced by an intense ultrashort laser pulse”. *Science* **298** (5598) 1596–1600 (2002).
DOI: 10.1126/science.1076782.
- [15] S. P. D. Mangles *et al.* “Monoenergetic beams of relativistic electrons from intense laser-plasma interactions”. *Nature* **431** (7008) 535–538 (Sept. 2004).
DOI: 10.1038/nature02939.
- [16] J. Faure *et al.*
“A laser-plasma accelerator producing monoenergetic electron beams”.
Nature **431** (7008) 541–544 (Sept. 2004). DOI: 10.1038/nature02963.
- [17] C. G. R. Geddes *et al.* “High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding”.
Nature **431** (7008) 538–541 (Sept. 2004). DOI: 10.1038/nature02900.
- [18] W. P. Leemans *et al.*
“GeV electron beams from a centimetre-scale accelerator”.
Nat. Phys. **2** (10) 696–699 (Oct. 2006). DOI: 10.1038/nphys418.
- [19] Hyung Taek Kim *et al.*
“Enhancement of electron energy to the multi-GeV regime by a dual-stage laser-wakefield accelerator pumped by petawatt laser pulses”.
Phys. Rev. Lett. **111** (16) 165002 (Oct. 2013).
DOI: 10.1103/PhysRevLett.111.165002.
- [20] W. P. Leemans *et al.*
“Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime”.
Phys. Rev. Lett. **113** (24) 245002 (Dec. 2014).
DOI: 10.1103/PhysRevLett.113.245002.
- [21] A. J. Gonsalves *et al.*
“Tunable laser plasma accelerator based on longitudinal density tailoring”.
Nat. Phys. **7** (11) 862–866 (Nov. 2011). DOI: 10.1038/nphys2071.
- [22] K. Schmid *et al.* “Density-transition based electron injector for laser driven wakefield accelerators”. *Phys. Rev. ST Accel. Beams* **13** (9) 091301 (Sept. 2010).
DOI: 10.1103/PhysRevSTAB.13.091301.
- [23] G. Golovin *et al.* “Tunable monoenergetic electron beams from independently controllable laser-wakefield acceleration and injection”.
Phys. Rev. ST Accel. Beams **18** (1) 011301 (Jan. 2015).
DOI: 10.1103/PhysRevSTAB.18.011301.

- [24] M. Hansson *et al.* “Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator”. *Phys. Rev. ST Accel. Beams* **18** (7) 071303 (July 2015).
DOI: 10.1103/PhysRevSTAB.18.071303.
- [25] L.-L. Yu *et al.* “Two-color laser-ionization injection”. *Phys. Rev. Lett.* **112** (12) 125001 (Mar. 2014).
DOI: 10.1103/PhysRevLett.112.125001.
- [26] C. W. Siders *et al.* “Plasma-based accelerator diagnostics based upon longitudinal interferometry with ultrashort optical pulses”. *IEEE Transactions on Plasma Science* **24** (2) 301–315 (1996).
DOI: 10.1109/27.509994.
- [27] A. Pukhov and J. Meyer-ter Vehn.
“Laser wake field acceleration: the highly non-linear broken-wave regime”. *Applied Physics B* **74** (4–5) 355–361 (2002). DOI: 10.1007/s003400200795.
- [28] F. Grüner *et al.* “Design considerations for table-top, laser-based VUV and X-ray free electron lasers”. *Applied Physics B* **86** (3) 431–435 (2007).
DOI: 10.1007/s00340-006-2565-7.
- [29] Matthias Fuchs *et al.* “Laser-driven soft-X-ray undulator source”. *Nat. Phys.* **5** (11) 826–829 (Nov. 2009). DOI: 10.1038/nphys1404.
- [30] S. Kneip *et al.* “X-ray phase contrast imaging of biological specimens with femtosecond pulses of betatron radiation from a compact laser plasma wakefield accelerator”. *Applied Physics Letters* **99** (9) 093701 (2011).
DOI: 10.1063/1.3627216.
- [31] Zhirong Huang, Yuantao Ding, and Carl B. Schroeder.
“Compact X-ray free-electron laser from a laser-plasma accelerator using a transverse-gradient undulator”. *Phys. Rev. Lett.* **109** (20) 204801 (Nov. 2012).
DOI: 10.1103/PhysRevLett.109.204801.
- [32] N. D. Powers *et al.* “Quasi-monoenergetic and tunable X-rays from a laser-driven Compton light source”. *Nat. Photon.* **8** (1) 28–31 (Jan. 2014).
DOI: 10.1038/nphoton.2013.314.
- [33] S M Hooker *et al.*
“Multi-pulse laser wakefield acceleration: a new route to efficient, high-repetition-rate plasma accelerators and high flux radiation sources”. *J. Phys. B* **47** (23) 234003 (2014).
URL: <http://stacks.iop.org/0953-4075/47/i=23/a=234003>.
- [34] A. Popp *et al.*
“All-optical steering of laser-wakefield-accelerated electron beams”. *Phys. Rev. Lett.* **105** (21) 215001 (Nov. 2010).
DOI: 10.1103/PhysRevLett.105.215001.
- [35] W. J. Brown *et al.*
“Experimental characterization of an ultrafast Thomson scattering x-ray source with three-dimensional time and frequency-domain analysis”. *Phys. Rev. ST Accel. Beams* **7** (6) 060702 (June 2004).
DOI: 10.1103/PhysRevSTAB.7.060702.

- [36] W. S. Graves *et al.* “MIT inverse Compton source concept”. *Nuclear Instruments and Methods in Physics Research Section A* **608** S103–S105 (2009).
doi: 10.1016/j.nima.2009.05.042.
- [37] W. S. Graves *et al.* “Compact x-ray source based on burst-mode inverse Compton scattering at 100 kHz”.
Phys. Rev. ST Accel. Beams **17** (12) 120701 (Dec. 2014).
doi: 10.1103/PhysRevSTAB.17.120701.
- [38] F. Albert *et al.* “Characterization and applications of a tunable, laser-based, MeV-class Compton-scattering gamma-ray source”.
Phys. Rev. ST Accel. Beams **13** (7) 070704 (July 2010).
doi: 10.1103/PhysRevSTAB.13.070704.
- [39] Cesar Jauregui, Jens Limpert, and Andreas Tunnermann.
“High-power fibre lasers”. *Nat. Photon.* **7** (11) 861–867 (Nov. 2013).
doi: 10.1038/nphoton.2013.273.
- [40] Andreas Tünnermann, Thomas Schreiber, and Jens Limpert.
“Fiber lasers and amplifiers: an ultrafast performance evolution”.
Appl. Opt. **49** (25) F71–F78 (Sept. 2010). doi: 10.1364/AO.49.000F71.
- [41] Jonathan C. Knight. “Photonic crystal fibres”.
Nature **424** (6950) 847–851 (Aug. 2003). doi: 10.1038/nature01940.
- [42] J. Limpert *et al.* “Extended single-mode photonic crystal fiber lasers”.
Opt. Express **14** (7) 2715–2720 (Apr. 2006). doi: 10.1364/OE.14.002715.
- [43] J. C. Knight. “Photonic crystal fibers and fiber lasers”.
J. Opt. Soc. Am. B **24** (8) 1661–1668 (Aug. 2007).
doi: 10.1364/JOSAB.24.001661.
- [44] Enrico Seise *et al.*
“Coherent addition of fiber-amplified ultrashort laser pulses”.
Opt. Express **18** (26) 27827–27835 (Dec. 2010). doi: 10.1364/OE.18.027827.
- [45] L. Daniault *et al.* “High peak-power stretcher-free femtosecond fiber amplifier using passive spatio-temporal coherent combining”.
Opt. Express **20** (19) 21627–21634 (Sept. 2012). doi: 10.1364/OE.20.021627.
- [46] Arno Klenke *et al.* “530 W, 1.3 mJ, four-channel coherently combined femtosecond fiber chirped-pulse amplification system”.
Opt. Lett. **38** (13) 2283–2285 (July 2013). doi: 10.1364/OL.38.002283.
- [47] Arno Klenke *et al.*
“22 GW peak-power fiber chirped-pulse-amplification system”.
Opt. Lett. **39** (24) 6875–6878 (Dec. 2014). doi: 10.1364/OL.39.006875.
- [48] K. Nakajima. “Plasma-wave resonator for particle-beam acceleration”.
Phys. Rev. A **45** (2) 1149–1156 (Jan. 1992). doi: 10.1103/PhysRevA.45.1149.
- [49] D A Johnson *et al.* “Plasma wakefield generation by multiple short pulses”.
Physica Scripta **T52** 77–81 (1994).
URL: <http://stacks.iop.org/1402-4896/1994/i=T52/a=013>.

- [50] D. Umstadter, E. Esarey, and J. Kim.
“Nonlinear plasma waves resonantly driven by optimized laser pulse trains”.
Phys. Rev. Lett. **72** (8) 1224–1227 (Feb. 1994).
doi: 10.1103/PhysRevLett.72.1224.
- [51] S. Dalla and M. Lontano. “Large amplitude plasma wave excitation by means of sequences of short laser pulses”.
Phys. Rev. E **49** (3) R1819–R1822 (Mar. 1994).
doi: 10.1103/PhysRevE.49.R1819.
- [52] Guy Bonnaud, Denis Teychenné, and Jean-Louis Bobin.
“Wake-field effect induced by laser multiple pulses”.
Phys. Rev. E **50** (1) R36–R39 (July 1994). doi: 10.1103/PhysRevE.50.R36.
- [53] D. Umstadter *et al.*
“Resonantly laser-driven plasma waves for electron acceleration”.
Phys. Rev. E **51** (4) 3484–3497 (Apr. 1995). doi: 10.1103/PhysRevE.51.3484.
- [54] C. B. Schroeder *et al.*
“Physics considerations for laser-plasma linear colliders”.
Phys. Rev. ST Accel. Beams **13** (10) 101301 (Oct. 2010).
doi: 10.1103/PhysRevSTAB.13.101301.
- [55] N. Bourgeois, J. Cowley, and S. M. Hooker.
“Two-pulse ionization injection into quasilinear laser wakefields”.
Phys. Rev. Lett. **111** (15) 155004 (Oct. 2013).
doi: 10.1103/PhysRevLett.111.155004.
- [56] C. Benedetti *et al.* “Numerical investigation of electron self-injection in the nonlinear bubble regime”. *Physics of Plasmas* **20** (10) 103108 (2013).
doi: 10.1063/1.4824811.
- [57] A. Pak *et al.* “Injection and trapping of tunnel-ionized electrons into laser-produced wakes”. *Phys. Rev. Lett.* **104** (2) 025003 (Jan. 2010).
doi: 10.1103/PhysRevLett.104.025003.
- [58] C. McGuffey *et al.*
“Ionization induced trapping in a laser wakefield accelerator”.
Phys. Rev. Lett. **104** (2) 025004 (Jan. 2010).
doi: 10.1103/PhysRevLett.104.025004.
- [59] C. E. Clayton *et al.* “Self-guided laser wakefield acceleration beyond 1 GeV using ionization-induced injection”.
Phys. Rev. Lett. **105** (10) 105003 (Sept. 2010).
doi: 10.1103/PhysRevLett.105.105003.
- [60] B. B. Pollock *et al.* “Demonstration of a narrow energy spread, ~ 0.5 GeV electron beam from a two-stage laser wakefield accelerator”.
Phys. Rev. Lett. **107** (4) 045001 (July 2011).
doi: 10.1103/PhysRevLett.107.045001.
- [61] E. Esarey *et al.*
“Electron injection into plasma wakefields by colliding laser pulses”.
Phys. Rev. Lett. **79** (14) 2682–2685 (Oct. 1997).
doi: 10.1103/PhysRevLett.79.2682.

- [62] M. Chen *et al.* “Electron injection and emittance control by transverse colliding pulses in a laser-plasma accelerator”.
Phys. Rev. ST Accel. Beams **17** (5) 051303 (May 2014).
doi: 10.1103/PhysRevSTAB.17.051303.
- [63] C. G. R. Geddes *et al.* “Plasma-density-gradient injection of low absolute-momentum-spread electron bunches”.
Phys. Rev. Lett. **100** (21) 215004 (May 2008).
doi: 10.1103/PhysRevLett.100.215004.
- [64] Alexander Buck *et al.*
“Real-time observation of laser-driven electron acceleration”.
Nat. Phys. **7** (7) 543–548 (July 2011). doi: 10.1038/nphys1942.
- [65] M. B. Schwab *et al.* “Few-cycle optical probe-pulse for investigation of relativistic laser-plasma interactions”.
Applied Physics Letters **103** (19) 191118 (2013). doi: 10.1063/1.4829489.
- [66] S. P. Le Blanc *et al.* “Single-shot measurement of temporal phase shifts by frequency-domain holography”. *Opt. Lett.* **25** (10) 764–766 (May 2000).
doi: 10.1364/OL.25.000764.
- [67] N. H. Matlis *et al.* “Snapshots of laser wakefields”.
Nat. Phys. **2** (11) 749–753 (Nov. 2006). doi: 10.1038/nphys442.
- [68] P Dong *et al.* “Holographic visualization of laser wakefields”.
New Journal of Physics **12** (4) 045016 (2010).
url: <http://stacks.iop.org/1367-2630/12/i=4/a=045016>.
- [69] Zhengyan Li *et al.* “Single-shot visualization of evolving laser wakefields using an all-optical streak camera”. *Phys. Rev. Lett.* **113** (8) 085001 (Aug. 2014).
doi: 10.1103/PhysRevLett.113.085001.
- [70] Zhengyan Li *et al.*
“Single-shot tomographic movies of evolving light-velocity objects”.
Nat. Commun. **5** (Jan. 2014). doi: 10.1038/ncomms4085.
- [71] J. P. Geindre *et al.*
“Frequency-domain interferometer for measuring the phase and amplitude of a femtosecond pulse probing a laser-produced plasma”.
Opt. Lett. **19** (23) 1997–1999 (Dec. 1994). doi: 10.1364/OL.19.001997.
- [72] C. W. Siders *et al.* “Laser wakefield excitation and measurement by femtosecond longitudinal interferometry”.
Phys. Rev. Lett. **76** (19) 3570–3573 (May 1996).
doi: 10.1103/PhysRevLett.76.3570.
- [73] Peng Dong *et al.*
“Formation of optical bullets in laser-driven plasma bubble accelerators”.
Phys. Rev. Lett. **104** (13) 134801 (Mar. 2010).
doi: 10.1103/PhysRevLett.104.134801.
- [74] Z. Li *et al.* “Frequency-domain streak camera for ultrafast imaging of evolving light-velocity objects”. *Opt. Lett.* **35** (24) 4087–4089 (Dec. 2010).
doi: 10.1364/OL.35.004087.

- [75] D. J. Kane and R. Trebino. “Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating”. *IEEE J. Quantum Electron.* **29** (2) 571–579 (Feb. 1993). DOI: 10.1109/3.199311.
- [76] Rick Trebino and Daniel J. Kane. “Using phase retrieval to measure the intensity and phase of ultrashort pulses: frequency-resolved optical gating”. *J. Opt. Soc. Am. A* **10** (5) 1101–1111 (May 1993). DOI: 10.1364/JOSAA.10.001101.
- [77] C. Iaconis and I. A. Walmsley. “Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses”. *Opt. Lett.* **23** (10) 792–794 (May 1998). DOI: 10.1364/OL.23.000792.
- [78] M. Galassi *et al.* *GNU Scientific Library Reference Manual*. 3rd ed. Jan. 2009. ISBN: 0954612078. URL: http://www.gnu.org/software/gsl/manual/html_node/.
- [79] Dennis C. Ghiglie and Mark D. Pritt. *Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software*. Wiley-Interscience, 1998. ISBN: 0471249351.
- [80] Richard M. Goldstein, Howard A. Zebker, and Charles L. Werner. “Satellite radar interferometry: two-dimensional phase unwrapping”. *Radio Science* **23** (4) 713–720 (1988). DOI: 10.1029/RS023i004p00713.
- [81] James M. Supplee, Edward A. Whittaker, and Wilfried Lenth. “Theoretical description of frequency modulation and wavelength modulation spectroscopy”. *Appl. Opt.* **33** (27) 6294–6302 (Sept. 1994). DOI: 10.1364/AO.33.006294.
- [82] Donna Strickland and Gerard Mourou. “Compression of amplified chirped optical pulses”. *Opt. Commun.* **56** (3) 219–221 (Dec. 1985). URL: <http://www.sciencedirect.com/science/article/pii/0030401885901208>.
- [83] Ian Walmsley, Leon Waxer, and Christophe Dorrer. “The role of dispersion in ultrafast optics”. *Review of Scientific Instruments* **72** (1) 1–29 (2001). DOI: 10.1063/1.1330575.
- [84] M. Raghuramaiah *et al.* “Simultaneous measurement of pulse-front tilt and pulse duration of a femtosecond laser beam”. *Optics Communications* **223** 163–168 (2003). DOI: 10.1016/S0030-4018(03)01638-9.
- [85] J. R. Fienup. “Phase retrieval algorithms: a comparison”. *Appl. Opt.* **21** (15) 2758–2769 (Aug. 1982). DOI: 10.1364/AO.21.002758.
- [86] S. I. Bajlekov *et al.* “Longitudinal electron bunch profile reconstruction by performing phase retrieval on coherent transition radiation spectra”. *Phys. Rev. ST Accel. Beams* **16** (4) 040701 (Apr. 2013). DOI: 10.1103/PhysRevSTAB.16.040701.

- [87] Patrick O'Shea *et al.*
 "Highly simplified device for ultrashort-pulse measurement".
Opt. Lett. **26** (12) 932–934 (June 2001). DOI: 10.1364/OL.26.000932.
- [88] SCHOTT AG. *Schott Optical Glass Data Sheets*. 2014.
 URL: http://www.schott.com/advanced_optics/english/abbe_datasheets/schott-datasheet-all-english.pdf.
- [89] Biswajit Pathak and Bosanta R Boruah. "Improved wavefront reconstruction algorithm for Shack-Hartmann type wavefront sensors".
Journal of Optics **16** (5) 055403 (2014).
 URL: <http://stacks.iop.org/2040-8986/16/i=5/a=055403>.
- [90] R. W. Gerchberg and W. O. Saxton. *Optik* **35** 237 (1972).
- [91] Simon Hooker and Colin Webb. *Laser Physics*. Oxford University Press, 2010.
 ISBN: 978-0-19-850692-8.
- [92] R. Rankin *et al.* "Refraction effects associated with multiphoton ionization and ultrashort-pulse laser propagation in plasma waveguides".
Opt. Lett. **16** (11) 835–837 (June 1991). DOI: 10.1364/OL.16.000835.
- [93] P. Chessa *et al.* "Temporal and angular resolution of the ionization-induced refraction of a short laser pulse in helium gas".
Phys. Rev. Lett. **82** (3) 552–555 (Jan. 1999).
 DOI: 10.1103/PhysRevLett.82.552.
- [94] M. Czerny and A. F. Turner.
 "Über den Astigmatismus bei Spiegelspektrometern".
Zeitschrift für Physik **61** (11–12) 792–797 (1930). DOI: 10.1007/BF01340206.
- [95] Alistair Lawrence-Douglas.
 "Ionisation Effects for Laser-Plasma Interactions by Particle-in-Cell Code".
 PhD thesis. University of Warwick, Sept. 2013.
 URL: http://wrap.warwick.ac.uk/57465/1/WRAP_THESIS_Lawrence-Douglas_2013.pdf.
- [96] Nicholas Matlis. Private communication. 2015.
- [97] Albert H. Nuttall. "Some windows with very good sidelobe behavior".
IEEE Transactions on Acoustics, Speech and Signal Processing **29** (1) 84–91 (Feb. 1981). DOI: 10.1109/TASSP.1981.1163506.
- [98] Xiaoshi Zhang *et al.* "Quasi-phase-matching and quantum-path control of high-harmonic generation using counterpropagating light".
Nat. Phys. **3** (4) 270–275 (Apr. 2007). DOI: 10.1038/nphys541.
- [99] Warren S. Warren, Herschel Rabitz, and Mohammed Dahleh.
 "Coherent control of quantum dynamics: the dream is alive".
Science **259** (5101) 1581–1589 (1993).
 DOI: 10.1126/science.259.5101.1581.
- [100] M. Boscolo *et al.*
 "Generation of short THz bunch trains in a RF photoinjector". *Nuclear Instruments and Methods in Physics Research Section A* **577** (3) 409–416 (2007).
 DOI: 10.1016/j.nima.2007.04.129.

- [101] Kazi Sarwar Abedin, Noriaki Onodera, and Masaharu Hyodo.
“Repetition-rate multiplication in actively mode-locked fiber lasers by higher-order FM mode locking using a high-finesse Fabry-Perot filter”.
Applied Physics Letters **73** (10) 1311–1313 (1998). DOI: 10.1063/1.121879.
- [102] Craig W. Siders *et al.* “Efficient high-energy pulse-train generation using a 2^n -pulse michelson interferometer”. *Appl. Opt.* **37** (22) 5302–5305 (Aug. 1998).
DOI: 10.1364/AO.37.005302.
- [103] Shengguang Liu and Yen-Chieh Huang. “Generation of pre-bunched electron beams in photocathode RF gun for THz-FEL superradiation”. *Nuclear Instruments and Methods in Physics Research Section A* **637** S172–S176 (2011).
DOI: 10.1016/j.nima.2010.02.050.
- [104] Govind Agrawal. *Nonlinear Fiber Optics*. 5th ed. Academic Press, 2013.
ISBN: 978-0-12-397023-7.
- [105] I. H. Malitson.
“Interspecimen comparison of the refractive index of fused silica”.
J. Opt. Soc. Am. **55** (10) 1205–1209 (Oct. 1965).
DOI: 10.1364/JOSA.55.001205.
- [106] David Milam. “Review and assessment of measured values of the nonlinear refractive-index coefficient of fused silica”.
Appl. Opt. **37** (3) 546–550 (Jan. 1998). DOI: 10.1364/AO.37.000546.
- [107] A. Aryshev *et al.*
Femtosecond response time measurements of a Cs₂Te photocathode. July 2015.
arXiv: 1507.03302 [physics.acc-ph].
- [108] B. Dromey *et al.* “Generation of a train of ultrashort pulses from a compact birefringent crystal array”. *Appl. Opt.* **46** (22) 5142–5146 (Aug. 2007).
DOI: 10.1364/AO.46.005142.
- [109] Tom Robinson *et al.* “Generation and control of ultrafast pulse trains for quasi-phase-matching high-harmonic generation”.
J. Opt. Soc. Am. B **27** (4) 763–772 (Apr. 2010).
DOI: 10.1364/JOSAB.27.000763.
- [110] Ryuzi Yano and Hideki Gotoh. “Tunable terahertz electromagnetic wave generation using birefringent crystal and grating pair”.
Japanese Journal of Applied Physics **44** 8470–8473 (2005).
URL: <http://stacks.iop.org/1347-4065/44/i=12R/a=8470>.
- [111] T. Robinson *et al.*
“Simple technique for generating trains of ultrashort pulses”.
Opt. Lett. **32** (15) 2203–2205 (Aug. 2007). DOI: 10.1364/OL.32.002203.
- [112] Roman Walczak. Private communication. 2015.
- [113] Gorachand Ghosh. “Dispersion-equation coefficients for the refractive index and birefringence of calcite and quartz crystals”.
Optics Communications **163** 95–102 (1999).
DOI: 10.1016/S0030-4018(99)00091-7.

- [114] Pierre Tournois. “Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems”. *Optics Communications* **140** (4–6) 245–249 (1997). DOI: 10.1016/S0030-4018(97)00153-3.
- [115] Antoine Monmayrant, Sébastien Weber, and Béatrice Chatel. “A newcomer’s guide to ultrashort pulse shaping and characterization”. *Journal of Physics B: Atomic, Molecular and Optical Physics* **43** (10) 103001 (2010). URL: <http://stacks.iop.org/0953-4075/43/i=10/a=103001>.
- [116] Andrew M. Weiner. “Ultrafast optical pulse shaping: A tutorial review”. *Optics Communications* **284** (15) 3669–3692 (2011). DOI: 10.1016/j.optcom.2011.03.084.
- [117] A. M. Weiner *et al.* “Programmable femtosecond pulse shaping by use of a multielement liquid-crystal phase modulator”. *Opt. Lett.* **15** (6) 326–328 (Mar. 1990). DOI: 10.1364/OL.15.000326.
- [118] H. Wang *et al.* “20-fs pulse shaping with a 512-element phase-only liquid crystal modulator”. *IEEE J. Sel. Top. Quantum Electron.* **7** (4) 718–727 (July 2001). DOI: 10.1109/2944.974244.
- [119] A. M. Weiner and D. E. Leaird. “Generation of terahertz-rate trains of femtosecond pulses by phase-only filtering”. *Opt. Lett.* **15** (1) 51–53 (Jan. 1990). DOI: 10.1364/OL.15.000051.
- [120] H. Dammann and K. Görtler. “High-efficiency in-line multiple imaging by means of multiple phase holograms”. *Optics Communications* **3** (5) 312–315 (1971). DOI: 10.1016/0030-4018(71)90095-2.
- [121] J. Jahns *et al.* “Dammann gratings for laser beam shaping”. *Optical Engineering* **28** (12) 1267–1275 (1989). DOI: 10.1117/12.7977142.
- [122] Dmitry Pestov, Vadim V. Lozovoy, and Marcos Dantus. “Multiple independent comb shaping (MICS): phase-only generation of optical pulse sequences”. *Opt. Express* **17** (16) 14351–14361 (Aug. 2009). DOI: 10.1364/OE.17.014351.
- [123] W H Press *et al.* *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 2007. ISBN: 978-0-521-88068-8.
- [124] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1996. ISBN: 9780585030944.
- [125] R. Paschotta *et al.* “Ytterbium-doped fiber amplifiers”. *IEEE J. Quantum Electron.* **33** (7) 1049–1056 (July 1997). DOI: 10.1109/3.594865.