

High Average Power Plasma Wakefield Acceleration and Dark Photon Searches

Adam Scaachi

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
of
University College London.

Department of Physics and Astronomy
University College London

November 18, 2025

I, Adam Scaachi, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

Particle accelerators have played an instrumental role in the creation of the Standard Model of particle physics. However, despite its success there remains to be physical phenomena in nature that the Standard Model does not explain. The need to search for physics beyond the Standard Model motivates the development of improved acceleration techniques. One such technique is plasma wakefield acceleration, where the acceleration gradient is induced by the electric fields within a plasma. The application of this technique towards particle physics experimentation requires that the plasma wakefield acceleration process can be continuously repeated, such that the acceleration signature of one event is identical to the next. The time taken for the system to return to its original state after an acceleration event determines the repetition rate of the accelerator. Measurements of the repetition rate of a plasma wakefield accelerator under a variety of plasma conditions are presented in this thesis. The application of such an acceleration scheme to search for the dark photon, a hidden sector particle beyond the Standard Model, is also investigated.

Impact Statement

Plasma wakefield acceleration has proven to be an effective technique for accelerating electrons to high energies. However, questions remain regarding the achievable repetition rate of such a scheme, as well as what applications it may have towards enhancing high energy physics research. This thesis addresses the refill of material in the capillary occurring on the millisecond timescale, thereby complementing existing work of plasma wakefield evolution on the nanosecond and microsecond timescales. As such, it contributes towards a more complete understanding of the limitations on the repetition rate of plasma wakefield acceleration, and has an impact on various fields including both particle physics and medical physics.

Also addressed is the use of plasma wakefield acceleration to search for a hypothetical hidden sector particle known as the dark photon. The use of plasma wakefield acceleration allows for currently unexplored regions of the parameter space to be probed, contributing towards filling the gap that existing experiments have so far been unable to reach. Discovery of the dark photon would provide evidence of physics beyond the Standard Model, representing a major breakthrough in our fundamental understanding of the universe.

The role taken by the author in this research includes contributing towards the collection of data at FLASHForward for the repetition rate study, as well as fully conducting the subsequent analysis of these datasets. In addition, it includes developing simulations to model the gas flow through the FLASHForward capillary, and developing simulations to optimise the configuration of a dark photon search

experiment and report its sensitivity.

The work contained in this thesis has been used in the following publications:

- A. Caldwell *et al.*, “Particle physics applications of the AWAKE acceleration scheme,” *arXiv:1812.11164*, 2018.
- E. Gschwendtner *et al.*, “The AWAKE Run 2 programme and beyond,” *Symmetry*, vol. 14, no. 8, p. 1680, 2022.

Acknowledgements

I would first like to thank Matthew for his supervision. I also wish to thank Tony, James, and Trupen for the help they provided me with throughout my research, as well as members of the AWAKE-UCL meetings for their useful feedback. Finally, I would like to thank my family for their unconditional support.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 22 |
| 1.1 | Conventional Acceleration | 22 |
| 1.2 | Plasma Wakefield Acceleration | 24 |
| 1.2.1 | Laser-Driven Plasma Wakefield Acceleration | 25 |
| 1.2.2 | Beam-Driven Plasma Wakefield Acceleration | 26 |
| 1.2.3 | Proton-Driven Plasma Wakefield Acceleration | 27 |
| 1.3 | Outline | 28 |
| 2 | The FLASHForward and AWAKE Experiments | 29 |
| 2.1 | FLASHForward | 29 |
| 2.1.1 | Experimental Facility | 30 |
| 2.1.2 | Internal Injection (X-1) | 31 |
| 2.1.3 | External Injection (X-2) | 34 |
| 2.1.4 | High Average Power (X-3) | 37 |
| 2.2 | AWAKE | 42 |
| 2.2.1 | Experimental Facility | 43 |
| 2.2.2 | Self-modulation | 44 |
| 2.2.3 | Electron Acceleration | 46 |
| 2.2.4 | Run 2 Programme | 48 |
| 2.2.5 | Particle Physics Applications | 50 |
| 3 | Repetition Rate of Plasma Wakefield Acceleration | 52 |
| 3.1 | Motivation | 52 |

| | | |
|----------|--|------------|
| 3.2 | Experimental Method | 53 |
| 3.3 | Spectrometer Image Analysis | 55 |
| 3.3.1 | Pixel Threshold Value | 55 |
| 3.3.2 | Connected Component Analysis | 58 |
| 3.3.3 | Light Leakage | 64 |
| 3.3.4 | Bunch Energies | 67 |
| 3.3.5 | Reset Signature | 69 |
| 3.4 | Results | 73 |
| 3.5 | Discussion | 87 |
| 4 | Simulation of Gas Filling in the FLASHForward Capillary | 91 |
| 4.1 | Motivation | 91 |
| 4.2 | Simulation | 92 |
| 4.2.1 | Theory | 92 |
| 4.2.2 | Configuration | 95 |
| 4.2.3 | Output | 95 |
| 4.3 | Gas Filling in the FLASHForward Capillary | 97 |
| 4.3.1 | Flow Rate Calibration | 98 |
| 4.3.2 | Gas Filling Curves | 100 |
| 4.4 | Discussion | 100 |
| 5 | Design of a Dark Photon Search Experiment | 104 |
| 5.1 | Dark Photons | 104 |
| 5.1.1 | Portal Interaction | 105 |
| 5.1.2 | Production & Decay | 106 |
| 5.1.3 | Current Limits | 107 |
| 5.1.4 | Proposed Experiments | 110 |
| 5.2 | Motivation and Design | 114 |
| 5.3 | Simulation | 115 |
| 5.3.1 | Exclusion Regions | 117 |
| 5.3.2 | Track Classification | 118 |

| | | |
|----------|---------------------------------------|------------|
| 5.3.3 | Background Heat Maps | 119 |
| 5.3.4 | Magnetic Field | 121 |
| 5.3.5 | Event Replay | 121 |
| 5.4 | Trackers and Decay Volume | 122 |
| 5.4.1 | Decay Volume Length | 123 |
| 5.4.2 | Tracker Size and Resolution | 125 |
| 5.4.3 | Acceptance | 128 |
| 5.5 | Target and Magnetic Field | 129 |
| 5.5.1 | Target Thickness | 129 |
| 5.5.2 | Magnetic Field Length | 132 |
| 5.5.3 | Optimal Configuration | 134 |
| 5.6 | Discovery Potential | 135 |
| 6 | Conclusions | 138 |
| | Bibliography | 141 |

List of Figures

| | | |
|-----|--|----|
| 2.1 | Schematic of the FLASH superconducting radiofrequency modules supplying electrons to the FLASHForward beamline. The FLASH1 and FLASH2 free-electron laser beamlines are also shown. | 30 |
| 2.2 | Schematic of the FLASHForward beamline updated to include the X-TDS and other features of the Phase-II upgrade. | 31 |
| 2.3 | (a) Energy spectra and (b) charge of the injected witness beam. The injection is switched off by inserting a beam block inside the injection laser path for shots 101-200 and 301-400. | 33 |
| 2.4 | (a) Waterfall plot demonstrating stable operation across 1885 shots. (b) Mean energy of the injected witness beam (blue) of (45 ± 5) MeV and rms energy spread (red) of 4.4%. | 34 |
| 2.5 | 3D parameter scan of the plasma density (1) versus notch position (2) as a function of notch width (3). The red circle indicates the determined optimal operating point. | 35 |
| 2.6 | Waterfall plot demonstrating high stability across 5000 consecutive shots, with an energy gain of 45.4 MeV stable within 3% rms (± 1.4 MeV). Histograms of the energy spread and efficiency are also included, showing fully preserved energy spread in 6.4% of the shots and an energy transfer efficiency distribution between 30% and 50%. | 36 |

| | | |
|------|--|----|
| 2.7 | Comparison of the incoming trailing bunch charge (grey) with the accelerated witness bunch charge (blue) in a tail-collimator scan. Full charge coupling is observed until the position is less than approximately $-360\ \mu\text{m}$, at which point some of the beam particles cross the axis and are defocused by the plasma electrons. | 37 |
| 2.8 | Schematic of the 10 Hz macro-pulse and MHz bunch train structure provided by FLASH to FLASHForward for high power experimentation. | 38 |
| 2.9 | The relative energy change and transverse size between bunches propagating through unperturbed and perturbed plasma. | 39 |
| 2.10 | Energy spectra and transverse distributions of the driving and trailing probe bunches after interaction with perturbed plasma. | 40 |
| 2.11 | Comparison of the mean energy evolution of the trailing and driving bunches with the equivalent simulated values. Grey data points correspond to measurements where some charge was lost. | 42 |
| 2.12 | Layout of the AWAKE experiment. The proton bunch and laser pulse propagate through a 10 m long column of rubidium vapour. The laser pulse ionises the rubidium to form a plasma which then interacts with the proton bunch, resulting in wakefields. Electrons are then injected into the plasma cell and become captured by the wakefields, accelerating them to high energy. | 43 |
| 2.13 | Streak camera images showing microbunches at low (top left) and high (bottom left) plasma densities. A longer time window is used to show the full microbunch train at low density (right). | 45 |
| 2.14 | Image of scintillator screen (a) showing an electron acceleration signal. A projection of the image is obtained (b) by integrating vertically over the observed charge. | 46 |
| 2.15 | Projections obtained from consecutive electron-injection events over a 2 hour running period. Each event is a vertical integration of a background subtracted spectrometer image. | 47 |

- 2.16 Measurements of the maximum energy at different plasma densities, with and without the use of a plasma density gradient. The gradients chosen were those observed to maximise the energy gain. 48
- 2.17 Layout for run 2 of the AWAKE experiment showing the new electron source, vapour source, and laser beam that are to be introduced. 49
- 3.1 Timeline of the experimental method, indicating when the second high voltage discharge and the beam injection occurs relative to the first high voltage discharge. The timescales of processes that govern plasma wakefield evolution are also shown. 54
- 3.2 Schematic of the FLASHForward gas system. High pressure gas is regulated to ~ 1 bar and a flow valve limits gas flow to the buffer volume. Gas can only exit the system via the ends of the plasma cell, resulting in a constant mass flow between the flow valve and vacuum. 54
- 3.3 An unprocessed spectrometer image of size 1056 by 890 pixels. The region selected for investigating the distribution of background counts is indicated by a red rectangle. 55
- 3.4 Distribution of pixel counts in the background region. 56
- 3.5 The drive (left) and witness (right) bunches shown after the first background subtraction stage. Pixels with values below the threshold value are set to zero. A threshold value of $x_T = \bar{x} + n \cdot x_{RMS}$ is used where $n = 0, 1, 2, 3$ from top to bottom. 57
- 3.6 Connected components with areas less than the two largest components are removed, resulting in only the drive and witness bunches left remaining. 59
- 3.7 Connected components with areas less than the two largest components are removed, resulting in the combined drive and witness bunch component left remaining, along with multiple background components (circled in red) each having the second largest area. 59

3.8 Connected components with areas less than the largest component are removed, resulting in only the connected drive and witness bunches left remaining. 60

3.9 A spectrometer image where the drive bunch has a larger mean pixel value than the witness bunch. Connected components with areas less than the two largest components are removed, resulting in only the drive and witness bunches left remaining. 61

3.10 A spectrometer image from a dataset where the drive bunches have larger mean pixel values than the witness bunches. Connected components with areas less than the two largest components are removed, resulting in the combined drive and witness bunch component left remaining, along with a single background component (circled in red). 62

3.11 Ratio of the product of total area and mean pixel value for spectrometer images where the drive and witness bunches are connected (blue) or separate (red). 63

3.12 Connected components with areas less than the largest component are removed, resulting in only the connected drive and witness bunches left remaining. 64

3.13 A spectrometer image where the light leakage overlaps with the witness bunch, resulting in them forming a single connected component. 64

3.14 A spectrometer image where the light leakage is located between the drive and witness bunches, resulting in its own connected component. 65

3.15 Isolated light leakage components from ten spectrometer images (a-j) are used to predict the light leakage in further spectrometer images (k). 66

3.16 Subtraction of the predicted light leakage component from the spectrometer image shown in figure 3.13. 67

3.17 A spectrometer image where the x-axis is converted to energy in MeV. 67

| | | |
|------|---|----|
| 3.18 | Relationship between the horizontal position on the spectrometer screen and energy. | 68 |
| 3.19 | Histogram of the sum of pixel values for each column of energy in the spectrometer image shown in figure 3.17. Bins corresponding to the drive and witness bunches are shown in red and blue respectively. | 69 |
| 3.20 | Waterfall plot of a set of reference measurements showing a drift in the measured energy of the drive (lower energy band) and witness (higher energy band) bunches. | 70 |
| 3.21 | Energies of the witness and drive bunches as a function of time between pulses of the high voltage discharge, with the reduced chi-squared statistics calculated using statistical uncertainty (top), estimated using the reference dataset (middle), and inferred a posteriori (bottom). | 71 |
| 3.22 | Data taken in an argon plasma capillary at 40.0 mbar backing pressure using HV discharge pulses of 15 kV. | 75 |
| 3.23 | Data taken in an argon plasma capillary at 40.0 mbar backing pressure using HV discharge pulses of 20 kV. | 76 |
| 3.24 | Data taken in an argon plasma capillary at 30.1 mbar backing pressure using HV discharge pulses of 15 kV. | 77 |
| 3.25 | Data taken in an argon plasma capillary at 30.1 mbar backing pressure using HV discharge pulses of 20 kV. | 78 |
| 3.26 | Data taken in an argon plasma capillary at 30.1 mbar backing pressure using HV discharge pulses of 25 kV. | 79 |
| 3.27 | Data taken in an argon plasma capillary at 20.6 mbar backing pressure using HV discharge pulses of 15 kV. | 80 |
| 3.28 | Data taken in an argon plasma capillary at 20.6 mbar backing pressure using HV discharge pulses of 20 kV. | 81 |
| 3.29 | Data taken in an argon plasma capillary at 20.6 mbar backing pressure using HV discharge pulses of 25 kV. | 82 |

| | | |
|------|---|----|
| 3.30 | Data taken in an argon plasma capillary at 15.4 mbar backing pressure using HV discharge pulses of 15 kV. | 83 |
| 3.31 | Data taken in an argon plasma capillary at 15.4 mbar backing pressure using HV discharge pulses of 20 kV. | 84 |
| 3.32 | Data taken in an argon plasma capillary at 15.4 mbar backing pressure using HV discharge pulses of 25 kV. | 85 |
| 3.33 | Data taken in a hydrogen plasma capillary at 17.5 mbar backing pressure using HV discharge pulses of 27 kV. | 86 |
| 3.34 | Measured reset times for different experimental configurations of the argon plasma capillary. The grey cell corresponds to a configuration where no measurements were taken. | 87 |
| 3.35 | Reset curves indicating the energy of the witness and drive bunches, and therefore the readiness of the plasma for acceleration, as a function of time between the current and previous discharge pulses. The curves are labelled according to the backing pressure of the configuration. | 88 |
| 3.36 | Reset curves indicating the energy of the witness bunch as a function of time between the current and previous discharge pulses for the 15.4 mbar (far left), 20.6 mbar (middle left), 30.1 mbar (middle right), and 40.0 mbar (far right) configurations. The curves are labelled according to the voltage parameter of the configuration. . . | 89 |
| 4.1 | 3D models of the FLASHForward capillary before meshing (a) and after meshing (b). | 95 |
| 4.2 | Simulation results corresponding to $8 \times 10^{-7} \text{ kg s}^{-1}$ of hydrogen gas entering the capillary. The pressure as a function of the longitudinal position along the length of the capillary is plotted for a variety of times between 0.01 ms (a) and 1 ms (f). | 96 |
| 4.3 | Simulation results corresponding to $8 \times 10^{-7} \text{ kg s}^{-1}$ of hydrogen gas entering the capillary. The pressure at the midpoint of the capillary as a function of time is plotted. | 97 |

| | | |
|-----|--|-----|
| 4.4 | Relationship between the backing pressure of the FLASHForward buffer system and the capillary pressure after achieving equilibrium for both argon and hydrogen. | 98 |
| 4.5 | Relationship between the capillary pressure at equilibrium and the flow rate for both argon (a) and hydrogen (b). | 99 |
| 4.6 | Simulation results showing gas filling curves for different gas species and backing pressures. | 101 |
| 4.7 | Experimental results showing reset curves for plasma wakefield acceleration in argon at varying backing pressures. | 101 |
| 5.1 | Current limits on the mass and mixing strength of dark photons as provided by a variety of prompt decay experiments (NA48/2, A1, LHCb, KLOE, BaBar, and CMS) and highly displaced vertex experiments (E774, NA64(e), E141, ν -Cal, CHARM, and E137). Excluded regions from supernovae (SN1987A) and the magnetic moment of the electron $(g-2)_e$ are also included. | 108 |
| 5.2 | Projections for proposed dark photons searches (MESA, Belle-II, FCC-ee, HL-LHC/FCC-hh, LHeC/FCC-eh, NA64(e) ⁺⁺ , AWAKE, HPS, LHCb, FASER, NA62, SeaQuest, and SHiP) are shown as coloured curves. Existing limits are shown as grey areas. A red line shows the coupling range of a new gauge boson X that could explain the ⁸ Be anomaly. | 111 |
| 5.3 | Exclusion regions plotted for an example configuration for the cases of 10^{14} , 10^{15} , and 10^{16} electrons fired at the target. | 114 |
| 5.4 | Schematic of the proposed experimental setup for an electron beam dump experiment used to detect the visible decay of dark photons. | 115 |
| 5.5 | Visualisation of the simulated experiment running in a background configuration. The target is shown in blue, the magnetic field in green, Mylar windows of the vacuum tube in white, and a general detector in red. | 116 |

- 5.6 Table output from the simulation running in the signal configuration, columns show the mixing strength, mass (GeV), and probability of measuring a dark photon signal per electron bunch (left). Sensitivity contour where the first column ($m_{A'} = 0.008$ GeV) corresponds to the values from the table output (right). 117
- 5.7 Positions of the hits in electromagnetic showers caused by dark photon decay products have been categorised into signal (blue) and background (red) by the original event classifier. 118
- 5.8 Positions of the hits in the same electromagnetic showers have been categorised using the new method, the absence of red markers indicates successful event classification. 119
- 5.9 Visualisation of the cell structured detector showing a run with 10 particle hits, the coordinates of the cells for each hit are obtained and printed. 120
- 5.10 Visualisation of a run where a cone shaped beam is fired at the top right quadrant of the detector (left), the detector is shown as a red grid, particle trajectories are shown as red lines, and hits are shown as yellow dots. The total energy entering each cell throughout the run is displayed as a heat map (right). 120
- 5.11 Visualisation of an event where a charged particle travels through a magnetic field of 0.1 T (top) and 0.3 T (bottom). 121
- 5.12 Visualisation of a run where an isometric source results in 5 events that interact with a detector plane, meeting the requirement for the random engine status to be saved (left). Visualisation of a separate run where the 5 events of interest are replayed by loading their random engine status (right). 122
- 5.13 Projected exclusion regions plotted for an example configuration with decay volume lengths of 1 m, 5 m, and 20 m. 123

- 5.14 Path length distributions for dark photons with $m_{A'} = 0.1$ GeV and $\varepsilon = 1 \times 10^{-4}$ (red), $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ (blue), and $m_{A'} = 0.3$ GeV and $\varepsilon = 3 \times 10^{-5}$ (black). 124
- 5.15 Relationship between the opening angle of the e^+e^- pair and the energy of the parent dark photon. The parameters of the generated dark photons are $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ 125
- 5.16 Separation distance between hits of e^+e^- pairs in a tracker located 1 m after the target for dark photons with $m_{A'} = 0.1$ GeV and $\varepsilon = 1 \times 10^{-4}$ (red), $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ (blue), and $m_{A'} = 0.3$ GeV and $\varepsilon = 3 \times 10^{-5}$ (black). 126
- 5.17 Separation distance between hits of e^+e^- pairs in a tracker located 20 m after the target for dark photons with $m_{A'} = 0.1$ GeV and $\varepsilon = 1 \times 10^{-4}$ (red), $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ (blue), and $m_{A'} = 0.3$ GeV and $\varepsilon = 3 \times 10^{-5}$ (black). 127
- 5.18 Acceptance for dark photons with $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ for both the larger size tracker and the higher resolution tracker placed at a variety of distances from the target. 128
- 5.19 Projected exclusion regions plotted for an example configuration with target thicknesses of 10 cm, 40 cm, and 150 cm. 130
- 5.20 Total energy (top) and number of particles (bottom) passing through the detectors, results are normalised to represent one electron bunch. 131
- 5.21 Number of background particles measured at the up-stream detectors for different configurations of target thicknesses and magnet lengths, results are normalised to represent one electron bunch. Grey cells correspond to configurations where no background is detected within the simulation runtime. 132

- 5.22 Signal to background ratio for different configurations of target thicknesses and magnet lengths, results are normalised to represent one electron bunch. Grey cells correspond to configurations where no background is detected within the simulation runtime. Dark photons are generated with parameters $m_{A'} = 0.1 \text{ GeV}$ and $\epsilon = 3 \times 10^{-5}$. 133
- 5.23 Number of electron bunches required to exclude dark photons of any given parameter point in the mass and mixing strength plane with 95% confidence for the configuration described in this chapter. 136
- 5.24 Current limits (top) and projected exclusion regions of proposed experiments (bottom) for dark photons in the mass and mixing strength plane. The configuration described in this chapter is indicated on both plots by a cyan curve labelled AWAKE. 137

List of Tables

| | | |
|-----|---|----|
| 3.1 | List of the total area and mean pixel value of the five largest connected components in the spectrometer image discussed above. Components are listed in descending order based on their total area. | 58 |
| 3.2 | List of the total area and mean pixel value of the nine largest connected components in the spectrometer image shown in figure 3.7. Components are listed in descending order based on their mean intensity. | 60 |
| 3.3 | List of the total area and mean pixel value of the two largest connected components in the spectrometer image shown in figure 3.9. Components are listed in descending order based on their total area. The product of the total area and mean pixel value is also provided for both components. | 62 |
| 3.4 | List of the total area and mean pixel value of the two largest connected components in the spectrometer image shown in figure 3.10. Components are listed in descending order based on their total area. The product of the total area and mean pixel value is also provided for both components. | 63 |
| 3.5 | List of the major axis length and orientation of the three connected components with the largest major axis lengths for the spectrometer image shown in figure 3.14. Components are listed in descending order based on their major axis length. | 65 |

| | | |
|-----|--|-----|
| 4.1 | Input flow rates for the gas filling simulation to correspond to a given backing pressure and gas species in the FLASHForward experiment. | 100 |
| 5.1 | Signal, background, and the signal to background ratio is provided for the following configurations: a 100 cm thick target followed by a 50 cm long magnetic field, a 130 cm thick target without any magnetic field, and a 100 cm thick target surrounded by a magnetic field of equal length. Dark photons are generated with parameters $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ | 135 |

Chapter 1

Introduction

The Standard Model (SM) of particle physics [1] remains the most complete description of all known elementary particles and how they interact with three of the four known fundamental forces. The construction of this framework has relied substantially on the increasingly powerful particle accelerators developed over the latter half of the 20th century. However, there exists several deficiencies of the SM, including but not limited to its inability to explain baryon asymmetry [2], the non-zero mass of the neutrino [3], the accelerating expansion of the universe [4], and the apparent preservation of CP-symmetry by the strong interaction. Furthermore, it fails to provide a sufficient candidate for dark matter [5], and some of its predictions show possible contention with experimental measurements [6, 7]. This motivates the need to search for physics beyond the Standard Model (BSM), thereby also motivating the need to develop more powerful particle accelerators.

1.1 Conventional Acceleration

Conventional particle accelerators use radio-frequency (RF) cavities to provide an oscillating electric field, exerting a force on a charged particle along its direction of propagation and thereby accelerating it. A limit is posed on the electric field strength that can be supported by an RF cavity, as too strong of a field would ionise electrons from its metallic surface, resulting in cavity 'breakdown'. Arranging a series of cavities along a linear path allows for continuous acceleration to indefinitely

high energies, however this becomes limited in practise by the cost of building and running an indefinitely long linear acceleration facility.

Arranging RF cavities in a circle and allowing charged particles to circulate through them repeatedly enables continuous acceleration within a finitely large circular acceleration facility. This requires manipulating the trajectory of charged particles, which is achieved by using a magnetic field to bend them with a radius of curvature, r , given by

$$r = \frac{p}{qB}, \quad (1.1)$$

where p and q are the momentum and charge of the particle respectively, and B is the strength of the magnetic field.

However, circular acceleration is limited by the loss of energy due to synchrotron radiation, an electromagnetic radiation emitted when relativistic charged particles are subject to an acceleration perpendicular to their velocity. The energy loss per orbit scales according to

$$\Delta E \propto \frac{\gamma^4}{r}, \quad (1.2)$$

where γ is the Lorentz factor. This is particularly problematic when accelerating lighter particles such as electrons or positrons, which have a power loss per orbit of over 10^{13} times that of the proton [8].

Consequently, the construction of an electron-positron collider with a centre of mass energy on the TeV scale would have to be linear; to keep the size and cost feasible the problem of cavity breakdown would have to be circumvented. One approach involves accelerating particles within plasma, a state of matter already ionised to such an extent that large accelerating gradients can be employed without negating the acceleration medium itself.

1.2 Plasma Wakefield Acceleration

Plasma refers to an electrically quasi-neutral medium consisting of unbound ions and electrons. Although unbound, moving charged particles generate electric currents and thereby influence the other charges, resulting in them behaving collectively.

If the plasma is perturbed, the electrons become displaced with respect to the ions and are then pulled back by the Coulomb force, this results in plasma oscillations where the charge density oscillates at the plasma frequency, ω_p , given by

$$\omega_p = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}}, \quad (1.3)$$

where n_e is the plasma electron density, e is the electron's charge, ϵ_0 is the permittivity of free space, and m_e is the effective mass of the electron. Such perturbations can be applied by the ponderomotive force from an intense laser pulse, or by the space-charge field of a particle bunch provided by an electron beam or a proton beam.

The charge density oscillations generate a strong longitudinal electric field capable of accelerating trailing particles. The strength of this electric field, E , is given by

$$E = \frac{m_e c \omega_p}{e}, \quad (1.4)$$

where c is the speed of light.

Particles can be injected into the perturbed plasma at the precise phase of the plasma oscillations where the electric field is directed in the same direction as the plasma wave's propagation. The injected particles are effectively trapped in this phase and experience sustained acceleration as they travel along with the oscillating

field.

Through this method, accelerating gradients of $\mathcal{O}(\text{GVm}^{-1})$ can be achieved, significantly higher than the $\mathcal{O}(\text{MVm}^{-1})$ accelerating gradients produced in RF cavities. This allows for particle accelerators to be built with a lower requirement of physical space and cost, while also bypassing the issue of energy loss due to synchrotron radiation. However, challenges remain regarding the beam quality, luminosity, and scalability of such a particle accelerator, and these challenges must be overcome before this technique can be applied for particle physics experiments.

1.2.1 Laser-Driven Plasma Wakefield Acceleration

These plasma oscillations, also referred to as plasma wakefields, began to gain traction as a potential particle acceleration structure in 1979 when Tajima and Dawson proposed using an intense laser pulse to perturb the plasma via the ponderomotive force, thereby inducing strong electric fields [9]. This idea became known as laser wakefield acceleration (LWFA) and was demonstrated experimentally by three separate teams in 2004 [10–12], with each team producing a monoenergetic electron beam at the 100 MeV energy scale.

However, this technique has several limitations, one of them being the distance over which particles can be accelerated. As the wakefield travels at the speed of the laser pulse, which is less than the speed of light in a vacuum, electrons being accelerated to velocities close to the speed of light will escape the accelerating region of the wakefield after a dephasing length, L_d , given by

$$L_d \approx \frac{\lambda_p^3}{\lambda^2}, \quad (1.5)$$

where λ_p is the wavelength of the plasma given by

$$\lambda_p \approx \frac{2\pi c}{\omega_p}. \quad (1.6)$$

and λ is the wavelength of the laser used to drive the wakefield. This results in a dephasing length of 60 mm for a 800 nm laser pulse propagating in a plasma of density 10^{18} cm^{-3} [13].

Another issue arises from the repetition rate of the high intensity laser pulses used to drive the wakefields, which often have a frequency of $\mathcal{O}(\text{Hz})$. This presents a limit on the achievable luminosity of a LWFA beam, further weakening its viability as an alternative to conventional acceleration. Additionally, dispersion of the laser pulses becomes more pronounced over longer distances, limiting the acceleration that can be achieved within a single stage. Experiments have so far only demonstrated acceleration up to 7.8 GeV in a single stage [14], with simulations indicating that single stage acceleration up to 10 GeV is possible [15].

1.2.2 Beam-Driven Plasma Wakefield Acceleration

The use of an electron beam as the wakefield driver rather than an intense laser pulse was proposed in 1985 [16, 17], marking the advent of beam-driven plasma wakefield acceleration (PWFA). The first demonstration of this idea was achieved in 1988 at the Argonne National Laboratory, where plasma wakefields were excited by a leading 'drive' bunch and probed by a following 'witness' bunch [18]. As the electron driver travels through the plasma at the same speed as the accelerated witness electrons, dephasing will no longer occur, significantly extending the distance over which particles can be accelerated. Furthermore, electron drive bunches can be provided at much higher repetition rates than high intensity lasers, greatly increasing the achievable luminosity of such an accelerated beam.

Further experimental demonstrations of the larger energy gain achievable using a beam driver were successful. In 2007, the energy doubling of a 42 GeV electron beam at the Stanford Linear Accelerator Center was achieved in a plasma wakefield accelerator of length 85 cm [19], albeit with a continuous energy spread. In 2016, FACET demonstrated an energy gain of 9 GeV for witness electrons driven using a

21 GeV drive beam, this time with an energy spread of less than 5% [20]. In 2020, the FLASHForward experiment [21] demonstrated per-mille energy spread preservation using a 1 GeV drive beam to accelerate witness electrons by 45 MeV [22].

However, the distance over which particles can be accelerated is still limited. The drive beam loses energy as it propagates through the plasma, when it depletes significantly it will no longer have the sufficient energy required to sustain the desired wakefield amplitude and provide effective acceleration for the witness beam. This means that dividing the acceleration process into multiple stages would still be required for reaching the TeV energy scale.

1.2.3 Proton-Driven Plasma Wakefield Acceleration

Proton beams, such as the ones provided by the Super Proton Synchrotron (SPS) and Large Hadron Collider (LHC), have energies far exceeding even the most energetic electron beams currently in existence. Such a beam can be used to drive wakefields, allowing for proton-driven plasma wakefield acceleration (PDPWA) [23].

The AWAKE experiment [24] has so far demonstrated single stage acceleration of a witness beam from 19 MeV to 2 GeV within 10 m using wakefields driven by 400 GeV protons from the SPS [25]. As proton beam drivers have much higher energies than their electron counterpart, the length of the accelerator can be scaled to such an extent that electron acceleration to the TeV scale would be possible within a single stage [26], greatly enhancing the use of particle accelerators for high energy physics research.

As the plasma used in the AWAKE experiment is currently produced using laser ionisation, it will eventually become necessary to explore methods of sustaining a plasma state along larger distances to fully exploit the advantages of PDPWA. One such method involves applying a high voltage across the gas, thereby ionising it and forming plasma. Limitations of the acceleration process imposed by the high

voltage discharge scheme, such as the rate at which it can be repeated, remains an active area of research and is addressed in this thesis. The potential application of this acceleration scheme to search for a hypothetical hidden sector particle known as the dark photon is also addressed.

1.3 Outline

This thesis details the work carried out by the author on the FLASHForward and AWAKE experiments.

Chapter 2 provides an overview of the FLASHForward and AWAKE experiments, including a description of their experimental facilities and a summary of their previous results and future plans.

Chapter 3 details an experimental investigation undertaken to determine the repetition rate of plasma wakefield acceleration under a variety of plasma conditions at FLASHForward.

Chapter 4 presents simulation results of the entry of gas into the FLASHForward plasma capillary. Comparisons are then made between the time taken to fill the capillary and the experimental measurements of the previous chapter.

Chapter 5 proposes a new search using the AWAKE scheme to produce and detect a hypothetical hidden sector particle known as the dark photon. Simulations are performed to optimise the experimental configuration and assess its discovery potential.

Chapter 6 summarises the work and offers suggestions for further study.

Chapter 2

The FLASHForward and AWAKE Experiments

With plasma wakefield acceleration providing a promising technique to complement conventional acceleration methods, a widespread effort is being made to further current understanding of the field and develop the technologies required to build plasma wakefield accelerators capable of addressing physics questions. The FLASHForward and AWAKE experiments both contribute towards this effort. A summary of both is provided in this chapter.

2.1 FLASHForward

The FLASHForward experiment [21] is an electron driven plasma wakefield accelerator located at the Deutsches Elektronen-Synchrotron (DESY) in Hamburg, Germany. The three primary areas of exploration are high brightness beam generation using internally injected electrons, high quality beam generation using externally injected electrons, and high average power beam generation using externally injected electrons accelerated at a high repetition rate. These goals are pursued by the X-1, X-2, and X-3 experiments respectively. An overview of the experimental facility is provided, followed by a discussion of the three main research goals of the collaboration.

2.1.1 Experimental Facility

The FLASH linear accelerator (see figure 2.1) is capable of providing ultra-short (~ 50 fs), low emittance ($2 \mu\text{m rad}$), and low energy spread (0.1%) bunches of 1.25 GeV electrons to FLASHForward through the use of its superconducting niobium accelerating cavities. The longitudinal current profiles of the bunches can be manipulated [27] to induce a linear correlation between the longitudinal position and the energy of the electrons within each bunch, this energy chirp enables the bisection of the electron bunch into a drive and witness bunch pair [28] during their extraction into the FLASHForward beamline. The beamline consists of seven sections: extraction and compression, matching and final focusing, central interaction, initial diagnostics, undulators, X-TDS, and the final dump.




Figure 2.1: Schematic of the FLASH superconducting radiofrequency modules supplying electrons to the FLASHForward beamline. The FLASH1 and FLASH2 free-electron laser beamlines are also shown. Image reproduced from Ref. [21].

The extraction and compression section extracts the bunches from FLASH into the FLASHForward beamline, resulting in a 4 m transverse separation from FLASH2. Two dipoles are used to correct the trajectory change from the extraction dipole, and a mask and wedge system known as the scraper is used to remove parts of beam, generating the drive and witness bunch pairs. The matching and final focusing section uses a series of quadrupoles to focus the beam down to a transverse size of $5 \mu\text{m}$, while also reducing the spatial and angular transverse jitter to $<10 \mu\text{m}$ and <0.5 mrad respectively. Both of these sections are equipped with differential pumping stages to provide a gradient from high pressure to the vacuum further upstream, this compensates for the vacuum degradation that occurs when gas is loaded into the plasma capillary.

The central interaction region is composed of a plasma capillary and diagnostics mounted upon a baseplate. Beneath the baseplate is a hexapod positioning system [29], which allows for precise rotational and translational movement of the capillary. Capillaries with lengths of 33 mm, 50 mm, and 195 mm are available, each with a diameter of 1.5 mm. The plasma is created by ionising either hydrogen or argon gas in the capillary by using either a 25 TW laser pulse or an electrical discharge with voltage up to 25 kV. A series of high strength quadrupoles are then used to capture and transport electrons of up to 2.5 GeV, the electrons then interact with a dipole magnet before hitting either a LYSO or a Lanex scintillator screen to measure their energy. The scintillator screen is then imaged using CCD cameras.

An X-band transverse deflecting structure, known as the X-TDS, was installed in 2019 to allow for reconstruction of the slice emittance in one transverse plane and the geometric emittance in the other. An electron spectrometer set-up was also installed to provide the dispersion required for measuring the longitudinal phase space of the driver and witness bunches [30]. A schematic of the upgraded beamline is shown in figure 2.2.




Figure 2.2: Schematic of the FLASHForward beamline updated to include the X-TDS and other features of the Phase-II upgrade. Image reproduced from Ref. [21].

2.1.2 Internal Injection (X-1)

Witness beams can be injected either from an external source, or internally by trapping the ambient plasma electrons. In the case of internal injection, the electron bunch is formed with a much smaller phase space, resulting in a witness bunch with much smaller normalised transverse emittance, tens to hundreds of picocoulombs

of charge, and femtosecond bunch durations [31–33]. These properties provide an opportunity to generate beams with a much larger brightness compared to conventional accelerators.

A density downramp can be used to reduce the phase velocity of the wake, allowing for fast plasma electrons to rephase into the accelerating region where they become trapped and form the witness bunch. The X-1 experiment [34] uses electron beams from the FLASH linac with ~ 800 pC of charge and a mean energy of up to 1.25 GeV to drive a wake inside the FLASHForward capillary. Two independently focused pulses of the 25 TW laser system are used to generate the plasma, one of which is focused along the electron-beam axis while the other is focused transverse to it. This combined laser ionisation technique results in a sharp density spike surrounded by a plateau region, providing the required density downramp.

The X-1 experiment accelerated a drive beam of charge (790 ± 4) pC to an energy of (1116 ± 6) MeV and compressed it to a peak current of 2.1 kA. The 50 mm long capillary was filled with a mixture of helium and argon gas at a 2:1 ratio, and a plateau plasma density of $1.4^{+0.4}_{-0.3} \times 10^{16} \text{ cm}^{-3}$ was achieved. The charge was measured using toroids in combination with stripline and cavity BPMs (beam position monitors) [35], the energy was measured using the spectrometer dipole and LANEX screen, and the plasma density was determined by comparing the measured and simulated energy losses of the drive beam.

An investigation was made into how the energy of the witness beam depends on the presence of the injection laser. Figure 2.3 shows an injected witness beam charge of (32.1 ± 9.6) pC dropping to (0 ± 5) pC when the injection laser path is blocked, a correlation can then be observed between the witness beam charge and the energy spectra, demonstrating that the density downramp induced by the injection laser enables the acceleration of internally injected electrons. The witness beam charge was measured by subtracting the charge measurements of a down-

stream BPM from an upstream toroid [36], and measured again independently using a charge-calibrated phosphor screen in the spectrometer for confirmation.

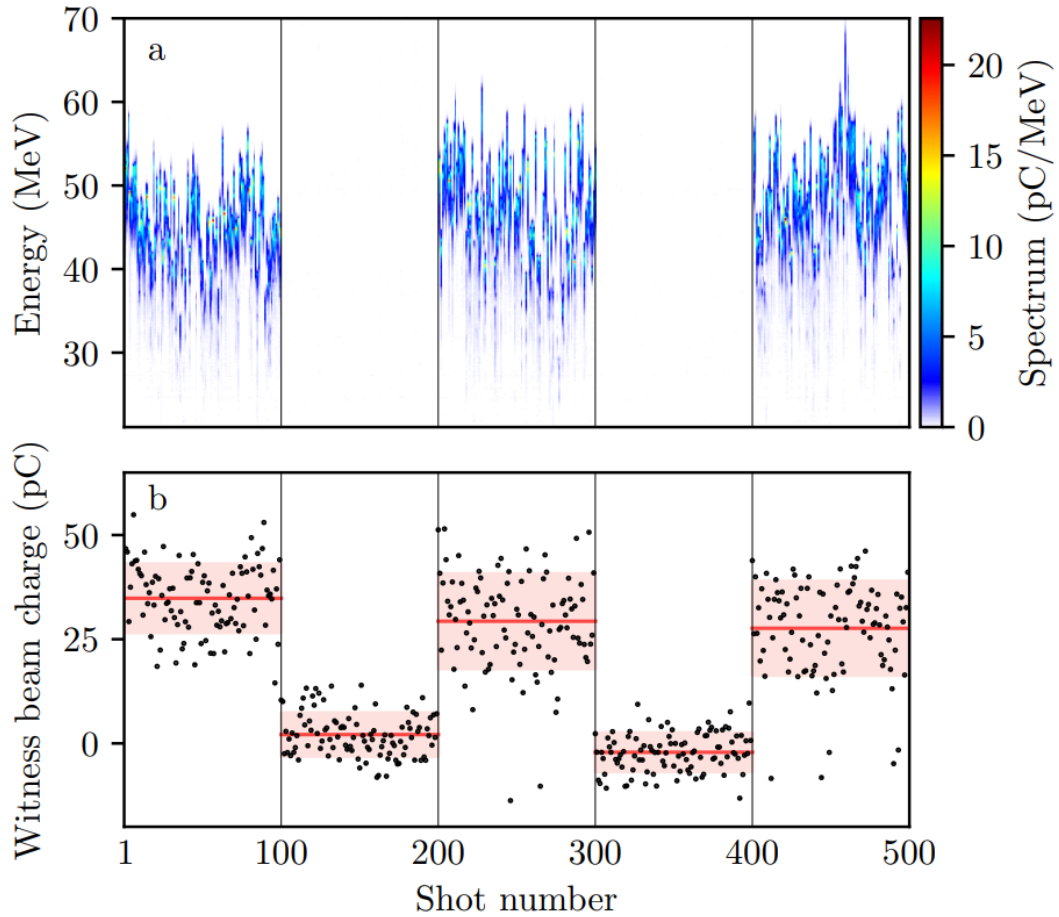


Figure 2.3: (a) Energy spectra and (b) charge of the injected witness beam. The injection is switched off by inserting a beam block inside the injection laser path for shots 101-200 and 301-400. Image reproduced from Ref. [34].

A dataset of 1885 consecutive shots were taken to investigate the reliability of the acceleration process. Figure 2.4 shows a waterfall plot of the projected witness beam energy spectra, along with the evolution of mean energy and rms energy spread. The injected witness beam was measured for $(95.4 \pm 2.5) \%$ of the shots attempted, with the measurements demonstrating a mean charge of (33 ± 10) pC, mean energy of (45 ± 5) MeV, and rms relative energy spread of 4.4%.

Although internal injection had been previously shown to work in princi-




Figure 2.4: (a) Waterfall plot demonstrating stable operation across 1885 shots. (b) Mean energy of the injected witness beam (blue) of (45 ± 5) MeV and rms energy spread (red) of 4.4%. Image reproduced from Ref. [34].

ple [31–33], the X-1 experiment demonstrates the first stable operation of this technique for beam driven wakefield accelerators. Further studies were made into the injected bunch properties as a function of laser alignment and laser energy, the results of which can be found in Ref. [34].

2.1.3 External Injection (X-2)

The X-2 experiment aims to demonstrate that the quality of externally injected electrons can be maintained by simultaneously preserving the emittance and energy spread while also operating at a high acceleration efficiency. It has been shown in principle that using high current trailing bunches with a trapezoidal current profile can flatten the wakefield such that all particles experience the same acceleration [37], enabling the kinetic energy of the wake to be efficiently extracted while minimising the energy spread.

Electron bunches with an energy of 1 GeV were provided by the FLASH linac and then divided into witness and driver bunch pairs using the wedge shaped notch collimator. A three-dimensional parameter scan of the plasma density, the position of the notch, and the width of the notch was performed and the transformer ratio (given by the mean energy gain of the witness bunch normalised by the maximum energy loss within the driver), the energy-transfer efficiency, and the energy-spread-to-gain ratio was calculated for each combination of parameters in order to locate the optimal beam loading operation point (see figure 2.5).

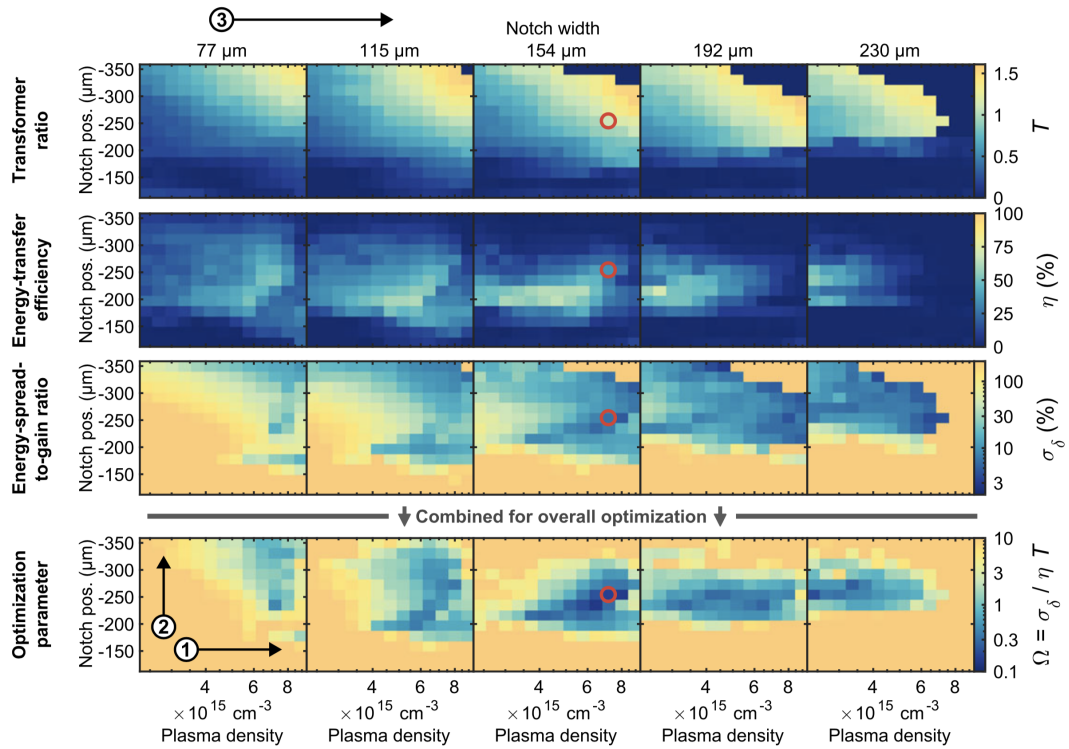


Figure 2.5: 3D parameter scan of the plasma density (1) versus notch position (2) as a function of notch width (3). The red circle indicates the determined optimal operating point. Image reproduced from Ref. [22].

A dataset of 5000 shots was obtained at the optimal operating point (see figure 2.6). High stability is demonstrated throughout the data taking period as the energy gain is shown to be (45.4 ± 1.4) MeV, stable to 3% rms. The histogram of energy spread shows 6.4% of the shots to have lower energy spread than the 0.16% FWHM energy spread seen when the plasma generation is off, meaning that energy spread has been fully preserved. The rest of the shots show a median energy spread of 0.2%

FWHM, corresponding to a relative increase of 28%. Energy transfer efficiency, defined by the ratio of energy gained by the witness bunch to energy lost by the drive bunch, is found to be distributed between 30% and 50% with an average of $(42 \pm 4)\%$.

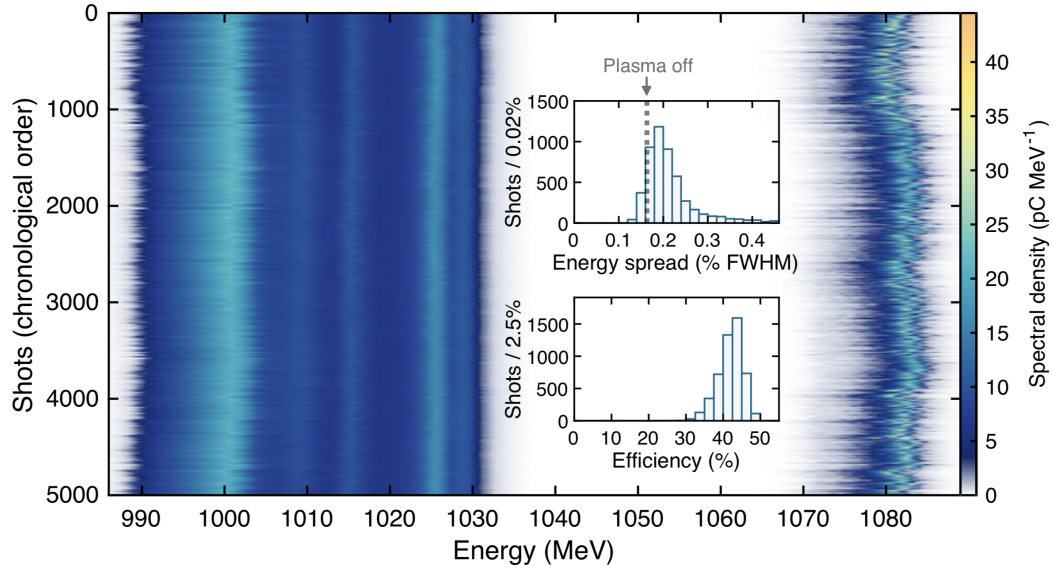


Figure 2.6: Waterfall plot demonstrating high stability across 5000 consecutive shots, with the energy gain of 45.4 MeV stable within 3% rms (± 1.4 MeV). Histograms of the energy spread and efficiency are also included, showing fully preserved energy spread in 6.4% of the shots and an energy transfer efficiency distribution between 30% and 50%. Image reproduced from Ref. [22].

For energy spread preservation to be conclusively demonstrated it is also necessary to show that no charge is lost. A scan of the tail-collimator position was performed and the charge of the incoming trailing bunch was compared to the charge of the accelerated witness bunch (see figure 2.7). Both measurements are in agreement until the tail of the trailing bunch becomes progressively removed from the wakefield cavity, at which point there is a strong defocusing by the high density of plasma electrons streaming around the cavity.

Having achieved an experimental demonstration of simultaneously preserving the energy spread while at a high energy transfer efficiency and with full charge coupling, the X-2 experiment aims to show further levels of control over the acceleration process. The next steps involve the simultaneous preservation of emittance,




Figure 2.7: Comparison of the incoming trailing bunch charge (grey) with the accelerated witness bunch charge (blue) in a tail-collimator scan. Full charge coupling is observed until the position is less than approximately $-360 \mu\text{m}$, at which point some of the beam particles cross the axis and are defocused by the plasma electrons. Image reproduced from Ref. [22].

as well as maintaining the quality of externally injected electrons while operating at larger energy gains.

2.1.4 High Average Power (X-3)

For plasma wakefield acceleration to meet the demands of photon science users and high energy physicists, it is necessary to demonstrate that operation at high average power is also possible. The FLASH gun and linac modules operate with a 10 Hz macro-pulse structure, with each pulse having a $800 \mu\text{s}$ long SRF flat top consisting of micro-pulses with up to 3 MHz frequency (see figure 2.8). At full capacity, the SRF modules can accommodate 2400 bunches per macro pulse, resulting in 24000 bunches per second. With each bunch having energy on the Joule-level, an average power of 24 kW is achieved, providing FLASHForward with an unmatched capability for conducting high average power plasma wakefield research.




Figure 2.8: Schematic of the 10 Hz macro-pulse and MHz bunch train structure provided by FLASH to FLASHForward for high power experimentation. Image reproduced from Ref. [38].

The X-3 experiment aims to probe plasma properties over a six order of magnitude temporal range, from the sub-ns level to the ms-level. This is important for plasma wakefield acceleration as unlike with conventional accelerators the accelerating medium now evolves over time in response to perturbations caused by the previously accelerated bunch. For timescales on the ps to ns level, movement of the plasma ions result in a non-uniform density profile before the regions of differing plasma density and temperature eventually equilibrate. On the ms timescale, the pressure gradient between the plasma in the capillary and the surrounding vacuum results in a net loss of gas which must be refilled. A review of investigations made into how ion motion affects the plasma wakefield acceleration process is provided below, and a more detailed study of how the loss of material from the capillary affects the plasma wakefield acceleration process is discussed later in this thesis.

Wakefields driven by an intense leading electron bunch were used to perturb the plasma. A probe bunch, bisected into a pair of bunches, then propagated through the perturbed plasma at varying times after the leading bunch, allowing for study into the effect of the perturbed plasma on the acceleration process after a given amount of time since the perturbation occurred. Comparisons were made with a

separate run where the plasma was unperturbed, with the relative energy change and the relative transverse size being used to quantify the difference between the perturbed and unperturbed case (see figure 2.9).




Figure 2.9: The relative energy change and transverse size between bunches propagating through unperturbed and perturbed plasma. Image reproduced from Ref. [39].

The relative energy change was defined as $\frac{\mu_{E,u} - \mu_{E,p}}{\Delta\mu_E}$ where $\mu_{E,u}$ and $\mu_{E,p}$ represent the mean energy of the unperturbed and perturbed bunches respectively, and $\Delta\mu_E$ represents the average energy gain (or loss) of the trailing probe bunch (or driving probe bunch). The relative transverse size was defined as $\frac{\sigma_{x,p} - \sigma_{x,u}}{\sigma_{x,u}}$ where $\sigma_{x,p}$ and $\sigma_{x,u}$ represent the transverse size of the trailing probe bunch in the perturbed and unperturbed scheme respectively. The convergence of the perturbed plasma to the unperturbed state was observed to occur at 63 ns.

After the initial perturbation, plasma ions are impelled towards the beam axis by the strong radial electric fields of the beam and the decaying plasma wave [40–43]. After the on-axis density reaches a maximum, the ion wave continues to propagate outwards. The non-uniformity of the radial density profile to

lowest order can be expressed as $n(r) = n_0(1 + \alpha r^2)$, where n_0 is the on-axis ion density, α is the curvature of the ion channel, and r represents the radial distance from the axis. To demonstrate that ion motion can plausibly explain the observed effect of the perturbation on the acceleration process, values for n_0 and α were first extracted from experimental signatures in the energy spectra and transverse distributions of the driving and trailing probe bunches (shown in figure 2.10). These values were then used to simulate the full evolution of the beam-plasma interaction, allowing for the comparison of predictions based on the ion motion model to be made with other features of the data.

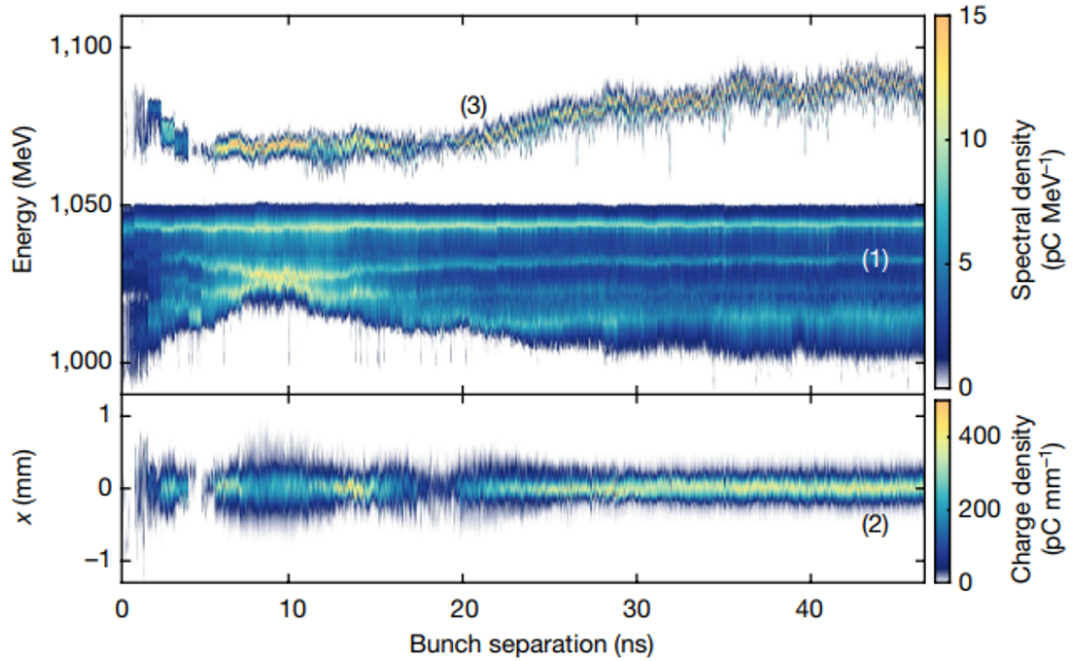


Figure 2.10: Energy spectra and transverse distributions of the driving and trailing probe bunches after interaction with perturbed plasma. Image reproduced from Ref. [39].

The three key experimental signatures observed in the energy spectra and transverse distributions are labelled as: (1) the betatron-mismatch bands, (2) transverse beam size oscillations, and (3) the mean energy evolution. A summary of how these features arise and how they were used to collectively confirm the ion motion model is provided:

- Betatron-mismatch bands: This arises from the betatron oscillations that elec-

trons experience within transverse wakefields. Slices of the bunch with different energies oscillate at different frequencies, resulting in bands of raised intensity that are independent of the on-axis density. The focusing force at a given longitudinal slice is modified by a factor $1 + \alpha r^2$ when in the presence of a parabolic transverse-plasma-density profile, allowing for the curvature of the channel to be extracted from the betatron bands.

- Transverse beam size oscillations: This arises from the longitudinal wakefield amplitude experienced by the trailing probe bunch decreasing and its betatron oscillation frequency being modified as the on-axis plasma density gradually decreases over time. As a result, the bunch accrues a betatron phase over the plasma length that varies with n_0 , leading to a differing divergence and transverse size at the plasma exit which appears in the data as transverse beam size oscillations. The transverse size is small compared to the parabolic channel, making this signature independent of α and allowing for the on-axis ion density to be extracted.
- Mean energy evolution: This arises from a combination of n_0 and α modifying the wakefield cavity which leads to an energy change in both the driver and trailing probes. The values of n_0 and α derived from the above effects were used to construct an evolving two-dimensional density map which was used as an input for particle-in-cell simulations, allowing for predictions of the mean energy evolution of the trailing and driving bunches to be compared to their measured value from the experimental data (see figure 2.11). The excellent agreement provides independent validation of the derivation procedure and verifies that the ion motion model is a plausible description of the physics occurring at this timescale.

The X-3 experiment also aims to investigate how the need to replenish any material expelled from the plasma affects the efficacy of plasma wakefield accelerators. This requires the introduction of a second high voltage discharge system to ensure that the plasma is refreshed to the same state for the second bunch as

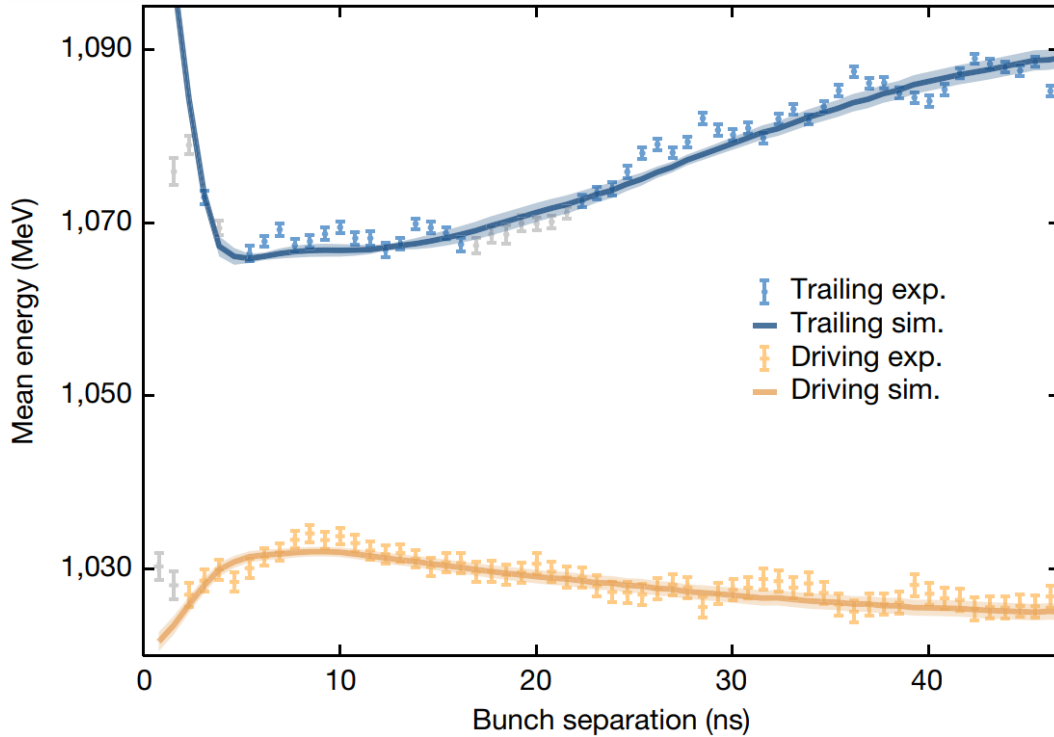


Figure 2.11: Comparison of the mean energy evolution of the trailing and driving bunches with the equivalent simulated values. Grey data points correspond to measurements where some charge was lost. Image reproduced from Ref. [39].

it was for the first bunch, an experimental demonstration of which is provided in chapter 3 of this thesis. As material is ejected from the capillary following the high voltage discharge, gas must be refilled into the capillary, the dynamics of which are explored in chapter 4 of this thesis. The final goal of this long timescale study involves scaling up the number of high voltage discharges, demonstrating that high average power acceleration can be maintained indefinitely.

2.2 AWAKE

The Advanced Wakefield Experiment (AWAKE) [24] is a proton driven plasma wakefield accelerator located at the European Organisation for Nuclear Research (CERN) in Geneva, Switzerland. An overview of the experimental facility is provided, followed by a review of recent experimental achievements and a discussion of future plans and potential applications of the scheme.

2.2.1 Experimental Facility

The AWAKE experimental set-up (shown in figure 2.12) consists of a 120 fs long laser pulse with a wavelength of 780 nm and an energy of up to 450 mJ. The pulse propagates through a 10 m long column of rubidium (Rb) vapour, thereby ionising it to produce plasma. The density of the plasma can be adjusted between 10^{14} cm^{-3} and 10^{15} cm^{-3} by controlling the temperature of the rubidium flasks (between $160 \text{ }^\circ\text{C}$ and $210 \text{ }^\circ\text{C}$) on both ends of the plasma cell. Heating the rubidium flasks to different temperatures allows for a plasma density gradient to be induced.




Figure 2.12: Layout of the AWAKE experiment. The proton bunch and laser pulse propagate through a 10 m long column of rubidium vapour. The laser pulse ionises the rubidium to form a plasma which then interacts with the proton bunch, resulting in wakefields. Electrons are then injected into the plasma cell and become captured by the wakefields, accelerating them to high energy. Image reproduced from Ref. [25].

Wakefields are driven by 19 kJ proton bunches from CERN's Super Proton Synchrotron (SPS), with each bunch containing between 2.5×10^{11} and 3.1×10^{11} protons, and with each proton having an energy of 400 GeV. Electron bunches with

a charge of (656 ± 14) pC are generated using a radio-frequency gun before being accelerated to (18.84 ± 0.05) MeV using a radio-frequency structure [25]. These electrons are then injected into the plasma cell, with their injection angle and focal point being controlled by magnets.

After acceleration in the plasma cell, the electron bunches can be measured using an electron spectrometer. This consists of two quadrupole magnets, one of which is used to focus the bunches vertically while the other is used to focus them horizontally. A 1 m long electromagnetic dipole with an adjustable magnetic field strength between 0.1 T and 1.4 T is then used to deflect the bunches by an amount proportional to their energy. The dipole directs the electron bunches to a 0.5 mm thick gadolinium oxysulfide scintillator screen, and a charge-coupled device (CCD) camera is used to detect the emitted light. The energy of the accelerated electrons can then be determined by the horizontal component of their hit positions on the screen.

2.2.2 Self-modulation

Driving wakefields with protons as opposed to electrons is beneficial as the larger achievable energy of proton beams such as the one provided by the SPS allows for wakefields to be driven over longer distances, meaning that electrons can then be accelerated to very high energies within a single acceleration stage. However, the bunch length of the drive beam is required to be short compared to the wavelength of the plasma, which is not the case for bunches produced at the SPS. Although this may seem problematic, a particle-plasma interaction known as self-modulation allows for long bunches to be split into a series of high density microbunches, facilitating their use in driving wakefields.

The transverse wakefields driven by a long proton bunch when entering the plasma periodically focuses and defocuses, resulting in regions of higher and lower bunch density respectively. The regions of focused protons are separated by the

plasma wavelength, self-modulating to form a train of microbunches that can resonantly drive large amplitude wakefields. To prevent the growth of unstable oscillations in the bunch centroid via a phenomena known as the hosing instability, the self-modulation must be seeded. This is achieved at AWAKE by using a laser pulse to create a relativistic ionisation front [24].

Microbunches have been observed for different plasma densities in the AWAKE experiment [44] using a 73 ps (~ 1 ps resolution) streak camera window (see figure 2.13). The top left plot shows the presence of microbunches at low plasma density ($2.457 \times 10^{14} \text{ cm}^{-3}$), with the time profile (green line) to the left also showing the microbunches and their periodicity. The bottom left plot shows the same but at higher plasma density ($6.994 \times 10^{14} \text{ cm}^{-3}$), and the plot on the right shows the full microbunch train with a plasma density of $2.190 \times 10^{14} \text{ cm}^{-3}$.

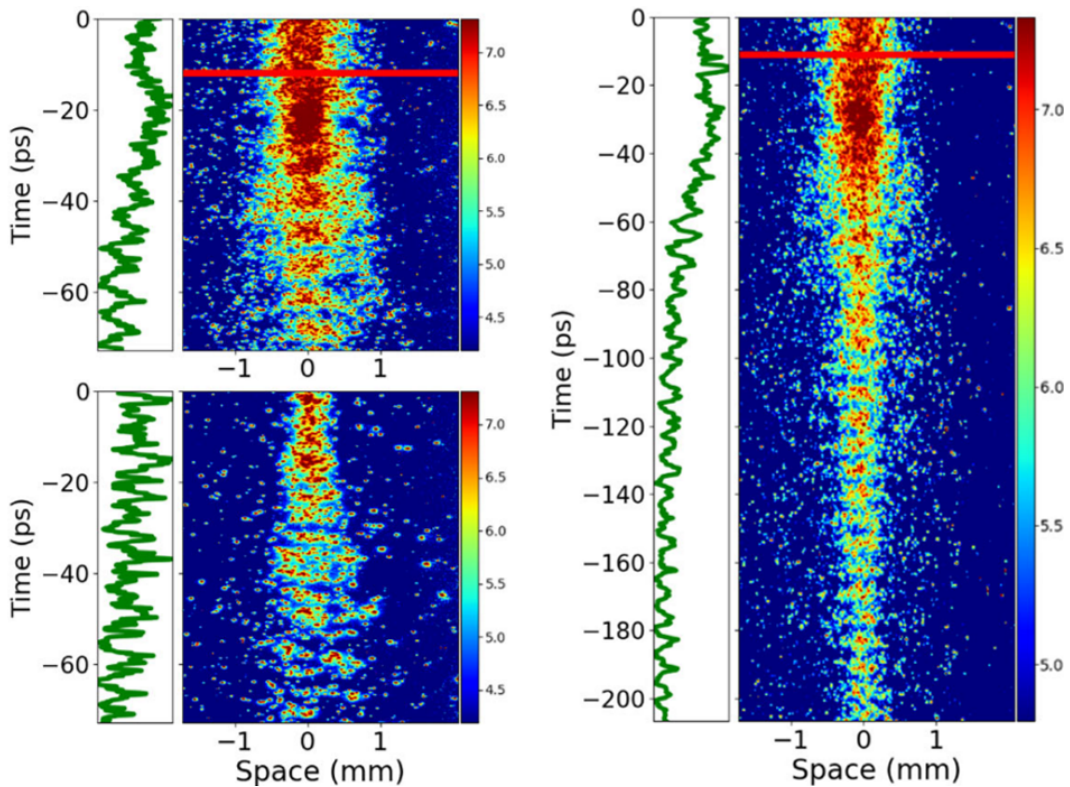


Figure 2.13: Streak camera images showing microbunches at low (top left) and high (bottom left) plasma densities. A longer time window is used to show the full microbunch train at low density (right). Image reproduced from Ref. [44].

2.2.3 Electron Acceleration

Electron acceleration using a self-modulated proton beam was achieved at the AWAKE experiment [25]. An image of an accelerated electron bunch on the scintillator screen is shown (see figure 2.14a). A projection (see figure 2.14b) was then obtained by integrating over a central region of the scintillator.

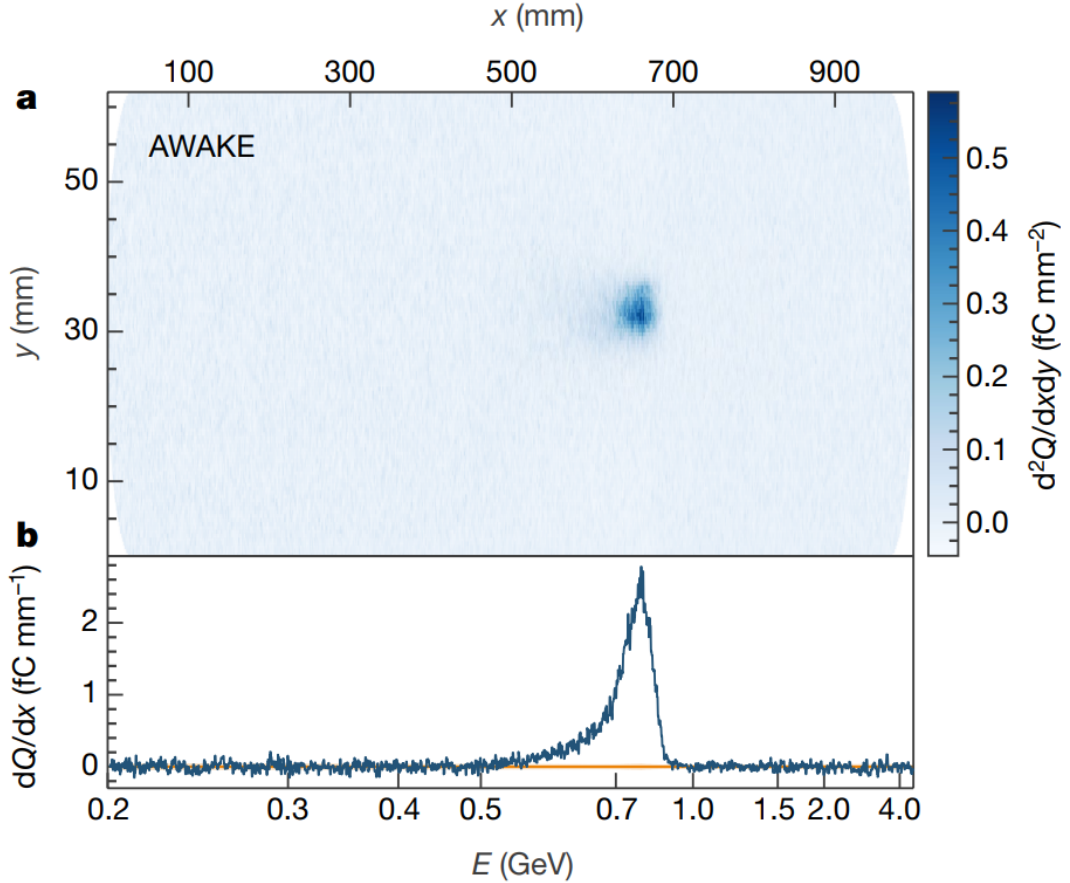


Figure 2.14: Image of scintillator screen (a) showing an electron acceleration signal. A projection of the image is obtained (b) by integrating vertically over the observed charge. Image reproduced from Ref. [25].

The event shown above corresponds to a plasma density of $1.8 \times 10^{14}\ cm^{-3}$, with a $(5.3 \pm 0.3)\%$ density increase over the 10 m plasma cell length in the direction of the propagation of the proton bunch. The mean energy of the accelerated electron bunch is (800 ± 40) MeV, with a FWHM of (137.3 ± 13.70) MeV and $\mathcal{O}(pC)$ charge.

The reproducibility of the electron bunch acceleration was investigated by plotting the projections from consecutive electron-injection events over a 2 hour running period (see figure 2.15). The quadrupoles were varied to focus over a range of approximately 460 to 620 MeV, and other parameters such as the proton-bunch population naturally vary each shot. Despite this, the electron bunches were shown to be consistently accelerated to an energy of approximately 600 MeV. The plasma density for this running period was also $1.8 \times 10^{14} \text{ cm}^{-3}$ however no density gradient was present, resulting in a different energy from the event shown in figure 2.14.




Figure 2.15: Projections obtained from consecutive electron-injection events over a 2 hour running period. Each event is a vertical integration of a background subtracted spectrometer image. Image reproduced from Ref. [25].

The effect of introducing a plasma density gradient was investigated by determining the maximum energy that could be achieved with and without a density gradient present (see figure 2.16). In each case, the density gradient was chosen

to maximise the peak energy that could be observed at a given plasma density. It can be seen that the highest obtainable energy increases with the density, reaching an energy of (2.0 ± 0.1) GeV at a density of $6.6 \times 10^{14} \text{ cm}^{-3}$ increasing by $(2.2 \pm 0.3)\%$ in the direction of the propagation of the proton bunch.




Figure 2.16: Measurements of the maximum energy at different plasma densities, with and without the use of a plasma density gradient. The gradients chosen were those observed to maximise the energy gain. Image reproduced from Ref. [25].

2.2.4 Run 2 Programme

Current seeding of the self-modulation process involves using a laser to place a relativistic ionisation front in the centre of the proton bunch, however this leaves the front of the bunch un-modulated which could in the case of long plasma cells interfere with the acceleration process. The self-modulation process can in principle be seeded by an electron bunch, thereby modulating the entire proton bunch and allowing the acceleration process to be scaled to larger distances. The current run of the AWAKE experiment [45] therefore consists of demonstrating seeding using an electron bunch.

In addition, the effect of imposing a density step in the plasma will be investigated. Simulations show that the amplitude of wakefields decays after their saturation, however the application of a density step can be used to maintain a near-saturation amplitude for a long distance. A new vapour source with additional observation ports and diagnostics will be exchanged with the current source, enabling direct measurements of the plasma wakefields.

As the self-modulation was found to saturate over a distance of 3 to 5 m [46], experiments involving a 10 m long self-modulator plasma followed by a 10 m long accelerator plasma will be conducted. To facilitate this, a second electron source and a second vapour source will be implemented into the beamline (as shown in figure 2.17). The second vapour source will be ionised using a second laser beam that is injected from the downstream end, propagating backwards with respect to the proton beam.




Figure 2.17: Layout for run 2 of the AWAKE experiment showing the new electron source, vapour source, and laser beam that are to be introduced. Image reproduced from Ref. [45].

Finally, the second plasma source will be exchanged with a different plasma technology. Creating plasma using laser ionisation becomes limited at longer distances by the depletion of energy as the laser pulse propagates through the vapour,

therefore plasma sources without length limitations will be investigated. The use of helicon argon plasma [47] and direct-current electrical discharge in noble gases [48] are currently under study.

2.2.5 Particle Physics Applications

The first considered application of electron bunches provided by the AWAKE scheme involves firing them at a metal target, potentially producing a hypothetical particle known as the dark photon. The collision of an electron beam with a fixed target requires less challenging beam parameters than other use cases such as an electron-positron collider, making it a suitable initial application while still having a strong and novel particle physics case. The design of a dark photon search experiment based on the AWAKE scheme is the focus of chapter 5 of this thesis.

Another considered initial application is the investigation of strong-field quantum electrodynamics (QED) in electron-laser collisions. As electrons pass through an intense laser pulse they experience a strong field within which QED becomes non-linear. Such experiments are capable of probing physics within extraordinary conditions, such as those at the surface of a neutron star or the event horizon of a black hole. As the collision rate is limited by the roughly 1 Hz repetition frequency of high power lasers, the requirement on the luminosity of the electron bunches is less stringent, resulting in this being a suitable initial application of the AWAKE scheme.

Collisions between electrons and protons provide the next potential application. As the emittance of the proton beam dominates, less stringent requirements are placed on the emittance of the electron beam, making it a suitable near future application. Collisions between $\mathcal{O}(50\text{ GeV})$ electrons and protons from the LHC, as proposed in the PEPIC (plasma electron-proton/ion collider) project [49], allow for investigation into quantum chromodynamics in a new kinematic regime. Collisions between 3 TeV electrons and protons from the LHC, as proposed in the VHEeP

collider [50], allow for investigation into hadronic cross sections and their potential saturation, while also being sensitive to physics beyond the standard model such as quark substructure and leptoquarks.

Lastly, the use of the AWAKE acceleration scheme is considered as a possible technology for the injector of a future electron-ion collider (EIC) at Brookhaven National Laboratory [51]. The EIC is expected to collide 275 GeV protons from its 4 km long proton accelerator with electrons of up to 20 GeV. The AWAKE scheme can provide an electron beam meeting the specifications of the EIC, and potentially provide even higher energy electrons, thereby extending the kinematic reach of the EIC in its investigation of parton dynamics in the proton [52].

Chapter 3

Repetition Rate of Plasma Wakefield Acceleration

Continuous repetition of the plasma wakefield acceleration process is necessary for the application of this technique to particle physics experimentation. The use of a second high voltage discharge system to ensure that the plasma is refreshed to the same state for the acceleration of a second electron bunch is demonstrated at the FLASHForward experiment, and the time taken for this to occur has been measured experimentally under a variety of plasma conditions.

3.1 Motivation

A gradient in the pressure between the plasma capillary and surrounding vacuum results in the expulsion of some of the plasma. Restoration of this lost material can be achieved by filling the capillary with gas and using it to create more plasma. Any plasma remaining in the capillary after the expulsion will undergo recombination, and therefore must also be converted into plasma again. The time taken for the plasma to achieve the same state for the acceleration of a subsequent electron bunch as it was for the previous electron bunch determines the rate at which the plasma wakefield acceleration process can be repeated.

In the case of plasma generation using a high voltage discharge system, contin-

uous plasma wakefield acceleration would require either stacking several discharge units, or developing a discharge unit with a switch capable of indefinitely reigniting the plasma at a sufficiently high frequency. The latter option, being more flexible and efficient, is planned to be employed at the FLASHForward experiment. Once installed, the validity of sustained plasma wakefield acceleration at high repetition rates can be tested. Prior to this, an upper limit on the repetition rate of unsustained plasma wakefield acceleration can be determined through the use of a single additional discharge unit identical to the current unit.

3.2 Experimental Method

Operation of the FLASHForward beamline with two identical and independently controlled discharge units allows for a second high voltage discharge to be fired through the gas at a variable time delay with respect to the first discharge. The acceleration of an electron bunch via the plasma created during the second discharge is characterised for different values of this time delay. The repetition rate can then be obtained by determining the minimum separation in time at which the acceleration signal is seen to be consistent with the signal observed when using only a single discharge.

The timeline of the experiment relative to the first high voltage discharge, as well as the typical timescales of processes that govern plasma wakefield evolution, are shown below (see figure 3.1). The time delay between the two discharges varies between $0.5 \mu\text{s}$ and 20 ms , and the electron bunch is injected shortly after the second discharge. This range encompasses the gas refill timescale, enabling measurements of the time taken for the lost material to be replenished.

Measurements in argon plasma were taken for each combination of backing pressures of 15.4, 20.6, 30.1, and 40.0 mbar, and discharges of 15, 20, and 25 kV, with the exception of an unstable configuration at 40.0 mbar and 25 kV where

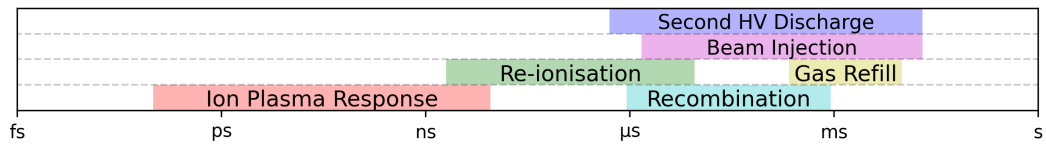


Figure 3.1: Timeline of the experimental method, indicating when the second high voltage discharge and the beam injection occurs relative to the first high voltage discharge. The timescales of processes that govern plasma wakefield evolution are also shown.

the acceleration signal fluctuated significantly. Backing pressure refers here to the absolute pressure of gas in the buffer volume of the FLASHForward gas system (see figure 3.2), and can be adjusted using a flow valve.

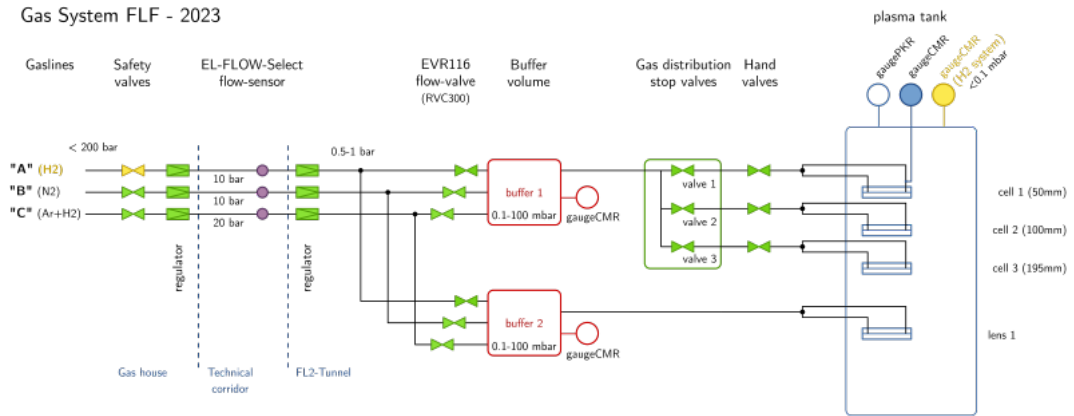


Figure 3.2: Schematic of the FLASHForward gas system. High pressure gas is regulated to ~ 1 bar and a flow valve limits gas flow to the buffer volume. Gas can only exit the system via the ends of the plasma cell, resulting in a constant mass flow between the flow valve and vacuum.

A ‘separation’ scan and a ‘reference’ dataset was taken for each configuration. The separation scan refers to measurements taken for a variety of time delays between the first and second discharge, and the reference dataset refers to the measurements taken using only one discharge. The reference images were taken to be compared with each time delay in the separation scan, allowing for the repetition rate to be determined. Measurements in hydrogen plasma were also taken using a configuration of 17.5 mbar and 27 kV.

3.3 Spectrometer Image Analysis

Interaction of background particles with the spectrometer screen results in the presence of pixels throughout the image with count rates arising from sources other than the signal bunches. These background counts are removed from the spectrometer image using various background subtraction techniques. This allows for the bunch energies to be determined, thereby enabling a plasma reset signature to be defined.

3.3.1 Pixel Threshold Value

A region of the spectrometer image that does not contain the drive and witness bunches is selected so that the distribution of background counts can be investigated. This region was selected as the first 300 rows in the spectrometer image shown below (see figure 3.3).




Figure 3.3: An unprocessed spectrometer image of size 1056 by 890 pixels. The region selected for investigating the distribution of background counts is indicated by a red rectangle.

The distribution of background counts is plotted (see figure 3.4). The mean pixel count in this region, \bar{x} , is determined to be 0.4711 pixels, and the root mean square of pixel count in this region, x_{RMS} , is determined to be 4.4967 pixels.

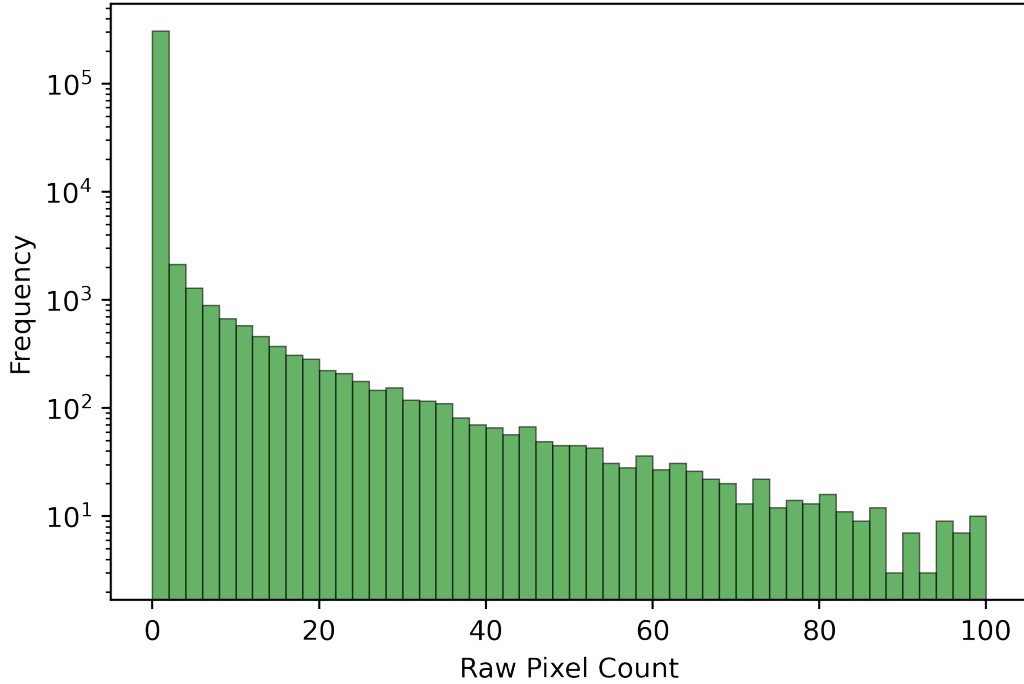


Figure 3.4: Distribution of pixel counts in the background region.

Defining a pixel threshold value, x_T , allows for pixels to be classified as background and therefore removed if their counts are below the threshold. This threshold is defined as

$$x_T = \bar{x} + n \cdot x_{RMS}, \quad (3.1)$$

where n is a positive integer selected such that as many background pixels are removed as possible without affecting the drive and witness bunches. Although removing background pixels below the threshold value does not sufficiently remove all background, it greatly increases the computation speed of the subsequent background subtraction techniques that are yet to be applied. The effect of defining the pixel threshold value for n ranging from 0 to 3 was investigated (see figure 3.5).

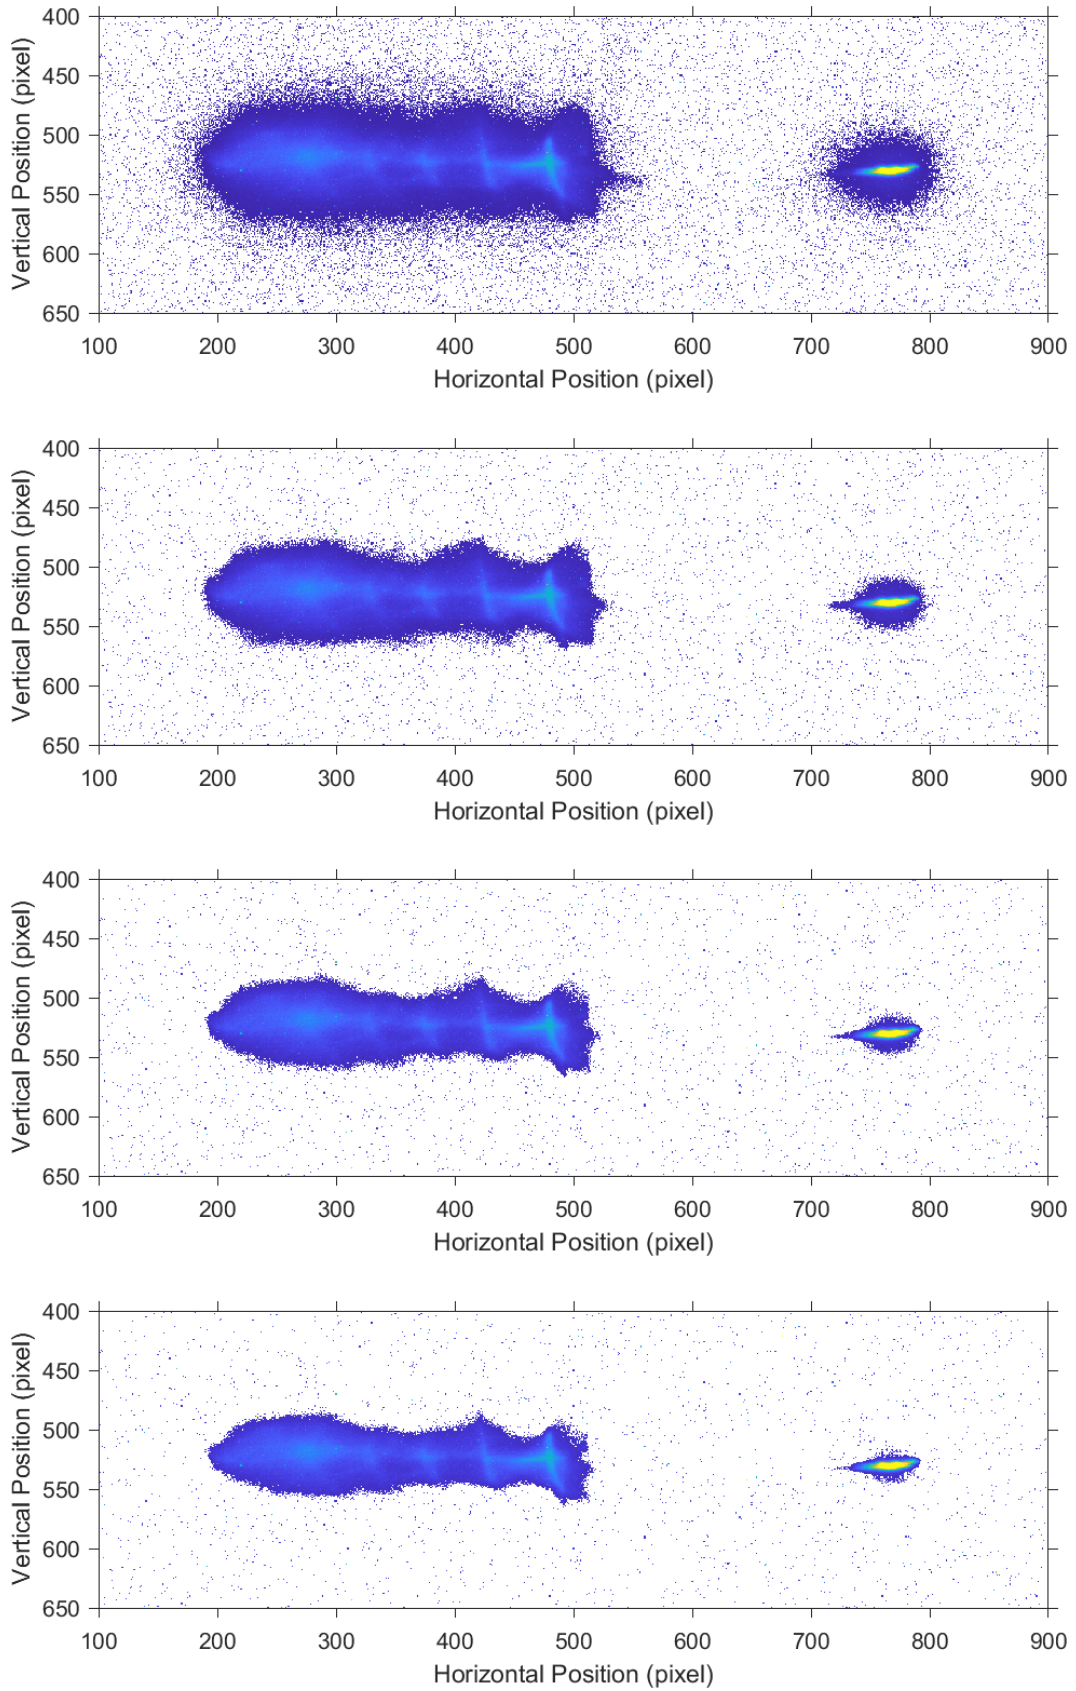


Figure 3.5: The drive (left) and witness (right) bunches shown after the first background subtraction stage. Pixels with values below the threshold value are set to zero. A threshold value of $x_T = \bar{x} + n \cdot x_{RMS}$ is used where $n = 0, 1, 2, 3$ from top to bottom.

Setting n equal to 3 was found to sufficiently speed up the background subtraction procedure of this spectrometer image while also leaving the drive and witness bunch signals unaffected. As images taken within the same run have similar beam parameters, the determined value of n is suitable for all spectrometer images within a given dataset. However, across multiple runs, it is possible that the value of n will need to be determined again.

3.3.2 Connected Component Analysis

A connected component refers to a collection of pixels, each of which being connected to at least one other pixel in the collection. Two pixels are part of the same collection if they are connected along the horizontal, vertical, or diagonal direction. The construction of these components allows for their comparison by properties such as total area or mean pixel value, leading to the identification of the drive and witness bunches among the background.

In the simplest scenario, as is the case for the spectrometer image mentioned thus far, the connected components can be listed in descending order based on their total area (see table 3.1).

Table 3.1: List of the total area and mean pixel value of the five largest connected components in the spectrometer image discussed above. Components are listed in descending order based on their total area.

| Total Area (pixel) | Mean Pixel Value (pixel) |
|--------------------|--------------------------|
| 16808 | 44.274 |
| 932 | 118.7 |
| 9 | 19.0 |
| 8 | 20.625 |
| 8 | 15.875 |

Removing all connected components other than the two components with the largest areas results in only the components corresponding to the drive and witness bunches left remaining (see figure 3.6).




Figure 3.6: Connected components with areas less than the two largest components are removed, resulting in only the drive and witness bunches left remaining.

An additional complication arises if the acceleration is low enough such that the drive and witness bunches are not sufficiently separated, resulting in them forming a single connected component. In this case, removing all components with areas less than the two largest components results in the combined drive and witness bunches left remaining, along with the next largest background component(s) as shown in figure 3.7.

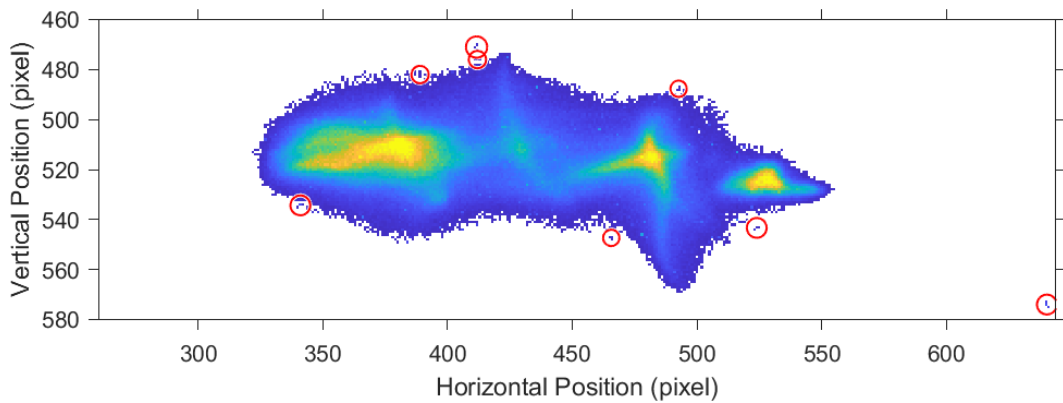


Figure 3.7: Connected components with areas less than the two largest components are removed, resulting in the combined drive and witness bunch component left remaining, along with multiple background components (circled in red) each having the second largest area.

A table listing the connected components in descending order based on their mean pixel value reveals that a number of components have the second largest area

of 3 pixels (see table 3.2).

Table 3.2: List of the total area and mean pixel value of the nine largest connected components in the spectrometer image shown in figure 3.7. Components are listed in descending order based on their mean intensity.

| Total Area (pixel) | Mean Pixel Value (pixel) |
|--------------------|--------------------------|
| 11547 | 93.88 |
| 3 | 28.333 |
| 3 | 22.667 |
| 3 | 22,0 |
| 3 | 20.667 |
| 3 | 18.0 |
| 3 | 16.333 |
| 3 | 15.667 |
| 3 | 15.333 |

The mean pixel values of the connected components can be used to determine whether the drive and witness form two distinct connected components. When they are separate, the largest component has a lower mean pixel value than the second largest component (as shown in table 3.1). When they are connected, the largest component also has the largest mean pixel value (as shown in table 3.2). With the spectrometer image now being classified as having a connected drive and witness bunch, all but the largest connected component can be removed (see figure 3.8).




Figure 3.8: Connected components with areas less than the largest component are removed, resulting in only the connected drive and witness bunches left remaining.

Identifying whether the drive and witness bunches are connected using this method relies on the assumption that the witness bunch has a larger mean pixel value than the drive bunch. Although this is true for every spectrometer image in the dataset discussed above, different runs with different beam parameters can result in the drive bunch having the larger mean pixel value. This is problematic as the largest connected component will also have the largest mean pixel value whether the drive and witness bunches are connected or not.

A dataset where the spectrometer images contain drive bunches with larger mean pixel values than the witness bunches is considered below. In the case that the drive and witness bunches form two distinct connected components, the two components with the largest areas are kept and all other components are removed (see figure 3.9).




Figure 3.9: A spectrometer image where the drive bunch has a larger mean pixel value than the witness bunch. Connected components with areas less than the two largest components are removed, resulting in only the drive and witness bunches left remaining.

A table listing the connected components in descending order based on their total area is used to calculate the product of the total area and mean pixel value for the two largest components to determine their ratio (see table 3.3). For this spectrometer image, the product of the area and mean pixel value for the largest and second largest components are found to be 907856 and 24491 respectively, resulting in a ratio of 37.

Table 3.3: List of the total area and mean pixel value of the two largest connected components in the spectrometer image shown in figure 3.9. Components are listed in descending order based on their total area. The product of the total area and mean pixel value is also provided for both components.

| Total Area (pixel) | Mean Pixel Value (pixel) | Product of Area and Pixel Value |
|--------------------|--------------------------|---------------------------------|
| 17441 | 52.053 | 907856 |
| 595 | 41.161 | 24491 |

Should the drive and witness bunches form a single connected component, removing all but the two largest area components will result in the combined drive and witness bunches left remaining, along with the next largest background component(s) as shown in figure 3.10.

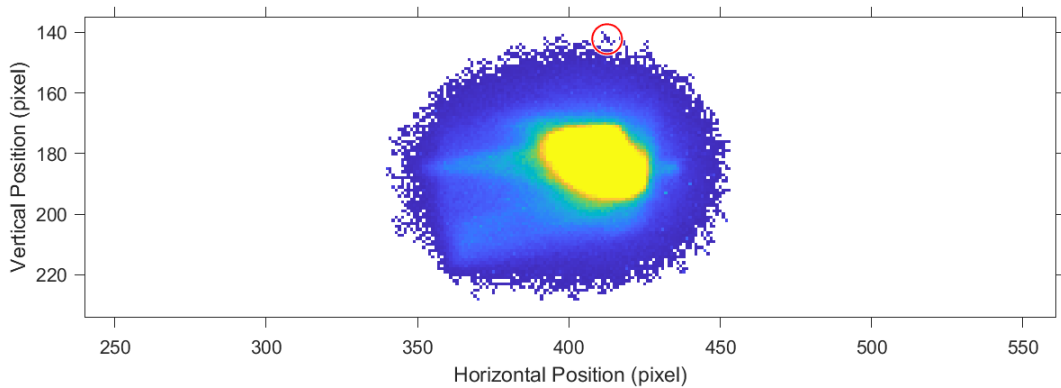


Figure 3.10: A spectrometer image from a dataset where the drive bunches have larger mean pixel values than the witness bunches. Connected components with areas less than the two largest components are removed, resulting in the combined drive and witness bunch component left remaining, along with a single background component (circled in red).

A table listing the connected components in descending order based on their total area is again used to calculate the product of the total area and mean pixel value for the two largest components to determine their ratio (see table 3.4). For this spectrometer image, the product of the area and mean pixel value for the largest and second largest components are found to be 1146202 and 97 respectively, resulting in a ratio of 11817.

Table 3.4: List of the total area and mean pixel value of the two largest connected components in the spectrometer image shown in figure 3.10. Components are listed in descending order based on their total area. The product of the total area and mean pixel value is also provided for both components.

| Total Area (pixel) | Mean Pixel Value (pixel) | Product of Area and Pixel Value |
|--------------------|--------------------------|---------------------------------|
| 6953 | 164.85 | 1146202 |
| 8 | 12.125 | 97 |

As the mean pixel value itself is no longer sufficient to classify whether the drive and witness bunches are separate or connected, the ratio of the products is used instead. Images where the drive and witness bunches form a single connected component have a much larger ratio than when the drive and witness bunches are sufficiently separated. This was verified by plotting the ratio for 200 spectrometer images that have been classified manually (see figure 3.11).

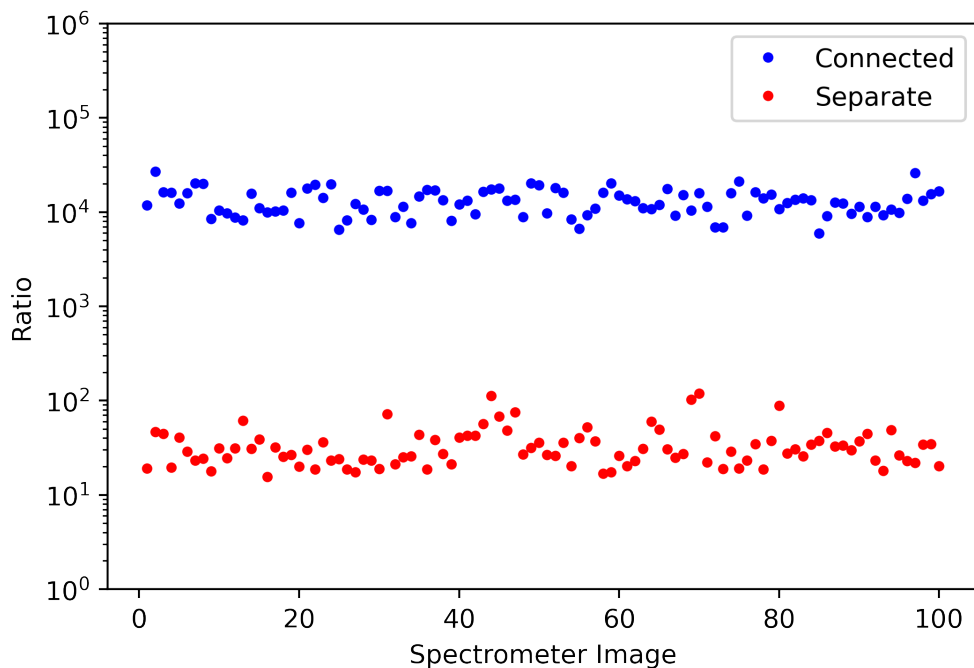


Figure 3.11: Ratio of the product of total area and mean pixel value for spectrometer images where the drive and witness bunches are connected (blue) or separate (red).

A criteria was defined whereby spectrometer images with a ratio above 10^3 are classified as having a connected drive and witness bunch, and those with a ratio

below 10^3 are classified as having sufficiently separated drive and witness bunches. With the recently considered spectrometer image now being classified as having a connected drive and witness bunch, all but the largest connected component can be removed (see figure 3.12).




Figure 3.12: Connected components with areas less than the largest component are removed, resulting in only the connected drive and witness bunches left remaining.

3.3.3 Light Leakage

Another source of background arises from light leakage between two adjacent spectrometer screens. This is particularly troublesome when the leakage overlaps with a signal bunch, resulting in a connected component consisting of both the signal bunch and the light leakage (see figure 3.13).

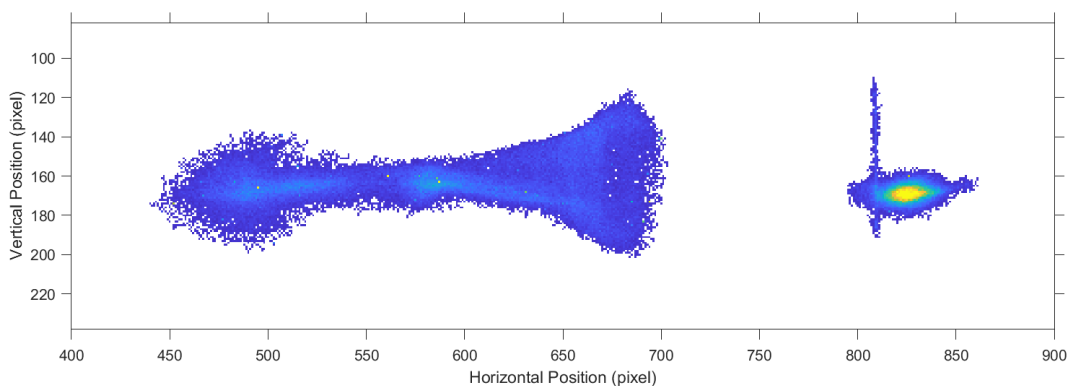


Figure 3.13: A spectrometer image where the light leakage overlaps with the witness bunch, resulting in them forming a single connected component.

When the acceleration is larger, the witness bunch is located further along the horizontal axis than the light leakage, thereby forming a distinct connected component (see figure 3.14). Such images can be used to construct a prediction of the how the light leakage will present itself on the spectrometer screen.




Figure 3.14: A spectrometer image where the light leakage is located between the drive and witness bunches, resulting in its own connected component.

The leakage component is first isolated by investigating the absolute value of the orientation for the three connected components with the largest major axis lengths (see table 3.5).

Table 3.5: List of the major axis length and orientation of the three connected components with the largest major axis lengths for the spectrometer image shown in figure 3.14. Components are listed in descending order based on their major axis length.

| Major Axis Length (pixel) | Orientation ($^{\circ}$) |
|---------------------------|----------------------------|
| 525.53 | -1.9605 |
| 347.65 | -88.568 |
| 178.45 | 2.3586 |

Components where the absolute value of the orientation is less than 85 are removed, resulting in an image consisting of only the leakage (see figure 3.15a). This is repeated to obtain a further nine such images (see figures 3.15b to 3.15j), allowing for a prediction of the leakage to be constructed (see figure 3.15k).

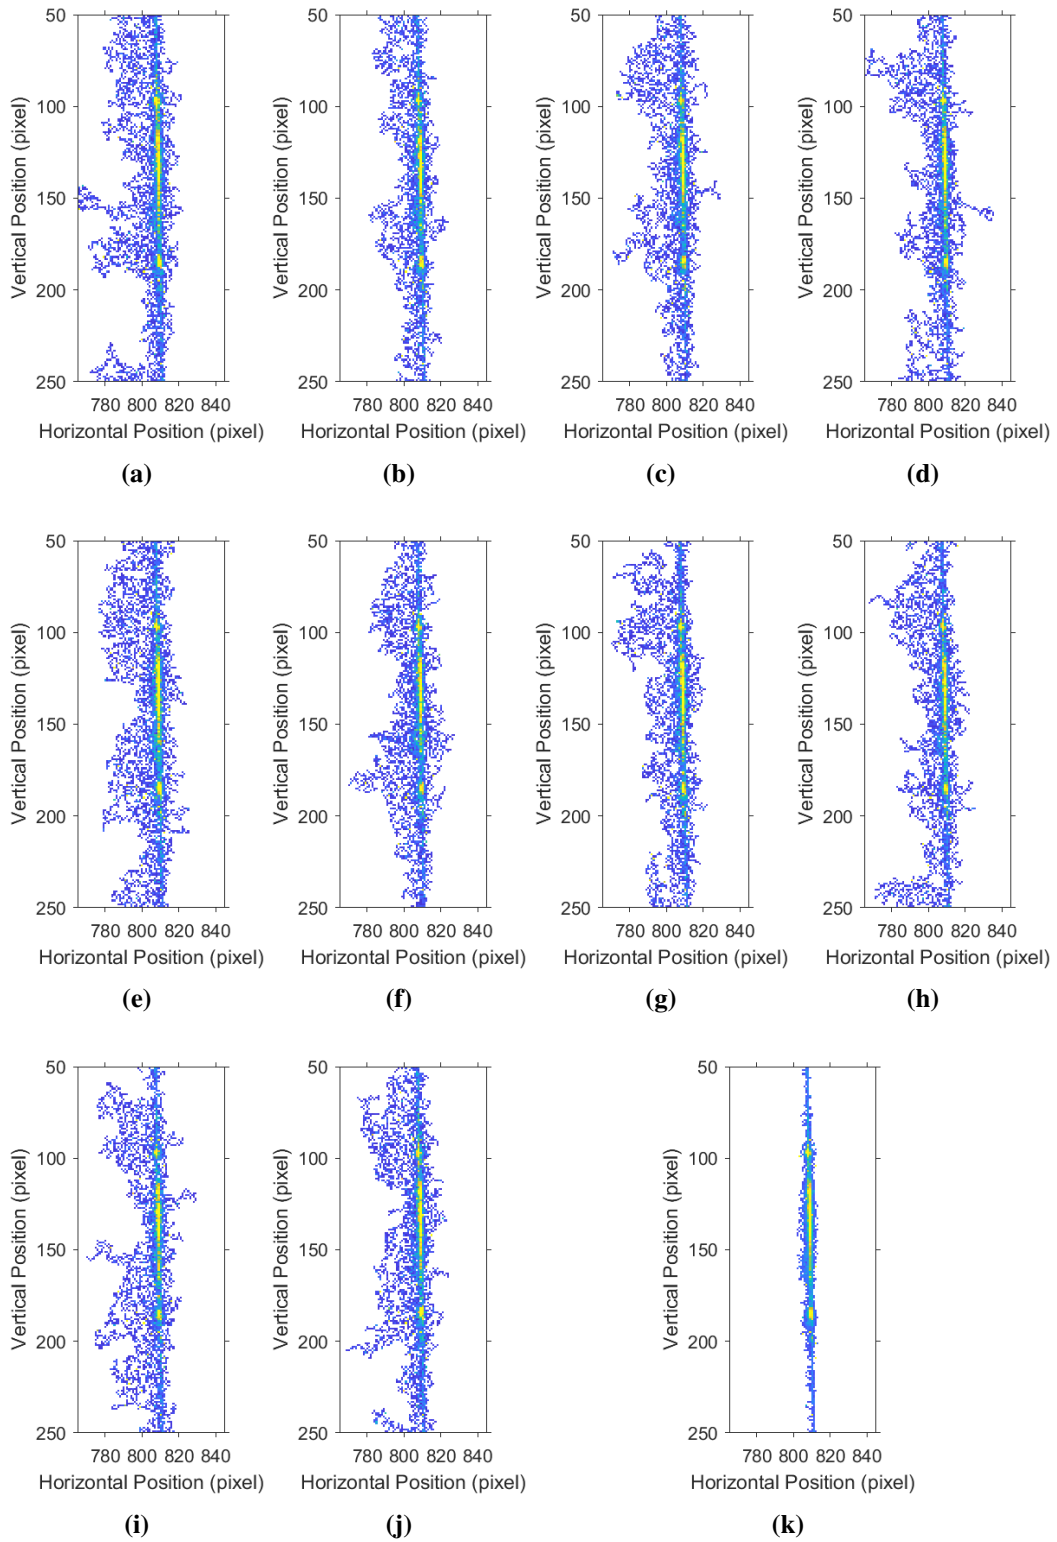


Figure 3.15: Isolated light leakage components from ten spectrometer images (a-j) are used to predict the light leakage in further spectrometer images (k).

The pixel threshold value is set to a lower value than usual to ensure that the entire leakage can be obtained for all ten of the spectrometer images. The predicted leakage image is constructed by averaging the ten images and removing pixels with a low count rate. This prediction can then be subtracted from any image where the leakage overlaps with a signal component. A demonstration of this subtraction using the spectrometer image shown in figure 3.13 is shown below (see figure 3.16).




Figure 3.16: Subtraction of the predicted light leakage component from the spectrometer image shown in figure 3.13.

3.3.4 Bunch Energies

Determining the energy of the drive and witness bunches requires first converting the horizontal position of the spectrometer screen into its corresponding energy, the spectrometer image shown in figure 3.6 is used to demonstrate this (see figure 3.17).




Figure 3.17: A spectrometer image where the x-axis is converted to energy in MeV.

This conversion is determined by calculating electron trajectories based on the dipole current recorded in the dataset, and therefore needs to be determined for each dataset. A plot showing how each column of pixels corresponds to energy in the dataset of the above spectrometer image is shown below (see figure 3.18). The trajectory model has been validated with calibration scans using non-interacted electron bunches of known energy.




Figure 3.18: Relationship between the horizontal position on the spectrometer screen and energy.

The bunch energies, μ_E , are then calculated for both the drive and witness bunches using the equation

$$\mu_E = \frac{\sum_i E_i \cdot dQ_i}{\sum_i dQ_i}, \quad (3.2)$$

where E_i is the energy corresponding to column i of the spectrometer image, and dQ_i is the sum of the pixel values within column i .

A histogram of the sum of pixel values for each energy is shown below (see figure 3.19). The minima of the histogram located between the first and last non-zero values is used to determine the range of columns used to calculate the drive bunch energy (shown in red), as well as the range of columns used to calculate the witness bunch energy (shown in blue). In the case that the drive and witness

bunches form a single connected component, the energy calculation is done using all columns to obtain the energy of the combined bunches.




Figure 3.19: Histogram of the sum of pixel values for each column of energy in the spectrometer image shown in figure 3.17. Bins corresponding to the drive and witness bunches are shown in red and blue respectively.

3.3.5 Reset Signature

The separation scan and reference datasets can be visualised using a waterfall plot. Each column of this plot is produced by rotating a spectrometer image 90° anti-clockwise before compressing it into a single column so that it can be inserted into the waterfall. A waterfall plot of the reference measurements corresponding to the dataset involving argon plasma with a backing pressure of 30.1 mbar and a high voltage discharge of 25 kV is shown below (see figure 3.20). It can be seen that the energy of both the drive and witness bunches steadily increases in the later shots, suggesting the presence of an unwanted drift in the electron energy during data acquisition.

Drift in the electron energy during a set of measurements presents a source of systematic uncertainty that must be accounted for. However, drift in the electron energy occurring between the separation scan dataset and reference dataset is particularly problematic as it inhibits their comparison. Consequently, the repetition





Figure 3.20: Waterfall plot of a set of reference measurements showing a drift in the measured energy of the drive (lower energy band) and witness (higher energy band) bunches.

rate is instead found by constructing reset signature plots (see figure 3.21) that compare the consistency of the final measurements in the separation scan. These plots show that as the time between discharge pulses increases, the energy transferred from the drive bunch to the witness bunch initially decreases as material is lost from the capillary. However, at longer times the energy transferred increases as the lost material is replenished, before a stable energy transfer is reached when the capillary is full.


The mean energy of the drive and witness bunches for each separation time are calculated for the purpose of establishing a reset signature. Outlier energies, defined as energies more than three scaled median absolute deviations from the median energy, are removed from the calculation of the mean. The uncertainty for each mean energy is initially (as seen in figure 3.21a) given by the standard error

$$\frac{\sigma_E}{\sqrt{n}}, \quad (3.3)$$


where σ_E and n are the standard deviation and the number of values of the energies comprising that mean energy.



(a) Witness (blue) and drive (red) bunch energies of the separation scan where the uncertainty is given by the standard error, the combined bunch energies (purple) are plotted where the witness and drive bunches can not be resolved (top). The reduced chi-squared statistics of horizontal fits to the drive and witness energies as a function of how many points are fitted (bottom).



(b) The reduced chi-squared statistics of horizontal fits to the drive and witness energies as a function of how many points are fitted. The uncertainty is estimated using the reference dataset.



(c) The reduced chi-squared statistics of horizontal fits to the drive and witness energies as a function of how many points are fitted. The uncertainty is inferred a posteriori using the reduced chi-squared statistics from figure 3.21a.

Figure 3.21: Energies of the witness and drive bunches as a function of time between pulses of the high voltage discharge, with the reduced chi-squared statistics calculated using statistical uncertainty (top), estimated using the reference dataset (middle), and inferred a posteriori (bottom).

For both the drive and witness energies, a horizontal line is fitted to the final k points, with k increasing from 2. The reduced chi-squared statistic of the drive and witness fits, χ_D^2 and χ_W^2 respectively, are calculated for each value of k . The largest value of k for which the reduced chi-squared statistics still indicate a good fit gives the upper limit on the repetition rate.

Due to the witness bunches carrying less charge than the drive bunches, the witness energy gain is larger than the drive energy loss. As a result, the witness bunch energies provide a more sensitive probe of the acceleration efficacy. Consequently, the witness energies are also more sensitive to drifts in the electron energy. As this source of systematic uncertainty has yet to be accounted for, the reduced chi-squared statistic of the witness fit suggests a poor fit to the data. Nonetheless, it can be seen for this particular separation scan that a drastic increase in the reduced chi-squared statistics occur when performing the fit with the final five points, indicating that the energies are relatively consistent among the final four points and that the ideal conditions for a subsequent acceleration (where the state of the plasma capillary is the same as it was for the previous acceleration) may be achieved within 2 ms. Furthermore, it is unlikely that the final eight points are instead what represent these conditions, as the fluctuation between each of these adjacent points is on average much smaller than the energy increase between the fourth-to-last and fifth-to-last point.

Investigations were made into whether a more representative uncertainty could be determined through comparison of the largest time delay of the separation scan with the reference dataset. Assuming that the ideal acceleration conditions are reestablished within 20 ms, any observed difference should arise entirely from the combined statistical and systematic uncertainties. The relative change in the mean energies between these two experimental configurations was taken as the corrected uncertainty on the mean energies of the separation scan, allowing for the reduced chi-squared statistics to be recomputed (see figure 3.21b). It can be seen that the

reduced chi-squared statistics now suggest that the model is over-fitting to the data, indicating that the errors are overestimated. This is likely because the drift in electron energy between the separation scan and reference measurement is larger than the drift in electron energy between measurements of different time delays within the separation scan.

An alternative approach of accounting for drifts in the electron energy is to use the reduced chi-squared statistics to obtain corrected uncertainty estimates, $\hat{\sigma}$, a posteriori. In the dataset considered here it is clear that the ideal acceleration conditions are reestablished within at least 5 ms, as this is the time between high voltage discharge pulses which corresponds to the largest witness bunch energy gain. As a horizontal fit between the final three points must model the data well, the corrected uncertainty estimates can be given by


$$\hat{\sigma} = \sqrt{\sigma^2 \chi^2}, \quad (3.4)$$

where σ are the initial uncertainty estimates given by the standard error, and χ^2 is the reduced chi-squared statistic calculated from a fit of the final three points. Applying this method and recomputing the reduced chi-squared statistics (see figure 3.21c) indicates that the horizontal fit models the data well, further suggesting that the corrected uncertainty estimates are representative of the true error. It can now be seen that performing the fit with the final four points results in $\chi_W^2 \lesssim 1$ and $\chi_D^2 \lesssim 1$, confirming with additional rigour that the ideal conditions for a subsequent acceleration are achieved within 2 ms.


3.4 Results

The method of estimating the uncertainties a posteriori is applied to all datasets and the time taken for the plasma to achieve the same state for the acceleration of a subsequent electron bunch is determined for each of the experimental configurations. Due to having less sensitivity, the reset signature is occasionally achieved

earlier by the drive bunch than the witness bunch, therefore the reset signature of the witness bunch is used to determine the reset time. The final waterfall and reset signature plots corresponding to the argon plasma configurations are shown below in descending order of pressure, with equal pressures further sorted in ascending order of voltage (see figures 3.22 to 3.32). Additionally, the plots corresponding to the hydrogen plasma configuration are also shown (see figure 3.33).




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ^2_W (middle) and χ^2_D (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.22: Data taken in an argon plasma capillary at 40.0 mbar backing pressure using HV discharge pulses of 15 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.23: Data taken in an argon plasma capillary at 40.0 mbar backing pressure using HV discharge pulses of 20 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ^2_W (middle) and χ^2_D (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.24: Data taken in an argon plasma capillary at 30.1 mbar backing pressure using HV discharge pulses of 15 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.25: Data taken in an argon plasma capillary at 30.1 mbar backing pressure using HV discharge pulses of 20 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.

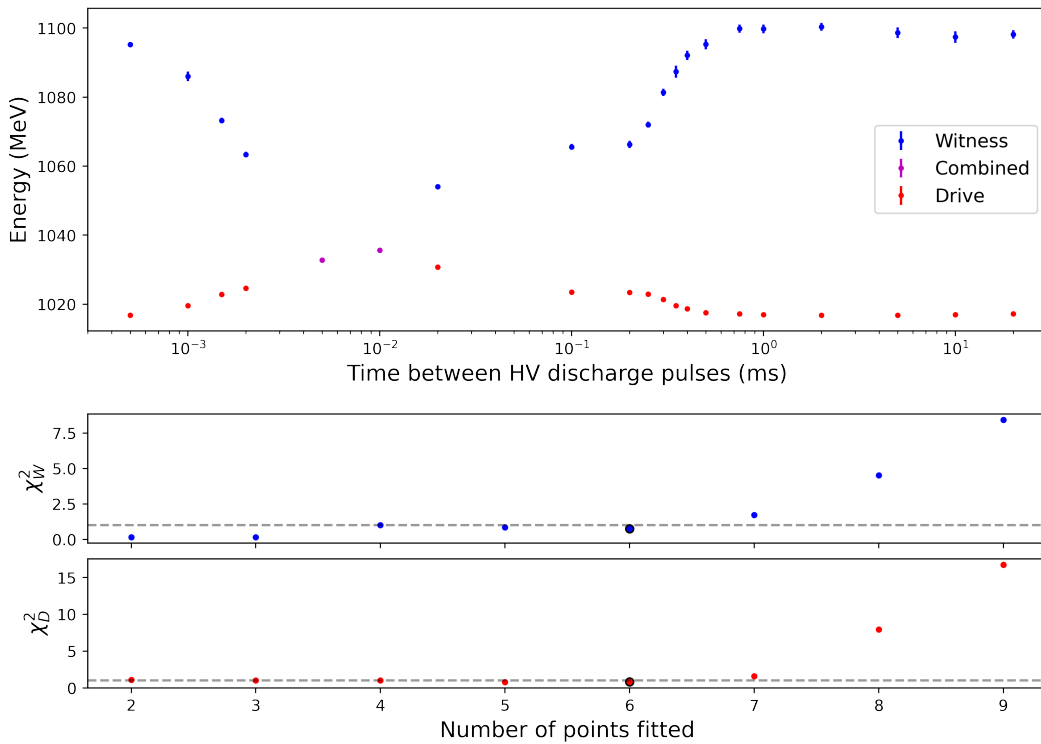


(b) Bunch energies for the separation scan (top). χ^2_W (middle) and χ^2_D (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.26: Data taken in an argon plasma capillary at 30.1 mbar backing pressure using HV discharge pulses of 25 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ^2_W (middle) and χ^2_D (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.27: Data taken in an argon plasma capillary at 20.6 mbar backing pressure using HV discharge pulses of 15 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.28: Data taken in an argon plasma capillary at 20.6 mbar backing pressure using HV discharge pulses of 20 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.29: Data taken in an argon plasma capillary at 20.6 mbar backing pressure using HV discharge pulses of 25 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.30: Data taken in an argon plasma capillary at 15.4 mbar backing pressure using HV discharge pulses of 15 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.31: Data taken in an argon plasma capillary at 15.4 mbar backing pressure using HV discharge pulses of 20 kV.




(a) Waterfall plot of the separation scan and reference measurements with the mean energy of the electron bunches and their standard error overlaid.




(b) Bunch energies for the separation scan (top). χ_W^2 (middle) and χ_D^2 (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.32: Data taken in an argon plasma capillary at 15.4 mbar backing pressure using HV discharge pulses of 25 kV.



(a) Waterfall plot of the separation scan with the mean energy of the electron bunches and their standard error overlaid.



(b) Bunch energies for the separation scan (top). χ^2_W (middle) and χ^2_D (bottom) against number of points fitted are plotted, with the reset time circled in black.

Figure 3.33: Data taken in a hydrogen plasma capillary at 17.5 mbar backing pressure using HV discharge pulses of 27 kV.

3.5 Discussion

The reset signature of the drive bunch is achieved either earlier or at the same time as the reset signature of the witness bunch for each configuration, indicating that while the drive bunch signature is less sensitive it still provides results consistent with the witness bunch signature. The reset times for each configuration of argon plasma are shown below (see figure 3.34), where the reset signature of the witness bunch is used to determine the reset time due to its higher sensitivity.

It can be seen that having too high of a backing pressure causes the reset time to increase substantially. Similarly, having too low of a pressure also causes a slight increase in the reset time. Changes in the discharge voltage have a less significant effect, however it can still be seen that higher voltages sometimes result in faster reset times. These results suggest that there is an optimal pressure value, between 15.4 and 30.1 mbar in the case of argon, where the reset time is minimised.




Figure 3.34: Measured reset times for different experimental configurations of the argon plasma capillary. The grey cell corresponds to a configuration where no measurements were taken.

Through consideration of the reset times in isolation, an alternative interpretation is that the reset dynamics have converged for pressures at and below 20.6 mbar, with any differences in the corresponding measurements being statistically insignificant. However, inspection of the reset curves (see figure 3.35) shows vastly different reset dynamics at 15.4 mbar (red dotted lines) when compared to 20.6 mbar (blue dotted lines), as indicated by the gap between these curves at smaller discharge pulse delays. This supports the claim that these configurations can be distinguished statistically despite the similarity between their reset times.




Figure 3.35: Reset curves indicating the energy of the witness and drive bunches, and therefore the readiness of the plasma for acceleration, as a function of time between the current and previous discharge pulses. The curves are labelled according to the backing pressure of the configuration.

Further inspection of the reset curves reveals that configurations with higher pressures take longer for the acceleration to begin. This suggests that the reset time is largely affected by the expulsion of material, as with more material to expel it takes a longer time before the capillary can begin to refill. The steeper gradients on the higher pressure curves indicate that the capillary refills faster at these higher pressures. However, despite this faster refill, the observation that higher pressure configurations have an overall slower reset time shows that it is the delay before the

refill can begin rather than the time taken to refill that dominates the reset time. This indicates material expulsion rather than material refill to be the dominant limiting mechanism, contrary to initial expectations.

The observed decrease in reset time at higher discharge voltages likely follows from a decreased expulsion time due to a stronger discharge, allowing for refill of the capillary to begin sooner. This is supported by the reset curves for 30.1 mbar and 40.0 mbar (see figure 3.36), where the curves corresponding to a larger voltage begin earlier in time. Although the reset curves for 20.6 mbar each begin at the same time of 0.02 ms, the higher voltage curves show more energy gain at this time, indicating that more material refill has already occurred. The plasma density remains large enough to always accelerate the second electron bunch in each of the 15.4 mbar configurations, therefore no comment can be made about how the time delay for acceleration to begin depends on the discharge voltage. However, it can be seen that the choice of discharge voltage is less significant for smaller pressures than for larger pressures, likely because any parameter that affects expulsion will be less pronounced the less material there is to expel.




Figure 3.36: Reset curves indicating the energy of the witness bunch as a function of time between the current and previous discharge pulses for the 15.4 mbar (far left), 20.6 mbar (middle left), 30.1 mbar (middle right), and 40.0 mbar (far right) configurations. The curves are labelled according to the voltage parameter of the configuration.

The fastest measured reset time of the argon plasma capillary, corresponding to

the 20.6 mbar configurations with 20 kV and 25 kV high voltage discharge pulses, was determined to be 0.50 ms. This corresponds to a repetition rate of 2.0 kHz, demonstrating that the acceleration of a subsequent electron bunch can be repeated at a rate of $\mathcal{O}(\text{kHz})$. The measured reset time of the hydrogen plasma configuration was determined to be 0.22 ms, corresponding to a repetition rate of 4.55 kHz. The lower mass of hydrogen molecules in comparison to argon molecules, resulting in faster expulsion and refill, provides a reasonable explanation for why a higher repetition rate is achieved using hydrogen.

It should be noted that as this study involves measurements of only two consecutive acceleration events, it is still unclear if further subsequent acceleration events can be replicated with the same repetition rate as is reported from these measurements. Further studies involving the subsequent acceleration of multiple electron bunches will be critical for establishing the rate limitations of consistent and sustained plasma wakefield acceleration.

On the other hand, accelerating multiple electron bunches within the same plasma lifetime, as opposed to generating a new plasma for each acceleration event as was done here, allows for further increase of the average power of the accelerator. The European X-ray Free Electron Laser [53] is expected to provide bunch trains of 2700 bunches with a train repetition rate of 10 Hz and a 220 ns inter-bunch separation. The International Linear Collider [54] is expected to provide bunch trains of 1312 bunches with a train repetition rate of 5 Hz and a 554 ns inter-bunch separation. Similarly, a mode of operation for plasma wakefield accelerators should be considered where bunch trains are provided with a train repetition rate of 4.55 kHz and with an inter-bunch separation of 63 ns. In this case, the train repetition rate (determined in this chapter) is limited by the expulsion and refill of material in the capillary, and the inter-bunch separation is limited by the ion motion of the plasma following the previous acceleration event (as discussed in section 2.1.4). Further study is needed to determine how many bunches would be within each bunch train.

Chapter 4

Simulation of Gas Filling in the FLASHForward Capillary

The effects contributing towards the limitation on the maximum achievable repetition rate are not well understood. In particular, it is unclear how the expulsion and the refill of material individually affect the repetition rate, as well as how these effects depend on parameters such as the backing pressure. Simulation of the entry of gas into the plasma capillary is used to provide a more complete understanding of the reset curve dynamics presented in the previous chapter.

4.1 Motivation

Initially it was assumed that the repetition rate of plasma wakefield acceleration was ultimately limited by the time taken to replenish lost material within the capillary, and so it was expected that higher backing pressures would increase the repetition rate. However, the experimental measurements performed in the previous chapter show that increasing the backing pressure can in fact decrease the repetition rate.

In order to gain further insight on how the experimental measurements should be interpreted, it would be useful to understand the dynamics involved with the filling of the capillary. Any discrepancy between this understanding and the experimental measurements would indicate the effect of other processes on the repetition

rate, either validating the conclusions made in the previous chapter or prompting their reassessment.

4.2 Simulation

Ansys Fluent [55] was used to simulate the filling of the FLASHForward capillary. Several equations describing the fluid dynamics of the system are solved computationally, thereby enabling the behaviour of the fluid flow to be described (see section 4.2.1). The geometry of the FLASHForward capillary was accurately modelled, allowing for simulations to be performed in alignment with the experimental configuration (see section 4.2.2). Finally, the pressure at the midpoint of the capillary as a function of time was determined to be an appropriate metric to gauge the filling of the capillary, and is therefore selected as the simulation output (see section 4.2.3).

4.2.1 Theory

The first of these equations is the continuity equation which ensures that mass is neither created nor destroyed, its general form is given by

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = S_m \quad (4.1)$$

where ρ is the fluid density, t is time, \vec{v} is the velocity of the fluid, and S_m is a source term representing mass addition. This equation states that the change in density at a given point in space is determined by the net flow of mass into or out of that point, accounting for any additional mass sources and sinks.

The second equation is the conservation of momentum equation and is given by

$$\frac{\partial}{\partial t}(\rho \vec{v}) + \nabla \cdot (\rho \vec{v} \otimes \vec{v}) = -\nabla p + \nabla \cdot (\bar{\bar{\tau}}) + \rho \vec{g} + \vec{F} \quad (4.2)$$

where ρ is the fluid density, t is time, \vec{v} is the velocity of the fluid, p is the static

pressure, $\bar{\tau}$ is the stress tensor, \vec{g} is the acceleration due to gravity, and \vec{F} represents external forces. This ensures that momentum within the fluid is conserved, taking into account any forces acting on the fluid due to the pressure gradient, internal friction, gravity and any potential external forces.

The third equation, used to model heat transfer, is the conservation of energy equation and is given by

$$\frac{\partial}{\partial t}(\rho E) + \nabla \cdot (\vec{v}(\rho E + p)) = \nabla \cdot \left(k_{\text{eff}} \nabla T - \sum_j h_j \vec{J}_j + (\bar{\tau}_{\text{eff}} \cdot \vec{v}) \right) + S_h \quad (4.3)$$

where ρ is the fluid density, t is time, E is the total energy, \vec{v} is the velocity of the fluid, p is the static pressure, k_{eff} is the effective thermal conductivity, T is the temperature, h_j is the specific enthalpy of species j , \vec{J}_j is the diffusion flux of species j , $\bar{\tau}_{\text{eff}}$ is the effective stress tensor, and S_h is a source term representing external heat sources. This equation shows that the change in energy density at a given point depends on the convective flux of total energy, as well as the conductive heat flux, internal heat sources and sinks, heat produced via internal friction, and any potential external heat sources and sinks.

The realizable $k - \varepsilon$ model [56] is used to model turbulent flow within the system, resulting in an additional two transport equations to be solved. The first of these is used to account for turbulent kinetic energy (k), the energy associated with turbulent fluctuations in the flow, and is given by

$$\begin{aligned} \frac{\partial}{\partial t}(\rho k) + \frac{\partial}{\partial x_j}(\rho k u_j) &= \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] \\ &+ G_k + G_b - \rho \varepsilon - Y_M + S_k \end{aligned} \quad (4.4)$$

where ρ is the fluid density, t is time, x_j represents the spatial coordinate in the j -th direction, u_j represents the velocity component in direction x_j , μ is the dynamic viscosity of the fluid, μ_t is the turbulent viscosity, σ_k is the turbulent Prandtl number for turbulent kinetic energy, G_k represents the generation of turbulent kinetic en-

ergy due to velocity gradients in the fluid, G_b represents the generation of turbulent kinetic energy due to buoyancy forces, ε is the dissipation rate of turbulent kinetic energy, Y_m represents the dissipation of turbulent energy caused by density fluctuations in compressible turbulence, and S_k is a source term representing external turbulence sources. This means that the change in turbulent kinetic energy depends on the transport, diffusion, generation, dissipation and damping of turbulent kinetic energy, alongside additional contributions from external sources.

The second of these equations is used to account for the dissipation rate of turbulent kinetic energy (ε), the rate at which turbulent kinetic energy is converted into internal energy, and is given by

$$\begin{aligned} \frac{\partial}{\partial t}(\rho\varepsilon) + \frac{\partial}{\partial x_j}(\rho\varepsilon u_j) = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_\varepsilon} \right) \frac{\partial \varepsilon}{\partial x_j} \right] + \rho C_1 S \varepsilon \\ - \rho C_2 \frac{\varepsilon^2}{k + \sqrt{\nu \varepsilon}} + C_{1\varepsilon} \frac{\varepsilon}{k} C_{3\varepsilon} G_b + S_\varepsilon \end{aligned} \quad (4.5)$$

where ρ is the fluid density, t is time, x_j represents the spatial coordinate in the j -th direction, u_j represents the velocity component in direction x_j , μ is the dynamic viscosity of the fluid, μ_t is the turbulent viscosity, σ_ε is the turbulent Prandtl number for the dissipation rate, C_1 is a coefficient that adjusts the turbulence dissipation based on the local strain rate, S is the mean strain rate, k is the turbulent kinetic energy, ν is the kinematic viscosity, G_b represents the generation of turbulent kinetic energy due to buoyancy forces, S_ε is a source term representing additional influences on the dissipation rate, and C_2 , $C_{1\varepsilon}$, and $C_{3\varepsilon}$ are constants established to ensure that the model performs well. This means that the change in dissipation rate depends on the transport, diffusion, generation and dissipation of the dissipation rate, alongside additional contributions from external sources.

Each of the above continuous equations can be discretised in time and space. Discretisation in time is achieved through the application of a time integration scheme to integrate each term of the differential equations over a time step Δt , and

discretisation in space is achieved by dividing the domain into a finite number of volumes. Via discretisation, a set of algebraic equations representing the conservation laws are obtained from the continuous equations. These algebraic equations can then be solved using numerical methods until the solution converges, meaning that the solution no longer changes significantly from one iteration to the next.

4.2.2 Configuration

To simulate the gas flow through the FLASHForward capillary, it is necessary to first define its geometry. A model was created (see figure 4.1a) consisting of a 50 mm long capillary with a radius of 0.75 mm. Additional tubes with a radius of 1.5 mm are included towards each end of the capillary, representing the inlet pipes used in the FLASHForward experiment. The model is then meshed (see figure 4.1b) in order to discretise the spatial domain, allowing for the comparatively complex geometry to be described by a collection of simpler volumes.




Figure 4.1: 3D models of the FLASHForward capillary before meshing (a) and after meshing (b).

4.2.3 Output

The choice of gas species and the rate at which the gas flows into the inlets are provided as input parameters. A demonstration of the simulation is examined using

hydrogen gas and a flow rate of $8 \times 10^{-7} \text{ kg s}^{-1}$. An assessment of how full the capillary is at a given moment in time can be obtained by inspecting the pressure within it. The pressure as a function of the longitudinal position along the length of the capillary is plotted (see figure 4.2) for a variety of times between 0.01 ms and 1 ms, where each pressure value is taken at the centre of the circular cross-section of the capillary.

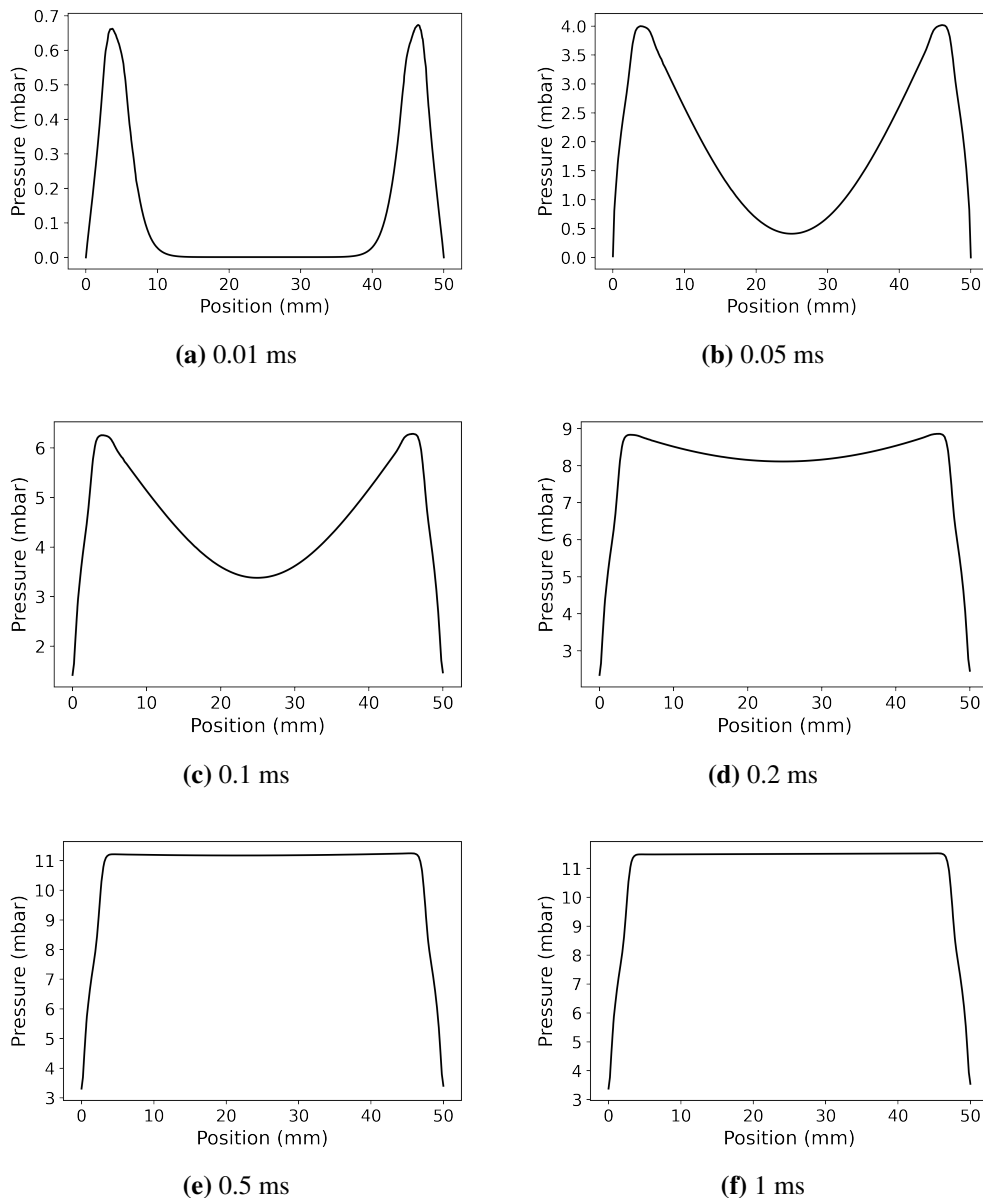


Figure 4.2: Simulation results corresponding to $8 \times 10^{-7} \text{ kg s}^{-1}$ of hydrogen gas entering the capillary. The pressure as a function of the longitudinal position along the length of the capillary is plotted for a variety of times between 0.01 ms (a) and 1 ms (f).

It can be seen that at smaller times the highest pressure is at the points of intersection between the inlet pipes and the capillary, whereas the lowest pressure is at the midpoint of the capillary. As the time increases the pressure at the midpoint of the capillary increases until it is equal to the pressure at the intersection points. This configuration achieves equilibrium after approximately 1 ms, and is limited by the time taken for the midpoint of the capillary to fill.

As the midpoint of the capillary is the last position to be filled, the pressure at this position can be plotted as a function of time (see figure 4.3). This provides a gauge of how full the capillary is at any given moment, and is obtained as an output from each simulated run.




Figure 4.3: Simulation results corresponding to $8 \times 10^{-7} \text{ kg s}^{-1}$ of hydrogen gas entering the capillary. The pressure at the midpoint of the capillary as a function of time is plotted.

4.3 Gas Filling in the FLASHForward Capillary

The relationship between the flow rate used in the simulation and the pressure used in the FLASHForward experiment must be found (see section 4.3.1). Once this is done, simulations that correspond to the experimental configurations discussed in

the previous chapter can be performed (see section 4.3.2).

4.3.1 Flow Rate Calibration

Experimental data was taken at backing pressures of 15.4, 20.6, 30.1 and 40.0 mbar for argon, as well as 17.5 mbar for hydrogen. Since the capillary opens to a vacuum, the pressure inside the capillary is lower than the backing pressure, and it is necessary to know the relationship between the backing pressure set during the experiment and the capillary pressure seen in the simulation. This information is readily available in the context of the FLASHForward capillary and gas system [57], and has been plotted below (see figure 4.4). A quadratic function is fitted to the data so that the capillary pressure resulting from a given choice of backing pressure can be obtained.






Figure 4.4: Relationship between the backing pressure of the FLASHForward buffer system and the capillary pressure after achieving equilibrium for both argon and hydrogen.

The relationship between the capillary pressure and the flow rate needed to achieve such a pressure can be determined by running the simulation with different flow rate inputs and then outputting the pressure once equilibrium is achieved (see figure 4.5a for argon and figure 4.5b for hydrogen). As it is only the equilib-

rium state that is of interest here, a steady-state simulation is performed rather than modelling the evolution of the system over time as was done before. A quadratic function is fitted to the data so that the flow rate needed to achieve a specific capillary pressure can be obtained.



(a) Argon



(b) Hydrogen

Figure 4.5: Relationship between the capillary pressure at equilibrium and the flow rate for both argon (a) and hydrogen (b).

The flow rate corresponding to a given gas species and backing pressure can then be determined by first using the fits from figure 4.4 to convert backing pressure into capillary pressure, and then using the fits from figure 4.5 to convert capillary pressure into flow rate. The flow rates corresponding to the experimental configurations discussed in the previous chapter are given below (see table 4.1).

Table 4.1: Input flow rates for the gas filling simulation to correspond to a given backing pressure and gas species in the FLASHForward experiment.

| Gas Species | Backing Pressure [mbar] | Flow Rate [kg/s] |
|-------------|-------------------------|-----------------------|
| Hydrogen | 17.5 | 1.68×10^{-7} |
| Argon | 15.4 | 1.03×10^{-6} |
| Argon | 20.6 | 1.56×10^{-6} |
| Argon | 30.1 | 2.70×10^{-6} |
| Argon | 40.0 | 4.15×10^{-6} |

4.3.2 Gas Filling Curves

The flow rates determined above were used as input parameters to simulate gas filling in the FLASHForward capillary. As different flow rates result in different equilibrium pressures, the pressure of the capillary is normalised to represent the fraction filled, allowing for ease of comparison. The fraction filled is defined as

$$\frac{P(t)}{P_{eq}} \quad (4.6)$$

where $P(t)$ is the pressure at the midpoint of the capillary at time t , and P_{eq} is the pressure at the midpoint after achieving equilibrium. Values for P_{eq} are obtained from steady-state simulations.

4.4 Discussion

The gas filling curves are plotted for each of the experimental configurations (see figure 4.6), and the reset curves shown in the previous chapter are presented again

(see figure 4.7) to facilitate comparison between the simulation and experimental data.




Figure 4.6: Simulation results showing gas filling curves for different gas species and backing pressures.




Figure 4.7: Experimental results showing reset curves for plasma wakefield acceleration in argon at varying backing pressures.

It can be seen from the gas filling curves that having a higher backing pressure results in faster filling of the capillary. This is supported by the gradient of the reset curves being steeper for higher pressures, which suggests that when the capillary does start to re-fill it can do so faster when the backing pressure is larger. The experimental observation that having too low a backing pressure causes an increase in the reset time can be explained by the increased time it takes for the gas filling to complete.

The gas filling curves show that the capillary can be considered full in less than 3 ms for each configuration. However, the reset curves show that for larger pressures it can take up to 10 ms to achieve the same acceleration signature. This supports the claim that the material expulsion mechanism itself contributes significantly to the achievable repetition rate, as gas filling alone can not account for the observed reset time.

The experimental measurements indicate that a reset time of 0.5 ms is achievable with a backing pressure of 20.6 mbar. However, the corresponding gas filling curve shows that the capillary can not completely fill within this time. This suggests that not all material is expelled at lower backing pressures, and that the gas filling starts with a significant amount of gas already present in the capillary. This is supported by the experimental observation that at the lowest tested backing pressure the plasma density remains high enough to accelerate the second electron bunch at all times.

In summary, the discrepancy between the gas filling simulation and the reset time measurements indicates that the expulsion of material plays a dominant role in determining the repetition rate of plasma wakefield acceleration. Agreement would suggest that each acceleration event is limited by the time it takes to fill an empty capillary with the expulsion process itself having no interference due to expelling all material within a negligible amount of time, an assumption previously consid-

ered to be reasonable. The disagreement highlights that for low backing pressure configurations the assumption of having to fill an empty capillary is false, and for high backing pressure configurations the assumption of expulsion not interfering with the gas refill is false.

Further understanding of the factors that determine the repetition rate of plasma wakefield acceleration could be obtained by simulating the effect of the high voltage discharge on the plasma, thereby modelling the expulsion of the material. Considering this alongside the gas filling study presented in this chapter would provide a better understanding of the physics processes involved, and a comparison between these further simulation studies and the experimental measurements presented in the previous chapter would confirm if the underlying processes have been fully understood.

Chapter 5

Design of a Dark Photon Search

Experiment

The search for a hypothetical hidden sector particle known as the dark photon may provide an interesting application for the recent advances in plasma wakefield acceleration. A new search is proposed using the AWAKE facility to provide a high number of electrons on target at a reasonably large energy. Simulations are performed to optimise the experimental configuration and assess its discovery potential.

5.1 Dark Photons

The Standard Model of particle physics [1] is a gauge quantum field theory based on the local $SU(3) \times SU(2) \times U(1)$ gauge symmetry. Despite describing the strong, weak, and electromagnetic forces remarkably well, there are several deficiencies that demand the existence of additional physics beyond the Standard Model. Of particular note, it provides no explanation for the origin of its symmetry, meaning that a more complete theory could very well include additional gauge interactions. Furthermore, it offers no insight into the nature of dark matter [5], an unknown substance known to make up more than 80% of the matter in the universe.

One interesting avenue for the discovery of physics beyond the Standard Model arises from the potential existence of a hidden sector, an additional sector of particles that are not charged directly under the forces of the Standard Model

but rather under at least one new gauge group [58–61]. Particles of this hidden sector may provide suitable candidates for dark matter, as they could experience gravitational interactions while lacking strong and electromagnetic ones. However, if their only interaction with our visible sector is through the gravitational force, they are unlikely to ever be detected. As a result, we must disregard the idea that all interactions are described by a gauge theory and instead rely on the existence of a portal through which the dark and visible sectors can interact.

5.1.1 Portal Interaction

The form such a portal would take can be classified according to its mediator field, with the best motivated cases being comprised of the vector (spin 1), neutrino (spin 1/2), Higgs (scalar), and axion (pseudo-scalar) mediators. Of particular interest is the vector portal as it involves one dark and one visible abelian gauge boson, the visible boson being the photon of the U(1) gauge group and the dark boson being a new gauge boson known as the dark photon, denoted A' , which arises from an additional U(1) symmetry. This is significant because the multiplication of the field strengths of two abelian gauge fields result in a dimension four operator, allowing the two gauge bosons to mix kinetically. The portal operator that gives rise to this kinetic mixing is given by

$$\frac{\varepsilon}{2} F_{\mu\nu} F^{\mu\nu'}, \quad (5.1)$$

where $F_{\mu\nu}$ is the photon field strength, $F^{\mu\nu'}$ is the dark photon field strength, and ε is the mixing parameter that determines the strength of the kinetic mixing [62].

Experimental efforts to detect the dark photon usually assume that the dark photon has mass, as this provides a direct coupling between the dark photon and the current of ordinary Standard Model matter. Further distinctions regarding the experimental approach are made based on the value of the dark photon mass, as if $m_{A'} > 2m_e$ the visible decay channel into Standard Model particles becomes available. The decay into hidden sector states χ can also occur but through an invisible

decay channel provided that $m_{A'} > 2m_\chi$, and this channel may even dominate if the coupling to the hidden sector states is large enough or if the visible decay channel is kinematically inaccessible.

Interaction between the massive dark photon and particles of the Standard Model is given by

$$-\epsilon e J^\mu A'_\mu, \quad (5.2)$$

where e is the coupling of the visible photon to the visible sector, J^μ is the electromagnetic current, and A'_μ is the dark photon field [63]. As the dark photon has a direct coupling to the electromagnetic current, it can be produced via the same mechanisms as the visible photon.

5.1.2 Production & Decay

The main production channels are:

- **Bremsstrahlung:** Electrons incident on a target material with charge Z are scattered, becoming off shell and emitting a dark photon. This production mechanism is utilised in beam-dump experiments, where the dark photon carries most of the beam energy and the electron is scattered at a much larger angle. This interaction in an electron beam experiment is given by $e^- Z \rightarrow e^- Z A'$.
- **Annihilation:** An electron-positron pair annihilates, producing one visible and one dark photon ($e^+ e^- \rightarrow \gamma A'$) in the non-resonant case and just a dark photon ($e^+ e^- \rightarrow A'$) in the resonant case. This production mechanism is favoured for detecting invisible decay channels if using either an $e^+ e^-$ collider experiment or a positron beam-dump experiment.
- **Meson decay:** Low mass dark photons may be produced in meson decays [64] as given by $M \rightarrow \gamma A'$, where the meson M is a π^0, η, η', K or D . They may also be produced in the rare meson decays $K \rightarrow \pi A', \phi \rightarrow \eta A'$, and

$D^* \rightarrow D^0 A'$. This production mechanism is often utilised in hadronic environments where mesons are produced in large quantities.

- Drell-Yan: A quark-antiquark pair annihilates, producing a dark photon as given by $q\bar{q} \rightarrow A'$ [65]. This production mechanism is employed in hadron collider experiments and proton beam-dump experiments.

The dark photon can then be detected via a variety of approaches. The case of a decay to visible final states such as lepton or hadron pairs ($A' \rightarrow \ell^+\ell^-$ or $A' \rightarrow h^+h^-$) is often used in collider and beam-dump experiments. The detection of dark photon decays are indicated by the presence of a resonance over the irreducible background. This bump hunt technique is usually accompanied by a reconstruction of the detected tracks, as the detection of a displaced vertex is indicative of a dark photon decay.

Bump hunting can also be used in the missing mass distribution to indirectly detect a dark photon decay into invisible final states. As this method is independent of the probability of the dark photon decaying, a larger sensitivity can be achieved for the same luminosity than in the visible final state case. However, this approach comes with more stringent requirements on the efficacy of the background rejection.

5.1.3 Current Limits

The experimental setup used to detect the dark photon decay greatly affects what region of the parameter space the experiment is sensitive to. In the case of the visible dark photon decay, which is henceforth the focus of this chapter, experiments assuming prompt decays (generally collider experiments) excel at detecting dark photons with larger mixing values ($\epsilon > 10^{-3}$) and up to higher masses (several tens of GeV), whereas experiments assuming decays with a highly displaced vertex (generally beam-dump experiments) excel when the mixing strength is below 10^{-3} and the mass is relatively low (up to a few GeV). Current limits are shown in figure 5.1, with each contributing source detailed below.




Figure 5.1: Current limits on the mass and mixing strength of dark photons as provided by a variety of prompt decay experiments (NA48/2, A1, LHCb, KLOE, BaBar, and CMS) and highly displaced vertex experiments (E774, NA64(e), E141, ν -Cal, CHARM, and E137). Excluded regions from supernovae (SN1987A) and the magnetic moment of the electron $(g-2)_e$ are also included. Image reproduced from Ref. [66].

Prompt Decay Experiments

- NA48/2: A search for dark photon production via the meson decay $\pi^0 \rightarrow \gamma A'$ followed by the dark photon promptly decaying into an electron and positron pair ($A' \rightarrow e^+e^-$) was performed using a sample of 1.69×10^7 reconstructed $\pi^0 \rightarrow \gamma e^+e^-$ decays [67].
- A1: The spectrometer setup of the A1 collaboration was used to search for electron and positron pairs arising from the possible decay of dark photons ($A' \rightarrow e^+e^-$) [68].
- LHCb: Dark photons produced in proton-proton collisions at 13 TeV and decaying via $A' \rightarrow \mu^+\mu^-$ were searched for in a data sample corresponding

to an integrated luminosity of 5.5 fb^{-1} [69].

- KLOE: Dark photon production via annihilation ($e^+e^- \rightarrow \gamma A'$) followed by the decay into a pion and antipion pair ($A' \rightarrow \pi^+\pi^-$) was searched for using 28 million $e^+e^- \rightarrow \gamma \pi^+\pi^-$ events [70].
- BaBar: A total of 514 fb^{-1} data was used to search for dark photon decay into an electron and positron pair or muon and antimuon pair ($A' \rightarrow e^+e^-$, $A' \rightarrow \mu^+\mu^-$) preceded by the $e^+e^- \rightarrow \gamma A'$ annihilation production interaction [71].
- CMS: The same search as LHCb was performed using 137 fb^{-1} of fully reconstructed data and 96.6 fb^{-1} of data with reduced trigger level information [72].

Highly Displaced Vertex Experiments

- E774: A total of 0.52×10^{10} electrons accelerated to 275 GeV were fired at a Tungsten target, utilising the Bremsstrahlung production channel ($e^- Z \rightarrow e^- Z A'$) and then searching for decays into an electron and positron pair ($A' \rightarrow e^+e^-$) [73].
- NA64(e): Also using the Bremsstrahlung production channel followed by the dark photon decay into an electron and positron pair, the NA64(e) experiment fired 8.4×10^{10} electrons of 150 GeV each at a Tungsten target [74].
- E141: A 9 GeV electron beam fired 2×10^{15} electrons at a Tungsten target, again making use of the Bremsstrahlung production channel followed by the dark photon decay into an electron positron pair [75].
- ν -Cal: A 70 GeV proton beam fired 1.71×10^{18} protons at a neutrino calorimeter, the data was reanalysed to set limits for the dark photon which may have been produced via proton Bremsstrahlung ($p Z \rightarrow p Z A'$) [76, 77].
- CHARM: A search for excess electrons in heavy neutrino decays ($\nu_h \rightarrow \nu e^+e^-$) using neutrinos originating from meson decays was reanalysed to set

limits for dark photons that may have been produced via $\eta \rightarrow \gamma A'$ and $\eta' \rightarrow \gamma A'$ before decaying into an electron and positron pair ($A' \rightarrow e^+e^-$) [78].

- E137: Data from a 20 GeV electron beam fired at an aluminium target [79] was reanalysed to look for signals of dark photon production via Bremsstrahlung ($e^- Z \rightarrow e^- Z A'$) followed by decay into electron and positron pairs ($A' \rightarrow e^+e^-$) [80]. The analysis was extended to include dark photon production via the annihilation of positrons produced in electromagnetic showers with atomic electrons in the target ($e^+e^- \rightarrow \gamma A'$ and $e^+e^- \rightarrow A'$) [81].

Other limits

- $(g-2)_e$: Should the dark photon exist it would contribute towards the value of the magnetic moment of the leptons. As a result, the precise experimental measurement of the magnetic moment of the electron $(g-2)_e$ allows for limits to be set on the mass and mixing strength of the dark photon [82].
- SN1987A: The production of dark photons in a supernova could result in contradictions between the observed cooling and theoretical models, as the more long-lived dark photons would be able to escape thereby generating excess cooling. This allows for limits to be set on the lower mixing strength dark photons [83].

5.1.4 Proposed Experiments

It can be seen above that the collider experiments approach the parameter space from the top and extend downwards, whereas the beam-dump experiments approach from the left and extend diagonally down and to the right. This results in a gap in the parameter space that neither experimental approach has so far been able to probe. A number of proposed experiments hope to address this gap, whereas others are more suited at extending the current limits further. The projected exclusion regions from such searches are shown in figure 5.2 and a brief overview of each experiment is given below:

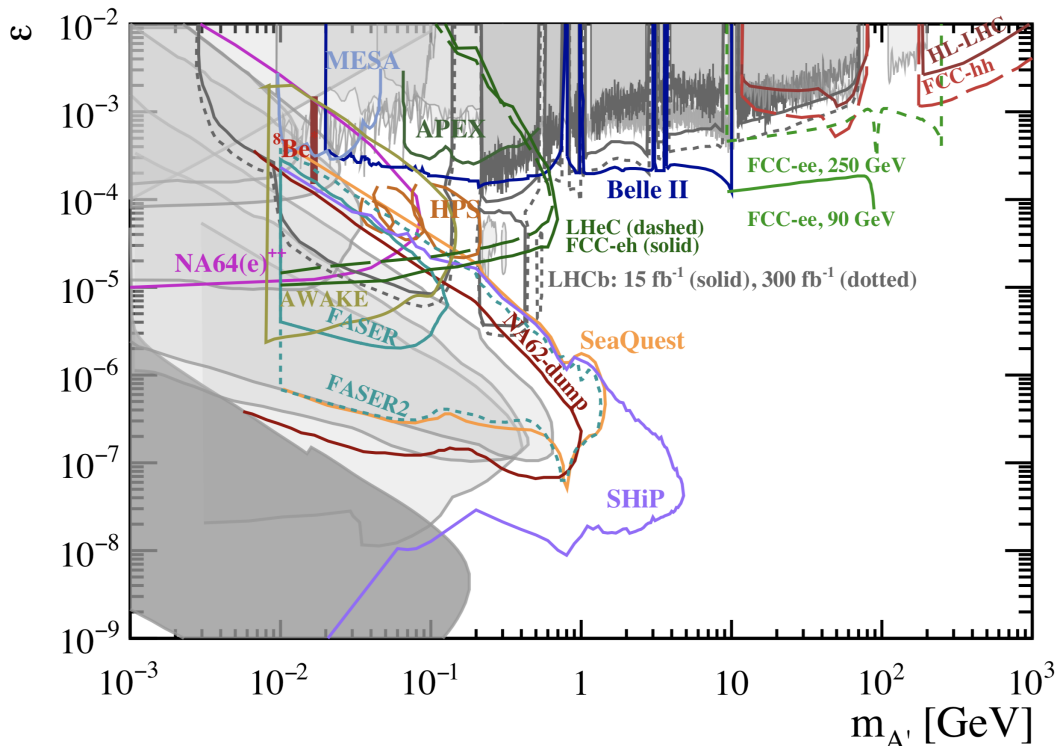


Figure 5.2: Projections for proposed dark photons searches (MESA, Belle-II, FCC-ee, HL-LHC/FCC-hh, LHeC/FCC-eh, NA64(e)⁺⁺, AWAKE, HPS, LHCb, FASER, NA62, SeaQuest, and SHiP) are shown as coloured curves. Existing limits are shown as grey areas. A red line shows the coupling range of a new gauge boson X that could explain the ^8Be anomaly. Image reproduced from Ref. [66].

- MESA: A 155 MeV electron beam fired at a gas target will be used to produce dark photons via the Bremsstrahlung interaction. A twin arm dipole spectrometer will be able to measure the subsequent decays into electron and positron pairs [84, 85].
- Belle-II: A search for dark photon production via annihilation followed by the decay into an electron and positron pair or a muon pair will be made based on the collection of 50 ab^{-1} of data [86].
- FCC-ee: A future circular electron-positron collider could make use of dark photon production via annihilation ($e^+e^- \rightarrow \gamma A'$) followed by decay into muon pairs ($A' \rightarrow \mu^+\mu^-$). Exclusion regions for $\sqrt{s} = 90 \text{ GeV}$ (150 fb^{-1}) and $\sqrt{s} = 250 \text{ GeV}$ (5 ab^{-1}) are shown [87].
- HL-LHC/FCC-hh: The high luminosity upgrade at the large hadron collider

and a future circular proton-proton collider could produce dark photons via the Drell-Yan process, before decaying into electron and positron pairs or muon pairs. Exclusion regions are plotted for the HL-LHC ($\sqrt{s} = 14$ TeV, 3 ab^{-1}) and the FCC-hh ($\sqrt{s} = 100$ TeV, 3 ab^{-1}) [88].

- LHeC/FCC-eh: A future circular electron hadron collider, or a so called large hadron electron collider built using the current large hadron collider, could produce dark photons via deep inelastic scattering (e^- parton $\rightarrow e^-$ parton A') which would then decay into lepton pairs. Exclusion regions are plotted for the LHeC ($\sqrt{s} = 1.3$ TeV, 1 ab^{-1}) and the FCC-eh ($\sqrt{s} = 3.5$ TeV, 3 ab^{-1}) [89].
- NA64(e)⁺⁺: An upgrade to the NA64 experiment, a run consisting of 5×10^{12} electrons on target will be performed [90].
- AWAKE: An experiment is proposed that aims to use plasma wakefield accelerated electron bunches provided by the AWAKE facility. The exclusion region provided in figure 5.2 does not include any evaluation of the background and so can not be compared fairly with the other regions. More information on the AWAKE facility was provided in section 2.2, and further details on the proposed dark photon search are provided in the remainder of this chapter (see section 5.2 onward).
- HPS: An electron beam with energy varying between 2.2 and 6.6 GeV will produce dark photons via the Bremsstrahlung interaction in a thin Tungsten target, the following electron and positron pair decay products will then be measured [91].
- LHCb: A total of 15 fb^{-1} data will be taken for dark photon searches using the $D^{*0} \rightarrow D^0 A'$ followed by decay into an electron and positron pair [92]. In addition, an inclusive approach where the A' production mechanism need not be specified can be performed by using LHCb's high sensitivity to the $A' \rightarrow \mu^+ \mu^-$ decay [93]. Plans to take a further 300 fb^{-1} of data in a further

run are also being considered.

- **FASER:** Long lifetime dark photons produced as a result of proton-proton collisions at the large hadron collider may decay into lepton pairs at distances far from the interaction point. FASER hopes to measure such events using 150 fb^{-1} of data, with FASER2 aiming to collect 3000 fb^{-1} [94].
- **NA62-dump:** A total of 10^{19} protons with an energy of 400 GeV will be used to create a kaon beam to search for a multitude of feebly-interacting particles. Limits can be set on the visible dark photon decay as dark photons may be produced in the dump itself or via the decay of mesons that are produced in the dump [95].
- **SeaQuest:** Using a 120 GeV proton beam to fire 10^{18} protons on target, the Drell-Yan, meson decay, and proton Bremsstrahlung mechanisms will be used to produce dark photons. The decay signal will be given by the detection of an electron and positron pair or muon pair [96].
- **SHiP:** Protons with an energy of 400 GeV will be used to produce dark photons via the Drell-Yan, meson decay, proton Bremsstrahlung, and QCD Compton scattering mechanisms. The experiment will be sensitive to visible final states producing a displaced vertex. A total of 2×10^{20} protons will be fired at the target [97].

Along with extending the current limits and addressing the gap in the parameter space, the proposed experiments will also cover the region corresponding to the potential existence of a 17 MeV vector gauge boson X , indicated in figure 5.2 by a vertical red line. Such a boson is motivated by the anomalous opening angle and invariant mass distributions of e^+e^- pairs measured in ${}^8\text{Be}$ nuclear transitions [98], where such a boson could be produced via ${}^8\text{Be}^* \rightarrow {}^8\text{Be} X$ and then decay via $X \rightarrow e^+e^-$.

5.2 Motivation and Design

Ongoing advances in the field of plasma wakefield acceleration pose an exciting opportunity for the development of a new electron beam-dump experiment designed to search for dark photons. The use of a bunch structured beam allows for a much higher number of electrons to be fired at the target when compared to single electron beam experiments operating for the same duration. The potential gain in sensitivity that results from providing more electrons is demonstrated by simulating an example configuration and plotting the exclusion regions corresponding to 10^{14} , 10^{15} , and 10^{16} electron on target, shown in figure 5.3.




Figure 5.3: Exclusion regions plotted for an example configuration for the cases of 10^{14} , 10^{15} , and 10^{16} electrons fired at the target.

However, using a bunch structured beam also introduces additional challenges when it comes to distinguishing between signal and background as many of the standard cuts are no longer applicable. As a result, the detector design must be optimised such that the number of background detections are at an acceptable level, while also maintaining a large proportion of the additional dark photon signal

gained through the higher beam luminosity.

The proposed experimental setup (shown in figure 5.4) consists of a tungsten target followed by a vacuum tube, three trackers, and an electromagnetic calorimeter. The electron beam is based on the beam parameters that may be achieved by a potential near-future AWAKE configuration, namely a 50 GeV electron beam with 10^9 electrons per bunch. Ideally, dark photons would escape the target without any interaction before decaying into e^+e^- pairs inside the vacuum tube. The detection of two tracks corresponding to an e^+e^- pair in all three trackers followed by at least one electromagnetic shower in the calorimeter provides the experimental signature of a visible dark photon decay. The case of decays into particles of the hidden sector is unexplored, as missing energy techniques require the tagging of individual events and therefore can not be used in conjunction with the pulsed beam structure produced using plasma wakefield accelerators.




Figure 5.4: Schematic of the proposed experimental setup for an electron beam dump experiment used to detect the visible decay of dark photons.

5.3 Simulation

Maximising the discovery potential of such an experiment requires a detailed understanding of how particles propagate through matter, therefore simulations have been performed using the GEANT4 toolkit [99–101]. These simulations have been used to inform the choice of particle tracking technology and the length between the target and detectors (see section 5.4), as well as the thickness of the target and the optimal configuration of magnetic fields (see section 5.5). Finally, the simulation is used to assess the discovery potential of the experimental design (see section 5.6).

A visualisation of the simulated experiment is shown in figure 5.5, specifically showing a simulation configuration for studying background particles. A configuration for studying the dark photon signal also exists and is based on the dark photon class provided by the NA64 collaboration but adapted by the AWAKE collaboration to account for bunched electrons; this configuration can be used to determine the number of electrons on target required to exclude the dark photon at the 95% confidence level (see section 5.3.1). Improvements to the track classification were made to better distinguish between what is considered signal or background (see section 5.3.2). As a bunched electron beam could result in the accumulation of many low energy particles at the same position in the detector, the ability to group real-time coincidental events is desirable and was therefore implemented (see section 5.3.3). A magnetic field was placed after the target to investigate how they affect the sensitivity of the experiment (see section 5.3.4). Lastly, a feature allowing for the repeat of specific events of interest for acquiring more information or subjecting it to stricter requirements was developed (see section 5.3.5).




Figure 5.5: Visualisation of the simulated experiment running in a background configuration. The target is shown in blue, the magnetic field in green, Mylar windows of the vacuum tube in white, and a general detector in red.

5.3.1 Exclusion Regions

The simulation requires the user to input a mass and mixing strength for the dark photon, and outputs the probability that a dark photon meets the signal criteria per electron bunch. For very few or zero signal events the 95% confidence level upper limit is given by

$$N_{95\%} = 3 \quad (5.3)$$

when the number of events observed $N_{\text{obs}} = 0$ [102]. Therefore, the number of electron bunches required to exclude a dark photon of a given mass and mixing strength with 95% confidence is given by taking the reciprocal of the probability and multiplying by 3.

A screenshot showing the table output of the simulation for a variety of mixing strengths at a mass of 0.008 GeV is shown below (see figure 5.6). A sensitivity contour is also shown where the first column at $m_{A'} = 0.008$ GeV corresponds to the the values in the table. It can be seen that the transition between the blue cells (high sensitivity) and the yellow cells (low sensitivity) follows the increase and decrease in the probabilities in the table, note that the order of the mixing strengths in the table are inverted compared to the plot.




Figure 5.6: Table output from the simulation running in the signal configuration, columns show the mixing strength, mass (GeV), and probability of measuring a dark photon signal per electron bunch (left). Sensitivity contour where the first column ($m_{A'} = 0.008$ GeV) corresponds to the values from the table output (right).

5.3.2 Track Classification

The classification of signal and background tracks was originally handled by comparing the track identification number of each particle hit, `trackID`, with the track identification numbers of the dark photon decay products, `electronID` and `positronID`. This correctly identifies the decay products themselves as signals, however their interactions can result in the creation of new particles which no longer match the required identification numbers and are therefore misclassified as background (see figure 5.7).




Figure 5.7: Positions of the hits in electromagnetic showers caused by dark photon decay products have been categorised into signal (blue) and background (red) by the original event classifier.

Correct event classification (see figure 5.8) was achieved by creating an algorithm that recursively finds the identification number of the previous particle, `parentID`. The `G4TrackingManager` class was used to create a map, pairing the IDs of all tracks with their parents. Defining a function `GetParentID(trackID)` to return the `parentID` from the map allows for any track to be traced back to its origin.




Figure 5.8: Positions of the hits in the same electromagnetic showers have been categorised using the new method, the absence of red markers indicates successful event classification.

Particle tracks can then be classified as signal if its `trackID` or `parentID` matches with either the `electronID` or `positronID`, if there are no matches then the ID of the previous track is obtained so that it can be tested as well. This process continues until the particle is either confirmed to have originated from a dark photon, or until `parentID` = 0, signifying that the current track is the initial beam and that all relevant tracks have been tested.

5.3.3 Background Heat Maps

To account for the possibility of many low energy particles escaping the target and mimicking a single higher energy particle, a simulation configuration was implemented to group nearby hits and sum their frequency and energy. A general detector was created with a 52×52 cell structure by replicating a single cell along the y-axis to create a column of cells, and then replicating the column along the x-axis to create the rows. This allows the column number of a hit to be obtained from the copy number of the cell volume, and the row number from the copy number of the cell's mother volume, a demonstration is shown in figure 5.9.




Figure 5.9: Visualisation of the cell structured detector showing a run with 10 particle hits, the coordinates of the cells for each hit are obtained and printed.

At the end of each event the total number of particles entering each cell is recorded, along with their total energy. A cone shaped beam angled towards the top right quadrant of the detector is used to demonstrate the summation of energies entering through each cell, see figure 5.10. Optionally, the user can choose to normalise the plot to represent a run of 10^9 electrons fired, giving the expected result for the background caused by an electron bunch provided by AWAKE.




Figure 5.10: Visualisation of a run where a cone shaped beam is fired at the top right quadrant of the detector (left), the detector is shown as a red grid, particle trajectories are shown as red lines, and hits are shown as yellow dots. The total energy entering each cell throughout the run is displayed as a heat map (right).

5.3.4 Magnetic Field

It is useful to simulate magnetic fields to assess whether they are capable of improving the sensitivity of the experiment, therefore a volume representing a magnetic field was added to the simulation. The `G4GenericMessenger` class was used to define a new command directory `/awake` for setting parameters such as the magnetic field. These features are displayed in figure 5.11 where a charged particle with all particle physics interactions disabled travels through the magnetic field (shown in green), the `/awake/field` command is used to demonstrate the additional bending caused by the 0.3 T field compared to the 0.1 T field.

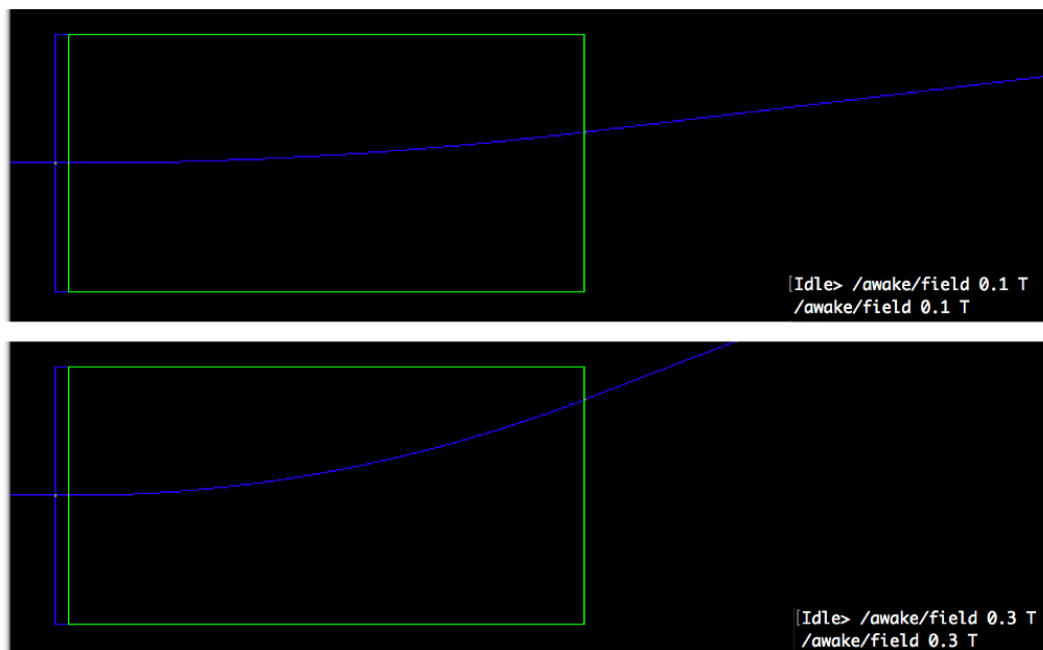


Figure 5.11: Visualisation of an event where a charged particle travels through a magnetic field of 0.1 T (top) and 0.3 T (bottom).

5.3.5 Event Replay

For geometry configurations where background events are uncommon, it could be useful to repeat previously simulated background events without using computing resources to re-simulate uninteresting events. A feature was implemented that saves the status of the random engine for any event where a particle interacts with a specified geometry. The random engine status can then be restored if any need to

repeat the event arises, for example re-simulating with a thicker target or acquiring additional information about the event. A demonstration is shown below (see figure 5.12) where a run involving an isotropic source results in 5 events that interact with a detector plane, the random engine status for these events are saved and then loaded in a separate run consisting only of those events of interest.




Figure 5.12: Visualisation of a run where an isotropic source results in 5 events that interact with a detector plane, meeting the requirement for the random engine status to be saved (left). Visualisation of a separate run where the 5 events of interest are replayed by loading their random engine status (right).

5.4 Trackers and Decay Volume

The signal criteria for a dark photon to be detected requires that both of the decay products pass through all of the trackers with a separation distance large enough such that they can be distinguished from one another. Trackers that are larger in size are better at ensuring that the decay products are detected, however having a larger tracker often means having a lower spatial resolution and therefore sacrificing the ability to distinguish between decay products that are detected close together. As a result, it is important to understand how the issue of a decay product missing a tracker compares to the issue of both decay products hitting a tracker close enough such that they are erroneously detected as one particle. As the position of a decay product hitting a tracker depends on the distance it travels, the decay volume length

(distance between the target and the first tracker) must be considered in tandem.

5.4.1 Decay Volume Length

An investigation was made into how the length of the decay volume impacts the sensitivity of the experiment. Projected exclusion regions for an example configuration with varied decay volume lengths are plotted (see figure 5.13), dark photons with less than 5 GeV energy are not considered in the analysis as the electromagnetic shower is poorly described at these energies [103].




Figure 5.13: Projected exclusion regions plotted for an example configuration with decay volume lengths of 1 m, 5 m, and 20 m.

It can be seen that longer decay volume lengths excel at detecting lower mixing strength dark photons. These dark photons have a longer lifetime and therefore a longer path length, meaning that they are more likely to escape the decay volume and avoid detection. This can be mitigated by extending the length of the decay volume, resulting in longer decay length configurations having a larger sensitivity for the low mixing strength region of the parameter space.

All configurations perform equally well at detecting dark photons with larger mixing strengths. These dark photons have a shorter lifetime and therefore a shorter path length, making them less likely to escape the decay volume. The path length distribution for dark photons with a mass of 0.1 GeV and a mixing strength of 3×10^{-5} is plotted (see figure 5.14), where it can be seen that the dark photon is unlikely to travel more than 1 m before decaying. Increasing the dark photon mass to 0.3 GeV, or increasing the mixing strength to 1×10^{-4} results in even shorter decay path lengths. Consequently, a longer decay volume length is not advantageous for dark photons in these regions.




Figure 5.14: Path length distributions for dark photons with $m_{A'} = 0.1$ GeV and $\epsilon = 1 \times 10^{-4}$ (red), $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$ (blue), and $m_{A'} = 0.3$ GeV and $\epsilon = 3 \times 10^{-5}$ (black).

Perhaps surprisingly, a shorter decay volume length does not seem to provide a significant advantage for higher mixing strength dark photons. Given that a longer decay volume provides more distance for the decay products to travel apart, it might be expected that they are less likely to hit the trackers. The observed absence in the loss of detection efficiency arises from the 5 GeV cut applied to the dark photon energy, as higher energy dark photons have a smaller opening angle between the

decay products (see figure 5.15). The lower energy dark photons removed by the energy cut therefore have a large overlap with the dark photons that would have avoided detection due to an increased decay volume length.




Figure 5.15: Relationship between the opening angle of the e^+e^- pair and the energy of the parent dark photon. The parameters of the generated dark photons are $m_{A'}$ = 0.1 GeV and $\epsilon = 3 \times 10^{-5}$.

5.4.2 Tracker Size and Resolution

To better understand how the compromise between tracker size and resolution affects a dark photon detection meeting the signal criteria, toy detectors that vary in size and resolution are defined with specifications resembling a straw tube detector and a micromegas detector.

Straw tube detectors [104] consist of many thin tubes filled with gas, with each tube having a wire down the centre. As a charged particle passes through a tube the gas is ionised, resulting in electron ion pairs. These pairs drift due to a potential difference between the wire and the walls of the tube, thereby producing a current. The micromegas detector [105] is also filled with gas, however it uses a microscopic mesh to form a thin amplification region. When a charged particle passes through the chamber the gas is ionised and the resulting electron is acceler-

ated towards the mesh. During this, further electron ion pairs are produced in an avalanche process, but due to the small amplification region the avalanche is more localised and controlled leading to a better spatial resolution. On the other hand, the compact design of the micromegas detector results in a reduced active area.

The specifications of the toy detectors considered in this study are given below:

- Larger size tracker (resembling a straw tube detector): Radius of 2 m and can distinguish particles at least 5 mm apart.
- Higher resolution tracker (resembling a micromegas detector): Radius of 0.5 m and can distinguish particles at least 1.75 mm apart.

The separation distance between the decay products at the first tracker for a decay volume length of 1 m is plotted, see figure 5.16.




Figure 5.16: Separation distance between hits of e^+e^- pairs in a tracker located 1 m after the target for dark photons with $m_{A'} = 0.1$ GeV and $\epsilon = 1 \times 10^{-4}$ (red), $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$ (blue), and $m_{A'} = 0.3$ GeV and $\epsilon = 3 \times 10^{-5}$ (black).

For dark photons with $m_{A'} = 0.1$ GeV and $\epsilon = 1 \times 10^{-4}$ (shown in red), and particularly for dark photons with $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$ (shown in blue),

a number of dark photon decays result in a separation distance between decay products of less than the 5 mm spatial resolution of the larger size tracker. This occurs more for the $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$ dark photons as the lower mixing strength can cause the decay to happen closer to the tracker, giving less distance for the decay products to separate. In both of these cases it is shown to be unlikely for the decay products to travel so far apart that they would miss the tracker entirely, meaning that the higher resolution tracker provides more sensitivity. For the dark photon with $m_{A'} = 0.3$ GeV and $\epsilon = 3 \times 10^{-5}$ (shown in black), all decay products are further apart than 5 mm, meaning the high resolution tracker offers no benefit in this case. The larger separation distances results from having a larger dark photon mass, which increases the opening angle between the decay products.

The separation distance between the decay products at the first tracker for a decay volume length of 20 m is also plotted, see figure 5.17.




Figure 5.17: Separation distance between hits of e^+e^- pairs in a tracker located 20 m after the target for dark photons with $m_{A'} = 0.1$ GeV and $\epsilon = 1 \times 10^{-4}$ (red), $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$ (blue), and $m_{A'} = 0.3$ GeV and $\epsilon = 3 \times 10^{-5}$ (black).

When extending the decay volume length to 20 m, the separation distance between the decay products increases to a large enough extent that the resolution

of the tracker is not a limiting factor for any of the parameters tested. Much like with the 1 m decay volume length, the dark photon with $m_{A'} = 0.3$ GeV and $\varepsilon = 3 \times 10^{-5}$ results in the largest separation distance between the decay products. Unlike before, the difference between the $m_{A'} = 0.1$ GeV and $\varepsilon = 1 \times 10^{-4}$ dark photon compared to the $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ dark photon is much less pronounced. This is because the difference in path lengths arising from these two dark photon parameters is small when compared to the 20 m decay volume length. For a decay volume this long, the larger size tracker provides more sensitivity.

5.4.3 Acceptance

The acceptance, defined as the percentage of dark photons with energy greater than 5 GeV whose decay products both hit the tracker with a large enough separation distance to be distinguished, is plotted for both types of trackers at a variety of decay volume lengths (see figure 5.18). The parameters that the dark photons are generated at are $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$, as this presents an interesting target region in the so far unexplored parameter space.




Figure 5.18: Acceptance for dark photons with $m_{A'} = 0.1$ GeV and $\varepsilon = 3 \times 10^{-5}$ for both the larger size tracker and the higher resolution tracker placed at a variety of distances from the target.

It can be seen that the higher resolution tracker performs better for shorter decay volumes and that the larger size tracker performs better for longer decay volumes, as expected from the simulations shown previously. However, the increase in sensitivity gained from the higher resolution tracker above the larger size tracker at short distances is relatively small, even for the $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$ dark photons which were shown to benefit from the higher resolution tracker more than the other investigated parameters. Conversely, the increase in sensitivity gained from the larger size tracker at long distances is much larger, enabling a configuration where much lower mixing strength dark photons can be probed with minimal sacrifice to the initially defined region of interest. As a result, a larger size tracker paired with a 20 m long decay volume has been selected. It should be noted that all simulations correspond to an initial beam energy of 50 GeV, and that higher energy beams creating higher energy dark photons which decay into electron and positron pairs with smaller opening angles (as is the case for NA64 and future potential higher energy AWAKE beams) will likely find additional value in the higher resolution tracker beyond that which has been demonstrated here.

5.5 Target and Magnetic Field

Alongside enabling the production of dark photons, the target is also responsible for shielding the detectors from particles that could potentially mimic a dark photon signal. Increasing the thickness of the target or applying a magnetic field can help to reduce the number of unwanted particles that make it to the detectors. However, both of these adjustments impact the measurement of the dark photon decay products as well, making it necessary to reach a balance that maximises the dark photon signal while still achieving a tolerable level of background noise.

5.5.1 Target Thickness

An investigation was made into how the thickness of the target impacts the sensitivity of the experiment. Projected exclusion regions for an example configuration

with varied target thicknesses are plotted (see figure 5.19). It can be seen that thinner targets excel at detecting higher mixing strength dark photons. These dark photons have a shorter lifetime, meaning that they are less likely to escape the target before decaying. Conversely, lower mixing strength dark photons mostly travel a long enough distance to escape the target before decaying, resulting in no significant loss in sensitivity from having a thicker target.




Figure 5.19: Projected exclusion regions plotted for an example configuration with target thicknesses of 10 cm, 40 cm, and 150 cm.

In all cases above, the projected exclusion regions assume that no events occur that are indistinguishable from a dark photon decay event, although the likelihood of this being true in practice decreases as the target thickness decreases. As the initial estimate of the achievable sensitivity of this experiment (discussed in section 5.1.4 and presented in figure 5.2) also makes this assumption despite using a 10 cm target, an assessment of the background for that configuration was performed by constructing heat maps of the total energy and number of particles reaching the detectors per electron bunch on target (see figure 5.20). As most cells record upwards of 1 TeV, with the central region recording upwards of 15 TeV, it is clear that any signal from a dark photon produced using a 50 GeV beam could not have been identified among the background. Furthermore, the large number of particles pass-

ing through the trackers would disable any attempt to reconstruct the dark photon decay vertex. As a result, it is necessary to adjust the target thickness and consider the use of magnetic fields for reducing the background level to an acceptable state.




Figure 5.20: Total energy (top) and number of particles (bottom) passing through the detectors, results are normalised to represent one electron bunch.

5.5.2 Magnetic Field Length

A magnetic field with a field strength of 2 T was placed immediately after the target. Investigations were made into how effective increasing the target thickness and the length of the magnetic field is at reducing the number of background particles detected (see figure 5.21). It can be seen that both increasing the target thickness and increasing the magnet length reduce the background, as expected. However, it is also shown that increasing the thickness of the target reduces the background much more effectively, requiring less physical space to do so.




Figure 5.21: Number of background particles measured at the up-stream detectors for different configurations of target thicknesses and magnet lengths, results are normalised to represent one electron bunch. Grey cells correspond to configurations where no background is detected within the simulation runtime.

As it is also important to limit the loss of signal while reducing the background, the signal to background ratio was also investigated (see figure 5.22). Signal here is defined as the probability of a dark photon detection per electron bunch on target, therefore this ratio does not represent the signal and background present in a dark

photon decay event but rather is used to assess how the probability of a detection varies alongside the background level for any given electron bunch fired at the target. The dark photons are generated with parameters $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$, corresponding to the region of interest in the parameter space identified during the previous discussion. A clear improvement can be seen as the target thickness increases, although it is less evident for increases in the magnet length for targets with a thickness of 20 or 50 cm, suggesting that the loss in signal is proportionally equal to the loss in background when extending the length of the magnetic field in these cases. Nonetheless, the 100 cm target length followed by a magnet length of 50 cm provides the best ratio of all tested configurations, although it is still plausible that this configuration is surpassed by even thicker targets with no magnetic field.




Figure 5.22: Signal to background ratio for different configurations of target thicknesses and magnet lengths, results are normalised to represent one electron bunch. Grey cells correspond to configurations where no background is detected within the simulation runtime. Dark photons are generated with parameters $m_{A'} = 0.1$ GeV and $\epsilon = 3 \times 10^{-5}$.

5.5.3 Optimal Configuration

The optimal configuration is not necessarily the one that provides the best signal to background ratio, but rather the one that provides the highest probability of detecting dark photons while presenting a small enough number of background particles such that there is no chance of misidentifying ordinary Standard Model events as a dark photon decay. Although an upper limit on the number of measured background particles required to prevent this is yet to be determined, an estimate of 100 background particles measured per bunch is made to proceed with the comparison of three potential configurations:

- 100 cm thick target followed by a 50 cm long 2 T magnetic field: Simulations performed above show this configuration to result in 80 ± 28 measured background particles, putting it below the upper limit of measured background particles defined above.
- 130 cm thick target without any magnetic field: The simulations above also indicate that increasing the target thickness may improve the signal to background ratio more effectively than configurations with a magnetic field, a 130 cm thick target without any magnetic field is estimated to result in a number of measured background particles below the upper limit defined above.
- 100 cm thick target surrounded by a 2 T magnetic field of equal length: Providing a magnetic field within the same physical space as the target may provide an interesting configuration as it limits the total distance over which dark photon decay products could be affected. Simulations assume that the magnetic field surrounding the target performs equivalently to a magnetic field placed after the target.

As each configuration is effective at limiting the number of background particles arriving at the detector, it is computationally demanding to acquire large statistics. Consequently, 10% of an AWAKE electron bunch was simulated for each configuration, with the results then normalised to an entire bunch of 10^9 electrons

on target (see table 5.1).

Table 5.1: Signal, background, and the signal to background ratio is provided for the following configurations: a 100 cm thick target followed by a 50 cm long magnetic field, a 130 cm thick target without any magnetic field, and a 100 cm thick target surrounded by a magnetic field of equal length. Dark photons are generated with parameters $m_{A'}$ = 0.1 GeV and $\epsilon = 3 \times 10^{-5}$.

| Configuration | Signal | Background | Ratio |
|-----------------------------|--------------------------------|-------------|----------------------------|
| Target followed by magnet | $(2.9 \pm 0.4) \times 10^{-6}$ | 80 ± 28 | $(4 \pm 1) \times 10^{-8}$ |
| Target without magnet | $(1.6 \pm 0.2) \times 10^{-6}$ | 40 ± 20 | $(4 \pm 2) \times 10^{-8}$ |
| Target surrounded by magnet | $(3.3 \pm 0.4) \times 10^{-6}$ | 80 ± 28 | $(4 \pm 2) \times 10^{-8}$ |

With the given statistics, the configurations seem to perform similarly. Although more precise comparisons would require further simulation to acquire more events, it is at least clear that all three configurations are capable of reducing the number of background particles to a reasonably low level. The configuration consisting of a 100 cm thick target surrounded by a magnetic field of equal length is henceforth used as its smaller total size provides an advantage when considering dark photons with other parameters where a larger proportion of the decays may occur between 100 cm and 150 cm from the start of the target.

5.6 Discovery Potential

The number of electron bunches required to exclude dark photons of any given parameter point in the mass and mixing strength plane was evaluated for the configuration determined above (see figure 5.23), where one electron bunch consists of 10^9 electrons of 50 GeV energy each. This can be used to determine projected exclusion regions based on the number of electron bunches on target for an experimental set-up consisting of a 100 cm thick target surrounded by a 2 T magnetic field, with a 20 m long decay volume preceding three trackers of radius 2000 mm that can distinguish particles at least 5 mm apart. Each electron bunch on target results in an expected number of background particle measurements of 80 ± 28 , all of which are assumed to be incapable of mimicking the dark photon decay signature.

Additionally, it is assumed that dark photons with energy below 5 GeV can not be reliably identified as signal events.




Figure 5.23: Number of electron bunches required to exclude dark photons of any given parameter point in the mass and mixing strength plane with 95% confidence for the configuration described in this chapter.

Comparison of the AWAKE dark photon search is made with the current limits (see top plot of figure 5.24) and limits from other proposed experiments (see bottom plot of figure 5.24), assuming that 10^7 electron bunches are provided during the lifetime of the experiment. The projected sensitivity of the configuration described in this chapter is indicated in both plots by a cyan curve. The original projected exclusion region for an AWAKE dark photon search (indicated by a sycamore curve) makes the assumption that no events occur that are indistinguishable from a dark photon decay event, an assumption that was shown here to be false. However, after optimisations were made to the initial configuration it was shown that this scheme is still capable of probing currently unexplored areas of the parameter space, specifically in the area that coincides with the gap between beam-dump and collider experiments.

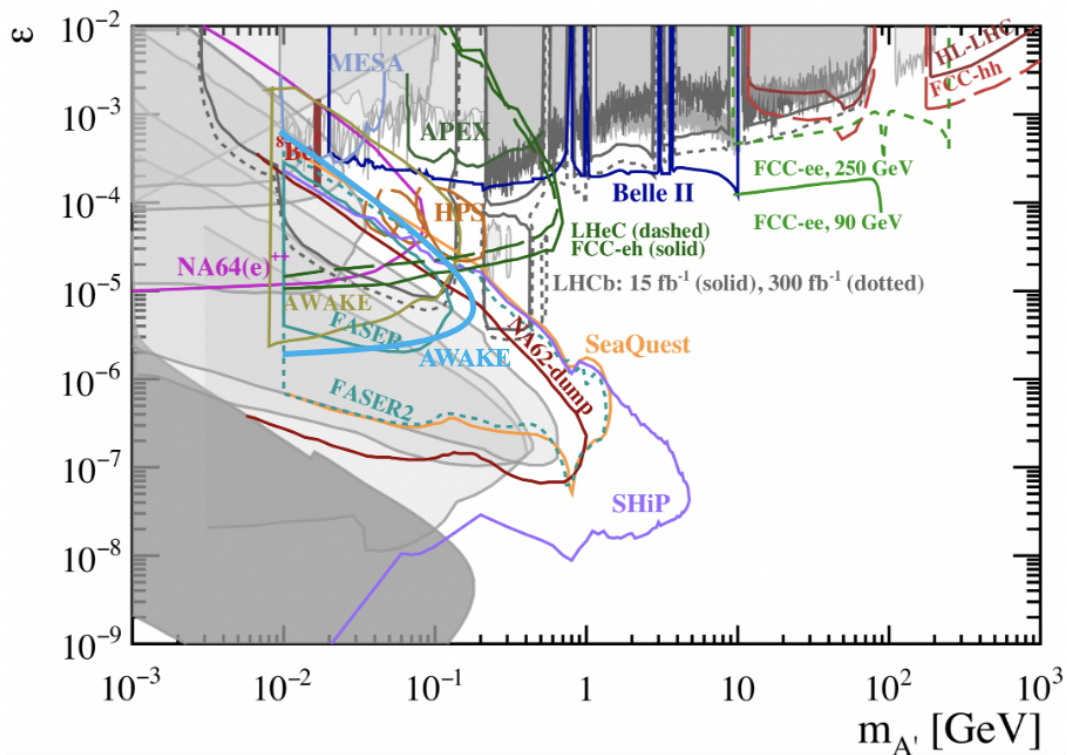



Figure 5.24: Current limits (top) and projected exclusion regions of proposed experiments (bottom) for dark photons in the mass and mixing strength plane. The configuration described in this chapter is indicated on both plots by a cyan curve labelled AWAKE. Images adapted from Ref. [66].

Chapter 6

Conclusions

The repetition rate of plasma wakefield acceleration was measured experimentally under a variety of plasma conditions. It was determined that in the case where the plasma is generated anew for each event the acceleration can be repeated with a rate of 2.0 kHz in argon plasma and 4.55 kHz in hydrogen plasma. The use of bunch trains, where multiple bunches are accelerated within the same plasma, can further increase the average power of the accelerator. In such a case, the time between each bunch train is determined by the measurements presented here. Simulations were performed to further understanding of the physics processes that limit the repetition rate, providing verification that the expulsion and refill of material following the high voltage discharge both contribute significantly.

The repetition rate measurements are determined by the time between two consecutive acceleration events. The extent to which a third acceleration event can retain the properties of the previous two is so far unknown, and further study should investigate any limitations that arise when accelerating numerous electron bunches continuously. For example, it is expected that heating of the plasma by the beam will increase. Whether this affects the time taken for the state of the plasma to reset, and whether a steady state of reproducible acceleration can be reached within a reasonable 'warming up' time remains to be seen. Furthermore, as the average power of the accelerator increases, so does the need to appropriately manage any leftover power in the system. This means that further study into improving the

power transfer efficiency from the drive to the witness beam should be conducted, and that an appropriate cooling scheme for the plasma cell should be developed.

The work presented in this thesis has indicated that material expulsion contributes significantly to the repetition rate of a plasma wakefield accelerator, however not much is known about the dynamics of this expulsion. The introduction of diagnostic tools to measure this expulsion, as well as the development of simulations to gain an improved understanding of this process, both provide interesting avenues for further study. Subsequently, improvements to the design of the capillary to minimise material loss could be achieved, thereby increasing the maximum achievable repetition rate.

A dark photon search based on a potential near-future 50 GeV electron beam with 10^9 electrons per bunch was proposed, with simulations performed to optimise the experimental configuration. This included comparing different tracker technology and decay volume length combinations, as well as different combinations of target thicknesses and magnetic field configurations, ensuring that the background level was tolerable and that the dark photon signal was maximised. The sensitivity was then assessed for the optimal configuration and compared to other experiments, where it was concluded that this search is capable of probing currently unexplored regions of the dark photon parameter space.

One limitation of this study is that only the decay of dark photons into e^+e^- pairs was considered. For dark photons with mass $m_{A'} > 2m_\mu$, the decay into a $\mu^+\mu^-$ pair is available too. Improvements to the dark photon sensitivity can be made by considering both decay channels. Furthermore, the sensitivity towards other hypothetical particles should be investigated and quantified. For example, interaction between high energy bremsstrahlung photons and the nuclei of the target could result in the production of axionlike particles (ALPs) [106] via the Primakoff effect [107]. These ALPs could then escape the target before decaying into two

photons, providing a detectable signature.

The extent to which a far-future TeV energy scale electron beam could further probe the parameter space remains to be seen and should be considered as a topic for further study. For an electron beam of higher energy more dark photons would be produced, and these dark photons would reach higher energies too. On the other hand, more background particles will also be produced, and these particles will also reach higher energies. The increased background can be mitigated with additional shielding, and so the experimental configuration should be re-optimised to account for the higher energy electron beam. The sensitivity of this new higher energy configuration can then be determined.

One of the difficulties associated with optimising the experimental configuration arises from quantifying the background. The computational cost of simulating background events in such an experiment is already large, and it would be considerably more expensive for a TeV energy scale configuration. As a result, further study would benefit from making use of advances in deep generative models to simulate particle showers [108], significantly increasing the speed at which these background events can be generated.

Finally, in addition to computing the exclusion regions, a more complete analysis could consider non-zero background estimates and explore how likely the experiment would be to observe a statistically significant excess of signal events if the dark photon model is indeed true. This can be quantified by the expected discovery significance, often defined as the median significance that would be observed over repeated experiments. An estimate of this value can be obtained using the Asimov approximation [109], where the expected background is constrained by a control measurement obtained via simulation.

Bibliography

- [1] M. D. Schwartz, *Quantum field theory and the standard model*. Cambridge University Press, 2014.
- [2] A. D. Sakharov, “Violation of CP-invariance, C-asymmetry, and baryon asymmetry of the Universe,” in *In The Intermissions... Collected Works on Research into the Essentials of Theoretical Physics in Russian Federal Nuclear Center, Arzamas-16*, pp. 84–87, World Scientific, 1998.
- [3] Y. Fukuda *et al.*, “Evidence for oscillation of atmospheric neutrinos,” *Phys. Rev. Lett.*, vol. 81, no. 8, p. 1562, 1998.
- [4] A. G. Riess *et al.*, “Observational evidence from supernovae for an accelerating universe and a cosmological constant,” *The Astronomical Journal*, vol. 116, no. 3, p. 1009, 1998.
- [5] F. Zwicky, “On the Masses of Nebulae and of Clusters of Nebulae,” in *A Source Book in Astronomy and Astrophysics, 1900–1975*, pp. 729–737, Harvard University Press, 1979.
- [6] D. Aguillard *et al.*, “Measurement of the positive muon anomalous magnetic moment to 0.20 ppm,” *Physical review letters*, vol. 131, no. 16, p. 161802, 2023.
- [7] The LHCb collaboration, “Comprehensive analysis of local and nonlocal amplitudes in the $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decay,” *J. High Energy Phys.*, vol. 2024, no. 9, 2024.

- [8] M. Conte and W. W. MacKay, *Introduction To The Physics Of Particle Accelerators*. World Scientific Publishing Company, 2008.
- [9] T. Tajima and J. M. Dawson, “Laser Electron Accelerator,” *Phys. Rev. Lett.*, vol. 43, pp. 267–270, 1979.
- [10] S. P. Mangles *et al.*, “Monoenergetic beams of relativistic electrons from intense laser–plasma interactions,” *Nature*, vol. 431, no. 7008, pp. 535–538, 2004.
- [11] C. Geddes *et al.*, “High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding,” *Nature*, vol. 431, no. 7008, pp. 538–541, 2004.
- [12] J. Faure *et al.*, “A laser–plasma accelerator producing monoenergetic electron beams,” *Nature*, vol. 431, no. 7008, pp. 541–544, 2004.
- [13] S. M. Hooker, “Developments in laser-driven plasma accelerators,” *Nature Photonics*, vol. 7, no. 10, pp. 775–782, 2013.
- [14] A. Gonsalves *et al.*, “Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide,” *Phys. Rev. Lett.*, vol. 122, no. 8, p. 084801, 2019.
- [15] W. Leemans *et al.*, “The BERkeley Lab Laser Accelerator (BELLA): A 10 GeV laser plasma accelerator,” in *AIP Conference Proceedings*, vol. 1299, pp. 3–11, American Institute of Physics, 2010.
- [16] P. Chen *et al.*, “Acceleration of electrons by the interaction of a bunched electron beam with a plasma,” *Phys. Rev. Lett.*, vol. 54, no. 7, p. 693, 1985.
- [17] R. D. Ruth *et al.*, “A plasma wake field accelerator,” *Part. Accel.*, vol. 17, no. SLAC-PUB-3374, p. 171, 1984.
- [18] J. B. Rosenzweig *et al.*, “Experimental observation of plasma wake-field acceleration,” *Phys. Rev. Lett.*, vol. 61, no. 1, p. 98, 1988.

- [19] I. Blumenfeld *et al.*, “Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator,” *Nature*, vol. 445, no. 7129, pp. 741–744, 2007.
- [20] M. Litos *et al.*, “9 GeV energy gain in a beam-driven plasma wakefield accelerator,” *Plasma Physics and Controlled Fusion*, vol. 58, no. 3, p. 034017, 2016.
- [21] R. D’Arcy *et al.*, “FLASHForward: plasma wakefield accelerator science for high-average-power applications,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 377, p. 20180392, 2019.
- [22] C. A. Lindstrøm *et al.*, “Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator,” *Phys. Rev. Lett.*, vol. 126, p. 014801, 2021.
- [23] A. Caldwell *et al.*, “Proton-driven plasma-wakefield acceleration,” *Nature Physics*, vol. 5, no. 5, pp. 363–367, 2009.
- [24] P. Muggli *et al.*, “AWAKE readiness for the study of the seeded self-modulation of a 400 GeV proton bunch,” *Plasma Physics and Controlled Fusion*, vol. 60, no. 1, p. 014046, 2017.
- [25] E. Adli *et al.*, “Acceleration of electrons in the plasma wakefield of a proton bunch,” *Nature*, vol. 561, no. 7723, pp. 363–367, 2018.
- [26] A. Caldwell and K. V. Lotov, “Plasma wakefield acceleration with a modulated proton bunch,” *Physics of Plasmas*, vol. 18, no. 10, p. 103101, 2011.
- [27] P. Piot *et al.*, “Generation and characterization of electron bunches with ramped current profiles in a dual-frequency superconducting linear accelerator,” *Phys. Rev. Lett.*, vol. 108, no. 3, p. 034801, 2012.
- [28] S. Schröder *et al.*, “Tunable and precise two-bunch generation at FLASHForward,” in *Journal of Physics: Conference Series*, vol. 1596, p. 012002, IOP Publishing, 2020.

- [29] J. Dale *et al.*, “Vorrichtung mit beweglicher Aufnahme für Vakuumkammern,” *Deutsche Patentanmeldung*, no. 10, p. 116, 2014.
- [30] R. D’Arcy *et al.*, “Longitudinal phase space reconstruction at FLASHForward using a novel X-band transverse deflection cavity (XTDC),” *Proceedings of IPAC’18, Vancouver, Canada*, p. 1567, 2018.
- [31] B. Hidding *et al.*, “Ultracold electron bunch generation via plasma photocathode emission and acceleration in a beam-driven plasma blowout,” *Phys. Rev. Lett.*, vol. 108, no. 3, p. 035001, 2012.
- [32] F. Li *et al.*, “Generating high-brightness electron beams via ionization injection by transverse colliding lasers in a plasma-wakefield accelerator,” *Phys. Rev. Lett.*, vol. 111, no. 1, p. 015003, 2013.
- [33] A. Martinez De La Ossa *et al.*, “High-quality electron beams from beam-driven plasma accelerators by wakefield-induced ionization injection,” *Phys. Rev. Lett.*, vol. 111, no. 24, p. 245003, 2013.
- [34] A. Knetsch *et al.*, “Controlled density-downramp injection in a beam-driven plasma wakefield accelerator,” *arXiv:2007.12639*, 2020.
- [35] D. Lipka *et al.*, “Development of Cavity BPM for the European XFEL,” in *proceedings of LINAC*, 2010.
- [36] D. Lipka *et al.*, “First experience with the standard diagnostics at the European XFEL injector,” in *IBIC 2016-5th International Beam Instrumentation Conference*, p. MOBL02, 2016.
- [37] M. Tzoufras *et al.*, “Beam loading in the nonlinear regime of plasma-based acceleration,” *Phys. Rev. Lett.*, vol. 101, no. 14, p. 145002, 2008.
- [38] R. D’Arcy *et al.*, “FLASHForward: plasma wakefield accelerator science for high-average-power applications,” *Philosophical Transactions of the Royal Society A*, vol. 377, no. 2151, p. 20180392, 2019.

- [39] R. D’Arcy *et al.*, “Recovery time of a plasma-wakefield accelerator,” *Nature*, vol. 603, no. 7899, pp. 58–62, 2022.
- [40] L. Gorbunov *et al.*, “Plasma ions dynamics in the wake of a short laser pulse,” *Phys. Rev. Lett.*, vol. 86, no. 15, p. 3332, 2001.
- [41] L. Gorbunov *et al.*, “Dynamics of a plasma channel created by the wakefield of a short laser pulse,” *Physics of Plasmas*, vol. 10, no. 4, pp. 1124–1134, 2003.
- [42] J. B. Rosenzweig *et al.*, “Effects of ion motion in intense beam-driven plasma wakefield accelerators,” *Phys. Rev. Lett.*, vol. 95, no. 19, p. 195002, 2005.
- [43] J. Vieira *et al.*, “Ion motion in self-modulated plasma wakefield accelerators,” *Phys. Rev. Lett.*, vol. 109, no. 14, p. 145005, 2012.
- [44] E. Adli *et al.*, “Experimental observation of proton bunch modulation in a plasma at varying plasma densities,” *Phys. Rev. Lett.*, vol. 122, no. 5, p. 054802, 2019.
- [45] E. Gschwendtner *et al.*, “The AWAKE Run 2 programme and beyond,” *Symmetry*, vol. 14, no. 8, p. 1680, 2022.
- [46] M. Turner *et al.*, “Experimental study of wakefields driven by a self-modulating proton bunch in plasma,” *Physical Review Accelerators and Beams*, vol. 23, no. 8, p. 081302, 2020.
- [47] B. Buttenschön *et al.*, “A high power, high density helicon discharge for the plasma wakefield accelerator experiment AWAKE,” *Plasma Physics and Controlled Fusion*, vol. 60, no. 7, p. 075005, 2018.
- [48] N. Torrado *et al.*, “Double pulse generator for unipolar discharges in long plasma tubes for the AWAKE experiment,” *Submitted for publication*, 2022.
- [49] E. Gschwendtner *et al.*, “AWAKE++: the AWAKE acceleration scheme for new particle physics experiments at CERN,” tech. rep., 2018.

- [50] A. Caldwell and M. Wing, “VHEeP: A very high energy electron–proton collider,” *The European Physical Journal C*, vol. 76, no. 8, p. 463, 2016.
- [51] A. Accardi *et al.*, “Electron-Ion Collider: The next QCD frontier,” *The European Physical Journal A*, vol. 52, pp. 1–100, 2016.
- [52] J. Chappell *et al.*, “A compact electron injector for the EIC based on plasma wakefields driven by the RHIC-EIC proton beam,” *arXiv:1907.01191*, 2019.
- [53] W. Decking *et al.*, “A MHz-repetition-rate hard X-ray free-electron laser driven by a superconducting linear accelerator,” *Nature photonics*, vol. 14, no. 6, pp. 391–397, 2020.
- [54] P. Bambade *et al.*, “The International Linear Collider: A Global Project,” *arXiv:1903.01629*, 2019.
- [55] Ansys Fluent, Release 18.0.
- [56] T. H. Shih *et al.*, “A new k-epsilon eddy viscosity model for high Reynolds number turbulent flows: Model development and validation,” tech. rep., 1994.
- [57] T. Parikh. Private Communication.
- [58] P. Fayet, “Effects of the spin-1 partner of the goldstino (gravitino) on neutral current phenomenology,” *Physics Letters B*, vol. 95, no. 2, pp. 285–289, 1980.
- [59] P. Fayet, “A la recherche d’un nouveau boson de spin un,” *Nuclear Physics B*, vol. 187, no. 1, pp. 184–204, 1981.
- [60] L. Okun, “Boundaries of electrodynamics: paraphotons,” *Zh. Eksp. Teor. Fiz.:(USSR)*, vol. 83, no. 9, 1982.
- [61] H. Georgi *et al.*, “Photon oscillations and cosmic background radiation,” *Nature*, vol. 306, no. 5945, pp. 765–766, 1983.

- [62] B. Holdom, “Two U(1)’s and Epsilon Charge Shifts,” *Physics Letters B*, vol. 166, no. 2, pp. 196–198, 1986.
- [63] D. Feldman *et al.*, “Stueckelberg Z-prime extension with kinetic mixing and millicharged dark matter from the hidden sector,” *Physical Review D*, vol. 75, no. 11, p. 115001, 2007.
- [64] H. Yukawa, “On the interaction of elementary particles,” *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, vol. 17, pp. 48–57, 1935.
- [65] S. D. Drell and T. M. Yan, “Massive Lepton-Pair Production in Hadron-Hadron Collisions at High Energies,” *Phys. Rev. Lett.*, vol. 25, pp. 316–320, 1970.
- [66] M. Fabbrichesi *et al.*, *The physics of the dark photon: a primer*. Springer, 2021.
- [67] J. Batley *et al.*, “Search for the dark photon in π^0 decays,” *Physics Letters B*, vol. 746, pp. 178–185, 2015.
- [68] H. Merkel *et al.*, “Search at the Mainz Microtron for light massive gauge bosons relevant for the muon g-2 anomaly,” *Phys. Rev. Lett.*, vol. 112, no. 22, p. 221802, 2014.
- [69] R. Aaij *et al.*, “Search for $A' \rightarrow \mu^+ \mu^-$ Decays,” *Phys. Rev. Lett.*, vol. 124, no. 4, p. 041801, 2020.
- [70] A. Anastasi *et al.*, “Limit on the production of a new vector boson in $e^+e^- \rightarrow U\gamma, U \rightarrow \pi^+\pi^-$ with the KLOE experiment,” *arXiv:1603.06086*, 2016.
- [71] J. Lees *et al.*, “Search for a Dark Photon in e^+e^- Collisions at BaBar,” *Phys. Rev. Lett.*, vol. 113, no. 20, p. 201801, 2014.
- [72] CMS Collaboration, “Search for a narrow resonance decaying to a pair of muons in proton-proton collisions at 13 TeV,” tech. rep., CMS-PAS-EXO-19-018, 2019.

- [73] A. Bross *et al.*, “Search for short-lived particles produced in an electron beam dump,” *Phys. Rev. Lett.*, vol. 67, no. 21, p. 2942, 1991.
- [74] D. Banerjee *et al.*, “Improved limits on a hypothetical X(16.7) boson and a dark photon decaying into e^+e^- pairs,” *Physical Review D*, vol. 101, no. 7, p. 071101, 2020.
- [75] E. Riordan *et al.*, “Search for short-lived axions in an electron-beam-dump experiment,” *Phys. Rev. Lett.*, vol. 59, no. 7, p. 755, 1987.
- [76] J. Blümlein and J. Brunner, “New exclusion limits for dark gauge forces from beam-dump data,” *Physics Letters B*, vol. 701, no. 2, pp. 155–159, 2011.
- [77] J. Blümlein and J. Brunner, “New exclusion limits on dark gauge forces from proton bremsstrahlung in beam-dump data,” *Physics Letters B*, vol. 731, pp. 320–326, 2014.
- [78] S. Gninenko, “Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays,” *Physics Letters B*, vol. 713, no. 3, pp. 244–248, 2012.
- [79] J. D. Bjorken *et al.*, “Search for neutral metastable penetrating particles produced in the SLAC beam dump,” *Physical Review D*, vol. 38, no. 11, p. 3375, 1988.
- [80] B. Batell *et al.*, “Strong constraints on sub-GeV dark sectors from SLAC beam dump E137,” *Phys. Rev. Lett.*, vol. 113, no. 17, p. 171802, 2014.
- [81] L. Marsicano *et al.*, “Dark photon production through positron annihilation in beam-dump experiments,” *Physical Review D*, vol. 98, no. 1, p. 015031, 2018.
- [82] M. Pospelov, “Secluded U(1) below the weak scale,” *Physical Review D*, vol. 80, no. 9, p. 095002, 2009.

- [83] J. H. Chang *et al.*, “Revisiting supernova 1987A constraints on dark photons,” *Journal of High Energy Physics*, vol. 2017, no. 1, pp. 1–33, 2017.
- [84] L. Doria *et al.*, “Dark Matter at the Intensity Frontier: the new MESA electron accelerator facility,” *arXiv:1908.07921*, 2019.
- [85] L. Doria *et al.*, “Search for light dark matter with the MESA accelerator,” *arXiv:1809.07168*, 2018.
- [86] E. Kou *et al.*, “The Belle II physics book,” *Progress of Theoretical and Experimental Physics*, vol. 2019, no. 12, p. 123C01, 2019.
- [87] M. Karliner *et al.*, “Radiative return capabilities of a high-energy, high-luminosity e^+e^- collider,” *Physical Review D*, vol. 92, no. 3, p. 035010, 2015.
- [88] D. Curtin *et al.*, “Illuminating dark photons with high-energy colliders,” *Journal of High Energy Physics*, vol. 2015, no. 2, pp. 1–45, 2015.
- [89] M. D’Onofrio *et al.*, “Searching for dark photons at the LHeC and FCC-he,” *Physical Review D*, vol. 101, no. 1, p. 015020, 2020.
- [90] D. Banerjee *et al.*, “Addendum to the NA64 Proposal: Search for the $A' \rightarrow$ invisible and $X \rightarrow e^+e^-$ decays in 2021.”
- [91] P. H. Adrian *et al.*, “Search for a dark photon in electroproduced e^+e^- pairs with the Heavy Photon Search experiment at JLab,” *Physical Review D*, vol. 98, no. 9, p. 091101, 2018.
- [92] P. Ilten *et al.*, “Dark photons from charm mesons at LHCb,” *Physical Review D*, vol. 92, no. 11, p. 115017, 2015.
- [93] P. Ilten *et al.*, “Proposed inclusive dark photon search at LHCb,” *Phys. Rev. Lett.*, vol. 116, no. 25, p. 251803, 2016.
- [94] J. L. Feng *et al.*, “ForwArd Search ExpeRiment at the LHC,” *Physical Review D*, vol. 97, no. 3, p. 035001, 2018.

- [95] NA62 Collaboration, “ADDENDUM I TO P326 Continuation of the physics programme of the NA62 experiment,” tech. rep., 2019.
- [96] A. Berlin *et al.*, “Dark sectors at the Fermilab SeaQuest experiment,” *Physical Review D*, vol. 98, no. 3, p. 035011, 2018.
- [97] M. Anelli *et al.*, “A facility to Search for Hidden Particles (SHiP) at the CERN SPS,” *arXiv:1504.04956*, 2015.
- [98] A. Krasznahorkay *et al.*, “Observation of anomalous internal pair creation in ^8Be : A possible indication of a light, neutral boson,” *Phys. Rev. Lett.*, vol. 116, no. 4, p. 042501, 2016.
- [99] S. Agostinelli *et al.*, “Geant4—a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.
- [100] J. Allison *et al.*, “Geant4 developments and applications,” *IEEE Transactions on Nuclear Science*, vol. 53, no. 1, pp. 270–278, 2006.
- [101] J. Allison *et al.*, “Recent developments in Geant4,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 835, pp. 186–225, 2016.
- [102] G. J. Feldman and R. D. Cousins, “Unified approach to the classical statistical analysis of small signals,” *Physical review D*, vol. 57, no. 7, p. 3873, 1998.
- [103] D. Banerjee *et al.*, “Search for a hypothetical 16.7 MeV gauge boson and dark photons in the NA64 experiment at CERN,” *Phys. Rev. Lett.*, vol. 120, no. 23, p. 231802, 2018.
- [104] A. Braghieri *et al.*, “The Straw Tube Tracker of the PANDA experiment,” in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, pp. 1013–1018, IEEE, 2009.

- [105] Y. Giomataris *et al.*, “MICROMEGAS: a high-granularity position-sensitive gaseous detector for high particle-flux environments,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 376, no. 1, pp. 29–35, 1996.
- [106] P. Svrcek and E. Witten, “Axions in string theory,” *Journal of High Energy Physics*, vol. 2006, no. 06, p. 051, 2006.
- [107] H. Primakoff, “Photo-production of neutral mesons in nuclear electric fields and the mean life of the neutral meson,” *Physical Review*, vol. 81, no. 5, p. 899, 1951.
- [108] ATLAS Collaboration, “Deep generative models for fast photon shower simulation in ATLAS,” *arXiv:2210.06204*, 2022.
- [109] G. Cowan, “Discovery sensitivity for a counting experiment with background uncertainty,” *en. In*, vol. 8, 2012.