DEMONSTRATION OF THE
HOLLOW CHANNEL PLASMA WAKEFIELD ACCELERATOR

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Spencer J. Gessner
September 2016

© Copyright by Spencer J. Gessner 2016
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Tor Raubenheimer) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Alex Chao)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Mark Hogan)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Michael Peskin)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Ron Ruth)

Approved for the Stanford University Committee on Graduate Studies

Abstract

A plasma wakefield accelerator is a device that converts the energy of a relativistic particle beam into a large-amplitude wave in a plasma. The plasma wave, or wakefield, supports an enormous electric field that is used to accelerate a trailing particle beam. The plasma wakefield accelerator can therefore be used as a transformer, transferring energy from a high-charge, low-energy particle beam into a high-energy, low-charge particle beam. This technique may lead to a new generation of ultra-compact, high-energy particle accelerators.

The past decade has seen enormous progress in the field of plasma wakefield acceleration with experimental demonstrations of the acceleration of electron beams by several gigaelectron-volts. The acceleration of positron beams in plasma is more challenging, but also necessary for the creation of a high-energy electron-positron collider. Part of the challenge is that the plasma responds asymmetrically to electrons and positrons, leading to increased disruption of the positron beam. One solution to this problem, first proposed over twenty years ago, is to use a hollow channel plasma which symmetrizes the response of the plasma to beams of positive and negative charge, making it possible to accelerate positrons in plasma without disruption.

In this thesis, we describe the theory relevant to our experiment and derive new results when needed. We discuss the development and implementation of special optical devices used to create long plasma channels. We demonstrate for the first time the generation of meter-scale plasma channels and the acceleration of positron beams therein.

Acknowledgments

When I started as a graduate student at FACET, my first task was to measure the dispersion of our energy spectrometer. I was assigned beam time – 12:00AM-8:00AM – my normal working hours at FACET. I showed up to the control room and told the operators what needed to be done, expecting them to carry out the task. Instead, they turned the controls over to me. I was now operating the world’s highest-energy electron beam, which I immediately crashed into the wall. That was a fun night.

The last five years working at FACET have been a lot like that night; frequent failure, sparse success, and a good time all along. I have my co-workers and advisors to thank for that. Mark Hogan threw me in the deep end and I never looked back. His trust gave me the confidence to fail and try again. Vitaly Yakimenko gave me guidance, support, and taught me how to make convincing arguments based on physics, which he chose to ignore anyway. My advisor Tor Raubenheimer advocated for me in every situation. He always encouraged me to dig deeper and helped me understand the theory behind the experiment.

When I first arrived at SLAC, Mike Litos sent me an e-mail with “everything I needed to know about FACET.” It was way too much and not nearly enough, much like our work. Erik Adli followed up with instructions for how to run plasma simulations and the lyrics to heavy-metal versions of Norwegian folk songs. Selina Green showed me the art of experimental assembly and how to handle the micron-thin \$5K YAG crystal without breaking it. Sebastien Corde arrived and unleashed a torrent of physics, the likes of which the world has never seen.

Then there was the tunnel. Dark, cramped, and way too hot, it was the crucible where students were formed into scientists. With the unflinching resolve a British Army Officer at the Battle of Balaclava, Christine Clarke sent her troops into the Valley of Death:

Forward, the FACET Brigade!
Charge for the laser!

Navid Vafaei-Najafabadi led the way, undeterred by heat exhaustion or lack-of-sleep. Ken Marsh brought some levity to the work. Rafal Zgadaj took it away. James Allen stood tough through the night. Reinforcements arrived in a Trojan Horse: Aihua Deng, Alex Knetsch, Grace Manahan,

Oliver Karger, Paul Sherkl, and Thomas Heinemann created an inescapable laser labyrinth. Brendan O'Shea and Oliver Williams helpfully mocked us while we worked.

Once we finished our tasks underground, we found little reprieve in the control room, where we suffered from MCCBS. Luckily, we had Carl Lindstrøm, the Matlab Mozart, and Antoine Doche, the laser wrangler, to help us with our code. Joel Frederico created indecipherable data structures for Chris Clayton to decipher. Nate Lipkowitz, Jerry Yocky, Tonee Smith and OPs always delivered, even if it wasn't what we ordered.

Away from the experiment, I had tons of help from the physicists at SLAC. I learned beam dynamics from Alan Fisher, Glen White, Uli Wienands, F.J. Decker, J.P. Delahaye, and Ron Ruth. Genady Stupakov, Karl Bane, Sami Tantawi, and Alex Chao taught me how to manipulate Maxwell's equations and calculate wakefields. Michael Peskin taught me why we build accelerators. Chan Joshi taught me how to write a paper. Dieter Walz taught me how to endure.

Over the years, I've found support for my physics education in surprising places. Dan Balick taught me Special Relativity when I was a freshman in high school. Hannah Siemann took a keen interest in my research and supported my work through the Siemann Fellowship. She is the only person I know outside of accelerated physics who is genuinely excited about accelerator physics. John "Dr. K" Kirtley gave me an opportunity to get involved in research by opening his lab to me while I was still in high school. Now a researcher myself, I am amazed by how generous he was with his time. More importantly, he suggested I check out UC Santa Barbara.

Writing this thesis would not have been possible without the support of my friends, who understood exactly how much it sucked because I told them repeatedly. I am lucky to still have them. From New York: Fox, Jordan, Maltz, Robins, Sean, and Vik. From UCSB: Clark, Katie, Mikey, Noah, Haque, Ron, Sam, Sarah, Slop, Spiff, and Wheels. From Stanford: Calvin, DerMen, James, Katie, Mason, Masha, Matt, Max, Miriam, PapaG, Sam, and Sibs.

From a young age, my parents Hal and Jessica encouraged me intellectually and enrolled me in every science program they could find. They gave me *The Hitchhiker's Guide to the Galaxy* and that convinced me I should be a physicist. I've relied on their love and support throughout my decades-long education. My younger brother Dylan always looked up to me when we were kids. Now I look up to him. My grandparents Shirley, Mickey and Jerry emphasized the importance of education. Every Saturday morning in the fall, my Grandpa would drive me to the city so I could attend sciences classes and then drive me back to Westchester for my football games. We talked about life, the universe, and everything during those car rides.

Finally, I want to thank my wife Kathleen, who has sustained me throughout the Ph.D. process. When we first met, she assumed I was a "dumb jock". I thought she was hot. To our surprise, we discovered an intellectual and emotional depth in each other that lies at the heart of our relationship. Her love pushed me through to the finish.

Contents

Abstract	v
Acknowledgments	vi
1 Introduction	1
1.1 Particle Physics and Politics	1
1.2 Limitations of Conventional Accelerators	2
1.3 Plasma Acceleration	3
1.3.1 Laser-Driven Plasma Acceleration	4
1.3.2 Beam-Driven Plasma Acceleration	5
1.3.3 FACET and the E200 Collaboration	5
1.4 The Problem with Positrons	6
1.4.1 Hollow Channel Plasma Wakefield Acceleration	7
1.5 Outline	7
2 Electron and Positron Acceleration in Uniform Plasmas	9
2.1 Introduction	9
2.2 Plasma wakefield acceleration in the linear regime	9
2.3 Temporal response of the plasma	11
2.4 Fields in the plasma	12
2.4.1 Longitudinal Field	14
2.4.2 Transverse Forces	18
2.5 Comment on the approximations used in this derivation	20
2.6 Limitations of the Linear Regime	21
2.7 Plasma wakefield acceleration in the non-linear regime	24
2.7.1 Electron beam loading in the plasma bubble	28
2.7.2 The QuickPIC particle-in-cell code	31
2.8 Positron acceleration in the non-linear plasma wakefield	32
2.8.1 Observation of high-gradient acceleration of positrons in a non-linear wakefield	33

2.8.2	Beam loading in non-linear plasma wakes	35
2.8.3	Outlook for positron acceleration in the non-linear regime	36
3	Hollow Channel Plasma Wakefield Acceleration	37
3.1	Introduction	37
3.2	Modal Description of the Plasma Channel	37
3.3	The hollow channel plasma as a waveguide	38
3.3.1	The TM accelerating mode	40
3.3.2	Single Particle Wake Excitation	45
3.4	Validity of the Linear Approximation	46
3.4.1	The strongly nonlinear regime	49
3.5	Beam Loading in the Hollow Channel Wakefield	51
3.6	The HEM Modes	52
3.6.1	General Solutions for HEM modes	52
3.6.2	The Dipole Mode	52
3.7	The Beam Breakup Instability	54
3.7.1	Derivation of BBU expression	55
3.7.2	BBU for an Offset Beam with Angular Misalignment	57
3.7.3	Implications of the BBU on Experiments and Future Colliders	58
4	Optics for Generating Hollow Channel Plasmas	60
4.1	Introduction	60
4.2	Fourier Optics	61
4.3	Bessel Optics	63
4.4	High-Order Bessel Optics	65
4.4.1	Stationary Phase Approximation	65
4.5	Numerical Code for Evaluating Optics at FACET	67
4.6	Kinoform Phase Gratings	68
4.7	Optical Bench Tests	70
4.8	Axilens Optics	71
4.9	Tunnel Ionization in Intense Laser Fields	73
5	FACET	77
5.1	Introduction	77
5.2	Particle Beam Dynamics and Notation	78
5.2.1	Coordinate System	78
5.2.2	Particle Transport	79
5.2.3	Beam Emittance	81

5.3	A Source-to-Sink Description of FACET	82
5.3.1	Thermionic Gun	82
5.3.2	Capture and Acceleration	83
5.3.3	Damping Rings	84
5.3.4	Compression and Acceleration	86
5.3.5	Positron Production and Return Line	88
5.3.6	Sector 20 Compression and Optics	89
5.4	Two-Bunch Generation	92
5.5	Experimental Area and Systems	94
5.5.1	Heat Pipe Oven	94
5.5.2	SYAG Energy Spectrometer	96
5.5.3	X-Band Deflecting Cavity	97
5.5.4	Electro-Optic Sampling	98
5.5.5	Transverse Profile Monitors	99
5.5.6	Dumpline Spectrometer	99
5.6	FACET Laser	100
5.7	Data Acquisition System	102
6	Experimental Results	105
6.1	Experimental Preparation	105
6.1.1	Choice of Kinoform	105
6.1.2	Preparation of the Laser	106
6.1.3	Preparation of the Lithium Heat Pipe Oven	108
6.1.4	Preparation of the Positron Beam	109
6.2	Determination of Hollow Channel Plasma Structure	110
6.2.1	Raster Scan Data from 2014	111
6.2.2	Channel Reconstruction Algorithm	113
6.3	Single-Bunch Energy Loss Measurements	114
6.3.1	Results from 2014	114
6.3.2	Determination of the Channel Length	115
6.3.3	Energy Loss Data from 2016	117
6.3.4	Discussion	118
6.4	Observation of Energy Gain with a Two-Bunch Setup	120
6.4.1	Finding the Accelerating Phase	120
6.4.2	Energy Correction	122
6.4.3	Energy Gain and Beam Loading	123
6.4.4	Discussion	124
6.5	Conclusion	125

7	Conclusions	127
7.1	Towards a Plasma-Based Linear Collider	127
7.1.1	Challenges Associated with a Hollow Channel PLC	127
7.1.2	Other Opportunities for Plasma Elements in a Linear Collider	129
7.2	Anti-Plasma	129
7.3	Outlook	130
A	Derivation of the Point Charge Wakefields	131
A.1	Maxwell's Equations	131
A.2	Hollow Channel Plasma Geometry	132
A.3	Sources	133
A.4	Fourier Decomposition	133
A.5	Solution for azimuthally symmetric case	134
A.5.1	Inside the Channel	134
A.5.2	The Channel Wall	135
A.5.3	Outside the channel	136
A.5.4	Boundary Conditions	136
A.5.5	Inverting the Transform	139
A.6	Solution for the dipole case	140
A.6.1	Inside the Channel	141
A.6.2	The Channel Wall	143
A.6.3	Outside the Channel	144
A.6.4	Boundary Conditions	145
A.6.5	Inverting the Transform	147
A.7	Comparison to Previous Work	147
A.8	Investigation of the Plane Wave Component	149
B	Bessel Function Identities	151
B.1	Bessel's Equations	151
B.2	Series and Integral Representations	152
B.3	Derivative Identities	152
B.4	Bessel Boundary Notation and Identities	153
C	Fourier Optics Code	154
	Bibliography	158

List of Tables

5.1	Table of past, present, and future PWFA facilities, showing beam species, energy, and peak current, which is a key figure of merit for experiments. Parameters taken from references [75, 25, 76, 26, 77, 78, 79].	78
-----	--	----

List of Figures

2.1	A plot of the density perturbation due to a relativistic beam with $N_b = 4 \times 10^8$ electrons and dimensions $\sigma_r = 20 \mu\text{m}$ and $\sigma_z = 15 \mu\text{m}$ propagating to the right in a plasma with density $n_0 = 8 \times 10^{16}$. The black lines are the radial and longitudinal on-axis lineouts of the plasma density.	13
2.2	Contour for the integral 2.38.	16
2.3	A plot of the longitudinal electric field for the same beam and plasma parameters used in Fig. 2.1. The longitudinal field lags the density perturbation by 90 degrees. The black lines are the radial and longitudinal on-axis lineouts of the field.	17
2.4	A plot of the transverse force experienced by a beam particle for the same beam and plasma parameters used in Fig. 2.1 and Fig. 2.3. The force is in phase with the density perturbation and leads the longitudinal field by 90 degrees. The black line is the radial lineout of the force at the center of the drive bunch.	19
2.5	A sketch comparing the accelerating and focusing fields in a plasma wakefield accelerator. The fields are simultaneously focusing and accelerating for 90 degrees per cycle, independent of beam charge. The beam propagates to the left. Reproduced from ref. [17].	22
2.6	The transverse wakefield field as a function of r for three different values of the beam normalized beam width $k_p\sigma_r$. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$	23
2.7	A sketch comparing the trajectory of plasma electrons (red dots) under the influence of dense electron (a) and positron (b) beams. In the plasma electrons form a bubble around the electron beam with stationary ions (blue dots) inside the bubble. The explicit form of the non-linear positron wake is more complicated and depends strongly on the positron beam parameters.	24
2.8	Transverse beam profiles of an electron beam after free space propagation (a) and after being focused in the plasma (b). Transverse beam profiles of a positron beam after free space propagation (c) and after being focused in the plasma (d). Note the appearance of the beam halo indicating emittance growth because some particles in the positron bunch are defocused. Figures from refs. [40] and [41].	25

2.9	Trajectories of plasma electrons under the influence of a drive beam moving to the left (not shown) with (a) density $n_b/n_0 = 1$ and (b) density $n_b/n_0 = 10$. Reproduced from ref. [35].	26
2.10	Left: QuickPIC simulation of a high-density electron beam with $n_b/n_0 \approx 30$, traveling to the right in plasma with density 10^{17} cm^{-3} . The black bubble region is devoid of plasma electrons. The plasma ions are not shown. Right: A transverse slice of the plasma density along $\xi = 60 \text{ } \mu\text{m}$. The width of the plasma sheath is much smaller than the bubble radius.	27
2.11	Left: The longitudinal field in the plasma bubble shown in Figure 2.10. The accelerating field at the back of the double exceeds 60 GV/m . Right: The transverse wakefield for the same bubble.	28
2.12	(a) A QuickPIC simulation of the wake due to a driving electron bunch with charge $N = 6.36 \times 10^9$ and bunch length $\sigma_z = 25 \text{ } \mu\text{m}$. The beam is propagating in a plasma with density $n = 5 \times 10^{16} \text{ cm}^{-3}$. The driving electron bunch is scalloped because it is not matched to the plasma focusing force. (b) The same simulation as in (a) with a loaded witness bunch with charge $N = 4.87 \times 10^9$ and bunch length $\sigma_z = 47 \text{ } \mu\text{m}$. The parameters of the drive beam, witness beam, and plasma match parameters from the FACET experiment. (c) The experimentally measured energy spectrum of the drive beam and witness beam in the absence of plasma. (d) The energy spectrum of the drive beam after deceleration in plasma. The maximum energy loss is over 3 GeV . (e) The energy spectrum of the witness beam after acceleration of plasma. The maximum energy gain is over 2 GeV . (f) A comparison of the measured energy spectrum with the results of energy spectrum simulated using QuickPIC in (b). Figure reproduced from ref. [45].	30
2.13	The time resolved energy spectrum of the positron beam after propagating through a 1.4 m long plasma at density $n = 1.8 \times 10^{14} \text{ cm}^{-3}$. The blue triangles are the mean time-resolved energy slices of the beam in the absence of plasma, and the red squares are the same quantity when the beam propagates through plasma. The bunch length is $700 \text{ } \mu\text{m}$. The shape of the energy gain and loss is indicative of a linear plasma wake. Figure from ref. [23].	33
2.14	The energy spectrum of the accelerated portion of a positron beamlet after propagation in 1.3 m long Lithium plasma oven. The beamlet has a spectral peak near 25 GeV , with 207 pC charge and an rms energy spread of 1.8% . Figure reproduced from ref. [56].	34

2.15	(a) The unloaded wakefield due to a truncated positron bunch. The wake has a bubble structure similar to wakes produced by an electron beam driver, but in this case the bubble trails the positron beam. (b) The loaded wakefield of a positron bunch after 130 cm of propagation in a plasma. The bunch structure has evolved considerably and is in equilibrium. The longitudinal field is flat along much of the bunch. The dotted orange line shows the beam dimensions at the start of the simulation. Figure reproduced from ref. [56].	35
3.1	a. A cutaway of the hollow channel plasma illustrating the tubular geometry. The drive beam propagates along the z -axis of the channel. b. A transverse slice of the hollow channel showing the on-axis vacuum region labeled '1', the annular plasma region labeled '2', and the vacuum outside the channel labeled '3'.	38
3.2	The normalized longitudinal electric field as a function of r for three different values of the beam energy γ . $\gamma \approx 40,000$ corresponds to the beam energy at FACET. The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$	42
3.3	The normalized radial force as a function of r for three different values of the beam energy γ . $\gamma \approx 40,000$ corresponds to the beam energy at FACET. The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$	43
3.4	Top Left: QuickPIC simulation of an electron beam with $N = 2 \times 10^9$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. The electron beam density has been enhanced by a factor of five for visibility. Bottom Left: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), showing excellent agreement. The beam current from QuickPIC (dashed blue) and theoretical calculation (dashed red) are also shown and the peak beam current is 1 kA. Top Right: QuickPIC simulation of a positron beam propagating to the right with the same beam and plasma parameters as the previous case. Bottom Right: Comparison of the on-axis longitudinal field from QuickPIC and theory, also showing excellent agreement.	46

3.5	Top Left: QuickPIC simulation of an electron beam with $N = 1 \times 10^{10}$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom Left: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), showing excellent agreement. The peak beam current is 5 kA. Top Right: QuickPIC simulation of a positron beam propagating to the right with the same beam and plasma parameters as the previous case. Bottom Right: Comparing the on-axis longitudinal fields from QuickPIC and theory, we observe a large discrepancy.	48
3.6	Top: QuickPIC simulation of an electron beam with $N = 2 \times 10^{10}$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), showing agreement through the first oscillation. The discrepancy between theory and simulation emerges after the plasma electrons undergo a full oscillation.	49
3.7	Top: QuickPIC simulation of an positron beam with $N = 2 \times 10^{10}$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), along with the beam current which reaches a maximum value of 10 kA at $z = 0 \mu\text{m}$. The model and simulation diverge after less than half of a single wake period. The field spike is due to a large concentration of on-axis plasma electrons.	50
3.8	Top: QuickPIC simulation of an electron beam drive bunch with $N = 6 \times 10^9$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ and a positron witness bunch with $N = 1 \times 10^9$, $\sigma_z = 13 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$, inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red). The peak current of the drive beam is 3 kA and the peak current of the witness beam is 1.5 kA. The accelerating gradient at the location of the positron witness bunch is 475 MV/m, although the model predicts 800 MV/m. The discrepancy is due to plasma electrons streaming into the vacuum channel at the location of the positron witness.	51
3.9	The normalized longitudinal electric field as a function of r for three different values of the beam energy γ . The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$	53

3.10	The normalized radial force as a function of r for three different values of the beam energy γ . The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$	54
4.1	An illustration of the coordinate system used in the Huygens-Fresnel and Fresnel diffraction equations.	62
4.2	A sketch of the axicon optic showing the focused rays converging at a constant angle β . The focus is the region where the rays cross.	64
4.3	Left: Image of the axicon focus in the transverse plane measured in the laser laboratory at FACET with an 800 nm Ti:sapphire laser. Right: The transverse profile along the line $y = 0 \mu\text{m}$ fit with the function $J_0^2(k_\perp r)$. The fit value for the convergence angle $\beta = 5.7 \text{ mrad}$ and the axicon design value is 5.6 mrad.	64
4.4	A spiral shaped optic produces a phase shift $\Phi(\phi) = m\phi$ where m is an integer. . . .	65
4.5	Left: Intensity distribution in r and z of an axicon plus spiral phase optic with convergence angle $\beta = 4.4 \text{ mrad}$ and $m = 7$. Right: A comparison of intensity distributions in r showing increasing contrast between the central minimum and first maximum with increasing m	67
4.6	Left: Intensity distribution in x and y of an axicon plus spiral phase optic with convergence angle showing the effect of a nearby aperture. Right: Intensity distribution in x and y of an axicon plus spiral phase optic with a veritcal astigmatism in the wavefront, represented by a Zernike polynomial with order $m = 2$, $n = 2$	68
4.7	Left: Etching mask for a $m = 7$, binary optic. Right: Etching mask for a $m = 7$, 8-layer staircase optic.	69
4.8	Left: Numerical calculation of the efficiency of the first-order diffraction versus number of etched layers. Right: Scanning electron microscope image of an $m = 6$, 8-layer staircase optic. Image courtesy of the manufacturer, NIL Technologies.	70
4.9	Left: Image of the $m = 8$ optic at a distance of 86 cm from the optic. Right: Image of the $m = 8$ optic at a distance of 86 cm from the optic.	71
4.10	Top: Rail measurement of the $m = 8$ kinoform optic for $86 < z < 362 \text{ cm}$. The peak intensity is normalized to 1. Bottom: Numerical simulation of the same optic. Diffraction effects are also observed and can be seen as faint wiggles near $z = 86 \text{ cm}$. The peak intensity is normalized to 1.	72
4.11	A comparison of the intensity along the line-focus for an axicon and axilens illuminated by a 100 fs, 150 mJ laser pulse. The intensity of both optics surpasses the hydrogen ionization threshold of $2 \times 10^{14} \text{ W/cm}^2$ from $z = 3 \text{ m}$ to $z = 4 \text{ m}$. The axilens has significant intensity overhead and therefore can be operated with reduced laser power to achieve the same result.	73

4.12	Top: Rail measurement of the an 8-layer axilens optic with $R = 2$ cm, $f_0 = 4$ m, and $\Delta_z = 2$ m. The FWHM length of the focus was 1.3 m. Bottom: Numerical simulation of the same optic. The FWHM length of the focus is 2 m.	74
4.13	Top: Intensity of the line-focus of an $m = 8$ optic, illuminated by a 40 mJ, 100 fs laser pulse. Bottom: The ionization contour of the intensity pattern in Lithium with a first ionization energy of 3.59 eV.	76
5.1	Schematic of the accelerator and compression systems at FACET. Thick yellow sections indicate the presence of RF accelerating structures in the linac. Black lines are used to highlight the path of electrons and the gray line indicates the path of positrons, starting from the target in Sector 19. The locations of the energy spread and bunch length diagnostics throughout the linac are highlighted in red and blue, respectively.	82
5.2	Cut-away view of the thermionic gun. Reproduced from reference [81].	83
5.3	Schematic of the SLAC beam source. Two guns are shown on the left. The thermionic gun (top) is used at FACET while the polarized gun (bottom) is decommissioned. The inset to the schematic shows the beam energy and bunch length after the gun, the sub-harmonic bunchers, s-band buncher, and K-02 accelerating section. The bunch length does not change after the K-02 section and prior to injection into the damping ring.	84
5.4	Left: Streak data showing 5 consecutive turns of the electron beam in the NDR as well as the projection of the electron beam current showing a 6.6 mm r.m.s. profile. Right: Bunch length versus RF amplitude showing a characteristic $V^{-1/2}$ dependence. Measurements performed by A. Fisher and J.M. Allen.	86
5.5	a.) Schematic depicting bunch compression. i.) A long uncorrelated beam enters an RF cavity near the zero crossing of the RF field. In this illustration, the beam travels to the right and the front of the beam is decelerated while the tail is accelerated. There is no change to the centroid beam energy. ii.) The chirped beam enters a magnetic chicane with positive R_{56} . The particles in the front of the bunch traverse a longer distance than particles in the back, which causes the overall bunch length to shrink. iii.) After full compression, the bunch size is reduced to a minimum value determined by the initial, uncorrelated energy spread and there is no residual chirp. b.) Schematic of the 4-bend Sector 10 bunch compressor, showing the pre-FACET electron arm (top) and the newly installed positron arm (bottom).	87

5.6	Left: Diagram of the positron production target. The flux-concentrator is the triangular device in the center of the image. Right: Field profiles as a function of z , the distance from the target in mm for the flux-concentrator (solid line) and tapered solenoid (dashed line). The scale for the flux-concentrator is the inner y -axis. Figures reproduced from references [90] and [89].	88
5.7	The original lattice optics for FACET Sector 20 including the W-chicane, final focus and spectrometer. The x -axis shows the longitudinal coordinate s in meters, as measured from the start of Sector 20. The four lattice functions plotted are $\sqrt{\beta_x}$ and $\sqrt{\beta_y}$ in $\text{m}^{1/2}$ on the left y -axis and D_x and D_y in m on the right y -axis. The IP is located at $s = 64$ m. The shape of the W-chicane is reflected by the D_x curve. At the start of the 2013 run, the β -functions were reduced everywhere in the chicane to reduce non-linear magnet errors. The optics shown here were in place during the 2011 and 2012 runs. Figure reproduced from reference [91].	90
5.8	Left: To ameliorate short-range wakefields, the lower-charge positron bunch leads the electron bunch down the linac by half an RF cycle, or roughly 5 cm. Right: Positrons traverse the upper arm of the sailboat chicane and electrons traverse the lower W arm. Both arms act to compress the bunch. The R_{56} of each chicane is tunable so that the difference in arrival time of the electrons and positrons at the entrance to the plasma can be made as small as needed [74].	91
5.9	Beamline diagram of the Sector 20 illustrating the concept of two-bunch generation and critical elements of the experimental area. a) beam notching device, b) transverse deflecting structure, c) initial spectrometer, d) final-focus quadrupole magnets, e) lithium plasma ionization laser, f) lithium vapour column, g) spectrometer imaging quadrupole magnets, h) spectrometer dipole magnet, and i) Cherenkov and phosphor screens. Bend dipole magnets in the W-shaped chicane are each labelled D. The arrow indicates the electron beams direction of motion. Figure and Caption reproduced from reference [45].	92
5.10	a) LiTrack simulation of the under-compressed beam at the end of Sector 20. Note that the beam naturally has a two bunch profile without any collimation. b) The same simulation, but with the center of the spectrum notched out and the low energy tail removed. Note that LiTrack uses z to designate the co-moving bunch coordinate. The beam travels to the left with the head of the beam at values of $z < 0$	93
5.11	Line drawing of the experimental setup showing all of the critical hardware and diagnostic devices. Inset is to scale. Reproduced from reference [30]	95

5.12	Left: The heat-pipe oven on the test bench at UCLA. Right: Bench measurement of the vapor density in the heat-pipe oven for the conditions used in the experiment. Only one half of the plasma density profile is measured, as it is assumed to be symmetric about the midpoint. The length of the vapor region is defined by the full-width, half-maximum (FWHM) of the density measurement and is 132 cm in this case. Image and data from K.A. Marsh and N. Vafaei-Najafabadi.	96
5.13	a) Illustration of the SYAG energy spectrometer. Note that the scintillating YAG crystal intercepts the x-rays emitted by the beam, but not the beam itself. b) Example of a dispersion measurement where the beam energy is varied by ± 150 MeV. The movement of the beam profile is used to calibrate the horizontal axis. c) Example of a beam energy spectrum using the calibration from the measurement shown in b). .	97
5.14	Left: Photograph of the x-band deflecting cavity. The SYAG wiggler magnets can be seen in the background (light blue). Right: Illustration of the principle of the deflecting cavity. At FACET, the cavity deflects vertically as opposed to horizontally. Reproduced from reference [94].	98
5.15	Sample EOS image showing a two-bunch beam profile with a peak-to-peak separation of $160 \mu\text{m}$. The black line is the projection of the image.	99
5.16	Left: Schematic of the FACET laser showing a) Coherent Vitaro oscillator, b) Coherent Legend regen, c) 4-pass pre-amp, d) 4-pass main-amp with Thales Saga green pump lasers e) vacuum relay imaging system and f) the periscope/telescope beam expander leading to the laser transport. Right: Schematic of the laser transport from the FACET laser room to the IP integration box in the tunnel.	101
5.17	a) Image of laser compressor. One of the large ($12 \text{ cm} \times 14 \text{ cm}$) gold compressor gratings is seen facing the camera in the center of the image. The input and output of the laser is depicted by the red arrow. b) Image of the IP box setup in 2014 showing the path of the positron beam (blue), laser (red) and their co-linear trajectory (purple) after the holed, gold folding mirror. A white arrow shows the motion of the kinoform holder for the raster scan. c) Image of the PB box in 2016 showing the path of the positron beam (blue), laser (red) and their co-linear trajectory (purple) after the holed, gold folding mirror. After the laser passes through the kinoform it is sent to a delay stage (not shown) that is used to vary the length of the plasma channel.	102
5.18	Diagram of the <code>facet_daq</code> network architecture in Sector 20. Three different flavors of the diagnostic IOCs are shown. In practice, there are 10 scalar data IOCs, 4 sCMOS IOCs and 4 GigE IOCs involved in the DAQ. Note that the NAS is bonded to the switch and mounted on each of the machines running the camera IOCs, so it acts as a local hard drive with read/write rates greater than the network capacity.	103

6.1	Schematic of the experimental setup. The laser passes through the kinoform and is coupled to the beam axis by a gold mirror with a small central hole. The top right inset shows the laser profile upstream of the lithium oven. A scintillating YAG screen 1.95 m downstream of plasma is used to measure the positron beam profile. A scintillating Lanex screen downstream of the dipole measures the beam energy spectrum. Reproduced from reference [30].	106
6.2	a) An example of an astigmatic high-order Bessel profile. b) The laser profile after correcting the astigmatism. c) An example of a poorly illuminated laser profile. d) An optimized profile after correcting for astigmatism and illumination.	107
6.3	Left: Laser intensity profile for different longitudinal positions of PB delay stage, referred to as a rail measurement. At each location, an image of laser profile is recorded and a vertical lineout is taken from the image. The vertical lineouts are stacked together to create the rail measurement. Right: A simulated rail measurement using the FACET OPTICS code.	108
6.4	A mosaic of images illustrating the principle behind the raster scan. Each subimage shows both the positron beam (central hot spot) and laser profile (ring) at low intensity reflecting off of a titanium foil. The laser is scanned in the transverse plane while remaining parallel to the beam trajectory. The arrows indicate the direction of force that the positron beam experiences for a plasma channel located at the position of the laser.	110
6.5	Left: An image from the upstream camera showing the laser profile and the fit to the central position. Middle: All of the shots from the entire raster scan in the transverse plane. They are binned according to the dotted lines. Right: 2D histogram of the binned shots showing uniform relatively uniform coverage of the scanned plane. . . .	111
6.6	Left: The kick map shows the magnitude and direction of the kick delivered to the beam averaged over approximately 5 shots as the channel location is scanned with respect to the beam trajectory. A net kick of (43.6,48.7) μrad in (x,y) is subtracted from the data. We superimpose two equal intensity contours of the central Bessel peak of the laser measured upstream of the plasma as red dashed lines. Right: Area of the positron beam measured on a YAG screen downstream of the plasma averaged over 10 shots as the channel location is scanned with respect to the beam trajectory. . . .	112
6.7	Left: The expected kicks for an ideal $J_7^2(k_\perp r)$ profile through the first minimum in r . Right: Results of our optimized 2016 raster scan. The shape of our measurements matches what is expected in the ideal case. However, the measurement is much “blurrier”, indicating the ionized region was larger than the maximum of the $J_7^2(k_\perp r)$, which is expected because the ionization process is a threshold effect. The non-zero beam size also contributes to the blurred measurement.	113

6.8	Left: The measured kick map for the optimized single bunch raster scan in 2016. Right: The reconstructed channel profile for this data.	114
6.9	The incoming energy, measured at the SYAG spectrometer, versus the outgoing energy, measured in the dumpline spectrometer. The laser-on shots lose an average of 20 MeV energy compared to the laser off case.	115
6.10	Left: The dumpline energy spectra for all 315 shots in the dataset, 10% of which are laser-off. The energy spectra are corrected for incoming energy jitter. Middle: The average laser-on and laser-off spectra, with error bars showing the local statistical error in the measurement at 20.15 and 20.55 GeV. The error bars have been enhanced by a factor of 5 to make them visible. Right: A histogram of the energy centroids for all shots, corrected for incoming energy. The mean energy loss is 18.9 MeV with a width of 3.2 MeV.	116
6.11	Left: Schematic of the kick measurement showing how beam offset is translated into kick. Right: The result of the kick measurement, where the beam displacement of the downstream YAG screen has been converted into an angle based on the distance of the screen from the plasma channel. The offset of the channel is measured by the laser imaging system.	117
6.12	Left: Change centroid energy versus channel length. The energy centroid is corrected for incoming energy jitter on a shot-by-shot basis. Shots that do not cleanly propagate through the channel are removed from the data. The energy loss for a channel length of 45 cm is 107.7 MeV, or an effective gradient of 239 MeV/m. The data was taken with a plasma density of $n = 1 \times 10^{17} \text{ cm}^{-3}$. Right: QuickPIC simulation of the hollow plasma channel using parameters from the 2014 experiment with $n = 8 \times 10^{16} \text{ cm}^{-3}$. The beam propagates to the left and the 1σ contour is shown by the white dashed line. The black dotted lines at $\pm 240 \mu\text{m}$ and $\pm 290 \mu\text{m}$ are the inner and outer radii of the plasma channel, respectively. Lineouts of the simulated and calculated on-axis E_z field ($r = 0 \mu\text{m}, z$) and the radial variation in E_z at the peak decelerating field ($r, z = 11 \mu\text{m}$), are shown with solid and dashed black lines, respectively. The maximum decelerating field is 220 MeV/m.	118
6.13	Top: QuickPIC simulation of the hollow channel plasma wakefield with a positron drive and witness beam using parameters from the experiment. The drive-witness separation is $160 \mu\text{m}$. Bottom: On-axis longitudinal field and bunch profiles from the simulation. The expected accelerating wakefield at the location of the witness bunch is 135 MeV/m.	119

6.14	Left: Waterfall plot of the EOS signal over the course of the phase ramp scan. The drive beam is the bright stripe on the right of the image and the witness beam extends out to the left. Right: Centroid-to-centroid bunch separation measured over the course of the scan.	120
6.15	Left: Charge in the drive and witness bunch over the course of the phase ramp scan. Right: Energy gain/loss of the witness and energy loss of the drive bunch over the course of the scan.	121
6.16	Change in energy of the drive and witness beams as a function of bunch separation. The energy gain and loss of the beams is scaled by the bunch charge at each step. .	122
6.17	Left: Typical result for matching a laser-on spectrum with a laser-off spectrum at the upstream spectrometer. Right: We compare the spectra of the same laser-on and laser-off shots to measure the energy change of the drive and witness bunches. . . .	123
6.18	Left: Energy change of the drive bunch after propagating in a 25 cm long hollow channel. The mean energy loss of the drive beam 11.5 MeV. Left: Energy change of the witness bunch. The mean energy gain of the witness beam is 19.6 MeV.	124
6.19	Left: The energy loss of the drive beam as a function of charge in the drive beam. The slope is $-1.5 \text{ MeV}/10^8$ particles. Right: The energy gain of the witness beam as a function of charge in the witness beam. The slope is $-4.0 \text{ MeV}/10^8$ particles. Assuming a linear wakefield and a drive beam with constant bunch length. We can infer that the beam loading effect due to the witness beam alone is $-2.5 \text{ MeV}/10^8$ particles.	125
6.20	Result of an OSIRIS-2D simulation in the $r - z$ plane. Left: The density of the beam and plasma from the simulation. The beam propagates along the channel axis to the left. The plasma is initiated with radially gaussian profile centered at $r = 250 \text{ } \mu\text{m}$ with $\sigma_r = 20 \text{ } \mu\text{m}$. The density of the plasma at the peak of the profile is $3 \times 10^{16} \text{ cm}^{-3}$. The beam contains 3×10^9 particles with a bunch length of $40 \text{ } \mu\text{m}$. Right: The longitudinal field associated with this wake.	126
7.1	Simplified schematic of a 1 TeV PWFA Linear Collider from reference [38].	128
A.1	Contour for the integral A.65.	140
A.2	Left: A plot of χ_0 and χ_1 as a function of the normalized outer plasma radius $k_p b$. We compare to the limiting values Ω_0 and Ω_1 . Right: A plot of $A_0/2$ and $A_1/2$ as a function of the normalized outer plasma radius $k_p b$. We compare to the limiting values κ_0 and κ_1 . For these calculations, we chose $k_p a = 5$ and vary $k_p b$ from 5.5 to 10.148	

A.3	The colored lines show the value of the geometric component of P_2 for several values of $r_a = k_p a$ as a function of $r_b = k_p b$ which is greater than r_a . The black dashed line shows the geometric component of P_2 as a function of r_a with a fixed annular width of 0.25. For all plots, $k_p r_0 = 1$	150
C.1	Output of attached code.	157

Chapter 1

Introduction

1.1 Particle Physics and Politics

Particle physicists seek answers to the questions “What are the fundamental forces of nature, what are we made of, and what is the origin of the universe?” These questions are broadly human and have been asked in one form or another since the beginning of civilization. Today, scientists from all over the world collaborate and compete to find new answers to these questions using particle accelerators; machines so large that they span international borders. However, their size and cost are an impediment toward progress. In the 1980s, the United States began developing a new machine, the Superconducting Super Collider (SSC), that would collide protons with a center-of-mass energy of 40 trillion electronvolts, roughly three times the energy of the currently operating Large Hadron Collider (LHC). The SSC project was plagued by cost overruns and mismanagement, but more importantly, public support for fundamental research was flagging in the twilight of the Cold War. Theoretical physicist and Nobel Laureate Steven Weinberg went on the Larry King Show to debate a congressman who opposed the project. Weinberg recalled the following exchange [1]:

I explained that the SSC was going to help us learn the laws of nature, and I asked if that didn’t deserve a high priority. I remember every word of his answer. It was “No.”

The project was canceled in 1992.

As physicists, it is our responsibility to educate the public and inspire new generations of scientists. Unfortunately, awe and wonder rarely add up to dollars and cents. We are forced to be pragmatic and to do more with less. When it comes to particle physics, this means developing technologies that shrink the size and cost of the particle accelerators while increasing their energy and efficiency. Although we fervently believe that our research deserves more funding, we acknowledge that the technology at the heart of these machines hasn’t changed much in the last few decades. We need a technological revolution.

1.2 Limitations of Conventional Accelerators

Particle accelerators come in two forms: linear and circular. Linear accelerators (linacs) accelerate particles using a string of radio-frequency (RF) cavities. The RF cavities store electromagnetic fields and the particles absorb energy as they pass through the cavities. The overall length of a linear accelerator is determined by the final energy W_f and the strength of the accelerating gradient E_z

$$L \propto \frac{W_f}{E_z}. \quad (1.1)$$

The 3 kilometer-long linear accelerator at SLAC was the longest in the world and accelerated particles up to a final energy of 50 gigaelectronvolts (GeV) with an effective gradient of 16 MeV/m*. The SLAC linac formerly fed the SLAC Linear Collider (SLC) with a center-of-mass collision energy of 91 GeV. The SLC was the world's first and only linear collider. The collider turned off in 1998, outgunned by the circular Large Electron-Positron collider (LEP) at CERN which could reach higher collision energies. In 2006, the linac was divided into the LCLS and FACET accelerators.

The final energy of the SLAC linac, and most warm-RF accelerators, is limited by the power output of the klystrons that pump energy into the RF cavities. A more fundamental limitation is the maximum field that can exist inside of the cavity. If the electric field is too large at the surface of the cavity, electrons will be ripped from the cavity walls in a process known as “breakdown” [2]. This process can be circumvented to some extent by cooling the cavity [3] or by using dielectric structures that are more resistant to breakdown [4] and there is on-going research on both of these techniques. However, the breakdown limitation is ultimately unavoidable for cavity-based acceleration, because the strength of the field cannot be made greater than the inter-atomic bond strength.

An alternative method for accelerating particle beams to extremely high energies is a circular machine called a synchrotron. In a linear accelerator the particles are only accelerated once as they pass through the cavities, but in a synchrotron, magnets are used to keep the beam on a circular orbit so that it passes through the same cavity many times. As a result, the accelerating gradient in the cavity can be small and still provide effective acceleration. The LHC, which produces beams with over 100 times the energy of beams produced at SLAC, uses RF cavities with a gradient of only 5 MeV/m, less than one quarter the gradient of the SLAC RF cavities.

There are two effects limiting the final energy that can be achieved in a circular machine. Charged particles produce synchrotron radiation as the magnets bend them around the ring. For each turn around the ring, the particles radiate away some of their energy

$$E_{loss} \propto \frac{\gamma^4}{\rho}, \quad (1.2)$$

*The effective gradient is the final energy produced by the linac divided by the length of the linac. It is smaller than the accelerating gradient inside the RF cavities (22 MeV/m at SLAC) because the linac has other elements along its length that do not contribute to the final beam energy. In addition, some of the RF cavities are left off. They serve as backups to other cavities whose RF power sources may fail.

where γ is the relativistic factor E/mc^2 and ρ is the bending radius of the magnets. The scaling of the synchrotron energy loss is extremely unfavorable for electrons and positrons due to the γ^4 dependence. For example, at LEP the maximum beam energy was 104.5 GeV, corresponding to a relativistic factor $\gamma = 200,000$, and the bending radius of the magnets in the LEP tunnel (now the LHC tunnel) was about 3.5 km. Despite the enormous size of the machine, the radiated energy loss per turn was almost 3 GeV, with a total radiated power of 18 megawatts. That's equivalent to the power consumption of about 5,000 households. In physics, we usually measure progress by orders of magnitude, but increasing the beam energy by a factor of 10 while keeping the bending radius constant would lead to power losses 10,000 times greater, which is a little less than half of the power consumption of the United States. Alternatively, one could choose to keep power loss constant by increasing the bending radius by a factor of 10,000, but the Fermilab, a hypothetical particle accelerator orbiting the Earth, is currently beyond the means of the world economy [5].

The world's highest energy particle colliders are also circular, but they accelerate protons instead of positrons and electrons. For the same relativistic factor γ , a proton stores roughly 2,000 times the energy of an electron due to its mass. Therefore, the proton radiates away only a small fraction ($\sim 8.8 \times 10^{-14}$) of the power that an electron would radiate at the same beam energy. Operating in the same tunnel as LEP, the LHC proton beam loses only 10 keV per turn, which is a measly 6,000 watts. The limitation for circular colliders that accelerate protons is not synchrotron radiation. Instead, it is the magnet strength that limits beam energy. The superconducting magnets installed at the LHC have a magnetic field strength of $B = 8.33$ Tesla and the bending radius for a 7 TeV beam is 2.8 km. Note that the scaling of beam energy with magnetic field strength is favorable compared to equation 1.2

$$E_{beam} \propto B\rho, \quad (1.3)$$

where ρ is again the bending radius. To achieve an order of magnitude increase in beam energy, one can increase the magnetic field strength by an order of magnitude or increase the radius by an order of magnitude or some combination thereof. Efforts are underway to design a Future Circular Collider (FCC) with a center-of-mass collision energy of 100 TeV, with magnet strengths between 16-20 T and a radius of about 16 km [6]. While this machine is technologically feasible, it is likely to cost tens of billion of dollars, and therefore unlikely to be built any time soon.

What are the alternatives? At the moment, there are none. But of the many concepts proposed for high-gradient particle acceleration, plasma acceleration has received the most attention.

1.3 Plasma Acceleration

The most highly-energetic particles observed on Earth are extra-galactic cosmic rays. These particles regularly impact the Earth with energies up to 100 million times greater than those produced by the LHC. Although the acceleration mechanism for the highest energy cosmic rays is unknown, there

is broad consensus that Fermi Acceleration in plasmas driven by supernova remnants accounts for most of the spectrum up to 1 million TeV [7]. We don't have (or want) supernovae on Earth to power our accelerators, but we note that plasmas are an effective medium for producing extremely high energy particles.

Plasmas are fluids composed of free electrons and ions. Under laboratory conditions at FACET, the plasma is diffuse and ephemeral, yet it can support accelerating fields of 100 gigavolts per meter, roughly five thousand times stronger than the 20 megavolts per meter available in copper radio-frequency (RF) cavities in the SLAC linac.

The SLAC RF cavities are not without their own merits. After all, they have been used to produce particle beams for world-class research for over five decades. The cavities are used to shape and control electromagnetic fields so that the particles experience the same accelerating gradient as they travel down the linac. The ability to shape and control fields is called into question when considering plasmas as accelerators. How do you control a fluid that changes shape in the presence of an electromagnetic field?

1.3.1 Laser-Driven Plasma Acceleration

In 1979, Tajima and Dawson recognized that the plasma's response to a high-intensity laser could be used to coax the plasma into a shape that is useful for particle acceleration [8]. We refer to this shape, and the associated fields, as the plasma wakefield. The laser transfers energy to the plasma electrons via the ponderomotive force, leading to longitudinal charge separation and large accelerating fields that scale with the plasma density n

$$E \approx 100\sqrt{n \text{ (} 10^{18} \text{ cm}^{-3}\text{)}} \text{ [GV/m]}. \quad (1.4)$$

This technique requires that the laser pulse length is shorter than the plasma wavelength, which is roughly 100 microns. At the time of their proposal, there were no lasers with the requisite intensity and pulse length, but after the development of chirped-pulse amplification technology [9] in the 1980s, a wave of experiments followed demonstrating the promise of laser wakefield acceleration (LWFA). In 2004, three groups simultaneously published the "dream-beam" results [10, 11, 12], demonstrating acceleration of near mono-energetic electron beams to energies greater than 100 MeV. These results were followed by acceleration above 1 GeV at the L'OASIS laboratory at LBL in 2006 [13], above 2 GeV with the Texas Petawatt laser at UT Austin in 2012 [14], and above 4 GeV with the BELLA laser at LBL in 2014 [15]. In all of these experiments, the accelerated electrons are generated within the plasma wakefield, and therefore have small a emittance. The rapid development of LWFA has led to international attention and serious discussions on using the technology as the source for free-electron lasers (FEL) operating in the soft x-ray regime [16].

1.3.2 Beam-Driven Plasma Acceleration

Five years after Tajima and Dawson's proposal, physicists at SLAC and UCLA developed a related idea called Beam-Driven Plasma Wakefield Acceleration (PWFA) [17, 18]. In a plasma wakefield accelerator, a high-energy, ultrarelativistic electron beam propagates into a plasma and the electric field of the beam forces the plasma electrons away from the beam axis. After the beam passes by, the plasma electrons return to the axis which sets up an oscillation, or wake. The scaling of the wakefields with the plasma density is the same as in equation 1.4.

The first demonstration of a plasma wakefield accelerator took place at Argonne National Laboratory in 1988 [19]. In addition to being the first proof-of-concept, the experiment illustrated a critical principle for PWFA called two-bunch, or "drive-witness" acceleration. In the two-bunch scheme, a high-charge drive beam enters the plasma and creates the wakefield. Energy is transferred from the drive beam to the plasma. A lower charge witness bunch enters the plasma some distance behind the drive bunch. If the witness bunch is at the correct phase of the wake, it will be accelerated and gain energy.

Shortly after the experiments at Argonne, experimenters and theorists recognized that non-linear plasma wakefields, the so-called blowout regime, could be extremely useful for accelerating electrons [20, 21]. But similar to the case of LWFA, there were no facilities capable of producing ultrarelativistic, high-charge, highly-compressed particle beams. That changed in 1997, when scientists from UCLA, USC, and SLAC formed the E157 collaboration to study PWFA using the 30 GeV SLAC electron beam at the Final Focus Test Beam facility (FFTB) [22]. Although the E157 experiment was not successful in its stated goal of achieving 1 GeV acceleration in a plasma, it laid the groundwork for its successor experiments (E162, E164 and E167), which demonstrated acceleration of positrons [23], 2.7 GeV acceleration of electrons [24], and energy doubling of a 42 GeV electron beam [25], respectively. In each of these experiments, the accelerated particles had a continuous energy spread up to the maximum energy, whereas most accelerator applications require beams with a narrow spread in energy. Accelerating beams with small energy spreads was the top priority for PWFA research heading into FACET, the successor facility to the FFTB. The results of these studies are described in Section 2.8.2.

1.3.3 FACET and the E200 Collaboration

The Facility for Advanced Accelerator Experimental Tests (FACET) [26] is a unique research facility for studying plasma wakefield acceleration. The FACET linac is the longest linear accelerator that is currently in operation. At 20.35 GeV, the electron and positron beams produced at FACET are the highest energy lepton beams being produced anywhere in the world, and FACET is the only facility offering positrons for PWFA research. FACET produces particle beams of extraordinary intensity, with peak currents greater than 19 kA, a key figure-of-merit for PWFA experiments. The peak current available at FACET is also the largest in the world.

Plasma wakefield acceleration experiments at FACET are carried out by the E200 Collaboration. E200 is primarily a collaboration between researchers from SLAC National Accelerator Laboratory and the University of California at Los Angeles, with contributions from the University of Oslo in Norway, École Polytechnique in France, the Max Planck Institute in Germany, and Tsinghua University in China. The E200 collaboration is also responsible for the E225 hollow channel plasma wakefield experiment described in this thesis. E200 is the most recent incarnation of the SLAC-UCLA partnership that has advanced the field of beam-driven plasma wakefield acceleration over the last three decades.

1.4 The Problem with Positrons

Despite the initial success with positron PWFA experiments in 2003, there was no progress toward demonstrating the viability of positron acceleration in plasma in the decade that followed. There are two issues that have impeded the progress of positron PWFA research. The first is practical; the only facilities capable of producing intense positron bunches for PWFA were the FFTB from 2000-2003 and FACET from 2014-2016. Positrons are created by sending high energy electrons into a heavy metal target. At FACET, a 20 GeV electron beam is utilized for positron generation. Other currently operating PWFA facilities produce 100 MeV electron beams which cannot be used for efficient positron generation. We also note that electrons accelerated via PWFA may have uses in FELs, medical linacs, and as γ -ray Compton sources, but positrons are only used in high-energy particle colliders. Therefore, positron PWFA research facilities are funded by the high-energy physics (HEP) community, and HEP has been reluctant to invest in positron PWFA due to concerns about the tractability of accelerating positrons in plasma. This brings us to the second issue plaguing positron PWFA: the asymmetry of the plasma response to electron and positron beams.

Conventional RF accelerators are beam-charge agnostic. They have been used to accelerate electrons, positrons, protons, and heavy ions. They may one day be used to accelerate muons [27]. The SLAC linac, when used as the source for the SLC, accelerated electrons and positrons on the same RF pulse by injecting the positron beam 180° out of phase with the electron beam. On the other hand, plasmas are composed of mobile electrons and immobile ions, which break the symmetry of the plasma as an accelerator. The shape of the plasma wakefield changes dramatically when sourced by an electron or positron beam that are identical, except for the sign their charge.

One simple solution to this problem is to develop a plasma source with a mobile positive species and an immobile negative species, but such plasmas are hard to come by. At the Advanced Accelerator Concepts Workshop in 2002, Professor Bob Siemann was presenting on the challenges of positron acceleration in plasma. Acknowledging that he had no solutions to the problem, he solicited advice from the audience. A physicist from USC raised his hand to offer a suggestion. Bob turned to him and said, “If you say anti-plasma, I will hit you.” The physicist lowered his hand.

The purpose of this thesis is to present a solution to the problem of accelerating positrons in plasma — one that would elicit a more positive reaction from Professor Siemann.

1.4.1 Hollow Channel Plasma Wakefield Acceleration

Hollow Channel Plasma Wakefield Acceleration (HC-PWFA), was proposed over twenty years ago as a technique for laser confinement in a laser wakefield accelerator [28]. It was quickly realized that the hollow channel has attractive properties for particle acceleration. The longitudinal fields are radially uniform, and there are no transverse forces due to background ions or on-axis plasma electrons. This circumvents a major issue for positron PWFA in uniform plasmas, where positrons are defocused by the background ions or experience a non-linear focusing force due to plasma electrons that are sucked in to the beam volume.

HC-PWFA holds exceptional theoretical promise, with dozens of publications on the topic over the last twenty years. However, the concept has never been realized experimentally because it is challenging to produce a hollow channel plasma. FACET is the first experimental facility that combines ultrarelativistic positron beams with a laser that is suitable for creating a hollow channel plasma [29]. Our first experimental run at FACET showed that it is indeed possible to generate large amplitude wakefields inside the channel while preserving the quality of the positron beam that drives the wake [30]. In subsequent runs, we demonstrated the acceleration of a trailing positron beam in the hollow channel wakefield [31].

1.5 Outline

In this thesis, we describe the motivation, theory, development and experimental results of the E225 Hollow Channel Plasma Wakefield Experiment.

In Chapter 2, we introduce the concept of beam-driven plasma wakefield acceleration in uniform plasmas. We review the theory of PWFA in the linear and nonlinear regime and highlight two important experimental results on PWFA in uniform plasmas at FACET. We also examine limitations of positron PWFA in uniform plasmas and motivate the study of the Hollow Channel Plasma Wakefield Accelerator.

In Chapter 3, we derive the theory of HC-PWFA in the linear regime. We compare our results with simulations and explore the transition to the nonlinear regime. We also discuss beam instabilities in the hollow channel and provide analytic estimates of their effect on the experiment.

In Chapter 4, we describe the development of the high-order Bessel “kinoform” optic. We review the theory of Fourier Optics and the Fresnel propagation integral. We describe a numerical code used to calculate laser profiles at FACET. Finally, we review the theory of tunnel ionization for laser generated plasmas.

In Chapter 5, we describe the FACET accelerator test facility. We provide a source-to-sink description of the FACET linac, as well as positron production and the FACET Titanium-Sapphire (Ti:Saph) laser system. We discuss longitudinal phase space dynamics and the generation of the two-bunch profile for beam loading experiments.

In Chapter 6, we review the results of the E225 Hollow Channel Plasma Wakefield Experiment. This was the first demonstration of HC-PWFA. We describe the novel experimental techniques that were developed to determine the shape of the channel. We use beam deflection data to determine the length of the channel and use energy loss data to measure the strength of the wakefield.

We provide a summary of this thesis in Chapter 7 and speculate on plasma-based linear colliders. Finally, we discuss future directions for this work.

All calculations are in CGS units unless otherwise noted.

Chapter 2

Electron and Positron Acceleration in Uniform Plasmas

2.1 Introduction

We begin this chapter by providing a theoretical description of plasma wakefield acceleration in uniform plasmas in the linear regime, following the treatment by Keinigs and Jones [32], with insights from Ruth and Chao [17], Blumenfeld [33], and Schroeder [34]. We discuss the limitations of PWFA in the linear regime and describe the motivation for PWFA in the non-linear “blowout” regime. Next, we provide a brief theoretical introduction to the theory of blowout wakes using Lu’s sheath description [35]. In our description of PWFA in the non-linear regime, we touch on the topic of beam loading and its application to recent experiments at FACET. Finally, we discuss positron acceleration in the non-linear regime, including recent experimental successes and known challenges which motivate research into hollow channel plasma acceleration.

2.2 Plasma wakefield acceleration in the linear regime

In this section, we describe the response of the plasma to the beam, the “wakefield”, using a fluid description of the plasma. The motion of the plasma fluid is describe by the continuity equation

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\vec{v}) = 0. \quad (2.1)$$

The continuity equation says that the local change in the plasma density n is equal to the amount of plasma flowing into or out of the region. This is equivalent to charge conservation. The response

of the plasma to fields is governed by the Lorentz force law

$$m \frac{\partial(n\vec{v})}{\partial t} = en \left(\vec{E} + \frac{\vec{v} \times \vec{B}}{c} \right), \quad (2.2)$$

and the fields evolve according to Maxwell's equations

$$\nabla \cdot \vec{E} = 4\pi\rho, \quad (2.3)$$

$$\nabla \cdot \vec{B} = 0, \quad (2.4)$$

$$\nabla \times \vec{E} = -\frac{1}{c} \frac{\partial \vec{B}}{\partial t}, \quad (2.5)$$

$$\nabla \times \vec{B} = \frac{4\pi}{c} \vec{J} + \frac{1}{c} \frac{\partial \vec{E}}{\partial t}. \quad (2.6)$$

These equations completely determine the evolution of the plasma and fields, but combining these equations produces a set of non-linear PDEs that cannot be solved analytically for arbitrary sources.

To make headway, we need to use some simplifying assumptions about the plasma and source. The plasma is initially uniform and neutral, with both the electron and ion number density equal to n_0 . The ions are much heavier than the plasma electrons and do not evolve on time scales relevant for this discussion. The dynamic quantity $n(r, z, t)$ refers to the plasma electron density. The ions are represented by a static, uniform background with positive charge and density n_0 .

The source is an ultrarelativistic particle beam with $v = c$ and density $n_b \ll n_0$. The requirement that the beam is ultrarelativistic allows us to ignore the evolution of the beam distribution in response to forces generated in the plasma. In this case, we say that the beam is “rigid.” The requirement that the beam density is much less than the plasma density means that we can treat the system perturbatively. The plasma electron density is given by $n(r, z, t) = n_0 + n_1(r, z, t)$, where n_1 is the perturbation to the plasma density. The plasma velocity $\vec{v}$, and fields $\vec{E}$ and $\vec{B}$ are also perturbations. Plugging our expression for n into equations 2.1 and 2.2 gives

$$\frac{\partial n_1}{\partial t} = -n_0 \nabla \cdot \vec{v}, \quad (2.7)$$

$$\frac{\partial \vec{v}}{\partial t} = \frac{en_0}{m_e} \vec{E}, \quad (2.8)$$

where we have only kept terms that are first-order in the perturbative quantities. Taking the divergence of equation 2.8 gives

$$\frac{\partial \nabla \cdot \vec{v}}{\partial t} = \frac{en_0}{m_e} \nabla \cdot \vec{E}. \quad (2.9)$$

We plug in for $\nabla \cdot \vec{E}$ using Gauss's law

$$\nabla \cdot \vec{E} = 4\pi e(n_1 + n_b). \quad (2.10)$$

Taking a time derivative of equation 2.7 allows us to combine the equations which gives

$$\frac{\partial^2 n_1}{\partial t^2} + \omega_p^2 n_1 = -\omega_p^2 n_b, \quad (2.11)$$

with plasma frequency

$$\omega_p = \sqrt{\frac{4\pi e^2 n_0}{m_e}}. \quad (2.12)$$

Experiments at FACET typically use plasma densities in the range of $1 \times 10^{16} - 1 \times 10^{17} \text{ cm}^{-3}$. The plasma frequency is greater than 10 THz and the wavelength is less than 500 μm , or about 200 times smaller than the wavelength used in SLAC's s-band cavities.

Experiments that use plasmas as the accelerating medium are challenging because they require temporal and spatial control of the drive beam at the fs/ μm scale. But the payoff is that the fields associated with the plasma oscillations are enormous. The characteristic field strength is

$$E_0 = \frac{m_e c \omega_p}{e} \quad (2.13)$$

or 30 GeV/m at $n_0 = 1 \times 10^{17} \text{ cm}^{-3}$, which is over a thousand times the gradient of the SLAC s-band cavity.

2.3 Temporal response of the plasma

One interesting aspect of equation 2.11 is that it is a driven harmonic oscillator without a damping term, *i.e.* the plasma is non-dissipative. This is because we assumed the plasma has zero temperature (all particles are at rest before the drive beam arrives) and it is a collisionless, perfect conductor.

We want to examine the response of the plasma to an ultrarelativistic ($\beta = 1$), radially symmetric drive beam with beam density given by

$$n_b(r, z, t) = R(r)Z(z - ct). \quad (2.14)$$

It is convenient to use the co-moving coordinate $\xi = z - ct$, where ξ represents the position along the bunch. Re-expressing equation 2.11 in terms of the new coordinate gives

$$\frac{\partial^2 n_1(r, \xi)}{\partial \xi^2} + k_p^2 n_1(r, \xi) = -k_p^2 R(r)Z(\xi), \quad (2.15)$$

where $k_p = \omega_p/c$ is the plasma wavenumber. The skin depth $k_p^{-1} = c/\omega_p$ is the characteristic length scale. Note that $n_1(r) = n_b(r)$.

We can explicitly calculate the response of the plasma using a delta function source in the co-moving coordinate $Z(\xi) = \delta(\xi)$. Integrating both sides of 2.15 from $\xi = -\epsilon$ to $\xi = \epsilon$ and taking the

limit as $\epsilon \rightarrow 0$ gives

$$\left. \frac{\partial n_1}{\partial \xi} \right|_{\xi=0} = -k_p^2 R(r). \quad (2.16)$$

For $\xi < 0$, there is no source and no plasma perturbation. For $\xi > 0$, the source has passed but the perturbation remains. The solution to equation 2.15 with no source is $A \sin k_p \xi + B \cos k_p \xi$. The delta function tells us the slope of our solution at $\xi = 0$. We immediately can infer that $B = 0$ and $A = -k_p R(r)/2\pi r$, so our Green's function solution is

$$n_1^G(r, \xi) = -k_p R(r) \sin(k_p \xi) \Theta(\xi), \quad (2.17)$$

where $\Theta(\xi)$ is the Heaviside step function. We can use the solution to derive the plasma perturbation for arbitrary source distributions $n_b = R(r)Z(\xi)$ by convolving the source with the Green's function. For instance, if the bunch distribution is Gaussian, we get

$$n_1(r, \xi) = -k_p \frac{N_b}{\sqrt{(2\pi)^3 \sigma_r^2 \sigma_z}} e^{-\frac{r^2}{2\sigma_r^2}} \int_{\xi}^{\infty} e^{-\frac{\xi'^2}{2\sigma_z^2}} \sin k_p(\xi - \xi') d\xi'. \quad (2.18)$$

Figure 2.1 shows the response of the plasma due to a bi-Gaussian drive beam with $n_b \ll n_0$.

2.4 Fields in the plasma

In the previous section, we showed that the response of the plasma to the drive beam is a sinusoidal oscillation in time with frequency ω_p . We would like to know the variation of the fields in time and space so that we can calculate their effect on the drive beam.

Returning to Maxwell's equations, we take the cross product of equation 2.5 and the time derivative of equation 2.6 and use the curl of a curl vector identity

$$\nabla \times \nabla \times \vec{E} = \nabla(\nabla \cdot \vec{E}) - \nabla^2 \vec{E} \quad (2.19)$$

to get

$$\nabla^2 \vec{E} - \frac{1}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} = \frac{4\pi}{c^2} \frac{\partial \vec{J}}{\partial t} + 4\pi \nabla \rho. \quad (2.20)$$

The sources can be divided into a component due to the drive beam and a component due to the plasma so $\rho = \rho_b + \rho_p$ and $\vec{J} = \vec{J}_b + \vec{J}_p$ with $\vec{J}_b = c\rho_b \hat{z}$. Equation 2.2 provides a relationship between $\vec{J}_p$ and $\vec{E}$. Assuming the effect of the B field is negligible, we can multiply both sides of the equation by e/m to get

$$\frac{\partial \vec{J}_p}{\partial t} = \frac{\omega_p^2}{4\pi} \vec{E}. \quad (2.21)$$

Figure 2.1: A plot of the density perturbation due to a relativistic beam with $N_b = 4 \times 10^8$ electrons and dimensions $\sigma_r = 20 \mu\text{m}$ and $\sigma_z = 15 \mu\text{m}$ propagating to the right in a plasma with density $n_0 = 8 \times 10^{16}$. The black lines are the radial and longitudinal on-axis lineouts of the plasma density.

Substituting this expression into equation 2.20 and rearranging gives

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \frac{\omega_p^2}{c^2} \right) \vec{E} = \frac{4\pi}{c} \frac{\partial \rho_b}{\partial t} \hat{z} + 4\pi \nabla (\rho_b + \rho_p). \quad (2.22)$$

The magnetic field is obtained directly from equation 2.22 by applying Faraday's law 2.5

$$\left(\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \frac{\omega_p^2}{c^2} \right) \vec{B} = -4\pi \nabla \times \rho_b \hat{z}. \quad (2.23)$$

The EM wave equations 2.22 and 2.23 together with the plasma response equation 2.11 give the full response of the fields and plasma to a non-evolving source.

2.4.1 Longitudinal Field

We now attempt to solve for the accelerating field E_z by using a Fourier transform in ξ

$$E_z(\xi) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{E}_z(k) e^{ik\xi} dk. \quad (2.24)$$

The Fourier transform is useful because it allows us to express derivatives with respect to t and z as multiplicative operators. Recalling that $\xi = z - ct$, we have

$$\frac{\partial}{\partial z} \rightarrow ik, \quad (2.25)$$

$$\frac{1}{c} \frac{\partial}{\partial t} \rightarrow -ik. \quad (2.26)$$

The Laplacian can be separated into a transverse and longitudinal component $\nabla^2 = \nabla_{\perp}^2 + \partial_z^2$ and plugged into equation 2.22

$$(\nabla_{\perp}^2 - k_p^2) \hat{E}_z = \frac{4\pi}{c} \frac{\partial \hat{\rho}_b}{\partial t} + 4\pi \frac{\partial}{\partial z} (\hat{\rho}_b + \hat{\rho}_p), \quad (2.27)$$

with $k_p = \omega_p/c$. Note that the derivatives with respect to z and t cancel. This is a consequence of using an ultrarelativistic source with $v = c$, which means the phase velocity $v_{\phi} = \omega_p/k_p = c$. If $v \neq c$, there is a factor of $1/\gamma^2$ that appears in the combined derivative. We will return to this point later in the chapter.

We specify the source to be a delta function in both r and ξ

$$\rho_b = e \frac{\delta(r - r_0)}{2\pi r} \delta(\xi). \quad (2.28)$$

We are still assuming radial symmetry, so the delta function in r means that we are considering a ring-shaped particle with radius r_0 . The Fourier transform in ξ of the source is simply

$$\hat{\rho}_b = e \frac{\delta(r - r_0)}{2\pi r}. \quad (2.29)$$

Transforming ρ_p and plugging the expression for $\hat{\rho}_b$ into 2.11 gives

$$\hat{\rho}_p = e \frac{k_p^2}{k^2 - k_p^2} \frac{\delta(r - r_0)}{2\pi r}. \quad (2.30)$$

We insert the expressions for $\hat{\rho}_b$ and $\hat{\rho}_p$ into 2.27 to get

$$(\nabla_{\perp}^2 - k_p^2) \hat{E}_z = 4\pi e \frac{ik k_p^2}{k^2 - k_p^2} \frac{\delta(r - r_0)}{2\pi r}. \quad (2.31)$$

The contributions from ρ_b and J_b canceled out. This is because an ultrarelativistic charge does not source an E_z field. Rather, the fields are “pancaked” and point radially outward. What remains is the field due to the density-modulated plasma distribution. We refer to this type of wave as a longitudinal wave (even though it has field components in all directions), similar to waves on a slinky.

The operator $\nabla_{\perp}^2$ is expressed in cylindrical coordinates as

$$\nabla_{\perp}^2 = \frac{1}{r} \frac{d}{dr} r \frac{d}{dr} + \frac{1}{r^2} \frac{d^2}{d\phi^2}. \quad (2.32)$$

Derivatives with respect to ϕ are zero because the source is radially symmetric. Expanding $\nabla_{\perp}^2$ in 2.27 gives

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \hat{E}_z}{\partial r} \right) - k_p^2 \hat{E}_z = 4\pi e \frac{ik k_p^2}{k^2 - k_p^2} \frac{\delta(r - r_0)}{2\pi r}. \quad (2.33)$$

The LHS of equation 2.33 is the modified Bessel equation of order zero. If there is no source the solution is

$$\hat{E}_z(r) = AI_0(k_p r) + BK_0(k_p r) \quad (2.34)$$

where I_m and K_m are the modified Bessel functions of order m . The Green’s function solution for a radial delta function is given by Jackson in section 3.11 [36]

$$g_0(r, r_0) = 4\pi I_0(k_p r_{<}) K_0(k_p r_{>}). \quad (2.35)$$

The notation $r_{<}$ means that $r < r_0$ and $r_{>}$ means $r > r_0$. Another way to express the same concept is

$$g_0(r, r_0) = 4\pi [I_0(k_p r) K_0(k_p r_0) \Theta(r_0 - r) + I_0(k_p r_0) K_0(k_p r) \Theta(r - r_0)]. \quad (2.36)$$

So the solution for the Fourier transformed E_z is

$$\hat{E}_z = 2e \frac{ik k_p^2}{k^2 - k_p^2} g_0(r, r_0). \quad (2.37)$$

Note that $g_0(r)$ is not a function of k , so we can invert the transform to find the E_z field due to a delta function source

$$E_z(r, \xi) = ek_p^2 \frac{i}{\pi} g_0(r, r_0) \int_{-\infty}^{\infty} \frac{k}{k^2 - k_p^2} e^{ik\xi} dk. \quad (2.38)$$

There are two poles located at $k = \pm k_p$. The integral is over the real axis and hits both poles, but this integral can be shown to be finite by evaluating it as a contour integral in the complex plane. We define the contour C with two parts C_1 and C_2 . C_1 is the original integral along the real axis and C_2 is an arc with a radius that goes to infinity. We use causality to choose C_2 such that its total contribution to the integral is zero. This is achieved by closing in the upper half plane, because in

Figure 2.2: Contour for the integral 2.38.

that case $ik\xi < 0$ for $\xi > 0$. The integral along C_2 is dominated by the exponential, which goes to zero as $ik\xi \rightarrow -\infty$. This imposes causality because we only have fields in the case that $\xi > 0$. The contour for this integration is shown in figure 2.2.

Our contour encloses two poles so we can use the Cauchy Integral Theorem to find the value of the integral

$$2\pi i f(z_0) = \oint \frac{f(z)dz}{z - z_0}. \quad (2.39)$$

We use partial fractions to make our integral take the form:

$$E_z = ek_p^2 \frac{i}{\pi} g_0(r, r_0) \oint \frac{ke^{ik\xi}}{k^2 - k_p^2} dk = ek_p^2 \frac{i}{\pi} g_0(r) \oint \frac{e^{ik\xi}}{2} \left[\frac{1}{k - k_p} + \frac{1}{k + k_p} \right] dk. \quad (2.40)$$

Since the C_2 contribution is zero, the contour integral is equal to the original integral 2.38. Applying the Integral Theorem we get

$$E_z^G(r, \xi) = -ek_p^2 g_0(r, r_0) \cos(k_p \xi) \Theta(\xi), \quad (2.41)$$

where the Heaviside function emerges due to causality based on our choice of contour.

At this point, we should comment that it is no surprise that the E_z field is proportional to cosine, since the density perturbation that sourced it is proportional to sine (equation 2.17) and Gauss's law tells us that $E_z = 4\pi e \partial_z n_1$. However, in deriving the fields with a δ -function source, we learn the explicit form of the radial Green's function 2.36, which will allow us to calculate the radial dependence of the fields with arbitrary beam shapes.

The expression for E_z 2.41 is often referred to as the single-particle wake function W_z . The E_z

Figure 2.3: A plot of the longitudinal electric field for the same beam and plasma parameters used in Fig. 2.1. The longitudinal field lags the density perturbation by 90 degrees. The black lines are the radial and longitudinal on-axis lineouts of the field.

field sourced by the beam distribution is the wake function convolved with the beam density

$$E_z(r, \xi) = -ek_p^2 \int_{\xi}^{\infty} \int_0^{\infty} R(r') g_0(r, r') Z(\xi') \cos k_p(\xi - \xi') r' dr' d\xi'. \quad (2.42)$$

The E_z field corresponding to the density perturbation in Figure 2.1 is plotted in Figure 2.3. Note that the E_z field lags the density perturbation by 90 degrees. The radial shape of the E_z field is the convolution of the gaussian charge distribution with the K_0 and I_0 functions. It is “wider” than the plasma perturbation.

2.4.2 Transverse Forces

The procedure we used to solve for E_z can also be used to solve for the transverse components of E and B . Returning to equations 2.22 and 2.23 we have

$$(\nabla_{\perp}^2 - k_p^2) \hat{E}_r = 4\pi \frac{\partial}{\partial r} (\hat{\rho}_b + \hat{\rho}_p), \quad (2.43)$$

$$(\nabla_{\perp}^2 - k_p^2) \hat{B}_{\phi} = 4\pi \frac{\partial}{\partial r} \hat{\rho}_b \quad (2.44)$$

where we have used the fact that the source is radially symmetric and has velocity $\vec{v} = c\hat{z}$. We are typically interested in the transverse force on the beam particles, which we refer to as $W_{\perp}$. The Lorentz force law tells us that

$$\vec{W}_{\perp} = E_r \hat{r} + \frac{\vec{v} \times B_{\phi} \hat{\phi}}{c}. \quad (2.45)$$

Plugging in $\vec{v} = c\hat{z}$ we have $\vec{W}_{\perp} = (E_r - B_{\phi})\hat{r}$. Combining equations 2.43 and 2.44 gives

$$(\nabla_{\perp}^2 - k_p^2) \hat{W}_r = 4\pi \frac{\partial \hat{\rho}_p}{\partial r}. \quad (2.46)$$

Again, we see that the beam source cancels out. For an ultrarelativistic beam propagating in vacuum, the repulsive space charge force is exactly canceled by the attractive magnetic force of co-propagating currents. In the derivation of the E_z field, we found that the longitudinal force W_z was sourced entirely by the plasma perturbation. The same is true of the radial force W_r .

Plugging in 2.30 for $\hat{\rho}_p$ we get

$$(\nabla_{\perp}^2 - k_p^2) \hat{W}_r = 4\pi e \frac{k_p^2}{k^2 - k_p^2} \frac{\partial \hat{\rho}_b}{\partial r}. \quad (2.47)$$

This expression looks a lot like 2.27, but there is an important difference. In Cartesian coordinates, the Laplacian acts on each component of the vector equally, *i.e.* $\nabla^2 \vec{A} = (\nabla^2 A_x, \nabla^2 A_y, \nabla^2 A_z)$. But in polar coordinates we have

$$(\nabla^2 \vec{A})_r = \nabla^2 A_r - \frac{A_r}{r^2} - \frac{2}{r^2} \frac{\partial A_{\phi}}{\partial \phi}, \quad (2.48)$$

$$(\nabla^2 \vec{A})_{\phi} = \nabla^2 A_{\phi} - \frac{A_{\phi}}{r^2} + \frac{2}{r^2} \frac{\partial A_r}{\partial \phi}, \quad (2.49)$$

$$(\nabla^2 \vec{A})_z = \nabla^2 A_z. \quad (2.50)$$

If there is variation ϕ , the r and ϕ components are coupled, which is a huge problem when trying

Figure 2.4: A plot of the transverse force experienced by a beam particle for the same beam and plasma parameters used in Fig. 2.1 and Fig. 2.3. The force is in phase with the density perturbation and leads the longitudinal field by 90 degrees. The black line is the radial lineout of the force at the center of the drive bunch.

to solve these equations. In our case, there is no variation in ϕ , so we can expand the LHS of 2.30

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \hat{W}_r}{\partial r} \right) - \left(\frac{1}{r^2} + k_p^2 \right) \hat{W}_r = 4\pi e \frac{k_p^2}{k^2 - k_p^2} \frac{\partial \hat{\rho}_b}{\partial r}. \quad (2.51)$$

The homogeneous solution to 2.51 are the modified Bessel functions of order one and the corresponding Green's function is

$$g_1(r, r_0) = 4\pi I_1(k_p r_0) K_1(k_p r). \quad (2.52)$$

The Fourier transformed, single particle, transverse wakefunction is therefore

$$\hat{W}_r = 2e \frac{k_p^2}{k^2 - k_p^2} g_1(r, r_0). \quad (2.53)$$

Inverting the transform gives

$$W_r = \frac{ek_p^2 g_1(r, r_0)}{\pi} \int_{-\infty}^{\infty} \frac{1}{k^2 - k_p^2} e^{ik\xi} dk. \quad (2.54)$$

This integral is evaluated with the same procedure used to find 2.41, but in that case the partial fractions yield a sine function instead of a cosine

$$W_r = -ek_p g_1(r, r_0) \sin(k_p \xi). \quad (2.55)$$

Convolving the wake function with the beam distribution gives

$$F_r(r, \xi) = -ek_p \int_{\xi}^{\infty} \int_0^{\infty} \frac{\partial R(r')}{\partial r} g_1(r, r') Z(\xi') \sin k_p(\xi - \xi') r' dr' d\xi'. \quad (2.56)$$

The transverse force due to the plasma perturbation in Figure 2.1 is plotted in Figure 2.4.

2.5 Comment on the approximations used in this derivation

In deriving the longitudinal field and transverse force experienced by the beam, we made two critical approximations. The first is that the beam density is much less than the plasma density, and therefore, the dynamic response of the plasma can be treated perturbatively. This approach is well established and accurately describes experiments carried out in the linear regime [19]. The second approximation is that the beam is ultrarelativistic traveling at speed c . This allows us to consider the beam distribution to be static, and unaffected by plasma forces. Indeed, for the ultrarelativistic beam at FACET, where $\gamma = 40,000$, the velocity of the beam particles $\beta = v/c$ is nearly constant, despite the fact that beam particles may lose up to 99% percent of their initial energy [37]. On the other hand, the transverse forces in the plasma are extremely large, exceeding 1 MT/m. These forces do affect the beam distribution, causing the beam to evolve on the length scale of a plasma-betatron wavelength*

$$\lambda_\beta = \sqrt{2\gamma} \lambda_p. \quad (2.57)$$

Considering a plasma density $n = 1 \times 10^{17} \text{ cm}^{-3}$, typical of experiments at FACET, the plasma wavelength is roughly 100 μm , while $\lambda_\beta = 3 \text{ cm}$. This tells us two things. First, the length scale for beam evolution is much larger than the plasma wavelength. This motivates the “quasi-static” approximation used in the QuickPIC code discussed in Section 2.7.2. Second, since we are interested in meter-scale plasmas, we see that the beam particles will undergo many oscillations as they traverse

*The plasma-betatron wavelength considered here is due to the force of the plasma ions in the blow-out regime. In the linear regime, the plasma-betatron wavelength is longer than the blowout regime by a factor of $\sqrt{n_0/\delta n}$.

the plasma, so we cannot ignore the effect of the plasma on the beam distribution.

We make one final note on the ultrarelativistic assumption. Mathematically, setting $v = c$ greatly simplifies solving the wave equation and executing the Fourier transforms. However, Keinigs and Jones point out that there is a danger in making this assumption before the inverse Fourier transform is performed, because it gives incorrect limiting forms for the radial fields as the plasma density is taken to zero. Nevertheless, the expressions derived above provide an extremely accurate description of the forces when compared to simulation results with $\gamma = 40,000$.

2.6 Limitations of the Linear Regime

An important aspect of the linear regime is that the plasma response is independent of the sign of the charge of the drive beam. The density perturbation is a “ripple” on the surface of the plasma, and both electrons and positrons can be accelerated in the crests and troughs of the plasma wave. In addition, the wake has both focusing and defocusing phases. Consider the sketch in Figure 2.5 from the original PWFA paper by Ruth and Chao [17] showing the cosine-like longitudinal field and sine-like transverse force. Irrespective of beam charge, there is a region of the wake that is both accelerating and focusing. In our vision for a linear collider that uses plasmas to accelerate particles [38], we anticipate using electron drive beams to create the wake and positron witness beams to be accelerated in the wake. Noting that the two species have opposite charge, we should place the electron drive bunch in the decelerating and focusing phase of the wake extending from 0 to $\pi/2$. In this region, the drive beam gives up energy to the wake, while being focused by the plasma. The plasma focusing ensures that the beam can be channeled through the plasma for lengths much greater than the vacuum β^* . The positron witness beam should be located in the region of the wake from $3\pi/2$ to 2π , where it will be accelerated and focused.

At this point in our analysis, the linear regime appears to be quite promising. However, a fundamental limitation arises when we consider the specific nature of the focusing force. For convenience, we approximate our Gaussian beam using a parabolic distribution [17]

$$R(r) = 1 - \frac{r^2}{\sigma_r^2}. \quad (2.58)$$

This distribution allows us to carry out the radial integral in eq. 2.56 with the result

$$F_r(r) \propto -\frac{1}{\sigma_r^2} \left(K_2(k_p \sigma_r) I_1(k_p r) - \frac{r}{k_p \sigma_r^2} \right). \quad (2.59)$$

The focusing force has a linear focusing component proportional to r/σ_r^2 and a non-linear defocusing component proportional to $I_1(k_p r)$. This implies that the beam radius should be small ($k_p \sigma_r \ll 1$) in order for the beam to only sample the linear part of the field.

Figure 2.5: A sketch comparing the accelerating and focusing fields in a plasma wakefield accelerator. The fields are simultaneously focusing and accelerating for 90 degrees per cycle, independent of beam charge. The beam propagates to the left. Reproduced from ref. [17].

We can now pose the question “Does this type of focusing force provide a stable equilibrium for the transverse beam size?” The transverse beam size evolves according to the beam envelope equation

$$\sigma_r'' + k^2 \sigma_r = \frac{\varepsilon^2}{\sigma_r^3} \quad (2.60)$$

where k is the wavenumber of a particle due to the force of a focusing field and ε is the beam emittance. We say that the beam is “matched” if the focusing k^2 exactly balances the pressure due to the beam emittance. This occurs for a beam size $\sigma_r = \sqrt{\beta\varepsilon}$, where $\beta = 1/k$ is the reduced wavelength of the focusing field. If the beam is mismatched, meaning that the focusing is too strong, the beam size will oscillate around the matched value.

In traditional accelerators, k is set by the linear focusing force of external magnetic fields while in the plasma, k is set by the strength of the transverse plasma wakefield. Let’s assume for the moment that $k_p \sigma_r \ll 1$, so the beam is only sampling the linear part of the field. In this case, $F_r \propto r/k_p \sigma_r^4$, meaning that the force on the beam depends on the beam size. Figure 2.6 shows the the transverse force experienced by narrow beams. Note the rapid increase in focusing strength (slope of F_r) near $r = 0$ with decreasing transverse beam size σ_r .

If there is mismatch between the beam emittance and plasma focusing strength, the beam size may decrease and the focusing force will increase. This leads to a run-away collapse of the beam in the plasma. When the collapsing beam particles are electrons, the process terminates when there are no more plasma electrons left in the beam volume. We call this the “blow-out” regime. In the blow-out regime, the beam particles are subjected to the strong transverse forces due to the stationary, uniform ion background. It is possible to match the beam into the plasma wakefield because the ions provide a linear focusing force. Compare this to the case of a collapsing positron beam, which

Figure 2.6: The transverse wakefield field as a function of r for three different values of the beam normalized beam width $k_p \sigma_r$. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$.

pulls plasma electrons into the beam volume as opposed to expelling them. The resulting force is highly non-linear and variable along the bunch coordinate z , inevitably leading to emittance growth. Figure 2.7 contrasts the trajectory of plasma electrons under the influence of intense electron and positron beams.

The collapse of the positron beam terminates when the transverse beam profile of the positron beam takes on the shape of a Bennett profile [39]

$$R(r) = \frac{1}{\left[1 + \frac{r^2}{a^2}\right]^2}, \quad (2.61)$$

where the parameter a is a function of the beam emittance and plasma density. Note that the Bennett profile is far from ideal; the r.m.s. of this distribution is undefined, implying that the emittance of this distribution is very large. Indeed, this was the conclusion of experiments at the Final Focus Test Beam facility (FFTB) at SLAC which studied positron beam propagation in a plasma [40, 41]. Figure 2.8 shows the effect of the plasma on electron and positron beams. The non-linearity of the plasma focusing force on the positron beam is evident as a halo of beam particles.

Figure 2.7: A sketch comparing the trajectory of plasma electrons (red dots) under the influence of dense electron (a) and positron (b) beams. In the plasma electrons form a bubble around the electron beam with stationary ions (blue dots) inside the bubble. The explicit form of the non-linear positron wake is more complicated and depends strongly on the positron beam parameters.

2.7 Plasma wakefield acceleration in the non-linear regime

As we saw in the previous section, plasma acceleration in the linear regime is unstable, with the beam naturally evolving to produce a non-linear wake. In this section, we discuss the theory of blow-out wakes and describe their advantageous properties for electron acceleration.

The significance of the blowout regime was recognized in the late 80's and early 90's [20, 21]. The coupled field and fluid equations are intractable in the non-linear regime, so the focused shifted to a Hamiltonian description of the relativistic plasma fluid [42, 43]. The Hamiltonian of a plasma electron is given by

$$H = \gamma m c^2 + \phi \quad (2.62)$$

where ϕ is the scalar potential and the particle energy is related to the particle momentum by

$$\gamma = \sqrt{1 + \frac{p^2}{m_e^2 c^2}}. \quad (2.63)$$

The particle momentum is related to the canonical momentum $\vec{P}$ by $\vec{p} = \vec{P} - q\vec{A}/c$, where $\vec{A}$ is the vector potential. We assume that the Hamiltonian depends on z and t through the co-moving coordinate $\xi = z - ct$. By definition, we have $dH/dt = \partial H/\partial t$ and

$$\frac{\partial H}{\partial t} = -c \frac{\partial H}{\partial \xi} = c \frac{dP_z}{dt}. \quad (2.64)$$

and there is a constant of motion

$$\frac{d}{dt}(H - cP_z) = 0. \quad (2.65)$$

Prior to the arrival of the drive beam, the plasma electrons are stationary and $H - cP_z = m_e c^2$. For

Figure 2.8: Transverse beam profiles of an electron beam after free space propagation (a) and after being focused in the plasma (b). Transverse beam profiles of a positron beam after free space propagation (c) and after being focused in the plasma (d). Note the appearance of the beam halo indicating emittance growth because some particles in the positron bunch are defocused. Figures from refs. [40] and [41].

times after the arrival of the beam, we have

$$H - cP_z = \gamma m_e c^2 + q\phi - cp_z - qA_z = m_e c^2. \quad (2.66)$$

We define the pseudo-potential $\psi = \phi - A_z$. This form of the potential emerges naturally from examining the problem in the co-moving frame [43]. Then we have

$$\gamma m_e c^2 - cp_z = m_e c^2 + e\psi, \quad (2.67)$$

where we have substituted $q = -e$. From Maxwell's equations, we can show

$$\nabla_{\perp}^2 \psi = -(c\rho - J_z), \quad (2.68)$$

$$\frac{d\psi}{d\xi} = E_z, \quad (2.69)$$

$$\frac{d\psi}{dr} = -(E_r - B_{\phi}) = W_{\perp}. \quad (2.70)$$

Clearly, ψ is the quantity of interest. Note that ψ is sourced entirely by the plasma charge density and axial current. The drive beam does not contribute to ψ because $J_{z_b} = c\rho_b$, so $c\rho_b - J_{z_b} = 0$. Therefore, we will describe the wakefield entirely if we are able to determine the plasma response ρ_p and J_{z_p} due to an axial drive beam current $\lambda(\xi)$.

Figure 2.9: Trajectories of plasma electrons under the influence of a drive beam moving to the left (not shown) with (a) density $n_b/n_0 = 1$ and (b) density $n_b/n_0 = 10$. Reproduced from ref. [35].

A critical insight into determining the plasma charge and current density in the non-linear regime comes from examining the transition from linear to non-linear plasma wakes [44]. The transition can be described in terms of the trajectories of plasma electrons displaced by the fields of an intense drive beam. In the linear regime, the trajectories represent a laminar flow; the plasma electron trajectories do not cross. In the non-linear regime, the trajectories do cross, leading to the development of an electron sheath. Figure 2.9 illustrates this transition by showing example trajectories of plasma electrons with different initial radii. We see that for a drive beam with density $n_b/n_0 = 1$, the kick due to the beam on the plasma electrons is not strong enough to induce sheath crossing and the flow remains laminar. At larger beam densities, the trajectories do cross and a sheath is formed, as seen in Figure 2.9(b). All of the plasma electrons expelled from a volume near the beam axis are contained in a sheath of radius r_b and width Δ . The value of Δ is assumed *a priori* to be much less than r_b .

The region free of plasma electrons is referred to as the plasma bubble. It is filled with a uniform background of ions. Figure 2.10 shows an example of an ultrarelativistic wake ($n_b \approx 30n_0$). The plasma sheath has a maximum radius of 60 μm and the width of the sheath at this location is approximately 8 μm .

Using a force balance equation, Lu was able to solve for the trajectory of the plasma electron sheath [35]. In the ultrarelativistic blowout regime, where $n_b \gg n_0$, it is possible to write down a second-order differential equation for the sheath

$$r_b \frac{d^2 r_b}{d\xi^2} + 2 \left[\frac{dr_b}{d\xi} \right]^2 + 1 = \frac{4\lambda(\xi)}{r_b^2}, \quad (2.71)$$

which defines an elliptical trajectory. Note the use of natural units ($c = e = m_e = 1$). The axial

Figure 2.10: Left: QuickPIC simulation of a high-density electron beam with $n_b/n_0 \approx 30$, traveling to the right in plasma with density 10^{17} cm^{-3} . The black bubble region is devoid of plasma electrons. The plasma ions are not shown. Right: A transverse slice of the plasma density along $\xi = 60 \mu\text{m}$. The width of the plasma sheath is much smaller than the bubble radius.

beam current $\lambda(\xi)$ is presumed to be known, and the sheath radius is found by directly integrating eq. 2.71. In the ultrarelativistic limit, the pseudo-potential is given by

$$\psi(r, \xi) = \frac{1 + \alpha}{4} r_b(\xi)^2 - \frac{1}{4} r^2, \quad (2.72)$$

where $\alpha = \Delta/r_b \ll 1$. Using equations 2.69 and 2.70, we can solve for the wakefield E_z and $W_\perp$ inside the bubble

$$E_z(r, \xi) = \frac{1}{2} r_b(\xi) \frac{dr_b}{d\xi}, \quad (2.73)$$

$$W_\perp(r, \xi) = -\frac{r}{2}. \quad (2.74)$$

Note the use of natural units where length scales are normalized to the plasma skin depth k_p^{-1} and fields are normalized to the characteristic field strength 2.13. Equations 2.73 and 2.74 illustrate why the blowout regime is ideal for electron acceleration. First, the longitudinal electric field is independent of r ; all particles in a given longitudinal slice see the same accelerating field. Second, the focusing force is linear and independent of ξ . This means that the beam emittance can be preserved as it propagates in a plasma bubble over distances many times larger than the betatron wavelength. Figure 2.11 shows the accelerating and transverse wakefields for the plasma bubble in Figure 2.10. The accelerating fields are over a thousand times greater than the accelerating fields achieved in SLAC's RF cavities. The transverse focusing force is roughly a million times the strength of fields achieved in state-of-the-art, superconducting quadrupoles, which operate at over 100 T/m.

Figure 2.11: Left: The longitudinal field in the plasma bubble shown in Figure 2.10. The accelerating field at the back of the double exceeds 60 GV/m. Right: The transverse wakefield for the same bubble.

2.7.1 Electron beam loading in the plasma bubble

The most significant result so far from experiments at FACET is the demonstration of the acceleration of a witness bunch inside the blowout wake of a driving beam [45]. We refer to this as the “two-bunch” experiment. The basic principle of PWFA as an accelerator technology is to use a driving bunch that creates and gives up energy to the plasma wakefield, followed by a witness bunch that rides the wake and gains energy. The witness bunch has a finite length that is shorter, but not much shorter, than the length of the wakefield, which implies that the longitudinal field $E_z(\xi)$ will vary over the length of the witness bunch, causing particles at different locations in ξ to gain energy at a different rate. This potentially disastrous feature of the plasma wakefield is countered by the beam loading effect, where the presence of the witness bunch “flattens” the wake.

It is possible to choose a current profile for the witness bunch such that the longitudinal field is constant over the length of the bunch. In the linear regime, the total accelerating field felt by the witness bunch is the superposition of the wake due to the drive beam and the self wake of the witness beam [46]

$$E_z(\xi) = E_{z0} \cos(k_p \xi) - ek_p^2 \int_{\xi_0}^{\xi} \lambda(\xi') \cos[k_p(\xi - \xi')] d\xi', \quad (2.75)$$

where ξ_0 is the front of the witness bunch, measured with respect to the local maximum of the wakefield. It is straightforward to show that if the current profile has the form $\lambda(\xi) = a\xi + b$, the accelerating field will have a constant value $E_z(\xi) = E_{acc}$ over the length of the witness bunch.

Plugging the current profile into eq. 2.75 and solving for a and b gives

$$a = -\frac{E_{z0}}{e} \cos(k_p \xi_0), \quad (2.76)$$

$$b = \frac{E_{z0}}{e} \left[\frac{\sin(k_p \xi_0)}{k_p} - \xi_0 \cos(k_p \xi_0) \right]. \quad (2.77)$$

The beam profile described here is triangular (or trapezoidal if the beam truncates early), with a sharply rising current at the front of the bunch followed by a linear fall off. While this is an encouraging result, a careful consideration of energy conservation shows that there is trade-off between the accelerating gradient and the efficiency of energy extraction from the wakefield [46]. Specifically, both the accelerating gradient and energy extraction efficiency depend on the starting position of the witness bunch, with $E_{acc} \propto \cos(k_p \xi_0)$ and efficiency $\eta \propto \sin^2(k_p \xi_0)$. As a compromise, one might try setting $k_p \xi_0 = \pi/4$, which would provide 70% of the maximum accelerating gradient with 50% energy extraction from the wake.

Remarkably, the trapezoidal beam current distribution described above also produces wake flattening in the non-linear regime, despite the fact that the principle of superposition does not apply. Plugging in the same current profile $\lambda(\xi) = a\xi + b$ into the RHS of equation 2.71 and using some mathematical approximations [47], we can find the constants a and b leading to a flat accelerating field E_{acc}

$$a = -\frac{E_{acc}}{e}, \quad (2.78)$$

$$b = \frac{E_{acc}}{e} \left[\xi_0 + \sqrt{\left(\frac{eE_{acc}}{m_e c \omega_p} \right)^4 + \left(\frac{\omega_p R_b}{2c} \right)^4} \right], \quad (2.79)$$

where R_b is the maximum bubble radius and E_{acc} is a function of ξ_0 . Another remarkable feature of the blowout regime is that the extraction efficiency is independent of the accelerating gradient. The energy stored per unit length in the bubble is

$$Q_{acc} E_{acc} = \frac{m_e c^2}{4^3 r_e} (k_p R_b)^4, \quad (2.80)$$

where r_e is the classical electron radius. Note that there is trade-off between the witness charge accelerated, Q_{acc} , and the accelerating gradient, but the efficiency can be made close to 100% for any choice of E_{acc} . QuickPIC simulations have been used to demonstrate up to 90% extraction efficiency with a transformer ratio $T = E_{acc}/E_{dec} \approx 1$ [48].

In the two-bunch experiment performed at FACET, we did not have access to the trapezoidal current profiles required for perfect wake flattening. Instead, we used roughly gaussian current profiles. The effect of using a gaussian current profile is shown in Figure 2.12. This QuickPIC simulation compares the loaded (a) and unloaded (b) wakes in the blowout regime. Note the change

Figure 2.12: (a) A QuickPIC simulation of the wake due to a driving electron bunch with charge $N = 6.36 \times 10^9$ and bunch length $\sigma_z = 25 \mu\text{m}$. The beam is propagating in a plasma with density $n = 5 \times 10^{16} \text{ cm}^{-3}$. The driving electron bunch is scalloped because it is not matched to the plasma focusing force. (b) The same simulation as in (a) with a loaded witness bunch with charge $N = 4.87 \times 10^9$ and bunch length $\sigma_z = 47 \mu\text{m}$. The parameters of the drive beam, witness beam, and plasma match parameters from the FACET experiment. (c) The experimentally measured energy spectrum of the drive beam and witness beam in the absence of plasma. (d) The energy spectrum of the drive beam after deceleration in plasma. The maximum energy loss is over 3 GeV. (e) The energy spectrum of the witness beam after acceleration of plasma. The maximum energy gain is over 2 GeV. (f) A comparison of the measured energy spectrum with the results of energy spectrum simulated using QuickPIC in (b). Figure reproduced from ref. [45].

in shape of the longitudinal field (red curve) in the loaded and unloaded cases. Although the field is not perfectly flattened by the long, gaussian witness bunch, it is sufficiently flat that the witness particles all gain energy at roughly the same rate, resulting in a peaked energy spectrum around 22 GeV after 36 cm of acceleration in a Lithium plasma. The simulation results match experimental observations, where the witness beam was observed to gain 2 GeV of energy in the plasma with a modest energy spread of 2%.

During the experiment, the ratio of the charge in the drive beam to the charge in the witness beam varied between 2:1 and 1:1. When the charge in the witness beam roughly equaled the charge in the drive beam, the energy extraction efficiency was measured to be greater than 50%. The large extraction efficiency and small energy spread represented a major breakthrough for PWFA research. A subsequent experiment in a 1.3 m long Lithium oven has demonstrated acceleration of the witness

beam up to 9 GeV [49].

Before moving on, it is worth mentioning that the beam loading effect described here only modifies the longitudinal force felt by the particles in the witness bunch. The transverse force is sourced entirely from the uniform ion background, and the beam particles do not experience a self-force due to space charge because they are ultrarelativistic. We will keep this fact in mind when discussing prospects for positron acceleration in uniform plasmas in Section 2.8.

2.7.2 The QuickPIC particle-in-cell code

In the previous two sections, we made reference to the 3D particle-in-cell code QuickPIC [50, 51]. A particle-in-cell code uses a kinetic (as opposed to fluid) model of the plasma. Plasma and beam particles are represented by macro-particles, where the macro-particle charge is scaled to the number of beam or plasma particles it represents. In a typical QuickPIC simulation, both the beam and plasma macro-particles represent roughly 1,000 beam or plasma particles. The macro-particles are then deposited onto a grid. A typical particle-in-cell code proceeds in a loop.

1. Calculate the fields due to the particle distribution and beam current at the grid points.
2. Interpolate the fields onto the positions of the macro-particles.
3. Apply a Lorentz force kick to the macro-particles.
4. Propagate the macro-particles to their new locations and return to step 1.

Each iteration of the loop represents a single time-step, which is typically a fraction of the plasma response time $\tau = 1/\omega_p$. The transverse size of the grid cells are a fraction of the plasma skin depth $k_p^{-1} = c/\omega_p$.

Particle-in-cell codes are computationally intensive, and full simulations may require hundreds of thousands of CPU-hours to complete. As an example, consider the typical input parameters for a simulation of a FACET beam interacting with a meter-scale plasma. The typical plasma density is $1 \times 10^{17} \text{ cm}^{-3}$ and the associated length scale is $d = 16 \text{ } \mu\text{m}$. For accurate simulation results, the grid size is set to $d/20 = 0.8 \text{ } \mu\text{m}$. A fully 3D PIC code, such as OSIRIS [52], requires over a million time-steps to simulate the interaction. Using UCLA's Hoffman2 CPU cluster, which provides up to 512 cores to its users, this simulation takes several days.

The QuickPIC code provides a significant reduction in the amount of time needed to simulate the interaction by simulating less. QuickPIC uses the quasi-static assumption to divide the problem of simulating the beam-plasma interaction into two parts. First, there is the plasma response to the beam and plasma fields, which occurs on the time-scales discussed above. Second, there is the response of the ultrarelativistic particle beam to the plasma fields. The beam evolves at a rate $\tau_\beta = \sqrt{2\gamma}/\omega_p$. At FACET, $\gamma = 40,000$, and the beam evolution time is roughly 300 times longer

than the plasma evolution time. Therefore, the simulation time is reduced by evolving the beam particles at a much slower rate than the plasma particles.

QuickPIC deposits the beam and plasma particles onto a co-moving grid. The longitudinal size of the simulation grid is determined by the extent of the beam-plasma response in co-moving frame that provides relevant data to the user. This is typically the length of the first plasma bubble, or about $200\text{ }\mu\text{m}$ at FACET. QuickPIC calculates the plasma response in the transverse plane for each longitudinal slice of the simulation grid, propagating from front to back along the bunch. The PIC loop described above is implemented as spatiotemporal steps, where the fields, particles, and current distribution are the inputs to the next trailing slice. If the longitudinal cell size is $0.8\text{ }\mu\text{m}$, as above, and longitudinal size of the grid is $200\text{ }\mu\text{m}$, the simulation will complete 250 spatiotemporal steps for each update of the beam particles. The rate at which the beam particles are updated is typically $\tau_\beta/5$, or once every $900\text{ }\mu\text{m}$ in this example. Note that if the beam particles are unchanged, the plasma response is unchanged as well. No work needs to be done “between” beam updates. A 1 meter long beam plasma interaction will require ~ 1100 beam time-steps to run to completion, or $1100 \times 250 = 275,000$ spatiotemporal steps, which is a factor of 4-5 improvement over the fully 3D code. Simulations which used to take days can now be completed in a few hours.

2.8 Positron acceleration in the non-linear plasma wakefield

Prior to experiments at FACET, little headway has been made on the topic of positron acceleration in the non-linear regime. There are several factors affecting the state of positron PWFA research. One significant hurdle is that SLAC is the only laboratory in the world capable of producing intense positron beams for PWFA research. Prior to 2014, the only experiment to study positron acceleration in plasma was E162 at FFTB in 2003. The beam and plasma parameters for this experiment placed it in the weakly non-linear regime. The time resolved energy measurement shown in Figure 2.13 maps the longitudinal electric field onto the beam. The shape of the energy gain and loss demonstrates the linearity of the wake. While the acceleration observed in this experiment was significant, the beam emittance was diluted, as seen in Figure 2.8 [40, 41].

A more fundamental obstacle for positron PWFA research was the lack of credible theoretical or simulation work that indicated a stable regime for positron acceleration. In the conceptual design of a plasma-based linear collider, positron beams are accelerated in the wake of electron drive beams [38]. Inspecting the unloaded wakefield in Figure 2.12(a), we see that there are some problems with this concept. First, the region where the blowout wake is both accelerating and focusing for plasma electrons is at the back of bubble where the plasma electron collapse back on axis. This region is much smaller than the region which is accelerating and focusing for electrons, which is the entire back half of the bubble. Second, the plasma electron density is not uniform in this region and as a result, both the focusing force and acceleration gradient vary transversely in this region [53]. Finally,

Figure 2.13: The time resolved energy spectrum of the positron beam after propagating through a 1.4 m long plasma at density $n = 1.8 \times 10^{14} \text{ cm}^{-3}$. The blue triangles are the mean time-resolved energy slices of the beam in the absence of plasma, and the red squares are the same quantity when the beam propagates through plasma. The bunch length is $700 \mu\text{m}$. The shape of the energy gain and loss is indicative of a linear plasma wake. Figure from ref. [23].

the slope of the accelerating field is not conducive to beam loading. In order to flatten the wake, the beam must be loaded ahead of the maximum accelerating field, which would mean inside the bubble where the fields are defocusing.

All of these factors taken together indicate that positron acceleration in the non-linear wake of electron beam will be extremely challenging. However, there is an alternative approach that uses plasma to accelerate positrons for linear collider applications called the plasma “afterburner” [54]. The energy reach of a linear collider based on conventional RF acceleration can be extended by adding a plasma cell to the end of the linac and using high-energy electron or positron beams to double the energy of a trailing electron or positron bunch. While this may be a cost-effective method for increasing the energy of a linear collider, the afterburner technique as applied to positrons was mostly conjecture with little theoretical work to back it up [55]. Experiments at FACET in 2014 sought to determine if the concept was indeed viable.

2.8.1 Observation of high-gradient acceleration of positrons in a non-linear wakefield

The experiment at FACET studied the wakes produced by a high-energy, highly compressed positron bunch passing through a 1.3 meter long Lithium oven with vapor density $n = 8 \times 10^{16} \text{ cm}^{-3}$. The 20.35 GeV positron beam contained $N = 1.4 \times 10^{10}$ particles with a bunch length of $\sigma_z \sim 40 \mu\text{m}$.

Figure 2.14: The energy spectrum of the accelerated portion of a positron beamlet after propagation in 1.3 m long Lithium plasma oven. The beamlet has a spectral peak near 25 GeV, with 207 pC charge and an rms energy spread of 1.8%. Figure reproduced from ref. [56].

The peak current of the beam $I = 6.7$ kA is a record for positron beams. The beam was brought to a focus near the start of the lithium oven with a transverse beam size $\sigma_r < 100$ μm . The energy spectrum of the positron beam was viewed downstream of the lithium oven after the beam passed through a dipole magnet. The complete experimental setup is described in Chapter 6. The energy spectrum of the positron beam varied from shot to shot due to fluctuations in the beam charge and compression, but on some shots we observed a surprising result—multi-GeV energy gain with a spectral peak.

Figure 2.14 shows an example of a shot with spectrally distinct peak at 25 GeV. The bulk of the positron beam loses energy in the plasma, while a small amount of charge, or beamlet, gains a considerable amount of energy. The beamlet in Figure 2.14 contains 1.3×10^9 positrons, or roughly 10% of the total beam charge. The rms energy spread is 1.8%. By comparing the energy lost by the decelerated positrons to the energy gained by the accelerated positrons, we determined that the extraction efficiency was up to 30%. Although, we expected to observe acceleration of *some* positrons, the low-energy spread and high charge of the accelerated beamlet were unanticipated. We subsequently gained an understanding of the wakefield structure through an exhaustive simulation study [56].

Figure 2.15: (a) The unloaded wakefield due to a truncated positron bunch. The wake has a bubble structure similar to wakes produced by an electron beam driver, but in this case the bubble trails the positron beam. (b) The loaded wakefield of a positron bunch after 130 cm of propagation in a plasma. The bunch structure has evolved considerably and is in equilibrium. The longitudinal field is flat along much of the bunch. The dotted orange line shows the beam dimensions at the start of the simulation. Figure reproduced from ref. [56].

2.8.2 Beam loading in non-linear plasma wakes

The beam loading mechanism in the positron beam driven wake is significantly more complicated than that of an electron beam driven wake. In the case of the electron driver, the shape of the wake depends only on the longitudinal beam current (so long as the transverse beam size is smaller than the plasma bubble). This can be observed in Figures 2.12 (a) and (b), where the shape of the wake is minimally affected by the presence of the electron witness beam. More importantly, the transverse force felt by the electrons inside the bubble is the same in each case. For the positron driver, the shape of the wake, and therefore the fields, depend strongly on the local beam size and density. Figure 2.15 demonstrates this effect. In panel (a), a truncated positron driver provides a strong kick to the plasma electrons directed radially inward toward the beam axis. The plasma electrons cross behind the positron driver and form a plasma sheath, as in the case of an electron beam driver. In panel (b), a long positron drive beam provides the same initial kick to the plasma electrons, but some plasma electrons are sucked into the positron beam, creating a strong, non-linear focusing field on axis. This causes positrons in the tail of the bunch to be focused toward the beam axis, which further increases the amount of plasma electrons trapped on axis. The evolution occurs over 20 cm, at which point the transverse profile of the positrons has evolved to match the plasma forces. In addition, the presence of the on-axis plasma electrons flattens the longitudinal field inside the wake, so the positrons in the rear of the bunch all gain energy at the same rate.

These simulations explain the experimental results, which were hard to anticipate because we typically consider drive beams that retain their gaussian form throughout the plasma acceleration process. The evolved beam has significant structure in both the transverse and longitudinal dimensions, similar to the shape of an arrow. This effect was observed in early simulation studies of positron beam propagation in plasma, but the effect was not appreciated at the time [40].

2.8.3 Outlook for positron acceleration in the non-linear regime

The positron acceleration results from FACET are encouraging because they open the door to an entirely new regime of PWFA. There is an enormous parameter space corresponding to possible shapes for positron drive and witness bunches. A critical question is whether or not these beams can be constructed in such a way that they are in equilibrium with the plasma wake at the start of the interaction and whether or not the trailing witness bunch emittance can be preserved. A significant simulation effort is underway at UCLA to try and answer these questions.

Although the feasibility of a positron beam driven afterburner is still an open question, positron acceleration in the blowout wake of an electron beam appears to be a non-starter. This is the primary motivation for an alternative technique that uses plasma to accelerate positrons in the wake of an electron beam—the hollow channel plasma wakefield accelerator.

Chapter 3

Hollow Channel Plasma Wakefield Acceleration

3.1 Introduction

In this chapter we describe the theory of wakefield excitation in a hollow channel plasma. Why study hollow channel plasmas? As we shall see, wakefields in the hollow channel geometry have a unique advantage over those in a uniform plasma; they provide accelerating fields without transverse forces that can disrupt the beam, so long as the particle beam propagates along the channel axis.

We model the plasma channel as a non-evolving structure with a dielectric constant given by $\epsilon = 1 - \omega_p^2/\omega^2$. The geometry of the channel is depicted in Figure 3.1. Previous analyses have examined similar geometries where the plasma outside the channel extends to infinity [28, 57, 34]. Here, we describe a plasma channel with an annular shape, as is the case in our experiment.

We derive the accelerating fields associated with the azimuthally symmetric TM mode and the transverse wakefields due to the dipole HEM mode. We discuss the effect of the transverse wakefield on the beam in terms of the beam-breakup instability. We also examine the limits of the linear theory and discuss prospects for non-linear wakes in the hollow channel geometry.

3.2 Modal Description of the Plasma Channel

The plasma channel confines electromagnetic waves with frequencies below the plasma frequency ω_p . In this sense, it is analogous to a dielectric waveguide. The geometry of the plasma channel determines the phase velocity, wavelength, and transverse shape of these waves. We are only interested in the waves that co-propagate with a relativistic drive beam. Therefore, we require the phase velocity of the wave to be $v_\phi = \beta c$, where βc is the beam velocity. The waves that satisfy this condition are

Figure 3.1: a. A cutaway of the hollow channel plasma illustrating the tubular geometry. The drive beam propagates along the z -axis of the channel. b. A transverse slice of the hollow channel showing the on-axis vacuum region labeled ‘1’, the annular plasma region labeled ‘2’, and the vacuum outside the channel labeled ‘3’.

the co-propagating modes of the channel. These modes come in three flavors. First, there is the azimuthally symmetric ($m = 0$) transverse magnetic mode (TM) with $B_z = 0$ and $E_z \neq 0$. This is the most important of the modes we will consider, because it provides the longitudinal electric field that can be used to accelerate particles. Second, there is the azimuthally symmetric transverse electric mode (TE) with $E_z = 0$ and $B_z \neq 0$. This mode is not excited by a relativistic particle beam traveling through the channel and we will therefore ignore it. Finally, there are the hybrid electromagnetic modes (HEM_m) with $B_z \neq 0$ and $E_z \neq 0$, which lack azimuthal symmetry ($m \neq 0$). These modes contribute to transverse wakefields that may disrupt the beam in the channel.

3.3 The hollow channel plasma as a waveguide

The hollow channel plasma is an annular tube with infinite extent in the longitudinal z direction. Plasma with density n is present between the inner radius a and outer radius b . The annulus divides the $r - \phi$ plane into three regions, denoted 1, 2, and 3 in Figure 3.1:

1. The vacuum region inside the channel.
2. The plasma annulus extending from $a < r < b$.
3. The vacuum region outside the plasma channel.

In each of these regions, we solve the sourceless wave equations

$$\nabla^2 E_z - \frac{\varepsilon}{c^2} \partial_t^2 E_z = 0, \quad (3.1)$$

$$\nabla^2 B_z - \frac{\varepsilon}{c^2} \partial_t^2 B_z = 0. \quad (3.2)$$

It can be shown that for waves propagating solely in the z direction, solving for E_z and B_z , and applying the appropriate boundary conditions determines all components of $\vec{E}$ and $\vec{B}$ [36].

The propagating wave has the form

$$E_z(r, \phi, z, t) = E_z(r, \phi) e^{i(k_z z - \omega t)}, \quad (3.3)$$

$$B_z(r, \phi, z, t) = B_z(r, \phi) e^{i(k_z z - \omega t)}. \quad (3.4)$$

$E_z(r, \phi)$ and $B_z(r, \phi)$ are complex functions with a relative phase of $\pi/2$. We associate this phase with the B field by convention. Plugging these expressions into eq. 3.3 and 3.4 gives

$$\left(\nabla_{\perp}^2 - k_z^2 + \frac{\varepsilon \omega^2}{c^2} \right) E_z(r, \phi) = 0, \quad (3.5)$$

$$\left(\nabla_{\perp}^2 - k_z^2 + \frac{\varepsilon \omega^2}{c^2} \right) B_z(r, \phi) = 0. \quad (3.6)$$

We consider only the modes that propagate at the same speed as the drive beam

$$v_{\phi} = \beta c = \frac{\omega}{k_z}. \quad (3.7)$$

Substituting $\omega = \beta c k_z$ into eq. 3.5 and 3.6 gives

$$(\nabla_{\perp}^2 + k_{\perp}^2) E_z(r, \phi) = 0, \quad (3.8)$$

$$(\nabla_{\perp}^2 + k_{\perp}^2) B_z(r, \phi) = 0, \quad (3.9)$$

with

$$k_{\perp}^2 = k_z^2 (\varepsilon \beta^2 - 1). \quad (3.10)$$

Once we have solved for E_z and B_z , we use the following expressions to find the transverse components of the fields

$$\vec{E}_{\perp} = \frac{i}{k_{\perp}^2} \left[k_z \nabla_{\perp} E_z - \frac{\omega}{c} \hat{z} \times \nabla_{\perp} B_z \right], \quad (3.11)$$

$$\vec{B}_{\perp} = \frac{i}{k_{\perp}^2} \left[k_z \nabla_{\perp} B_z + \mu \varepsilon \frac{\omega}{c} \hat{z} \times \nabla_{\perp} E_z \right], \quad (3.12)$$

as well as the transverse force on the particles

$$F_r = e(E_r + \beta \hat{z} \times B_\phi). \quad (3.13)$$

Note that $\mu = 1$ in vacuum and plasma.

Equations 3.8 and 3.9 are transverse Helmholtz equations. In cylindrical coordinates, the fields $E_z(r, \phi)$ and $B_z(r, \phi)$ can be decomposed into azimuthal Fourier components

$$E_z(r, \phi) = \sum_{m=0}^{\infty} E_m(r) e^{im\phi}, \quad (3.14)$$

$$B_z(r, \phi) = \sum_{m=0}^{\infty} B_m(r) e^{im\phi}, \quad (3.15)$$

with each component E_m, B_m solving the Helmholtz equation.

3.3.1 The TM accelerating mode

The TM mode has no azimuthal dependence ($m = 0$) and the longitudinal field has the form $E_z(r, z, t) = E_{z0}(r) e^{i(k_z z - \omega t)}$. We will solve equation 3.8 in each of the three regions described above and apply the appropriate boundary conditions to produce a complete description of the fields in the TM mode. This treatment is analogous to the methods used to determine the fields in dielectric structures [58]*.

The dielectric constant varies as we move from vacuum to plasma and back to vacuum

$$\varepsilon = \begin{cases} 1 & r < a, \\ 1 - \frac{\omega_p^2}{\omega^2} & a < r < b, \\ 1 & r > b, \end{cases} \quad (3.16)$$

where ω_p is the plasma frequency. Plugging in the expression for the dielectric into equation 3.10, and using the substitution $\omega = \beta c k_z$ produces the following expressions for the transverse wavenumber

$$k_\perp^2 = \begin{cases} -\frac{k_z^2}{\gamma^2} & r < a, \\ -\frac{k_z^2}{\gamma^2} - k_p^2 & a < r < b, \\ -\frac{k_z^2}{\gamma^2} & r > b, \end{cases} \quad (3.17)$$

*It is worth pointing out one significant difference between the modes in a hollow channel plasma and in a dielectric structure. The dielectric structure permits radial solutions proportional to oscillating Bessel functions J and Y . As a result, for the $m = 0$ case, there are multiple modes which satisfy the boundary conditions and they are enumerated as TM_{0n} . In the case of the plasma channel, there is only one mode that satisfies the boundary conditions and we refer to it simply as TM.

where k_p is the plasma wavenumber. Note that $k_\perp^2$ is negative in each of the three regions. Rather than work with imaginary values for the transverse wave number we express equation 3.8 as

$$(\nabla_\perp^2 - |k_\perp|^2) E_z(r) = 0. \quad (3.18)$$

As we saw in Section 2.4.1, the transverse Helmholtz equation with negative transverse wavenumber has a modified Bessel solution. Since, we are solving this equation in three regions, we have the solutions

$$E_z(r) = A_i I_0(|k_{\perp i}|r) + B_i K_0(|k_{\perp i}|r) \quad (3.19)$$

where $i = 1, 2, 3$ denotes the radial region. The transverse wavenumber in each region is $|k_{\perp 1,3}| = k_z/\gamma$ and $|k_{\perp 2}| = \sqrt{k_z^2/\gamma^2 + k_p^2}$. Using equation 3.11 and 3.12 we can immediately deduce the transverse field components

$$E_r(r) = \frac{ik_z |k_{\perp i}|}{k_{\perp i}^2} [A_i I_1(|k_{\perp i}|r) - B_i K_1(|k_{\perp i}|r)], \quad (3.20)$$

$$B_\phi(r) = \frac{i\beta\epsilon k_z |k_{\perp i}|}{k_{\perp i}^2} [A_i I_1(|k_{\perp i}|r) - B_i K_1(|k_{\perp i}|r)]. \quad (3.21)$$

We now need to apply the boundary conditions to eliminate five of the six constants $A_{1,2,3}$ and $B_{1,2,3}$ in favor of the on-axis field amplitude $A_1 = E_{z0}$. The following boundary conditions are necessary and sufficient:

1. The field is finite at $r = 0$. Therefore $B_1 = 0$.
2. The E_z field is continuous at the vacuum-plasma boundary at $r = a$.
3. The E_z field is continuous at the plasma-vacuum boundary at $r = b$.
4. The radial displacement $D_r = \epsilon E_r$ is continuous at $r = a$.
5. The radial displacement $D_r = \epsilon E_r$ is continuous at $r = b$.
6. The field is zero at $r = \infty$. Therefore $A_3 = 0$.

Boundary conditions 2, 3, 4, and 5 are expressed as a linear system of equations

$$E_{z0} I_0(|k_{\perp 1}|a) = A_2 I_0(|k_{\perp 2}|a) + B_2 K_0(|k_{\perp 2}|a) \quad (3.22)$$

$$A_2 I_0(|k_{\perp 2}|b) + B_2 K_0(|k_{\perp 2}|b) = B_3 K_0(|k_{\perp 3}|b) \quad (3.23)$$

$$\frac{ik_z |k_{\perp 1}|}{k_{\perp 1}^2} E_{z0} I_1(|k_{\perp 1}|a) = \frac{i\epsilon k_z |k_{\perp 2}|}{k_{\perp 2}^2} [A_2 I_1(|k_{\perp 2}|a) - B_2 K_1(|k_{\perp 2}|a)] \quad (3.24)$$

$$\frac{i\epsilon k_z |k_{\perp 2}|}{k_{\perp 2}^2} [A_2 I_1(|k_{\perp 2}|b) - B_2 K_1(|k_{\perp 2}|b)] = -\frac{ik_z |k_{\perp 3}|}{k_{\perp 3}^2} B_3 K_1(|k_{\perp 3}|b) \quad (3.25)$$

Figure 3.2: The normalized longitudinal electric field as a function of r for three different values of the beam energy γ . $\gamma \approx 40,000$ corresponds to the beam energy at FACET. The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$.

which is more conveniently solved as a matrix equation

$$0 = \begin{pmatrix} I_0(k_{\perp 1}a) & -I_0(k_{\perp 2}a) & -K_0(k_{\perp 2}a) & 0 \\ 0 & I_0(k_{\perp 2}b) & K_0(k_{\perp 2}b) & -K_0(k_{\perp 3}b) \\ \frac{-1}{|k_{\perp 1}|}I_1(k_{\perp 1}a) & \frac{\varepsilon}{|k_{\perp 2}|}I_1(k_{\perp 2}a) & \frac{-\varepsilon}{|k_{\perp 2}|}K_1(k_{\perp 2}a) & 0 \\ 0 & \frac{-\varepsilon}{|k_{\perp 2}|}I_1(k_{\perp 2}b) & \frac{\varepsilon}{|k_{\perp 2}|}K_1(k_{\perp 2}b) & \frac{-1}{|k_{\perp 3}|}K_1(k_{\perp 3}b) \end{pmatrix} \begin{pmatrix} E_{z0} \\ A_2 \\ B_2 \\ B_3 \end{pmatrix}. \quad (3.26)$$

This set of equations has a solution if the determinant of the matrix is equal to zero. Furthermore, the equation $\det(M) = 0$, where M is the matrix in equation 3.26, constrains the wavenumber k_z .

We can solve equation 3.26 for the constants A_2 , B_2 , B_3 , and k_z . However, it is neither illuminating or convenient to write down the algebraic solutions. Instead, we plot numerical solutions for $E_z(r)$ and $F_r(r)$ in Figures 3.2 and 3.3, respectively. For small values of γ , there is large variation of the E_z field with r inside the channel. This is not ideal for particle acceleration, because the bunch has a non-zero radial size, and particles further from the axis will experience a different accelerating or decelerating field. As γ increases to a moderately relativistic value of 5, the radial dependence is suppressed. At $\gamma = 40,000$, the beam energy at FACET, the variation of the field in the channel is

Figure 3.3: The normalized radial force as a function of r for three different values of the beam energy γ . $\gamma \approx 40,000$ corresponds to the beam energy at FACET. The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$.

$$\Delta E_z/E_0 < 1.5 \times 10^{-9}.$$

The transverse force shows a similar improvement with increasing γ . At small γ , the field is slightly defocusing inside the channel, while at $\gamma = 40,000$ the force is negligible for all values of $r < a$. The transverse force increases rapidly and becomes focusing at the edge of the channel. At first glance, it appears that the plasma boundary acts as a wall, bouncing wayward particles back toward the center of the channel. Indeed, this was the conclusion of some early research on the topic of the hollow channel plasma accelerator [59]. However, this analysis only considered the forces associated with the TM mode, meaning that the beam centroid must remain centered on-axis as the particles diverge and reach the walls. A more detailed analysis reveals that the forces associated with the higher-order HEM mode lead to beam defocusing and transverse instabilities. We will address this point in Section 3.7.

We return to our analysis of the TM mode in the ultra-relativistic regime. This is the most useful regime for particle acceleration, as demonstrated in Figures 3.2 and 3.3, and this is the regime in which we operate at FACET. Taking the limit $\gamma \rightarrow \infty$, $\beta \rightarrow 1$, we have $ck_z = \omega$, $k_{\perp 1} = k_{\perp 3} = 0$,

and $k_{\perp 2} = k_p$. Solving for ε in $\det(M) = 0$ gives

$$\varepsilon = \frac{k_p a}{2} \frac{B_{00}(a, b)}{B_{10}(a, b)}, \quad (3.27)$$

and

$$k_z = k_p \left[\frac{2B_{10}(a, b)}{2B_{10}(a, b) - k_p a B_{00}(a, b)} \right]^{1/2}, \quad (3.28)$$

where we have introduced “Bessel-boundary function” notation to keep track of the copious I s and K s that arise

$$B_{ij}(x, y) = I_i(k_p x) K_j(k_p y) + (-1)^{i-j+1} I_j(k_p y) K_i(k_p x). \quad (3.29)$$

We note that our expression 3.28 is identical in structure to Schroeder’s expression (2) in [57] for the longitudinal wavenumber of zeroth-order mode for a hollow channel surrounded by plasma extending to infinity[†]

$$k_0 = k_p \left[\frac{2K_1(k_p a)}{2K_1(k_p a) + k_p a K_0(k_p a)} \right]^{1/2}. \quad (3.30)$$

In Appendix A.7 we show explicitly that our solution reduces to equation 3.30 in the limit $b \rightarrow \infty$.

Continuing with our analysis, the limiting expressions for A_2 , B_2 , and B_3 provide a complete description of the TM mode in the ultra-relativistic regime

$$E_z(r) = \begin{cases} E_{z0} & r < a, \\ E_{z0} \frac{B_{00}(r, b)}{B_{00}(a, b)} & a < r < b, \\ 0 & b < r, \end{cases} \quad (3.31)$$

$$E_r(r) = \begin{cases} -ik_z E_{z0} \frac{r}{2} & r < a, \\ -ik_z E_{z0} \frac{1}{k_p} \frac{B_{10}(r, b)}{B_{00}(a, b)} & a < r < b, \\ -ik_z E_{z0} \frac{ab}{2r} \frac{B_{10}(b, b)}{B_{10}(a, b)} & b < r, \end{cases} \quad (3.32)$$

$$B_\phi(r) = \begin{cases} -ik_z E_{z0} \frac{r}{2} & r < a, \\ -ik_z E_{z0} \frac{a}{2} \frac{B_{10}(r, b)}{B_{10}(a, b)} & a < r < b, \\ -ik_z E_{z0} \frac{ab}{2r} \frac{B_{10}(b, b)}{B_{10}(a, b)} & b < r. \end{cases} \quad (3.33)$$

Note that $B_\phi = \varepsilon E_r$ in each of the three regions, so the transverse wake $W_\perp = E_r - \beta B_\phi$ with $\beta = 1$

[†]We have taken some liberties in reproducing this expression to emphasize the similarities between the two equations.

is

$$W_{\perp}(r) = \begin{cases} 0 & r < a, \\ (1 - \varepsilon)E_r & a < r < b, \\ 0 & b < r. \end{cases} \quad (3.34)$$

Therefore, an ultrarelativistic beam will not experience a transverse wakefield resulting from the TM mode while propagating in the channel.

3.3.2 Single Particle Wake Excitation

So far, we have been discussing the shape of the TM mode in a hollow channel under the assumption that it has been oscillating for all time. In practice, this mode is excited by a relativistic beam of finite duration. In order to find the response of the channel to the beam, we derive a Green's function for a single particle with charge q traveling at speed c down the center of the channel. We provide a detailed derivation in Appendix A and reproduce the result here

$$E_z^G(\xi) = -E_{z0} \cos(\chi k_p \xi) \Theta(\xi), \quad (3.35)$$

where E_{z0} is the single particle wake amplitude

$$E_{z0} = \frac{4qk_p^2 B_{00}(a, b)}{k_p a [2B_{10}(a, b) - k_p a B_{00}(a, b)]} = \frac{4qk_z^2 \varepsilon}{(k_p a)^2} \quad (3.36)$$

and $k_z = \chi k_p$ is the wavenumber of the mode as described in equation 3.28.

We find the total wake response by convolving the Green's function with the beam current

$$W_z(\xi) = \int_{\xi}^{\infty} I(\xi') E_z^G(\xi - \xi') d\xi', \quad (3.37)$$

where $I(\xi)$ is the beam current. As an example, consider the wakefield produced by a gaussian beam current long after the beam has passed. We can extend the bottom limit of the integral to infinity to find the maximum wake amplitude

$$\max(W_z) = \int_{-\infty}^{\infty} \frac{N}{\sqrt{2\pi}\sigma_z} E^{-\frac{\xi'^2}{2\sigma_z^2}} E_{z0} \cos[\chi k_p(\xi - \xi')] d\xi' = N E_{z0} e^{-\frac{1}{2}k_z^2 \sigma_z^2}. \quad (3.38)$$

Using the parameters from the simulation with a positron beam driver shown in Figure 3.4, we find that the maximum accelerating field behind the bunch is 305 MV/m. Although a gradient less than 1 GV/m is considered small for PWFA experiments, we note that the gradient can be made larger by moving towards higher beam currents, higher plasma densities, and a smaller inner channel radius.

Figure 3.4: Top Left: QuickPIC simulation of an electron beam with $N = 2 \times 10^9$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. The electron beam density has been enhanced by a factor of five for visibility. Bottom Left: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), showing excellent agreement. The beam current from QuickPIC (dashed blue) and theoretical calculation (dashed red) are also shown and the peak beam current is 1 kA. Top Right: QuickPIC simulation of a positron beam propagating to the right with the same beam and plasma parameters as the previous case. Bottom Right: Comparison of the on-axis longitudinal field from QuickPIC and theory, also showing excellent agreement.

3.4 Validity of the Linear Approximation

In this section, we use simulation results from QuickPIC to demonstrate the accuracy of our theoretical model in the linear regime and explore the transition to the nonlinear regime. Figure 3.4 shows QuickPIC simulations of identical electron and positron beams driving a wake in an annular plasma channel. The beam parameters are $N = 2 \times 10^9$, $\sigma_z = 40 \mu\text{m}$, and $\sigma_x = \sigma_y = 20 \mu\text{m}$, and the plasma parameters are $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. The width of the annulus is two plasma skin depths. The lower plots compare the on-axis longitudinal field output from QuickPIC with our calculation of W_z . There is excellent agreement between the theory and the simulation.

When should we expect the linear approximation to provide an accurate description of the hollow channel wake? The assumption used in the linear theory is that the walls of the channel are a fixed boundary. Of course, plasma electrons are free to move. The inner boundary layer of the channel

will respond to the beam fields by being pushed into the plasma annulus if the initial kick is due to a driving electron beam, or by being pulled into the vacuum channel if the initial kick is due to a driving positron beam. The weak beam driver used in the simulations shown in Figure 3.4 does not cause a significant perturbation to the boundary layer over the length of the simulation and therefore the theory accurately describes the wake. However, a more intense beam driver will perturb the boundary. If the displacement of the boundary plasma electrons is on the order of a plasma skin depth k_p^{-1} , we expect that the linear theory will no longer hold.

Consider the charge per unit length in the boundary layer of the plasma annulus

$$\frac{Q_p}{L} = n_0\pi(a + \Delta R)^2 - n_0\pi a^2 \approx n_0 2\pi a \Delta R, \quad (3.39)$$

where we have assumed that the width of the boundary layer ΔR is small compared to the radius of the inner boundary. The width of ΔR depends on the number of plasma electrons that are required to shield the beam charge per unit length N/σ_z

$$\Delta R = \frac{N}{2\pi a \sigma_z n_0}. \quad (3.40)$$

For the simulation parameters in Figure 3.4, we calculate $\Delta R = 0.8 \mu\text{m}$, which is less than the transverse cell size of $1.4 \mu\text{m}$ used in simulations. The boundary layer remains fixed at a constant radius throughout the simulation.

For cases of interest, the bunch length is on the order of a plasma skin depth. In the simulations shown in this section, we have set $k_p \sigma_z = 2$. We can solve for the total charge in the drive beam satisfying the inequality $\Delta R < k_p^{-1}$. Under these conditions

$$N < \frac{a}{r_0}, \quad (3.41)$$

where r_0 is the classical electron radius. Plugging in $a = 130 \mu\text{m}$, we find $N < 4 \times 10^{10}$ particles per bunch, which is an incredibly large bunch charge. This scaling is favorable for wakefield generation in the linear regime, because the single particle wakefunction 3.36 is proportional to a^{-2} for small a . If we reduce the beam charge according to 3.41 while decreasing the inner radius of the channel, we expect the accelerating gradient to increase as a^{-1} .

There are two caveats to this argument that need to be addressed. First, the shielding scenario applies when the plasma electrons are pushed into the plasma annulus by the kick of an electron beam driver. In this case, the plasma boundary changes by ΔR . But if, instead, the plasma electrons receive a radially inward kick from a positron drive beam, they will be pulled into the channel vacuum and the radius of the boundary layer will change by a large amount. Figure 3.5 shows the wakes due to identical electron and positron beam drivers with $N = 1 \times 10^{10}$ particles per bunch, five times the bunch charge of the previous simulation. All other parameters are held fixed. As we expect from

Figure 3.5: Top Left: QuickPIC simulation of an electron beam with $N = 1 \times 10^{10}$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom Left: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), showing excellent agreement. The peak beam current is 5 kA. Top Right: QuickPIC simulation of a positron beam propagating to the right with the same beam and plasma parameters as the previous case. Bottom Right: Comparing the on-axis longitudinal fields from QuickPIC and theory, we observe a large discrepancy.

the shielding argument, the wakefield driven by the electron beam driver is well modeled by linear theory. For the positron drive beam, there is good agreement between the simulation and model over the first half-wavelength of the wakefield, but discrepancy arises at the tail end of the beam, roughly $50 \mu\text{m}$ behind the beam centroid, where the plasma electrons are observed to stream into the channel.

The second caveat relates to the longitudinal region where the model is valid. Figure 3.6 shows a QuickPIC simulation of an electron beam driver with $N = 2 \times 10^{10}$ particles per bunch, with all other parameters identical to the previous simulations. The shielding argument implies that the linear approximation should still provide an accurate description of the wakefield. We see that this is indeed the case over the first wavelength, but at a point roughly $200 \mu\text{m}$ behind the beam centroid, the plasma electrons leave the boundary and stream into the vacuum channel and the model breaks down. This comes as no surprise since the plasma electrons receive the same kick as with a positron beam driver, but 180° out of phase. The initial kick is stored as electrostatic potential energy over the first half of the wavelength, and converted into kinetic energy in the second

Figure 3.6: Top: QuickPIC simulation of an electron beam with $N = 2 \times 10^{10}$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), showing agreement through the first oscillation. The discrepancy between theory and simulation emerges after the plasma electrons undergo a full oscillation.

half of the wavelength. We see that our model of wakefields driven by electron and positron beams are essentially the same, but the description of the electron beam-driven wakefield is valid for a full wavelength of the wake.

3.4.1 The strongly nonlinear regime

As a final example, we examine the wake due to an intense positron drive beam with $N = 2 \times 10^{10}$ particles per bunch, ten times the charge used in the first example. Figure 3.7 shows the results of the QuickPIC simulation. Our model diverges from the simulation results after less than half of a wake period. The plasma electrons cross the axis of the channel near $z = 150 \mu\text{m}$ behind the bunch centroid, causing a sharp spike in the longitudinal field. The field is sharply decelerating, and then accelerating for a positron witness beam placed behind the spike, but the presence of on-axis plasma electrons produces a nonlinear focusing force.

Several authors have claimed that the nonlinear hollow channel wakes are suitable for accelerating

Figure 3.7: Top: QuickPIC simulation of an positron beam with $N = 2 \times 10^{10}$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$ with inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red), along with the beam current which reaches a maximum value of 10 kA at $z = 0 \mu\text{m}$. The model and simulation diverge after less than half of a single wake period. The field spike is due to a large concentration of on-axis plasma electrons.

positron beams [60, 54, 61, 29]. However, in our simulations of nonlinear wakefields in the hollow channel, we find no evidence to support these claims. A strong accelerating field for positrons always coincides with on-axis plasma electrons, leading to a nonlinear focusing force on-axis. Furthermore, the slope of the accelerating field is negative in the accelerating region, which is unfavorable for beam loading compensation.

We cannot make a decisive claim that the nonlinear hollow channel wake is unsuitable for positron acceleration. Further work is needed, and perhaps different channel geometries may reveal a solution which provides large accelerating gradients without transverse focusing forces. One critical step in this direction would be the extension of Lu's sheath theory [44] to include the complicated sheath patterns encountered in nonlinear hollow channel wakes. The challenge in this case is to make an accurate *a priori* assumption about the shape of sheath. Initial work on this topic [62] has been used to describe the first half-period of the wake, which is useful for electron acceleration, but the assumptions used in this theory are not extended to the second half of the wake where there is an accelerating field for positrons.

Figure 3.8: Top: QuickPIC simulation of an electron beam drive bunch with $N = 6 \times 10^9$, $\sigma_z = 40 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ and a positron witness bunch with $N = 1 \times 10^9$, $\sigma_z = 13 \mu\text{m}$, $\sigma_x = \sigma_y = 20 \mu\text{m}$ propagating to the right in an annular plasma channel with density $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$, inner radius $a = 130 \mu\text{m}$ and outer radius $b = 170 \mu\text{m}$. Bottom: The on-axis longitudinal field from QuickPIC (solid blue) and theoretical calculation (solid red). The peak current of the drive beam is 3 kA and the peak current of the witness beam is 1.5 kA. The accelerating gradient at the location of the positron witness is 475 MV/m, although the model predicts 800 MV/m. The discrepancy is due to plasma electrons streaming into the vacuum channel at the location of the positron witness.

3.5 Beam Loading in the Hollow Channel Wakefield

The principle motivation for the development of HC-PWFA is to demonstrate a technique for accelerating positrons in the wake of an electron beam driver. Figure 3.8 shows that this is indeed possible. Note that the simulated accelerating gradient at the location of the loaded positron bunch is less than the value expected from linear theory. As discussed in Section 3.4, the linear theory becomes inaccurate one full wavelength behind the electron beam driver, which is where the wake is accelerating for positrons. At this point in the wake, plasma electrons stream into the vacuum channel and the field is reduced. The presence of the positron witness bunch amplifies this effect by providing a radially inward kick to the streaming plasma electrons. This effect can be mitigated by moving to larger channel diameters, but in that case the gradient is also reduced.

3.6 The HEM Modes

High order modes are excited by asymmetric or off-axis beams propagating in the plasma channel. With the exception of the radially symmetric case $m = 0$, all modes require both electric and magnetic longitudinal field components to satisfy the boundary conditions. These hybrid modes give rise to transverse wakefields that have deleterious effects on the beam emittance. After deriving the general form of the HEM modes in the hollow channel plasma, we will discuss in particular the $m = 1$, HEM₁ dipole mode which is the source of dipole beam-breakup instability (BBU).

3.6.1 General Solutions for HEM modes

We solve for the fields associated with the HEM modes using the same approach that we detailed in Section 3.3 for finite γ , now with $m > 0$. Equations 3.8 and 3.9 have solutions

$$E_z(r) = A_i I_m(|k_{\perp i}|r) + B_i K_m(|k_{\perp i}|r), \quad (3.42)$$

$$B_z(r) = C_i I_m(|k_{\perp i}|r) + D_i K_m(|k_{\perp i}|r), \quad (3.43)$$

with $k_{\perp i}$ given by equation 3.17 and $i = 1, 2, 3$ as shown in Figure 3.1. The transverse field components are derived from E_z and B_z using equations 3.11 and 3.12. Again, we require the field to be finite on axis and at infinity. Therefore, the fields are proportional to I_m inside the channel wall and K_m outside the channel wall, and superposition of the two in the plasma region. We apply boundary conditions on the fields to eliminate seven of the eight constants $A_2, B_{2,3}, C_{1,2}, D_{2,3}$ in favor of A_1 , the amplitude of the E_z field inside the channel. Note that the B field is continuous across the plasma boundary because there are no surface currents. Furthermore, we rescale A_1 such that

$$A_1 = \frac{E_m}{I_m(|k_{\perp 1}|a)}. \quad (3.44)$$

This is a general but useful normalization of the field inside the channel because

$$\lim_{\gamma \rightarrow \infty} E_z(r) = E_m \frac{r^m}{a^m}, \quad (3.45)$$

which allows us to directly compare wake amplitudes at $r = a$.

3.6.2 The Dipole Mode

The most significant of the higher order modes is the HEM₁ dipole mode. The transverse wakefield associated with this mode drives the beam-breakup instability which tends to drive it towards the wall of the plasma channel.

Figures 3.9 and 3.10 show the numerically calculated fields for three different values of γ . The field converges rapidly to the ultrarelativistic limit and we only plot the fields up to $\gamma = 10$ rather

Figure 3.9: The normalized longitudinal electric field as a function of r for three different values of the beam energy γ . The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$.

than $\gamma = 40,000$ corresponding to the beam energy at FACET.

In Appendix A.6, we calculate all components of the $m = 1$ fields starting from Maxwell's equations in the limit $v = c$. The salient features of the fields are

- $E_z \propto B_z \propto r$ inside the channel,
- $E_z = B_z = 0$ outside the channel,
- The transverse fields are proportional to r^2 inside the channel,
- The transverse fields are proportional to r^{-2} plus a constant outside the channel,
- The transverse wakefield $W_\perp = E_r - B_\phi$ is constant inside the channel,
- The transverse wakefield is zero outside the channel.

The fact that there is a constant field extending to infinity outside the channel seems surprising at first glance. In the $v < c$ case, the transverse fields outside the plasma are proportional to $K_2(kr/\gamma)$. For the parameters used in Figures 3.9 and 3.10, we find that the decay length $r = \gamma/k$ of the field outside the channel is 1.3 meters; an extremely large distance. The strength of this field outside

Figure 3.10: The normalized radial force as a function of r for three different values of the beam energy γ . The inner radius and outer radius of the plasma channel is denoted by black dashed lines at $a = 100 \mu\text{m}$ and $b = 150 \mu\text{m}$ respectively. The plasma density is $1 \times 10^{17} \text{ cm}^{-3}$.

the channel is linearly proportional to the offset of the beam in the channel and may prove to be a useful beam alignment diagnostic. We aim to explore this idea further in future work.

Before moving on to a discussion of the beam breakup instability, we reproduce here the Green's function for the transverse wakefield derived in equation A.138

$$W_{\perp}^G(\xi) = -\frac{8qr_0k_p^3}{(k_p a)^3} \frac{B_{11}(a, b)}{\chi [4B_{21}(a, b) - k_p a B_{11}(a, b)]} \sin(\chi k_p \xi) \Theta(\xi). \quad (3.46)$$

The strength of the wake scales linearly with the beam offset r_0 and has a strong dependence on the inner channel radius. In order to mitigate the effect of the transverse wakefield, the beam should be kept as close to the axis as possible and we should choose the largest inner radius a of the channel which still provides an acceptable accelerating gradient.

3.7 The Beam Breakup Instability

In this section, we follow Chao's derivation of the beam breakup instability [63] and use it to estimate the growth length of the instability. The beam breakup instability, acting on a bunch that is much

shorter than the wavelength of the transverse wakefield, can be conceptualized as follows:

- The head of an off-axis beam excites the $m = 1$ dipole mode with a strength proportional to r_0 , the offset of the bunch.
- The dipole mode creates a transverse wakefield that is constant in r and points in the same direction as the beam offset, away from the axis.
- Particles trailing the head of the beam experience a force from the transverse wake that grows in strength along the bunch.
- This force pushes the trailing particles away from the beam axis, exciting the mode even further.
- Particles toward the tail of the bunch experience a very large transverse force and are accelerated rapidly towards the channel wall.

The result of this process is a “banana-ing” of the bunch, which degrades the beam emittance. The BBU is present in all particle accelerators, not just hollow channel plasmas. We will diagnose the severity of the BBU in the hollow channel plasma by comparing the growth length of the instability to the length of the plasma channel. Roughly speaking, if the growth length is longer than the accelerating section, a bunch can be accelerated without disrupting the beam emittance. If the growth length of the instability is shorter than the accelerating section, the beam emittance will be diluted or the beam will reach the channel wall and be destroyed.

3.7.1 Derivation of BBU expression

We calculate the evolution of the beam shape as a function of linac coordinate s and beam coordinate ξ . We assume that the beam starts with an initial vertical offset y_0 in the channel traveling parallel to the channel axis. We use a perturbative expansion to express the transverse beam coordinate

$$y(s, \xi) = \sum_{n=0}^{\infty} y_n(s, \xi). \quad (3.47)$$

The first term in this series is the constant y_0 . The higher order terms in this series are determined from the lower order terms according to Chao’s equation 3.62 [63]

$$\frac{\partial^2 y_{n+1}}{\partial s^2} + k_\beta^2 y_{n+1} = -\frac{Nr_e}{\gamma} \int_{\xi}^{\infty} d\xi' \rho(\xi') \hat{W}_\perp(\xi - \xi') y_n(s, \xi'), \quad (3.48)$$

where k_β is the focusing experienced by the beam, γ is the beam energy, r_e is the classical radius of the electron, and $\rho(\xi)$ is the longitudinal charge distribution, and $\hat{W}_\perp = W_\perp^G/(qr_0)$. Equation 3.48 is an expression of the force acting on the y_{n+1} distribution due to an external focusing force k_β

and the transverse wake due to the term y_n . The perturbative approach works well for most beam distributions where the first one or two terms account for most of the charge in the beam. We assume that the beam energy γ is constant over the length of the hollow channel plasma (a good assumption for the hollow channel experiments at FACET where the actual change in beam energy is about 0.1%) and that the beam experiences no external focusing, so $k_\beta = 0$.

We perform this calculation for a flat-top bunch of length σ_z that is short compared to the wavelength of the wakefield. This makes the calculation more tractable because we can assume that the wakefield varies linearly over the bunch:

$$\rho(\xi) = \begin{cases} \frac{1}{\sigma_z} & |\xi| < \sigma_z/2 \\ 0 & |\xi| > \sigma_z/2 \end{cases}, \quad (3.49)$$

$$\hat{W}_\perp(\xi) \approx \begin{cases} \hat{W}_{\perp 0} \xi & 0 > \xi > -\sigma_z \\ 0 & \xi > 0 \text{ or } \xi < -\sigma_z \end{cases}, \quad (3.50)$$

where we have expanded $\hat{W}_\perp(\xi)$ at short distances to find

$$\hat{W}_{\perp 0} = \frac{8k_p^4}{(k_p a)^3} \frac{B_{11}(a, b)}{[4B_{21}(a, b) - k_p a B_{11}(a, b)]} \xi. \quad (3.51)$$

Plugging our expressions for ρ , $\hat{W}_\perp(\xi)$, and y_0 into equation 3.48 we have:

$$\frac{\partial^2 y_1}{\partial s^2} = -\frac{Nr_e}{\gamma} \int_{\xi}^{\sigma_z/2} d\xi' \hat{W}_{\perp 0} \frac{(\xi - \xi')}{\sigma_z} y_0. \quad (3.52)$$

Before evaluating this integral, we make the following variable substitution:

$$u(\xi') = \frac{\xi'}{\sigma_z} - \frac{1}{2}, \quad (3.53)$$

so that $\xi' = \sigma_z u + \sigma_z/2$ and $d\xi' = \sigma_z du$. Plugging into the integral we have:

$$\frac{\partial^2 y_1}{\partial s^2} = -\frac{Nr_e \sigma_z \hat{W}_{\perp 0}}{\gamma} y_0 \int_{\frac{\xi}{\sigma_z} - \frac{1}{2}}^0 du \left[\left(\frac{\xi}{\sigma_z} - \frac{1}{2} \right) - u \right]. \quad (3.54)$$

Note that the quantity in brackets comes from rearranging the term $\rho(\xi') \hat{W}_\perp(\xi - \xi')$ in the original integral 3.48. In other words, this part of the integral is independent of the function $y_n(s, z')$ that appears in each order of the series. Next, we define $b = \frac{\xi}{\sigma_z} - \frac{1}{2}$ and $A = Nr_0 \sigma_z \hat{W}_{\perp 0} / \gamma$. Carrying out the integral gives

$$\frac{\partial^2 y_1}{\partial s^2} = A y_0 \int_0^b du [b - u] = A y_0 \frac{b^2}{2}. \quad (3.55)$$

We integrate equation 3.55 twice with respect to s to find

$$y_1(s, \xi) = Ay_0 \frac{b^2}{2 \cdot 1} \frac{s^2}{2 \cdot 1} = Ay_0 \frac{b^2}{2 \cdot 1} \frac{s^2}{2 \cdot 1}. \quad (3.56)$$

We now have our solution for $y_1(s, \xi)$, which is source for the next term in the series. Repeating this process for the y_2 term gives

$$\frac{\partial^2 y_2}{\partial s^2} = A \int_0^b du(b-u)y_1 = A^2 y_0 \frac{s^2}{2 \cdot 1} \int_0^b du(b-u) \frac{u^2}{2 \cdot 1} = A^2 y_0 \frac{s^2}{2 \cdot 1} \frac{b^4}{4 \cdot 3 \cdot 2 \cdot 1}, \quad (3.57)$$

and integrating twice with respect to s we get

$$y_2(s, \xi) = A^2 y_0 \frac{s^4}{4!} \frac{b^4}{4!}. \quad (3.58)$$

We are able to identify the following pattern

$$y_n(s, \xi) = y_0 A^n \frac{s^{2n}}{(2n)!} \frac{b^{2n}}{(2n)!}, \quad (3.59)$$

which miraculously can be expressed as

$$y(s, \xi) = \sum_{n=0}^{\infty} y_n(s, \xi) = y_0 \sum_{n=0}^{\infty} \frac{(A^{1/2} s b)^{2n}}{((2n)!)^2} = \frac{y_0}{2} \left[J_0(2A^{1/4} s^{1/2} b^{1/2}) + I_0(2A^{1/4} s^{1/2} b^{1/2}) \right] \quad (3.60)$$

where we have used the identities for the Bessel series found in equations B.5 and B.6.

We define the growth length as the value of s which makes the argument of the Bessel function equal to one. For particles at the end of the bunch, $\xi = -\sigma_z/2$ and $b = 1$ we have

$$y(s, -\sigma_z/2) = \frac{y_0}{2} \left[J_0(2A^{1/4} s^{1/2}) + I_0(2A^{1/4} s^{1/2}) \right], \quad (3.61)$$

and therefore

$$L_g = \frac{1}{4} \sqrt{\frac{\gamma}{Nr_e \sigma_z \hat{W}_{\perp 0}}}. \quad (3.62)$$

Using parameters for the 2014 experiment detailed in Section 6.1.4, we find that the growth length for the particles at the tail of the bunch is about 5 cm. We operated our experiment with a relatively short (8 cm) channel to increase the likelihood that the beam was transported through the channel.

3.7.2 BBU for an Offset Beam with Angular Misalignment

The generic orbit for a beam propagating through a drift is

$$y(s, \xi) = y_0 + y'_0 s. \quad (3.63)$$

We can include the effect of the initial angular offset of the beam by calculating the BBU drift for the angular offset alone and adding it to equation 3.61.

The calculation proceeds identically to the one we just carried out, except that the lowest order orbit is proportional to s instead of a constant. By inspection, we find

$$y_n(s, \xi) = y'_0 A^n \frac{s^{2n+1}}{(2n+1)!} \frac{b^{2n}}{(2n)!}, \quad (3.64)$$

and summing over all orders we find

$$y(s, \xi) = \sum_{n=0}^{\infty} y'_0 A^n \frac{s^{2n+1}}{(2n+1)!} \frac{b^{2n}}{(2n)!} = y'_0 s^{1/2} \left[\frac{J_1(2A^{1/4}s^{1/2}b^{1/2}) + I_1(2A^{1/4}s^{1/2}b^{1/2})}{2A^{1/4}b^{1/2}} \right], \quad (3.65)$$

where we have again used the identities B.5 and B.6. Note that the growth length for a beam with an angular offset is the same as that of a beam displaced from the axis with no angular offset. The initial angle y'_0 is simply added to the angular deflection induced by the BBU.

The total orbit evaluated for the tail of the bunch ($\xi = -l/2$ and $b = 1$) is the sum of equations 3.61 and 3.65:

$$y(s, -l/2) = \frac{y_0}{2} \left[J_0(2A^{1/4}s^{1/2}) + I_0(2A^{1/4}s^{1/2}) \right] + \frac{y'_0 s^{1/2}}{2A^{1/4}} \left[J_1(2A^{1/4}s^{1/2}) + I_1(2A^{1/4}s^{1/2}) \right]. \quad (3.66)$$

In typical experiments at FACET, the initial offset of the beam in the channel tends to drive the instability.

3.7.3 Implications of the BBU on Experiments and Future Colliders

The growth length of the beam breakup instability is extremely short in the case of the hollow channel plasma. The growth length scales roughly linearly with the inner radius a of the plasma channel. Increasing the channel radius by a factor of ten will increase the usable channel length by the same amount. Unfortunately, the accelerating gradient in the channel is roughly proportional to a^{-2} , so increasing the channel radius by a factor of ten implies an increase in the channel length by a factor of 100 in order to achieve the same energy gain. This does not solve the problem.

In conventional linear accelerators, the BBU is mitigated by using external focusing magnets around or between the accelerating structures to focus the beam as it is accelerated. It is not clear how to create an external focusing system that will work with hollow channel plasmas. One issue is that the strength of the deflecting wakefield is much stronger than the strength of state-of-the-art quadrupoles, which operate at about 100 T/m. Another issue is that a strong quadrupole field breaks the symmetry of the hollow channel plasma. Alternatively, one could use a solenoid as the external focusing element. This would preserve the symmetry of the hollow channel, but solenoidal fields can only be produced with strengths of a few Tesla.

Despite these challenges, hollow channel plasmas are an exciting possibility for accelerating positrons in plasma. Experiments at FACET can help us to understand the tolerances on beam offsets in the channel which is a critical step toward understanding the applicability of this technique for a future linear collider.

Chapter 4

Optics for Generating Hollow Channel Plasmas

4.1 Introduction

In order to test the concept of the hollow channel plasma wakefield acceleration, we first need to generate a hollow plasma. Plasmas are fluids, but the hollow channel is a *structure*. At first glance, it seems improbable that we can form a plasma into a cylinder with well-defined boundaries, but that is exactly what is required in order to test the theory of hollow channel PWFA. Previous experiments have demonstrated that this is possible [64]. The plasma channel is formed by ionizing an annular region in an otherwise neutral metal vapor or gas using a high-intensity, shaped laser pulse. The channel holds its form for a only few trillionths of a second, but this is more than enough time for an ultrarelativistic beam to propagate through the channel and create a wakefield. The method described in this chapter improves upon earlier work which created channels that were 8 millimeters long with a 4 μm radius. The hollow channel plasma experiments at FACET require channels that are tens of centimeters to a meter long and a width of a few hundred microns.

We begin this chapter by describing the theory and design of high-order Bessel optics that are used to shape the ionizing laser into an annulus. We discuss the fabrication of these optics and laboratory results that demonstrate their performance. We also discuss the development of advanced optics called axilenses for experiments using Hydrogen gas as the plasma source. Finally, we review the theory of ionization by intense laser fields and its application to experiments at FACET.

4.2 Fourier Optics

The theoretical foundation for the design of the high-order Bessel optic is called Fourier Optics. An excellent introduction to the theory is found in Goodman's text [65]. Our goal is to describe how to produce a light wave with the desired transverse intensity profile by diffracting the wave through an aperture that modulates the phase and intensity of the initially uniform plane wave.

The electric field of a monochromatic, linearly polarized light wave is given by

$$\vec{E} = U(x, y, z)e^{i(kz - \omega t)}\hat{x}, \quad (4.1)$$

with $k = 2\pi/\lambda$ and $\omega = 2\pi c/\lambda$. In this case, the field is polarized along the x direction and propagates in the z direction. The diffracting aperture does not couple components of the E and B fields. In this case, the linearly polarized light remains linearly polarized after it has passed through the aperture*. Therefore, we can describe the evolution of the field component $E_x = U(x, y, z)e^{i(kz - \omega t)}$ with a scalar wave equation

$$\nabla^2 U - \frac{n^2}{c^2} \frac{\partial^2}{\partial t^2} U = 0 \quad (4.2)$$

where n is index of refraction of the medium in which the waves propagates. We will only consider propagation in vacuum where $n = 1$.

After applying the time derivative, the wave equation reduces to the Helmholtz equation

$$\nabla^2 U + k^2 U = 0. \quad (4.3)$$

which can be solved using Green's Theorem [65] to produce the Huygens-Fresnel equation

$$U(x, y, z) = -\frac{ik}{2\pi} \int_{\sigma} U(x', y', 0) \frac{e^{ik|\vec{r}|}}{|\vec{r}|} \cos \theta dx' dy'. \quad (4.4)$$

Equation 4.4 describes the field at a point in space separated by a distance z from an aperture σ in the $x' - y'$ plane. The vector $\vec{r} = (x - x')\hat{x} + (y - y')\hat{y} + z\hat{z}$ is directed from a point on the aperture plane to the field point, and $\cos \theta = z/|\vec{r}|$. Figure 4.1 illustrates the geometry of the integral. The aperture σ may take on an arbitrary shape in the $x' - y'$ plane.

We are interested in the field distribution at a point far from the aperture where $z^2 \gg (x - x')^2 + (y - y')^2$. The magnitude of $\vec{r}$ is

$$|\vec{r}| = z \sqrt{1 + \frac{(x - x')^2}{z^2} + \frac{(y - y')^2}{z^2}} \approx z \left[1 + \frac{1}{2} \frac{(x - x')^2 + (y - y')^2}{z^2} \right]. \quad (4.5)$$

*The waves that we consider here are focused by the aperture and have transverse variation in intensity and phase. Strictly speaking, these waves do not solve $\nabla \cdot \vec{E} = 0$ if they have pure transverse polarization. There is a z component of the E -field with magnitude $E_z \approx E_x \sin^2 \beta$, where β is the angle of the focused rays. In our optical design, $\beta = 5$ mrad and $E_z/E_x = 2.5 \times 10^{-5}$. Ultimately, we are interested in the intensity of light field $I \propto E_x^2 + E_z^2$, where the E_z component contributes less than one part in a billion, so we neglect it altogether.

Figure 4.1: An illustration of the coordinate system used in the Huygens-Fresnel and Fresnel diffraction equations.

Before we plug the approximation into equation 4.4, we note that we only need to retain the terms proportional to $(x - x')$ and $(y - y')$ in the exponent, where they are multiplied by the wavenumber k . Outside of the exponent, we drop those terms, which yields

$$U(x, y, z) = -\frac{ike^{ikz}}{2\pi z} \int_{\sigma} U(x', y', 0) e^{\frac{ik}{2z}[(x-x')^2 + (y-y')^2]} dx' dy'. \quad (4.6)$$

This is the Fresnel Diffraction integral and it is accurate if the condition $k[(x - x')^2 + (y - y')^2] \ll z$ is satisfied. Written in this form, equation 4.6 can be recognized as the convolution between the field at the aperture and the convolution kernel

$$h(x, y) = -\frac{ike^{ikz}}{2\pi z} e^{\frac{ik}{2z}[x^2 + y^2]}. \quad (4.7)$$

Alternatively, we can write the expression as a two-dimensional Fourier transform. Pulling out the terms that do not depend on the integration variables x' and y' gives

$$U(x, y, z) = -\frac{ike^{ikz\left[1 + \frac{x^2 + y^2}{z^2}\right]}}{2\pi z} \int_{\sigma} \left[U(x', y', 0) e^{\frac{ik}{2z}[x'^2 + y'^2]} \right] e^{-i(\lambda_x x' + \lambda_y y')} dx' dy', \quad (4.8)$$

where $\lambda_x = kx/z$ and $\lambda_y = ky/z$. This turns out to be an incredibly useful expression for computational purposes. If we know the field at the aperture $U(x', y', 0)$, we can use Fast Fourier Transform (FFT) algorithms to quickly compute the field in the $x - y$ plane. Moreover, if we know the desired shape of the field in the $x - y$ plane, we can use expression 4.8 to propagate the field back to the aperture, and then design the aperture to provide the desired field. In our case, we already know the desired field shape at the aperture due to prior work by Andreev [66] and Fan [64]. Nevertheless, the theory described here turned out to be extremely useful as we developed our optics and improved

upon previous designs.

4.3 Bessel Optics

Traditional optical systems use spherical or parabolic lenses and mirrors to focus light, and the intensity distribution near the focus has a gaussian profile. The study of these systems is called Gaussian Optics. Most laser wakefield acceleration (LWA) experiments use gaussian optics to produce tightly focused beams with on-axis intensities exceeding 10^{18} W/cm². The drawback of these systems is that the tight focus achieved over a limited distance, referred to as the Rayleigh length

$$z_R = \frac{\pi w_0^2}{\lambda}, \quad (4.9)$$

where w_0 is the radius of the laser at the focus, also referred to as the waist. A typical LWA experiment uses an 800 nm Titanium-Sapphire (Ti:Saph) laser focused to spot size of 20 μ m, so $z_R = 1.5$ mm.

PWFA experiments operate with meter-scale plasmas produced by intense laser pulses that are focused for over a meter. If we were to use lenses to produce a long focus, the beam waist would have to be 500 μ m wide in order to have a Rayleigh length of 1 m. This presents two issues. First, if the laser focus is large, the laser energy is spread over a large area and the laser intensity may be below the ionization threshold. Second, the laser beam has a finite, collimated size before it reaches the lens. At FACET the radius of the laser is 20 mm. The angle of convergence is $w_0/z_R = 0.5$ mrad, which means the focusing optic would have to be placed 40 meters upstream of the focus. This is too large a distance to be contained in the FACET experimental area.

In light of these issues, we choose a different type of focusing optic, called an axicon, which is commonly referred to as a Bessel optic, or a “diffraction-free” optic[†]. The axicon is a conically shaped optic that produces a long, tight focus while being located only a meter away from the start of the plasma, which matches the experimental constraints at FACET [67]. The focal spot produced by the axicon has a transverse intensity profile proportional to $J_0(x)^2$. Figure 4.2 illustrates the principle of axicon focusing. The focused rays all have the same convergence angle β . The high intensity region extends from the axicon to $z_{max} = r_0/\tan\beta$, where r_0 is the initial beam radius. We refer to this region as the line-focus. At FACET, we typically employ axicons with convergence angles $\beta > 5$ mrad, which allows us to produce a line-focus up to 4 meters long.

We can calculate the field and intensity of the line-focus by recognizing the axicon as an aperture that modifies the phase, but not the amplitude of the incoming light. The axicon produces a phase

[†]The term ‘diffraction-free’ refers to the fact that the transverse shape of the focus does not change over the length of the focus.

Figure 4.2: A sketch of the axicon optic showing the focused rays converging at a constant angle β . The focus is the region where the rays cross.

Figure 4.3: Left: Image of the axicon focus in the transverse plane measured in the laser laboratory at FACET with an 800 nm Ti:sapphire laser. Right: The transverse profile along the line $y = 0 \mu\text{m}$ fit with the function $J_0^2(k_\perp r)$. The fit value for the convergence angle $\beta = 5.7 \text{ mrad}$ and the axicon design value is 5.6 mrad .

shift that is linearly proportional to the distance from the center of the optic

$$\Phi_{ax}(r) = -k_\perp r, \quad (4.10)$$

with $k_\perp = k \sin \beta \approx k\beta$ is the perpendicular wavenumber. Assuming that the axicon is uniformly illuminated, the intensity along the line-focus is given by

$$I(r, z) = I_0 2\pi \beta^2 k z J_0^2(k_\perp r) \quad z < z_{max} \quad (4.11)$$

where I_0 is the incident intensity on the axicon. We provide a complete derivation of this expression in Section 4.4.1. The intensity rises linearly in z until the line-focus ends at $z = z_{max}$. The transverse intensity profile is proportional to $J_0^2(k_\perp r)$ along the line-focus. Figure 4.3 shows an example of an axicon profile measured at FACET.

Figure 4.4: A spiral shaped optic produces a phase shift $\Phi(\phi) = m\phi$ where m is an integer.

4.4 High-Order Bessel Optics

We now seek a method for producing a laser focus in which the transverse profile is constant along the focus and there is zero laser intensity on axis. The axicon optic produces a line-focus where the transverse profile does not depend on z , but the intensity is maximum on axis. Is there a simple modification to the axicon that retains the longitudinally-invariant transverse profile but with an intensity minimum on-axis?

The solution is to add angular momentum to the beam. This can be achieved with a spiral phase optic shown in Figure 4.4. The combined phase shift of the axicon and spiral phase optic is

$$\Phi_m(r, \phi) = -k_{\perp}r + m\phi, \quad (4.12)$$

where m is an integer. The resulting fields along the line-focus are proportional to $J_m(k_{\perp}r)$ and the intensity goes as $J_m^2(k_{\perp}r)$, providing the desired on-axis minimum. We now derive these results using the stationary phase approximation.

4.4.1 Stationary Phase Approximation

The axicon and spiral optics combine to form an aperture that shifts the phase of an incoming wavefront that is uniform in intensity and phase. The field at the aperture, immediately after the two optics is

$$U(r', \phi', 0) = U_0 e^{i(-k_{\perp}r' + m\phi')}. \quad (4.13)$$

We will use equation 4.6 to propagate the field to a location z downstream of the aperture, but first we need to rewrite the expression in polar coordinates. Making the substitution $x = r \cos \phi$, $y = r \sin \phi$ and $x' = r' \cos \phi'$, $y' = r' \sin \phi'$, we have

$$U(r, \phi, z) = -\frac{ike^{ikz}}{2\pi z} \int_{\sigma} U(r', \phi', 0) e^{\frac{ik}{2z}[r^2 + r'^2 - 2rr' \cos(\phi - \phi')]} r' dr' d\phi'. \quad (4.14)$$

Plugging in for $U(r', \phi', 0)$ and separating the integral into its radial and azimuthal components gives

$$U(r, \phi, z) = -U_0 \frac{ike^{ik\left(z+\frac{r^2}{2z}\right)}}{2\pi z} \int_0^{r_{max}} e^{i\left(-k_{\perp}r' + \frac{kr'^2}{2z}\right)} r' dr' \int_0^{2\pi} e^{i\left(m\phi' - \frac{kr r'}{z} \cos(\phi - \phi')\right)} d\phi'. \quad (4.15)$$

The integral representation of the Bessel function is

$$J_m(x) = \frac{1}{2\pi} \int_0^{2\pi} e^{i(m\phi - x \cos \phi)} d\phi. \quad (4.16)$$

Comparing this expression with the integral 4.15, we assign $x = krr'/z$. The argument of the cosine introduces a phase shift proportional to ϕ . Since we are ultimately interested in the intensity profile, and not the field itself, we can set our observation point $\phi = 0$, which gives

$$U(r, \phi = 0, z) = -2\pi U_0 \frac{ike^{ik\left(z+\frac{r^2}{2z}\right)}}{2\pi z} \int_0^{r_{max}} J_m\left(\frac{kr r'}{z}\right) e^{i\left(-k_{\perp}r' + \frac{kr'^2}{2z}\right)} r' dr'. \quad (4.17)$$

The exponential term only contributes to the integral when its argument is near a minimum or maximum

$$\frac{d}{dr'} \left(-k_{\perp}r' + \frac{kr'^2}{2z} \right) = 0, \quad (4.18)$$

$$\rightarrow r' = \beta z, \quad (4.19)$$

where we have used $k_{\perp} = k \sin \beta \approx k\beta$. The ratio $r'/z = \beta$ is the geometric convergence angle of the axicon rays. This gives the leading order term in the stationary phase approximation [68]. We replace $r' = \beta z$ everywhere in the integral except in the exponent

$$U(r, z) = -U_0 ik\beta e^{ik\left(z+\frac{r^2}{2z}\right)} J_m(k\beta r) \int_0^{r_{max}} e^{i\left(-k\beta r' + \frac{kr'^2}{2z}\right)} dr'. \quad (4.20)$$

Extending the bounds of the integral to $\pm\infty$ and completing the square gives

$$U(r, z) = -U_0 ik\beta e^{ik\left(z+\frac{r^2}{2z}-2\beta^2 z\right)} J_m(k\beta r) \sqrt{\frac{\pi z}{k}} (1+i). \quad (4.21)$$

The justification for extending the bounds of the integral is again the stationary phase approximation. Since the exponential only contributes to the integral near $r' = \beta z$, extending the bounds to unphysical domains makes the integral analytically tractable without changing its value. Finally, we calculate the intensity of the field

$$I(r, z) = \frac{c}{2} |U(r, z)|^2 = I_0 2\pi \beta^2 k z J_m^2(k\beta r), \quad (4.22)$$

Figure 4.5: Left: Intensity distribution in r and z of an axicon plus spiral phase optic with convergence angle $\beta = 4.4$ mrad and $m = 7$. Right: A comparison of intensity distributions in r showing increasing contrast between the central minimum and first maximum with increasing m .

where $I_0 = cU_0^2$ is the uniform incident intensity just before the aperture. Equation 4.22 has the desired properties for generating a hollow channel plasma. There is an intensity minimum on axis and the transition from the minimum to the maximum can be made sharp by selecting $m > 1$, as seen in Figure 4.5. It can be seen that equation 4.11, which describes the intensity of the axicon without the spiral phase optic, is just a special case of this derivation with $m = 0$.

4.5 Numerical Code for Evaluating Optics at FACET

The derivation of expression 4.11 confirms our strategy for selecting optics for generating a hollow channel plasma. However, the analytic approach is not particularly useful for describing optical fields in systems that have multiple apertures and optical aberrations, as is the case at FACET. We have developed a numerical code package called FACET OPTICS [69] to calculate intensity profiles for experimental conditions encountered at FACET.

The software accepts input from the user which defines the optical aperture. Typically, the aperture is represented by a top-hat amplitude function with a phase shift corresponding to a lens, axicon, or axicon plus spiral phase optic. The field is generated at the aperture plane $z = 0$, and the Fresnel Diffraction integral in the form of a 2D FFT is used to propagate the field to a position z downstream. The program proceeds in a loop, re-scaling the field at $z = 0$ for propagation to new positions in z . A sample algorithm with top-hat amplitude and an axicon-plus-spiral phase aperture is shown in Appendix C.

Additional apertures may occur downstream of the optic where mirrors with central holes are used to fold the laser on to the particle beam trajectory. The left panel of Figure 4.6 shows the effect of a nearby aperture which “wrinkles” the intensity. Aberrations are introduced using the

Figure 4.6: Left: Intensity distribution in x and y of an axicon plus spiral phase optic with convergence angle showing the effect of a nearby aperture. Right: Intensity distribution in x and y of an axicon plus spiral phase optic with a vertical astigmatism in the wavefront, represented by a Zernike polynomial with order $m = 2$, $n = 2$.

Zernike polynomials. Figure 4.6 shows an example of a vertical astigmatism commonly encountered at FACET due to drifts in the laser transport system, represented as a Zernike polynomial of order $m = 2$, $n = 2$.

The FACET OPTICS package, originally developed to address the specific problem of hollow channel plasma generation at FACET, was written with sufficient generality that it can be applied to other problems as well. FACET OPTICS was used in the design of the axilens optic, discussed in Section 4.8, and by groups at UCLA, DESY, and the University of Strathclyde for PWFA and LWA projects not related to FACET.

4.6 Kinoform Phase Gratings

The axicon plus spiral phase optic described in Section 4.4 represent the ideal optic for generating high-order Bessel profiles but in practice, there are some limitations with implementing the optics in this way. First, the spiral phase optic shown in Figure 4.4 is incredibly difficult to manufacture accurately, due to the discontinuity at $r = 0$. Second, the combination of a spiral phase optic and an axicon will produce a very thick optic. This is a problem for high-intensity laser pulses, which undergo self-phase modulation due to the non-linear Kerr effect as they pass through thick optics. The challenge is to produce a thin optic with the same phase properties as the axicon plus spiral optic system, thereby minimizing the self-phase modulation which degrades the quality of the focus downstream of the optic.

We solved this problem by developing a diffracting phase grating, called a kinoform, which approximates the phase of the axicon plus spiral optic system described by equation 4.12 [66, 64, 29].

Figure 4.7: Left: Etching mask for a $m = 7$, binary optic. Right: Etching mask for a $m = 7$, 8-layer staircase optic.

The kinoform is a piece of fused silica that is etched with the spiral phase pattern according to

$$\Phi(r, \phi) = \text{mod}(\Phi_m(r, \phi), \pi). \quad (4.23)$$

The etching was done by electron beam lithography on a fused silicon wafer 1 mm thick and 44 mm by 44 mm wide. The etched portion of the wafer was 40×40 mm, divided into pixels of size $3.33 \mu\text{m}$. The etch depth is set by the index of refraction of fused silica, the wavelength of the incident light, and the number of etching layers. In our first kinoform design, we used a single etching layer corresponding to a phase shift of π radians. Each pixel was assigned a value 0 or π according to equation 4.23. Pixels with phase value π are etched to a depth of 883 ± 9 nm, corresponding to a half-wavelength phase shift $\pi = \lambda/2(n - 1)$. The error on the depth measurement is reported by the manufacturer NIL Technologies. Errors at the 1% level have negligible effect on the quality of the focus. The left panel of Figure 4.7 shows the design of the binary masks used in the etching process.

In previous work, the spiral kinoforms were also designed with a binary etch pattern [66, 64], but a binary phase grating has limited diffraction efficiency into the first (focusing) diffraction order. The diffraction efficiency is defined as the fraction of light transmitted by the aperture that is focused into a given diffraction order. Numerical calculations, later confirmed by experimental measurements, showed that the first-order efficiency was only 44% with the binary kinoform.

We were able to improve the performance of the kinoform by increasing the number of layers in the grating, forming a staircase pattern. Using the FACET OPTICS package, we designed and calculated first-order diffraction efficiencies for optics with up to 20 layers. The efficiency asymptotically approaches 100% as the number of layers tends toward infinity. The trend is shown in the left panel of Figure 4.8. For our improved kinoform design, we chose 8-layer optics as a compromise between

Figure 4.8: Left: Numerical calculation of the efficiency of the first-order diffraction versus number of etched layers. Right: Scanning electron microscope image of an $m = 6$, 8-layer staircase optic. Image courtesy of the manufacturer, NIL Technologies.

fabrication cost and diffraction efficiency. The etch pattern is given by

$$\Phi(r, \phi) = \text{mod} \left(\Phi_m(r, \phi), \frac{\pi}{4} \right) \quad (4.24)$$

The diffraction efficiency of the 8-layer optic is 96%. In the 8-layer design, each layer shifts the phase of the light by $\pi/4$ radians relative to an adjacent layer. The etch depth of each layer is 220.75 ± 4 nm. We procured four sets of kinoforms with $m = 5, 6, 7$, and 8 designed to produce Bessel profiles with maxima at radii 150 μm , 200 μm , 250 μm , and 300 μm respectively.

4.7 Optical Bench Tests

Optical tests of the kinoforms were carried out in the FACET laser room to check that the devices met their design standards. The FACET laser is described in detail in Section 5.16. The tests were performed with a low power pulse that was telescoped and collimated so that it fully illuminated the optic with a uniform phase front. A 14-bit CCD camera with a pixel size of 3.75 μm was used to directly image the attenuated laser. The camera was mounted on a rail so that its position relative to the kinoform could be varied along the focus, while keeping the laser centered on the CCD chip. Figure 4.9 shows images of the Bessel focus of the $m = 8$ optic at 86 cm and 186 cm, respectively. Close to the optic, diffraction effects due to the aperture and the finite resolution of the grating are apparent. Further from the optic, the Bessel profile is uniform, aside from intensity modulations due to amplitude and phase imperfections in the illuminating laser wavefront. These imperfections can be corrected to some extent in the laser transport, as described in Section 6.1.2.

Each image from the optical rail measurement is analyzed and an intensity lineout is extracted

Figure 4.9: Left: Image of the $m = 8$ optic at a distance of 86 cm from the optic. Right: Image of the $m = 8$ optic at a distance of 86 cm from the optic.

from the center of the image. The lineouts are compiled into a two-dimensional matrix with each column of the matrix corresponding to the intensity profile at a given z location. The results from the rail measurement are shown in Figure 4.10.

The rail measurement of the $m = 8$ optic was in good agreement with the predictions from design. The optic produced a uniform, hollow focus over a length of at least 3 meters. Intensity variation along the focus is caused by uneven illumination of the optic, despite our efforts to illuminate the optic as uniformly as possible. The width of the peak-to-peak spacing of the Bessel maxima varies by $\pm 30 \mu\text{m}$ over the length of the focus. The variation in width is due to the fact that the illuminating phase front is not perfectly flat. The variation in the width of the channel could have a significant effect on the shape of the wake in the channel. However, in the experiments describe in Chapter 6, we choose to ionize a much shorter channel length of up to 25 cm, and in that case there is almost no apparent variation in the channel width.

4.8 Axilens Optics

The success of the kinoform design program led us to consider alternative phase functions for focusing optics. The axicon is useful in that produces a long line-focus suitable for creating meter-scale plasmas. However, with uniform illumination, the axicon intensity rises along the line-focus, depositing more energy further downstream. At the same time, laser imperfections due to errors in the phase front become more pronounced at locations further from the optic, which leads to a degraded and misshaped focus. These issues were deemed insignificant for experiments that used lithium vapor as the plasma source. Lithium is relatively easy to ionize, so we moved the lithium source

Figure 4.10: Top: Rail measurement of the $m = 8$ kinoform optic for $86 < z < 362$ cm. The peak intensity is normalized to 1. Bottom: Numerical simulation of the same optic. Diffraction effects are also observed and can be seen as faint wiggles near $z = 86$ cm. The peak intensity is normalized to 1.

closer to the optic where the intensity of the laser was low but uniform. In 2015, experiments began using hydrogen as the plasma source, which has a much larger first ionization energy. In order to ionize hydrogen, we had to extend the distance between the axicon and the ionization region, at which point we were stymied by issues relating to the quality of the Bessel focus.

The solution was to develop a special optic called an axilens which made more efficient use of the laser energy while still retaining some of the useful properties of the axicon. Axilenses are not a new idea [70], but the experiment at FACET is the first instance of an axilens being used to focus a high-power laser for plasma generation.

The phase function that defines the axilens is

$$\Phi(r) = \frac{kR^2}{2\Delta_z} \log \left(f_0 + \frac{\Delta_z}{R^2} r^2 \right), \quad (4.25)$$

where R is the radius of the optic, f_0 is the distance from the optic to the start of the focus, and Δ_z is the length of the focus. The axilens can be thought of as an optic with a lens-like, parabolic phase function near the center, and axicon-like, linear phase away from center. This creates a line-focus with a well defined start and end, so that all of the laser energy is used efficiently. A “high-order axilens”, similar to a high-order Bessel optic, can be achieved by adding an azimuthal phase term

Figure 4.11: A comparison of the intensity along the line-focus for an axicon and axilens illuminated by a 100 fs, 150 mJ laser pulse. The intensity of both optics surpasses the hydrogen ionization threshold of $2 \times 10^{14} \text{ W/cm}^2$ from $z = 3 \text{ m}$ to $z = 4 \text{ m}$. The axilens has significant intensity overhead and therefore can be operated with reduced laser power to achieve the same result.

$m\phi$ to expression 4.25. Figure 4.11 illustrates the usefulness of this optic by comparing it to the axicon. Both optics are illuminated by the same laser pulse, but the axilens has at least four times the intensity of the axicon in the region $3 \text{ m} < z < 4 \text{ m}$.

An 8-layer axilens was designed according to the same principle as the high-order Bessel kinoforms. Rail test measurements were carried out and the results are shown in Figure 4.12. The measurement revealed that the axilens produced a line focus that was significantly shorter than the length calculated using the **FACET OPTICS** package. The full-width, half-maximum (FWHM) of the design focus was 2 meters, but the measured FWHM is only 1.3 meters. The source of the discrepancy is still not fully understood. Nevertheless, the axilens project was successful in creating a optic that made efficient use of the laser energy while forming a meter-scale focus. The efficiency is critical for ionizing large volumes of hydrogen gas, because the laser can be operated with lower energy per pulse, which in turn reduces the effect of self-phase modulation in transmitting optics.

4.9 Tunnel Ionization in Intense Laser Fields

Throughout this chapter, we have discussed techniques for controlling and shaping high-intensity laser pulses for the purposes of ionizing metal vapor or gas, thereby generating a plasma. Here, we discuss the physical principle behind high-field laser ionization and determine the required laser

Figure 4.12: Top: Rail measurement of the an 8-layer axilens optic with $R = 2$ cm, $f_0 = 4$ m, and $\Delta_z = 2$ m. The FWHM length of the focus was 1.3 m. Bottom: Numerical simulation of the same optic. The FWHM length of the focus is 2 m.

intensity for ionization of Lithium and Hydrogen, the two atoms most commonly used as plasma sources at FACET.

The FACET Ti:Sapph laser operates at a central wavelength of 800 nm. The laser photons have a mean energy of 1.5 eV, but the first ionization energies of Lithium and Hydrogen are 5.39 eV and 13.6 eV, respectively, well above the photon energy. This precludes direct photo-ionization of Lithium and Hydrogen. A related ionization mechanism, called multi-photon ionization (MPI), occurs when an atomic electron absorbs multiple photons in a short period of time. This process occurs at too low a rate for Lithium (and at near-zero rate for Hydrogen) to be useful for generating plasma.

The ionization mechanism exploited at FACET is called tunnel ionization. When an atom is exposed to a large, external field, its energy levels are shifted from their nominal value. This is called the Stark effect. If the external field is large enough, the shift in energy levels may be sufficient for electrons to tunnel out of the atomic potential, ionizing the atom.

We are interested in the likelihood, or rate, at which electrons are freed from their atoms. The rate can be calculated by numerical integration of the Time-Dependent Schrödinger Equation with an external, oscillating field. However, it has long been known that semi-classical models which incorporate information on the ground state of the atom can be used to accurately estimate the ionization rate, even for atoms with many electrons [71]. Ammosov, Delone, and Krainov (ADK) developed a semi-classical theory that determines the ionization rate for an effective principle

quantum number in an AC electric field [72]. The effective principle quantum number is defined by

$$n^* = \sqrt{\frac{m_e e^4}{16\pi^2 \epsilon_0^2 \hbar^2} \frac{Z}{\sqrt{\zeta_i}}}, \quad (4.26)$$

(note MKS units) where Z is the charge of the ion after ionization and ζ_i is the energy of the ground state. Starting from the ADK model, physicists working on the topics of LWFA and PWFA have developed simplified engineering formulas for the ionization rate of commonly used atomic species [73]. For Hydrogen, where $n_H^* = 1$, we have

$$W_H \approx \frac{8.5 \times 10^{19}}{E(\text{GV/m})} \exp \left[-\frac{342.55}{E(\text{GV/m})} \right], \quad (4.27)$$

and for Lithium with $n_{Li}^* = 1.59$, we have

$$W_{Li} \approx \frac{3.6 \times 10^{21}}{[E(\text{GV/m})]^{2.18}} \exp \left[-\frac{85.5}{E(\text{GV/m})} \right], \quad (4.28)$$

In each of these expressions, W is the ionization rate in s^{-1} . A value of $W = 1$ indicates complete ionization after one second, and values of $W > 1$ mean that it takes less than one second to completely ionize the species. The field strength E oscillates rapidly at the laser frequency and varies slowly with the pulse envelope. The E field in these expressions refers to the time average of the field over the fast oscillation. Note that due to the exponential factor, the ionization process is a threshold effect. For instance, exposing Lithium to a field with strength 4 GV/m for 100 fs produces less than 1% ionization. Increasing the field strength to 5.3 GV/m produces 100% ionization in the same time window.

Ultimately, we are interested in the ionization fraction of the species $F = W\tau$ after the passage of a laser pulse of length τ . The local electric field strength depends on the pulse length through the field-intensity relationship

$$I = \frac{H(r)}{\tau} = \frac{c\epsilon_0}{2} |E|^2, \quad (4.29)$$

where $H(r)$ is the local fluence, assuming an axisymmetric focal spot. Therefore, $E \propto \tau^{-1/2}$ and we can solve for τ numerically by setting the ionization fraction $F = 1$. In practice, we choose to operate in a regime where the laser intensity is sufficiently large that the ionization fraction depends weakly on τ , which is the case for $I > 1 \times 10^{13}$ W/cm² in Lithium and $I > 2 \times 10^{14}$ W/cm² in Hydrogen. Figure 4.13 shows the intensity pattern and ionization contours along the line-focus of an $m = 8$ optic illuminated by a 40 mJ, 100 fs laser pulse. Note the rapid transition between the ionized and un-ionized region at the edge of the first intensity maximum due to the threshold nature of the ADK ionization rate. When operating above the ionization threshold, we are less sensitive to small variations in laser intensity along the focus, so these channels can be realized experimentally

Figure 4.13: Top: Intensity of the line-focus of an $m = 8$ optic, illuminated by a 40 mJ, 100 fs laser pulse. Bottom: The ionization contour of the intensity pattern in Lithium with a first ionization energy of 3.59 eV.

using laser pulses with small imperfections.

The ADK approximation can also be used to calculate ionization due to the fields of intense electron and positron beams. We describe the production of these beams in the following chapter.

Chapter 5

FACET

5.1 Introduction

The Facility for Advanced Accelerator Experimental Tests (FACET) at SLAC National Accelerator Laboratory is a state of the art research center for studying plasma wakefield acceleration [26]. Since 2011, FACET has produced the highest energy electron and positron beams in the world. In 2016, FACET completed its five-year mission and efforts are currently underway to prepare for FACET-II, the successor facility to FACET.

The Conceptual Design Report for FACET described the motivation for building this facility at SLAC [74]:

SLAC is the only place in the world with the high peak current, high energy electron and positron beams required to continue the development of beam driven plasma wakefield acceleration.

FACET was envisioned to address a number of “second-generation issues” in PWFA research including accelerating both electron and positron beams to high energies with a narrow energy spread, preserving beam emittance, generating small emittance beams, and efficient energy extraction from the plasma wake. FACET produced experimental results on each of these topics over the course of its five-year run.

FACET distinguishes itself from previous PWFA projects by producing beams with extremely large peak currents. The peak current is a key figure of merit for a PWFA facility because the plasma wake amplitude scales with the peak current. In addition, FACET is one of only two facilities in history capable of studying positron PWFA. The previous facility, called the FFTB, was the predecessor to FACET at SLAC. Table 5.1 compares FACET with previous and planned PWFA facilities.

In this chapter, we give a brief overview of particle beam physics and notation. We provide a

Facility	Laboratory	Dates	Beam Type	Energy [GeV]	Peak Current [kA]
AWA	ANL	1988–1998	e^-	0.02	0.325
FFTB	SLAC	1999–2006	$e^\pm$	< 42	17.9
ATF	BNL	2002–2013	e^-	0.06	0.075
FACET	SLAC	2011–2016	$e^\pm$	20.35	19.1
FlashForward	DESY	2017–	e^-	1.25	2.4
AWAKE	CERN	2017–	p^+	400	0.05
FACET II	SLAC	2019–	$e^\pm$	10	< 35

Table 5.1: Table of past, present, and future PWFA facilities, showing beam species, energy, and peak current, which is a key figure of merit for experiments. Parameters taken from references [75, 25, 76, 26, 77, 78, 79].

source-to-sink description of FACET, including the positron source, and discuss the longitudinal phase space manipulations used to generate extremely high peak current bunches. We also describe the facilities and hardware for PWFA experiments at FACET, including the Terawatt-class Ti:Saph laser, beam collimation system, experimental diagnostics, and data acquisition system.

5.2 Particle Beam Dynamics and Notation

In this section, we provide a brief introduction to beam dynamics in a linear accelerator. For a more in-depth treatment, see for example reference [80].

5.2.1 Coordinate System

The electron or positron beam particle is described by six coordinates $\vec{x} = (x, x', y, y', \xi, \delta)$. The x and y coordinates are the transverse position of the beam particle and $\xi = z - \beta ct$ is the longitudinal position of the particle in the co-moving frame. The transverse angles are $x' = p_x/p_\xi$ and $y' = p_y/p_\xi$. The relative longitudinal momentum of the particle is $\delta = (p_\xi - p_0)/p_0$, with p_0 the nominal beam momentum. All coordinates are measured relative to an origin defined by a test particle with the desired trajectory and energy, often referred to as the “golden orbit”.

At FACET, the beam energy is 20.35 GeV and the particles are ultrarelativistic ($\gamma \approx 40,000$ and $\beta \approx 1$). The relationship between the energy and momentum of the particle is given by

$$E = \sqrt{(\vec{p} \cdot \vec{p})c^2 + m_e^2 c^4} \approx \sqrt{p_z^2 c^2 + m_e^2 c^4} \approx p_z c. \quad (5.1)$$

The equivalence between energy and momentum is nearly exact for ultrarelativistic particles, so we usually refer to δ as the relative energy of the particle.

5.2.2 Particle Transport

As the beam travels through the accelerator it encounters a number of beamline elements that affect the trajectory of the beam particles. These include drifts, quadrupole focusing and defocusing magnets, dipole bend magnets, and RF accelerator cavities. Each element affects the orbit of the beam as well as the distribution of particles within the bunch. For instance, RF cavities produce a net affect on the beam by accelerating the particles as they travel through the linac, and they are also used to “chirp” the beam by creating a correlation in $\xi - \delta$. Similarly, dipole magnets deflect the beam as a whole and create dispersion which is a correlation in $x(y) - \delta$. We are primarily interested in describing the intrabeam effects of these elements, *i.e.* the variation around the reference trajectory:

- A drift correlates incoming angle $x'(y')$ with outgoing position $x(y)$,
- A quadrupole correlates incoming position $x(y)$ with outgoing angle $x'(y')$,
- A dipole correlates incoming energy δ with outgoing position $x(y)$,
- An RF cavity correlates incoming position ξ with outgoing δ .

Here, we have considered simple examples that produce linear transformations of the beam coordinates. In general, the effect of beamline elements on the final coordinates is a function of an arbitrary number of the initial coordinates to any power. To first order, the quadrupole correlates particle position with angle, but an important second order effect called chromaticity is the correlation of particle position and energy with the outgoing angle. Therefore, the total effect of a beamline element on a particle is given by the tensor equation

$$x_1^i = R_{ij}x_0^j + T_{ijk}x_0^jx_0^k + U_{ijkl}x_0^jx_0^kx_0^l + \dots \quad (5.2)$$

where $\vec{x}_0$ is the initial coordinate vector, $\vec{x}_1$ is the final coordinate vector and x^i is an element of the vector. Equation 5.2 can be used to describe all possible actions of a beamline element on a particle, with the exception of collective effects such as wakefields.

We now focus our attention on the first-order transport matrix R , which is sufficient for describing most of the beam dynamics at FACET. The action of the 6×6 matrix on the coordinate vector $\vec{x}$ is

$$\begin{pmatrix} x_1 \\ x'_1 \\ y_1 \\ y'_1 \\ \xi_1 \\ \delta_1 \end{pmatrix} = \begin{pmatrix} R_{11} & R_{12} & R_{13} & R_{14} & R_{15} & R_{16} \\ R_{21} & R_{22} & R_{23} & R_{24} & R_{25} & R_{26} \\ R_{31} & R_{32} & R_{33} & R_{34} & R_{35} & R_{36} \\ R_{41} & R_{42} & R_{43} & R_{44} & R_{45} & R_{46} \\ R_{51} & R_{52} & R_{53} & R_{54} & R_{55} & R_{56} \\ R_{61} & R_{62} & R_{63} & R_{64} & R_{65} & R_{66} \end{pmatrix} \begin{pmatrix} x_0 \\ x'_0 \\ y_0 \\ y'_0 \\ \xi_0 \\ \delta_0 \end{pmatrix}. \quad (5.3)$$

The elements of R project the initial coordinates into the final coordinates. We stated above that

a drift element correlates incoming angle with outgoing position. Therefore, the elements R_{12} and R_{34} will be non-zero for the transport matrix representing a drift element. Also note that because the FACET linac is operated as a separate-function accelerator, the particle motion in the x plane is decoupled from the y plane. Therefore, for most elements of the FACET linac R is a 2×2 block diagonal matrix.

The transport matrix is a linear transformation of the beam coordinates, and we can therefore concatenate these matrices together to determine the effect of several consecutive elements on the beam particle. In the case of n elements, we have

$$\vec{x}_n = R_n R_{n-1} R_{n-2} \dots R_2 R_1 \vec{x}_0, \quad (5.4)$$

where R_1 is the most upstream element and acts first on the particle, and R_n is the most downstream element acting last on the particle.

Before moving on, we consider two examples of transport matrices with particular relevance to FACET. First, bunch compression is achieved by chirping the beam in an RF cavity and then sending it through a series of dipole bending magnets and drifts. A simple example of such a system is the symmetric, four-bend chicane in Sector 10 at FACET shown in Figure 5.5b). The R matrix for this system is given by

$$R_{BC} = R_B R_{D1} R_{-B} R_{D2} R_{-B} R_{D1} R_B. \quad (5.5)$$

where the elements R_D represent drift and R_B represent bends. The four dipole magnets in the chicane are of equal strength, with the first and fourth magnets bending the beam in the positive x direction and the second and third magnets bending the beam in the negative x direction. There are three drifts in the chicane, with the first and third of equal length. The goal of the system is to compress the bunch longitudinally without inducing a final transverse position or angle that depends on the energy of the particle. This condition is expressed mathematically by $R_{16} = 0$, $R_{26} = 0$, and $R_{56} \neq 0$, which results in a system of equations that can be solved by a judicious choice of three parameters representing two drift lengths and the bend strength. We often describe bunch compression systems solely in terms of the magnitude of R_{56} .

The second example is a system of elements that images the beam in the x -plane. This type of system is implemented in the downstream spectrometer at FACET and is composed of a drift, a focusing quadrupole, a drift, a defocusing quadrupole, and a drift

$$R_{img} = R_{D3} R_{QD} R_{D2} R_{QF} R_{D1}. \quad (5.6)$$

The purpose of the imaging system is to transport the beam such that the final transverse positions of the beam particles do not depend on their initial angles. This is achieved by adjusting the strengths of the quadrupole magnets in order to satisfy the imaging condition $R_{12} = 0$. We can always satisfy the imaging condition for a narrow range of particle energies, but the beam energy

spectrum downstream of the plasma can be quite large. In this case, we scan the values of the magnets in order to image the beam at several different energies.

5.2.3 Beam Emittance

So far, we have discussed the dynamics of single particles traveling through the accelerator. We are interested in the properties of the beam as a whole. For example, the transverse beam size is given by

$$\sigma_x^2 = \int_{-\infty}^{\infty} (x - \langle x \rangle)^2 f(x) dx \quad (5.7)$$

where $f(x)$ is the normalized transverse beam distribution, which we assume to be gaussian. The quantities $\sigma_x'^2$, σ_y^2 , and $\sigma_y'^2$ are derived the same way and we also assume that their respective distributions are gaussian. The quantities σ_ξ^2 and σ_δ^2 can also be defined this way but the distributions $f(\xi)$ and $f(\delta)$ tend to be non-gaussian and the definition is less useful as a result.

The volume of phase space occupied by the beam particles is called the emittance. We define a separate emittance for each transverse plane

$$\varepsilon_x = \sqrt{\sigma_x^2 \sigma_{x'}^2 - \sigma_{xx'}^2}, \quad (5.8)$$

where

$$\sigma_{xx'}^2 = \int_{-\infty}^{\infty} (x - \langle x \rangle)(x' - \langle x' \rangle) f(x, x') dx dx'. \quad (5.9)$$

The beam emittance is a measure of the beam quality. A beam with a small emittance can be brought to tighter focus than a beam with a large emittance in a given focusing system.

The beam emittance depends on the beam energy through $x' = p_x/p_\xi$. As the beam is accelerated in the linac, p_ξ increases but p_x does not, and as a result x' decreases. This effect is called adiabatic damping. When discussing beam emittances at FACET, we typically refer to the normalized beam emittance $\varepsilon_n = \gamma\varepsilon$. The normalized beam emittance is conserved as the beam is accelerated in the linac, assuming there are no alignment errors or magnetic field errors. In practice, there is always some dilution of the beam emittance due to these types of errors. The accelerator operations team works to optimize the beam orbit to minimize emittance growth.

The beam size throughout the accelerator is given by

$$\sigma_x = \sqrt{\beta(z)\varepsilon_x}, \quad (5.10)$$

where $\beta(z)$ is the betafunction and is one of the three Courant-Snyder parameters α , β , γ . The betafunction defines the beam envelope in the linac and evolves according to

$$\beta''(z) + 2k^2(z)\beta(z) - \frac{2}{\beta(z)} = 0, \quad (5.11)$$

Figure 5.1: Schematic of the accelerator and compression systems at FACET. Thick yellow sections indicate the presence of RF accelerating structures in the linac. Black lines are used to highlight the path of electrons and the gray line indicates the path of positrons, starting from the target in Sector 19. The locations of the energy spread and bunch length diagnostics throughout the linac are highlighted in red and blue, respectively.

where $k(z)$ is the wavenumber of the periodic focusing lattice. The betafunction reaches a minimum value, referred to as β^* , at a waist chosen to be near the interaction point (IP). In free space, the evolution of the betafunction is given by

$$\beta(z) = \beta^* + \frac{s^2}{\beta^*}. \quad (5.12)$$

Near the IP, β^* determines the divergence length of the beam focal spot. It is analogous to the concept of the Rayleigh length applied to laser beams as discussed in Section 4.3.

5.3 A Source-to-Sink Description of FACET

FACET uses the first two-thirds of the original 3 kilometer-long SLAC linac. This includes the original thermionic source, damping rings, and s-band linac which operates at a frequency of 2.856 GHz. A four-bend chicane for positron compression at Sector 10 and the final W-chicane were installed specifically for FACET. Figure 5.1 shows an overview of the facility.

5.3.1 Thermionic Gun

The beam source for FACET is a thermionic electron gun that was developed and commissioned for the Stanford Linear Collider (SLC) in 1981 [81]. The gun consists of heated tungsten cathode and DC grid at 120 kV. The grid is driven by an avalanche pulser and produces 4×10^{10} electrons in a 2.5 ns long pulse. The radial beam size is about 1 cm. With a kinetic energy of 120 keV, the electrons are traveling at roughly 60% the speed of light. Figure 5.2 shows a cross-section of the thermionic gun. The beam is accelerated to the right.

Figure 5.2: Cut-away view of the thermionic gun. Reproduced from reference [81].

5.3.2 Capture and Acceleration

The bunch leaving the gun has sufficient charge for our experiments but the bunch length is much longer than the s-band bucket length. The bunch length is reduced by velocity compression in two sub-harmonic buncher cavities [82]. These cavities operate at the sixteenth sub-harmonic of the nominal s-band frequency (178.5 MHz). Each buncher has a short 3.5 cm gap where the field is applied. The first buncher is followed by a 108 cm long drift space and the second buncher is followed by a 36 cm long drift space. The purpose of the second sub-harmonic buncher is primarily to “tuck-in” the beam particles that sample the non-linear field away from the zero-crossing of the first buncher.

The beam enters the buncher near the zero-crossing of the RF field. The electrons near the front of the bunch experience a decelerating field and the electrons toward the back are accelerated. The energy modulation is only a few keV, but this is enough to reduce the bunch length from 2.5 ns to 0.25 ns in the short space allotted. A large solenoid encloses the sub-harmonic buncher to prevent the low-energy beam from diverging during the bunching process.

The sub-harmonic bunchers are immediately followed by a short, 10 cm long, s-band buncher. The s-band buncher further compresses the beam by another factor of 10 down to 25 ps, and accelerates the bunch to 2 MeV. At this point, the electrons are traveling 98% the speed of light and compression by velocity modulation is no longer possible. The s-band buncher is part of a larger accelerating structure called K-02 which accelerates the beam to 38 MeV with solenoidal focusing.

Figure 5.3: Schematic of the SLAC beam source. Two guns are shown on the left. The thermionic gun (top) is used at FACET while the polarized gun (bottom) is decommissioned. The inset to the schematic shows the beam energy and bunch length after the gun, the sub-harmonic bunchers, s-band buncher, and K-02 accelerating section. The bunch length does not change after the K-02 section and prior to injection into the damping ring.

Figure 5.3 shows the capture system from the gun through the K-02 section.

The K-02 structure is followed by a traditional FODO lattice. The beam is accelerated in four more sections up to 220 MeV before it reaches a vertical chicane where positrons from the return line are coupled into the linac at Sector 1. In Sector 1, the injected positrons and electrons are accelerated up to 1.19 GeV. Note that with a bunch length of 25 ps, the electron beam occupies 25 degrees of the s-band bucket, so the beam has a large energy spread when it is transferred to the damping ring via the Linac to Ring chicane (LTR) at the end of Sector 1. Only 70% of the initial charge at the gun is accelerated and transported through the LTR and into the damping ring.

5.3.3 Damping Rings

The beam produced by the thermionic gun is inadequate for experiments at FACET. In particular, the bunch length is too long, the energy spread is too large, and the transverse emittance is too large. The typical normalized transverse emittance for electron beams entering the damping ring is about 100 mm·mrad and about 1000 mm·mrad for the positron beam. The damping rings are used to shrink the beam's transverse and longitudinal phase space so that it can be used in experiments.

The mechanism for phase space reduction is synchrotron damping. As a particle travels around the ring, it radiates away energy in the form of synchrotron radiation [83]. The radiation carries

away both energy and transverse momentum from the particle. Two RF accelerating cavities in the ring are used to replace the energy lost to synchrotron radiation. In doing so, they only restore the longitudinal component of the momentum. The transverse momentum that was radiated away is not replaced, and therefore the transverse phase space is reduced. If the beam is slightly off energy, it will perform a synchrotron oscillation about the stationary phase. This oscillation is also damped by the radiation mechanism. Both the transverse and longitudinal phase space volumes would damp to zero, if not for quantum effects [84].

The damping rings at FACET were designed to damp pairs of electron and positron bunches for 120 Hz collisions at the SLC. The repetition rate was a key factor in determining the size of the damping rings. Specifically, the repetition rate sets the store time for the electron beams in the damping ring to be 8.3 ms, or 99.6% of the total time per cycle at the SLC. From this constraint, the transverse damping time was designed to be 3.5 ms, so the beam emittance is reduced by roughly a factor of 10.

FACET operates at a relatively low repetition rate of 10 Hz with a 3.2 nC charge in the electron beam and a 2.4 nC charge in the positron beam. The store time in the damping ring is extended to 16.6 ms to allow additional damping. In principle, this achieves a reduction in the transverse emittance by a factor of over 100. However, the equilibrium beam emittances of 30 mm·mrad in x and 3 mm·mrad in y are achieved after only 12 ms of store time.

Low transverse emittances are critical for producing a tightly focused beam in the FACET experimental area, but the energy spread and bunch length of the beam leaving the ring are also important parameters in determining the final bunch length. The equilibrium r.m.s. energy spread of the beam in the ring is $\sigma_\delta = 7.4 \times 10^{-4}$ is set by the beam energy and the geometry of the ring. The equilibrium r.m.s. bunch length is $\sigma_z = 5.2$ mm and is a function of the beam energy, RF amplitude and frequency, and the momentum compaction factor α of the ring [85].

In the first two years of FACET, there was concern that the bunch length of the beam leaving the ring was considerably longer than what was expected from theoretical estimates and what had been measured previously at the SLC. A series of streak camera measurements were made to determine the bunch length. The measurement uses synchrotron light emitted by the beam in one of the dipoles which reflects off an off-axis mirror and is transmitted to the streak camera via an optical image relay. The streak camera converts the light into electrons and uses an internal RF field to convert the temporal signal into a spatial streak. Figure 5.4 shows the beam profile from 5 consecutive turns in the ring. The beam profile is streaked along the y -axis and an additional slower streaking is used to separate consecutive turns along the x -axis. The measured bunch length was 6.6 mm r.m.s. at the nominal RF amplitude of 980 kV at during FACET operation. This is about 25% longer than what was observed in previous measurements [86]. The discrepancy was never fully understood, but the leading hypothesis is that a waveguide window in the ring RF system that was installed after the SLC era lowers the effective amplitude of the RF in the two ring cavities. The end result is that

Figure 5.4: Left: Streak data showing 5 consecutive turns of the electron beam in the NDR as well as the projection of the electron beam current showing a 6.6 mm r.m.s. profile. Right: Bunch length versus RF amplitude showing a characteristic $V^{-1/2}$ dependence. Measurements performed by A. Fisher and J.M. Allen.

the beam is longer than expected when it leaves the ring, and this is compensated by increasing the compressor amplitude in the Ring to Linac chicane (RTL).

5.3.4 Compression and Acceleration

The beam is extracted from the damping ring with a fast kicker magnet and sent to the main linac via the RTL. The RTL is the first of three magnetic compression sections where the bunch length is reduced. An s-band compressor cavity is located immediately after the exit of the ring. The beam crosses the cavity close to zero crossing, with the head of the beam gaining energy with respect to the tail. The resulting $z - \delta$ distribution is correlated, and we refer to this correlation as the chirp. The beam traverses the RTL chicane which has a negative value of R_{56} , meaning that high energy particles take a longer path than low energy particles. Depending on the amplitude of the compressor cavity RF, and therefore the strength of the chirp, the beam can be under-compressed, fully-compressed, or over-compressed when it exits the RTL [87]. Typically, we operate FACET with a slight over-compression in order to match the beam chirp out of the RTL with the chirp that is applied to the beam in Sectors 2-10. However, depending on the final desired beam profile, under-compression and full-compression are sometimes used as well. The left subplot of Figure 5.5 illustrates the principle of bunch compression via a chirp and magnetic chicane.

The over-compressed beam is roughly 1.2 mm long r.m.s. when it enters the main linac. There is residual chirp, with the energy increasing from head to tail. The 1.19 GeV beam is accelerated to 9 GeV over 800 meters of s-band linac from Sector 2 to Sector 10. The cavities in Sector 7 and 8 are unpowered. The beam is accelerated off-crest at an RF phase of $-20^\circ \pm 1^\circ$ degrees. The off-crest

Figure 5.5: a.) Schematic depicting bunch compression. i.) A long uncorrelated beam enters an RF cavity near the zero crossing of the RF field. In this illustration, the beam travels to the right and the front of the beam is decelerated while the tail is accelerated. There is no change to the centroid beam energy. ii.) The chirped beam enters a magnetic chicane with positive R_{56} . The particles in the front of the bunch traverse a longer distance than particles in the back, which causes the overall bunch length to shrink. iii.) After full compression, the bunch size is reduced to a minimum value determined by the initial, uncorrelated energy spread and there is no residual chirp. b.) Schematic of the 4-bend Sector 10 bunch compressor, showing the pre-FACET electron arm (top) and the newly installed positron arm (bottom).

chirp is consistent with the over-compressed chirp from the RTL. Two sections of RF cavities in Sector 9 are used as an energy feedback to keep the beam energy fixed at 9 GeV when it enters the Sector 10 magnetic chicane. Energy feedbacks are important because the RF phase of all the sectors drift throughout the day, which leads to net change in the energy gain.

The second stage of bunch compression occurs at the end of Sector 10 in the Linac Bunch Compressor Chicane (LBCC). The electron arm of the LBCC was installed in 2002 for the pre-LCLS Sub-Picosecond Photon Source (SPPS) experiments [88]. The positron arm of the chicane was installed as an accelerator upgrade for FACET in 2011. The LBCC is a 4-bend chicane with positive R_{56} . The schematic of the chicane is shown in the right subplot of Figure 5.5. The low energy particles at the front of the bunch take a longer path than the high energy particles in back. At full compression, the bunch length after the Sector 10 chicane is about $50 \mu\text{m}$. The RF phase in Sectors 2-10 is adjusted for over-compression and under-compression depending on the final bunch profile that is desired for experiments.

The compressed beam entering Sector 11 has a very large peak current, typically greater than 7 kA. The high-current beam drives a strong, longitudinal wakefield in the RF cavity structures over the final kilometer of the FACET linac. The wakefield decelerates particles in the tail of the bunch. The chirp is linear over most of the bunch, with the head of the beam at higher energy than the tail. The beam is accelerated on crest in Sectors 11-19 to a final energy of 20.35 GeV. The acceleration in Sectors 17 and 18 are used as part of a feedback system to control the energy of the beam into the FACET W-chicane and the electron scavenger line.

Figure 5.6: Left: Diagram of the positron production target. The flux-concentrator is the triangular device in the center of the image. Right: Field profiles as a function of z , the distance from the target in mm for the flux-concentrator (solid line) and tapered solenoid (dashed line). The scale for the flux-concentrator is the inner y -axis. Figures reproduced from references [90] and [89].

5.3.5 Positron Production and Return Line

For positron production, a kicker magnet is used to divert electron bunches from the main linac into the scavenger line at the start of Sector 19. The scavenger line bends the electron beam towards a rotating, high-Z tungsten alloy target. The target is roughly 6 radiation lengths thick. Positrons are generated by pair conversion of gamma rays produced from electron brehmstrahlung in the target.

The positrons produced in the target have a large energy spread and divergence and are captured by a “flux-concentrator” [89]. The flux-concentrator is a pulsed, solenoidal magnet which produces axial magnetic fields of almost 6 T. The enormous solenoidal field prevents the high emittance positron beam from diverging too rapidly. At the end of the 2016 run, the flux-concentrator began to produce large vibrations during pulsed operation, leading to a vacuum leak. We operated without the flux-concentrator for the final two weeks of the FACET run which reduced the positron throughput by 30%.

An additional tapered solenoid is used to couple the beam out of the flux concentrator and into the first accelerating section. The accelerating section uses high-gradient, constant impedance s-band cavities to bring the beam up to 200 MeV. The s-band section is enclosed by a 0.5 T solenoid. Figure 5.6 shows a diagram of the target section including the flux-concentrator as well as the field profile of the flux-concentrator and tapered solenoid.

After being accelerated to 200 MeV, the positron beam is focused in a FODO lattice and reaches a vertical turn-around section where the beam is raised from the ground to the ceiling. The vertical turn-around, also referred to as the East turn-around (ETA) has a trombone feature which is used to adjust the path length of the positrons for timing purposes. The positron beam is transported through Sectors 19 to 1 in reverse in a large beam pipe along the ceiling of the linac. The beam is

focused by a weak (large β) FODO lattice. The energy spread of the beam in the return line may be as large as 5% r.m.s.

At the start of Sector 1, the positrons reach the West Turn Around (WTA) and are injected into the main linac via the positron inflector, located at the same position as the four bend vertical chicane traversed by the electrons at the end of Sector 0. The injected positrons are accelerated in Sector 1 up to 1.19 GeV before being sent into the South Damping Ring (SDR) via the South Linac to Ring chicane (SRTL). The positron yield at injection into the SDR was never better than 1.5×10^{10} during FACET operation. Following injection into the SDR, the acceleration and compression scheme for positrons is identical to the process described for electrons in the preceding sections.

5.3.6 Sector 20 Compression and Optics

Sector 20 is the end of the FACET beam line. Prior to the FACET project, Sector 20 was a 100 meter long section of RF structures. Now it contains the final bunch compressor, called the W-chicane, the final focus system, the experimental area, the imaging spectrometer, and the beam dump.

The W-chicane was built for FACET using repurposed magnets from the SLC final focus and FFTB beamlines. The R_{56} is adjusted by changing the strengths of the quadrupoles in the chicane, which provides an additional handle for controlling beam chirp and compression. For the runs in 2014 and on, the R_{56} was fixed at -7 mm, since it was more practical to control bunch compression by adjusting the chirp in the linac. Note that the polarity of the W-chicane magnets are switched when changing between experiments with electron beams and positron beams. The beam energy entering the W-chicane is 20.35 GeV with the head of the beam at higher energy than the tail. The R_{56} in the chicane compresses the chirped beam down to its final bunch length of roughly $20 \mu\text{m}$. For a beam charge of 2×10^{10} particles per bunch, the beam peak current is greater than 19 kA at full compression.

Following bunch compression in the W-chicane, the beam enters the final focus system. The final focus is composed of five sets of quadrupole magnets which are used to match and focus the beam down to a transverse spot size of $20 \times 20 \mu\text{m}$ for values of $\beta^* < 1$ m. The maximum β -function in the final focus system is about 100 m, which is larger than the original design, but generally acceptable because the final focus magnets have low field aberrations. Figure 5.7 shows the lattice functions of the W-chicane from the start of S20 to the dump. Note that the lattice functions shown here reflect an earlier implementation of the Sector 20 optics that used larger β -functions throughout the chicane. The beam has a very large energy spread due to the wakefield chirp, and this combined with the large β s in the chicane led to strong non-linear distortions from higher-order magnetic fields during the first two years of experiments at FACET. The non-linear distortions prevented the beam from being tightly focused at the interaction point (IP). The solution was to reduce the incoming energy spread from the linac and to reduce the β -functions throughout the chicane, with the acceptable trade-off of increasing them in the final focus. In addition, two of the three pairs

Figure 5.7: The original lattice optics for FACET Sector 20 including the W-chicane, final focus and spectrometer. The x -axis shows the longitudinal coordinate s in meters, as measured from the start of Sector 20. The four lattice functions plotted are $\sqrt{\beta_x}$ and $\sqrt{\beta_y}$ in $\text{m}^{1/2}$ on the left y -axis and D_x and D_y in m on the right y -axis. The IP is located at $s = 64$ m. The shape of the W-chicane is reflected by the D_x curve. At the start of the 2013 run, the β -functions were reduced everywhere in the chicane to reduce non-linear magnet errors. The optics shown here were in place during the 2011 and 2012 runs. Figure reproduced from reference [91].

of sextupoles in the W-chicane were remounted onto remotely controlled movers. The sextupoles are used to correct chromaticity in the final focus and their positions are actively tuned to cancel dispersion and minimize the beam size at the IP.

After the beam is brought to a focus in the IP area, it rapidly diverges and needs to be refocused in the imaging spectrometer. The imaging spectrometer, composed of a pair of quadrupole magnets and a strong, vertical dipole, is used to determine the beam's energy and emittance after interacting with the plasma. Previous spectrometer designs for PWFA experiments at the FFTB were simpler in the sense that they only utilized a dipole placed close to the exit of the plasma to measure the beam energy. The benefit of such a design is that the optics are simple and it is possible to measure the full beam energy spectrum on every shot. The imaging spectrometer at FACET is not able to image the full energy spectrum of the beam after high-gradient acceleration in a uniform plasma because it is not possible to refocus a beam with 100% spread in energy. This means that the quadrupoles

Figure 5.8: Left: To ameliorate short-range wakefields, the lower-charge positron bunch leads the electron bunch down the linac by half an RF cycle, or roughly 5 cm. Right: Positrons traverse the upper arm of the sailboat chicane and electrons traverse the lower W arm. Both arms act to compress the bunch. The R_{56} of each chicane is tunable so that the difference in arrival time of the electrons and positrons at the entrance to the plasma can be made as small as needed [74].

in the spectrometer must be scanned in energy in order to reconstruct the beam energy spectrum, and this is a disadvantage compared to the dipole-only spectrometer design. However, there are several distinct advantages to the imaging spectrometer design employed at FACET over previous dipole-only designs. First, the imaging spectrometer can be set to image different locations in the experimental area, which allows us to distinguish energy changes from vertical kicks. Second, an imaging spectrometer is critical for resolving MeV scale changes to the beam energy after propagation in a hollow plasma channel, which is achieved by setting $R_{34} = 0$, as discussed in 5.2.2. Lastly, the chromatic effects associated with large energy-spread beams lets us measure the slice emittance of the beam in the x plane using the “butterfly effect”. The butterfly refers to the pattern made by the large energy-spread beam around the focused energy. The beam is pinched at the focused energy position in y and expands rapidly above and below this location. A fit to the transverse beam size as a function of energy provides a measure of the horizontal emittance.

After passing through the dipole spectrometer, the beam propagates for 10 meters in an expanding beam pipe. The beam exits the beam pipe through a steel window into air at an optical table just upstream of the beam dump where we measure the beam energy spectrum using scintillating screens and Cherenkov radiation generated by the ultrarelativistic beam propagating through air. Finally, the beam reaches the beam dump. The beam dump, designed by the indomitable Dieter Walz, is made of copper and tungsten and is surrounded by polyethylene to absorb neutrons. It is rated for 11 kW but average beam operation at FACET delivers less than 1 kW.

Figure 5.9: Beamline diagram of the Sector 20 illustrating the concept of two-bunch generation and critical elements of the experimental area. a) beam notching device, b) transverse deflecting structure, c) initial spectrometer, d) final-focus quadrupole magnets, e) lithium plasma ionization laser, f) lithium vapour column, g) spectrometer imaging quadrupole magnets, h) spectrometer dipole magnet, and i) Cherenkov and phosphor screens. Bend dipole magnets in the W-shaped chicane are each labelled D. The arrow indicates the electron beams direction of motion. Figure and Caption reproduced from reference [45].

The Sailboat Chicane

The design of the W-chicane is unique to FACET. At the FFTB, the precursor to FACET, a traditional dogleg chicane was used for the final beam compression. For FACET, the more complicated W design was chosen to accommodate a “sailboat chicane” capable of transmitting simultaneous electron and positron beams for electron-drive, positron-witness PWFA experiments. Figure 5.8 shows the original lattice design of the sailboat chicane system. Using an electron beam to drive the plasma wakefield for positron acceleration is likely necessary for the design of a plasma-based linear collider, because the creation of positron beams is energy intensive. However, the positron arm of the sailboat chicane was never built and the complicated W-chicane remains. We remain cautiously optimistic that the sailboat chicane will be built for the FACET-II project.

5.4 Two-Bunch Generation

In order to perform two-bunch, or drive-witness, PWFA experiments you need to have two bunches separated by distance less than the plasma wavelength. The plasma wavelength is typically a few hundred microns, but the wavelength of the s-band RF in the linac is 10 cm. We can rule out the possibility of using bunches in consecutive RF buckets for the two-bunch experiments, because our chicane with $R_{56} = 7$ mm cannot close the 10 cm gap. Instead, we use a single bunch and divide it into two bunches by using a system of collimators in conjunction with the beam chirp.

Following the first bend of the W-chicane, the beam comes to a point of high energy contrast,

Figure 5.10: a) LiTrack simulation of the under-compressed beam at the end of Sector 20. Note that the beam naturally has a two bunch profile without any collimation. b) The same simulation, but with the center of the spectrum notched out and the low energy tail removed. Note that LiTrack uses z to designate the co-moving bunch coordinate. The beam travels to the left with the head of the beam at values of $z < 0$.

where the dispersion is large and β_x is small. This means that the transverse position of the beam particles is well correlated with their energy. A notch-collimator, which is a thin blade made of tantalum, is inserted into the beam and scatters particles from the center of the energy spectrum. Jaw-collimators are inserted from the left and the right to remove high and low energy particles, respectively. The fact that the beam is chirped means that particles removed from a given location in the beam energy spectrum are also removed from a correlated position in the beam’s longitudinal profile. Those particles that survive the collimation process are further compressed as they continue through the W-chicane. This process is illustrated in Figure 5.9.

Note that the illustration demonstrates “over-compression”, where the high energy particles initially at the head of the bunch end up at the back of the bunch at the end of the W-chicane. This method was used for the 2013 run that produced the two-bunch *Nature* result [45]. However, from a beam-tuning perspective, the over-compressed setup was unreliable and prevented timely delivery to the experiments. Through an extensive simulation search using the codes LiTrack [92] and Lucretia [93], it was determined that setting the chicane $R_{56} = 7$ mm could be used for an under-compressed setup that naturally provided a two-bunch longitudinal profile without the use of collimators. This setup was implemented in 2014 and proved to be successful, although the collimation system is still used to enhance contrast between the drive and witness bunches. Figure 5.10a) shows the bunch at the end of Sector 20 without the use of collimators, and b) shows the increased contrast when collimators are used. The bunch travels to the left in this simulation.

LiTrack

The **LiTrack** code is a 2D, longitudinal phase space, particle tracking code implemented in **MATLAB**. The code has been in use at SLAC for several decades, first implemented by Karl Bane as **LTrack** and ported to **MATLAB** as **LiTrack** by Paul Emma. The code runs relatively quickly and includes a graphical user interface (GUI). The code was traditionally used to test and verify wakefield models and beam evolution in the SLAC linac.

For FACET, **LiTrack** was used as a design tool for optimizing the three chicanes and linac phases. The design study required thousands of iterations of the **LiTrack** algorithm to test different parameters. Here, a speedy code was a necessary to complete the optimization quickly and expand the parameter space. Over the course of study, we discovered a number of inefficiencies in the code, including a nested loop in the wakefield calculation that ran over the wakefield bins N_b . By vectorizing the calculation, we increased the speed of the code by N_b^2 . Since N_b is typically 128 or 256, this amounted to a significant speed up and allowed for new applications. One of these was a “flight-simulator” for FACET that allowed for real-time bunch length predictions [94]. My version of the code (**LiTrack4FACET**) has been adopted by other users at SLAC and DESY.

5.5 Experimental Area and Systems

For PWFA experiments, the location of the beam focus is set to the start of the Lithium heat-pipe oven, about 3.5 meters after the end of the last magnet in the final focus system. The region between the final focus and the imaging spectrometer is the experimental area. In addition to the Lithium oven, most of the experimental hardware is located in this region, including the laser compression and integration chambers, OTR foils and YAG crystals for beam diagnostics, as well as components for other experiments such as dielectric wakefield acceleration. Figure 5.11 shows a simplified schematic of Sector 20, highlighting the critical diagnostics. The experimental area is shown as an inset.

5.5.1 Heat Pipe Oven

The heat-pipe oven plasma source is the central feature of the experimental area. The heat-pipe oven is a stainless steel beam pipe with a wire mesh interior. Prior to installation in the beamline, the wire mesh is evenly coated with lithium. The steel pipe is wrapped externally with three resistive heating coils that can be individually controlled to vary the length of the vapor region. There are external water cooling jackets at each end of the oven to ensure a sharp temperature gradient between the hot oven and room temperature beam pipe. The entire apparatus measures 1.5 meters long and is enclosed in ceramic thermal installation and a stainless steel box.

A pair of thin beryllium windows are used to separate the oven volume from the high-vacuum conditions of the linac. Helium gas is pumped into the oven beamline at densities up to 30 Torr.

Figure 5.11: Line drawing of the experimental setup showing all of the critical hardware and diagnostic devices. Inset is to scale. Reproduced from reference [30]

The heating coils raise the temperature of oven to up to 1000°C , and the lithium melts and then evaporates. The hot lithium vapor atoms meet the room temperature helium atoms at either end of the oven near the cooling jackets. Here, they collide and exchange energy with the helium atoms in a 5 cm long region. The cold lithium condenses on the wire mesh and is wicked back to the center of the oven where it is reheated and re-evaporated. The left image in Figure 5.12 shows the heat pipe oven in the laboratory at UCLA.

The vapor density inside the oven is controlled by adjusting the pressure of the buffer gas. The amount of lithium vapor in the oven increases with heater power (oven temperature) until the pressure of the lithium vapor matches the pressure of the buffer gas. Increasing the heater power further does not increase the density of the lithium vapor. Instead, it causes the vapor to expand out from the center of the oven, producing a long column of nearly pure lithium gas [95]. The vapor density profiles are inferred from temperature measurements using a thermocouple that is pulled through the center of the oven. The right image in Figure 5.12 shows the vapor density measurement for one half of the plasma oven using the same parameters as the experiment. The density profile is symmetric about the center of the oven.

The choice of lithium as the plasma medium is motivated by several factors. First, it has a very low first-ionization energy, making it easy to ionize with a relatively low intensity laser or low current

Figure 5.12: Left: The heat-pipe oven on the test bench at UCLA. Right: Bench measurement of the vapor density in the heat-pipe oven for the conditions used in the experiment. Only one half of the plasma density profile is measured, as it is assumed to be symmetric about the midpoint. The length of the vapor region is defined by the full-width, half-maximum (FWHM) of the density measurement and is 132 cm in this case. Image and data from K.A. Marsh and N. Vafaei-Najafabadi.

particle beam. Second, it has the largest gap between the first and second-ionization energies of any atomic species, so the plasma density is insensitive to intensity of the ionizing laser after the first ionization. Finally, it is safer to handle than other easily ionized atoms like rubidium and cesium.

5.5.2 SYAG Energy Spectrometer

The SYAG energy spectrometer is located upstream of the experimental area in the final arm of the W-chicane. ‘SYAG’ is presumably a SLACism for YAG spectrometer, where YAG stands for yttrium aluminum garnet ($\text{Y}_4\text{Al}_5\text{O}_{12}$) which is the type of scintillating crystal used to measure the spectrum. The SYAG spectrometer is a critical diagnostic for the hollow channel plasma wakefield experiments because it measures the incoming particle energy and beam energy spectrum, upstream of the plasma interaction region, on every shot. This “non-destructive” measurement is performed by deflecting the beam with a half-period vertical wiggler magnet [96]. The beam trajectory through the wiggler causes the beam to bend sharply and emit synchrotron x-rays. The x-rays are intercepted by the YAG crystal that is inserted just above the beam trajectory, about 50 cm downstream of the final wiggler magnet. The total energy emitted by each particle in the beam is about 250 KeV, or 0.001% of the beam energy. The wiggler is located at a point where the horizontal dispersion η_x is large and β_x is small, which provides high energy contrast, similar to what is found at the location of the collimators. The intensity of the x-rays detected on the YAG crystal is proportional to the number of beam particles at each location in x , so the measured intensity distribution on the YAG screen corresponds to the beam energy spectrum. Figure 5.13a) provides a schematic of the diagnostic.

Figure 5.13: a) Illustration of the SYAG energy spectrometer. Note that the scintillating YAG crystal intercepts the x-rays emitted by the beam, but not the beam itself. b) Example of a dispersion measurement where the beam energy is varied by ± 150 MeV. The movement of the beam profile is used to calibrate the horizontal axis. c) Example of a beam energy spectrum using the calibration from the measurement shown in b).

The dispersion at the SYAG can change as the chicane sextupoles are moved for spot-size tuning. For this reason, we regularly measure the dispersion at the SYAG by making small changes to the incoming beam energy and recording the shift in the horizontal position of the spectrum. Figure 5.13b) shows an example of this measurement. The measured value of $\eta_x = \Delta x / \Delta E$ is used to calibrate the horizontal axis of the SYAG intensity profile and derive the energy spectrum shown in Figure 5.13c). Typical values for the dispersion at the SYAG range from 100 to 200 mm.

5.5.3 X-Band Deflecting Cavity

The X-band deflecting cavity, or TCAV, is located just upstream of the SYAG spectrometer. Unlike the linac cavities, which use TM modes to accelerate the beam, the TCAV supports a traveling wave HEM mode that deflects the beam. The TCAV is used as a vertical streaking device; electrons with different longitudinal positions z in a bunch experience different transverse deflecting fields. The particles are displaced vertically according to

$$\Delta Y = R_{34} \frac{eV_{\text{RF}}}{E_0} \sin kz \quad (5.13)$$

and R_{34} is the matrix element between the TCAV and the location where the streaked beam is imaged on a screen. β_y is at a maximum at the TCAV so that the kick due to the TCAV dominates over the natural beam divergence. The TCAV phase is calibrated such that the position $z = 0$ in the beam corresponds to the zero crossing of the deflecting wave. The phase is kept close to zero so that the correlation between y and z remains linear and to avoid deflecting the beam into the beam pipe. A $y - z$ calibration is performed by shifting the phase of the TCAV relative to the arrival time of the beam in small increments and measuring the net deflection [97]. The left panel of Figure 5.14

Figure 5.14: Left: Photograph of the x-band deflecting cavity. The SYAG wiggler magnets can be seen in the background (light blue). Right: Illustration of the principle of the deflecting cavity. At FACET, the cavity deflects vertically as opposed to horizontally. Reproduced from reference [94].

shows the principle of operation.

The TCAV was a critical diagnostic for measuring the bunch length of the ultrashort beam and the bunch separation when operating in two-bunch mode [45]. For the majority of the FACET program, the TCAV was the only diagnostic that has been used experimentally to resolve femtosecond beam structures. However, the TCAV is a destructive diagnostic, so it is not used while taking experimental data with plasma.

5.5.4 Electro-Optic Sampling

Recently, a new diagnostic technique called electro-optical sampling (EOS) has been used to measure the longitudinal beam profile in a non-destructive manner at FACET. The EOS system uses a birefringent GaP crystal that responds to the intense beam fields as the beam passes near to the crystal [98]. We use a linearly polarized pulse from the FACET laser (see Section 5.6) to illuminate the crystal. The laser pulse passes through the crystal and then a polarizer before being imaged onto a CCD camera. In the absence of a particle beam, the polarizer is set such that no light from the laser is transmitted to the camera. When the particle beam is present, the strong beam fields induce birefringence in the GaP crystal which rotates the polarization of the laser pulse passing through the crystal, making the beam visible on the CCD camera.

The EOS crystal is positioned at angle above the path of the beam. The geometry of the system is such that the temporal information relating to the passage of the beam is encoded spatially on the laser pulse as it passes through the crystal. Figure 5.15 shows the EOS signal for a two-bunch distribution. The resolution of the EOS system is sufficient to determine the peak-to-peak separation of the beams, but not high enough to measure the lengths of the individual bunches.

Figure 5.15: Sample EOS image showing a two-bunch beam profile with a peak-to-peak separation of 160 μm . The black line is the projection of the image.

5.5.5 Transverse Profile Monitors

There are a number of transverse profile monitors throughout the experimental area. The profile monitors are used to measure the transverse beam profile, or longitudinal profile if the TCAV is active. Most of these profile monitors are thin metal foils that generate optical transition radiation (OTR) as the beam passes through them. The foils are inserted into the beamline at an angle of 45 degrees with respect to the beam trajectory, acting as a mirror that sends the OTR light out of the beam pipe through a small window. The OTR foils also reflect low intensity laser light and it is possible to simultaneously image reflected laser light and OTR light from the foil so that the particle beam can be aligned to the laser. High intensity laser light, or a highly compressed beam, can damage or destroy the foils, a common occurrence at FACET.

Another type of profile monitor is a thin YAG crystal, like the one used in the energy spectrometer. When used as a profile monitor, the YAG crystal intercepts the beam directly and emits green light. The crystal is only 20 microns thick and has minimal impact on beam quality when used downstream of the plasma interaction. It is particularly useful when imaging low-charge beams, such as those used in the hollow channel plasma experiment.

5.5.6 Dumpline Spectrometer

The final diagnostic at FACET is the dumpline spectrometer. As discussed in Section 5.3.6, the dumpline spectrometer is an imaging spectrometer. The bend magnet deflects the beam vertically

and the vertical β -function is minimized at the dump diagnostics in order to provide high resolution imaging of the beam energy spectrum. This is analogous to the situation at the Sector 20 collimators and SYAG spectrometer, where the energy contrast is maximized in the transverse plane. The dispersion at the dump table is about 60 mm for the nominal 20.35 GeV setting of the bend magnet.

Three different camera systems are used to measure the final energy spectrum. The first camera images the beam on a Lanex screen immediately after the stainless steel dump line exit window. The exit window is relatively thick (5 mm aluminum) and scatters the beam, so the Lanex screen at this location provides the best opportunity to measure the beam's emittance. In addition, the camera at this location has a small field of view, and therefore measures the energy spectrum with a high resolution. This is the most important camera for the E225 Hollow Channel Plasma Wakefield experiment, because it can be used to measure small changes in the beam energy.

Downstream of the Lanex screen, a high-sensitivity DRZ screen intercepts the beam. The camera imaging the DRZ screen has a large field of view and can image particles with energies down to 1 GeV. This camera is particularly useful for the ionization injection experiments, where small amounts of charge are accelerated from zero energy up to a few GeV.

Finally, a set of silicon wafers are used to capture Cherenkov light created by the beam as it passes through air [99]. The Cherenkov system is the most robust measurement of the full beam energy spectrum, and it has the added advantage of being insensitive to synchrotron radiation produced by the beam in the spectrometer dipole. The scientific CMOS camera used to image the Cherenkov light and the GigE camera used to image the SYAG spectrometer are the only two cameras at FACET that were in continuous operation from 2013 to 2016.

5.6 FACET Laser

The FACET laser is a 10 terawatt-class system based on titanium-sapphire (Ti:sapph) amplifying medium operated at a central wavelength of 800 nm with a 60 nm bandwidth [67]. The laser oscillator is locked in phase with the reference RF from the FACET linac. Laser pulses from the mode-locked oscillator are sent to a regenerative amplifier (regen) at 120 Hz. The pulses are amplified up to 1 mJ and are stretched before being sent to the pre-amplifier. Prior to being stretched, the pulse width is shorter than 35 fs. The process of stretching the pulse before amplification and compressing the pulse after amplification is referred to as chirped-pulse amplification (CPA).

The pre-amplifier (pre-amp) and main-amplifier (main-amp) are both 4-pass Ti:sapph amplifiers pumped by 532 nm YAG lasers. The pre-amp uses an 8 mm diameter crystal to raise the pulse energy to 30 mJ and the main-amp uses a 20 mm diameter crystal to raise the pulse energy to up to 800 mJ. The main amplifier is normally operated with 500-600 mJ/pulse to avoid damaging optics in the laser transport system. A schematic of the FACET laser is shown in the left panel of Figure 5.16.

Figure 5.16: Left: Schematic of the FACET laser showing a) Coherent Vitara oscillator, b) Coherent Legend regen, c) 4-pass pre-amp, d) 4-pass main-amp with Thales Saga green pump lasers e) vacuum relay imaging system and f) the periscope/telescope beam expander leading to the laser transport. Right: Schematic of the laser transport from the FACET laser room to the IP integration box in the tunnel.

The laser is imaged from the output of the main amplifier to a periscope/telescope where it enters the FACET laser transport. The transport is 28 meters long and directs the laser from the FACET laser room in the gallery to the tunnel. It has 6 bends and a relay imaging system between bend 3 in the gallery and bend 4 in the tunnel. The transport is in vacuum from bends 1 to 4 and in air between bends 4 and 5 and bend 6 and the compressor box. The beam diameter is roughly 40 mm after bend 4 and through to the compressor. The schematic of the FACET laser transport is shown in the right panel of Figure 5.16.

The compressor box contains a set of large gold diffraction gratings for two-pass compression. The stretched pulse reflects off both gratings before hitting a retro-reflector and taking a second pass through the system. The first order grating efficiency is better than 90% for the horizontally polarized laser and the total efficiency of the compressor is about 75%. The gratings are mounted on a translation stage that is used to vary their separation in order to optimize the final pulse length. A minimum compressed pulse length of 56 fs was measured using a single-shot autocorrelator for a low intensity pulse.

Figure 5.17(a) shows the compressor and the path of the laser. The laser exits the compressor into the interaction point (IP) integration box. A window separates the compressor vacuum volume from the IP area plasma volume. In 2014 and 2015, the window was anti-reflective (AR) coated fused silica, which has a large nonlinear index of refraction. The compressed pulse experienced strong self-phase modulation (SPM) as it was transmitted through the window, which limited operation to

Figure 5.17: a) Image of laser compressor. One of the large ($12\text{ cm} \times 14\text{ cm}$ gold compressor gratings is seen facing the camera in the center of the image. The input and output of the laser is depicted by the red arrow. b) Image of the IP box setup in 2014 showing the path of the positron beam (blue), laser (red) and their co-linear trajectory (purple) after the holed, gold folding mirror. A white arrow shows the motion of the kinoform holder for the raster scan. c) Image of the PB box in 2016 showing the path of the positron beam (blue), laser (red) and their co-linear trajectory (purple) after the holed, gold folding mirror. After the laser passes through the kinoform it is sent to a delay stage (not shown) that is used to vary the length of the plasma channel.

low intensity. The AR-coated fused silica window was later replaced with an uncoated, 3 mm-thick CaF_2 window with a much lower nonlinear index.

After exiting the compressor, the laser enters the IP box and is focused by the kinoform (or axicon) optic. The optic is mounted on a 2D translation stage that can be actuated in the horizontal and vertical planes to change the position of the line-focus. A translation of the lens results in a translation of the line-focus, but does not change the angle of the focus. The laser is coupled on to the beam axis by a holed-gold folding mirror, as shown in Figure 5.17b). The positron (or electron) beam has a transverse size of about $100\text{ }\mu\text{m}$ at this location and easily passes through the 4 mm wide hole in the mirror.

In 2015, a new, larger IP box was installed at FACET. This box is referred to as the “Picnic Basket”, or PB chamber. The PB chamber adds an additional level of control to the system with a 30 cm long delay stage. The delay stage is used to vary the longitudinal positron of the line focus, and therefore control where the plasma starts and ends. The setup in the PB chamber for the 2016 run is shown in Figure 5.17c), but the delay stage is not captured in the image.

5.7 Data Acquisition System

The Experimental Physics and Industrial Control System (EPICS) is used to monitor and control the diagnostics in Sector 20 [100]. The diagnostics fall into two categories. The first type of diagnostics produces scalar data. The most common examples are BPMs, which record the beam position, toroids, which record the beam charge, and pyrometers, which are sensitive to the beam current. These devices have a local Input/Output Controller (IOC) that records and exports the data to the FACET Controls Network via the Channel Access (CA) system. Timing information and triggers

Figure 5.18: Diagram of the `facet_daq` network architecture in Sector 20. Three different flavors of the diagnostic IOCs are shown. In practice, there are 10 scalar data IOCs, 4 sCMOS IOCs and 4 GigE IOCs involved in the DAQ. Note that the NAS is bonded to the switch and mounted on each of the machines running the camera IOCs, so it acts as a local hard drive with read/write rates greater than the network capacity.

are provided by Event Receiver (EVR) hardware in the IOCs, which in turn receive their timing information from the Sector 20 Event Generator (EVG). The EVG is a SLAC-specific device that is used to bridge the gap between the legacy SLAC timing system called the Master Pattern Generator (MPG) and the relatively new timing systems implemented in EPICS. Each linac pulse has an identifier, or Pulse ID, which is embedded in the data.

The second set of diagnostics are cameras that image the particle beam and laser. FACET operates over 30 cameras in Sector 20. This places a heavy load on the FACET network because each camera image is 1-5 MB and the beam rate is 10 Hz, resulting in a maximum data rate of 5 Gbps, which is in excess of the 1 Gbps network capacity. During beam tuning, the image data only crosses the network if requested via CA, so the data rate is typically well below 1 Gbps assuming the operators query only a few cameras at a rate of 1 Hz or less. However, during the experiment we require synchronized data acquisition (DAQ) across all cameras, and it is not possible to transmit this amount of data via the network.

At the start of FACET, there was no DAQ system that satisfied the requirements of the E200 experiment, so we created a solution called ‘`facet_daq`’. `facet_daq` originally referred to a software package but has since come to mean the entire FACET data acquisition system for the PWFA experiments. The `facet_daq` required a re-thinking of how data is acquired when many cameras are simultaneously recording data. The core principles are:

- Local data storage
- Standard camera interfaces
- Beam synchronous acquisition (BSA)

In order to implement our system, we installed a local server in the Sector 20 gallery that would be used to control and monitor the `facet_daq`. We added Network-Attached Storage (NAS) drives that acted as an external hard drive to the machines running the camera IOCs, and we requisitioned EPICS `areaDetector` software for all of our camera IOCs. `areaDetector` is a set of camera drivers that provide a standard interface irrespective of the underlying hardware [101].

The beam synchronous image acquisition system implemented at FACET enabled a wide array of experimental results which would not have been possible otherwise. An excellent example of this is the energy correction techniques described in Section 6.4.2, which used synchronous measurements from the BPMs and the images from the upstream and downstream energy spectrometers. The lessons learned in designing the `facet_daq` will be applied to a new data acquisition system at FACET-II.

Chapter 6

Experimental Results

In this chapter, we discuss the preparation, execution, and analysis of the T504/E225 Hollow Channel Plasma Wakefield Experiment conducted at FACET. The experiments took place over three separate runs from 2014 to 2016. The first run in 2014 utilized a lithium vapor plasma source and a single-bunch positron drive beam. The second run in 2015 utilized a hydrogen plasma source with an electron drive beam. Field-ionization of hydrogen requires twenty times more laser intensity as compared to the lithium case. For the 2015 experimental run, we were unable to achieve this intensity without inducing significant self-phase modulation in the transmitting optics and the experiment was unsuccessful. In the final run in 2016, we again utilized the lithium vapor plasma source with a two-bunch positron drive and witness beam, and we demonstrated for the first time the acceleration of a positron beam in a hollow channel plasma wakefield.

6.1 Experimental Preparation

The experimental setup is sketched in Figure 6.1. It shows the critical components and diagnostics used in the E225 experiment. FACET is the first facility to combine a plasma source, terawatt-class laser, and positron beam, all of which are necessary ingredients for the hollow channel plasma wakefield experiment.

6.1.1 Choice of Kinoform

A number of kinoforms were fabricated for experiments at FACET. The smallest diameter kinoforms used an $m = 4$ spiral to produce a high-order Bessel focus with a peak-to-peak diameter of $150\text{ }\mu\text{m}$. At the other end of the spectrum are the $m = 8$ kinoforms with a $600\text{ }\mu\text{m}$ diameter. The wakefield amplitude scales roughly as the second power of the channel radius (at fixed plasma density), so the $m = 4$ kinoform can produce accelerating gradients that are 16 times larger than the $m = 8$ case for

Figure 6.1: Schematic of the experimental setup. The laser passes through the kinoform and is coupled to the beam axis by a gold mirror with a small central hole. The top right inset shows the laser profile upstream of the lithium oven. A scintillating YAG screen 1.95 m downstream of plasma is used to measure the positron beam profile. A scintillating Lanex screen downstream of the dipole measures the beam energy spectrum. Reproduced from reference [30].

an equivalent drive beam. However, the narrow-diameter channel also requires a low-emittance beam that can be focused through the channel with a β^* that is longer than the channel itself. Taking all of these factors into consideration, we opted to use the $m = 7$ kinoform with a peak-to-peak diameter of $500 \mu\text{m}$. The expected inner diameter of the plasma channel is about $400 \mu\text{m}$.

6.1.2 Preparation of the Laser

The laser was operated with a small fraction of the available pulse energy for generating the hollow plasma channel. Operating at low energy is possible when we use lithium as our plasma source, and it significantly reduces the risk of damaging our optics and other experimental hardware. In addition, we avoid deleterious effects due to self-phase modulation that occur when the laser pulse passes through a transmitting optics, such as the window between the compressor and IP box and the kinoform itself.

Figure 6.2: a) An example of an astigmatic high-order Bessel profile. b) The laser profile after correcting the astigmatism. c) An example of a poorly illuminated laser profile. d) An optimized profile after correcting for astigmatism and illumination.

During the 2014 run, we sent approximately 34 mJ per pulse through the transport to the kinoform. We used a binary-type kinoform which had 46% diffraction efficiency so the focused laser energy into the plasma was about 15 mJ. In 2016, we used an 8-layer kinoform optic with 96% diffraction efficiency and we typically sent 10-15 mJ through to the kinoform.

We spend a significant amount of time optimizing the high-order Bessel laser profile before the start of the experiment. Using very low laser pulse energy (< 1 mJ), we illuminate the kinoform and send the focused light through a window to a camera outside the IP box called the “Rail Camera”. Figure 6.2(a) and (c) show typical examples of the laser profile prior to correction. The aberration in Figure 6.2(a) is primarily due to astigmatic focusing in the laser transport system and is compensated by tilting the lens at the top of the periscope/telescope (labeled ‘f’ in Figure 5.16). Figure 6.2(b) shows the laser profile after correcting for the astigmatism. Another common aberration is shown in Figure 6.2(c). The incomplete ring is caused by uneven illumination of the kinoform. This can

Figure 6.3: Left: Laser intensity profile for different longitudinal positions of PB delay stage, referred to as a rail measurement. At each location, an image of laser profile is recorded and a vertical lineout is taken from the image. The vertical lineouts are stacked together to create the rail measurement. Right: A simulated rail measurement using the `FACET OPTICS` code.

be corrected by shifting the trajectory of the laser onto the kinoform, while keeping the angle of the incident laser constant. We typically steer the laser in the main-amp as well in order to change the nearfield profile and the intensity distribution on the kinoform. Figure 6.2(d) shows the high-order Bessel profile after the astigmatism and illumination have been corrected.

In 2016, we modified this procedure slightly by viewing the laser profile on an OTR foil just upstream of the plasma, which improved the accuracy of our corrections. We also installed a delay stage inside the PB chamber that allowed us to shift the end of the line-focus upstream and downstream from the center of the heat-pipe oven. The kinoform was held in mount with a 32 mm diameter circular steel mask that was used as an aperture for the laser. The mask also determines the end of the line focus. The convergence angle for the $m = 7$ optic is 4.4 mrad, so the 32 mm mask ends the line-focus 3.6 meters downstream of the optic.

We use the PB chamber delay stage to measure the intensity of the high-order Bessel profile as a function of z , the distance from the optic. At each value of z , we record the transverse laser profile and take a vertical lineout from the center of the profile. The lineouts are stacked together to form the rail measurement. Figure 6.3 shows excellent agreement between the measured and predicted rail profile. The laser intensity varies slowly until the end of the focus, and abruptly ends at a z -location of 350 cm.

6.1.3 Preparation of the Lithium Heat Pipe Oven

For the 2014 run, we operated with a 1.3 meter long lithium heat-pipe oven at a density of $8 \times 10^{16} \text{ cm}^{-3}$. The laser was operated at low intensity, and because the intensity of line-focus increases

linearly with the distance from the optic, only the downstream end of the lithium vapor was ionized. For this data, we had to use transverse kick measurements to determine the length of the ionized channel (see Section 6.3.2).

For the 2016 run, we operated the same 1.3 meter long oven, but we used only the central heater coil so that the length of the vapor region was 50 cm. In addition, we used a delay stage in the PB chamber to move the end of the line-focus along the length of the vapor column. This allowed us to directly control the length of the plasma. Unless otherwise noted, the 2016 data was acquired with a 25 cm long plasma channel with a density of $3 \times 10^{16} \text{ cm}^{-3}$.

6.1.4 Preparation of the Positron Beam

For the hollow channel plasma experiment to work, it is critical that the beam itself does not generate plasma as it propagates through the lithium vapor. The choice of lithium as our plasma medium is a double-edged sword; it is easy for the laser to ionize, but it is also easy for the beam to ionize. We use the ADK method detailed in Section 4.9 to calculate the ionization rate, replacing the laser field with the beam field. The peak field of positron beam must be well below the ionization threshold, which is about 5.3 GV/m for a 100 fs long beam [102]. We also require that the beam is cleanly transmitted through the channel, which means that it needs to be tightly focused, and the beam must be short compared to wavelength of the hollow channel plasma wakefield. In order to satisfy all of these constraints, we use a positron beam with significantly reduced charge compared to nominal operation in uniform plasmas.

For the 2014 experiments, we used a single positron bunch with $N = 5.34 \times 10^9$ and an r.m.s. bunch length of $\sigma_z = 35 \pm 5 \text{ } \mu\text{m}$. The mean energy of the beam was 20.35 GeV with 350 MeV energy spread (FWHM). The beam was brought to a focus near the start of the plasma channel with β_x^* and β_y^* of 50 cm and 75 cm, respectively. The focal spot was circular with $\sigma_x = \sigma_y = 50 \text{ } \mu\text{m}$. β^* was longer than the expected channel length in both dimensions, the transverse spot size remained roughly constant through the channel. The peak electric field was calculated to be 1.6 GV/m, well below the 5.3 GV/m field ionization threshold for lithium.

For the 2016 experiments, we operated with the two-bunch setup with drive charge $N_d = 2 - 3 \times 10^9$ and witness charge $N_w = 0.1 - 1 \times 10^9$. The bunch spacing was varied throughout the experiment to probe the wake structure and find the accelerating phase. We were able to focus the beam to a smaller spot size using larger β -functions, implying that the emittance of the positron beam had improved since 2014. Spot sizes as small as $\sigma_x = \sigma_y = 20 \text{ } \mu\text{m}$ were measured at the wire scanner upstream of the plasma oven for β -functions as large as 1 meter.

Figure 6.4: A mosaic of images illustrating the principle behind the raster scan. Each subimage shows both the positron beam (central hot spot) and laser profile (ring) at low intensity reflecting off of a titanium foil. The laser is scanned in the transverse plane while remaining parallel to the beam trajectory. The arrows indicate the direction of force that the positron beam experiences for a plasma channel located at the position of the laser.

6.2 Determination of Hollow Channel Plasma Structure

As discussed in Section 3.7, a positron beam propagating off-axis in the channel excites a high-order dipole mode that deflects the tail of the beam away from the center of the channel. Intuitively, this is because the positron beam is attracted to the plasma electrons in the wall of the channel. For example, if the beam is offset vertically above the channel axis, it will be kicked towards the top of the channel. If the beam is outside the channel wall, it will experience a kick towards the nearest part of the plasma annulus. We make use of this attraction to probe the shape of the channel by scanning the position of the laser with respect to the beam trajectory and recording the resulting kick. We call this the “raster scan”.

To illustrate this principle, Figure 6.4 shows a mosaic of images recorded on an OTR screen just upstream of the plasma. Each subimage shows the low-intensity laser as the ring and the positron beam as the central hotspot. The kinoform position is scanned by $\pm 500 \mu\text{m}$ both horizontally and vertically with respect to the positron beam vector.

In an ideal raster, the center of the kinoform defines the position of the line-focus, and translating the kinoform stage shifts the line-focus without affecting the angle. In practice, there is laser pointing jitter which affects both the relative angle and position of the line-focus with respect to the positron

Figure 6.5: Left: An image from the upstream camera showing the laser profile and the fit to the central position. Middle: All of the shots from the entire raster scan in the transverse plane. They are binned according to the dotted lines. Right: 2D histogram of the binned shots showing uniform relatively uniform coverage of the scanned plane.

beam trajectory. The pointing jitter of the laser was measured to be $25.4 \mu\text{rad}$ in x and $12.8 \mu\text{rad}$ in y RMS. The channel is formed roughly 3.5 meters downstream of the optic, so the angular jitter results in a transverse displacement of the channel by up to $90 \mu\text{m}$.

The pointing jitter of the laser means that we do not know the position of the laser based on the location of the kinoform translation stage alone. Therefore, we need a diagnostic that records the position of the laser on every shot. Downstream of the plasma, we use a holed-gold folding mirror to couple the spent laser light out of the beam pipe while transmitting the positron beam to the diagnostics section. The setup is analogous to the way the laser is coupled to the beam trajectory in the IP box. The spent laser light exits the beam pipe through an AR-coated fused silica window and passes through a lens and then through a set of three beam splitters. The beam splitters send the light to three cameras, which image the laser profile at the upstream, middle, and downstream end of the heat-pipe oven. These cameras are used to track the laser throughout the raster scan and allows us to recover the relative position of the laser and channel on a shot by shot basis. We developed an algorithm to find the center of the laser profile on each shot. An example output of the algorithm is shown in the left image of Figure 6.5. During the raster scan, the kick data are binned according to the position of the laser. The middle and right images of Figure 6.5 show all recorded laser positions over the course of a scan and their spatial binning, respectively.

6.2.1 Raster Scan Data from 2014

We begin the raster scan by aligning the laser to the beam trajectory. We compared the laser position to the beam position on two Optical Transition Radiation (OTR) foils upstream and downstream of the heat-pipe oven, separated by 1.84 meters. We then scan a region of interest around the aligned orbit and record the kick delivered to the beam on a scintillating YAG screen 1.95 meters downstream of the interaction region.

We observed an average kick in the $+x$, $+y$ direction of $43.6 \mu\text{rad}$ and $48.7 \mu\text{rad}$, respectively,

Figure 6.6: Left: The kick map shows the magnitude and direction of the kick delivered to the beam averaged over approximately 5 shots as the channel location is scanned with respect to the beam trajectory. A net kick of $(43.6, 48.7)$ μrad in (x, y) is subtracted from the data. We superimpose two equal intensity contours of the central Bessel peak of the laser measured upstream of the plasma as red dashed lines. Right: Area of the positron beam measured on a YAG screen downstream of the plasma averaged over 10 shots as the channel location is scanned with respect to the beam trajectory.

during the scan. We believe that this was due to an alignment error that occurred when the lithium heat-pipe oven was actuated into the beamline. The resulting vector field, or “kick map”, is shown in the left image of Figure 6.6, after subtracting off the average kick. The vectors represent the magnitude and direction of the kick in μrad averaged over approximately 5 shots for a given relative position of the laser with respect to the beam. In the experiment, the location of the laser is varied with respect to the beam trajectory, but for clarity we plot the kick vectors as though the laser is fixed and the beam trajectory is changed.

The kick map is consistent with our expectations for an annular ionized region. When the beam is outside the channel, the kicks point radially inwards toward the ionized region. When the beam is inside the channel, the kicks point radially outward toward the walls of the channel. We observe an annular region at a radius of $250 \mu\text{m}$ where the kick felt by the beam is minimized. In this region, the beam is propagating nearly parallel to and within the ionized channel wall, and therefore does not experience a net kick. The right image in Figure 6.6 shows the change in the beam area as measured on the same YAG screen. The transverse wakefield increases in strength longitudinally throughout the bunch. Particles towards the back of the bunch see a larger kick than particles in the front. The beam profile stretches in the direction of the kick, and the beam area increases as a result. The largest beam area growth occurs when the beam propagates along or through the ionized annulus. Both the kick map and beam area map show an annular feature with some asymmetry, which may be attributed to variations in the laser intensity around the first maximum of the Bessel profile, resulting in uneven ionization of the plasma annulus.

Figure 6.7: Left: The expected kicks for an ideal $J_7^2(k_\perp r)$ profile through the first minimum in r . Right: Results of our optimized 2016 raster scan. The shape of our measurements matches what is expected in the ideal case. However, the measurement is much “blurrier”, indicating the ionized region was larger than the maximum of the $J_7^2(k_\perp r)$, which is expected because the ionization process is a threshold effect. The non-zero beam size also contributes to the blurred measurement.

6.2.2 Channel Reconstruction Algorithm

The process of collecting and analyzing the thousands of shots needed for the raster scan measurement has improved over time. In 2014, we required hours to perform a single iteration of the raster, while in 2016 we were able to record and analyze the data in under ten minutes. This allowed us to perform scans and laser corrections iteratively to produce a hollow channel plasma with a nearly ideal shape.

Since the plasma channel attracts the beam, we can think of the channel as a potential well, and the kicks as the derivative of the potential. We can then use the kick map to reconstruct the shape of the plasma channel.

We start by assuming that the ionized region has the form $\phi \propto J_7^2(k_\perp r)$, and the transverse kicks are given by $\Delta_{x,y} = \partial_{x,y}\phi$. The amplitude of the total kick is therefore

$$\Delta = \sqrt{\partial_x \phi^2 + \partial_y \phi^2}. \quad (6.1)$$

The left images of Figure 6.7 show the expected kicks due to the ideal potential $\phi \propto J_7^2(k_\perp r)$. In practice, the ionized region extends beyond the maximum of the Bessel function. The right images of Figure 6.7 show the measured kicks. The blurring is due in part to the extended plasma and non-zero beam size. However, the shape of the measured kicks matches our expectations from the ideal case.

Figure 6.8: Left: The measured kick map for the optimized single bunch raster scan in 2016. Right: The reconstructed channel profile for this data.

We reconstruct the plasma shape by integrating the kicks over the scan according to

$$\phi_{\text{meas}} = \int \Delta_{\text{meas}} \cdot d\vec{r}. \quad (6.2)$$

Figure 6.8 shows the result of the reconstruction. The ionized region is not a perfect annulus, even after our extensive optimization efforts. Nevertheless, the subsequent energy loss measurements show that the annular plasma channel still supports longitudinal fields with the same magnitude as one would expect from a radially uniform channel.

6.3 Single-Bunch Energy Loss Measurements

After we complete the raster scan, we use the kick map data to align the laser and positron beam trajectories within the error of the laser jitter. The quadrupoles in the dumpline spectrometer are set to image the beam waist, located near the start of the plasma channel in the heat-pipe oven, to the LANEX screen, 9.6 meters downstream of the spectrometer dipole.

6.3.1 Results from 2014

We measure the beam energy spectrum and centroid on every shot at both the upstream SYAG spectrometer and downstream dumpline spectrometer. We measured an 18.1 MeV r.m.s. incoming energy jitter from the linac and we used this information to correct the energy measurement on the downstream spectrometer. We took data at a rate of 10 Hz, with the laser firing at a rate of 9 Hz. There should be no plasma in the laser-off case, so we can use this data to calibrate our spectrometers. Figure 6.9 shows the incoming and outgoing energy centroids.

Figure 6.9: The incoming energy, measured at the SYAG spectrometer, versus the outgoing energy, measured in the dumpline spectrometer. The laser-on shots lose an average of 20 MeV energy compared to the laser off case.

The energy resolution of the dumpline spectrometer is 3.1 MeV when imaging is set to focus the 20.35 GeV particles. Additional systematic errors in the energy measurement may arise from orbital effects. In particular, we checked for correlations between the energy loss and vertical kicks from the plasma channel, but none were observed. The observed kicks from the plasma have a maximum amplitude of $30 \mu\text{rad}$, producing a vertical displacement of $3.3 \mu\text{m}$, or equivalently 1 MeV at the LANEX screen, which is too small to measure.

The left image of Figure 6.10 shows all of dumpline energy spectra recorded for this dataset, corrected for incoming energy jitter. There are 315 shots in the dataset, 10% of which are laser off. The middle image shows the average laser-on and laser-off energy spectra, with the local statistical measurement error at 20.15 and 20.55 GeV. The histogram on the right shows the measured energy centroids for all shots, corrected for incoming energy. We measure a mean shift in the centroid beam energy of $18.7 \pm 3.37(\text{stat}) \pm 3.55(\text{syst})$ MeV.

6.3.2 Determination of the Channel Length

During the 2014 experimental run, we did not have direct control over the length of the plasma. The laser was operated at low intensity, and only ionized a plasma channel at the downstream end of the oven. In order to determine the gradient of the hollow channel plasma wakefield, we needed an estimate of the channel length. One way to assess the length of the channel is to use the kick

Figure 6.10: Left: The dumpline energy spectra for all 315 shots in the dataset, 10% of which are laser-off. The energy spectra are corrected for incoming energy jitter. Middle: The average laser-on and laser-off spectra, with error bars showing the local statistical error in the measurement at 20.15 and 20.55 GeV. The error bars have been enhanced by a factor of 5 to make them visible. Right: A histogram of the energy centroids for all shots, corrected for incoming energy. The mean energy loss is 18.9 MeV with a width of 3.2 MeV.

data and compare it to our model for the transverse wakefield. This concept is illustrated in the left image of Figure 6.11.

The growth length for particles in the center of the beam is given by

$$L_g = \frac{1}{2} \sqrt{\frac{\gamma}{Nr_e \sigma_z \hat{W}_{\perp 0}}} = 10.3 \text{ cm} \quad (6.3)$$

for beam energy γ , beam charge N , bunch length σ_z , classical electron radius r_e , normalized wake strength $W_{\perp 0}$. This equation differs from equation 3.62 by a factor of two because we are considering particles in the center of the bunch as opposed to the tail. For the 2014 data, we estimated that the channel would be about 10 centimeters long, given estimates from ADK theory based on the laser power. From equation 3.61, we find that the drift trajectory of the beam center is given by

$$y(s) = \frac{y_0}{2} \left[I_0 \left(\sqrt{\frac{s}{L_g}} \right) + J_0 \left(\sqrt{\frac{s}{L_g}} \right) \right] \quad (6.4)$$

where s is the propagation distance, y_0 is the initial transverse offset, and I_0 and J_0 are zeroeth order Bessel functions. Note that the growth of this instability is relatively slow. For an initial offset and angle of $5 \mu\text{m}$ and $20 \mu\text{rad}$ respectively (the RMS beam orbit jitter and laser pointing jitter measured at FACET), the beam can propagate for over a meter in the channel before the tail of the bunch reaches the channel wall.

We measured the vertical kick on the downstream YAG screen, and correlated the kick with the position of the laser as measured by our laser imaging system. We fit the result and compared it to the derivative of Eq. 6.4 evaluated at $s = L_c$, where L_c is the length of the channel. The measured kick relative to the incoming beam orbit is $0.25 \mu\text{rad}$ per μm offset, corresponding to a channel length of $L_c = 8.3^{+1.3}_{-1.0} \text{ cm}$, where the dominant contribution to the uncertainty comes from

Figure 6.11: Left: Schematic of the kick measurement showing how beam offset is translated into kick. Right: The result of the kick measurement, where the beam displacement of the downstream YAG screen has been converted into an angle based on the distance of the screen from the plasma channel. The offset of the channel is measured by the laser imaging system.

the uncertainty in bunch length. Note that the beam drift due to an initial angle y'_0 also contributes to the kick felt by the beam, but can be neglected if $y_0 \gg y'_0 L_c$, as is the case here. Therefore, the gradient we measure is $227^{+32}_{-30} \text{ MeV m}^{-1}$. This value is in excellent agreement with the theory and simulation shown in Figure 6.12.

6.3.3 Energy Loss Data from 2016

In the 2014 data, we were able to make an accurate estimate of the decelerating wakefield using our transverse kick model, but the uncertainty on the measurement was fairly large because we did not have a precise measurement of the bunch length at the time. We were able to improve on our result in 2016 by controlling the length of the channel and measuring the energy loss as a function of channel length.

As we mentioned in Section 6.1.2, we controlled the end of the kinoform line-focus with a 32 mm aperture mask, and measured the end of the focus to be 3.5 meters downstream of the optic. Next, we limited the length of the lithium vapor column to 50 cm by operating the heat-pipe oven with only the central heater. Finally, we used the PB chamber delay stage to change the path length between the kinoform and oven, so that the end of the focus could be translated through the lithium vapor. Note that for this measurement, the lithium vapor density was $n = 1 \times 10^{17} \text{ cm}^{-3}$, slightly higher than the density used in the 2014 experiment.

Figure 6.12 shows the results of the plasma length scan. The plasma length is changed in increments of 5 cm and the corrected energy loss of the centroid is recorded at each point. Note that for the longer channel lengths, relatively few shots exit the channel cleanly. Those shots that

Figure 6.12: Left: Change centroid energy versus channel length. The energy centroid is corrected for incoming energy jitter on a shot-by-shot basis. Shots that do not cleanly propagate through the channel are removed from the data. The energy loss for a channel length of 45 cm is 107.7 MeV, or an effective gradient of 239 MeV/m. The data was taken with a plasma density of $n = 1 \times 10^{17} \text{ cm}^{-3}$. Right: QuickPIC simulation of the hollow plasma channel using parameters from the 2014 experiment with $n = 8 \times 10^{16} \text{ cm}^{-3}$. The beam propagates to the left and the 1σ contour is shown by the white dashed line. The black dotted lines at $\pm 240 \mu\text{m}$ and $\pm 290 \mu\text{m}$ are the inner and outer radii of the plasma channel, respectively. Lineouts of the simulated and calculated on-axis E_z field ($r = 0 \mu\text{m}, z$) and the radial variation in E_z at the peak decelerating field ($r, z = 11 \mu\text{m}$), are shown with solid and dashed black lines, respectively. The maximum decelerating field is 220 MeV/m.

are deflected show substantially more energy loss (200-300 MeV) and are removed from the data. As a result, the error on the measurement increases with increasing channel length. The energy loss measured for a channel length of 45 cm is $107.7 \pm 35 \text{ MeV}$, which gives an effective gradient of $239 \pm 78 \text{ MeV/m}$, in good agreement with both the 2014 result and the QuickPIC simulation shown on the right of Figure 6.12.

6.3.4 Discussion

In the field of laser wakefield acceleration (LWFA), the use of shallow plasma channels for laser guiding has become a common technique for extending the interaction length [10]. The hollow channel concept is different in that the center of the channel is devoid of plasma. One critical aspect of the experiment is to ensure that the effects we have measured are not due to residual, on-axis, low-density plasma. First, we note that the measured energy change has a distinct on-off effect related to the presence of the laser. From this we infer that the positron beam is not creating a plasma in absence of the laser. Second, we have measured the contrast of the laser on the bench to be better than 60:1 for binary kinoform used in 2014. In the 2016 experiment, we used an 8-layer kinoform which improved the contrast by at least a factor of 2. The ADK ionization rate [72] for laser pulses below this intensity level is at least 100 times too small to ionize lithium. Finally, the transverse profile of the positron beam as measured on the downstream YAG profile monitor shows

Figure 6.13: Top: QuickPIC simulation of the hollow channel plasma wakefield with a positron drive and witness beam using parameters from the experiment. The drive-witness separation is $160 \mu\text{m}$. Bottom: On-axis longitudinal field and bunch profiles from the simulation. The expected accelerating wakefield at the location of the witness bunch is 135 MeV/m .

no change when the beam is propagating down the center of the channel, compared to when there is no plasma channel present. Previous work on positron beam propagation through uniform low density ($n_0 \approx 10^{13} \text{ cm}^{-3}$) plasma columns has shown the formation of a halo of charge surrounding a central intense spot due to transversely nonlinear focusing forces acting on different slices of the beam [40, 41]. No such halo formation was observed here. This shows that there are no measurable focusing forces on the positron bunch, as expected from propagation in a hollow plasma channel. Finally, we note that the inferred gradient of 227 MeV/m in the 2014 data and 239 MeV/m in the 2016 data is in excellent agreement with both our theoretical prediction and simulation result of 220 MeV/m , as seen in Fig. 6.12.

Figure 6.14: Left: Waterfall plot of the EOS signal over the course of the phase ramp scan. The drive beam is the bright stripe on the right of the image and the witness beam extends out to the left. Right: Centroid-to-centroid bunch separation measured over the course of the scan.

6.4 Observation of Energy Gain with a Two-Bunch Setup

We conclude this chapter with the recent, exciting observation of the acceleration of a positron witness bunch in a hollow channel plasma wakefield in 2016. For this data, we operated the heat-pipe oven in the 50 cm mode with a lithium vapor density of $3 \times 10^{16} \text{ cm}^{-3}$. The PB delay stage was set such that the laser ionized a channel 25 cm long. Note that we attempted the experiment at longer channel lengths, but the transverse wakefield, which increases at larger separations of the drive and witness bunch, was too strong and in general the witness beam was not cleanly transmitted through the channel when it was longer than 25 cm.

Prior to the experiment, we ran a number of QuickPIC simulations to understand the beam setup that optimally loads the hollow channel wake. For a plasma density of $2.5 \times 10^{16} \text{ cm}^{-3}$ with the inner channel radius located at $r = 200 \text{ μm}$, we found that a drive beam with 3×10^9 particles and witness beam with 1.5×10^9 particles produces a relatively flat wake when the centroid-to-centroid bunch separation is 160 μm . The results of this simulation are shown in Figure 6.13. The expected accelerating gradient for the witness beam is 135 MeV/m .

6.4.1 Finding the Accelerating Phase

We began the experiment using 5:2 drive:witness charge ratio with a bunch separation of 125 μm . Our choice of bunch separation was motivated both by the QuickPIC analysis described above and the fact that we were operating with a plasma density of $3 \times 10^{16} \text{ cm}^{-3}$ rather than $2.5 \times 10^{16} \text{ cm}^{-3}$.

For bunch separations of 125 μm or less, we only observed deceleration. We attempted to increase the bunch separation by scanning the RF phase of the positrons entering the linac (phase ramp)

Figure 6.15: Left: Charge in the drive and witness bunch over the course of the phase ramp scan. Right: Energy gain/loss of the witness and energy loss of the drive bunch over the course of the scan.

in order to change the chirp into Sector 20. We used the EOS diagnostic to monitor the bunch separation during the scan. Figure 6.14 shows the results from the phase ramp scan.

We start to observe acceleration of the witness beam at a bunch separation of 250 μm . Using the EOS diagnostic, we are able to measure the mean energy loss or gain of the witness bunch as a function of bunch separation. The right plot in Figure 6.15 shows that the witness bunch is maximally accelerated at a centroid-to-centroid bunch separation of roughly 400 μm , more than twice the separation we expect from simulation. We will address possible causes for this discrepancy in Section 6.4.4.

During the phase ramp scan, we also adjusted the collimation in Sector 20 to produce well separated drive and witness beams. The collimation affects the relative charge of the drive and witness beams, as shown on the left of Figure 6.15. The charge of the drive beam varied between $2.5 - 3 \times 10^9$ particles per bunch over the scan, and the witness charge varied between $0.5 - 1 \times 10^9$. Variations in beam charge lead to variations in wake strength, so the phase ramp scan is not a one-to-one map of the longitudinal wakefield. We can correct this effect to some extent by scaling the energy loss of the drive beam and energy gain of the witness beam by their average charge during the first step of the scan. The average charge of the drive beam was 2.7×10^9 and the average charge of the witness beam was 9.4×10^8 in the first step of the scan. The result of this correction is shown in Figure 6.16.

The drive beam energy loss should remain constant throughout the scan after correcting for charge, but it decreases by nearly a factor of two at the largest bunch separations. We suspect that this is primarily the result of bunch lengthening in the drive beam as we adjust the chirp and collimation of the two-bunch setup during the scan. The amplitude of the wake is inversely proportional to bunch length, which implies that the length of the drive beam has nearly doubled by the end of the scan. We observe changes in the width of the drive beam signal using the EOS diagnostic over the course of the scan, but the measured widths are not consistent with a simple

Figure 6.16: Change in energy of the drive and witness beams as a function of bunch separation. The energy gain and loss of the beams is scaled by the bunch charge at each step.

doubling of the bunch length. This may be because the EOS crystal does not have sufficient resolution to resolve the bunch profile. Another factor contributing to changes is in the drive beam energy is that at the smallest values of the bunch separation, the drive and witness beam overlap. The presence of additional witness beam charge in the decelerating phase of the wake contributes to the energy loss of all particles in that phase of the wake.

The corrected energy gain curve for the witness bunch crosses from the decelerating to accelerating phase of the wake at a bunch separation of roughly $240 \mu\text{m}$. The energy gain of the witness bunch is surprisingly flat for bunch separations ranging from 325 – $425 \mu\text{m}$. We set the bunch separation to be roughly $375 \mu\text{m}$ for the remainder of the data described in this chapter.

6.4.2 Energy Correction

Energy correction for the two-bunch data is significantly more complicated than the single bunch case. This is because the spectrum measured at the upstream SYAG spectrometer does not show a one-to-one correlation with the dumpline spectrum, even in the laser off case, due to non-linear dispersion. As a result, we are not able to simultaneously correct the drive and witness beam energy using the technique described in Section 6.3.1.

We developed an alternative approach to beam energy correction by recording a library of upstream and downstream energy spectra with the laser off and no plasma present. For shots where

Figure 6.17: Left: Typical result for matching a laser-on spectrum with a laser-off spectrum at the upstream spectrometer. Right: We compare the spectra of the same laser-on and laser-off shots to measure the energy change of the drive and witness bunches.

the laser is on and the beam is affected by the plasma, we compare the upstream spectrum and the beam charge, as measured by the dumpline BPM, to find the closest matching spectrum the library of laser-off shots. We then compare the downstream energy spectrum for the matched shots to infer the energy gain or loss of the beams. We refer to this technique as “lookup-table correction”. Figure 6.17 shows a typical result of the match.

6.4.3 Energy Gain and Beam Loading

Figure 6.18 shows the energy gain of the witness bunch and energy loss of the drive bunch for all shots in the dataset, after performing the look-up table energy correction. The average acceleration of the witness beam is 19.6 MeV over the 25-cm long channel, and the average deceleration of the drive bunch is only 11.5 MeV over the same length. The transformer ratio is 1.7. This is the first demonstration of high transformer ratio PWFA for a drive-witness setup.

The highest energy-gain shot in this dataset is 33.4 MeV, observed for a very low-charge witness beam with 5×10^7 particles, or one twentieth of the drive beam charge. The accelerating gradient for this shot is 133 MeV/m. The average energy gain of 19.6 MeV for the witness bunch corresponds to an accelerating gradient of 78 MeV/m, roughly half the value expected from the QuickPIC simulation shown in Figure 6.13. The average energy loss of 11.5 MeV for the drive bunch corresponds to a decelerating gradient of 46 MeV/m, also smaller than the value of the field predicted from simulation. The bunch length of the drive beam in the experiment might be significantly longer than the value used in the simulation. This could account for some, but not all, of the discrepancy between the expected and measured values of the gradient.

We also observe a very strong beam-loading effect where the acceleration of the witness bunch

Figure 6.18: Left: Energy change of the drive bunch after propagating in a 25 cm long hollow channel. The mean energy loss of the drive beam 11.5 MeV. Left: Energy change of the witness bunch. The mean energy gain of the witness beam is 19.6 MeV.

is reduced by increasing the charge of the witness bunch. Figure 6.19 shows the effect for both the drive and the witness beams. Increasing the drive beam charge increases the strength of the wake which in turn decelerates the drive bunch. Increasing the witness bunch charge loads the wake, and as a result, the net accelerating field is decreased. Note that in this experiment, the charge in the drive and witness beam are not independently tuned (*i.e.* more charge in the drive bunch implies less charge in the witness bunch and vice-versa). We are able to distinguish the beam loading effect from changes in wake amplitude by observing that the energy loss of the driver and the energy gain of the witness change at different rates with the respect to their bunch charge. Assuming that the drive bunch length does not change significantly with bunch charge, we can infer by superposition that the beam loading effect due to the witness beam alone is $-2.5 \text{ MeV}/10^8$ particles.

6.4.4 Discussion

In Section 6.4.1, we found that the bunch separation where we observe acceleration is longer than the value expected from simulations by over a factor of two. In addition, the accelerating phase of the wake is long and flat. These two features of the wake cannot be accounted for with the theoretical treatment provided in Section 3.3.2, even with the uncertainty in measurements of the bunch length of the beams. The vapor density of the lithium oven is known at the level of a few percent, so we do not believe that this is the source of the discrepancy either.

The most likely source of the discrepancy between theory and measurement is that the shape of the channel is significantly different from the sharp-walled channel used in the dielectric waveguide model described in Section 3.3. Our vacuum measurements of the laser profile shown in Figure 6.3, indicate that the channel should indeed have a sharp boundary. However, it is possible that phase

Figure 6.19: Left: The energy loss of the drive beam as a function of charge in the drive beam. The slope is $-1.5 \text{ MeV}/10^8 \text{ particles}$. Right: The energy gain of the witness beam as a function of charge in the witness beam. The slope is $-4.0 \text{ MeV}/10^8 \text{ particles}$. Assuming a linear wakefield and a drive beam with constant bunch length. We can infer that the beam loading effect due to the witness beam alone is $-2.5 \text{ MeV}/10^8 \text{ particles}$.

errors occur in the laser as it propagates through plasma, which could blur the boundaries of the channel. We believe that the plasma channel has a high degree of azimuthal symmetry based on the results of the raster scan measurements presented in Section 6.2.2, but we cannot probe the sharpness of the boundary using this beam based technique.

Figure 6.20 shows a simulation of a beam propagating through a channel where the radial plasma density is described by a gaussian centered at $r = 250 \text{ } \mu\text{m}$. The parameters of the simulation are similar to those found in the experiment. The simulation is performed using the OSIRIS-2D code [52] which simulates azimuthally symmetric beam plasma interactions in $r-z$ space. The beam propagates along the channel axis to the left.

The evolution of the field behind the drive bunch shows many complicated features including a prolonged accelerating phase similar to what we observe in the 2016 data. However, the location of the accelerating phase does not match our experimental results. The radially gaussian plasma profile is one of several hypotheses currently being tested in simulations. This work is ongoing.

6.5 Conclusion

Over the course of two runs in 2014 and 2016, we both demonstrated that we can create meter-scale hollow channel plasmas and use the positron beam as a sensitive diagnostic to explore the channel shape and wakefield. Our raster scan method was a novel technique for probing the channel structure, and this type of measurement has since been used by the E201 experiment to measure

Figure 6.20: Result of an OSIRIS-2D simulation in the $r - z$ plane. Left: The density of the beam and plasma from the simulation. The beam propagates along the channel axis to the left. The plasma is initiated with radially gaussian profile centered at $r = 250 \mu\text{m}$ with $\sigma_r = 20 \mu\text{m}$. The density of the plasma at the peak of the profile is $3 \times 10^{16} \text{ cm}^{-3}$. The beam contains 3×10^9 particles with a bunch length of $40 \mu\text{m}$. Right: The longitudinal field associated with this wake.

dipole wakefields in their dielectric structures.

We performed a suite of measurements to determine the amplitude of the plasma wakefield and we confirmed that the short-range wakefields match our expectations from theory. Finally, our observation of accelerated positrons in the hollow channel wakefield, 20 years since the idea was first proposed, is an important validation of the concept.

Chapter 7

Conclusions

In this thesis, we have demonstrated the acceleration of a positron beam in a hollow channel plasma. Hollow channel plasma acceleration has been a topic of theoretical interest for decades, and we have taken an important step forward by realizing this concept experimentally. But the motivation for this research is to develop a technique that can be used to accelerate positrons in a plasma-based linear collider. How do our results advance us towards this goal?

7.1 Towards a Plasma-Based Linear Collider

As part of the P5/Snowmass process, we performed a design study of a multi-TeV plasma-based linear collider (PLC) [38]. The goal of the study was to understand how staged plasma acceleration can be used to generate very high energy particle beams. Figure 7.1 shows the straw-man design.

In this design, the collider has a 4.5 kilometer-long footprint and achieves a collision energy of 1 TeV. By contrast, the proposed ILC is 31 kilometers-long with half the collision energy. A PLC can reduce the size of the overall facility, remove the need for expensive superconducting cavities, and accelerate particles with a high beam-to-plasma-to-beam efficiency, all of which reduce the construction and operation costs. However, there are no plans to build a PLC because of the many remaining challenges not addressed in our straw-man design.

7.1.1 Challenges Associated with a Hollow Channel PLC

One of the key assumptions of the design presented in 7.1 is that the electron and positron arms are treated equally. Given our current understanding of positron plasma wakefield acceleration, it is not possible to achieve equal operation of the two arms if the beams are undergoing acceleration in a uniform plasma. On the other hand, hollow channel plasmas could be used to accelerate beams of either charge in a symmetric fashion, but this statement comes with some caveats. Symmetric

Figure 7.1: Simplified schematic of a 1 TeV PWFA Linear Collider from reference [38].

operation implies that if the electron arm uses an electron drive and witness beam, then the positron arm should use positrons for both the drive and witness beams. Yet it would be impractical to use a positron drive beam to accelerate a positron witness beam because of the large amount of energy required to create positrons. If we choose instead to use an electron drive beam in the positron arm, the relative phase of the drive and witness beams will differ by 180° between the two arms. For hollow channel wakes excited in the linear regime (no charge separation at the boundary), this phase shift is trivial. But if the wake is nonlinear, or the hollow channel plasma does not have a well defined boundary as we saw in the previous chapter, the longitudinal shape of the wake may evolve significantly over the first wake period. This will effect the usable accelerating gradient as well as the shape and charge of the positron witness beam.

The second caveat is related to the beam breakup instability. Without the use of external focusing elements, the growth length of this instability is much shorter than the length of the accelerator. If external focusing elements are used, they must be as strong or stronger than the transverse wakefield, which may be as strong as a mega-Tesla per meter. No such focusing elements exist, and if they did it is unclear what effect they would have on the wake itself.

The final caveat worth mentioning is that the beams used in linear colliders have peak densities in excess of 10^{21} cm^{-3} . The fields near these beams are strong enough to ionize any neutral gas in their vicinity. Therefore, a hollow channel plasma, as implemented in the context of a linear collider, would need to be a vacuum along the axis of beam propagation. It is not clear how to create this type of plasma structure.

Despite these challenges, there is general excitement in the field regarding the use of hollow channel plasmas for accelerating positrons. Short of anti-plasma, it may be the best solution available.

7.1.2 Other Opportunities for Plasma Elements in a Linear Collider

In addition to using plasmas as a replacement for RF accelerating cavities, plasmas may one day replace the injectors, damping rings and final focusing systems in a linear collider. These systems account for a significant fraction of the infrastructure in traditional linear collider designs and they account for most of the length of the PLC straw-man design.

At FACET, we successfully demonstrated the injection of low emittance beams in a plasma [103]. This technology could be used to replace the injectors and damping rings at a linear collider. For the ILC, the total cost of these systems is on the order of \$1 billion. While we have shown that this concept works for generating electron beams, there is no analogous solution for positron beams.

Plasma lenses are another technology that could play an important role in future plasma colliders. The final focusing system in our straw-man design uses conventional magnetic focusing and takes up 3.5 km of total 4.5 km footprint. Replacing the magnetic focusing system with a plasma focusing system would significantly reduce the size and cost of the machine. Plasma lenses have already been demonstrated with both electron and positron beams [104], but they come with their own set of issues. The focusing in a plasma lens tends to be aberrative and dilutes the beam emittance, but plasma lenses continue to be an active research topic [105].

7.2 Anti-Plasma

In the introduction, we alluded to the fact that an anti-plasma source might be useful for positron plasma wakefield acceleration, but this idea is not taken seriously. Why not?

Let's consider the case for a plasma afterburner for the ILC. The goal of this device is to extend the energy reach of the machine once it has completed its useful lifetime at a design energy of 1 TeV. In this case, the positron beam energy is 500 GeV when it enters the afterburner. A plasma with density $3 \times 10^{17} \text{ cm}^{-3}$ will provide an accelerating gradient of 50 GeV/m. We therefore need a 10 meter long plasma cell to double the beam energy. Let's further assume that the region occupied by the anti-plasma is 1 cm^2 by 1000 cm, quite narrow compared to the setup at SLAC, but more than wide enough to accommodate the plasma wake which has a radius on the order of $100 \text{ }\mu\text{m}$. This means that in total we need 3×10^{20} anti-protons and 3×10^{20} positrons for our anti-plasma. Low energy positrons are easy to come by, so the challenge is to produce (and contain) 3×10^{20} anti-protons.

The CERN Anti-Proton Decelerator creates anti-protons by bombarding an iridium target with a 26 GeV proton beam containing 1.5×10^{13} particles. From the resulting spray of particles, only 5×10^7 anti-protons are captured. Therefore it takes about 8×10^6 GeV of energy to make a single anti-proton (roughly 1.25 mJ). Now we need 3×10^{20} anti-protons, and it will take $3.6 \times 10^{17} \text{ J}$ to make them. How much will this cost? A kilowatt-hour is equivalent to 3.6 megajoules and costs a little over 10 cents in the United States. Therefore, we will spend \$10 billion on energy alone for the

anti-protons for our anti-plasma afterburner. This is roughly the price of the ILC project as whole! And we haven't considered the cost of the infrastructure needed to make the anti-protons or how to contain them. Nevertheless, it is an interesting idea and it may become more attractive if we do not find a viable alternative for accelerating positrons in plasma.

7.3 Outlook

Beam-driven plasma wakefield acceleration is a burgeoning field. In the next few years, several new PWFA-capable facilities will come online. These include FlashForward at DESY, AWAKE and CALIPHES at CERN, ATF-II at BNL, and FACET-II at SLAC. FACET-II is a third-generation facility for demonstrating PWFA and will provide electron and positron beams with enormous peak currents as well as unprecedented control over the longitudinal bunch shape. If progress in PWFA research continues at the current rate, the next generation will be a fully functioning plasma wakefield accelerator.

Appendix A

Derivation of the Point Charge Wakefields

We follow Chao Chapter 2 [63] and solve for the resistive wall wakefields excited by an ultrarelativistic point charge directly from Maxwell's equations. We will first calculate the fields for the $m = 0$ case. The longitudinal electric field for the $m = 0$ case is of primary interest because it describes the accelerating wake for the ideal case of on-axis drive and witness beams. Next, we will calculate the fields for the $m = 1$ case. The transverse fields for the $m = 1$ case describe the dipole wakefield, which is typically the most important contribution to the transverse wake experienced by particles trailing an off-axis drive beam.

A.1 Maxwell's Equations

We begin by writing down Maxwell's equations in CGS units:

$$\nabla \cdot \vec{E} = 4\pi\rho, \tag{A.1}$$

$$\nabla \cdot \vec{B} = 0, \tag{A.2}$$

$$\nabla \times \vec{E} = -\frac{1}{c} \frac{\partial \vec{B}}{\partial t}, \tag{A.3}$$

$$\nabla \times \vec{B} = \frac{4\pi}{c} \vec{j} + \frac{1}{c} \frac{\partial \vec{E}}{\partial t}. \tag{A.4}$$

Chao expands the equations in cylindrical coordinates (r, ϕ, z) and writes them out term-by-term:

$$\frac{1}{r} \frac{\partial}{\partial r}(rE_r) + \frac{1}{r} \frac{\partial E_\phi}{\partial \phi} + \frac{\partial E_z}{\partial z} = 4\pi\rho, \quad (\text{A.5})$$

$$\frac{1}{r} \frac{\partial}{\partial r}(rB_r) + \frac{1}{r} \frac{\partial B_\phi}{\partial \phi} + \frac{\partial B_z}{\partial z} = 0, \quad (\text{A.6})$$

$$\frac{1}{r} \frac{\partial E_z}{\partial \phi} - \frac{\partial E_\phi}{\partial z} + \frac{1}{c} \frac{\partial B_r}{\partial t} = 0, \quad (\text{A.7})$$

$$\frac{\partial E_r}{\partial z} - \frac{\partial E_z}{\partial r} + \frac{1}{c} \frac{\partial B_\phi}{\partial t} = 0, \quad (\text{A.8})$$

$$\frac{1}{r} \frac{\partial}{\partial r}(rE_\phi) - \frac{1}{r} \frac{\partial E_r}{\partial \phi} + \frac{1}{c} \frac{\partial B_z}{\partial t} = 0, \quad (\text{A.9})$$

$$\frac{1}{r} \frac{\partial B_z}{\partial \phi} - \frac{\partial B_\phi}{\partial z} - \frac{1}{c} \frac{\partial E_r}{\partial t} = \frac{4\pi}{c} j_r, \quad (\text{A.10})$$

$$\frac{\partial B_r}{\partial z} - \frac{\partial B_z}{\partial r} - \frac{1}{c} \frac{\partial E_\phi}{\partial t} = \frac{4\pi}{c} j_\phi, \quad (\text{A.11})$$

$$\frac{1}{r} \frac{\partial}{\partial r}(rB_\phi) - \frac{1}{r} \frac{\partial B_r}{\partial \phi} - \frac{1}{c} \frac{\partial E_z}{\partial t} = \frac{4\pi}{c} j_z. \quad (\text{A.12})$$

Our goal is to solve these equations for a specified source charge and current (in this case an ultrarelativistic particle) and according to the boundary conditions of the hollow channel plasma.

A.2 Hollow Channel Plasma Geometry

The hollow channel plasma is radially symmetric with plasma density

$$n = \begin{cases} 0 & r < a \\ n_0 & a < r < b \\ 0 & b < r \end{cases} \quad (\text{A.13})$$

The wavenumber in the plasma is k_p and the frequency-dependent dielectric constant is

$$\varepsilon = 1 - \frac{k_p^2}{k^2}. \quad (\text{A.14})$$

We will be studying modes with $|k| < k_p$ which implies that $\varepsilon < 0$.

A.3 Sources

The source for our case is an ultrarelativistic “ring charge” given by

$$\rho = \sum_{m=0}^{\infty} \rho_m, \quad (\text{A.15})$$

$$\rho_m = \frac{qr_0^m}{\pi r_0^{m+1}(1 + \delta_{m0})} \delta(z - ct) \delta(r - r_0) e^{im\phi}, \quad (\text{A.16})$$

and $\vec{j} = c\rho\hat{z}$. Here, r_0 is the radius of the ring charge and q is its charge. As in previous sections, we will use the co-moving coordinate $\xi = z - ct$ to describe the longitudinal variation of the beam and wakefield. In the following derivation, we will specifically focus on the fields sourced by a radially symmetric ring charge ($m = 0$) and a ring charge with dipole variation ($m = 1$). The associated ring charges are

$$\rho_0 = \frac{q}{2\pi r_0} \delta(\xi) \delta(r - r_0), \quad (\text{A.17})$$

$$\rho_1 = \frac{q}{\pi r_0} \delta(\xi) \delta(r - r_0) e^{i\phi}. \quad (\text{A.18})$$

In addition to the ultrarelativistic ring charge inside the channel, the fields in the plasma layer also have a source due to the plasma current. The Lorentz Force Law 2.2 allows us to relate the electric field to the current in the plasma

$$\frac{\partial \vec{j}_p}{\partial t} = \frac{c^2 k_p^2}{4\pi} \vec{E}, \quad (\text{A.19})$$

where $\vec{j}_p$ is the current in the plasma and k_p is the plasma wavenumber.

A.4 Fourier Decomposition

Just as the ring charge can be composed into discrete azimuthal modes, so can the fields:

$$\vec{E}(r, \phi, z, t) = \sum_m \vec{E}_m(r, z, t) e^{im\phi}, \quad (\text{A.20})$$

$$\vec{B}(r, \phi, z, t) = \sum_m \vec{B}_m(r, z, t) e^{im\phi}. \quad (\text{A.21})$$

As discussed in Chao’s book, the azimuthal modes are decoupled from one another and we can solve Maxwell’s equations for each of the m modes separately. Furthermore, as we saw in Chapter 2 equation 2.24, the azimuthal modes can be written as a Fourier Transform over their longitudinal

modes

$$\vec{E}_m(r, z, t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{E}_m(r, k) e^{ik\xi} dk, \quad (\text{A.22})$$

$$\vec{B}_m(r, z, t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{B}_m(r, k) e^{ik\xi} dk. \quad (\text{A.23})$$

The transformed fields are especially useful when dealing with the δ -function source charge because the transformed source charge is a constant:

$$\hat{\rho}_m = \frac{qr_0^m}{\pi r_0^{m+1}(1 + \delta_{m0})} \delta(r - r_0). \quad (\text{A.24})$$

In addition, using the Fourier decomposed fields, most of the derivative operators become multiplicative operators:

$$\frac{\partial}{\partial \phi} = im, \quad (\text{A.25})$$

$$\frac{\partial}{\partial z} = ik, \quad (\text{A.26})$$

$$\frac{1}{c} \frac{\partial}{\partial t} = -ik. \quad (\text{A.27})$$

After substituting in the Fourier transformed fields into Maxwell's equations, we carry out all derivatives in ϕ , z , and t which leaves us with coupled differential equations in r . As we shall see, these equations are exactly solvable for the cylindrical hollow channel geometry.

A.5 Solution for azimuthally symmetric case

The $m = 0$ fields are sourced by the radially symmetric ring charge. In free space, the field outside of an ultrarelativistic ring charge moving in the z direction is zero for all components except for E_r and B_ϕ . Inside the channel, these fields act together to generate an E_z component. Note that in the case of the radially symmetric driver, there are no E_ϕ and B_r fields, which means that there will be no B_z field in the channel. In the cylindrical geometry, the $m = 0$ mode is the only TM mode. All higher order modes are HEM with E_z and B_z components.

A.5.1 Inside the Channel

Examining Maxwell's equations for the $m = 0$ mode, we find that the three equations A.5, A.8, and A.10 are sufficient to solve for our three non-zero components E_r , B_ϕ , and E_z^* within the channel. After replacing the derivatives with respect to z and t with A.26 and A.27, and substituting

*Note that we have dropped the hat on the Fourier transformed components.

in $\rho = \rho_0$, we have:

$$\frac{1}{r} \frac{\partial}{\partial r} (r E_r) + i k E_z = \frac{2q}{r_0} \delta(r - r_0), \quad (\text{A.28})$$

$$i k E_r - \frac{\partial E_z}{\partial r} - i k B_\phi = 0, \quad (\text{A.29})$$

$$-i k B_\phi + i k E_r = 0. \quad (\text{A.30})$$

Therefore, $E_r = B_\phi$ and $E_z = \text{const} = A$. Plugging in $E_z = A$ in equation A.28 and integrating with respect to r gives

$$E_r = -i \frac{k A}{2} r + \frac{2q}{r}. \quad (\text{A.31})$$

So the fields inside the channel are given by:

$$E_z = A, \quad (\text{A.32})$$

$$E_r = B_\phi = -i \frac{k A}{2} r + \frac{2q}{r}. \quad (\text{A.33})$$

A.5.2 The Channel Wall

Next, we examine the fields in the wall of the hollow channel where the plasma is present. There are no free source charges or currents, but there is the plasma current described by equation A.19. We assume that the transformed plasma current has the same temporal dependence as the fields, which gives

$$\vec{j} = i \frac{c k_p^2}{4\pi k} \vec{E}. \quad (\text{A.34})$$

We use the same three equations to solve for our field components, but this time there is j_r contribution from equation A.10:

$$\frac{1}{r} \frac{\partial}{\partial r} (r E_r) + i k E_z = 0, \quad (\text{A.35})$$

$$i k E_r - \frac{\partial E_z}{\partial r} - i k B_\phi = 0, \quad (\text{A.36})$$

$$-i k B_\phi + i k E_r = i \frac{k_p^2}{k} E_r. \quad (\text{A.37})$$

The relationship between the transverse field components is therefore

$$B_\phi = \left(1 - \frac{k_p^2}{k^2} \right) E_r = \varepsilon E_r, \quad (\text{A.38})$$

which gives the following pair of coupled differential equations:

$$\frac{\partial E_z}{\partial r} = \frac{ik_p^2}{k} E_r, \quad (\text{A.39})$$

$$\frac{1}{r} \frac{\partial}{\partial r} (r E_r) = -ik E_z. \quad (\text{A.40})$$

We can decouple the equations to find the following second order differential equation for E_z

$$r^2 \frac{\partial^2 E_z}{\partial r^2} + r \frac{\partial E_z}{\partial r} - k_p^2 r^2 E_z = 0, \quad (\text{A.41})$$

which is the zeroth order modified Bessel equation with solutions

$$E_z(r) = CI_0(k_p r) + DK_0(k_p r). \quad (\text{A.42})$$

The transverse field components follow from equations A.38, A.39, and the identities found in Section B.3

$$E_r(r) = -i \frac{k}{k_p} [CI_1(k_p r) - DK_1(k_p r)], \quad (\text{A.43})$$

$$B_\phi(r) = -i \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2} \right) [CI_1(k_p r) - DK_1(k_p r)]. \quad (\text{A.44})$$

A.5.3 Outside the channel

Outside the channel, there is no plasma and no sources from free charges or currents. We again have $E_r = B_\phi$ in this region and it follows that $E_z = \text{const} = G$. However we require that $E_z(r = \infty) = 0$ and therefore $G = 0$. Under these conditions, equation A.5 gives

$$\frac{1}{r} \frac{\partial}{\partial r} (r E_r) = 0, \quad (\text{A.45})$$

and we find

$$E_r = B_\phi = \frac{F}{r}. \quad (\text{A.46})$$

A.5.4 Boundary Conditions

We have solutions for the fields everywhere in terms of four constants A , B , C , and F . We use the boundary conditions on the fields at the inner and outer edges of the plasma channel at $r = a$ and $r = b$, respectively, to eliminate the constants and solve for A , the magnitude of the E_z field in the channel.

E_z is continuous across the plasma boundary, so applying the condition $E_z(r = b) = 0$ gives

$$CI_0(k_pb) + DK_0(k_pb) = 0, \quad (\text{A.47})$$

$$\rightarrow D = -C \frac{I_0(k_pb)}{K_0(k_pb)}. \quad (\text{A.48})$$

Using continuity at the inner boundary where $E_z(r = a) = A$ gives

$$CI_0(k_pa) - C \frac{I_0(k_pb)}{K_0(k_pb)} K_0(k_pa) = A, \quad (\text{A.49})$$

$$\rightarrow C = A \frac{K_0(k_pb)}{B_{00}(a, b)}, \quad (\text{A.50})$$

and I have introduced the Bessel boundary notation

$$B_{ij}(x, y) = I_i(k_px)K_j(k_py) + (-1)^{i-j+1}I_j(k_py)K_i(k_px), \quad (\text{A.51})$$

which for $B_{00}(a, b)$ gives

$$B_{00}(a, b) = I_0(k_pa)K_0(k_pb) - I_0(k_pb)K_0(k_pa). \quad (\text{A.52})$$

Plugging in our solutions for C and D into equations A.43 and A.43 gives the solutions for the transverse fields in the plasma boundary in terms of the field magnitude A

$$E_r = -iA \frac{k}{k_p} \frac{B_{10}(r, b)}{B_{00}(a, b)}, \quad (\text{A.53})$$

$$B_\phi = -iA \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2}\right) \frac{B_{10}(r, b)}{B_{00}(a, b)}. \quad (\text{A.54})$$

We use continuity of $D_r = \varepsilon E_r$ across the outer boundary to find F

$$F = -iAb \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2}\right) \frac{B_{10}(b, b)}{B_{00}(a, b)}. \quad (\text{A.55})$$

We now have solutions for the fields everywhere in terms of the field magnitude A :

$$E_z(r) = \begin{cases} A & r < a \\ A \frac{B_{00}(r,b)}{B_{00}(a,b)} & a < r < b, \\ 0 & r > b \end{cases} \quad (\text{A.56})$$

$$E_r(r) = \begin{cases} -i \frac{kA}{2} r + \frac{2q}{r} & r < a \\ -iA \frac{k}{k_p} \frac{B_{10}(r,b)}{B_{00}(a,b)} & a < r < b, \\ -iA \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2}\right) \frac{B_{10}(b,b)}{B_{00}(a,b)} \frac{b}{r} & r > b \end{cases} \quad (\text{A.57})$$

$$B_\phi(r) = \begin{cases} -i \frac{kA}{2} r + \frac{2q}{r} & r < a \\ -iA \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2}\right) \frac{B_{10}(r,b)}{B_{00}(a,b)} & a < r < b. \\ -iA \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2}\right) \frac{B_{10}(b,b)}{B_{00}(a,b)} \frac{b}{r} & r > b \end{cases} \quad (\text{A.58})$$

Finally, we solve for A by using the boundary condition $D_r = \varepsilon E_r$ at the inner boundary

$$-ikA \frac{a}{2} + \frac{2q}{a} = -ikA \left(1 - \frac{k_p^2}{k^2}\right) \frac{B_{10}(a,b)}{B_{00}(a,b)}, \quad (\text{A.59})$$

$$\rightarrow A(k) = \frac{4iqk_p^2 B_{00}(a,b)}{k_p a [2B_{10}(a,b) - k_p a B_{00}(a,b)]} \frac{k}{k^2 - \chi^2 k_p^2}, \quad (\text{A.60})$$

with

$$\chi = \sqrt{\frac{2B_{10}(a,b)}{2B_{10}(a,b) - k_p a B_{00}(a,b)}}. \quad (\text{A.61})$$

We note that $A(k)$ has resonances at $k = \pm \chi k_p$ which is the frequency we found in equation 3.28. Plugging in our value for k into the dielectric constant we find

$$\varepsilon = 1 - \frac{1}{\chi^2} = -\frac{k_p a}{2} \frac{B_{00}(a,b)}{B_{10}(a,b)} \quad (\text{A.62})$$

and the above equations for E_r and B_ϕ simplify

$$E_r(r) = \begin{cases} -i \frac{kA}{2} r + \frac{2q}{r} & r < a \\ -iA \frac{k}{k_p} \frac{B_{10}(r,b)}{B_{00}(a,b)} & a < r < b, \\ \frac{iAkab}{2r} \frac{B_{10}(b,b)}{B_{10}(a,b)} & r > b \end{cases} \quad (\text{A.63})$$

$$B_\phi(r) = \begin{cases} -i \frac{kA}{2} r + \frac{2q}{r} & r < a \\ \frac{iAka}{2} \frac{B_{10}(r,b)}{B_{10}(a,b)} & a < r < b. \\ \frac{iAkab}{2r} \frac{B_{10}(b,b)}{B_{10}(a,b)} & r > b \end{cases} \quad (\text{A.64})$$

A.5.5 Inverting the Transform

Since $A(k)$ is the Fourier Transform of $E_z(\xi)$, we need to invert the transform to find the point-charge wake function

$$E_z(\xi) = \int_{-\infty}^{\infty} \frac{dk}{2\pi} e^{ik\xi} A(k) = \frac{4iqk_p^2 B_{00}(a, b)}{k_p a [2B_{10}(a, b) - k_p a B_{00}(a, b)]} \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{ke^{ik\xi}}{k^2 - \chi^2 k_p^2}. \quad (\text{A.65})$$

The integral is over the real axis and hits both poles, but this integral can be shown to be finite by evaluating it in the complex plane. In this case the integral is written as

$$I = \oint_C \frac{dk}{2\pi} \frac{ke^{ik\xi}}{k^2 - \chi^2 k_p^2}. \quad (\text{A.66})$$

We define the contour C with two parts C_1 and C_2 . C_1 is the original integral along the real axis and C_2 is an arc with a radius that goes to infinity. We use causality to choose C_2 such that its total contribution to the integral is zero ($\xi = z - ct$ and there are no fields for $\xi > 0$). The integral along C_2 is dominated by the exponential, so we choose to close the loop in the upper half plane

$$I = \int_{C_1} \frac{dk}{2\pi} \frac{ke^{ik\xi}}{k^2 - \chi^2 k_p^2} + \int_{C_2} \frac{dk}{2\pi} \frac{ke^{-|k||\xi|}}{k^2 - \chi^2 k_p^2}. \quad (\text{A.67})$$

The integral along C_2 goes to zero as $|k| \rightarrow \infty$, so that leaves

$$I = \int_{C_1} \frac{dk}{2\pi} \frac{ke^{ik\xi}}{k^2 - \chi^2 k_p^2}. \quad (\text{A.68})$$

which is equivalent to our original integral. Note that the integral along C_1 includes two half loops in order to include the real poles. The purpose of these loops is to make sure that the poles are within the contour. The radius of these loops go to zero so that C_1 approximates the integral along the real axis from $\pm\infty$.

Finally we are in a position to evaluate this integral. Our contour encloses two poles so we use the Cauchy Integral Theorem to find the value of the integral

$$2\pi i f(z_0) = \oint \frac{f(z)dz}{z - z_0}, \quad (\text{A.69})$$

and we use partial fractions to make our integral take the form

$$I = \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{ke^{ik\xi}}{k^2 - \chi^2 k_p^2} = \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{e^{ik\xi}}{2} \left[\frac{1}{k - \chi k_p} + \frac{1}{k + \chi k_p} \right], \quad (\text{A.70})$$

Figure A.1: Contour for the integral A.65.

which, after application of the Integral Theorem, gives

$$I = i \frac{e^{i\chi k_p \xi} + e^{-i\chi k_p \xi}}{2} = i \cos(\chi k_p \xi). \quad (\text{A.71})$$

Plugging this result into our expression for E_z we find

$$E_z(\xi) = -\frac{4qk_p^2 B_{00}(a, b)}{k_p a [2B_{10}(a, b) - k_p a B_{00}(a, b)]} \cos(\chi k_p \xi) \Theta(\xi). \quad (\text{A.72})$$

Note that the sign of E_z is negative, indicating that the field is decelerating in the phase of the wake immediately trailing the beam particle.

A.6 Solution for the dipole case

For $m = 1$, the ring source has azimuthal variation and non-zero E_ϕ and B_r components which source a B_z field inside the channel. Therefore, we need to solve for all six components of $\vec{E}$ and $\vec{B}$ to find a complete description of the fields.

A.6.1 Inside the Channel

We will use equations A.7, A.8, A.10, and A.11 to find relations between E_r and B_ϕ , E_ϕ and B_r , and solve for E_z and B_z . After application of the derivatives in ϕ , z , and t , we have

$$\frac{1}{r}E_z = k(E_\phi + B_r), \quad (\text{A.73})$$

$$\frac{\partial E_z}{\partial r} = ik(E_r - B_\phi), \quad (\text{A.74})$$

$$\frac{1}{r}B_z = -k(E_r - B_\phi), \quad (\text{A.75})$$

$$\frac{\partial B_z}{\partial r} = ik(E_\phi + B_r), \quad (\text{A.76})$$

and combining these equations gives a coupled differential equation for E_z and B_z

$$\frac{\partial E_z}{\partial r} = -\frac{i}{r}B_z, \quad (\text{A.77})$$

$$\frac{\partial B_z}{\partial r} = \frac{i}{r}E_z. \quad (\text{A.78})$$

We decouple the equations by taking another derivative which gives

$$\left[r^2 \frac{\partial^2}{\partial r^2} + r \frac{\partial}{\partial r} - 1 \right] \{E_z, B_z\} = 0, \quad (\text{A.79})$$

with solutions

$$E_z = Ar + Br^{-1}, \quad (\text{A.80})$$

$$B_z = iAr - iBr^{-1}. \quad (\text{A.81})$$

We require that the fields are finite on-axis, and therefore set $B = 0$.

Our solutions for E_z and B_z provide the following relationships for the transverse field combinations $E_r - B_\phi$ and $E_\phi + B_r$

$$E_r - B_\phi = -i\frac{A}{k}, \quad (\text{A.82})$$

$$E_\phi + B_r = \frac{A}{k}. \quad (\text{A.83})$$

The quantity $W_\perp = E_r - B_\phi$ is the transverse force experienced by an ultrarelativistic particle traveling in the channel. We note that this force is constant in r for the $m = 1$ mode.

We use equations A.5 and A.6 to solve for the remaining field components

$$\frac{1}{r} \frac{\partial}{\partial r}(rE_r) + \frac{i}{r} E_\phi + ikE_z = \frac{4q}{r_0} \delta(r - r_0), \quad (\text{A.84})$$

$$\frac{1}{r} \frac{\partial}{\partial r}(rB_r) + \frac{i}{r} B_\phi + ikB_z = 0, \quad (\text{A.85})$$

which, after plugging in for E_z and B_z and using relationships A.82 and A.83 gives coupled differential equations in E_r and B_r

$$\frac{\partial}{\partial r}(rE_r) - iB_r = -iAk(r^2 + k^{-2}) + \frac{4q}{r_0} r \delta(r - r_0), \quad (\text{A.86})$$

$$\frac{\partial}{\partial r}(rB_r) + iE_r = Ak(r^2 + k^{-2}). \quad (\text{A.87})$$

We take a second derivative to decouple the equations which yields

$$\frac{\partial}{\partial r} \left[r \frac{\partial}{\partial r}(rB_r) \right] - B_r = 2Akr^2 - \frac{4iq}{r_0} r \delta(r - r_0). \quad (\text{A.88})$$

To solve the equation, we expand the derivative and define $U_r = \partial_r B_r$

$$r^2 \frac{\partial U_r}{\partial r} + 3rU_r = 2Akr^2 - \frac{4iq}{r_0} r \delta(r - r_0). \quad (\text{A.89})$$

Next, we observe that $r^3 \partial_r U_r + 3r^2 U_r = \partial_r [r^3 U_r]$, so multiplying both sides by r and integrating gives

$$r^3 U_r = Ak \frac{r^4}{2} - 4iqr_0 \Theta(r - r_0) + C_1. \quad (\text{A.90})$$

After dividing out the r^3 , we integrate again to find

$$B_r = Ak \frac{r^2}{4} + 2iqr_0 \left[\frac{1}{r^2} - \frac{1}{r_0^2} \right] \Theta(r - r_0) + \frac{C_1}{r^2} + C_2. \quad (\text{A.91})$$

The expression must be valid at $r = 0$ which means that $C_1 = 0$. We choose $C_2 = A/2k + C_3$ where C_3 is a free parameter. This choice of constant provides a symmetry in the expression of the

transverse field components

$$E_r = -iAk\frac{r^2}{4} + 2qr_0 \left[\frac{1}{r^2} + \frac{1}{r_0^2} \right] \Theta(r - r_0) - i\frac{A}{2k} + iC_3, \quad (\text{A.92})$$

$$B_r = +Ak\frac{r^2}{4} + 2iqr_0 \left[\frac{1}{r^2} - \frac{1}{r_0^2} \right] \Theta(r - r_0) + \frac{A}{2k} + C_3, \quad (\text{A.93})$$

$$E_\phi = -Ak\frac{r^2}{4} - 2iqr_0 \left[\frac{1}{r^2} - \frac{1}{r_0^2} \right] \Theta(r - r_0) + \frac{A}{2k} - C_3, \quad (\text{A.94})$$

$$B_\phi = -iAk\frac{r^2}{4} + 2qr_0 \left[\frac{1}{r^2} + \frac{1}{r_0^2} \right] \Theta(r - r_0) + i\frac{A}{2k} + iC_3. \quad (\text{A.95})$$

A.6.2 The Channel Wall

We return to equations A.7, A.8, A.10, and A.11, now with the additional plasma currents to find

$$\frac{i}{r}E_z - ikE_\phi - ikB_r = 0, \quad (\text{A.96})$$

$$ikE_r - \frac{\partial E_z}{\partial r} - ikB_\phi = 0, \quad (\text{A.97})$$

$$\frac{i}{r}B_z - ikB_\phi + ikE_r = i\frac{k_p^2}{k}E_r, \quad (\text{A.98})$$

$$ikB_r - \frac{\partial B_z}{\partial r} + ikE_\phi = i\frac{k_p^2}{k}E_\phi. \quad (\text{A.99})$$

Solving for E_r and B_ϕ gives

$$E_r = \frac{k}{k_p^2} \left[\frac{1}{r}B_z - i\frac{\partial E_z}{\partial r} \right], \quad (\text{A.100})$$

$$E_\phi = \frac{k}{k_p^2} \left[\frac{1}{r}E_z + i\frac{\partial B_z}{\partial r} \right]. \quad (\text{A.101})$$

Combining these expressions with equation A.7 (note that there is no free charge in the plasma)

$$\frac{1}{r}\partial(rE_r) + \frac{i}{r}E_\phi + ikE_z = 0, \quad (\text{A.102})$$

produces a Modified Bessel equation of order 1 for E_z

$$r\partial_r(r\partial_r E_z) - (1 + k_p^2 r^2)E_z = 0, \quad (\text{A.103})$$

with solutions

$$E_z = C_E I_1(k_p r) + D_E K_1(k_p r). \quad (\text{A.104})$$

The same approach can be applied to equation A.9 to find the solution for B_z

$$B_z = C_B I_1(k_p r) + D_B K_1(k_p r). \quad (\text{A.105})$$

The transverse field components follow from expressions A.100 and A.101. We will write them down after eliminating the constants C_E , C_B , D_E , and D_B .

A.6.3 Outside the Channel

Solving Maxwell's equations outside the channel yields the same expressions for E_z and B_z as inside the channel, but we select the decaying solution

$$E_z = G r^{-1}, \quad (\text{A.106})$$

$$B_z = -i G r^{-1}. \quad (\text{A.107})$$

This solution appears to satisfy the boundary condition that $E_z(r = \infty) = 0$, but we will see in a moment that $G = 0$ from boundary conditions on the other field components.

Repeating the procedure we used to produce equation A.88, we find

$$\frac{\partial}{\partial r} \left[r \frac{\partial}{\partial r} (r B_r) \right] - B_r = -2kG. \quad (\text{A.108})$$

Integrating this equation yields

$$B_r = -Gk \log(r) + \frac{P_1}{r^2} + P_2. \quad (\text{A.109})$$

At this point, it is not entirely clear how the boundary conditions at infinity should be applied. From an energy conservation standpoint, both the logarithmic term and the constant term yield infinite energy in the fields. On the other hand, the ultrarelativistic particle with speed c has infinite energy, and we need at least two constants to satisfy the boundary conditions. In light of this, it seems reasonable to set $G = 0$ so that E_z is zero outside the channel, consistent with vacuum fields for a

relativistic dipole. The resulting fields outside the channel wall are

$$E_z = 0, \quad (\text{A.110})$$

$$B_z = 0, \quad (\text{A.111})$$

$$E_r = -i \frac{P_1}{r^2} + iP_2, \quad (\text{A.112})$$

$$B_r = + \frac{P_1}{r^2} + P_2, \quad (\text{A.113})$$

$$E_\phi = - \frac{P_1}{r^2} - P_2, \quad (\text{A.114})$$

$$B_\phi = -i \frac{P_1}{r^2} + iP_2. \quad (\text{A.115})$$

A.6.4 Boundary Conditions

We have solutions for the fields everywhere in terms of seven constants A , C_3 , C_E , C_B , D_E , D_B , and P_1 . We would usually expect to have eight constants at this point, but we used an additional boundary condition on B_r to eliminate G in the previous section. We now use the boundary conditions on the plasma channel to eliminate the constants and solve for A .

We first use continuity of E_z and B_z to eliminate C_E , C_B , D_E , and D_B in favor of A . The process is the same as the one used in Section A.5.4, which yields expressions for E_z and B_z fields in the channel wall

$$E_z = Aa \frac{B_{11}(r, b)}{B_{11}(a, b)}, \quad (\text{A.116})$$

$$B_z = iAa \frac{B_{11}(r, b)}{B_{11}(a, b)}. \quad (\text{A.117})$$

We have recovered the relation $B_z = iE_z$. We use these expressions and the identities in Section B.3 to write down the transverse field components in the channel wall

$$E_r = -iAa \frac{k}{k_p} \frac{B_{21}(r, b)}{B_{11}(a, b)}, \quad (\text{A.118})$$

$$B_r = Aa \left[\frac{1}{kr} \frac{B_{11}(r, b)}{B_{11}(a, b)} + \frac{k}{k_p} \frac{B_{21}(r, b)}{B_{11}(a, b)} \right], \quad (\text{A.119})$$

$$E_\phi = -Aa \frac{k}{k_p} \frac{B_{21}(r, b)}{B_{11}(a, b)}, \quad (\text{A.120})$$

$$B_\phi = iAa \left[\frac{1}{kr} \frac{B_{11}(r, b)}{B_{11}(a, b)} - \frac{k}{k_p} \left(1 - \frac{k_p^2}{k^2} \right) \frac{B_{21}(r, b)}{B_{11}(a, b)} \right]. \quad (\text{A.121})$$

Next we can use continuity of $D_r = \varepsilon E_r$ to solve for P_1 and P_2 . Noting that $B_{11}(b, b) = 0$, we

find

$$P_1 = \frac{Aa}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)} \left(2 - \frac{k_p^2}{k^2} \right) b^2, \quad (\text{A.122})$$

$$P_2 = \frac{Aa}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)}. \quad (\text{A.123})$$

This gives us the complete description of the field outside the channel in terms of A

$$E_z = 0, \quad (\text{A.124})$$

$$B_z = 0, \quad (\text{A.125})$$

$$E_r = i \frac{Aa}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)} \left[1 - \left(2 - \frac{k_p^2}{k^2} \right) \frac{b^2}{r^2} \right], \quad (\text{A.126})$$

$$B_r = + \frac{Aa}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)} \left[1 + \left(2 - \frac{k_p^2}{k^2} \right) \frac{b^2}{r^2} \right], \quad (\text{A.127})$$

$$E_\phi = - \frac{Aa}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)} \left[1 + \left(2 - \frac{k_p^2}{k^2} \right) \frac{b^2}{r^2} \right], \quad (\text{A.128})$$

$$B_\phi = i \frac{Aa}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)} \left[1 - \left(2 - \frac{k_p^2}{k^2} \right) \frac{b^2}{r^2} \right]. \quad (\text{A.129})$$

We will examine the consequences of a finite value in the transverse fields at $r = \infty$ in Section A.8.

Finding the fields inside the channel, we use continuity of B_ϕ at $r = a$ to solve for the remaining constant C_3

$$C_3 = -kA \frac{a^2}{4} + aA \frac{k}{k_p} \frac{B_{21}(a, b)}{B_{11}(a, b)} + \frac{A}{2k} - 2iqr_0 \left[\frac{1}{a^2} - \frac{1}{r_0^2} \right]. \quad (\text{A.130})$$

Finally, we solve for A using continuity of $D_r = \varepsilon E_r$ on the inner boundary

$$-ikA \frac{a^2}{2} + \frac{4qr_0}{a^2} + iAa \frac{k}{k_p} \frac{B_{21}(a, b)}{B_{11}(a, b)} = -iAa \left(1 - \frac{k_p^2}{k^2} \right) \frac{k}{k_p} \frac{B_{21}(a, b)}{B_{11}(a, b)} \quad (\text{A.131})$$

$$\rightarrow A(k) = \frac{8iqr_0 k_p^4}{(k_p a)^3} \frac{B_{11}(a, b)}{4B_{21}(a, b) - k_p a B_{11}(a, b)} \frac{k}{k^2 - \chi^2 k_p^2}, \quad (\text{A.132})$$

with

$$\chi = \sqrt{\frac{2B_{21}(a, b)}{4B_{21}(a, b) - k_p a B_{11}(a, b)}}. \quad (\text{A.133})$$

A.6.5 Inverting the Transform

As in the previous section, we need to invert the transform, but in this case we are interested in the quantity

$$W_{\perp} = E_r - B_{\phi} = -i \frac{A}{k}. \quad (\text{A.134})$$

Plugging in our expression for $A(k)$, we have

$$W_{\perp}(\xi) = \frac{8qr_0k_p^4}{(k_p a)^3} \frac{B_{11}(a, b)}{4B_{21}(a, b) - k_p a B_{11}(a, b)} \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{e^{ik\xi}}{k^2 - \chi^2 k_p^2}. \quad (\text{A.135})$$

The integral is nearly identical to the one solved for equation A.65, except that the partial fraction decomposition is

$$I = \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{e^{ik\xi}}{k^2 - \chi^2 k_p^2} = \int_{-\infty}^{\infty} \frac{dk}{2\pi} \frac{e^{ik\xi}}{2\chi k_p} \left[\frac{1}{k - \chi k_p} + \frac{1}{k + \chi k_p} \right]. \quad (\text{A.136})$$

Carrying out the integral on the same contour as before and applying the Integral theorem gives

$$I = -\frac{\sin(\chi k_p \xi)}{\chi k_p}, \quad (\text{A.137})$$

and plugging into equation A.135 yields

$$W_{\perp}(\xi) = -\frac{8qr_0k_p^3}{(k_p a)^3} \frac{B_{11}(a, b)}{\chi [4B_{21}(a, b) - k_p a B_{11}(a, b)]} \sin(\chi k_p \xi). \quad (\text{A.138})$$

We find that $W_{\perp}$ is independent of r for the dipole mode (as stated previously), that it is 90° out of phase with E_z , and that its strength is proportional to a^{-3} .

A.7 Comparison to Previous Work

Schroeder [57] calculates the same quantities for the case of a hollow channel surrounded by infinite plasma ($r_b = \infty$). We can compare our calculations to Schroeder's and take the limit $r_b \rightarrow \infty$ to show that we recover the same results.

First, we'll compare the resonant frequency which Schroeder calls Ω

$$\Omega_m^2 = \frac{(1 + \delta_{m0})(m+1)K_{m+1}(k_p a)}{2(m+1)K_{m+1}(k_p a) + k_p a K_m(k_p a)}, \quad (\text{A.139})$$

Figure A.2: Left: A plot of χ_0 and χ_1 as a function of the normalized outer plasma radius $k_p b$. We compare to the limiting values Ω_0 and Ω_1 . Right: A plot of $A_0/2$ and $A_1/2$ as a function of the normalized outer plasma radius $k_p b$. We compare to the limiting values κ_0 and κ_1 . For these calculations, we chose $k_p a = 5$ and vary $k_p b$ from 5.5 to 10.

to our values for χ . For the two cases of interest, $m = 0$ and $m = 1$, we have

$$\Omega_0^2 = \frac{2K_1(k_p a)}{2K_1(k_p a) + k_p a K_0(k_p a)}, \quad (\text{A.140})$$

$$\Omega_1^2 = \frac{2K_2(k_p a)}{4K_2(k_p a) + k_p a K_1(k_p a)}, \quad (\text{A.141})$$

which has a similar form to our expressions

$$\chi_0^2 = \frac{2B_{10}(a, b)}{2B_{10}(a, b) - k_p a B_{00}(a, b)}, \quad (\text{A.142})$$

$$\chi_1^2 = \frac{2B_{21}(a, b)}{4B_{21}(a, b) - k_p a B_{11}(a, b)}. \quad (\text{A.143})$$

Recalling that

$$B_{00}(a, b) = I_0(k_p a)K_0(k_p b) - I_0(k_p b)K_0(k_p a), \quad (\text{A.144})$$

$$B_{10}(a, b) = I_1(k_p a)K_0(k_p b) + I_0(k_p b)K_1(k_p a), \quad (\text{A.145})$$

$$B_{11}(a, b) = I_1(k_p a)K_1(k_p b) - I_1(k_p b)K_1(k_p a), \quad (\text{A.146})$$

$$B_{21}(a, b) = I_2(k_p a)K_1(k_p b) + I_1(k_p b)K_2(k_p a), \quad (\text{A.147})$$

we see that in the limit $b \rightarrow \infty$, the terms proportional to $K_0(k_p b)$ and $K_1(k_p b)$ go to zero. The terms proportional to $I_0(k_p b)$ and $I_1(k_p b)$ are divergent, but they appear in both the numerator and

denominator and cancel out, so

$$\lim_{b \rightarrow \infty} \chi_m^2 = \Omega_m^2. \quad (\text{A.148})$$

We plot the limiting behavior for χ_0 and χ_1 in the left subplot of Figure A.2.

Next, we compare Schroeder's expression for the loss factors κ_m

$$\kappa_0 = \frac{2K_0(k_p a)}{k_p a [2K_1(k_p a) + k_p a K_0(k_p a)]}, \quad (\text{A.149})$$

$$\kappa_1 = \frac{4K_1(k_p a)}{k_p a [4K_2(k_p a) + k_p a K_1(k_p a)]}, \quad (\text{A.150})$$

to our own values for $A/2$

$$\frac{A_0}{2} = \frac{2B_{00}(a, b)}{k_p a [2B_{10}(a, b) - k_p a B_{00}(a, b)]}, \quad (\text{A.151})$$

$$\frac{A_1}{2} = \frac{4B_{11}(a, b)}{k_p a [4B_{21}(a, b) - k_p a B_{11}(a, b)]}, \quad (\text{A.152})$$

and as in the previous case we find $\lim_{b \rightarrow \infty} A_m/2 = \kappa_m$. The limiting behaviors for the loss factors is shown in the right subplot of Figure A.7.

A.8 Investigation of the Plane Wave Component

In our solution for the fields outside of the channel in the $m = 1$ case, we found that the transverse fields have a plane wave component with amplitude

$$P_2(k) = A(k) \frac{a}{2} \frac{k}{k_p} \frac{B_{21}(b, b)}{B_{11}(a, b)}. \quad (\text{A.153})$$

Plugging in for $A(k)$ and inverting the transform gives

$$P_2(\xi) = q k_p^2 \frac{k_p r_0}{(k_p a)^2} \frac{4\chi B_{21}(b, b)}{4B_{21}(a, b) - k_p a B_{11}(a, b)} \sin(\chi k_p \xi). \quad (\text{A.154})$$

Written in this way, we see that the strength of this field is proportional to q and k_p^2 multiplied by geometric factors. The first geometric factor

$$P_2 \propto \frac{k_p r_0}{(k_p a)^2} \quad (\text{A.155})$$

says that the field is linearly proportional to the bunch offset r_0 and inversely proportional to the inner channel radius a squared. In typical experiments at FACET, we have $k_p r_0 < 1$ and $k_p a \approx 5$. Therefore, the amplitude of this field is suppressed by keeping the beam near the axis and having a wide channel. The second geometric factor gives the field dependence on the outer radius b . It

Figure A.3: The colored lines show the value of the geometric component of P_2 for several values of $r_a = k_p a$ as a function of $r_b = k_p b$ which is greater than r_a . The black dashed line shows the geometric component of P_2 as a function of r_a with a fixed annular width of 0.25. For all plots, $k_p r_0 = 1$

is a complicated combination of Bessel functions, but its main effect is to attenuate the external field when $k_p(b - a) > 1$, *i.e.* when the width of the plasma annulus is greater than a skin depth. Figure A.3 shows the geometric dependence of P_2 as a function of $r_a = k_p a$ and $r_b = k_p b$ with fixed bunch offset $k_p r_0 = 1$.

Let's numerically estimate the strength of this field for the following FACET-like parameters: $n_0 = 7 \times 10^{16} \text{ cm}^{-3}$, $r_0 = 20 \text{ } \mu\text{m}$, $a = 100 \text{ } \mu\text{m}$, $b = 120 \text{ } \mu\text{m}$, and $N = 10^{10}$ particles. This plasma density gives $k_p^{-1} = 20 \text{ } \mu\text{m}$ and therefore $k_p r_0 = 1$, $k_p a = 5$, and $k_p b = 6$. In SI units we have

$$|P_2| = \frac{N e k_p^2}{\pi \epsilon_0} \frac{k_p r_0}{(k_p a)^2} \frac{\chi B_{21}(b, b)}{4 B_{21}(a, b) - k_p a B_{11}(a, b)} = 41.25 \text{ MV/m.} \quad (\text{A.156})$$

Appendix B

Bessel Function Identities

Bessel functions appear frequently in this work because they arise in relation to the cylindrical geometries associated with the hollow channel plasma and axicon calculations. In this appendix, we list some true facts about the Bessel functions.

B.1 Bessel's Equations

Bessel's equation is

$$x^2 \frac{d^2 y}{dx^2} + x \frac{dy}{dx} + (x^2 - \alpha^2)y = 0, \quad (\text{B.1})$$

and the solutions are:

$$Y(x) = AJ_\alpha(x) + BY_\alpha(x). \quad (\text{B.2})$$

Both $J_\alpha(x)$ and $Y_\alpha(x)$ are oscillatory functions.

The Modified Bessel equation is

$$x^2 \frac{d^2 y}{dx^2} + x \frac{dy}{dx} - (x^2 + \alpha^2)y = 0, \quad (\text{B.3})$$

and the solutions are:

$$Y(x) = AI_\alpha(x) + BK_\alpha(x). \quad (\text{B.4})$$

$I_\alpha(x)$ grows almost exponentially and $K_\alpha(x)$ decays almost exponentially.

B.2 Series and Integral Representations

The series expansion of $J_\alpha(x)$ is given by

$$J_\alpha(x) = \sum_{m=0}^{\infty} \frac{(-1)^m}{m! \Gamma(m + \alpha + 1)} \left(\frac{x}{2}\right)^{2m+\alpha}, \quad (\text{B.5})$$

and the series expansion of $I_\alpha(x)$ is given by

$$I_\alpha(x) = \sum_{m=0}^{\infty} \frac{1}{m! \Gamma(m + \alpha + 1)} \left(\frac{x}{2}\right)^{2m+\alpha}. \quad (\text{B.6})$$

$J_\alpha(x)$ can also be represented in integral form

$$J_n(x) = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{i(n\tau - x \sin(\tau))} d\tau. \quad (\text{B.7})$$

B.3 Derivative Identities

The following are useful identities for dealing with derivatives of Modified Bessel functions:

$$\frac{\partial I_n}{\partial x} = \frac{1}{2} (I_{n-1}(x) + I_{n+1}(x)), \quad (\text{B.8})$$

$$\frac{\partial K_n}{\partial x} = -\frac{1}{2} (K_{n-1}(x) + K_{n+1}(x)), \quad (\text{B.9})$$

$$I_{-n}(x) = I_n(x), \quad (\text{B.10})$$

$$K_{-n}(x) = K_n(x), \quad (\text{B.11})$$

$$K_{n+1}(x) = K_{n-1}(x) + \frac{2n}{x} K_n(x), \quad (\text{B.12})$$

$$I_{n+1}(x) = I_{n-1}(x) - \frac{2n}{x} I_n(x). \quad (\text{B.13})$$

Alternative derivative identities are

$$\frac{\partial I_n}{\partial x} = I_{n+1}(x) + \frac{n}{x} I_n(x), \quad (\text{B.14})$$

$$\frac{\partial K_n}{\partial x} = -K_{n+1}(x) + \frac{n}{x} K_n(x). \quad (\text{B.15})$$

B.4 Bessel Boundary Notation and Identities

When solving Maxwell's equations at multiple boundaries in cylindrical coordinates, we frequently encounter the following pattern of Bessel functions

$$B_{m,n}(x, y) = I_m(kx)K_n(ky) + (-1)^{m-n+1}I_m(ky)K_n(kx). \quad (\text{B.16})$$

It is always the case that one or both of the arguments x and y are constants. If x is a variable, we sometimes encounter derivatives of this expression

$$\frac{\partial}{\partial x} B_{m,n}(x, y) = k B_{m+1,n}(x, y) + \frac{m}{x} B_{m,n}(x, y). \quad (\text{B.17})$$

Appendix C

Fourier Optics Code

Sample code with output for computing $I(r, z)$ for a kinoform with convergence angle $\beta = 4.4$ mrad and $m = 7$. The code takes about a minute to run on a modern laptop.

For the complete FACET OPTICS code package, visit the FACET GitHub page.

```
1 %% Computational Parameters
2
3 % Number of computation points
4 n_pts    = 2048;
5
6 % Parameters for plots
7 mid = n_pts/2 + 1;
8 window = 150;
9 low = mid-window;
10 high = mid+window;
11
12 %% All laser parameters in microns except where noted
13
14 res      = 8;          % resolution in image plane (microns)
15 R        = 20.0e3;     % illuminated radius of mask (microns)
16 nz       = 101;       % number of points in z
17 zmin     = 2.00e6;     % upstream z location
18 zmax     = 5.00e6;     % downstream z location
19
20 % longitudinal sample points
```

```

21 zs = linspace(zmin,zmax,nz);
22
23 % Ti:Saph parameters
24 lambda = 0.800;           % wavelength in microns
25 k       = 2*pi/lambda;    % microns^-1
26 t0      = 100e-15;        % s
27 e0      = 100e-3;         % J
28 r0      = 2.0;            % R/1e4; % radius in cm
29 I0      = e0/(t0*pi*r0^2); % on-optic intensity
30
31 %% Optics Parameters
32
33 % axicon/kinoform parameters
34 beta = 0.0044;           % Convergence Angle
35 kp   = k*sin(beta);      % microns^-1
36 m    = 7;                % azimuthal order (set to zero for axicon)
37
38 % illumination parameters
39 sig2 = r0*1e4;           % radius of top hat illumination in microns
40 ord  = 30;               % super-gaussian order of top hat
41
42 %% Fresnel Calculation (main loop)
43
44 i_max = zeros(nz,1);     % Maximum on axis intensity
45 lo_cen = zeros(nz,n_pts); % Central lineout
46 Energy = zeros(nz,1);    % total energy per z-slice
47
48
49 for i = 1:nz;
50     i
51     z = zs(i);           % propagation distance in microns
52     r_mask = lambda*z/(2*res); % radius of mask in microns
53
54     % Mask coordinates
55     x = linspace(-r_mask,r_mask,n_pts);
56     dx = x(2)-x(1);
57     xx = repmat(x,n_pts,1);

```

```

58     yy = rot90(xx,1);
59     r2 = xx.^2+yy.^2;
60     r  = sqrt(r2);
61     phi = atan2(yy,xx);
62     phi(round((n_pts/2+1)):n_pts,:) = phi(round((n_pts/2+1)):n_pts
        ,:)+2*pi;
63
64     % Axis at image plane
65     fx = linspace(-lambda*z/(2*dx),lambda*z/(2*dx),n_pts);
66     dfx = fx(2)-fx(1);
67     Dfx = fx(end)-fx(1);
68
69     % axicon phase
70     PHI = -kp*r+m*phi;
71
72     % Top hat illumination
73     top_hat_env = exp(-(r.^ord)/(2*(sig2)^ord));
74
75     %Field after mask
76     u_ideal = top_hat_env.*exp(1i*PHI);
77
78     % fresnel term
79     fresnel = exp(1i*k*z)*exp(1i*k*r2/(2*z))/(1i*lambda*z);
80
81     % fresnel diffraction
82     img_ideal = fft2(u_ideal.*fresnel);
83     img_ideal = dx^2*fftshift(img_ideal);
84     i_ideal = img_ideal.*conj(img_ideal);
85     i_ideal = abs(I0*i_ideal);
86
87     i_max(i) = i_ideal(mid,mid);
88     lo_cen(i,:) = i_ideal(mid,:);
89
90
91 end
92
93 %% Plots

```



```

94
95 figure(1);
96 surf(fx(low:high),zs/1e4,lo_cen(:,low:high)/1e13); shading flat;
    axis tight;
97 set(gca,'fontsize',13);
98 xlabel('R [\mu m]','fontsize',14);
99 ylabel('Z [cm]','fontsize',14);
100 zlabel('I [10^{13} W/cm^2]','fontsize',14);
101
102
103 figure(2);
104 contourf(fx(low:high),zs/1e4,lo_cen(:,low:high)); c = colorbar;
105 ylabel(c,'I [W/cm^2]');
106 set(gca,'fontsize',13);
107 xlabel('R [\mu m]','fontsize',14);
108 ylabel('Z [cm]','fontsize',14);

```

Figure C.1: Output of attached code.

Bibliography

- [1] S. Weinberg. The Crisis of Big Science. *New York Review of Books*, 5 2012.
- [2] J. W. Wang and G. A. Loew. Field emission and RF breakdown in high gradient room temperature linac structures. In *Frontiers of accelerator technology. Proceedings, Joint US-CERN-Japan International School, Hayama and Tsukuba, Japan, September 9-18, 1996*, pages 768–794, 1997.
- [3] V. A. Dolgashev, J. R. Lewandowski, D. W. Martin, S. G. Tantawi, S. P. Weathersby, and A. D. Yermian. Study of RF Breakdown in Normal Conducting Cryogenic Structure. *Proceedings, 3rd International Conference on Particle accelerator (IPAC 2012)*, C1205201:3368–3370, 2012.
- [4] M. C. Thompson, H. Badakov, A. M. Cook, J. B. Rosenzweig, R. Tikhoplav, G. Travish, I. Blumenfeld, M. J. Hogan, R. Ischebeck, N. Kirby, R. Siemann, D. Walz, P. Muggli, A. Scott, and R. B. Yoder. Breakdown limits on Gigavolt-per-meter electron-beam-driven wakefields in dielectric structures. *Physical Review Letters*, 100(21):214801, 5 2008.
- [5] Wikipedia. List of accelerators in particle physics — wikipedia, the free encyclopedia, 2016. [Online; accessed 7-August-2016].
- [6] M. Benedikt and F. Zimmermann. Status and Challenges of the Future Circular Collider Study. 1 2016. CERN-ACC-2016-0005.
- [7] E. G. Berezhko and D. C. Ellison. A Simple Model of Nonlinear Diffusive Shock Acceleration. *The Astrophysical Journal*, 526(1):385–399, 11 1999.
- [8] T. Tajima and J. M. Dawson. Laser Electron Accelerator. *Physical Review Letters*, 43(4):267–270, 7 1979.
- [9] D. Strickland and G. Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 56(3):219–221, 12 1985.
- [10] C. G. R. Geddes, C. S. Toth, J. Van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding. *Nature*, 431(7008):538–41, 9 2004.

- [11] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka. A laser-plasma accelerator producing monoenergetic electron beams. *Nature*, 431(7008):541–4, 9 2004.
- [12] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. Monoenergetic beams of relativistic electrons from intense laser-plasma interactions. *Nature*, 431(7008):535–8, 9 2004.
- [13] W. P. Leemans, B. Nagler, A. J. Gonsalves, Cs. Tóth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker. GeV electron beams from a centimetre-scale accelerator. *Nature Physics*, 2(10):696–699, 9 2006.
- [14] X. Wang, R. Zgadzaj, N. Fazel, Z. Li, S. A. Yi, X. Zhang, W. Henderson, Y.-Y. Chang, R. Korzekwa, H.-E. Tsai, C.-H. Pai, H. Quevedo, G. Dyer, E. Gaul, M. Martinez, A. C. Bernstein, T. Borger, M. Spinks, M. Donovan, V. Khudik, G. Shvets, T. Ditmire, and M. C. Downer. Quasi-monoenergetic laser-plasma acceleration of electrons to 2GeV. *Nature Communications*, 4:1988, 1 2013.
- [15] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, Cs. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey. Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime. *Physical Review Letters*, 113(24):245002, 12 2014.
- [16] Z. Huang, Y. Ding, and C. B. Schroeder. Compact x-ray free-electron laser from a laser-plasma accelerator using a transverse-gradient undulator. *Physical Review Letters*, 109(20):204801, 11 2012.
- [17] R. D. Ruth, A. W. Chao, P. L. Morton, and P. B. Wilson. A Plasma Wake Field Accelerator. *Part. Accel.*, 17:171, 1985.
- [18] P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas. Acceleration of Electrons by the Interaction of a Bunched Electron Beam with a Plasma. *Physical Review Letters*, 54(7):693–696, 2 1985.
- [19] J. B. Rosenzweig, D. B. Cline, B. Cole, H. Figueroa, W. Gai, R. Konecny, J. Norem, P. Schoesow, and J. Simpson. Experimental Observation of Plasma Wake-Field Acceleration. *Physical Review Letters*, 61(1):98–101, 7 1988.
- [20] J. B. Rosenzweig. Nonlinear plasma dynamics in the plasma wake-field accelerator. *Physical Review Letters*, 58(6):555–558, 2 1987.

- [21] J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su. Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields. *Physical Review A*, 44(10):R6189–R6192, 11 1991.
- [22] M. J. Hogan, R. Assmann, F.-J. Decker, R. Iverson, P. Raimondi, S. Rokni, R. H. Siemann, D. Walz, D. Whittum, B. Blue, C. E. Clayton, E. Dodd, R. Hemker, C. Joshi, K. A. Marsh, W. B. Mori, S. Wang, T. Katsouleas, S. Lee, P. Muggli, P. Catravas, S. Chattopadhyay, E. Esarey, and W. P. Leemans. E-157: A 1.4-m-long plasma wake field acceleration experiment using a 30 GeV electron beam from the Stanford Linear Accelerator Center Linac. *Physics of Plasmas*, 7(5):2241, 5 2000.
- [23] B. E. Blue, C. E. Clayton, C. L. O’Connell, F.-J. Decker, M. J. Hogan, C. Huang, R. Iverson, C. Joshi, T. C. Katsouleas, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, R. Siemann, and D. Walz. Plasma-wakefield acceleration of an intense positron beam. *Physical Review Letters*, 90(21):214801, 2003.
- [24] M. J. Hogan, C. D. Barnes, C. E. Clayton, F. J. Decker, S. Deng, P. Emma, C. Huang, R. H. Iverson, D. K. Johnson, C. Joshi, T. Katsouleas, P. Krejcik, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, C. L. O’Connell, E. Oz, R. H. Siemann, and D. Walz. Multi-GeV energy gain in a plasma-wakefield accelerator. *Physical Review Letters*, 95(5):054802, 7 2005.
- [25] I. Blumenfeld, C. E. Clayton, F.-J. Decker, M. J. Hogan, C. Huang, R. Ischebeck, R. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, E. Oz, R. H. Siemann, D. Walz, and M. Zhou. Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator. *Nature*, 445(7129):741–4, 2 2007.
- [26] M. J. Hogan, T. O. Raubenheimer, A. Seryi, P. Muggli, T. Katsouleas, C. Huang, W. Lu, W. An, K. A. Marsh, W. B. Mori, C. E. Clayton, and C. Joshi. Plasma wakefield acceleration experiments at FACET. *New Journal of Physics*, 12(5):055030, 5 2010.
- [27] K. Yonehara. Recent progress of RF cavity study at Mucool Test Area. *Journal of Physics: Conference Series*, 408:012062, 2 2013.
- [28] T. C. Chiou, T. Katsouleas, C. Decker, W. B. Mori, J. S. Wurtele, G. Shvets, and J. J. Su. Laser wakefield acceleration and optical guiding in a hollow plasma channel. *Physics of Plasmas*, 2(1):310–318, 1 1995.
- [29] W. D. Kimura, H. M. Milchberg, P. Muggli, X. Li, and W. B. Mori. Hollow plasma channel for positron plasma wakefield acceleration. *Physical Review Special Topics - Accelerators and Beams*, 14(4):041301, 4 2011.

- [30] S. Gessner, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. P. Delahaye, J. Frederico, S. Z. Green, C. Hast, M. J. Hogan, C. Joshi, C. A. Lindstrøm, N. Lipkowitz, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, B. OShea, N. Vafaei-Najafabadi, D. Walz, V. Yakimenko, and G. Yocky. Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator. *Nature Communications*, 7:11785, 6 2016.
- [31] S. J. Gessner. A Demonstration of the Hollow Channel Plasma Wakefield Accelerator. In *International Partical Accelerator Conference*, Busan, Korea, 2016. THPPA01.
- [32] R. Keinigs and M. E. Jones. Two-dimensional dynamics of the plasma wakefield accelerator. *Physics of Fluids*, 30(1):252, 1987.
- [33] I. Blumenfeld. *Scaling of the longitudinal electric fields and transformer ratio in a non-linear Plasma Wakefield Accelerator*. PhD thesis, Stanford University, 2009.
- [34] C.B. Schroeder. *Plasma-based accelerator structures*. PhD thesis, University of California, Berkeley, 1999.
- [35] W. Lu, C. Huang, M. Zhou, M. Tzoufras, F. S. Tsung, W. B. Mori, and T. Katsouleas. A nonlinear theory for multidimensional relativistic plasma wave wakefields. *Physics of Plasmas*, 13(5):056709, 5 2006.
- [36] John David Jackson. *Classical Electrodynamics Third Edition*. Wiley, 1999.
- [37] S. Corde, E. Adli, J. M. Allen, W. An, C. I. Clarke, B. Clausse, C. E. Clayton, J. P. Delahaye, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, N. Vafaei-Najafabadi, D. Walz, and V. Yakimenko. High-field plasma acceleration in a high-ionization-potential gas. *Nature Communications*, 7:11898, 6 2016.
- [38] E. Adli, J.-P. Delahaye, S. J. Gessner, M. J. Hogan, T. Raubenheimer, W. An, C. Joshi, and W. Mori. A Beam Driven Plasma-Wakefield Linear Collider: From Higgs Factory to Multi-TeV. 8 2013. arXiv:1308.1145.
- [39] J. B. Rosenzweig, P. Schoessow, B. Cole, C. Ho, W. Gai, R. Konecny, S. Mtingwa, J. Norem, M. Rosing, and J. Simpson. Demonstration of electron beam self-focusing in plasma wake fields. *Physics of Fluids B: Plasma Physics*, 2(6):1376, 1990.
- [40] M. J. Hogan, C. E. Clayton, C. Huang, P. Muggli, S. Wang, B. E. Blue, D. Walz, K. A. Marsh, C. L. OConnell, S. Lee, R. Iverson, F.-J. Decker, P. Raimondi, W. B. Mori, T. C. Katsouleas, C. Joshi, and R. H. Siemann. Ultrarelativistic-Positron-Beam Transport through Meter-Scale Plasmas. *Physical Review Letters*, 90(20):205002, 5 2003.

- [41] P. Muggli, B. E. Blue, C. E. Clayton, F. J. Decker, M. J. Hogan, C. Huang, C. Joshi, T. C. Katsouleas, W. Lu, W. B. Mori, C. L. OConnell, R. H. Siemann, D. Walz, and M. Zhou. Halo Formation and Emittance Growth of Positron Beams in Plasmas. *Physical Review Letters*, 101(5):055001, 7 2008.
- [42] P. Sprangle, E. Esarey, and A. Ting. Nonlinear interaction of intense laser pulses in plasmas. *Physical Review A*, 41(8):4463–4469, 4 1990.
- [43] P. Mora and T. M. Antonsen, Jr. Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas. *Physics of Plasmas*, 4(1):217, 1 1997.
- [44] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas. Nonlinear theory for relativistic plasma wakefields in the blowout regime. *Physical Review Letters*, 96(16):165002, 4 2006.
- [45] M. Litos, E. Adli, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. P. Delahaye, R. J. England, A. S. Fisher, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, N. Vafaei-Najafabadi, D. Walz, G. White, Z. Wu, V. Yakimenko, and G. Yocky. High-efficiency acceleration of an electron beam in a plasma wakefield accelerator. *Nature*, 515(7525):92–95, 11 2014.
- [46] S. Wilks, J. M. Dawson, T. C. Katsouleas, and J. J. Su. Beam Loading Efficiency in Plasma Accelerators. *Part. Accel.*, 22:81–99, 1987.
- [47] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. Beam loading by electrons in nonlinear plasma wakes. *Physics of Plasmas*, 16(5):056705, 5 2009.
- [48] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. Beam loading in the nonlinear regime of plasma-based acceleration. *Physical Review Letters*, 101(14):145002, 10 2008.
- [49] M. Litos, E. Adli, J. M. Allen, W. An, C. I. Clarke, S. Corde, C. E. Clayton, J. Frederico, S. J. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, M. Schmeltz, N. Vafaei-Najafabadi, and V. Yakimenko. 9 GeV energy gain in a beam-driven plasma wakefield accelerator. *Plasma Physics and Controlled Fusion*, 58(3):034017, 3 2016.
- [50] C. Huang, V. K. Decyk, M. Zhou, W. Lu, W. B. Mori, J. H. Cooley, T. M. Antonsen, B. Feng, T. Katsouleas, J. Vieira, and L. O. Silva. QuickPIC: a highly efficient fully parallelized PIC code for plasma-based acceleration. *Journal of Physics: Conference Series*, 46(1):190–199, 9 2006.
- [51] W. An, V. K. Decyk, W. B. Mori, and T. M. Antonsen. An improved iteration loop for the three dimensional quasi-static particle-in-cell algorithm: QuickPIC. *Journal of Computational Physics*, 250:165–177, 10 2013.

- [52] R. G. Hemker. *Particle-In-Cell Modeling of Plasma-Based Accelerators in Two and Three Dimensions*. PhD thesis, University of California, Los Angeles, 2000. arXiv:1503.00276.
- [53] K. V. Lotov. Acceleration of positrons by electron beam-driven wakefields in a plasma. *Physics of Plasmas*, 14(2):023101, 2 2007.
- [54] S. Lee, T. Katsouleas, P. Muggli, W. B. Mori, C. Joshi, R. Hemker, E. S. Dodd, C. E. Clayton, K. A. Marsh, B. Blue, S. Wang, R. Assmann, F. J. Decker, M. Hogan, R. Iverson, and D. Walz. Energy doubler for a linear collider. *Physical Review Special Topics - Accelerators and Beams*, 5(1):011001, 1 2002.
- [55] X. Wang, P. Muggli, T. Katsouleas, C. Joshi, W. B. Mori, R. Ischebeck, and M. J. Hogan. Optimization of positron trapping and acceleration in an electron-beam-driven plasma wakefield accelerator. *Physical Review Special Topics - Accelerators and Beams*, 12(5):051303, 5 2009.
- [56] S. Corde, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, J. P. Delahaye, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, N. Lipkowitz, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, M. Schmeltz, N. Vafaei-Najafabadi, D. Walz, V. Yakimenko, and G. Yocky. Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield. *Nature*, 524(7566):442–445, 8 2015.
- [57] C. B. Schroeder, D. H. Whittum, and J. S. Wurtele. Multimode Analysis of the Hollow Plasma Channel Wakefield Accelerator. *Physical Review Letters*, 82(6):1177–1180, 2 1999.
- [58] B. D. O’Shea. *Gigavolt-per-Meter Wakefields in Annular Dielectric Structures*. PhD thesis, University of California, Los Angeles, 2014.
- [59] T. C. Chiou and T. Katsouleas. High Beam Quality and Efficiency in Plasma-Based Accelerators. *Physical Review Letters*, 81(16):3411–3414, 10 1998.
- [60] S. Lee, T. Katsouleas, R. G. Hemker, E. S. Dodd, and W. B. Mori. Plasma-wakefield acceleration of a positron beam. *Physical Review E*, 64:045501, 2001.
- [61] A. A. Sahai and T. C. Katsouleas. Optimal positron-beam excited plasma wakefields in Hollow and Ion-Wake channels. 12 2015. arXiv:1512.08013.
- [62] J. Thomas, I. Yu. Kostyukov, J. Pronold, A. Golovanov, and A. Pukhov. A non-linear theory for the bubble regime of plasma wake fields in tailored plasma channels. 10 2015. arXiv:1510.09012.
- [63] A. W. Chao. *Physics of collective beam instabilities in high energy accelerators*. Wiley, New York, 1993.

- [64] J. Fan, E. Parra, I. Alexeev, K. Y. Kim, H. M. Milchberg, L. Ya. Margolin, and L. N. Pyatnitskii. Tubular plasma generation with a high-power hollow Bessel beam. *Physical Review E*, 62(6):R7603–R7606, 12 2000.
- [65] J. W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill, 1996.
- [66] N. E. Andreev, S. S. Bychkov, V. V. Kotlyar, L. Ya. Margolin, L. N. Pyatnitskii, and P. G. Serafimovich. Formation of high-power hollow Bessel light beams. *Quantum Electronics*, 26(2):126–130, 2 1996.
- [67] S. Z. Green, E. Adli, C. I. Clarke, S. Corde, S. A. Edstrom, A. S. Fisher, J. Frederico, J. C. Frisch, S. Gessner, S. Gilevich, P. Hering, M. J. Hogan, R. K. Jobe, M. Litos, J. E. May, D. R. Walz, V. Yakimenko, C. E. Clayton, C. Joshi, K. A. Marsh, N. Vafaei-Najafabadi, and P. Muggli. Laser ionized preformed plasma at FACET. *Plasma Physics and Controlled Fusion*, 56(8):084011, 8 2014.
- [68] R. M. Herman and T. A. Wiggins. Production and uses of diffractionless beams. *Journal of the Optical Society of America A*, 8(6):932, 6 1991.
- [69] S. Gessner. https://github.com/aard-pwfa/facet_optics, 2015. [Online; accessed 7-August-2016].
- [70] J. Sochacki, Z. Jaroszewicz, A. Kołodziejczyk, and S. Bará. Phase retardation of the uniform-intensity axilens. *Optics Letters*, 17(1):7, 1 1992.
- [71] L.V. Keldysh. Ionization in the Field of a Strong Electromagnetic Wave. *Sov. Phys. JETP*, 20(5):1307, 1965.
- [72] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field. *Sov. Phys. JETP*, 64:1191–1194, 1986.
- [73] D. L. Bruhwiler, D. A. Dimitrov, J. R. Cary, E. Esarey, W. Leemans, and R. E. Giacone. Particle-in-cell simulations of tunneling ionization effects in plasma-based accelerators. *Physics of Plasmas*, 10(5):2022, 4 2003.
- [74] FACET Conceptual Design Report. Technical report, SLAC National Accelerator Laboratory, 2009. SLAC-R-930.
- [75] N. Barov, J. B. Rosenzweig, M. E. Conde, W. Gai, and J. G. Power. Observation of plasma wakefield acceleration in the underdense regime. *Physical Review Special Topics - Accelerators and Beams*, 3(1):011301, 1 2000.

- [76] V. Yakimenko, I. V. Pogorelsky, I. V. Pavlishin, I. Ben-Zvi, K. Kutsche, Yu. Eidelman, T. Hirose, T. Kumita, Y. Kamiya, J. Urakawa, B. Greenberg, and A. Zigler. Cohesive acceleration and focusing of relativistic electrons in overdense plasma. *Physical Review Letters*, 91(1):014802, 7 2003.
- [77] A. Aschikhin, C. Behrens, S. Bohlen, J. Dale, N. Delbos, L. di Lucchio, E. Elsen, J.-H. Erbe, M. Felber, B. Foster, L. Goldberg, J. Grebenyuk, J.-N. Gruse, B. Hidding, Zhanghu Hu, S. Karstensen, A. Knetsch, O. Kononenko, V. Libov, K. Ludwig, A.R. Maier, A. Martinez de la Ossa, T. Mehrling, C.A.J. Palmer, F. Pannek, L. Schaper, H. Schlarb, B. Schmidt, S. Schreiber, J.-P. Schwinkendorf, H. Steel, M. Streeter, G. Tauscher, V. Wacker, S. Weichert, S. Wunderlich, J. Zemella, and J. Osterhoff. The FLASHForward facility at DESY. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 806:175–183, 1 2016.
- [78] E. Gschwendtner, E. Adli, L. Amorim, R. Apsimon, R. Assmann, A.-M. Bachmann, F. Batsch, J. Bauche, V.K. Berglyd Olsen, M. Bernardini, R. Bingham, B. Biskup, T. Bohl, C. Bracco, P.N. Burrows, G. Burt, B. Buttenschön, A. Butterworth, A. Caldwell, M. Cascella, E. Chevalley, S. Cipiccia, H. Damerau, L. Deacon, P. Dirksen, S. Doebert, U. Dorda, J. Farmer, V. Fedosseev, E. Feldbaumer, R. Fiorito, R. Fonseca, F. Friebe, A.A. Gorn, O. Grulke, J. Hansen, C. Hessler, W. Hofle, J. Holloway, M. Hüther, D. Jaroszynski, L. Jensen, S. Jolly, A. Joulaei, M. Kasim, F. Keeble, Y. Li, S. Liu, N. Lopes, K.V. Lotov, S. Mandry, R. Martorelli, M. Martyanov, S. Mazzoni, O. Mete, V.A. Minakov, J. Mitchell, J. Moody, P. Muggli, Z. Najmudin, P. Norreys, E. Öz, A. Pardons, K. Pepitone, A. Petrenko, G. Plyushchev, A. Pukhov, K. Rieger, H. Ruhl, F. Salveter, N. Savard, J. Schmidt, A. Seryi, E. Shaposhnikova, Z.M. Sheng, P. Sherwood, L. Silva, L. Soby, A.P. Sosedkin, R.I. Spitsyn, R. Trines, P.V. Tuev, M. Turner, V. Verzilov, J. Vieira, H. Vincke, Y. Wei, C.P. Welsch, M. Wing, G. Xia, and H. Zhang. AWAKE, The Advanced Proton Driven Plasma Wakefield Acceleration Experiment at CERN. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 2 2016.
- [79] FACET-II Conceptual Design Report. Technical report, SLAC National Accelerator Laboratory, 2015. SLAC-R-1067.
- [80] H. Wiedemann. *Particle accelerator physics: Third edition*. Springer Berlin Heidelberg, 2007.
- [81] R. Koontz, R. Miller, T. McKinney, and A. Wilmunder. SLAC Collider Injector, RF Drive Synchronization and Trigger Electronics and 15 AMP Thermionic Gun Development. *IEEE Transactions on Nuclear Science*, 28(3):2213–2215, 1981.
- [82] M.B. James. *Production Of High Intensity Electron Bunches for the SLAC Linear Collider*. PhD thesis, Stanford University, 1987. SLAC-R-319.

- [83] M. Sands. The Physics of Electron Storage Rings: An Introduction. *Conf.Proc.*, C6906161:257–411, 1970.
- [84] T. O. Raubenheimer. *The Generation and Acceleration of Low Emittance Flat Beams for Future Linear Colliders*. PhD thesis, Stanford University, 1991. SLAC-R-387.
- [85] K.L. Bane. Bunch Lengthening in the SLC Damping Rings, 1990. SLAC-PUB-5177.
- [86] R.L. Holtzapple. *Longitudinal Dynamics at the Stanford Linear Collider*. PhD thesis, Stanford University, 1996. SLAC-R-487.
- [87] F.J. Decker, T. Raubenheimer, and R. Holtzapple. Over compression, a method to shape the longitudinal bunch distribution for a reduced energy spread, 1994. SLAC-PUB-6604.
- [88] M. Cornacchia, R. Tatchyn, J. Schmerge, R. Carr, A. Wootton, J. Arthur, J. Galayda, I. Lindau, J. Safranek, P. Emma, L. Bentson, J. Stohr, and P. Krejcik. A Subpicosecond photon pulse facility for SLAC, 2001. SLAC-PUB-8950.
- [89] J.E. Clendenin. High Yield Positron Systems for Linear Colliders. *Conf.Proc.*, C8903201:1107, 1989.
- [90] J.E. Clendenin, A. Breakstone, A. Odian, R. H. Helm, M.C. Ross, J.B. Truher, B.A. Barnett, A.V. Kulikov, C.G. Yao, S. D. Ecklund, G. Bartha, H.C. DeStaebler, and L.P. Keller. SLC Positron Source Startup, 1988. SLAC-PUB-4704.
- [91] Y. Nosochkov, A. Seryi, R. Erickson, M. Hogan, C. Spencer, W. Wittmer, L. Bentson, N. Li, and J. Seeman. Optics Design for FACET, 2009. SLAC-PUB-13603.
- [92] K. L. F. Bane and P. Emma. LiTrack: A Fast longitudinal phase space tracking code with graphical user interface. *Particle accelerator. Proceedings, Conference, PAC’05, Knoxville, USA, May 16-20, 2005*, C0505161:4266, 2005.
- [93] P. Tenenbaum. Lucretia: A Matlab-Based Toolbox for the Modelling and Simulation of Single-Pass Electron Beam Transport Systems. *11th International Conference on Elastic and Diffractive Scattering: Towards High Energy Frontiers: The 20th Anniversary of the Blois Workshops, 17th Rencontre de Blois (EDS 05) Chateau de Blois, Blois, France, May 15-20, 2005*, C0505161:4197, 2005.
- [94] A. Scheinker and S. Gessner. Adaptive method for electron bunch profile prediction. *Physical Review Special Topics - Accelerators and Beams*, 18(10):102801, 10 2015.
- [95] P. Muggli, K.A. Marsh, S. Wang, C.E. Clayton, S. Lee, T.C. Katsouleas, and C. Joshi. Photoionized lithium source for plasma accelerator applications. *IEEE Transactions on Plasma Science*, 27(3):791–799, 6 1999.

- [96] J. Seeman, W. O. Brunk, Richard A. Early, M. C. Ross, E. Tillmann, and Dieter R. Walz. SLC Energy Spectrum Monitor using Synchrotron Radiation. In *Proceedings, 1986 Linear Accelerator Conference*, pages 3–6, 1986.
- [97] R. Akre, P. Emma, P. Krejcik, and L. Bentson. Bunch length measurements using a transverse RF deflecting structure in the SLAC linac, 2002. SLAC-PUB-9241.
- [98] G. Berden, W. A. Gillespie, S. P. Jamison, E.-A. Knabbe, A. M. MacLeod, A. F. G. van der Meer, P. J. Phillips, H. Schlarb, B. Schmidt, P. Schmüser, and B. Steffen. Benchmarking of Electro-Optic Monitors for Femtosecond Electron Bunches. *Physical Review Letters*, 99(16):164801, 10 2007.
- [99] E. Adli, S.J. Gessner, S. Corde, M.J. Hogan, and H.H. Bjerke. Cherenkov light-based beam profiling for ultrarelativistic electron beams. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 783:35–42, 5 2015.
- [100] M. Knott, D. Gurd, S. Lewis, and M. Thuot. EPICS: A control system software co-development success story. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 352(1-2):486–491, 12 1994.
- [101] M. Rivers. areaDetector: <http://cars9.uchicago.edu/software/epics/areaDetector.html>. [Online; accessed 7-August-2016].
- [102] C. L. OConnell, C. D. Barnes, F.-J. Decker, M. J. Hogan, R. Iverson, P. Krejcik, R. Siemann, D. R. Walz, C. E. Clayton, C. Huang, D. K. Johnson, C. Joshi, W. Lu, K. A. Marsh, W. Mori, M. Zhou, S. Deng, T. Katsouleas, P. Muggli, and E. Oz. Plasma production via field ionization. *Physical Review Special Topics - Accelerators and Beams*, 9(10):101301, 10 2006.
- [103] N. Vafaei-Najafabadi, W. An, C. E. Clayton, C. Joshi, K. A. Marsh, W. B. Mori, E. C. Welch, W. Lu, E. Adli, J. Allen, C. I. Clarke, S. Corde, J. Frederico, S. J. Gessner, S. Z. Green, M. J. Hogan, M. D. Litos, and V. Yakimenko. Evidence for high-energy and low-emittance electron beams using ionization injection of charge in a plasma wakefield accelerator. *Plasma Physics and Controlled Fusion*, 58(3):034009, 3 2016.
- [104] J. S. T. Ng, P. Chen, H. Baldi, P. Bolton, D. Cline, W. Craddock, C. Crawford, F. J. Decker, C. Field, Y. Fukui, V. Kumar, R. Iverson, F. King, R. E. Kirby, K. Nakajima, R. Noble, A. Ogata, P. Raimondi, D. Walz, and A. W. Weidemann. Observation of Plasma Focusing of a 28.5 GeV Positron Beam. *Physical Review Letters*, 87(24):244801, 11 2001.

- [105] S. Steinke, J. van Tilborg, C. Benedetti, C. G. R. Geddes, C. B. Schroeder, J. Daniels, K. K. Swanson, A. J. Gonsalves, K. Nakamura, N. H. Matlis, B. H. Shaw, E. Esarey, and W. P. Lee-mans. Multistage coupling of independent laser-plasma accelerators. *Nature*, 530(7589):190–193, 2 2016.