Alternative approaches to rapid acceleration of ion beams – harmonic ratcheting for fast RF acceleration and laser driven acceleration of gas jet targets

A Dissertation Presented

by

Nathan Michael Cook

to

The Graduate School

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

in

Physics

Stony Brook University

December 2014

Stony Brook University

The Graduate School

Nathan Michael Cook

We, the dissertation committee for the above candidate for the

Doctor of Philosophy degree, hereby recommend

acceptance of this dissertation

Axel Drees - Dissertation Advisor
Professor, Department of Physics and Astronomy

Peter W. Stephens - Chairperson of Defense
Professor, Department of Physics and Astronomy

Steve Peggs - Committee Member
Professor, Department of Physics and Astronomy

Michael Zingale - Committee Member
Professor, Department of Physics and Astronomy

Brian D. Sheehy - External Member
Physicist, Brookhaven National Laboratory

Igor Pogorelsky - External Member
Physicist, Brookhaven National Laboratory

This dissertation is accepted by the Graduate School

Charles Taber
Dean of the Graduate School

Abstract of the Dissertation

Alternative approaches to rapid acceleration of ion beams – harmonic ratcheting for fast RF acceleration and laser driven acceleration of gas jet targets

by

Nathan Michael Cook

Doctor of Philosophy

in

Physics

Stony Brook University

2014

Energetic ion beams have vast potential in medicine, energy, and basic science, providing significant advantages in applications of radiation therapy, nuclear energy, and high energy physics. Conventional acceleration means are inefficient and costly, imposing stringent requirements on space, power, and speed of the machines designed to address these applications. This thesis considers two contrasting approaches to improving ion beam acceleration: fast acceleration using radio-frequency (RF) technology and laser driven acceleration of ion beams from over-dense plasma.

We first consider fast acceleration in a synchrotron using conventional RF cavities. We introduce a ferrite based RF scheme for a rapid cycling synchrotron known as "harmonic ratcheting." By systematically decreasing the harmonic number in steps during the acceleration cycle, a reduction in the required frequency range is

achieved. Two cavities alternately provide the accelerating voltage to allow tuning. A ratcheting approach allows for a doubling of gap voltage for fixed cavity length and input power. Simulations performed using a 65 m synchrotron design demonstrate the feasibility of the scheme for acceleration of C^{6+} ions to 400 MeV/nucleon at a 15 Hz repetition rate.

Next, we investigate the acceleration of ions through the interaction of an intense CO_2 laser and over-dense plasma. Brookhaven National Laboratory's Accelerator Test Facility possesses a TW-class CO_2 laser, with the unique capability to produce a single intense 5 ps pulse at 10 μm. This allows for the use of high purity gas jet targets at densities which exceed the critical density of the laser light. We demonstrate the repeatable acceleration of ions using a two pulse technique. A pre-pulse containing a few % of the main pulse energy arrives 25 ns prior to the main pulse, driving a hydrodynamic blast wave into the gas target. The main pulse then drives an electrostatic shock into the shaped plasma, producing ion beams with higher peak energies than predicted by other approaches. We observe accelerated beams only for a narrow range of pre-pulse energies, indicating the importance of the target density profile in enabling the acceleration.

To my parents and to my grandparents.

Contents

List of Figures

x

xiii

xvii

List of Tables

Acknowledgements

My path through graduate school has been long and strange, and I have been greatly helped by the contributions of many people along the way. I am sure to miss some people, but everyone mentioned or missed here has helped me along the way.

I'd like to begin by giving special thanks to my two advisors, Axel Drees and Todd Satogata. Todd is the reason I embarked on an accelerator physics track, and despite moving to Jefferson Lab midway through my tenure, he has managed to keep in touch with me and provide last minute advice on everything from conference works to papers. Axel showed tremendous faith in me by selflessly taking me on as his student following Todd's departure despite countless responsibilities with his own research and in the University Provost's office.

I also cannot ignore the considerable support I have received from my "host" at BNL, Dejan Trbojevic. I am grateful for his enthusiasm and encouragement throughout my time with the Collider-Accelerator department, as well as his eagerness to help me and include me in all manner of interesting projects. Also from C-AD, Steve Peggs provided me with great insight into how to think scientifically and how to improve my writing. His support was invaluable when preparing my first paper. From the Accelerator Test Facility, I must thank Igor Pogorelsky for the excellent treatment I have had at ATF, and for the strong support and respect he has shown me at every turn. Thanks as well to Vitaly Yakimenko for initially permitting me to participate in ATF research. I'd like to also thank Ilan Ben-Zvi for the unwavering support he has shown to students of accelerator physics, and for his willingness to provide guidance and answer questions at any time. Ilan's leadership is instrumental to any success, past, present, and future, of Stony Brook's accelerator science education initiative.

I am especially grateful for the help that I received from the postdocs and scientists at the ATF who contributed to the ion generation experiment. Olivier Tresca, Misha Polyanskiy, Nick Dover, and Chakra Maharjan were all exceptionally patient with me, willing to help me learn, and responsive

Chapter 1

Introduction

Charged particle beams are a fundamental tool in modern physics, enabling the investigation of high energy particle physics and nuclear physics. To this end, particle accelerators are designed to provide beams tailored to the needs of a given scientific campaign. Recent achievements in accelerator science have, on one hand, emphasized ion beam interactions at increasing peak energies, as in the case of the Large Hadron Collider, while others have stressed improvements in electron beam quality and precision, as in the case of next generation light sources. Nevertheless, there remain many applications within science, medicine, and industry for which controlled, beams of protons and ions are uniquely suited. These applications have given rise to a need for new accelerator solutions to which the next generation of accelerators will be suited.

The interest in improved production of controlled ion beams is motivated in large part by their specific energy loss properties in matter. At MeV energies, ionization losses dominate the interaction of energetic charged particles with matter. The energy loss per unit distance, or stopping power, for an incident particle with speed β, Lorentz factor γ, and charge z, is given by the Bethe formula [16]

$$\frac{dE}{dx} \sim K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \left(\frac{2 m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} \right) - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \qquad (1.1)$$

where T_{max} is the maximum energy transfer in a single collision, Z is the charge, A is the atomic mass, I is the mean excitation energy, and δ is a density correction for the target material, respectively. Equation 1.1 reveals that ionization energy remains rather constant at high energies, but grows dramatically for small velocities. For heavy charged particles, the resultant energy deposition curve is distinguished by a Bragg peak at the end of the particle's range. Due to their small mass, m_e, electrons are relativistic even at modest energies, and as such they do not exhibit a peak in their energy loss. Instead they lose energy at a nearly constant rate as they traverse the target [17]. Photons in the therapeutic range of a few hundred keV to several MeV, on the other hand, lose energy primarily through Compton scattering, and to a lesser effect pair production and photoelectric effects, resulting in an exponential decline in the linear energy transfer (LET) dE/dx to the material with distance.

Figure 1.1 shows the energy loss curves for various particles. It should be noted that ions also demonstrate sustained deposition beyond the end of their range, due to nuclear fragmentation, which can be a concern for some high precision applications. Nonetheless, the property of localized energy transfer inherent to heavy charged particles makes possible a number of valuable ap-
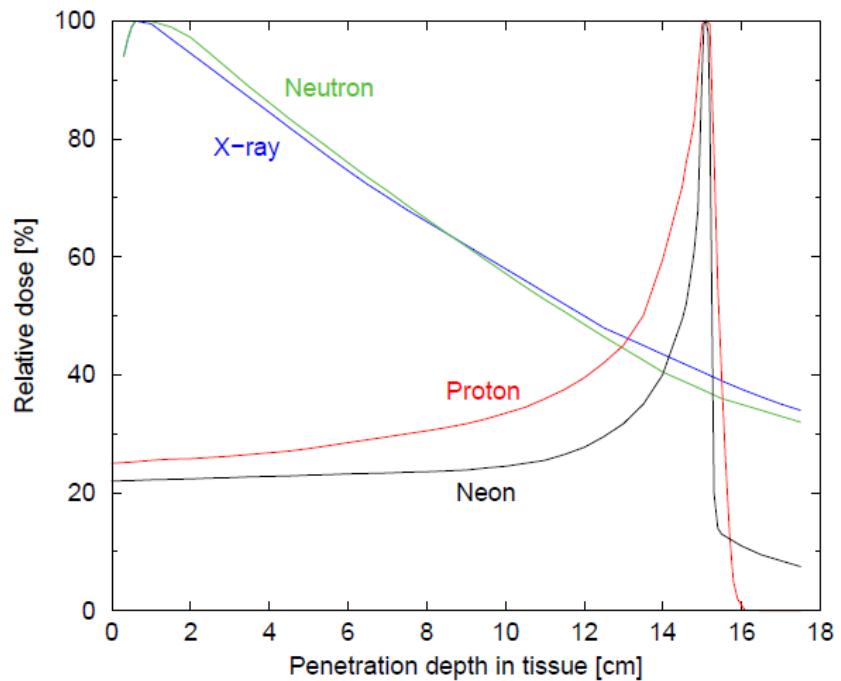
Figure 1.1: The corresponding dose-depth curves for various charged particles. Ions exhibit a Bragg Peak in their energy deposition near the end of their range [1].

plications in both medicine and energy. These applications can be separated on the basis of beam requirements into two categories: those requiring high precision, and those requiring high beam power.

1.1 High precision applications for ion beams in medicine

Ionizing radiation has been a major component of tumor treatment, as both an alternative and companion to surgical techniques, for more than 50 years. With worldwide cancer rates rising and projected to increase further, there is a considerable need for better and more efficient treatments. Particle accelerators offer the promise of targeted radiation therapy, making use of the Bragg peak feature of ion beam energy loss to localize a specified quantity of radiation (a "dose") to a specified position. Advancements in accelerator technology have made possible the controlled delivery of beams at time scales that are reasonable for treatment. The first hospital based proton therapy facility opened at Loma Linda University Medical Center (LLUMC) in 1990; today there are more than a dozen in operation in the US, more than 50 in the world, and many more in various stages of planning [18].

Protons have garnered considerable interest as they are lighter than other ions, and possess the highest possible charge to mass ratio Z/A, which increases their range in the patient for a given particle energy, thus reducing the acceleration required for therapeutic application. Most machines are designed to deliver beam up to ≈ 27 cm into a patient, corresponding to peak proton energies of ~ 206 MeV. This provides a significant cost advantage to proton machines relative to heavier ions. Nonetheless, these benefits are leveraged by a relative reduction in both efficacy and accuracy of treatment. The former is characterized by the "relative biological effecitiveness" (RBE), which quantifies the ratio of the absorbed dose between a reference radiation source (typically either 250 keV X-rays or Co^{60} gamma rays) and that of a specified source which correspond to the same survival fraction in a given tissue set. RBE is thus a disease-specific metric, and values for a given particle will vary with the specific application [19]. Generally, protons exhibit high RBE values for many types of treatment, while light ions such as helium, carbon, and neon are particularly effective for radiation resistant tumors [20].

Treatment accuracy is also difficult to define, but physical processes such as transverse scattering and longitudinal straggling provide insight into the limitations of a given beam in isolating a specific volume within a body. Protons undergo greater range straggling than heavier ions as a result of their lower

Application	Facility/proposal	Species	Beam energy [GeV/u]	Beam power [kW]
Therapy	iRCMS [23]	p^+	$0.07 - 0.33$	< 50
	iRCMS [23]	C^{6+}	$0.08 - 0.4$	< 50
Spallation	ESS [24]	p^+	2.5	5×10^3
ADSR	Jacobs [25]	p^+	1	$3 - 4 \times 10^3$
Isotopes	FRIB [26]	$p^+ \to^{238} U$	0.2	< 400

Table 1.1: Approximate ion beam parameters for various applications based on planned facilities.

mass, and similarly exhibit a higher degree of transverse multiple Coloumb scattering. This scattering provides a fundamental limitation on the accuracy of treatment at depth in tissue [21], and favors the use of heavier ions for many tumor locations. However, the absence of a fragmentation tail in the proton dose-depth curve enables precise, aggressive radiation doses for tumors anterior to vital organs.

Aside from radiation therapy, the properties of the Bragg peak can also be employed for imaging purposes. While x-ray attenuation lengths are on the order of centimeters, ion beams may travel much further. For a thin target, ion beams will exit the target with minimal energy loss, which can be quantified via 1.1 for a given material if the particle energy and target density and composition are known. Alternatively, information about the target can be extracted by sending beams of known energy through an unknown target, as with standard x-ray imaging. This is the premise of proton radiography, which is very promising for medical imaging. Proton radiography has also been extended for the purpose of measuring transient electromagnetic fields in high density plasmas [22].

1.2 Applications with high current and high energy

The applications discussed above rely on precise control of the beam energy, size, and focus, but require only small total beam flux. There remain, however, many additional applications of ion beams for which high current and high beam energy is a prerequisite. At highest energies, these applications fall mostly in the realm of high energy and nuclear physics, and the accelerators designed for these purposes are primarily large colliders. However, GeV scale ion beams are also extremely attractive for generating isotope beams [26], as spallation neutron sources, and for nuclear energy applications.

Spallation sources produce low energy neutrons, which are useful probes of material under a wide range of conditions. In contrast to charged particles, neutrons do not generate strong ionization losses along their path, acting instead to probe nuclear structure. As a result, they can be used to study sensitive materials without depositing large quantities of heat, making them especially well suited to reveal the structure of biological compounds and polymers. Spallation sources are designed as user-facilities; as such they emphasize high beam flux and stable operation, which in turn places requirements on the proton beam driving spallation [24].

Proton accelerators may also be used as fast neutron source to power a subcritical nuclear reactor. At GeV energies, protons impinging targets within the reactor core can produce fast neutrons with MeV kinetic energies. This additional neutron production mediates the reactor energy output, further safeguarding the core from meltdown. These accelerator driven subcritical reactors (ADSR) require high beam power, to within 1% of the desired reactor output power depending upon the operating criticality point. Moreover, beam interruptions induce considerable stress on the mechanical components of the reactor, as they trigger variations in heat output [27]. Thus, ADSRs offer considerable promise, but also pose design challenges which remain open.

An additional application in nuclear energy is the concept of fusion by fast ignition (FI) [28]. A compressed fuel pellet is subjected to a localized, intense source of radiation, sufficient to trigger fusion. This approach is thought to lessen the requirements on energy delivery in comparison to fusion by inertial confinement. Proton beams may be an attractive means of delivering the focused energy necessary to trigger fusion, and various schemes have been proposed which could even utilize laser accelerated protons [29].

1.3 Motivating alternative acceleration methods

While there is tremendous potential in these ion beam applications, the beam parameters as outlined in Table 1.1 impose steep requirements not simply on energy but on beam current and efficiency. For the past 80 years, accelerators have progressed exceptionally quickly, with peak particle energies increasing by a factor of 10 every 6-8 years [30]. However, this period appears to be coming to an end, due to the physical limitations of conventional radio frequency (RF) technology to produce higher sustained voltages for increasingly demanding applications. Moreover, the size and cost of state-of-the-art accelerators designed to explore these applications continue to increase. In response

to these challenges, a broad array of alternative acceleration techniques are being investigated with the goal of providing more efficient and better scaling acceleration for ion beams in the MeV range. These strategies may employ conventional RF accelerator technology, or they may employ entirely different physical interactions. This thesis considers two distinct, novel approaches to acceleration of ion beams with the long term goal of providing feasible solutions to these challenges.

In the first part of the thesis, we consider the problem of operating a small scale synchrotron using radio frequency cavity technology to achieve fast and efficient acceleration of ions across a range of energies required for medical applications. Chapter 2 provides the basic framework for a synchrotron, including a discussion of transverse dynamics, lattice design, and the longitudinal dynamics which support the design and operation of such a machine.

In Chapter 3.6, we consider parametrization and efficacy of conventional RF cavities in designing an RF ramp needed to achieve multi-Hz fast repetition rates in an ion synchrotron, and discuss the challenges associated with large dynamic range in tuning these cavities. We introduce a novel approach, harmonic ratcheting, to using ferrite-loaded RF cavities to achieve this rapid cycling condition with greater efficiency than can be achieved using a standard single bucket acceleration scheme [31].

For the second part of this thesis, we consider the frontier of particle acceleration using intense lasers. An intense laser incident on a material can ionize it, creating a plasma supporting large electric and magnetic fields. If the plasma is transparent to radiation at the laser wavelength, transmission of the incident fields produces a trailing wake in the plasma, capable of capturing and accelerating small emittance beams of electrons to high energies over a short distance. However, in the case of a laser incident on plasma which is overdense or opaque to the laser, coupling of the laser energy to the target may trigger a variety of response, capable of generating large electrostatic fields, which may subsequently accelerate ions.

In Chapter 4, we review the basic physics of laser plasma interactions which are pertinent to this regime. We will discuss the coupling of laser energy to the plasma through a variety of heating mechanisms, while also surveying the known mechanisms of laser driven ion acceleration and their specific strengths and weaknesses.

In Chapter 5, we consider the experimental setup at Brookhaven National Laboratory's Accelerator Test Facility (ATF) which has been designed to investigate ion generation in overdense plasmas using a unique, high intensity mid-IR ($\lambda \approx 9-11\mu$m) CO_2 amplified laser pulse. Particular attention is given to the in-situ optical plasma diagnostics, which take advantage of the relatively

7

low density gas targets to provide spatially resolved density information about the plasma at different times during the plasma evolution.

Chapter 6 discusses the nuances of particle detection using a scintillator based imaging diagnostic, and presents the results of a study to better characterize the performance of several candidate scintillators under beam conditions.

Finally, Chapter 7 presents some of the results from the ongoing experiment at the ATF to accelerate ions in an overdense plasma through the combination of an intense laser and gas jet target. The hydrodynamic evolution of the target gas under the incidence of an intense controllable pre-pulse is first discussed. Then, observations of ion acceleration and plasma evolution under different pre-pulse conditions are presented. These observations are supported by particle-in-cell simulations, which confirm the generation of a collisionless electrostatic shock capable of reflecting ions to MeV energies. Observed ion energies and spectral properties are found to be in good agreement with these simulations. Simulations also provide strong evidence of electron heating and filamentation commensurate with experimental results. In all, these results emphasize the importance of target shaping to achieve consistent ion acceleration. We then preview some potential future optimizations of this work which aim to improve target shaping methods and take advantage of upgrades to the CO_2 laser peak power.

Chapter 2

Periodic Accelerators

Modern accelerators conventionally fall into one of two categories: linear or circular. In a linear accelerator (linac), a particle makes a single pass through the beamline, whereas in a circular accelerator, a particle makes many passes, each pass imparting an incremental amount of the particle's final momentum. The decision to construct a linear or circular machine depends on the output beam requirements as well as the constraints of cost, space, and application. Generally speaking, linacs are utilized for lower energy applications which emphasize beam quality over peak energy. Additionally, for high energy electron accelerators, linacs have a significant advantage in minimizing the radiation electrons emit when accelerated radially (aptly named "synchrotron radiation"). The radiation for a given particle scales as

$$P_{rad} \propto \frac{E^4}{\rho^2} \propto \frac{1}{m^4} \tag{2.1}$$

where E is beam energy, ρ the bend radius, and m the mass of the radiating particle [32]. Due to the dependence on mass, synchrotron radiation is particularly significant for electron beams. Reducing power loss from this radiation requires increasing the machine radius, which can be costly and inefficient. An example of a prominent linac includes the Stanford Linear Accelerator, the longest linear accelerator in the world. The 3 km linac, originally designed for high energy physics applications, now provides beam for a host of applications, including plasma wakefield applications as well as providing the drive electron beam for the Linear Coherent Light Source (LCLS).

Conversely, circular accelerators are useful when high beam energy or total beam current is needed. Circular accelerators are effective at both low energies, where hadron cyclotrons (circular accelerators with static magnetic fields) can be particularly efficient, and at high energies, where particle velocities do not change considerably with additional energy increase. Thus they are ubiquitous in high energy science, with prominent examples including the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC). Both require millions of turns to accelerate particles to their collision energy, and need to maximize total beam current in the process, making synchrotrons the only feasible alternative. The downside is that at in-between energies, the revolution frequency of the particles changes very rapidly, requiring RF sources to rapidly adjust their operational frequency to remain synchronous with particles. Nonetheless, synchrotrons remain omnipresent at the cutting edge of medicine and nuclear applications of the kind described in Chapter 1.

Regardless of type, all accelerators share the same basic principles. Particles are accelerated by phase synchronous application of electromagnetic waves, overwhelmingly in the radio frequency region. As will be shown, phase

stable acceleration groups particles into bunches, which are steered along the lattice via (at minimum) a combination of dipole and quadrupole magnets. The arrangement of these magnets, cavities, and instrumentation comprises the accelerating lattice.

2.1 Transverse Dynamics

A synchrotron is a periodic accelerator with a fixed lattice through which the particles must maintain a (nearly) constant trajectory per turn. This requirement necessitates that the electric and magnetic fields used to accelerate, focus, and guide the beam be adjusted synchronously with the beam's increasing momentum. We first consider the transverse dynamics of a particle in a synchrotron, as determined by linear dipole and quadrupole magnetic field elements.

A particle of mass m_0 and charge q in an electromagnetic field experiences a force as described by the Lorentz equation,

$$\frac{d\boldsymbol{p}}{dt} = q\left[\boldsymbol{E}(\boldsymbol{r}_i, t) + \boldsymbol{v}_i \times \boldsymbol{B}(\boldsymbol{r}_i, t)\right] \tag{2.2}$$

In absence of an electric field, and assuming a magnetic field is applied perpendicular to the particle motion, the ideal particle instead revolves in a plane perpendicular to $\boldsymbol{B}$ with radius ρ. This centripetal motion can be equated to the particle momentum via

$$qvB = \frac{\gamma m_0 v^2}{\rho} \tag{2.3}$$

resulting in a formula for the radius and revolution frequency for the particle

$$\rho = \frac{p}{qB} \tag{2.4}$$

$$\omega = \frac{qB}{\gamma m_0} \tag{2.5}$$

Equation 2.4 is often given in terms of the product $B\rho$, also known as the magnetic rigidity of the particle. $B\rho$ provides a meaningful scaling for a given synchrotron design. For a given particle type, particle momentum, and radius of orbit, the required magnetic field is

$$B[T] = 3.34 \frac{p[GeV/c]}{Z\rho[m]} \tag{2.6}$$

where Z is the charge state of the particle. Equation 2.6 describes the fundamental relationship between particle energy and magnetic field, which must be maintained at all times in order to follow a fixed trajectory circular accelerator such as a synchrotron.

2.1.1 Equations of Motion
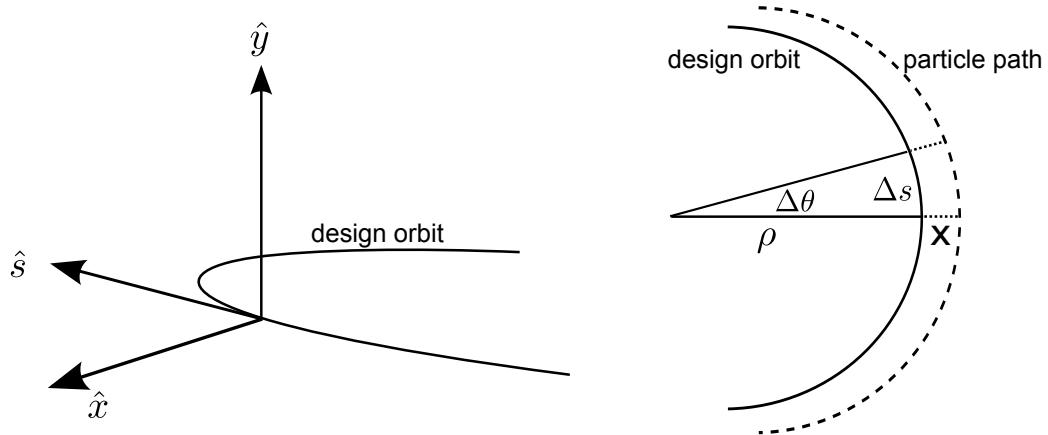
Figure 2.1: Coordinate system in the frame of reference of a particle travelling an arbitrary trajectory relative to the design path.

Now consider the equations of motion describing the periodic motion of a particle in a circular accelerator. First, specify a coordinate system based on the design orbit of the accelerator, which maintains a fixed radius ρ about its center. Then define the coordinate s by the longitudinal progression of the particle along the design trajectory. The transverse motion is described by x and y coordinates, in which x lies in the plane of the design path, such that the actual distance from center is given by $r = \rho + x$. Figure 2.1 demonstrates this coordinate system. It can furthermore be shown that s satisfies

$$ds = \rho d\theta \tag{2.7}$$

$$v_s = r\frac{dr}{d\theta} \tag{2.8}$$

$$\frac{ds}{dt} = v_s\frac{\rho}{r} \tag{2.9}$$

where v_s is the particle velocity along its individual (e.g. non-ideal) trajectory. Again we assume zero electric field, and consider the magnetic field effects as the particle is guided in a circle. Transformation of Equation 2.2 into this

12

frame yields a pair of differential equations governing the transverse motion.

$$x'' = -\frac{B_y}{B\rho}\left(1+\frac{x}{\rho}\right)^2 + \frac{\rho+x}{\rho^2} \tag{2.10}$$

$$y'' = \frac{B_x}{\beta\rho}\left(1+\frac{x}{\rho}\right)^2 \tag{2.11}$$

Although the principal magnetic field at work is the (constant) dipole field B_0, satisfying Equation 2.6, we can expand the fields about s, assuming that there is some first order variation in the fields, corresponding to a focusing field. This is a realistic assumption, as all dipole magnets produce fringe fields at their ends which contribute to higher order effects. Thus we expand

$$B_y(s) = B_0 + \frac{dB(s)}{dx}x \tag{2.12}$$

$$B_x(s) = \frac{dB(s)}{dy}y \tag{2.13}$$

We can then reduce the above equations to

$$x'' + \left[\frac{1}{\rho^2} + \frac{1}{\beta\rho}\frac{dB(s)}{dx}\right]x = 0 \tag{2.14}$$

$$y'' - \frac{1}{\beta\rho}\frac{dB(s)}{dy} = 0 \tag{2.15}$$

These represent the equations of motion for transverse oscillations in a circular accelerator, assuming constant energy and linearized magnetic field gradients. We can express these equations generally as a form of Hill's equation

$$\frac{d^2x}{ds^2} + K(s)x = 0 \tag{2.16}$$

Hill's equation is difficult to solve for arbitrary $K(s)$. Generally speaking, this equation describes periodic motion of a harmonic oscillator (explicitly so for $K(s) = C > 0$), and so we should expect sinusoidal solutions. We can find explicit solutions if we assume a piecewise-constant $K(s) = K$, corresponding to constant fields for a given element. Assuming initial conditions $x(0) = x_0$

and $x'(0) = x'_0$, we find

$$x(s) = x_0 \cos\left(\sqrt{K}s\right) + \frac{x'_0}{\sqrt{K}} \sin\left(\sqrt{K}s\right) \tag{2.17}$$

$$x'(s) = -x_0\sqrt{K} \sin\left(\sqrt{K}s\right) + x'_0 \cos\left(\sqrt{K}s\right) \tag{2.18}$$

Equation 2.18 provides a useful solution for the motion of a particle through a quadrupole (or weak focusing dipole). The coupled relationship between $x(s)$ and $x'(s)$ forms the basis for transfer matrix solutions for a series of such elements [33]. For now, we consider a periodic solution to a lattice with s dependence, and so must impose an additional constraint, that $K(s) = K(s + C)$, where C is the circumference of the orbit. In this case, we look for a solution of the form

$$x(s) = Aw(s) \cos(\psi(s) + \delta) \tag{2.19}$$

where $w(s) = w(s + C)$ retains the periodicity of the lattice, but $\psi(s)$ need not be. This form satisfies Hill's equaion, with coefficients satisfying

$$\psi' = \frac{k}{w(s)^2} \tag{2.20}$$

$$k^2 = w(s)^3 \left(w(s)'' + Kw(s)\right) \tag{2.21}$$

Here, k is a constant of integration. The solution to $w(s)$ is expressed via the Courant-Synder, or Twiss, parameters

$$\beta(s) = \frac{w(s)^2}{k} \tag{2.22}$$

$$\alpha(s) = -\frac{1}{2}\frac{d\beta(s)}{ds} \tag{2.23}$$

$$\gamma(s) = \frac{1 + \alpha^2}{\beta} \tag{2.24}$$

The Twiss-parameters are a product of lattice design, and can be used to generally describe the periodic motion of a beam in that lattice [33]. Our solution from Equation 2.19 takes the form:

$$x(s) = \sqrt{2J\beta(s)} \sin(\psi(s) + \phi_0) \tag{2.25}$$

where J is known as the Courant-Synder Invariant (or action); it and ϕ_0 are constants of motion dependent on the initial conditions. The function $\beta(s)$ is termed the betatron ("beta") function, and provides information about the

amplitude and phase of particle oscillations in given lattice. The phase advance is given by $\psi(s)$, which follows from Equations 2.21 and 2.24 to give

$$\psi(s) = \int_0^s \frac{ds}{\beta(s)} \tag{2.26}$$

We define the betatron tune $\nu_{i=x,y}$ as the number of betatron oscillations that a particle makes during each orbit. For a circular machine, we have

$$\nu_i = \frac{1}{2\pi} \int \frac{ds}{\beta_i} \simeq \frac{R}{\langle \beta_i \rangle} \tag{2.27}$$

2.1.2 Emittance

The action J has a particularly important interpretation. It can be shown that for a given s, $x(s)$ and $x'(s)$ form an ellipse in phase space, given by

$$\gamma x^2 + 2\alpha x x' + \beta x'^2 = A \tag{2.28}$$

corresponding to the same value of J but different phase.

Figure 2.2 illustrates the properties of the ellipse. The maximum amplitudes in both x and x' are functions of J and the Twiss parameters. Moreover, the area of the ellipse is given by $A = 2\pi J$, signifying that the area in phase space is also a constant of motion. This result, which could have been surmised on the basis of Liouville's theorem, applies generally to the entire 6-D phase space comprising the particle distribution function in both transverse planes, as well as in energy and time [34]. Although these relations were derived for a set of particles with equal action, they can be applied to a beam in which a cutoff value in the distribution is chosen to contain a particular fraction of the beam (e.g. 95%). The phase space area filled by that portion of the beam is termed the beam emittance, ϵ.

Equation 2.25 then describes the maximum beam envelope in a lattice defined by the given Twiss parameters, and with initial emittance and phase given by J and ϕ_0. The emittance characterizes many important beam properties, including the beam size and divergence. For a Gaussian beam, we have

$$\sigma_i = \sqrt{\beta_i(s)\epsilon_i} \tag{2.29}$$

$$\theta_i = \sqrt{\frac{\epsilon_i}{\beta_i(s)}} \tag{2.30}$$

For an accelerating beam, the emittance is no longer invariant, and instead is naturally reduced by the reduction in divergence corresponding with longitudi-
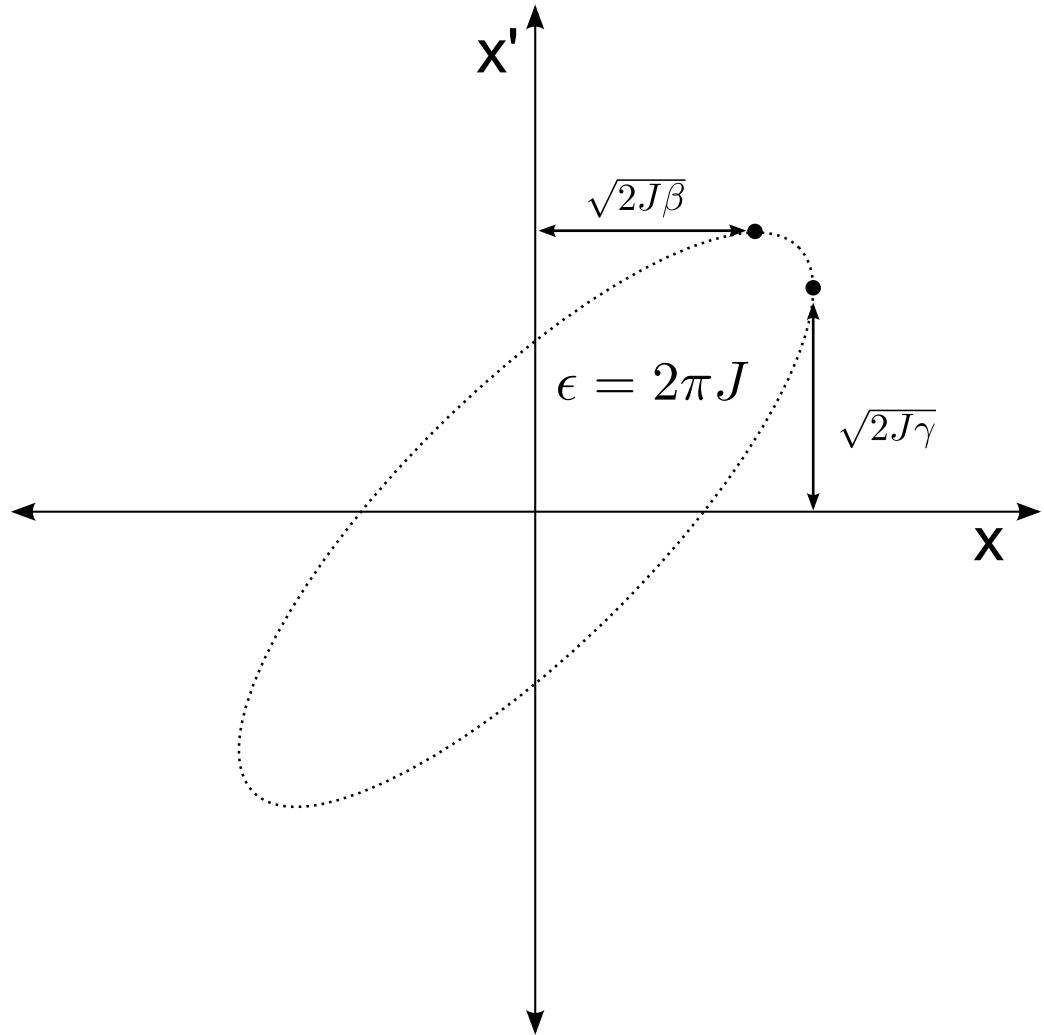
Figure 2.2: Particles with the same value of action J fall on an ellipse in phase space. Particle properties such as maximum displacement $x_{max} = \sqrt{2J\beta}$ and maximum momentum $x'_{max} = \sqrt{2J\gamma}$ are described using Twiss parameters.

nal acceleration. The emittance of the beam decreases linearly with increasing particle momentum, and as such a normalized emittance is defined according to

$$\epsilon_n = \epsilon \beta \gamma \tag{2.31}$$

where $\beta = v/c$ and γ are the relativistic Lorentz factors. We note that the product $\beta\gamma$ is proportional to particle momentum, and as such the normalized emittance is preserved during acceleration.

2.1.3 Dispersion

We have so far considered all particles in our lattice to have equal energy. Of course this is not a realistic assumption. Consider a particle with fractional momentum variation

$$\delta = \frac{\Delta p}{p_0} \tag{2.32}$$

from the design value. This momentum offset introduces an inhomogeneous correction to Hill's equation 2.16. The general solution follows from

$$x(s) = x_\beta(s) + D_x(s)\delta \tag{2.33}$$

where x_β, the homogeneous solution given by Equation 2.25, is offset by the value $D_x(s)\delta$. The function $D_x(s)$ is known as the dispersion function. The dispersion function satisfies the Hill equation

$$D'' + K(s)D = \frac{1}{\rho(s)} \tag{2.34}$$

for which the general solutions are quite complex, but can be estimated for individual elements using transfer matrices. Dispersion has the immediate consequence of increasing the beam size. Inclusion of a dispersion function $D_x(s)$ adjusts Equation 2.29 to be

$$\sigma_x(s) = \left[\frac{\epsilon_x \beta_x(s)}{\beta\gamma} + (D_x(s)\delta)^2 \right]^{1/2} \tag{2.35}$$

To minimize beam size, dispersion is often suppressed in relevant sections of a lattice, such as experimental interaction points, as well as injection and extraction points. This may be achieved by designing the bending sections of the accelerator using achromats, for which $D(s)$ and $D'(s)$ are zero at the entrance and exit. The simplest such arrangement is the double bend achromat, composed of a central quadrupole with symmetric dipoles on each side. The

quadrupole is designed to produce a 180° phase advance in the beam between the two dipoles, causing the contributions towards the dispersion from each to be equal and opposite. Variations on this design with additional bends and nonlinear elements are often used for practical reasons, but the basic idea remains to adjust quadrupole strengths such that the sum contribution of the dipoles to the dispersion is cancelled [32]. Single pass beam lines can be much more complex and still maintain their achromaticity [35].

2.1.4 Lattice Design

Modern accelerators are built upon the notion of strong focusing through the application of alternating quadrupole gradients in the transverse plane [36, 37]. This follows directly from the optical principle that a combination of thin focusing and defocusing lenses achieves a net focusing effect in both transverse planes. A simple form of this strong focusing principle is the FODO cell, comprised of a two quadrupoles, one focusing (F) and one defocusing (D), separated by drifts (O). The first quadrupole focuses in one plane while the other defocuses in that plane. For simplicity, the FODO cell is often described symmetrically by splitting the focusing magnet in half, producing a sequence of the form (1/2F - O - D - O - 1/2F). The symmetric FODO has the appealing property of producing a symmetric periodic betatron function, and is well described analytically using the thin lens approximation for magnet transfer elements. Figure 2.3 shows the β functions for an example symmetric FODO.

The FODO structure is used commonly in lattice bends as well, either with alternating quadrupoles spaced between dipole magnets, or by using combined function magnets. Iron-dominated combined function magnets are designed such that the two poles of a simple dipole are slanted, producing a field gradient that acts to focus or defocus the beam. Although combined function magnets cannot achieve fields quite as high as conventional normal conducting quadrupoles, they can reduce the total space required for a given bend.

Figure 2.4 shows a basic synchrotron lattice design for a ~ 55 m synchrotron, designed to accelerate both proton and fully stripped Carbon ions to therapeutic energies (270 MeV protons and 400 MeV/u C^{6+} ions). The accelerator uses a racetrack style design, with RF cavities along with injection and extraction septums grouped along the straights. Each lattice component is decomposed into a FODO cell structure. The straight sections are composed of two consecutive FODO lattices. Each arc consists of three dipole girders —each girder is separated into combined function magnets in a FODO arrangement with constant dipole field. The bending radius $\rho = 4.78$ m is minimized following Equation 2.6, subject to the achievable field strengths in the dipoles of $B_0 = 1.33$ T. The arcs are also designed to be achromatic, pro-
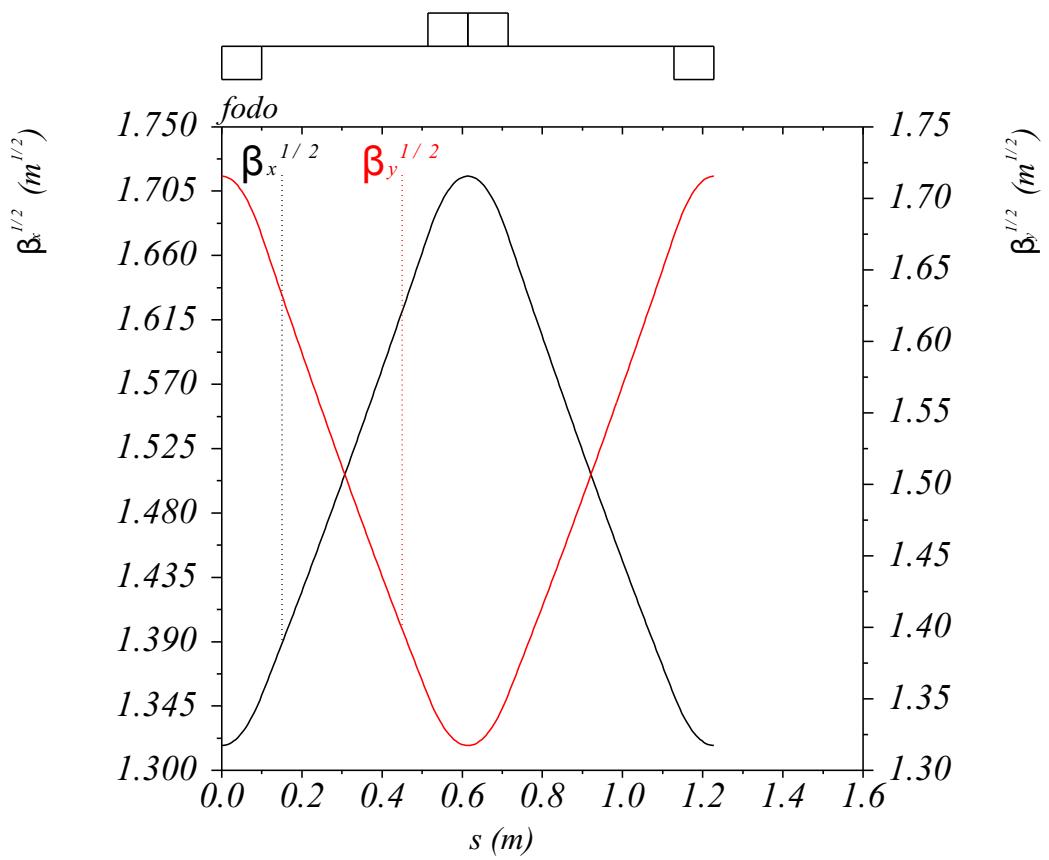
Figure 2.3: An example symmetric FODO, featuring a circular beam and zero dispersion.
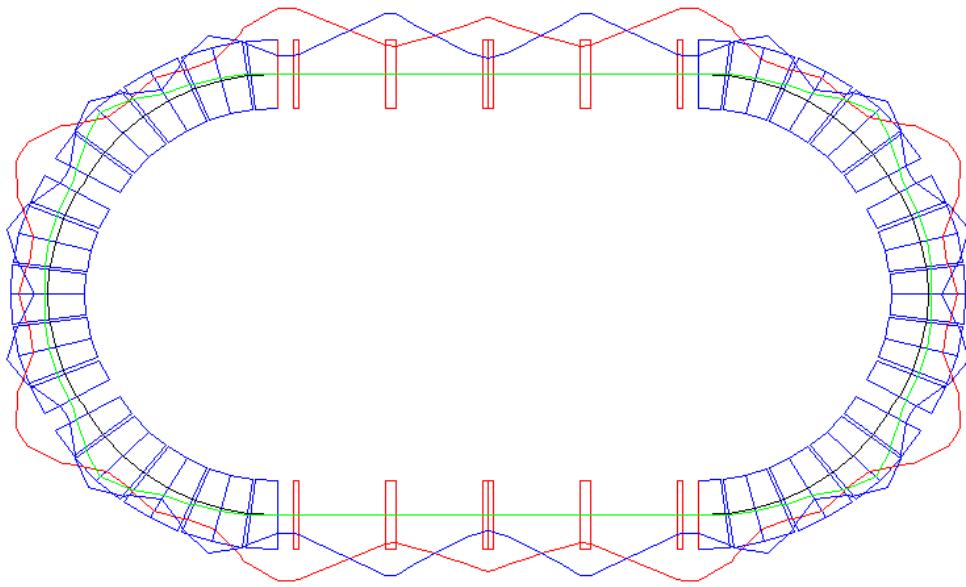
Figure 2.4: An entire lattice depicted graphically with parameters β_x (blue), β_y (red), and D_x(green) superimposed.

ducing zero dispersion in the straight sections. This minimizes beam size for injection, extraction and RF cavity passage. Nonetheless, the beam remains oblate in the x-plane as a result of the additional focusing in needed to reduce the dispersion.

2.2   Longitudinal Dynamics

In a simple linear lattice such as that described above, the transverse dynamics generally do not pose a problem to machine design. This is especially true for non-storage ring machines, which may retain the beam for many orders of magnitude fewer terms. For example, a gold ion beam in RHIC can make hundreds of millions of turns before being dumped, while a Carbon ion in a fast cycling synchrotron of the type discussed later only needs 72,000 turns to reach maximum energy. This significantly reduces the threat of resonances induced by tune-shift in the beam. Unfortunately, the relative ease of the lattice design is countered by the challenges posed by accelerating the beam quickly and efficiently in a such a small machine. We now review the basic longitudinal dynamics of particles in a circular accelerator to illustrate these difficulties.
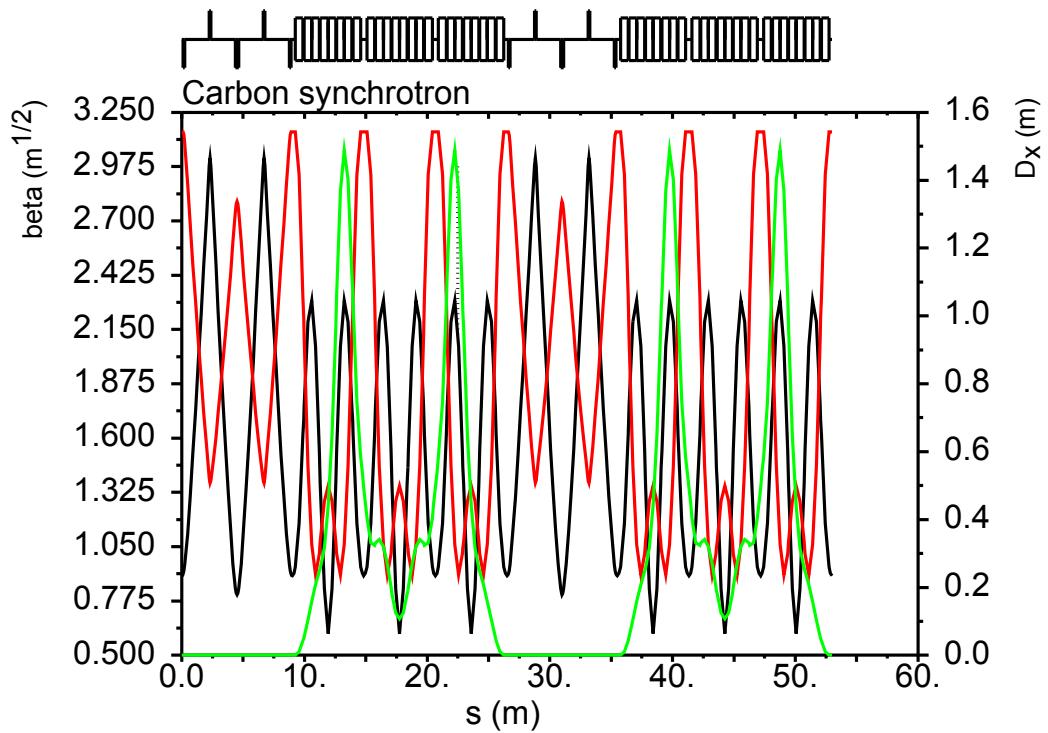
Figure 2.5: The lattice parameters β_x, β_y, and D_x are plotted versus s for the entire synchrotron.

2.2.1 Energy gain in a synchrotron

Consider a circular accelerator with circumference C. A particle of mass m and charge $q = Ze$ is injected with initial momentum p_0 and is accelerated to final momentum p_f, corresponding to a range in speeds

$$\beta_{min} \leq \beta(t) \leq \beta_{max} \tag{2.36}$$

The revolution frequency for an accelerating particle sweeps over the corresponding range

$$\frac{\beta_{min}\, c}{C} \leq f_{rev}(t) \leq \frac{\beta_{max}\, c}{C} \tag{2.37}$$

where c_0 is the speed of light and C is the synchrotron circumference. Acceleration is achieved through application of sinusoidally varying electric field

$$V_g(t) = V_0 sin(2\pi f_{RF} t + \phi_s) \tag{2.38}$$

where V_g specifies the voltage across an RF cavity, known as the "gap voltage", and ϕ_s is the synchronous phase, seen by the ideal particle which follows the design path of the synchrotron. We will assume for simplicity that the accelerating gap length is short compared to the distance travelled during one RF period, $g << C$. For optimum acceleration, the RF frequency is constrained to be an integer multiple of the revolution frequency,

$$f_{RF} = h\, f_{rev} \tag{2.39}$$

where h is the harmonic number. For simplicity, temporarily assume $h = 1$ for the following discussion. We can describe the one turn energy gain of the synchronous particle through

$$\delta U_s = qV_g \sin(\phi_s) \tag{2.40}$$

A test particle separated in time by $\delta t = t - t_s$ from the synchronous particle sees a different phase $\phi \neq \phi_s$ in the accelerating voltage. We introduce the variable

$$\varphi = \omega_{RF}\delta t = \phi - \phi_s \tag{2.41}$$

to describe the phase variation of this particle. The corresponding energy gain for the general particle is thus

$$\delta U = qV_g \sin(\phi) = qV_g \sin(\varphi + \phi_s) \tag{2.42}$$

The relative change in energy per turn between the synchronous and generic particle is

$$\Delta(\delta U) = qV_g \left[\sin(\varphi + \phi_s) - \sin(\phi_s) \right] \tag{2.43}$$

For small relative energy gains per turn, Equation 2.43 can be written as a differential equation

$$\frac{d(\delta U)}{dt} = \frac{qV_g}{2\pi} \omega_{RF} \left(\sin(\phi) - \sin(\phi_s) \right) \tag{2.44}$$

Before moving further, we make a change of variables from (t, U) to the system (φ, W)

$$W = \frac{-\delta U}{\omega_{RF}} \tag{2.45}$$

We can then rewrite Equation 2.44 in our new coordinates

$$\frac{dW}{dt} = \frac{qV_g}{2\pi} \left(\sin(\phi_s) - \sin(\varphi + \phi_s) \right) \tag{2.46}$$

2.2.2 Small amplitude longitudinal oscillations

Equation 2.46 describes the energy evolution in our new system. For a complete description of particle motion, we must also determine the variation of relative phase φ with time and W. Taking the particle revolution period to be given by Equation 2.37, it follows that the

$$\frac{dT}{T} = -\frac{d\omega}{\omega} = \left(\alpha_c - \frac{1}{\gamma^2} \right) \frac{dp}{p} \tag{2.47}$$

where α_c is the momentum compaction factor, which relates the path length difference between off momentum particles and the reference orbit. The quantity

$$\eta_c = \alpha - \frac{1}{\gamma^2} \tag{2.48}$$

is termed the "slip factor" because it relates the slip in time directly to the momentum offset. The momentum compaction factor relates to an energy

$$\gamma_t = \frac{1}{\sqrt{\alpha}} \tag{2.49}$$

which is known as the transition energy. As an accelerating particle passes this energy threshold, the slip factor changes sign from positive to negative; this can be heuristically explained as the particle behavior becoming relativistic. At

low energies, the time change is dominated by velocity perturbations, while at higher energies it is path length deviations arising from dispersion that dominate (Note: the slip factor may be presented with opposite sign [38], so care must be taken to note the order of terms in the expression). The momentum compaction factor is a function of lattice parameters, given by

$$\alpha = \frac{1}{C} \int \frac{D(s)}{\rho(s)} ds \tag{2.50}$$

The transition energy thus depends on the design parameters. In general, for a synchrotron the momentum compaction factor is inversely related to the horizontal tune $\alpha \approx \nu_x^{-2}$, and as such the transition energy goes as $\gamma_t \approx \nu_x$. A larger ring thus possesses a higher transition energy. Nonetheless, for ions, even small rings still require GeV-scale particle energies to approach transition.

Equation 2.47 can be re-written to describe the variation in phase coordinate φ with time. We have

$$\frac{d\varphi}{dt} \simeq -\frac{\omega_{RF}^2 \eta_c}{\beta^2 U_s} W \tag{2.51}$$

where substitution of W for $\mathrm{d}p/p$ follows from the relativistic energy-momentum relation. Equations 2.51 and 2.46 may be combined to form a single second order differential equation in φ describing the motion of the particle in phase space.

$$\frac{d^2\varphi}{dt^2} + \frac{h\omega_s^2 \eta_c q V_g}{2\pi \beta^2 U_s} \left(sin(\phi_s) - sin(\varphi + \phi_s) \right) = 0 \tag{2.52}$$

For small amplitude oscillations ($\varphi << \phi_s$), this equation simplifies to

$$\frac{d^2\varphi}{dt^2} + \frac{h\omega_s^2 \eta_c q V_g \cos(\phi_s)}{2\pi \beta^2 U_s} \varphi = 0 \tag{2.53}$$

which describes a linear oscillator with oscillation frequency

$$\Omega_s^2 = \frac{h\omega_s^2 \eta_c q V_g \cos(\phi_s)}{2\pi \beta^2 U_s} \tag{2.54}$$

Ω_s is the synchrotron oscillation frequency. As with transverse betatron oscillations, we define a corresponding synchrotron tune $Q_s = \Omega_s/\omega_s$. The synchrotron tune is usually relatively small, $Q_s < 1$, as particle energies generally well exceed the applied RF voltages, which are limited to the kV range. In order for the motion to be phase-stable about the synchronous phase, we must have $\Omega_s^2 > 0$, which is equivalent to requiring $\eta_c \cos(\phi_s) > 0$. Thus parti-

cles below transition must lie on the up-slope of the Voltage curve defined by $(0, \pi/2)$, while those above transition lie on the down-slope $(\pi/2, \pi)$. This coincides nicely with the description of momentum compaction described above.

2.2.3 Large Amplitude Oscillations and the Separatrix

In case of large amplitude oscillations, Equation 2.52 has no exact analytical solution, and numerical methods must be employed for arbitrary initial conditions. Thus we will first consider the simple case of the non-accelerating particle, for which $\phi_s = 0$. Equation 2.52 then has a solution of the form

$$\frac{1}{\Omega_s}\frac{d\phi}{dt} = \pm\sqrt{2(\cos(\phi) - \cos(\phi_m)}$$ (2.55)

where $\phi = \pm\phi_m$ corresponds to $d\phi/dt = 0$, the unstable fixed points where the particle path in phase space crosses the ϕ axis. If we choose $\phi_m = \pi$, then we arrive at special case solution

$$\frac{d\phi}{dt} = 2\Omega_s \cos\left(\frac{\phi}{2}\right)$$ (2.56)

Equation 2.56 describes the *separatrix*, which defines the boundary between stable and unstable orbits in phase space for a given Ω_s. Within the separatrix, stable solutions exist which are characterized by oscillation about ϕ_s as discussed in the previous section. The separatrix can be written in terms of the canonical variable W via Equation 2.51, and the corresponding dimensions can be calculated in meaningful units. The area within the separatrix is referred to as the *bucket area*, and follows from

$$A_s = 2\int_{-\phi_m}^{\phi_m} W \, d\phi$$ (2.57)

$$= 8\frac{R}{\pi hc}\left[\frac{qV_gU_s}{2\pi h|\eta_s|}\right]^{1/2}$$ (2.58)

The maximum permissible W is given by the bucket half-height, ΔW_s, follows, as does the bucket length $\Delta\phi$ via

$$\Delta W_s = \frac{A_s}{8}sin\left(\frac{\phi_m}{2}\right) = \frac{A_s}{8}$$ (2.59)

$$\Delta\phi = 2\phi_m = 2\pi$$ (2.60)

Equations 2.60 and 2.59 effectively define the momentum acceptance and maximum bunch length which can be injected into the bucket, respectively. They can also be applied to orbits of particles in the bucket with arbitrary maximum phase ϕ_m. Thus, the longitudinal emittance of a bunch can be deduced from the initial phase-space coordinates of the bunch, as the particles will follow phase space orbits with maxima described by Equations 2.59 and 2.60. Figure 2.6 shows a stationary bucket with various stable phase orbits.

As a corollary, observe that if a bunch which is injected into this bucket with an initial phase distribution that does not conform to a fixed emittance stable orbit, the bunch will spread out as the individual particles follow their own stable orbits. This results in filamentation of the bunch and an increase in emittance. Consider a bunch with zero momentum spread

$$\frac{dp_0}{p} = \delta W_0 = 0 \tag{2.61}$$

occupying the ϕ axis in phase space with a maximum at ϕ_m. Over time, the bunch will begin to filament, ultimately filling an area in phase space described by Equation 2.55, with corresponding maximum energy spread given by Equation 2.59. The end result is a distinct increase in total beam emittance. Figure 2.7 illustrates an example of a thin mismatched bunch injected into a stationary bucket.

For accelerating particles with arbitrary choice of ϕ_s, a numerical solution must be used. Nonetheless, analysis of the phase equation reveals some general properties of the separatrix. In particular, the unstable fixed points of the moving bucket along the phase axis are $(\phi_1, 0)$ and $(\phi_2, 0)$ where ϕ_1 and ϕ_2 satisfy

$$\phi_1 = \pi - \phi_s \tag{2.62}$$
$$0 = \cos(\phi_2) + \cos(\phi_s) + (\phi_2 + \phi_s - \pi)\sin(\phi_s) \tag{2.63}$$

Unlike the stationary case, the accelerating bucket is asymmetric about ϕ_s. The bucket half-height ΔW and bucket area A_{bk} are given by

$$\Delta W = \frac{A_s}{8\sqrt{2}} \left(\pi \sin(\phi_s) - 2\phi_s \sin \phi_s - 2\cos \phi_s\right)^{1/2} \tag{2.64}$$
$$A_{bk} = \alpha(\phi_s)A_{bk} \tag{2.65}$$
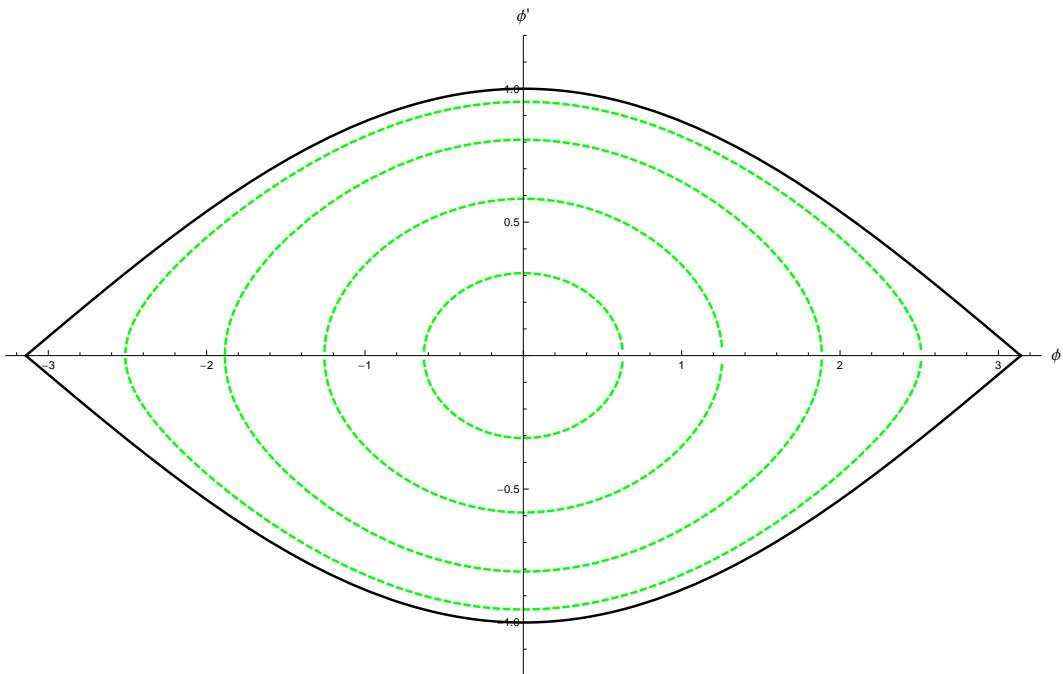
Figure 2.6: The stability of a bunch in an RF bucket is depicted in (ϕ, ϕ') phase space. The stationary RF bucket (solid black line) is accompanied by stable orbits (dashed green lines).
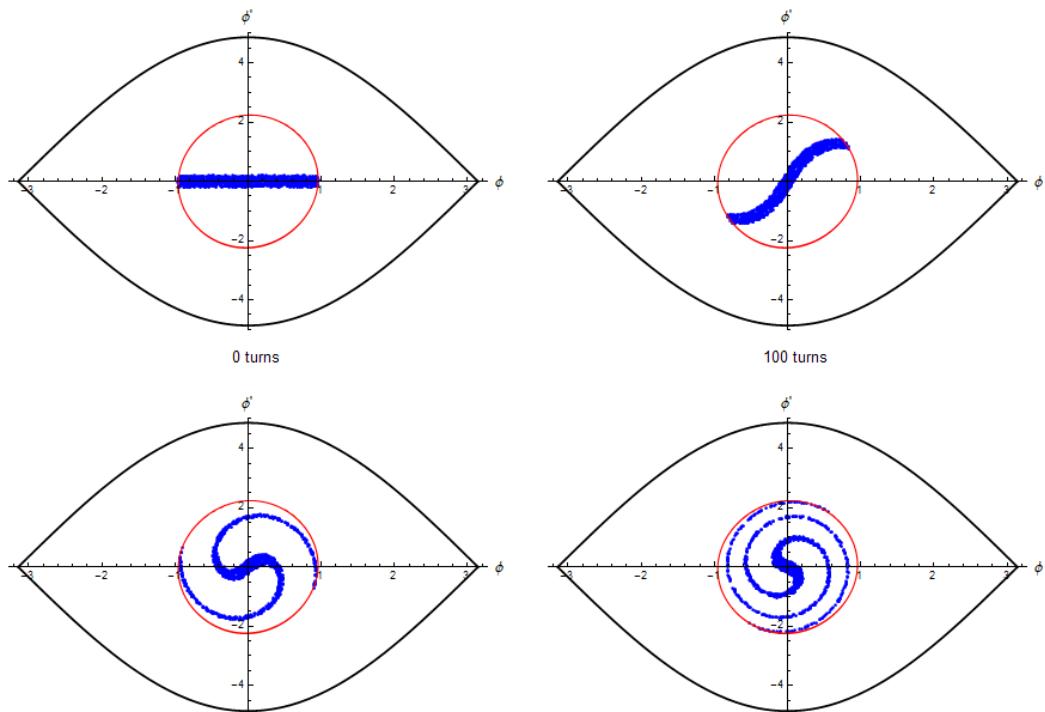
Figure 2.7: The injection of a bunch into a bucket is depicted through evolution in the particle phase space (ϕ, ϕ'), where ϕ is the relative phase coordinate. A "mismatched" bunch characterized by a long bunch with minimal momentum variation quickly filaments and eventually fills the phase region defined by the orbit of the particle with maximum emittance, depicted by the red line.

where $\alpha(\phi_s)$ is given by [39]

$$\alpha(\phi_s) = \frac{\sqrt{2}}{8} \int_{\phi_2}^{\phi_1} [(\pi - \phi_s - \phi)\sin\phi_s - \cos\phi - \cos\phi_s]^{1/2} \, d\phi$$
$$\approx \frac{1 - \sin\phi_s}{1 + \sin\phi_s} \tag{2.66}$$

Again note the role of applied voltage V_g in increasing the available bucket area, however when particles are accelerated, the bucket area is smaller by a factor of α_s. Thus, for a fixed cavity voltage, there is a clear trade off between bucket size and accelerating voltage. Synchronous phases in the range of $30° - 35°$ are often chosen as they provide a practical compromise. Figure 2.8 shows a moving bucket solution with $\phi_s = \pi/6$.
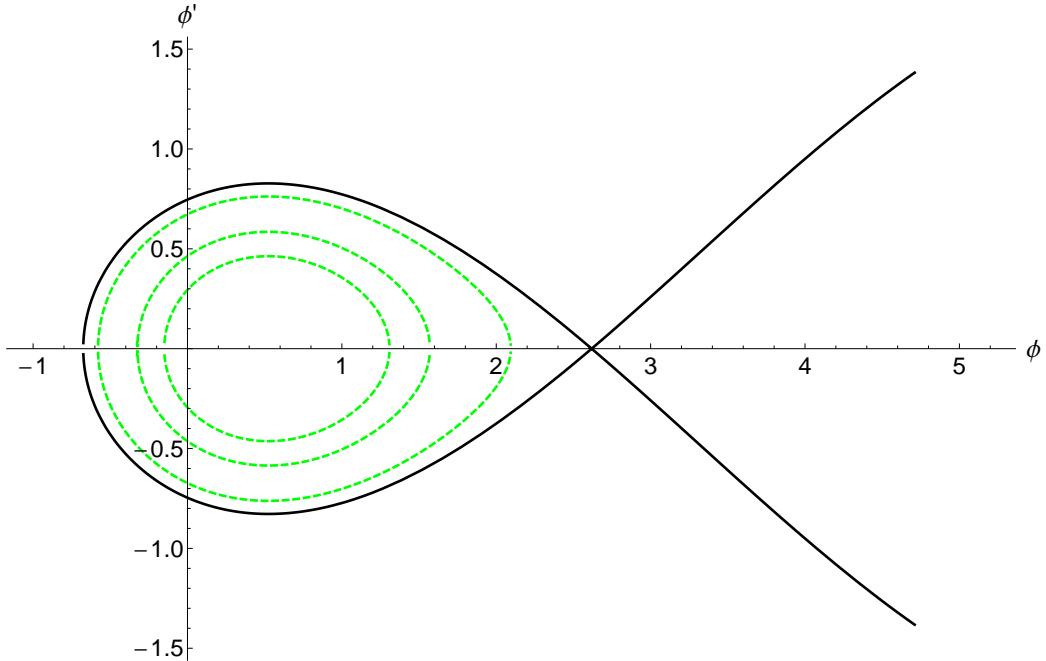
Figure 2.8: The accelerating bucket (solid black line) for synchronous phase $\phi_s = \pi/6$ along with stable orbits (green, dashed lines) is depicted in (ϕ, ϕ') phase space.

Chapter 3

Harmonic Ratcheting for Rapid Acceleration of Ions

Despite considerable improvements to RF technology and engineering, conventional acceleration of low β charged particles remains expensive and inefficient. Nonetheless, the potential for faster acceleration of low energy ion beams remains rich, with a host of applications proposed to take advantage of rapid cycling synchrotrons (RCS) or fixed field alternating gradient machines (FFAGs) [40–42]. The primary challenge of accelerating such particles is accommodating the large change in frequency during accelerating, which requires robust and flexible tuning regimens. Using speciality ferrites, a broad tuning range can be obtained at the cost of efficiency [43]. In many cases, the ferrite materials suffer dramatic loss effects when pushed to large biasing fields, rapid biasing changes, and large amplitude AC fluxes [44, 45]. These difficulties limit the advancement of fast and efficient accelerators for low energy ion beams.

This chapter will provide some insight into the performance metrics of ferrite-tuned RF cavities with an emphasis on the basics of cavity design and the limitations on performance at high repetition rates. A possible solution, harmonic ratcheting, is proposed as a means to limit the relative frequency range of cavity operation and to improve cavity efficiency for high voltage operation. The harmonic ratcheting scheme represents one possible solution to meeting the RF needs for a fast cycling machine for heavy ion therapy. To illustrate effectiveness of ratcheting, a realistic synchrotron design for carbon ions is considered, and an optimal ratcheting solution is presented. The discussion of the ratcheting approach, and its application to a medical synchrotron, is further discussed in Cook et al. [31].

3.1 Ferrite-loaded RF cavities

For conventional accelerators, RF cavities remain the only means of achieving the high energies (hundreds of MeV) needed for proton and ion applications. As noted previously, proton acceleration is challenging due to the rapid change in speed accompanying particle acceleration at low energy. This necessitates a frequency change which renders conventional cavities inefficient, a property which I will elucidate with a basic description of a conventional cavity.

3.1.1 RF cavity basics

The core purpose of an RF cavity is to deliver energy to the particles passing through it at a specified frequency. To deliver this energy, electromagnetic waves are driven at the desired frequency and intensity, localized to a region of particle passage. A cavity provides the boundary conditions along with thermal and electrical properties needed to support these fields.

Consider a cylindrically symmetric cavity of radius R and length L oriented along the z-axis. In a vacuum (or more generally a region devoid of charges and currents), Maxwell's equations reduce to the general wave equation for $\boldsymbol{E}$ and $\boldsymbol{B}$, respectively,

$$\nabla^2 \boldsymbol{E} - \frac{1}{c^2}\partial_t^2 \boldsymbol{E} = 0. \tag{3.1}$$

We will solve for the z-component of the electric field, as we wish to accelerate our particles in that direction. In the case of our cylindrical cavity, the wave equation reduces to

$$\frac{\partial^2 E_z}{\partial z^2} + \frac{1}{r}\frac{\partial}{\partial r}(r\frac{\partial E_z}{\partial r}) + \frac{1}{r^2}\frac{\partial^2 E_z}{\partial \phi^2} - \frac{1}{c^2}\partial^2{}_t \boldsymbol{E} = 0. \tag{3.2}$$

The time-dependent component of our solution is simply

$$E_t(t) = e^{i\omega t} \tag{3.3}$$

We may then propose a time-independent solution of the form

$$E_z(r, \phi, z) = R(r)\Phi(\phi)Z(z) \tag{3.4}$$

The z-dependent component is satisfied by $Z(z) = e^{-ikz}$ while the ϕ equation similarly reads $\Phi(\phi) = e^{-in\phi}$. We are left with a radial equation,

$$\frac{\partial^2 E_z}{\partial z^2} + \frac{1}{r}\frac{\partial E_z}{\partial r} + (\frac{\omega^2}{c^2} - k_z{}^2 - \frac{m^2}{r^2})E_z = 0. \tag{3.5}$$

Define the parameter k_c by

$$k_c^2 = \frac{\omega^2}{c^2} - k_z^2 \tag{3.6}$$

The solutions to this equation are given by Bessel functions J_n,

$$E_z = E_0 J_n((k_c r))e^{i(\omega t - n\phi - kz)}. \tag{3.7}$$

The cavity defines radial boundary conditions according to $J_n(k_c R) = 0$, meaning that the propagating modes are zeroes of the Bessel function solution. If we assume that the cavity is closed on both ends, we limit wave propagation to only standing wave modes, those with $E(z = 0, L) = 0$. Our solution is modified by the boundary condition $ik_z = \frac{\pi p}{L}$ for $p = 0, 1, \ldots$. We then have

a lowest order fundamental accelerating mode,

$$E_z = E_0 J_0 \left(k_c \frac{r}{R} \right) \cos \left(\omega_0 t \right) \tag{3.8}$$

$$\omega_0 = J_{01} \frac{c}{R} \tag{3.9}$$

where $J_{01} \simeq 2.405$ is the first zero of J_0. We refer to this mode as the TM_{010} ("transverse magnetic") mode. Figure 3.1 shows the radial version of this lowest order solution for E_z and B_ϕ.
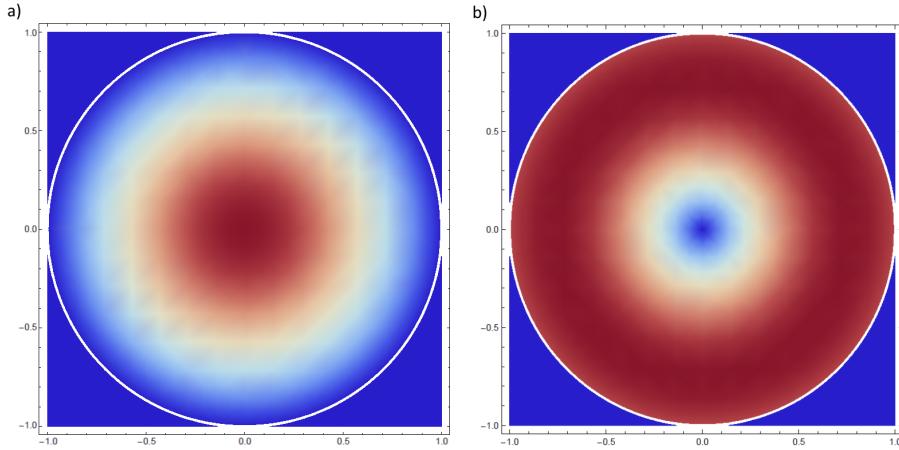
Figure 3.1: The fundamental accelerating mode TM_{01} is depicted according to (a) the field component E_z as a function of r, and (b) the corresponding magnetic field component B_ϕ. On axis, the electric field is at a maximum, while the magnetic field is zero.

To gain further insight into cavity performance, introduce a simple model of the cavity as an RLC circuit consisting of an inductor and capacitor in parallel. Assume that the inductor has some loss represented by a resistor in series. Figure 3.2(a) shows this "lumped circuit" model of a cavity.

The impedance of this circuit can be written as

$$Z(\omega) = \left(\frac{1}{Z_C} + \frac{1}{Z_R + Z_L} \right)^{-1} \tag{3.10}$$

The individual impedances can be written $Z_C = i/\omega C$, $Z_R = R$, and $Z_L = i\omega L$. We then have

$$Z(\omega) = \frac{R + i\omega L}{(1 - \omega^2 LC) + i\omega RC}. \tag{3.11}$$

We may then define the resonant frequency $\omega_0 = 1/\sqrt{LC}$ as it maximizes the
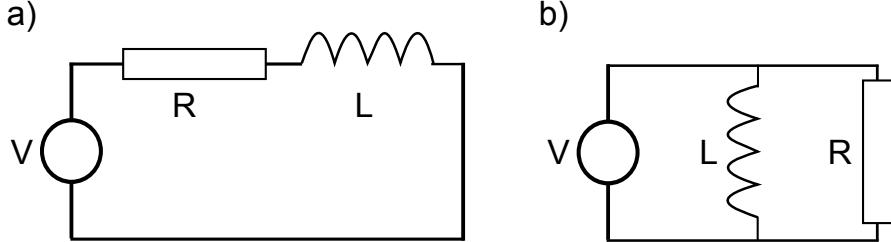
Figure 3.2: Two equivalent lumped circuit models use a driven RLC circuit to represent an RF cavity: (a) the series model and (b) parallel model. In each, the resistance is assumed to stem from the inductor. Transformations between each system can be performed assuming the impedance, and subsequently the quality factor, are equal.

impedance of the circuit, minimizing current flow to these elements (reducing them to zero in the $R = 0$ case). The real component, $\text{Re}(Z(\omega))$, defines a Lorentzian about ω_0 with full width at half maximum $\Delta\omega$. In the case of large impedance, $\Delta\omega/\omega \ll 1$ and we find

$$\frac{\Delta\omega}{\omega} = \frac{R}{\sqrt{L/C}}. \tag{3.12}$$

We'll introduce two figures of merit to better describe cavity performance. The first is the *quality factor* of cavity, defined by

$$Q = 2\pi \frac{\text{Stored Energy}}{\text{Energy lost per cycle}} = \omega \frac{U}{P} \tag{3.13}$$

Again we can use our RLC circuit model to better define Q. Assume we are driving our cavity at resonance with a current of the form $I = I_0 \cos(\omega t + \phi)$. In our model, energy is lost only to the resistor, and we have a time-averaged loss of $0.5I_0^2 R$. The total energy in the circuit is equal to the peak energy through the inductor, $U = 0.5I_0^2 L$. Thus we have

$$Q = \omega_0 \frac{L}{R} = \frac{\omega_0}{\Delta\omega} \tag{3.14}$$

We see that the quality factor is inversely related to the bandwidth of the cavity. For a particular resonance, greater efficiencies may be achieved at the cost of bandwidth and vice-versa. Thus, we conclude that efficient cavity performance is only achievable over a very narrow range of frequencies. Figure 3.3 demonstrates the this relationship for a sample circuit.

An additional figure of merit is the shunt impedance, R_p, of the cavity.
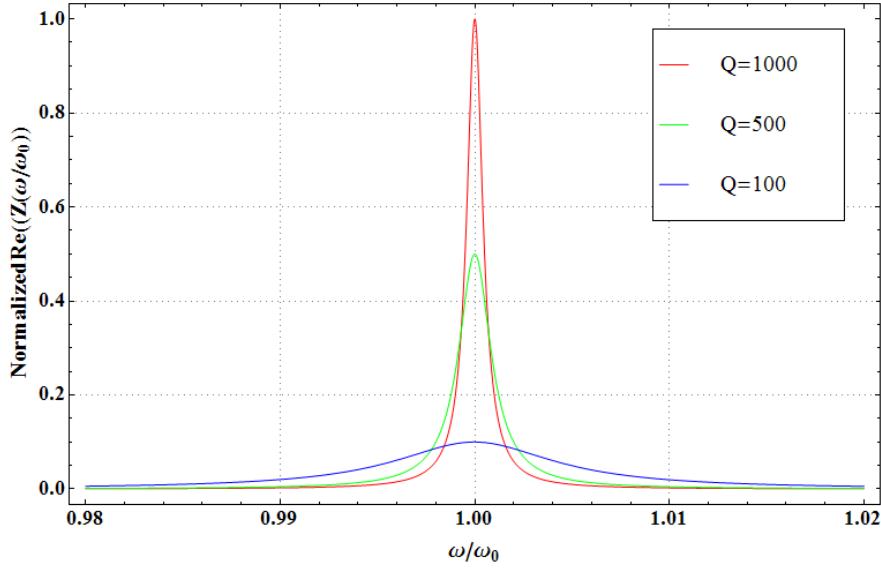
Figure 3.3: A plot of the real component of the impedance $Z(\omega)$ for several different Q values demonstrates the inverse relationship between Q and bandwidth.

This quantity relates the power dissipated per cycle for a given accelerating voltage produced by the cavity. The power lost per cycle is

$$P = \frac{|V_0^2|}{2R_p} \tag{3.15}$$

We then have

$$R_p = \frac{V_0^2}{2P} = \frac{QV_0^2}{\omega U} = Q\sqrt{\frac{L}{C}}, \tag{3.16}$$

where C and L are defined by our RLC circuit representation. The shunt impedance is often characterized per unit-length of the cavity due to its dependence on cavity geometry.

3.1.2 Ferrites

A ferromagnetic material is a material capable of exhibiting spontaneous or permanent magnetization in absence of an external field. This is unlike paramagnetic and diamagnetic materials, which only exhibit fields in response to an external field. Moreover the field strengths reached by ferromagnets can be many orders of magnitude stronger than those of a para/diamagnetic material.

Soft ferrites are a subset of ferromagnetic materials commonly used for

high frequency applications, as they exhibit minimal hysteresis losses over a wide range of frequencies. These materials commonly consist of NiZn or MnZn based iron oxides of the form $Ni_xZn_yFe_2O_4$, where the subscripts x and y detail the proportionality of Ni (or Mn) to Zn. NiZn ferrites exhibit high resistance and are thus preferable for high power operation, but are limited to frequencies below a few MHz. MnZn ferrites are more easily tuned to frequencies upwards of 10 MHz [46].
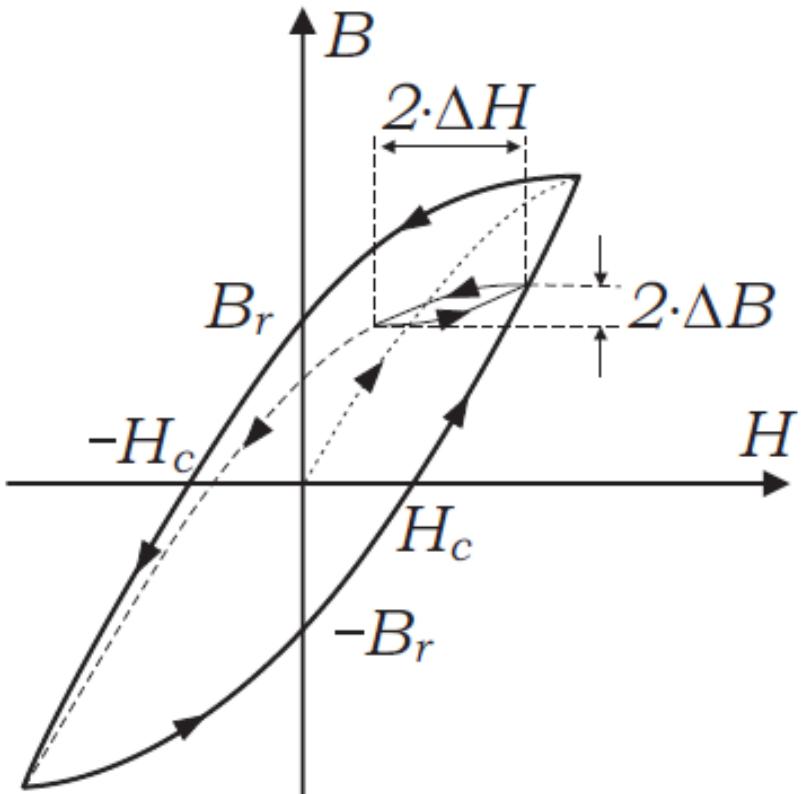
Figure 3.4: A plot showing an applied field H and the resultant induced field B within the ferrite. As the applied field to the ferrite is reduced, the ferrite's internal magnetization lags behind, leading to some remanent field $B \leq B_r$. The continuous loop describing a full cycle of decreasing and increasing the applied field is known as a hysteresis loop.

In the presence of an applied field, the ferrite exhibits an inductance according to $B = \mu H$, where

$$\mu = \mu' - i\mu'' \tag{3.17}$$

is the complex permeability of the material. The quantity μ is often expressed by a ratio $\mu_r = \mu/\mu_0$, where μ_0 is the vacuum permeability of free space. Un-

fortunately, the permeability is quite sensitive to the frequency of the applied field as well as the initial state of the ferrite. The latter relationship is represented by a Hysteresis loop, as shown in Figure 3.4. Beginning in an initially neutral state, an external field is applied and increased, resulting in a linear increase in the induction according to $B = \mu H$. Eventually, the magnetic domains within the material completely align and the inductance saturates. At this point, reducing the applied field does not result in a linear reduction in inductance along the same path; instead, the inductance lags behind. Even after reducing the applied field to 0, a residual inductance B_r remains. The applied field must be reduced to some negative value $-H_c$ known as the coercive field in order to reduce the inductance to zero. "Soft" ferrites have small values of H_c, making them easier to demagnetize in comparison with "hard" ferrites [46, 47].

Consider the response of such a material when the inductance is not driven to saturation, but rather stopped at a value B_1, before being reduced by ΔB to a value of B_2. We may then define the differential permeability $\mu_\Delta = dB/dH$ according to the change in applied field needed to bring about the corresponding change in inductance. In the limit of small changes this quantity represents the effective slope of the local hysteresis loop. Two important caveats regarding the hysteresis loop: permeability is a complex quantity for which both the real and complex components are frequency dependent. Thus, the response curves will vary considerably across a range of frequencies. Secondly, the work done (per unit volume) to the system follows from $W = V \int H dB$. Thus, the work done traversing the hysteresis loop with an external field is equal to the area within the hysteresis loop. As previously noted, softer ferrites with small values of H_c often exhibit smaller losses under cyclical biasing.

3.1.3 Ferrite-Loaded RF cavities

The discussion from Section 3.1.1 makes clear the role of cavity geometry in determining the resonant frequency and bandwidth of a classical cavity. To construct a cavity which is tunable over a broad range of frequencies, we will make use of the strong permeability and low losses of soft ferrites to manipulate a cavity's native resonant frequency. Figure 3.5 shows a wire-model of a basic parallel-biased cavity. Ferrites are manufactured into toroidal discs and stacked along the beam pipe, often balanced about a single capacitive break such as a ceramic insert. This segment is known as the cavity gap and is the point at which the generated voltage couples to the particles. A current is used to create an alternating magnetic field through the ferrite material according to $B = \mu \dfrac{I_{tot}}{2\pi r}$. The rate of change of flux across the cavity gap
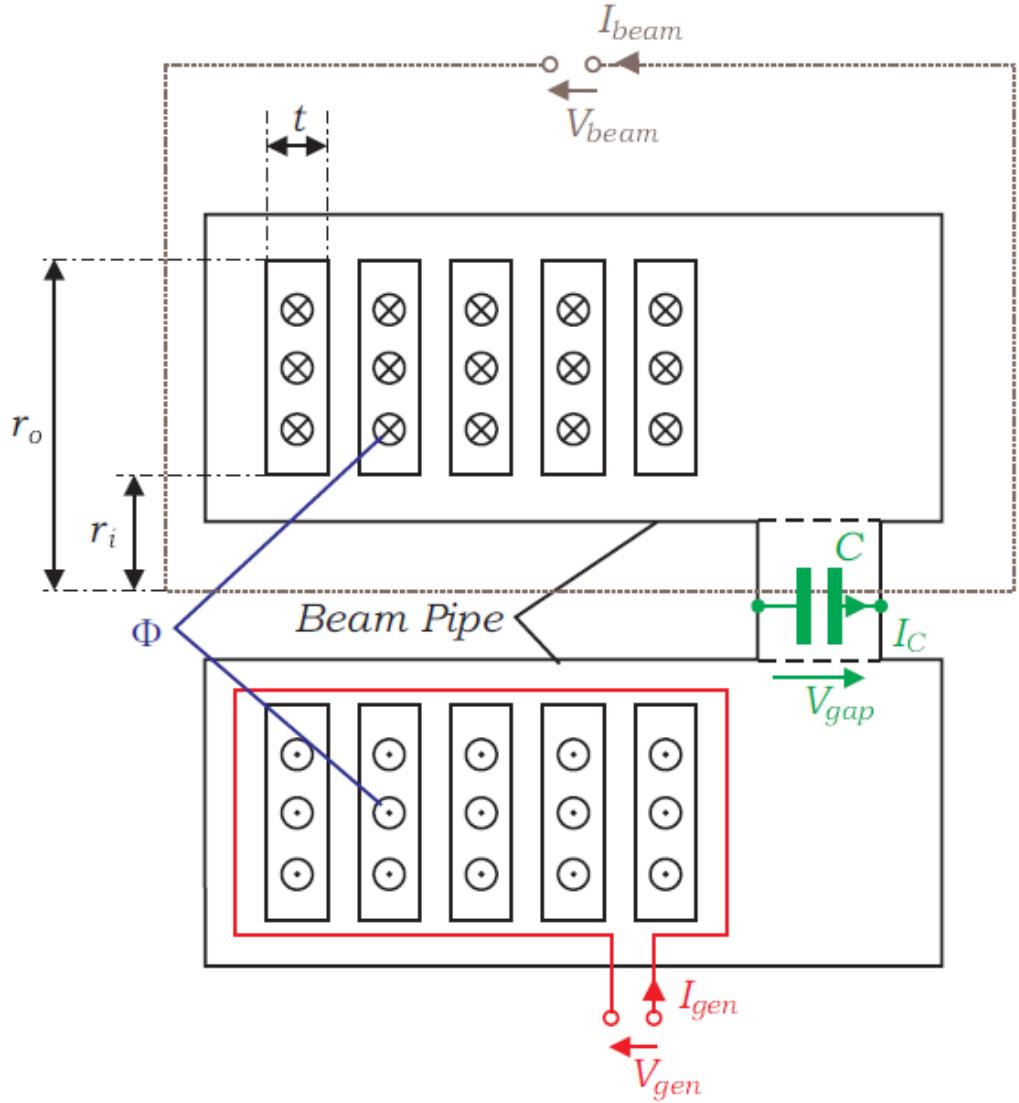
Figure 3.5: A standard parallel biasing configuration for a ferrite cavity is shown [2]. Ferrite rings surround the beam pipe on each side of a gap with effective capacitance C. A tuning circuit (not shown) is operated at high frequency to adjust the resonant frequency of the cavity while a generator circuit, providing current I_{gen}, is operated at the RF frequency to drive a corresponding voltage V_{gap} across the gap. The beam current may be treated as part of a circuit which closes outside of the cavity, and counteracts the generator current current, reducing V_{gap}.

determines the gap voltage via $V_{gap} = j\omega\Phi$. The flux through a single ferrite core is

$$\Phi_1 = \int \boldsymbol{B} \cdot d\boldsymbol{S} = \int_r Bldr = \frac{l\mu I_{tot}}{2\pi} \ln \frac{r_b}{r_a} \tag{3.18}$$

Recall that in the case of a real, lossy ferrite, the permeability μ is a complex quantity given by Equation 3.17. Applying that form to our expression, we find that the voltage across the cavity gap is

$$V_{gap} = i\omega\Phi_{tot} = NI_{tot}\frac{l\omega}{2\pi} \ln \frac{r_b}{r_a}(i\mu' + \mu'') \tag{3.19}$$

The voltage can also be defined according to our series circuit model by

$$V_{gap} = I_{tot}(i\omega L + R) \tag{3.20}$$

thus we deduce

$$L_s = \frac{Nl}{2\pi} \ln \frac{r_b}{r_a}\mu' \tag{3.21}$$

$$R_s = \frac{Nl\omega}{2\pi} \ln \frac{r_b}{r_a}\mu''. \tag{3.22}$$

We can then use Equation 3.14 to find

$$Q = \frac{\mu'}{\mu''} \tag{3.23}$$

This cavity design is naturally represented by the "series" RLC circuit as shown in Figure 3.2(a) with total impedance given by Equation 3.11. However, it is often convenient to model such a cavity with a parallel RLC circuit. We can glean further insight into the properties of the cavity by using a parallel representation. We will then apply a parallel circuit model, shown in Figure 3.2(b). We can equate the two models by assuming that the total impedances of each circuit are equal. In order to distinguish between the two, we will use the subscript s to indicate the series circuit quantity, and p to indicate parallel. We have

$$Z_s = Z_p \frac{1}{R_s + i\omega L_s} + i\omega C = \frac{1}{R_p} + \frac{1}{i\omega L_p} + i\omega C \frac{R_s - i\omega L_s}{R_s{}^2 + (\omega L_s)^2} = \frac{1}{R_p} - \frac{i}{\omega L_p} \tag{3.24}$$

Setting the real and imaginary parts equal yields expressions for R_p and L_p

$$R_p = \frac{R_s{}^2 + (\omega L_s)^2}{R_s} = R_s(1 + Q^2) \tag{3.25}$$

$$\omega L_p = \frac{R_s{}^2 + (\omega L_s)^2}{\omega L_s} = L_s(1 + \frac{1}{Q^2}) \tag{3.26}$$

Using these equations and Equation 3.14 we find

$$Q = \frac{\omega L_s}{R_s} = \frac{R_p}{\omega L_p} \tag{3.27}$$

We can further establish the properties of the parallel circuit if we redefine the complex permeability with parallel components according to $\mu = \mu_p{}' - i\mu_p{}''$. We can then rewrite Equations 3.26 and 3.14 as

$$\mu_p' = \mu_s'(1 + \frac{1}{Q^2}) \tag{3.28}$$

$$\mu_s'' = \frac{\mu_s'}{Q} = \frac{\mu_p'Q}{1 + Q^2} \tag{3.29}$$

We may now fully simplify our expression for R_p given in Equation 3.25,

$$R_p = R_s(1 + Q^2) = (1 + Q^2)\mu_s''\frac{Nl\omega}{2\pi}\ln\frac{r_b}{r_a}$$
$$R_p = \mu_p'QfNl\ln\frac{r_b}{r_a} \tag{3.30}$$

We see that the parallel shunt impedance R_p is proportional to the product $\mu_p'Qf$, which is often written simply as the μQf product. The loss properties of the ferrite are thus determined by this μQf product, and as such it is often a quoted figure of merit of the ferrite material.

3.1.4   Cavity parameters and performance

The cavity gap voltage can be related to the geometrical and magnetic parameters of the ferrite as well as the coil configuration and current generating the magnetic field. In order to tune the frequency of the cavity voltage, we must make use of the ferrite permeability. The value of μ effectively determines the inductance of the cavity, which determines the resonant frequency according to $\omega_0 = 1/\sqrt{LC}$. Applying a DC or slowly varying bias current induces a quasi-static magnetic field in the ferrite which changes the differential perme-

ability μ_Δ of the material. By adjusting this bias current, one can change the permeability and thus the effective resonant frequency of the cavity. Of particular note, because the voltage is proportional to the rate of change of flux, the voltage is naturally proportional to the operating frequency ω of the cavity. Thus operating a cavity at a higher frequency increases both V_{gap} and Q [2].

The method used to tune the cavity has important implications in the cavity performance. Two established methods of delivering the bias current are parallel and perpendicular biasing. In a parallel biasing setup, ferrites are prepared in a toroidal shape and lined coaxially about the beam pipe within the cavity, as in Figure 3.5. A small capacitative gap is made in the beam pipe using a ceramic insert about which the ferrites are situated. A generator circuit encompassing the ferrites provides the rapidly changing current for the purposes of creating the high frequency accelerating voltage, while a secondary circuit is wound about the ferrites to provide the biasing current. Perpendicular biasing is more difficult to achieve; the ferrite rings are stacked away from the beamline and coupled to the cavity. The biasing requires greater spacing in the ferrite stack for cooling and more complex current windings to achieve field saturation. Different ferrites are better suited to different biasing schemes.
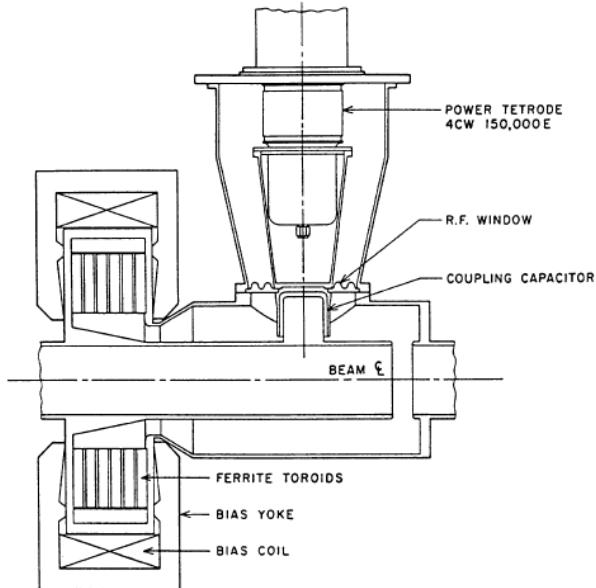
Figure 3.6: A diagram of a perpendicularly biased cavity as designed at Los Alamos National Laboratory. The cavity is designed to operate in the 50-84 MHz range, over which the ferrite rings show much higher Q values using perpendicular biasing [3] ©1985 IEEE.

It is difficult and inefficient to rapidly tune RF cavities over a wide range of frequencies. Operational frequencies are specific to individual ferrite materials. At sustained high fields, ferrites demonstrate considerable losses; this is known as the high-loss or Q-loss effect. Moreover, rapid tuning of ferrite biasing current results in strong dynamic losses up to 50% of the static Q value. For the purposes of accelerating protons and heavy ions to midrange energy levels suitable for radiation therapy (5 to 430 MeV/u carbon, 7 to 330 MeV protons), this problem is especially prominent, as a single machine might need to undergo a 5-fold frequency swing [45].
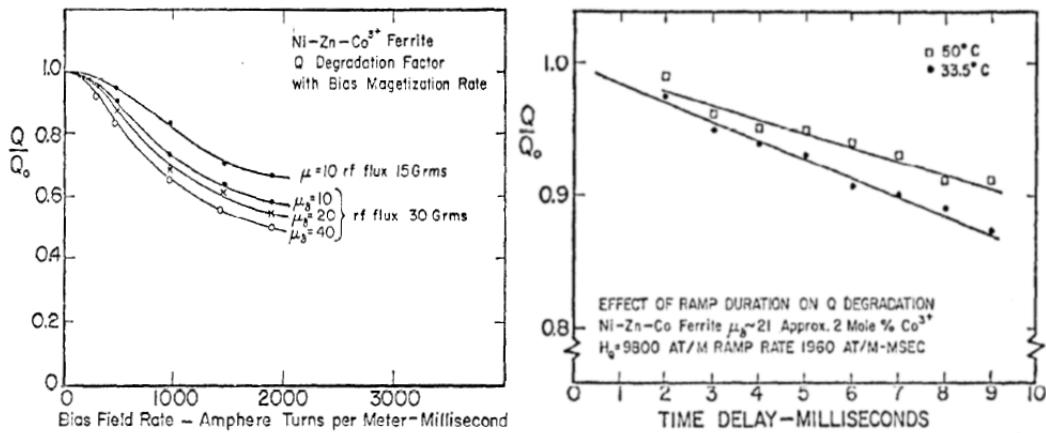
Figure 3.7: At left, The value of Q clearly declines with increasing bias field rate. This is known as f-dot loss, and can result in a 50% reduction in Q value for high bias field rates. At right, the value of Q is shown to decline when ramped over a long period of time. At lower operating temperatures, a 10ms ramp can lead to a 15% decline in Q. ©1979 IEEE

3.2 Harmonic ratcheting

Several different approaches have been considered as means of improving broadband RF performance. In lieu of ferrite, magnetic alloy (MA) cavities have been designed and used for high power rapid cycling synchrotrons [48]. Another solution is to chose the energy gain per turn such that cavities are spaced by an integer multiple of the RF period, thus allowing accelerating bunches to skip RF buckets. This is the premise of the "harmonic number jump" scheme, providing a dynamical adjustment of h to extend the viable dynamical range of a fixed-frequency cavity [49, 50]. Alternately, the required RF frequency range can be significantly reduced below the revolution frequency range by

decreasing h in steps as the ions accelerate and f_{rev} increases. This is the motivation and the basic method of harmonic ratcheting.

Suppose that two sets of RF cavities take turns accelerating the beam — one turns on when the other turns off, at different RF frequencies — so that the RF frequency is always constrained to remain in the range

$$f_{min} \leq f_{RF} \leq f_{max} \tag{3.31}$$

where f_{min} and f_{max} are externally determined design parameters. It is possible to make the transition back or forth between harmonic numbers

$$h = n \tag{3.32}$$
$$\text{and} h = n + \Delta$$

where n and Δ are positive integers, if

$$\frac{f_{max}}{n + \Delta} > \frac{f_{min}}{n} \tag{3.33}$$

that is, if

$$\Delta < nr \tag{3.34}$$

where

$$r \equiv \frac{f_{max}}{f_{min}} - 1 > 0 \tag{3.35}$$

is the "ratcheting parameter". Equation 3.34 shows that Δ has a maximum permissible value, which must be greater than 1 if a harmonic transition is to be possible. This is guaranteed for all values of n in the case that $r > 1$, when the RF frequency swings by more than a factor of two. If the swing ratio is less than 2, then $r < 1$ and transition is only possible if

$$n > \frac{1}{r} \qquad \text{and} \qquad r < 1 \tag{3.36}$$

Many harmonic transitions can take place during one acceleration ramp. Figure 3.8 shows an example with $r = 0.50$, where the revolution frequency increases from 0.61 MHz to 3.35 MHz, while the RF frequency is constrained to lie between 5.5 MHz and 8.3 MHz. Figure 3.9 shows the RF frequency and total gap voltage as a function of time for the same ramp.
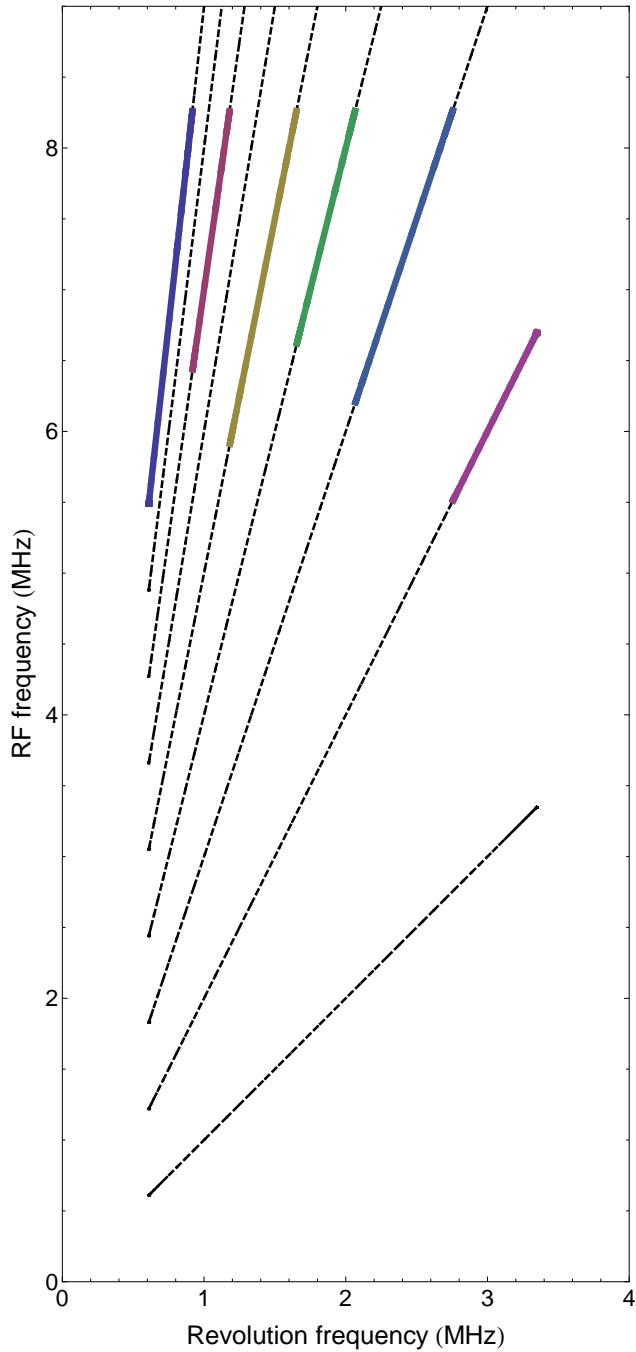
Figure 3.8: An example of harmonic ratcheting with $r = 0.50$, in the case where the revolution frequency increases from 0.61 MHz to 3.35 MHz, while the RF frequency is constrained between 5.5 MHz and 8.25 MHz. Color coded segments show one possible solution — with the harmonic number sequence $h = 9, 7, 5, 4, 3, 2$ — that maximizes the RF frequency while minimizing the number of ratchets.
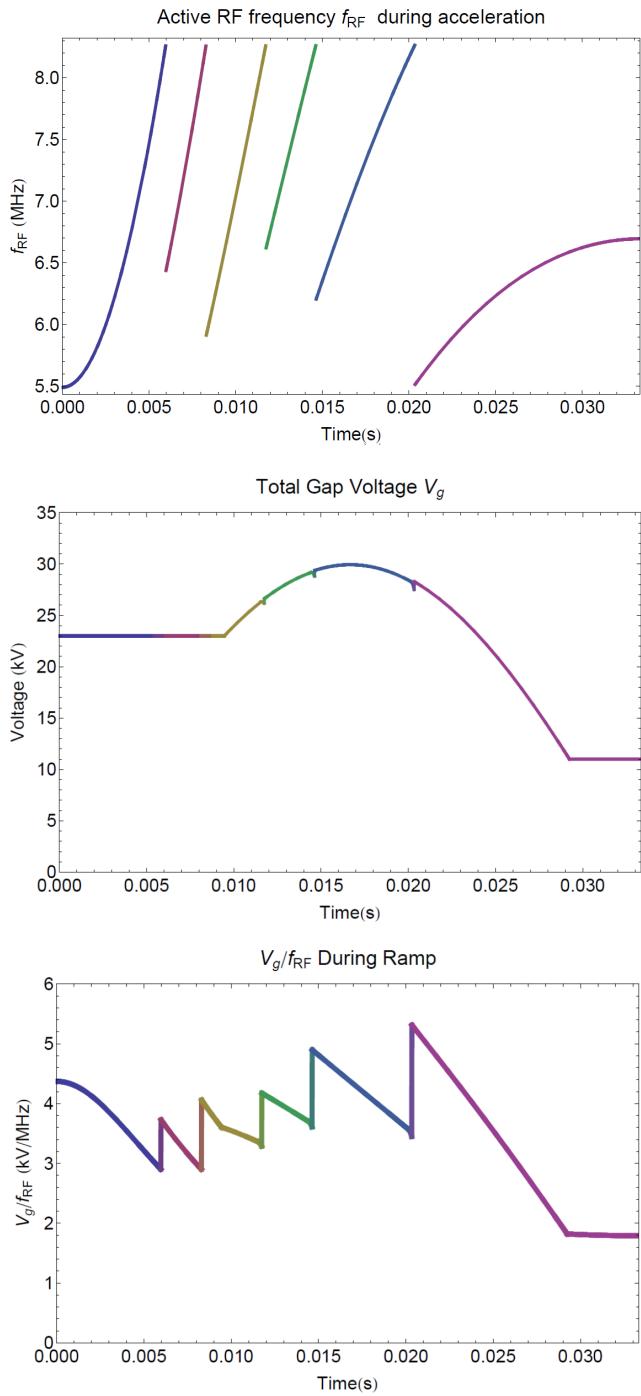
Figure 3.9: Active cavity parameters for a 65 m Carbon synchrotron operating at 15 Hz with ratcheting sequence of h=9,7,5,4,3,2. Top, active cavity voltage is shown versus ramp time. Middle, the active cavity frequency is shown. Bottom, the ratio of cavity voltage to frequency is plotted; this ratio remains consistent throughout acceleration.

3.2.1 Designing a ratcheting ramp

One constraint on the initial harmonic number, h_0, in the ratcheting ramp is imposed by the minimum required bucket length, t_{inj} to accept particles into a single bucket at injection. This length sets a maximum possible initial harmonic. Moreover, the energy acceptance at injection also scales inversely with $h_{max}{}^2$. These requirements can be summarized by

$$h_{max} \ \le \ \frac{1}{t_{inj}f_{rev}(0)} \tag{3.37}$$

$$\text{and } h_{max}{}^2 \ \le \ \frac{2QV_g\beta^2 E_s}{\pi|\eta_s||\Delta E_0|^2}$$

Here, $Q = Ze$ is particle charge, E_s is synchronous energy, and η_s is the phase slip factor. As the initial harmonic increases, a corresponding increase in voltage is necessary to keep the longitudinal acceptance the same.

3.2.2 Emittance growth at ratcheting transition

During the harmonic transition, the voltage on the cavity operating near f_{max} at harmonic $h = n$ is reduced while the voltage on a second cavity operating at harmonic $h = n + \Delta$ is raised to provide the desired gap voltage at the new harmonic. Assume that this change happens quickly (non-adiabatically) and that both the synchronous phase and total accelerating voltage are smoothly varying and/or fixed when the ratcheting takes place. The accelerating bucket area, A, scales with the harmonic number as

$$A \ \propto \ \frac{1}{h^{\frac{3}{2}}} \tag{3.38}$$

Assume the ramp undergoes a harmonic change according to Equation 3.32, with the harmonic number being reduced from $n + \Delta$ to n, and the total voltage does not change appreciably. The bucket area grows from A_i to A_f according to

$$A_f \ = \ \left(\frac{n+\Delta}{n}\right)^{\frac{3}{2}} A_i \tag{3.39}$$

Because the synchronous phase remains the same, it is only the slope of the voltage curve with respect to particle position that changes, and not the total magnitude of accelerating voltage. Thus, the transition from a larger to smaller harmonic does not result in a net shift in bunch position, but instead relaxes the voltage curve defining the stable longitudinal phase space. In order

to achieve the same restoring voltage slope, a particle with phase displacement θ_i under the previous harmonic $n + \Delta$ extends to a new phase coordinate

$$\theta_f = \frac{n + \Delta}{n} \theta_i \tag{3.40}$$

This results in a change in the RMS momentum spread, σ, from σ_i to σ_f where

$$\sigma_f^2 = \left(\frac{n + \Delta}{n}\right) \sigma_i^2 \tag{3.41}$$

The 95% longitudinal emittance of the beam, assumed Gaussian, follows from

$$S = K \left(\frac{\sigma_p}{p_s}\right)^2 \tag{3.42}$$

$$\text{with } K = \frac{3}{h f_{rev}} \sqrt{\frac{-2\pi h \eta E_s{}^3 \beta^2}{e V_g \cos \phi_s}} \tag{3.43}$$

where p_s is the synchronous particle momentum, E_s is the synchronous particle energy, and e is the elementary charge of an electron [33]. Using Equations 3.41 and 3.43 the new emittance becomes

$$S_f = \left(\frac{n + \Delta}{n}\right)^{\frac{3}{2}} S_i \tag{3.44}$$

Thus the emittance growth due to the harmonic number transition in Equation 3.44 matches the bucket area growth due to the same transition found in Equation 3.39. We have

$$\frac{S_i}{A_i} = \frac{S_f}{A_f} \tag{3.45}$$

For slow changes in harmonic number, emittance growth is precisely checked by an increase in bucket area, making the ratcheting process inherently stable. This property can be seen in the Figure 3.13 for the example ramp in which the transition takes place non-adiabatically.

3.3 Cavity performance with harmonic ratcheting

The resonant frequency of a ferrite-loaded cavity follows from the discussion in Section 3.1.1

$$f = \frac{1}{2\pi} \frac{1}{\sqrt{LC}} \tag{3.46}$$

where L, the inductance of the cavity, is adjustable between minimum and maximum value L_{min} and L_{max}, corresponding to frequencies f_{max} and f_{min}, respectively. The minimum inductance that is required is given by solving the equation

$$\frac{f_{max}}{f_{min}} = 1 + r = \sqrt{\frac{L_{max} + L_{stray}}{L_{min} + L_{stray}}} \tag{3.47}$$

Here, L_{stray} refers to the stray inductance of the cavity, which sets a minimum on the unloaded cavity properties. When r is small, L_{min} is large, and it is possible to ignore the stray inductance, L_{stray} (which is usually on the order of several μH); however in the case of a large frequency swing the entire range of plausible inductances must be tuned. Figure 3.10 demonstrates the reliance upon small stray inductance for large r values, using Phillips 4L2 ferrite for tuning. It is clear that ratcheting permits larger stray inductances.

Consider a cavity with a total ferrite length of l, made of many rings with inner and outer radii r_a and r_b. The inductance is given by

$$L = \frac{l}{2\pi} \mu' \ln \frac{r_b}{r_a} \tag{3.48}$$

where we take μ' to be the real component of the incremental complex permeability μ_Δ. The component μ' and the inductance of the cavity are tuned by biasing the ferrite with a pseudo-constant azimuthal magnetic field driven by a tuning current. Ignoring the stray inductance, the required dynamic range of L and μ' is

$$\frac{L_{min}}{L_{max}} = \frac{\mu'_{min}}{\mu'_{max}} = \frac{1}{(1+r)^2} \tag{3.49}$$

The voltage across the gap of a cavity follows

$$V_{gap} = 2\pi f \, B_{max} \, l \, r_a \ln \frac{r_b}{r_a} \sim f \, l \tag{3.50}$$

where $B_{max} \approx 0.01$ T is the maximum incremental RF magnetic field that is allowed at the inner radius of the ferrite rings, oscillating at the frequency f.

For a fixed ferrite ring geometry (r_a and r_b), and a fixed maximum field
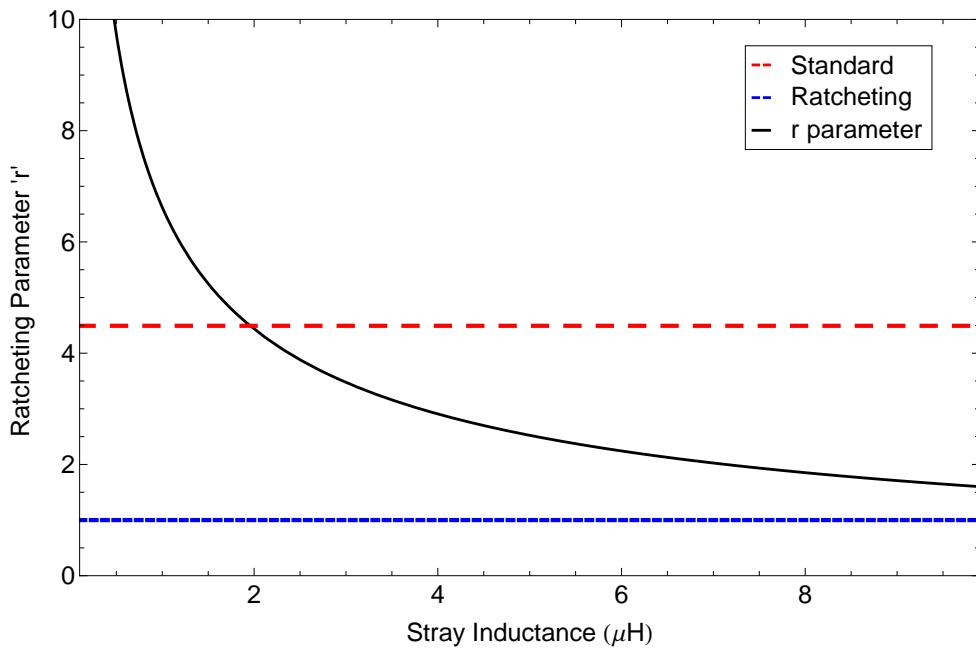
Figure 3.10: The achievable ratcheting parameter r is plotted (black) for the ferrite Phillips 4L2 using fitting data collected at BNL [4]. The plot demonstrates that, for a fixed capacitance, a low stray inductance is needed to achieve the high frequency ratio required by a single harmonic ramp (red). For a ratcheting solution (blue) with $f_{max}/f_{min} < 2$, inductances of several μH are acceptable.

B_{max}, then the gap voltage scales with ferrite length and cavity frequency. The ferrite length l can be reduced if the RF frequency in the cavity is increased. By inflating cavity frequency, harmonic ratcheting allows a given acceleration waveform to be achieved with shorter ferrite — and shorter cavities.

The total length of ferrite required with a ratcheting parameter r, relative to a non-ratcheting scheme in which ions are accelerated over the full dynamic range

$$ D \;=\; \frac{\beta_{max}}{\beta_{min}} \;=\; \frac{f_{rev,max}}{f_{rev,min}} \tag{3.51} $$

is given approximately by

$$ \frac{l_{ratchet}}{l_{non-ratchet}} \;=\; \frac{f_{min,non-ratchet}}{f_{min,ratchet}} \;\approx\; 2\,\frac{1+r}{D} \tag{3.52} $$

where the factor of 2 recognizes that only half the cavities are active at any one time. For example, Figure 3.8 illustrates acceleration over a dynamic range of $D = 5.5$, using a ratcheting parameter of $r = 0.50$. In this case ratcheting (ideally) makes it possible to decrease the total length of ferrite by a factor of 0.55. Conversely, the gap voltage can be almost doubled for a fixed amount of RF cavity real estate, in turn almost doubling the potential repetition rate of a rapid cycling synchrotron.

3.4 Example: a rapid cycling medical synchrotron

Consider a rapid cycling synchrotron for acceleration of C^{6+} ions for radiation therapy applications, with a 15 Hz repetition rate and a circumference of $C \simeq 65$ m. Ions are accelerated through a range 8 to 400 MeV/u, with a nominal bunch current of 10 pA.

Once the initial harmonic number is determined, the initial RF frequency is set at $f_{RF}(0) = h_0\,\beta c/C$. The frequency ceiling can then be chosen. In this example, we choose a ratcheting parameter of $r = 0.5$ via a final harmonic transition of $h = 3$ to $h = 2$. The resultant frequency swing is shown in Figure 3.9, with the harmonic sequence shown in Figure 3.11.

When the frequency ceiling is reached after acceleration at one particular harmonic, RF power is transferred to the second cavity system at a lower harmonic number. Each time the harmonic number is reduced, the RF frequency is reduced by a factor $h_{n+\Delta}/h_n$ where h_n is the new harmonic, and $h_{n+\Delta}$ is the previous value. The maximum frequency swing during the ramp is determined by the maximum value of this ratio, or equivalently the ratcheting parameter
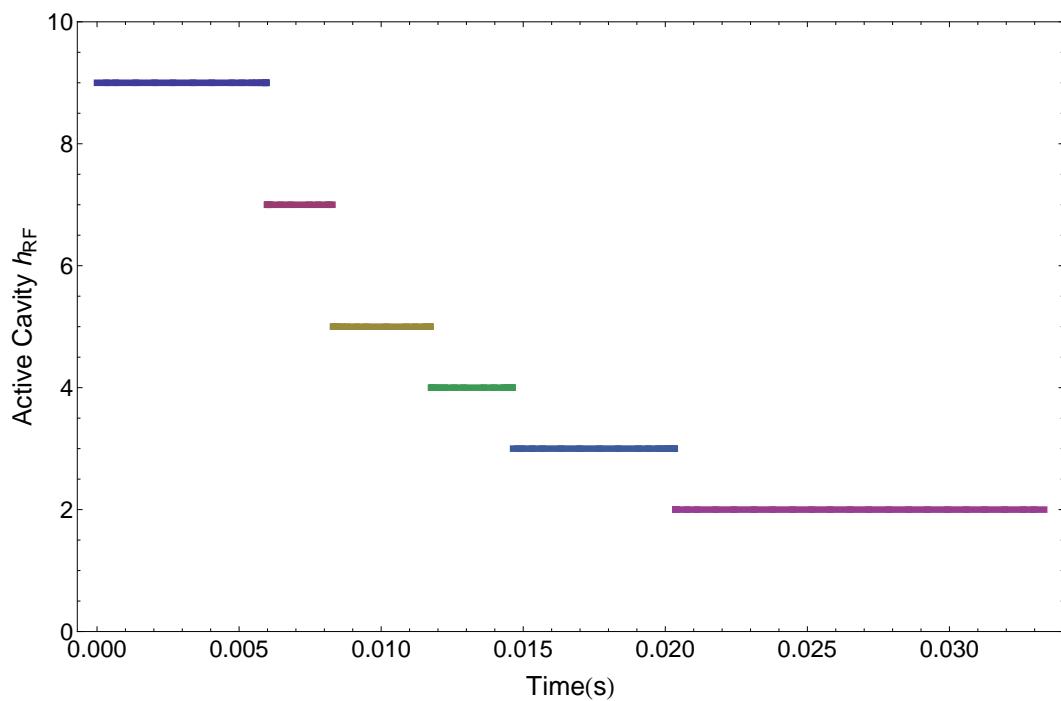
Figure 3.11: Active cavity harmonic for the 65 m carbon synchrotron harmonic ratcheting ramp. The cavity begins at h=9 and ratchets in sequence through h=7,5,4,3,2.

r. For example, if the largest transition in a particular ratcheting scheme is from $h = 3$ to $h = 2$, then the ratcheting parameter is $r = 0.5$, and that particular ramp is limited to a 50% frequency swing.
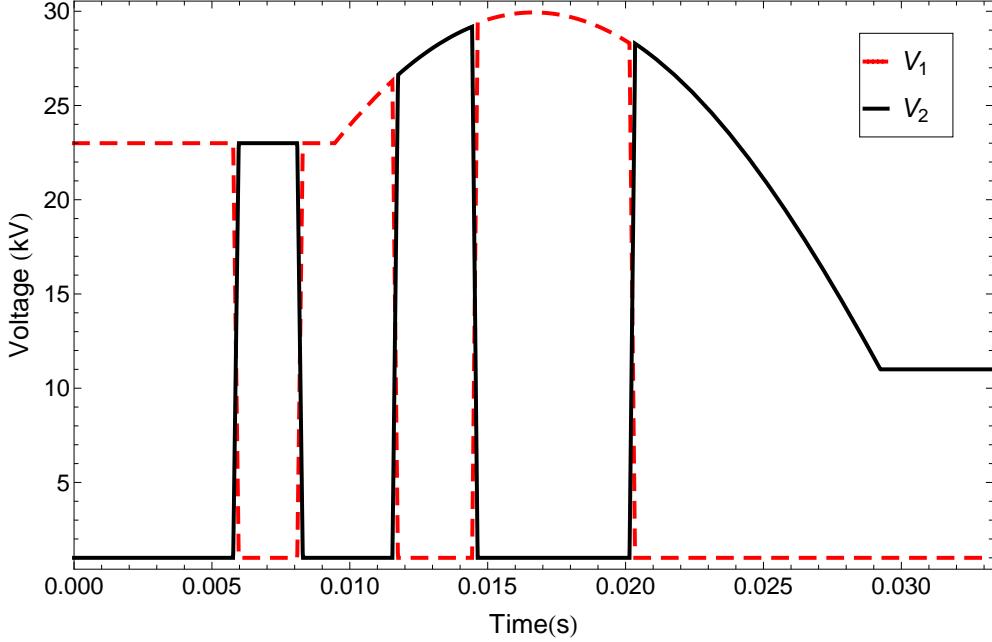
Figure 3.12: Output voltage for each cavity system is depicted for a test ratcheting sequence. A ramping time of 100 μs is chosen for the transition periods, and total voltages are calculated to meet the constraint $\phi_s \leq 33°$. A minimum voltage of 1 kV is assumed in order to allow RF tuning and phase synchronization for the "offline" cavities.

Figure 3.12 shows the output voltage needed for each cavity system in the example ratcheting ramp. The transfer of power in a harmonic transition to the inactive cavity cannot be performed instantaneously. The voltage across the cavity must be raised over many turns of the accelerator. This transition may be on the order of 100 μs depending upon the voltage level and frequency of the cavity. For a fixed total gap voltage, a change in h constitutes a change in the synchronous phase ϕ_s of the system according to

$$\phi_{s,f} = \phi_{s,i} \frac{n}{n + \Delta} \tag{3.53}$$

If the phase is changed too abruptly, the bunch center will oscillate about the new value of ϕ_s, and the bunch will filament, causing strands of particles to separate outward from the edge of the bunch and increasing emittance. Assume the transfer of voltage takes place over some time period τ_r. A particle

will make

$$n_t = \beta c \tau_r / C \tag{3.54}$$

turns in that time. If $\tau_r = 100$ μs, a $\beta = 0.1$ particle makes 46 turns of a 65 m synchrotron, thus smoothing the dynamics present during the transition period.
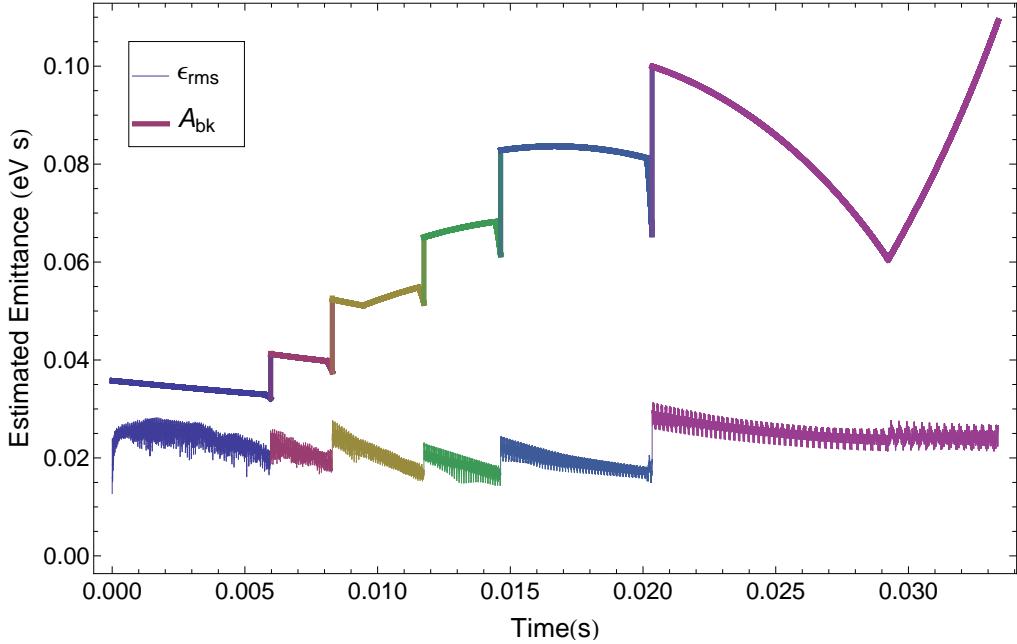
Figure 3.13: RMS bunch emittance, ϵ_{rms}, and bucket-area, A_{bk}, are shown for a simulated ramp with Carbon ions and a fixed voltage profile. The emittance spikes correspond to the harmonic transitions inducing a quick change in synchronous phase.

An alternative to the fixed voltage profile is to vary the voltage in order to produce a fixed synchronous phase over the course of acceleration. Maintaining a fixed phase reduces emittance growth and bunch filamentation, while reducing power requirements on the RF system. Detailed feedforward and diagnostic systems are required in order to maintain the proper voltage. The emittance evolution for our example synchrotron can be seen in Figure 3.13.

Using the harmonic sequence discussed above, Table 3.1 reports an estimate of the new acceleration parameters for the 15 Hz C^{6+} acceleration scheme, using Phillips 4L2 and Ferroxcube 8C12 ferrites [4]. The results underscore the significant reduction in magnetic field, power density, and the extreme decline in bias field required to tune the ferrite across the smaller relative frequency range.

Table 3.1: Summary of physical parameters and requirements for RF cavity as well as power numbers, assuming $L_{stray} = 0.5\mu H$, for C^{6+} acceleration at 15 Hz. We assume the use of two RF cavities with the same geometry, $\mu Q f$ values, and biasing scheme constructed from either Phillips 4L2 or Ferroxcube 8C12 ferrite. Both ratcheting and non-ratcheting ramps are considered for each ferrite.

Cavity Parameter		4L2	4L2 ratcheting	8C12	8C12 ratcheting
Harmonic Sequence		1	$9, 7, 5, 4, 3, 2$	1	$9, 7, 5, 4, 3, 2$
RF Frequency	[MHz]	0.65 to 3.57	5.5 to 8.25	0.65 to 3.57	5.5 to 8.25
Injection Bunch Length	[ns]	500	60	500	60
Maximum B-field	[mT]	1.33	0.96	1.33	0.96
Maximum Gap Voltage	[kV]	27.5	30	27.5	30
Max Cavity Power Loss	[kW]	8.4	6.7	8.4	6.7
Permeability Range, μ_r		4.93 to 200	86.9 to 200	24.8 to 200	86.9 to 200
Maximum Bias H	[A/m]	4485	342	968	320

3.5 Ratcheting challenges

There remain many challenges to effectively implementing a harmonic ratcheting system. These lie primarily in meeting the costs and constraints of building and tuning a flexible multi-cavity system which can stably operate at the desired performance level for a given feed-forward RF program.

3.5.1 Anomalous effects

Use of a harmonic ratcheting scheme may also reduce losses in the ferrites due to anomalous effects. At sustained high fields, ferrites demonstrate what is known as the high-loss or Q-loss effect (QLE). This can result in reductions of 15-20% of maximum field. High bias fields can increase this effect by reducing the threshold of Q-loss onset. However, it has been documented that at high bias rates, the threshold for this effect is significantly raised [44]. Given these considerations, a ratcheting scheme should mitigate the QLE by reducing bias fields and maintaining a high rate of change of bias field. High bias field rates may induce a dynamic loss effect in the ferrite. These losses may reduce a cavity's Q value by up to 40%. Unlike the QLE, this effect appears at varying levels of RF excitation and is more dependent on the temperature and duration of the biasing. By decreasing the time spent biasing before reaching resonance at high fields, the effect can be reduced considerably. It must be stressed that these effects are still poorly understood on a theoretical level, and their effects can vary considerably for different materials.

3.5.2   Magnetic alloy cavities

An alternative to ferrite-loaded cavities is made possible by the use of magnetic alloys, which feature high permeability values, resilience under high magnetic fields, and large inductances. These properties allow MA cavities to achieve low Q values, meaning they can support fast, broadband acceleration using feedforward systems to achieve the desired frequency band operation, as seen in the J-PARC 3 GeV RCS [51]. This low Q value combined with the robust nature of the material allows MA cavities to achieve high gradients, but at the cost of high power requirements, especially at low frequencies wherein the $\mu Q f$ product can be orders of magnitude lower than in ferrite cavities [48]. Despite the resilience of MA cores to high fields and temperatures, thermal stress can result in impedance drops and subsequent core damage. The so-called "core buckling" phenomenon requires proactive maintenance schedules requiring the replacement of 2 sets of cores each year at J-PARC [52]. Recently, such incidents have been found to be reducible by changing the epoxy resin used in core construction. In all, MA cavities provide a number of advantages for high frequency, high gradient operation for RCSs at the expense of power loss and maintenance costs. These tradeoffs are directly correlated to cavity performance, and as such a ratcheting approach would only reduce costs through a similar reduction in performance. Use of MA cavities, as with ratcheting, is useful for particular classes of problems, but they do not complement one another.

3.6   Summary

The harmonic ratcheting system is unique amongst RF schemes for proton and ion synchrotrons due to the active adjustment of harmonics throughout the ramp. Some other situations do exist in which multiple harmonics are used simultaneously. These include barrier bucket systems, re-bucketing systems, and bunch stretching RF systems. However, in each of these situations, the higher harmonic is commonly applied at a fixed energy ($\phi_s = 0$) system. The additional harmonic is used to stretch, squeeze, or redistribute the bunch in a way that is favorable for improving luminosity, reducing eddy currents, or extracting beam to a different system [53–56]. Harmonic ratcheting utilizes a multi-harmonic system to adjust the primary harmonic number during acceleration. In this way harmonic ratcheting is a distinctive scheme for RF acceleration of protons and ions.

Harmonic ratcheting presents an scheme for the RF acceleration of protons and ions. The scheme requires careful timing and extreme tuning of

ferrite materials in order to determine the optimum frequency range and harmonic transition times for acceleration. Considerable efforts will be needed to test prospective ferrite materials. Because of individual variations in materials as well as differing needs for cavities, successful implementation of harmonic ratcheting will require further investigation into generalizing a ratcheting solution to suit the needs of specific projects. Application of the ratcheting scheme to real cavity design, including ferrite selection, bias current circuitry, and RF power amplifiers are necessary to gather the experience needed to create an array of general solutions.

Chapter 4

A Laser Plasma Primer

In comparison to RF cavities, intense laser pulses have a number of appealing properties. First and foremost, electric fields generated by a laser pulse are many orders of magnitude higher than those seen in RF cavities, with peak fields in excess of 10^{12} V/m [57]. That these fields can be generated with a relatively small footprint when compared with large circular accelerators is attractive for many applications. Moreover, the wall plug power of ultra-intense laser systems has improved to be upwards of 5–20 % depending upon the gain medium, which though still lower than superconducting RF cavities, is compensated by the increase in gradient. All of these factors, combined with the inherent flexibility provided by a laser system in terms of output timing, pulse power, and polarization suggests that lasers could be ideal means of accelerating particles.

Unfortunately the transfer of energy to particles is not so easy. It turns out that in its simplest form, the interaction between a laser field and an energetic free particle is not conducive to acceleration, an observation first made by P.M. Woodward and J.D. Lawson when studying laser fields in the vicinity of an aperture [58]. Their conclusions are now referred to as the Lawson-Woodward theorem, which states that a free particle in a laser field experiences no net acceleration, subject to the following assumptions [59]:

1. The laser field is in vacuum with no boundaries

2. The particle is highly relativistic along the acceleration path

3. No static fields are present

4. The region of interaction is infinite

5. Ponderomotive effects are neglected

It is easy to see that when these conditions are satisfied, the particles will not experience any net acceleration. Consider a highly relativistic electron in a vacuum subject to an intense laser pulse. Assume the laser propagates in the z-direction, with the laser field linearly polarized along the x-axis. The electric field E_z is given by

$$E_z = \frac{1}{2\pi} \int dk_x \int dk_y \hat{E}_z(k_x, k_y) \exp[i(k_x x + k_y y + k_z z - \omega t)] \qquad (4.1)$$

Poisson's equation gives us an expression for the field component, given our assumptions,

$$\hat{E}_z = -\frac{1}{k_z}(k_x \hat{E}_x + k_y \hat{E}_y) = -\frac{k_x}{k_y}\hat{E}_x \qquad (4.2)$$

We may then calculate the energy gain ΔU,

$$\Delta U = q \int E_z dz = \frac{q}{2\pi} \int dk_x \int dk_y \frac{k_x}{k_y} \hat{E}_x(k_x, k_y) \int dz \exp[(ik_z z - \omega t)] \quad (4.3)$$

Since the interaction must be infinite, we integrate over all z, obtaining

$$\Delta U = \frac{q}{2\pi} \int dk_x \int dk_y \frac{k_x}{k_y} \hat{E}_x(k_x, k_y) \delta(k_z - k) \quad (4.4)$$

Note that $\delta(k_z - k) \neq 0$ only when $k_z = k$. But this implies that $k_x = 0$. Thus we have $\Delta U = 0$. This result has an intuitive physical interpretation. The equivalence $k_z = k$ implies that the particle is traveling at the speed of light, and cannot be further accelerated. However, for all $k_z \neq k$, the particle is traveling at $v < c$ and thus there will be a natural phase slippage between the laser field and the particle, which over the infinite interaction region averages to zero net acceleration.

Though the Lawson-Woodward theorem at first seems to be an obstacle, the nature of the constraints serves to guide the design of an experiment in which lasers can accelerate particles. A number of vacuum acceleration schemes have been proposed [60, 61] which cleverly elude the requirements of Lawson to obtain net acceleration. However, a natural solution to the theorem is simply to violate the first condition, that of vacuum. By introducing a medium, the laser has a means of finite energy transfer. Therein lies the role of plasma.

4.1 A plasma primer

A plasma can be defined as a quasi-neutral system of charged particles in great enough number to exhibit collective behavior. Because the constituents of a plasma are charged, they interact through the Coloumb force as well as through collisions. Plasmas found in nature are quasi-neutral, meaning the total charge is zero, but locally secular fields may exist.

Plasma frequency

The plasma frequency is the fundamental oscillation frequency of electrons in a plasma in response to an electromagnetic perturbation.

Consider a 1D plasma with uniform electron density n_e, in which at some time t, a slab of plasma is displaced from its quasi-neutral position by some distance δ_x. Applying Poisson's equation gives the electric field generated by

the charge separation

$$E(x, t) = -\frac{n_e e}{\epsilon_0} \delta_x \tag{4.5}$$

The resulting equation of motion is then

$$\frac{d^2 E(x, t)}{dt^2} = -\omega_p{}^2 \tag{4.6}$$

where the plasma frequency, ω_p is

$$\omega_p{}^2 = \frac{n_e e^2}{m_e \epsilon_0} \tag{4.7}$$

The ion plasma frequency follows the same form, with the ion mass and density substituted, and so is usually small enough to be ignored.

Debye length

The plasma frequency sets a lower limit on the time-scale of collective (i.e. plasma-like) behavior according to the plasma period $2\pi/\omega_p$. The distance corresponding to this time-scale is known as the Debye length, and is given by the distance travelled by an average particle during one period,

$$\lambda_D = \frac{v_t}{\omega_p} \tag{4.8}$$

where v_t is the thermal velocity of the particle. For a plasma with average electron temperature T_e, and substituting Equation 4.7, the Debye length is

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n_e e^2}} \tag{4.9}$$

The Debye length can be approximated by [62]

$$\lambda_D[\text{cm}] \approx 7.43 \times 10^2 \sqrt{\frac{T_e[eV]}{n_e[cm^{-3}]}} \tag{4.10}$$

For example, if $T_e = 10$ keV and $n_e \approx 10^{19}$ cm^{-3}, $\lambda_D \simeq 0.23$ μm. The Debye length gives the average distance between free electrons in a plasma. Moreover, it provides a maximum distance over which Coloumb interactions affect particle dynamics. It can be thought of as a maximum impact parameter in this regard.

Plasma parameter

We may further describe a plasma with the use of the dimensionless plasma parameter, defined by

$$\Lambda = 4\pi n_e \lambda_D{}^3 = \frac{4\pi \epsilon_0{}^{3/2}}{e^3} \frac{T_e{}^{3/2}}{n_e{}^{1/2}} \tag{4.11}$$

where T_e is given in eV. The plasma parameter not only describes the number of particles in the Debye sphere (a sphere of radius λ_D), but also the ratio of the maximum impact parameter, λ_D, to the minimum impact parameter under Coloumb scattering, b_0. That is,

$$\Lambda = \frac{\lambda_D}{b_0} \tag{4.12}$$

If $\Lambda \gg 1$, then the Debye sphere is heavily populated, and electrostatic interactions between individual particles tend to be shielded. The result is a weakly coupled plasma. Conversely, if $\Lambda \ll 1$, then long range electrostatic forces dominate over kinetic interactions, leading the plasma to be strongly coupled. Most laser generated plasmas (ICF plasmas, capillary plasmas, gas target plasmas, and foams) are weakly coupled, and we will concern ourselves with this regime for the remainder of the discussion.

4.1.1 Collisions

In a weakly coupled plasma, collisions are dominated by long range Coloumb interactions. At large Λ, collective effects of shielding and scattering are minimal. However, because Coloumb collisions usually result in small angle scattering, a 'collision' between plasma particles is defined to be a 90° scattering. For weakly coupled plasmas, the plasma parameter offers a convenient means of quantifying the "rate of scattering". Recall that Λ describes the ratio of average particle separation to minimum impact parameters. The larger this ratio, the smaller the rate of collisions should be. The electron-ion collision frequency follows from

$$\nu_{ei} \sim 2.6 \times 10^{-6} \frac{Z n_e \ln \Lambda}{T_e{}^{3/2}} \tag{4.13}$$

where n_e is given in centimeters. The electron-electron and ion-ion collision frequencies follow from

$$\nu_{ee} = \frac{1}{\sqrt{Z}}\nu_{ei} \tag{4.14}$$

$$\nu_{ii} = Z^2 \left(\frac{m_e}{m_i}\right)^2 \nu_{ei} \tag{4.15}$$

The mean free path, in this case the average distance between 90° scattering events, follows from

$$\lambda_{ee} = \frac{v_{th,e}}{\nu_{ee}} \tag{4.16}$$

Consider a plasma with temperatures $T_e = 10$ keV and $T_i = 10$ eV. The corresponding mean free paths are on the scale of hundreds of millimeters. Thus we see that even at modest temperatures, mm-scale plasmas can be considered collisionless.

4.1.2 Describing plasma dynamics

Consider a plasma of N charged particles. The most fundamental description of a plasma can be obtained by solving for the individual dynamics of each of the plasma's particle constituents. Doing so amounts to solving the equation of motion of each particle under the influence of the electric fields $\boldsymbol{E}(\boldsymbol{r}_i, t)$ and magnetic fields $\boldsymbol{B}(\boldsymbol{r}_i, t)$ generated by the remaining plasma particles. This equation of motion is given by the Lorentz Force according to,

$$\boldsymbol{a_i} = \frac{d\boldsymbol{v}_i}{dt} = \frac{q}{m_i}\left[\boldsymbol{E}(\boldsymbol{r}_i, t) + \boldsymbol{v}_i \times \boldsymbol{B}(\boldsymbol{r}_i, t)\right] \tag{4.17}$$

where $\boldsymbol{r}_i$ and $\boldsymbol{v}_i$ are the respective position and velocity vectors for the plasma particles. The fields $\boldsymbol{E}(\boldsymbol{r}_i, t)$ and $\boldsymbol{B}(\boldsymbol{r}_i, t)$ are described by Maxwell's Equations:

$$\nabla \cdot \boldsymbol{E} = \frac{\rho}{\epsilon_0} \tag{4.18}$$

$$\nabla \cdot \boldsymbol{B} = 0 \tag{4.19}$$

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t} \tag{4.20}$$

$$\nabla \times \boldsymbol{B} = \mu_0 \left(\boldsymbol{j} + \epsilon_0 \frac{\partial \boldsymbol{E}}{\partial t}\right) \tag{4.21}$$

Here, ρ is the charge density and $\boldsymbol{j}$ the charge current.

The kinetic approach

While the approach outlined above provides the most complete description of the plasma, it is simply untenable for a large collection of particles. A typical plasma must be treated macroscopically instead. One means of doing so is through introducing a distribution function f_s for the system, which describes the six dimensional phase-space density of each particle species, according to

$$n_s = \int f_s(\boldsymbol{r}, \boldsymbol{v}, t) d\boldsymbol{v} \tag{4.22}$$

where n_s is the particle density for particle species s, and the zeroth moment of $f_s(\boldsymbol{r}, \boldsymbol{v}, t)$. The first moment of the distribution function describes the bulk velocity, $\boldsymbol{u}$ of a given species,

$$u_s = \frac{1}{n} \int \boldsymbol{v} f_s(\boldsymbol{r}, \boldsymbol{v}, t) d\boldsymbol{v} \tag{4.23}$$

Equivalently, the zeroth and first moment nicely relate the species charge density, ρ_s, and current density, $\boldsymbol{j}$, respectively

$$\rho_s = q \int f_s(\boldsymbol{r}, \boldsymbol{v}, t)) d\boldsymbol{v} \tag{4.24}$$

$$\boldsymbol{j}_s = q \int \boldsymbol{v} f_s(\boldsymbol{r}, \boldsymbol{v}, t) d\boldsymbol{v} \tag{4.25}$$

As f_s describes a physical ensemble, continuity of the phase space must hold. From this assumption follows the Boltzmann kinetic equation,

$$\frac{\partial f_s}{\partial t} + \boldsymbol{v} \cdot \nabla f_s + \boldsymbol{a}_s \cdot \nabla_v f_s = \left(\frac{\partial f_s}{\partial t} \right)_c \tag{4.26}$$

where $\boldsymbol{a}_s$ describes the force operator on the system, and the term $(\partial f_s/\partial t)_c$ describes the collisional forces between particles in the distribution. For a plasma, $\boldsymbol{a}_s$ is simply the Lorentz Force operator as seen in Equation 4.17. Furthermore, since we are interested in collisionless plasmas, we may set the collisional term to 0,

$$\left(\frac{\partial f_s}{\partial t}\right)_c = 0 \tag{4.27}$$

$$\frac{\partial f_s}{\partial t} + \boldsymbol{v} \cdot \nabla f_s + \boldsymbol{a}_s \cdot \nabla_v f_s = 0 \tag{4.28}$$

This is the Vlaslov equation, and when combined with Equations 4.18 – 4.21 produces the Vlaslov-Maxwell description of a kinetic plasma. Though more tractable than a microscopic analysis, analytical solutions to the model are extremely limited in their scope. In practice, numerical approaches are used to solve this system of equations, subject to realistic boundary conditions. The most well known application is the Particle-in-Cell approach, in which the system of equations is solved for pseudo-particles interacting with electric and magnetic fields defined on a fixed mesh [57].

The hydrodynamic approach

A complementary strategy is to describe the plasma as a fluid of charge, with macroscopic properties. We may begin again with the same distribution function f along with the Boltzmann equation. However, instead of solving the equations to produce individual particle solutions, we solve for the evolution of the ensemble plasma properties.

Recall that the moments of f correspond to observable properties which provide a description of the plasma "fluid." The particle density, n_s, is given by the zeroth moment (Equation 4.22), the bulk velocity, $\boldsymbol{u}_s$, is given by the first moment (Equation 4.23), and the pressure tensor, $\boldsymbol{p}_s$, is given by the second moment,

$$\boldsymbol{p}_s(\boldsymbol{r}, t) = m_s \int (\boldsymbol{v} - \boldsymbol{u}_s)(\boldsymbol{v} - \boldsymbol{u}_s) f_s(\boldsymbol{r}, \boldsymbol{v}, t)) d\boldsymbol{v} \tag{4.29}$$

Note that the pressure tensor is a measure of thermal motion, and if the particles all move with a steady velocity, then $\boldsymbol{p_s} = 0$ There is also a third order moment, known as the heat flux density, $\boldsymbol{q_s}$,

$$\boldsymbol{q}_s(\boldsymbol{r}, t) = \frac{1}{2} m_s \int \|\boldsymbol{v} - \boldsymbol{u}_s\|^2 (\boldsymbol{v} - \boldsymbol{u}_s) f_s(\boldsymbol{r}, \boldsymbol{v}, t)) d\boldsymbol{v} \tag{4.30}$$

The evolution of these quantities are described by the fluid equations, which themselves arise from moments of the Boltzmann equation (Equation 4.26), each of which relates a conserved quantity of the system. The zeroth order moment produces the continuity equation

$$\frac{dn_s}{dt} - \nabla \cdot (n_s \boldsymbol{u_s}) = 0 \tag{4.31}$$

which describes the conservation of mass. The first order moment produces the momentum transport equation, expressing conservation of momentum,

$$m_s n_s \left[\frac{\partial \boldsymbol{u}_s}{\partial t} + (\boldsymbol{u} \cdot \nabla) \boldsymbol{u} \right] = q_s n_s (\boldsymbol{E} + \boldsymbol{u}_s \times \boldsymbol{B}) - \nabla \cdot \boldsymbol{p}_s + \Phi_s \tag{4.32}$$

where Φ_s is the total collisional friction force exerted on a given species in the system.

As a brief example, consider a system in thermodynamic equilibrium with temperature T. The distribution function is given by the standard Maxwellian distribution,

$$f(\boldsymbol{v}) = n \left(\frac{m_s}{2\pi k_B T} \right)^{3/2} \exp \left(-\frac{m_s \boldsymbol{v}^2}{2k_B T} \right) \tag{4.33}$$

We note that solving for the pressure tensor (Equation 4.29) gives the expected relation $P_s = n_s k_B T$ for an ideal gas.

4.2 Laser-plasma interactions

The origin of the plasmas we are considering in this work is the interaction of an intense laser pulse with neutral matter. It is valuable then to consider some of the basic parameters of the laser-interaction that ionizes and heats the target, as well as the properties of the pulse as it propagates into the plasma.

4.2.1 Ionization mechanisms

In most instances, the target of an intense laser pulse is not pre-ionized, meaning that the laser is responsible for forming the plasma as well as driving the interaction mechanism of interest. To that end, it is worthwhile to summarize the basic ionization mechanisms which are triggered by an intense laser interacting with a neutral system.

Multi-photon ionization

Hydrogen-like atoms in the ground state have ionization potentials of the form $U(Z) \approx 13.6 \times Z^2 [\text{eV}]$. As most lasers operate in the optical range, photon energies tend to be less than a few eV. Thus, it generally requires more energy

than can be delivered by a single photon to ionize them. Fortunately, if a large number of photons are absorbed in a short period of time, then ionization may still occur. This is known as "multi-photon ionization" [63].

Collisional ionization

Free electrons in a plasma may collide with other electrons. In a partially ionized plasma, it is possible that a sufficiently energetic electron may collide with a bound electron, imparting enough energy to ionize the atom while remaining free itself [63]. This process is known as collisional ionization. The rate of ionizing collisions, ν_c follows from

$$\nu_c \simeq n_e v_e 4\pi a_b{}^2 \left(\frac{U_H{}^2}{U_k k_B T_e} \right) \ln \left(\frac{k_B T_e}{U_k} \right) \tag{4.34}$$

where a_b is the Bohr radius, U_H is the ionization potential of the Hydrogen atom in the ground state, and U_k is the ionization potential of the constituent plasma species being ionized.

Barrier suppression and tunnel ionization

There are additional routes to ionization which aren't reliant upon direct energy transfer to electrons. In particular, an external laser field can perturb the Coloumb potential of an atom, providing a field-dependent modification to the Coloumb potential

$$V(x) = -Ze/|x| - E_x x \tag{4.35}$$

If the field is strong enough, an electron may escape with less kinetic energy than the ionization potential. This is known as barrier suppression ionization, and it occurs if the laser field exceeds the threshold set by,

$$I_{BS} = \frac{U_k{}^4 c\pi\epsilon_0{}^3}{2q^2 e^6} \approx \frac{4 \times 10^9}{q^2} \frac{U_k^4}{[eV]} [Wcm^{-2}] \tag{4.36}$$

For a hydrogen atom in its ground state, $I_{BS} \approx 10^{14}$ Wcm^{-2}. However, even if the laser field is below this thresold, suppression of the Coloumb barrier increases the probability of ionization through tunneling of the elctron. The rate of tunnel ionization is given by the Ammosov-Delone-Krainov (ADK) formula [64],

$$\nu_{ADK} \approx 6.6 \times 10^{16} \frac{q^2}{n_{ef}^{4.5}} \exp\left(-\frac{2q^3}{3n_{ef}^3}\left[\frac{E_a}{E(t)}\right]^{1/2}\right)\left(10.87\frac{q^3}{n_{ef}^4}\left[\frac{E_a}{E(t)}\right]^{1/2}\right)^{(2n_{ef}-1.5)}$$

$$(4.37)$$

Keldysh parameter

Keldysh and Perelomov [65] introduced a parameter γ that separates the multiphoton and tunneling ionization regimes. The parameter is given by

$$\gamma = \omega_L \sqrt{\frac{2U_k}{I_L}} \tag{4.38}$$

where I_L is the laser intensity. For $\gamma > 1$, multiphoton ionization dominates, while for $\gamma < 1$, tunneling (and barrier suppression) ionization prevail.

4.2.2 Free electron motion

Due to the significant mass difference, electron motion is preferentially excited during laser-plasma interactions. Consider a free electron in plane wave propagating in $\hat{x}$ with frequency ω_L. The electric field is given by $\boldsymbol{E} = E_0\boldsymbol{\epsilon}\exp(kx - i\omega_L t)$ where $\boldsymbol{\epsilon}$ is the polarization vector. The equation of motion of the electron is given by Equation 4.17. The solution, to lowest order in $\boldsymbol{v}$ is simply

$$\boldsymbol{v} = \frac{e}{m_e\omega_L}\boldsymbol{E} \tag{4.39}$$

We define the dimensionless quantity a_0 to be the normalized peak oscillation amplitude $a_0 = v_0/c \equiv eE_0/m_e\omega_p c$ of the electron in the field. For non-relativistic motion, $a_0 \ll 1$ by definition. A consideration of relativistic motion leads to the more complete form

$$a_0 = \frac{\gamma v_0}{c} = \frac{p_0}{m_e c} \approx 0.85\left(\frac{I_{18}\lambda^2}{[\text{W cm}^{-2}]}\right) \tag{4.40}$$

where $I_{18} = I_L/10^{18}$ is the peak laser intensity in units of 10^{18} Wcm^{-2} and λ the laser wavelength $\lambda = 2\pi c/\omega_L$ in microns. Note that a_0 represents the peak relativistic momentum of the electron as it oscillates in the laser field.

Now we consider electron motion to second order to gain further insight into an exact solution. Again we begin with the Lorentz equation of motion,
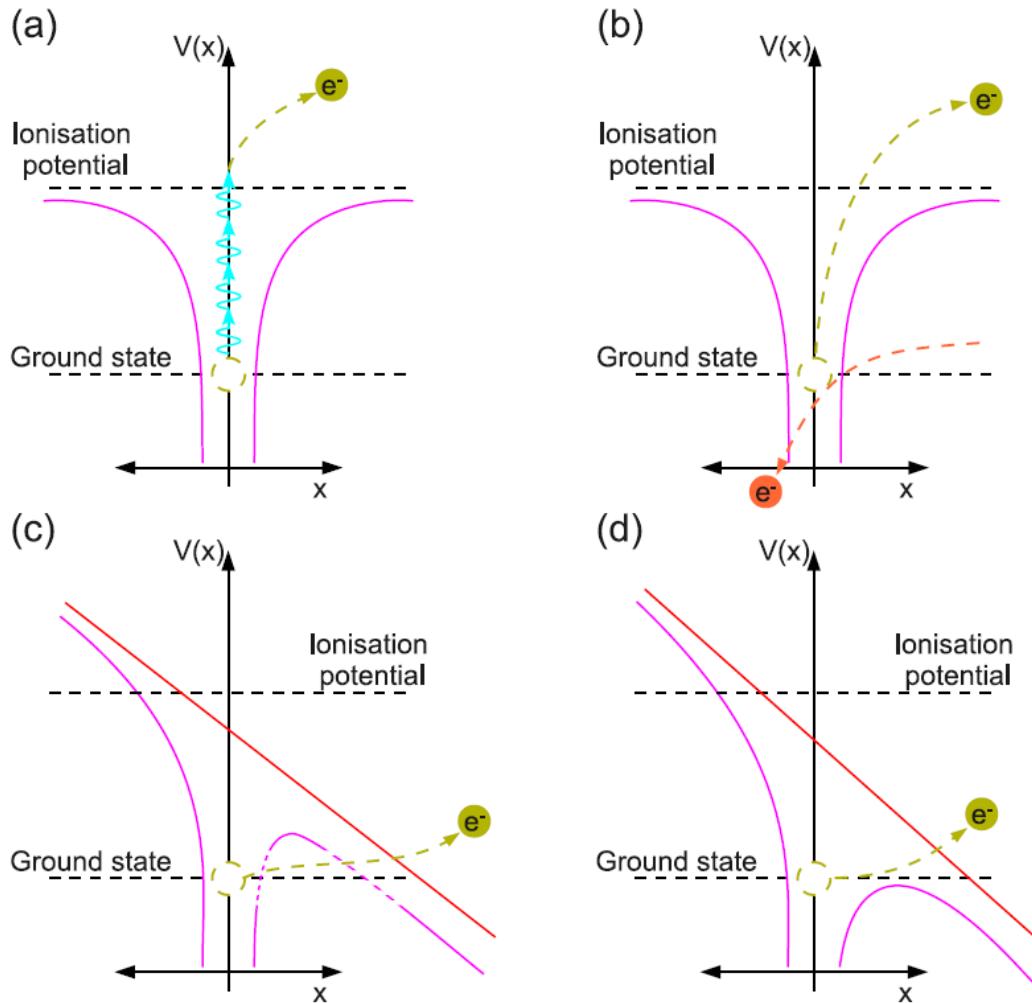
Figure 4.1: Different ionization mechanisms are illustrated The red line indicates the laser field while magenta lines show the binding field of the atom. (a) Multi-photon ionization. (b) Collisional ionization. (c) Tunnel Ionization. (d) Barrier suppression ionization. Courtesy Olivier Tresca.

Equation 4.17 alongside the total electron energy

$$\frac{d}{dt}(m_e \gamma c^2) = -e\boldsymbol{v} \cdot \boldsymbol{E} \tag{4.41}$$

Assume an elliptically polarized plane wave travelling in the x-direction, given by vector potential,

$$\boldsymbol{A} = A_0 \left[\hat{y}\delta cos(kx - \omega t) + \hat{z}(1 - \delta^2)sin(k - \omega t) \right] \tag{4.42}$$

where $\delta \leq 1$ specifies the polarization of the wave. We note that $\boldsymbol{A} \cdot \hat{x} = 0$, and immediately obtain from Equation 4.17 and equation for the transverse momentum $\boldsymbol{p}_\perp$

$$p_{y,z} = \frac{e}{c}A_{y,z} - p_{0y,0z} \tag{4.43}$$

From the energy equation along with the x-component of the Lorentz equation of motion, we obtain a relationship for the

$$m_e \gamma c - p_x = -\alpha \tag{4.44}$$

where α is a constant of motion, subject to initial conditions. If we assume the electron was initially at rest, and the corresponding vector potential was zero, we have $p_{0y,0z} = 0$ and $\alpha = -m_e c$. We can then solve for the the longitudinal momentum in terms of the vector potential.

$$p_x = m_e c(\gamma - 1)\frac{p_y{}^2 + p_z{}^2}{2m_e c} = \frac{1}{2m_e c}\left(\frac{e\boldsymbol{A}}{c}\right)^2 \tag{4.45}$$

We then obtain momentum equations for the particle in the laboratory frame

$$p_x = \frac{a_0{}^2}{4}\left[1 + (2\delta^2 - 1)\cos(2\phi)\right] \tag{4.46}$$

$$p_y = \delta a_0 \cos(\phi) \tag{4.47}$$

$$p_z = (1 - \delta^2)^{1/2}a_0 \sin(\phi) \tag{4.48}$$

Thus, for linear polarization ($\delta = 0,1$) the electron oscillates along the x-axis with frequency $w\omega$, whereas for circular polarization ($\delta = \pm\frac{1}{\sqrt{2}}$) the motion at 2ω is suppressed. Regardless, the electron exhibits average motion in the x direction according to $p_x \propto a_0{}^2/4$, which corresponds to the momentum transfer $p_x = m_e c(\gamma - 1)$.

Ponderomotive force

Of course realistic laser pulses are not well approximated by monochromatic plane waves. An envelope function with a finite width and duration better describes a pulsed beam. In this case, we can still draw conclusions about the motion of a particle subject to this pulse averaged over many cycles.

Assume that such a pulsed beam is described by the product of a slowly varying spatial envelope function $\tilde{E}(\boldsymbol{r}, t)$ with an oscillating temporal profile, as described by

$$\boldsymbol{E}(\boldsymbol{r}, t) = \tilde{E}(\boldsymbol{r}, t) \exp(-i\omega_L t) \tag{4.49}$$

Consider an electron sitting on the central axis of a laser at the position of peak intensity. As we've shown above, the electron oscillates in the polarization plane at the laser frequency. During the first half of the laser period, the electron is pushed away from the central axis. During the second half of the oscillation, the electron is pushed back to the center. However, in this time the laser intensity has diminished slightly due to the presence of the finite envelope. As the restoring force decreases with the envelope amplitude, the electron continues to drift from the central axis. The force responsible for this net drift is known as the ponderomotive force (or the nonlinear force). It results from the cycle averaged gradient in the field seen by the electron, and can be defined in the non-relativistic case by

$$F_p = -\frac{e^2}{4m_e\omega_L{}^2} \nabla \left\langle \boldsymbol{E^2} \right\rangle \equiv -\nabla \Phi_p \tag{4.50}$$

where Φ_p is termed the ponderomotive potential. For relativistic intensities, the ponderomotive potential is given by $\Phi_p = m_e c^2(\gamma - 1)$. We note that this potential is equivalent to that gained from the average motion in the plane wave solution given in Equation 4.45, relating the time-averaged plane wave solution and to "slow" component of the electron motion averaged over a wave envelope. Thus, the ponderomotive potential is in fact the cycle-averaged oscillation energy of the electron. A natural consequence of the ponderomotive force is that charged particles will be expelled from regions of high field, although the mass-dependence will preferentially select electrons.

The ponderomotive force is also referred to as the nonlinear force, as it results from the second order motion of an electron. Taking the solution from Equation 4.39 and introducing second order terms leads to the force on an electron of

$$f_2 \equiv f_p(1 + (\delta^2 - 1)\cos(2\omega_L t)) \tag{4.51}$$
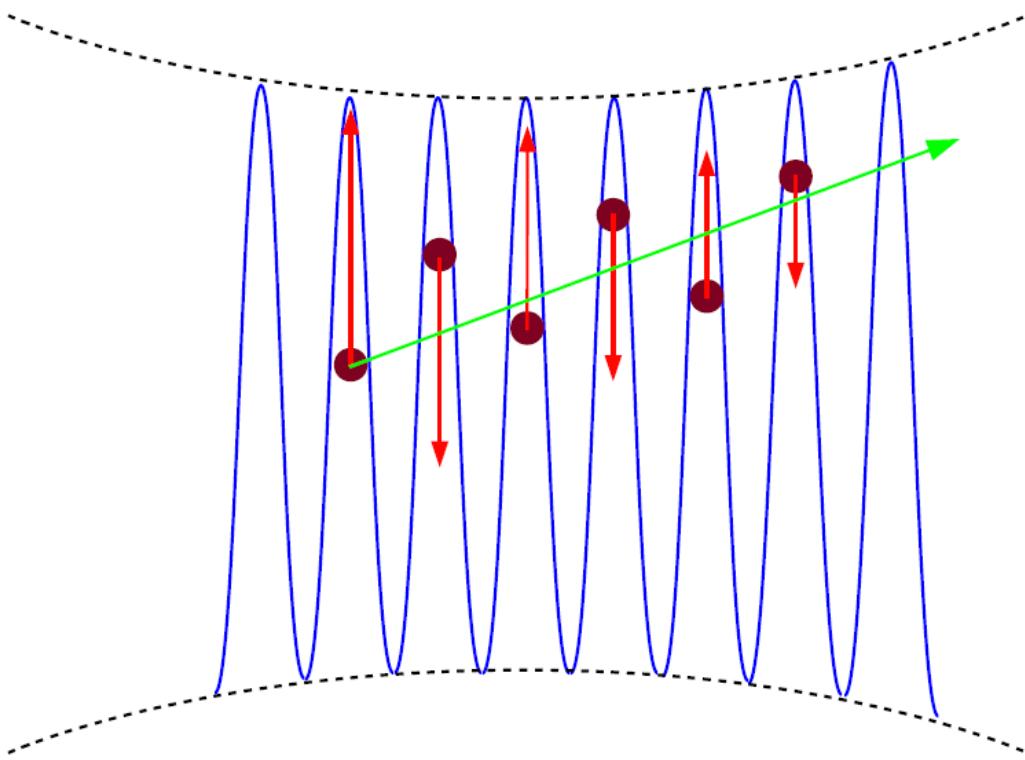
Figure 4.2: A free electron is depicted in a pulsed laser field. The electron experiences kicks (red) due to the laser oscillations in time, which over many cycles lead to a net drift (green), which is described by the ponderomotive force.

where δ is a polarization magnitude, in which $\delta = 0, \pm 1$ corresponds to linear polarization and $\delta = \pm 1/\sqrt{2}$ corresponds to circular polarization. The motion of the electrons at twice the laser frequency is known as quiver motion and plays an important role in plasma heating. Notice that when averaged over a laser cycle, Equation 4.51 reduces to the ponderomotive force.

4.2.3 Propagation of waves in plasma

Consider a plane wave of the form $\boldsymbol{E}_L = \boldsymbol{E}(x)\exp(-i\omega_L t)$ propagating along $\hat{x}$ in an un-magnetized collisionless plasma. We will assume that ω_L is large enough to consider the ions in the plasma as stationary on the timescale of the laser period. Following Equation 4.32 to first order in $\boldsymbol{u}$, the equation of motion for the fluid is

$$\frac{\partial \boldsymbol{u}}{\partial t} \simeq -\frac{e}{m_e}\boldsymbol{E}_L \tag{4.52}$$

Recall from Equations 4.23 and 4.25 that the current density is $\boldsymbol{j} = -n_e e \partial \boldsymbol{u}/\partial t$. We have,

$$\frac{\partial \boldsymbol{j}}{\partial t} = -\frac{n_e e^2}{m_e}\boldsymbol{E}_L = \boldsymbol{\sigma}\boldsymbol{E} \tag{4.53}$$

where $\boldsymbol{\sigma} = \dfrac{i\omega_p{}^2}{4\pi\omega_L}$ is the plasma conductivity. We can now describe the propagation of the laser field in the plasma. Application of Ampere's Law and Faraday's Law (Equations 4.21 and 4.20) along with some cross product identities produce the wave equation,

$$\nabla^2 \boldsymbol{E} + \frac{\omega_L}{c^2}\epsilon\boldsymbol{E} = 0 \tag{4.54}$$

where

$$\epsilon = 1 - \frac{\omega_p{}^2}{\omega_L{}^2} \tag{4.55}$$

Assuming plane wave solutions of the form $\boldsymbol{E}_x = \exp(ikx)$, Equation 4.54 gives

$$k^2 c^2 = \omega_L{}^2 - \omega_p{}^2 \tag{4.56}$$

It is clear that for $\omega_L < \omega_p$, the wavenumber k is not real, implying that a physical wave cannot propagate. Thus, ω_p sets a minimum frequency for propagation of light in plasma. Conversely, we may define a maximum plasma density above which light may not penetrate. This is known as the critical density n_c, and is given by

$$n_c = \frac{m_e \omega_p{}^2}{4\pi e^2} \approx \frac{1.1 \times 10^{21}}{\lambda_{\mu m}{}^2} \ [\text{cm}^{-3}] \qquad (4.57)$$

where $\lambda_{\mu m}$ is the vacuum wavelength of the laser in μm. In the case of an infrared pulse with $\lambda = 1$ μm, the critical density is $n_c = 10^{21}$ cm^{-3}. The critical density effective separates a plasma into two regions: an underdense $(n_e < n_c)$ region in which the laser propagates through the plasma, and an overdense $(n_e > n_c)$ region which the laser cannot penetrate. The boundary between these two regions has $n_e = n_c$ and is known as the critical surface. The critical surface marks the point in the plasma where the laser is effectively reflected, and is characterized by a very sharp spike in density due to the pile-up of energetic electrons under the influence of the laser field. This is a good assumption for most intense laser plasma interactions (LPIs) because the ponderomotive potential of the laser typically exceeds the thermal pressure $n_e T_e$ exerted by the overdense plasma [66].

At relativistic intensities $(a_0 \gtrsim 1)$, the increase in electron mass produces a complementary increase in n_e, allowing the laser to propagate further. This is known as "self-induced" transparency, leading to the relativistically adjusted formula for critical density

$$n_c = \gamma(a_0)n_{c0} \approx (1 + \frac{a_0{}^2}{2})^{1/2} n_{c0} \qquad (4.58)$$

where $\gamma(a_0)$ corresponds to the relativistic quiver motion of the electrons in a linearly polarized laser field [67, 68]. For circular polarization, $a_0^2/2 \to a_0^2$.

Skin depth

Beyond the critical surface, the laser field may still penetrate through a skin effect, resulting in evanescent waves of magnitude

$$\boldsymbol{E} = E_0 \exp(-x/l_s) \qquad (4.59)$$

where l_s is the collisionless skin depth [63] defined by

$$l_s = \frac{c}{\omega_p}\left(1 - \frac{\omega^2}{\omega_p{}^2}\cos^2(\theta)\right)^{1/2} \qquad (4.60)$$

Here, θ is the angle of incidence of the pulse with respect to the surface normal. For highly overdense plasmas, $l_s \approx c/\omega_p$.

Relativistic self-focussing

From Equation 4.55 we may define the index of refraction of the plasma

$$n_r = \sqrt{\epsilon} = \sqrt{1 - \frac{\omega_p^2}{\omega_L^2}} = \sqrt{1 - \frac{n_e}{n_c}} \qquad (4.61)$$

where we've substituted Equation 4.7 in to achieve the final form. We note that there is an inverse relationship between n_r and n_e. Consider an intense laser pulse propagating in the plasma. The ponderomotive force of the laser will expel electrons from high field regions, creating a reduced density channel along the laser axis. This channel acts as a focussing lens for the laser. As the laser naturally defocuses and exits the reduced density region, light refracts back towards the laser axis. If the intensity of the laser is strong enough to create a steep transverse density gradient, then total internal reflection may occur, and the laser may be further focussed. This effect is known as relativistic self-focussing. For a Gaussian beam penetrating a uniform density plasma, the threshold for self-focussing is given by

$$P_c = 17.5 \left(\frac{\omega_L}{\omega_p}\right)^2 [GW] \qquad (4.62)$$

Assuming a beam with power P, focal spot size σ_0, ideal Rayleigh length $z_R = \pi w_0^2 / \lambda$, the beam radius follows from

$$\langle r^2 \rangle = \sigma_0^2 \left[1 + \frac{z^2}{z_R^2}\left(1 - \frac{P}{P_c}\right)\right] \qquad (4.63)$$

For $P > P_c$, the beam collapses to zero radius in a distance

$$z_c = \frac{z_R}{\left(\frac{P}{P_c} - 1\right)^{1/2}} \qquad (4.64)$$

The proper derivation of this value makes use of the nonlinear Schrodinger equation, from which conservation of the Hamilton of the system as well as total beam power places a constraint on the beam size above a certain threshold, corresponding to self-focussing [63]. Such a derivation requires assumptions of a Gaussian beam of only weak relativistic amplitude ($a_0 \simeq 1$) incident on a smoothly varying (nearly uniform) density profile [57]. For few-cycle pulses and extreme intensities, Equation 4.62 does not hold.

4.2.4 Energy absorption mechanisms

The discussion above makes clear that the propagation of the laser into a plasma is a messy process, and that in overdense plasmas the laser should be reflected after minimal penetration. Nonetheless, significant heating is observed during laser plasma interactions. The laser interaction often produces two distinct Maxwellian populations, one with low temperature ($\sim$10s of eV), and another with very high temperature ($\sim$ 10s of keV to MeV). The processes through which the electrons are heated remains a subject of much research, with many creative and effective explanations developed over the past 20 years. We briefly introduce several of the most important mechanisms for this energy transfer.

Collisional absorption

The most basic means of energy transfer to the plasma are collisional absorption, via long range Coloumb interactions — also referred to as inverse-Bremsstrahlung [69]. However, as we we have seen in Section 4.1.1, collision rates follow a $\nu \sim T^{-3/2}$ relationship. This means that as the plasma absorbs energy through collisions, the cross section for further collisions is reduced. We conclude that at high laser intensities, collisional absorption is suppressed. The true source of hot electrons must be collisionless in nature.

Resonant absorption

One means of achieving collisionless heating is through the excitation of resonant plasma oscillations at $\omega_p = \omega_L$. This requires a laser with a field component parallel to to the target surface. In particular, a p-polarized electromagnetic wave incident at an angle to the surface may excite oscillations along the critical surface of the plasma. A caveat to this mechanism is that when at angle θ_L with respect to the surface normal, the laser is only able to penetrate to a density of $n_c\cos^2(\theta_L)$ due to the diminished field. As a result, exciting resonance at the "true" critical plane requires excitement by evanescent waves beyond the laser penetration distance; maximizing resonant absorption requires choosing the laser angle of incidence to maximize the perpendicular electric field component of these evanescent waves at the target critical surface.

Vacuum heating

There is another form of "not so resonant" heating, proposed by Brunel [70]. In the presence of very short-scale length density gradients (such as a solid-vacuum interface), electrons oscillating in the plasma may be pushed from the

plasma within a single laser wavelength, returning with energy roughly equal to the oscillation energy. Alternatively, electrons may be pushed beyond the critical surface threshold through the same means, sending hot electrons each laser cycle into the overdense region of the plasma. Because the electron is driven across the barrier within a single oscillation period, the excitation is non-resonant. Similarly to resonant absorption, this effect requires excitation of plasma oscillations normal to the target surface, and as such necessitates p-polarized laser light incident at an angle θ_L with respect to the surface normal. For modest laser intensities absorption fractions of $\eta \approx 80\%$ can be achieved and vacuum heating is a primary mechanism of energy deposition. But for high laser intensities and extremely thin targets, absorption tends to saturate at 15%, suggesting that yet other mechanisms may be responsible for the observed hot electron populations [63].

$\boldsymbol{j} \times \boldsymbol{B}$ heating

The above absorption mechanisms cannot explain the production of hot electrons from non p-polarized waves at normal incidence. However, the ponderomotive force can also produce heating at a similar scale to the Brunel effect [71, 72]. Recall from Equation 4.48 that electron quiver motion occurs at twice the laser frequency. These electrons may escape the laser field into neighboring overdense regions of the plasma, producing a hot electron population outside of the laser region. This is known as $\boldsymbol{j} \times \boldsymbol{B}$ (or ponderomotive) heating because the $2\omega_L$ term in 4.48 arises from the contribution of the $\boldsymbol{j} \times \boldsymbol{B}$ term in electron equation of motion. For relativistic laser intensities, $\boldsymbol{j} \times \boldsymbol{B}$ heating is responsible for the hot electron population observed in many experiments with foil targets [72]. We can estimate this hot electron temperature by assuming the electrons acquire energy roughly equal to the ponderomotive potential from which they escaped [73]. Thus,

$$T_{hot} = (\gamma - 1)\, m_e c^2 \approx \left[\sqrt{1 + 0.731 I_{18} \lambda_{\mu m}^2} - 1 \right] m_e c^2 \qquad (4.65)$$

$$T_{hot} \approx 0.47 a_0^{2/3} \qquad (4.66)$$

where the numerical approximation comes from substituting 4.40 in to the ponderomotive potential. We note that this mechanism is polarization dependent. For circular polarization, the oscillating term vanishes, and heating is suppressed. Circularly polarized light is thus preferred when attempting to minimize electron temperature in the plasma. Unlike vacuum heating and resonant absorption, ponderomotive heating may occur in absence of p-

polarization and with normal laser incidence. In these instances, as well as for relativistic laser intensities, ponderomotive heating tends to be the dominant mechanism [63].

While ponderomotive scaling provides an intuitive explanation for hot electron production, experimental measurements show deviations from the predicted spectrum, with hot electron temperatures instead following [73, 74]

$$T_{hot} = 215 \left(I_{18}\lambda_{\mu m}^2\right)^{1/3} [\text{keV}] \qquad (4.67)$$

While Equations 4.66 and 4.67 agree to within 30% within a broad range of laser intensities ($a_0 < 30$), the Ponderomotive scaling consistently overestimates the hot electron temperature seen in experiments. Using a one dimensional "black-box" model, Haines et al. [73] were able to explain the lower scaling based on relativistic momentum transfer between the laser and target in a steady state "long pulse" regime. Their model produces scalings in agreement with Equation 4.67. The reduction in hot electron temperature can be explained heuristically by the motion of a plasma electron in the overdense target; at low intensities, the electron will travel a distance much greater than the collisionless skin depth of the target in a single laser cycle, thus preventing it from acquiring the full ponderomotive potential. As the laser intensity becomes highly relativistic ($a_0 \gtrsim 180$), the skin depth gradually increases and ponderomotive scaling takes over. Their model predicts the possibility of laser absorbtion fractions of up to 80–90%, which has been observed in recent high intensity laser studies [75]. Electron heating remains a subject of great interest for studies of fast ignition and hot, dense plasma behavior [76].

It should be noted that hot electrons cannot be limitlessly produced within a plasma, regardless of intensity or mechanism. There is in fact an upper limit for current propagation in a uniform quasi-neutral plasma is known as the Alfvèn limit, above which self generated magnetic fields within the plasma counteract the flow of charge [77, 78]. The absorption mechanisms of the kind described above predict the generation of a hot electron current which can exceed that limit. Thus, there must exist a counter-propagating current which reduces the net flow of charge through the plasma. In a plasma target, cool electrons in the upstream plasma are responsible for generating this return current, allowing hot electrons to penetrate and heat the target upstream. However, the presence of two opposing electron currents can give rise to a variety of instabilities (Buneman, Rayleigh-Taylor, and Weibel instabilities) which can have significant implications on acceleration dynamics [63, 76].

4.3 Ion acceleration mechanisms

The mechanisms responsible for accelerating beams of particles from a plasma are remarkably diverse, coinciding with the array of laser conditions, plasma properties, and target choices available in experiment. Moreover, the different mechanisms often occur in combination, making a pure study of a single process quite difficult. Nonetheless, simple models have been developed to describe several of the most important acceleration modes utilized in modern laboratories, and a brief overview of the particulars of each will be presented. Of particular note, it will be shown how the target structure and basic laser properties are essential to preferentially triggering a given mechanism. The discussion will be limited to the acceleration of ions from overdense or nearly overdense plasmas, although it should be noted that acceleration of electrons in underdense plasmas via wakefield acceleration remains the most common laser acceleration application [79].

4.4 Target normal sheath acceleration

The most popular intense lasers utilize Ti:Sapphire as the gain medium, owing to its exceptional bandwidth and durability [80]. The Ti:Sapphire gain bandwidth is centered at ~ 800 nm wavelengths, corresponding to a critical density of $n_c \approx 1.4 \times 10^{21}$ cm^{-3}. As a result, in order to create an overdense plasma, solid materials such as metal foils must be used to achieve necessary densities. The acceleration mechanism which often dominates the interaction between a laser and foil is known as target normal sheath acceleration (TNSA).

Figure 4.3 diagrams the basic TNSA process. During the initial laser interaction, laser energy is absorbed through resonant heating mechanisms, creating a hot electron population. Characteristically, these electrons are hot enough to be considered collisionless over the target length. They travel to the rear surface of the target, eventually escaping into vacuum. There, the hot electrons form a charged cloud, resulting in a strong charge separation field between the electron cloud and the ions on the target surface. The electrons temperature follows a Maxwellian distribution, and as such the resulting sheath field is correlated to electron temperature, according to

$$E_{\text{sheath}} \simeq \frac{k_B T_e}{e \lambda_D} \tag{4.68}$$

For typical laser intensities, hot electron temperatures can reach the MeV-range, resulting in sheath fields on the order of TV/m. The resulting energy spectrum is modulated by the Maxwellian electric field distribution, and fol-
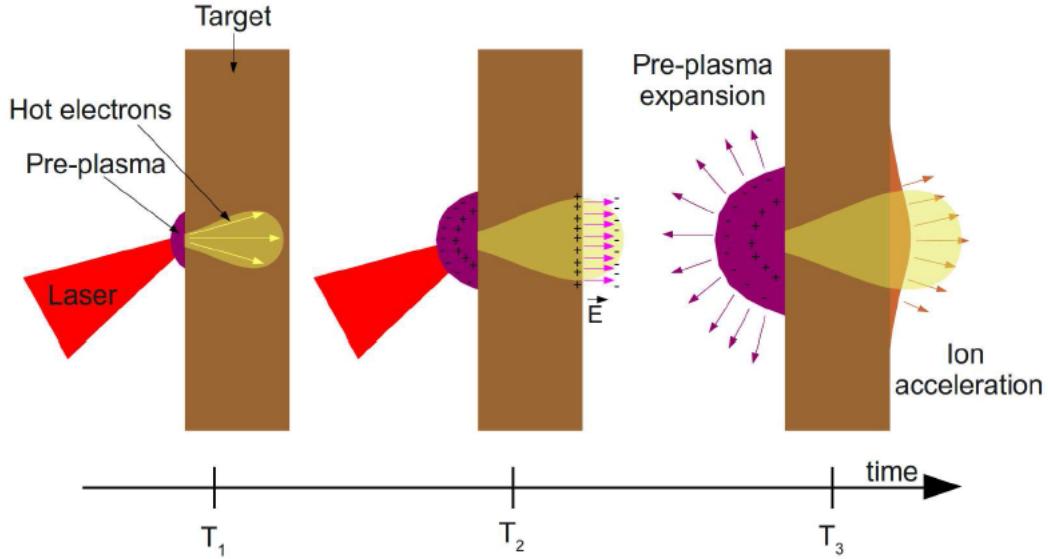
Figure 4.3: An illustration of the TNSA process. Hot electrons escaping into the vacuum behind the target generate a quasi-electrostatic field, which accelerates ions from the rear of the target.

lows from

$$N_i(\epsilon) = N_{i0}e^{\dfrac{-\epsilon}{k_B T_e}} \tag{4.69}$$

Recall that the peak electron temperatures follow from pondermotive heating, which scales with laser intensity (or irradiance) as $T_e \propto \sqrt{a_0} \sim \sqrt{I\lambda^2}$. As such, peak ion energies also scale as $\epsilon \propto \sqrt{a_0} \sim \sqrt{I\lambda^2}$ [81, 82]. Moreover, the transverse distributon of electrons along the surface of the foil is Gaussian. A consequence of this is that field strengths are significantly higher on the laser axis, resulting in beams that are spatially correlated in energy and divergence. Ions accelerated from the central axis peak field region have highest energies and lowest divergence, while the opposite is true of ions accelerated at a large distance from the laser axis.

Regardless of target choice, protons beams dominate the output of TNSA in most circumstances, due to the presence of hydrocarbon contamination on the surface of foils under standard vacuum conditions. Proton contaminants are preferentially accelerated owing to their high charge to mass ratio. Such contaminants can be removed using careful heating and laser ablation techniques [83]. Nonetheless, this feature remains a primary challenge to creating pure ion beams using the TNSA mechanism. Figure 4.4 shows an example

output spectrum of protons accelerated via TNSA from an experiment at the ATF [5].
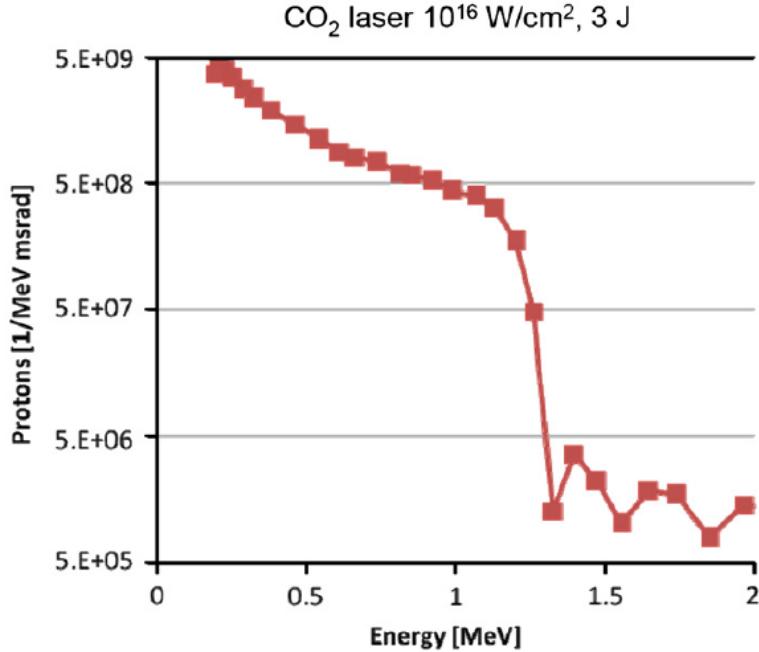
Figure 4.4: A characteristic proton acceleration spectrum from the interaction between a laser and aluminum foil. Experiments were performed at the Accelerator Test Facility using a 10.3 μm CO_2 laser with 1 TW peak power and a 50 μm foil target. The thermal distribution of proton energies is characteristic of TNSA [5].

4.5 Radiation pressure acceleration

While TNSA is conveniently achieved using current technologies, the drawbacks of beam quality, purity, and energy scaling have piqued interest in alternative mechanisms. One strategy is to utilize the momentum carried by the laser fields to efficiently transfer energy to the plasma. This is the premise behind radiation pressure acceleration (RPA).

For thick targets at $n_e \geq n_c$, the radiation pressure acts to push electrons into the target, creating a critical surface within the plasma. Over time, the laser pressure will continue to push the critical surface into the plasma, causing acceleration of ions due to charge separation at the laser front. This mechanism is known as "hole-boring" [84]. To better understand the dynamics, we will

consider a simple steady-state model and determine the motion of the critical surface under the influence of an intense laser.

The radiation pressure of a normally incident laser, in the plane wave approximation, is given by

$$P_{rad} = (1 + R - T)\frac{I_L}{c} \tag{4.70}$$

where R and T are the reflectance and transmittance of the target, respectively.

We will first consider the target to be a perfectly reflecting plasma mirror in the non-relativistic regime, wherein $v_{hb} \ll c$. The plasma will be assumed to be collisionless, which is a reasonable assumption. In this case, the plasma is seen to be reflecting from the laser interaction point at velocity v_{hb}, imparting momentum $2n_i m_i v_{hb}^2$, while the the laser radiation pressure is given by the ideal form $P_{rad} = 2I/c$. We have

$$\frac{2I}{c} = 2n_i m_i v_{hb}^2 \tag{4.71}$$

We define the dimensionless piston parameter Π

$$\Pi = \frac{I_L}{m_i n_i c^3} \propto a_0{}^2 \tag{4.72}$$

and determine the hole-boring velocity as

$$v_{hb} = \sqrt{\frac{I}{m_i n_i c}} = c\sqrt{\Pi} = a_0 c\sqrt{\frac{Z m_e n_c}{A m_i n_i}} \tag{4.73}$$

In the lab frame, the reflected ions have accelerated to $v = 2v_{hb}$, gaining energy

$$\epsilon_{max} = \frac{2I}{n_i c} = 2m_i c^2 \Pi \tag{4.74}$$

Thus the peak velocity achievable through hole-boring scales with I/n. This equation is correct for non-relativistic motion. As v_{hb} approaches c, some slight modifications must be made [85]. The RHS of 4.71 must include a factor of γ_{hb} to adjust for relativistic momentum. Also, the laser intensity in the reference frame of the plasma mirror, I_m, is no longer equal to that of the rest frame, I, due to the relativistic doppler shift

$$\frac{I_m}{I} = \frac{1 - (v_{hb}/c)}{1 + v_{hb}/c)} \tag{4.75}$$

The doppler shift of the laser frequency during hole-boring at high intensities has been a subject of considerable study, as it can provide a useful indirect

measurement of the hole-boring velocity with strong temporal resolution [86]. These relativistic adjustments lead to a refinement of the hole-boring velocity and resulting reflected ion energy

$$v_{hb} = c\frac{\Pi^{1/2}}{1 + \Pi^{1/2}} \tag{4.76}$$

$$\epsilon_{max} = 2m_p c^2 \frac{\Pi}{1 + \Pi^{1/2}} \tag{4.77}$$

This expression is valid for relativistic intensities, the so-called "laser-piston" regime. In comparison with TNSA, it is evident that peak energies scale linearly with laser intensity rather than with the square root of intensity (a_0 v. $a_0{}^2$), which is an immediate benefit to accessing the RPA regime.

While the above model provides a good estimation of peak accelerated ion energies, it is worth briefly discussing the dynamics of the charge separation layer. Again we consider a plane wave incident normal to the target surface. The radiation pressure acts to push electrons further into the target, resulting in a depletion layer extending from $x = 0$ to $x = d$ consisting of the heavier ions, and a corresponding compression region of width l_s. This creates a charge separation field which pulls the ions in the depletion layer. Simultaneously, the increase in electron density in the compression region generates a counterbalancing electric field which permeates the overdense plasma a distance of $\sim l_s$. This situation is depicted in Figure 4.5 (I). As time continues, the radiation pressure of the laser continues to squeeze the compression layer. It can be shown that to a linear approximation, the density and electric field within the compression region are described by a self-similar solution, which predicts that the ions will continue to pile up until a singularity in ion density is reached at position $x = d + l_s$. When the singularity is approached, the compression region collapses ("wave-breaking") and the corresponding ions in the compression layer will be accelerated downstream to a maximum energy of $2v_{hb}$. The specific spectral properties of the beam depend on the specific dynamics of the collapse, and are often inferred through simulation for varying laser and target parameters [87]. If the laser pulse is long enough, an equilibrium will be re-established, and the entire process will repeat itself, leading to pulsed acceleration of target ions.

The assumption of perfect reflectivity means that the electron density profile will be coherently steepened, producing a uniform critical surface to which laser energy is perfectly coupled. This is clearly unrealistic, as non-uniformities in the laser and target profile can drive a variety of plasma instabilities, particularly in any underdense ramp preceding the critical surface. The laser
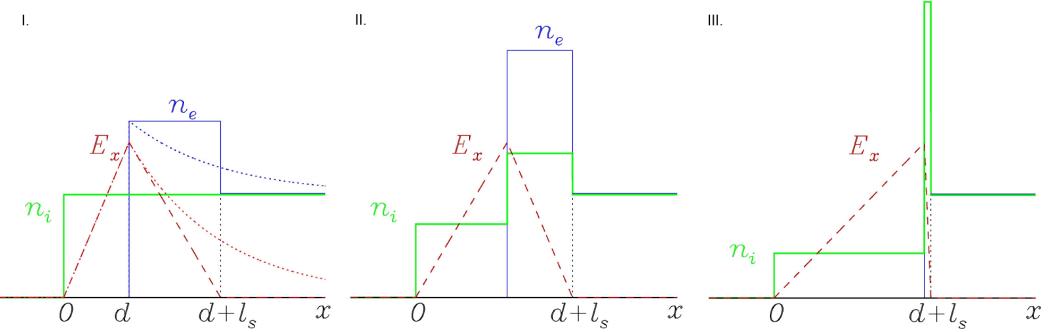
Figure 4.5: Illustration of the hole-boring process. (I) The laser field pushes electrons into the plasma, which build up at the critical surface. (II) As more electrons are pushed into the plasma, ions are correspondingly dragged in to balance the latent electric field. (III) Electrons and ions are eventually compressed into a thin layer of extremely high density at the laser critical surface with a correspondingly steep gradient in electric field, capable of accelerating ions to high energies.

polarization is thus of critical importance to driving a long term hole-boring process. For linear laser polarizations, electron heating triggers target expansion, strong laser absorption into hot electrons, and uneven propagation of the laser into the target. For all but ultra-relativistic laser intensities, the resulting ion acceleration will be dominated by the TNSA mechanism [88, 89]. It is thus common to use circular polarization to achieve RPA acceleration, as this suppresses $\boldsymbol{j} \times \boldsymbol{B}$ heating. Experiments performed at the ATF in 2009 demonstrated RPA driven acceleration using a circularly polarized CO_2 laser and near-critical density gas jet target [6]. Figure 4.6 shows the resulting proton spectra for several shots.

For realistic intensities and foil target structures, sub-micron scale thicknesses are optimal for efficient acceleration of ions by radiation pressure. For ultra relativistic intensities ($a_0 \gg 1$) the onset of transparency limits the target thickness to $d_{min} \simeq \lambda a_0 (n_c/n_e)$ [90]. Thus, for modest density foils, optimum target thickness at an intensity of $a_0 = 10$ is somewhere around $d \approx \lambda/4 = 250\mu m$ for a $1\mu m$ laser.

In the limit of these thin targets and extremely high laser intensities, the laser pressure can push forward the entire target structure coherently, in what is known as the light sail regime. Target electrons act as a relativistic mirror, reflecting the laser pulse energy with minimal parasitic absorption of laser energy, while dragging along ions via a strong charge separation field. This is possible only at highly relativistic laser intensities, in which the target can be accelerated to near the speed of light within a single laser period. As the laser
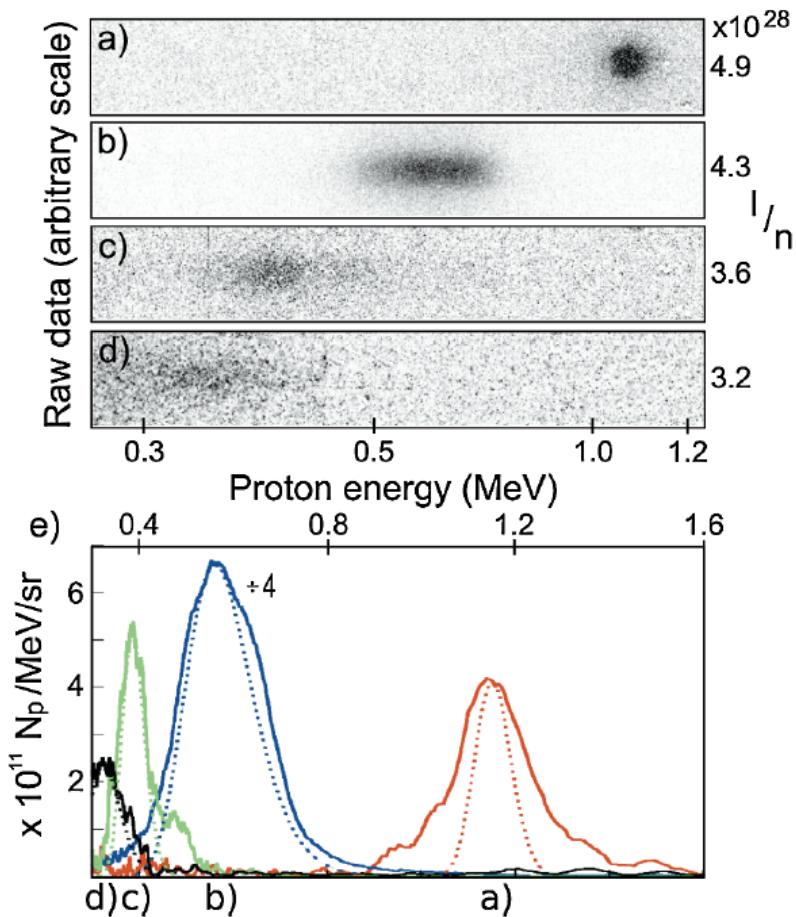
Figure 4.6: Results from experiments performed at the Accelerator Test Facility using a 10.3 μm CO_2 laser and near critical density hydrogen gas jet demonstrate signatures of radiation pressure acceleration using the hole-boring mechanism [6]. Raw and processed proton spectra are shown for varying peak density n and laser intensity I. Peak proton energies scale linearly with I/n, as predicted by hole-boring models.

interacts with the target, the group velocity of the light is slowed, providing a "phase stable acceleration over" a period longer than the laser pulse duration. This regime represents the single most efficient means of laser energy transfer, capable of accelerating particles to many GeV over a short distance. It is particularly efficient for accelerating heavy ions, as multi-species targets can be accelerated to the same energies [89]. However, this regime remains out of reach for modern laser and target technologies.

4.6   Collisionless shock acceleration

Electromagnetic shocks have been of great interest to astronomers for some time. The observation of disturbances surrounding the Earth's magnetosphere led to the description of the planetary bow shock [91], while the measurement of high energy cosmic rays was hypothesized to be the result of reflection from shocks generated by supernovae [92]. Though most atmospheric shocks dissipate energy through collisions across the highly viscous shock front, interstellar shocks have been observed to form across small length scales compared to the highly collisionless plasmas they traverse [93]. As a corollary, these cold, highly collisionless plasmas tend to be highly magnetized, and are termed high β plasmas, as their magnetic energy density is much larger than their thermal energy density. This is a notable distinction from the hot laser generated plasmas, which for most target structures we consider, tend to be dominated by thermal pressure and electric field effects.

As a result, the shocks of interest for laser particle acceleration are collisionless electrostatic shocks (CES). These shocks are characterized by a discontinuity (shock front) across which there is a positive jump in the electrostatic potential Φ in the plasma. The shock travels at the shock speed v_{sh}, effectively forming a moving potential barrier. Ions ahead of the shock may be reflected from the shock, gaining a velocity of $2v_{sh}$ in the direction of the shock.

We will review some basic theory regarding the generation of a collisionless shock to provide proper context for the remaining discussion. We'll restrict our discussion to purely electrostatic shocks in 1D, though there are many similarities in the treatment of each [94]. In our collisionless, unmagnetized plasma, we'll assume the initial electron population is isothermal at temperature T_e, and that its density can be described by a Boltzmann equilibrium $n_e = n_0 \exp\left(e\Phi/T_e\right)$, and furthermore that the electric field satisfies Poisson's equation (4.18) $d_x^2 \Phi = 4\pi e(n_e - Zn_i)$. The evolution of the plasma can be described by the continuity (Equation 4.31) and fluid (4.32) equations.

We look for solutions which characterize a shock moving at velocity v_{sh}, and as such move to the reference frame of the shock. In this frame, ions move

towards the shock at velocity v_{sh}, but the shock wave itself is stationary. As such, our fluid equations reduce to

$$n_i u_i = - \left(\frac{n_0}{Z} \right) v_{sh} \tag{4.78}$$

$$\frac{m_i u_i^2}{2} + Ze\Phi = \frac{m_i v_{sh}^2}{2} \tag{4.79}$$

where Φ is the potential across the shock, assuming $\Phi = 0$ in the unshocked region. We can eliminate u_i to obtain

$$n_i = \frac{n_0}{Z} \frac{v_{sh}}{\left(v_{sh}^2 - \frac{2Ze\Phi}{m} \right)^{1/2}} \tag{4.80}$$

We see that the solution is only real-valued if $0.5 m_i v_{sh}^2 > Ze\Phi$. That is, ions pass through the shock assuming their kinetic energy exceeds the electrostatic potential of the shock. This gives us a reflection condition for the shock,

$$Ze\Phi > 1/2 m_i v_i^2 \tag{4.81}$$

Introducing the dimensionless variables $\xi = x/\lambda_d$, $\phi = \exp\dfrac{\Phi}{T_e}$, $M = v_s h/c$, substitution of Equation 4.80 into Poisson's equation yields a differential equation which may be integrated to give

$$\frac{(d_\xi \phi)^2}{2} + U_1(\phi) = 0 \tag{4.82}$$

$$U_1(\phi) = - \left[e^\phi + M \left(M^2 - 2\phi \right)^{1/2} \right] + 1 + M^2 \tag{4.83}$$

with boundary condition $U_1(0) = 0$. The above equation effectively describes the motion of a particle in a potential $U_1(\phi)$, for which the motion can be described classically. This analogy was used by Sagdeev in solving for an array of acoustic wave solutions, and the construction and analysis of a funciton $U_1(\phi)$ is known as the pseudopotential method [95]. This method predicts distinct classes of wave solutions; each arises from a particular class of pseudopotential. If we first assume real-valued solutions to Equation 4.80, then we find a solution in which the pseudoparticle has zero energy, and is constrained to a potential in which it performs a single orbit with infinite period. The corresponding solution $\phi(x)$ is known as a soliton, a dispersionless wave propagation through the medium. The ideal soliton is symmetric, and the fluid state of the plasma is returned to its pre-soliton state after passage [57].

The complex nature of the Sagdeev potential makes the determination of a complete parameter range of shock solutions extremely difficult. Instead,

86

one-dimensional solutions are sought for various plasma states by assuming specific electron and ion distribution functions [96]. For an isothermal plasma, supersonic soliton solutions may arise for $1 < M < M_c$, where M_c is the critical Mach number corresponding to the onset of particle reflection. For the system described above, Equation 4.83 gives a critical Mach number corresponding to $e^{M^2/2} - M^2 - 1 = 0$, which yields $M_c = 1.585$. The shock wave solution to this problem is naturally dissipative, corresponding to an asymmetrical potential in which downstream particles are reflected from the shock front. A shock may only arise for $M > M_c$, and the corresponding transition layer is on the order of $c/\omega_p \ll \lambda_{\mathrm{mfp}}$ in the collisionless plasma. Figure 4.7 illustrates the potential characteristics of these two solutions.
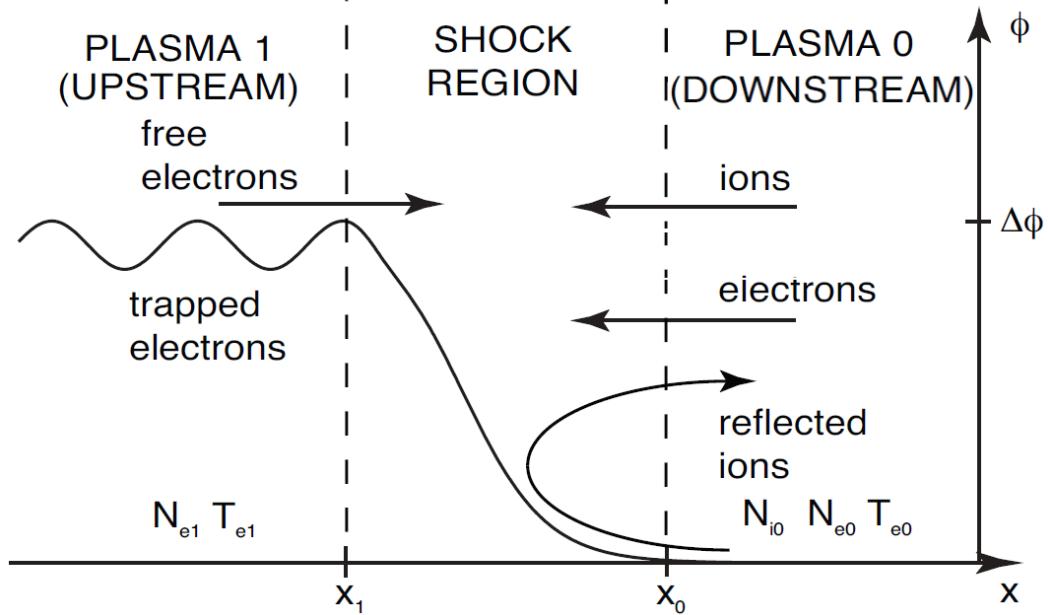
Figure 4.7: Diagram showcasing two characteristic features of electrostatic shocks - periodic potential wells behind the shock, leading to trapping, and ion reflection off the leading edge of the shock [7].

Generation of a CES in the plasma can occur through a variety of means. The potential for an intense short pulse laser to drive a shock has been considered for some time, but it is only with recent technology that the opportunity to generate such shocks in a laboratory has become possible [71, 97, 98]. Recall that the radiation pressure of the laser can act as a piston to drive forward a high density critical surface into the plasma with velocity u_{hb}. If $u_{hb} \gg c_s$, a shock may be generated which travels at this recessional velocity. Note though the requirement that $M > M_c$, and as such for hot plasmas with large c_s, a

high hole-boring velocity is needed to exceed the critical Mach number and launch a shock [99].

Alternatively, the laser may be used to generate a hot electron population at the focal position. Heating of a local electron population can create large gradients in temperature and density. Instead of the laser radiation pressure, the electron thermal pressure expanding into the surrounding plasma acts to generate a fast moving shock. For an expansion shock with Mach number M, the shock speed will strongly depend on the temperature, T_e, of the plasma, via the sound speed c_s, leading to a similar dependence of the reflected ion energy on temperature via $KE_i = 0.5 m_i v_{sh}^2 \sim 2ZM^2 T_e$.

For large ratios in density, the shock may achieve mach numbers $M \simeq \sqrt{2T_e/T_0}$, where T_e is the electron temperature of the hot, expanding plasma, and T_0 is the temperature in the downstream plasma [7]. For reasonable laser parameters, this means that a CES may be generated with speeds even greater than the hole-boring velocity. Figure 4.8 illustrates the potential dynamics for a shock that outpaces the critical surface. Regardless of the shock impetus, the theoretically achievable shock velocity scales favorably with laser intensity a_0, via $v_{hb} \propto a_0$ or $T_h \propto a_0$ for high intensities.

One of the signatures of shock acceleration is the persistence of reflection beyond the timescale of the laser interaction. This is an important distinction between the shock interaction and a hole-boring process, as the CES will be able to traverse a significantly larger portion of a target structure. Shock propagation is primarily affected through energy dissipation, for which reflection is a primary culprit. This has the particular consequence that for high reflection, the shock may accelerate a large flux of particles, but lose energy quickly, leading to a broad spread in energy. Conversely, in a hot uniform plasma, the shock may reflect only the least energetic downstream ions while travelling with a nearly constant velocity. Assuming a large enough thermal pressure driving the shock, this can produce a quasi-monoenergetic spectrum, but with relatively low total charge [100]. The characteristics of the plasma downstream of the shock is thus crucial to the long term shock dynamics. Figure 4.9 compares shock behavior for weakly reflective shocks in a hot isothermal plasma to that of highly reflective shocks in cold dense plasmas.

While the generation of a hot electron population is integral to achieving high plasma temperatures, it may subsequently lead to the formation of sheath fields at the rear of the target. In thin targets, and targets with short scale-length density ramps,the sheath fields strongly modulate the reflected ion population, leading to a broad TNSA spectrum [7, 101].
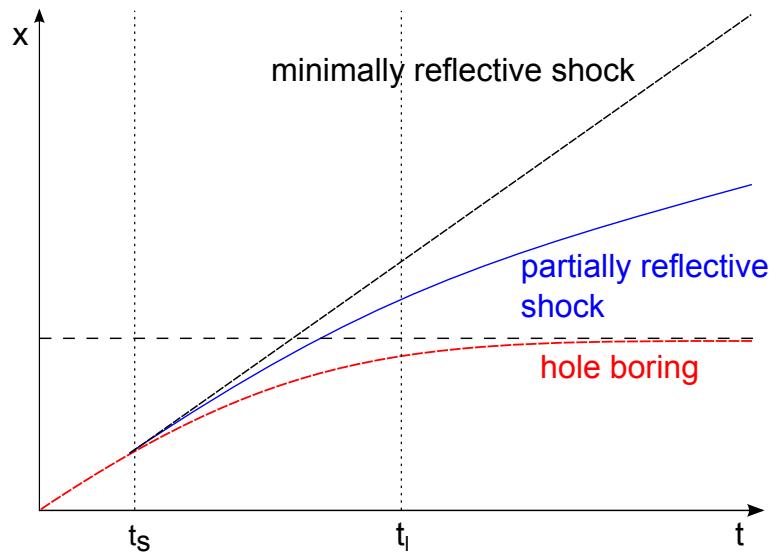
Figure 4.8: Illustration of the possible variation between critical surface and shock evolution for a laser driven shock. The laser is present from $t = 0$ to $t = t_l$, during which it bores a hole in the target, corresponding to the red dashed line. As the laser intensity wanes, v_{hb} is reduced, and as the laser terminates the critical surface eventually comes to rest. At $t = t_s$, a CES forms, propagating ahead of the shock into the downstream plasma. Depending upon the reservoir of pressure behind the shock, as well as the degree of reflection, the shock may continue with minimal dissipation (black line), or in the case of high dissipation, beam loading may reduce the shock velocity with time (blue line).
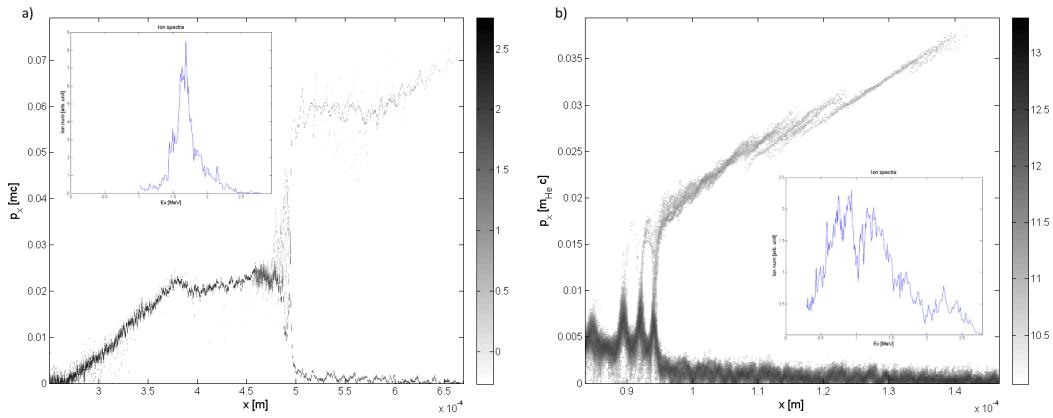
Figure 4.9: Phase space (x-px) diagrams from PIC simulations performed by the author illustrate two different outcomes of collisionless shock generation. In (a), the shock is generated by expansion into a uniformly hot plasma, and reflects a small fraction of ions, generating a quasi-monoenergetic spectrum (inset). In (b), a laser driven shock traverses a cold dense plasma, inducing very high reflection. This dissipation reduces the speed of the shock, broadening the output spectrum of the ions (inset).

90

Chapter 5

Experimental Configuration at BNL Accelerator Test Facility

Brookhaven National Laboratory's Accelerator Test Facility (ATF) provides a testing ground for a variety of advanced accelerator schemes. The facility features an electron beam linear accelerator in conjunction with Nd:YAG and CO_2 lasers for use in tandem or alone. The ion generation experiment makes use of the CO_2 laser as a drive beam, and the YAG for pulse slicing and diagnostics. The discussion of the experimental setup begins with the ATF CO_2 laser.

5.1   ATF Terrawatt CO_2 Laser

The ATF features a world class Terrawatt (TW) CO_2 laser system uniquely capable of producing a single 5 ps intense pulse at mid-IR wavelengths, between 9 and 11 μm. This system is the focal point of laser ion acceleration research at BNL, as well as a host of other experiments which make use of the advantages conferred by longer wavelength pulses, such as nonlinear Thomson scattering [102], inverse Compton scattering [103], inverse free electron lasers [104], and staged electron laser acceleration [105]. Producing ultrashort pulses is particularly challenging at 10μm, and the unique approach necessary to achieve it merits a brief introduction.

Isotopic amplification

The production of a short laser pulse is predicated on the gain medium possessing a smooth spectral bandwidth. The FWHM bandwidth $\delta\nu$ required to amplify a pulse of FWHM τ is given by

$$\tau_p = \beta/\delta\nu \tag{5.1}$$

where for a Gaussian pulse $\beta = 2\ln 2/\pi \approx 0.441$ [80]. The most popular gain mediums in high intensity lasers (e.g. Ti:Sapphire) feature large bandwidths and a high intensity damage threshold.

Although the CO_2 molecule has a long history of use in generating high power CW lasers for industrial applications [106, 107], the use of a gas medium gives rise to a number of unique design challenges. For short pulse lasers, a primary challenge is the spectral modulation of the CO_2 gain profile due to rotational excitations. At 10 μm, this rotational structure splits the gain profile into 1.3 to 1.8 cm^{-1} segments. This modulation prohibits the amplification of a single short pulse, producing instead a train of pulses.

A common means of improving the bandwidth of a gas laser is through collisional broadening. By increasing the pressure in the gas, the effective lifetime of the pumped state is reduced, as collisions stimulate radiative transitions on

timescales shorter than the natural lifetime. This results in a broadening of the spectral profile. For a CO_2 laser containing a mixture of CO_2, N_2, and He gases, the broadening is given by [108]

$$\delta\nu_p = 0.1149\,P_{CO2} + 0.0794\,P_{N2} + 0.0598\,P_{He} \qquad (5.2)$$

Unfortunately, at typically achievable pressures (≤ 10 bar), pressure broadening remains smaller than the rotational line spacing. To combat this feature, the ATF laser system uses an isotopic mixture of CO_2 molecules. Substitution of an ^{16}O atom with an ^{18}O breaks the symmetry of the molecule, doubling the density of rotational lines and thus increasing the gain bandwidth. At equilibrium, the ATF regenerative amplifier contains a mixture of the isotopologues ^{16}O:^{12}C:^{16}O, ^{16}O:^{12}C:^{18}O, and ^{18}O:^{12}C:^{18}O in a [1:2:1] ratio. In conjunction with pressure broadening, the gain profile broadens enough to support the amplification of a single pulse of 5 ps FWHM. Figure 5.1 illustrates the achievable pulse configurations under varying amplifier pressure and isotopic combinations.
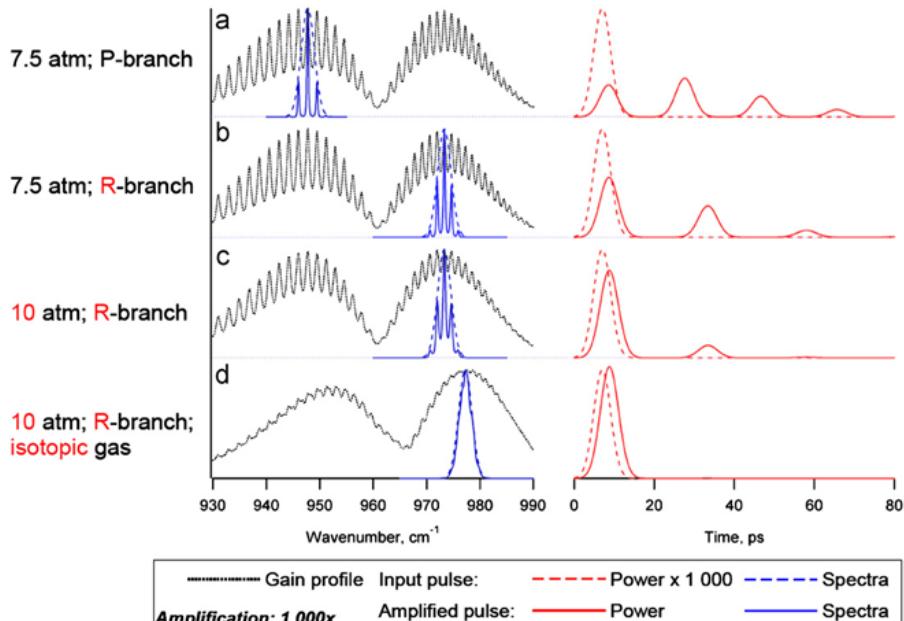
Figure 5.1: A comparison of achievable bandwidth and corresponding simulated pulse profiles for different CO_2 amplification schemes. Only using a combination of isotopic and pressure broadening can a single ps-scale pulse be achieved at experimental laser intensities [8] ©2011 Elsevier.

Pulse slicing

Figures 5.2 and 5.3 present a diagram of the ATF CO_2 laser system and the nominal pulse length and energy after each major component, respectively [109]. A hybrid transversely excited atmospheric (TEA) seed laser produces an initial 200 ns 20 mJ pulse at 10.3 μm. A Pockels cell is used to reduce the duration of this pulse to 10 ns before injection into a 3-bar UV pre-ionized preamplifier, which outputs 5 mJ of energy in a 10 ns pulse. In order to achieve amplification of a 5 ps pulse, this long input pulse must first be shortened; this is done through a two step pulse slicing technique.
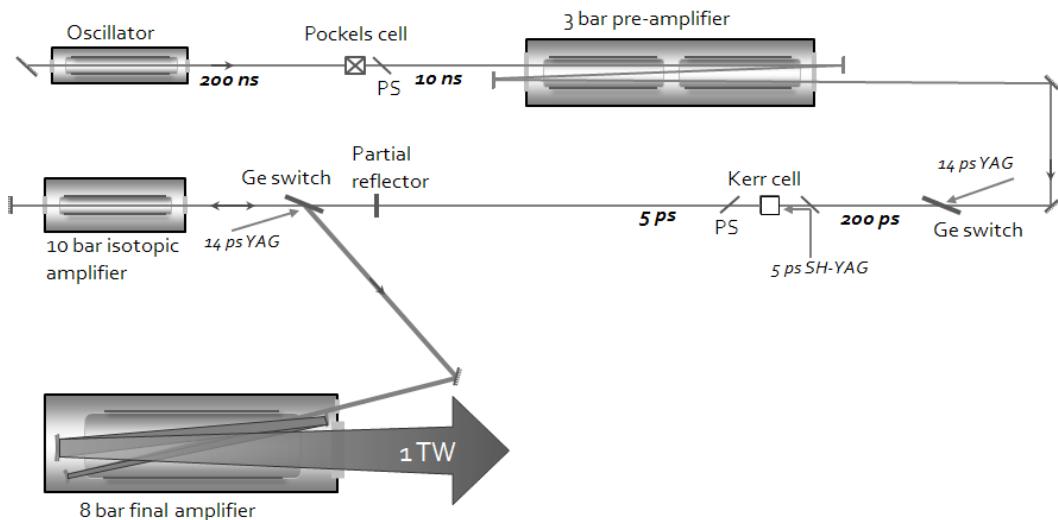
Figure 5.2: Schematic of the ATF TW-class CO_2 laser system. [9]
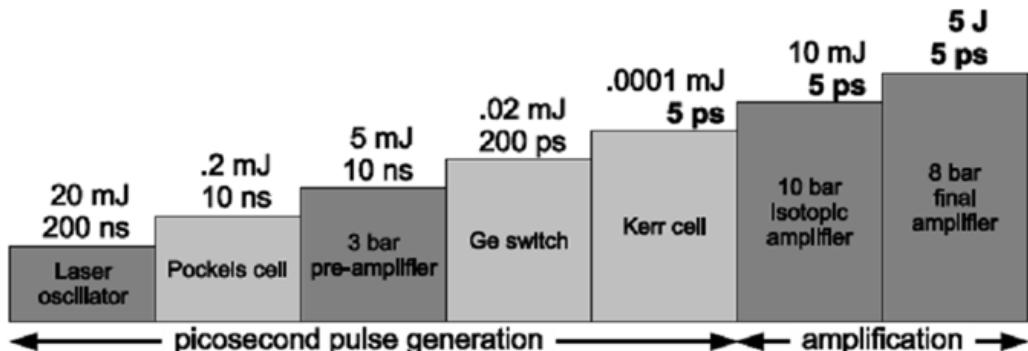
Figure 5.3: Chart of pulse properties during amplification in ATF CO_2 laser chain. [9]

The pulse is first shortened through the use of an optically triggered ger-

manium semiconductor switch. A $1\,\mu$m Nd:YAG pulse with 14-ps duration incident on the plate modulates the reflection and transmission properties of the germanium. The $1\,\mu$m light has photon energy greater than the bandgap of the semiconductor, and thus excites electrons into the conducting band. The result is a large increase in the free electron density, creating a effective electron-hole plasma on the surface of the plate. The surface plasma density depends on the fluence of the trigger beam; for high fluences, the density may exceed the critical density for $10\,\mu$m light, causing reflection of the subsequent CO_2 pulse, which would otherwise propagate through the semiconductor [110]. For $n \sim n_c$, reflectivity is around $\sim 36\%$, but at increased carrier densities, reflectivity of greater than 95% can be achieved [111]. The Germanium switch is placed at Brewster's angle so as to minimize background reflectivity from the p-polarized CO_2 pulse.

The plasma mirror effect persists beyond the control pulse over the period of the electron carrier diffusion time, which is of the order of 200 ps [109]. The result is that the use of a single switch produces a pulse with a very sharp rise time but very long tail. A second stage is required to further reduce pulse length. This may be accomplished using a second Ge switch with variable trigger pulse timings [112, 113]. However, this imposes steep requirements on the duration, timing, and photon fluence of the secondary pulse to achieve consistent pulse shapes and minimize losses.

An alternative approach is to gate the pulse using the optical Kerr effect [114]. Many materials exhibit an intensity dependent index of refraction, given by

$$n = n_0 + n_2 I \tag{5.3}$$

where I is the cycle-averaged pulse intensity given by [115]

$$I = \frac{n_0 c}{2\pi} |E(\omega)|^2 \tag{5.4}$$

The interaction of a short, intense short pulse with a medium featuring a high n_2 value induces a temporary birefringence. A synchronized, co-propagating witness pulse of wavelength λ will experience a phase shift roughly equal to the B integral, given by

$$\Delta\phi = \frac{2\pi}{\lambda} \int_L n_2 I \mathrm{d}z \tag{5.5}$$

By adjusting the length of the Kerr cell and/or the intensity of the probe pulse, the phase change of the witness beam can be controlled. Finally, a polarizer can be used to pick out only the portion of the pulse which overlapped

with the passage of the pump beam. This results in a sliced pulse of duration equal to that of the pump beam, limited to the response time of the nonlinear medium. The ATF uses a Kerr cell filled with liquid CS_2, which has an n_2 value of 3×10^{-14} cm^2/W and a relaxation time of 2 ps [116]. A 5 ps Nd:YAG pulse which is $\pi/4$ out of phase with the witness CO_2 beam pumps the Kerr cell, while the ~ 200 ps p-polarized CO_2 output of the germanium switch described above acts as the witness beam [109]. The final output is a 5 ps pulse with approximately 100 nJ of energy. Figure 5.4 illustrates the slicing process.
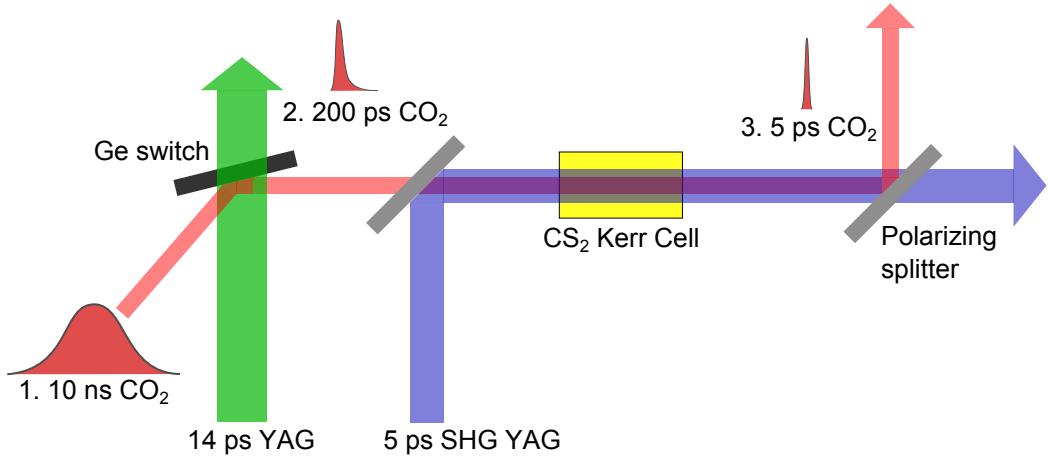
Figure 5.4: The ATF uses a two stage pulse slicing technique to generate a 5 ps FWHM CO_2 pulse. First, a germanium switch is activated by a 14 ps YAG, causing reflection of a ~ 200 ps p-polarized CO_2 pulse. Next, the pulse passes through a CS_2 filled Kerr cell pumped by a synchronized 5 ps frequency doubled YAG pulse. A polarization filter picks out the 5 ps, polarization rotated CO_2 pulse.

Final amplification

The 5 ps pulse is coupled into the 10-bar isotopic regenerative amplifier, with a volume of 80 cm^3, containing a mixture of CO_2, N_2, and He gases in a $0.5 : 0.5 : 9$ ratio. The amplifier is UV-preionized and transverse electric discharge pumped. The pulse makes 8-12 passes en route to amplification by a factor of 10^5 up to 10 mJ of energy. Extraction of the pulse is triggered via an additional Nd:YAG pulse and germanium plate at Brewster's angle. Amplification to the TW scale is achieved in a large aperture multi-pass 8 bar CO_2 amplifier. The 8 m^3 volume is filled with non-isotopic CO_2, along with N_2 and He in a $2 : 1 : 28$ ratio, and is x-ray pre-ionized and transversely discharge

Table 5.1: Summary of nominal output parameters for the ATF CO_2 laser system as used during the 2013 ion generation experimental campaign.

Laser Parameter	Units	Value
Laser Wavelength, λ	[μm]	10.3
Pulse Energy	[J]	5
Pulse Length, FWHM, τ	[ps]	5
Peak Power	[W]	9.3×10^{11}
Nominal Peak Intensity	[W/cm^2]	1.42×10^{16}
Peak Normalized Intensity, a_0		1.075
Beam radius ($1/e^2$), w_0	[μm]	65
Beam FWHM	[μm]	76.5
Measured Rayleigh Length, z_R	[mm]	0.400
Effective Beam Quality Parameter, M^2		1.6

pumped. Field broadening of the CO_2 gain structure minimizes pulse splitting, despite the lack of isotopically enriched gas.

The output pulse contains up to ~ 11 J of energy in 5 ps FWHM for a peak power of 2.2 TW. Pulse duration and sequence has been diagnosed through the use of a streak camera to a resolution of $1 - 2$ ps [9]. The pulse travels to the experimental chamber, where it is focused by an $f/3$ off-axis parabola down to a spot size of $w_0 = 65\,\mu$m (FWHM $= 76.5\,\mu$m). This corresponds to a maximum achievable intensity of 2×10^{16} W/cm^2, or equivalently a normalized intensity of $a_0 = 1.6$. The beam quality parameter, M^2 relates the divergence θ of the real beam to the ideal divergence θ_0 of a perfect Gaussian, according to [117]

$$\theta = M^2\theta_0. \tag{5.6}$$

The real beam will diverge more quickly than an ideal beam, thus leading to values of $M^2 > 1$. Measurement of the Rayleigh length and focal spot of the CO_2 beam produce a value of $M^2 = 1.6$. Table 5.1 displays the nominal laser output parameters as described above.

5.1.1 Controlled pre-pulse generation

The use of a Ge plate as an active output coupler for the regenerative amplifier has the consequence of enabling the production of pre-pulses prior to the output of the main pulse. While low intensity pulses pass through the semiconductor, high intensity pulses may trigger surface plasma formation and subsequently result in partial reflection of the pulse. This self-induced

reflectance couples a small amount of light out of the regenerative amplifier during each pass above a minimum intensity, resulting in a train of low intensity 5 ps pulses spaced at the round trip time of the amplifier cavity (25 ns). As the reflection is intensity dependent, each subsequent pulse is exponentially larger in amplitude, so long as the amplifier remains in a linear amplification regime. However, if the amplifier nears saturation, adjacent pulses may have similar amplitudes, resulting in much lower contrast between the full pulse energy and the passively reflected pulse. At peak amplitudes from the regenerative amplifier, self-induced reflectance may transmit pulses with upwards of 10% of the laser energy. By varying the high voltage discharge within the amplifier, the level of saturation can be adjusted and the resultant pre-pulse energy can be controlled independently of the main pulse energy. This allows for the creation of a 2 pulse sequence with variable pre-pulse energy between a value much less than 1% and 10% of the main pulse energy. Figure 5.5 illustrates this process.

5.2 Experimental chamber and laser plasma interaction diagnostics

The target chamber contains the final focusing of the CO_2 beam, in addition to the gas jet, target mount, and Thomson parabola spectrometer for ion diagnostics. The chamber is pumped to a vacuum of better than 10^{-4} Torr to minimize laser filamentation and nonlinear focussing effects. A basic illustration of the chamber is shown in Figure 5.6.

Gas jet target

High pressure gas jets offer a flexible option for an array of laser plasma experiments. They are popular in an array of underdense, electron acceleration experiments, in which channelled jets can provide densities on the order of 10^{17} cm^{-3} over a cm-scale channel [118, 119]. Alternatively, with extremely high backing pressures they can achieve near critical densities of 10^{21} cm^{-3} even for 1 μm light [120]. Jets can be operated at Hz repetition rates, limited mostly by the vacuum requirements of the chamber. They also provide a high purity target for a variety of elements, free from the adsorption and contamination seen in foil targets.

The ATF uses a cylindrical gas jet with a 1 mm nozzle [121]. The jet produces a plume of gas that is cylindrically symmetric about the nozzle axis, and produces a longitudinal density profile with triangular shape at a range of backing pressures. Although the backing pressure was varied, all of the data
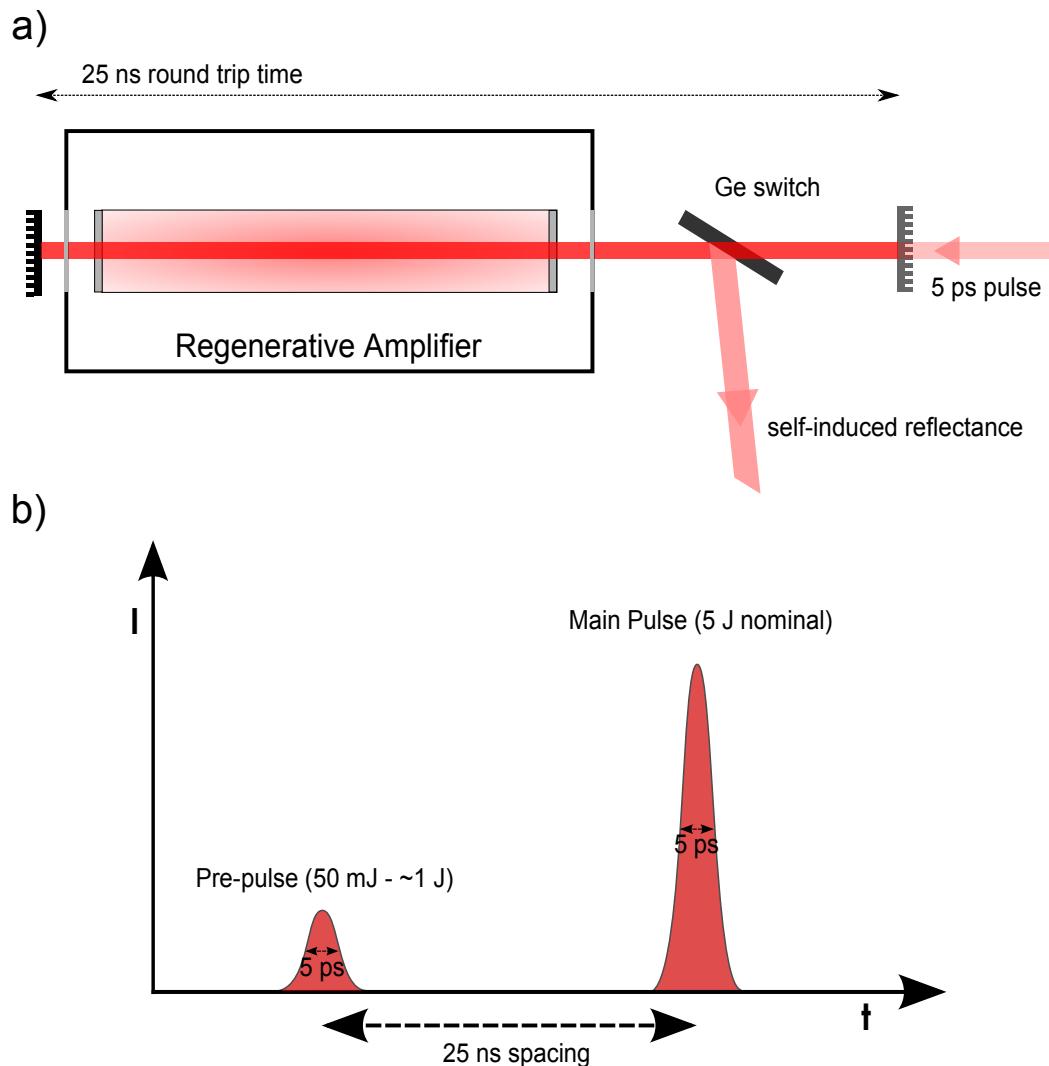
Figure 5.5: (a) Generation of an energetic pre-pulse 25 ns prior to the main pulse is achieved through self-induced reflectance from the germanium switch used to extract the main-pulse. (b) The resulting "train" of 5-ps pulses is illustrated.
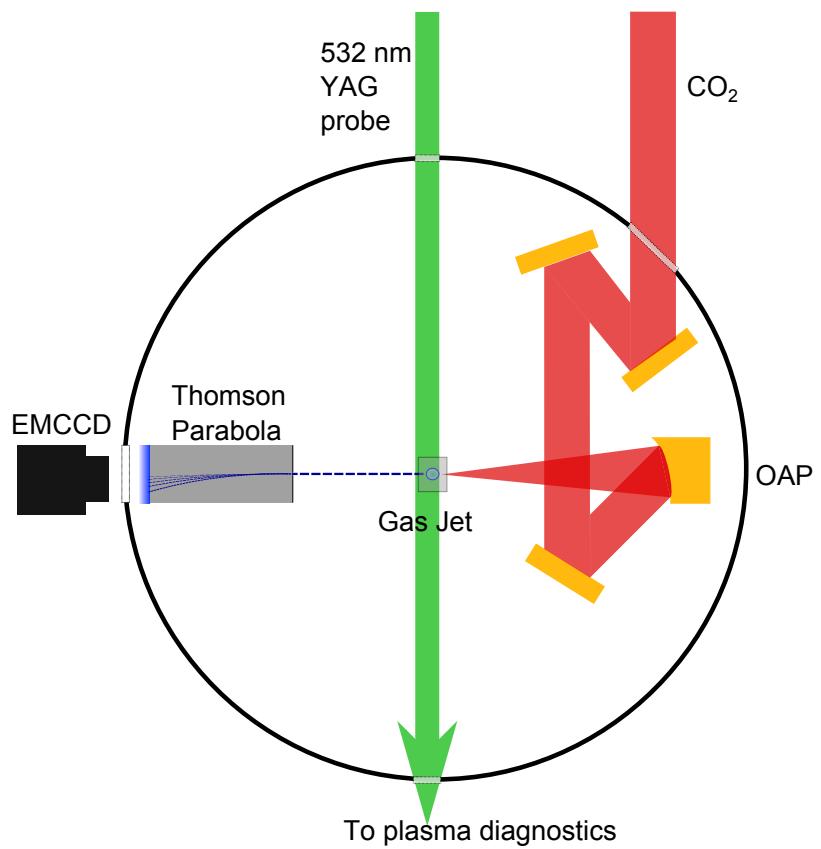
Figure 5.6: Overhead illustration of the ATF ion generation experimental chamber.
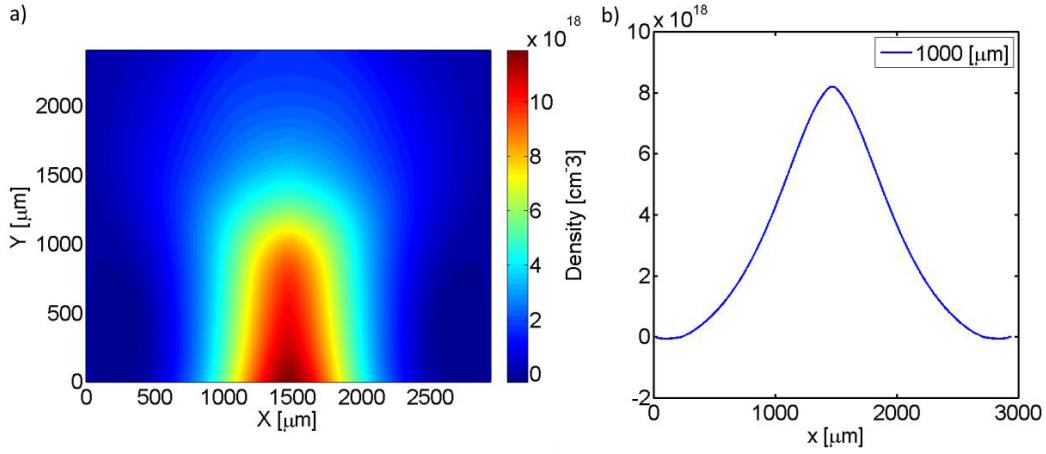
Figure 5.7: (a) A 2D profile of the cylindrically symmetric helium neutral density from a 1 mm nozzle with 12 bar backing pressure. (b) A lineout along the laser axis reveals a near linear ramp of approximately 800 μm in length to a peak neutral density of 0.8×10^{18} cm^{-3}, equivalent to a fully ionized plasma density of about $1.6n_c$.

presented in this thesis was taken with 12 bar backing pressure. Figure 5.7 shows a typical 2D slice longitudinal helium neutral density profile along the laser axis, approximately 1 mm above the jet. The profile is triangular, ramping up over approximately 800 μm to a peak neutral density of 0.8×10^{18} cm^{-3}, corresponding to a fully ionized plasma density of $\sim 1.6n_c$. The gas profile has been characterized at three different heights above the jet, with varying backing pressure, as shown in Figure 5.8.

The gas jet sits on a 3-axis mechanical stage to allow for variation of laser focal position within the gas target. A Newport ESP301 motion controller is used to adjust and maintain the jet position to within 10 μm.

Optical probing and plasma characterization

Imaging of the gas or plasma density requires optical probing with a synchronized laser pulse. A short pulse Nd:YAG laser system is used for this purpose at the ATF. As the YAG pulse is used for slicing the CO$_2$, the relative timing of the two pulses is known, and can be used as a basis for achieving synchronization. For diagnostic ease, the original 14 ps 1068 nm pulse is frequency doubled to 532 nm in a KDP crystal via second harmonic generation. The resulting 10 ps pulse is split into channels of $\hat{s}$ and $\hat{p}$ polarizations. One channel enters an adjustable delay stage allowing up to 2 ns delays with < 1 ps minimum increments. The probe beams are then sent into the chamber, where they intersect the target perpendicular to the laser axis. A basic diagram is
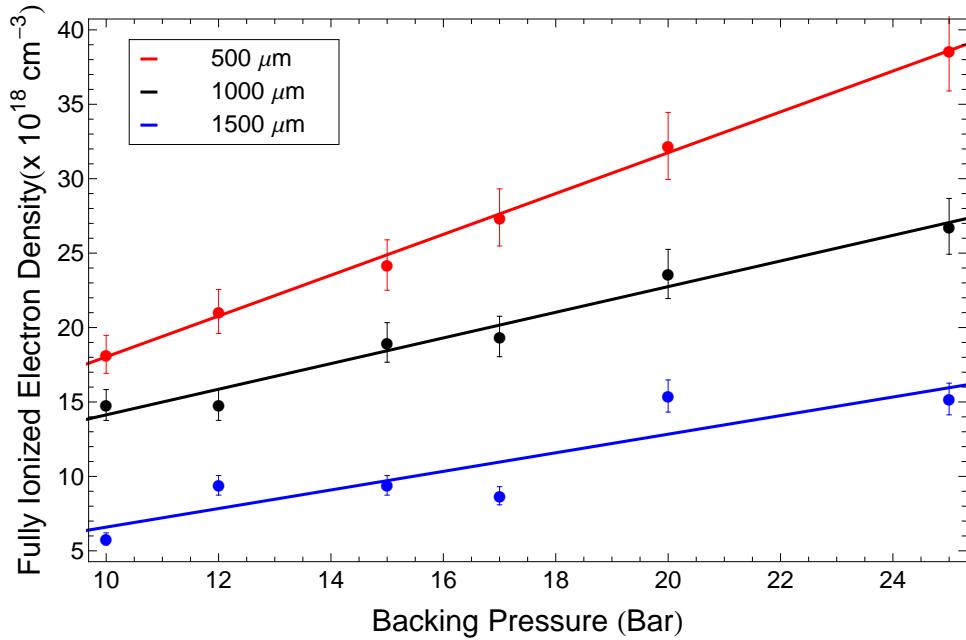
Figure 5.8: A comparative plot from scans of the helium neutral density for varying backing pressure demonstrates a fairly linear increase in density with backing pressure for the same distance from the nozzle. Errors are estimated from resolution of the camera along with shot to shot variation in gas pressure.

shown in Figure 5.9.
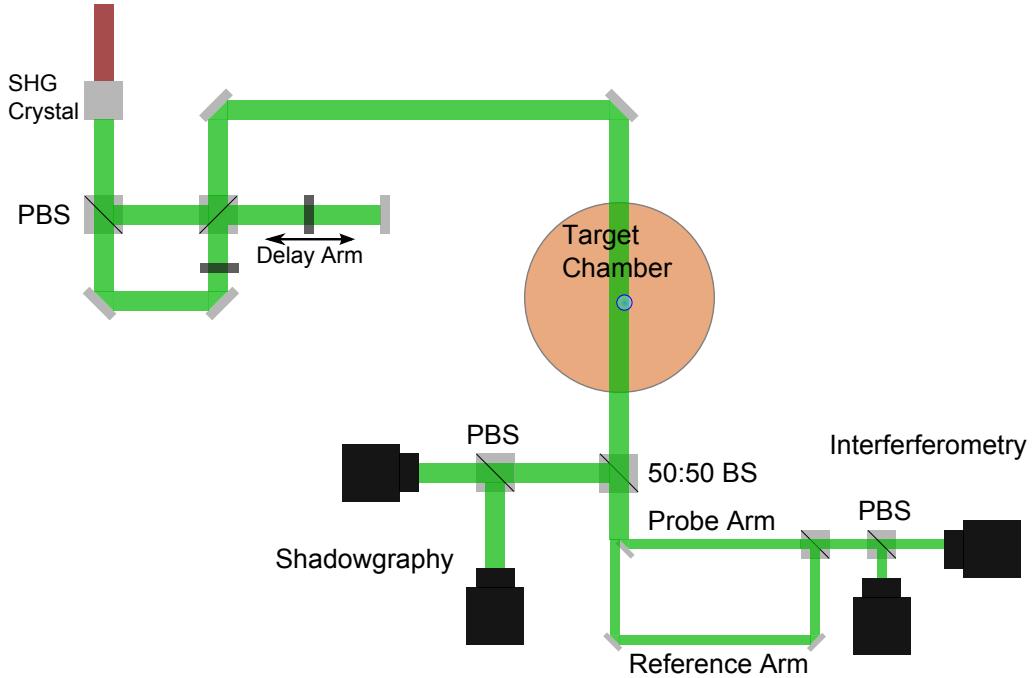
Figure 5.9: Schematic of the 532 nm Nd:YAG optical probing line at the BNL ATF.

Recall from Chapter 4 that the critical density of a plasma scales as $1/\lambda^2$. Because the gas jet is operated near the critical density of $10\,\mu$m light, it remains transparent to light in the visible spectrum. Light passing through the plasma experiences a phase advance due to the index of refraction $n_r = c/v_{ph}$. Careful measurement of this phase advance allows for a reconstruction of the plasma profile and serves as the basis for in situ imaging of the plasma before and after the laser plasma interaction [122].

Consider a ray of light with wavelength λ travelling in the x-direction. If that light passes through a material with index of refraction $n_r(z)$, along the path s, it accumulates a phase shift $\Delta\phi$ of

$$\Delta\phi = \frac{2\pi}{\lambda} \int_s n_r(x)\mathrm{d}x \tag{5.7}$$

The index of refraction of the plasma is given by Equation 4.61. This can be expressed in terms of the plasma density n_e and probe beam critical density, n_p via Equation 4.7. For $n_e << n_p$, we have

$$n_r = \left(1 - \frac{n_e}{n_p}\right)^{1/2} \approx 1 - \frac{1}{2}\frac{n_e}{n_p} \tag{5.8}$$

The 532 nm YAG probe has

$$n_c \sim 2 \times 10^{21}\,\text{cm}^{-3} \gg n_c(10\,\mu\text{m}) \tag{5.9}$$

so the expression holds. We then can write the phase shift as a function of density, via

$$\Delta\phi = \frac{2\pi}{\lambda} \int_s 1 - \frac{1}{2}\frac{n_e}{n_p}\mathrm{d}x \tag{5.10}$$

Unfortunately, this equation provides no information regarding the specific shape of the density profile, as the total phase shift only gives information about the cumulative amount of material that the light has passed through. However, if we make assumptions about the shape of the density, then it can be calculated numerically. In particular, we assume that the density profile in question is cylindrical about the drive laser axis ($\hat{z}$), such that $n_e = n_e(r, z)$. Then we can express the above equation through the polar transformation to the radial coordinate r, which gives [123]

$$\Delta\phi = \frac{2\pi}{\lambda} \int_y 1 - \frac{1}{2}\frac{n_e(r,z)}{n_p}\frac{r}{\sqrt{r^2 - y^2}}\mathrm{d}r \tag{5.11}$$

using the polar coordinate system shown in Figure 5.10.

Equation 5.11 provides a direct relationship between the path-dependent phase accrual of the laser and the density of the material generating the phase advance. Thus, by measuring the phase advance of a laser passing through the plasma, relative to a reference ray, the radial density profile of the plasma can be calculated.

To achieve this, a Mach-Zender interferometer was constructed in which a reference beam passes through one leg, and the probe beam down the other. A diagram of the interferometer is shown in Figure 5.9. The two beams are crossed at an angle to introduce a baseline fringe pattern in the image. The phase shift of the probe beam corresponds to additional fringe shifts from the neutral pattern. The phase change just due to the plasma follows from the difference in phase shift of each path, resulting in

$$\Delta\phi = \frac{2\pi}{\lambda} \int_y \frac{1}{2}\frac{n_e(r,z)}{n_p}\frac{r}{\sqrt{r^2 - y^2}}\mathrm{d}r \tag{5.12}$$

104
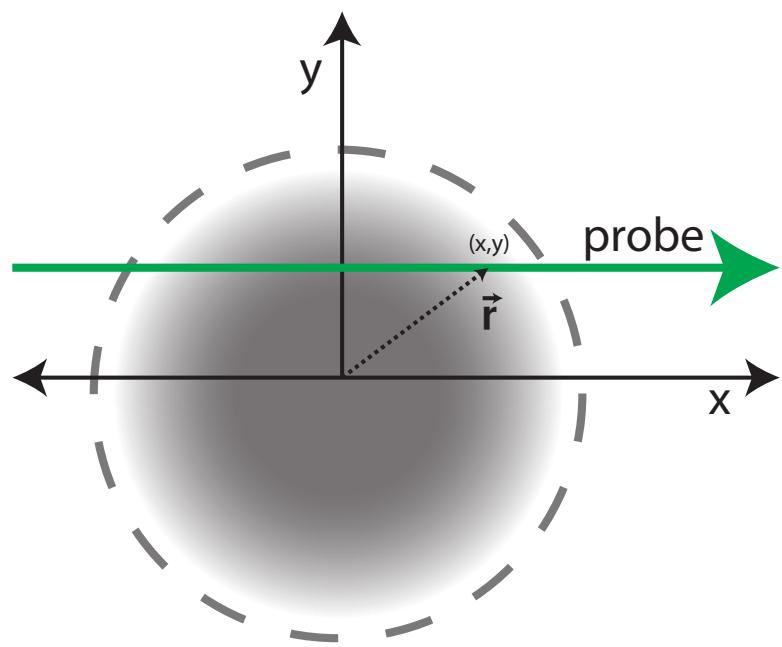
Figure 5.10: Polar coordinate system for radially symmetric density being probed transversely by shorter wavelength probe pulse.

since the index of refraction of air is $n_r = 1$. Equation 5.12 is in the form of an Abel transform of a cylindrically symmetric function,

$$F(y) = 2 \int_y^\infty \frac{f(r)r}{\sqrt{r^2 - y^2}} dr \qquad (5.13)$$

Performing the inverse of the Abel transform gives

$$f(r) = -\frac{1}{\pi} \int_r^\infty \frac{dF(y)}{dy} \frac{dx}{\sqrt{r^2 - y^2}} \qquad (5.14)$$

Thus, to recover the radial density function $n_e(r, z)$, an Abel inversion can be performed to give

$$n_e(r, z) = \frac{\lambda n_p}{\pi^2} \int_r^\infty \frac{d\phi(y, z)}{dy} \frac{dx}{\sqrt{r^2 - y^2}} \qquad (5.15)$$

Given an appropriate two-dimensional phase map, the value of $\frac{d\phi(y,z)}{dy}$ can be computed, and the resulting radial density map can be calculated. For neutral density gases, the density is related to the index of refraction via the Clausius-Mossotti relation [124], $n_r(gas) \sim 1 + Kn/n_0$, where n_0 is the gas density at standard temperature and pressure (STP). Substitution of this relation into the model described above allows for similar calculation of the neutral gas density for a cylindrically symmetric gas jet via Abel inversion.

The interferogram analysis is limited by the critical assumption inherent to the Abel transform, that the density be cylindrically symmetric about a particular axis. As a result, extracted density profiles effectively average over asymmetric variations in phase along the probe beam axis. Moreover, systematic density gradients across the central axis may be ignored if symmetry is assumed and a single hemisphere is chosen to produce the radial phase function $\phi(y, z)$. This problem can be alleviated by splitting the phase map into separate hemispheres about the probe axis, and performing two individual transforms. The resulting density profiles must necessarily be continuous along the laser axis, and thus interpolations can be performed at small r values to achieve a smooth result. This technique has been used in particular to combat the monotonic decrease in neutral gas density with increasing distance from the nozzle, producing a naturally asymmetric plasma density above and below the laser axis.

Another shortcoming of the Abel inversion is its mathematical impotency for small radii. Equation 5.15 diverges for small r, producing a non-numerical result for density along the laser axis. While described technique of interpola-
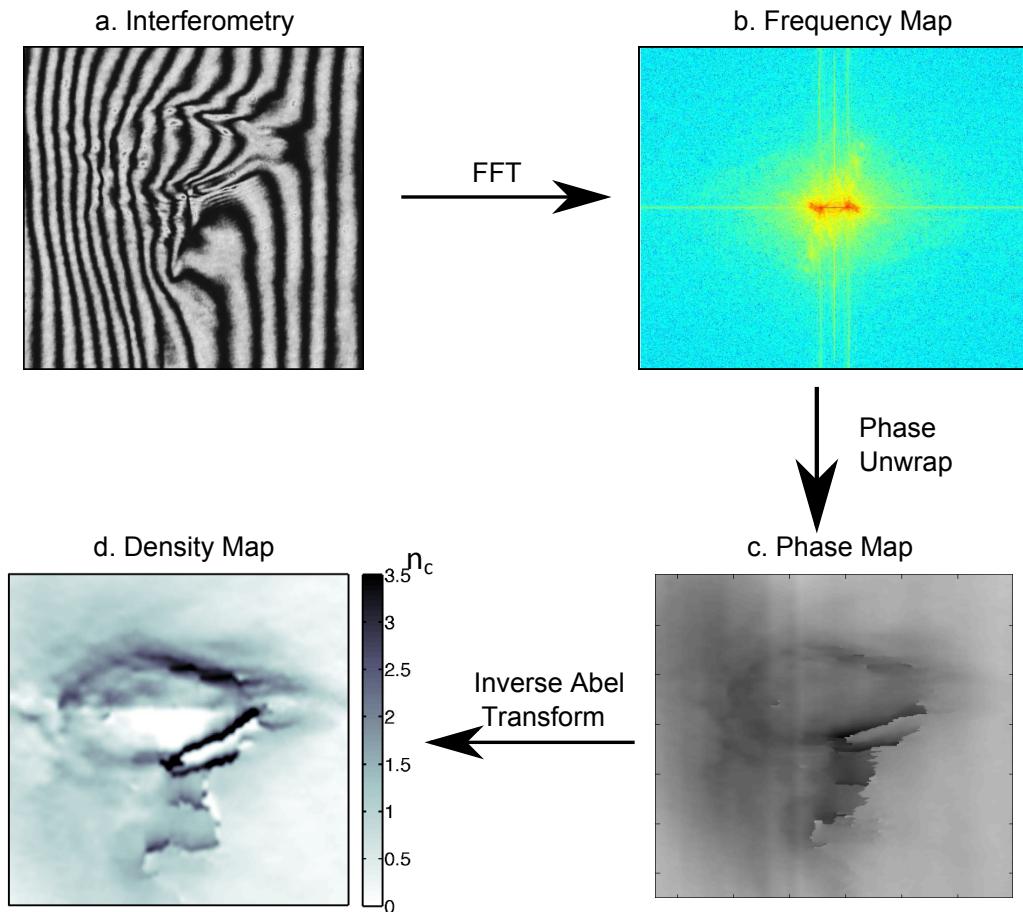
Figure 5.11: An illustration of the interferometry analysis process. The gas plume enters from below and the CO_2 laser enters from the left. Beginning with (a), the original interferogram (plus an additional reference image), a 2D frequency map of the image (b) is made via discrete Fourier transform. A subset of the frequency components is chosen, and a phase extraction is run to unwrap the phase (c) of the interferogram. Lastly, an inverse Abel transformation is performed on the phase map, resulting in the final extracted density profile (d).

tion at small r values does provide a continuous density map, it may obscure sharp on-axis density features, which are of particular importance to plasma acceleration mechanisms. As a check on the verisimilitude of the extracted density profiles, a series of simulated interferograms were constructed to properly match the observed phase shift.

To simulate an interferogram, a 3D plasma density is designed to best approximate the extracted profile. This can be done through mathematical approximation of 3D paraboloid shapes to fit the blast wave structure of the laser produced plasma. Alternatively, the profile extracted from data can be interpolated, and modified as desired to obtain the best fit phase profile. This approach has been taken for analysis of relevant density profiles, and the generated profiles are shown alongside results in the next chapter.

Shadowgraphy

In addition to interferometry, a fraction of the probe beam energy is reserved for shadowgraphy. The light which has passed through the dense plasma is directly imaged, and corresponding patterns of low and high intensity are formed as a result of the angular deviation of the light from its normal path due to the index of refraction of the plasma. The resulting intensity profile is a function of the second derivative of the index of refraction of the plasma, and as a result the shadowgraph provides a simple visual confirmation of very sharp changes in the density profile of the plasma [125]. Figure 5.12 shows an example shadowgraph image along with the corresponding interferogram.
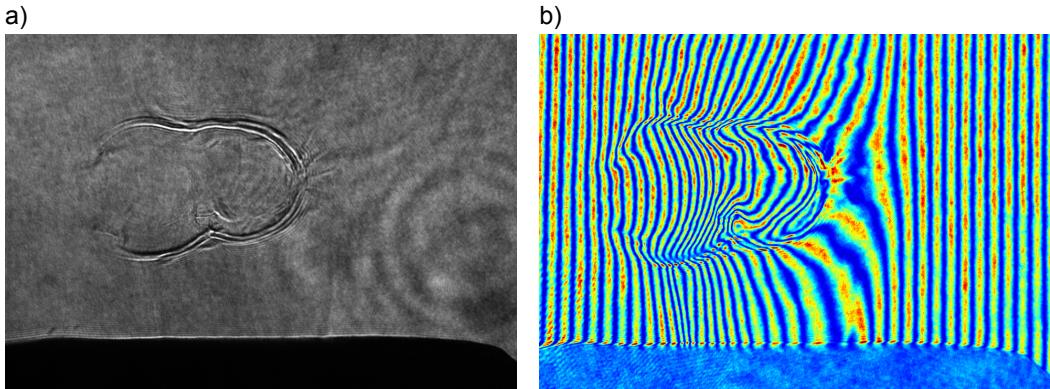
Figure 5.12: (a) An example shadowgraph is shown, taken 300 ps after the interaction of the CO_2 beam with a helium gas jet, alongside (b) the corresponding interferogram.

5.3 Ion beam diagnostics

Ion beams accelerated from gas jet targets may have very small divergences, and as such uncovering the spectral information of the ion beam requires the use of additional filters. One such method is to use a stacked array of ion sensitive detectors. As the incoming ions will lose the most energy per unit length within the Bragg Peak at the end of their range, each detector in the array preferentially detects ions within a range of energy corresponding to its depth [126].

An alternative approach is to disperse the accelerated ions according to their momentum and charge-to-mass ratio. This is the principle of the Thomson parabola spectrometer, first used in 1913 to identify isotopes of neon. The spectrometer uses a combination of electric and magnetic fields to bend particles along specified axes. The resulting deflection can easily be calculated using the equation of motion of a particle in an electromagnetic field. Assume an ion beam propagates along the z direction. An electric field, $\boldsymbol{E} = E\hat{\mathbf{x}}$, is generated over a distance L_E, while a magnetic field, $\boldsymbol{B} = B\hat{\mathbf{x}}$, is generated over a distance L_B. If a detection screen is placed a distance d_E from the endpoint of the electric field segment, and d_B from the endpoint of the magnetic field segment, the resulting drifts are given by,

$$
\Delta E = \frac{qEL_E}{mv^2}\left(1/2L_E + d_E\right)\hat{x}
$$
$$
\Delta B = \frac{qBL_B}{mv}\left(1/2L_B + d_B\right)\hat{y} \tag{5.16}
$$

The trace along the detector forms the shape of a parabola, given by

$$
\Delta y = \frac{EmC_E}{q(BC_B)^2}\Delta x^2 \tag{5.17}
$$
$$
C_E = L_E(L_E/2 + d_E) \tag{5.18}
$$
$$
C_B = L_B(L_B/2 + d_B \tag{5.19}
$$

The shape of the parabola depends on the charge-to-mass ratio q/m of the ion species, as well as the specific parameters of the spectrometer components. Permanent dipole magnets with fixed size and channel width were used in the experiment. Measurements were taken with a Hall probe to map the three field components. Particle paths can then be simulated for varying drift distances and alignments. Figure 5.13 shows paths calculated for protons of varying kinetic energy.
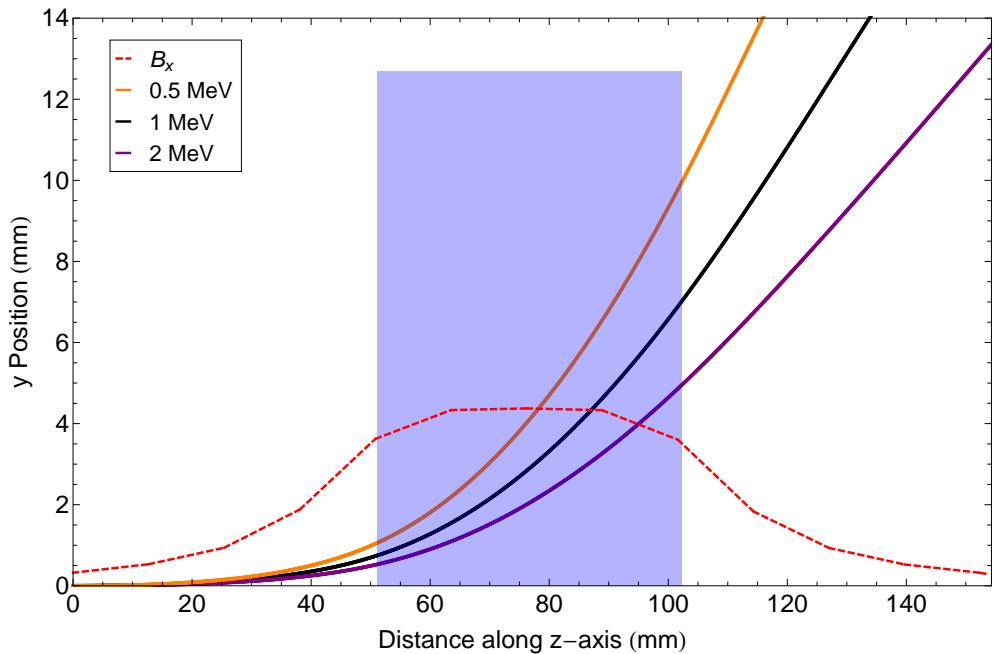
Figure 5.13: A plot of the effective magnetic field (dashed red line) from one of the dipole magnets used for the experiment, extrapolated from Hall probe measurements. Proton paths are shown for varying energies, assuming the particles enter along the central axis of the magnet.

Light output from the scintillator is imaged with a 16-bit Princeton ProEM+ CCD camera. The camera features a 1024x1024 CCD with 13 um pixel size, and optional electron-multiplication up to 1000x. Figure 5.14 shows an example trace of helium ions accelerated in experiment. The red line trace corresponds to singly ionized helium, while the green trace corresponds to the predicted path for doubly ionized helium, and the black line to the projected path for protons. A strong neutral signal is also observed (top left), the product of transmitted laser photons, photon emission from the plasma, and neutral particles resulting from recombination at the rear of the target. X-ray emission contributes to signal noise, resulting in large bright spots across the scintillator. For large signals, the corresponding signal spills into neighboring pixels as the CCD register transfers charge horizontally. This produces the horizontal smearing seen in Figure 5.14. X-ray production and CCD noise was observed consistently during the helium run, and is the primary contributor to background signal.
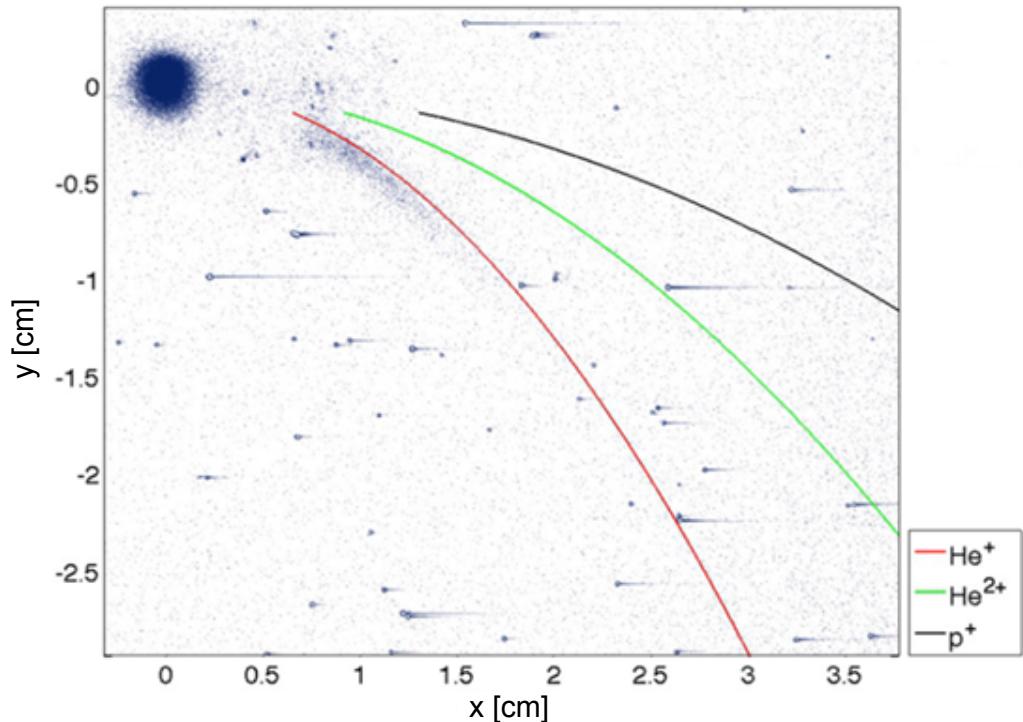
Figure 5.14: Example Helium ion spectra as viewed through the Thomson parabola, with projected particle traces plotted.

The scintillator diagnostic combined with a Thomson parabola spectrometer provides a nearly real-time single shot measurement of ion number and

energy spectrum. However, the choice of which scintillating material to use, and the precise calibration of scintillator yield and subsequent ion number calculations required considerable study. Ultimately, the plastic scintillator BC-408 was employed for use in ion acceleration experiments. The testing and analysis required to make that decision are the subject of the next chapter.

Chapter 6

Diagnostic Tools for Imaging Laser Accelerated Ions

Laser driven acceleration using a gas jet source offers the possibility of high repetition rate ion bunch acceleration. Nonetheless, the current understanding and execution of this technique is imprecise and unreliable. Improvements in technology, specifically with regards to target manipulation, pulse length, and laser intensity may provide orders of magnitude advances in power and efficiency, while continued experimental testing should result in a better control over individual interactions. The evolving nature of this field necessitates flexible diagnostic techniques suited to measure a myriad of experimental outputs.

To this end, we have adopted the strategy of imaging accelerated ion bunches using scintillating materials rather than using resins, films, or image plates to track impinging ions. Beam size and particle number can both be obtained through the image alone. With the use of filters or an electromagnetic spectrometer, measurements of emittance, mean energy, and energy spread can also be made. High resolution imaging can be performed at a high repetition rate, upwards of 100 Hz depending upon the camera and image properties. In order for imaging to be effective, substantial calibrations must be performed with the scintillating material chosen. Most publications using these methods of imaging make use of particle species, energy, or fluences that are insufficient for our needs [10, 127–129]. Figure 6.1 shows the results of comparison tests done by Forck et al. [10], illustrating the significant differences in scintillation yield for the same incident beam. With this in mind, a series of tests were performed on a variety of candidate materials to obtain meaningful calibrations and draw conclusions for use in acceleration experiments.

The goal of these tests was to find a material that produces an adequate and consistent amount of light under expected laser ion bunch conditions ($10^4 - 10^6$ protons with energies $1 - 20$ MeV). The material should have sufficient resolution to determine beam properties such as bunch size and correspondingly the emittance. Lastly, the material should be relatively inexpensive and/or robust for extended use when exposed to high radiation flux at a repetition rate near 1 Hz.

This chapter provides a discussion of the classes of single shot imaging diagnostics. The candidate scintillators for experiment at the ATF are then introduced. The results of a series of yield calibration tests are presented, along with a detailed analysis of secondary effects such as afterglow, scattering effects, and radiation damage. The discussion follows many of the key details as presented in Cook et al. [130].
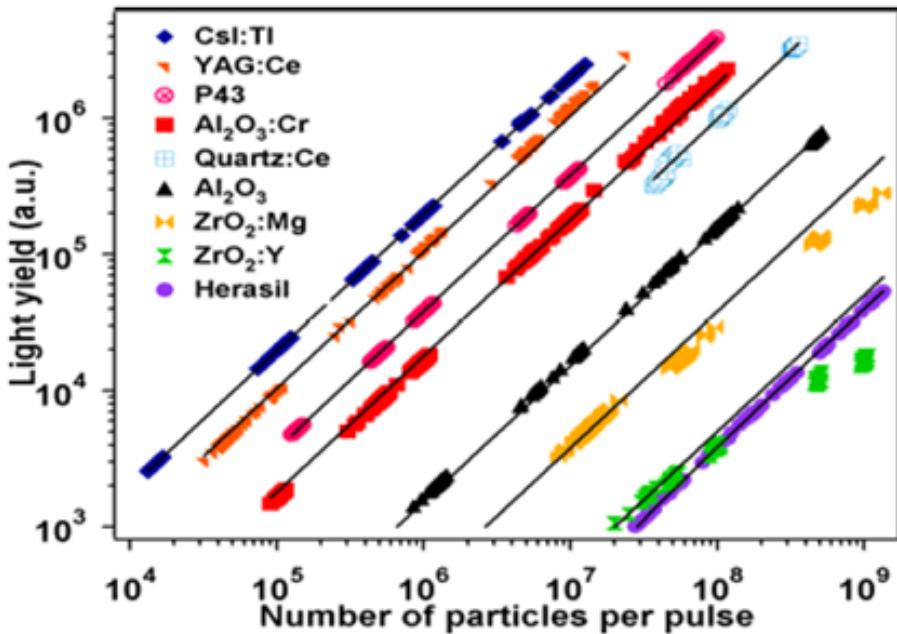
Figure 6.1: Results of a study by Forck et al. [10] demonstrate the orders of magnitude variation in scintillation response for the same beam conditions.

6.1 Film and polymer detectors

Allyl diglycol carbonate, known as CR-39, is a thermoset plastic polymer that is used for the detection of ionizing radiation. When CR-39 is damaged, it becomes more susceptible to chemical dissolving or "etching" in the presence of bases. The polymer is exposed to a solution of NaOH and water. The NaOH slowly erodes the bulk surface at the "bulk etching rate" v_b, but more quickly dissolves the damaged plastics at the "track etching rate" v_t. If the ratio of the two rates (v_t/v_b) is high enough, then the polymer may be etched in such a way that the resulting CR-39 displays conical pits corresponding to the impact of ionizing radiation [131]. The track response has been found to be a function of restricted energy loss (REL). REL is the fraction of energy loss that produces electrons with $E < E_{max}$. E_{max} is target specific; for CR-39 $E_{max} \sim 200$ KeV. As particle velocity decreases, REL increases to unity as the limit for energy transfer in any given collision becomes less than E_{max} [132]. Thus, as particles slow down they become more effective at forming etch-able tracks in CR-39. As a result, high Z ions are more easily etched, as are low energy beams.

Once the CR-39 has been etched, the pits can be viewed using a microscope and analyzed using software. Position sensitive comparisons between different

layers are made in order to find coincidences that represent true particle interactions, and to sort of out background radiation. In this way, the number of incoming particles may be determined. By stacking thin layers of CR-39, energy resolution comparable to the energy loss per unit length for a given thickness can also be obtained. Thus, for a given particle bunch, the total number of particles and the respective energy spread may be determined to better than 1% accuracy [131]. Figure 6.2 shows an array of CR-39 images after proton beam incidence and etching.
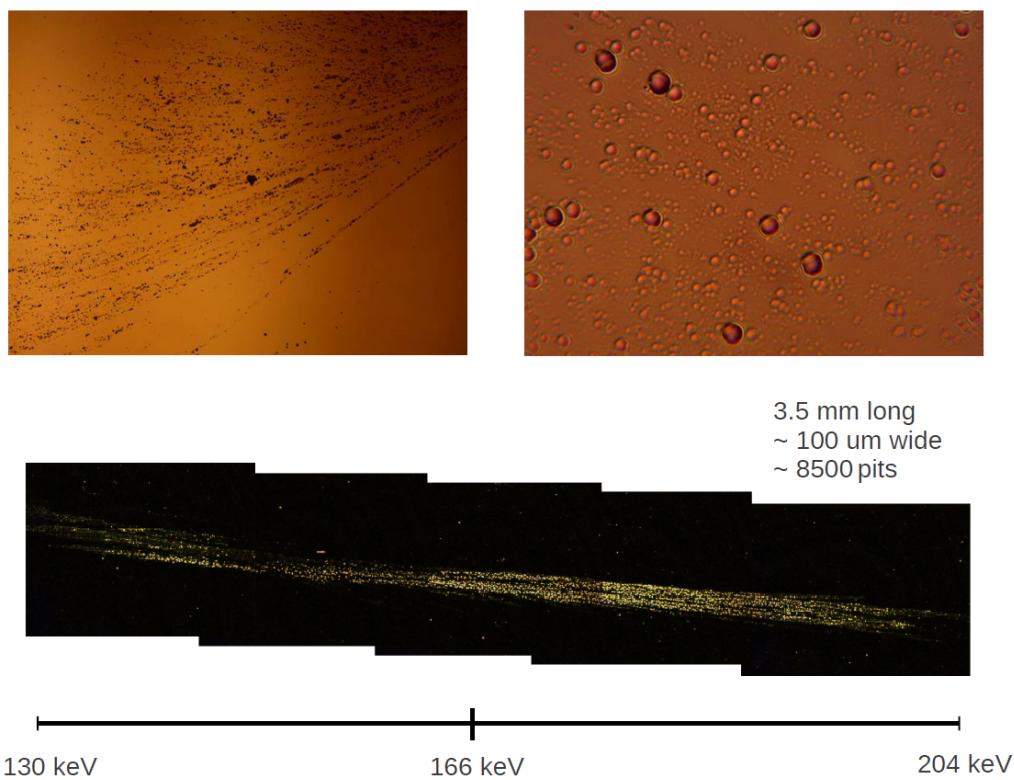
Figure 6.2: Two different examples of CR-39 images are shown. Top, an ion track (left) with pits visible under magnification (right). Pit diameters reach 5 μm after 20 minutes of etching with 6 molar NaOH at 80°C. Bottom, a track made by ions propagating through a magnetic spectrometer prior to incidence with the CR-39.

The primary disadvantage of such a technique is the time needed for etching and visualization. Microscopes with calibrated counting software must be used to identify, measure, and count the pits for each layer of CR-39. For higher energy particles, many layers comprise a single stack as is needed to slow or stop incoming radiation. Etching rates, and thus processing time, depends on

116

NaOH concentration as well as temperature, but times of a few hours or more are common [131]. In addition to CR-39, radiochromic film remains a popular one-time-use diagnostic which allows for individual particle identification as well as energy analysis, but requires considerable time to develop the film and compile measurements [133].

Fujifilm image plates, composed of layers of photo-stimulable phosphors, are also popular radiation detection diagnostics. Incident radiation excites and traps electrons in the band gap of the phosphor layer. Trapped electrons can be "imaged" by further stimulating them to the conduction band of the layer; the corresponding decay emits light which can be readout by calibrated scanners. This process can be repeated assuming trapped charges do not escape, a major advantage for image plates. However, there is a small probability for spontaneous recombination, yielding a characteristic decline in signal over time. Thus image plates produce time-sensitive photographs. Moreover, highly specific scanner systems are required to read an image plate and reproduce a useful image. Depending on the specific plate and scanner system parameters (resolution, scale, dimensions), read-out can take anywhere from a few minutes to more than an hour per square foot of image plate [133].

6.2   Micro-channel plate detectors

A micro-channel plate detector (MCP) acts as an analog signal enhancement device and is sensitive to a wide range of radiation. An MCP consists of millimeter thick glass plate with an array of channels running through the glass perpendicular to the face of the plate. The channels are small in diameter compared to the length of the plate (about 2–25 μm), and they are coated with a high resistance material, the standard for which is nickel-chrome alloy [134]. The mechanism of signal enhancement is secondary electron emission, akin to that in photomultipliers. When radiation strikes the face of the plate, some of it enters the micro-channels and interacts with the thin metallic coating. If the radiation has enough energy, the metal will emit secondary electrons normal to the channel surface.

A voltage is applied across the length of the plate, ensuring that these electrons strike the channel wall opposite of their release point with enough energy to trigger an avalanche effect in the channel. The buildup of secondary electrons is limited by the geometry of the plates as well as the applied voltage. When enough electrons have been freed, they produce a space-charge driven field strong enough to combat the applied voltage. New secondaries do not gain enough energy to further produce and saturation is reached [134].

MCPs may be constructed in a variety of shapes and configurations, but

for a standard single plate detector, signal gains of 10^4 electrons per input particle or photon can be achieved at a standard operating voltage of 1000 V to 1400 V, depending upon the geometry of the plate. Gain is a function of operating voltage, the plate length-channel diameter ratio α, channel bias angle ϕ, as well as the open-air ratio of the MCP, as determined by the channel-to-channel spacing and channel diameter. Assuming a standard secondary electron yield of δ per incident electron, the gain through an entire channel is $G = \delta^n$, where n is the number of times a secondary electron collide with the walls with enough energy to free δ additional secondaries. The expression for secondary yield from a thinly layered material with a constant stopping power is given by

$$\delta \sim \frac{P}{E} \frac{\lambda}{\cos \theta} \frac{dE}{dx} \tag{6.1}$$

where P is the probability of electron escape from the surface, E is the energy needed to produce such an electron, λ is the characteristic length of electron diffusion in the metal, and θ is the angle of incidence of the electron with respect to the normal of the channel surface [135]. Thus we see that the channel coating largely determines the secondary yield, though the angle of incidence of the beam has some effect. The quantity n is usually written as $n = g\alpha$, where g is a parameter reflecting the kinematics of produced secondaries (energy and direction with respect to the channel wall) as well as the applied voltage V across the MCP [135].
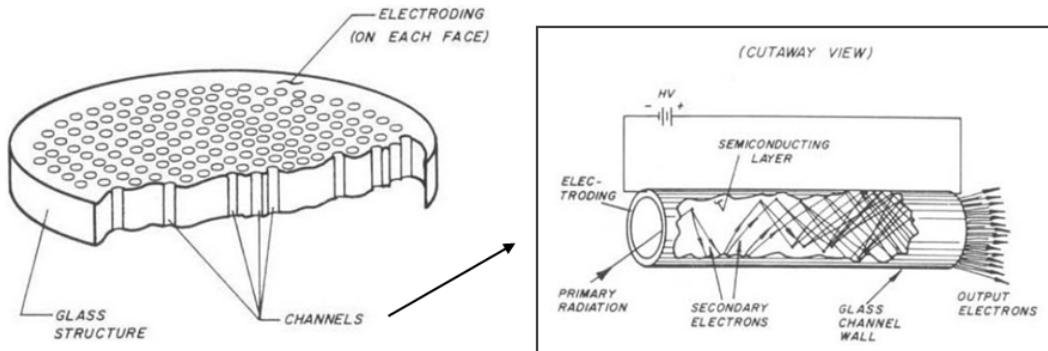
Figure 6.3: A cutaway of the basic single stage MCP with the channel structure diagram.

MCPs must be operated with very specific conditions or else damage to the channel coatings may occur. During operation, vacuums of 10^{-6} Torr are commonly required, and voltages must be carefully increased to the operating level in 100 V increments or else voltage breakdown, channel deterioration, or

arcing are possible. Storage of an MCP requires an oil-free vacuum of 10^{-4} Torr or better [136]. Improper storage may similarly damage individual channels or reduce performance and lifetime. The lifetime of an MCP is primarily a function of the charge extracted through use. Most MCPs will maintain more than 50% of nominal gain for up to 50C of extracted charge. This limit varies considerably depending upon the channel size, spacing, and length. Detection of beams with high instantaneous current may also pose a problem for a high gain MCP. If enough incoming particles interact with a given channel, saturation is reached, resulting in a nonlinear response with particle flux.

6.3 Scintillator candidates

We investigated a variety of scintillating materials with the right combination of properties as an alternative to films. Three scintillators were chosen for testing with proton beams as they have a combination of properties which make them preferred for detecting low energy protons. All scintillators operate from the same basic principle: incoming light is absorbed by the material, prompting a transition to some excited molecular or atomic energy level. The state quickly relaxes back to ground state through a combination of photon emission (fluorescence, phosphorescence) and non-radiative processes. It is this fast photon emission that is the defining property of scintillation. However, differences in materials can result in significant variations in the absorption and emission process. In this respect, scintillators are generally placed into one of two broad categories: organic and inorganic scintillators.

6.3.1 Organic scintillators

Organic scintillators, composed of molecules that contain carbon, sometimes specified to require a C-H bond as well, exhibit strong scintillation properties. Organic scintillators are valued for their simplicity, strong yield, and fast transition times. The reason for this is their simple construction and the direct absorption-emission responsible for the scintillation. Due to their molecular structure, they feature tightly spaced electronic and vibrational energy levels that allow for a range of absorption wavelengths to different vibrational levels associated with the first electronic excited state. Once excited to one of these vibrational states, the molecule quickly relaxes to its lowest vibrational state. This process of excitation and vibrational relaxation can take place in a fraction of a nanosecond. The remaining excited electronic state will eventually decay back to the ground electronic state, usually in a few nanoseconds, resulting in photon emission. Figure 6.6 sketches the common processes [12].

119

The simplicity and speed of the organic scintillator relies on the Stokes shift, which describes the effective redshift in the fluorescence spectrum of a molecule compared with its absorption spectrum. The vibrational relaxations at sub-fluorescence timescales are partially responsible for this shift, as they reduce excitation energy non-radiatively. As a result, most photons emitted tend to be of longer wavelength than those absorbed. The Frank Condon principle further contributes to this wavelength shift. The principle states that during electronic transitions, a vibrational state change is more likely to occur if the wave functions overlap. In effect, changes between vibrational modes that have similar nuclear coordinates are favoured during fluorescence. For example, during a transition $S_1 \rightarrow S_0$, vibrational states (denoted by v) may change from $v = 0$ to $v' = 2$, further reducing emission wavelength [137]. In addition to the Stokes shift, several other processes may inhibit or distort the scintillation; the most important will be discussed later.
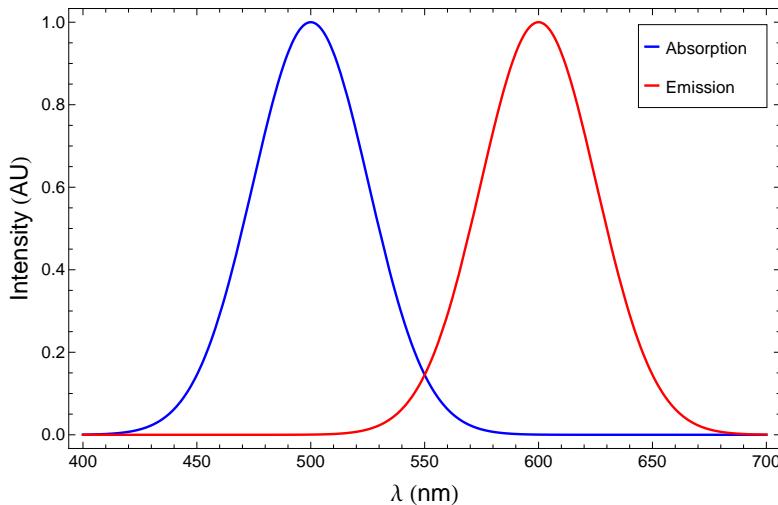
Figure 6.4: Visualization of the stokes shift, wherein the wavelength of fluorescence is much longer than that of absorption. Some fluorescent pigments exhibit extremely high shifts, up to 180 nm difference in spectral peaks [11].

6.3.2 Polyvinyl toluene

Polyvinyl toluene (PVT) consists of long chains covalently bonded vinyl toluene molecules. It is commonly used as the primary solvent for a family of plastic scintillators. While PVT itself scintillates, in most cases a mixture of solute molecules, usually fluorophores, such as p-terphenyl(TP) and 2,5-diphenyloxazole (PPO) are added to the PVT base to form a solution with more desired yield

characteristics. The choice of solute molecule and concentration can reduce emission wavelengths, improve response time, and increase yield [138]. Some variants also features a longer mean emission wavelength, which is favorable for CCD cameras [139].

PVT scintillators are commonly used as large volume detectors for time-of-flight (TOF) particle identification owing to their sensitivity to a wide range of radiation [129, 140]. They are flexible in composition and may be constructed in large arrays for bulk detection. This flexibility comes at the expense of robustness. PVT is a low Z organic compound comprised of covalently bonded vinyl toluene chains, and as such it is easily damaged by high LET radiation.

BC-408 and BC422-Q, PVT-based scintillators from St. Gobain, were chosen for the calibration tests. Both scintillators emit light through several channels, due to the presence of multiple solvents. BC-408 has a maximum emission at 424 nm with a stated decay time is 2.1 ns, while BC-422Q emission peaks at about 375 nm with a stated decay time of 1.6 ns [139, 141]. Both scintillators were 5 cm × 5 cm squares, coated with 50 μm of aluminum on one side. The BC-408 square was 0.5 mm thick and new, while the BC-422Q sheet was 1 mm thick and had been used previously in ion experiments. The effect of the difference in age on scintillator output is discussed in the analysis.

6.3.3   Inorganic scintillators

There are many models describing the interaction of radiation with an inorganic crystalline material and the subsequent light emissions, but most rely upon the same idea: ionizing radiation striking a crystal produces electrons. These electrons may travel a short distance before losing their energy to bremsstrahlung and producing additional electron-hole (e-h) pairs. These e-h pairs are coupled by virtue of their production within the valence and conducting band of the crystal, respectively. The pairs may travel much longer than free thermal electrons, and ultimately may reach an activator site. Once at the activator site, the e-h pair may recombine, driving a sub-band gap transition in the activator and subsequently the emission of light (e.g Tl^+ $3P_{0,1} \rightarrow 1S_0$).

The light is emitted at the wavelength associated with the transition to ground state and on a timescale associated with the relaxation of the energetic state [142]. Figure 6.5 illustrates this basic model. Emission is reduced by several non-radiative capture and quenching processes. In transit, the e-h pairs may become uncoupled and recaptured, or may become trapped in crystal defects. Once at the activator site, it is also possible that the dopant ion, rather than emitting a photon, quenches instead, emitting phonons into the lattice.
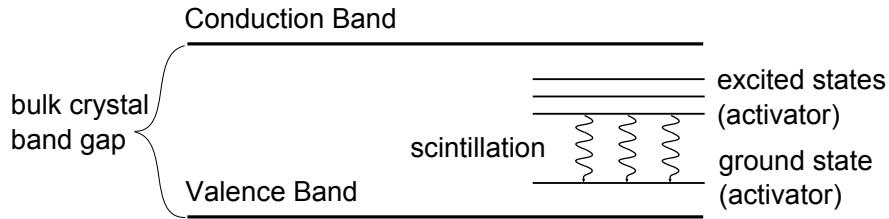
Figure 6.5: An illustration of the relative excitation energies of "activators" compared to the band gap of the crystal.

6.3.4 Additional properties

Several factors may impair the efficiency of a scintillator, reducing the light yield relative to the incident flux. Incoming particles may damage bonds or ionize atoms, creating defects in the material that absorb photons released by the scintillator, trapping or releasing them non-radiatively. Radiation also commonly creates depletion regions that trap electrons and may dissociate e-h pairs before they reach an activator site. Adjusting the dopant concentration may improve a material's resistance to these defects by reducing the relative abundance of these color centers [143]. Damage thresholds tend to be lower for organic scintillators, since their structure relies on weak interactions and covalent bonds to create the necessary vibrational energy transition structure. Fortunately, evidence suggests that emission characteristics such as wavelength and temporal response are not affected by damage [140].

Even undamaged scintillators have features that diminish the scintillation yield. Depending upon the construction of the scintillator, emitted light may be reflected by boundaries between different segments of material, or it may become trapped in lattice defects created during the material's construction. Similarly, these defects may capture or dissipate e-h pairs. In the event that e-h pairs make it to a ground-state activator site, it is still not guaranteed that scintillation occurs. Although it is most likely that absorption results in scintillation at the characteristic transition, it is possible for the excited electron to undergo a spin-orbit coupling, resulting in a spin flip that transfers the excitation to a triplet state of similar energy. This is known as intersystem crossing. Once in the triplet state, the system may still decay back to the singlet and release radiation. However, the time scale of this relaxation is many orders of magnitude greater than that of the usual scintillation transition, resulting in extremely delayed light response. This is responsible for the phenomena of phosphorescence. It is also possible that the triplet state electron relaxes non-radiatively, producing vibrational energy in the form of phonons through a process known as quenching. Intersystem crossing is more common

in high-Z molecules and in molecules with close vibrational level spacing. In all, the presence of these mechanisms further reduces scintillation yield and in the case of lattice defects and phosphorescence, producing a long-lived yield component or afterglow.
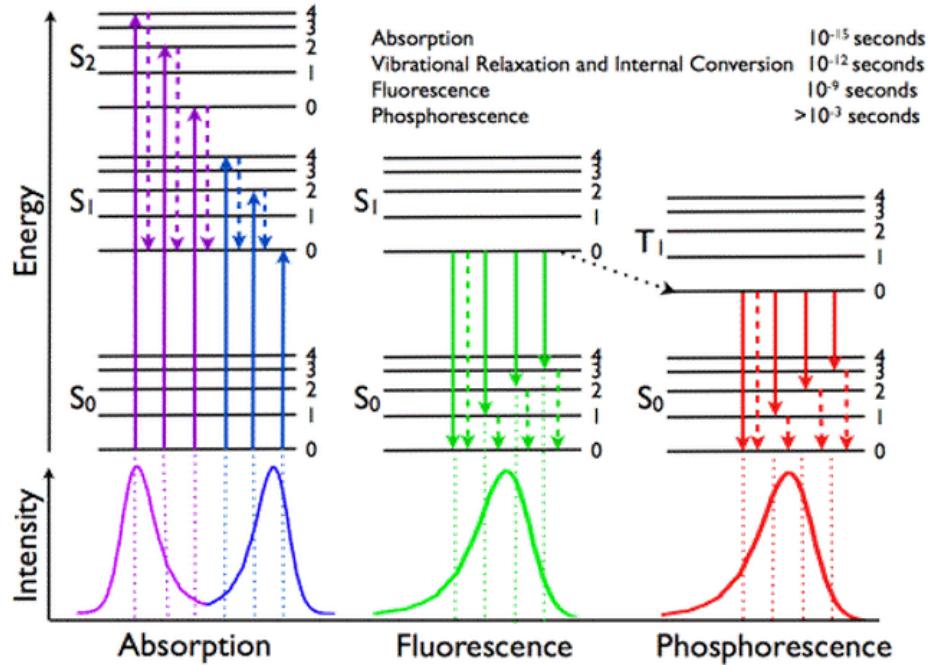
Figure 6.6: A Jablonski energy diagram denotes sample electronic and vibrational transitions that may occur after photon absorption and the relative intensity with which they occur. [12]

The appearance of afterglow among the scintillators we tested warrants additional discussion. Like phosphorescence, afterglow is the persistence of emission beyond the standard scintillation relaxation time. Unlike phosphorescence, the same fluorescent transition responsible for scintillation produces the light, as a result afterglow maintains a similar emission spectrum rather than being redshifted. A common mechanism responsible for afterglow is the trapping of electrons or holes in defects in a scintillator. The electron or hole is released when they acquire enough thermal energy to escape. This process delays the e-h pairs from reaching activators, and lengthens the time it takes for all e-h pairs to be dissipated within the lattice. Because afterglow emission is driven by thermal excitation of the pairs, the timescale and magnitude of the afterglow is highly dependent on the temperature of the material. At low temperature, trapped carriers are less likely to escape and contribute to the afterglow. Similarly, continued exposure to saturating beam fluxes further

increases the likelihood of trapped pairs. Crystal defects can be caused by impurities in the lattice or radiation damage [144].

6.3.5 Al_2O_3 :Cr_2O_3 — chromox

Consisting of 99.4% Alumina (Al_2O_3) doped with 0.5% Chromium Oxide (Cr_2O_3), chromox was developed by materials scientists at CERN looking for a robust and efficient scintillator to use as beam targets. The end result is a sturdy compound that can withstand high radiation doses ($10^{20}p^+/cm^2$ at flux levels of 10^{15} p$^+$/cm^2 per 0.5 μs pulse) yet still produce light yields comparable with standard scintillators even for energetic protons and heavy ions. Unlike CsI:Tl, Chromox is not hygroscopic, meaning that residual water vapor in storage or in the detection environment should not result in any defects or loss of performance. The ceramic has also been reported to operate stably between room temperature and 450C [145]. Although its ceramic construction improves toughness,the grain structure introduces many lattice impurities and boundaries that further increase afterglow and reduce spatial resolution.
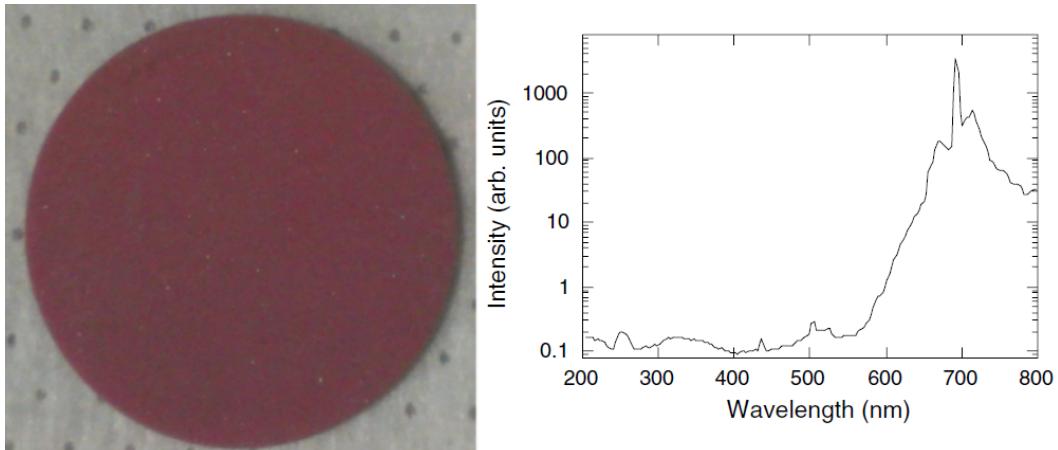
Figure 6.7: At left, an image of a 1 inch diameter chromox disc. To the right, the emission spectrum of scintillation light is shown, with a peak at 693 nm.

Chromox operates through a similar scintillation process to CsI:Tl, with the activator sites occupied by Cr^{3+} ions in about 0.5% abundance. Ionizing radiation ultimately produces e-h pairs which diffuse to the activator sites. If they make it, they excite the chromium ions that eventually decay from their lowest excited state ^{2}E to the ^{4}A$_2$ state and emit photons at the R lines. At room temperature, these occur at 692.2 nm and 694.3 nm with a decay time of 3.4 ms. Scintillation efficiency again depends on stopping power, but with its higher threshold for damage chromox remains a suitable choice for proton

detection. We obtained a 2.5 cm diameter, 0.82 mm thick chromox disc from Morgan Advanced Ceramics. We evaporated an approximately 225 nm thick aluminum coating on the disc; such coatings are commonly used to diminish noise from laser light and plasma emission.

6.4 Experimental methods

6.4.1 Stony Brook tandem van de graaff

Tandem Van de Graaff generators have been used as accelerating structures for more than half a century. They can provide stable, low energy accelerators for an array of ion species. A Van de Graaff operates by transporting charge on an insulating belt or chain along a loop through a conductive terminal. The charges are removed from the belt via a metal brush (in case of a belt) or pulley (in case of a chain) and collect at the surface of the conductive terminal, creating a very high voltage between the terminal and ground. Negative ions enter the Tandem at the positive terminal, and are accelerated towards ground.

Stony Brook uses a "King Tandem" built by the High Voltage Engineering Corporation [14, 146]. The Tandem, as evidenced by its name, utilizes two consecutive Van de Graaff accelerators. To accomplish this, the second Van de Graaff is positioned following the first, but charges the terminal negatively. A foil stripper is placed along the beamline in between the two accelerators. Negatively charged ions are accelerated away from the positive terminal through the foil strippers, which remove electrons from the ions, creating fully stripped positive ions. These ions are then further accelerated from ground (at the stripper) to the negatively charged terminal at the end of the second Van de Graaff. The Stony Brook Tandem operates using a Laddertron chain that improves peak voltage and stability [147]. Figure 6.8 shows the charging and acceleration schematic for an FN-type tandem Van de Graaff.

Completed in 1968, the Tandem offers considerable beam flexibility and stability. By using varying source materials, ions of many different species may be created, although protons comprise the majority of standard use. Ions may be accelerated to total energies of 2 MeV to 15 MeV with 0.01% accuracy. Beam currents are flexible from 200 pA to 8.0 nA. The tandem feeds into a target area composed of several beam lines with quadrupole focusing upstream of target chambers. Using the quadrupoles, focused beams of area 1.5 mm^2 can be created on target. Lastly, the tandem can be modified to create pulsed beams of varying length and current density, as will be described in the following paragraph. Table 6.1 summarizes the basic tandem operating parameters.
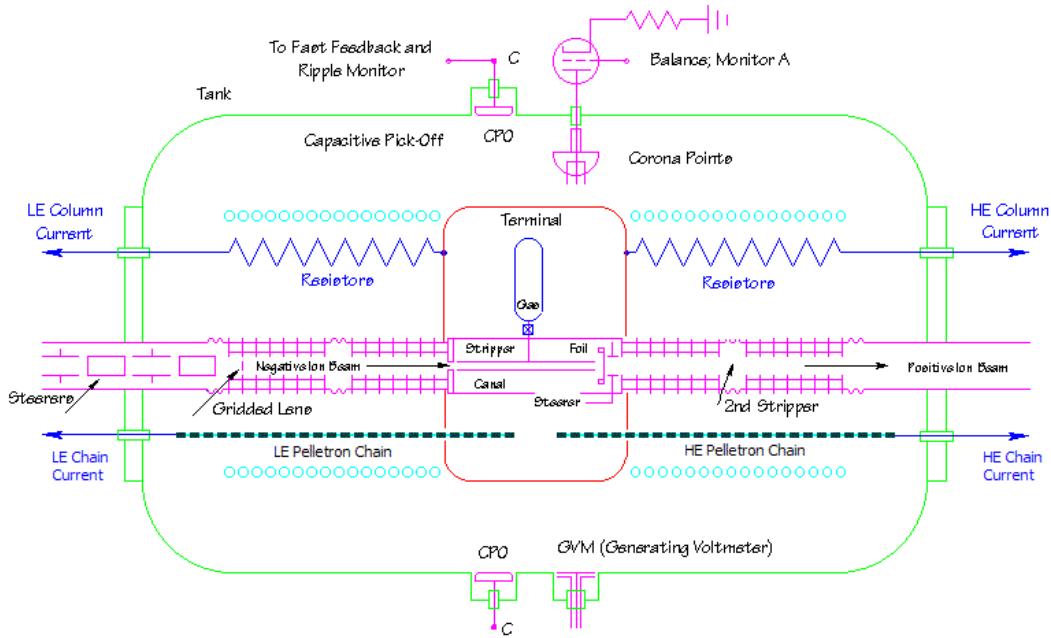
Figure 6.8: Charging schematic for an FN type tandem Van de Graaff [13]. The tandem at Stony Brook is an FN-type "King" machine, and uses a Laddertron rather than a pelletron chain, but the basic operational principles remain the same [14].

Beam Parameters	Minimum	Maximum	Relative Uncertainty
Proton Beam Energy (MeV)	2.00	15.00	$< 0.01\%$
Beam Current (nA)	0.2	8.0	$< 10\%$
Pulse Length (ms)	25	CW	$< 5\%$
1σ beam size (mm^2)	1.5	100	$< 10\%$

Table 6.1: Summary of Stony Brook Tandem Van de Graaff beam parameters delivered to target.

The tandem is by design a device for producing continuous beams of ions. However, we can benefit from using ion bunches with appropriate total charge and length. To accomplish this, a gating device is placed at the origin of the beam so as to effectively prevent passage of particles except at desired intervals. An electromagnet is placed before the aperture of the ion source. Increasing the voltage further deflects ions to the outside of the aperture, effectively blocking their entry to the tandem.A pulsed beam is created from by periodically blocking the continuous source beam. Using this method, pulses as short as 300 μs with a triangle wave intensity profile over time can be produced. Pulses shorter than 1 μs can be created using a fast switching circuit.

6.4.2 CCD camera

We used a Basler scA1400-17gm monochrome CCD camera, featuring a 58% quantum efficiency at 545 nm, a 14 pixel detection threshold, and a resolution of 1392 × 1040 pixels at 6.45μm x 6.45μm pixel size. The camera supports a 12-bit image depth. Maximum capture rate is 17 fps for full size images but can be increased by reducing image size through area of interest (A.O.I.) specification.

The camera is mounted with a 75 mm manual C-Mount TV lens from Computar. The lens has a maximum aperture of $f/1.4$ and paired with the CCD has a minimum focal distance of approximately 55 cm, corresponding to a magnification of $M = 0.16$ and an effective f-number of $f\sharp = 1.62$. All images were taken with maximum aperture as follows from $NA = 1/2f\sharp$. This equates to a solid angle at maximum aperture of [148]

$$\Omega = 2\pi(1 - \cos(\arcsin NA)) = 0.31 \text{ sr} \tag{6.2}$$

A monochromator was used to test the lens transparency over the wavelength range of 190 nm to 800 nm, which should cover the effective range of any potential scintillator. Baseline tests were performed to verify the accuracy of the machine, and two scans taken with the lens position adjusted to maximize the light entering the collecting sphere in the monochromator. Figure 6.9 shows the results of the scan.

6.4.3 The European Machine Vision Association model

Captured images are analyzed following the European Machine Vision Association (EMVA) Model 1288. The model provides a rigorous structure for identifying camera and lens properties and using them to establish proper rela-
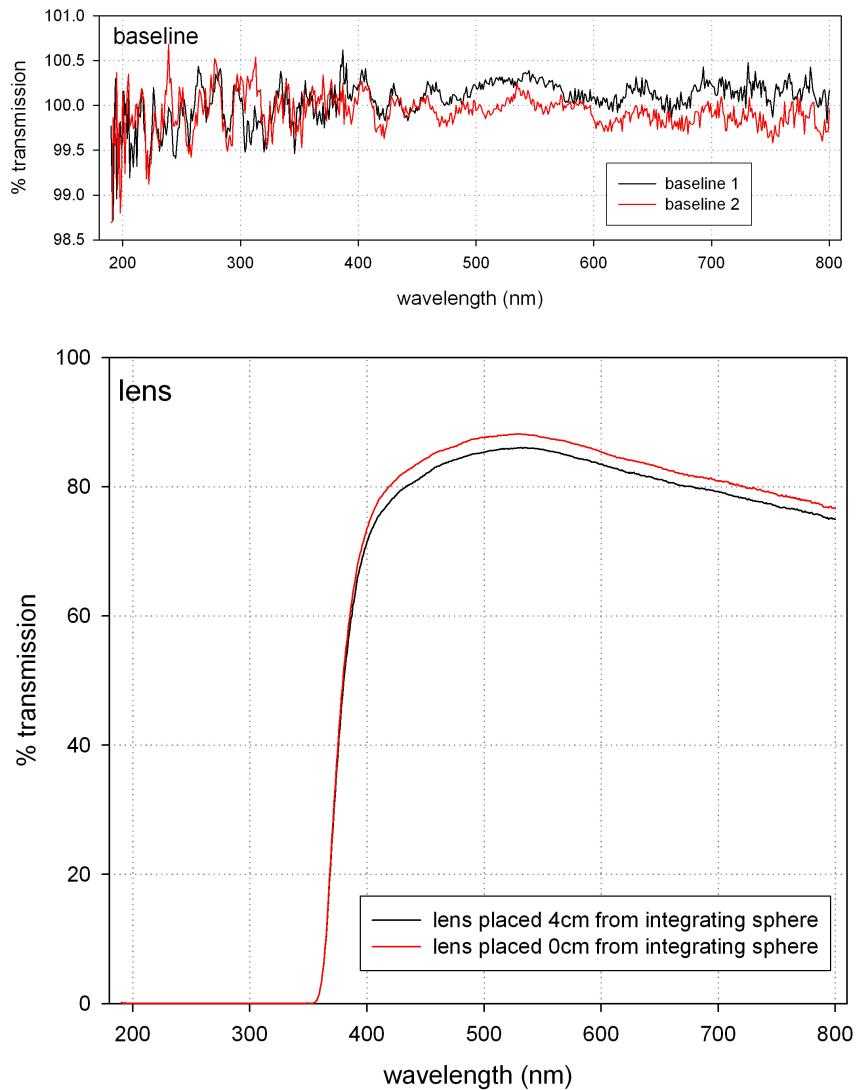
Figure 6.9: Below, lens transparency scan results at varying distance from collection sphere. Above, baseline tests for each scan. Tests were performed using a monochromator courtesy of Sean Stoll and BNL Physics Department.

128

tionships between digital image properties and source properties. To summarize, we are interested in calculating the yield of the scintillator, Y, in photons per proton. This yield represents the number of photons which are emitted from the surface of the scintillator, and thus takes the form of Equation 6.3,

$$Y = \frac{\mu_\gamma}{f_\gamma N_p} \tag{6.3}$$

where f_γ is the fraction of photons emitted which reach the CCD, and N_p is the number of protons incident on the scintillator, and μ_γ is the number of photons collected by the CCD, according to

$$\mu_\gamma = \frac{n_{pixels}(\mu_y - \mu_{y.dark})}{\eta K} \tag{6.4}$$

where n_{pixels} are the number of pixels in the A.O.I., μ_y is the average pixel grey value, $\mu_{y.dark}$ is the average background grey value, η is the quantum efficiency of the CCD at the wavelength of light captured, and K is the camera digital conversion gain. Lastly, we define f as

$$f_\gamma = \frac{\Omega}{4\pi} t_w t_l \cos\theta \tag{6.5}$$

where Ω is solid angle seen by the lens, θ is the angle between the emitting surface and the CCD, t_w and t_l are the window and lens transparencies, respectively. Figure 6.10 provides an overview of the process [15].
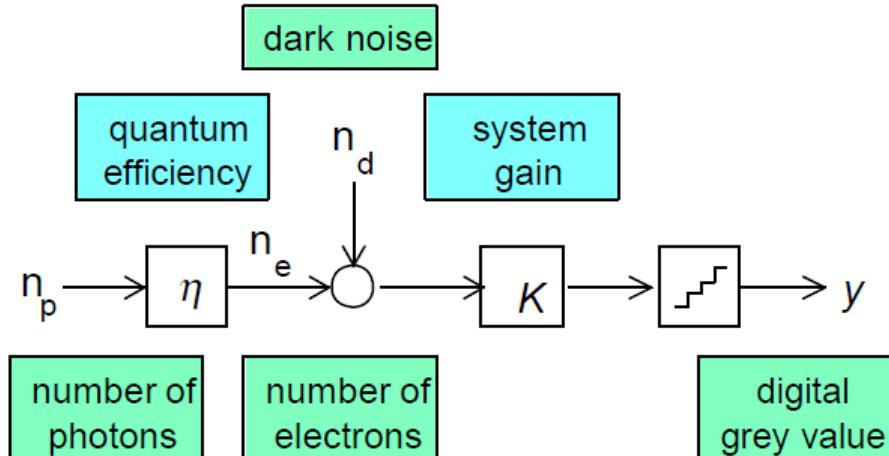
Figure 6.10: CCD image capture and processing protocol, specified by the European Machine Vision Association [15].

6.5 Results and analysis

The goal of these tests is to quantify the most important properties of the material while making note of qualities that may promote or inhibit the use of a particular scintillator. In particular, the yield, spatial resolution, and temporal response of each scintillator are the focal points of the investigation.

6.5.1 Yield analysis

While the models for light emission do have individual differences, the primary factors determining scintillation yield include incident beam energy, beam energy loss properties within the material (this takes into account both particle and material properties), and light transmission properties within the material. In performing our analysis, we take into consideration each factor and make the necessary adjustments to provide a compatible metric for each material.

Tests were performed in transmission, with the scintillator's rear face imaged. For each scintillator, proton energy was adjusted through a range of relevant energies, limited by the Tandem's abilities and by the thickness of the scintillator. Current was adjusted as needed to produce sufficient signal strength prior to capturing images, and was not intentionally varied during capture. Images were taken at each energy and current setting in groups of 300 images. Camera settings such as exposure time and gain were adjusted to provide excellent signal-noise; as such the gain was kept to a minimum. Additionally, a 200×200 pixel A.O.I. was defined according to the position of the beam prior to capturing images. This both streamlines and normalizes the capture and analysis process, as images are compared across the same number of pixels. Moreover, a smaller A.O.I. allows for a higher camera frame rate. Each image is tagged with its time of capture as well to properly catalogue the images. Control of these various parameters is maintained using a C++ program that uses the Basler Pylon application programming interface (API) for communicating with the camera.

Once the images are taken, analysis is done using Mathematica code, that reads in the image and processes grey values for each, performing the necessary sampling and calculations as outlined by the EMVA model to arrive at an average yield figure for any given set of images.

Measured yields

Figure 6.11 compares the yield for all three candidate scintillators, along with fits. Tests for the PVT-based discs were limited to 8 MeV due to proton penetration depth being larger than the detector thickness. Uncertainties are

dominated by current regulation instability on the order of ± 10%. Yields increase monotonically with energy as expected, although the nonlinear response at high energies requires further consideration. At 12 MeV, sparking within the tandem resulted in additional current instabilities, relegating the 12 MeV data point for Chromox as an outlier.
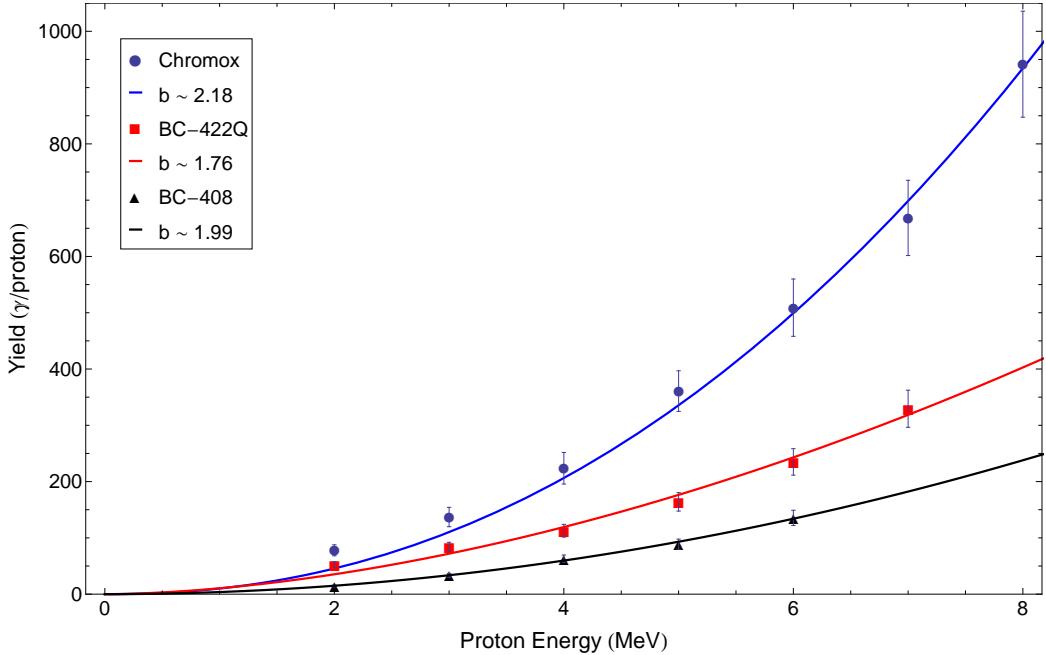
Figure 6.11: Comparison of raw yields $Y(E)$ for each scintillator in transmission, chromox (blue), BC-422Q (red), and BC-408(black). Fits are made of the form $Y(E) = aE^b$, and exponents are shown.

Chromox exhibits a high flux threshold for damage, and as such does not exhibit a characteristic decline in performance at increasing energies or dose rates. However, plastic scintillators exhibit a much lower threshold for damage, and this damage has been shown to result in a significant reduction in yield due to the onset of nonradiative quenching as well as irreversible damage to the molecular composition of the scintillator [129]. Quenching is found to be correlated to molecular damage to the scintillator as well as the presence of absorptive impurities.

In the absence of strong quenching or damage, scintillation light yield scales linearly with proton flux and particle energy [129]. To further investigate, we placed the PVT under a continuous 8 MeV beam at 1 nA and examined the yield over time; we recorded a 30% decline in mean intensity during 2 minutes of exposure, as seen in Figure 6.12. Afterwards, the beam was blocked for 5 minutes, at which point the average yield was found to be consistent with

the two minute value. As no significant afterglow was present, we conclude that there is a long term reduction in yield corresponding to damage to the scintillator. This type of long term response is indicative of the permanent damage associated with quenching at high radiation dose [129].
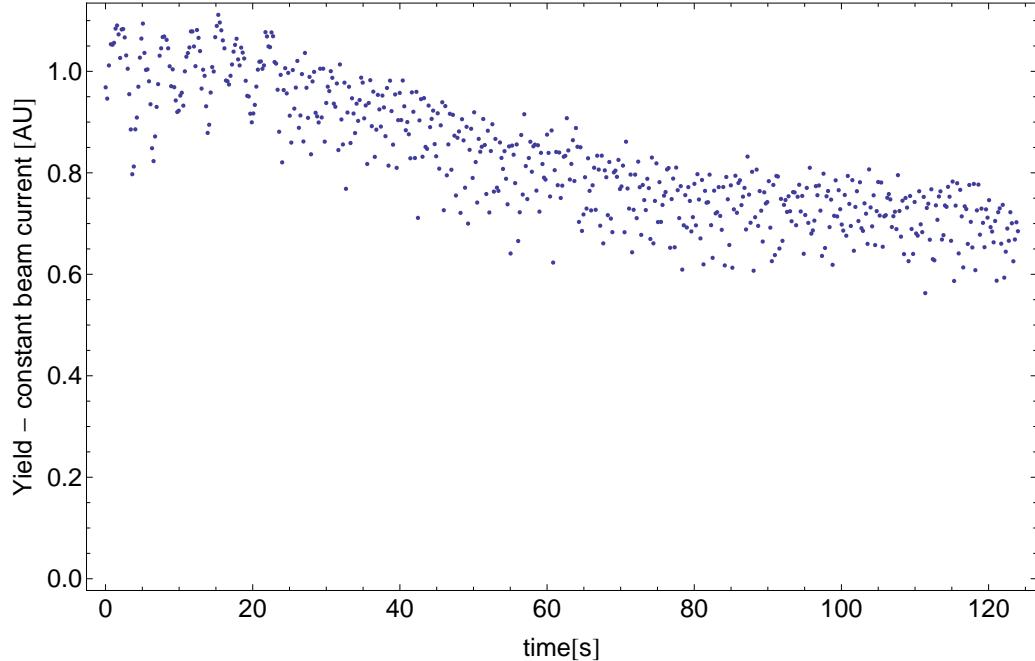
Figure 6.12: PVT yield decay under CW proton irradiation. The output decays monotonically over the course of the exposure, corresponding to a 32% drop in mean intensity.

Birks model and damage in scintillators

A model developed by Birks has been successful in describing the response of scintillators to particles of varying energy [149]. The model formulates the luminescent yield per unit track length, dL_B/dx as a function of the specific energy loss of the ionizing particle, dE/dx, according to

$$\frac{dL_B}{dx} = \frac{S\frac{dE}{dx}}{1 + kB\frac{dE}{dx}} \tag{6.6}$$

where S is the scintillation efficiency, while kB describes the quenching effect within the scintillator. In particular, the B parameter refers to the density of damaged molecules exposed to the ionizing radiation with specific energy loss dE/dx, while the k parameter gives the fraction of these molecules which then

exhibit quenching. The result is a kB parameter describing how much energy is lost due to quenching for a given dE/dx along the path of the particle. An empirical extension to this model was used by Craun and Smith [150] to fit data obtained by Smith et al [151]. They include include a parameter, c, which is second order in dE/dx,

$$\frac{dL_B}{dx} = \frac{S\frac{dE}{dx}}{1 + kB\frac{dE}{dx} + c(\frac{dE}{dx})^2} \tag{6.7}$$

In the limit of small dE/dx, yield tends towards a linear relationship, with the scintillation efficiency approaching unity. This is the case for electrons up to a few hundred keV and other low LET particles [152].

Best fit parameters using Equation 6.7 for the raw yields of each scintillator are shown in Table 6.2. In finding parameters that best fit the data, the Stopping Range of Ions in Matter database was used to produce stopping power curves over the range of energies for each material [153]. Note that a negative c parameter is observed for materials that exhibit a stronger response at low energies [154].

Table 6.2: Comparison between Birk's model (Equation 6.7) fits and transparency adjusted model fits (Equation 6.9). The right column gives the goodness of fit according to the residual sum of squares (RSS).

Material	Metric	kB (g/MeV cm^2)	c ((g/MeV cm^2)^2)	λ (cm)	RSS
Chromox	Birks	5.375×10^{-3}	-1.5×10^{-6}	-	0.0381
	Transparency	2.33×10^{-2}	0	0.058	0.0168
BC-422Q	Birks	2.25×10^{-2}	1.15×10^{-5}	-	0.0766
	Transparency	1.25×10^{-2}	1.05×10^{-6}	1.5	0.0638
BC-408	Birks	2×10^{-2}	2×10^{-7}	-	0.00586
	Transparency	2×10^{-2}	2×10^{-7}	210	0.00586

Transparency considerations

The transparency of each scintillator is also considered as a means to explain nonlinearities in the yields, as well as possible deviations from expected performance. Scattering and absorption at grain boundaries in chromox necessarily plays a role in the attenuation of emitted scintillation light, introducing additional nonlinearities to the yield. Moreover, damaged plastic scintillators demonstrate reduced transparency [140]. We introduce a transmission coefficient $T(x)$, which defines the fraction of scintillation light emitted at a given

position within the material that escapes the rear side. Assuming homogeneity holds over the attenuation length of the material, λ_t, we define

$$T(x) = e^{\frac{-(w-x)}{\lambda_t}} \tag{6.8}$$

where w is the width of the scintillator. This transmission fraction is then applied to Birk's formula for dL/dx along the path length of the ionizing particle. We have

$$\frac{dL_t}{dx} = T(x)\frac{dL_B}{dx} \tag{6.9}$$

Figure 6.13 shows a fit of this model to the chromox data. The transparency adjustment improves the original fit without the use of the empirical c parameter, suggesting that at the energies probed, there is no second order quenching effect. This is in agreement with the high radiation resistance of chromox, and the lack of observable damage.
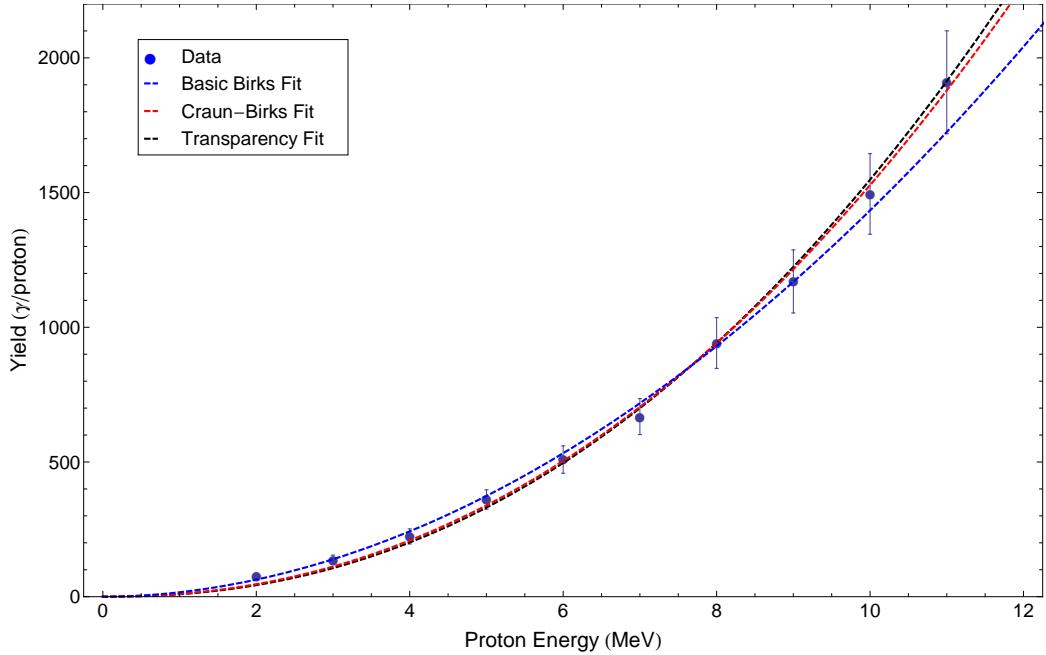
Figure 6.13: Chromox yield shown fit to a transparency adjusted Birk's model, along with measured data and the best fit. With a calculated attenuation length of $\lambda_t = 1.1$ mm, the model produces a fit with $kB = 0.018$ g/MeV cm^2 and $c = 0$.

This model may be extended to BC-422Q, wherein molecular damage leads to an increase in scattering and diminished transparency [140]. Table 6.2 shows

fitting parameters both with and without transparency adjustments. Fitting the transparency metric suggests an attenuation length of 1.5 mm in the damaged scintillator, compared to lengths of greater than 1 meter for undamaged PVT-based compounds. No additional scattering effects are observed in the output of BC-408, and as such we would not expect to see an improvement in fit by introducing a transparency correction. This is confirmed by the best fit parameters.

6.5.2   Scattering effects

The previous discussion of transparency presupposes the existence of significant scattering effects in both chromox and damaged plastic scintillators, specifically BC-422Q. To further evaluate the role of scattering, we introduce a simple model to qualify scattering in chromox. We then extend this model to plastic scintillators and demonstrate that scattering effects resulting from damage to the scintillator can introduce significant changes to the beam profile. As before, we consider the chromox material composed of grains with $w_g = 12\,\mu\text{m}$. We assume a beam of photons with Gaussian spatial and angular distributions σ_x and σ_θ, respectively. Such an assumption provides a reasonable approximation to experimental data, as accelerated ions will pass through a circular pinhole, producing a Gaussian profile. We observe this profile in images taken with a fresh sheet of BC-408 as shown in Figure 6.14.

We assume that at each grain boundary, a given photon will scatter according to a scattering distribution Θ, and define a scattering parameter $\alpha = w/w_g$ representing the effective number of scattering occurrences. The cumulative scattering produces tails in the outgoing photon spatial distribution. The tails fit a Lorentzian distribution, as seen in Figure 6.15. The results of this simple chromox model are confirmed experimentally by images taken of laser-accelerated protons. Figure 6.15 confirms that the resultant spectrum fits a Lorentzian profile.

This model may be extended to a plastic scintillator, wherein structural damage to the solvent plays the role of grain boundaries, causing scintillation light to scatter as it traverses the plastic. The older PVT based BC-422Q showed significant signs of damage-induced reductions in transparency. This damage also resulted in scattering, producing a Lorentzian distribution of light from an incoming Gaussian profile, as shown in Figure 6.14. Conversely, an undamaged sheet of BC-408 plastic scintillator does not exhibit significant scattering, thus reproducing the expected Gaussian profile of the accelerated ion beam. These results are especially significant because an accurate measurement of the beam's Gaussian width is used to determine the energy spread of the accelerated ion beam.
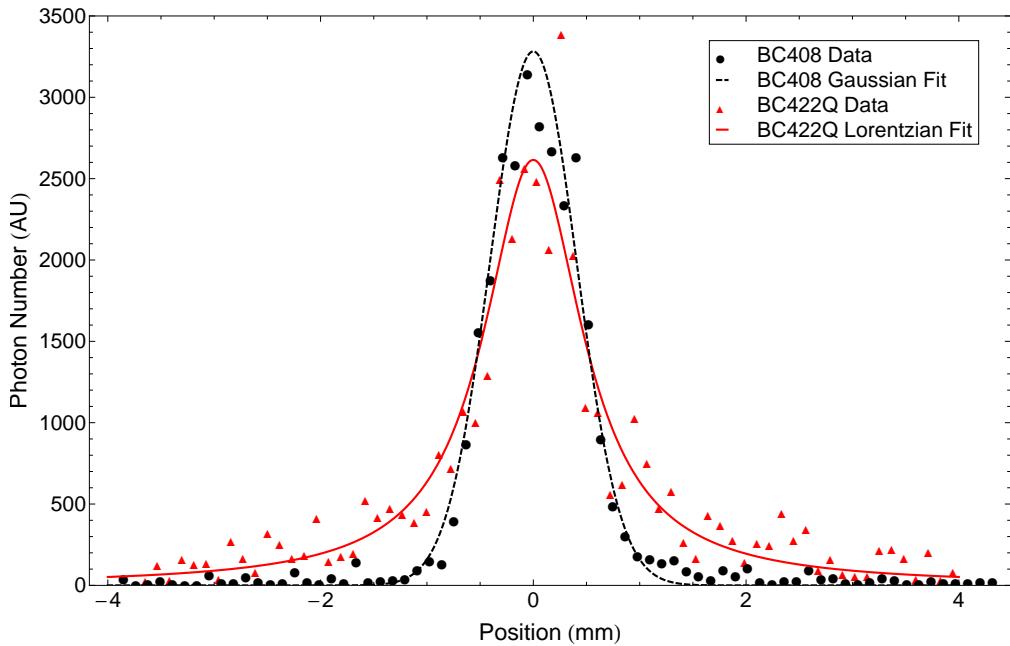
Figure 6.14: An undamaged sheet of BC-408 exposed to a Gaussian proton beam reproduces a Gaussian profile, while a damaged BC-422Q sheet produces a Lorentzian profile with signficant scattering.

6.5.3 Afterglow

Though boasting impressive scintillation response to ions and high LET particles, chromox suffers from a considerable afterglow effect. This is in part due to quenching and defects as in CsI:Tl. However, chromox is fundamentally different in its physical constitution than CsI:Tl and other crystalline scintillators. Chromox is constructed from 10-15 μm diameter grains. Thus, incoming light may be refracted or reflected as it crosses the boundary between grains. The light may also be trapped in lattice defects that are more abundant at these boundaries. As a result of this, the afterglow effect is enhanced, as more light is trapped for a longer period of time as quenching and reflection happen more frequently. This effect has been measured in our chromox sample both under continuous and pulsed beam conditions.

We generated a beam of 6.0 MeV protons at a current of 500 pA, and imaged the disc in reflection, so that light emitted from the front surface (the surface of beam incidence) was collected. The beam was shutoff after an incidence period of several minutes, and the camera was set to take pictures continuously at a 1 ms exposure time at approximately a 1 Hz repetition rate. We found that there is noticeable afterglow for more than a minute after beam
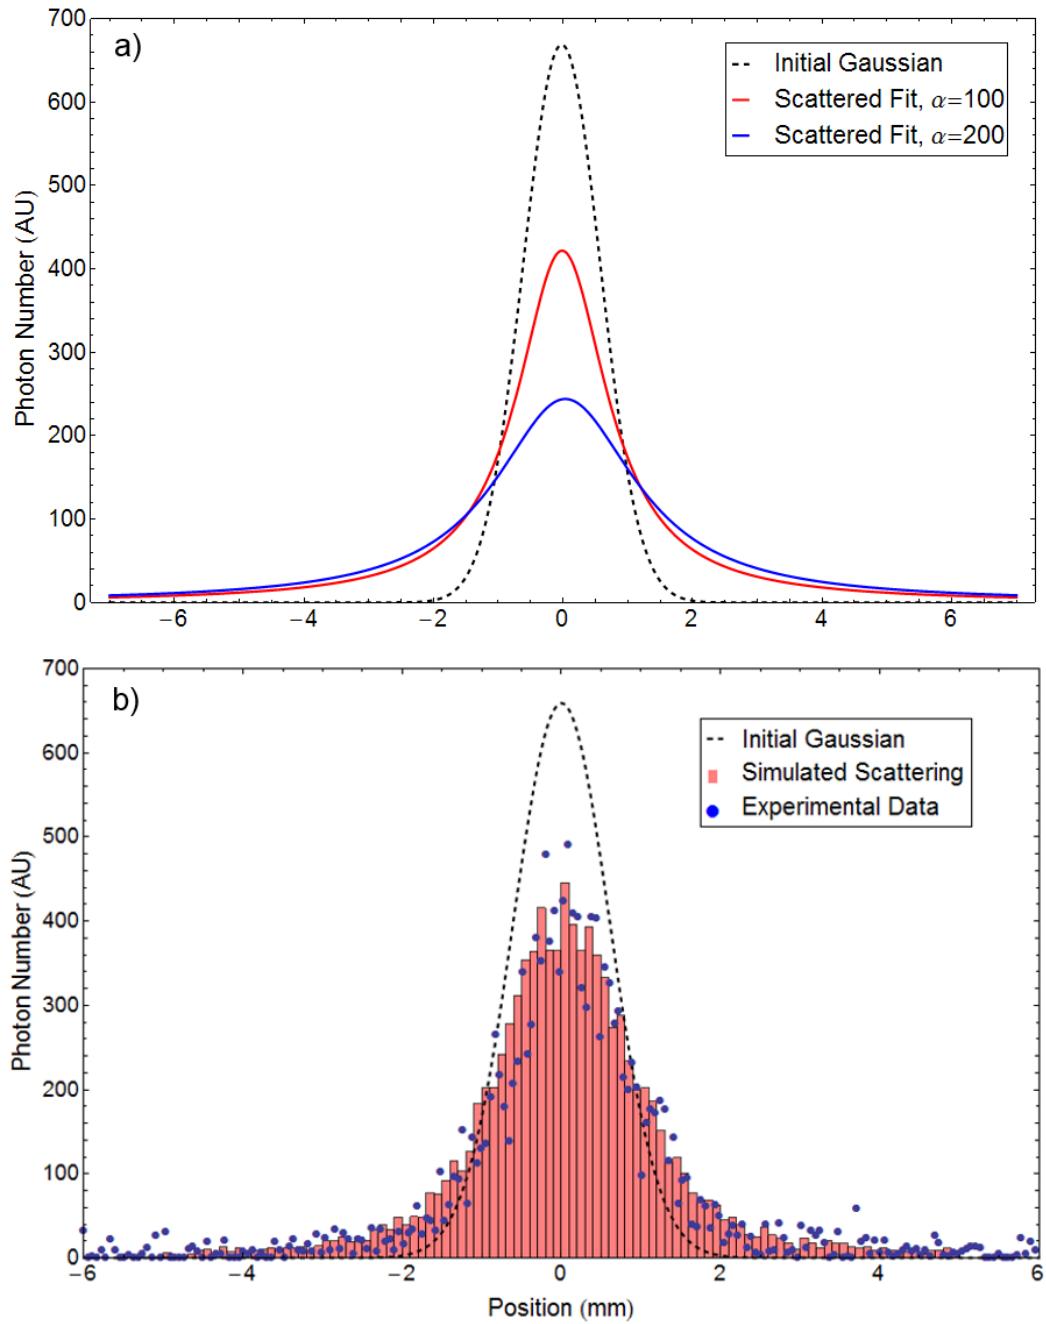
Figure 6.15: (a) The simple scattering model applied to a 1 mm chromox disc for varying scattering parameter α. (b) A photon spectrum taken from a laser-accelerated proton beam imaged with chromox.

is shut off. The sample exhibited a strong immediate afterglow, 28% of the value under beam incidence 1 second after beam shutoff. After 10 seconds, the afterglow intensity was less than 10% of the intensity under incidence. Even after 90 seconds, an afterglow of approximately 1.5% of the intensity under incidence was observed.

Using these measurements, we were able to compare several fits to learn the behavior of the afterglow at long time scales. Previous studies [127] suggest that a stretched exponential of the form

$$Y(t) = e^{-(\frac{t}{\tau})^{\beta}} \tag{6.10}$$

is a natural fit for afterglow, where the natural decay of the signal due to release of quenched photons is further modulated by the dispersive factor β. However, our data suggests that a normalized sum of standard and stretched exponentials,

$$Y(t) = \frac{1}{2}(e^{-\frac{\tau}{\delta}} + e^{-(\frac{t}{\tau})^{\beta}}) \tag{6.11}$$

hereafter referred to as the combination fit, fits better in both the short and long time scales of the afterglow. Our best fit suggests $\delta = 4.3$ s, $\tau = 23.3$ s, and $\beta = 0.672$; Figure 6.16(a) shows a comparison between best fits of the standard exponential, stretched exponential, and combination fit. The combination fit follows the long term behavior while also improving the short term fit.

As our laser acceleration scheme produces sub-ns pulses at about a 1 Hz repetition rate, the response to a pulsed beam is more applicable to a laser acceleration experiment. By modulating a beam aperture, we generated 25 ms pulses at 3 Hz repetition rate. Peak chromox response was measured, and the afterglow image intensity is plotted versus time in Figure 6.16 (b). It can be seen that in the case of pulsed beam, the light output falls much more quickly, reaching less than 3% of its nominal value after 0.1 seconds. Again, a combination fit seems to provide the best candidate, with values of $\delta = 5$ ms, $\tau = 0.7$ ms, and $\beta = 0.21$, although a stretched exponential is also effective. This fit is characteristic of phosphorescence and suggests that chromox is suitable for pulsed use at Hz-range repetition rates.

6.6   Conclusion

We have measured the response of several candidate scintillators for use in a high repetition rate laser driven ion acceleration experiment. Yield and scattering models were applied to obtain a general characterization of the scintil-
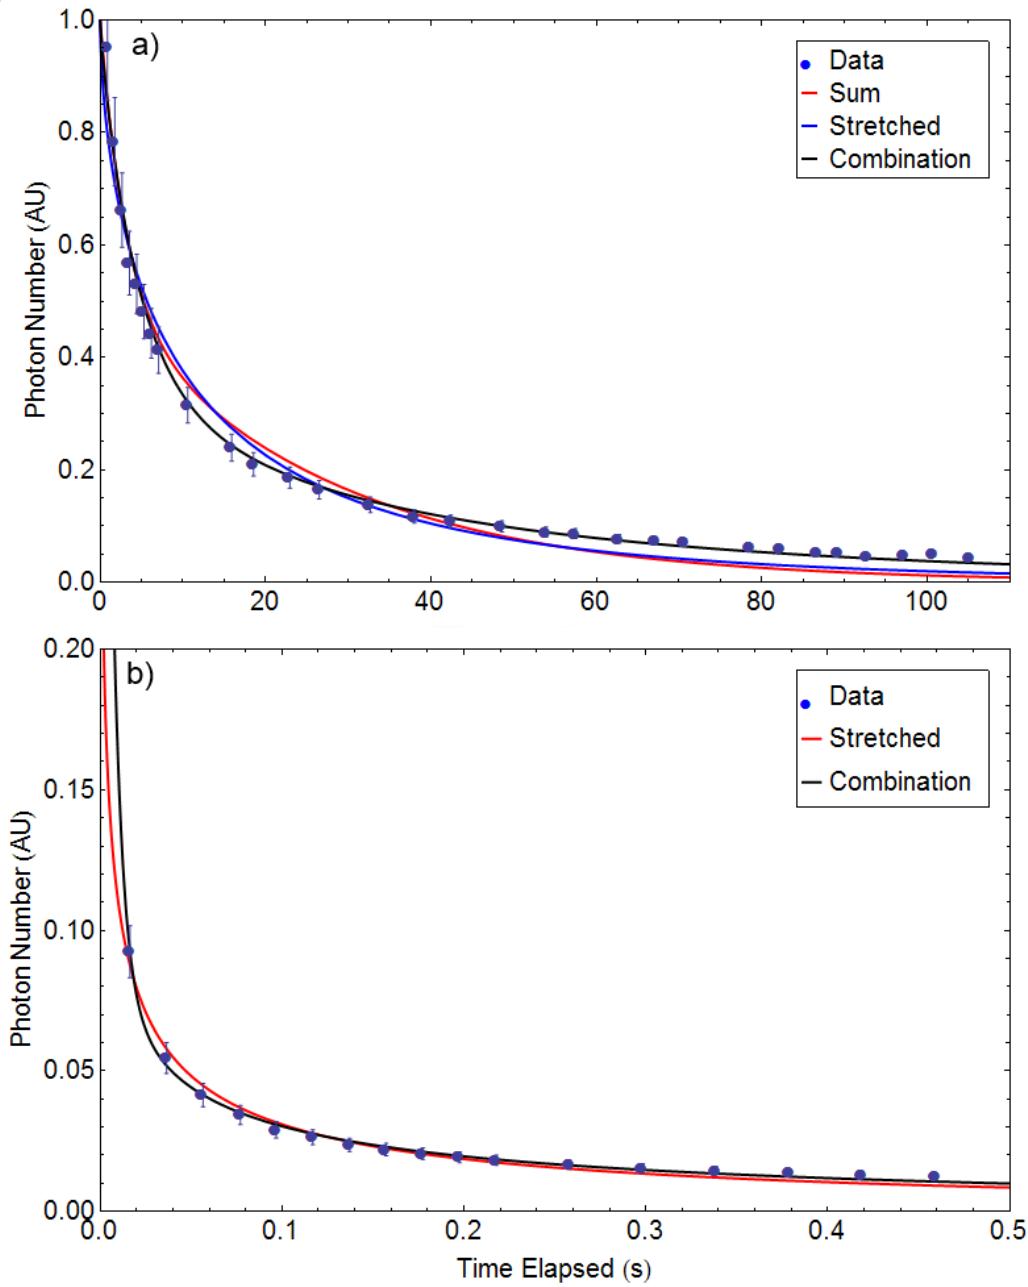
Figure 6.16: (a) A comparison of fits to the long term afterglow behavior of chromox under CW proton beam exposure. (b) A 5 nA proton beam is pulsed for 25 ms on, 325 ms off.

lator response. Our findings suggest that chromox has the strongest response to protons in the 2-10 MeV range. Moreover, the afterglow in response to a pulsed beam is not significant on a 1 second time scale, similar to the current maximum laser repetition rate at the ATF. Natural opacity significantly reduces yield, depending on the thickness of the scintillator. Likewise, quenching remains a concern at high energies, but chromox demonstrated the best radiation hardness among candidates. However, at low energies, photon scattering at grain boundaries reduces both yield and resolution for chromox. Assuming comparable disc thickness, a PVT-based scintillator provides a superior response at energies less than 4 MeV. However, damage due to cumulative exposure significantly reduces both yield and resolution in plastic scintillators. This is a critical issue for diagnostics relying on output distributions to quantify beam properties such as energy and charge state, and in this case plastic scintillators should be avoided for high energy and intensity ion beam diagnostics. For future high energy and high repetition rate imaging, chromox will provide a valuable diagnostic tool. Nonetheless, the resolution of $BC - 408$ at low energies makes it the clear choice for ~ 1 MeV ion detection at low fluxes, and is the scintillator of choice for current ion acceleration experiments.

Chapter 7

Laser driven shock wave acceleration of ions at BNL Accelerator Test Facility

This chapter emphasizes the key findings regarding the observation and analysis of helium ion beam generation using the experimental arrangement described in chapter 5. First, a discussion of basic hydrodynamics prefaces the interaction of a low intensity pre-pulse with the neutral gas jet target. Next, experimental results with helium targets are presented, with an emphasis on the effect of the pre-pulse on the longitudinal density profile of the target. A series of particle-in-cell simulations have been performed to analyze the action of the primary intense pulse on the shaped plasma profile. The simulations are discussed and evidence of collisionless shock acceleration is presented and evaluated. An overview of some of the key results, as well as a discussion of experiments with hydrogen gas jets, may be found in several conference proceedings [155] and a submitted manuscript [156].

The unique generation of a 5 ps pre-pulse at a fixed time prior to the main intense pulse is a defining feature of the approach to this experiment. As will be shown, the impact of the pre-pulse is critical to determining the nature of the main pulse interaction. Due to the nature of a two pulse chain with relatively long time scales, as well as the natural variation in neutral density profile achievable by the gas jet, the specific target dynamics are extremely complex. A discussion of the basic hydrodynamic evolution of the jet under the influence of the pre-pulse provides an excellent starting point for understanding the subsequent laser plasma interaction.

7.1 Hydrodynamic target shaping

A laser pulse interacting with the gas target with intensities below the appearance threshold for He^+ can significantly heat the target through collisional processes. Energy is absorbed along the laser axis, with a considerable fraction deposited near the focal position of the laser. This energy drives a shock which expands radially into the target. Assume that the unshocked gas has flow velocity u_1, density ρ_1, pressure P_1, and internal energy ϵ_1, while the shocked gas is characterized by u_2, ρ_2, P_2, and ϵ_2. The relation between gas properties ahead of and behind the shock are defined by the shock jump conditions, which follow from the fluid equations describing continuity, conservation of momentum and energy. Assuming a steady state, planar shock, and ignoring viscosity effects outside of the transition layer (the shock itself), one obtains the Rankine-Hugoniot jump conditions [137],

$$\rho_1 u_1 = \rho_2 u_2 \tag{7.1}$$
$$\rho_1 u_1^2 + P_1 = \rho_2 u_2^2 + P_2 \tag{7.2}$$
$$\frac{1}{2}u_1^2 + \epsilon_1 + \frac{P_1}{\rho_1} = \frac{1}{2}u_2^2 + \epsilon_2 + \frac{P_2}{\rho_2} \tag{7.3}$$

The shock is characterized by its Mach number, $M = u_1/c_s$, where c_s is the sound speed in the upstream gas, given for an ideal gas by

$$c_s = \sqrt{\frac{\gamma P}{\rho}} = \sqrt{\frac{\gamma R T}{m}} \tag{7.4}$$

where R is the gas constant, m is the molecular mass, T is temperature, and γ is the adiabatic constant of the gas, the ratio of specific heats $\gamma = C_p/C_v$. This constant depends on the degrees of freedom in molecular motion; as the number increases, $\gamma \to 1$. For a monoatomic gas such as helium, $\gamma = 5/3$, but for diatomic molecules such as Hydrogen gas (H_2), $\gamma \approx 1.4$.

The jump conditions can be re-written in terms of the Mach number,

$$\frac{\rho_2}{\rho_1} = \frac{(\gamma + 1)M^2}{(\gamma - 1)M^2 + 2} \tag{7.5}$$
$$\frac{P_2}{P_1} = \frac{2\gamma M^2 - (\gamma - 1)}{\gamma + 1} \tag{7.6}$$
$$\frac{T_2}{T_1} = \frac{((\gamma - 1)M^2 + 2)(2\gamma M^2 - (\gamma - 1))}{(\gamma - 1)M^2} \tag{7.7}$$

For a strong shock ($M \gg 1$), these conditions reduce to

$$\frac{\rho_2}{\rho_1} \approx \frac{\gamma + 1}{\gamma - 1} \tag{7.8}$$
$$\frac{P_2}{P_1} \approx \frac{2\gamma}{\gamma + 1}M^2 \tag{7.9}$$
$$\frac{T_2}{T_1} \approx \frac{2\gamma(\gamma - 1)}{(\gamma + 1)^2}M^2 \tag{7.10}$$

For a strong shock, the density ratio at the transition layer peaks at ρ_2/ρ_1 4 for $\gamma = 5/3$. The pressure within the shocked region is much larger than the ambient pressure of the gas; this describes the blast wave regime of a shock.

The dynamics of the blast wave are described by Sedov' self-similar solution. The self-similarity variable ξ is defined by

$$\xi = \frac{r}{(Et^2/\rho_1)^{1/5}} \tag{7.11}$$

where E is the deposited energy driving the shock. The radius r_{sh} and rate of expansion D_{sh} of the blast wave are given by

$$r_{sh}(t) = \xi_0 \left(\frac{Et^2}{\rho_1}\right)^2 \tag{7.12}$$

$$D_{sh}(t) = \frac{2}{5}\xi_0 \left(\frac{E}{\rho_1 t^3}\right)^2 = \frac{2}{5}\xi_0 \left(\frac{E}{\rho_1 t^3}\right)^{1/2} R_{sh}^{-3/2} \tag{7.13}$$

where ξ_0 is a dimensionless parameter which normalizes the self-similarity variable [137]. Its value can be calculated numerically by requiring that energy conservation hold for the solution ($\xi_0 = 1.17$ for $\gamma = 5/3$). Equation 7.13 is invalid at $t = 0$, but otherwise provides a good approximation of the shock speed. It should also be noted that this solution is predicated on the deposition of energy into a point-like source over a short timescale. As the ATF CO_2 pulse is very short (5 ps) relative to the expansion time (25 ns), this assumption is reasonable for most of the pre-pulse energies considered.

While Sedov produced exact analytical solutions to the density, pressure, and velocity within the shocked region, the expressions are complicated parametric functions of an internal parameter α [157]. Taylor produced an approximation [158, 159] which provides reasonable agreement within 5% [160], according to

$$n(\gamma) = \frac{7\gamma - 1}{\gamma^2 - 1} \tag{7.14}$$

$$p(\gamma) = \frac{2(\gamma + 5)}{7 - \gamma} \tag{7.15}$$

$$\rho_T(r, t, \gamma) = \left(\frac{r}{r_{sh}}\right)^{\frac{3}{\gamma-1}} \left(\frac{\gamma + 1}{\gamma} - \frac{\left(\frac{r}{r_{sh}}\right)^{n(\gamma)-1}}{\gamma}\right)^{-p(\gamma)} \tag{7.16}$$

Figure 7.1 shows the predicted axial density profile behind a spherical shock according to the solution given by Equation 7.16.

144
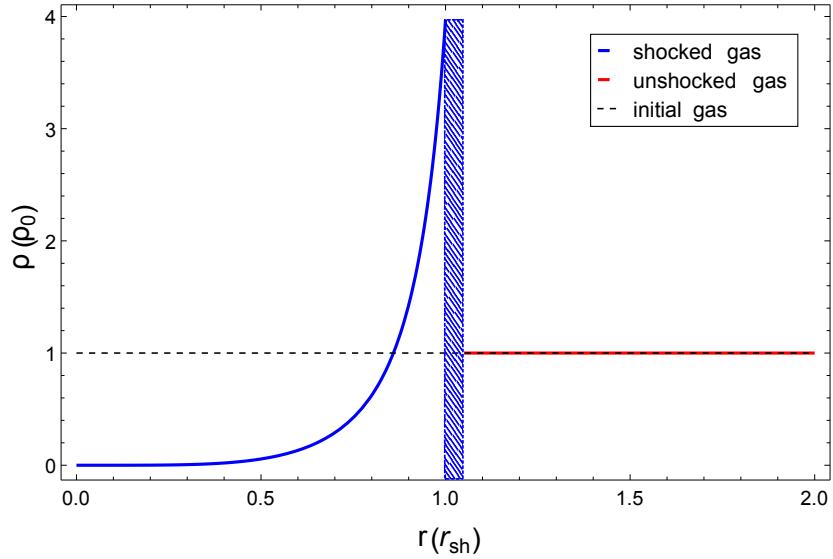
Figure 7.1: Sedov-Taylor approximation of longitudinal density profile of a blast wave propagating into a uniform gas of density ρ_0. The blue line is the density within the shocked region, while the red line depicts the upstream, unshocked gas. The transition region is denoted by the narrow region between the two, the thickness of which is realistically several mean free paths.

7.1.1   FLASH simulations

Two dimensional, cylindrically-symmetric, hydrodynamic simulations were performed using the FLASH hydrodynamics code to simulate the deposition of pre-pulse energy into the gas target (assumed to approximate an ideal gas). The longitudinal density profile of the neutral gas was used, but variations in density along the vertical axis were ignored. Energy was deposited instantaneously into a cylinder of length 600 μm and radius 80μm centered at the laser focus, and the target was allowed to evolve for 25 ns. While FLASH efficiently tracks particle density and temperature, the simulations performed did not consider ionization. However, assuming the gas is in local thermal equilibrium (LTE), we may apply the Saha equation to determine the ionization state of the gas as a function of temperature [137].

For a free gas in LTE, the ratio of particles in ionization state $q + 1$ to those in ionization state q is given by

$$\frac{n_{q+1}n_e}{n_q} = \frac{(2\pi m k_B T)^{3/2}}{h^3}\frac{2g_{q+1}}{g_q}e^{-\frac{\chi}{k_B T}} \tag{7.17}$$

where h is Planck's constant, g_q and g_{q+1} are the statistical weights (degen-

eracy) of the ground states of the different ionization levels, and χ is the ionization potential for state q_{i+1}. For helium, $g_1/g_0 = 2$ and $g_2/g_1 = 0.5$. For multiple ionizations, the quantities n_{q+1}, n_q, and n_e satisfy conservation conditions

$$\sum_i n_i = n_{He} \tag{7.18}$$

$$\sum_i q_i n_i = n_e \tag{7.19}$$

As the ionization level is extremely sensitive to temperature, we expect to see stronger ionization farther behind the shock, where the density is lower and the temperature is higher. Figure 7.2 shows a sample ionization profile behind a strong shock.
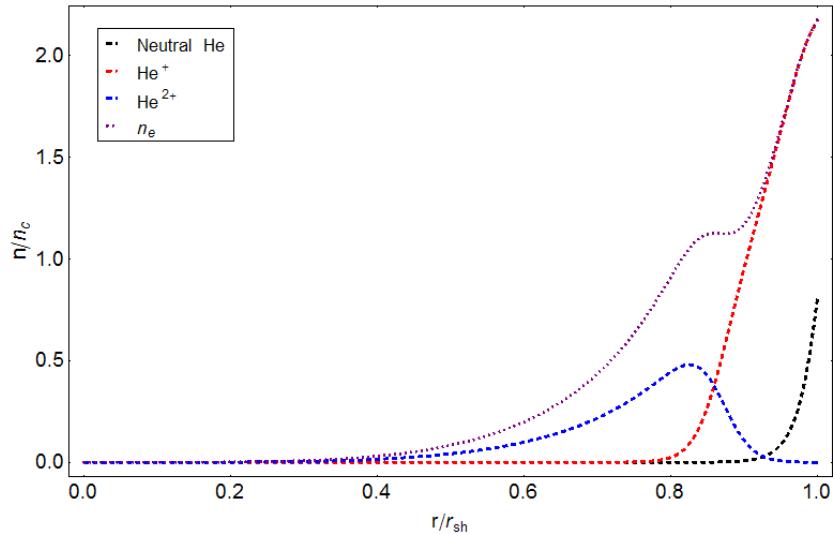
Figure 7.2: Predicted ionization states behind a spherical shock in the Sedov-Taylor approximation assuming about 150 mJ of deposited laser energy.

Figure 7.3 shows the resulting simulated electron density profile in comparison with an experimentally extracted density profile from interferometry taken just prior to the arrival of the main pulse. As the pre-pulse only minimally ionizes the gas, phase extraction produces a noisy image, requiring additional smoothing to produce a physical profile; the resulting density profile is inaccurate at short length scales. However, the size of the region of ionization and the approximate level of ionization agree between simulation and experiment.

Examination of the total particle density reveals more information about
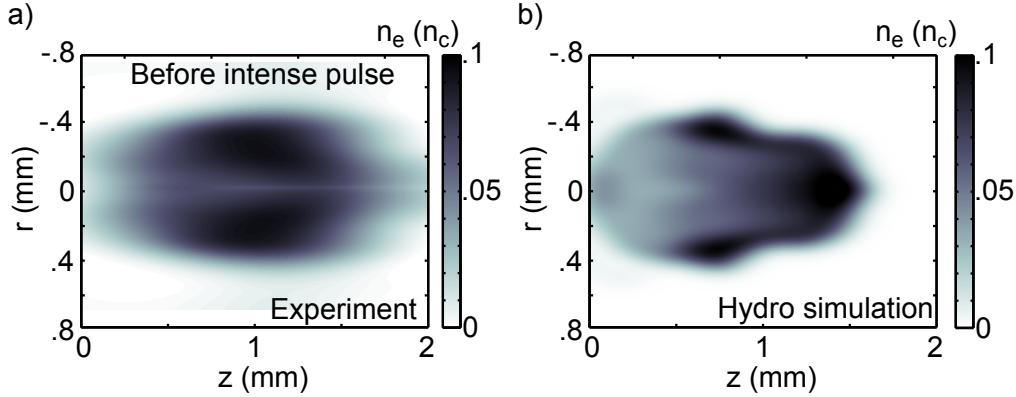
Figure 7.3: (a)Experimentally extracted electron density profile from interferometry just prior to main pulse arrival. (b) Electron density calculated from 2D FLASH code simulating hydrodynamic expansion of gas and estimating the resultant ionization using the Saha equation (Equation 7.17). The initial gas profile is centered at $z = 1.5$ mm. Density profiles are given in units of the critical plasma density n_c for 10.3 μm light.

the hydrodynamic evolution of the neutral gas profile under the influence of the pre-pulse. Figure 7.4 (a) shows the total particle density after 25 ns of expansion. Here, the characteristic blast wave profile expanding radially into the plasma generates a cavity wall with peak densities commensurate with that of the strong shock condition in Equation 7.10. Unfortunately, this wall cannot be seen prior to the interaction of the main pulse due to limited ionization of the high density region. Instead, interferometry taken approximately 300 ps after the main pulse arrival is used as an indicator of the blast wave conditions. Since the rate of expansion of the cavity is on the order of the signal speed in helium (10^4 ms^{-1}), the 300 ps time difference can be ignored. Figure 7.4 (b) shows the extracted density profile from one interferometry image; striking similarities in the size, shape, and peak densities can be seen between this profile and the hydrodynamic simulation. The egg shape of the blast wave is typical of the envelope solutions to inhomogeneous propagation of a shock, as described by Kompaneets and others [161, 162].

This provides strong evidence that the influence of the pre-pulse is predominantly hydrodynamic. Moreover, it is clear that the interferometry provides an indication of pre-pulse strength via the size and scale length of the high density cavity. Figure 7.5 shows the density profile within the shocked region according to FLASH simulations, alongside a comparison with the Taylor approximation described by Equation 7.16 for the equivalent axial expansion length r_{sh}.
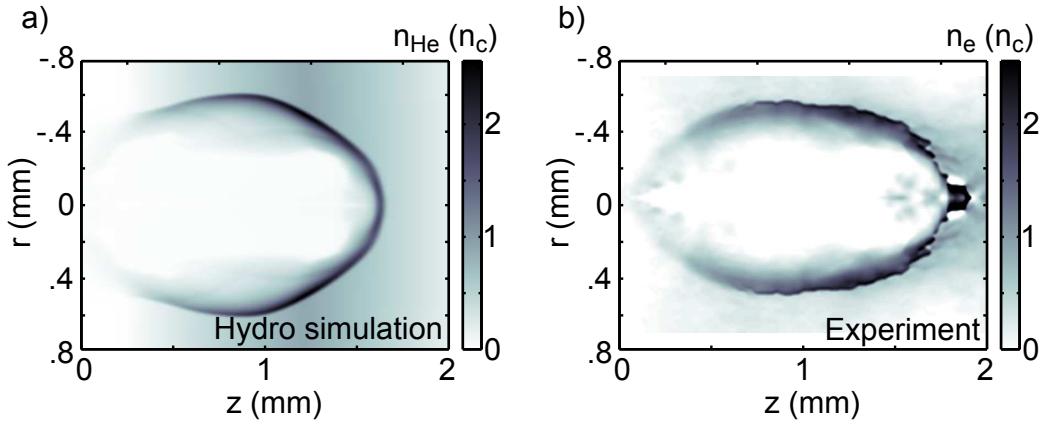
Figure 7.4: (a) Total helium particle density calculated from 2D FLASH code simulating hydrodynamic expansion of gas for 25 ns. (b) Experimentally extracted plasma density profile from interferometry 300 ps after the main pulse arrives reveals the shape of the entire blast wave profile. The initial gas profile is centered at $z = 1.5$ mm. Density profiles are given in units of the critical plasma density n_c for 10.3 μm light.
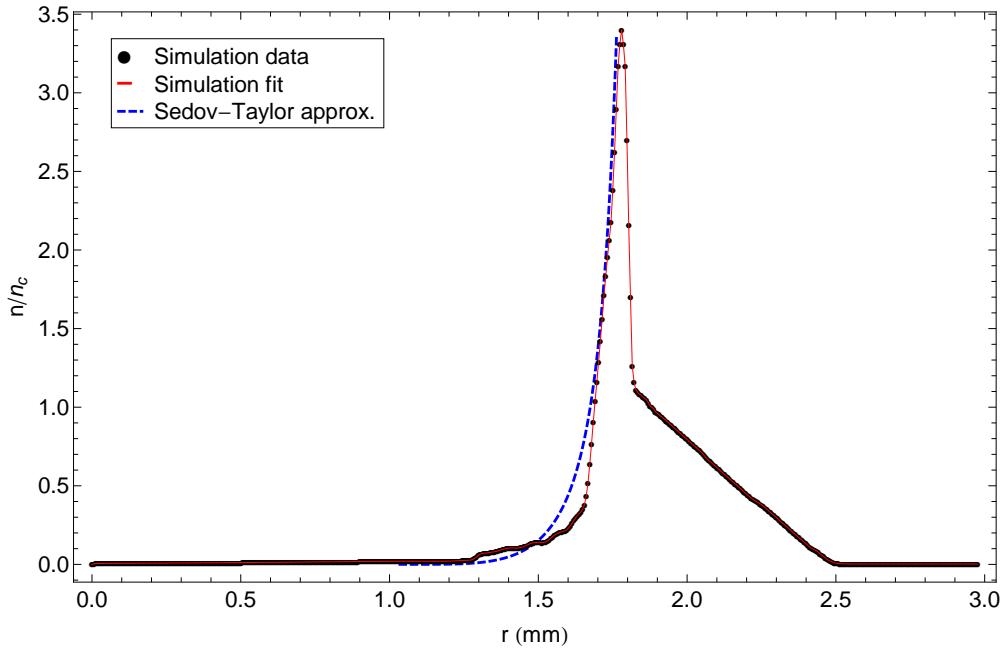
Figure 7.5: A comparison of the axial profile of the shocked region between FLASH simulations and the Taylor approximation, assuming consistent shock positions.

7.2 Helium ion acceleration

We now present some of the results taken with a helium gas jet target. For the majority of experimental trials, the laser was focused at the center of the 1 mm jet, approximately 800 μm above the nozzle. Variation from this focal position produced considerably fewer instances of successful ion detection. Figure 7.6 (a) shows the range of pre-pulse energy and main pulse intensities in which accelerated helium ions are observed. It is immediately apparent that an accelerated ion beam is observed only for a narrow range of pre-pulse energies, from approximately 110 mJ to 200 mJ, and for main pulse normalized intensities of greater than $a_0 = 1.2$. Each shot was taken with the same relative laser focal position within the gas jet, at the center of nozzle. The dashed line shows the testable limits of the pre-pulse, main pulse energy parameter space given by empirically observed laser performance. Scanning the entire space is unrealistic, because the pre-pulse and main pulse energies are weakly coupled by virtue of the self-induced reflectance method of pre-pulse generation.
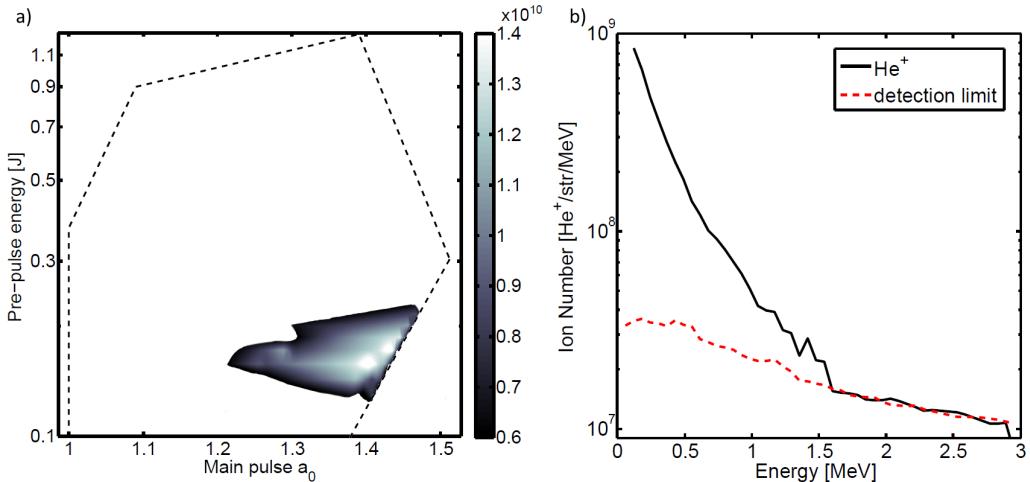
Figure 7.6: Left, average flux of experimentally observed helium ions as a function of incident pre-pulse energy and main pulse normalized intensity. The dashed black line separates the experimentally tested parameter space. Right, characteristic helium spectra for $a_0 = 1.5$ and $E_{pp} \approx 150$ mJ are shown in black. The spectra demonstrates a detection edge at approximately 1.5 MeV, above which the signal remains consistent with the 3σ level of detector background, shown by the dashed red line.

Figure 7.6 (b) shows a typical helium spectrum taken with $a_0 = 1.4$ and $E_{pp} \approx 150$ mJ. The spectrum is consistent with all shots taken, exhibiting large energy spreads (100%) and peak observed energies between 1 and 2 MeV. The

total flux of helium ions remains consistent from shot to shot, ranging from $0.8 - 1.4 \times 10^{10}$ ions/sr, corresponding to a total particle yield of 1–2 $\times 10^5$ ions per shot. The scintillator yield for helium is extrapolated from the proton calibrations by using the mean scaling factor 0.3 ± 0.1 (helium yield: proton yield), which has been experimentally observed in comparative studies with differing ion types [163]. The incoming flux of ions corresponds to a very narrow localized emittance of $\epsilon < 0.5$ mm-mrad (normalized emittance of $\epsilon_n \approx 5$ nm-rad), a factor of 10 better than most RF sources. A more proper treatment of the emittance would require multiple shots with better energy resolution.

Unfortunately, the total number of ions in the high end of the energy spectrum lies just at the detection limit for our diagnostic, and the high energy tail of the spectrum is poorly resolved. The detection limit for the ion signal was determined by the three sigma threshold above the mean background signal at each discrete energy step. The dispersive properties of the Thomson parabola spectrometer limit the dynamic range of our system due to background signal at high energies. This is mostly the result of noise driven by background light, camera readout noise, and scattering within the scintillator. For future experiments, adjustments to the spectrometer positioning will be made for greater separation at high energies.

7.2.1 Pre-pulse regimes

The remarkable feature of these results is the sensitivity of the acceleration process to the pre-pulse energy. The reproducibility of accelerated ion beams for a select range of pre-pulse energies greatly exceeded results from a 2011 experiment using a pulse train output from a non-isotopic CO_2 amplifier [6, 164]. We have defined three regimes of pre-pulse energies — low, mid, and high — each of which can be associated with characteristic density profiles and laser interaction properties.

High pre-pulse regime

The high pre-pulse regime is identified by $E_{pp} > 200$ mJ, for which no ion acceleration was observed. Figure 7.7 shows the plasma density before the main pulse arrival, and 300 ps afterwards for a shot with $E_{pp} \approx 1$ J.

In this instance, the energetic pre-pulse causes rapid expansion of a blast wave into the gas jet. By the time the main pulse arrives, the blast wave has driven most of the residual gas away from the laser axis. The remaining gas has very low density, and the laser channels through the entire target, unable to generate the critical surface necessary for ion reflection. This further supports

150
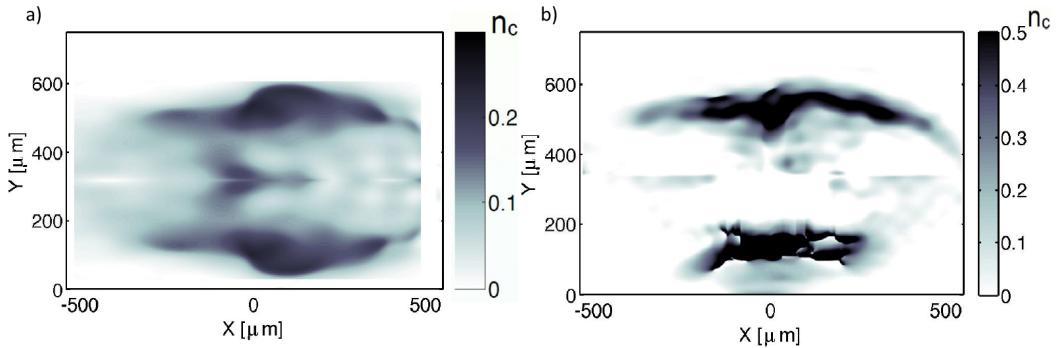
Figure 7.7: Density extracted from interferograms taken (a) immediately before the main pulse and (b) 300 ps after the main pulse, for a shot with $E_{pp} \approx 1\,\mathrm{J}$ and main pulse normalized intensity $a_0 = 1.4$.

the conclusion that the pre-pulse evacuates gas from the laser axis, as no coinciding increase in plasma density was seen along the low density channel. Even at pre-pulse energies closer to 200 mJ, optical probing shows considerable expansion of the blast wave beyond the focal point of the laser, and no ion acceleration is observed. This suggests that the main pulse will not efficiently couple energy to the overdense region of the plasma at focus.

Low pre-pulse regime

The low pre-pulse regime is defined by $E_{pp} < 100$ mJ, for which no ion acceleration was observed. Figure 7.8 shows the plasma density at 300 ps after the main pulse arrival for a shot with $E_{pp} \approx 0$ mJ. As the pre-pulse energy is quite low, very little plasma is generated prior to the arrival of the main pulse. As a result, no phase map could be extracted from the first interferogram taken, so the pre-pulse effects must be inferred from a single image.

The main pulse energy is well above the self-focussing threshold $P_c = 17.5\ n_e/n_c$ GW. There is some evidence of self-focussing in the early portion of the jet. The corresponding approximate focal length for a laser intensity of $a_0 = 1.2$ is $z_c \approx 230\ \mu$m, in good agreement with the length of the channel observed in the interferometry. The action of the main pulse is to first ionize the gas and deplete electrons from the high intensity region of the laser. This results in the formation of a narrow channel with sharp edges parallel to the laser axis. However, the radius of the generated cavity is only about $\sim 50\ \mu$m, making it clear that in absence of a pre-pulse, the laser deposits much of its energy in a short region of underdense gas at the foot of the jet. This interaction is commensurate with published results from past experiments, in which the formation of density perturbations transverse to the laser axis led to
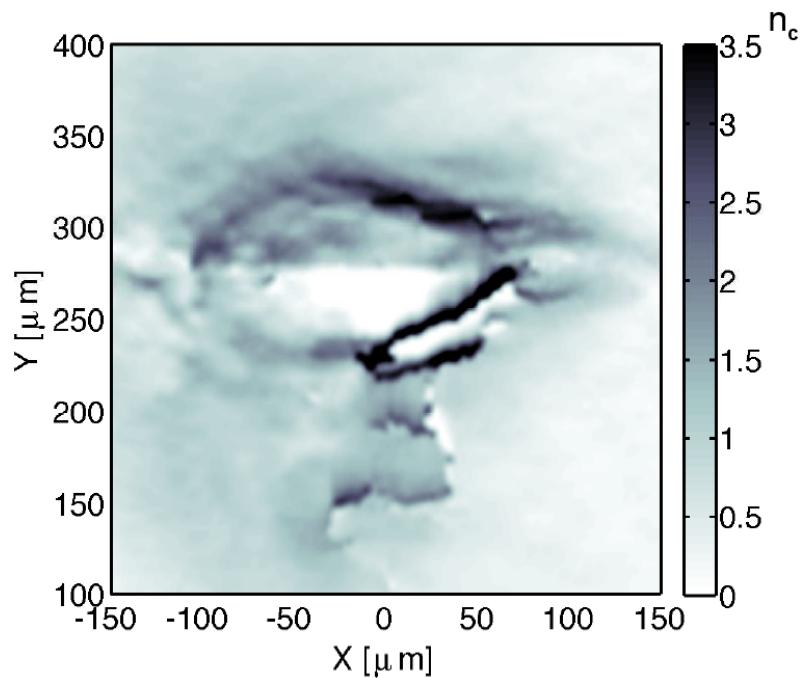
Figure 7.8: Density calculated from interferograms taken 300 ps after the main pulse for a shot with $E_{pp} \approx 0\,\text{mJ}$ and main pulse normalized intensity $a_0 = 1.2$. The ionized density prior to the main pulse is noise-dominated, preventing phase-extraction and density calculation.

the observation of expansion shocks perpendicular to laser propagation [165].
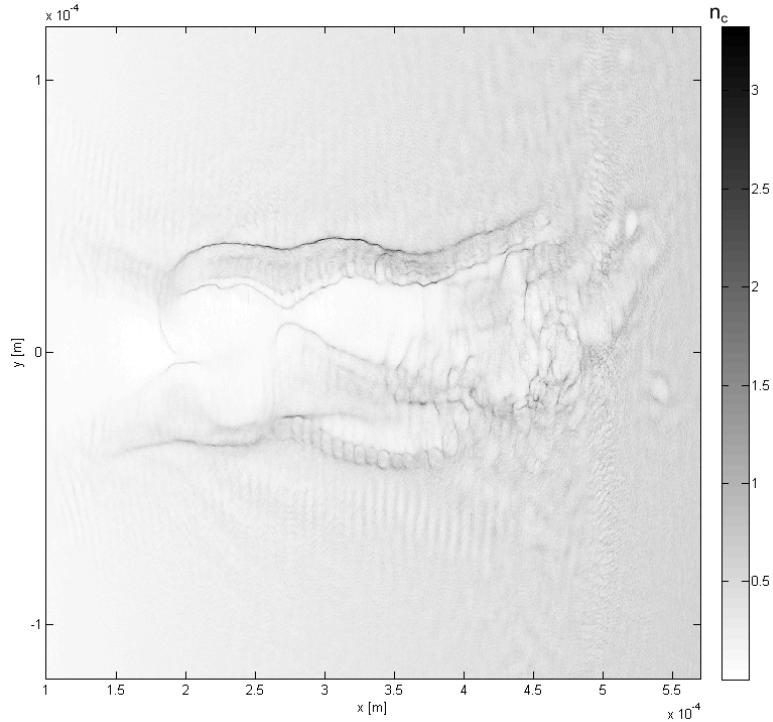
Figure 7.9: (a) Density profile extracted from 2D PIC simulations without pre-pulse shaping, taken at 12 ps. The laser penetrates only 400 μm into the jet, dispersing in the underdense region through a combination of self-focussing, ponderomotive channeling, and electron heating. No considerable forward ion acceleration is seen.

This behavior is confirmed by 2D particle-in-cell simulations. Simulations performed using the native density profile of the gas jet, with no pre-pulse shaping, demonstrate the deposition of laser pulse energy over a 200–300 μm region of underdense plasma. Channel formation on the order of the laser FWHM is seen in Figure 7.9, with sharp density perturbations of $n_{He} = 2-3n_c$ forming at the edge of the channel, transverse to the laser axis. Ion phase space distributions reveal trace ion acceleration to energies below 0.1 MeV, which would be unlikely to be observed.

Mid pre-pulse regime

The mid pre-pulse regime is defined by $110\,\text{mJ} < \text{E}_{pp} < 200\,\text{mJ}$, for which ion acceleration was consistently observed. Figure 7.10 shows the plasma density

before the main pulse arrival, and 300 ps afterwards for a shot with $E_{pp} \approx$ 150 mJ.
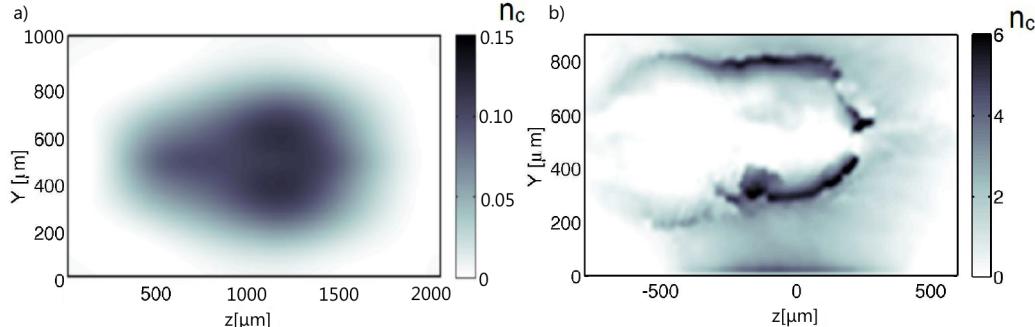

Figure 7.10: Density extracted from interferograms taken (a) immediately before the main pulse and (b) 300 ps after the main pulse, for a shot with $E_{pp} \approx 150 \, \text{mJ}$ and main pulse normalized intensity $a_0 = 1.4$.

The sharp high-density feature seen along the laser axis after the main pulse interaction indicates that a high density gradient necessary to drive an electrostatic shock was present at the time of the main pulse arrival, given the slower time scale of hydrodynamic expansion. As the Abel inversion is extremely sensitive at small distances from its axis of symmetry, a simulated interferogram was constructed to confirm the experimental profile. A three dimensional density profile was created to best match the phase advance required to generate the observed interferogram. Figure 7.11 shows the results of the exercise. The interferogram closely matches the observed image, and the corresponding density profile has a clear on-axis peak at approximately $n = 6.5 \, n_c$. Again, this density corresponds to a fully ionized helium profile with a peak neutral density of $4n_{i0} \approx 3.2n_c$. Within the cavity, the profile steepens over a 100 μm length, while the plasma profile ahead of the shock declines over a longer scale length due to electron heating during the laser plasma interaction. Transverse variations in density are seen near the laser axis as well, providing evidence of electron beam filamentation, which is discussed below.

7.3   Particle-in-cell simulations

While optical probing provides considerable insight into the plasma behavior, temporal resolution is limited by the duration of the probe pulse. The YAG probe arrives at the target with a 10 ps duration after being up-converted to
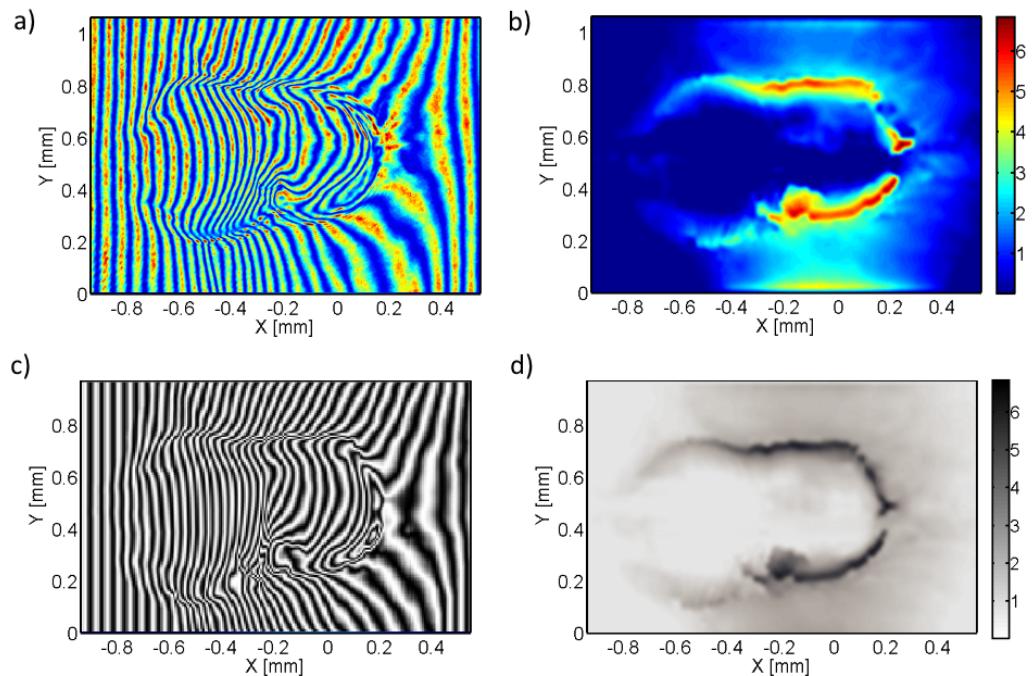
Figure 7.11: A manually constructed interferogram for a shot with $E_p p \approx$ 150 mJ taken 300 ps after the main pulse interaction.Above, (a) and (b) show the experimental interferograms and extracted density profiles, respectively, while (c) and (d) show the corresponding simulated profiles.

532 nm wavelengths. As a result, probing within the timescale of the 5ps LPI is impossible. Thus simulations must be used to gain further insight.

A series of PIC simulations were performed using the open source code EPOCH to investigate the LPI on the timescale of the laser pulse. EPOCH is a multi-dimensional, fully electromagnetic, relativistic code which was developed by EPSRC-funded Collaborative Computation Plasma Physics consortium of 30 UK researchers. The EPOCH code was developed under UK Engineering and Physics Sciences Research Council grants EP/G054940/1, EP/G055165/1 and EP/G056803/1 [166]. The simulations were performed in 2D with 30 particles per cell and square cells of $\lambda/50$ length. The plasma consisted of fully ionized He^{2+} and e^- populations, and was initialized cold with $T_e = T_i = 0\,\mathrm{eV}$. The laser is modelled as a Gaussian with 5 ps FWHM; the pulse reaches its maximum intensity at 7 ps and has a 14 ps duration, equivalent to a reduction in intensity of more than $3\,\sigma$ from the maximum before termination.
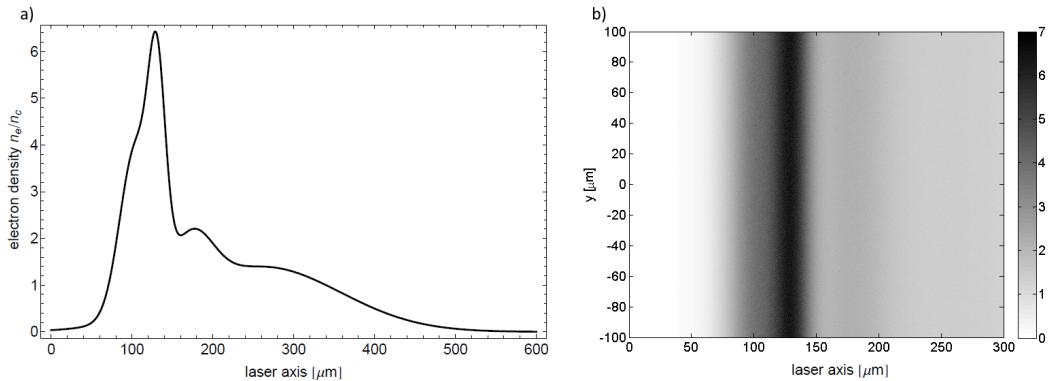
Figure 7.12: 2D PIC sims are initialized with electron and ion densities corresponding to experimentally extracted lineous taken across the blast wave profile slightly off axis from the drive laser.

The density profile and laser energy were modelled to fit the profile seen from a shot taken with 150 mJ pre-pulse and main pulse intensity $a_0 = 1.4$. The density profile was obtained from a lineout taken from the Abel inversion shown in Figure 7.10. The use of an "on-axis" profile is discouraged both by the presence of electron filamentation and motion of the plasma due to the main laser pulse. Instead, the density profile is taken slightly off-axis, still normal to the blast wave, and fit with a sum of Gaussians. The resulting profile and its implementation in the simulations are shown in Figure 7.12. Variations in y-axis are ignored for the following series of simulations.

Specifically consider the case of a laser incident with normalized intensity of $a_0 = 1.4$ on the profile described, falling squarely within the mid pre-pulse regime for which ion acceleration was most frequently observed. Other

extracted profiles have been simulated, as well as minor variations to this profile, all with similar results. A sequence of timesteps are presented to illustrate the shock dynamics during its three phases - formation, propagation, and breakdown.

Shock formation

Figure 7.13 shows the state of the plasma at the peak of laser intensity, corresponding to $t = 7$ ps in our simulation. The two dimensional ion density reveals the extent of laser propagation into the plasma, with a high density wall having formed at the head of the pulse. Examination of the electron density and corresponding electric field along the laser axis reveals the formation of a critical surface corresponding to the transition from underdense to opaque plasma. The field has created a sharp ramp in electric potential, capable of reflecting the cold downstream ions. The ion phase space reveals the onset of reflection from the leading edge of the shock, corresponding to acceleration of downstream ions to energies in excess of 2 MeV. The formation of the shock at 7 ps is in good agreement with published simulations predicting formation of a collisionless shock on the timescale of a few multiples of $2\pi/\omega_{pi}$, where ω_{pi} is the ion plasma frequency [97]. If we assume $n_i \approx 4n_c$, then $2\pi/\omega_{pi} \approx 1.5$ ps, and the observed formation of a shock at 7 ps is consistent.

Behind the shock, the drop in potential along with modulation from the still incident laser field has produced acceleration of a separate group of ions. The presence of this ion population is characteristic of long pulse interactions, for which $\tau_l > 1/\omega_{pi}$. If the shock formation is weak enough, a portion of this ion population will not be trapped, instead following an unstable orbit in phase space, eventually passing the shock and entering the downstream plasma. The specific dynamics of this process are not yet well understood; for the experimental parameters considered here, this secondary bunch is clearly an impediment to generating a cleanly reflected bunch. However the total flux of this bunch is small relative to the shock accelerated population.

Shock propagation

Once the shock has formed, reflection continues uninterrupted, with the shock propagating in absence of a strong ponderomotive push from the laser. Figure 7.14 reveals the dynamics at $t = 12$ ps for the simulation. The shock has advanced ahead of the critical surface, travelling further into the density ramp of the target. The shock exhibits a clear variation in potential behind the front, which is in line with the significant trapping seen behind the shock.

It is clear from the ion phase space that the momentum to which ions are
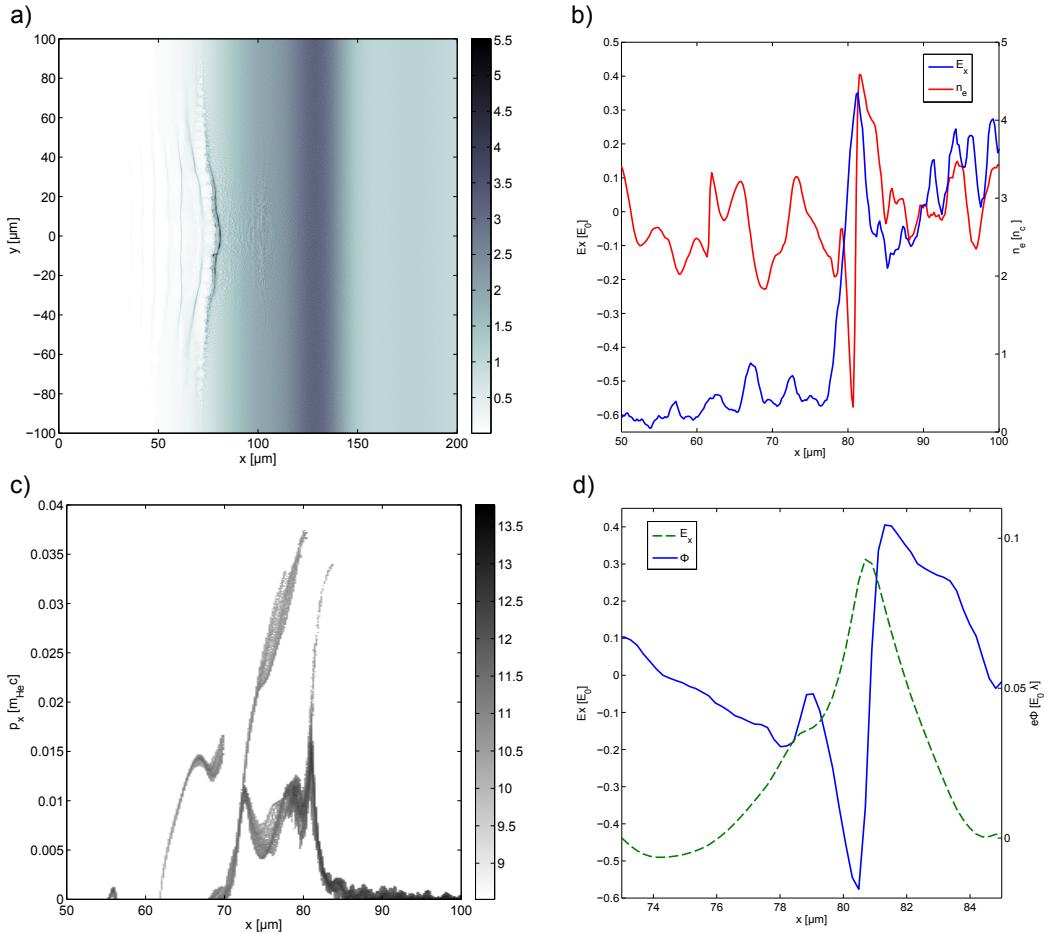
Figure 7.13: Snapshots of the plasma state at $t = 7\,\mathrm{ps}$ during a simulation with experimental laser and target parameters: (a) the ion density n_i, (b) the electric field E_x in units of $E_0 = m_e\omega_l c/q$ and electron density n_e in units of n_c, (c) the ion phase space $x - p_x$ with momentum in units of $m_{He}c$, and (d) the electrostatic potential Φ in units of $E_0\lambda_l$. The ion phase space is extracted from a region spanning ± 5 microns about the $y = 0$ axis in the simulation grid. Electric field, electron density, and electrostatic potential are averaged over a the same region in $\hat{y}$.
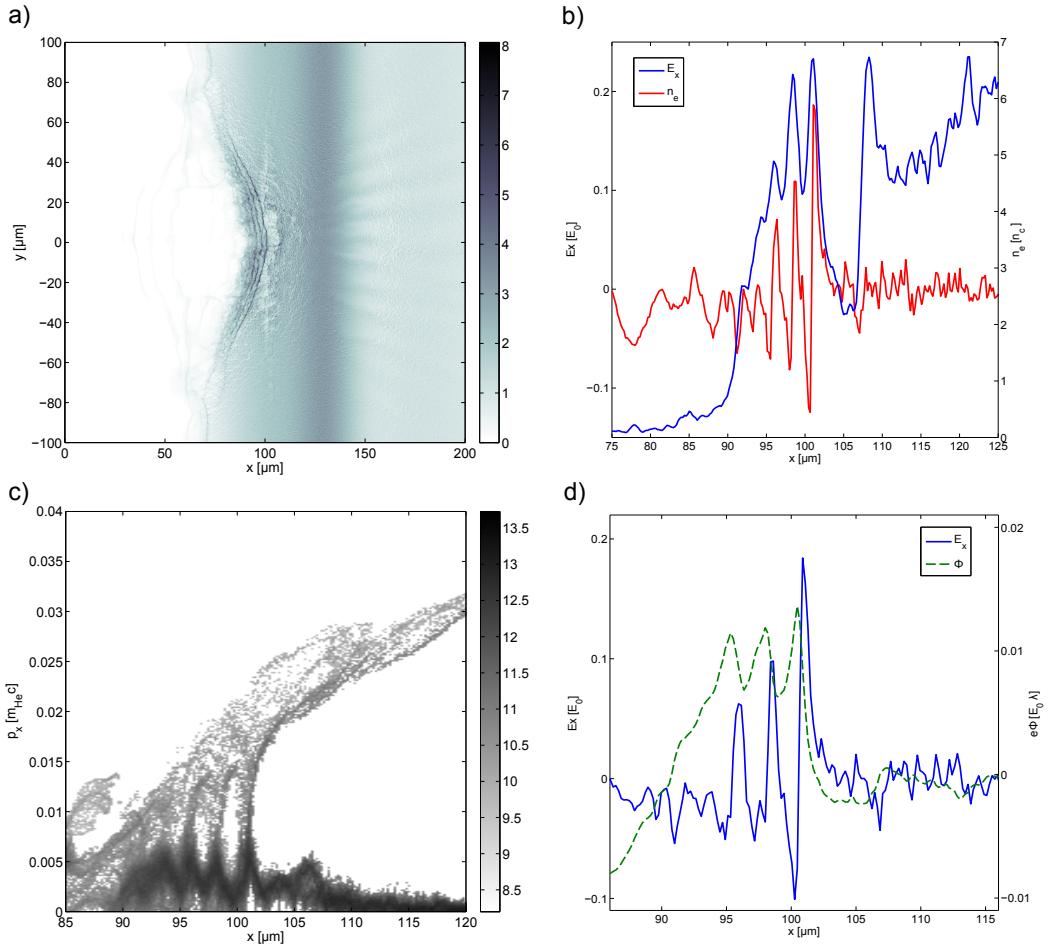
Figure 7.14: Snapshots of the plasma state at $t = 12\,\mathrm{ps}$ revealing the propagation of the shock and high reflection: (a) the ion density n_i, (b) the electric field E_x in units of $E_0 = m_e \omega_l c / q$ and electron density n_e in units of n_c, (c) the ion phase space $x - p_x$ with momentum in units of $m_{He}c$, and (d) the electrostatic potential Φ in units of $E_0 \lambda_l$. The ion phase space is extracted from a region spanning ± 5 microns about the $y = 0$ axis in the simulation grid. Electric field, electron density, and electrostatic potential are averaged over a the same region in $\hat{y}$.

reflected later in the shock evolution are considerably lower, and the resulting spectra has already spread, with ions reflected at the 12 ps time step only reflected to approximately 0.5 MeV. As we expect ions in our relatively cold plasma to be reflected to a velocity of $2v_{sh}$, it follows that the shock must have slowed. To test this hypothesis, the recessional velocity of the critical surface was tracked, along with the motion of the shock into the plasma.

The position of the critical surface and shock front are plotted in Figure 7.15. The shock quickly separates from the critical surface, travelling with initial speed $v_{sh} = 5.07$ μm/ps. In time, the shock velocity slows; the reduction in shock velocity is in good agreement with the ion phase space at a given time. For example, at the time of formation, ion reflection to $2v_{sh}$ predicts momenta of $0.034c$ (2.2 MeV), coinciding with what is seen in Figure 7.13. Likewise, the shock velocity at 12 ps is calculated to be $v_{sh} = 2.5$ μm/ps, corresponding to reflected momenta of $0.017m_{He}c$, again in good agreement with the phase space presented in Figure 7.14.

Conversely, the trailing critical surface is seen to move with near constant speed $v_{cs} = 3.43\mu$m/ps through the central portion of the pulse envelope, between 4 and 10 ps. This speed agrees with the hole-boring model described by Equation 4.73; for a laser of amplitude $a_0 = 1.4$ incident on $n_e = n_c$, $v_{hb} = 3.5\mu$m/ps [88]. Moreover, the critical surface is seen to effectively stop beyond approximately 12 ps, as the incident laser intensity has been reduced to less than 1% of its peak value, and terminates completely by $t = 15$ ps.

7.3.1 Shock breakdown

The duration of shock reflection depends heavily on the target plasma conditions, as the degree of reflection and dissipation in the shock are strongly influenced by plasma temperature and electron motion. For the simulation parameters under inspection, the shock shows clear signs of deterioration at $t = 18$ ps. Figure 7.16 shows some corresponding plasma parameters. The density profile shows a strong damping of the shock density perturbation as it enters the peak of the shaped target density. Heating and expansion of the target density profile can be clearly seen across the 50 μm segment of electron density lineout in Figure 7.16 (b). The corresponding shock potential has weakened significantly.

Figure 7.17 shows the progression of shock potential and shock velocity. A clear decline in shock potential along with a corresponding increase in ion temperature ahead of the shock agree with the observed behavior. At early times, the shock clearly satisfies the reflection conditions $Ze\Phi > 1/2mv_i^2$. Beyond 14 ps, the shock has lost considerable energy, and downstream ion kinetic energy begins to exceeds the shock potential. Moreover, the heated

160
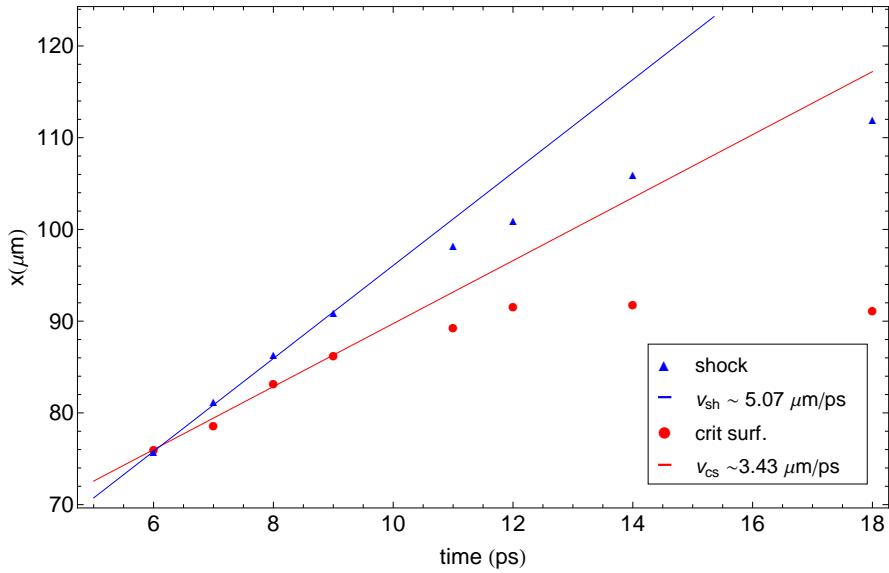
Figure 7.15: Evolution of critical surface and shock positions are plotted versus time from 2D PIC simulations. The shock forms approximately at the critical surface, and quickly separates from it, travelling at $v_{sh} \approx 5 \ \mu$m/ps, while the critical surface recedes at $v_{cs} \approx 3.43 \ \mu$m/ps.

electron population in the neighborhood of the shock can also overcome the potential. Thus, thermal motion of the plasma begins to dominate over that of the shock. Inspection of the electron density profile (Figure 7.16 (b)) shows a clear washing out of the previously trapped electron population behind the shock.

To summarize the shock evolution as described above, Figure 7.18 presents a series of three illustrations of the ion phase space, each corresponding to a separate phase of the shock. Each diagram emphasizes the defining characteristics of ion behavior associated with the shock properties. The onset of the shock, shown in part (a), is characterized by a localized reflection to twice the shock velocity, in agreement with Figure7.13 (c). As the shock is only just forming, the associated particle trapping has not fully developed. As the shock propagates, its phase space may be characterized by part (b) of the figure. A well defined shock front reflects ions to twice the instantaneous shock velocity, while well defined trapping is seen behind the shock. In part (c), the shock potential has fallen, and the trapping behind the shock dissipates, freeing the trapped ions to drift in phase space. This drift results in a shearing of the phase space profile, and can be seen across the entire ion population.
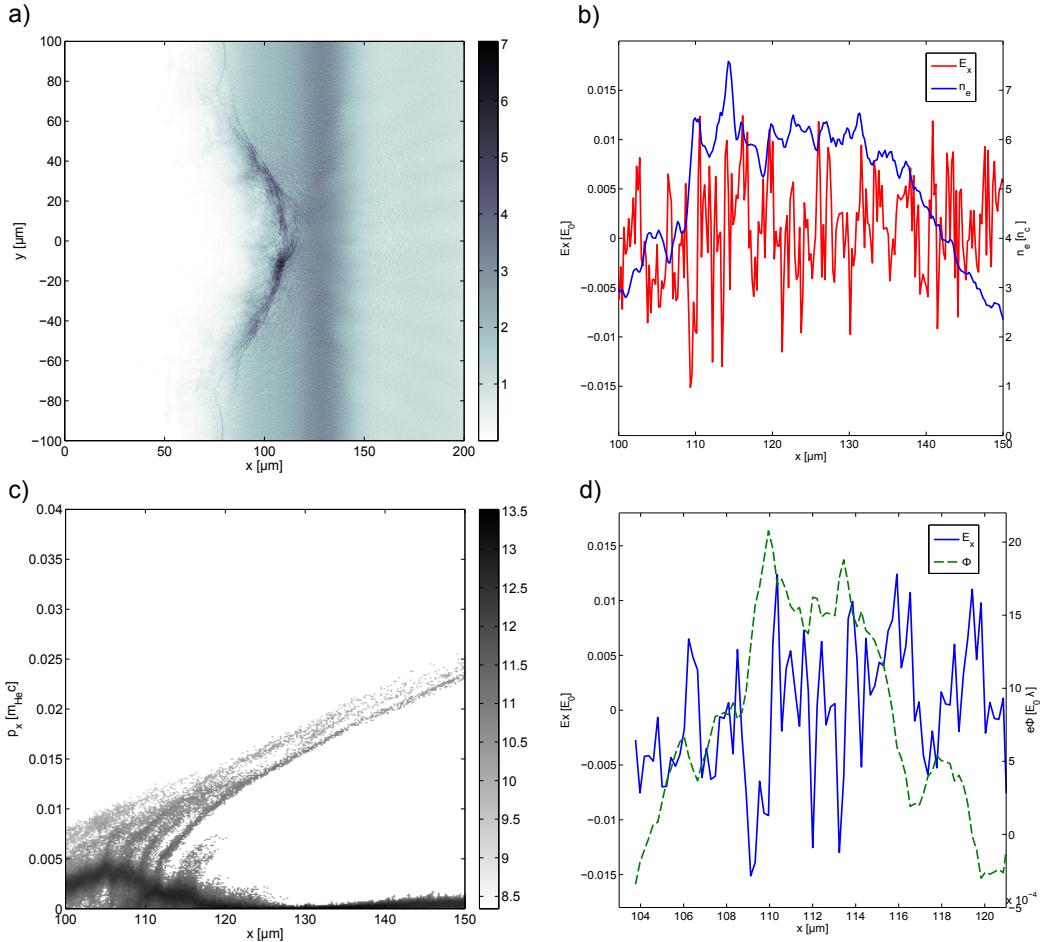
Figure 7.16: Snapshots of the plasma state at $t = 18 \, \mathrm{ps}$ during which significant deterioration of the plasma is observed: (a) the ion density n_i, (b) the electric field E_x in units of $E_0 = m_e \omega_l c / q$ and electron density n_e in units of n_c, (c) the ion phase space $x - p_x$ with momentum in units of $m_{He}c$, and (d) the electrostatic potential Φ in units of $E_0 \lambda_l$. The ion phase space is extracted from a region spanning ± 5 microns about the $y = 0$ axis in the simulation grid. Electric field, electron density, and electrostatic potential are averaged over a the same region in $\hat{y}$.
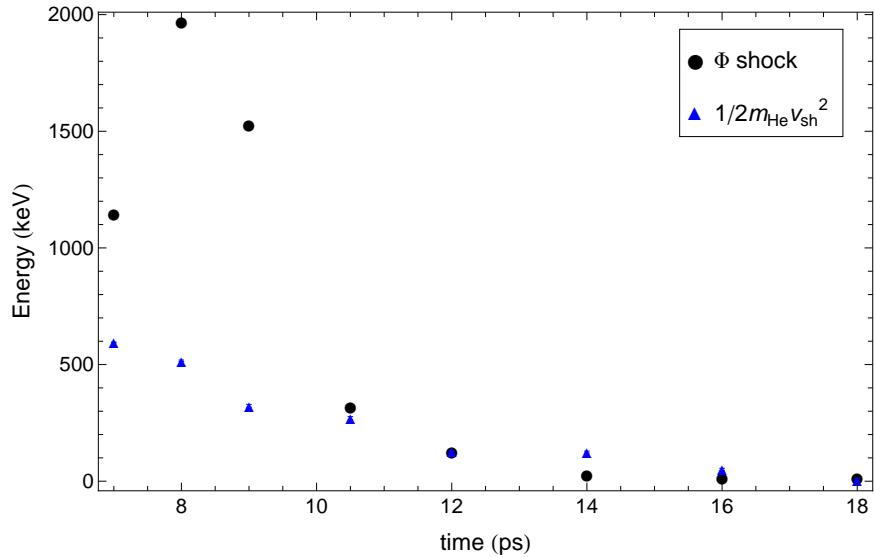
Figure 7.17: Evolution of shock potential and corresponding upstream ion kinetic energy (in the frame of the shock). The potential initially clearly satisfies the reflection condition $Ze\Phi > \frac{1}{2}mv_i^2$, but after several ps falls sharply.
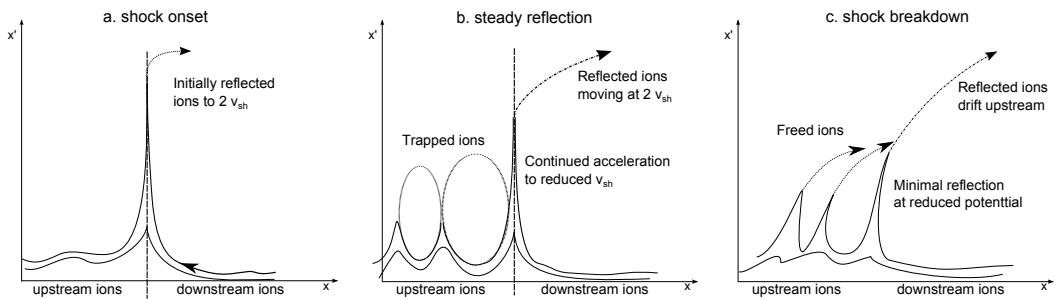
Figure 7.18: An illustration of the ion phase space throughout the three phases of shock evolution characterizes the main features of the shock. In (a), the shock has just formed, reflecting ions to twice the shock velocity. After some time (b), ions continue to be reflected, while particles behind the shock potential are trapped in the trailing potential. Finally, as the shock potential drops below the incoming ion kinetic energy, reflection is suppressed (c), and the phase space exhibits shearing associated with the drift of the plasma particles.

Axial variations in density modelling

The parameter space defined by the possible target and laser configurations is extremely large, and a thorough investigation through simulations would require many more processor hours than are available. Thus, variations in the density along the jet axis are ignored for the majority of simulations and analysis. The deposition of laser energy is also certainly asymmetric, and the resulting perturbations in gas dynamics reduce the reproducibility of experimental results. The use of a cylinder-like deposition of energy in FLASH allows for some investigation of radial variations when modelling the target density profile. A set of 2D-PIC simulations have been performed with radial variations included, and produce effectively identical shock behavior. Under these conditions, the simulated interaction exhibits stronger radial confinement, and the resulting shock does produce density perturbations on the micron scale about the laser axis which are well supported by shadowgraphy. Figure 7.19 shows a comparison between the FLASH produced simulation output and a shadowgraphy frame taken 300 ps after the LPI. The on-axis bowing of the density profile is distinct, as is the presence of two distinct density perturbations slightly off axis. The forward most perturbation could represent the residual shock structure following the long term expansion.
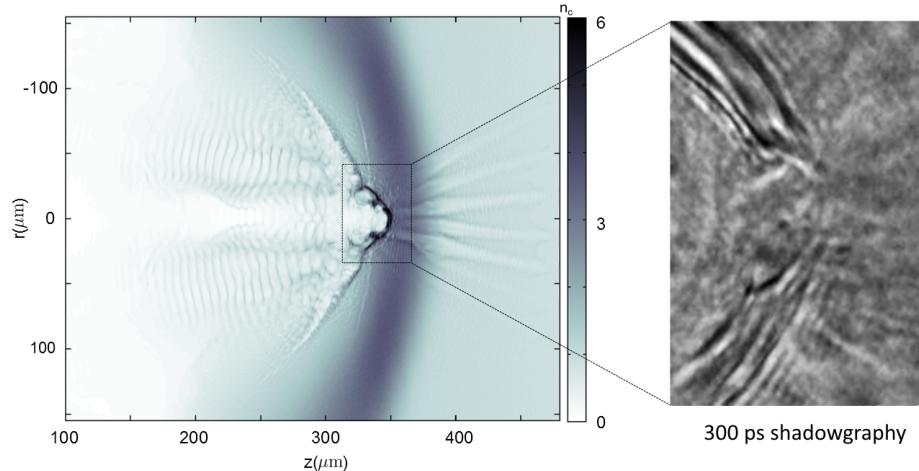
Figure 7.19: Helium ion density from a 2D PIC simulation using a radially adjusted density profile obtained from FLASH simulations shows strong radial confinement of the shock along the laser axis. The corresponding bowing of the density profile is observed in shadowgraphy images taken 300 ps after an LPI.

Electron heating and filamentation

The laser interaction is expected to produce hot electrons with ponderomotive scaling, according to Equation 4.66. Thus, for $a_0 = 1.4$, one might expect to see hot electron generation to energies on the order of $\epsilon = m_e c^2$, in the hundreds of keV range. We can estimate the effective fast electron beam current with a simple absorbtion metric. Assume η is the fraction of laser energy absorbed by a population of fast electrons travelling with speed v_h, and with energy ϵ, the density of fast electrons follows from

$$\eta I_l = n_b v_h \epsilon \tag{7.20}$$

$$n_b = \frac{\eta I_L}{v_h \epsilon} \tag{7.21}$$

If we assume hot electron populations of 300 keV and $\eta = 0.4$ ([70]), then the resulting electron density is $n_b \sim 4 \times 10^{18}$ cm^{-3}. If this beam current lies within an area of πw_0^2, where w_0 is the beam waist, then we obtain currents in excess of 1 MA! This current is greatly exceeds the Alfvèn limit which follows from $I_A \approx 17\beta\gamma$ kA = 20 kA for 300 keV electrons [78]. This necessitates the formation of counterpropagating return current, and forms the basis for a number of possible plasma instabilities. In the particular instance of large velocity distributions, magnetic field perturbations may give rise to the formation of electron beam filaments. The growth of these filaments can be attributed to the Weibel instability, particularly in high density plasmas for which laser filamentation is not present [167].

The growth rate of this instability can be approximated by [76]

$$\gamma_w = \omega_p \left(\frac{n_b}{\gamma_p n_e} \right)^{1/2} \tag{7.22}$$

where n_b is the hot electron beam density. For our model calculations above, we can assume $n_e \sim 4n_c$, in line with early simulation results (see Figure 7.13 (b)), and we have $n_b \sim 0.05$. We would then expect growth periods of $\frac{1}{\gamma_W} < 10$ fs. Thus, after some picoseconds of heating by the laser, filamentation is expected. In fact, shadowgraphy clearly identifies filamentation across an array of different shots; Figure 7.20 (a) shows one such instance. Simulations confirm the formation of filaments on ps timescales, as seen in Figure 7.14. The magnetic field extracted from simulations shows a growing quasi-static mode ahead of the shock in agreement with the formation of filaments. Moreover, the filamentation is not seen for low pre-pulse shots in which the laser does not see a strong overdense plasma - Figure 7.20 (b) shows the lack of filamentation

for a shot with no pre-pulse.
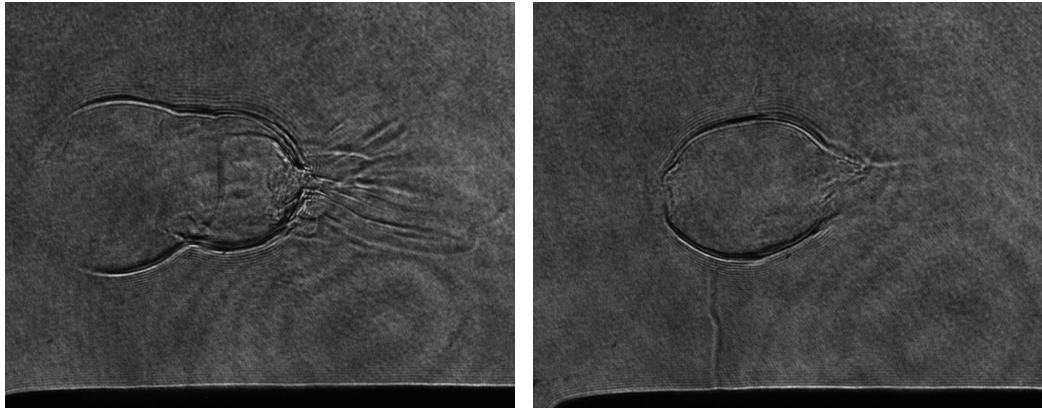
Figure 7.20: Left shadowgraphy reveals filamentation structures 300 ps after the LPI in the instance of a shaped high density plasma. Right, few filaments are seen in conjunction with low intensity pre-pulses which do not generate strong axial target shaping.

The fast electron temperature is estimated from momentum distributions extracted from simulations. The resulting distribution is non-Maxwellian, with peak forward electron energies reaching a few MeV. A rough estimate of this hot electron tail in the distribution places the electron temperature at approximately 300 keV around the peak of laser intensity. This temperature is indicative of the scaling suggested by Beg et al. [74], for which hot electron temperatures were empirically fitted to follow Equation 4.67 rather than a ponderomotive scaling, and has been observed in other experiments for similar laser irradiance (10^{18} W/λ^2cm^2) [168]. A more complete approximation of the temperature would require greater resolution in the output momentum space, and correspondingly smaller output timesteps to allow averaging over a few plasma periods. Regardless, it is clear that the bulk electron temperature remains nonrelativistic, indicating that one may not expect many of the theories concerning shocks in "hot relativistic plasmas" to apply directly [7, 169]. Study of the role of electron filamentation on the shock propagation remains ongoing. While other investigations have shown that density variations can affect the spatial extent of the filaments [170], it remains unclear how strongly they influence ion reflection and shock propagation at early timescales. Filamentation may be responsible for mediating shock formation [171], but the laser energy threshold for this mechanism remains greater than the peak energies achievable at ATF .

Scaling with laser intensity

A series of simulations has also been performed to investigate the scaling of ion energies and spectra with laser intensity. The same interferogram-derived density profile as in Figure 7.12 is used, with the density scaled according to Equation 4.58 to adjust for relativistic transparency effects. Figure 7.21 demonstrates the scaling of output spectra with laser intensity, as well as a comparison to modelled hole-boring velocity. Shock formation and reflection is seen in each case; the shock propagates ahead of the critical surface, reflecting ions to energies greater than predicted by a simple hole-boring model. As the shock continues to propagate, the shock slows and eventually breaks down. Peak energies clearly scale linearly with a_0, while the resulting ion spectra exhibit the same large spread in energies observed during experiment. It is clear that although this scaling is favorable in comparison to $\epsilon \sim \sqrt{a_0}$ scaling seen in TNSA approaches [82, 172], further adjustments to the target structure are needed to achieve a small energy variation in the output beam.
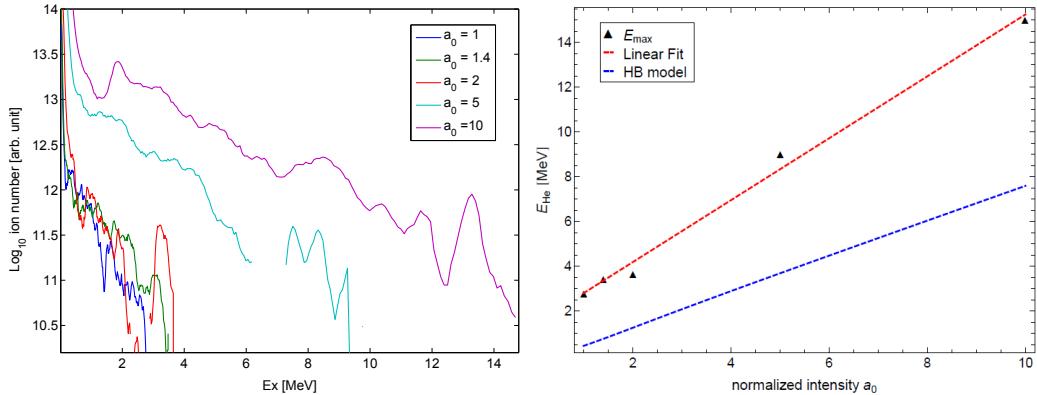
Figure 7.21: Left, 2D PIC simulations show consistent output spectra with increasing laser intensity (a_0), incident on targets with the same density profile. At right, peak energies scale linearly with laser intensity, and clearly exceed the hole-boring model predictions, assuming 50% laser absorption.

7.4 Next steps

Experiment and simulation analysis make clear the difficulty in generating ideal output beams with the tested target structures. Even with increases in laser a_0, the generated ion spectra exhibit similarly large spreads in energy. Simulations, as shown above, clearly predict the formation of a shock on scale lengths shorter than the steepest density profiles we have extracted during

experiment. It is the goal of future experiments to emphasize target dynamics to achieve beam optimization. Amongst the various target parameters, density scale length is of the utmost importance.

Recent theoretical and computational investigations of laser driven CES theory [7] suggest that for an isothermal plasma at relativistic temperatures, the ideal target length scales as

$$L \propto \frac{\lambda}{2} \left(\frac{m_i}{m_e} \right)^{1/2} \tag{7.23}$$

For helium gas, Equation 7.23 suggests that targets on the order of 0.4 mm would be ideal for generating a hot isothermal electron population. This is more than a factor of two smaller than the hydrodynamically shaped profiles that have been discussed above. One possible solution to achieving a target of this length is through the use of a knife edge to shear the gas flow.
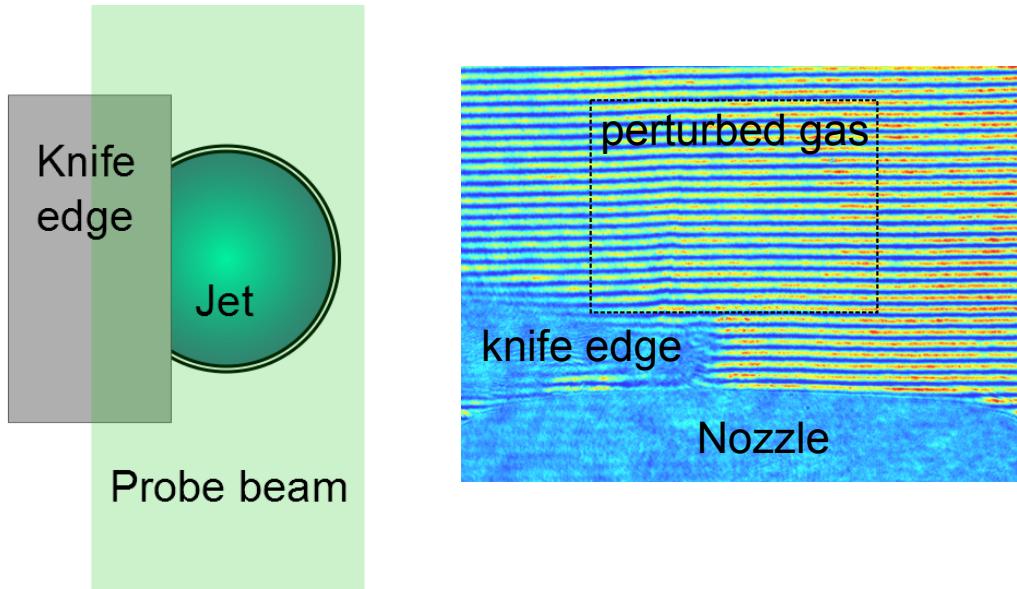
Figure 7.22: Left, the basic set up for using a thin blade to create a shock in gas density profile along the laser axis. Right, interferometry of the configuration illustrating the longitudinal arrangement of the blade and nozzle.

A fixed impediment has the advantage of reliability and serves to further decouple the pre-pulse energies from that of the main pulse. The knife edge can be aligned with the nozzle of the jet so as to create an oblique shock stemming from the edge, steepening the neutral density profile along the laser axis. Figure 7.22 illustrates a top down view of the arrangement - the knife

edge can be placed facing either direction along the laser axis, allowing for tailoring of either, or both, ramps of the neutral gas density.
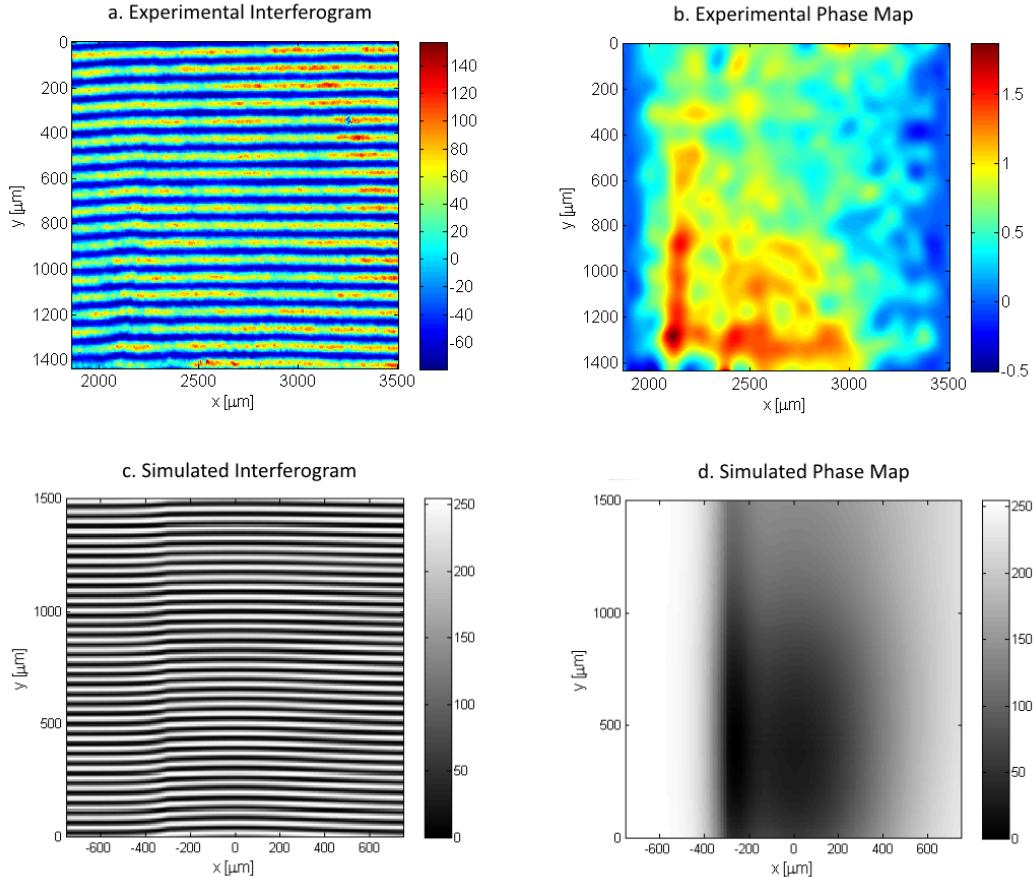
Figure 7.23: Top, interferometry (a) and phase extraction (b) of helium gas perturbed by a knife edge. The knife placement is bottom left. Bottom, the simulated interferogram (c) and resultant phase profile (d) are shown.

This approach has been preliminarily investigated to consider the types of profiles which are achievable. By varying knife position and jet backing pressure, the available parameter space was explored in search of an ideal arrangement. Figure 7.23 shows the interferogram and extracted phase map for one such configuration, with 23 bar backing pressure.

Abel inversion no longer provides a suitable estimate of the density profile as the gas profile is asymmetric after interacting with the knife. Instead, a simulated interferogram was constructed to produce an approximate density profile. The resulting profile shows the expected steepening along the knife side, and a reduction in the scale length of the target. Further steepening

could be achieved by sending a pre-pulse incident on this pre-steepened profile. Alternatively, the knife edge could be placed opposite of the laser, to allow for further shortening of the profile. Figure 7.24 shows calculated profiles based on these potential nozzle and knife edge arrangements. Variation in pre-pulse energy and delay timing may allow for even further compression, making it possible to reduce target size below 500 μm.
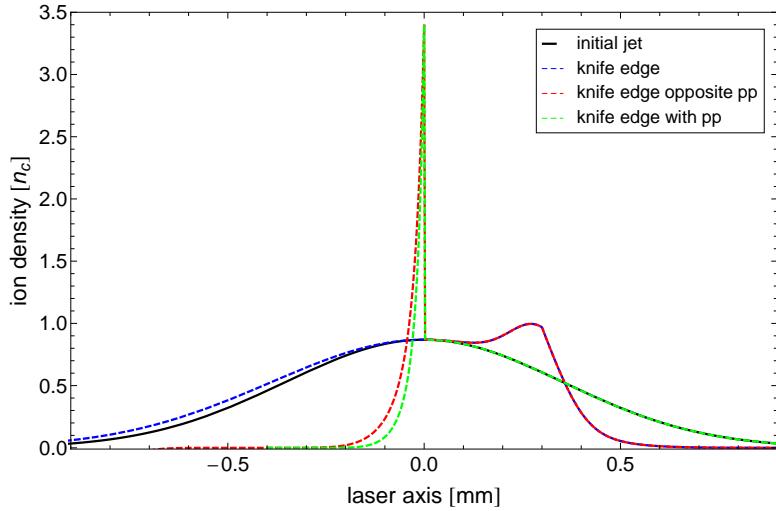
Figure 7.24: A comparison of the longitudinal neutral density profiles achievable using varying combinations of knife edge and pre-pulse shaping. The initial jet (black) may be perturbed at the front or the rear (blue) of the jet. If the knife edge perturbs the front of the target, additional steepening can be achieved (green). If the rear of the jet is perturbed, the subsequent pre-pulse incidence produces an overall shortening of the target (red).

7.5 Conclusions

We have observed the acceleration of He$^+$ ion beams to MeV energies via the interaction of an intense, two-pulse sequence of 10 μm light with a near critical density gas jet target. The generation of a variable pre-pulse allows for the manipulation of the neutral density profile through heating and hydrodynamic expansion of the gas over the 25 ns inter-pulse period. Within a range of pre-pulse energies, we consistently observe the generation of ion beams with energies greater than 1 MeV. We conclude that the appearance of accelerated ion beams is extremely sensitive to the longitudinal target density.

With an appropriately steepened profile, 2D PIC simulations predict the formation of an electrostatic collisionless shock capable of reflecting ions, in

agreement with the experimental observations. In the absence of steepening, the laser bores a channel in the underdense plasma, self focussing and eventually filamenting before building a strong enough energy density gradient in the plasma to produce a shock. Simulations further reveal the continued progression of the shock into dense plasma results in significant heating and reflection of downstream ions. Dissipation in the plasma reduces the shock velocity and potential, resulting in a broad range of reflected ion energies. Rather than a quasi-monoenergetic peak, the resulting spectrum exhibits a large energy spread.

The problem of generating monoenergetic beams from shock reflection is challenging, and achieving the proper combination of laser pulse characteristics and target design continue to be elusive with modern technology. Nonetheless, the unique opportunities conferred by a flexible long wavelength laser along with dynamic gas targets are an exciting source of new and interesting physics, with the promise of eventually achieving a high repetition rate, low emittance, variable ion source. In conjunction with advances in conventional accelerator technology, a hybrid approach featuring laser generated ion beams injected into a high repetition rate accelerator could one day provide a feasible alternative to current synchrotron designs for applications in science, medicine, and industry.

Bibliography

[1] C. Gardner and S. Peggs editors. Conceptual design of the rcms. 2003.

[2] H. Klingbeil. Ferrite cavities. Jan 2012, 1201.1154v1.

[3] W. R. Smythe, T. G. Brophy, R.D. Carlini, C.C. Friedrichs, D.L. Grisham, G. Spalek, and L. C. Wilkerson. Rf cavities with transversely biased ferrite tuning. *Nuclear Science, IEEE Transactions on*, 32(5):2951–2953, Oct.

[4] Y. Zhao. A study of ferrite cavity. Technical Report CAP-341-RF-R, April 2002.

[5] Igor Pogorelsky, Peter Shkolnikov, Min Chen, Alexander Pukhov, Vitaly Yakimenko, Paul McKenna, David Carroll, David Neely, Zulfikar Najmudin, Louise Willingale, Daniil Stolyarov, Elena Stolyarova, and George Flynn. Proton and ion beams generated with picosecond co2 laser pulses. *AIP Conference Proceedings*, 1086(1):532–537, 2009.

[6] Charlotte A. J. Palmer, N. P. Dover, I. Pogorelsky, M. Babzien, G. I. Dudnikova, M. Ispiriyan, M. N. Polyanskiy, J. Schreiber, P. Shkolnikov, V. Yakimenko, and Z. Najmudin. Monoenergetic proton beams accelerated by a radiation pressure driven shock. *Physical Review Letters*, 106(1):014801, 2011.

[7] F. Fiuza, A. Stockem, E. Boella, R. A. Fonseca, L. O. Silva, D. Haberberger, S. Tochitsky, W. B. Mori, and C. Joshi. Ion acceleration from laser-driven electrostatic shocksa). *Physics of Plasmas (1994-present)*, 20(5), 2013.

[8] I. V. Pogorelsky, V. Yakimenko, M. Polyanskiy, P. Shkolnikov, M. Ispiryan, D. Neely, P. McKenna, D. Carroll, Z. Najmudin, and L. Willingale. Ultrafast co2 laser technology: Application in ion acceleration. *Nuclear Instruments and Methods in Physics Research Section*

A: Accelerators, Spectrometers, Detectors and Associated Equipment, 620(1):67–70, 2010.

[9] Mikhail N. Polyanskiy, Igor V. Pogorelsky, and Vitaly Yakimenko. Picosecond pulse amplification in isotopic co2 active medium. *Optics Express*, 19(8):7717–7725, Apr 2011.

[10] E. Gütlich, P. Forck, W. Ensinger, and B. Walasek-Höhne. Scintillation screen investigations for high-current ion beams. *IEEE Transactions on Nuclear Science*, 57(3):1414–1419, 2010.

[11] GT. Hanson, R. Aggeler, D. Oglesbee, M. Cannon, RA. Capaldi, RY. Tsien, , and SJ. Remington. Investigating mitochondrial redox potential with redox-sensitive green fluorescent protein indicators. *Journal of Biological Chemistry*, 279(13):13044–53, 2004.

[12] Antoine J.W.G. Visser and Olaf J. Rolinski. Basic photophyics. dec 2010.

[13] Fn tandem van de graaff accelerator. Technical report, 2012.

[14] Gene Sprouse. Experimental nuclear physics at stony brook: Past, present, and future. 2006.

[15] European Machine Vision Association. *EMVA 1288 Standard for Characterization of Image Sensors and Cameras*, 3.0 edition, 2010.

[16] K.A. Olive and Particle Data Group. Review of particle physics. *Chinese Physics C*, 38(9):090001, 2014.

[17] Kenneth R. Hogstrom and Peter R. Almond. Review of electron beam therapy physics. *Physics in Medicine and Biology*, 51(13):R455, 2006.

[18] Particle Therapy Cooperative Group. Hadron therapy patient statistics, 2014.

[19] Dieter Schardt, Thilo Elsässer, and Daniela Schulz-Ertner. Heavy-ion tumor therapy: Physical and radiobiological benefits. *Rev. Mod. Phys.*, 82:383–425, Feb 2010.

[20] S.E. Combs. Clinical indications for carbon ion radiotherapy and radiation therapy with other heavier ions. In U. Linz, editor, *Ion Beam Therapy: Fundamentals, Techonology, Clinical Applications*. Springer-Verlag, Berlin Heidelberg, 2012.

[21] S.G. Peggs. Fundamental limits to stereotactic proton therapy. *Nuclear Science, IEEE Transactions on*, 51(3):677–681, June 2004.

[22] A. J. Mackinnon, P. K. Patel, M. Borghesi, R. C. Clarke, R. R. Freeman, H. Habara, S. P. Hatchett, D. Hey, D. G. Hicks, S. Kar, M. H. Key, J. A. King, K. Lancaster, D. Neely, A. Nikkro, P. A. Norreys, M. M. Notley, T. W. Phillips, L. Romagnani, R. A. Snavely, R. B. Stephens, and R. P. J. Town. Proton radiography of a laser-driven implosion. *Phys. Rev. Lett.*, 97:045001, Jul 2006.

[23] D. Trbojevic, C. Gardner, and N. Cook editors. ion rapid cycling medical synchrotron : Pre-conceptual design report. 2011.

[24] Steve Peggs, editor. *European Spallation Source Technical Design Report*. 2013.

[25] V. Ashley et al. The accelerator-driven thorium reactor power station. *Proceedings of the Institution of Civil Engineers*, pages 127–135, 2011.

[26] J. Wei et al. Progress towards the facility for rare isotope beams. *Proceedings of PAC2013*, pages 1453–1457, 2013.

[27] Ali Ahmad, Suzanne L. Sheehy, and Geoffrey T. Parks. The effect of beam interruptions on the integrity of {ADSR} fuel pin cladding: A thermo-mechanical analysis. *Annals of Nuclear Energy*, 46(0):97 – 105, 2012.

[28] Max Tabak, James Hammer, Michael E. Glinsky, William L. Kruer, Scott C. Wilks, John Woodworth, E. Michael Campbell, Michael D. Perry, and Rodney J. Mason. Ignition and high gain with ultrapowerful lasers*. *Physics of Plasmas (1994-present)*, 1(5):1626–1634, 1994.

[29] M. Roth, T. E. Cowan, M. H. Key, S. P. Hatchett, C. Brown, W. Fountain, J. Johnson, D. M. Pennington, R. A. Snavely, S. C. Wilks, K. Yasuike, H. Ruhl, F. Pegoraro, S. V. Bulanov, E. M. Campbell, M. D. Perry, and H. Powell. Fast ignition by intense laser-accelerated proton beams. *Phys. Rev. Lett.*, 86:436–439, Jan 2001.

[30] W.H. Panofsky. The evolution of particle accelerators & colliders, 1997.

[31] N. Cook, J. M. Brennan, and S. Peggs. Harmonic ratcheting for fast acceleration. *Phys. Rev. ST Accel. Beams*, 17:042001, Apr 2014.

[32] H. Wiedemann. *Particle Accelerator Physics*. Springer, Berlin, 2007.

[33] D. A. Edwards and M. J. Syphers. *An Introduction to the Physics of High Energy Accelerators.* Wiley-VCH Verlag GmbH, 2008.

[34] M. Conte and W. MacKay. *An Introduction to the Physics of Particle Accelerators.* World Scientific, Singapore, 1991.

[35] N. Tsoupas, V. Litvinenko, D. Kayran, and W.W. MacKay. Design of an achromatic and uncoupled medical gantry for radiation therapy. *Proceedings of 2011 Particle Accelerator Conference*, pages 2163–2165, 20011.

[36] Ernest D. Courant, M. Stanley Livingston, and Hartland S. Snyder. The strong-focusing synchroton¯a new high energy accelerator. *Phys. Rev.*, 88:1190–1196, Dec 1952.

[37] E. D. Courant, M. S. Livingston, H. S. Snyder, and J. P. Blewett. Origin of the "strong-focusing" principle. *Phys. Rev.*, 91:202–203, Jul 1953.

[38] E.D. Courant. Computer Studies Of Phase-Lock Acceleration. pages 201–204, 1961.

[39] C.J. Gardner. Rf capture and acceleration of gold ions in booster. Technical Report 7, November 1999.

[40] K.J. Peach et al. Conceptual design of a nonscaling fixed field alternating gradient accelerator for protons and carbon ions for charged particle therapy. *Physical Review Special Topics - Accelerators and Beams*, 16(030101), 2013.

[41] S. Peggs, T. Satogata, and J. Flanz. A survey of hadron therapy accelerator technologies. *Proceedings of PAC07*, pages 115–119, 2007.

[42] R.C. Arnold and J. Meyer-ter Vehn. Inertial confinement fusion driven by heavy-ion beams. *Reports on Progress in Physics*, 50(5), 1987.

[43] J.M. Brennan. Ferrite loaded cavities. In A.W. Chao and M. Tigner, editors, *Handbook of Accelerator Physics and Engineering*, pages 524–526. World Scientific, 1999.

[44] K. Kaspar, H. G. König, and T. Winnefeld. Studies on the maximum rf voltages in ferrite-tuned accelerating cavitiess. *Proceedings of EPAC 2004*, pages 985–987, 2004.

[45] J.E. Griffin and G. Nicholls. A review of some dynamic loss properties of ni-zn accelerator rf system ferrite. *Nuclear Science, IEEE Transactions on*, 26(3):3965–3967, June.

[46] Ferroxcube. *Soft Ferrites: Introduction*, Sept 2008.

[47] Amikam Aharoni. Micromagnetics: past, present and future. *Physica B: Condensed Matter*, 306(14):1 – 9, 2001. Proceedings of the Third International Symposium on Hysteresis an d Micromagnetics Modeling.

[48] C. Ohmori and others. Development of a high gradient rf system using a nanocrystalline soft magnetic alloy. *Physical Review Special Topics - Accelerators and Beams*, 16(112002), 2013.

[49] A.G. Ruggiero. Rf acceleration with harmonic number jump. *Physical Review Special Topics - Accelerators and Beams*, 9(100101), 2006.

[50] T. Planche et al. Harmonic number jump acceleration of muon beams in zero-chromatic ffag rings. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 632:7–17, 2011.

[51] F. Tamura et al. Multiharmonic rf feedforward system for beam loading compensation in wide-band cavities of a rapid cycling synchrotron. *Physical Review Special Topics - Accelerators and Beams*, 14(051004), 2011.

[52] M. Nomura et al. The origin of magnetic alloy core buckling in j-parc 3 gev rcs. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 623:903–909, 2010.

[53] Joseph M. Kats. Bucket parameters and bunching factor for double rf system. Technical Report 378, August 1993.

[54] E.N. Shaposhnikova. Bunched beam transfer matrices in single and double rf systems. Technical Report 94-19, August 1994.

[55] Y.H. Chin. Double rf system for bunch shortening. Technical Report LBL-29622, November 1990.

[56] C. Carli and M. Chanel. Creation of hollow bunches by redistribution of phase space surfaces. Technical Report 2001-073, 2002.

[57] A. Macchi. *A Superintense Laser-Plasma Interaction Theory Primer.* Springer, London, 2013.

[58] P.M. Woodward. A method of calculating the field over a plane aperture required to produce a given polar diagram. *Journal of Institution of Electrical Engineers*, 93:1554–1558, 1947.

[59] J.D. Lawson. Lasers and accelerators. *IEEE Transactions on Nuclear Science*, 26(3):4217–4219, 1979.

[60] T. Plettner, R. L. Byer, E. Colby, B. Cowan, C. M. S. Sears, J. E. Spencer, and R. H. Siemann. Visible-laser acceleration of relativistic electrons in a semi-infinite vacuum. *Phys. Rev. Lett.*, 95:134801, Sep 2005.

[61] D. Cline, L. Shao, X. Ding, Y. Ho, Q. Kong, and P. Wang. First observation of acceleration of electrons by a laser in a vacuum. *Journal of Modern Physics*, 4:1–6, 2013.

[62] J.D. Huba. NRL plasma formulary. Technical report, 2007.

[63] P. Gibbon. *Short Pulse Laser Interactions with Matter: An Introduction.* World Scientific, London, 2005.

[64] M.V. Ammosov, N.B. Delone, and V.P. Krainov. Tunnel ionization of complex ions and of atomic ions in an alternating electromagnetic field. *Journal of Experimental Theoretical Physics*, 64(6):1191–1194, 1986.

[65] L.V. Keldysh. Ionization in the field of a strong electromagnetic wave. *Journal of Experimental Theoretical Physics*, 20(5):1307–1314, 1965.

[66] Andrea Macchi, Marco Borghesi, and Matteo Passoni. Ion acceleration by superintense laser-plasma interaction. *Rev. Mod. Phys.*, 85:751–793, May 2013.

[67] Predhiman Kaw and John Dawson. Relativistic nonlinear propagation of laser beams in cold overdense plasmas. *Physics of Fluids (1958-1988)*, 13(2):472–481, 1970.

[68] Claire Max and Francis Perkins. Strong electromagnetic waves in overdense plasmas. *Phys. Rev. Lett.*, 27:1342–1345, Nov 1971.

[69] L. Schlessinger and J. Wright. Inverse-bremsstrahlung absorption rate in an intense laser field. *Phys. Rev. A*, 20:1934–1945, Nov 1979.

[70] F. Brunel. Not-so-resonant, resonant absorption. *Phys. Rev. Lett.*, 59:52–55, Jul 1987.

[71] J. Denavit. Absorption of high-intensity subpicosecond lasers on solid density targets. *Phys. Rev. Lett.*, 69:3052–3055, Nov 1992.

[72] Scott C. Wilks and W.L. Kruer. Absorption of ultrashort, ultra-intense laser light by solids and overdense plasmas. *Quantum Electronics, IEEE Journal of*, 33(11):1954–1968, Nov 1997.

[73] M. G. Haines, M. S. Wei, F. N. Beg, and R. B. Stephens. Hot-electron temperature and laser-light absorption in fast ignition. *Phys. Rev. Lett.*, 102:045008, Jan 2009.

[74] F. N. Beg, A. R. Bell, A. E. Dangor, C. N. Danson, A. P. Fews, M. E. Glinsky, B. A. Hammel, P. Lee, P. A. Norreys, and M. Tatarakis. A study of picosecond lasersolid interactions up to 1019 wcm2. *Physics of Plasmas (1994-present)*, 4(2):447–457, 1997.

[75] Y. Ping, R. Shepherd, B. F. Lasinski, M. Tabak, H. Chen, H. K. Chung, K. B. Fournier, S. B. Hansen, A. Kemp, D. A. Liedahl, K. Widmann, S. C. Wilks, W. Rozmus, and M. Sherlock. Absorption of short laser pulses on solid targets in the ultrarelativistic regime. *Phys. Rev. Lett.*, 100:085004, Feb 2008.

[76] A.P.L. Robinson, D.J. Strozzi, J.R. Davies, L. Gremillet, J.J. Honrubia, T. Johzaki, R.J. Kingham, M. Sherlock, and A.A. Solodov. Theory of fast electron transport for fast ignition. *Nuclear Fusion*, 54(5):054003, 2014.

[77] D.A Jarosynzski, R. Bingham, and R.A. Cairns. *Laser-Plasma Interactions*. CRC Press, Boca Raton, 2009.

[78] M. Honda. On the maximum current for a self-focusing relativistic electron beam. *Physics of Plasmas (1994-present)*, 7(5):1606–1608, 2000.

[79] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. *Rev. Mod. Phys.*, 81:1229–1285, Aug 2009.

[80] Weiner. A.M. *Ultrafast Optics*. John Wiley & Sons, Hoboken, 2003.

[81] S. C. Wilks, A. B. Langdon, T. E. Cowan, M. Roth, M. Singh, S. Hatchett, M. H. Key, D. Pennington, A. MacKinnon, and R. A. Snavely. Energetic proton generation in ultra-intense lasersolid interactions. *Physics of Plasmas (1994-present)*, 8(2):542–549, 2001.

[82] P. Mora. Plasma expansion into a vacuum. *Phys. Rev. Lett.*, 90:185002, May 2003.

[83] M. Hegelich, S. Karsch, G. Pretzler, D. Habs, K. Witte, W. Guenther, M. Allen, A. Blazevic, J. Fuchs, J. C. Gauthier, M. Geissel, P. Audebert, T. Cowan, and M. Roth. Mev ion jets from short-pulse-laser interaction with thin foils. *Phys. Rev. Lett.*, 89:085002, Aug 2002.

[84] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon. Absorption of ultra-intense laser pulses. *Phys. Rev. Lett.*, 69:1383–1386, Aug 1992.

[85] A P L Robinson, P Gibbon, M Zepf, S Kar, R G Evans, and C Bellei. Relativistically correct hole-boring and ion acceleration by circularly polarized laser pulses. *Plasma Physics and Controlled Fusion*, 51(2):024004, 2009.

[86] M. Zepf, M. Castro-Colin, D. Chambers, S. G. Preston, J. S. Wark, J. Zhang, C. N. Danson, D. Neely, P. A. Norreys, A. E. Dangor, A. Dyson, P. Lee, A. P. Fews, P. Gibbon, S. Moustaizis, and M. H. Key. Measurements of the hole boring velocity from doppler shifted harmonic emission from solid targets. *Physics of Plasmas (1994-present)*, 3(9):3242–3244, 1996.

[87] Andrea Macchi, Federica Cattani, Tatiana V. Liseykina, and Fulvio Cornolti. Laser acceleration of ion bunches at the front surface of overdense plasmas. *Phys. Rev. Lett.*, 94:165003, Apr 2005.

[88] A P L Robinson, M Zepf, S Kar, R G Evans, and C Bellei. Radiation pressure acceleration of thin foils with circularly polarized laser pulses. *New Journal of Physics*, 10(1):013021, 2008.

[89] Andrea Macchi, Tatiana V. Liseikina, Sara Tuveri, and Silvia Veghini. Theory and simulation of ion acceleration with circularly polarized laser pulses. *Comptes Rendus Physique*, 10(23):207 – 215, 2009. Laser acceleration of particles in plasma Acclration laser de particules dans les plasmas.

[90] V. A. Vshivkov, N. M. Naumova, F. Pegoraro, and S. V. Bulanov. Nonlinear electrodynamics of the interaction of ultra-intense laser pulses with a thin foil. *Physics of Plasmas (1994-present)*, 5(7):2727–2741, 1998.

[91] P. J. Kellogg. Flow of plasma around the earth. *Journal of Geophysical Research*, 67(10):3805–3811, 1962.

[92] Roger Blandford and David Eichler. Particle acceleration at astrophysical shocks: A theory of cosmic ray origin. *Physics Reports*, 154(1):1 – 75, 1987.

[93] B. T. Draine and C. F. McKee. Theory of interstellar shocks. *Annual Review of Astronomy and Astrophysics*, 31:373–432, 1993.

[94] D. W. Forslund and J. P. Freidberg. Theory of laminar collisionless shocks. *Phys. Rev. Lett.*, 27:1189–1192, Nov 1971.

[95] R. Z. Sagdeev. Cooperative Phenomena and Shock Waves in Collisionless Plasmas. *Reviews of Plasma Physics*, 4:23, 1966.

[96] R. A. Cairns, R. Bingham, P. Norreys, and R. Trines. Laminar shocks in high power laser plasma interactions. *Physics of Plasmas (1994-present)*, 21(2):–, 2014.

[97] Luís O. Silva, Michael Marti, Jonathan R. Davies, Ricardo A. Fonseca, Chuang Ren, Frank S. Tsung, and Warren B. Mori. Proton shock acceleration in laser-plasma interactions. *Phys. Rev. Lett.*, 92:015002, Jan 2004.

[98] G. Sorasio, M. Marti, R. Fonseca, and L. O. Silva. Very high mach-number electrostatic shocks in collisionless plasmas. *Phys. Rev. Lett.*, 96:045005, Feb 2006.

[99] N. P. Dover, C. A. J. Palmer, M. Babzien, A. R. Bell, A. E. Dangor, T. Horbury, M. Ispiriyan, M. N. Polyanskiy, J. Schreiber, S. Schwartz, P. Shkolnikov, V. Yakimenko, I. Pogorelsky, and Z. Najmudin. Optical probing of shocks driven into overdense plasmas by laser hole-boring. *arXiv*, 2012.

[100] Dan Haberberger, Sergei Tochitsky, Frederico Fiuza, Chao Gong, Ricardo A. Fonseca, Luis O. Silva, Warren B. Mori, and Chan Joshi. Collisionless shocks in laser-produced plasma generate monoenergetic high-energy proton beams. *Nat Phys*, 8(1):95–99, 2012.

[101] F. Fiuza, A. Stockem, E. Boella, R. A. Fonseca, L. O. Silva, D. Haberberger, S. Tochitsky, C. Gong, W. B. Mori, and C. Joshi. Laser-driven shock acceleration of monoenergetic ion beams. *arXiv*, 2012.

[102] Marcus Babzien, Ilan Ben-Zvi, Karl Kusche, Igor V. Pavlishin, Igor V. Pogorelsky, David P. Siddons, Vitaly Yakimenko, David Cline, Feng Zhou, Tachishige Hirose, Yoshio Kamiya, Tetsuro Kumita, Tsunehiko Omori, Junji Urakawa, and Kaoru Yokoya. Observation of the second harmonic in thomson scattering from relativistic electrons. *Phys. Rev. Lett.*, 96:054802, Feb 2006.

[103] P. Oliva, M. Carpinelli, B. Golosio, P. Delogu, M. Endrizzi, J. Park, I. Pogorelsky, V. Yakimenko, O. Williams, and J. Rosenzweig. Quantitative evaluation of single-shot inline phase contrast imaging using an inverse compton x-ray source. *Applied Physics Letters*, 97(13):–, 2010.

[104] W.D. Kimura, N.E. Andreev, M. Babzien, I. Ben-Zvi, D.B. Cline, C.E. Dilley, S.C. Gottschalk, S.M. Hooker, K.P. Kusche, S.V. Kuznetsov, I.V. Pavlishin, I.V. Pogorelsky, A.A. Pogosova, L.C. Steinhauer, A. Ting, V. Yakimenko, A. Zigler, and F. Zhou. Inverse free electron lasers and laser wakefield acceleration driven by co2 lasers. *Phil. Trans. R. Soc. Lond. A*, 364(1840):611–622, 2006.

[105] W. D. Kimura, L. P. Campbell, C. E. Dilley, S. C. Gottschalk, D. C. Quimby, A. van Steenbergen, M. Babzien, I. Ben-Zvi, J. C. Gallardo, K. P. Kusche, I. V. Pogorelsky, J. Skaritka, V. Yakimenko, D. B. Cline, P. He, Y. Liu, L. C. Steinhauer, and R. H. Pantell. Detailed experimental results for laser acceleration staging. *Phys. Rev. ST Accel. Beams*, 4:101301, Oct 2001.

[106] C.K.N. Patel. High power carbon dioxide lasers. *Scientific American*, pages 23–24, 1968.

[107] P.K. Cheo. CO$_2$ lasers. In A.K. Levine and A.J. DeMaria, editors, *Lasers*. Marcel Dekker, Inc., New York, 1971.

[108] R.K. Brimacombe and J. Reid. Accurate measurements of pressure-broadened line widths in a transversely excited co$_2$ discharge. *IEEE Jounral of Quantum Electronics*, 19(11):1668–1673, 1983.

[109] M.N. Polyanskiy and M. Babzien. Ultrashort pulses. In D.C. Dumitras, editor, *Ultrashort Pulses, CO2 Laser - Optimisation and Application*. InTech, 2012.

[110] A. J. Alcock and P. B. Corkum. Ultra-fast switching of infrared radiation by laser-produced carriers in semiconductors. *Canadian Journal of Physics*, 57(9):1280–1290, 1979, http://dx.doi.org/10.1139/p79-176.

[111] M. I. Gallant and H. M. van Driel. Infrared reflectivity probing of thermal and spatial properties of laser-generated carriers in germanium. *Phys. Rev. B*, 26:2133–2146, Aug 1982.

[112] Claude Rolland and P. B. Corkum. Generation of 130-fsec midinfrared pulses. *J. Opt. Soc. Am. B*, 3(12):1625–1629, Dec 1986.

[113] S. Ya. Tochitsky, R. Narang, C. Filip, C. E. Clayton, K. A. Marsh, and C. Joshi. Generation of 160-ps terawatt-power co2 laser pulses. *Opt. Lett.*, 24(23):1717–1719, Dec 1999.

[114] Robert A. Fisher, P. L. Kelley, and T. K. Gustafson. Subpicosecond pulse generation using the optical kerr effect. *Applied Physics Letters*, 14(4):140–143, 1969.

[115] R.W. Boyd. *Nonlinear Optics*. Academic Press, San Diego, 2003.

[116] Catalin V. Filip, Ritesh Narang, Sergei Ya. Tochitsky, Christopher E. Clayton, and Chandrashekhar Joshi. Optical kerr switching technique for the production of a picosecond, multiwavelength co2 laser pulse. *Appl. Opt.*, 41(18):3743–3747, Jun 2002.

[117] Chad Nelson and Jordan Crist. Predicting laser beam characteristics. *LTJ*, 1:36–39, January 2012.

[118] A. Maksimchuk, S. Reed, N. Naumova, V. Chvykov, B. Hou, G. Kalintchenko, T. Matsuoka, J. Nees, P. Rousseau, G. Mourou, and V. Yanovsky. Energy scaling of quasi-monoenergetic electron beams from laser wakefields driven by 40-tw ultra-short pulses. *Applied Physics B*, 89(2-3):201–207, 2007.

[119] J. Osterhoff, A. Popp, Zs. Major, B. Marx, T. P. Rowlands-Rees, M. Fuchs, M. Geissler, R. Hörlein, B. Hidding, S. Becker, E. A. Peralta, U. Schramm, F. Grüner, D. Habs, F. Krausz, S. M. Hooker, and S. Karsch. Generation of stable, low-divergence electron beams by laser-wakefield acceleration in a steady-state-flow gas cell. *Phys. Rev. Lett.*, 101:085002, Aug 2008.

[120] D. Kaganovich, M. H. Helle, D. F. Gordon, and A. Ting. Measurements and simulations of shock wave generated plasma-vacuum interface. *Physics of Plasmas (1994-present)*, 18(12):–, 2011.

[121] S. Semushin and V. Malka. High density gas jet nozzle design for laser target production. *Review of Scientific Instruments*, 72(7):2961–2965, 2001.

[122] M. Borghesi, A. Giulietti, D. Giulietti, L. A. Gizzi, A. Macchi, and O. Willi. Characterization of laser plasmas for interaction studies: Progress in time-resolved density mapping. *Phys. Rev. E*, 54:6769–6773, Dec 1996.

[123] A. Giulietti, M. Borghesi, C. Danson, D. Giulietti, L.A. Gizzi, A. Macchi, and D. Neely. Picosecond interferometry of plasmas before and after short laser pulse propagation. In O. Svelto, S. De Silvestri, and G. Denardo, editors, *Ultrafast Processes in Spectroscopy*, pages 299–303. Springer US, 1996.

[124] K. Markov and L Preziosi. *Heterogeneous Media: Modelling and Simulation.* Springer, Boston, 1999.

[125] P.K. Panigrahi and K. Muralidhar. *Schlieren and Shadowgraph Methods in Heat and Mass Transfer.* Springer, 2012.

[126] P. R. Bolton, M. Borghesi, C. Brenner, D. C. Carroll, C. De Martinis, A. Flacco, V. Floquet, J. Fuchs, P. Gallegos, D. Giove, J. S. Green, S. Green, B. Jones, D. Kirby, P. McKenna, D. Neely, F. Nuesslin, R. Prasad, S. Reinhardt, M. Roth, U. Schramm, G. G. Scott, S. Ter-Avetisyan, M. Tolley, G. Turchetti, and J. J. Wilkens. Instrumentation for diagnostics and control of laser-accelerated proton (ion) beams. *Physica Medica: European Journal of Medical Physics*, 30(3):255–270, 2014.

[127] Kieran J. McCarthy, J. García López, F. Martín Hernández, B. Zurro, A. Baciero, and M. A. Respaldiza. The response of a chromium doped alumina screen to kev and mev ions. *Journal of Nuclear Materials*, 321(1):78–83, 2003.

[128] Y. Larochelle, L. Beaulieu, B. Djerroud, D. Dor, P. Gendron, E. Jalbert, R. Laforest, J. Pouliot, R. Roy, M. Samri, and C. St-Pierre. Energy-light relation for csi(t1) scintillators in heavy ion experiments at intermediate energies. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 348(1):167–172, 1994.

[129] L. Torrisi. Plastic scintillator investigations for relative dosimetry in proton-therapy. *Nuclear Instruments and Methods in Physics Research*

Section B: Beam Interactions with Materials and Atoms, 170(0):523–530, 2000.

[130] N. Cook, O. Tresca, and R. Lefferts. Scintillator diagnostics for the detection of laser accelerated ion beams. *Journal of Instrumentation*, 9(09):P09004, 2014.

[131] D.L. Henshaw. Applications of cr-39 nuclear track detector in medicine and technology. *Physics in Technology*, 13(6):266, 1982.

[132] P. Baiocchi, S. Cecchini, H. Dekhissi, V. Garutti, G. Giacomelli, G. G. Giani, E. Katsavounidis, G. Iori, L. Patrizii, V. Popa, P. Serra, V. Togo, U. Valdre, and E. Vilela. Calibration with relativistic and low velocity ions of a cr39 nuclear track detector. *Radiation Measurements*, 25(14):145–150, 1995.

[133] S. Reinhardt. *Detection of laser-accelerated protons*. PhD thesis, Ludwig Maximilian University of Munich, 2012.

[134] Joseph Ladislas Wiza. Microchannel plate detectors. *Nuclear Instruments and Methods*, 162:587 – 601, 1979.

[135] R. Prasad, D. Doria, S. Ter-Avetisyan, P. S. Foster, K. E. Quinn, L. Romagnani, C. M. Brenner, J. S. Green, P. Gallegos, M. J. V. Streeter, D. C. Carroll, O. Tresca, N. Dover, C. A. J. Palmer, J. Schreiber, D. Neely, Z. Najmudin, P. McKenna, M. Zepf, and M. Borghesi. Calibration of thomson parabolamcp assembly for multi-mev ion spectroscopy. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 623(2):712–715, 2010.

[136] Photonis USA. *Storage, Handling, and Operation of Microchannel Plates.*

[137] Y.B. Zel'dovich and Y.P. Raizer. *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*. Dover Publications, Inc., Mineola, 2002.

[138] Kelly D. Rakes. Evaluating the response of polyvinyl toluene scintillators used in portal detectors. Master's thesis, Air Force Institute of Technology, mar 2008.

[139] Saint-Gobain Crystals. Bc-400, bc-404, bc-408, bc-412, bc416 premium plastic scintillators. Technical report.

184

[140] Zhao Li, Wu Chong, Heng Yuekun, Zhao Xiaojian, Shi Feng, Sun Zhijia, Wu Jinjie, An Zhenghua, Zhao Yuda, Zhang Ziping, and Wang Yifang. Properties of plastic scintillators after irradiation. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 552(3):449–455, 2005.

[141] Saint-Gobain Crystals. Bc-422q ultra-fast timing plastic scintillator. Technical report.

[142] P. A. Rodnyi, P. Dorenbos, and C. W. E. van Eijk. Energy loss in inorganic scintillators. *physica status solidi (b)*, 187(1):15–29, 1995.

[143] C. Greskovich and S. Duclos. Ceramic scintillators. *Annual Review of Materials Science*, 27(1):69–88, 1997.

[144] C. Greskovich and S. Duclos. Ceramic scintillators. *Annual Review of Materials Science*, 27(1):69–88, 1997.

[145] C.D. Johnson. The development and use of alumina ceramic fluorescent screens. Technical report, European Laboratory for Particle Physics, July 1990.

[146] John W. Noé. Operation of the stony brook tandem/linac accelerator system. *Review of Scientific Instruments*, 57(5):757–760, 1986.

[147] R.G. Herb. Pelletron accelerators for very high voltage. *Nuclear Instruments and Methods*, 122(0):267 – 276, 1974.

[148] CVI Melles Griot. *Fundamental Optics*, 2009.

[149] J.B. Birks. *The Theory and Practice of Scintillation Counting*. Macmillan, New York, 1964.

[150] Craun R. and Smith D. Analysis of response data for several organic scintillators. *Nuclear Instrumentation and Methods in Physics Research*, 80:239–244, 1970.

[151] Smith D., Polk R., and Miller T. Measurement of the response of several organic scintillators to electrons, protons and deuterons. *Nuclear Instrumentation and Methods in Physics Research*, 64:157–166, 1969.

[152] G.V ORielly et al. The response of plastic scintillator to protons and deuterons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors, and Associated Equipment*, 368:745–749, 1996.

[153] J.F. Ziegler, M.D. Ziegler, and J.P. Biersack. Srim - the stopping range of ions in matter (2010). *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 268:1818–1823, 2010.

[154] B.H. Daub et al. Response of bc-418 plastic scintillator low-energy protons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors, and Associated Equipment*, 701:171–175, 2013.

[155] N. Cook, N.P. Dover, C.M. Maharjan, Z. Najmudin, I. Pogorelsky, M.N. Polyanskiy, and P. Shkolnikov. Spectral broadening of ions accelerated by a radiation pressure driven shock. *Proceedings of PAC2013*, pages 144–146, 2013.

[156] O. Tresca, N.P. Dover, N. Cook, C. Maharjan, M. N. Polyanskiy, Z. Najmudin, P. Shkolnikov, and I. Pogorelsky. Controlled shock acceleration of helium ions by laser irradiation of hydrodynamically shaped gas jets. 2014.

[157] L.I. Sedov. Propagation of intense (strong) blast waves. *Prikl Mat Mekh*, 10:241, 1946.

[158] G.I. Taylor. The air wave surrounding an expanding sphere. *Proceedings of the Royal Society A*, 186:273, 1946.

[159] G.I. Taylor. The formation of a blast wave by a very intense explosion, i. *Proceedings of the Royal Society A*, 201:159–174, 1950.

[160] O. Petruk. Approximations of the self similar solution for blastwave in a medium with power-law density variation. 2000, astro-ph/0002112.

[161] A.S. Kompaneets. A point explosion in an inhomogeneous atmosphere. *Soviet Physics Doklady*, 5:46–48, 1960.

[162] J. Lukkezen and L. Rademaker. Shock waves through inhomogeneous media. September 2006.

[163] G. Blasse and B.C. Grabmaier. *Luminescent Materials*. Springer, Berlin, 1994.

[164] Z. Najmudin, C. A. J. Palmer, N. P. Dover, I. Pogorelsky, M. Babzien, A. E. Dangor, G. I. Dudnikova, P. S. Foster, J. S. Green, M. Ispiriyan,

D. Neely, M. N. Polyanskiy, J. Schreiber, P. Shkolnikov, and V. Yakimenko. Observation of impurity free monoenergetic proton beams from the interaction of a co2 laser with a gaseous target a). *Physics of Plasmas (1994-present)*, 18(5):–, 2011.

[165] M. S. Wei, S. P. D. Mangles, Z. Najmudin, B. Walton, A. Gopal, M. Tatarakis, A. E. Dangor, E. L. Clark, R. G. Evans, S. Fritzler, R. J. Clarke, C. Hernandez-Gomez, D. Neely, W. Mori, M. Tzoufras, and K. Krushelnick. Ion acceleration by collisionless shocks in high-intensity-laser˘underdense-plasma interaction. *Phys. Rev. Lett.*, 93:155003, Oct 2004.

[166] K. Bennett, C. Brady, H. Schmitz, and C. Ridgers. Epoch users manual. July 2012.

[167] Burton D. Fried. Mechanism for instability of transverse plasma waves. *Physics of Fluids (1958-1988)*, 2(3):337–337, 1959.

[168] Hui Chen, S. C. Wilks, W. L. Kruer, P. K. Patel, and R. Shepherd. Hot electron energy distributions from ultraintense laser solid interactions. *Physics of Plasmas (1994-present)*, 16(2):–, 2009.

[169] A. Stockem, F. Fiuza, A. Bret, R. A. Fonseca, and L. O. Silva. Exploring the nature of collisionless shocks under laboratory conditions. *Sci. Rep.*, 4, 2014.

[170] N.P. Dover, O. Tresca, N. Cook, C. Maharjan, M. N. Polyanskiy, P. Shkolnikov, I. Pogorelsky, and Z. Najmudin. Observation of filamentation instability from energetic electrons propagating through near-critical density plasma. manuscript in preparation. 2014.

[171] F. Fiuza, R.A. Fonseca, J. Tonge, W.B. Mori, and L.O. Silva. Weibel-instability-mediated collisionless shocks in laboratory with ultraintense lasers. *arXiv*, 2012.

[172] Daido Hiroyuki, Nishiuchi Mamiko, and S. Pirozhkov Alexander. Review of laser-driven ion sources and their applications. *Reports on Progress in Physics*, 75(5):056401, 2012.