To the Graduate Council:
I am submitting herewith a dissertation written by Martin Joseph Kay entitled "Pointing Control and Stabilization of the High-Energy UV Laser for Laser-Assisted Charge Exchange." I have examined the final paper copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Physics.

Yuri Efremenko, Major Professor

We have read this dissertation
and recommend its acceptance:

Yuri Efremenko

Sarah Cousineau

Nadia Fomin

Maik Lang

Abdurahim Oguz

Accepted for the Council:

Dixie Thompson

Vice Provost and Dean of the Graduate School

To the Graduate Council:
I am submitting herewith a dissertation written by Martin Joseph Kay entitled "Pointing Control and Stabilization of the High-Energy UV Laser for Laser-Assisted Charge Exchange." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Physics.

Yuri Efremenko, Major Professor

We have read this dissertation
and recommend its acceptance:

Yuri Efremenko

Sarah Cousineau

Nadia Fomin

Maik Lang

Abdurahim Oguz

Accepted for the Council:

Dixie Thompson

Vice Provost and Dean of the Graduate School

(Original signatures are on file with official student records.)

Pointing Control and Stabilization of the High-Energy UV Laser for Laser-Assisted Charge Exchange

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Martin Joseph Kay

May 2023

Acknowledgements

This research used resources of the Spallation Neutron Source, which is a DOE Office of Science User Facility. This work has also been partially supported by the U.S. DOE Grant No. DE-FG02-13ER41967. This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy (DOE). The U.S. government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan [1].

Abstract

Laser-Assisted Charge Exchange (LACE) is an experimental method of charge exchange injection into a proton accumulator ring that is being developed at the Spallation Neutron Source (SNS) in Oak Ridge National Laboratory (ORNL) as an alternative to hazardous injection foils. The current scheme of LACE requires a high-energy, low-repetition-rate UV (355 nm) laser beam (140 mJ pulses at 10 Hz) to be transported over 65 meters to the laser-particle interaction point (IP) in a high-radiation area of the accelerator. Thermal effects and other disturbances along the free-space laser transport line cause the beam to slowly drift away from the IP and jitter at a frequency comparable to the pulse repetition rate. A control system was designed, simulated, and constructed to stabilize the pointing of the laser beam to allow stable operation of the experiment. The laser pointing stabilization system is based on feedback between Complementary Metal-Oxide-Semiconductor (CMOS) cameras and a steering mirror with piezoelectric actuators. A PC running custom-made LabVIEW software acts a controller in open- and closed-loop modes, as well as a diagnostic tool. An analytical model of the system was used for optimization of the control law, and the system performs as well in the field as it did in laboratory tests. The laser pointing stabilization system eliminates the slow drift by keeping the beam aligned at the IP for an indefinite amount of time, and the jitter is reduced to the level of the pulse-to-pulse fluctuations.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The research reported on in this thesis took place at the Spallation Neutron Source (SNS) accelerator at Oak Ridge National Laboratory (ORNL). ORNL is one of 17 laboratories overseen by the U.S. Department of Energy (DOE) which address a wide range of scientific and technological challenges of our time. ORNL is managed and operated by UT-Battelle, LLC, a non-profit limited liability company formed in 50-50 partnership between the University of Tennessee and Batelle Memorial Institute for the purpose of delivering the DOE's research mission at ORNL. This chapter gives an overview of the SNS accelerator and the motivation for development of Laser-Assisted Charge Exchange (LACE), a promising technique that could be used as a new mode of charge exchange injection into the next generation of high-power synchrotrons. A brief history of LACE is outlined, and then the main concepts used in the SNS scheme are discussed.

1.1 The Spallation Neutron Source

The SNS project began in 1996 as a collaboration between five DOE national laboratories: Lawrence Berkeley National Laboratory (LBNL), Los Alamos Natonal Laboratory (LANL), Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL) and ORNL [2]. Each laboratory was responsible for the design and

construction of a different part of the accelerator (Fig. 1.1). LBNL was responsible for the H^- Linac Injector System, colloquially named the "front-end" of the accelerator. LANL designed the 1 GeV linear accelerator (linac) and the associated radiofrequency (RF) systems. BNL was responsible for beam transport lines and the proton accumulator ring. ANL constructed the neutron scattering instrumentation around the spallation target. ORNL was tasked with the installation and integration of different components of the accelerator, commissioning of the beam, and ultimately operating the facility. In 2000, Thomas Jefferson National Accelerator Facility (Jefferson Laboratory) became the sixth member of the collaboration when plans were made to upgrade the linac to include a superconducting section. Construction of the accelerator began in 1998 and was completed in 2006 when the proton beam was first delivered to the liquid mercury spallation target. Construction costs remained under the $1.4B budget, and the project was completed on-schedule. During the design phase of the SNS, the requirement was set that the facility should be able to initially operate with greater than 90% user availability, with the ultimate goal being 95%.

1.1.1 Overview

The SNS was designed to initially produce a 1 MW, 1 GeV proton beam of $\sim$1 μs pulses at 60 Hz that would be delivered to the spallation target. The accelerator and associated infrastructure were also designed to be upgraded some time after initial commissioning, and the option was reserved to increase the proton beam power on target up to 2 MW. Presently, the SNS routinely delivers a 1.5 MW proton beam in 1 ms long pulses at 1.3 GeV to the target with around 95% availability. The Proton Power Upgrade (PPU) project is underway now with plans to increase the proton beam power to 2.8 MW, where 2 MW will be directed to the existing target and 0.8 MW will go to the future Second Target Station (STS) [3]. The PPU is slated to be completed in 2025, and the STS project is still in the design phase. At the time of

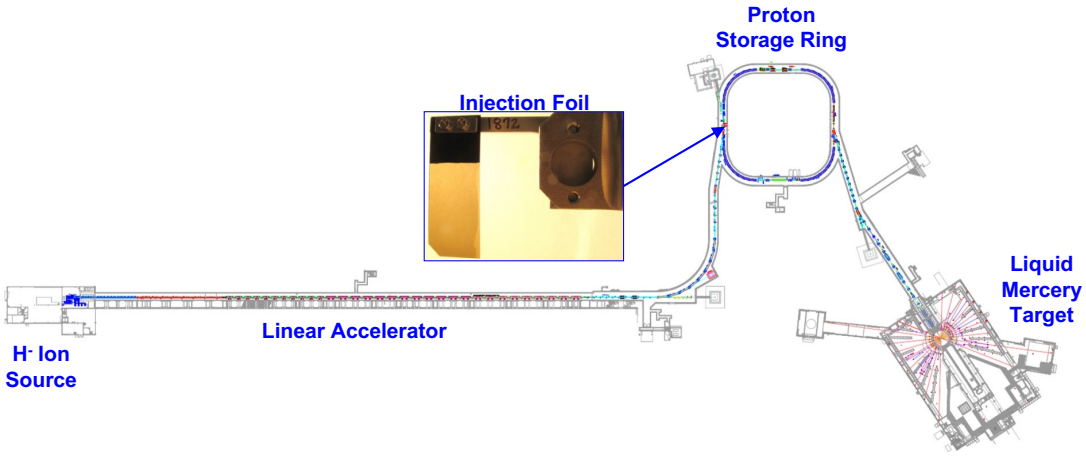

Figure 1.1: Layout of the SNS accelerator.

this writing, the SNS boasts both the most intense proton beam and highest peak brightness neutron beam in the world.

The four main components of the SNS are the front-end systems, linac, proton accumulator ring, and liquid mercery target. The front-end accelerator systems comprise the H^- ion source, low-energy beam transport line (LEBT), a radiofrequency quadrupole (RFQ) accelerator, and the medium-energy beam transport line (MEBT). The ion source outputs a continuous, ∼50 mA H^- beam at 65 keV that is focused and guided into the RFQ by a series of electrostatic "lenses" in the LEBT (Fig. 1.2). The RFQ accelerates the ion beam from 65 keV to 2.5 MeV while performing the initial bunching of the beam. The RFQ creates micropulses from the continuous beam at 402.5 MHz with 2.48 ns gaps and bunches them into 1 ms macropulses with 38 mA current. The MEBT houses the chopper system which completes the bunching process. The 1 ms H^- macropulses are chopped at a frequency of ∼1 MHz into 645 ns minipulses with 300 ns gaps which are then injected into the linac (Fig. 1.3).

The linac begins with the normal-conducting section, which consists of six 402.5 MHz drift-tube linac (DTL) cavities followed by four 805 MHz coupled-cavity linac (CCL) modules. The DTL section accelerates the beam from 2.5 MeV to 87 MeV, and the CCL modules accelerate the H^- beam to 186 MeV. The final part of the linac is the superconducting linac (SCL), which consists of medium beta and high beta sections. The medium beta section is made up of 11 cryomodules, each of which is 4.9 m long and contains 3 superconducting radiofrequency (SRF) cavities at 805 MHz. The high beta section has 15 cryomodules, which are each 5.8 m in length and contain 4 SRF cavities. The medium beta section accelerates the beam from 87 MeV to 387 MeV, and the high beta section accelerates the H^- to the final energy of 1.3 GeV. The High Energy Beam Transport (HEBT) line transports the ion beam 169 m from the linac to the ring injection site without any additional acceleration (Fig. 1.4). At the ring injection site, the H^- beam is merged with the circulating proton beam just before the point where both beams pass through the injection foil. The electrons

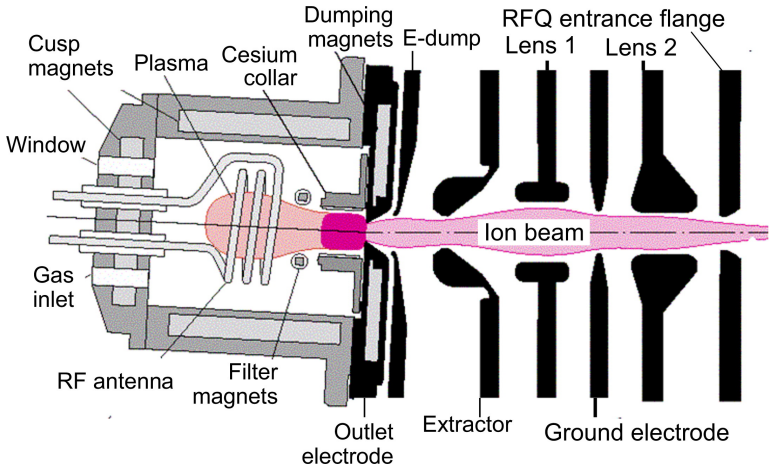

Figure 1.2: The H^- ion source and LEBT line. Reproduced from [2].

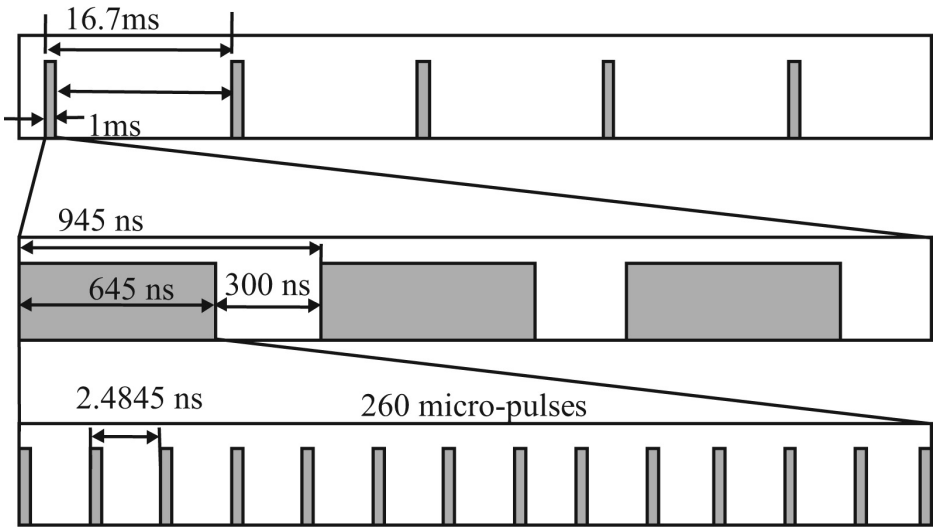

Figure 1.3: Macropulse temporal structure of the ion beam. Reproduced from [2].

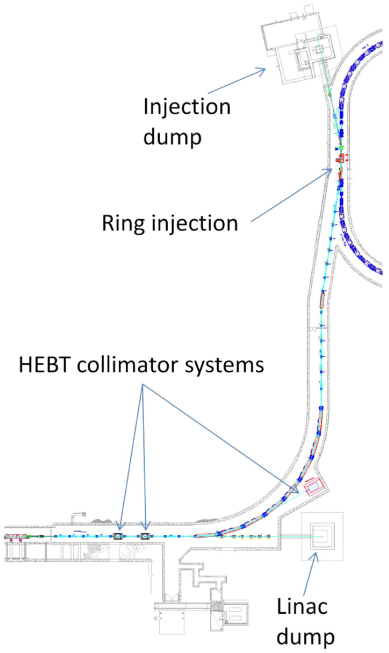

Figure 1.4: The HEBT line from the linac exit to the ring injection site. Reproduced from [2].

are stripped from the H^- particles by the injection foil and then join the other protons circling the ring. The process is called H^- charge exchange injection (CEI), and it is used to build up an intense proton beam in the accumulator ring by repeating the process many times. The ring is 248 m in circumference, and the revolution frequency of the proton beam is close to 1 MHz. Proton bunches are extracted from the ring at a rate of 60 Hz, and after extraction, the 1.5 MW proton beam travels 150 m down the Ring to Target Beam Transfer (RTBT) line to the liquid mercury target (Fig. 1.5). At the target, the proton beam has an RMS transverse size of $\sigma = 7$ (H) $\times$ 20 (V) cm and peak current density of 180 mA/m^2. Spallation occurs when the proton beam makes contact with the target (Fig. 1.6), and the resulting neutrons are guided to the 17 different neutron scattering instruments in the experimental hall.

1.1.2 H^- Charge Exchange Injection Issues

H^- CEI using injection foils is the standard method for producing a high intensity proton beam in an accumulator ring. It was invented by Alvarez in 1951 [5], and it allows for multi-turn injection into an accumulator ring or synchrotron with limited emittance growth and low beam loss [6]. Limited emittance growth of the circulating beam can be achieved with multi-turn CEI through the process of phase space painting, which allows for an intense beam to be built up in the ring. Accelerators that do not use CEI for multi-turn injection lose around 10% of the beam due to the inefficiency of direct injection. That amount of beam loss may be tolerable for low injection powers, but it would cause an unacceptable level of activation in the injection region at high power proton accelerators. At the SNS, the foils are 17 mm $\times$ 30 mm in area so that only 2% of the ions are expected to miss the foil for the nominal RMS beam size of $\sigma = 1.6$ mm $\times$ 1.9 mm, and another 3% of the ions are expected to not be fully stripped for the foil thickness of 300 μg/cm^2 (Fig. 1.7). Therefore, the injection foils nominally convert 95% of the H^- into protons.

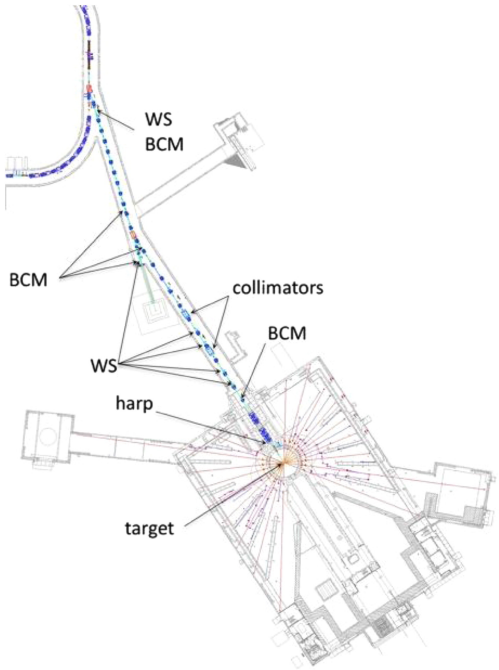

Figure 1.5: The RTBT and target building. WS: wire scanner, BCM: beam current monitor. Reproduced from [2].

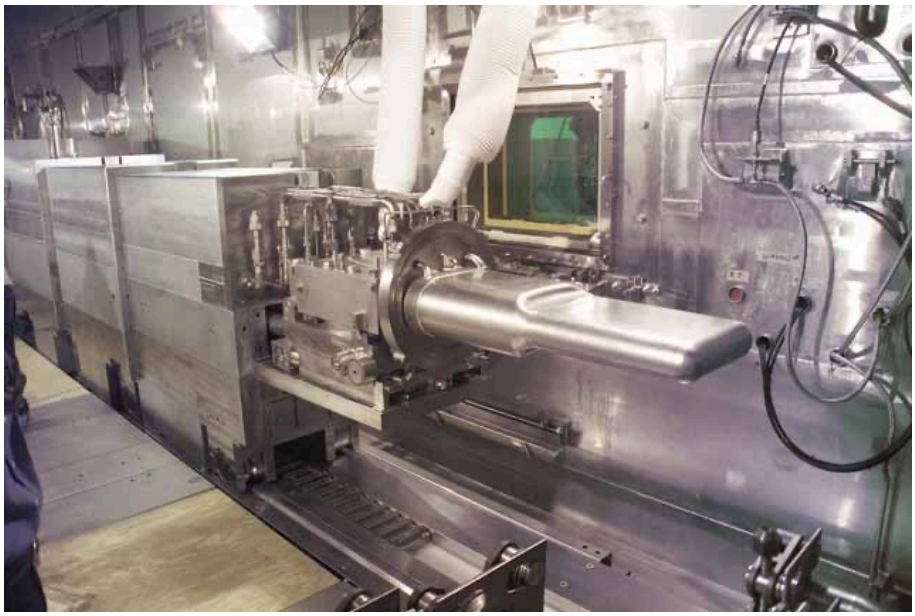

Figure 1.6: The SNS mercury target. Reproduced from [4].

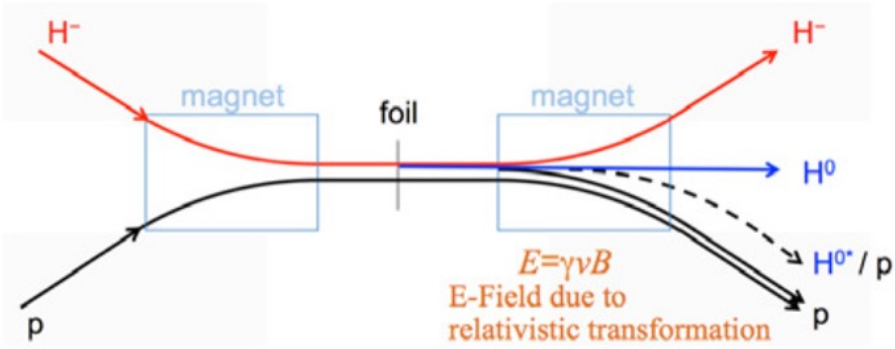

Figure 1.7: H^- CEI scheme and beam losses. The H^- beam from the linac is merged with the proton beam (p) circulating in the ring just before the injection foil. Most of the H^- are fully stripped and join the protons in the ring, but 3% are expected to be un-stripped (H^-) or partially stripped (H^0). Some of the partially stripped beam will be in an excited state (H^{0^*}), and if the electrons are stripped by the electric field in their rest frame more than a few millimeters past the foil, they will end up as losses. Reproduced from [7].

Complications that arise from injection foils include (1) beam loss from unstripped (H^-) and partially stripped (H^0) beam, (2) uncontrolled beam loss from partially stripped beam in an excited state (H^{0^*}), (3) control of stripped electrons, (4) uncontrolled beam loss and activation caused by foil scattering, and (5) foil lifetime [7]. Problems arising from (1)–(3) have been mitigated with engineering solutions. The H^- and H^0 that pass through the injection foil are converted into protons by a thicker, secondary foil and steered into the injection beam dump. If the H^{0^*} are stripped by the electric field in their rest frame due to the Lorentz-transformed magnetic field more than a few millimeters past the foil, the newly-created protons cannot enter the ring and will end up as losses. These losses were mitigated by placing the injection foil in a magnetic field at a location where the H^{0^*} would be stripped immediately or not at all. The energetic electrons leftover after stripping can become trapped in the magnetic field and come back to strike the stripper foil, the stripper foil holder, or the vacuum chamber. Damage from stripped electrons was prevented by redesigning the vacuum chamber and installing an electron catcher.

The remaining issues with injection foils are (4) and (5), and the two are closely related to the proton beam power. Foil scattering leads to uncontrolled beam loss and activation in the ring, and the frequency of scattering events scales with the proton beam power. Energy deposited in the foil by the proton beam causes it to heat up, eventually leading to buckling/wrinkling and hole-burning (Fig. 1.8). The sublimation rate of the foil grows exponentially with its temperature, and it is believed that a foil will not last more than a day in production if the temperature reached at thermal equilibrium exceeds 2000 K (Fig. 1.9). For the 1.4 MW proton beam, it is estimated that thermal equilibrium occurs between 1550 K and 1650 K [8]. Therefore, due to increased losses from foil scattering and the reduced lifetime, injection foils will no longer be a practical mode of CEI above a certain power, perhaps an order of magnitude greater than at the SNS. Studies are ongoing to precisely model foil lifetime vs. proton beam power.

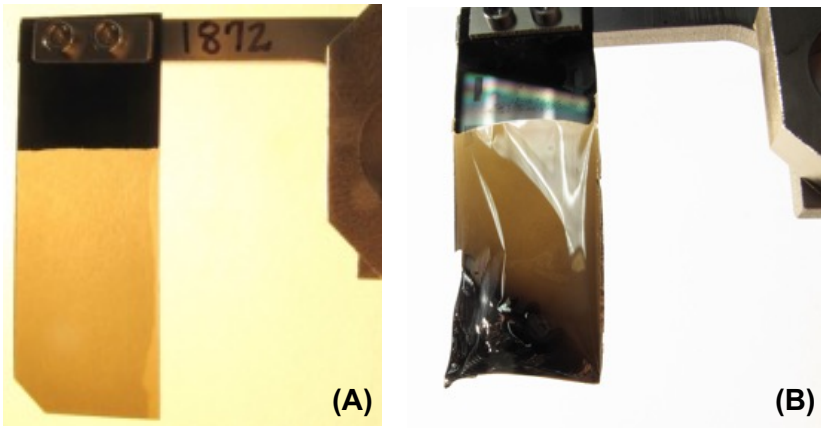

Figure 1.8: Injection foils damage from thermal effects. (A): New injection foil. (B): Damaged injection foil.

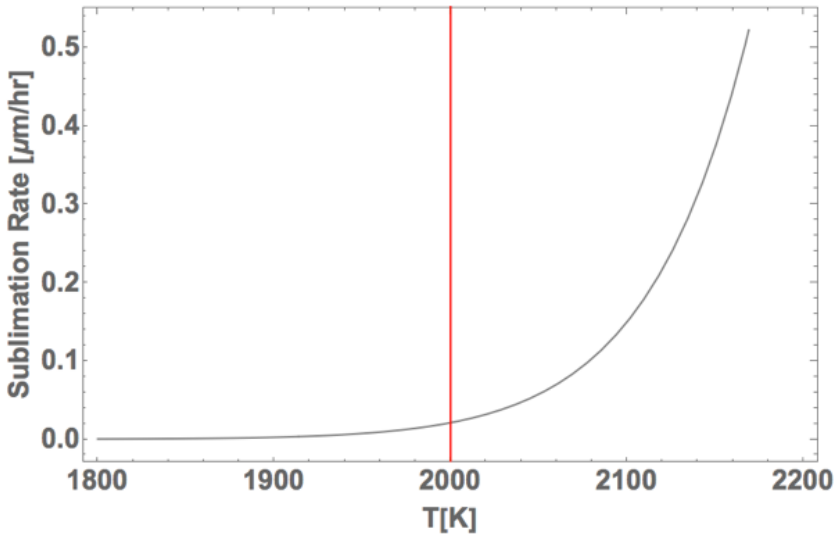

Figure 1.9: Estimated sublimation rate vs. foil temperature. The red line represents the thermal equilibrium temperature at which foils will no longer be practical.

1.2 Laser-Assisted Charge Exchange

LACE is material-free method for converting a relativistic beam of H^- into protons that is being developed at the SNS. The technique avoids many of the hazards and limitations of injection foils, and it could potentially be used as the mode of CEI into the next generation of high power proton synchrotrons.

1.2.1 Background

The history of LACE began in 1984 when Zelenskiy *et al.* proposed a method of charge exchange without foils [9]. In their method, a relativistic beam of H^- is converted into polarized protons through photodetachment by a laser. It would have been a viable method of CEI for high-intensity accumulator rings, such as at the SNS, except it would require an impractically high power laser for sufficient photodetachment. Yamane built upon the idea, and in a 1998 paper he proposed a scheme involving an ensemble of two magnets and a laser [10]. In the three-step scheme, the less tightly bound electrons are stripped by the first magnet, then the laser resonantly excites the electrons in the neutralized beam before the second magnet so that they are possible to Lorentz strip. The scheme Yamane proposed called for less laser power, but it would still be impossible to excite every electron in the bunch due to the energy spread of the ion beam, which would cause the laser photons to be out of resonance in the rest frame of the H^0 due to the relativistic Doppler effect. Danilov *et al.* proposed to solve the problem using a diverging laser beam to cover the spread of resonant frequencies in the laboratory frame [11]. Danilov's solution made high efficiency charge exchange possible for a short-duration ion beam. In 2006, the first LACE experiment was conducted at the SNS as a proof-of-principle (POP) demonstration where a 6 ns long portion of a 700 ns H^- minipulse at ~870 MeV was converted into protons with 90% efficiency [12].

The POP experiment was an important first step, but if LACE is to ever replace injection foils at the SNS, it must be as efficient or more than the foils and scalable

to the 6% duty factor of the production ion beam (1 ms macropulses at 60 Hz). For now, the laser power required to strip the 1 ms ion beam prohibits the operational feasibility of LACE. Scaling the scheme used in the POP experiment to the full production duty factor of the 1 GeV ion beam would require an average power of ∼600 kW for the 355 nm laser, which is not available with state-of-the-art technology. Ongoing experimental efforts at the SNS are focused on circumventing this problem, and several breakthroughs have been made since the POP experiment in 2006.

1.2.2 Concept

The LACE Team at the SNS uses a scheme based on the one proposed by Yamane. The three-step process of converting the H^- beam into a proton beam consists of: (1) stripping of the loosely bound electron ($H^- \rightarrow H^0$), (2) photoexcitation of the remaining electron by the laser ($H^0 \rightarrow H^{0^*}$), and (3) stripping of the excited electron ($H^{0^*} \rightarrow p$). The scheme is illustrated in Fig. 1.10. Steps (1) and (3) are accomplished when the particles pass through the fields of the two permanent magnets. The laser resonantly excites the remaining electron in the ground state to reduce its binding energy before it can be stripped by the second magnet. The efficiency of LACE is the product of the efficiency of all three steps.

Using a magnetic field to remove electrons from a relativistic particle beam is called Lorentz stripping because the electrons are stripped by the Lorentz force from the electric field that is experienced in the rest frame of the particles. The electric field in the rest frame of a particle E due to the magnetic field in the laboratory frame B is

$$E = \beta\gamma cB, \tag{1.1}$$

where β is the ratio of particle velocity to the speed of light c, and $\gamma = 1/\sqrt{1-\beta^2}$. An empirical formula which gives the lifetime τ of an H^- ion in an electric field was

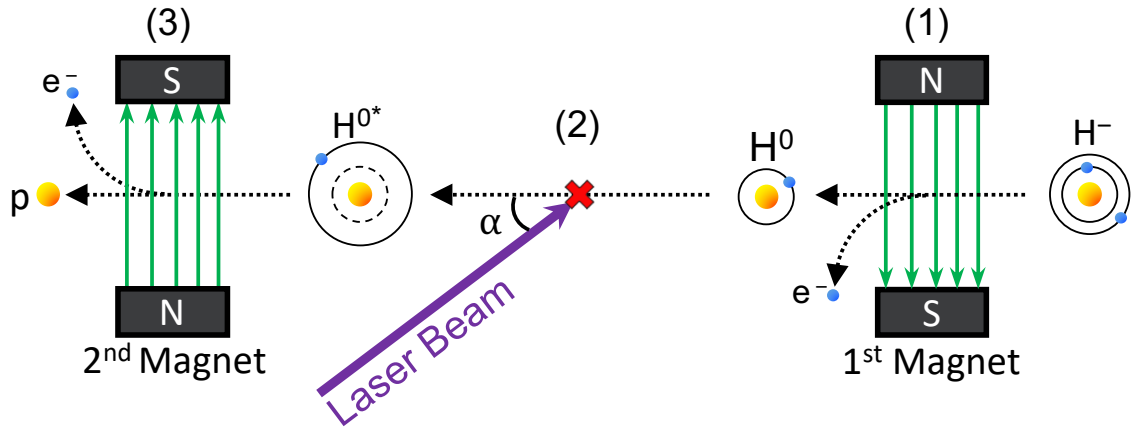

Figure 1.10: The 3-step LACE processes: (1) Lorentz stripping of loosely bound electron, (2) photoexcitation by the laser, and (3) Lorentz stripping of the final, excited electron.

found by Scherk [13] and Jason *et al.* [14]. That formula is

$$\tau = \left(\frac{A_1}{E}\right) \exp\left(\frac{A_2}{E}\right), \tag{1.2}$$

where $A_1 = 2.47 \times 10^{-6}$ V·s/m ($\pm 4\%$) and $A_2 = 4.49 \times 10^9$ V/m ($\pm 0.25\%$). The probability of Lorentz stripping of the H^- is the inverse of the lifetime (τ^{-1}). In his paper, Yamane compared the probability of Lorentz stripping for H^- and H^0 and showed they were identical at a certain electric field strength if the H^0 is in an excited, $n = 3$ electronic state (Fig. 1.11). Therefore, identical magnets could be used to provide the fields needed for steps (1) and (3), as long as the electrons are excited to $n = 3$ in step (2). The electric field strength for which the probabilities of Lorentz stripping are equal is close to $E = 6 \times 10^8$ V/m and corresponds to a magnetic field strength of $B = 1.1$ T in the laboratory frame for a 1 GeV beam.

The electrons of the neutralized beam (H^0) are assumed to be in the ground state ($n = 1$) after Lorentz stripping of the first electron. The laser excites the electron to $n = 3$ via the Rabi oscillation, a processes where a bound electron oscillates between two quantum states when driven by a harmonic potential (in this case the laser field). The resonant excitation $n = 1 \rightarrow 3$ corresponds to a reduction in binding energy of 12.1 eV and requires a photon of frequency $\nu = 3 \times 10^{15}$ Hz ($\lambda = 100$ nm) to produce. The relativistic Doppler effect is used to produce the resonant photons in the rest frame of the H^0 from laser light that is $\lambda = 355$ nm in the laboratory frame. The relativistic Doppler equation gives the frequency of light in the H^0 rest frame due to the Lorentz transformation. It is given by

$$\nu = \gamma(1 + \beta \cos\alpha)\nu_0 \tag{1.3}$$

where ν and ν_0 are the laser frequencies in the H^0 rest frame and laboratory frame, respectively, and α is angle at which the laser and ion beams cross. The value of α is chosen based on the energy of the particles to produce the desired frequency in the

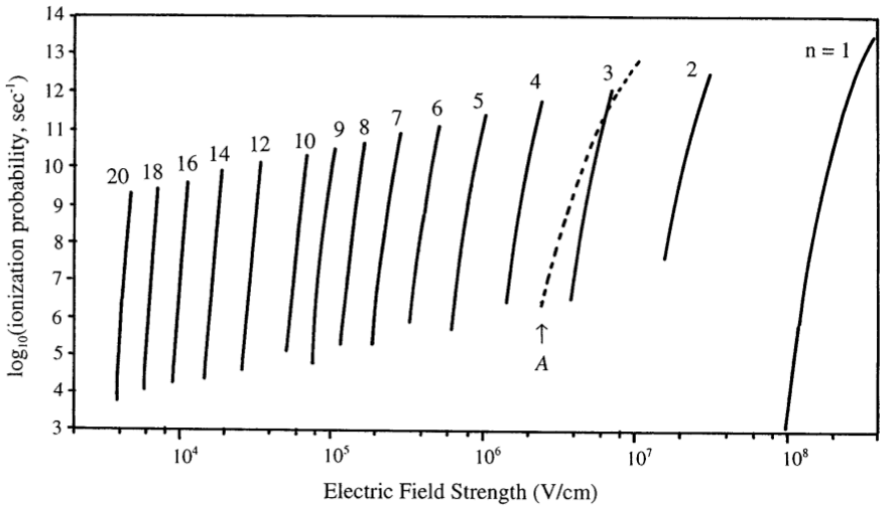

Figure 1.11: Probability of Lorentz stripping for H^0 and H^-. Probabilities for H^0 are shown for electronic energy levels ranging from $n = 1$ to $n = 20$. The probability for H^- is indicated by "A". Reproduced from [10].

rest frame of the H^0. The first four energy levels of H^0 and the Doppler shift for a 1 GeV hydrogen atom are shown in Fig. 1.12. The lifetime of the $n = 3$ state is $\tau = 5.3 \times 10^{-9}$ s [10], and the hydrogen particles will be Lorentz stripped with high probability if they are excited when they pass through the second magnetic field.

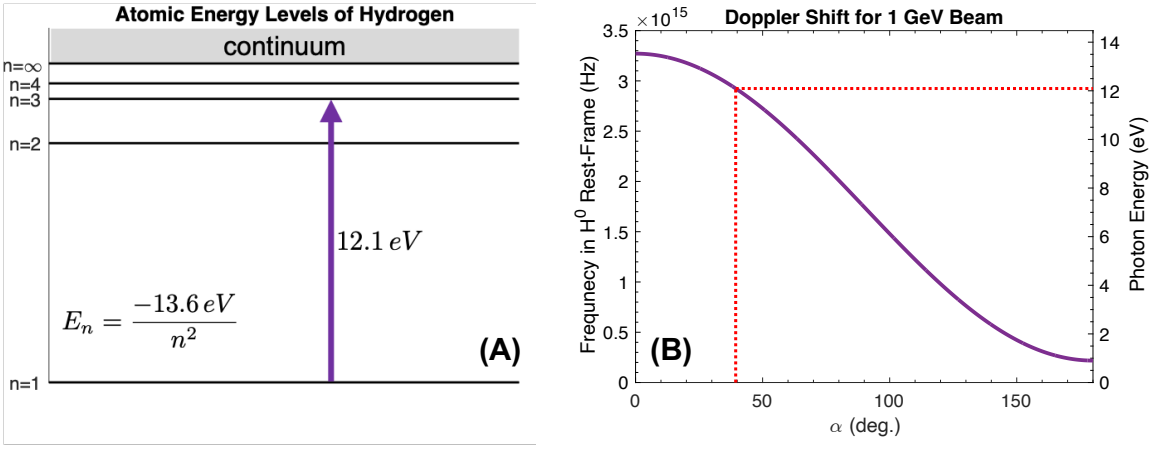

Figure 1.12: Atomic energy levels of H^0 and Doppler shift for a 1 GeV beam. (A): The first four energy levels of hydrogen and photon energy required for the $n = 1 \rightarrow 3$ excitation. (B): Frequency of the 355 nm laser beam in the rest frame of the H^0 and α required for resonance.

Chapter 2

The 10 μs Experiment

It was stated in the Introduction that the goal of the LACE project at the SNS is to strip the full duty factor ion beam with efficiency equivalent to, or greater than, that of the foils. The primary obstacle is the generation of the high-energy UV laser pulses that would be required to strip the entire 1 ms H^- macropulse. In [12], three ideas were proposed to increase the duty factor of the stripped beam with the limited laser power available. Those methods were employed in a 2016 proof-of-practicality experiment in which 10 μs long pulses were converted into protons with high efficiency, a factor of 1000 times longer than in the POP experiment [15, 16]. This chapter details the experimental design, laser power saving schemes, and results of the 10 μs Experiment. Results showed the highest recorded LACE efficiency was >95%, but the distribution of all efficiencies measured shows large variance that was caused by instability of critical parameters of the experiment.

2.1 Experimental Design

The experimental design of LACE at the SNS was changed after the POP experiment in 2006. LACE now takes place in the straight section of the HEBT before the ring injection area (Fig. 2.1). A vacuum chamber was specially designed for the 10 μs Experiment and installed around the ion beamline between two quadrupole magnets

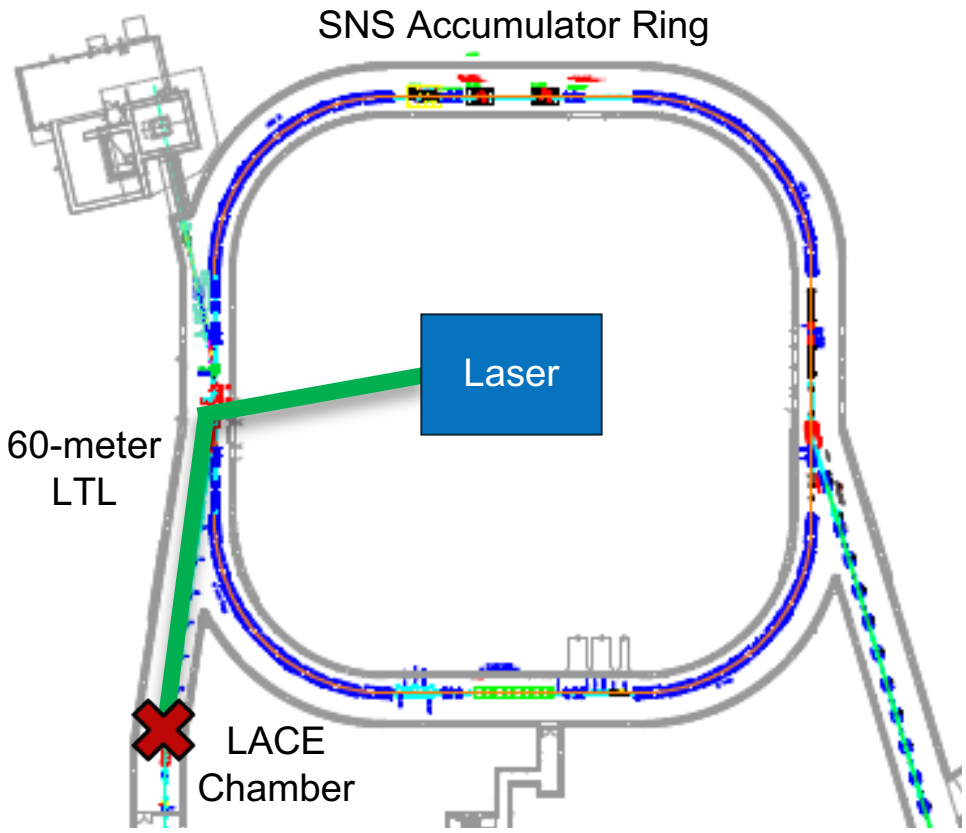

Figure 2.1: Footprint of the LACE experiment. The laser is housed in the above-ground Ring Service Building (blue) and the beam is transported through a 60-meter laser transport line (green) to the LACE chamber in the HEBT.

(Fig. 2.2). The LACE chamber contains the magnets and has vacuum windows for the laser beam to enter at the design α. The beam current monitor (BCM) after the chamber is used to measure the current of the proton beam emerging from the chamber. The wire scanner measures the position of the ion beam at the laser-H^0 interaction point (IP). The two permanent magnets are in a Halbach cylindrical array [17] and are designed to accommodate the SNS ion beam and requirements for LACE (Fig. 2.3). The magnetic field strength within the aperture of an infinitely long Halbach cylindrical array is

$$B = B_r \cdot \ln \frac{R_2}{R_1} \tag{2.1}$$

where R_1 and R_2 are the inner and outer radius of the magnet and B_r is the remnant field of the permanent magnet material. In this case, $B_r = 1.2$ T and R_1 and R_2 are equal to 14.5 mm and 60 mm, respectively. The magnetic field strength and gradient along the beam path within the LACE chamber is plotted in (Fig. 2.4). The magnets are mounted on a retractable flange that is lowered during regular accelerator operations so as not to interfere with neutron production.

The laser uses a master oscillator power amplifier (MOPA) configuration to produce a pulsed, UV ($\lambda = 355$ nm) beam at 10 Hz with megawatt peak-power [18]. The four main components are the fiber seed laser, pulse picker, solid-state amplifiers, and two harmonic conversion crystals (Fig. 2.5). The seed laser is the "front-end" of the laser system, and it can be interchanged with other fiber lasers that produce $\lambda = 1064$ nm light. The pulse picker is an acousto-optic modulator (AOM), which modulates the seed beam with bursts of tunable length at a repetition rate of 10 Hz. The amplification stage consists of six Nd:YAG rods in series, each pumped by a xenon flash lamp. The flash lamps are triggered at the same rate as the AOM so that one laser pulse is amplified per shot of the flash lamps. The final stage is generation of the third harmonic of the IR beam through a nonlinear process by two lithium triborate (LBO) crystals. The first LBO crystal converts the fully

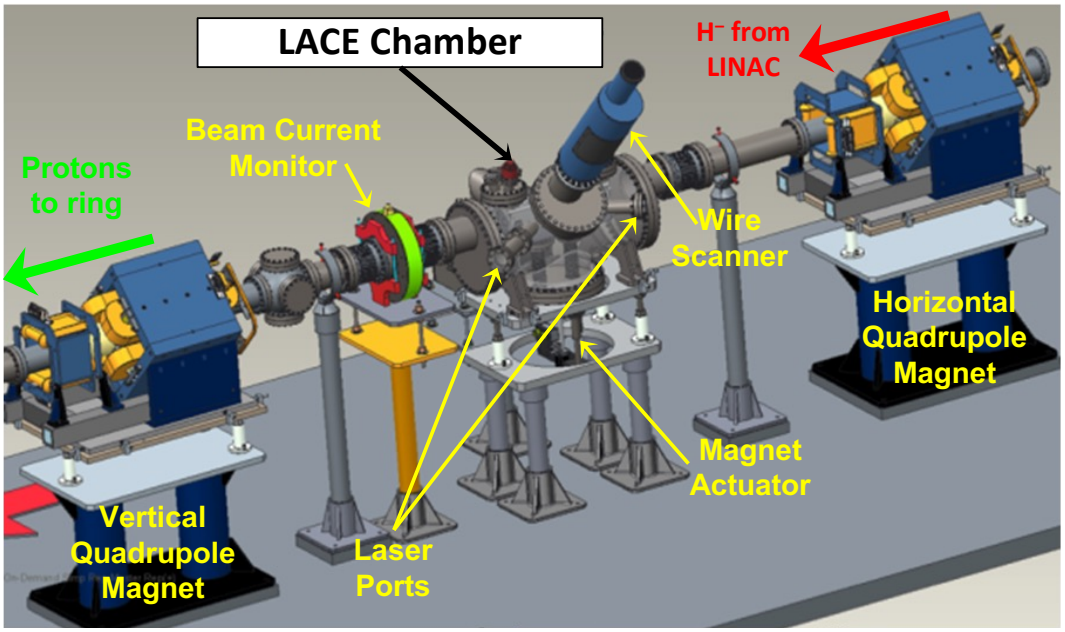

Figure 2.2: The LACE chamber and associated diagnostic instruments. Vacuum ports allow the laser beam to enter at the correct angle. The wire scanner is used to align the laser and ion beams. The beam current monitor measures the current of the stripped beam. Actuator retracts stripping magnets during neutron production.

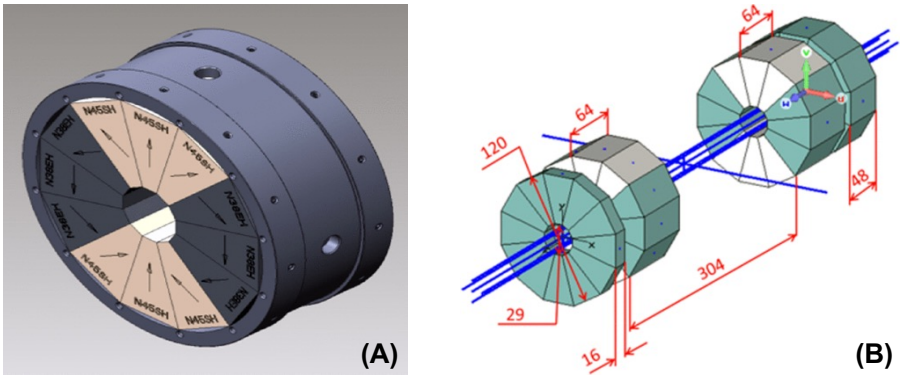

Figure 2.3: Stripping magnets inside the LACE chamber. (A): Individual magnet in the Halbach array. The color corresponds to permanent magnet material used, and the arrows show magnetization direction. (B) Arrangement of the two magnets inside the LACE chamber; dimensions are in millimeters. Reproduced from [16].

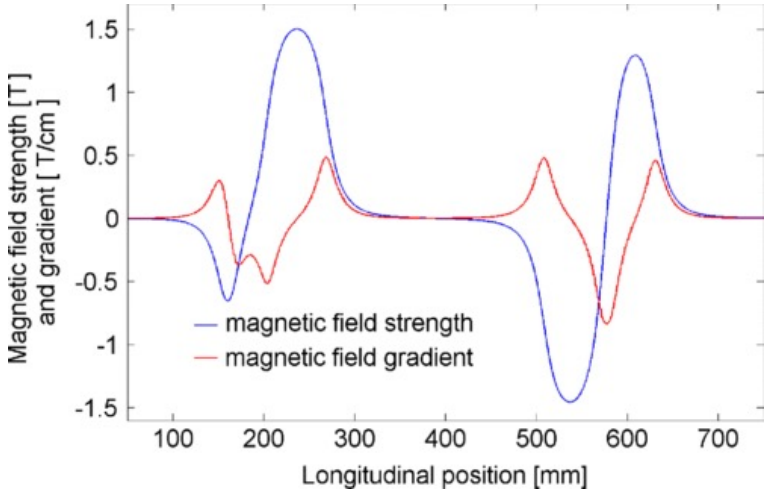

Figure 2.4: The magnetic field strength and gradient along the beam path inside the LACE chamber. Reproduced from [16].

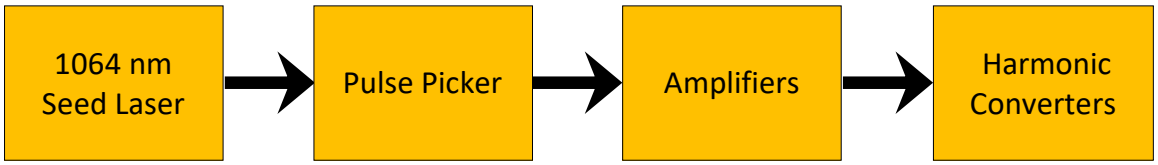

Figure 2.5: Block diagram of the MOPA scheme for the UV laser.

amplified 1064 nm light to 532 nm (green), and then the second crystal converts the green light to 355 nm UV light. After harmonic conversion the beam is sent to the local optics tables just before the LACE chamber.

The section of the accelerator tunnel where the LACE chamber is located is a high radiation area (0.8 – 1 rem/hour). In the absence of a radiation-shielded laser room built into the tunnel, the decision was made to house the laser system in the above-ground Ring Service Building (RSB) and retrofit a 60-meter, free-space laser transport line (LTL) around the existing infrastructure to transport the laser beam to the IP (Fig. 2.6). A 20 m cable chase leading from the RSB to the ring injection area was repurposed for the LTL, and a periscope-tower was constructed to guide the beam around the ring injection area. The rest of the LTL is enclosed by a 6" diameter steel pipe connecting the laser table in the RSB to Local Optics Table 1 (LOT1) before the LACE chamber. A total of 8 mirrors are required to traverse the 7 bends at oblique angles, and the rigid structure extends over the two separate building foundations of the RSB and HEBT tunnel. Although its entrance and exit are sealed with vacuum windows, the LTL cannot be efficiency made a vacuum because the vacuum pressure in the chase section will cause dirt and other debris to be sucked in. Relay imaging is also prohibited by the inaccessibility of critical parts of the LTL and difficulty fitting the necessary optics inside the structure. Thus, a collimated, non-Gaussian beam must be propagated most of the 65-meter to the LACE chamber.

2.2 Laser Power Saving Schemes

The first laser power saving scheme was to address Doppler spreading due to the energy spread of the ion beam in a more fundamental way than with the diverging laser beam used in [12], which was not an economical use of laser power. Danilov *et al.* derived an expression that would result in $d\nu/d\gamma = 0$ at the IP as a function of design parameters [11]. The angle the laser beam crosses the trajectory of an off-momentum

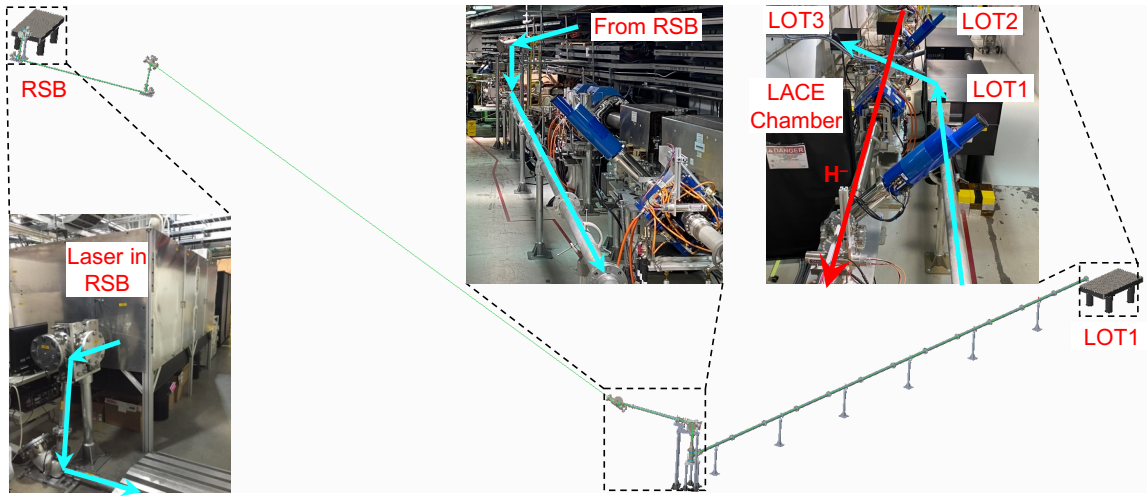

Figure 2.6: The 60-meter LTL used to transport the UV laser beam to the LACE chamber.

H^0 particle in the horizontal plane is given by

$$\alpha = \alpha_0 - x', \tag{2.2}$$

where α_0 is the angle for the reference energy particle, $x' = D'dp/p_0$ for a particle with momentum deviation dp/p_0, and D' is the derivative of the dispersion function. Taking the derivative of Eq. (1.3) with respect to γ gives

$$\frac{d\nu}{d\gamma} = \left(\frac{\beta + \cos\alpha}{\beta} - \frac{d\alpha}{d\gamma}\gamma\beta\sin\alpha \right) \nu_0, \tag{2.3}$$

where from $d\alpha = -x'$ we get

$$\frac{d\alpha}{d\gamma} = -D'/\gamma\beta^2. \tag{2.4}$$

Plugging Eq. (2.4) into Eq. (2.3) gives

$$\frac{d\nu}{d\gamma} = \left(\frac{\beta + \cos\alpha}{\beta} + D'\frac{\sin\alpha}{\beta} \right) \nu_0, \tag{2.5}$$

and $d\nu/d\gamma = 0$ for

$$D' = -\frac{\beta + \cos\alpha}{\sin\alpha}. \tag{2.6}$$

For the 1 GeV beam and the design laser-H^0 crossing angle of $\alpha = 37.5°$, the value of D' for which $d\nu/d\gamma = 0$ at the IP is D' = -2.74 rad. The effect dispersion tailoring had on the required peak laser power and the expected LACE efficiency was analyzed using a computational model of the laser-H^0 interaction [19]. Fig. 2.7 shows the peak laser power required for 90% stripping is reduced as D' approaches the optimal value and that the expected LACE efficiency is expected to increase by about ∼15%.

The second laser power saving scheme was to optimize the size and divergence of the ion beam. After dispersion tailoring, the residual frequency spread is due to the divergence of ion beam in the horizontal plane, given by

$$\Delta x' = \sqrt{(1 + \alpha_h)/\beta_h}, \tag{2.7}$$

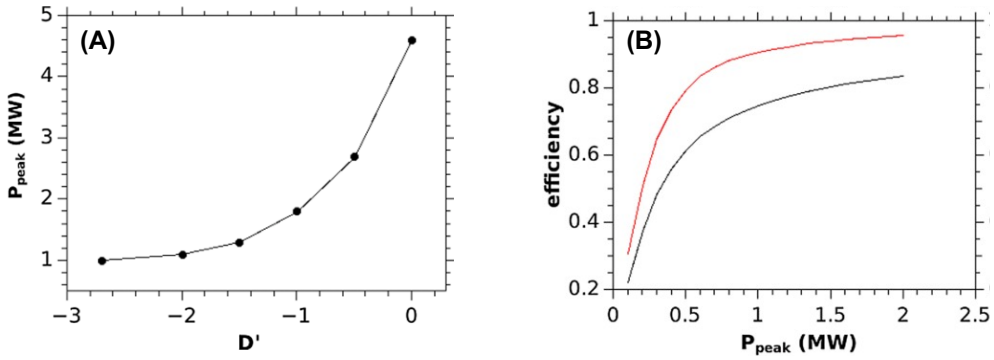

Figure 2.7: The effect of dispersion tailoring. (A): Peak laser power required for 90% stripping as a function of D'. (B): LACE efficiency as a function of peak laser power with (red) and without (black) dispersion tailoring. Reproduced from [16].

where α_h and β_h are the horizontal Twiss parameters. During the experiment, $\Delta x'$ was minimized by using magnets upstream to make $\alpha_h \approx 0$ and β_h very large. Since the photoexcitation efficiency is proportional to photon density, sufficient excitation can be achieved with less power if the laser spot is focused to a small size at the IP. Of course, the laser pulse must still overlap the H^0 if they are to be excited, so the longitudinal and vertical size of the ion beam were minimized to also shrink the laser pulse. The ion beam is guaranteed to be fully overlapped by a 50 ps (FWHM) laser pulse if $\sigma_l < 12.5$ ps and $\sigma_v < 1$ mm. The photoexcitation efficiency for different configurations of the laser beam size and divergence at given in Fig. (2.8). Optimization of the ion beam size and divergence resulted in a factor of 2–5$\times$ reduction in average laser power.

The third and last laser power saving scheme implemented in the 10 μs Experiment was on the laser side. Recall the macropulse temporal structure of the ion beam in Fig. 1.3. No stripping can occur in the gaps between micropulses, and so the pulse length of the laser was extended to 10 μs without sacrificing peak power by matching its temporal structure to the ion beam at the micropulse level. Fig. 2.9 shows the temporal structures of the laser beams used in the POP and 10 μs Experiments compared to the ion beam. The two-level temporal structure of the laser beam was formed using a mode-locked seed laser that produced a 1064 nm beam of 35–50 ps pulses at 402.5 MHz and then chopping the beam with the AOM to create 10 μs minipulses at 10 Hz. The IR beam was then amplified to 8 MW peak-power and converted to its third harmonic ($\lambda = 1064$ nm $\rightarrow$ 355 nm). After harmonic conversion, the UV beam would have a peak power of 2 MW. Fig. 2.10 shows the architecture of the MOPA laser for the 10 μs experiment. Table 2.1 lists relevant laser parameters used in the experiment. The working principle and configuration of the laser system for the 10 μs Experiment are described in detail in [20]. Temporal matching of the laser and H^- micropulses resulted in a further 70$\times$ reduction in the average laser power required for high efficiency.

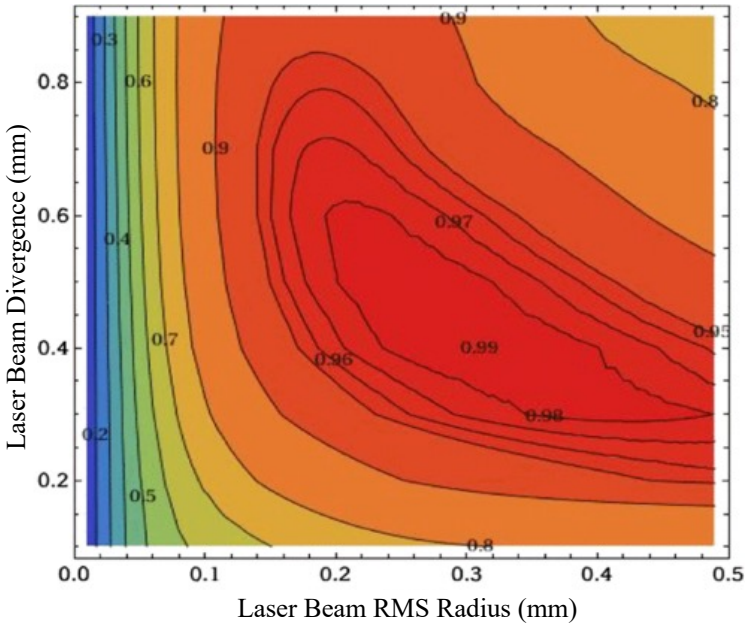

Figure 2.8: Topological map of photoexcitation efficiency for the design ion beam parameters as a function of laser beam size and divergence at the IP. Reproduced from [16].

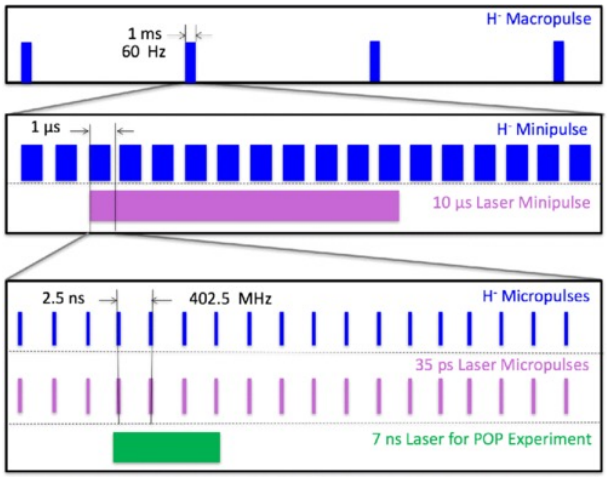

Figure 2.9: Temporal structure of the ion beam and laser beam from the POP and the 10 μs Experiments. Reproduced from [15].

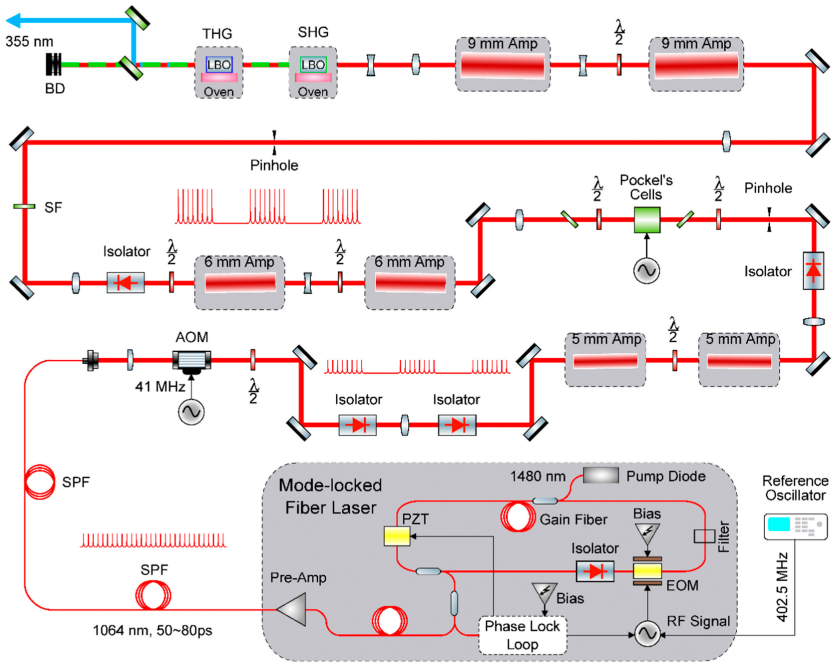

Figure 2.10: Laser architecture for the 10 μs Experiment. Reproduced from [20].

Table 2.1: Required vs. delivered laser parameters.

Parameter	Required	Delivered
Macropulse length (μs)	10	12
Micropulse length (ps)	>30	30-50
Laser transport line efficiency (%)	>60	70
Max. peak power at IP (MW)	>1.0	2.0
Full vertical divergence at IP (mrad)	1.6–2.6	2.3–3.0
Vertical beam size (4σ) at IP (mm)	1.0–1.4	1.1

2.3 Results

The laser power saving schemes discussed above resulted in 95% efficiency stripping for a 10 μs long portion of the H^- macropulse at 1 GeV. The efficiency of LACE is calculated by measuring the beam current before and after stripping. The LACE efficiency ε is the ratio

$$\varepsilon = \left| \frac{I_p}{I_{H^-}} \right|, \tag{2.8}$$

where I_p and I_{H^-} are the proton and H^- beam currents, respectively. Fig. 2.11 shows 8 superimposed waveforms of protons beam current and the averaged H^- beam current. The laser pulse stripped a 10 μs portion of 11 μs H^- macropulses. In the analysis, the amplitude of each minipulse seen in Fig. 2.11 was averaged, and the charge exchange efficiency for each minipulse was calculated using Eq. (2.8). The maximum charge exchange efficiencies from Fig. 2.11 were around 95%.

The 10 μs Experiment was a successful demonstration of the practicality of LACE for the production ion beam. The H^- beam was stripped to protons with efficiency equivalent to the foils, and the macropulse length was extended by over three orders of magnitude since the POP experiment. The laser power limitation is the fundamental obstacle to overcome in the development of LACE, but there are other challenges that prohibit use of LACE as an operational system, regardless of whether or not the full duty factor ion beam can be stripped. In the experiment, a total of 456 macropulses were stripped to protons, each containing 10 minipulses. The proton waveforms shown in Fig. 2.11 represent the best data from the experiment. The LACE efficiencies for all 4560 minipulses measured in the experiment are shown in Fig 2.12. It is clear that 95% stripping does not represent the typical results of the experiment. For LACE to replace injection foils at the SNS, not only does LACE need to be capable of stripping the full duty factor ion beam (1 ms macropulses at 60 Hz), it must strip *every* pulse with $\geq$95% efficiency.

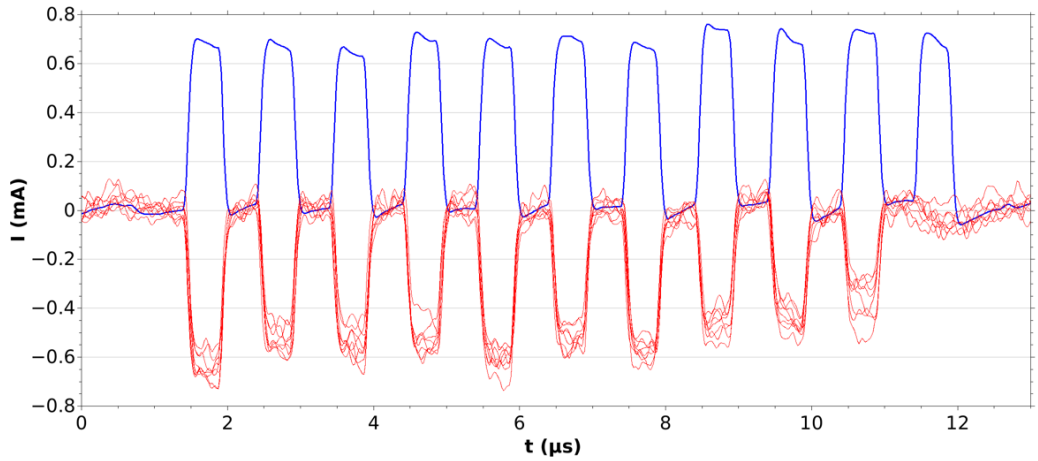

Figure 2.11: Experimental results of the 10 μs LACE experiment. Eight shots of proton signal (red) are shown relative to the H^- signal (blue). Reproduced from [15].

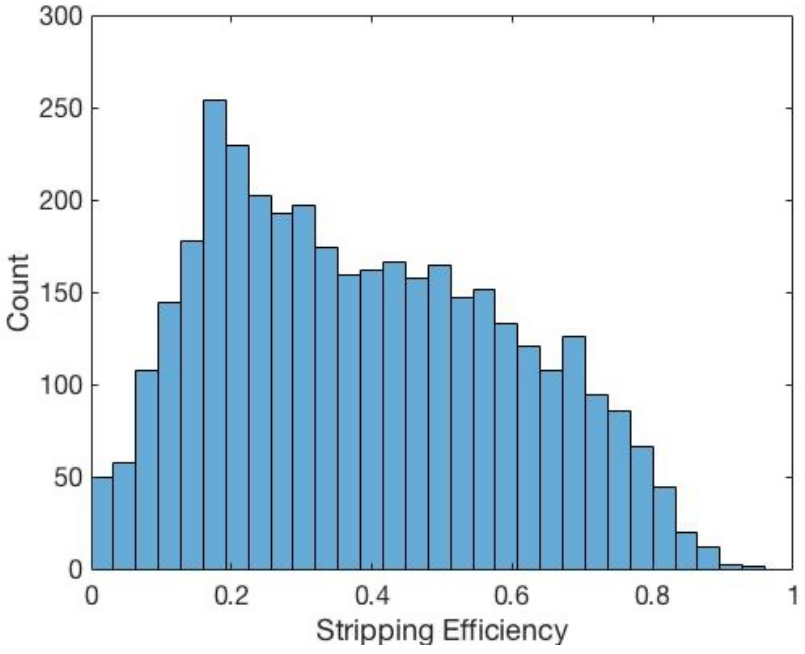

Figure 2.12: Distribution of all efficiencies measured in the 10 μs Experiment. Most pulses were stripped with below 95% efficiency. Reproduced from [16].

Chapter 3

The Cause of Stripping Instability

The overall efficiency of LACE is the product of the efficiency of each of the three steps. Lorentz stripping by the permanent magnetic fields in steps (1) and (3) is very stable and close to 100% efficient, so it is believed that instability in step (2), photoexcitation by the laser, is the source of variance in LACE efficiency seen in Fig. 2.12. The primary parameters in step (2) that the photoexcitation efficiency is most sensitive to are the peak laser power delivered to the IP and the alignment of the two beams. The 10 μs Experiment was designed so that there would be more photons in the core of the laser pulse than are necessary to excite every particle. This was done so that pulse-to-pulse fluctuations in laser power would not lead to insufficient photoexcitation if the laser beam is properly aligned to the ion beam. Additionally, the laser micropulses are longer than the ion beam micropulses, so there should be good overlap in the temporal-domain as well. This leaves instability in the horizontal crossing angle and imperfect spatial overlap of the ion beam by the laser as the most likely causes of the variance in LACE efficiency observed in the 10 μs Experiment. In this chapter, stability studies of the laser and ion beam are detailed, and it is concluded that laser pointing instability at the IP was the cause of variance in LACE efficiency. Additional studies showed the laser pointing instability is sourced primarily from the 65-meter transport through the free-space LTL rather than the laser itself.

3.1 Laser and Ion Beam Stability Studies

Since it is the relative alignment of the laser and ion beams that is important, instability of either could cause variance in the LACE efficiency. The spatial stability of the ion beam was measured using a beam position monitor (BPM) upstream of the LACE chamber. The 2-dimensional position of each of the 10 minipulses in a 10 μs ion beam macropulse was measured for 100 macropulses. The mean position of each minipulse and 95% confidence interval are shown in Fig. 3.1. Each minipulse is expected to be within ± 0.1 mm of the design trajectory. For the transverse size of the laser beam given in Table 2.1, the ion beam should be fully overlapped at the IP if the laser is assumed to be aligned correctly. Thus, ion beam instability does not account for the large spread of efficiencies observed in the 10 μs Experiment.

Laser power fluctuations and ion beam instability cannot explain the large variance in the LACE efficiency. That leaves laser pointing instability as the last remaining suspect. After the 10 μs Experiment was concluded, additional diagnostics were installed to monitor the position and angle of the laser beam at the IP for the next phase of LACE experiments. In 2022, a study was done to correlate the position and angle of the laser beam at the IP with LACE efficiency. A LACE experiment was set up for the study using different ion and laser beam parameters than the 10 μs Experiment, and it was not optimized for high efficiency stripping. Still, the data recorded in the experiment was adequate to show the sensitivity of the LACE efficiency to laser alignment. Twenty minutes of LACE efficiency data was measured simultaneously with the position and angle of the laser beam at the IP. The waveforms in Fig. 3.2 clearly show the LACE efficiency is highly correlated with the motion of the laser beam. The horizontal angle and vertical position of the laser beam at the IP are the most sensitive laser parameters due to the laser and ion beams crossing in the horizontal plane and the reliance on the relativistic Doppler shift to produce the resonant photons. Fig. 3.3 shows the LACE efficiency sensitivity to those parameters. The colormap clearly shows a region where the efficiency is optimized and that the

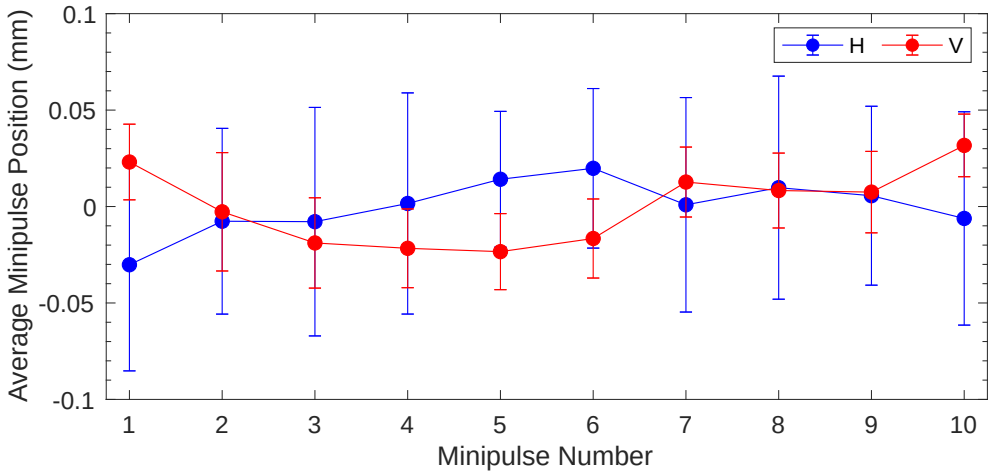

Figure 3.1: Ion beam stability measured with a BPM upstream of the LACE chamber. Error bars represent the 95% confidence interval.

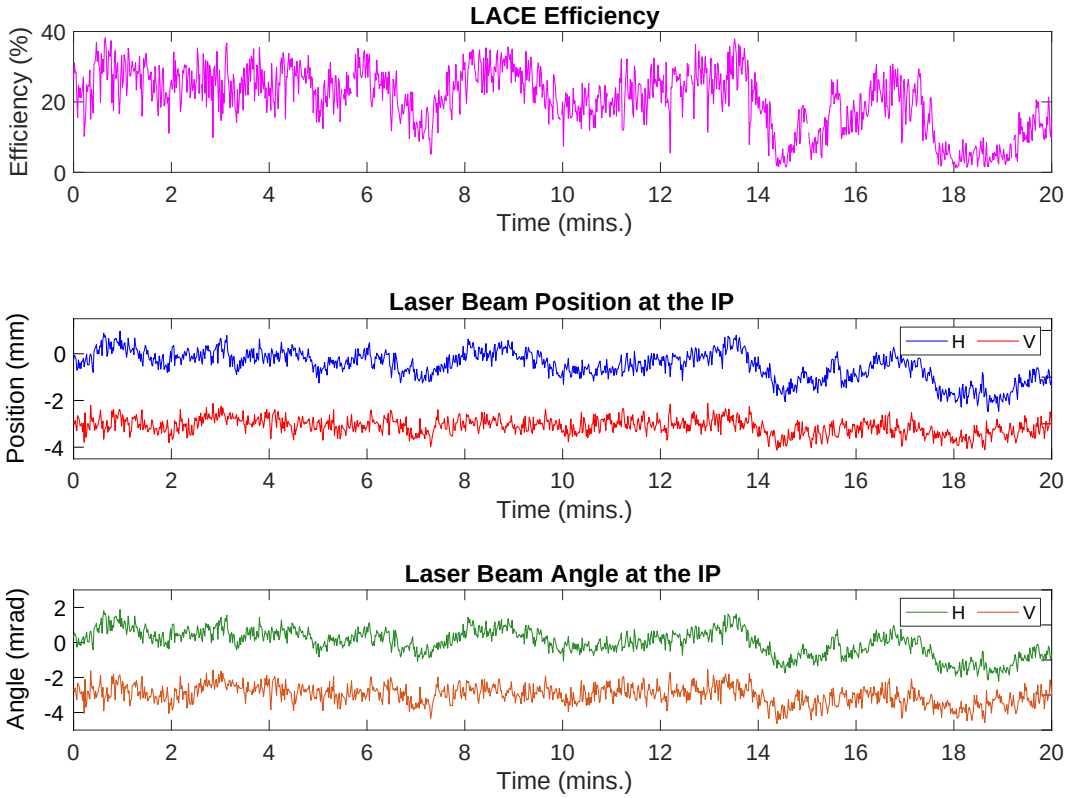

Figure 3.2: LACE efficiency compared to laser pointing stability over 20 minutes. Efficiency is highly correlated with the motion of the laser.

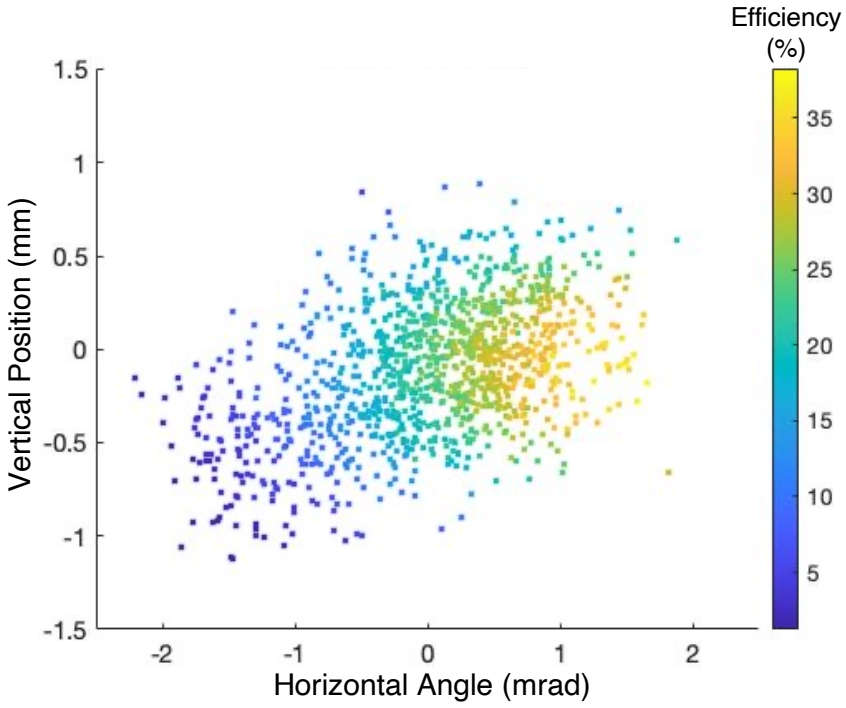

Figure 3.3: Correlation of LACE efficiency with the sensitive laser parameters. The colormap shows a region where the efficiency was maximized and the efficiency falls off with distance from that region.

lower than expected efficiencies were the result of the laser beam wandering out of alignment. It also shows the value of each efficiency data point is very close to that of its neighbors in the parameter space of the plot, suggesting there are no other parameters besides the horizontal angle and vertical position of the laser beam whose fluctuations are causing significant variance in the LACE efficiency. These studies conclude that laser pointing instability at the IP was the primary source of the large variance in LACE efficiency.

3.2 Sources of Laser Pointing Instability

The pointing instability of the laser can be sourced from both its architecture and the LTL used to transport the beam to the LACE chamber. Thermal effects from the environment can alter the alignment of internal optics and, combined with thermal lensing in the six Nd:YAG rods pumped by flash lamps, lead to some combination of shift and tilt of the output beam [21]. The relative alignment of the two LBO crystals was found to have an especially strong effect on the short-term fluctuations of the beam pointing and intensity profile. The retrofit on the LTL is also non-ideal for transporting a high-energy UV beam over 65-meters. The LTL is constructed from several sections of stainless steel pipe and is rigidly bolted to the floor. As a result, seismic vibrations are transmitted from the ground and cause the transport mirrors to vibrate. Additionally, any relative motion between the foundations of the RSB and ring tunnel will result in some degree of laser misalignment at the IP. The LTL is also not a vacuum, so dust and condensation accumulate inside the structure, and air currents caused by the temperature gradient between the RSB and tunnel effect the beam as it propagates. High energy beams are also vulnerable to thermal lensing in the atmosphere, which results in the profile becoming distorted or breaking up after a long enough distance.

Studies were done to identify whether the pointing instability originates inside the laser itself or is the result of propagation through the LTL. In the following analysis,

the motion of the laser beam will be broken down into its fast and slow components. The fast component of the motion will henceforth be called *jitter*, and the slow motion will be referred to as *drift*. The change in position or pointing angle that occurs between consecutive laser pulses (at 10 Hz) will be defined specifically as *pulse-to-pulse jitter*. Jitter is caused by high frequency disturbances, such as mechanical vibrations or air currents, and drift is mostly due to thermal effects, such as expansion of transport optics or internal components of the laser.

3.2.1 Laser Beam Drift Analysis

The position and angle of the laser beam were measured for a period of 1 hour before and after the LTL, on the laser table in the RSB and on LOT1 (Fig. 3.4). It is immediately clear that there was significantly more drift observed after the LTL than before it. The colormap shows the average position and angle of the beam was drifting significantly over the course of the measurement on LOT1, but in the RSB there was relatively little drift compared to the amplitude of the jitter. These data can also be analyzed in the frequency domain by examining the Fourier amplitude spectrum. Fig. 3.5 shows the motion of the beam picked up significant low frequency components after propagating through the LTL, which dominate the spectrum on LOT1. In the position and angle spectra from before the LTL, there is a clear signal at around 3 mHz, but the motion is dominated by components within the frequency band between 10 mHz and 100 mHz . These signals are still present in the spectra after propagation through the LTL, however the low frequency components dominate. It can be concluded from this analysis that slow frequency disturbances are picked up in the LTL, and thus the LTL is the primary source of drift.

3.2.2 Laser Beam Jitter Analysis

It is also desirable to to know if the LTL contributes to the jitter as well. The jitter component of the data in Fig. 3.4 was isolated by detrending the horizontal and

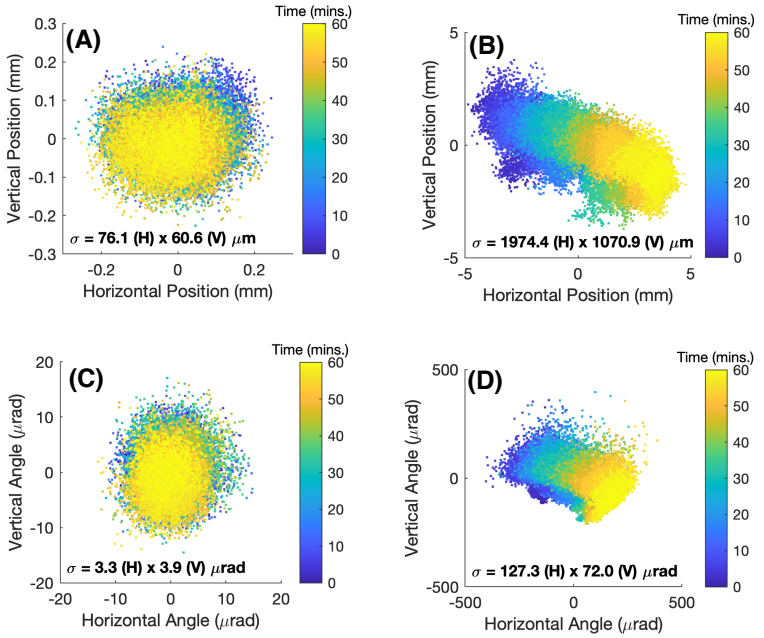

Figure 3.4: Laser pointing stability before and after the LTL measured over 1 hour at a 10 Hz sampling rate. Top: the position of the laser beam at locations (A) before and (B) after the LTL. Bottom: pointing angle of the beam (C) before and (D) after the LTL. The RMS value of the data is represented by σ.

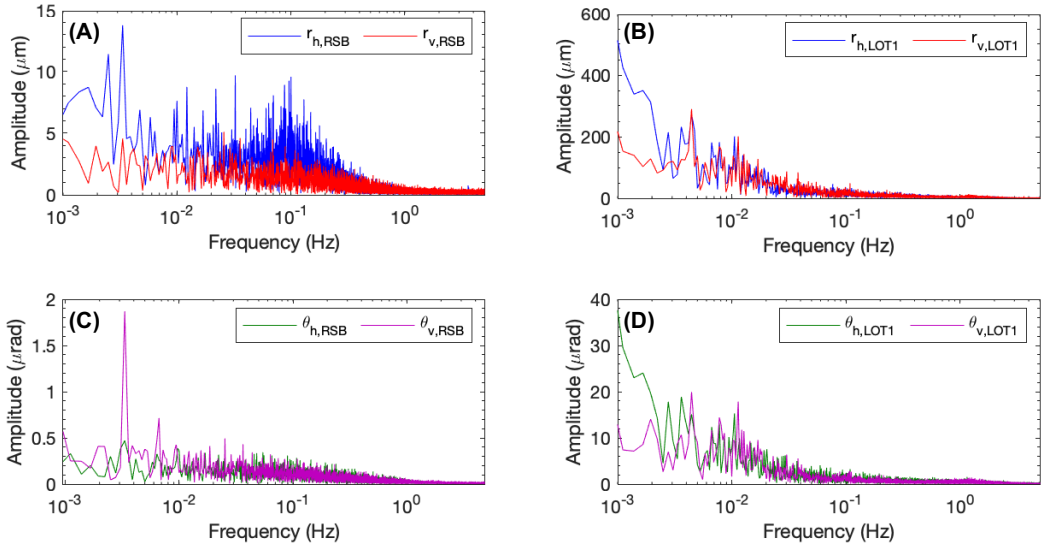

Figure 3.5: Fourier amplitude spectra. Top: the position spectrum (A) before and (B) after the LTL. Bottom: pointing angle spectrum (C) before and (D) after the LTL.

vertical components of the time series with separate linear fits. Then, the pulse-to-pulse jitter was calculated in each dimension. The pulse-to-pulse jitter in the horizontal dimension is given by

$$\begin{pmatrix} \Delta r_{h,k} \\ \Delta\theta_{h,k} \end{pmatrix} = \begin{pmatrix} r_{h,k} \\ \theta_{h,k} \end{pmatrix} - \begin{pmatrix} r_{h,k-1} \\ \theta_{h,k-1} \end{pmatrix}, \tag{3.1}$$

where $r_{h,k}$ and $\theta_{h,k}$ are the horizontal position and angle of the beam at discrete-time index k, and an equivalent equation exists for the vertical dimension. To simplify the following analysis, only the magnitude of the pulse-to-pulse jitter will be considered. The magnitude of the pulse-to-pulse jitter is then given by

$$\Delta r_k = \sqrt{(\Delta r_{h,k})^2 + (\Delta r_{v,k})^2} \tag{3.2a}$$

$$\Delta\theta_k = \sqrt{(\Delta\theta_{h,k})^2 + (\Delta\theta_{v,k})^2}. \tag{3.2b}$$

The magnitude of the pulse-to-pulse jitter recorded over the hour is plotted in Fig. 3.6. If the jitter originates from the laser, there should be a strong correlation between the magnitude of the pulse-to-pulse jitter in the RSB and on LOT1. Fig. 3.7 shows there is a very weak correlation between the pulse-to-pulse jitter measured at both locations, suggesting most of the jitter originates from the LTL as well. The autocorrelation function can also be used to analyze the correlation in the time domain. Fig. 3.8 shows the autocorrelation functions for $\Delta r_{k,RSB}$ and $\Delta\theta_{k,RSB}$ show there is a very weak correlation between the magnitude of the pulse-to-pulse jitter for different time lags, besides for the pulse that immediately following the one in question (times k and $k-1$). On the other hand, the autocorrelation functions for $\Delta r_{k,LOT1}$ and $\Delta\theta_{k,LOT1}$ show there is a significant correlation for time lags up to about a minute. This suggests there are transient disturbances, lasting a few seconds to about a minute, that affect the LTL and increase the jitter of the beam.

The increase in pulse-to-pulse jitter due to propagation through the LTL will now be quantified using the data in Fig. 3.4. According to Ray Optics [22], the

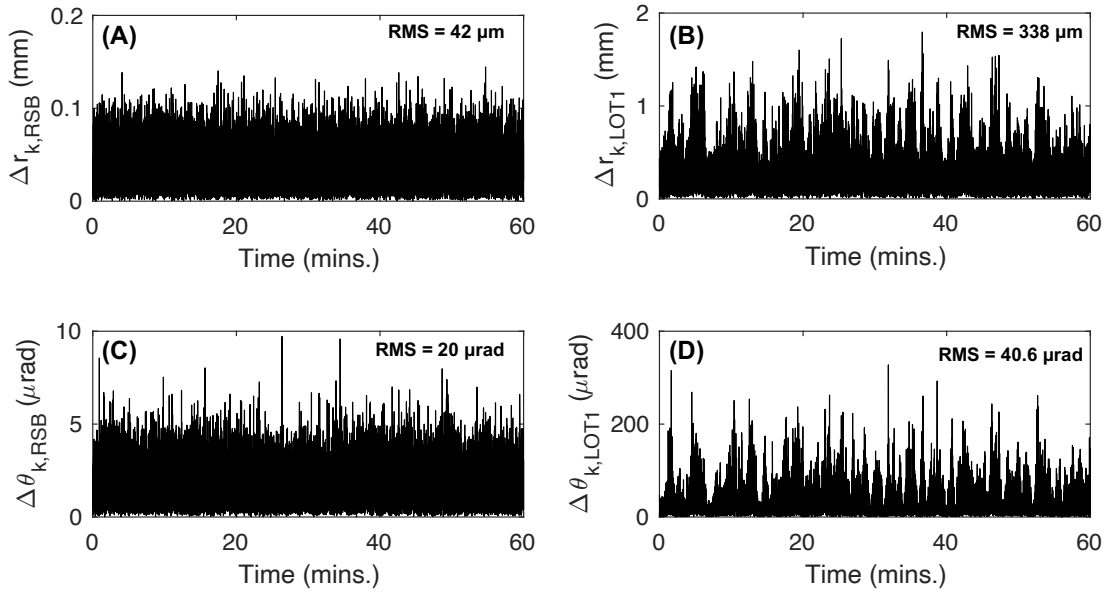

Figure 3.6: Pulse-to-pulse jitter measured before and after the laser transport line. (A): Position before LTL. (B): Position after LTL. (C): Angle before LTL. (D): Angle after LTL.

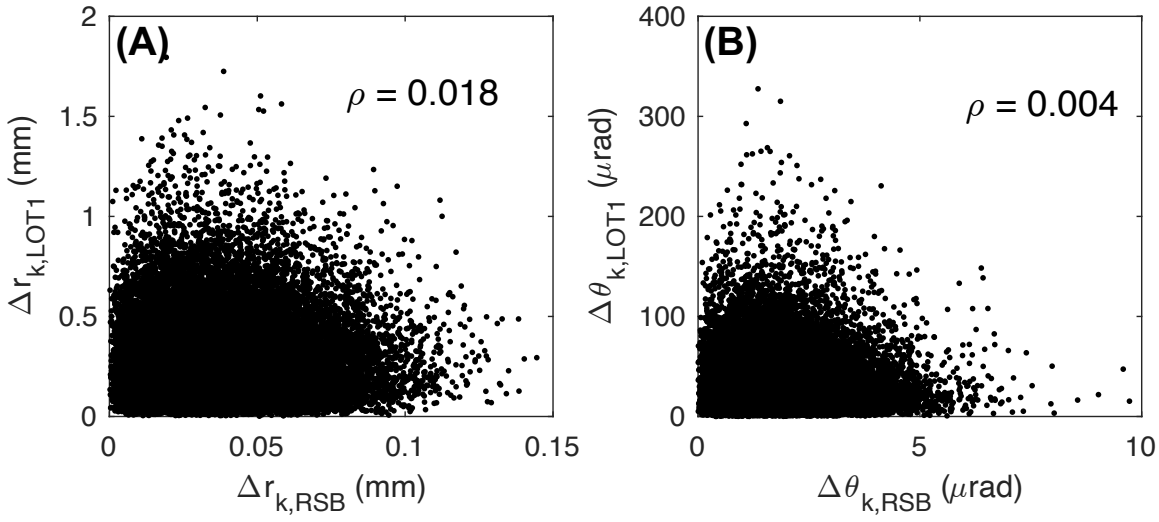

Figure 3.7: Correlation between the magnitude of the pulse-to-pulse jitter of the (A) position and (B) angle before and after the LTL.

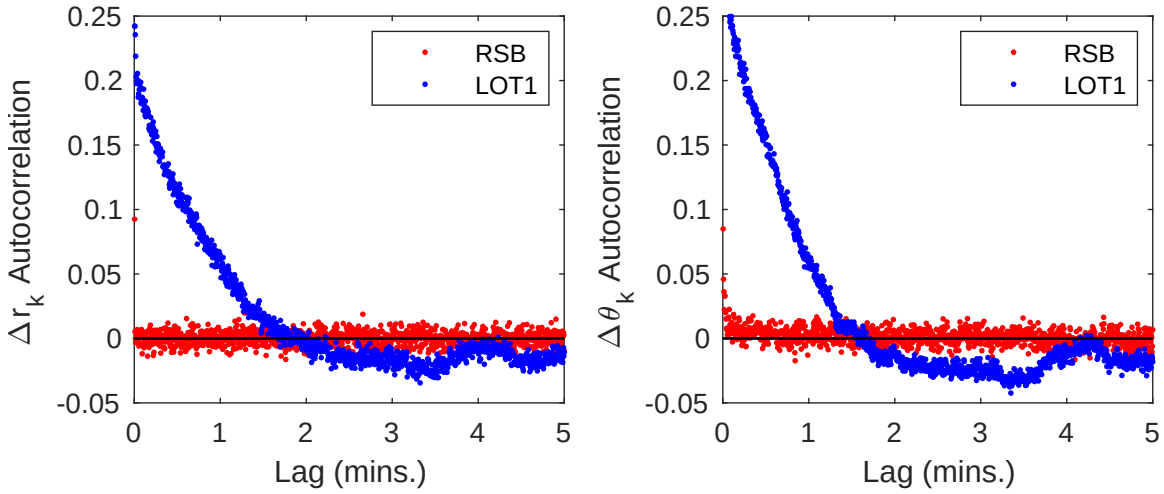

Figure 3.8: Autocorrelation between the magnitude of the pulse-to-pulse jitter before and after the LTL.

2-dimensional pulse-to-pulse jitter expected on LOT1 due the pulse-to-pulse jitter measured in the RSB is

$$\begin{pmatrix} \Delta r_{h,k} \\ \Delta \theta_{h,k} \\ \Delta r_{v,k} \\ \Delta \theta_{v,k} \end{pmatrix}_{LOT1} = \begin{pmatrix} 1 & L & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & L \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \Delta r_{h,k} \\ \Delta \theta_{h,k} \\ \Delta r_{v,k} \\ \Delta \theta_{v,k} \end{pmatrix}_{RSB} , \tag{3.3}$$

where L is the distance between the measurement locations. For the placement of the cameras used for this study, $L = 52.2$ m. Eq. (3.3) will be used to calculate the expected jitter on LOT1 based on the jitter measurements in the RSB. Fig. 3.9 shows distribution of the expected and measured jitter on LOT1. The widths of the distributions of the measured jitter on LOT1 are greater than expected from measurements in the RSB. This means that not all the jitter observed on LOT1 can be explained by the physics model of Eq. (3.3), so the additional jitter must be due to external disturbances in the LTL. Table 3.1 lists the RMS values of the expected and measured jitter. The external disturbances in the LTL increases the RMS of the position jitter by 2-3×, and the RMS of the pointing angle jitter by a factor of 20-30×. It is clear from the analysis in this section that the LTL is the primary source of laser pointing instability.

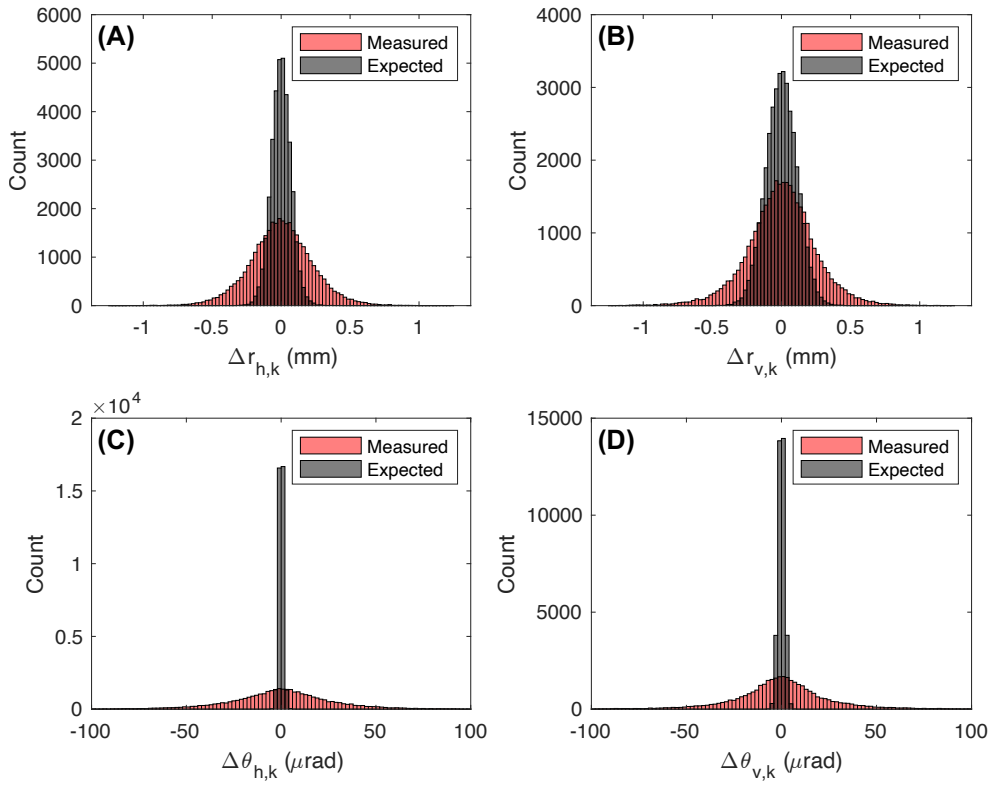

Figure 3.9: Measured Vs. Expected jitter on LOT1 after propagation through the LTL. The widths of the distributions show the measured jitter is much greater than is expected, suggesting the LTL is a significant source of jitter. Top: pulse-to-pulse jitter of the (A) horizontal and (B) vertical position. Bottom: pulse-to-pulse jitter of the (C) horizontal and (D) vertical angle.

Table 3.1: Measured Vs. Expected Jitter on LOT1.

Parameter	Expected Jitter	Measured Jitter
Position (RMS, μm)	73 (H) $\times$ 113 (V)	229 (H) $\times$ 249 (V)
Angle (RMS, μrad)	1.1 (H) $\times$ 1.7 (V)	29.0 (H) $\times$ 28.4 (V)

Chapter 4

Laser Pointing Stabilization Systems

High power lasers with complex architecture are known to be susceptible to beam pointing fluctuations. Beam pointing fluctuations in bulk lasers originate from mechanical vibrations and thermal effects, which can affect the alignment of internal optics and lead to a combination of shift and tilt of the output beam. When a laser beam with high peak power is propagated over long distances, there is certain to be a degradation in pointing stability. Solutions for stabilizing the pointing of a laser beam can be broadly categorized as either passive or active. In this chapter, an overview of both categories will be given. The laser pointing stabilization system for LACE uses an active scheme based on feedback, and a model of the system is used to analyze its performance under realistic disturbance conditions. The SNS Laser Wire Scanner (LWS) uses a similar feedback system to stabilize the pointing of a high-energy laser beam. An overview of the LWS will be given, and a study of the laser pointing stability that results with feedback OFF and ON will be discussed.

4.1 Passive Stabilization

Laser pointing stabilization systems are referred to as "passive" when they improve the stability without the need to monitor the beam or requiring an input of energy to function. Passive stabilization systems avoid unnecessary exposure to disturbances or dampen their effect through the use of specialized equipment or techniques. The most effective passive stabilization systems are included in the design phase of an experiment, as they can significantly increase the footprint of the optical system. First order passive stabilization solutions include commercial products for enhanced stability or vibration isolation, such as pedestal posts, pneumatic isolators or sorbothane bumpers. Heat sinks can also be bought off-the-shelf to help mitigate thermal effects. Modern optics tables include passive stabilization systems in the form of pneumatic legs, to isolate the surface of the table from seismic vibrations, and a monolithic surface to ensure all optics are moving together if the table does get disturbed. It is also common practice to enclose the path of the laser beam to minimize air currents which perturb the beam as it propagates through the atmosphere. Ideally, a LTL will also be evacuated to keep the optics clean and prevent moisture buildup. For long-distance transport, it is best to minimize the number of transport optics and mount mirrors on a monolithic structure with vibration isolation. Relay imaging should be used for long-distance transport, as opposed to direct propagation of the laser beam. There are also more advanced tabletop techniques that suitable for certain applications or classes of laser [23, 24], and many of the most elaborate passive stabilization systems are found in dedicated laser facilities [25, 26].

4.2 Active Stabilization

Stabilization systems are referred to as "active" when they require some additional input of energy to function. Active stabilization implies a control system that regulates the output of a dynamic system in the presence of a disturbance through

a control action in order to keep the output of the system at a desired value, or *setpoint*. The laser pointing stabilization system described in this dissertation uses a kinematic mirror with piezoelectric (PZT) actuators to keep the beam pointed at a user-defined setpoint on the sensor of a Complementary Metal–Oxide–Semiconductor (CMOS) camera located after the LTL. In this application, the measured output is the position of the beam on the camera sensor, the control action is the shift in position caused by the feedback mirror (FBM) tilting the beam upstream, and the disturbance is the shift caused by perturbations in the LTL. The output at any given moment in time is the sum of the disturbance and control action. The model is given in the discrete-time domain by

$$y_k = x_k + u_k, \tag{4.1}$$

where y_k, x_k and u_k are the output, disturbance, and control action signals at discrete-time index k, respectively. The value of u_k is chosen to counteract x_k in order to keep y_k at a setpoint.

Control systems are categorized as open-loop, closed-loop (feedback), or feedforward [27]. For open-loop control systems, u_k is independent of measurements of x_k or y_k. An example of a system where open-loop control is used is a clothes dryer that is set to run for a specified amount of time. After the time is up, the dryer will stop regardless of whether or not the clothes are dry. An open-loop laser pointing stabilization system would amount to a system that allows a human to shift the position of the beam manually when he or she notices it has drifted away from the setpoint. In this case, that would mean a human manually adjusting the control signal to the PZT actuators of the FBM to steer the beam. Feedback and feedforward control systems are employed to automate a process. Both categories rely on measurements of y_k or x_k and a control law to determine what the value of u_k should be. The laser pointing stabilization system for LACE uses a control scheme based on feedback from the CMOS camera to the FBM.

4.2.1 Feedback Control

In feedback control systems, a controller automatically determines the control action based on a control law that is a function of the difference in the setpoint and the measured output. This quantity is referred to as the error signal e_k, and it is given by

$$e_k = y_k - y_{sp}, \tag{4.2}$$

where y_{sp} is the setpoint. An example of a closed-loop control system is a thermostat that is programmed to measure the temperature of a room periodically and turn the furnace on or off depending on the sign of the error signal. Most active pointing stabilization systems are based on one [28, 29, 30, 31, 32] or two [33, 34, 35, 36, 37, 38] feedback loops with a photosensor linked to a kinematic mirror. The performance of a feedback system is largely influenced by the rate at which it can update the control action. Because feedback requires an error signal, the rate at which y_k is sampled often limits the rate of feedback. Therefore, well-designed feedback systems for continuous wave (CW) lasers can possibly make corrections at hundreds of kHz, limited by the sampling rate of the sensor or response time of other hardware components. Such systems generally yield much better results than those applied to laser beams with low pulse repetition rates. The challenge in stabilizing beams with low repetition rates in accelerator environments stems from the limited rate of feedback and the highly stochastic nature of the disturbances. Despite these challenges, there are several reported instances where low repetition rate IR [28, 29, 30] and UV [33] laser beams were transported and stabilized in accelerator environments at locations 10's to 100's of meters away from the laser.

The most popular control law used in feedback systems is that of the proportional-integral-derivative (PID) controller. PID control takes into account the past, present and predicted future error signal to bring the output back to the setpoint with minimal delay or overshoot. The discrete PID control law is given in the discrete-time domain

by

$$u_{k+1} = K_P \cdot e_k + K_I \cdot \sum_{n=0}^{k} e_n + K_D \cdot (e_k - e_{k-1}), \tag{4.3}$$

where K_P, K_I and K_D are the gain coefficients associated with the proportional, integral and derivative terms, respectively. The proportional term updates u_k based on the present value of e_k, the integral term updates u_k based on the history of the error signal, and the derivative term updates u_k based on the predicted change in the error signal at discrete time index $k+1$. The three gain coefficients can be "tuned" to produce a satisfactory feedback response for a wide range of disturbances. The derivative term is known to incorrectly predict the error if the disturbance is particularly noisy, which will reduce the performance of the system or even cause it to become unstable, especially in the case of discrete systems. In this application, use of the derivative term is ruled out because the disturbance is highly random due to the pulse-to-pulse jitter.

The optimal performance of a PI controller was investigated using a MATLAB simulation of the feedback system based on the model given by Eq. (4.1). The PI control law is given by the first two terms in Eq. (4.3). A 20-minute disturbance signal that was previously measured in the field with the un-stabilized UV laser was used as x_k in the simulation to estimate the performance of the feedback system under realistic conditions (Fig. 4.1). Stability will be quantified by the RMS value of the error signal, and the performance of the feedback system is optimal when the RMS value has been minimized for a given disturbance signal. For the simulation study, y_{sp} was set to zero, and the RMS of y_k was calculated for different combinations of K_P and K_I. There is no feedback when K_P and K_I are both zero because $u_k = 0$ for all k, therefore $y_k = x_k$ in that special case. The results of the PI study are shown in Fig. 4.2, where the RMS value of x_k is 0.633. The white area represents combinations of K_P and K_I that resulted in worse stability than without feedback. By inspecting the Fig. 4.2, it can be seen that an equivalent level of stability can be achieved with just the integral (I) term as with the PI controller. Since I and PI controllers are equally

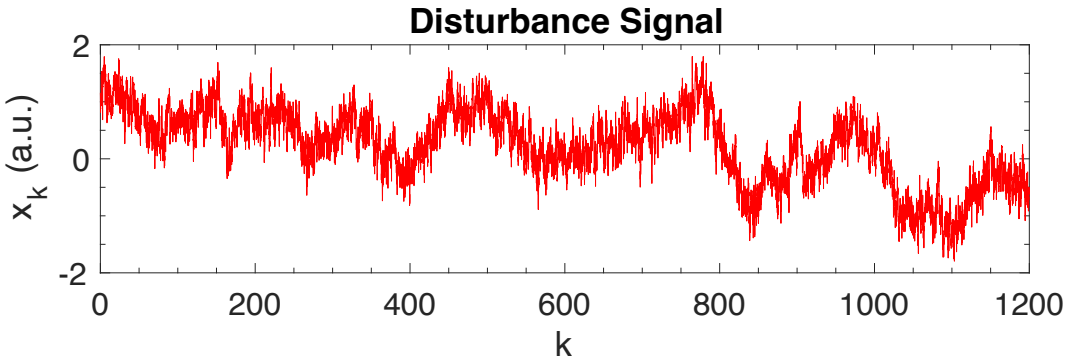

Figure 4.1: Disturbance signal used in the PI controller simulation study. The waveform was taken from real beam position data for the un-stabilized UV laser.

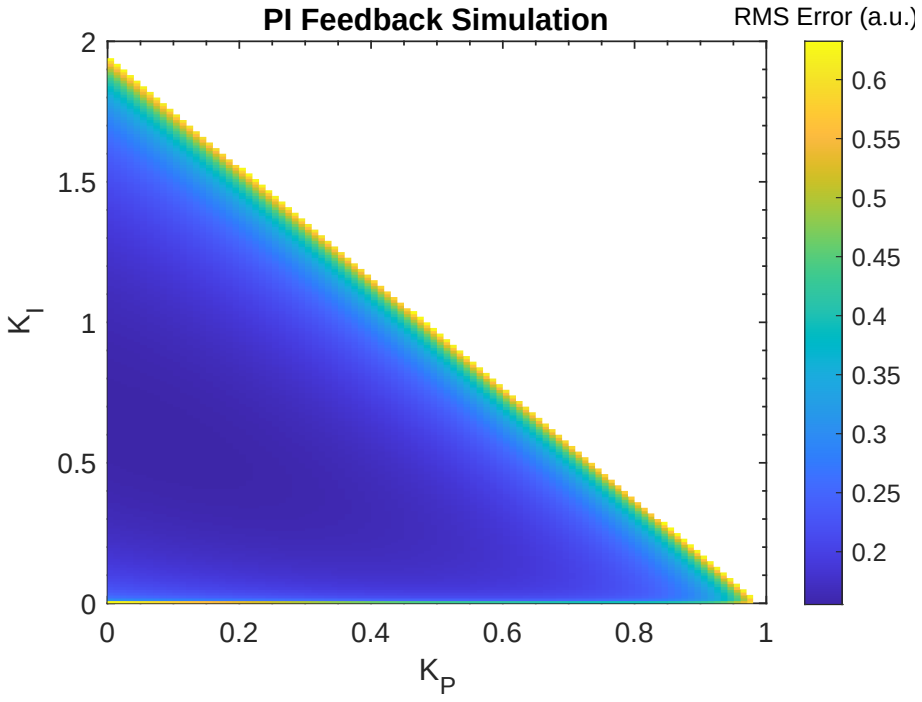

Figure 4.2: Simulation to find to optimal performance of a PI controller. Equivalent stability can be achieve with PI- and I-controllers.

as effective for the disturbances present in the real system, an integral controller is used to reduce the complexity of the feedback system and make it easier to tune the controller.

The discrete-time domain model of the feedback system will be used to derive an equivalent model in the z-domain in order to formally analyze the stability and frequency response of the system. Because the feedback system is used to stabilize the laser beam at a single position on the camera sensor, y_{sp} is a constant, and so x_k is the only dynamic input. Thus, each axis of the FBM constitutes a single-input-single-output (SISO) feedback system that can be modeled in the frequency domain by a transfer function. It is easier to derive the transfer function if the integral control law is expressed as

$$u_{k+1} = u_k - K \cdot e_k, \tag{4.4}$$

where K is positive definite and the subscript "I" in Eq. (4.3) is dropped. To find the transfer function, Eq. (4.4) and Eq. (4.1) are combined into a single expression for y_k as a function of x_k and constants, and then the time-shift property of the z-transform is used. In standard closed-loop form, the transfer function is given by

$$\frac{Y(z)}{X(z)} = \frac{1 - z^{-1}}{1 + (K-1) \cdot z^{-1}}, \tag{4.5}$$

and the block diagram showing the flow of signals in the z-domain is shown in Fig. 4.3.

The transfer function can be used to determine the values of K that are guaranteed to produce a stable feedback response. Notice the pole in Eq. (4.5) where the term in the denominator $P(z) = (K-1) \cdot z^{-1}$ is equal to -1. If the complex contour traced out for $P(z)$ by the locus of points in the z-plane corresponding to real frequencies encloses the point -1, the feedback system will exhibit a resonance-like response that will cause y_k to become unbounded. For discrete systems, real frequencies are mapped onto the unit circle in the z-plane by the transformation $z = \exp(i2\pi f/f_s)$, where f_s is the sampling frequency (10 Hz, in this case). The Nyquist Stability Criterion

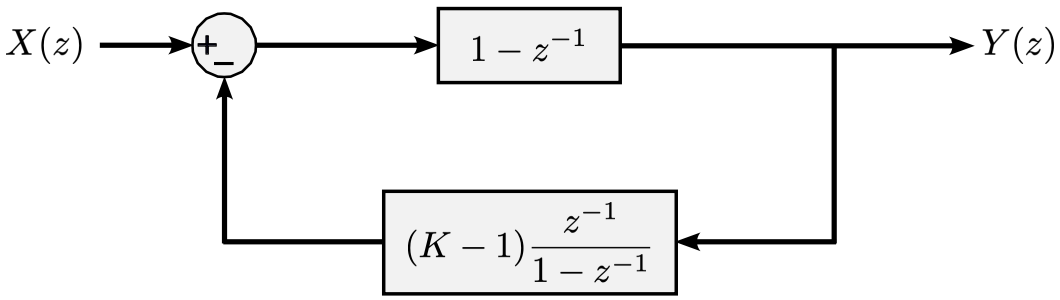

Figure 4.3: Feedback control block diagram in the z-domain.

[39] states that the feedback system is stable if the number of clockwise encirclements of -1 is opposite the number of poles of $P(z)$ outside the unit circle in the z-plane and less than or equal to zero, otherwise it is unstable. Since $P(z)$ has no poles outside the unit circle, the Nyquist contour must not enclose -1. Nyquist contours for several values of K are plotted in Fig. 4.4. The Nyquist contours are traced out in the clockwise direction as f is increased from $-f_s/2$ to $f_s/2$, so all points on or encircled by the contours are enclosed. The contours for $K = 2$ and 2.5 both enclose the point -1, while the contours for $K < 2$ do not. By examining $P(f)$, it is easy to see that the Nyquist contour is always a circle of radius $(K - 1)$, and -1 will be enclosed for $K \geq 2$, as well as for $K = 0$ (no feedback). Therefore, the range of K for which the feedback system is absolutely stable is determined to be $0 < K < 2$.

The effect the magnitude of the feedback gain has on the output for a general disturbance signal can be understood by analyzing the frequency response of the feedback system for several values of K. For linear, time-invariant systems such as this, y_k in response to a single-frequency (sinusoidal) x_k is also sinusoidal, but with different amplitude and phase. The frequency response is given by two plots which show the amplitude gain and phase shift of y_k relative to a sinusoidal x_k at different frequencies. In Fig. 4.5, the frequency response curves for the stable values of K already examined ($K = 0.1, 0.5, 1$ and 1.5) are plotted. From the amplitude gain response, we see that low frequency x_k are well suppressed, but for each value of K, there is a certain crossover frequency above which x_k is amplified. When considering a realistic x_k, which contains Fourier components across the whole measurement band, there is a trade-off between reducing its low frequency components and amplifying the high frequency ones. As K increases, the feedback system becomes more efficient at suppressing low frequency components, but the amplification of high frequency ones also increases until the system eventually becomes unstable. The optimal value of K for a realistic disturbance balances its high and low frequency components to yield the best overall stability of the output. The trade-off is illustrated in Fig. 4.6,

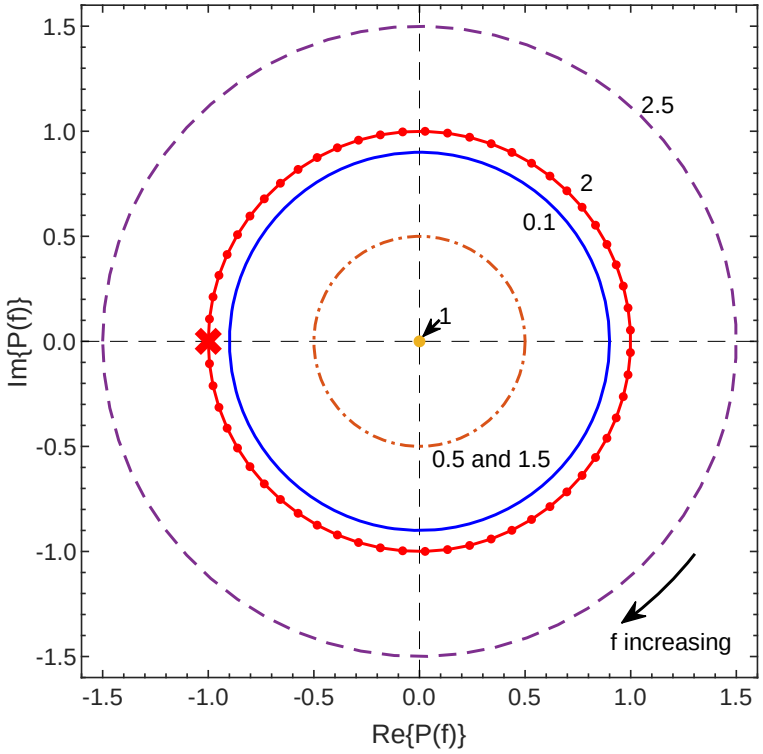

Figure 4.4: Nyquist contours for $K = 0.1, 0.5, 1, 1.5, 2$, and 2.5 are plotted. Contours that enclose the point -1 indicate the feedback system is unstable. Reproduced from [40].

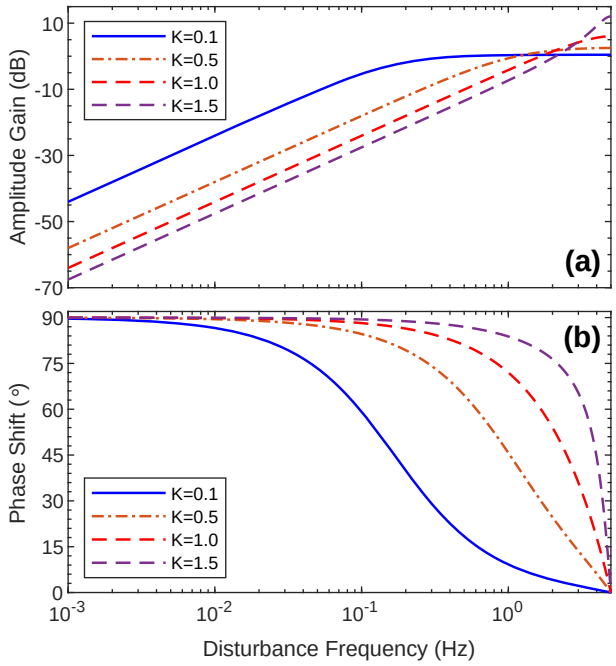

Figure 4.5: Frequency response of the feedback system for different values of K showing the (a) amplitude gain and (b) phase shift of the output relative to a sinusoidal disturbance input. Reproduced from [40].

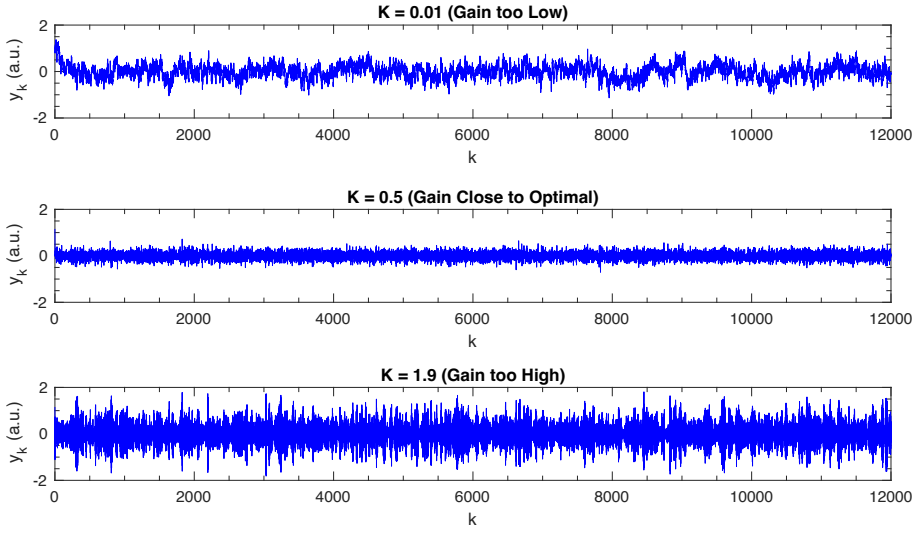

Figure 4.6: Simulation of y_k showing the trade-off between jitter and drift when tuning K. Top: drift is not eliminated when K is too low. Middle: drift is eliminated and jitter is not noticeably increased when K is optimized. Bottom: beam does not drift, but jitter is amplified above an acceptable level when K is too great.

which shows several waveforms for y_k that resulted from feedback simulations where x_k from Fig. 4.1 was stabilized with different K.

4.2.2 Feedforward Control

A few comments about feedforward control will be made as it pertains to the LACE system. Feedforward control systems are distinct from feedback systems because the control law does not relay on measurement of the error signal given by Eq. (4.2). Instead, feedforward control systems rely on direct measurement of the disturbance signal and a model of the dynamic system which is used to predict the effect of the disturbance on the output ahead of time. In this way, there is no latency (caused by the sampling time) between the measured output and update to the control action, as is the case with feedback. Feedforward is the most sophisticated category of control system, and the primary challenge in its implementation is accurately modeling both the dynamic system and disturbance signal. For the LACE optical system, the disturbance signal x_k at the location of the feedback cameras is truly the result of many individual disturbances affecting the beam as it propagates through the LTL. Because the disturbance is the result of vibrations, air currents and thermal effects, a model of the disturbance would require continuous measurements of these parameters at many locations, as well as a reliable model of the propagation of the high-energy laser pulse through the atmosphere. Due to the complexity and highly stochastic nature of the LACE optical system, predicting x_{k+1} is essentially impossible. Therefore, feedback is the most suitable, and certainly most economical, mode of control in this case.

4.3 Laser Wire Scanner Feedback System

The laser wire scanner (LWS) at the SNS is another system that requires the long-distance transport of a pulsed, high-energy laser beam. The LWS uses a feedback-based system to stabilize the position of an IR (1064 nm) beam of 50 mJ pulses at 30 Hz at distances up to 250 m. Parameters of the laser are listed in Table 4.1. The footprints of the LWS and LACE systems are shown in Fig. 4.7. The optical transport system for the LWS is less complex than that of LACE and uses a dedicated LTL with superior passive stabilization. This section is a brief detour away from the LACE system that will help illustrate the usefulness of feedback systems for laser pointing stabilization applications. The section includes an overview of the LWS, the setup of the feedback-based stabilization system, and the stability of the beam that results with feedback OFF and ON.

4.3.1 Overview

The LWS is a nonintrusive diagnostic system used to measure both the transverse and longitudinal emittance and spatial profile of ion beam minipulses at different points in the linac [41, 42, 43]. In the spatial profiling system, a high-power, Q-switched Nd:YAG (1064 nm) laser is used to photo-neutralize H^- where the laser beam intersects the minipulse. The detached electrons are steered in the vertical direction by a magnetic field and then collected in a Faraday cup downstream. The number of electrons captured by the Faraday cup and the known efficiency of the photo-detachment process are used to calculate the number of H^- particles that were neutralized. Each minipulse is assumed to be identical across macropulses, and a complete scan is performed by incrementally shifting the laser beam to different positions along the profile of the minipulse. Then, the transverse profile is reconstructed by fitting the H^- density data to a Gaussian curve (Fig. 4.8).

The calculation for the H^- density from the number of photo-detached electrons measured by the Faraday cup follows from [44]. Consider the interaction of (3D)

Table 4.1: Laser parameters for the LWS.

Parameter	Value
Wavelength	1064 nm
Pulse length (FWHM)	10 ns
Pulse energy	50 mJ
Pulse repetition rate	30 Hz
LTL length	250 m

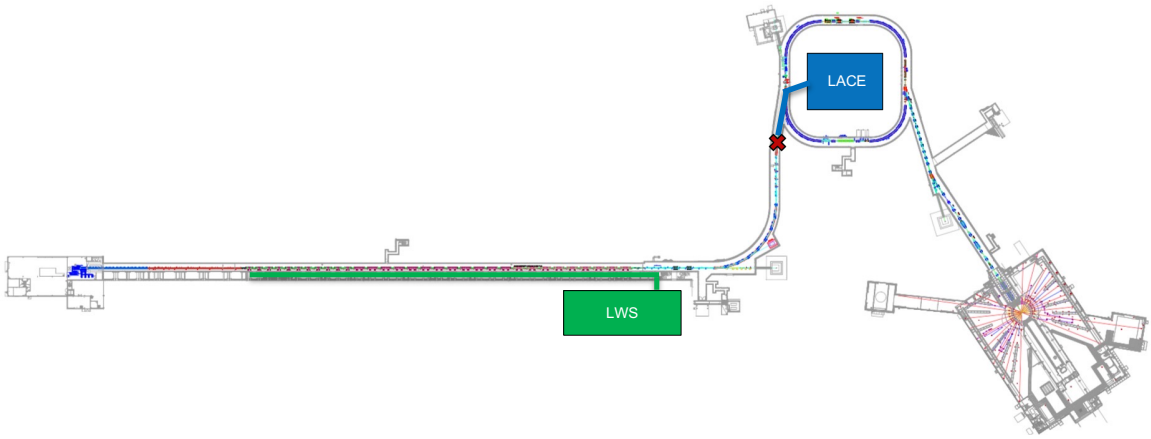

Figure 4.7: Footprint of the LWS system compared to LACE. The transport distances for LACE and the LWS are 65 m and 250 m, respectively.

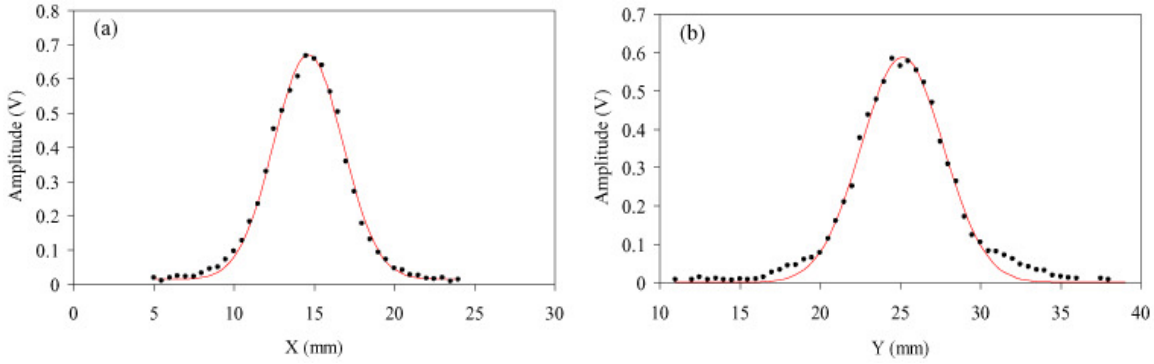

Figure 4.8: H^- density profile scan fitted to a Gaussian curve. (A): Horizontal H^- density profile. (B): Vertical H^- density profile. Reproduced from [28].

Gaussian-shaped laser and ion pulses crossing at an angle θ (Fig. 4.9). Assume the ion beam is traveling along the x-axis, the y-axis lies in the plane in which both beams propagate, and the z-axis is orthogonal to both. The rate of photo-neutralization is expressed as

$$\frac{dN_{det}}{dt} = c(1 - \beta\cos\theta) \cdot G \cdot \Sigma_0, \tag{4.6}$$

where N_{det} is the number of photo-detached electrons, c is the speed of light, β is the particle velocity as a faction of the speed of light, Σ_0 is the photo-neutralization cross section, and G is the geometric overlap factor. G is expressed as

$$G = \int n_l \cdot n_b \, dV, \tag{4.7}$$

where n_l and n_b are the photon and H^- density functions, respectively. Because the laser and ion beams are pulsed, both n_l and n_b depend explicitly on time. The total yield of photo-detached electrons, N_{det}, can be found by integrating Eq. (4.6) over the total interaction time,

$$N_{det} = c(1 - \beta\cos\theta)\Sigma_0 \int\int n_l(t) \cdot n_b(t) \, dV \, dt. \tag{4.8}$$

The H^- density as a function of time is expressed (in the lab frame) as

$$n_b(t) = n_{b,0} \cdot \exp\left(-\frac{(x - \beta ct)^2 + y^2 + z^2}{2(\sigma_{b,x}^2 + \sigma_{b,y}^2 + \sigma_{b,z}^2)}\right) \tag{4.9}$$

where $n_{b,0} = \frac{N_b}{(2\pi)^{\frac{3}{2}}\sigma_{b,x}\sigma_{b,y}\sigma_{b,z}}$ is the H^- density at the center of the bunch and $\sigma_{b,x}, \sigma_{b,y}$ and $\sigma_{b,z}$ are the ion beam RMS sizes along the respective axes. $N_b = \frac{I_b}{ef_{RF}}$ is the number of H^- in the bunch expressed as a function of beam current I_b, RF frequency f_{RF}, and the particle charge magnitude e. The laser pulse travels along the x-axis of a coordinate system rotated about the z-axis. The two coordinate systems are related by the transformation

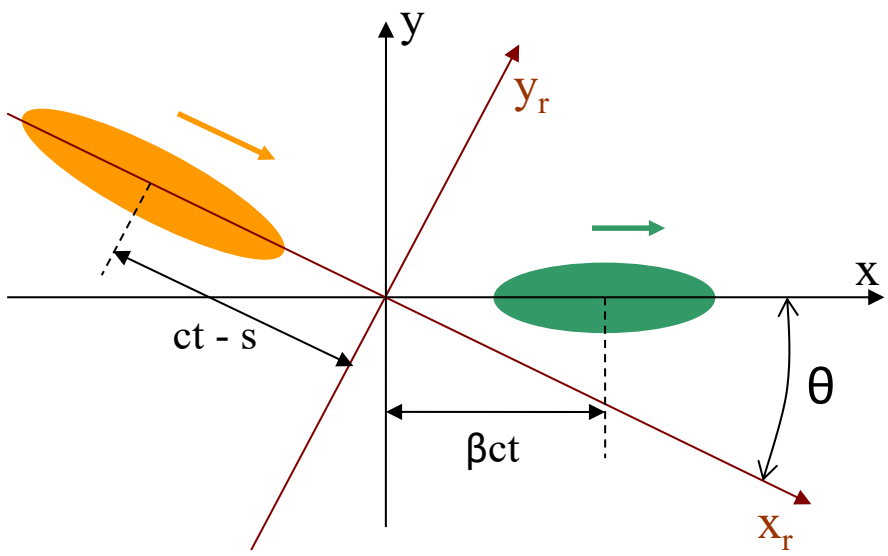

Figure 4.9: Interaction of 3D Gaussian laser (orange) and H^- (green) pulses. Reproduced from [44].

$$
\begin{aligned}
x_r &= x\cos\theta - y\sin\theta \\
y_r &= x\sin\theta + y\cos\theta \\
z_r &= z.
\end{aligned}
\tag{4.10}
$$

The photon density in the laser pulse is expressed as

$$
n_l(t) = n_{l,0}\cdot\exp\left(-\frac{(x_r - ct - s)^2 + y_r^2 + (z_r - d)^2}{2(\sigma_{l,x}^2 + \sigma_{l,y}^2 + \sigma_{l,z}^2)}\right),
\tag{4.11}
$$

where $n_{l,0} = \frac{N_l}{(2\pi)^{\frac{3}{2}}\sigma_{l,x}\sigma_{l,y}\sigma_{l,z}}$ is the photon density at the center of the laser pulse and $\sigma_{l,x}$, $\sigma_{l,y}$ and $\sigma_{l,z}$ are the laser beam RMS sizes along the respective axis. The number of photons in a laser pulse is given by $N_l = \frac{W_l\lambda}{hc}$, where λ is the photon wavelength and W_l is the energy contained in the pulse. The vertical shift of the laser beam is d, s is the laser pulse delay (centers of both beams meet at origin when $s = 0$), and h is Planck's constant. Substitution of Eqs. (4.9) and (4.11) into Eq. (4.8) gives

$$
N_{det} = \frac{1}{2\pi hce}\cdot\frac{I_b\lambda W_l}{f_{RF}}\cdot\frac{(1-\beta\cos\theta)}{\beta\sin\theta}\cdot\frac{\Sigma_0}{\sigma_{eff}\sqrt{\sigma_{b,z}^2 + \sigma_{l,z}^2}}\exp\left(\frac{s^2}{2\sigma_{b,z}^2} - \frac{d^2}{2(\sigma_{b,z}^2 + \sigma_{l,z}^2)}\right)
\tag{4.12}
$$

where σ_{eff} is expressed as

$$
\sigma_{eff} = \sqrt{\frac{\sigma_{b,x}^2\sin^2\theta + \sigma_{b,y}^2(\beta - \cos\theta)^2 + \sigma_{l,x}^2\beta^2\sin^2\theta + \sigma_{l,y}^2(1-\beta\cos\theta)^2}{\beta^2\sin^2\theta}}.
\tag{4.13}
$$

Maximum photo-neutralization occurs when $d = s = 0$. In that case, Eq. (4.12) reduces to

$$
N_{det} = \frac{1}{2\pi hce}\cdot\frac{I_b\lambda W_l}{f_{RF}}\cdot\frac{1-\beta\cos\theta}{\beta\sin\theta}\cdot\frac{\Sigma_0}{\sigma_{eff}\sqrt{\sigma_{b,z}^2 + \sigma_{l,z}^2}}.
\tag{4.14}
$$

By identifying N_b in Eq. (4.14), the photoneutralization efficiency η can be written as

$$\eta = \frac{N_{det}}{N_b} = \frac{\lambda W_l}{2\pi hc} \cdot \frac{1 - \beta\cos\theta}{\beta\sin\theta} \cdot \frac{\Sigma_0}{\sigma_{eff}\sqrt{\sigma_{b,z}^2 + \sigma_{l,z}^2}}. \tag{4.15}$$

Using η for the given configuration and N_{det} measured by the Faraday cup, the number of H^- in the region of intersected by the laser pulse can be expressed as

$$N_b = \frac{N_{det}}{\eta} \tag{4.16}$$

Finally, the the H^- density in that region is calculated using the known size of the laser pulse.

In the SNS implementation of the LWS, the transverse profiling stations have the ability to perform either a horizontal or vertical scan of the ion beam profile (Fig. 4.10). Part of the laser beam is picked off from the LTL by a partially-reflective mirror (M), and then a flipper mirror (FM) steers the beam to select between a horizontal or vertical scan. During a scan, the laser beam is shifted to different positions along the minipulse profile by a translatable steering mirror (SM) mounted together with a lens (L), which focuses the laser beam near the IP. The high-power laser beam is expanded to a diameter of approximately 2 cm (4σ) before the transport line. The lens, of focal length 200 mm, was placed as close as possible to the vacuum window so as to minimize laser-induced damage to the window (W) as light enters the chamber. After the interaction, the light leaves the chamber through an exit window. A photodiode (PD) is placed at the center of each beam dump (BD) to confirm the presence of the laser pulse.

There are a total of nine profile measurement stations along the length of the SCL (Fig 4.11). The measurement stations can be operated independently, or simultaneously to scan the profile of several adjacent minipulses at once [45]. There is a 30 ns window in which the laser pulse phase can be tuned so that a single pulse intersects a H^- minipulse at all nine stations simultaneously (Fig. 4.12). The

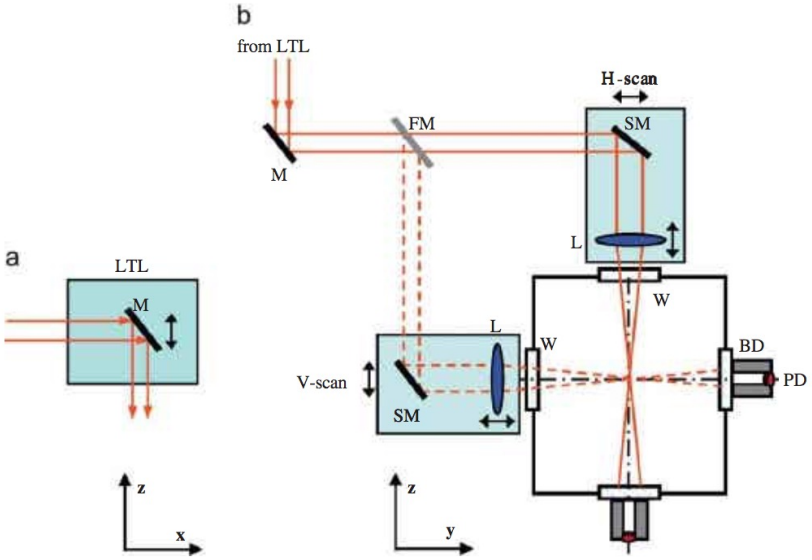

Figure 4.10: LWS transverse profiling station. (A): Part of the laser beam is picked off from the LTL by a partially-reflective mirror (M) and steered to the profile measurement station. (B): A flipper mirror (FM) is used to steer the beam to make a horizontal or vertical profile scan. The steering mirror (SM) and lens (L) are mounted on a pico-motor stage that translates along the profile of the ion beam. Reproduced from [41].

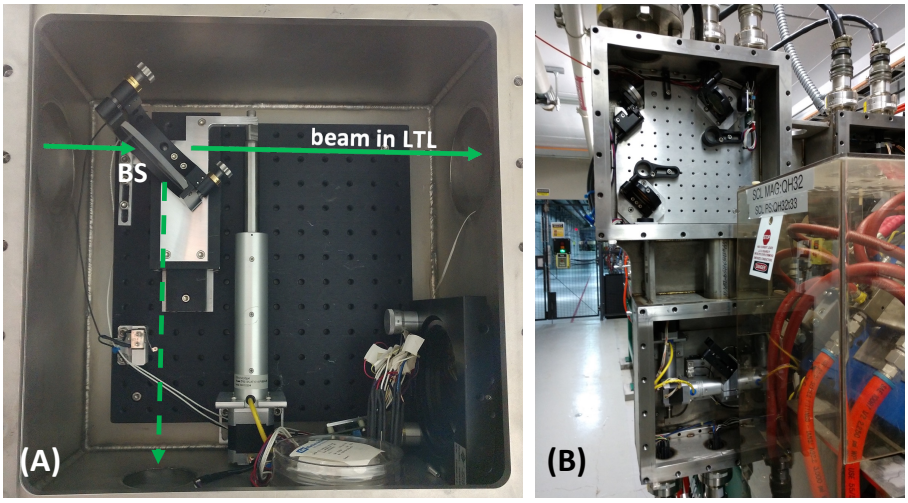

Figure 4.11: LWS transverse profile measurement station in the SCL. (A): A BS picks off part of the laser beam from the LTL. (B): Measurement station between cryomodules in the SCL.

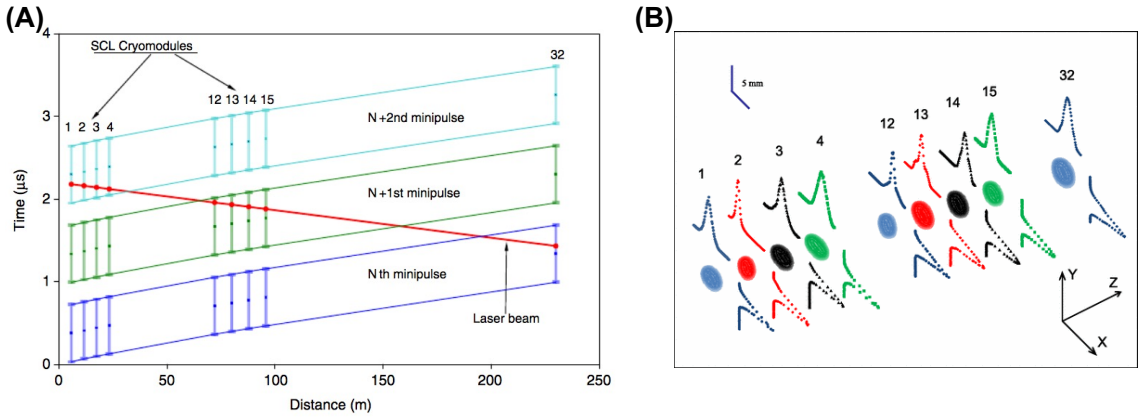

Figure 4.12: Simultaneous profile scan at different LWS stations. (A): Timing window for a simultaneous scan. (B): Result of a simultaneous scan. Reproduced from [45].

simultaneous scan allows operators and physicists to more conveniently track the evolution of the beam through the SCL.

4.3.2 Pointing Stabilization System Setup and Equipment

The layout of the laser pointing stabilization system for the LWS is shown in Fig. 4.13. The length of the SCL requires a 250 m long, free-space LTL to transport the high-energy laser pulse to each LWS station. The turbulent accelerator environment and high heat load on the optics cause the laser beam to slowly drift out of alignment at the furthest stations. A laser pointing stabilization system uses feedback between a digital camera and PZT mirror to keep the beam aligned.

The laser and FBM with open-loop PZT actuators (New Focus 8301) are located in a laser room in the HEBT service building outside of the accelerator tunnel, and the final mirror before the straight of the LTL is located 25 m after the laser. The LTL for the LWS was constructed at the same time as the SNS tunnel and runs along the ceiling of the linac. A CMOS camera (Allied Vision GC750) is used to measure the position of the laser beam, and a viewscreen (VS) is used to image the beam because of its large size and to prevent damage to the CMOS sensor from the high-energy laser pulses (Fig. 4.14). The magnification of the zoom lens must be measured in order to correctly calculate the shift of the real beam on the VS from the camera measurements. The relationship between the shift of the real beam on the VS and the image on the CMOS sensor is given by

$$\Delta x_{VS} = M^{-1} \cdot \Delta x_{CMOS} \tag{4.17}$$

where M is the magnification, and Δx_{VS} and Δx_{CMOS} is the shift on the VS and CMOS sensor, respectively.

Different PCs are used for the camera and the FBM. Separate LabVIEW programs were custom-made for the image acquisition and feedback routines (Fig. 4.15). The role of the camera program is the acquire the raw images, subtract background signal,

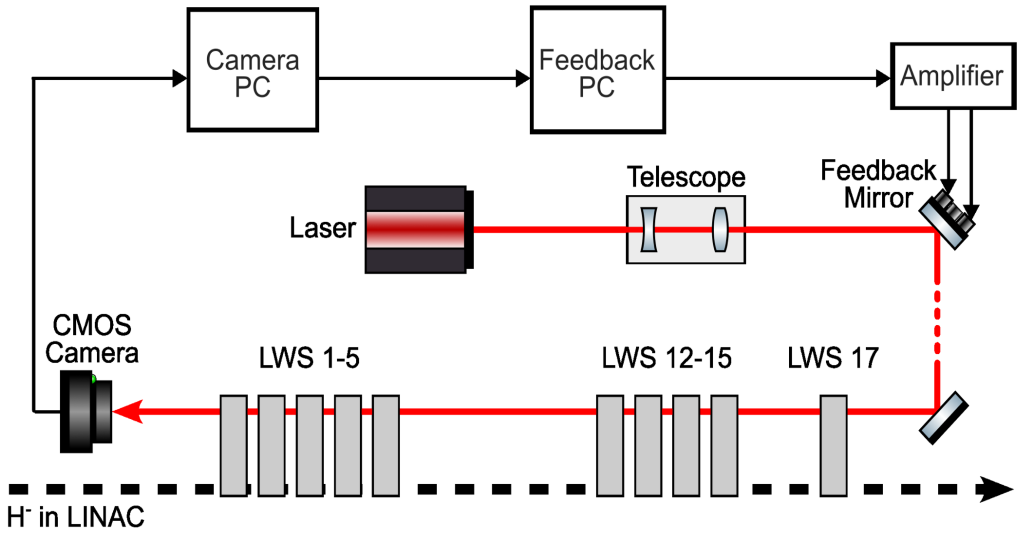

Figure 4.13: Layout of the laser pointing stabilization system for the LWS. A CMOS camera in the SCL measures the position of the laser beam and a PZT feedback mirror keeps the beam from drifting away from the furthest stations.

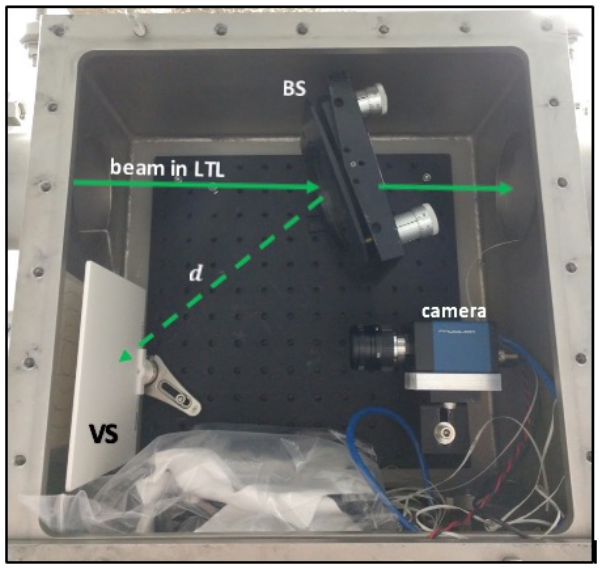

Figure 4.14: CMOS camera monitoring the image of the beam on a VS.

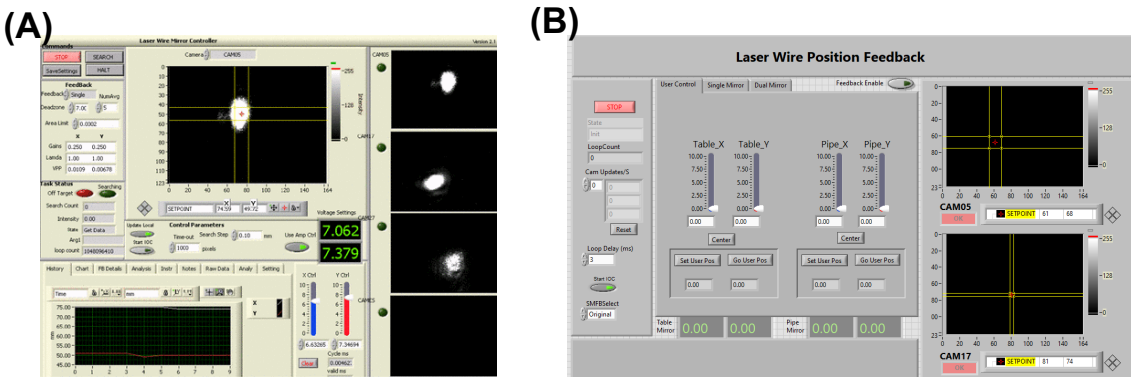

Figure 4.15: LabVIEW software for the cameras and feedback mirror. (A): Camera program GUI. (B): Feedback program GUI

and calculate the beam centroid coordinates in the reference frame of the camera. The centroid coordinates are then sent to the feedback machine over the internal Industrial Control Systems (ICS) network. The feedback program constructs the error signal from the centroid and the pre-programmed setpoint in the coordinate system of the camera. It then uses the error signal to carry out the feedback routine and updates the voltage to each axis of the FBM (horizontal and vertical). The LWS feedback system uses the integral control law described in the previous section. The PC outputs the DC voltage through an analog-out Peripheral Component Interconnect Express (PCIe) card (National Instruments 6052E) that was installed in the machine. The signal output by the PC is amplified by the FBM controller (Physik Instrumente GmbH E616.SS0) before it is applied to the PZT actuators. The feedback system runs at the maximum 30 Hz, equal to the pulse repetition rate of the laser.

4.3.3 Pointing Stabilization Study

The laser pointing stability that results with feedback OFF and ON was measured and will be briefly analyzed. The analysis will again be broken down into the drift (slow) and jitter (fast) components of the pointing fluctuations. Fig. 4.16 shows the beam drift with the feedback system OFF and ON at a position 143 m from the laser over 1 hour. With feedback OFF, the position of the beam was drifting significantly over the course of the measurement. The colormap shows that the beam was drifting horizontally from left to right and oscillating in the vertical dimension at a more-or-less constant frequency. With feedback ON, the beam did not drift, and it remained jittering about the same average position for the duration of the measurement. It is clear that the feedback system efficiently controls the slow drift of the beam, but that there is some residual jitter that the feedback system cannot control. The pulse-to-pulse jitter with feedback OFF and ON is shown in Fig. 4.17. At a glace, it appears that the feedback system did not have much effect on the pulse-to-pulse jitter of the beam. The feedback OFF and ON data were taken on different

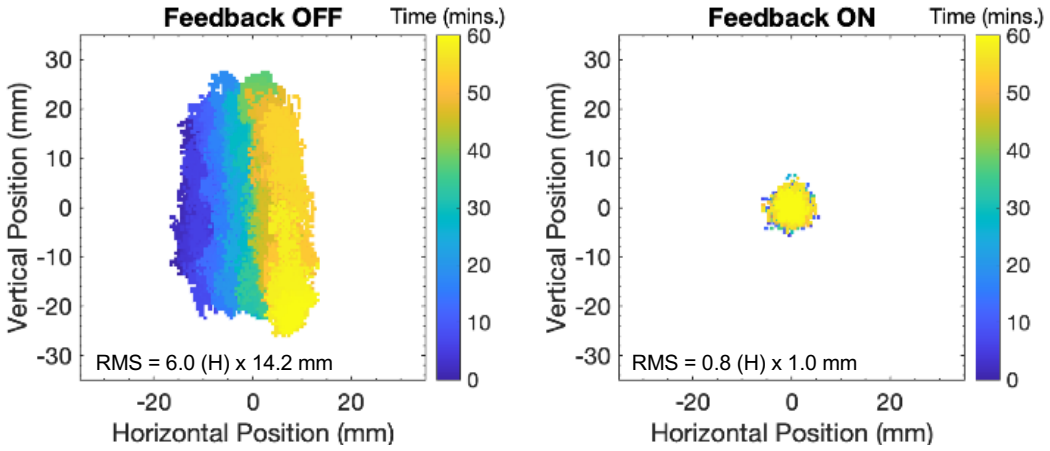

Figure 4.16: Laser beam drift with feedback OFF and ON measured for 1 hour.

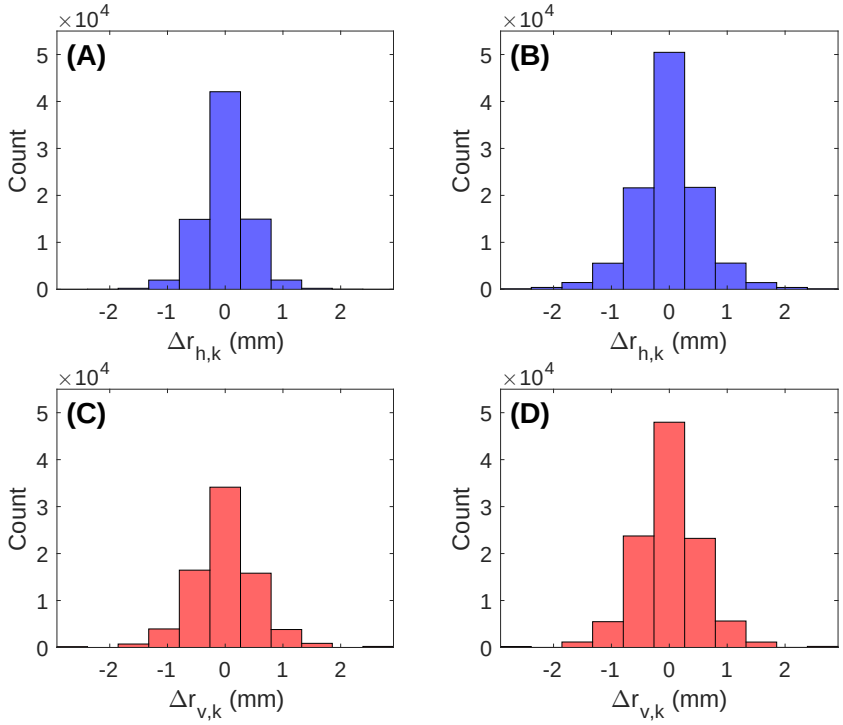

Figure 4.17: Pulse-to-pulse jitter in the horizontal dimension with feedback (A) OFF and (B) ON, and in the vertical dimension with feedback (C) OFF and (D) ON.

days, and so any different in the pulse-to-pulse jitter can be assumed to be due to small, seasonal changes in the optical system or condition of the laser. The laser used for the LWS has a service lifetime of 1000 hours, and it should be expected that the pointing stability, as well as other parameters of the beam, will degrade over time. The full range of the drift and STD of the pulse-to-pulse jitter are listed in Table 4.2.

Table 4.2: Laser pointing stability results.

Dataset	Drift Full Range (mm)	Pulse-to-Pulse Jitter STD (μm)
Feedback OFF	29.68 (H) $\times$ 53.40 (V)	431 (H) $\times$ 649 (V)
Feedback ON	11.13 (H) $\times$ 12.00 (V)	584 (H) $\times$ 602 (V)

Chapter 5

Stabilization System Development in the Laboratory

This chapter details the development of the laser pointing stabilization feedback system used for LACE. The system was developed and tested in a laboratory environment before it was deployed to the field to stabilize the UV laser beam. The feedback system for LACE differs from that of the LWS because it is designed to stabilize the horizontal angle and vertical position of the laser beam at the location of the laser-particle IP. It also uses more advanced image processing techniques to extract the position of the laser beam with software that is streamlined and localized to a single PC. First, the scheme of the system is elaborated upon and the optical setup in the laboratory is shown. Next, the instrumentation used in the system is introduced, and the custom-made LabVIEW program that functions as the feedback controller and data acquisition system is also detailed. Then, the feedback gain parameter is optimized prior to a laboratory test with a realistic disturbance using a MATLAB simulation. Finally, the results of the laboratory test are presented and compared to the expected results from the simulation.

5.1 Setup in the Laboratory

As was explained in Section 3.1, the laser parameters the LACE experiment is most sensitive to are the horizontal angle at which the laser and ion beams cross and the vertical distance between the centers of the two beams. As such, the stabilization system was designed to minimize the RMS error of the horizontal angle and vertical position of the laser beam. Since it is only necessary to stabilize one parameter in each dimension, a single FBM was used in the system, whereas two would be required to completely control the trajectory of the beam. The FBM was placed just after the laser, and two CMOS cameras were set up to make independent measurements of the position and angle of the laser beam used in the test after an artificial disturbance was introduced. The scheme is illustrated in Fig. 5.1. The LTL was simulated on an optics table by a 16-meter path, and a second kinematic mirror was placed near the middle to generate the artificial disturbance. The disturbance mirror (DM) was used to reproduce a real disturbance signal that was previously measured in the field with the un-stabilized UV laser to test the system under realistic conditions. At the end of the path, the Position Camera monitored the position of the beam, and the Angle Camera was setup to measure the pointing angle (Fig. 5.2).

The cameras directly measure the position of the beam on their sensors, however an optical technique was used to calculate the pointing angle before the lens from the position of the beam measured by the Angle Camera. The techniques requires the beam to be collimated, which is the purpose of the telescope (T) before the FBM. This technique is explained by the Ray Optics theory of light [22]. Ray Optics can be used to calculate the transverse position and angle of a laser beam, modeled as a ray, at a chosen location along its path if the position and angle are known at some other location. An optical system is modeled in a single transverse dimension by a 2×2 matrix, and the laser beam is propagated between two points by multiplying its state vector $\vec{o} = (r, \theta)^T$ by the matrix describing the optical system between them.

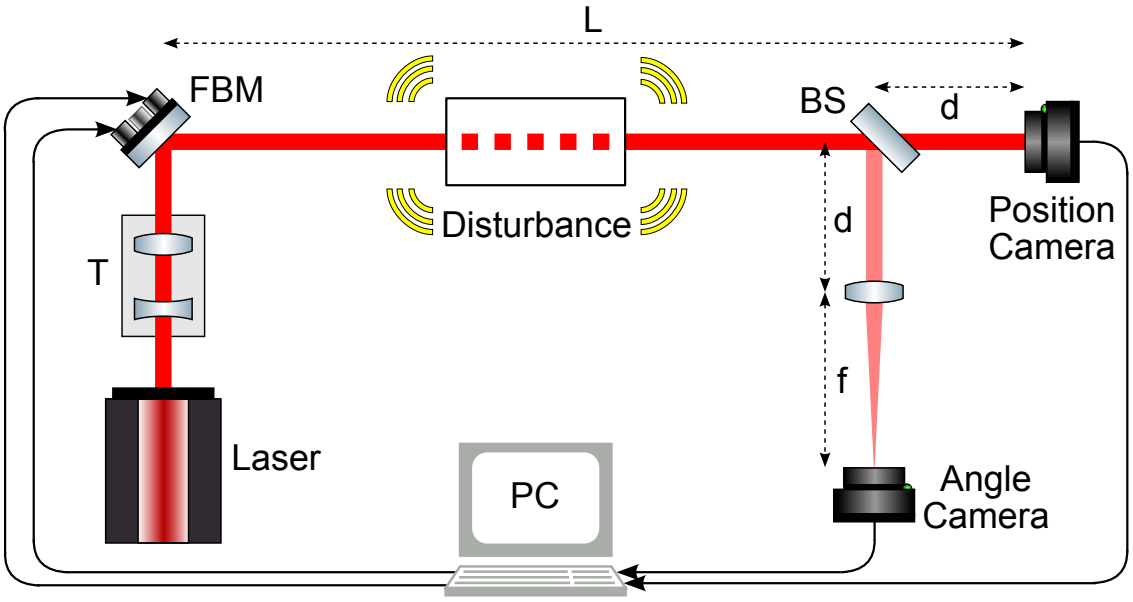

Figure 5.1: Scheme of the feedback-based laser pointing stabilization system. The FBM stabilizes the horizontal angle of the beam, measured by the Angle Camera, and the vertical position, measured by the Position Camera.

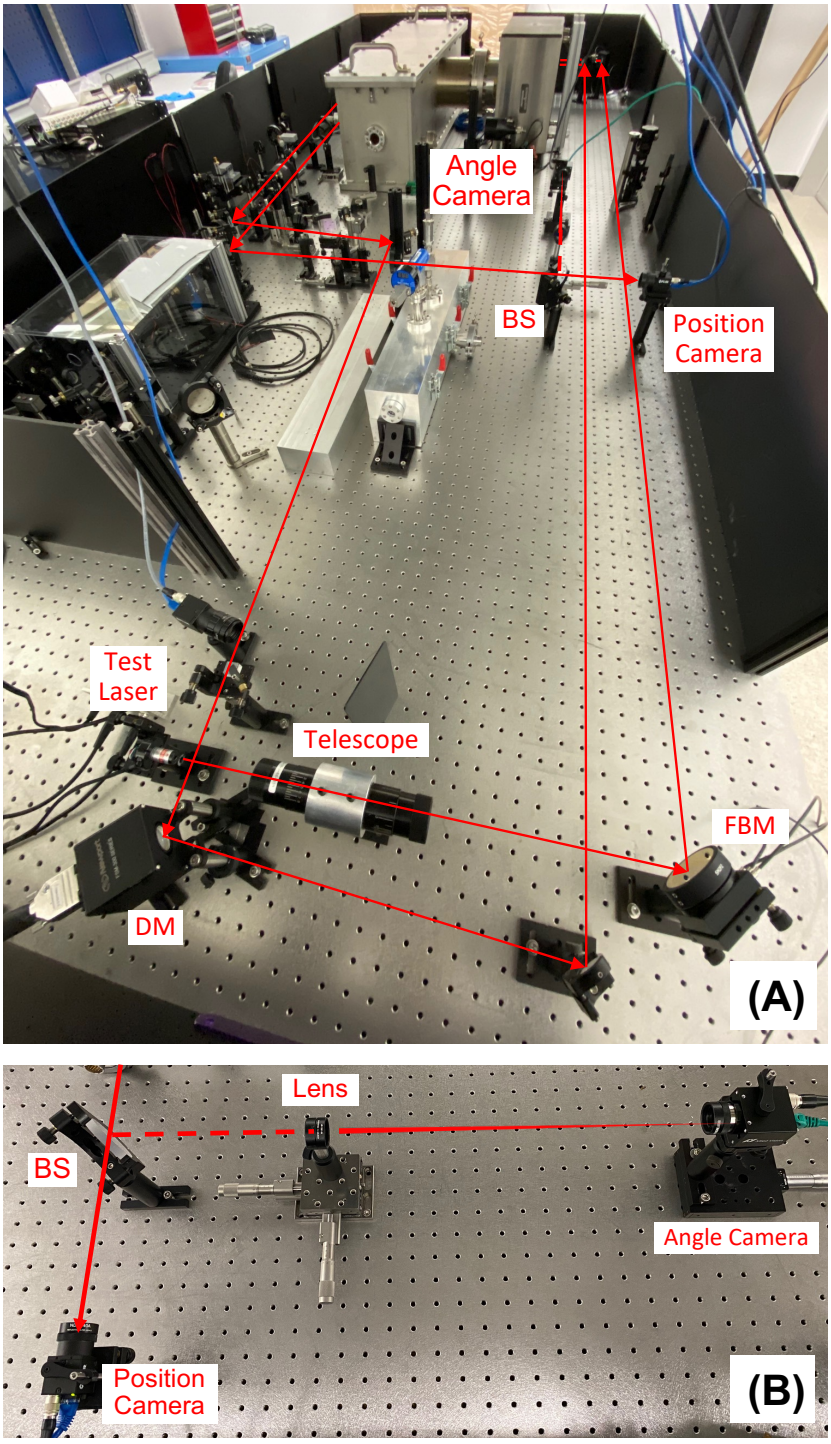

Figure 5.2: Optical setup in the laboratory. (A): 16-meter path. (B): Position and angle measurements.

The general formulation is given by

$$\vec{o}_2 = O_{sys} \cdot \vec{o}_1, \tag{5.1}$$

where O_{sys} is the matrix describing the optical system between points 1 and 2 along the optical axis. O_{sys} is constructed from the ordered product of matrices describing the individual elements that make up the optical system, given by

$$O_{sys} = O_n \cdot ... \cdot O_2 \cdot O_1, \tag{5.2}$$

where the subscripts indicate their order along the path of the beam.

In general, the angle of a laser beam at one point in an optical system is a linear combination of its position and angle at another point upstream. However, the pointing angle of a laser beam can be calculated from a measurement of its position downstream, without coupling to the position upstream, if the downstream position measurement is made in the focal plane of a convex lens. The position of a laser beam on the sensor of a camera placed a distance d after a lens is given by

$$\begin{aligned}
\vec{o}_{cam} &= O_d \cdot O_{lens} \cdot \vec{o}_{lens} \\
&= \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix} \cdot \begin{pmatrix} r_{lens} \\ \theta_{lens} \end{pmatrix} \\
&= \begin{pmatrix} 1 - d/f & d \\ -1/f & 1 \end{pmatrix} \cdot \begin{pmatrix} r_{lens} \\ \theta_{lens} \end{pmatrix} \\
&= \begin{pmatrix} (1 - d/f) \cdot r_{lens} + d \cdot \theta_{lens} \\ -r_{lens}/f + \theta_{lens} \end{pmatrix} \\
\rightarrow r_{cam} &= (1 - d/f) \cdot r_{lens} + d \cdot \theta_{lens},
\end{aligned} \tag{5.3}$$

where f is the focal length of the lens, and the subscripts "$lens$" and "cam" indicate values for the beam incident on the lens and sensor of the Angle Camera, respectively.

If the CMOS sensor is positioned in the focal plane of a convex lens such that $d = f$, then this first term in the last line of Eq. (5.3) drops out, and the angle of the beam before the the lens can be calculated directly from the position measured by the camera as

$$\begin{aligned} r_{cam} &= f \cdot \theta_{lens} \\ \rightarrow \theta_{lens} &= \frac{r_{cam}}{f}. \end{aligned} \tag{5.4}$$

With this technique, the angle of the beam before the lens could be measured with a single camera.

The two cameras are set up to simultaneously measure the position and angle at the same location. This is possible because the beam incident on the lens before the Angle Camera is an image, or "copy", of the beam on the Position Camera sensor, meaning it shares the same state vector. The imaging technique can be explained by showing that the state vector of the laser beam transmitted through the BS onto the sensor of the Position Camera is equal to the state vector of the reflected beam incident on the lens. Starting with the state vector of the beam incident on the BS, the calculation is as follows

$$\begin{aligned} O_d \cdot \vec{o}_{BS} &= O_d \cdot O_{refl} \cdot \vec{o}_{BS} \\ \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} r_{BS} \\ \theta_{BS} \end{pmatrix} &= \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} r_{BS} \\ \theta_{BS} \end{pmatrix} \\ \begin{pmatrix} r_{BS} + d \cdot \theta_{BS} \\ \theta_{BS} \end{pmatrix} &= \begin{pmatrix} r_{BS} + d \cdot \theta_{BS} \\ \theta_{BS} \end{pmatrix} \\ \vec{o}_{cam} &= \vec{o}_{lens} \\ \rightarrow r_{cam} &= r_{lens} \end{aligned} \tag{5.5}$$

where O_d and O_{refl} are the ray transfer matrices for free-space propagation over the distance d and reflection off a flat surface, and the subscripts "BS", "$lens$" and

"*cam*" indicate values for the beam incident on the BS, lens and sensor of the Position Camera, respectively. These two optical techniques allow complete characterization of the state vector (in both transverse dimensions) at the end of the optical path.

5.1.1 Instrumentation

The three essential components of the laser pointing feedback system are the CMOS cameras, PZT mirror, and the PC with associated interface cards. CMOS cameras are a type of digital camera based on semiconductor technology which has replaced charged-coupled-device (CCD) cameras in recent decades due to their low manufacturing cost and power-efficiency. A CMOS sensor consist of a 2-dimensional array of light-sensitive pixels, each consisting of a small photodiode coupled to a metal-oxide-semiconductor field-effect transistor (MOSFET). The sensor provides a 2-dimensional intensity profile of the incident light by reading the signal from each pixel simultaneously. CMOS cameras are commonly used to measure the transverse intensity profiles of laser beams (Fig. 5.3). The specific camera model used in the feedback system is the Prosilica GC655 made by Allied Vision (Fig. 5.4). They are members of a class of high resolution cameras with high frame rates know as Gigabit Ethernet (GigE). GigE cameras are named as such because they are capable of transmitting data over Ethernet at a rate of 1 Gigabit per second. These cameras were chosen for the feedback system due to their high dynamic range, pixel size, and frame rate. Some relevant specifications are listed in Table 5.1. The GC655 cameras were also found to be relatively radiation tolerant and have remained functional after exposure to radiation doses in excess of 100 Gy.

The FBM used in the laboratory system was a 2 inch diameter mirror on a tip/tilt mount with open-loop PZT actuators (Thorlabs KC1-T-PZ). The working principle of the FBM is illustrated Fig. 5.5. PZT crystals generate a voltage when they are deformed, and conversely, an applied voltage will cause the material to expand or contract. The strain of the material increases with the magnitude of the DC voltage

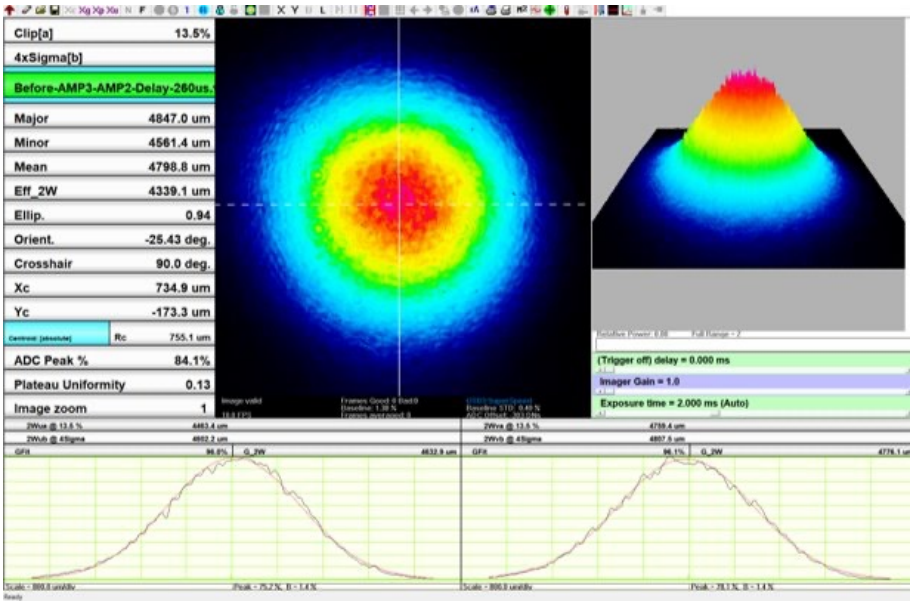

Figure 5.3: Laser beam profile measurement with a CMOS camera.

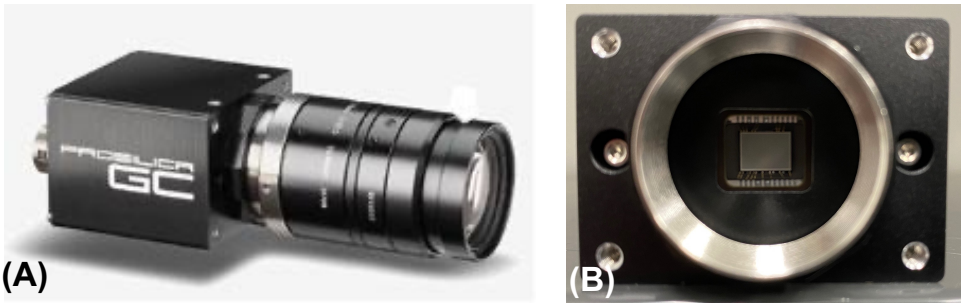

Figure 5.4: GC655 camera. (A): Zoom lens attachment. (B): CMOS sensor.

Table 5.1: GC655 camera specifications.

Parameter	Value
Pixels	659 (H) × 493 (V)
Pixel Size	9.9 (H) × 9.9 (V) μm
Sensing Area	6.4 (H) × 4.8 (V) mm
Bit-depth	12-bit
Frame Rate	90 fps

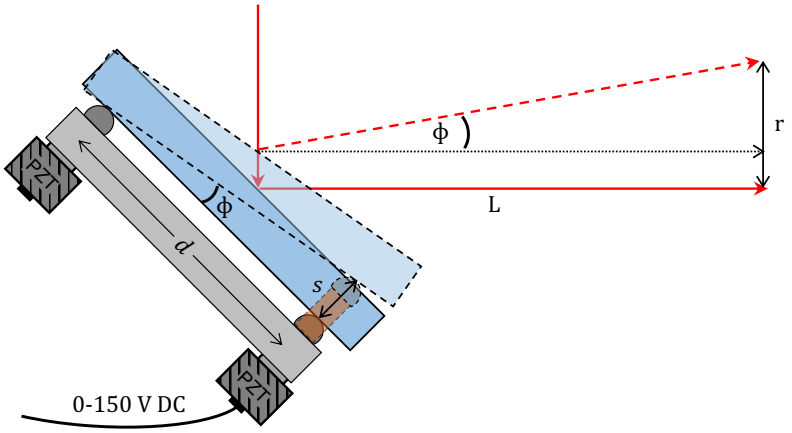

Figure 5.5: Working principle of the FBM. The mirror tilts when the PZT actuator extends in response to a DC control signal.

V applied over it. The actuators are assembled from a stack of PZT wafers, which are wired in parallel so as to receive the same voltage (Fig. 5.6). The tilt of the FBM about each axis was controlled by a $0 - 150$ V DC control signal from its associated amplifier (Thorlabs MDT693D). Because the extension of the actuator s is very small relative to the lever arm d, the small angle approximation can be used. In the approximation, the tilt of the FBM ϕ is directly proportional to s such that $\phi \approx s/d$, and because s is directly proportional to V, the relationship between $\Delta\phi_k$ and ΔV_k can then be written as

$$\Delta\phi_k \simeq \chi \cdot \Delta V_k \tag{5.6}$$

where k is the discrete-time index and χ is the measured microradians-per-volt coefficient that characterizes the response of the FBM for small ΔV_k. The response coefficients for the horizontal and vertical axes of the FBM were found to be $\chi_h = 28.0$ μrad/V and $\chi_v = 21.6$ μrad/V, respectively.

It should be mentioned that PZT materials have two distinct nonlinearities: hysteresis and creep (Fig. 5.7). The hysteresis is the effect where the strain of the PZT lags behind the amplitude of the voltage as it increases or decreases. The creep is a time dependent phenomena where the strain continues to change while the voltage remains constant. More sophisticated PZT actuators have a built in strain sensor, which forms a feedback loop with the PZT amplifier to ensure a linear response of the actuator. The DM (Newport FSM-300) worked in much the same way as the FBM, but it had its own amplifier, as well as an internal strain sensor to ensure a linear response. The voice coil actuators of the DM can respond to higher frequency control signals, but the response has higher uncertainty than PZT-based ones. The control signal to the DM was also supplied by the PC through two additional analog-out channels.

The feedback PC was purchased as a dedicated machine for the feedback system. It was necessary to choose a machine capable of handling the processing demands to execute the control routine at a rate of 10 Hz. The processor is the Intel(R) Core(TM)

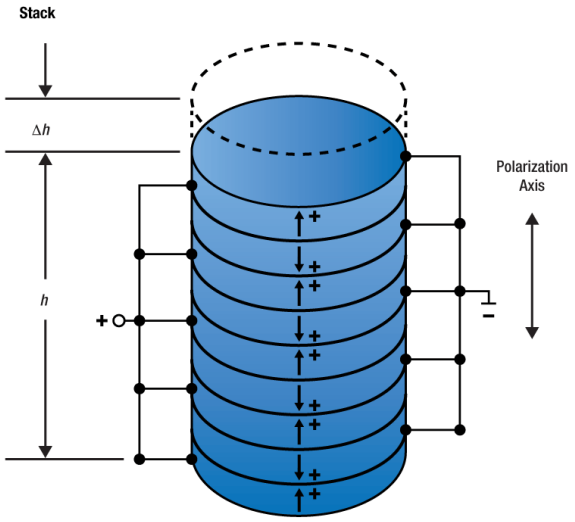

Figure 5.6: Working principle of a PZT stack actuator. PZT wafers are stacked ontop of each other and wired in parallel to receive the same control signal. Reproduced from [46].

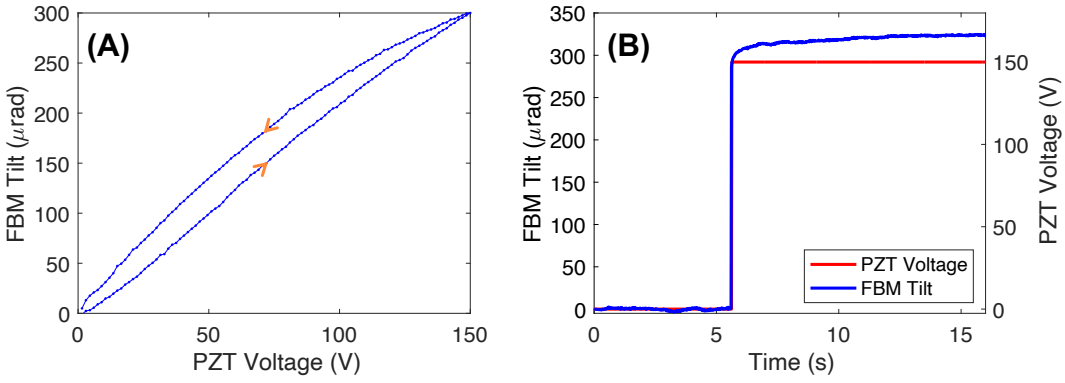

Figure 5.7: PZT nonlinearities. (A): Hysteresis. (B): Creep.

i7-10700K CPU @ 3.80 GHz with 8 cores. It runs a 64-bit operating system with 64 GB of RAM and 943 GB of SSD storage. Windows was chosen as the operating system because of its compatibility with LabVIEW and other common scientific instruments. A GigE-compatible network interface card (Intel I350-T4) with 4 ports was installed in the machine to interface with multiple cameras, and an analog-out PCIe card (National Instruments PCIe-6738) was also installed to output a $0-10$ V DC control signal to the FBM amplifier (Fig. 5.8).

5.1.2 Controller Software

The controller software was written in LabVIEW and relies on the Vision Acquisition Software [47] add-on to communicate with the GigE cameras. LabVIEW programs have two components, a graphical user interface (GUI) and the block diagram, which is the "code" for the program (Fig. 5.9). The GUI allows the user to observe the beam in real-time and adjust dynamic controls for the system. The three critical tasks the software must perform each iteration of the control routine are (1) acquire and process the raw images from the cameras, (2) calculate the position of the beam centroid in the coordinate system of the cameras, and (3) update and and output the control signal (Fig. 5.10). The software allows the user to adjust the signal gain, exposure time, and external trigger delay for each camera independently through the GUI. The image processing component consists of six steps: a user-selected region-of-interest (ROI), background subtraction, a zeroth-order filter, a particle filter, a linear filter of variable window size to smooth the image, and re-scaling of the image signal. Of these five steps, background subtraction is the only essential one in ideal conditions, but the ROI and particle filter are used to exclude unwanted signal artifacts that may be picked up by the camera. The zeroth-order filter is used to improve the accuracy of the centroid calculation by removing dead or damaged pixels that accumulate in the field due to radiation exposure. The smoothing filter should be used in conditions of

Figure 5.8: GigE and PCIe cards installed in the PC.

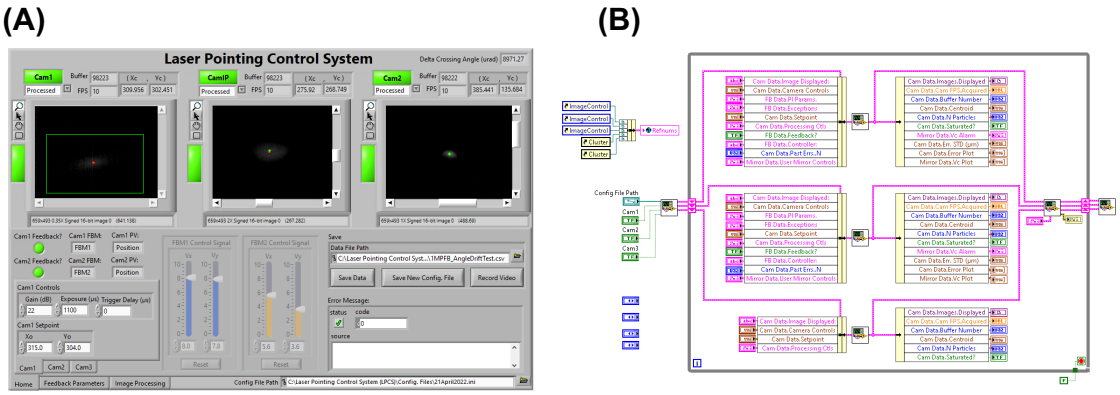

Figure 5.9: Control software (A) GUI and (B) block diagram.

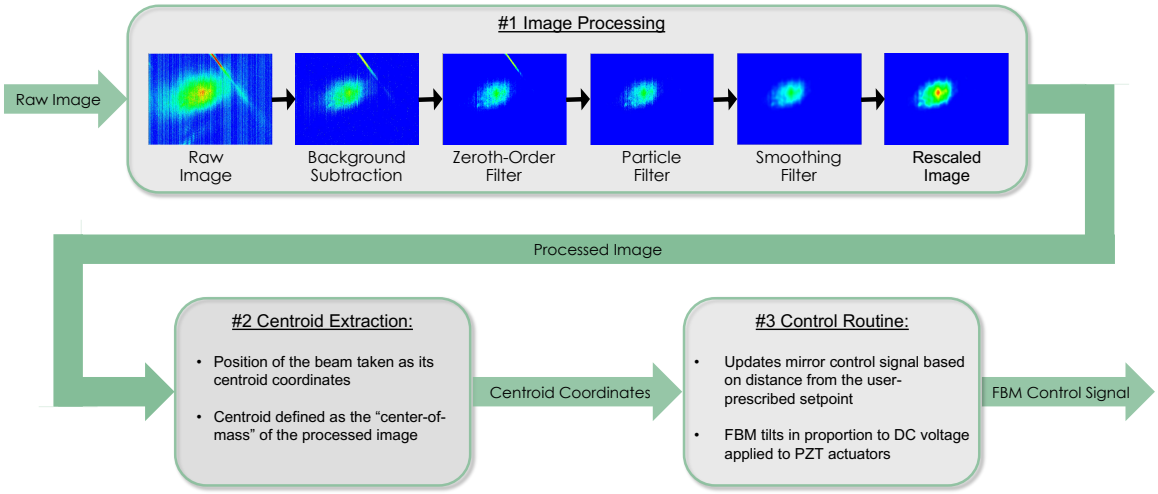

Figure 5.10: Three steps of the feedback routine. The controller software processes the camera images, extracts the beam centroid coordinates, and then updates the control signal.

poor beam quality, and re-scaling is should be done in situations where the signal-to-background ratio is low. The images of the beam shown in Fig. 5.10 were taken after the camera had been exposed to significant radiation and with an artifact present to demonstrate the robustness of the image processing routine. After processing, the beam centroid coordinates are calculated using the first-order moment, or "center-of-mass" definition. The formula is given by

$$(r_h, r_v) = \left(\frac{\sum_{i=1}^{N}\sum_{j=1}^{M} i \cdot p_{i,j}}{\sum_{i=1}^{N}\sum_{j=1}^{M} p_{i,j}}, \frac{\sum_{i=1}^{N}\sum_{j=1}^{M} j \cdot p_{i,j}}{\sum_{i=1}^{N}\sum_{j=1}^{M} p_{i,j}} \right) \tag{5.7}$$

where M and N are the number of rows and columns of pixels in the sensor, respectfully, i and j are the column index and row index, and $p_{i,j}$ is the intensity of an individual pixel. The error signal $e_k = r_k - r_{sp}$ is then calculated for each dimension, where r_k is equivalent to r_h or r_v in Eq. (5.7) at the k^{th} iteration of the control routine. The program then updates the control signal V_k according to the control law given by Eq. (4.4). The control law can be written for either dimension in terms of the voltage to the actuators as

$$V_{k+1} = V_k + K \cdot \chi^{-1} \Delta\phi_k. \tag{5.8}$$

where $\Delta\phi_k$ is the tilt required to shift the beam back to the setpoint. $\Delta\phi_k$ is calculated from e_k, and its functional form depends on whether position or angle of the beam is the process variable. When position is the process variable, $\Delta\phi_k = -e_k/L$, where L is the distance from the FBM to the Position Camera, and $L >> e_k$. When angle is the process variable $\Delta\phi_k = -e_k/f$, where f is the focal length of the lens before the Angle Camera, and $f >> e_k$.

For this system, the horizontal axis of the FBM stabilizes the angle and the vertical axis stabilizes the position. Thus, the control law can be written out for the horizontal

axis of the FBM as

$$V_{h,k+1} = V_{h,k} - K_h \cdot \chi_h^{-1} \frac{r_{h,k} - r_{h,sp}}{f}, \tag{5.9}$$

where $r_{h,k}$ and $r_{h,sp}$ are the horizontal components of the beam centroid and setpoint for the Angle Camera, respectively. The control law for the vertical axis of the FBM is

$$V_{v,k+1} = V_{v,k} - K_v \cdot \chi_v^{-1} \frac{r_{v,k} - r_{v,sp}}{L}, \tag{5.10}$$

where $r_{v,k}$ and $r_{v,sp}$ are the vertical components of the beam centroid and setpoint for the Position Camera, respectively. For each dimension, the feedback gain K is chosen to minimizes the RMS of e_k, and its optimal value depends on the specific disturbance signal present in the system.

5.2 Laboratory Test

The feedback system was tested under realistic conditions in the field with the DM reproducing a disturbance signal previously measured with the un-stabilized UV laser. Before the test, the feedback gain was optimized for each dimension using a simulation with the disturbance signals used in the test. The test was done using the optimal feedback gain for each axis, and the RMS error for the horizontal angle and vertical position measured in the test were very close to the values predicted by the simulation.

5.2.1 Feedback Gain Optimization

A MATLAB simulation was used to optimize the feedback gain for the horizontal angle and vertical position disturbance signals generated by the DM by comparing the RMS error of each parameter for different values of K. The simulation is based on the model given by Eq. (4.1), in which the process variable y_k is the sum of the disturbance signal x_k and the control action u_k. The control law used to calculate u_k in the simulation is given by Eq. (4.4). Prior to the simulation, each axis of the DM generated a 20-minute disturbance signal with feedback OFF that was measured by

the cameras at a sampling rate of 10 Hz. The horizontal axis of the DM generated x_k for the horizontal angle, measured by the Angle Camera, and the vertical axis of the DM generated x_k for the vertical position, measured by the Position Camera. The disturbance signals measured in the laboratory were then used in the feedback simulation, and the RMS error of each parameter was calculated for different values of K. The simulation was done for the stable values of K, determined in Section 4.2.1, ranging between 0 (feedback OFF) to 2 (unstable). The value that resulted in the smallest RMS error of each parameter was taken as the optimal value that would be used in the laboratory test.

The results of the feedback simulation are shown in Fig. 5.11. The RMS error of both parameters is normalized to the RMS of the disturbance signal (feedback OFF) because the ratio is independent of the amplitude of x_k used in the simulation. The simulation shows the RMS error of both the horizontal angle and vertical position is reduced for most values of K tried. It also shows the RMS errors that result from values of K in the neighborhood of the optimal value are nearly identical. However, as K approaches 2, the RMS error becomes greater than it would be with feedback OFF, indicating feedback reduced the stability of the system. It was explained in Section 4.2.1 that the amplification of high frequency components of x_k increases with K. Above a certain value of K, determined by the specific disturbance signal, the increased jitter that results from feedback will lead to reduced overall stability, even though the beam does not drift. From the simulation, it could be concluded that the optimal value of K in the field would likely be below 1, and that the performance of the feedback system would not be especially sensitive to small deviations from it. The optimal values of K found for each process variable and the normalized RMS errors predicted by the simulation are given in Table 5.2.

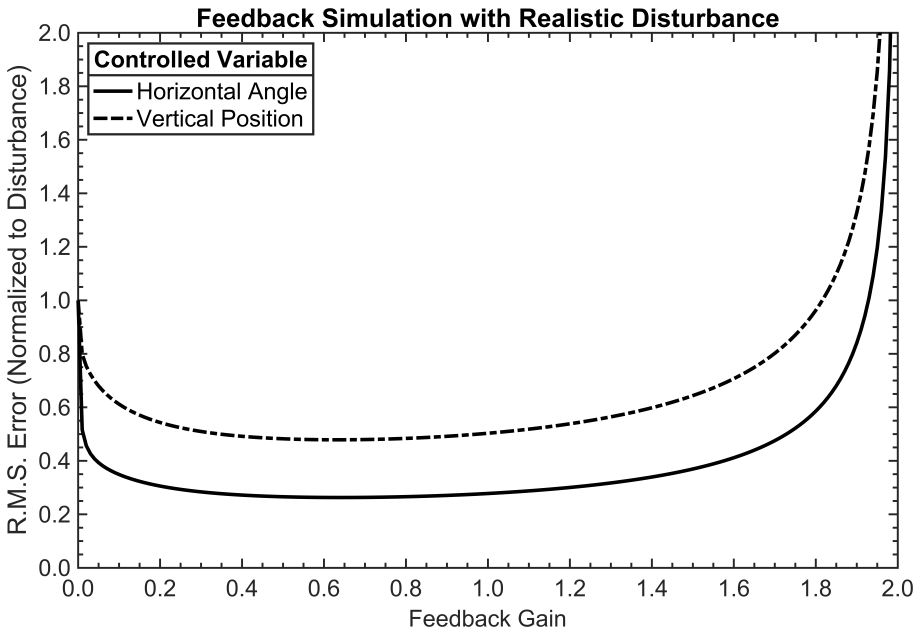

Figure 5.11: Simulation to find the optimal feedback gain. The optimal value of the feedback gain is the one that minimizes the RMS error of the controlled variable for a given disturbance signal.

Table 5.2: Results of the laboratory simulation.

Process Variable	Optimal K	RMS Error (Normalized)
Horizontal Angle	0.64	0.48
Vertical Position	0.63	0.26

5.2.2 Results

For the laboratory test, the DM was again programmed to generate the artificial disturbances that were used in the simulation, but this time feedback was turn ON and the optimal gains were used. The results of the laboratory test are shown in Fig. 5.12. The horizontal angle and vertical position with feedback OFF (red) and ON (black) are plotted alongside one another to show the effect of the feedback system. Comparing the feedback OFF and ON waveforms, it is clear the feedback system effectively eliminated the slow drift, and the residual motion of the beam was purely jitter caused by the high frequency components of the disturbance signal. The Fourier amplitude spectra of the waveforms confirm this conclusion. The low frequency components seen in the feedback OFF spectra, associated with drift, were reduced to negligible amplitudes with feedback ON. The spectra also show the amplitudes of frequency components above ∼1 Hz were increased, which is to be expected based on the frequency domain analysis of Section 4.2.1. The feedback ON spectra show the "leftover" high frequency components, which combine to manifest as jitter in the time domain, are mostly confined to the frequency band between 0.1 Hz and the Nyquist frequency (5 Hz).

The feedback OFF waveforms plotted in Fig. 5.12 were measured first and used as the disturbance signals in the simulation. The RMS of those waveforms are what the RMS of the horizontal angle and vertical position with feedback ON were normalized to. Table 5.3 lists the un-normalized RMS values with feedback ON for the experimental data and those predicted from the simulation. The predicted and measured RMS values closely match, suggesting the optimal K found from the the simulation were correct, or at least very close. Possible sources of error in the laboratory test include the values of the FBM response coefficients χ, which were not a parameter of the simulation, and the angular resolution of the FBM and DM. Based on comparison of the measured and predicted stability with feedback ON, the results of the laboratory test are representative of the optimal performance of the feedback

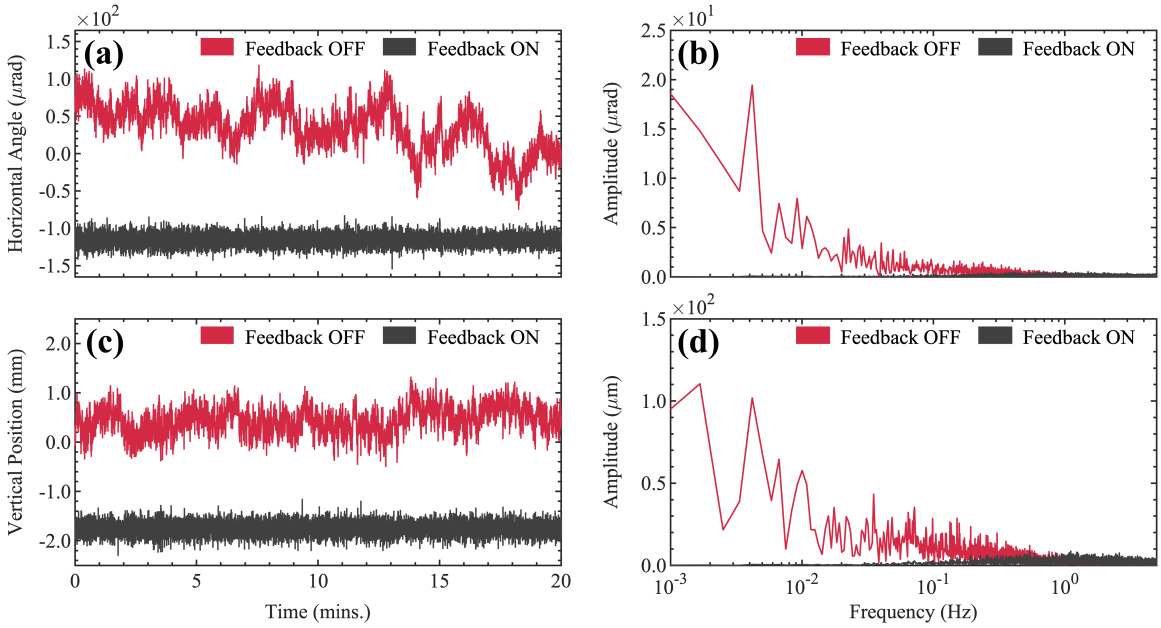

Figure 5.12: Results of the laboratory test. (a) waveforms and (b) spectra of the horizontal angle measured by the Angle Camera with feedback OFF and ON. (c) waveforms and (b) spectra of the vertical position measured by the Position Camera with feedback OFF and ON

Table 5.3: Measured and predicted stability with feedback ON.

Process Variable	Measured RMS Error	Predicted RMS Error
Horizontal Angle (μrad)	8.1	8.2
Vertical Position (μm)	135	129

system for typical disturbances in the LACE optical system. The stabilization system could be expected to effectively control drift when deployed to the field, but it became apparent that the best overall stability achievable with the UV laser would be defined by the jitter.

Chapter 6

Stabilization System Results in the Field

The laser pointing feedback system was successfully deployed to the field after it was developed and tested in the laboratory. Since its deployment, it has been used to stabilize the UV laser beam for two LACE experiments conducted in 2022 and 2023. As discussed in Chapter 2, the 10 μs Experiment in 2016 was a milestone that extended the length of the stripped ion beam pulse by a factor of 1000 from the POP experiment in 2006. However, it is still not possible to scale LACE for the full ion beam duty factor used for neutron production due to the limited laser power available with current technology. The two recent LACE experiments were demonstrations of additional techniques that would allow further scalability and advancement of the project. In this chapter, an overview of both LACE experiments is given. Then, the optical configurations of the laser beam transport and pointing stabilization systems are detailed. After that, results are shown for the pointing stability of the UV beam at the laser-particle IP and compared to the results of the laboratory test. Last, the impact the system had on the variance in LACE efficiency is shown by comparing distributions of efficiency measurements made with feedback OFF and ON.

6.1 Overview of Recent Experiments

The first new LACE experiment the pointing stabilization system was used for was the Sequential Excitation Experiment. It was a successful demonstration of the concept outlined in [48], whereby the energy per laser pulse required for efficient photoexcitation of the neutralized ion beam (H^0) is reduced by splitting the resonant excitation from $n = 1$ to $n = 3$ into two steps ($n = 1 \rightarrow 2 \rightarrow 3$). The experiment was done with the existing hardware from 2016, and it could not be optimized for high efficiency. A detailed report of the Sequential Excitation Experiment is given in [49]. The second LACE experiment is the Crab Crossing Experiment, which reduces the required laser pulse energy through manipulation of the ion beam. The concept is to rotate the ion beam in the horizontal plane by a certain angle so that complete overlap of the H^0 will still occur with a shorter laser pulse, thus reducing the energy per laser pulse while retaining the required peak power [50]. The results of the Crab Crossing Experiment are still preliminary, and a paper is expected later in 2023.

6.1.1 The Sequential Excitation Experiment

The sequential excitation scheme proposed in [48] is to excite the electron of the H^0 atom from the ground state to $n = 3$ with two sequential resonant excitations from a green laser beam recycled inside an optical cavity constructed around the IP (Fig. 6.1). The sequential excitation scheme has two advantages over the single step excitation. First, the scheme is more flexible because different laser wavelengths can be used other than 355 nm, such as 515 nm or 532 nm. The green wavelengths are generated from the 2nd harmonic of the fundamental wavelength, as opposed to the 3rd harmonic in the case of 355 nm. Compared to the UV laser, it will be easier to create a high quality green laser beam with the required parameters, transport it to the LACE chamber, and then recycle it inside an optical cavity. Second, the total laser power required to excite the electron from 1s to 2p and then from 2p to 3d is less than is required for the 1-step excitation from 1s to 3p used in the

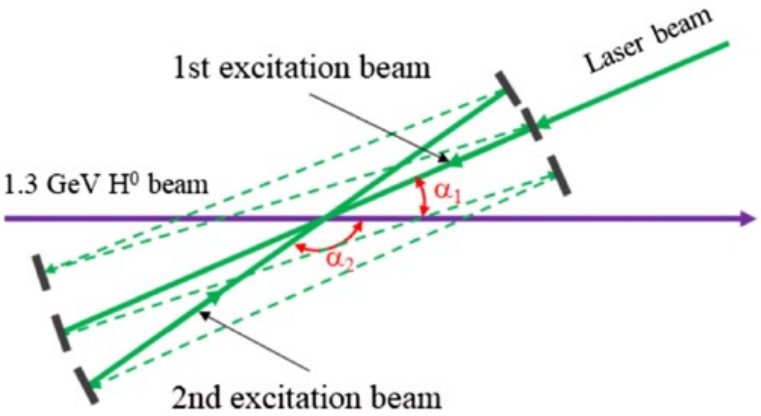

Figure 6.1: Sequential excitation with a green laser beam inside an optical cavity. Reproduced from [48].

10 μs Experiment. The less stringent laser power requirement is due to the stronger electric transition dipole compared with the single step. Another sequential excitation scheme that could be done with the existing UV laser is to use 355 nm light for the first excitation and the fundamental 1064 nm wavelength for the second. The relative laser power requirements for each scheme are given in Fig. 6.2, in which the color purple indicates 355 nm light, green indicates 515 nm or 532 nm, and red indicates 1064 nm. Sequential excitation with the optical cavity is estimated to result in an order of magnitude laser power savings.

A proof-of-concept experiment that could be done with the existing chamber and UV laser was also suggested in [48], and the experiment was completed in late 2022. Sequential excitation was accomplished by splitting the UV laser beam before the chamber to create two beams that overlapped at the IP (Fig. 6.3). A photomultiplier tube (PMT) sensitive to the wavelength of the fluorescent photons from the $n = 2 \rightarrow 1$ de-excitation was installed to measure the efficiency of the first excitation, and the full LACE efficiency was measured with the BCM after the chamber (Fig. 2.2). The existing chamber was designed to optimize the LACE efficiency for the 10 μs Experiment, and a redesign of the chamber for sequential excitation would require a significant investment of time and resources. The experimental parameters that could not be altered in the design are the laser wavelength ($\lambda = 355$ nm), the laser-ion crossing angles, $\alpha_1 = 37.5^o$ and $\alpha_2 = 142.5^o$, and the strength of the two permanent magnets. The maximum LACE efficiency achievable with these parameters is below the desired >95% but is sufficient to demonstrate the method. Due to the constrains on the laser wavelength and laser-ion crossing angles, the ion beam energy was lowered from 1 GeV to 718 MeV, which caused the vertical size of the beam to increase by a factor of about 5 at the IP. Because the transverse size of the laser beam was increased to fully overlap the ion beam, the photon density was reduced, which resulted in a decrease in photoexcitation efficiency of about 50% for the first excitation. After the beam energy was set for the $n = 1 \rightarrow 2$ excitation, there were no free parameters left to precisely tune the $n = 2 \rightarrow 3$ excitation from the second laser beam. However, the

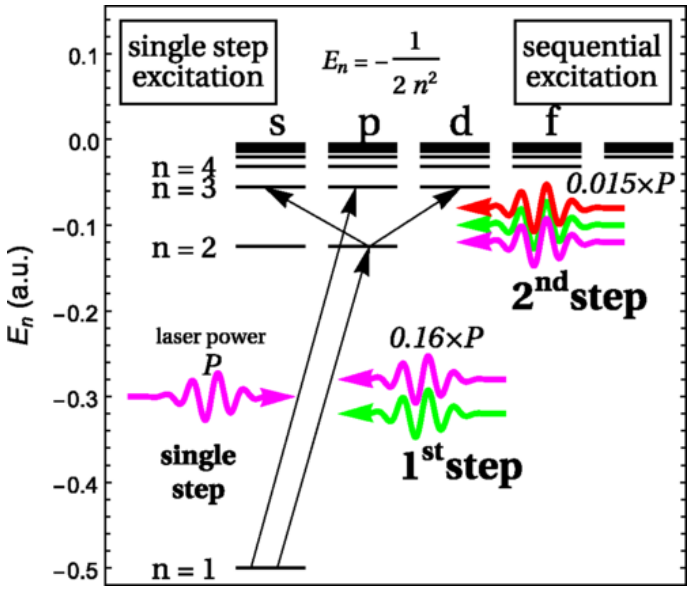

Figure 6.2: Laser power required for single and sequential excitation schemes. Reproduced from [48].

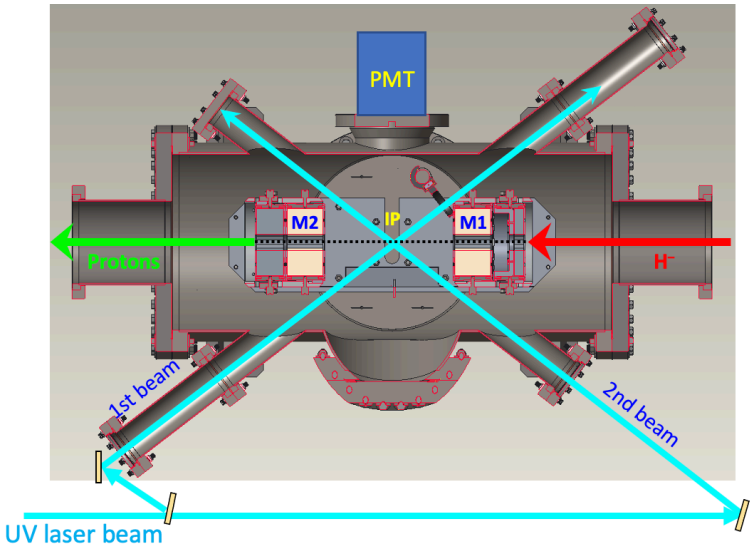

Figure 6.3: Sequential excitation inside the existing LACE chamber. The UV beam is split before the LACE Chamber and the two beams are aligned at the IP. The PMT is used to measure the efficiency of the first excitation from fluorescent photons.

$n = 2 \rightarrow 4$ excitation could be accomplished by modifying α_2 with a system of optical prisms inserted into the chamber. The transition efficiency of $n = 2 \rightarrow 4$ is about $8\times$ less than for $n = 2 \rightarrow 3$, which contributed further to the reduction in LACE efficiency. Parameters of the ion beam used in the experiment are listed in Table 6.1.

The overall scheme of the MOPA laser system for the Sequential Excitation Experiment is the same as that given in Section 2.1, but some modifications are to be discussed. The mode-locked seed laser from the 10 μs Experiment was replaced with a single-frequency, narrow linewidth (<5 kHz) fiber laser with a wavelength around 1064.5 nm (Fig. 6.4). The new seed laser produced a 50 mW CW beam that was fiber coupled to an intensity modulator to produce the necessary pulse structure. The intensity modulator used two electro-optic lithium niobate Mach-Zehnder modulators (EOMs) to produce 50 ns pulses at 350 kHz. The pulse length and repetition rate were determined by the signal from an external arbitrary waveform generator (AFG). The waveform of the signal sent to the EOMs by the AFG was programmed to produce a square laser pulse shape. The output of the intensity modulator was pre-amplified by a polarization maintaining (PM) fiber amplifier before being sent to a second fiber amplifier to give the output beam an average power of 3 mW. After the second fiber amplifier, the rest of the laser architecture is the same. Important parameters of the laser beam used for the Sequential Excitation Experiment are given in Table 6.2.

6.1.2 The Crab Crossing Experiment

The photon density inside a laser pulse of a given energy can be increased by reducing its transverse size or length. The laser pulse cannot be smaller than the ion beam if complete overlap is to occur, and so the minimum transverse size of the laser beam is determined by the size of the ion beam at the IP. The transverse size of the ion beam can be compressed to the limit determined by its emittance using quadrupole magnets upstream of the IP. However, there are no longitudinal focusing elements in the ∼240 m section of the HEBT between the exit of the linac and the LACE

Table 6.1: Ion beam parameters for the Sequential Excitation Experiment.

Parameter	Value
Energy	718.5 MeV
Beam Current	30 mA
Pulse Length	1 μs
Pulse Repetition Rate	1-5 Hz
Transverse Emittance (RMS)	0.5 μm
Vertical Size at IP (RMS)	0.5 mm
Horizontal Size at IP (RMS)	1.0 mm
Momentum Spread (RMS)	0.5-1.0 $\times$ 10^{-3}
Horizontal Angular Spread (RMS)	1.0 mrad

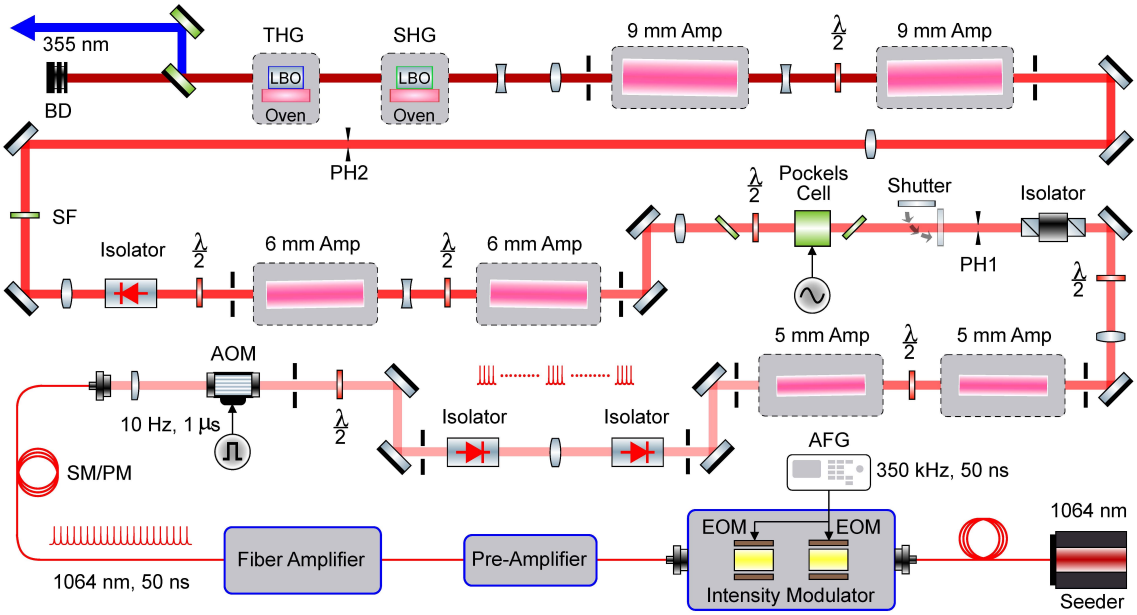

Figure 6.4: Configuration of the UV laser for the Sequential Excitation Experiment. Reproduced from [49].

Table 6.2: Laser beam parameters at the IP for the Sequential Excitation Experiment.

Parameter	Value
Wavelength	355 nm
Pulse Width (FWHM)	50 ns
Pulse Repetition Rate	10 Hz
Peak Power (total for both beams)	0-1.3 MW
1st Beam Vertical Size (RMS)	0.6 mm
1st Beam Horizontal Divergence	1.4 mrad
2nd Beam Vertical Size (RMS)	0.9 mm
2nd Beam Horizontal Divergence	0.08 mrad

chamber, so the ion beam expands longitudinally due to the energy spread of its constituent particles (Fig. 6.5). A dipole magnet is used to guide the beam through the 90^o bend in the HEBT (Fig. 1.4). After the bend, particles with different energies follow different trajectories, defined by the dispersion function of the beamline. The beamline optics are configured to provide a nominal dispersion function of zero from the exit of the bend to the ring injection site. The zero dispersion function keeps the longitudinal axis of the beam aligned with its momentum vector, but if there is a non-zero dispersion function at the IP, the bunch will rotate in the horizontal plane while its momentum vector remains pointed in the same direction (Fig. 6.6). This technique is called crab crossing because the head-to-tail direction of the bunch is different from its direction of motion, resembling the sideways walk of a crab. With a proper rotation angle, crab crossing at the micropulse-level will allow efficient overlap of the ion beam with a shorter laser pulse. For the experiment, the mode-locked seed laser from the 10 μs Experiment was reinstalled to produce the macropulse structure of the laser beam (Table 6.3). The Crab Crossing Experiment is ongoing, and the preliminary results are promising. The final results of the experiment will be reported in a paper expected to be published in late 2023.

6.2 Optical Configuration in the Field

The configurations of the optical transport and laser pointing stabilization systems for LACE are shown in Fig. 6.7. The setup is identical for the Sequential Excitation and Crab Crossing Experiments, except the second laser beam (entering the chamber from the right) is not used for the latter. The LACE chamber is installed between two quadrupole magnets (QH28 & QV29) in the straight section of the HEBT upstream of the ring injection area. As discussed in Section 2.1, the laser is located in the RSB to protect it from radiation-induced damage, and the beam is transported to the LACE chamber over a total distance of 65 m. Most of that distance is covered by the 60-meter, free-space LTL that connects the laser table in the RSB to LOT1. In

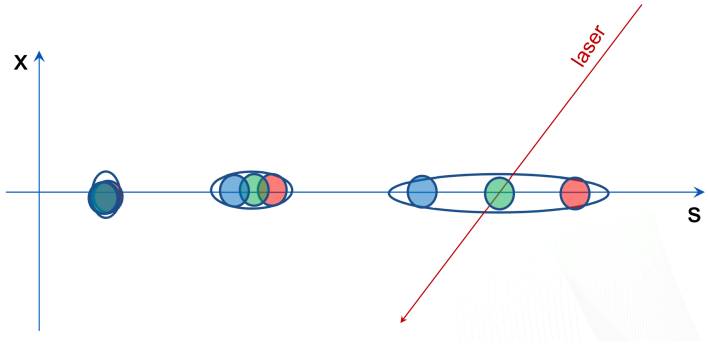

Figure 6.5: Longitudinal expansion of the ion beam due to the energy spread of the constituent particles. Red: high energy particles. Green: nominal energy particles. Blue: low energy particles. Reproduced from [50].

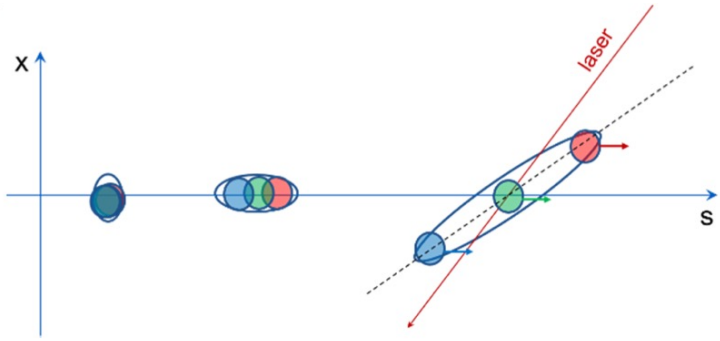

Figure 6.6: Crab crossing due to non-zero dispersion function. The ion beam rotates in the horizontal plane while its momentum vector stays pointed in the same direction. Reproduced from [50].

Table 6.3: Laser pulse structure for the Crab Crossing Experiment.

Parameter	Micropulse Value	Macropulse Value
Pulse length (FWHM)	~50 ps	1.5 μs
Repetition Rate	402.5 MHz	10 Hz
Energy	100 μJ	60 mJ

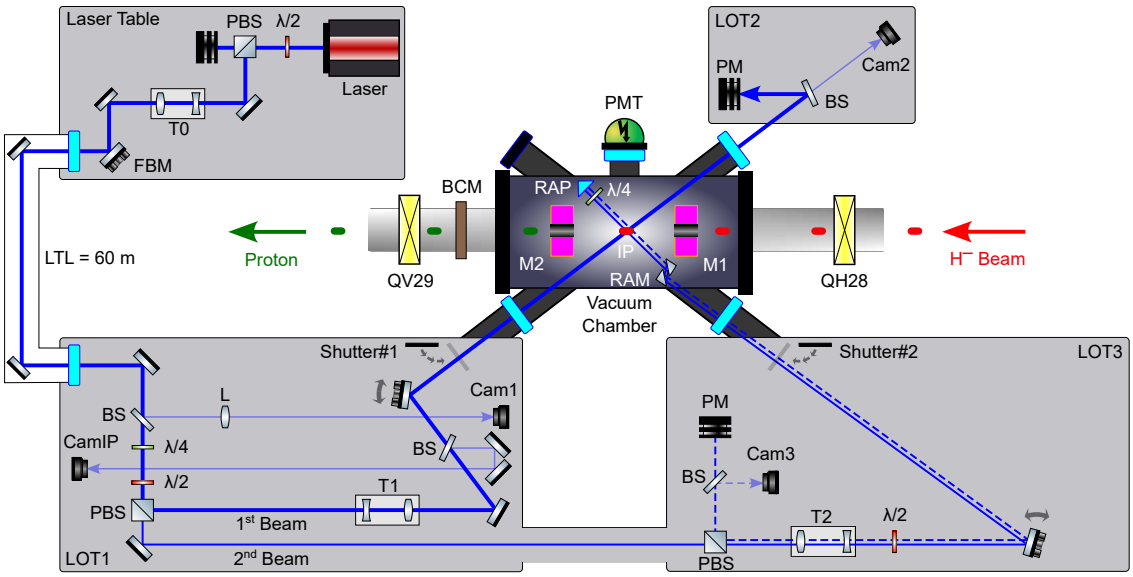

Figure 6.7: Experimental configuration of the optical transport and pointing stabilization systems. The dashed lines represent the retro-reflected light out of the vacuum chamber from the second beam. The horizontal angle and vertical position of the first beam are stabilized at the IP by feedback between the FBM and Cam1 (angle) and CamIP (position). PBS: polarized beam splitter, FBM: feedback mirror, T0, T1, T2: telescopes, LTL: laser transport line, BS: beam splitter, PMT: photo-multiplier tube, M1, M2: magnets, QH28, QV29: quadrupole magnets, BCM: beam current monitor, IP: interaction point, RAP: right angle prism, RAM: right angle mirror, $\lambda/2$: half-wave plate, $\lambda/4$: quarter-wave plate. Reproduced from [40].

the RSB, a telescope (T0) expands and collimates the beam before the LTL to avoid damaging any of the mirrors that cannot be easily accessed. The transport efficiency from the RSB to the IP was measured to be about 45%, and the beam losses come mainly from absorption and scattering on the mirror surfaces and air. The half-wave plate ($\lambda/2$) before the polarizing beam splitter (PBS) on LOT1 is used to control the amount of light going to the first and second laser beam. Telescopes (T1 & T2) before the LACE chamber shape each laser beam to give them the size and divergence for the interaction with the H^0 called for by the experiment. The optical system of right angle mirrors (RAM), a right angle prism (RAP), and a quarter-wave plate ($\lambda/4$) inside the chamber modifies angle of the second beam to produce the $n = 2 \rightarrow 4$ resonant excitation. After the interaction with the H^0, the second beam is retro-reflected back out of the chamber so that its position and power can be monitored by Cam3 and a power meter (PM). For the Sequential Excitation Experiment, around two thirds of the laser power was directed to the first beam, and the rest was directed to the second beams. For the Crab Crossing Experiment, the half-wave plate is configured so that all laser power is to sent to the first beam. The BCM after the chamber is used to measure the total LACE efficiency, and the PMT was used to measure the efficiency of the $n = 1 \rightarrow 2$ excitation in the Sequential Excitation Experiment.

The feedback system is configured to stabilize the horizontal angle of the main beam before it is split by the PBS and the vertical position of the first beam at the IP. The FBM is the final mirror before the LTL entrance window, and the two feedback cameras are located on LOT1. The horizontal angle of the main beam is stabilized because the resonant excitations from both laser beams in the Sequential Excitation Experiment are sensitive to the laser-H^0 crossing angle. Both the position and angle are measured using the optical technique described in Section 5.1. For the angle measurement, a fraction of the main beam is sampled by a BS before it is split by the PBS, and the sensor of Cam1 is placed in the focal plane of a $f = 1000$ mm lens (L). For the position measurement, the first beam is split by a BS after T1, and CamIP is placed where the reflected beam is an image of the transmitted beam

at the IP. Laser power is attenuated before the cameras with neutral density filters (not shown) to prevent damage and keep their sensors slightly below the saturation limit. The angle of the first beam at the IP is reconstructed offline from the positions measured by CamIP and Cam2 and the known distance between the IP and Cam2. Fig. 6.8 shows the correlation of the angles measured at Cam1 and reconstructed at the IP. The correlation coefficient ρ is close to 1 in each dimension, indicating a strong correlation between the angles at both locations. The strong correlation justifies the method of stabilizing the angle of the main beam to indirectly stabilize the angles of the first and second beams. The slope m of the linear fit is explained by the magnifying effect T1 has on the angular fluctuations of the main beam.

The feedback PC is located on a rack in the RSB, and it can be remotely connected to through any other PC on the SNS internal network. Long Ethernet cables were pulled through a chase adjacent to the one repurposed for the LTL for the cameras and routed to the PC. The control signal for each axis of the FBM is channeled from the PC to the PZT amplifier near the laser table through 50-meter BNC cables. The stabilization system is able to capture every laser pulse and execute the feedback routine at the maximum rate of 10 Hz, despite the large amount of data to process and wide area covered.

6.3 Results

Results for the pointing stability of the UV laser and the impact the stabilization feedback system had on the variance in LACE efficiency are reported in this section. First, the performance of the stabilization system in the field is reported on and compared to the results of the laboratory test in Section 5.2.1. Then, the effectiveness of the stabilization system for controlling the variance in LACE efficiency is analyzed by comparing distributions of efficiency measurements taken with and without feedback.

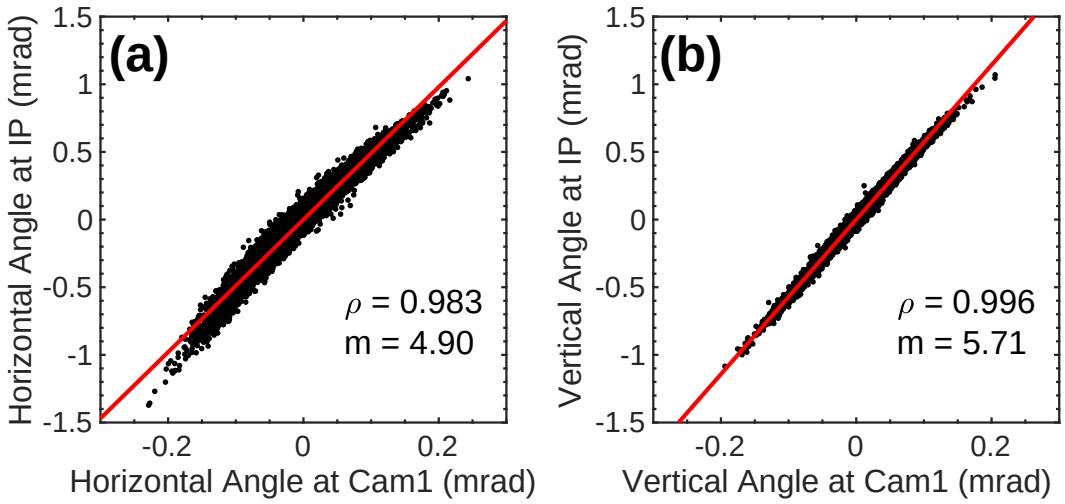

Figure 6.8: Correlation between the beam angles measured at the IP and by Cam1. (A): Horizontal correlation. (B): Vertical correlation. Reproduced from [40].

6.3.1 Laser Pointing Stability

First results with the laser pointing stabilization system were obtained shortly after it was installed during an accelerator shutdown period. The results discussed here are also reported on in [40]. The laser beam is expected to be more stable under shutdown conditions compared to during neutron production because magnets and other instruments in the tunnel that create disturbances are shut off. The feedback gain was tuned by incrementally increasing the value from zero (feedback OFF) while visually observing the response of the system for around 30 seconds at each increment. The feedback gain was increased until the jitter was amplified above an acceptable level. Then, the feedback gain was returned to what was believed to be the optimal value, and 20-minute data sets were recorded for that and several other values in the neighborhood to compare and asses the sensitivity of the system. The RMS values for the 2-dimension position and angle of the beam during the tuning are plotted in Fig. 6.9. As was discussed in Section 5.2.1, there is a range of values of the feedback gain for which the performance of the stabilization system is nearly equivalent. Each value tried reduced the RMS values for the 20-minute data sets by a factor of about 2 from the feedback OFF case. Even as disturbance conditions evolve over time, the feedback gain can easily be tuned by observing the response of the system on the software GUI. For the conditions under which this test was conducted, the optimal values for the feedback gain for the horizontal and vertical axes of the FBM were found to be 0.1 and 0.4, respectively.

The results for the horizontal angle and vertical position of the beam at the IP, recorded for 20 minutes, are shown in Fig. 6.10. The performance of the system in the field is very similar to in the laboratory test. With feedback ON, drift was eliminated, and the beam remained jittering about the setpoints for the duration of the measurement. The histograms for the 20-minute data sets show the RMS error shrank $\sim 2\times$ from the feedback OFF case. The Fourier spectra show the amplitudes of the low frequency components present in the feedback OFF spectra, associated with

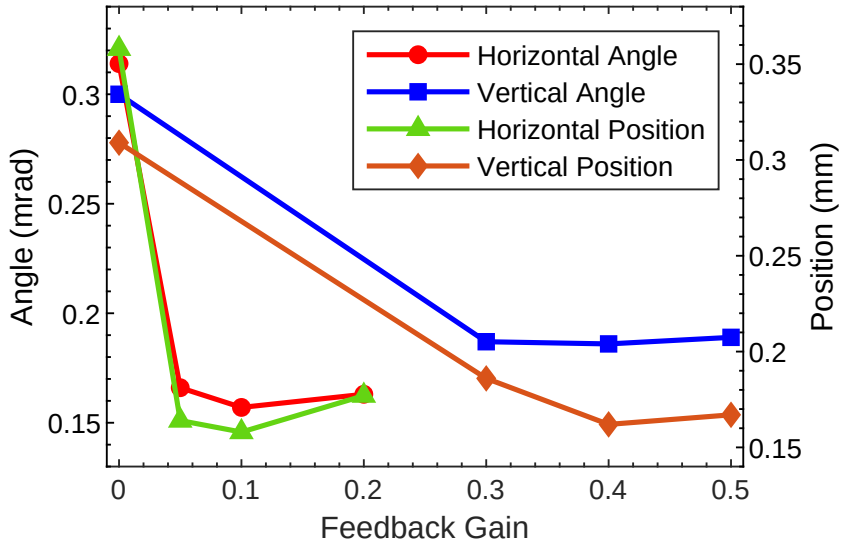

Figure 6.9: Tuning the feedback system by comparing the stability for several values of the feedback gain. Reproduced from [40].

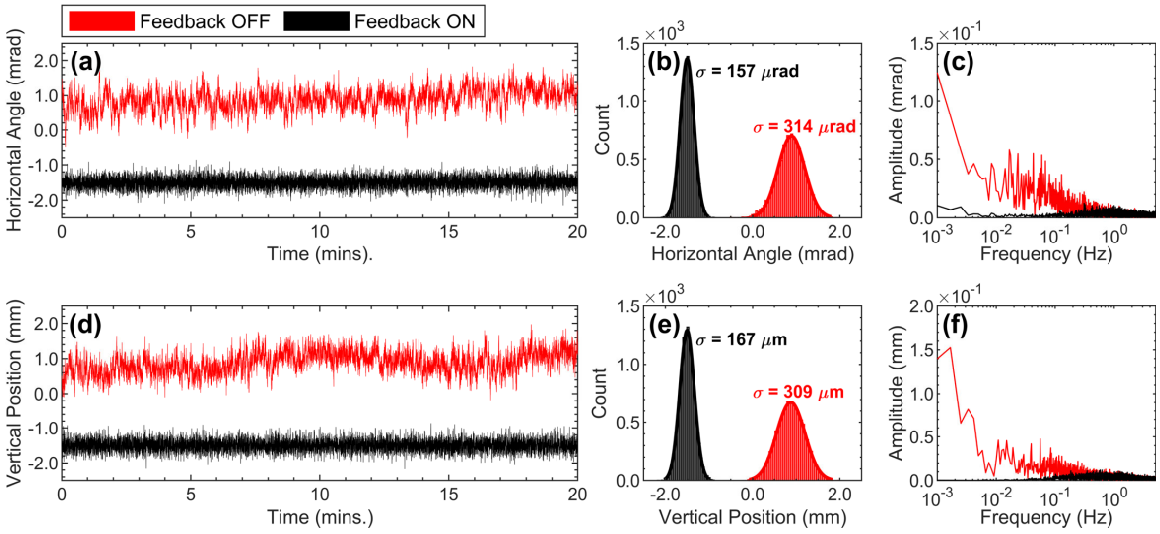

Figure 6.10: Stability of the sensitive laser pointing parameters. Top: the (a) waveforms, (b) distributions and (c) spectra of the horizontal angle of the first beam at the IP with feedback OFF and ON. Bottom: the (d) waveforms, (e) distributions and (f) spectra of the vertical position of the first beam at the IP with feedback OFF and ON. Reproduced from [40].

drift, were greatly reduced, while the components above ∼0.1 Hz, associated with the jitter, remained when feedback was turned ON. The laboratory and field tests both showed residual frequency components within the ∼0.1 Hz to 5 Hz band, and components above about 1 Hz were amplified, as expected from the analysis in Section 4.2.1. Fig 4.6 shows the trade-off between suppressing drift and amplifying jitter. In this application, the optimal value of the feedback gain is that which balances those two effects to minimize the RMS error of the stabilized parameter. The 2-dimensional position and angle of the laser beam at the IP with feedback OFF and ON are shown in Fig. 6.11, and the RMS values of each parameter is given in Table 6.4. Although only 20 minute data sets are reported, the feedback system will keep the beam centered on the setpoints for an indefinite amount of time. If longer feedback OFF data could be collected without the beam wandering out of the field-of-view of the cameras, a reduction in the RMS error at least an order of magnitude greater than a factor of 2 would have been observed due to the slow drift of the laser beam. The RMS error with feedback ON remains constant for data sets longer than ∼20 minutes (at a sampling rate of 10 Hz) due to the collection of enough samples for reliable statistics.

6.3.2 Laser-Assisted Charge Exchange Experiments

Setting up a LACE experiment is a time consuming process, and aligning the laser to the ion beam is the most challenging part. The setpoints for the feedback system are recorded after the laser beam is aligned through the LACE chamber by hand during a maintenance, so they only approximately represent the correct alignment. The procedure for aligning the laser beam at the beginning of an experiment starts with finding the exact position of the ion beam with the wire scanner, inserting the stripping magnets to neutralize the ion beam while the wire scanner is left in place, and then manually steering the laser beam with pico-motor mirrors so as to maximize the signal of thermally-induced electrons from the interaction of the laser beam with the wire. Several years had past since the 10 μs Experiment, and, because setting

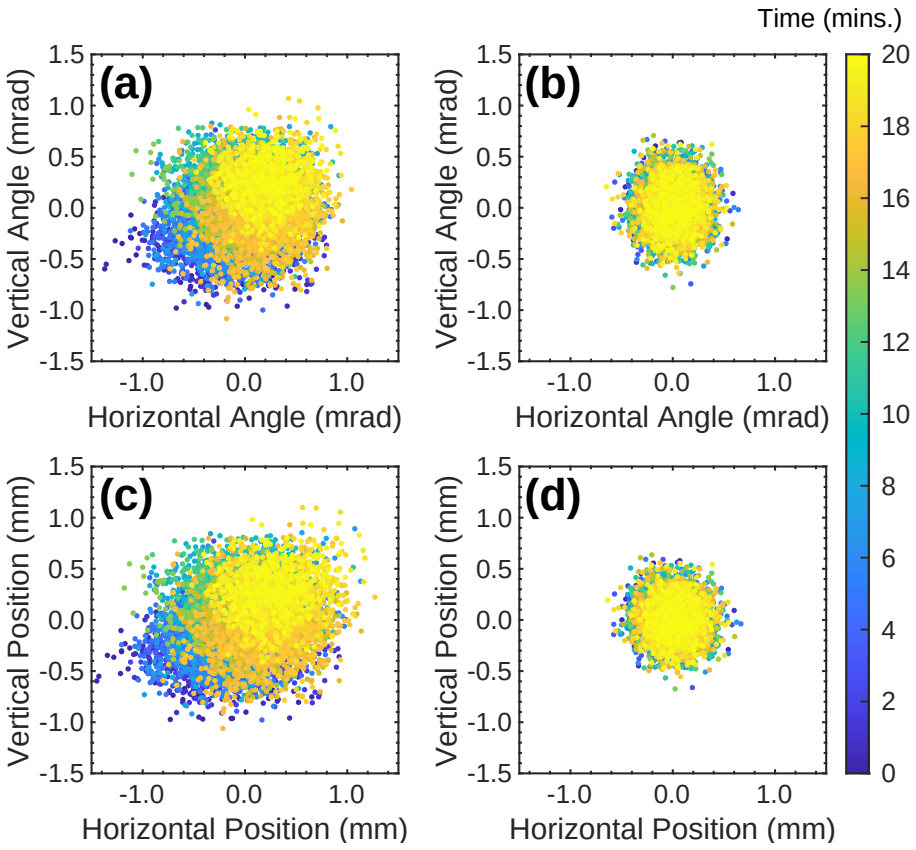

Figure 6.11: Laser pointing stability at the IP with feedback (a,c) OFF and (b,d) ON. Reproduced from [40].

Table 6.4: Laser pointing stability in the field for the feedback ON and OFF cases for 20-minute measurements.

Parameter	RMS Position (μm)	RMS Angle (μrad)
Feedback OFF	358 (H) $\times$ 309 (V)	314 (H) $\times$ 300 (V)
Feedback ON	158 (H) $\times$ 167 (V)	157 (H) $\times$ 189 (V)

up the Sequential Excitation Experiment requires the alignment of two laser beams, the LACE Team decided it was best to relearn the procedure by repeating a 1-step excitation experiment. Although the experiment was not set up for optimal efficiency, it was the first chance to observe the effect the feedback stabilization system had on the variance in efficiency.

LACE efficiency was recorded during the 1-step experiment for about 20 minutes with feedback OFF and ON to compare the variance. The widths of the efficiency distributions show the stabilization system significantly reduced the variance of the measurements (Fig. 6.12). The STD of the measured efficiencies with feedback OFF was 8.53% and fell to 3.68% with feedback ON, representing more than a factor of 2 improvement. Timestamps from the SNS universal timing system were used to match LACE efficiency and laser pointing data that were taken simultaneously. The waveforms for the LACE efficiency and laser pointing parameters (in the laser reference frame) for the feedback ON data set are shown in Fig. 6.13. The corresponding feedback OFF data set was already shown in Fig. 3.2 to justify the need for a laser pointing stabilization system. As expected, the laser beam was not drifting in the feedback ON data set, and the remaining efficiency variance is caused primarily by the jitter of the laser beam. With feedback ON, the RMS of the position and angle of the laser beam at the IP were measured to be 195 (H) $\times$ 199 (V) μm and 263 (H) $\times$ 315 (V) μrad, respectively. The corresponding RMS values for the feedback OFF data set were 623 (H) $\times$ 347 (V) μm and 744 (H) $\times$ 535 (V) μrad. Comparing the RMS values, the feedback system improved the stability of the laser beam 2–3$\times$ for all parameters. Fig. 3.3 was used to show the sensitivity to the horizontal angle and vertical position of the laser beam at the IP, and 6.14 shows the stability of those parameters with feedback ON. The colormap in both plots clearly show a region where the LACE efficiency was optimal. The stability of the measured LACE efficiency and sensitive laser parameters with feedback OFF and ON is given in Table 6.5. It should also be mentioned that the laser beam is likely to eventually drift away from the ion beam without the stabilization system, which would necessitate a repeat of the

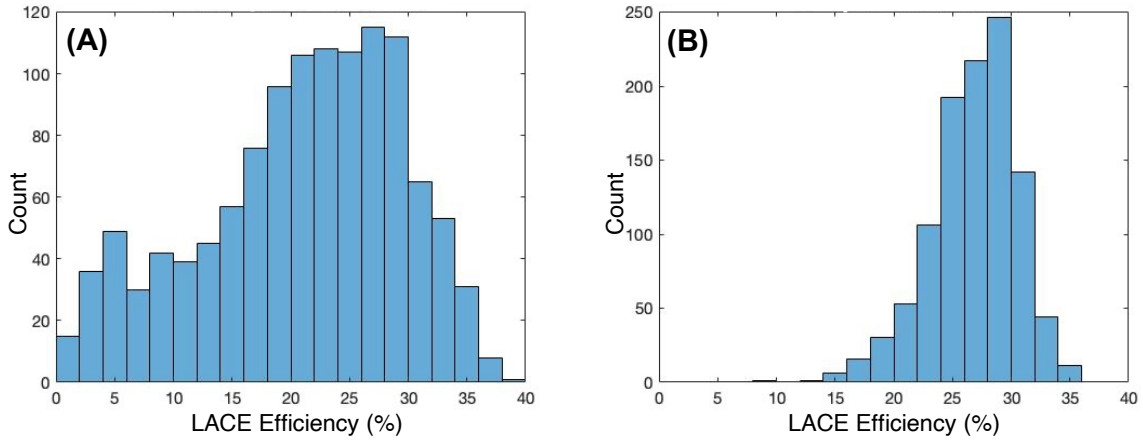

Figure 6.12: LACE efficiencies measured with feedback (A) OFF and (B) ON during the 1-step excitation experiment.

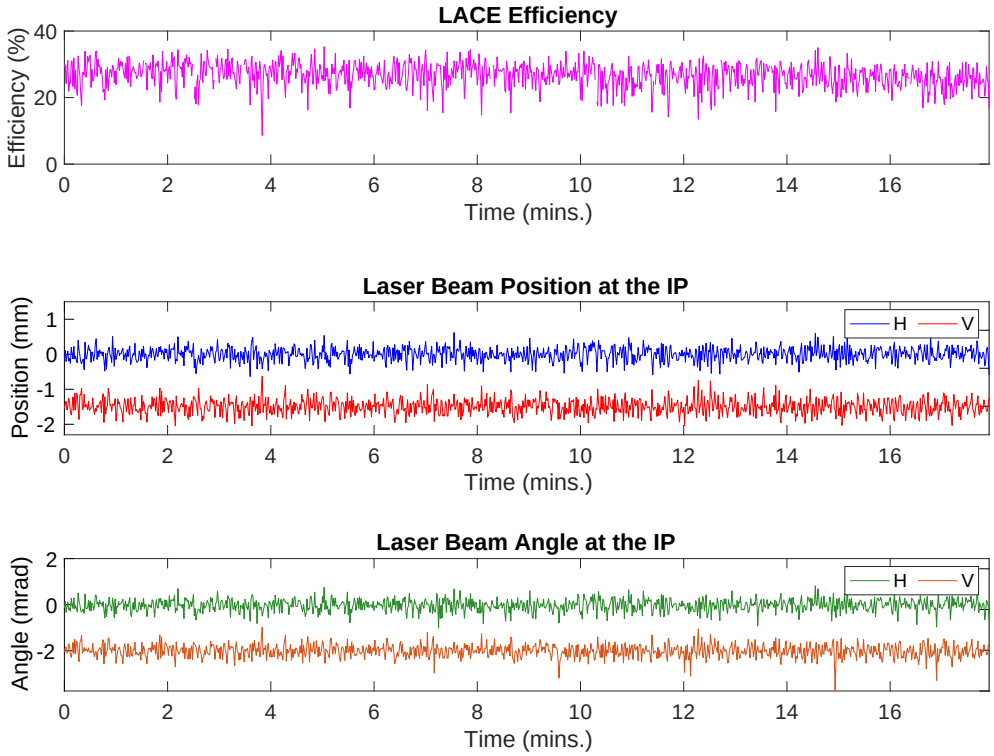

Figure 6.13: LACE efficiency and laser pointing stability during the 1-step experiment with feedback ON.

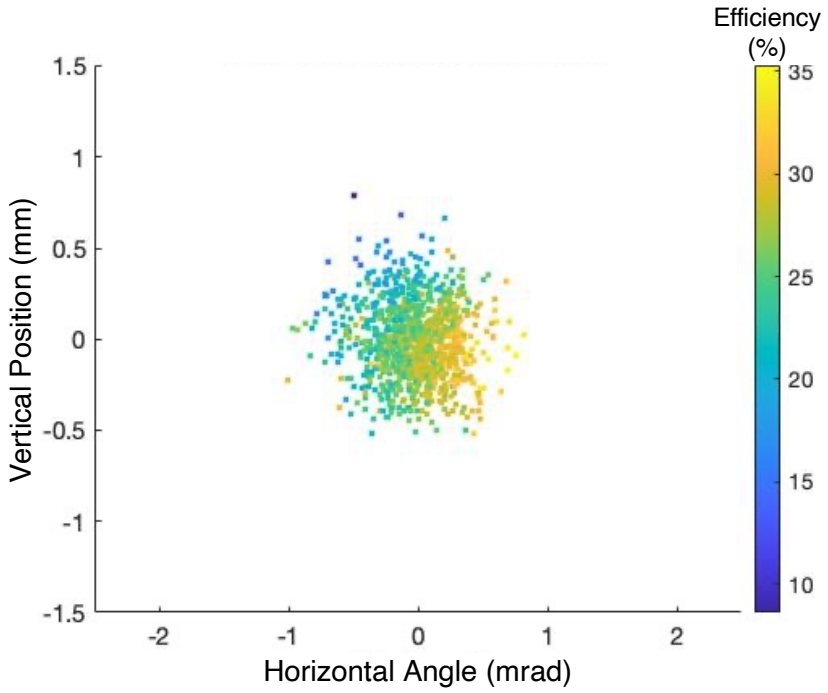

Figure 6.14: LACE efficiency as a function of the horizontal angle and vertical position of the laser beam.

Table 6.5: Stability with feedback OFF and ON during the 1-step experiment.

Parameter	Feedback OFF Value	Feedback ON Value
LACE Efficiency STD (%)	8.53	3.68
Horizontal Angle RMS (μrad)	744	263
Vertical Position RMS (μm)	347	199

alignment procedure outlined in the preceding paragraph. The stabilization system is a useful tool during experiments, and one like it would be necessary to sustain the operation of any LACE system that may be deployed in the future as a working CEI system.

Another study to evaluate the performance of the laser pointing stabilization system was conducted during the Sequential Excitation Experiment. Because the existing LACE chamber was not designed for sequential excitation, the position and angle of the second laser beam could not be measured precisely, and so the stability of the second beam will not be reported on. LACE efficiency distributions measured with feedback OFF and ON are shown in Fig. 6.15, and the waveforms for the feedback ON data set are shown in Fig. 6.16. The results of the feedback OFF/ON study are quantified in Table 6.6. The effect of the stabilization system is less obvious for the Sequential Excitation Experiment because the laser beam did not drift significantly with feedback OFF on the day of the experiment, and drift is the primary source of the large LACE efficiency variance over time. The LACE efficiency recorded with feedback ON is shown as a function of the horizontal angle and vertical position of the first laser beam in Fig. 6.17. The correlation is not as strong as in the 1-step experiment because the efficiency is sensitive to the horizontal angle and vertical position of the second beam as well. Still, it can be seen that there are, on average, higher efficiencies on one side of the distribution. Thorough optimization of the alignment of both beams later on in the experiment resulted in higher efficiencies. The distribution of efficiencies from the data set containing the highest recorded LACE efficiency of 12% is shown in Fig. 6.18. These results were recorded with feedback ON, however, the optimal alignment of the first beam turned out to be such that it was blocked by the wall of the back flange of the LACE chamber before it reached Cam2. As a result, the laser beam angle could not be measured, but it should be assumed the stability was equivalent to that shown in Fig. 6.16.

Preliminary results of the Crab Crossing Experiment will now be discussed briefly. LACE efficiency was measured for two different crab crossing angles to demonstrate

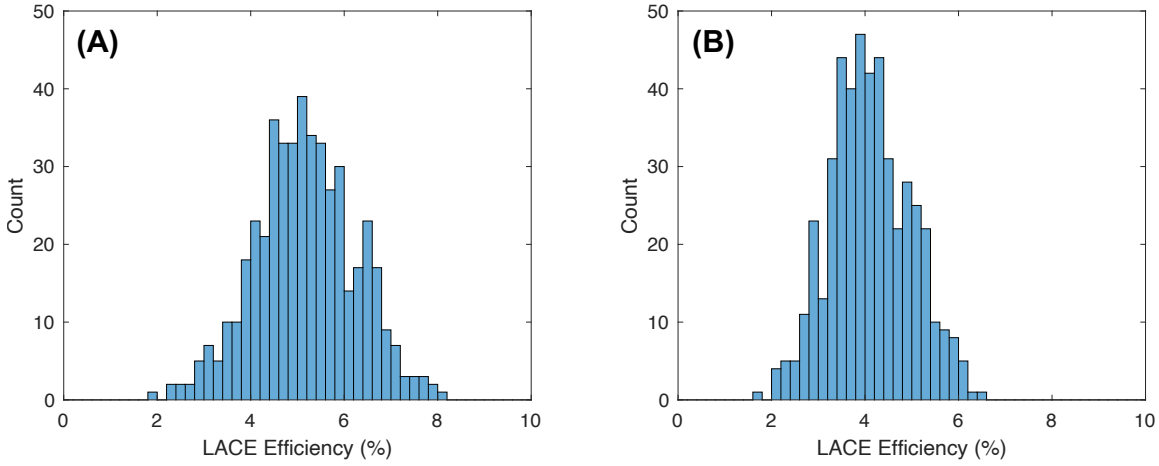

Figure 6.15: LACE efficiencies measured during the Sequential Excitation Experiment with feedback (A) OFF and (B) ON.

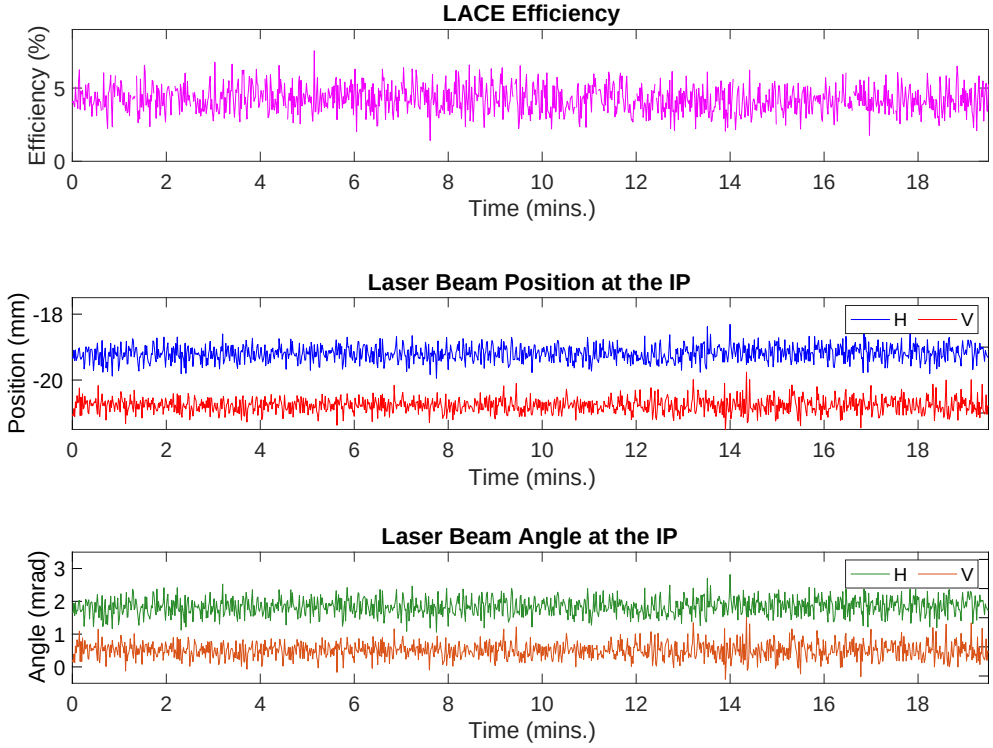

Figure 6.16: LACE efficiency and pointing stability of the first laser beam during the Sequential Experiment with feedback ON.

Table 6.6: Stability with feedback OFF and ON during the Sequential Excitation Experiment.

Parameter	Feedback OFF Value	Feedback ON Value
LACE Efficiency STD (%)	1.07	0.86
Horizontal Angle RMS (μrad)	426	265
Vertical Position RMS (μm)	488	280

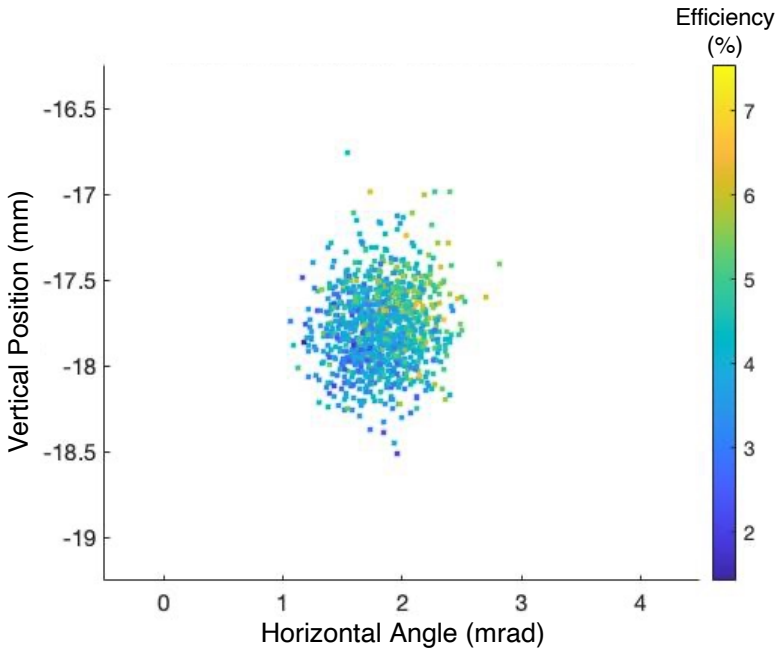

Figure 6.17: LACE efficiency as a function of the sensitive parameters of the first laser beam.

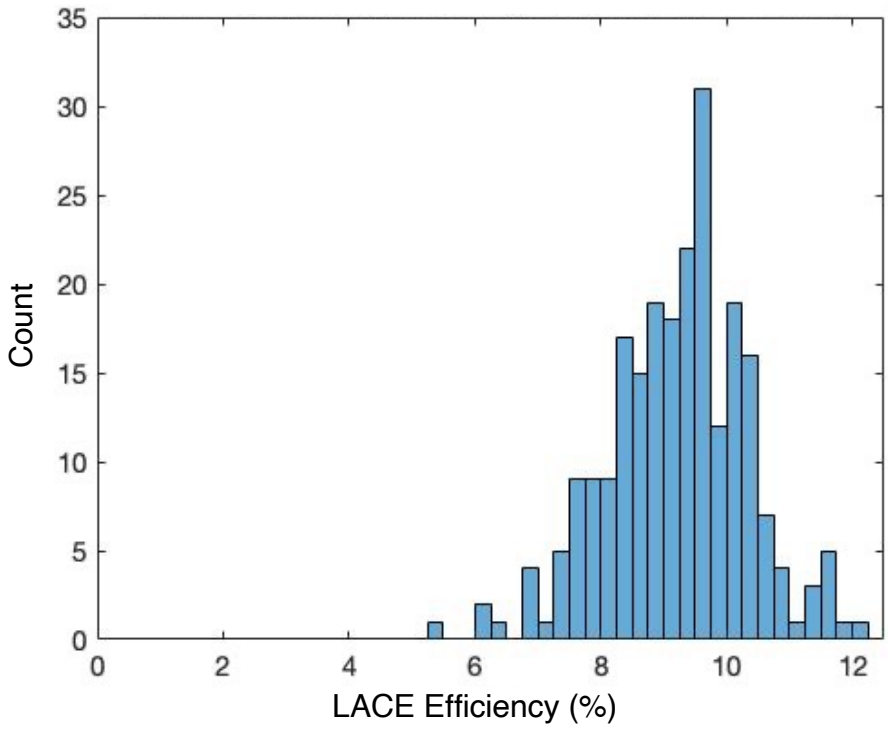

Figure 6.18: Highest LACE efficiency data recorded during the Sequential Excitation Experiment.

its effect. Efficiency distributions for rotations of 12^o and 19^o in the horizontal plane are shown in Fig. 6.19. The mean efficiency and maximum recorded in each data set are listed in Table 6.7. The mean LACE efficiency for the 19^o crab crossing angle was about $4\times$ greater than for 12^o because of the better spatial overlap of the laser and ion beam at the micropulse-level. This shows crab crossing can be used to reduce the average laser power required for high efficiency LACE by allowing sufficient overlap with a shorter laser pulse. The width of the distribution for the 19^o angle is significantly larger than for the 12^o angle due to the geometry of the collision (Fig. 6.6). The cross-sectional area of the ion beam in reference to the laser decreases as the crab crossing angle is increased, making it more likely that imperfect overlap will occur. Thus, crab crossing increases the sensitivity to laser jitter. Measurements of the position of the laser beam at the IP show the stability was unchanged between both data sets and confirm that greater sensitivity is due to geometry (Fig. 6.20). The RMS of the laser position with feedback ON was measured to be 252 (H) $\times$ 215 (V) μm over 45 minutes, which is comparable to other experiments. The angle of the beam at the IP could not be measured because the alignment of the laser beam resulted in it being blocked before Cam2, but it can be assumed the angular stability was comparable to previous experiments as well.

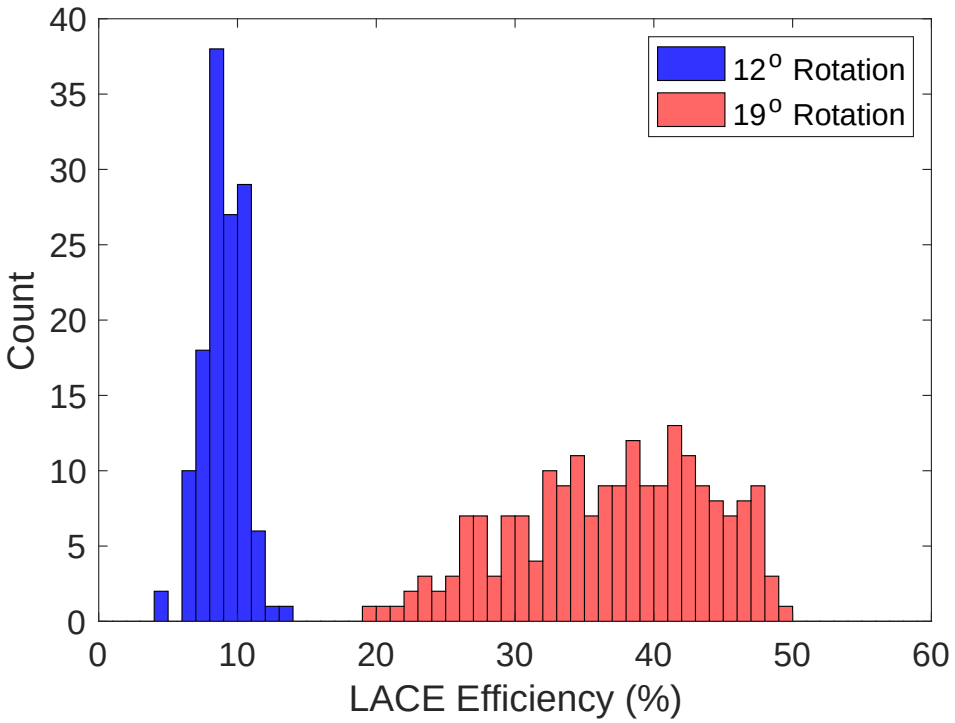

Figure 6.19: LACE efficiency for different crab crossing angles. Crab crossing makes the efficiency more susceptible to jitter by reducing the cross-sectional area of the ion beam in reference to the laser.

Table 6.7: Summary of preliminary results of the Crab Crossing Experiment.

Parameter	12^o Rotation	19^o Rotation
Maximum Efficiency (%)	13.2	49.8
Mean Efficiency (%)	8.97	37.0

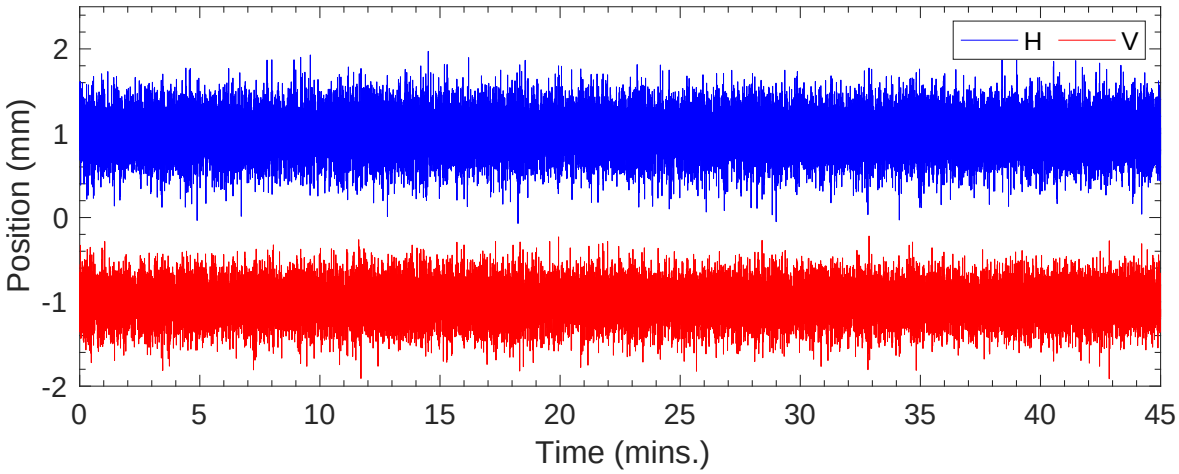

Figure 6.20: Laser position at the IP during the Crab Crossing Experiment.

Chapter 7

Conclusion

The ultimate goal of the LACE project at the SNS is to convert the production ion beam into protons with efficiency equivalent to injection foils. The three steps in the LACE process are (1) Lorentz stripping of one electron by a permanent magnet, (2) resonant excitation of the remaining electron from the ground state to $n = 3$ by the laser, and (3) Lorentz stripping of the excited electron by a second permanent magnet. The primary obstacle that prevents direct scaling of the original method used in the 2006 POP experiment to the full duty factor ion beam is the average laser power required to strip the ion beam with high efficiency. Ongoing experimental efforts at the SNS are focused on validating and testing methods which make more economical use of the limited available laser power. The 10 μs Experiment in 2016 employed several laser power saving techniques that allowed high efficiency LACE for a 10 μs long H^- pulse, a factor of 1000 times longer than in the POP experiment. The Sequential Excitation Experiment in 2022 was a demonstration of a method that reduces the peak laser power required for photoexcitation by breaking it up into two sequential resonant excitations from two laser beams. The Crab Crossing Experiment in 2023 has confirmed a method which allows a shorter laser pulse to be used through manipulation of the ion beam. Future experiments will center around combining these techniques after upgrades to the experimental design.

The overall efficiency of LACE is the product of the efficiencies of steps (1), (2) and (3). Lorentz stripping in steps (1) and (3) is a very stable process that is close to 100% efficient, and so it is step (2), photoexcitation by the laser, that determines the overall LACE efficiency. The variance in LACE efficiency measured during experiments is the direct result of variance in the photoexcitation efficiency. Sub-optimal photoexcitation is the result of the laser beam not fully overlapping the ion beam at the IP or not intersecting it at the correct angle to produce resonant photons in the rest frame of the H^0. The geometry of the laser-particle interaction makes the horizontal angle and vertical position of the laser beam at the IP the laser pointing parameters that LACE is most sensitive to. Disturbances along the LTL and thermal effects from the high energy laser beam cause it to slowly drift and jitter about an average trajectory after propagating 65 m to the IP. During experiments, drift causes the LACE efficiency to eventually fall to zero, and jitter is the source of shot-to-shot variation.

A laser pointing stabilization system was designed to stabilize the horizontal angle and vertical position of the laser beam at the IP. The system is based on feedback between two CMOS cameras and a PZT steering mirror before the LTL. One camera measures the angle of the beam after the LTL, and the other monitors an image of the beam at the IP. A PC running custom-made LabVIEW software acts as the controller, as well as a data acquisition system. Each 10 Hz iteration of the feedback loop, the controller software processes the raw camera images, extracts the coordinates of the beam centroid, and then updates and outputs the control signal going to the FBM. The feedback gain is the only free parameter of the control law, and it is tuned to optimize the performance of the system for a given disturbance signal. The feedback system is stable if the value of the feedback gain is between zero (feedback OFF) and 2 (unstable). When tuning the feedback gain, there is always a trade-off between suppressing low frequency components of the disturbance signal and amplifying high frequency ones. A feedback-based stabilization system is incapable of reducing the jitter when feedback is applied at or below the pulse repetition rate of the laser.

When the system is optimally tuned, drift will be eliminated, and the beam will stay jittering about the setpoint for an indefinite amount of time.

The laser pointing stabilization system is an improvement in the experimental design of LACE. It reliably stops the laser beam from drifting away from the ion beam and eliminates the need for repeated realignment during experiments. It also makes it easier to confirm the optimal alignment of the laser beam by providing a more stable signal for the stripped beam to use as a benchmark. However, jitter must also be eliminated if LACE is to be used as a working system for CEI. There are three options that could reduce the jitter with the existing laser and LTL. First is to use passive stabilization techniques, such as those described in Section 4.1, to avoid or suppress high frequency disturbances in the LTL. This option is the most likely to be successful but would require a significant investment of time and resources. The second option is to stabilize a CW or high repetition rate pilot laser beam that is aligned to overlap the optical path of the UV laser. With a pilot beam, the rate of feedback could potentially be increased to the kHz range, but there is no guarantee that its stability will correlate well with that of the UV beam, especially over an extended period of time. A third option to reduce the jitter is to develop a feedforward control scheme capable of reliably predicting the pointing fluctuations and correcting for them ahead of time. However, there is no obvious physics model for the pulse-to-pulse jitter that be constructed with the available information, and machine learning techniques are likely to yield limited success given the seasonal changes in the disturbance and amount of data needed for training. Even if all three options are combined successfully, LACE is still unlikely to match injection foils in terms of stability. A future LACE system that could be used for CEI at the SNS will likely need to use a more stable laser, such as the green lasers proposed in [48], and a transport distance several times less than the 65 meters in the existing system.

Bibliography

[1] The U.S. Department of Energy. DOE public access plan. http://energy.gov/downloads/doe-public-access-plan, 2022. iii

[2] S. Hendersen *et al.* The Spallation Neutron Source accelerator system design. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 763:610–673, 2014. 1, 5, 6, 8

[3] M. S. Champion, C. N. Barbier, M. S. Connell, E. Daly, N. J. Evans, J. Galambos, D. J. Harding, M. P. Howell, G. D. Johns, S.-H. Kim, J. S. Moss, B. W. Riemer, and K. S. White. Progress on the proton power upgrade project at the spallation neutron source. In *Proc. 31st International Linear Accelerator Conference (LINAC'22)*, number 31 in International Linear Accelerator Conference, pages 390–393. JACoW Publishing, Geneva, Switzerland, 09 2022. 2

[4] Oak Ridge National Laboratory. How SNS Works. https://neutrons.ornl.gov/content/how-sns-works. 8

[5] L. W. Alvarez. Energy doubling in dc accelerators. *Review of Scientific Instruments*, 22(9):705–706, 1951. 7

[6] S. Y. Lee. *Accelerator Physics*. WORLD SCIENTIFIC, 4th edition, 2019. 7

[7] M. A. Plum. H^- Charge Exchange Injection Issues at High Power. In *Proc. of ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB'16), Malmö, Sweden, July 3-8, 2016*, pages 304–309, 2016. 9, 10

[8] M. A. Plum, J. Holmes, R. W. Shaw, and C. S. Feigerle. SNS stripper foil development program. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 590(1):43–46, 2008. 10

[9] A. N. Zelenskiy, S. A. Kokhanovskiy, V. M. Lobashev, N. M. Sobolevskiy, and E. A. Volferts. A method of polarizing relativistic proton beams by laser radiation. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 227(3):429–433, 1984. 12

[10] I. Yamane. H^- charge-exchange injection without hazardous stripping foils. *Phys. Rev. ST Accel. Beams*, 1:053501, Sep 1998. 12, 16, 17

[11] V. Danilov, A. Aleksandrov, S. Assadi, S. Henderson, N. Holtkamp, T. Shea, A. Shishlo, Y. Braiman, Y. Liu, J. Barhen, and T. Zacharia. Three-step h^- charge exchange injection with a narrow-band laser. *Phys. Rev. ST Accel. Beams*, 6:053501, May 2003. 12, 24

[12] V. Danilov, A. Aleksandrov, S. Assadi, J. Barhen, W. Blokland, Y. Braiman, D. Brown, C. Deibele, W. Grice, S. Henderson, J. Holmes, Y. Liu, A. Shishlo, A. Webster, and I. N. Nesterenko. Proof-of-principle demonstration of high efficiency laser-assisted h^- beam conversion to protons. *Phys. Rev. ST Accel. Beams*, 10:053501, May 2007. 12, 19, 24

[13] L. R. Scherk. An improved value for the electron affinity of the negative hydrogen ion. *Canadian Journal of Physics*, 57(4):558–563, 1979. 15

[14] A. J. Jason, D. W. Hudgings, and O. B. van Dyck. Neutralization of h- beams by magnetic stripping. *IEEE Transactions on Nuclear Science*, 28(3):2703–2706, 1981. 15

[15] S. Cousineau, A. Rakhman, M. Kay, A. Aleksandrov, V. Danilov, T. Gorlov, Y. Liu, M. Plum, A. Shishlo, and D. Johnson. First demonstration of laser-assisted charge exchange for microsecond duration h^- beams. *Phys. Rev. Lett.*, 118:074801, 2017. 19, 29, 32

[16] S. Cousineau, A. Rakhman, M. Kay, A. Aleksandrov, V. Danilov, T. Gorlov, Y. Liu, C. Long, A. Menshov, M. Plum, A. Shishlo, A. Webster, and D. Johnson. High efficiency laser-assisted h^- charge exchange for microsecond duration beams. *Phys. Rev. Accel. Beams*, 20:120402, 2017. 19, 22, 23, 27, 29, 32

[17] K. Halbach. Design of permanent multipole magnets with oriented rare earth cobalt material. *Nuclear Instruments and Methods*, 169(1):1–10, 1980. 21

[18] C. Huang, C. Deibele, and Y. Liu. Narrow linewidth picosecond uv pulsed laser with mega-watt peak power. *Opt. Express*, 21(7):9123–9131, Apr 2013. 21

[19] T. V. Gorlov. Laser Stripping H^- Charge Exchange Injection by Femtosecond Lasers. In *Proc. of ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB'16), Malmö, Sweden, July 3-8, 2016*, number 57 in ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams, pages 212–215, Geneva, Switzerland, Aug. 2016. JACoW. doi:10.18429/JACoW-HB2016-MOPL011. 26

[20] Y. Liu, A. Rakhman, A. Menshov, A. Webster, T. Gorlov, A. Aleksandrov, and S. Cousineau. Laser and optical system for laser assisted hydrogen ion beam stripping at sns. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 847:171–178, 2017. 28, 30

[21] R. Paschotta. The rp photonics encyclopedia. https://www.rp–photonics.com/beam_pointing_fluctuations.html, 2008. 37

[22] B. E. A. Salah and M. C. Teich. *Fundamentals of Photonics*, chapter Ray Optics, pages 1 – 40. Wiley, 1991. 40, 74

[23] Y. Wu, D. French, and I. Jovanovic. Passive beam pointing stabilization. *Opt. Lett.*, 35:250–252, 2010. 46

[24] K. Tyszka and M. Dobosz. Laser beam angular stabilization system based on a compact interferometer and a precise double-wedge deflector. *Review of Scientific Instruments*, 89:085121, 2018. 46

[25] H.-H. Chu, S.-Y. Huang, and L.-S. *et al.* Yang. A versatile 10-tw laser system with robust passive controls to achieve high stability and spatiotemporal quality. *Appl. Phys. B*, 79:193–201, 2004. 46

[26] S. Borneis, T. Laštovička, M. Sokol, T.-M. Jeong, F. Condamine, O. Renner, V. Tikhonchuk, H. Bohlin, A. Fajstavr, J.-C. Hernandez, and et al. Design, installation and commissioning of the eli-beamlines high-power, high-repetition rate hapls laser beam transport system to p3. *High Power Laser Science and Engineering*, 9:e30, 2021. 46

[27] J. Bechhoefer. *Control Theory for Physicists*. Cambridge University Press, 1st edition, 2021. 47

[28] R. Hardin, Y. Liu, C. Long, A. Aleksandrov, and W. Blokland. Active beam position stabilization of pulsed lasers for long-distance ion profile diagnostics at the spallation neutron source (sns). *Opt. Express*, 19:2874–2885, 2011. 48, 59

[29] L. K. Nguyen, A. J. Curcio, W. J. Eisele, A. V. Fedotov, A. Fernando, W. Fischer, P. Inacker, J. P. Jamilkowski, D. Kayran, K. Kulmatycski, D. Lehn, T. A. Miller, M. G. Minty, and A. Sukhanov. Active Pointing Stabilization Techniques Applied

to the Low Energy RHIC Electron Cooling Laser Transport at BNL. In *Proc. NAPAC'19*, pages 938–941, 2019. 48

[30] P. Cinquegrana, S. Cleva, A. Demidovich, G. Gaio, R. Ivanov, G. Kurdi, I. Nikolov, P. Sigalotti, and M. B. Danailov. Optical beam transport to a remote location for low jitter pump-probe experiments with a free electron laser. *Phys. Rev. ST Accel. Beams*, 17:040702, 2014. 48

[31] T. Kanai, A. Suda, S. Bohman, M. Kaku, S. Yamaguchi, and K. Midorikawa. Pointing stabilization of a high-repetition-rate high-power femtosecond laser for intense few-cycle pulse generation. *Applied Physics Letters*, 92:061106, 2008. 48

[32] R. Singh, K. Patel, J. Govindarajan, and A. Kumar. Fuzzy logic based feedback control system for laser beam pointing stabilization. *Appl. Opt.*, 49:5143–5147, 2010. 48

[33] F. Kaiser, S. Köhler, F. Peters, L. Winkelmann, and I. Hartl. Uv laser beam stabilization system for the european xfel electron injector laser beamline. In *CLEO: 2015*, page JTh2A.99, 2015. 48

[34] S. Grafström, U. Harbarth, J. Kowalski, R. Neumann, and S. Noehte. Fast laser beam position control with submicroradian precision. *Optics Communications*, 65:121–126, 1988. 48

[35] C. W. Siders, E. W. Gaul, M. C. Downer, A. Babine, and A. Stepanov. Self-starting femtosecond pulse generation from a ti:sapphire laser synchronously pumped by a pointing-stabilized mode-locked nd:yag laser. *Review of Scientific Instruments*, 65:3140–3144, 1994. 48

[36] A. Stalmashonak, N. Zhavoronkov, I. V. Hertel, S. Vetrov, and K. Schmid. Spatial control of femtosecond laser system output with submicroradian accuracy. *Appl. Opt.*, 45:1271–1274, 2006. 48

[37] G. Genoud, F. Wojda, M. Burza, A. Persson, and C.-G. Wahlström. Active control of the pointing of a multi-terawatt laser. *Review of Scientific Instruments*, 82:033102, 2011. 48

[38] M. Kay, A. Rakhman, and S. Cousineau. Active pointing control of pulsed, high-power laser beam after 65-meter transport. In *2022 IEEE 17th International Conference on Control & Automation (ICCA)*, pages 565–570, 2022. 48

[39] J. DiStefano, A. Stubberud, and I. Williams. *Schaum's Outline of Feedback and Control Systems*, chapter Nyquist Analysis, pages 246 – 298. McGraw-Hill Education, 1990. 53

[40] M. Kay and A. Oguz. Pointing Stabilization of 140 mJ, 10 Hz UV Laser for Laser-Assisted Charge Exchange. *submitted to Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 2023. 54, 55, 104, 107, 108, 109, 111

[41] Y. Liu, A. Aleksandrov, S. Assadi, W. Blokland, C. Deibele, W. Grice, C. Long, T. Pelaia, and A. Webster. Laser wire beam profile monitor in the spallation neutron source (sns) superconducting linac. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 612(2):241–253, 2010. 57, 64

[42] Y. Liu, A. Aleksandrov, C. Long, A. Menshov, J. Pogge, A. Webster, and A. Zhukov. Nonintrusive emittance measurement of 1 gev h beam. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 675:97–102, 2012. 57

[43] S. Assadi. SNS transverse and longitudinal laser profile monitor design, implementation, results, and improvement plans. In *Proc. LINAC'06*, 2006. 57

[44] A. Aleksandrov. Calculation of the neutralization degree in the h- beam due to interaction with laser beam. SNS-Note-AP-44. 57, 61

[45] Y. Liu, C. Long, C. Huang, R. Dickson, and A. Aleksandrov. Simultaneous ion beam profile scan using a single laser source. *Phys. Rev. ST Accel. Beams*, 16:012801, Jan 2013. 63, 65

[46] Thorlabs. Piezoelectric Tutorial. 83

[47] National Instruments. Vision Acquisition Software, 2020. 84

[48] T. Gorlov, A. Aleksandrov, S. Cousineau, Y. Liu, A. Rakhman, and A. Shishlo. Sequential excitation scheme for laser stripping for a h^- beam. *Phys. Rev. Accel. Beams*, 22:121601, 2019. 95, 96, 97, 98, 124

[49] A. Aleksandrov, S. Cousineau, T. Gorlov, Y. Liu, A. Oguz, A. Shishlo, A. Zhukov, and M. Kay. Experimental demonstration of sequential excitation scheme for h^- laser assisted charge exchange. *accepted by Physical Review Accelerators and Beams for publication*, 2023. 95, 100

[50] A. Aleksandrov, S. Cousineau, T. Gorlov, Y. Liu, A. Rakhman, and A. Shishlo. A Crab-Crossing Scheme for Laser-Ion Beam Applications. In *Proc. NAPAC'19*, number 4 in North American Particle Accelerator Conference, pages 639–641. JACoW Publishing, Geneva, Switzerland, 10 2019. 95, 103

Vita

Born in Chicago, Illinois, Martin Joseph Kay grew up in Valparaiso, Northwest Indiana and graduated from Valparaiso High School in 2012. He began his post-secondary education at Purdue University, Indiana, where he had his first research experiences and received a Bachelor of Science degree in Physics in Spring 2016. The following Summer, he began working at the Spallation Neutron Source accelerator in Oak Ridge National Laboratory as an intern for the Accelerator Physics and Beam Instrumentation groups. He continued his work at the Spallation Neutron Source as a Graduate Research Assistant after he entered the Doctor of Philosophy program in Physics at The University of Tennessee, Knoxville in Fall 2016. He received a Master of Science degree in Physics in August 2019 and graduated with a Doctor of Philosophy degree in Physics in May 2023. His research interests intersect the fields of optics, accelerator physics, and control systems engineering. After graduation, he will begin a new position as a Postdoctoral Research Associate at Los Alamos National Laboratory, New Mexico. He is grateful for all the support and guidance he received from his mentors as he begins his new career.